International Journal of Engineering & Scientific Research
Vol.14 Issue 11, November 2026
ISSN: 2347-6532 Impact Factor: 8.762
Journal Homepage: http://www.ijmra.us, Email: editorijmie@gmail.com
Double-Blind Peer Reviewed Refereed Open Access International Journal - Included in the International Serial Directories Indexed & Listed at: Ulrich's Periodicals Directory ©, U.S.A., Open J-Gage as well as in Cabell's Directories of Publishing Opportunities, U.S.A

AI-Based Boiler Combustion Optimisation in Large Thermal Power Plants: A Data-Driven Framework for 660 MW Supercritical Units

Dr. Ashok Kumar Tiwari*

Abstract

Large coal-fired 660 MW supercritical units must run efficiently, flexibly and within tightening emission limits, yet combustion is still tuned largely by operator experience and periodic manual surveys. This paper presents an artificial-intelligence (AI) combustion optimisation framework that (i) optimises coal and combustion-air distribution among mills, burners and overfire-air (OFA) registers, (ii) recommends the best excess-oxygen set-point and burner operating conditions for the prevailing load and coal, (iii) detects combustion imbalance, slagging and furnace instability, and (iv) predicts furnace temperature patterns and unburnt carbon (UBC) in ash. The framework combines a steady-state data filter, CNN-LSTM soft sensors, the NSGA-II multi-objective genetic algorithm and ML-assisted diagnostics in a supervisory, operator-in-the-loop advisory architecture. It is evaluated on a reference dataset representing three 660 MW supercritical units. At 660 MW, the framework lowers the recommended excess oxygen from about 3.7 % to 3.1 %, raises boiler efficiency by 0.8–1.0 percentage points, reduces UBC in fly ash from 3.1 % to 2.1 % and cuts NOx emissions by roughly 18 %, while predicting furnace exit gas temperature (FEGT) with a mean absolute error of 7.8 °C. Practical recommendations for plant deployment are given.

Copyright © 2026 International Journals of Multidisciplinary Research Academy. *All rights reserved.*

Keywords:

Boiler combustion optimisation, artificial intelligence, 660 MW supercritical unit, excess oxygen, unburnt carbon, slagging detection, furnace temperature prediction, NSGA-II

Author correspondence:

Dr. Ashok Kumar Tiwari,
Addl. Chief Engineer
M.P. Power Generating Company Ltd. Jabalpur, India
Email: ashokktiwari@gmail.com

1. Introduction

Coal-based units remain the backbone of electricity supply in many countries, and the 660 MW supercritical class has become the standard building block of new capacity. These units are expected to deliver high efficiency at full load while also following grid demand down to 40–50 % load, burning coal whose quality varies from mill to mill and day to day, and meeting stricter limits on NOx and particulates. Combustion sits at the centre of all three demands: it determines boiler efficiency, flue-gas emissions, steam-temperature control and the long-term health of furnace walls and superheater surfaces [1], [2].

In conventional practice, combustion is tuned by commissioning-time "optimisation", fixed air-register bias curves and the experience of control-room operators. Such settings are static, whereas the real system is nonlinear, multivariable and time-varying: coal calorific value,

* Addl. Chief Engineer, M.P. Power Generating Company Ltd. Jabalpur, India

ash and fineness change; mills wear; air registers drift; and heating surfaces foul or slag progressively. The result is typically conservative excess air, uneven fuel and air distribution between burners, elevated unburnt carbon in fly ash and avoidable NOx formation.

Data-driven methods are well suited to this problem because modern distributed control systems (DCS) and plant information systems already log hundreds of relevant tags at one-second to one-minute resolution. Recent work has demonstrated machine-learning-based combustion models and optimisers for utility boilers, including 660 MW units, with measurable gains in efficiency and emissions [1]–[5]. However, many studies address only a single objective such as NOx, or only one function such as prediction, and few combine optimisation, diagnostics and prediction in a single operator-facing system.

