Engineering 59 (2026) 197–214

Contents lists available at ScienceDirect

Engineering

journal homepage: www.elsevier.com/locate/eng

Research
Low Carbon Transformation for Conventional Energies—Article

Building Digital Twin for 3D Multi-Field Reconstruction and Optimization of Industrial-Scale Combustion Systems

Linzheng Wang *, Yaojun Li, Sili Deng *

Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

ARTICLE INFO

Article history:
Received 12 February 2025
Revised 28 July 2025
Accepted 19 August 2025
Available online 22 August 2025

Keywords:
Digital twin
Multi-field reconstruction
Multi-objective optimization
Biomass combustion
Machine learning

ABSTRACT

In pursuit of a low-carbon energy transition, biomass and other carbon-neutral fuels are increasingly utilized in modern combustion facilities. However, controlling these systems remains challenging due to their complex geometries, dynamic interactions, and diverse operating conditions. Data-driven digital twins have emerged as powerful tools for optimizing performance and minimizing emissions in industrial combustion systems. Their core functions include reconstructing multi-physical combustion fields and predicting and optimizing key performance metrics, such as efficiency and pollutant emissions. Despite recent advancements, existing approaches typically treat reconstruction and optimization as separate tasks, limiting their efficiency and scalability. Furthermore, developing digital twins for real-world industrial applications requires extensive high-fidelity data, which is often impractical to obtain. To address these limitations, we propose the multi-field reconstruction net (MFRNet) framework, which integrates dimension expansion, variable extension, and feature fusion techniques to enhance data efficiency and predictive accuracy. Using an industrial-scale biomass grate furnace as a case study, we construct a comprehensive dataset, consisting of 288 low-fidelity 2D cases (covering eight physical fields) and 48 high-fidelity 3D cases (covering eleven physical fields). The MFRNet achieves high-precision multi-field reconstruction under complex conditions while significantly reducing the reliance on costly 3D simulations. By leveraging intermediate features pre-trained during reconstruction, the model enhances scalar predictions, notably improving CO and NO emission accuracy, even with limited high-fidelity data. The trained model is then directly applied for multi-objective optimization under varying operating conditions, demonstrating robust predictive accuracy and reliable optimization guidance. This scalable and data-efficient digital twin framework is easily adapted for other combustion systems, offering an intelligent paradigm for active control, real-time optimization, and enhanced operational efficiency in modern combustion facilities.

1. Introduction

In context of low-carbon energy transitions, biomass and other carbon-neutral fuels are increasingly used in combustion facilities as alternatives to fossil fuels [1]. This shift is driven by the urgent global need to reduce greenhouse gas emissions and achieve carbon neutrality targets across power generation, heating, and industrial sectors [2]. However, these alternative fuels present unique challenges for combustion systems due to their diverse compositions and heterogeneous properties. As environmental regulations become more stringent and the demand for sustainable energy solutions increases, there is a pressing need to optimize combustion processes to minimize environmental impact while maintaining high performance [3–5].

Traditional combustion control strategies, originally designed for fossil fuels, struggle to maintain stable operation and optimal performance when handling such variable fuel characteristics. For many years, extensive experimental research and the development of simulation methods have laid the groundwork for our understanding and controlling of combustion systems [6]. With the increasing need to accommodate multiple fuel types and dynamic capacity requirements in a unified energy system [7], new operational challenges have emerged. The ability to achieve fast and accurate state reconstruction and operational adjustment has become particularly critical [8].

* Corresponding authors.
E-mail addresses: wanglz@mit.edu (L. Wang), silideng@mit.edu (S. Deng).

https://doi.org/10.1016/j.eng.2025.08.020

Digital twins are virtual replicas of physical systems, enabling real-time monitoring, predictive maintenance, and optimization by leveraging data from advanced modeling techniques or on-site monitoring [9–11]. Digital twins can help combustion facilities navigate many challenges, such as variable fuel properties that affect combustion stability, capacity fluctuations driven by dynamic energy demands, and operational variations in parameters such as air–fuel ratios. Through real-time reconstruction and predictive analytics, digital twins enable adaptive control strategies to optimize combustion processes across these diverse scenarios, ensuring facilities maintain both efficiency and satisfy emission standards [12].

