energies

MDPI

Article

Graph Neural Operator-Based Surrogate Modelling of Multi-Field CFD Results in Biomass Boiler

Przemysław Motyl [1,*], Danuta Król [2] and Sławomir Poskrobko [3]

[1] Faculty of Mechanical Engineering, Casimir Pulaski Radom University, 26-600 Radom, Poland
[2] Astwa Sp. z o.o., 15-102 Białystok, Poland; dankrol@wp.pl
[3] Faculty of Civil Engineering and Environmental Sciences, Białystok University of Technology, 15-351 Białystok, Poland; drposkrobko@wp.pl
* Correspondence: p.motyl@urad.edu.pl

Abstract

Computational fluid dynamics provides detailed spatial distributions of physical fields in biomass boiler combustion, but the computational cost of each simulation limits its application in parametric studies and near-real-time workflows. This work investigates whether a Graph Neural Operator (GNO) can serve as a fast surrogate model that maps boiler operating parameters to six coupled CFD field distributions simultaneously. The reference case is a 10 kW wood-pellet boiler with internal flue gas recirculation (FGR), described and experimentally validated in an earlier publication by the authors. CFD data were collected on the symmetry plane of the combustion chamber for 80 operating points defined by the thermal load ratio (P/P_0) and the excess air ratio λ. A GNO surrogate was trained on 64 cases to predict temperature, velocity magnitude, static pressure, and mole fractions of CO, O_2, and CO_2 at each node of an unstructured spatial graph. On a held-out validation set of 16 operating cases, the model achieved R^2 values of 0.988 for temperature, 0.919 for velocity magnitude, 0.982 for pressure, 0.999 for CO, 0.992 for O_2, and 0.985 for CO_2. After training, each prediction is generated in a single forward pass, providing a computationally efficient approximation compared to the full CFD solver. A dedicated generalisation study on independent off-grid CFD cases confirmed that the surrogate interpolates within the parameter domain with essentially no loss of accuracy and degrades only moderately when extrapolated towards a higher thermal load and leaner mixtures. The results demonstrate that a baseline GNO surrogate can capture the spatial structure of coupled thermo-fluid and species fields in a realistic combustion geometry within the investigated parameter range and suggest applicability to digital-twin-oriented workflows where repeated parametric queries of boiler operation are required.

Keywords: graph neural operator; surrogate model; computational fluid dynamics; biomass combustion; digital twin; multi-field prediction; biomass boiler; flue gas recirculation; machine learning for CFD

Academic Editor: Jiaqiang E

Received: 22 June 2026
Revised: 10 July 2026
Accepted: 11 July 2026
Published: 14 July 2026

1. Introduction

Low-emission combustion of biomass in small-scale domestic boilers is an important topic in the context of European energy policy and the increasing use of renewable solid fuels. Raising the efficiency and reducing the emissions of this conversion process is part of the wider effort to decarbonise the heat supply, and any tool that lowers the cost of optimising and controlling boiler operation therefore has direct bearing on the performance

of the distributed energy systems into which these units are embedded. The design and optimisation of such boilers require a detailed understanding of the coupled thermo-fluid processes and reaction kinetics inside the combustion chamber. Modelling this type of equipment is particularly demanding because the flow field, temperature distribution, and gaseous species concentrations are tightly coupled through turbulence, radiation heat transfer, and global combustion reactions. Flue gas recirculation (FGR), used to reduce NO_x emissions and improve combustion efficiency, further increases the physical complexity of the system.

Computational fluid dynamics provides a well-established framework for modelling these processes. Three-dimensional CFD models, solved with turbulence closures such as the realisable k-ε model and radiation models such as the discrete ordinates method, can resolve the spatial structure of temperature, velocity, pressure, and species concentration fields with a level of detail that is difficult to obtain by experiment alone. The reference case used in this work is a 10 kW wood-pellet boiler with an innovative furnace geometry and internal FGR channels, described and experimentally validated by the authors in [1]. That work showed good agreement between CFD results and measured temperature and emission values under full-load conditions.

The main limitation of CFD is its computational cost. A single simulation of the full 3D boiler model with approximately 945,000 polyhedral cells requires on the order of 20–30 min of computation per operating point. This makes direct use of CFD in parametric studies, design optimisation, or near-real-time digital-twin applications impractical. Data-driven surrogate models trained on a limited number of CFD simulations offer a way to obtain fast approximate predictions of field distributions for new operating conditions.

Classical surrogate approaches based on multilayer perceptrons or convolutional neural networks have been applied to flow field approximation [2,3], but they require data on structured grids and generalise poorly when the mesh changes. The Neural Operator paradigm, introduced by Li et al. [4,5], addresses this limitation by learning a mapping between function spaces rather than between fixed finite-dimensional vectors. The Fourier Neural Operator [6] is a widely used spectral instance of this framework, and neural operators have more broadly been reviewed as a tool for accelerating scientific simulation and design [7]. Neural operators can be evaluated on arbitrary discretisations of the input domain and are therefore naturally suited to unstructured CFD meshes.

The Graph Neural Operator (GNO), proposed in [4], realises the kernel integration of the operator approximation through message passing on a spatial graph constructed from the CFD mesh points. Each node represents a spatial location in the domain and edges connect neighbouring nodes according to a proximity criterion. The operator kernel is parameterised by a neural network that takes edge features as input. This framework is compatible with unstructured point clouds and does not require any mesh regularisation.

Several recent studies have demonstrated the application of neural operators to combustion and boiler problems. Wu et al. [8] used a factorised Fourier Neural Operator to reconstruct temperature and species fields in turbulent jet flames at different resolutions. Chen et al. [9] applied proper orthogonal decomposition as a reduced-order model for temperature and species prediction in a pulverised coal boiler. Wang et al. [10] developed a flow-field-informed POD-based digital twin for a coal-fired boiler. Fourier neural operators have likewise been applied to the fast prediction of multi-physics sensor responses [11]. Particle transport modelling in biomass combustors has also been addressed through improved Lagrangian tracking formulations [12]. Recent work has further explored neural operator architectures that unify graph-based and attention-based representations [13]. These works establish the broader context for data-driven field prediction in combustion

engineering. However, published examples of the GNO formulation applied to realistic boiler geometries with simultaneous prediction of multiple coupled fields remain limited.

The aim of this paper is to evaluate the feasibility of a baseline GNO surrogate model for the simultaneous prediction of six coupled CFD output fields in the biomass boiler case from [1]. The novelty lies in the application of a GNO-based neural surrogate to a realistic combustion-related CFD case involving coupled thermo-fluid and species-transport fields, rather than in proposing a new GNO architecture. The paper does not present an ablation study of architectural variants and does not claim superiority of the implemented model over other neural operator formulations. The implemented surrogate is based on the standard GNO framework and is trained and evaluated on the boiler CFD dataset within the investigated parameter range.

2. Biomass Boiler with Flue Gas Recirculation—CFD Reference Case

The physical system studied in this work is a 10 kW prototype boiler fuelled by wood biomass pellets with a diameter of 3 mm and a length of approximately 10–15 mm. The boiler uses a retort-type burner with a central cylindrical combustion chamber (Figure 1). An innovative feature of the design is the arrangement of internal flue gas recirculation channels surrounding the firebox. The exhaust gases pass through a four-pass convective circuit before reaching the flue. This configuration promotes intensive heat transfer, extends the residence time of combustion products, and supports low-emission operation meeting the requirements of the ECODESIGN Directive [1].