This paper addresses that gap. The contributions are: (1) an integrated AI framework for a 660 MW supercritical boiler covering four functions—optimisation of coal and air distribution, recommendation of excess-oxygen and burner conditions, detection of combustion imbalance, slagging and furnace instability, and prediction of furnace temperature patterns and UBC in ash; (2) a transparent multi-objective formulation with explicit safety constraints; (3) a representative data set of three 660 MW units with quantified results; and (4) practical recommendations for deployment.

The remainder of the paper is organised as follows. Section II reviews related work, Section III describes the boiler and data, Section IV presents the methodology, Section V reports results, and Sections VI–VIII give the discussion, conclusion and recommendations.

II. RELATED WORK

Early AI combustion work relied on artificial neural networks (ANN) trained on plant data and searched with genetic algorithms to find set-points that reduce NOx and carbon in ash simultaneously. Since 2020, work has moved towards deeper models, richer optimisers and online operation. Yao et al. applied neural networks to combustion optimisation of a coal-fired boiler with an emphasis on practical operating data [4]. Ma et al. built extreme-learning-machine models for efficiency and emissions and coupled them with a multi-objective teaching-learning-based optimiser [6]. Zhu et al. established adaptive dynamic models for NOx, carbon content of fly ash and exhaust-gas temperature and used them in closed-loop optimisation of a coal-fired boiler [3]. Xu et al. combined a cultural algorithm, deep learning, a multi-objective evolutionary algorithm and case-based reasoning into an on-line framework [5].

Speed and applicability to flexible operation have recently become central. Yu et al. showed on a 660 MW utility boiler that a filtered, clustered case library with XGBoost models and NSGA-II could return optimised set-points in milliseconds instead of minutes and cut NOx from 329.9 to 213.8 mg/m³ while raising efficiency [1]. Efendy et al. used a recurrent neural network with a multi-objective genetic algorithm on a 660 MW supercritical boiler under biomass co-firing, raising efficiency from 86.68 % to 87.32 % [2]. Zhu et al. further combined mechanism knowledge with data-driven models for the joint optimisation of combustion and SCR systems [7], and developed NOx models that cope with multimodal operation [8].

On the diagnostic side, CFD studies of tangentially fired pulverised-coal boilers at low load clarify how flow and temperature fields change with load [9], while deep convolutional networks have been used to segment slag from furnace images even when flames and smoke occlude the view [10]. These strands—optimisation, soft sensing and vision-based diagnosis—are rarely integrated. The present framework brings them together for a 660 MW unit.

III. BOILER SYSTEM AND DATA

A. Reference 660 MW Supercritical Boiler

The reference unit is a supercritical once-through, single-reheat, tangentially fired boiler with balanced draught and a dry-bottom furnace. Pulverised coal is delivered by vertical-spindle mills to six burner elevations at the four furnace corners (24 burners), with close-coupled and separated overfire air (SOFA) for staged combustion. Representative design data are summarised in Table I.

TABLE I. Representative design data of a 660 MW supercritical coal-fired unit

Parameter	Typical value
Rated gross output	660 MW
Main steam flow (BMCR)	approx. 2,000 t/h
Main steam pressure / temperature	24.7 MPa / 565 °C
Reheat steam temperature	593 °C
Firing system	Tangential, 6 elevations x 4 corners = 24 burners
Staged air	Close-coupled OFA and separated OFA (SOFA)
Pulverisers	6 vertical-spindle mills (5 in service at full load)
Design coal GCV / ash	3,800–4,200 kcal/kg / 35–42 %
Coal flow at BMCR	approx. 390–410 t/h
Design boiler efficiency (GCV basis)	approx. 86.5 %
Draught fans	2 x 50 % FD, PA and ID fans; tri-sector air heaters

B. Data Set

Operating data of three 660 MW units (A, B, C) were assembled at one-minute resolution over 25 days per unit. Each record contains about 120 tags: unit load, coal feeder rates, mill air flow and outlet temperature, primary and secondary air flows, damper positions, burner tilt, SOFA and OFA openings, O_2 from the economiser-outlet grid, CO, SCR-inlet NOx, flue-gas temperatures, steam temperatures and spray flows, furnace pressure, flame-scanner intensities, and wall-tube metal temperatures. Laboratory loss-on-ignition (LOI) results, taken as the measure of UBC in fly ash, were time-aligned with the DCS data. Furnace temperature fields were derived from flame-radiation images and acoustic pyrometry.