Despite rapid progress, key technical challenges remain. The primary obstacle is acquiring the high-fidelity data required to ensure model accuracy [13–15]. Combustion processes are inherently complex due to high-dimensional condition spaces and the coupling of multiple physical fields [16]. In practice, the state of combustion equipment is influenced by a multitude of factors, such as operating capacity, fuel type, and air distribution strategy. These factors often have nonlinear and complex impacts on various aspects (e.g., flame distribution, combustion efficiency, and pollutant formation) [17]. Furthermore, the combustion process involves interactions of fluid flow, heat transfer, and chemical reactions. To understand the combustion state within the equipment, it is essential to monitor the spatiotemporal distribution of multiple physical quantities, including temperature, pressure, and concentrations of chemical species [18]. Generating sufficient data to capture these complexities is time consuming and resource intensive. Excessive acquisition and training costs compromise the practicality of a digital twin.

In this case, effective integration of heterogeneous data sources is essential to reduce dependency on large data sets while maintaining model fidelity [19]. Proven strategies include integrating low-dimensional and high-dimensional simulation data [20], combining low-resolution and high-resolution simulation data [21], and utilizing both simplified and detailed reaction mechanisms in simulations [22]. However, these approaches have yet to be effectively applied to system-level digital twins for industrial-scale combustion.

Meanwhile, system-level digital twins are expected to provide real-time feedback on the combustion state and adjust the operating mode accordingly. Therefore, the models should possess comprehensive functionalities, including field reconstruction, parametric optimization, and control [23–25]. Yet recent studies often treat reconstruction and optimization as separate tasks based on fundamentally different methods. Reconstruction generally use model reduction techniques or end-to-end training methods to build a surrogate solver [26–28]. Inputs typically include measurable information or operational parameters from the combustion system to obtain complete spatial distributions of specific physical quantities. For instance, Savarese et al. [27] leveraged sparse sensing techniques to construct digital twins of combustion systems. They utilized chemical reactor network models as soft sensors to reconstruct full-field distributions for selected variables. Zhang and Li [28] adopted clustering segmentation to enhance reduced-order models, which rapidly predicted the temperature field inside a coal-fired boiler.

Conversely, the primary strategy for optimization tasks involves establishing a correspondence between operational parameters (e.g., airflow positions, fuel types) and target outcomes (like efficiency and emission level). This setup facilitates the identification of optimal operating conditions in response to varying scenarios [29–32]. For example, Sungur et al. [31] optimized the design of a biomass pellet burner pot by altering the supply airflow position within a pellet stove. Their model, based on Gaussian process regression, determined the relationship between the input and output parameters of the stove, aiming to enhance efficiency. Xu et al. [32] proposed a data-driven digital twin for a circulating fluidized bed boiler. Their approach determines benchmarks for controllable variables by analyzing historical operational data, thereby conserving energy.

Consequently, existing digital twin frameworks tend to have restricted capabilities, lacking the integration necessary to address the multifaceted challenges of combustion processes. The separation of reconstruction and optimization limits the potential of digital twin systems to provide holistic solutions for combustion system management.

To address these challenges, we propose a machine learning framework, multi-field reconstruction net (MFRNet), tailored for building digital twins of complex combustion systems. Our workflow maximizes data efficiency, reducing reliance on extensive high-fidelity datasets without sacrificing accuracy. By leveraging multi-fidelity data fusion techniques, the model aims to achieve 3D multi-field reconstruction within high-dimensional condition spaces using minimal high-fidelity data. The reconstruction model facilitates the extraction of latent features through representation learning. These extracted features can be shared and expanded across multiple tasks, including scalar prediction and multi-objective optimization, thus enabling function extension and improving the versatility of the digital twin.