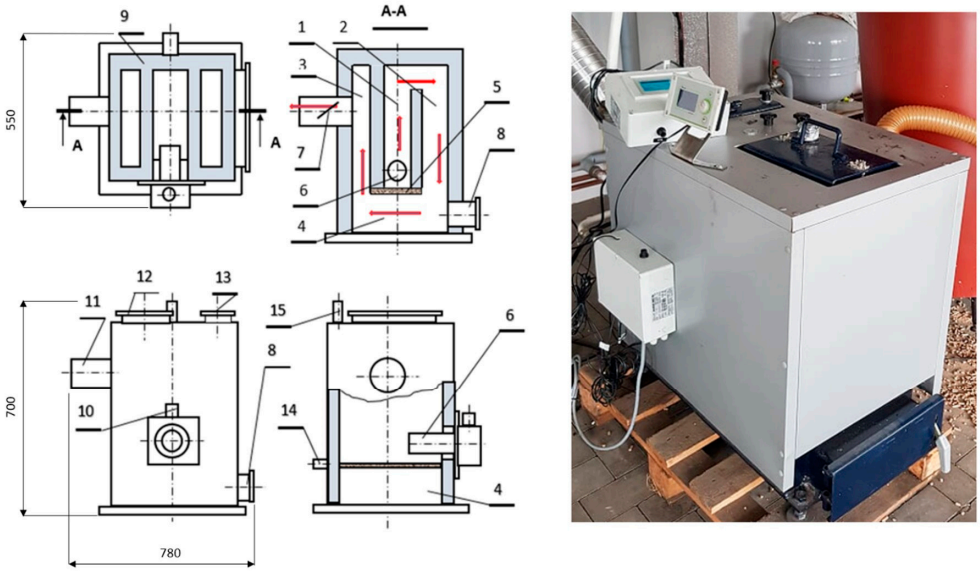

Figure 1. View of the boiler with 10 kW flue gas circulation: 1—combustion chamber, 2—convection chamber I, 3—convection chamber II, 4—combustion gas circulation chamber, 5—airtight sliding grate closing the combustion chamber, 6–10 kW retort burner, 7—thrust adjustment slide, 8—ash removal cleaning duct, 9—water jacket (flue gas/water heat exchanger), 10—fuel feed (pellet), 11—flue (exhaust outlet), 12—inspection hatch of the furnace and convection chamber I, 13—inspection hatch for convection chamber II, 14—return water connector, and 15—feed water connector. Reproduced from [1].

The combustion process begins with partial gasification of the pellet fuel in the retort chamber, producing a syngas mixture with the following dry-state composition: 15.1% CO_2, 18.4% CO, 14.2% H_2, 3.3% CH_4, and 49.0% N_2 by volume, with a lower heating value of approximately 4857 kJ/Nm^3 [1]. The syngas is subsequently burned in the main combustion chamber, where secondary air is supplied through peripheral inlets.

The CFD model of the full boiler was built in ANSYS Fluent 2025 R2 using a numerical grid of approximately 945,000 polyhedral cells (Figure 2). The realisable k-ε turbulence model with enhanced wall treatment was used. Radiation heat transfer was modelled with the discrete ordinates method. Combustion reactions were represented by a three-step global kinetics scheme (CH_4 oxidation, CO oxidation, H_2 oxidation) using the Finite-Rate/Eddy-Dissipation hybrid model [1]. Boundary conditions included a syngas inlet temperature of 1123 K, an air inlet temperature of 300 K, and wall cooling by a water jacket at 338 K with a convection coefficient of 2500 W/(m^2·K).

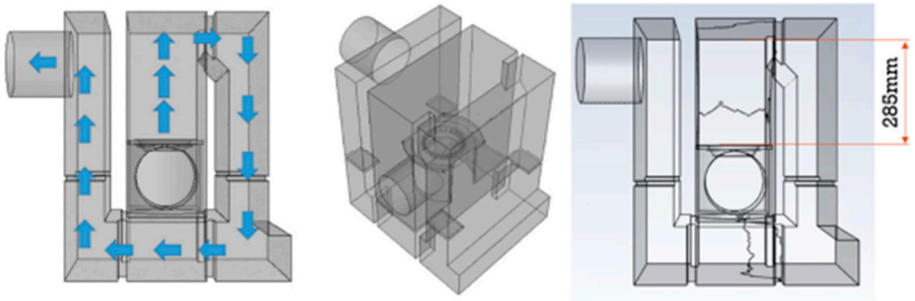

Figure 2. Boiler modelling geometry used in the CFD simulations. Reproduced from [1].

The CFD model was experimentally validated at the experimental facility of Poznań University of Technology. Measured average emission values of CO = 91 mg/Nm^3 and NO_x = 197 mg/Nm^3 were obtained under full-load conditions, with a combustion chamber temperature of 450–500 °C and a flue gas temperature of 157–197 °C. The boiler efficiency was determined to be 92%. These results confirmed the validity of the CFD model at the design operating point [1].

In the present study, the validated CFD model served as the reference solver for training dataset generation. Multiple CFD simulations were performed for different combinations of operating parameters. For each simulation, spatial distributions of temperature, velocity magnitude, static pressure, and mole fractions of CO, O_2, and CO_2 were exported from the Fluent solver on the two-dimensional symmetry plane of the combustion chamber.

3. Dataset Preparation and Graph Representation

3.1. Operating Parameter Space and CFD Simulation Plan

Two scalar parameters were adopted as independent variables defining the boiler operating point: the thermal load ratio P/P_0, expressing the current thermal output relative to the nominal 10 kW load, and the excess air ratio λ, controlling the amount of combustion air supplied relative to stoichiometric. These two parameters govern the syngas inlet velocity V_fuel and the air inlet velocity V_air through the boundary conditions of the CFD model according to scaling relations in Equation (1) derived from the full-load operating data of [1]:

$$V_{fuel} = \frac{P}{P_0} \cdot V_{fuel,0};\ V_{air} = \frac{\lambda}{\lambda_0} \cdot \frac{P}{P_0} \cdot V_{air,0} \quad (1)$$

where $\lambda_0 = 2$ is the reference excess air ratio at which the boiler was operated and experimentally validated in [1], and the reference inlet velocities V_fuel,0 = 1.08 m/s and V_air,0 = 2.34 m/s are derived from the full-load volumetric flow rates reported in [1] (Q_syngas = 0.00206 Nm^3/s and Q_air = 0.00431 Nm^3/s, both corresponding to $\lambda = \lambda_0$) divided by the respective inlet cross-sectional areas of the Fluent geometry defined in [1]. The parameter space covered P/P_0 values from 0.70 to 1.40 and λ values from 1.25 to 3.00, resulting in a structured grid of 80 CFD simulations.

3.2. Graph Construction from CFD Point Cloud

Each exported CFD case was transformed into a spatial graph. Nodes of the graph correspond to the centroid coordinates of the mesh cells on the symmetry plane. The node set is identical across all 80 simulations because the computational mesh does not change between operating conditions.

The edge connectivity was determined by a hybrid approach combining Delaunay triangulation and k-nearest neighbour (k-NN) search. First, a Delaunay triangulation of the node positions was computed and all triangle edges were included in the graph. Then, for each node, the 15 nearest neighbours were identified using a kd-tree and bidirectional edges to these neighbours were added. The union of Delaunay and k-NN edges was used as the final edge set, with duplicates removed. This approach ensures that the graph reflects the local spatial structure of the CFD mesh and that boundary nodes with long Delaunay edges remain connected to their spatial neighbours.

Each directed edge (i, j) was assigned an eight-dimensional edge attribute vector encoding the geometric relationship between its endpoints: the unit direction vector (two components), the normalised Euclidean distance, the square of the normalised distance, the inverse normalised distance, the signed angle to the horizontal axis in normalised form, and the sine and cosine of the angle.

3.3. Node Features and Output Representation

Each node was assigned an input feature vector of four components: the two global operating parameters (P/P_0 and λ), which are the same for all nodes in a given graph, and the two normalised spatial coordinates (x_norm and y_norm). Spatial coordinates were normalised to the range $[-1, 1]$ using global minima and maxima across all simulations. The total input dimension per node was therefore four.

The output target for each node was a six-dimensional vector containing the physical values of temperature T [K], velocity magnitude $|v|$ [m/s], static pressure p [Pa], and mole fractions of CO, O_2, and CO_2. Each output channel was normalised independently using a min–max transformation to the range $[-1, 1]$. Normalisation statistics were computed exclusively from the training set and applied to both training and validation data to prevent data leakage.

3.4. Train and Validation Split

The 80 available CFD cases were divided into a training set of 64 cases and a validation set of 16 cases using a random 80/20 split with a fixed random seed for reproducibility (seed = 42). The validation set contained operating conditions not seen during training, providing an estimate of model generalisation within the investigated parameter range.

4. Baseline Graph Neural Operator Surrogate

4.1. Architecture Overview

The surrogate model is based on the Graph Neural Operator formulation. In the GNO framework, the surrogate approximates a nonlinear operator that maps input functions defined on the spatial domain to output functions on the same domain through an iterative kernel integration realised by message passing on the spatial graph. The architecture in this work consists of three sequential stages: a lifting layer, a stack of message passing layers, and an output projection (Figure 3).

For each graph node, the lifting layer maps the four-dimensional input feature vector to a hidden representation of dimension 192. This representation is then processed by eight successive GNO message passing layers, each aggregating information from neighbouring nodes according to the edge attributes. After the message passing stack, a graph-level field

aggregation step is applied to couple the representations of different physical fields before the output projection. The projection layer maps the 192-dimensional hidden representation to six output values corresponding to the six predicted CFD fields.