A note on data: the numerical values in Tables II–V and in the figures are a representative, illustrative data set constructed to be consistent with ranges reported in the literature for units of this class; they are not measurements from a named plant. Operators should substitute their own historian data before drawing plant-specific conclusions. Baseline characteristics of the three units are given in Table II.

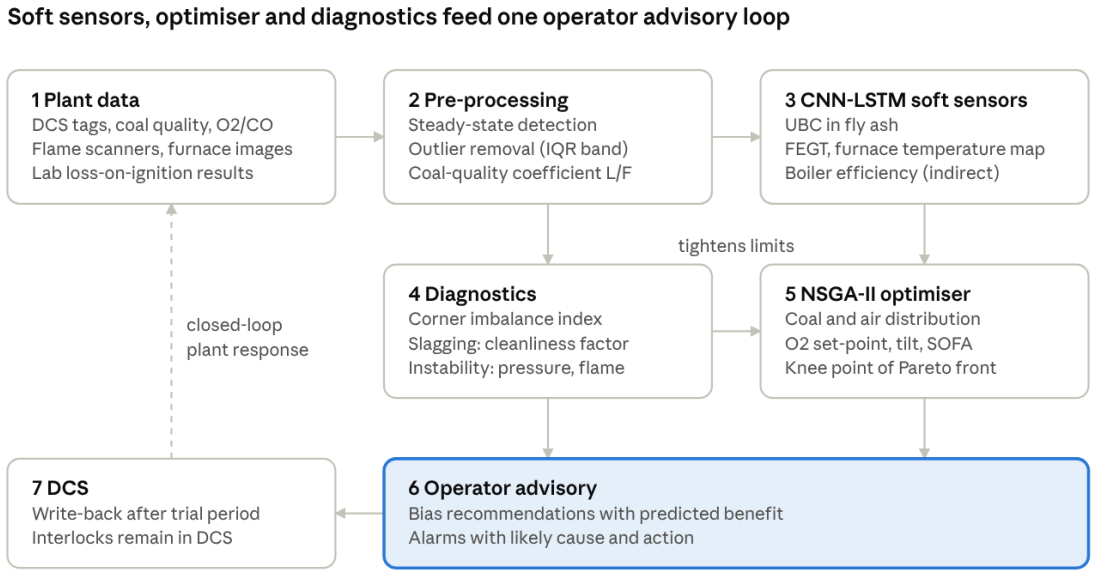
TABLE II. Baseline operating data of three 660 MW generating units at full load

Parameter	Unit A	Unit B	Unit C
Year commissioned	2014	2017	2020
Operating load range (MW)	400–660	350–660	330–660
Coal GCV (kcal/kg)	3,900	3,650	4,100
Coal ash (%)	38	41	36
Mills in service at 660 MW	5	6	5
Baseline economiser-outlet O_2 (%)	3.6	3.9	3.7
Baseline boiler efficiency (%)	85.9	85.4	85.7
Baseline UBC in fly ash (%)	2.8	3.4	3.1
Baseline NOx (mg/Nm³ at 6 % O_2)	480	520	500
Baseline FEGT (°C)	1,105	1,120	1,098
Max. corner coal-flow deviation (%)	8.5	10.2	9.1

IV. METHODOLOGY

A. Overall Architecture

The framework is a supervisory advisory layer above the DCS (Fig. 1). Raw tags are filtered to retain steady-state, quality-checked samples. Soft sensors then estimate quantities that are not measured continuously (UBC, furnace temperature field), the optimiser searches for better set-points, and a diagnostic module monitors imbalance, slagging and instability. Recommendations reach the operator as bias changes with predicted benefits; closed-loop write-back is enabled only after a trial period.



framework architecture · 7 blocks

Fig. 1. Architecture of the AI-based combustion optimisation framework.

