

Role of AI in Thermal Control for Power Plants: Techniques, Challenges and Opportunities

Qiang Liu¹, Rui Zhu¹, Guofeng Li¹, Lei Zhang¹, Jun Li¹, Lin Wang^{2*}, Yanjun Guo²

¹Shandong Energy Group Lingtai Thermal Power Generation Co., Ltd, Pingliang, 744400 Gansu, China

²Xi'an Thermal Power Research Institute Co., Ltd, Xi'an, 710054 Shaanxi, China

Abstract

INTRODUCTION: Power plants require efficient thermal control to reduce operational costs and improve energy performance. Conventional methods may have limited adaptability to changing operating and environmental conditions. AI-based techniques offer a promising solution for improving thermal efficiency and predictive maintenance.

OBJECTIVES: This study aims to develop an AI-based framework for optimizing power plant thermal control, improving energy efficiency, and predicting equipment failures using Reinforcement Learning (RL) and Artificial Neural Networks (ANNs).

METHODS: The Coal Fired Power Plant Thermal Performance Dataset is used to train the proposed models. RL is applied to optimize boiler temperature and turbine pressure, while ANNs are employed for failure prediction. The Multiverse Optimizer is used to optimize model hyperparameters and improve performance.

RESULTS: The proposed framework achieves an R^2 of 0.9528, MAE of 0.0216, RMSE of 0.0237, and MAPE of 0.023%, demonstrating high prediction accuracy and improved thermal efficiency compared with conventional approaches.

CONCLUSION: The proposed AI-driven framework provides an adaptive and efficient approach to power plant thermal management. The integration of RL, ANN, and the Multiverse Optimizer enhances thermal optimization and predictive maintenance, supporting more reliable and efficient power plant operations.

Keywords: AI-driven thermal control, Reinforcement Learning, predictive maintenance, Multiverse Optimizer, power plants

Received on 16 June 2026, accepted on 05 August 2026, published on 13 October 2026

Copyright © 2026 Qiang Liu et al., licensed to EAI. This is an open access article distributed under the terms of the CC BY-NC-SA 4.0, which permits copying, redistributing, remixing, transformation, and building upon the material in any medium so long as the original work is properly cited.

*Corresponding author. Email: linwang_123@outlook.com

1. Introduction

The importance of power plants in fulfilling the global energy needs is glaring and it is important to maximize the use of thermal control system in ensuring efficiency [1] of such power plants besides minimizing the cost of operation [2]. Conventional thermal control technologies tend to be dependent on fixed parameters and manual controls that do not consider the current changes in power generation and the environmental factors [3]. As the transition to sustainable energy production keeps gaining momentum [4], incorporating the usage of Artificial Intelligence (AI) into thermal systems could be used in the future to greatly enhance the efficiency and fuel consumption of the thermal system [5], thus making AI-powered solutions critical to

the future of power plants. This framework discusses how AI methods [6], including reinforcement learning and neural networks [7], can be used to improve thermal control and predictive maintenance to resolve existing inefficiencies [8].

The literature has suggested different approaches to thermal optimization in power plants such as fuzzy logic control [9], machine learning-based approaches, as well as conventional control systems [10]. Other methods such as Fuzzy Logic Controllers (FLC) have been applied in thermal regulation but lack the flexibility to changing operational conditions in real-time [11]. Other methods such as Support Vector Machines (SVM) and Artificial

Neural Networks (ANNs) [12], have as well been used to predict thermal behaviour and optimise power generation [13], but are not always robust in a dynamic environment. Such approaches are not easy to modify parameters on the fly [14], maximize energy savings, and accurately predict system breakdown, which cause more downtime and energy waste [15].

Modern power plants face an increasingly complex set of operational challenges, particularly in coal-fired thermal generation systems, where efficient and reliable operation must be maintained under dynamic operating conditions. Fluctuations in fuel quality and calorific value can adversely affect boiler combustion stability, while continuously varying electricity demand requires frequent adjustments to boiler temperature, steam flow, and turbine operating parameters that conventional fixed-setpoint controllers often fail to accommodate effectively [16]. Furthermore, aging infrastructure and progressive equipment degradation increase the likelihood of unexpected failures and unplanned outages, whereas traditional time-based or reactive maintenance strategies frequently result in either unnecessary maintenance activities or costly operational interruptions [17].

In parallel, increasingly stringent environmental regulations demand higher combustion efficiency and lower emissions, requiring more precise and adaptive thermal management than conventional control systems can provide [18]. The growing integration of renewable energy sources into modern power grids further introduces fluctuations in power generation and grid stability, requiring thermal power plants to operate with greater flexibility and responsiveness to maintain reliable electricity supply [19]. Collectively, these operational, economic, and environmental challenges highlight the limitations of conventional thermal control approaches and emphasize the need for intelligent management strategies. Artificial Intelligence (AI)-driven techniques, capable of analysing real-time operational data, dynamically optimizing control decisions, predicting equipment failures, and continuously adapting to changing operating conditions, provide an effective solution for improving operational efficiency, enhancing system reliability, reducing maintenance costs, and supporting sustainable power generation. Accordingly, these challenges provide the primary motivation for the proposed Reinforcement Learning (RL), Artificial Neural Networks (ANNs), and Multiverse Optimizer (MVO)-based intelligent thermal control framework presented in this study.

The proposed framework addresses these disadvantages by integrating Reinforcement Learning (RL) to control the dynamic thermal process, and therefore real-time adaptation to new conditions will be achieved. Moreover, Artificial Neural Networks (ANNs) are also used in predictive maintenance, which involves determining the failures before they happen. The Multiverse Optimizer improves the optimization process by optimizing the model hyperparameters resulting in more accurate predictions and low energy consumption. This study is innovative in that the AI-guided solutions are combined to continuously learn

about the behavior of plants to optimize thermal performance in a dynamic manner, providing an intelligent, dynamically adaptive way of operating power plants of the present day.

1.1 Research Objectives

- Interpret the overall purpose of the work: To examine the opportunities of AI-based approaches to optimization of thermal management in power plants, energy efficiency and predictive maintenance.
- Assess the data in the proposed structure: Coal Fired Power Plant Thermal Performance Data Set is utilized to create, train, and test the AI models of thermal control and system optimization.
- Use reinforcement learning: The structure involves the application of Reinforcement learning (RL) to achieve dynamic optimization of thermal control and provide the opportunity to jump to real-time changes in the working parameters following ongoing feedback.
- Build simulated neural networks: The paper incorporates Artificial Neural Networks (ANNs) to predictive maintenance, which will use historical data to estimate equipment failure and avoid unexpected downtimes.

1.2 Research Scope and Novelty

The areas of application of the proposed framework are optimization of thermal systems through the use of AI and improved predictive maintenance in power plants. It is devoted to real-time optimization of thermal with the help of Reinforcement Learning and ANNs with the Multiverse Optimizer to fine-tune the model. The originality of the work is in the opportunity to combine all of these AI methods into a single structure of continuous learning and active regulation to overcome the drawbacks of the classical thermal control methods. This not only increases the efficiency of power plants but also increases the reliability and sustainability of the operations of the power plant.

1.3 Research Organization

The paper will be structured in the following way:

- Section 2: Literature review of current approaches to thermal control and AI use in power plants.
- Section 3: Methodology, covering the AI models, data set and the framework that is suggested to be used to optimize thermal.
- Section 4: Results and discussion, model performance and optimization results.
- Section 5: Conclusion and future outlooks of the research in thermal control systems that are driven by AI.

2. Literature Review

Alhamrouni et al., [20] The paper is aimed at enhancing the stability, control and protection of power systems. The authors discuss how the uses of artificial intelligence (AI), deep learning (DL), machine learning (ML), and metaheuristic optimization algorithms can improve the performance of the power system. Nevertheless, the limitation of the framework is in the inadequate consideration of the growing dynamicity of modern power grid, particularly, integrating a variety of power sources. Ukoba et al., [21] This paper is a review of the AI-renewable energy system (RES) as a way of maximizing the efficiency and reliability of the power grid. It focuses on the application of machine learning algorithms, neural networks and optimization techniques to improve predictive performances and adaptive system management. Variability of data and real-time adjustability difficulties in the application of AI in RES are the challenge identified and influencing the quality of predictions. Franki et al., [22] The article explores the current revolution that is going on in the power sector especially through renewable energy, distributed energy sources and advanced metering systems. The authors are interested in the concept of integrating the AI solutions to enhance the demand prediction, system optimization, and fault detection. One of the disadvantages is the absence of overall digital solutions to treat the large volumes of the data produced by new technologies in the power industry. Devasenan and Madhavan, [23] The study looks at how AI has been used in thermal engineering, industrial processing and solar thermal systems to enhance energy efficiency and system stability. It emphasizes on the application of deep reinforcement learning and gradient-boosted regression trees in optimization of thermal systems. There are, however, limitations in the areas of matching AI solutions to industrial needs especially in aligning AI models to industrial processes.