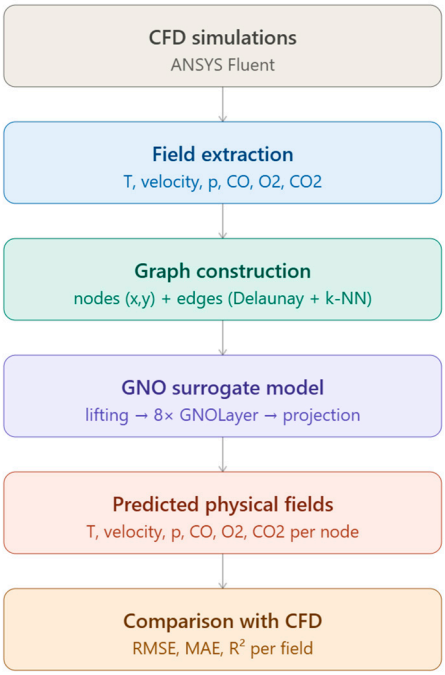

Figure 3. Workflow of the baseline GNO surrogate, from CFD reference data and graph construction to per-node prediction of the six physical fields and comparison with CFD.

4.2. Lifting Layer

The lifting layer transforms per-node input features into the hidden representation space. It is implemented as a three-layer multi-layer perceptron (MLP) with layer normalisation and GELU activation functions after each linear layer except the last. The input dimension is four and the output dimension is the model width of 192. A dropout rate of 0.1 is applied after the first hidden layer.

4.3. GNO Kernel Integration

The surrogate model is built on the neural operator framework introduced in [4]. In that framework, a nonlinear operator mapping input functions to output functions is approximated by a stack of iterative kernel integration steps. For each iteration t, the hidden node representation v(x) is updated according to the continuous formulation in Equation (2):

$$v^{(t+1)}(x) = \sigma\left(Wv^{(t)}(x) + \int_{B(x,r)} \kappa_\phi(x,\, y,\, a(x),\, a(y))v^{(t)}(y)\, dy\right) \tag{2}$$

where σ is a pointwise nonlinear activation, W is a learnable weight matrix, κϕ is a kernel network parameterised by ϕ that maps the pair (x, y, a(x), a(y)) to a matrix in $R^{(n\times n)}$, and the integral is taken over a ball B(x, r) of radius r centred at x. This formulation is motivated by the Green's function representation of the solution operator for elliptic PDEs [4]. On a discrete spatial graph, the integral is replaced by a Monte Carlo sum over the neighbourhood N(i) of each node i:

$$v_i^{(t+1)} = \sigma\left(Wv_i^{(t)} + \frac{1}{|N(i)|}\sum_{j\,\in\, N(i)} \kappa_\phi\left(e_{ij}\right)v_j^{(t)}\right) \tag{3}$$

where e_{ij} is the edge attribute vector encoding the geometric relationship between nodes i and j. This graph-based discretisation retains mesh independence because the kernel $\kappa\phi$ shares parameters across all node pairs regardless of the underlying grid resolution [4]. The implementation in this work realises Equation (3) through the pre-norm message passing described in Section 4.4.

4.4. Message Passing Layers

Each GNO message passing layer follows the pre-norm formulation [4]. Let $\widetilde{h}_i^{(t)} = \mathrm{LN}(h_i^{(t)})$ denote the layer-normalised input. For an edge $(j \to i)$, the message is computed by a two-layer MLP with internal batch normalisation (MsgMLP) taking the concatenation of $\widetilde{h}_i^{(t)}$ and the edge-modulated neighbour representation $\widetilde{h}_j^{(t)} + \mathrm{EdgeLin}(e_{ij})$. The messages are aggregated by mean pooling over N(i). The update step concatenates $\widetilde{h}_i^{(t)}$ with the aggregated messages and passes them through a second two-layer MLP with internal batch normalisation (UpdateMLP). Layer normalisation LN is applied to the UpdateMLP output before adding the residual. The layer update is given by Equation (4):

$$h_i^{(t+1)} = h_i^{(t)} + \alpha \cdot ln\left(UpdateMLP\left(\tilde{h}_i^{(t)}, \frac{1}{|N(i)|} \cdot \sum_{j \in N(i)} MsgMLP\left(\tilde{h}_i^{(t)}, \tilde{h}_j^{(t)} + EdgeLin(e_{ij})\right)\right)\right) \quad (4)$$

where $\widetilde{h}_i^{(t)} = \mathrm{LN}(h_i^{(t)})$ is the pre-normalised representation, LN is applied to the UpdateMLP output only (not the residual sum), N(i) is the neighbour set, and $\alpha = 0.3$. Eight such layers are stacked. This is consistent with the kernel integration view of GNO [4,5].

4.5. Field Aggregation Module

After the message passing stack, a graph-level field aggregation step is applied before the output projection. This module uses a learnable set of six field tokens, one per predicted output field, to exchange information across the spatial domain. In the first stage, each field token attends to the full set of node representations via multi-head attention, aggregating global information about each field from all nodes. In the second stage, the updated field tokens are used as keys and values in a second attention step in which the node representations serve as queries, distributing the aggregated field context back to each node. Both attention steps are followed by feed-forward sub-layers with layer normalisation. This mechanism is included as a practical implementation choice for multi-output operator learning and is not presented as an architectural contribution. Its individual effect on accuracy has not been quantified in this study and is identified as a topic for future ablation study.

4.6. Output Projection

The output projection maps the 192-dimensional hidden node representation to the six output values. It is implemented as a single four-layer MLP with skip connection. The main path consists of linear layers with layer normalisation, GELU activation, and dropout of 0.1, with the intermediate width reduced to 96 before the final linear layer producing the six outputs. A linear skip connection from the input to the output is added with a weight of 0.3. The six outputs correspond to the normalised values of T, $|v|$, p, CO, O_2, and CO_2, respectively.

4.7. Hyperparameter Summary

The complete set of hyperparameters of the baseline GNO surrogate is summarised in Table 1.

Table 1. Hyperparameters of the baseline GNO surrogate.

Value	Hyperparameter
192	Model width (hidden dimension)
8	Number of GNO message passing layers
8	Edge attribute dimension
4 (P/P_0, λ, x_{norm}, y_{norm})	Node input dimension
6 (T, $\|v\|$, p, CO, O_2, CO_2)	Node output dimension
15	k-NN neighbours for edge construction
Mean pooling	Message aggregation
None (raw normalised coordinates)	Positional encoding
1 (single shared MLP)	Output projection heads
0.3	Residual scale α
0.1	Dropout rate
approx. 4.5×10^6	Total trainable parameters

5. Training and Evaluation Procedure

5.1. Loss Function

The model was trained to minimise a weighted sum of per-field losses. For each field, the base loss was the mean squared error (MSE) between predicted and reference normalised values. The weight assigned to each field was chosen to reflect its engineering importance and typical dynamic range after normalisation. The weights were: temperature 800, pressure 500, CO 3000, O_2 2000, velocity magnitude 100, CO_2 400. These field weights were held constant throughout training and were normalised to unit sum before forming the total loss.

For the CO and O_2 fields, an additional asymmetric penalty was included to control the operationally critical error directions. Let $\hat{y}$ and y denote predicted and reference normalised values and let τ_f be a field-specific threshold in normalised units ($\tau\_CO = -0.95$, $\tau\_O_2 = -0.80$, both close to the lower bound of the normalised range). For nodes where the reference value lies at or below the threshold ($y \leq \tau\_f$), predictions exceeding the threshold are penalised by $w\_fa \cdot \max(0, \hat{y} - \tau\_f)^2$, suppressing false indications of elevated concentration in regions where the reference field is near its minimum. For the remaining nodes ($y > \tau\_f$), predictions falling below the threshold are penalised by $w\_miss \cdot \max(0, \tau\_f - \hat{y})^2$. The penalty factors were $w\_fa = 5$ and $w\_miss = 2$ for CO, and $w\_fa = 2$ and $w\_miss = 1$ for O_2. The penalty terms were ramped linearly from 10% to 100% of their nominal magnitude over the first 20 training epochs and held constant thereafter, so that the penalties do not dominate the optimisation before the bulk MSE has decreased. The per-field loss is the mean over nodes of the squared error plus the penalty terms, and the total loss is the weighted sum of the per-field losses as described above.