We demonstrate MFRNet on an industrial-scale biomass grate furnace, a system of particular relevance for carbon reduction and renewable energy production [33]. Our method is highly adaptable and can be extended to a wide range of combustion and reacting flow systems. The remainder of this paper is organized as follows. Section 2 details the computational methodology, including the overall framework, network architectures, and dimension- and variable-expansion strategies, as well as our data-generation procedure. Details can be found in Section S1 in Appendix A. Section 3 evaluates MFRNet's multi-field reconstruction and scalar-prediction performance (including NO_x variables), and finally presents multi-objective optimization toward high combustion efficiency and low NO_x emissions.

2. Methods

Developing digital twin models for combustion systems presents two key challenges. First, generating high-fidelity 3D field datasets is computationally expensive, thereby limits training data diversity, especially in industrial applications with numerous operating parameters. Second, conventional methods typically decouple reconstruction from optimization, relying on black-box models that lack physical insight. Our approach addresses these issues by leveraging low-fidelity data to efficiently explore the high-dimensional operating space and extract essential system features.

We propose a unified machine learning framework, MFRNet, which enables high-accuracy multi-field reconstruction, scalar prediction, and operational optimization while minimizing dependence on high-fidelity data. The framework is validated on an industrial-scale biomass grate furnace, where a digital twin is developed for multi-field reconstruction and multi-objective optimization across an 11-parameter condition space, including capacity, equivalence ratio, air distribution, and grate movement velocities. Although fuel properties, such as composition and kinetic parameters, are undeniably important, incorporating such variability would considerably expand the dimensionality of the simulation space and the associated computational cost. Thus, to demonstrate the viability of our digital twin framework, fuel-related parameters are held constant. Future studies will aim to explicitly integrate fuel diversity into the model development process.

The core model architecture consists of several key modules that enable its functionality, and the workflow is shown in Fig. 1. Firstly, to accommodate the 11-dimensional condition space, we construct a multi-fidelity dataset. Using Latin hypercube sampling (LHS), we select 288 cases, each representing a unique operating condition defined by a set of parameters. These cases are simulated using a highly simplified 2D numerical model to establish a low-fidelity dataset. Subsequently, we select 48 representative cases from this larger dataset to create a smaller 3D high-fidelity dataset. Details of case sampling are demonstrated in our previous work [20].

With this data foundation, we develop a digital twin system based on MFRNet. The core architecture is a neural network featuring a multi-modal input layer and a multi-channel dual-net structure, enabling end-to-end multi-field reconstruction using either operational parameters or sensor data. The training process begins with pre-training on the low-fidelity dataset to establish the initial 2D MFRNet. To overcome the inherent limitations of 2D simplifications in combustion system analysis, we finetune the pre-trained model on the small high-fidelity dataset using dimension expansion and variable extension modules. Dimension expansion refers to the process of extrapolating or inferring 3D spatial fields from lower-dimensional data, typically using pre-trained models on 2D slices to reconstruct full 3D distributions. In contrast, variable extension focuses on inferring additional physical fields not originally present in the dataset. For example, deriving pollutant fields such as NO_x from primary combustion variables through latent feature fusion. Specifically in the network structure, the dimension expansion module extrapolates z-direction heterogeneity, effectively capturing boundary effects and enhancing 3D system representation. The variable extension module addresses discrepancies between low- and high-fidelity simulations by leveraging physical field interdependency, ensuring accurate learning of new variables despite limited high-fidelity data availability. These processes yield a refined 3D MFRNet.

Our model framework culminates in a dynamic feature fusion mechanism that adapts latent representations extracted from multi-field reconstruction. This mechanism employs attention-based merging techniques, further improving the accuracy in downstream scalar prediction tasks, and ultimately accomplishing reconstruction and optimization towards complex conditions. By integrating these components into a unified learning framework, our approach significantly enhances the accuracy and efficiency of digital twin models for industrial combustion systems. The following sections detail the model architecture, data generation, preprocessing, and training configurations.

2.1. Multi-field reconstruction

2.1.1. Multi-modal inputs

A digital twin must represent the system's state. In real-world combustion systems, multiple sources of information, such as operational parameters from control panels, internal temperature sensor data, and visual imagery of flames, collectively reflect the system's combustion state. Effectively integrating these diverse data sources enables a comprehensive and adaptive response to varying operational conditions [34]. Critically, the ability to process multimodal input ensures that the digital twin can function effectively even with partial information. This capability is particularly valuable for tasks like intelligent fault diagnosis, where the system must identify and respond to anomalies by correlating data from multiple sources [35]. Therefore, multi-modal input serves as a crucial component in digital twin systems, enabling them to integrate diverse data sources and enhance their practical applicability effectively.