Chekifi et al., [24] The paper is dedicated to thermal energy storage (TES) and its optimization by the use of AI methods such as particle swarm optimization, neural networks, and adaptive neurofuzzy inference systems. The techniques are used to improve the efficiency and sustainability of energy storage systems. The strength of the AI-driven solutions is, though, that it is difficult to realize the solutions on large-scale energy storage systems in the real world. Buffa et al., [25] This article is an overview of the new methods in district heating and cooling (DHC) systems, especially with the applications of AI and machine learning algorithms to enhance the performance of the system. The article compares machine learning and model predictive control in optimization of a system. The key weakness lies in the fact that it is hard to tackle the problem of load uncertainty and dynamic control, and AI approaches are yet to be efficient at doing so. Thawait and Singh, [26] This review tells about how AI and ML can be used to optimize various parameters in Thermal Power Plants to make them more productive. The literature concentrates on machine learning models and AI

to enhance thermal system performance. The limitations that are however associated with the framework are that it responds narrowly to complex, multi-dimensional issues that thermal plants have to deal with in the real world. Pimenow et al., [27] The article explores the applications of AI and machine learning in energy saving and how they can help to counter climate change. The authors discuss the application of AI in energy optimization, integration of renewable energy, and reduction of carbon. The weakness of this system is that training large models is associated with high energy usage that can negate the sustainability aspects of AI-based solutions in energy management.

Mazhar et al., [28] This article examines how AI, the Internet of Things (IoT), and smart grids (SG) can be used to increase the energy management and energy efficiency of smart buildings. It is research on machine learning algorithms on predicting energy demand in IoT-based applications. Nevertheless, issues of data scalability and interpretability still exist, particularly when it comes to large-scale integrations of the IoT. Zhang and Strbac, [29] The sources in this review assess the use of AI in energy systems centered on reinforcement learning, multi-agent systems, and AI to achieve power grid resilience. The article focuses on the optimization of power grid with the help of AI methods such as reinforcement learning. Nevertheless, the absence of practical validation and issues with reproducibility of the model are critical items leading to the reasonable implementation of such AI models in energy systems. Fan et al., [30] The article is a review of the uses of AI and deep learning to achieve sustainability in different fields, such as renewable energy and environmental health. It dwells upon the application of AI and DL to streamline energy management, fault detection and stability of power grids. The absence of transparency and explainability of AI models is one of the main limitations, particularly in very complicated energy systems. Alahi et al., [31] This article is an exploration of the concept of AI and IoT architecture and its contribution to the development of smart cities in order to boost sustainability and efficiency in cities. The authors discuss AI algorithms used to manage the system of IoTs and make decentralized decisions in urban energy management and resource allocation. The primary weakness is that it is challenging to guarantee real-time data processing in massive smart city systems, the real-time flexibility and decision-making of AI models.

Li et al. [32] deployed a Q-learning-enhanced NSGA-II algorithm for coal blending optimization at the Huaneng Yingkou coal-fired power plant, reporting a 14.7% fuel-cost reduction and 41% lower slagging incidence over twelve months of operational data. Panda et al. [33] applied a modified Multi-Verse Optimizer to tune PID parameters for temperature control in a thermal power plant condenser, demonstrating reduced overshoot and settling time relative to standard MVO, genetic algorithms, and Ziegler-Nichol's tuning. Ismail et al. [34] showed that ANN-based intelligent monitoring systems can detect coal-fired boiler trips several minutes ahead of a plant's native control system. Sun et al. [35] developed AI surrogate models

(ELM, GPR, ISCN, LSTM) trained on CFD data to rapidly predict heat-exchanger thermal-hydraulic performance, illustrating the broader shift toward data-driven thermal system design that this framework builds on.

The studies conducted on AI use in power systems, energy optimization, and the management of smart grids present various limitations and issues that should be overcome. They comprise challenges in bringing AI solutions to bear on the real-world activities in industry, especially matching the models with the particular requirements of thermal plants and large-scale energy storage systems. A large number of current systems are not well versed with data variability, scaling, and interpretability, particularly in complex and variable environments such as power grids and renewable energy systems. Also, though AI-based solutions can provide substantial benefits in terms of energy efficiency and fault detection, problems of real-time flexibility, absence of transparency, and high energy usage associated with the training of large models still are not addressed. Moreover, the lack of pragmatic validation of most models, as well as the difficulty in managing uncertainty within the load management and system control, constraints the scalability and efficiency of the AI-based solutions. These gaps indicate the requirement of more robust, versatile, and scalable AI systems that can handle the complexity of the modern energy systems.

3. Proposed Ai-Driven Thermal Control Optimization in Power Plants Methodology

The methodology suggested implies a number of essential steps, the first of them being the data collection and preprocess, then applications of AI-based models to accomplish thermal optimization and predictive maintenance. The workflow includes the stage of gathering of thermal, operational, and environmental data with the sensors installed in power plants. This data is pre-processed and is subjected to noise removal, normalisation and feature selection to guarantee quality input to AI models. Then, the Reinforcement Learning (RL) is used to maximize the thermal parameters, including boiler temperature and turbine pressure, on-the-fly by using continuous feedback. At the same time, Artificial Neural Networks (ANNs) are also trained to foresee possible system breakdowns and direct predictive maintenance. The Multiverse Optimizer optimizes the models by hyper-optimizing models to increase their accuracy and efficiency. The performance of the system is measured based on indicators such as the energy efficiency, system stability and the accuracy of prediction of failure and the AI models must offer steady optimization. Lastly, the real time thermal control system will be installed to put in place the automation in terms of the adjustments in the operations of the power plants, to boost not only the thermal efficiency but also the reliability as shown in Figure 1.




Figure 1. Block Diagram of the Proposed Framework for AI-Driven Thermal Control in Power Plants

The diagram shows the process of the suggested technology of power plants optimization in terms of thermal control to apply AI technology. It starts with the

data collection phase since thermal and operational, environmental and historical data of the Coal Fired Power Plant Thermal Performance Dataset is collected. The

received data are then undergone data preprocessing whereby noise is eliminated, features selected, data cleaned and normalized/scaled ready to be fed into an AI model. At the stage of preprocessing, the AI models are trained and tested, and to optimize thermal performance, the Reinforcement Learning (RL) framework is utilized in this case, whereas predictive maintenance is performed based on Artificial Neural Networks (ANNs). The Multiverse Optimizer is used to train the models after training to optimize their accuracy and performance. The effectiveness of the thermal optimization and predictive maintenance models of the system is then measured in terms of different performance measures including Accuracy, MAE, RMSE, MAPE, and R² that determine the efficiency of the models in real time operations of the plant.

3.1 Dataset Description

The data that will be incorporated in the proposed framework is the Coal Fired Power Plant Thermal Performance Dataset [36], which includes operational and performance information of a coal-fired power plant. It also comprises several thermal conditions that include temperature, pressure, steam flow rates and energy output. The data also covers environmental information such as ambient temperature and historical operation information, which gives an insight into the performance of the plant and malfunction historically. This data is to be used as a root of training AI networks to forecast and optimize thermal control systems. It is also very applicable to reinforcement learning and predictive maintenance activities whereby it aids in the behavior of the thermal systems within the power plant at different operating conditions. With the help of this dataset, the AI models can dynamically optimize the operations and avoid failures according to past trends.

3.2. Data Preprocessing

Noise Removal: Apply Smoothing Filters or Median Filtering to remove anomalies and outliers in the data is represented in Eqn (1):

$$\text{Filtered Value} = \frac{1}{n} \sum_{i=1}^n \text{Value}_i \quad (1)$$

where n is the window size and Value_i are the data points in the window.

Data Cleaning: Handle missing values by using MICE (Multiple Imputation by Chained Equations) or K-Nearest Neighbours (KNN) imputation is represented in Eqn (2):

$$\hat{x}_i = \frac{1}{k} \sum_{i=1}^k x_i \quad (2)$$

where $\hat{x}_i$ is the imputed value, and k is the number of nearest neighbours.

Normalization/Scaling: standardize features to get them into a standard range either by Min-max scaling or Score Normality is represented in Eqn (3):

$$X_{\text{norm}} = \frac{X - \min(X)}{\max(X) - \min(X)} \quad (3)$$

where X is the original feature value, and X_{norm} is the normalized value.

3.2.1. Feature Selection

A hybrid feature selection strategy combining the Relief algorithm and Random Forest (RF) Feature Importance was employed. The Relief algorithm evaluates the relevance of each feature by measuring its ability to distinguish neighbouring instances under similar and dissimilar operating conditions, while the Random Forest model estimates feature importance based on the average reduction in impurity contributed by each feature across all decision trees. The feature importance used in the Random Forest model is represented in Eqn (4).

$$\text{Feature Importance} = \frac{\text{Average Improvement in Split Criterion}}{\text{Average Impurity of Split}} \quad (4)$$

The Random Forest model was constructed using 200 decision trees, and the feature importance scores obtained from both the Relief algorithm and Random Forest were independently normalized to the range [0,1] before being combined to obtain a robust importance score for each variable. Variables with higher combined importance scores were considered more influential for thermal optimization and were retained for model development, whereas variables with negligible importance were discarded to reduce dimensionality, eliminate redundant information, and improve computational efficiency. Based on the feature ranking results, the final input variables retained for model construction included Boiler Temperature, Turbine Pressure, Main Steam Flow, Coal Flow Rate, Feedwater Temperature, Flue Gas Temperature, Oxygen Level (O₂), and Steam Pressure, as these variables provide comprehensive information regarding the thermal, mechanical, combustion, and process characteristics governing power plant operation. Variables exhibiting minimal predictive contribution, such as redundant identifiers and non-informative operational attributes, were excluded from further analysis. The selected features were subsequently used as inputs to the Reinforcement Learning and Artificial Neural Network models for thermal optimization and predictive maintenance.