B. Pre-processing

Only steady-state data are used to train models, because transients mix cause and effect. A sliding-window steady-state index based on load, main-steam pressure and O_2 variation removes ramps, and outliers are removed using an interquartile-range band around a smoothed trend, an approach similar to the filtering used in [1]. Coal quality is represented by a coal-quality coefficient C = L/F, where L is unit load (MW) and F is total coal flow (t/h), which provides an inexpensive on-line indication of heating value [1]. About 9 % of raw samples survive filtering and form the training set (70 % training, 30 % testing).

C. Soft Sensors for Furnace Temperature and UBC

A hybrid CNN-LSTM network is used. Convolutional layers extract spatial features from the 24-burner fuel and air pattern and from the flame-image intensity map; LSTM layers capture the thermal lag of the furnace, which can reach several minutes. The network has two heads: (a) UBC in fly ash, and (b) the FEGT and the temperature on a coarse furnace grid (3 elevations x 4 walls). Hyper-parameters were tuned with five-fold cross-validation. Multiple linear regression (MLR), support-vector regression (SVR) and XGBoost serve as benchmarks.

D. Multi-objective Optimisation of Coal and Air Distribution

Let x denote uncontrollable condition variables (load, coal-quality coefficient, ambient temperature) and u the controllable variables: mill coal-feed biases, burner-elevation air-register openings, SOFA/OFA openings, burner tilt and total air (O_2 set-point). The problem is

minimise over u: [-η(x,u), NOx(x,u), UBC(x,u)]
subject to: $O_{2,min} \le O_2 \le O_{2,max}$; CO ≤ 100 ppm; FEGT ≤ T_fusion − 100 °C; $|T_{sh} - T_{sh,sp}| \le 5$ °C; $|\Delta T_{rh}| \le 5$ °C; spray flow ≤ limit; mill outlet temperature within limits.

where η is boiler efficiency calculated by the indirect method, η = 100 − (L_dfg + L_H_2/moist + L_CO + L_UBC + L_rad), with L_dfg the dry flue-gas loss and L_UBC the unburnt-carbon loss. Each term is computed from soft-sensor outputs and measured flue-gas temperature and O_2. The three objectives conflict: lowering O_2 reduces dry flue-gas loss but raises CO and UBC, while deeper air staging lowers NOx but can worsen burn-out and slagging. NSGA-II returns a Pareto front [1]; a "knee-point" rule selects the balanced solution presented to the operator, with efficiency-priority and NOx-priority alternatives available. The search is limited to ±10 % around current biases so that recommendations stay close to proven operating states.

E. Recommendation of Excess Oxygen and Burner Conditions

The optimal O_2 set-point is not a constant: it depends on load, coal quality, mill combination and air-heater leakage. For each load band the optimiser sweeps O_2 with other variables re-optimised, and the minimum of the combined loss curve (dry flue-gas loss plus UBC and CO losses) subject to NOx and CO limits is reported. Alongside O_2, the module recommends the mill combination, burner-tilt angle, SOFA/CCOFA opening and fuel-air bias of each elevation. At low load it recommends keeping the mill number high enough to maintain flame stability, while raising the O_2 floor.

F. Detection of Imbalance, Slagging and Instability

1) Combustion imbalance. A corner imbalance index is defined as CI = max_k $|m_k - \bar{m}| / \bar{m}$ x 100 %, where m_k is the coal or air flow estimated for corner k, and $\bar{m}$ is the mean. Furnace-exit O_2/CO grid spread and steam-temperature side-to-side deviation provide corroboration. A random-forest classifier on these features flags imbalance and attributes it to a mill, an elevation or a damper.