Various techniques can be applied to align and integrate data from disparate sources for modality matching and feature fusion [36]. In this work, we focus on two primary types of input: operational parameters ($\boldsymbol{U}$) and internal temperature measurements ($\boldsymbol{S}$) from a furnace. Specifically, we select the capacity ratio CR, equivalence ratio ER, distribution of primary air in each wind section α_i (i = 1, 2, ..., 5), and velocities of the grate on each grate section v_j (j = 1, 2, 3, and 4) as the parameters in the condition vector:

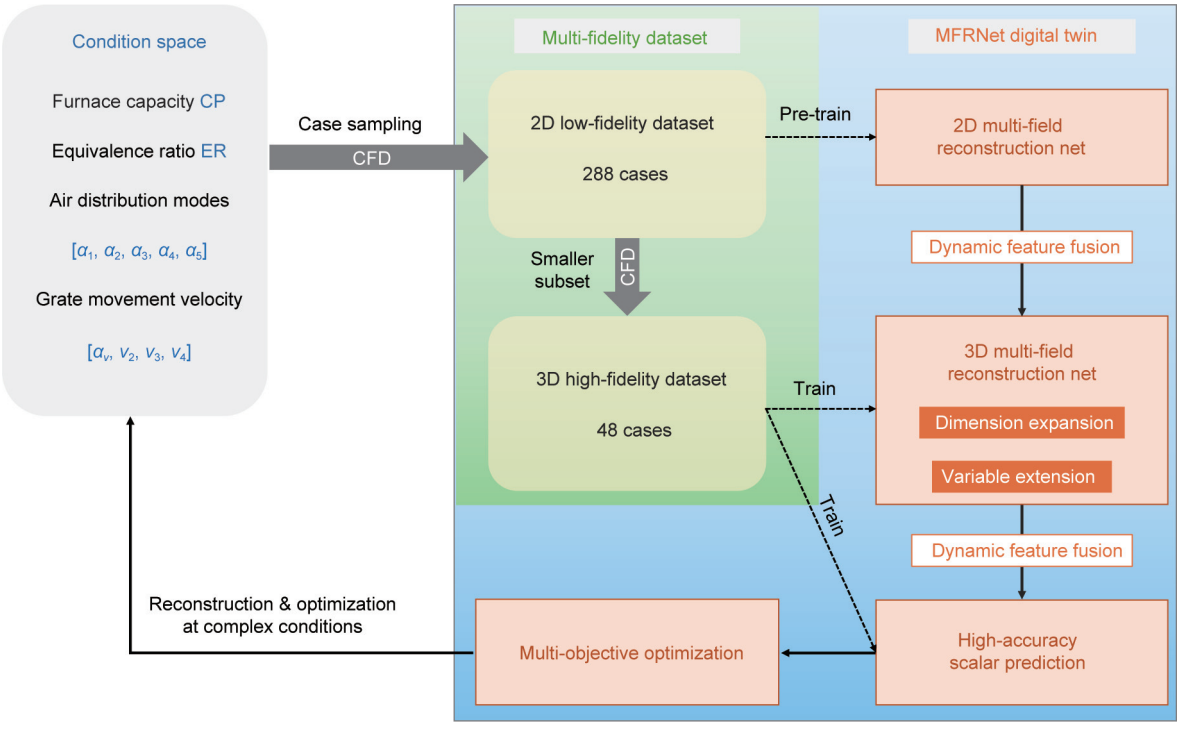

Fig. 1. The general workflow of building digital twin for multi-field reconstruction, scalar prediction and multi-objective optimization based on MFRNet.

$$\boldsymbol{U} = [\mathrm{CR}, \mathrm{ER}, \alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5, v_1, v_2, v_3, v_4]^{\mathrm{T}} \tag{1}$$

The simulated temperature data T are generated from 25 random positions $\{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{25}\}$:

$$\boldsymbol{S} = [T(\boldsymbol{x}_1), \ldots, T(\boldsymbol{x}_{25})] \tag{2}$$

These inputs are processed using contrastive learning techniques to achieve effective multi-modal integration, as shown in Fig. 2.