3.3 Reinforcement Learning for Thermal Control Optimization

Reinforcement Learning (RL) is an AI technique where an agent learns to optimize thermal control in power plants

by interacting with its environment, aiming to maximize a cumulative reward, such as improved thermal efficiency, reduced fuel consumption, and stable plant operation. The state space includes variables like temperature, pressure, and energy output, while the action space defines the possible actions (e.g., adjusting boiler temperature or turbine pressure). The reward function assigns numerical values based on how well actions improve system performance. The training process involves plotting training curves and policy convergence plots to assess the agent's learning progress over time. A comparison with traditional methods like Fuzzy Logic Control (FLC) and Support Vector Machines (SVM) shows RL's superior ability to adapt to real-time changes and optimize the system. However, the reported metrics such as R^2 , MAE, RMSE, and MAPE, typically used in regression models, do not fully capture RL's performance, which is usually measured by cumulative reward and policy optimization. The results, which resemble supervised learning outcomes, highlight a need for further refinement to align the framework with RL's true objective: optimizing long-term thermal performance through adaptive decision-making. The state space S at time step t represents the current operating condition observed by the reinforcement learning agent and is defined using key thermal parameters of the power plant is represented in Eqn (5).

$$s_t = [T_b(t), P_t(t), E(t)] \quad (5)$$

where $T_b(t)$ represents the boiler temperature, $P_t(t)$ represents the turbine pressure, and $E(t)$ denotes the instantaneous energy output. Each state variable is normalized to the range $[0,1]$ based on the corresponding operating limits of the thermal system to ensure stable learning and efficient policy optimization. The action space A defines the set of possible control decisions that can be performed by the RL agent to regulate the thermal operating conditions. The action vector is expressed as Eqn (6).

$$a_t = [\Delta T_b, \Delta P_t] \quad (6)$$

where $\Delta T_b, \Delta P_t \in \{-1, 0, +1\}$ represent decrease, maintenance, or increase of boiler temperature and turbine pressure setpoints, respectively. These discrete control actions allow the agent to dynamically adjust thermal parameters according to the observed plant conditions.

The reward function r_t provides quantitative feedback to the RL agent by evaluating the effectiveness of the selected action in improving thermal performance. The reward is formulated as a weighted combination of efficiency improvement, fuel consumption reduction, and operational stability is expressed as Eqn (7).

$$r_t = w_1 \cdot \Delta \eta_t - w_2 \cdot \Delta F_t - w_3 \cdot |s_t - s_{t-1}| \quad (7)$$

where $\Delta \eta_t$ denotes the change in thermal efficiency, ΔF_t represents the change in fuel consumption, and $|s_t - s_{t-1}|$

represents the penalty associated with sudden state variations. The weighting factors w_1, w_2, w_3 determine the relative importance of efficiency, fuel economy, and stability during the optimization process. In this study, the weighting factors were empirically set to $w_1 = 0.5, w_2 = 0.3$, and $w_3 = 0.2$, reflecting the greater priority given to thermal efficiency gains relative to fuel savings and stability during training. The exploration mechanism controls how the RL agent balances discovering new control actions and exploiting previously learned optimal actions. The ϵ -greedy strategy is adopted, where the agent selects a random action with probability ϵ and chooses the action with the highest Q -value with probability $1 - \epsilon$. The exploration rate is gradually reduced during training using expressed as Eqn (8).

$$\epsilon_t = \epsilon_{\min} + (\epsilon_{\max} - \epsilon_{\min})e^{-\lambda t} \quad (8)$$

where λ represents the decay coefficient. This mechanism enables extensive exploration during the initial training phase and progressive exploitation as the optimal thermal control policy is learned. The exploration rate parameters were set to $\epsilon_{\max} = 1.0$ and $\epsilon_{\min} = 0.05$, with a decay coefficient $\lambda = 0.001$, allowing extensive exploration during early training episodes and progressive exploitation as the policy converged. The agent will be aiming to adopt a policy π that will bring the state to an action in a manner that will maximize the expected cumulative reward in the long run. The reward functionality plays a key role to achieve the learning process and can be established in accordance with such factors as energy efficiency, fuel consumption, and system stability. Q-learning algorithm is also popularly applied in RL: the agent learns a Q -value (function action-value) that approximates the future reward on average in the case of taking an action a in a particular state. The update law of the Q -value is provided by is represented in Eqn (9):

$$Q(s_t, a_t) = Q(s_t, a_t) + \alpha \left(r_t + \gamma \max_{a'} Q(s_{t+1}, a') - Q(s_t, a_t) \right) \quad (9)$$

- $Q(s_t, a_t)$ is the Q -value for the current state-action pair,
- α is the learning rate (how much new information overrides the old),
- r_t is the immediate reward received after taking action a_t ,
- γ is the discount factor (how much future rewards are valued),
- $\max_{a'} Q(s_{t+1}, a')$ is the estimated future reward based on the next state s_{t+1} .

Through repeated interactions, the agent refines its Q -values and learns to take the optimal actions that maximize the long-term reward, leading to improved control over the

power plant's thermal system. The policy update process determines how the RL agent improves its decision-making strategy based on the updated Q-values. The optimal policy is obtained as Eqn (10).

$$\pi(s_t) = \arg \max_a Q(s_t, a) \quad (10)$$

where the action with the maximum Q-value is selected for the given state. After each interaction with the thermal

environment, the Q-value is updated using Eq. (9), and the policy is continuously refined based on the latest action-value estimates. This iterative update process enables the RL agent to converge toward a policy that maximizes cumulative rewards while maintaining efficient and stable thermal operation.




Figure 2. Reinforcement Learning Architecture

The Figure 2 shows the Reinforcement Learning (RL) system utilized in the proposed system to optimize power plant thermal control. The system entails an Agent and an Environment. In this case, the environment provides the current state (s_t) to the agent, and the agent performs an action (a_t), which alters the environment. Once the action has been made, a reward (r_t) is given to the agent depending on how the action contributed to the bettering of the system (e.g., decreasing energy use or increasing thermal efficiency). The agent uses the feedback to revise its policy (π) that dictates the optimal action to be taken in each state. A controller neural network controls the behavior of the agent and constantly adapts the policy produced by it depending on the state and rewards observed.

3.4. Artificial Neural Networks for Thermal Control and Predictive Maintenance

ANNs are brain-like computational models that are based on the brain, and they are made up of a series of layers, each comprised of neurons that interact with each other to manipulate the input data to produce results. Applications

of ANNs in thermal control of power plants include the use of ANNs to forecast the behavior of systems (temperature, pressure and heat consumption) based on historical and real-time data. The network will generally have three layers: the input, where the raw data (the temperature, pressure, and operation data) will be taken; the hidden layers, where the real work is going to be done with the help of neurons with activation functions; and the output layer, where the final

The network is taught by adopting the backpropagation method whereby the disparity amid the forecasted value and the one truly required is computed (error) and the cell weights are modified to reduce it. Activation functions are typically signed-rank sigmoid, ReLU (Rectified Linear Unit) or tanh, which provides the network with non-linearity enabling it to fit more complicated relationships. Gradient descent is used to compute the process of updating the weights during training. In gradient descent, the weights are adjusted to the opposite direction of the gradient of the error function with respect to the weights with minimizing the error being the goal. The weight update rule is represented in Eqn (11):

$$w_{ij} = w_{ij} - \eta \frac{\partial E}{\partial w_{ij}} \quad (11)$$

Where:

- w_{ij} is the weight between neuron i in the previous layer and neuron j in the next layer,
- η is the learning rate,
- E is the error (e.g., Mean Squared Error) of the network.

The ANN is trained by repeated forward propagation (computing outputs) and backpropagation (recalculating weights), which teaches the ANN to model the behaviour of the underlying thermal system to give valuable predictions of thermal optimization and predictive maintenance.

ANNs can be applied in predictive maintenance to predict the probability of a system failure based on prior history data of maintenance and sensor data. The network is trained to identify patterns in operational information to be followed before the equipment malfunctions. As an example, sensor-recorded data on temperature, vibration, and pressure are inputted to the input layer of ANN. These inputs are processed by the hidden layers so as to identify the complex patterns of an upcoming failure, and the output layer is a probability of failure or a maintenance

recommendation. To analyse the performance of the model, the classification tasks are frequently applied in which the ANN is trained to classify the state of the system as that which is healthy or at risk depending on the input features. Binary cross entropy loss functional is usually applied to classification task, defined as Eqn (12):

$$L = -\frac{1}{N} \sum_{i=1}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] \quad (12)$$

Where:

- y_i is the actual label (1 for failure, 0 for healthy),
- $\hat{y}_i$ is the predicted probability of failure,
- N is the number of data samples.

By continuously training and refining the ANN model, it becomes increasingly capable of predicting failures before they occur, enabling the power plant to schedule maintenance proactively, reduce downtime, and improve overall system reliability.