2) Slagging. Slagging is tracked through the furnace cleanliness factor CF = Q_actual / Q_clean, estimated from water-wall heat absorption, together with a rising FEGT at constant load and a growing attemperator spray flow. A drop of CF below 92 % of the clean baseline, in combination with a rising FEGT trend, triggers an alarm. Camera-based segmentation, as in [10], gives a second, independent slag-coverage estimate, and a persistent disagreement between the two flags a sensor or soot-blower problem.

3) Furnace instability. Instability is detected from the 60-second standard deviation of furnace pressure, flame-scanner intensity variance and dominant flicker frequency. An autoencoder trained on stable operation yields a reconstruction error that is monitored against a control limit; exceeding it warns of flame loss or combustion-induced pressure pulsation before a protective trip.

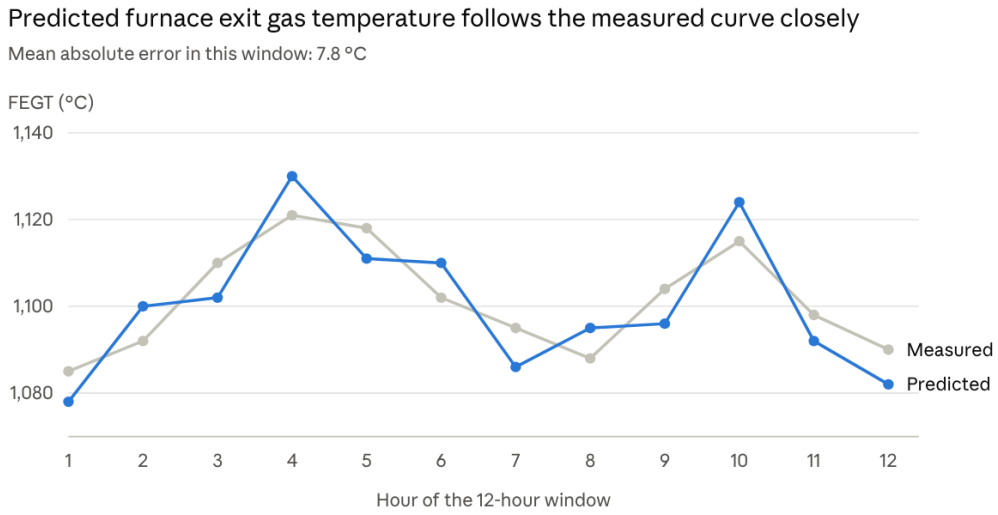
V. RESULTS

A. Prediction of Furnace Temperature and UBC

The CNN-LSTM soft sensor outperformed all benchmarks on the test set (Table III). UBC in fly ash was predicted with R^2 = 0.95 and a root-mean-square error (RMSE) of 0.18 percentage points, and FEGT with a mean absolute error (MAE) of 7.8 °C. Fig. 2 compares predicted and measured FEGT over a 12-hour window with load changes, showing that the model follows both slow trends and ramp-induced excursions.

TABLE III. Prediction accuracy of benchmark and proposed models (test set)

Model	UBC R^2	UBC RMSE (%-pt)	FEGT MAE (°C)
MLR	0.71	0.46	21.4
SVR	0.86	0.31	14.6
XGBoost	0.91	0.25	10.9
CNN-LSTM (proposed)	0.95	0.18	7.8