In the training process, the input vectors $\boldsymbol{U}$ and $\boldsymbol{S}$ are mapped into latent space using separate multilayer perceptrons (MLPs) (M_{U} and M_{S}) designed as projection heads.

$$M_{\mathrm{U}}(\boldsymbol{U}) = \boldsymbol{Y}_{\mathrm{U}} \tag{3}$$

$$M_{\mathrm{S}}(\boldsymbol{S}) = \boldsymbol{Y}_{\mathrm{S}} \tag{4}$$

where $\boldsymbol{Y}_{\mathrm{U}}$ and $\boldsymbol{Y}_{\mathrm{S}}$ are both designed to be 25-dimensional vectors. We define $\boldsymbol{Y}_{\mathrm{U}}$ and $\boldsymbol{Y}_{\mathrm{S}}$ as a positive pair if they correspond to the same condition $\boldsymbol{U}$, whereas all other combinations are considered negative pairs. We denote $\boldsymbol{y}_{\mathrm{u},i}$ as the latent embedding from $\boldsymbol{Y}_{\mathrm{U}}$, and $\boldsymbol{y}_{\mathrm{s},i}$ as the corresponding latent embedding from $\boldsymbol{Y}_{\mathrm{S}}$. Thus, for a positive pair, ($\boldsymbol{y}_{\mathrm{u},i}$, $\boldsymbol{y}_{\mathrm{s},i}$) are paired by the same operational condition, whereas for negative pairs, we consider combinations ($\boldsymbol{y}_{\mathrm{u},i}$, $\boldsymbol{y}_{\mathrm{s},j}$) with $i \neq j$. Let N_{pos} be the total number of positive pairs and N_{neg} be the total number of negative pairs in the training batch. Then, the contrastive loss L_{con} for a set of positive and negative pairs can be formulated as [37]:

$$L_{\mathrm{pos}} = \frac{1}{N_{\mathrm{pos}}} \sum_{i=1}^{N_{\mathrm{pos}}} \|y_{\mathrm{u},i} - y_{\mathrm{s},i}\|^2 \tag{5}$$

$$L_{\mathrm{neg}} = \frac{1}{N_{\mathrm{neg}}} \sum_{i=1}^{N_{\mathrm{neg}}} \max\left(0, d_{\mathrm{m}} - \|y_{\mathrm{u},i} - y_{\mathrm{s},j}\|^2\right) \tag{6}$$

$$L_{\mathrm{con}} = \omega_{\mathrm{pos}} \cdot L_{\mathrm{pos}} + \omega_{\mathrm{neg}} \cdot L_{\mathrm{neg}} \tag{7}$$

where ω_{pos} and ω_{neg} are weights applied to the positive and negative loss components, respectively. We set both to be 1.0 in our cases. This equal weighting was chosen to ensure a balanced contribution from both the positive and negative loss terms, particularly since each term is normalized by the number of corresponding pairs. This symmetric setting helps maintain model sensitivity by avoiding over-emphasis on either term, thus leading to more robust and valid outcomes during training. d_{m} is the learnable margin parameter that dictates how far apart the negative pairs should be.

2.1.2. The structure of neural network for reconstruction

The reconstruction neural network forms the core of the MFRNet framework. We design the reconstruction neural network to reconstruct multiple physical fields from multi-modal input data. The general structure is an extension of the dual net previously utilized in our work [20] and can be considered as a variant of the DeepONet architecture [38], as shown in Fig. 2.