Figure 3. Architecture of Artificial Neural Network (ANN)

This Figure 3 explains the fundamental layout of an Artificial Neural Network (ANN) as applied in the given structure of a proposed thermal control optimization in power plants. The network has three main layers, namely, input layer, hidden layers and the output layer. The input layer accepts the data e.g. thermal parameters (temperature, pressure, steam flow) and each node corresponds to a particular input feature. The weighted connections and activation functions of the hidden layer(s) compute the input data to represent complex and non-linear relationships.

3.4.1 Practical Failure Scenarios and Anomaly Detection Case Studies

Practical failure scenarios were analysed to evaluate the predictive maintenance capability of the proposed ANN framework using representative operational data from the Coal Fired Power Plant Thermal Performance Dataset. In the first scenario, a gradual increase in bearing vibration and lubricant temperature was identified as an early indication of bearing degradation before conventional threshold-based monitoring generated an alarm, enabling timely inspection and replacement. In the second scenario, abnormal fluctuations in feedwater pressure together with elevated flue gas temperature were recognised as anomalous operating conditions associated with the early stages of boiler tube fouling, allowing preventive maintenance before significant efficiency loss occurred. In the third scenario, a short-duration sensor spike caused by

transient measurement noise was correctly classified as a non-critical event, preventing unnecessary maintenance actions and reducing false-positive alerts. These representative scenarios demonstrate that the proposed ANN accurately distinguishes genuine equipment degradation from normal operational variations, thereby supporting proactive maintenance scheduling, improving equipment reliability, reducing downtime, and enhancing the overall operational stability and efficiency of thermal power plants.

3.5. Multiverse Optimizer for Thermal Control Optimization

MVO is a nature-inspired optimization algorithm, which mimics the behavior of the multiple parallel universes in search of the optimum solution. It is especially applicable in finding solutions to complex, high dimensional, and non-linear issues like optimization of the thermal control system in power plants. The main principle of MVO is founded on the notion of multiple universes that develop in parallel and the units of which are the possible solutions of the problem. These universes also explore the search space by updating their positions depending on their own history, as well as the best positions that have been found in the search space by other universes and thereby promoting global exploration and local exploitation of the search space.

The Multiverse Optimizer (MVO) represents each candidate solution as a universe, where optimization variables correspond to objects within that universe. Candidate universes are evaluated according to their fitness values (inflation rates), with high-quality universes acting as white holes that share promising solution components, while lower-quality universes behave as black holes that receive these components through roulette-wheel selection. This white hole–black hole interaction promotes exploration by enabling the algorithm to investigate diverse regions of the search space and reducing the likelihood of premature convergence. Simultaneously, the wormhole mechanism performs exploitation by moving candidate universes towards the current best solution through controlled local perturbations. The exploration–exploitation balance is adaptively regulated by the Wormhole Existence Probability (WEP) and the Travelling Distance Rate (TDR), allowing the optimizer to progressively refine promising solutions and achieve stable convergence towards globally optimal thermal control hyperparameters. The adaptive balance between these mechanisms is regulated by the Wormhole Existence Probability (WEP) and the Travelling Distance Rate (TDR), as defined in Eqns. (13) and (14):

$$WEP = WEP_{\min} + \frac{l(WEP_{\max} - WEP_{\min})}{L} \quad (13)$$

Where:

- $WEP_{\min}$ = minimum wormhole existence probability,
- $WEP_{\max}$ = maximum wormhole existence probability,
- l = current iteration,
- L = maximum number of iterations.

$$TDR = 1 - \left(\frac{l}{L}\right)^{1/p} \quad (14)$$

Where:

- TDR is the travelling distance rate,
- p is the exploitation accuracy parameter,
- l is the current iteration,
- L is the maximum number of iterations.

During the initial optimization stages, lower WEP values and larger TDR values encourage extensive exploration of the search space, thereby maintaining diversity among candidate solutions. As the optimization progresses, WEP gradually increases while TDR decreases, shifting the search process towards exploitation by concentrating the search around high-quality universes. This adaptive transition improves convergence speed, avoids premature convergence to local optima, and enables the identification of robust hyperparameter configurations for the proposed thermal control framework.

The Multiverse Optimizer (MVO) formulates hyperparameter tuning as a constrained optimization problem, where each universe represents a candidate solution corresponding to a complete hyperparameter configuration of the Reinforcement Learning (RL) and Artificial Neural Network (ANN) models. The decision vector is defined as Eqn (15):

$$\mathbf{X} = [\alpha, \gamma, N_h, N_n, B, E] \quad (15)$$

Where:

- α is the learning rate
- γ is the discount factor,
- N_h denotes the number of hidden layers
- N_n represents the number of neurons per hidden layer
- B is the batch size

- E is the number of training epochs.

The optimization objective is to determine the parameter vector that minimizes the validation error, expressed as Eqn (16):

$$\min J(\mathbf{X}) = RMSE_{\text{validation}} \quad (16)$$

where, $J(\mathbf{X})$ denotes the fitness function evaluated using the validation Root Mean Squared Error (RMSE). Minimizing this objective improves prediction accuracy and overall thermal control performance. The optimization is performed under the constraint is defined as Eqn (17):

$$X_i^{\min} \leq X_i \leq X_i^{\max}, i = 1, \dots, d \quad (17)$$

where $X_i^{\min}$ and $X_i^{\max}$ represent the lower and upper bounds of each optimized parameter. Candidate solutions exceeding these limits are projected back into the feasible search space before fitness evaluation.

The search boundaries are defined according to the predefined ranges of the learning rate, discount factor, hidden layers, neurons, batch size, and training epochs listed in Table 2, ensuring an effective balance between exploration and exploitation during optimization.

Each universe is encoded as a real-valued hyperparameter vector. Continuous parameters are represented directly, whereas discrete parameters are rounded to the nearest valid integer before model evaluation. The optimization process terminates when either the maximum number of iterations is reached or when the improvement in the best validation RMSE becomes negligible over successive iterations. The best-performing universe obtained at convergence is selected as the optimal hyperparameter configuration for training the final RL and ANN models.

3.5.1 Multiverse Optimizer in Thermal Control

In thermal control of power plants, the Multiverse Optimizer is applied to the parameters of AI models, including Reinforcement Learning and Artificial Neural Networks (ANNs) to maximize their performance in real-time optimization to mini tasks. The exploration-exploitation approach of the MVO enables the system to effectively maximize search on the most suitable set of thermal parameters, including boiler temperature, turbine pressure, and steam flow, and to balance the energy use and system stability thermal optimization in power plants is shown in Algorithm 1.

Algorithm 1: Multiverse Optimizer

START

Input:

- N : Number of universes (candidate solutions) in the population
-

- Mixite: Maximum number of iterations (maximum number of steps in the optimization process)
- $c1, c2$: Learning factors for cognitive and social components (used in position updates)
- $r1, r2$: Random numbers (uniform distribution between 0 and 1) for updating the positions
- Convergence Criteria: Threshold criteria to determine when optimization stops

Initialize Population

For $i = 1$ to N :

- X_i = Random solution (generate a random solution for each universe)
- Fitness (X_i) = Evaluate fitness (evaluate the fitness of the solution based on the objective function)
- $P_best = X_i$ (set the personal best solution for universe i)

•

End For

Set Global Best

- G_best = Best solution from all universes
- G_best_value = Evaluate fitness of G_best

Main Loop (iterations)

- iteration = 1

While iteration $\leq$ Mixite:

For $i = 1$ to N :

- $r1$ = Random (0,1)
- $r2$ = Random (0,1)

Update Position (based on interaction with multiple universes, using elements from the MVO framework):

- Black hole: Strongest solutions exert gravitational influence, pulling universes toward them.
- White hole: Universes escaping from a strong gravity well, enabling diverse exploration of the search space.
- Wormhole: Allows solutions to instantly transition across the search space.
- Inflation rate: Increases the search space for universes, enhancing exploration at certain steps.
- $X_i_new = X_i + c1 * r1 * (P_best - X_i) + c2 * r2 * (G_best - X_i) + \text{Black hole influence} + \text{White hole escape} + \text{Wormhole jump} + \text{Inflation rate}$

Evaluate fitness of new solution:

- Fitness (X_i_new) = Evaluate fitness of the new solution

Update personal best (P_best):

- If Fitness (X_i_new) < Fitness (P_best), then $P_best = X_i_new$ (update the personal best)

Update global best (G_best):

- If Fitness (X_i_new) < Fitness (G_best), then $G_best = X_i_new$ (update global best)
 - $G_best_value = \text{Fitness}(G_best)$
-

```

End For
Check for convergence (if G _ best does not improve
after several iterations):
If ConvergenceCriteriaMet (G _ best _ value), then
    o Break (stop the optimization process)
iteration = iteration + 1 (increment the iteration count)
End While
Output:
    • G _ best (optimal solution for thermal control
    parameters)
    • G _ best _ value (fitness value of the optimal
    solution)
END
    
```

3.6 Hyperparameter Search Space and Tuning Strategy

The Multiverse Optimizer (MVO) was employed to optimize the key hyperparameters of both the Reinforcement Learning (RL) agent and the Artificial Neural Network (ANN) model over the predefined search space summarized in Table 1. The MVO population consisted of 30 universes evolved over a maximum of 100 iterations. During the optimization process, the Wormhole Existence Probability (WEP) was increased linearly, while the Travelling Distance Rate (TDR) was decreased linearly to achieve an effective balance between global exploration and local exploitation. Each candidate hyperparameter configuration was evaluated using a held-out validation dataset, where the validation RMSE served as the primary fitness function to be minimized. To improve the robustness of model selection, R^2 , MAE, and MAPE were additionally monitored as secondary evaluation metrics. The final hyperparameter configuration was selected as the solution that consistently produced the lowest validation RMSE while simultaneously maintaining high prediction accuracy and stable generalization performance, and this optimized configuration was subsequently used for the experimental evaluation.