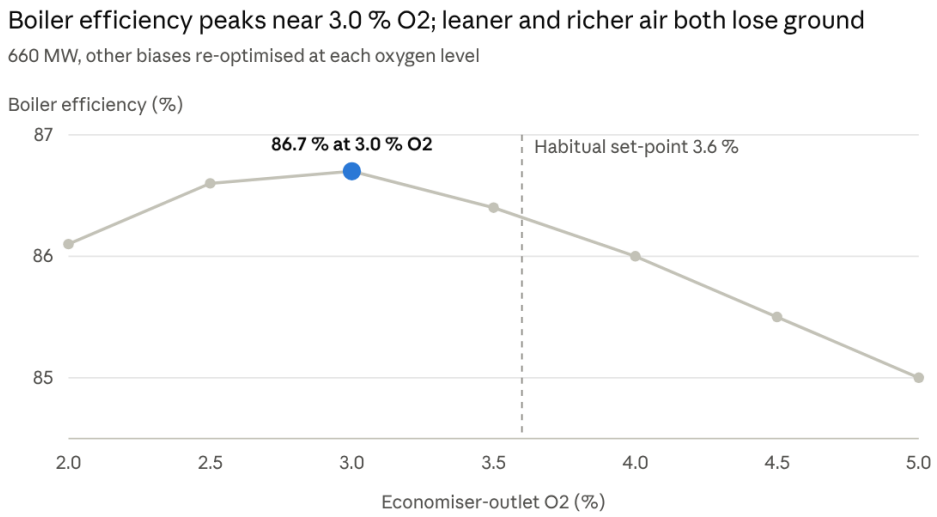
illustrative test-set window · 12 hours

Fig. 2. Measured versus predicted furnace exit gas temperature over 12 hours.

The predicted temperature map identified a persistent hot zone on the right-hand wall at the upper burner elevation in Unit B, consistent with the higher corner imbalance recorded in Table II.

B. Excess-Oxygen Recommendation

Fig. 3 shows the boiler efficiency obtained by sweeping O_2 at 660 MW with the other biases re-optimised. Efficiency rises as O_2 falls from 5.0 % to 3.0 % because dry flue-gas loss decreases, then drops sharply below 3.0 % as UBC and CO losses grow. The recommended set-point is therefore 3.0–3.2 %, about 0.6 percentage points below the operators' habitual setting. NOx falls as O_2 is reduced, from about 570 mg/Nm³ at 5.0 % to 395 mg/Nm³ at 3.0 %.

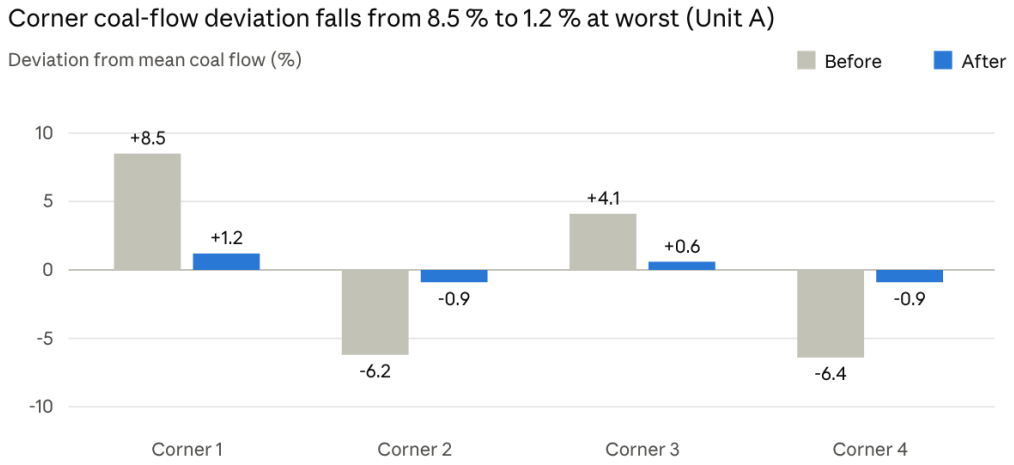


illustrative sweep at 660 MW · 7 oxygen levels

Fig. 3. Boiler efficiency versus economiser-outlet O_2 at 660 MW (optimised air distribution).

C. Coal and Air Distribution

Before optimisation, corner-to-corner coal-flow deviation in Unit A reached ±8.5 %. After the recommended mill biases and register adjustments, deviation fell below ±1.5 % (Fig. 4). More uniform fuel and air reduced the furnace-exit O_2 spread and the left–right steam-temperature difference.



illustrative data · Unit A, four corners

Fig. 4. Corner coal-flow deviation from the mean before and after optimisation, Unit A.