Specifically, the reconstruction neural network is structured to include a set of branch net sub-channels ($\{\mathbf{BR}_i\}_{N_{\mathrm{f}}}$, $i = 1, \ldots, N_{\mathrm{f}}$), corresponding to N_{f} different physical fields. These fields include temperature T, pressure P, x-velocity v_x, y-velocity v_y, and the mass fractions of O_2, CO_2, H_2O, and CO. Biomass undergoes a devolatilization stage in which volatile compounds (e.g., CH_4 and similar hydrocarbons) are generated, but our computational fluid dynamics (CFD) analysis indicates that their concentrations are comparatively low, and these species are rapidly consumed near the fuel grate. This simplification enables us to focus on demonstrating the effectiveness of our reconstruction framework while maintaining computational tractability. Importantly, the modular design of MFRNet can be readily extended to include arbitrary volatile species of concern as long as the training CFD data is available.

A trunk net **TR** is employed to process spatial or coordinate information, which is shared across different physical fields. The branch net sub-channels receive a unified representation $\boldsymbol{R}$ that combines input condition vectors with temperature data points through a transformation applied by a projection head.

$$\boldsymbol{R} = 0.5 \times \boldsymbol{Y}_{\mathrm{U}} + 0.5 \times \boldsymbol{Y}_{\mathrm{S}} \tag{8}$$

Each physical field has its dedicated sub-channel, which outputs a m-dimensional vector representing the coefficients of basis functions for that field [38]. In this work we set $m = 50$.

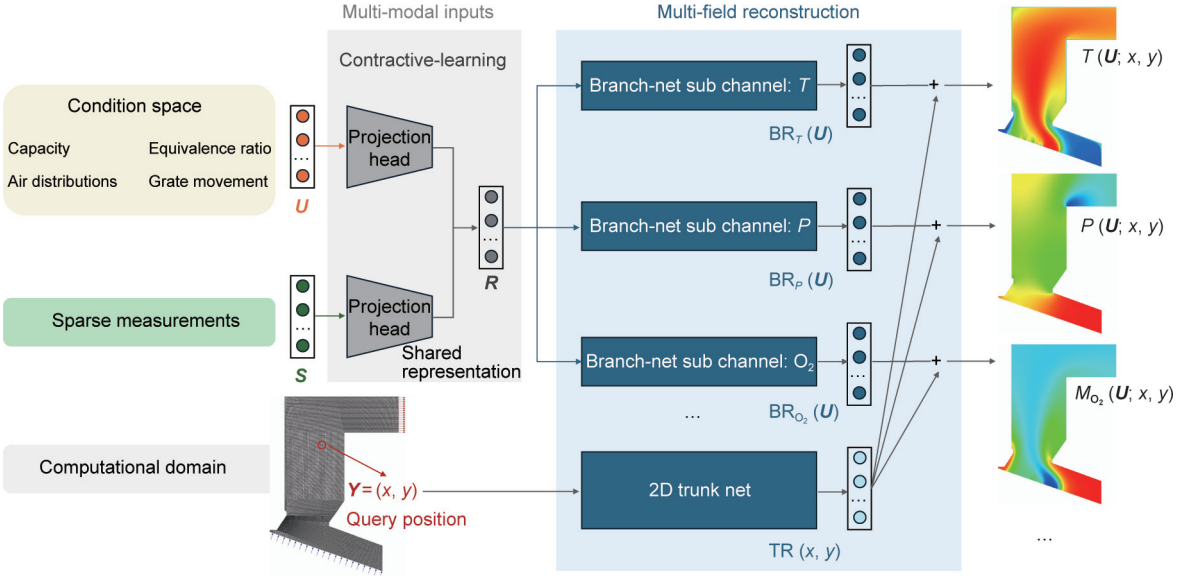

Fig. 2. The core neural network structure to achieve multi-field reconstruction. The model consists of a multi-modal input layer and a multi-channel dual-net architecture, enabling end-to-end full-field reconstruction based on either operational parameters or sensor data. $\boldsymbol{Y}$: the coordinate of a query position; $\boldsymbol{R}$: the unified representation from multi-modal inputs; **BR**: branch net; **TR**: trunk net. The predicted T, P, and M_{O_2} are temperature, pressure, and mass fraction of oxygen, respectively.