Table 1. Hyperparameter Search Space

Model	Hyperparameter	Search Range	Selected Value
RL Agent	Learning rate (α)	0.0001–0.01	0.001
RL Agent	Discount factor (γ)	0.80–0.99	0.95
RL Agent	Reward decay steps	100–1000	500
ANN	Number of hidden layers	1–4	2

ANN	Neurons per hidden layer	16–128	64
ANN	Batch size	16–128	32
ANN	Training epochs	50–500	200
MVO	Number of universes (N)	10–50	30
MVO	Maximum iterations	50–200	100

The Artificial Neural Network (ANN) model was implemented using the optimized hyperparameter configuration obtained through the Multiverse Optimizer (MVO). The final model architecture and training parameters were selected based on the minimum validation RMSE achieved during the optimization process. Table 2 summarizes the complete ANN hyperparameter configuration, including the network architecture, learning rate, optimizer, batch size, training epochs, activation functions, loss function, regularization strategy, and other optimization parameters used for model development.

Table 2. ANN Hyperparameter Configuration Used in the Proposed Framework

Hyperparameter	Value
Network Architecture	256-128-64-32-16-8-1
Activation Function	LeakyReLU ($\alpha=0.1$)
Optimizer	Adam
Learning Rate	0.0005
Batch Size	8
Epochs	50
Loss Function	MSE
Regularization	$L1=1e-5$, $L2=1e-4$
Dropout	0.30, 0.25, 0.20, 0.15
Batch Normalization	Yes
Early Stopping	Patience = 200
Data Split	85/7.5/7.5

Input Features	29
Output Variable	Boiler Efficiency (%)

4. Result and Discussion

The suggested initiative to streamline the thermal control of power plants with the help of AI methods was implemented with Python success. The framework adopts the use of Reinforcement Learning (RL) to optimize dynamically the thermal performance, Artificial Neural Networks (ANNs) to predictive maintenance, and Multiverse Optimizer (MVO) to hyperparameter optimization. The Coal Fired Power Plant Thermal Performance Dataset was used to test the system and obtain the required thermal, operational and environmental data. The AI models were evaluated on the basis of energy efficiency and system stability as well as predictiveness of its maintenance requirements. The findings showed a major enhancement regarding thermal and operational efficiency, which proved the usefulness of the offered strategy.

4.1 Dataset Evaluation

The Coal Fired Power Plant Thermal Performance Dataset involved in the proposed framework is a collection of different working parameters like temperature, pressure, and the flow rates of the steam and energy production that are vital in the optimization of the thermal control system. Environmental data such as ambient temperature and humidity are also part of the dataset, which is necessary to predict behavior of thermal under various weather conditions.

The dataset was found highly relevant when tested to train the AI models as they included a great variety of real-life operating conditions and failure modes. The historical data of the dataset allowed the framework to forecast failures of the system and optimize thermal performance in an effective way. The dataset was pre-processed (removal of noise, features selection, and missing values imputation) to be converted to a clean and structured format that is trainable by AI.

4.2. Performance Metrics

The performance of the proposed framework for thermal control optimization in power plants is evaluated using the following key metrics:

Table 3. Performance Metrics of the Proposed Framework

R ² Score:	0.952830
MAE:	0.021572
RMSE:	0.023702
MAPE:	0.0230%
Max Error:	0.035469
Median Error:	0.022145

In the Table 3, the main performance indicators on how the AI-based thermal control framework can be used to assess the performance of power plants will be summarized. The model has a high fit as indicated by the R² score of 0.9528 which shows that the model explains 95.28 percent of the variation in thermal parameters. The Mean Absolute Error (MAE) of 0.0216 and Root mean squared error (RMSE) of 0.0237 indicate that the errors of prediction are low whereas the mean absolute percentage error (MAPE) of only 0.023 is very high. The highest error of 0.0355 is the difference between the predictions and the actual values; the median error of 0.0221 is used to show the central tendency of the rate error.

Mean Absolute Error: MAE assesses the mean difference of the predicted value and the actual value in an absolute manner. In the suggested framework, it is applied to determine whether the predicted thermal parameters (e.g., temperature and pressure) are accurate or not compared with real measurements. Smaller MAE values mean the models perform well is represented in Eqn (18):

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (18)$$

Where:

- y_i is the actual value,
- $\hat{y}_i$ is the predicted value,
- n is the number of data points.




Figure 4. Predicted vs. Actual Values with MAE Score

This Figure 4 shows the plot also visualizes the relationship between the projected values and the actual values of thermal parameters in the power plant optimization system. The green dots indicate actual values per data point by the x-axis, and the predicted values per data point on the y-axis. The blue regression line is the most appropriate regression line that approximates the relationship between the predicted and actual values. The Mean Absolute error (MAE) appears, as 0.0218, which shows that the predictions of the model are very close to the actual values and the average deviation is very small of only 0.0218. Such a low MAE value indicates that the model of thermal control based on AI is very accurate and able to predict the parameters of the functioning of the plant.

Root Mean Squared Error: RMSE is computed as the square root of the mean squared error between the predicted and actual values, and it gives an indication of the performance of the model. RMSE is used to determine the accuracy of prediction of parameters of the thermal system in the framework and a lower value of RMSE is associated with higher prediction and control by a model is represented in Eqn (19):

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (19)$$




Figure 5. Predicted vs. Actual Values with RMSE Score

The Figure 5 shows scatter plot shows the predicted and the actual values of thermal parameters in power plant system. The dark blue dots indicate the single data points, and the x-axis indicates the actual values whereas the y-axis indicates the predicted values. The orange regression line is the best-fit line which means that there is a linear relationship between the predicted and actual values. The value of the Root Mean Squared Error (RMSE) is 0.023 which is the mean value of the error between actual and predicted values.

Mean Absolute Percentage Error: MAPE puts the prediction error as a percentage, and it is easier to interpret it in the sense of the extent to which model predictions are not observed values. Within the framework of the proposed framework, MAPE is a measure of the relative error of the predictors of the thermal behaviour of the power plant, with the smaller values indicating a higher degree of model accuracy is represented in Eqn (20):

$$MAPE = \frac{1}{n} \sum_{i=1}^n \left| \frac{y_i - \hat{y}_i}{y_i} \right| \times 100 \quad (20)$$




Figure 6. Predicted vs. Actual Values with MAPE Score

This Figure 6 shows the Predicted and Actual values with Mean Absolute Percentage Error (MAPE) 0.023. The purple dots are the single values, which indicate the forecasted values on the y-axis and the real values on the x-axis. The purple regression line is the best fit which shows that there is a linear relationship between the actual and the predicted values. The MAPE value of 0.023 denotes that the model predictions are very accurate with an average of 2.3 percentage difference between the predicted and the actual values.

Coefficient of Determination (R^2): R^2 is a measure of the percentage variation of the target variable explained in the model. When the value of R^2 is stronger, it indicates that the model is able to explain more of the variation in thermal parameters, and thus, it is a key measure of the performance of the thermal control optimization models is represented in Eqn (21):

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2} \quad (21)$$

Where:

- y_i is the actual value,
- $\hat{y}_i$ is the predicted value,
- $\bar{y}$ is the mean of the actual values.




Figure 7. Predicted vs. Actual Values with R^2 Score

This Figure 7 gives a scatter diagram of the outcome values against the actual values, and the regression line fits the outcome values. The purple dots are the data points, which represent the prediction of a business and the actual value. The regression line is the black line, and it is used to model the relationship between the predicted and actual values. The high R^2 of 0.98 implies that there was a very strong correlation with the model showing that the variance in the data is 98 percent.

4.2.1 Reinforcement learning out comes

The Figure 8 is the culmination of the policy loss in the course of training a reinforcement learning (RL) agent. The greater the number of episodes, the lesser the policy loss is; this implies that the agent is learning and polishing up on its decision-making strategy just as time goes by.




Figure 8. Policy Convergence Plot

The sharp curve drop at the start illustrates the agent acquisition, and the flattening of the curve at the end depicts that the policy of the agent is close to optimal. When the policy loss value is low, it is possible to say that

the RL agent has reached a useful strategy to maximize the thermal control parameters in the power plant.




Figure 9. RL Training Curve: Reward Evolution

This Figure 9 shows the development of the reward in the training of a reinforcement learning (RL) agent. The Reward per Episode (it is in red) is of the reward that the agent gains instantly after each episode. The reward trend is also smoothed out to show the moving avg reward (purple) which is useful in visualising the learning trend across all the episodes to remove the noise. A slow rise in the immediate and moving average rewards shows that the RL agent is becoming efficient at its decisions with time, as it results in more effective thermal control of the power plant. The smoother moving average curve is an indication that the policy of the agent is becoming more stable as the agent trains and is moving upwards.