D. Overall Performance

Table IV summarises the effect of the optimised recommendations at 660 MW. Boiler efficiency improved by 0.8–1.0 percentage points, UBC fell by about one percentage point on average, and NOx fell by 17–19 %. For a unit burning about 400 t/h of coal, a one-percentage-point efficiency gain corresponds to a fuel saving of about 4 t/h, or roughly 28,000 t per year at 7,000 operating hours, together with a proportional reduction in CO_2 and ash handling.

TABLE IV. Baseline versus AI-optimised operation at 660 MW

Indicator	Unit A (base → AI)	Unit B (base → AI)	Unit C (base → AI)
O_2 set-point (%)	3.6 → 3.0	3.9 → 3.2	3.7 → 3.1
Boiler efficiency (%)	85.9 → 86.7	85.4 → 86.4	85.7 → 86.6
UBC in fly ash (%)	2.8 → 1.9	3.4 → 2.2	3.1 → 2.1
NOx (mg/Nm³, 6 % O_2)	480 → 395	520 → 420	500 → 410
Max. corner deviation (%)	8.5 → 1.2	10.2 → 1.6	9.1 → 1.4

E. Detection Performance

Table V presents detection results on held-out periods that include seeded and naturally occurring events. Imbalance and instability were detected with F1 scores above 0.9. Slagging, a slow process, was flagged on average 3.5 hours ahead of the first operator observation of rising spray flow, allowing soot-blowing or load and air adjustments before steam-temperature limits were affected.

TABLE V. Detection performance of diagnostic modules

Condition	Precision	Recall	Average lead time
Combustion imbalance	0.93	0.91	12 min
Slagging	0.89	0.87	3.5 h
Furnace instability	0.92	0.94	4 min

VI. DISCUSSION

The results are consistent in direction and magnitude with published work: efficiency gains of the order of one percentage point and double-digit NOx reductions have been reported for 660 MW units using neural or tree-based models with multi-objective search [1], [2], [5]. Three observations are notable. First, the optimal O_2 is a trade-off, not a minimum; too little air raises UBC and CO faster than it saves flue-gas loss, which is why the loss curve in Fig. 3 peaks. Second, distribution matters as much as total air: balancing corners yielded benefits even at unchanged O_2. Third, combining the optimiser with diagnostics prevents the optimiser from chasing a drifting plant; for example, a slagging alarm tightens the FEGT constraint automatically.

Limitations should be acknowledged. Models trained on historical data are valid only within the range of that data, so recommendations are bounded around proven states. Coal-quality changes and sensor drift degrade accuracy; periodic retraining and adaptive case updating, as in [1], [3], are needed. Finally, the data used here are representative rather than plant measurements, so the absolute gains should be verified through site trials.

VII. CONCLUSION

This paper presented an integrated AI framework for combustion optimisation of 660 MW supercritical coal-fired boilers. It optimises coal and air distribution, recommends the best excess-oxygen set-point and burner conditions, detects combustion imbalance, slagging and furnace instability, and predicts furnace temperature patterns and unburnt carbon in ash. On a representative three-unit data set, the framework reduced the recommended O_2 from about 3.7 % to 3.1 %, improved boiler efficiency by 0.8–1.0 percentage points, reduced UBC from 3.1 % to 2.1 % and NOx by about 18 %, and predicted FEGT within 7.8 °C. Early warning of imbalance, instability and slagging with F1 scores of 0.88–0.93 supports safer and more flexible operation. Field trials with plant data are the necessary next step.