5.2. Optimisation

The model was optimised using the AdamW algorithm with a learning rate of 2×10^{-4} and a weight decay of 1×10^{-4}. A cosine annealing learning rate schedule was applied over 200 epochs, reducing the learning rate from the initial value to a minimum of 1×10^{-6}. A linear warm-up phase was applied over the first 10 epochs, increasing the learning rate linearly from zero to the initial value. Gradient norms were clipped to 1.0. The batch size was two graphs per step.

Early stopping with a patience of 50 epochs was applied based on the validation loss. In this implementation, the model weights from the final training epoch were used for all reported evaluations. The best-validation checkpoint was saved to a disk during training but was not reloaded prior to evaluation; the reported metrics therefore correspond to the final-epoch model, not the minimum-loss checkpoint.

5.3. Evaluation Metrics

Model performance was evaluated on the held-out validation set of 16 CFD cases using three scalar metrics computed for each output field separately: the coefficient of determination R^2, the root mean squared error (RMSE), and the mean absolute error (MAE). RMSE and MAE are reported in physical units after inverse min–max normalisation of the model output. All metric values were computed by concatenating node-level predictions and targets across all 16 validation graphs before computing aggregate statistics.

6. Results

6.1. Quantitative Validation Metrics

Table 2 presents the R^2, RMSE, and MAE values for each of the six predicted fields on the held-out validation set. All metric values are computed in physical units after denormalisation. The R^2 values for all six fields exceed 0.91, indicating that the surrogate accounts for more than 91% of the spatial variance of each field. The highest accuracy is obtained for the CO mole fraction (R^2 = 0.999), followed by O_2 (R^2 = 0.992), temperature (R^2 = 0.988), CO_2 (R^2 = 0.985), pressure (R^2 = 0.982), and velocity magnitude (R^2 = 0.919). These metrics are derived from a single fixed random split of 16 validation cases (seed = 42); no cross-validation was performed. The relatively lower accuracy for velocity magnitude is discussed on physical grounds in Section 7. Generalisation uncertainty arising from the small validation set size is addressed in Section 8.

Table 2. Validation metrics of the baseline GNO surrogate on the held-out validation set (16 cases). RMSE and MAE are in physical units after inverse min–max normalisation.

Field	Units	MSE	MAE	RMSE	R^2
Temperature	K	770.9	18.96	27.77	0.988
Velocity magnitude	m/s	0.0543	0.162	0.233	0.919
Pressure	Pa	0.0286	0.11	0.169	0.982
Mole fraction CO	[–]	2.97×10^{-6}	5.19×10^{-4}	1.72×10^{-3}	0.999
Mole fraction O_2	[–]	1.34×10^{-5}	2.27×10^{-3}	3.66×10^{-3}	0.992
Mole fraction CO_2	[–]	3.24×10^{-5}	3.89×10^{-3}	5.69×10^{-3}	0.985

6.2. Temperature Field

The temperature field achieved R^2 = 0.988, RMSE = 27.77 K, and MAE = 18.96 K on the validation set (Figure 4). Given that the temperature range in the combustion chamber spans from approximately 340 K at the inlets to over 1650 K near the flame front, the RMSE of approximately 28 K corresponds to a mean relative error of less than 2% of the total dynamic range. The model correctly reproduces the high-temperature combustion zone near the retort burner outlet, the temperature gradients along the convective channels, and the cooling effect of the water-jacketed walls.

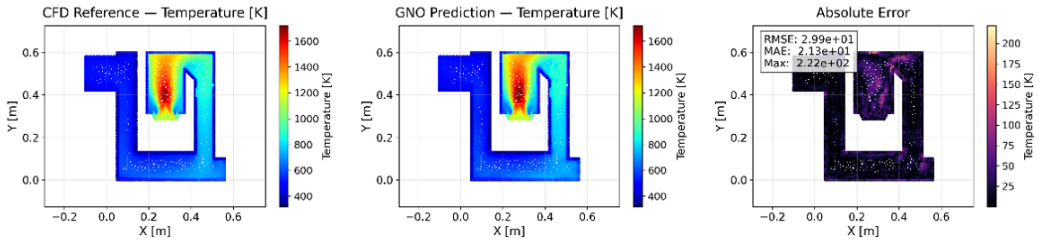

Figure 4. Comparison of the temperature field T [K]: CFD reference (**left**), GNO surrogate prediction (**centre**), and pointwise absolute error $|T_{pred} - T_{true}|$ (**right**) for a representative validation case (P/P_0 = 1.0, λ = 2.0). The error statistics displayed within the figure refer to this individual case; aggregate validation metrics are reported in Table 2.

6.3. Velocity Magnitude Field

The velocity magnitude field showed R^2 = 0.919, RMSE = 0.233 m/s, and MAE = 0.162 m/s (Figure 5). The velocity range in the domain spans from zero at solid walls to approximately 7 m/s in the highest-velocity jet regions. The relatively lower R^2 for velocity is consistent with the behaviour observed in other data-driven flow field surrogates and is discussed in Section 7.

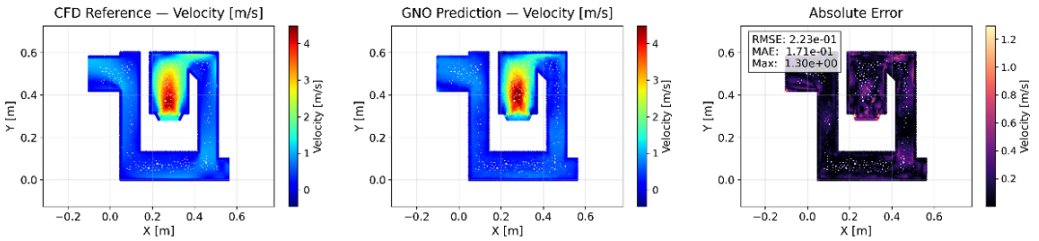

Figure 5. Comparison of the velocity magnitude field $|v|$ [m/s]: CFD reference (**left**), GNO prediction (**center**), and absolute error (**right**) for the same representative validation case. Case-specific error statistics are shown within the figure; aggregate validation metrics are reported in Table 2.

6.4. Pressure Field

The pressure field achieved R^2 = 0.982, RMSE = 0.169 Pa, and MAE = 0.110 Pa (Figure 6). The gauge static pressure in the computational domain varies over a small absolute range (approximately 11 Pa across the validation set, with gauge values between about −16 Pa and −5 Pa), so even small absolute prediction errors affect R^2. The model captures the overall pressure gradient driven by the fan pressure difference between the burner inlet and the outlet and reproduces the local pressure variations associated with flow through the FGR passages.

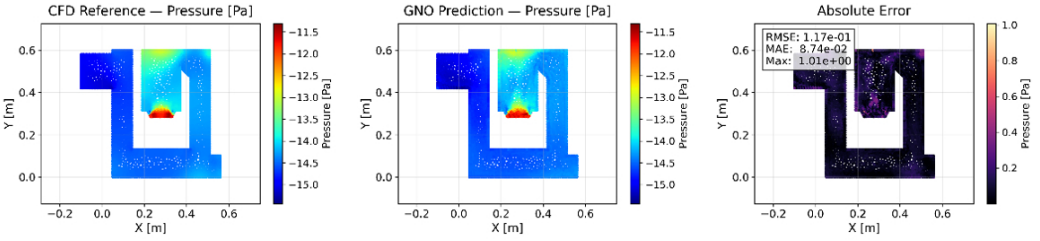

Figure 6. Comparison of the gauge static pressure field p [Pa]: CFD reference (**left**), GNO prediction (**center**), and absolute error (**right**) for the same representative validation case. Case-specific error statistics are shown within the figure; aggregate validation metrics are reported in Table 2.

6.5. CO Mole Fraction Field

The CO mole fraction field was predicted with the highest accuracy among all six fields (Figure 7) (R^2 = 0.999, RMSE = 1.72 × 10^{-3}, MAE = 5.19 × 10^{-4}). CO concentration

in the combustion chamber is highly spatially structured, with high values near the flame zone and near-zero values in the well-mixed post-flame region. The physical coupling between CO and temperature, which is strong in the combustion zone, is reflected in the similarly high accuracy for temperature.