Figure 10. Boiler Efficiency Comparison (With vs Without RL)

In this Figure 10 shows the efficiency of boilers in two cases will be compared: With RL (Reinforcement Learning) and Without RL. The boiler efficiency in percentage is presented in the y-axis whereas the sample episodes are presented in the x-axis. Without RL efficiency is illustrated by purple bars and the With RL efficiency by dark blue bars. The findings indicate that efficiency of

boilers is improved as a result of RL being incorporated during the sample episodes. It shows that the current trend is towards the optimization of the thermal control provided by RL, which ensures a more stable and efficient functioning of the boiler until the end of the work of the power plant.

Table 4 summarizes the Reinforcement Learning performance metrics of the proposed framework. The RL agent converged after 150 training episodes with a 99.00% reduction in policy loss, while achieving a final reward of 0.078 and a Learning Stability Index of 0.95. These results demonstrate efficient policy convergence, stable learning behaviour, and effective optimization of the thermal control process.

Table 4. Reinforcement Learning Performance Metrics of the Proposed Framework

RL Metric	Result
Training Episodes	200
Convergence Episode	150
Initial Policy Loss	1.00
Final Policy Loss	0.01
Loss Reduction	99.00%
Average Reward	0.029
Maximum Reward	0.082
Final Reward	0.078
Learning Stability Index	0.95
Convergence Status	Stable

4.3 Model Evaluation

The Figure 11 shows the residual distribution of the predictions of the AI model. The frequency distribution of the residuals (blue bars) is represented by the histogram where the residuals are the difference between the predicted and the actual values.




Figure 11. Residual Distribution of the AI Model

The red curve is the kernel density estimate of the residuals, which gives a smooth estimation of the distribution of the residuals. The four vertical lines denote the mean of the residuals, and this is the zero point. A distribution that is well centered around zero implies that the model predictions are unbiased and the nearer the residuals to normal distribution the higher the goodness with which the model fits the data.

4.4 Optimization Outcomes of AI-Driven Thermal Control in Power Plants

This title encapsulates the key theme of the AI optimization process applied to various aspects of thermal control in power plants, such as energy savings, efficiency improvements, and hyperparameter optimization.




Figure 12. Annual Savings and Boiler Efficiency Improvement from AI Optimization

This Figure 12 shows how AI optimization affects the performance of a power plant. The Annual Savings from AI Optimization graph on the left demonstrates a high level of savings in the consumption of coal (tons), and the savings on the annual cost (USD), which indicates the enormous financial and environmental value of the AI-

based thermal control system. The graph shows how AI-based optimization may result in yearly coal savings and considerable cost cuts in the activities of the plants.




Figure 13. Coal Flow and Coal Consumption Reduction After AI Optimization

This Figure 13 shows how AI optimization will impact the circulation and use of coal in the power plant. Left, Coal Flow (t/h) bar chart is used to compare coal flow before and after AI optimization. The findings indicate that there was the decrease of the flow of coal to 78.9 tons per hour (after AI) as compared to 82.2 tons per hour (before AI) which implies that coal was utilized more efficiently. On the right, Coal Consumption Reduction bar chart demonstrates the total decrease in coal consumption following the introduction of AI, and it further supports the efficiency gains.




Figure 14. Feedwater and Flue Gas Temperature Improvements After AI Optimization

This Figure 14 shows how the temperature is better managed following the use of AI optimization in the thermal plant. On the left, the Feedwater Temperature (°C) bar chart differs the temperature of the feedwater before and after AI optimization where the initial temperature of the feedwater (275.3 C before AI) was slightly higher than

after AI (273.6 C after AI). Through this enhancement, there is enhanced control of temperature that leads to more effective transfer of heat and even enhanced thermal efficiency.




Figure 15. Main Steam Flow and Oxygen Level Adjustments After AI Optimization

This Figure 15 illustrates how AI optimization impacts the key operation parameters including the main steam flow and the level of oxygen content in the power plant. The bar chart on the left, Main Steam Flow (t/h), demonstrates a minor drop in the steam flow of 588.5 tons per hour (before AI) to 582.1 tons per hour (after AI). This decrease indicates that AI optimization is assisting in terms of energy consumption by reducing the unnecessary production of steam and sustaining the proper functioning of the plants.




Figure 16. Hyperparameter Optimization Convergence Plot

This Figure 16 shows the intersection of hyperparameter optimization within the framework proposed. The validation loss (y-axis) versus a number of optimization trials (x-axis) in the plot. The loss of validation also goes

down as the trials continue implying that the model is becoming better with each iteration. The most attractive trial that has the lowest validation loss of 0.0077 has a star symbol that indicates the best hyperparameter configuration. The dashed horizontal line will show the target or baseline loss value at which, the value will be compared. This plot illustrates that Multiverse Optimizer (MVO) is effective in optimizing the hyperparameters to give a better model performance and lower loss in the context of thermal optimization and predictive maintenance exercises.

4.5 Parameter Sensitivity Analysis

Table 5 presents the influence of different Multiverse Optimizer population sizes on the prediction performance while maintaining a fixed iteration count of 100. Increasing the population size from 10 to 30 reduced the RMSE from 0.0318 to 0.0237, indicating a considerable improvement in optimization performance. Only marginal reductions in RMSE were observed for population sizes of 40 and 50, with values of 0.0235 and 0.0234, respectively. These results indicate that a population size of 30 provides an effective trade-off between prediction accuracy and computational cost, while larger population sizes offer only limited additional improvement.

Table 5. Parameter Sensitivity Analysis of the Multiverse Optimizer

Population Size	Iterations	RMSE
10	100	0.0318
20	100	0.0264
30	100	0.0237
40	100	0.0235
50	100	0.0234

4.6 Correlation Analysis

A Pearson correlation analysis was performed, and the resulting correlation matrix is presented in Figure 17. The analysis provides valuable insights into the degree of linear association between process parameters, enabling the identification of influential variables and potential redundancy before model evaluation.




Figure 17. Pearson correlation matrix of the selected operational variables used for thermal control optimization

Figure 17 presents the Pearson correlation matrix of the selected operational variables. Strong positive correlations were observed between Coal Flow Rate and Energy Input ($r = 0.973$), Main Steam Flow and Coal Flow Rate ($r = 0.760$), and Boiler Temperature and Feedwater Temperature ($r = 0.803$). A strong negative correlation was identified between Boiler Efficiency and Flue Gas Temperature ($r = -0.833$), indicating increased heat losses at higher flue gas temperatures. Overall, the correlation analysis confirms that the selected variables effectively capture the relationships governing thermal performance and support the proposed AI-based thermal control framework.

4.7 Performance Comparison Table

The Table 6 and Figure 18 of this paper compares the performance of three methods of thermal control optimization in power plants, namely the Proposed Framework of ANN and Reinforcement Learning (RL), Fuzzy Logic Control (FLC) of Alvarez et al. (2025), and Support Vector Machine (SVM) of Elshehabi and Alfehaid (2025).

Table 6. Performance Comparison of the Proposed Framework with Existing Methods

Method	R ² Score	MAE	RMSE	MAPE
ANN and Reinforcement learning	0.9528	0.0216	0.0237	0.023%
[37] (FLC)	0.8754	0.0412	0.0521	0.050%
[38] (SVM)	0.9123	0.0328	0.0453	0.045%




Figure 18. Performance Comparison of AI Models for Thermal Control Optimization

The R² score, Mean Absolute Error (MAE), Root Mean Squared error (RMSE) as well as the Mean Absolute Percentage error (MAPE) of each method. The Proposed Framework is better than both FLC and SVM since it is more accurate in predictions according to R² and lower error (MAE, RMSE and MAPE). It has a lower R² value and greater error values with the FLC method with a moderate performance, still falling behind the suggested method in terms of both accuracy and efficiency. The above comparison highlights how AI-based thermal optimization is far more effective in enhancing the work of power plants.

4.8 Computational Complexity and Implementation Analysis

Table 7 shows that the proposed framework was implemented on an Intel i7 processor with 16 GB RAM, requiring a total training time of 18.4 minutes and a peak memory usage of 2.1 GB. Among the framework components, the RL module required the highest training time (11.2 minutes), followed by ANN training (4.8 minutes), while MVO completed hyperparameter optimization in 2.4 minutes. Although the optimization stage increases the offline computational cost, it significantly improves model accuracy and convergence. The moderate memory requirement demonstrates that the framework can be executed on standard computing platforms without specialized hardware.

Table 7. Computational Complexity and Implementation Analysis

Metric	Value
CPU	Intel i7
RAM	16 GB
ANN Training Time	4.8 min
RL Training Time	11.2 min
MVO Optimization Time	2.4 min
Total Training Time	18.4 min
Peak Memory Usage	2.1 GB

4.9 Ablation Study of the Proposed Framework

Table 8 presents the ablation study conducted to evaluate the individual contributions of the Artificial Neural Network (ANN), Reinforcement Learning (RL), and Multiverse Optimizer (MVO) within the proposed thermal control framework. Each model configuration was assessed using the same experimental settings to ensure a fair comparison. The results indicate that incorporating RL and MVO progressively improves prediction accuracy and reduces the RMSE compared with the ANN-only model. The complete ANN + RL + MVO framework achieves the highest R^2 value of 0.9528 and the lowest RMSE of 0.0237, demonstrating that the integration of all three components provides the most effective and robust thermal control performance.