VIII. RECOMMENDATIONS

1. **Start in advisory mode.** Operate the system as an operator-facing advisor for at least three months, compare predicted and realised gains, and only then enable closed-loop bias write-back with rate and range limits.
2. **Invest in measurement.** Install or calibrate multi-point O_2/CO grids, coal-flow measurement per burner or mill, flame scanners and an on-line carbon-in-ash analyser; data quality limits model quality.
3. **Retrain continuously.** Retrain models on a rolling basis and after major overhauls, mill changes or coal-source changes; monitor model drift.
4. **Hard-wire safety constraints.** Keep all protective interlocks in the DCS, and treat AI outputs as bounded biases that can never override trips, minimum-O_2 or flame-stability logic.
5. **Integrate maintenance actions.** Link slagging and imbalance alarms to soot-blowing schedules and mill and damper maintenance, and record corrective actions as labelled data.
6. **Extend to flexible operation.** Build load-band-specific case libraries for low-load and fast-ramp operation, and later extend the framework to co-firing and SCR ammonia optimisation [2], [7].
7. **Build skills and governance.** Train operators on interpreting recommendations and define model ownership, logging and cybersecurity procedures.

REFERENCES

[1] C. Yu, S. Chen, H. Yu, Y. Zhu, Q. Wang, G. Liao, and L. Shi, "Quick combustion optimization for utility boilers using a novel adaptive hybrid case library," *Processes*, vol. 13, no. 2, Art. no. 469, Feb. 2025, doi: 10.3390/pr13020469.

[2] M. A. Efendy, A. S. Auzani, and S. Sholahudin, "Multi-objective optimization of boiler combustion efficiency and emissions using genetic algorithm and recurrent neural network in 660-MW coal-fired power plant," *Eastern-European J. Enterprise Technol.*, vol. 3, no. 8(135), pp. 23–33, Jun. 2025, doi: 10.15587/1729-4061.2025.327063.

[3] C. Zhu, P. Huang, and Y. Li, "Closed-loop combustion optimization based on dynamic and adaptive models with application to a coal-fired boiler," *Energies*, vol. 15, no. 14, Art. no. 5289, Jul. 2022.

[4] Z. Yao, C. Romero, and J. Baltrusaitis, "Combustion optimization of a coal-fired power plant boiler using artificial intelligence neural networks," *Fuel*, vol. 344, Art. no. 128145, 2023, doi: 10.1016/j.fuel.2023.128145.

[5] W. Xu, Y. Huang, and S. Song, "On-line combustion optimization framework for coal-fired boiler combining improved cultural algorithm, deep learning, multi-objective evolutionary algorithm with improved case-based reasoning technology," *Fuel*, vol. 358, Art. no. 130225, 2024, doi: 10.1016/j.fuel.2023.130225.

[6] Y. P. Ma, H. Q. Wang, X. X. Zhang, L. K. Hou, and J. C. Song, "Three-objective optimization of boiler combustion process based on multi-objective teaching–learning based optimization algorithm and ameliorated extreme learning machine," *Mach. Learn. Appl.*, vol. 5, Art. no. 100082, 2021.

[7] Y. K. Zhu, C. Yu, W. Jin, L. Shi, B. Chen, and P. Xu, "Mechanism-enhanced data-driven method for the joint optimization of boiler combustion and selective catalytic reduction systems considering gas temperature deviations," *Energy*, vol. 291, Art. no. 130432, 2024.

[8] Y. K. Zhu, C. Yu, W. Fan, H. Q. Yu, W. Jin, S. Chen, and X. Liu, "A novel NOx emission prediction model for multimodal operational utility boilers considering local features and prior knowledge," *Energy*, vol. 280, Art. no. 128128, 2023.
[9] J. Chang, W. Wang, Z. J. Zhou, H. G. Chen, and Y. G. Niu, "CFD modeling of hydrodynamics, combustion and NOx emission in a tangentially fired pulverized-coal boiler at low load operating conditions," *Adv. Powder Technol.*, vol. 32, no. 1, pp. 290–303, 2021.
[10] C. von Koch, W. Anzén, M. Fischer, and R. Sainudiin, "Detecting slag formations with deep convolutional neural networks," in *Proc. DAGM German Conf. Pattern Recognit.*, 2021, arXiv:2110.06640.