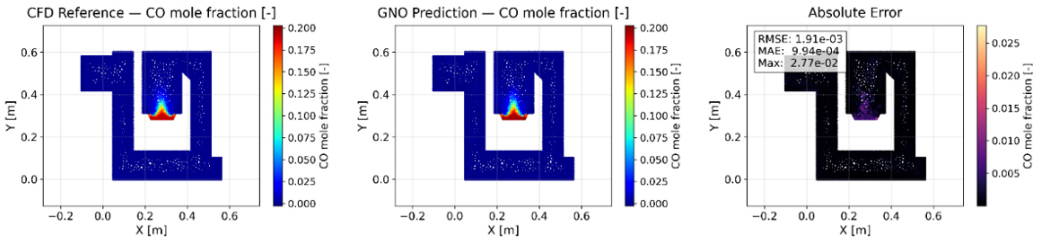

Figure 7. Comparison of the CO mole fraction field X^c_o [–]: CFD reference (**left**), GNO prediction (**center**), and absolute error (**right**) for the same representative validation case. Case-specific error statistics are shown within the figure; aggregate validation metrics are reported in Table 2.

6.6. O_2 and CO_2 Mole Fraction Fields

O_2 prediction yielded $R^2 = 0.992$ and CO_2 prediction yielded $R^2 = 0.985$. Both fields show smooth spatial distributions with most variation concentrated in the near-burner region (Figures 8 and 9) where mixing between primary combustion air, secondary air, and combustion products takes place. These fields are closely coupled to the temperature field through combustion stoichiometry and the excess air ratio λ, which is one of the two model input parameters.

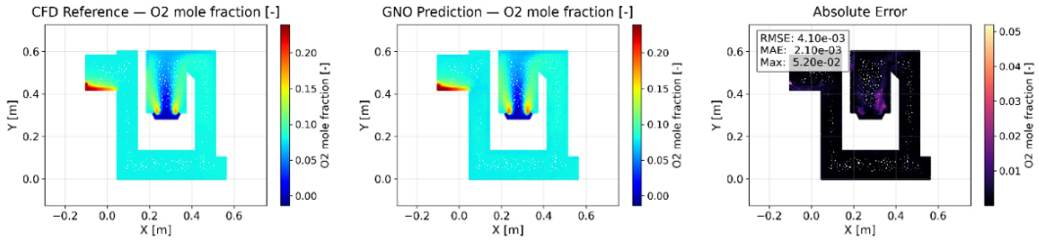

Figure 8. Comparison of the O_2 mole fraction field X_{o2} [–]: CFD reference (**left**), GNO prediction (**center**), and absolute error (**right**) for the same representative validation case. Case-specific error statistics are shown within the figure; aggregate validation metrics are reported in Table 2.

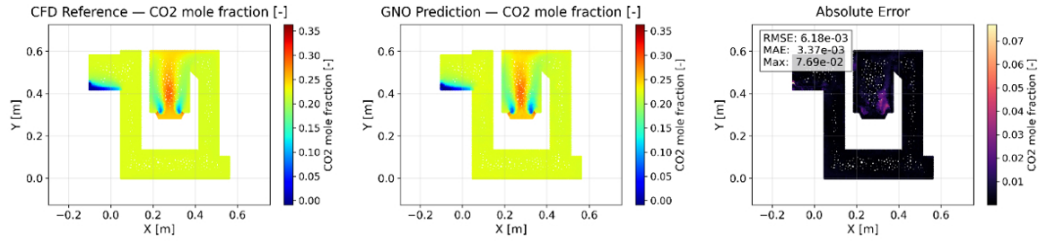

Figure 9. Comparison of the CO_2 mole fraction field X^c_{o2} [–]: CFD reference (**left**), GNO prediction (**center**), and absolute error (**right**) for the same representative validation case. Case-specific error statistics are shown within the figure; aggregate validation metrics are reported in Table 2.

6.7. Parity Plots

Figure 10 shows parity (scatter) plots of predicted versus reference values for all six fields on the full validation set. For every variable, the bulk of the point cloud is concentrated along the line of perfect agreement, with no systematic bias observed across the operating space. The largest scatter occurs for the velocity magnitude field, consistent with its lower R^2 value. For species concentration fields, the scatter plots show tight clustering along the ideal line, confirming the high accuracy reported in Table 2.

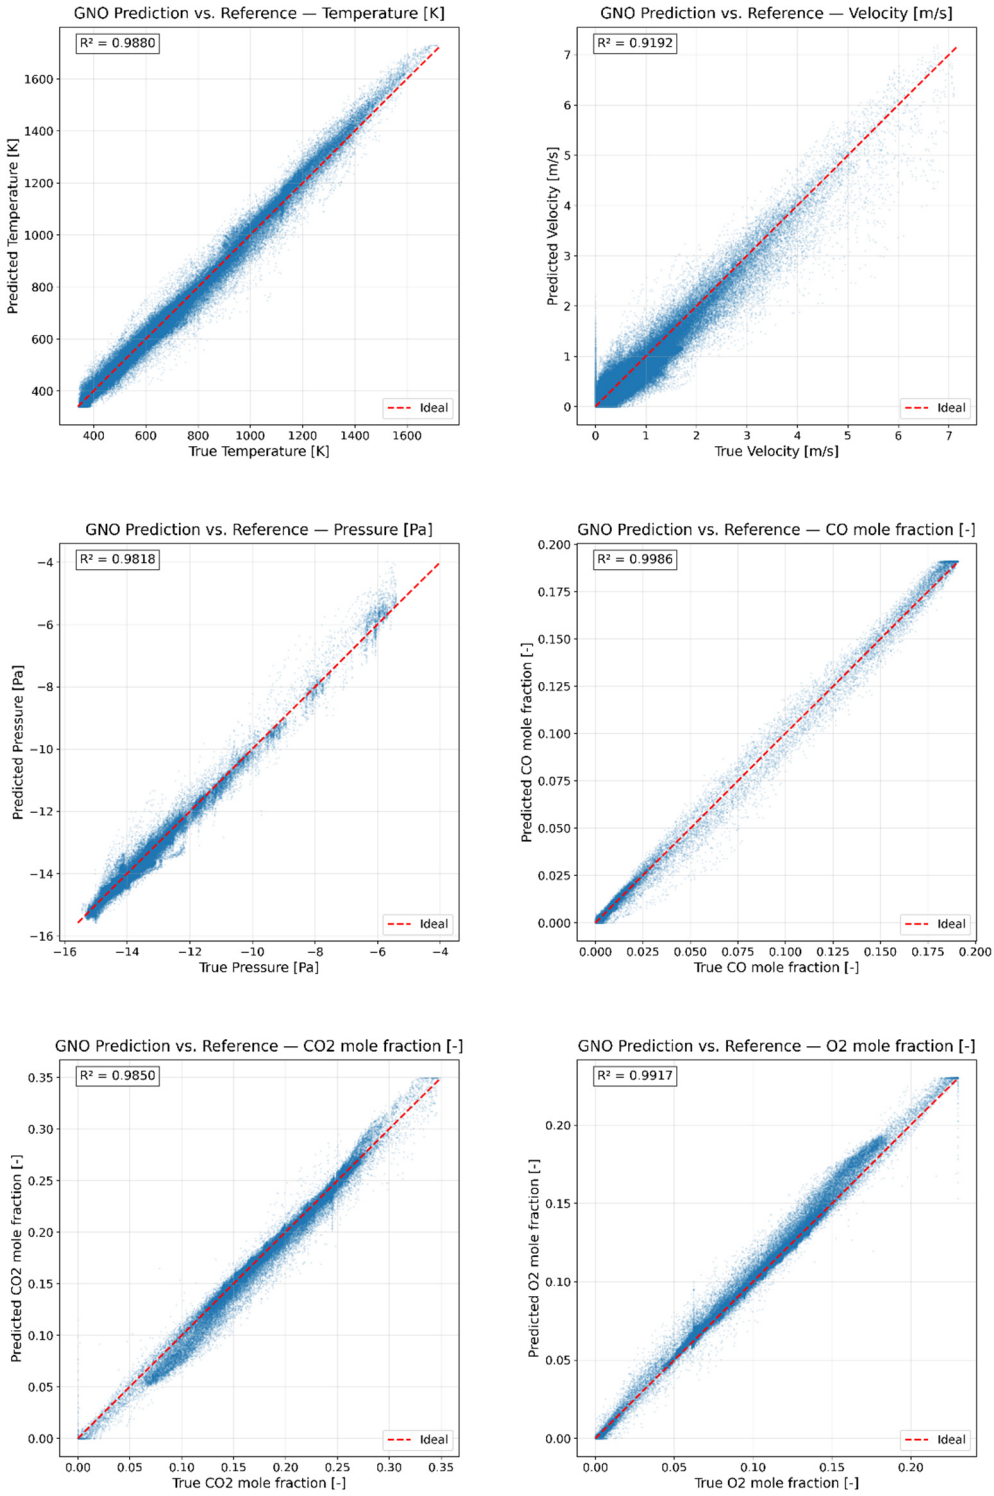

Figure 10. Parity plots of GNO surrogate predictions vs. CFD reference values for all six output fields on the full validation set (16 cases). The dashed red line indicates perfect agreement. R^2 values are indicated in each panel.