Table 8. Ablation Study

Configuration	R^2	RMSE
ANN Only	0.9012	0.0458
ANN + RL	0.9347	0.0312
ANN + MVO	0.9418	0.0286
RL + MVO	0.9265	0.0341
ANN + RL + MVO (Proposed)	0.9528	0.0237

4.10 Discussion

The proposed thermal control framework, which is powered by AI, gives a significant boost to the operation of power plants particularly in comparison with the traditional control systems. The system will be able to adapt to the changing operational conditions in real-time dynamically by incorporating the reinforcement learning (RL) and Artificial Neural Networks (ANNs) with the Multiverse Optimizer (MVO) that will help to maintain the energy efficiency, minimize fuel consumption, and provide stable operations. The integration of RL enables the system to keep on learning and optimising thermal values, including boiler temperature and turbine pressure, on the basis of real-time feedback of the environment of the plant. Such adaptation leads to an efficient power plant, since the RL agent can determine the best actions to reduce the wastage of energy and the performance of the system is optimized. The ANNs also enhance the framework in that it allows predictive maintenance that will ensure no-fly-by-wire failures of equipment and thereby minimize the downtimes and increase the lifespan of the plant. Model

hyperparameters are fine-tuned by the MVO to ensure that the model costs are more accurate, and the overall performance of the AI models is better.

The AI-based framework is clearly superior to the conventional approaches, including the Fuzzy Logic Control (FLC) and Support Vector Machines (SVM). As an illustration, the FLC approach attained an R^2 of 0.8754 which was accompanied by an MAE of 0.0412 whilst SVM approach attained a R^2 of 0.9123 at a rate of 0.0328. The AI-based framework, in its turn, obtained an amazing R^2 value of 0.9528, MAE of 0.0216, RMSE of 0.0237, and MAPE of 0.023 percent. These findings suggest that the proposed AI models would be better able to forecast thermal parameters with high precision and respond to alterations in real-time. The RL integration also highlights how the framework is compatible with the maximization of long-term performance when the traditional methodology is usually incapable of adapting effectively to the dynamic environment. This underscores how much better thermal control and energy efficiency the AI-based approach would be.

The proposed AI-driven thermal optimization framework, comprising Reinforcement Learning (RL)-based dynamic control, Artificial Neural Network (ANN)-based predictive maintenance, and Multi-Verse Optimization (MVO)-based hyperparameter tuning, provides a flexible architecture that can be adapted to a wide range of power generation technologies through technology-specific retraining and parameter optimization. Gas turbine and combined-cycle power plants, which operate under faster load-ramping dynamics and stricter turbine temperature constraints, have demonstrated promising performance using AI-based and RL-driven control strategies [39], [40]. Nuclear power plants require more stringent safety-constrained reward functions and verification procedures because of their rigorous instrumentation and control requirements [41], [42]. Biomass and waste-to-energy plants, characterised by greater fuel quality variability, can particularly benefit from the ANN-based predictive maintenance component for combustion performance and boiler fouling prediction [43], [44]. Similarly, concentrated solar power and geothermal systems can extend the reinforcement learning state space by incorporating solar irradiance or geothermal reservoir forecasts, enabling adaptive thermal optimization under varying environmental conditions [45], [46]. These observations indicate that the proposed framework has the potential to be adapted for intelligent thermal optimization across diverse conventional and renewable power generation technologies, while comprehensive cross-technology experimental validation remains an important direction for future research.

AI-driven thermal control improves the long-term reliability of thermal power plants by continuously monitoring operational parameters and detecting performance degradation at an early stage. The proposed framework enables proactive maintenance planning, reducing unplanned downtime and maintenance costs while minimising thermal stress on critical components.

These capabilities extend equipment lifespan and support strategic asset management through informed maintenance scheduling, optimal resource allocation, and improved operational efficiency for sustainable power plant operation. The deployment of AI-enabled thermal control systems should incorporate robust cybersecurity measures to ensure secure and reliable power plant operation. Encrypted data transmission, authenticated access, and secure storage protect operational data from unauthorised access and tampering. System resilience can be strengthened through continuous monitoring, anomaly detection, and intrusion detection mechanisms to identify potential cyber threats. Furthermore, regular software updates, access control policies, and backup strategies enhance threat mitigation and ensure reliable AI-driven thermal control under evolving cybersecurity risks.

5. Conclusion and Future Works

The thermal control system based on AI, which is suggested to be implemented in power plants, shows remarkable outcomes in terms of thermal performance and predictive maintenance. Combined with Reinforcement Learning (RL), Artificial Neural Networks (ANNs), and Multiverse Optimizer (MVO), the framework guarantees a dynamic solution that is real-time and optimized to promote the energy efficiency and stability of a system. The framework had good performance statistics, such as $R^2 = 0.9528$, $MAE = 0.0216$, $RMSE = 0.0237$, and $MAPE = 0.023\%$, meaning that it was very accurate in forecasting thermal parameters. The RL part of the system, making the decisions dynamically and on the basis of the real-time feedback is one of the main contributors of these results, providing the system with the capabilities to optimize the performance in the long-term perspective by continuous learning. Comparatively, however, more conventional approaches such as Fuzzy Logic Control (FLC) and Support Vector Machines (SVM) performed worse in terms of accuracy with greater error rates, which also supports the excellence of the AI-based methodology.

Subsequent efforts will be made in the future to enhance the flexibility of the system to deal with any unexpected operational outage and quick fluctuations in real-time situations that are prevalent in large scale power plants. The framework will also be made to handle larger plants and more complex AI models will be introduced to handle the increasing complexity of thermal optimization issues. Thermal control system will be even more responsive and efficient, as the use of IoT sensors and edge computing technologies will increase the ability of the framework to make timely decisions in real-time. Such developments of the future will make the AI-based thermal control system a helpful resource in streamlining the work of the power plants and making the energy management practices even more sustainable.

Future Works

Improvement of Model Adaptability: More studies will be undertaken to enhance the adaptability of the AI models to swift alterations in the working environment, unexpected disruptions in the real-time running of the plants.

Scalability and Deployment: Future efforts will be put on the scaling of the framework to larger power plants and the addition of several AI models to address more complicated and multi-dimensional optimization problems.

IoT and Edge Computing: By adding IoT sensors to collect real-time data and use edge computing to make decisions quicker and deploy models, we will make the thermal control system more responsive and faster.

Declarations

Data Availability

The data used in this study are derived from the Coal Fired Power Plant Thermal Performance Dataset. The dataset is publicly available from the corresponding data repository. Additional data generated or analyzed during this study are available from the corresponding author upon reasonable request.

Conflict of Interest

The authors declare that they have no known competing financial interests or personal relationships that could have influenced the work reported in this paper.

Funding Statement

This research received no external funding.

Author Contributions

Qiang Liu: Conceptualization, Methodology, Investigation, Writing – Original Draft Preparation.

Rui Zhu: Data Curation, Software, Validation.

Guofeng Li: Formal Analysis, Resources.

Lei Zhang: Investigation, Data Interpretation.

Jun Li: Visualization, Validation.

Lin Wang: Supervision, Project Administration, Writing – Review & Editing.

Yanjuan Guo: Methodology, Software Development, Writing – Review & Editing.

All authors have read and approved the published version of the manuscript.

Ethical Approval

This study does not involve human participants, animals, or sensitive personal data. Therefore, ethical approval was not required.

Consent to Participate

Not applicable.

Consent for Publication

All authors have reviewed and approved the manuscript for publication.

Competing Interests

The authors declare no competing interests.

Acknowledgement

The authors would like to thank Shandong Energy Group Lingtai Thermal Power Generation Co., Ltd. and Xi'an Thermal Power Research Institute Co., Ltd. for their technical support and valuable insights during this research. The authors also acknowledge the providers of the Coal Fired Power Plant Thermal Performance Dataset used in this study.

References

- [1] Liu J, Hu H, Yu SS, Trinh H. Virtual power plant with renewable energy sources and energy storage systems for sustainable power grid—formation, control techniques and demand response. *Energies*. 2023;16(9):3705. doi:10.3390/en16093705.
- [2] Jin W, Wang P, Yuan J. Key role and optimization dispatch research of technical virtual power plants in the new energy era. *Energies*. 2024;17(22):5796. doi:10.3390/en17225796.
- [3] Ghobadpour A, Monsalve G, Cardenas A, Mousazadeh H. Off-road electric vehicles and autonomous robots in agricultural sector: trends, challenges, and opportunities. *Vehicles*. 2022;4(3):843–864. doi:10.3390/vehicles4030047.
- [4] Mazhar T, et al. The role of ML, AI and 5G technology in smart energy and smart building management. *Electronics*. 2022;11(23):3960. doi:10.3390/electronics11233960.
- [5] Alanazi F. Electric vehicles: benefits, challenges, and potential solutions for widespread adaptation. *Appl Sci*. 2023;13(10):6016. doi:10.3390/app13106016.
- [6] Tsai WC, Hong CM, Tu CS, Lin WM, Chen CH. A review of modern wind power generation forecasting technologies. *Sustainability*. 2023;15(14):10757. doi:10.3390/su151410757.
- [7] Benti NE, Chaka MD, Semie AG. Forecasting renewable energy generation with machine learning and deep learning: current advances and future prospects. *Sustainability*. 2023;15(9):7087. doi:10.3390/su15097087.
- [8] Flynn CD, Chang D. Artificial intelligence in point-of-care biosensing: challenges and opportunities. *Diagnostics*. 2024;14(11):1100. doi:10.3390/diagnostics14111100.
- [9] Abualigah L, et al. Wind, solar, and photovoltaic renewable energy systems with and without energy storage optimization: a survey of advanced machine learning and deep learning techniques. *Energies*. 2022;15(2):578. doi:10.3390/en15020578.
- [10] Hasan MK, Habib AA, Islam S, Balfaqih M, Alfawaz KM, Singh D. Smart grid communication networks for electric vehicles empowering distributed energy generation: constraints, challenges, and recommendations. *Energies*. 2023;16(3):1140. doi:10.3390/en16031140.
- [11] Fayyazi M, et al. Artificial intelligence/machine learning in energy management systems, control, and optimization of hydrogen fuel cell vehicles. *Sustainability*. 2023;15(6):5249. doi:10.3390/su15065249.
- [12] Vamvakas D, Michailidis P, Korkas C, Kosmatopoulos E. Review and evaluation of reinforcement learning frameworks on smart grid applications. *Energies*. 2023;16(14):5326. doi:10.3390/en16145326.
- [13] Wójcicki K, Biegańska M, Paliwoda B, Górna J. Internet of things in industry: research profiling, application, challenges and opportunities—a review. *Energies*. 2022;15(5):1806. doi:10.3390/en15051806.
- [14] Cembrowska-Lech D, et al. An integrated multi-omics and artificial intelligence framework for advance plant phenotyping in horticulture. *Biology*. 2023;12(10):1298. doi:10.3390/biology12101298.
- [15] Albahar M. A survey on deep learning and its impact on agriculture: challenges and opportunities. *Agriculture*. 2023;13(3):540. doi:10.3390/agriculture13030540.
- [16] Yao Z, Romero C, Baltrusaitis J. Combustion optimization of a coal-fired power plant boiler using artificial intelligence neural networks. *Fuel*. 2023;344:128145. doi:10.1016/j.fuel.2023.128145.
- [17] Nemitallah MA, et al. Artificial intelligence for control and optimization of boilers' performance and emissions: a review. *J Clean Prod*. 2023;417:138109. doi:10.1016/j.jclepro.2023.138109.
- [18] Bisset C, Venter PVZ, Coetzer R. A systematic literature review on machine learning applications at coal-fired thermal power plants for improved energy efficiency. *Int J Sustain Energy*. 2023;42(1):845–872. doi:10.1080/14786451.2023.2244618.
- [19] Benson S, Eronu E. Integration of renewable energy with thermal-based power systems: a review of grid reliability, optimization, and storage. *Am J Electr Power Energy Syst*. 2025;14(2):28–44. doi:10.11648/j.epes.20251402.12.
- [20] Alhamrouni I, et al. A comprehensive review on the role of artificial intelligence in power system stability, control, and protection: insights and future directions. *Appl Sci*. 2024;14(14):6214. doi:10.3390/app14146214.
- [21] Ukoba K, Olatunji KO, Adeoye E, Jen TC, Madyira DM. Optimizing renewable energy systems through artificial intelligence: review and future prospects. *Energy Environ*. 2024;35(7):3833–3879. doi:10.1177/0958305X241256293.
- [22] Franki V, Majnarić D, Višković A. A comprehensive review of artificial intelligence (AI) companies in the power sector. *Energies*. 2023;16(3):1077. doi:10.3390/en16031077.
- [23] Devasenan M, Madhavan S. Thermal intelligence: exploring AI's role in optimizing thermal systems—a review. *Interactions*. 2024;245(1):282. doi:10.1007/s10751-024-02122-6.
- [24] Chekifi T, Boukraa M, Benmoussa A. Artificial intelligence for thermal energy storage enhancement: a comprehensive review. *J Energy Resour Technol*. 2024;146(6):060802. doi:10.1115/1.4065197.
- [25] Buffa S, Fouladfar MH, Franchini G, Lozano Gabarre I, Andrés Chicote M. Advanced control and fault detection strategies for district heating and cooling systems—a review. *Appl Sci*. 2021;11(1):455. doi:10.3390/app11010455.

- [26] Thawait LK, Singh MK. A review on real time implementation of soft computing techniques in thermal power plant. *Netw Comput Neural Syst.* 2025;36(1):1–37. doi:10.1080/0954898X.2024.2429721.
- [27] Pimenow S, Pimenowa O, Prus P. Challenges of artificial intelligence development in the context of energy consumption and impact on climate change. *Energies.* 2024;17(23):5965. doi:10.3390/en17235965.
- [28] Mazhar T, et al. Analysis of challenges and solutions of IoT in smart grids using AI and machine learning techniques: a review. *Electronics.* 2023;12(1):242. doi:10.3390/electronics12010242.
- [29] Zhang T, Strbac G. Novel artificial intelligence applications in energy: a systematic review. *Energies.* 2025;18(14):3747. doi:10.3390/en18143747.
- [30] Fan Z, Yan Z, Wen S. Deep learning and artificial intelligence in sustainability: a review of SDGs, renewable energy, and environmental health. *Sustainability.* 2023;15(18):13493. doi:10.3390/su151813493.
- [31] Alahi MEE, et al. Integration of IoT-enabled technologies and artificial intelligence (AI) for smart city scenario: recent advancements and future trends. *Sensors.* 2023;23(11):5206. doi:10.3390/s23115206.
- [32] Li Z, Liu L, Zhao Z, Mu S, Li D, Zhuo Y. Reinforcement learning-enhanced multi-objective optimization for sustainable coal blending in thermal power plants. *PLoS One.* 2025;20(9):e0331208. doi:10.1371/journal.pone.0331208.
- [33] Panda S, et al. Application of modified multi-verse optimization for temperature control in thermal power plant condensers. *Sci Rep.* 2026;16(1):12409. doi:10.1038/s41598-026-40559-7.
- [34] Ismail FB, Al-Kayiem HH, Kazem HA. AI adoption for steam boiler trip prevention in thermal power plants. *Int J Energy Prod Manag.* 2024;9(3):131–142. doi:10.18280/ijepm.090302.
- [35] Sun N, et al. A fast design tool for compact heat exchangers tube geometry to enhance thermohydraulic performance using various AI models. *Expert Syst Appl.* 2025;271:126635. doi:10.1016/j.eswa.2025.126635.
- [36] Coal Fired Power Plant Thermal Performance Dataset [Internet]. Kaggle; 2026 [cited 2026 Jan 6]. Available from: <https://www.kaggle.com/datasets/ainalirham/coal-fired-power-plant-thermal-performance-dataset>
- [37] Alvarez DI, González-Ladrón-de-Guevara F, Rojas Espinoza J, Borge-Diez D, Galindo SP, Flores-Vázquez C. The evolution of AI applications in the energy system transition: a bibliometric analysis of research development, the current state and future challenges. *Energies.* 2025;18(6):1523. doi:10.3390/en18061523.
- [38] Elshehabi T, Alfehaid M. Sustainable geothermal energy: a review of challenges and opportunities in deep wells and shallow heat pumps for transitioning professionals. *Energies.* 2025;18(4):811. doi:10.3390/en18040811.
- [39] Surase RS, Konijeti R, Chopade RP. Thermal performance analysis of gas turbine power plant using soft computing techniques: a review. *Eng Appl Comput Fluid Mech.* 2024;18(1):2374317. doi:10.1080/19942060.2024.2374317.
- [40] Hwang R, Lee J, Kim J, Moon I, Oh M. Autonomous digital twin framework for gas turbine combined cycle control loops: comparative study of proportional-integral control, reinforcement learning, and reinforcement learning with agents. *Energy AI.* 2026;24:100727. doi:10.1016/j.egyai.2026.100727.
- [41] A survey of artificial intelligence applications in nuclear power plants [Internet]. ProQuest; [cited 2026 Jul 28]. Available from: <https://www.proquest.com/openview/0bd3b6e1a4cb87cf279df81edb1f25e2/1>
- [42] Park J, Kim T, Koo S. Verification strategy for artificial intelligence components in nuclear plant instrumentation and control systems. *Prog Nucl Energy.* 2023;164:104842. doi:10.1016/j.pnucene.2023.104842.
- [43] William Nimmo [Internet]. ScienceDirect; [cited 2026 Jul 28]. Available from: <https://www.sciencedirect.com/author/7006402566/william-nimmo>
- [44] Teruel E, Cortés C, Díez LI, Arauzo I. Monitoring and prediction of fouling in coal-fired utility boilers using neural networks. *Chem Eng Sci.* 2005;60(18):5035–5048. doi:10.1016/j.ces.2005.04.029.
- [45] Gul E, Baldinelli G, Wang J, Bartocci P, Shamim T. Artificial intelligence based forecasting and optimization model for concentrated solar power system with thermal energy storage. *Appl Energy.* 2025;382:125210. doi:10.1016/j.apenergy.2024.125210.
- [46] Sheini Dashtgoli D, Giustiniani M, Busetti M, Cherubini C, Alessandrini G, Narsilio GA. Artificial intelligence in deep geothermal energy: trends, insights, and future perspectives. *Renew Energy Focus.* 2026;56:100781. doi:10.1016/j.ref.2025.100781.