6.8. Training Convergence

Figure 11 shows the training and validation loss curves as a function of epoch number. Both losses decrease steadily over most of the schedule, and the validation loss reaches its minimum at approximately epoch 152. Because the validation loss continued to improve over most of the 200-epoch budget, the early-stopping criterion (patience of 50 epochs, Section 5.2) was not triggered and training ran to the full schedule. Between the minimum

and the final epoch, the validation loss remains essentially flat, so the final-epoch weights used for all reported evaluations are practically equivalent to the best-validation checkpoint. The cosine annealing schedule maintains smooth convergence and avoids oscillations in the later training phase. No evidence of overfitting is observed within the training range.

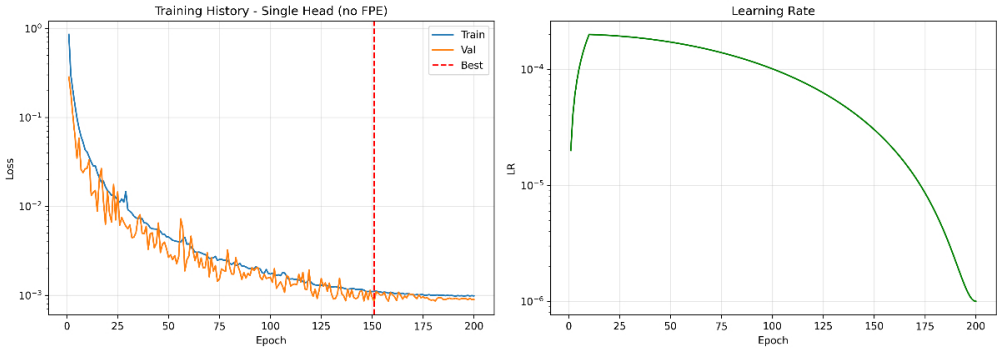

Figure 11. Training and validation loss curves for the GNO surrogate (**left panel**) and the cosine annealing learning rate schedule (**right panel**). Blue curve: training loss; orange curve: validation loss. The dashed red line indicates the epoch at minimum validation loss (saved as checkpoint, not used for evaluation; see Section 5.2).

6.9. Computational Efficiency

Each CFD simulation of the 10 kW boiler required approximately 20–30 min per operating point on a workstation equipped with a multi-core CPU and 32 GB RAM. After training, the surrogate model generates a prediction for a new operating point in a single forward pass. Each forward pass on a validation graph required less than 1 s of wall-clock time on an NVIDIA RTX 4500 GPU (24 GB VRAM). For deployments without a dedicated GPU, the same forward pass executed on a multi-core CPU required approximately 8–15 s per operating point, still representing a reduction of two orders of magnitude relative to the full CFD run. Offline training on 64 CFD cases required approximately 20–30 min on the same GPU and is a one-time cost. For parametric studies requiring, for example, 100 new operating points, the surrogate provides predictions in approximately 100 forward passes, compared to approximately 2000 min of CFD computation for the same set of cases.

6.10. Generalisation: Interpolation and Extrapolation Beyond the Training Grid

To probe the generalisation behaviour of the surrogate beyond the random validation split, an additional set of 14 CFD simulations was generated at operating points that were not part of the structured training grid. Eleven points were placed in the interior of the parameter domain ($P/P_0 \in [0.735, 1.225]$, $\lambda \in [1.375, 2.563]$) to assess interpolation, and three points were placed beyond the upper boundary of the training envelope (P/P_0 up to 1.47 and λ up to 3.175) to assess extrapolation towards a higher thermal load and leaner combustion. For each new operating point, an independent ANSYS Fluent 2025 R2 simulation was carried out on the same numerical mesh, and the GNO prediction was compared node-by-node against the CFD reference ($N \approx 12{,}930$ nodes per case). The very low load regime ($P/P_0 < 0.70$ combined with $\lambda < 1.25$) corresponds to a deep turn-down of the boiler towards flame extinction and lies outside the intended steady-operation envelope of the surrogate; it is therefore excluded from the present evaluation and discussed in Section 8.

The generalisation metrics for both the interpolation and extrapola-tion regimes are summarised in Table 3. Within the interior of the parameter domain, the surrogate reproduces all six fields with an accuracy that is statistically indistinguishable from the held-out validation set of Table 2 (for example, $R^2 = 0.989$ for temperature and $R^2 = 0.999$ for CO),

confirming that the random-split metrics are representative of genuine interpolation rather than memorisation of the training grid. Under extrapolation towards a higher thermal load and leaner mixtures the accuracy degrades only moderately and monotonically: the temperature, pressure and species fields retain $R^2 \geq 0.96$, whereas the velocity magnitude field is the most sensitive, its mean R^2 decreasing from 0.928 to 0.866. The spatial structure of all fields remains correctly reproduced even at the most extreme retained point ($P/P_0 = 1.47$, $\lambda = 3.00$), with prediction errors primarily concentrated in the core reaction zone and the thin flame-front region, as shown in Figure 12.

Table 3. Generalisation metrics of the GNO surrogate evaluated against independent CFD simulations at off-grid operating points: interpolation (11 interior points) and extrapolation towards higher P/P_0 and λ (3 points). Values are means over the points in each regime; RMSE is reported in physical units.

Field	Units	RMSE (Ext.)	RMSE (Int.)	R^2 (Ext.)	R^2 (Int.)
Temperature	K	32.4	25.7	0.979	0.989
Velocity magnitude	m/s	0.53	0.201	0.866	0.928
Pressure	Pa	0.569	0.129	0.966	0.97
Mole fraction CO	[–]	2.7×10^{-3}	1.6×10^{-3}	0.996	0.999
Mole fraction O_2	[–]	4.3×10^{-3}	3.4×10^{-3}	0.987	0.988
Mole fraction CO_2	[–]	7.7×10^{-3}	6.1×10^{-3}	0.964	0.966

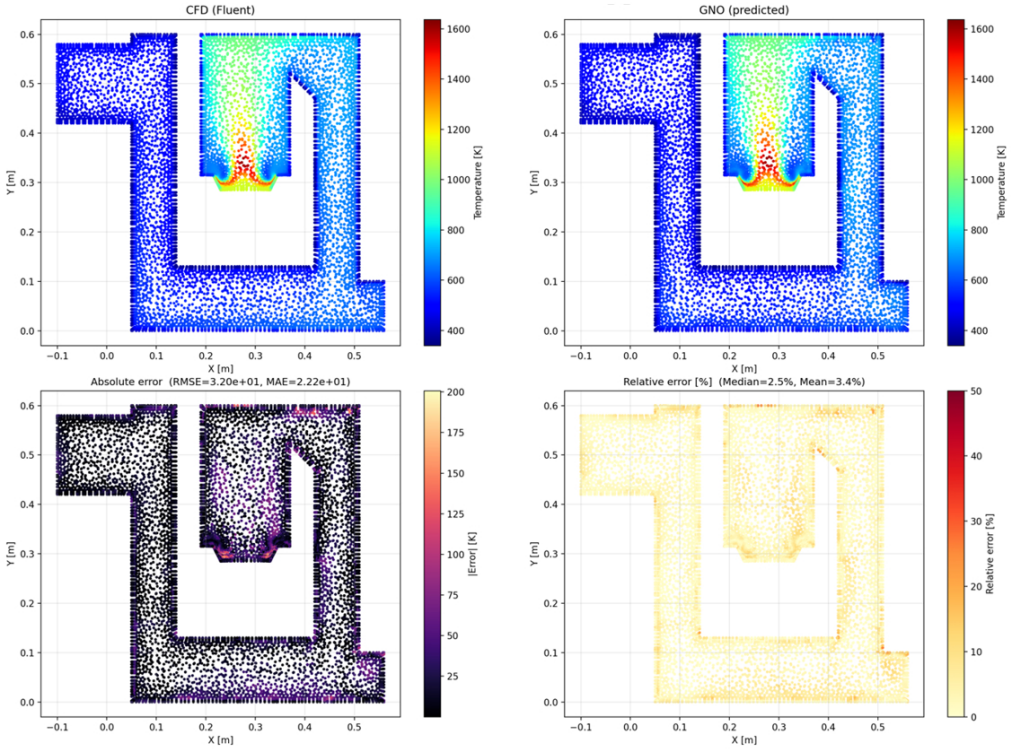

Figure 12. GNO surrogate prediction versus CFD reference for an extrapolation point above the training envelope ($P/P_0 = 1.47$, $\lambda = 3.00$): temperature field. The layout shows the CFD reference (**top-left**), the GNO prediction (**top-right**), the absolute error (**bottom-left**), and the relative error (**bottom-right**). Prediction errors are primarily concentrated in the core reaction zone and the flame-front region, with a median relative temperature error of 2.5%.

7. Discussion

The results in Section 6 show that a baseline GNO surrogate trained on 64 CFD cases of the 10 kW biomass boiler with FGR can reproduce the spatial distributions of six coupled

physical fields with R^2 values between 0.919 and 0.999 on held-out validation cases. This demonstrates the feasibility of using a GNO-based neural surrogate for multi-field CFD prediction in a realistic combustion engineering application.

The accuracy obtained for the temperature field (R^2 = 0.988, RMSE = 27.77 K) is relevant from a combustion engineering perspective. Temperature governs combustion kinetics, thermal NO_x formation, and the heat load distribution in the boiler. An RMSE of approximately 28 K over a field range exceeding 1300 K corresponds to a normalised error of about 2%, which is acceptable for rapid parametric screening and digital-twin monitoring applications.

The relatively lower accuracy of the velocity magnitude field (R^2 = 0.919) is physically explicable. The velocity field in a combustion chamber with FGR contains complex features including a high-velocity central jet, recirculation zones, and secondary flows driven by the convective channels. Reproducing these features from only two input parameters and spatial coordinates is a demanding task for any data-driven model. The result is nonetheless consistent with what has been reported for data-driven flow-field surrogates in comparable studies [3,9].

The high accuracy for the CO mole fraction (R^2 = 0.999) may be related to the tight coupling between CO concentration and the input parameter λ, which directly controls the air-to-fuel ratio. The CO distribution is strongly correlated with the temperature field and the excess air condition, both of which the model captures from the spatial graph and the scalar input parameters. Additionally, the asymmetric loss penalty for CO applied during training discouraged both spurious indications of elevated CO in low-concentration regions and the suppression of genuinely non-zero CO values towards the lower bound of the normalised range.

The GNO framework is well matched to the problem structure. CFD data on an unstructured mesh is naturally represented as a graph, with nodes corresponding to cell centroids and edges encoding spatial proximity. The GNO message passing operation approximates the kernel integration of the neural operator on this graph, allowing the learned mapping to be applied to the CFD point cloud without requiring a structured grid. This property is particularly valuable for combustion CFD data, where the mesh is refined in the flame zone and near wall boundaries, resulting in a highly non-uniform point distribution.

The practical utility of the surrogate extends beyond prediction accuracy. Once trained, the model generates field predictions for any combination of (P/P_0, λ) within the trained parameter range in a single forward pass. This makes it suitable for rapid parametric scans, optimisation studies with an approximate objective function, and supporting digital-twin workflows. For example, given measured values of P and λ, the surrogate can provide an estimate of the full spatial temperature and species distribution inside the combustion chamber, which would otherwise require either a full CFD run or an invasive in situ measurement campaign.

The generalisation analysis in Section 6.10 provides a quantitative validity envelope for the surrogate, which is essential for the digital-twin applications outlined above. Two features are noteworthy. First, the accuracy under interpolation matches that of the held-out validation set, which confirms that the model has learned the underlying parameter-to-field mapping rather than the discrete training grid. Second, extrapolation is markedly asymmetric: degradation towards a higher thermal load and leaner mixtures is gradual and field-dependent, with the velocity magnitude being the most affected, whereas the composition fields governed by stoichiometry (CO, O_2) remain robust because λ encodes the air-to-fuel ratio directly. This behaviour is physically consistent with the structure of the problem, in which the scalar transport is more directly constrained by the input

parameters than the momentum field. The practical implication is that the surrogate can be queried with confidence within and modestly beyond the upper training boundary, while predictions near the lower load limit should be treated with caution, as discussed in Section 8.

Beyond its role as a computational accelerator, the surrogate addresses functions that bear directly on the operation and design of energy-conversion equipment. Its two inputs, the thermal-load ratio and the excess-air ratio, are the operational levers that govern combustion completeness and thermal efficiency in a solid-biomass boiler. Because the model returns the full temperature and species fields, including the CO distribution that directly indicates unburned-fuel loss and hence conversion efficiency, it enables an efficiency- and emission-aware search of the operating envelope that would be prohibitively expensive with the full CFD solver. The surrogate is therefore a tool for optimising the energy-conversion process itself, not merely for reproducing individual flow solutions.

The investigated operating range deliberately spans part-load and modulating operation, the regime in which small domestic boilers spend most of their service life and in which both efficiency and emissions typically deteriorate. Fast prediction across this range supports load-following control and the integration of biomass heat generators into hybrid and flexible heating configurations, alongside heat pumps, thermal storage, or district-heating networks, where operating points must be re-evaluated frequently in response to demand and to the availability of other energy carriers. In this sense the surrogate contributes to the dispatch and system-level coordination of distributed renewable heat, and not only to component-level combustion analysis.

The short inference time makes real-time, physics-based supervisory functions realistic. Given only the two measured operating quantities, the model reconstructs the internal thermo-chemical state of the combustion chamber that would otherwise require either a full CFD run or an invasive measurement campaign. It can thus act as a soft sensor within an energy-management system or a digital twin of the boiler, enabling condition monitoring, model-predictive control, and early detection of off-design operation. This is the concrete mechanism by which neural-operator surrogates can embed high-fidelity combustion physics into the operational layer of an energy system.

The approach is also relevant to the broader decarbonisation of the heating sector. Solid biomass remains a major source of renewable heat, and low-emission, high-efficiency operation of domestic biomass appliances is a stated objective of energy and ecodesign policy. By lowering the computational barrier to efficiency- and emission-aware design and control, the surrogate supports the deployment of cleaner biomass heat. Because the graph formulation is mesh- and geometry-agnostic, the same methodology transfers to other thermal energy-conversion systems, extending its potential contribution beyond the single reference unit studied here.

The results should be interpreted in the context of the assumptions and constraints described in Section 8. The surrogate is a data-driven approximation of the CFD solver output and not a physics solver. It does not enforce mass conservation, momentum balance, or chemical equilibrium constraints explicitly. Its predictions are reliable within the training parameter range and for the specific geometry studied.

8. Limitations and Future Work

Several limitations of the current study should be noted. First, the training dataset is limited to 80 CFD cases on a structured grid in the two-dimensional (P/P_0, λ) parameter space. The density of coverage may be insufficient to capture highly nonlinear operating regimes, particularly in the low-λ region where CO production increases sharply. Extending the dataset by targeted sampling using active learning or space-filling designs

would improve coverage. Second, the validation set consists of only 16 cases from a single random split (seed = 42). The reported R^2 and RMSE values carry unquantified sampling uncertainty—bootstrap confidence intervals or k-fold cross-validation over the full 80-case dataset would give a more robust estimate of generalisation performance and are recommended before operational deployment. The dedicated generalisation study reported in Section 6.10 partially addresses this point by demonstrating accurate interpolation and graceful degradation under upward extrapolation on independent CFD cases. It also defines the boundaries of applicability: the surrogate is validated for steady operation within and modestly above the training envelope, but not for the deep turn-down regime towards flame extinction ($P/P_0 < 0.70$ combined with $\lambda < 1.25$), where the collapse of the dynamic pressure signal and the change in flow regime make the surrogate unreliable for the momentum and pressure fields. Extending the training set to cover such transient near-extinction conditions, or coupling the surrogate with a regime classifier, is left for future work.

Third, the CFD data were extracted from the two-dimensional symmetry plane of the combustion chamber, not from the full three-dimensional domain. The 2D representation captures the main features of the flow but loses information about azimuthal secondary flows and three-dimensional phenomena in the FGR channels. Extension to the full 3D domain would require substantially larger graphs, increased VRAM, and a scalable graph partitioning strategy for mini-batch training. Geometry-informed neural operators have been proposed for precisely such large-scale three-dimensional PDE problems [14].

Fourth, the surrogate model was evaluated only on CFD data without comparison to spatially resolved experimental measurements. The CFD reference model was validated at the design operating point in [1], but the surrogate inherits all modelling assumptions and errors of the CFD model. An additional validation against in situ measurements such as thermocouple arrays or gas sampling would provide a more complete assessment.

Future work could address these limitations in several directions. A full ablation study comparing the baseline GNO with variants that include Fourier positional encoding based on random Fourier features [15], multiple field-specific projection heads, or physics-informed loss terms would clarify the contribution of each component. Uncertainty quantification using probabilistic extensions, such as dropout-based ensembles or Bayesian neural operators, would allow the model to report prediction confidence alongside field estimates. Transfer learning between boiler geometries and extension to parametric geometry variations would increase the range of applicability. Finally, integration of the surrogate into a closed-loop digital twin system, in which the surrogate is updated online as new CFD or measurement data become available, is a longer-term research direction. As the reference unit is a domestic boiler fired with standardised wood pellets and without on-line fuel diagnostics, fuel-quality variation (e.g., moisture content) and transient ash deposition were outside the present scope; the former could, however, be explored offline as an additional CFD input parameter in future work.

9. Conclusions

This paper presented an application of a baseline Graph Neural Operator surrogate for the simultaneous prediction of six coupled CFD field distributions in a 10 kW biomass boiler with internal flue gas recirculation. The training dataset consisted of 80 steady-state CFD simulations performed in ANSYS Fluent 2025 R2 across a defined range of thermal load ratio and excess air ratio. The spatial CFD data were represented as an unstructured graph with node features encoding the operating parameters and spatial coordinates, and edge attributes encoding the geometric relationships between neighbouring mesh nodes.

The trained surrogate achieved R^2 values of 0.988 for temperature, 0.919 for velocity magnitude, 0.982 for pressure, 0.999 for CO mole fraction, 0.992 for O_2 mole fraction, and 0.985 for CO_2 mole fraction on a held-out validation set of 16 operating cases. These results demonstrate that a baseline GNO architecture, without Fourier positional encoding or domain-specific multi-head projection, can capture the spatial structure of coupled thermo-fluid and species fields in a realistic combustion geometry to a level of accuracy relevant for parametric analysis. A dedicated generalisation study on independent off-grid CFD cases further showed that this accuracy is preserved under interpolation and degrades only moderately when the model is extrapolated towards a higher thermal load and leaner mixtures, while delineating the deep turn-down regime towards flame extinction as the boundary of reliable application.

The principal practical outcome is a surrogate model that generates full spatial field predictions in a single forward pass, replacing per-case CFD simulations that require approximately 20–30 min per case with inference time of less than 1 s on a GPU or approximately 8–15 s on a CPU. This speed ratio makes the surrogate suitable for rapid parametric scans and digital-twin-oriented workflows.

The study is positioned as an application-oriented feasibility demonstration rather than a final optimisation of the GNO architecture for this problem. The novelty lies in the application of the GNO framework to a realistic multi-field combustion CFD case. Future work should include a full ablation study of architectural components, extension to three-dimensional CFD data, experimental validation of field distributions, and evaluation of uncertainty quantification methods for the surrogate predictions.

Author Contributions: Conceptualisation, P.M.; Methodology, P.M., D.K. and S.P.; Software, P.M.; Formal analysis, P.M.; Validation, D.K. and S.P.; Writing—original draft preparation, P.M.; Writing—review and editing, D.K. and S.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The CFD simulation data, trained model weights, and source code generated during this study are not publicly available because they relate to a patented boiler design under Polish utility-model protection and are subject to confidentiality obligations.

Acknowledgments: The authors gratefully acknowledge the technical support of the Faculty of Mechanical Engineering, Casimir Pulaski Radom University, in providing computational resources for CFD data generation and neural-operator training.

Conflicts of Interest: Danuta Król was employed by the Astwa Sp. z o. o., Białystok, Poland. The remaining author declare that the research was conducted in the conducted in the absence of any commercial or financial relationships that could be construed as potential conflict of interest.

References

1. Motyl, P.; Król, D.; Poskrobko, S.; Juszczak, M. Numerical Modelling and Experimental Verification of the Low-Emission Biomass Combustion Process in a Domestic Boiler with Flue Gas Flow around the Combustion Chamber. *Energies* **2020**, *13*, 5837. [CrossRef]
2. Guo, X.; Li, W.; Iorio, F. Convolutional Neural Networks for Steady Flow Approximation. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*; Association for Computing Machinery: New York, NY, USA, 2016; pp. 481–490.
3. Bhatnagar, S.; Afshar, Y.; Pan, S.; Duraisamy, K.; Kaushik, S. Prediction of Aerodynamic Flow Fields Using Convolutional Neural Networks. *Comput. Mech.* **2019**, *64*, 525–545. [CrossRef]
4. Li, Z.; Kovachki, N.; Azizzadenesheli, K.; Liu, B.; Bhattacharya, K.; Stuart, A.; Anandkumar, A. Neural Operator: Graph Kernel Network for Partial Differential Equations. *arXiv* **2020**, arXiv:2003.03485. [CrossRef]

5. Kovachki, N.B.; Li, Z.; Liu, B.; Azizzadenesheli, K.; Bhattacharya, K.; Stuart, A.M.; Anandkumar, A. Neural Operator: Learning Maps Between Function Spaces With Applications to PDEs. *J. Mach. Learn. Res.* **2023**, *24*, 1–97.
6. Li, Z.; Kovachki, N.; Azizzadenesheli, K.; Liu, B.; Bhattacharya, K.; Stuart, A.; Anandkumar, A. Fourier Neural Operator for Parametric Partial Differential Equations. *arXiv* **2021**, arXiv:2010.08895. [CrossRef]
7. Azizzadenesheli, K.; Kovachki, N.; Li, Z.; Liu-Schiaffini, M.; Kossaifi, J.; Anandkumar, A. Neural Operators for Accelerating Scientific Simulations and Design. *Nat. Rev. Phys.* **2024**, *6*, 320–328. [CrossRef]
8. Wu, S.; Liang, W.; Luo, K.H. Resolution-Invariant Temperature and Species Field Reconstruction Using Factorised Fourier Neural Operator for Turbulent Jet Flames. *Combust. Flame* **2025**, *283*, 114591. [CrossRef]
9. Chen, X.; Zhong, W.; Li, T. Fast Prediction of Temperature and Chemical Species Distributions in Pulverised Coal Boiler Using POD Reduced-Order Modelling for CFD. *Energy* **2023**, *276*, 127663. [CrossRef]
10. Wang, T.; Zhong, W.; Chen, X.; Zhou, G.; Shi, J.; Zhang, B. A Digital Twin of Coal-Fired Boiler for Physical Fields Prediction Using Flow-Field-Informed POD-XGBoost Reduced-Order Modelling. *Energy* **2025**, *329*, 136732. [CrossRef]
11. Sayghe, A.; Mousa, M.; Batiyah, S.; Husawi, A. Fourier Neural Operators for Fast Multi-Physics Sensor Response Prediction. *Sensors* **2026**, *26*, 1165. [CrossRef] [PubMed]
12. Yin, C. A Novel Accurate Model for Tracking Irregular Particles: Development, Implementation, and Impact on Biomass Combustors. *Bioresour. Technol.* **2025**, *429*, 132519. [CrossRef] [PubMed]
13. Ryu, J.; Cho, N.; Hwang, H.J. A Neural Operator Unifying Graph Neural Networks and Point Transformers. *IEEE Access* **2025**, *13*, 124264–124277. [CrossRef]
14. Li, Z.; Kovachki, N.; Choy, C.; Li, B.; Kossaifi, J.; Otta, S.; Nabian, M.A.; Stadler, M.; Hundt, C.; Azizzadenesheli, K.; et al. Geometry-Informed Neural Operator for Large-Scale 3D PDEs. In *Advances in Neural Information Processing Systems (NeurIPS 2023)*; Curran Associates, Inc.: Red Hook, NY, USA, 2023; Volume 36. [CrossRef]
15. Tancik, M.; Srinivasan, P.P.; Mildenhall, B.; Fridovich-Keil, S.; Raghavan, N.; Singhal, U.; Ramamoorthi, R.; Barron, J.T.; Ng, R. Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*; Curran Associates, Inc.: Red Hook, NY, USA, 2020; pp. 7537–7547. [CrossRef]