$$\boldsymbol{c}_i = \mathbf{BR}_i(\boldsymbol{R}) = \left[c_i^1, \ldots, c_i^m\right]^{\mathrm{T}}, i = 1, \ldots, N_{\mathrm{f}} \tag{9}$$

The trunk net accepts coordinates of the query position $\boldsymbol{Y} = (x, y)$ as input and produces a 50-dimensional output vector representing the values of the basis functions at the given position [38].

$$\boldsymbol{p} = \mathbf{TR}(\boldsymbol{Y}) = [p_1, \ldots, p_m]^{\mathrm{T}} \tag{10}$$

The final prediction for the ith physical field ϕ_i at a query position. $\boldsymbol{Y}$ is obtained by taking the dot product of the output from the corresponding branch net sub-channel and the trunk net [38].

$$\phi_i(\boldsymbol{Y}) = \boldsymbol{p} \cdot \boldsymbol{c}_i = \sum_{j=1}^{m} \left(p_j \cdot c_i^j\right) \tag{11}$$

We utilize fully connected neural network architecture to build the branch and trunk nets. Rectified linear unit (ReLU) activation functions are used, and a sigmoid activation function is applied to the final layer. This network is the foundation for subsequent dimension expansion and variable extension. Key hyperparameters such as layer counts, neuron numbers, and dropout rates have been tuned to optimize performance. Details are provided in Section S2 in Appendix A.

While the reconstruction neural network is theoretically capable of 3D field reconstruction, the primary limitation is the size of the dataset necessary for training. Building a comprehensive 3D dataset covering various operational conditions requires substantial computational resources. Thus, the reconstruction neural network is preliminarily trained using a 2D combustion dataset. This dataset provides a complex yet manageable scenario to test and refine the network architecture. The pre-trained model will be further extended through the implementation of dimension expansion and variable extension modules on a smaller 3D dataset.

2.1.3. Dimension expansion

Building upon the pre-trained 2D reconstruction net, we seamlessly fine-tune it to accommodate 3D datasets. The pre-trained 2D network is employed as the basis for dimension expansion into a 3D representation. The parameters of this pre-trained net are frozen during the extension phase to preserve the learned feature representations, allowing the subsequent training process to focus on adapting these features for 3D reconstruction and variable extension tasks. This extension is particularly significant for addressing the effects of wall-bounded flows and heat transfer, which inherently exhibit 3D characteristics.

Specifically, the goals of the dimensional expansion are: ① handling non-uniformity along the z-direction, and ② accuracy refinement in x- and y-directions. To achieve these objectives, we introduce the extended trunk net for z-direction $\mathbf{TR}_z$ and a set of accuracy refinement nets $\{\mathbf{AR}\}_{N_{\mathrm{f}}}$, as illustrated in Fig. 3. These components enable the flexible and accurate modeling of 3D physical fields while maintaining computational efficiency. For each combustion-related physical field ϕ_i, its value at a specific query position $\boldsymbol{Y} = (x, y, z)$ in the 3D computational domain is formulated as:

$$\phi_i(x, y, z) = (\boldsymbol{p} \cdot \boldsymbol{p}_z) \cdot \boldsymbol{c}_i \cdot \mathrm{Corr}_i(x, y, z; \boldsymbol{U}) \tag{12}$$

where $\boldsymbol{p}$ and $\boldsymbol{c}_i$ are generated by the pre-trained 2D MFRNet, and $\boldsymbol{p}_z$ is the output of $\mathbf{TR}_z(z)$. $\mathrm{Corr}_i(x, y, z; \boldsymbol{U})$ is the correction factor specific to the ith field, output by the auxiliary network $\mathbf{AR}_i$. This network also comprises a branch net and a trunk net. The detailed settings are given in Section S2.

2.1.4. Variable extension

The core concept of variable extension lies in the intrinsic correlations among the spatiotemporal distributions of different physical variables in a multi-physics coupled system. Because the pre-trained model has already extracted features for core combustion variables, it can accurately predict unseen variables from small datasets by reusing these features. In this study, we aim to extend the learning of combustion-related physical fields to include the 3D distribution of NO_x-related variables, namely, the concentrations of NO, HCN, and NH_3, using small datasets.

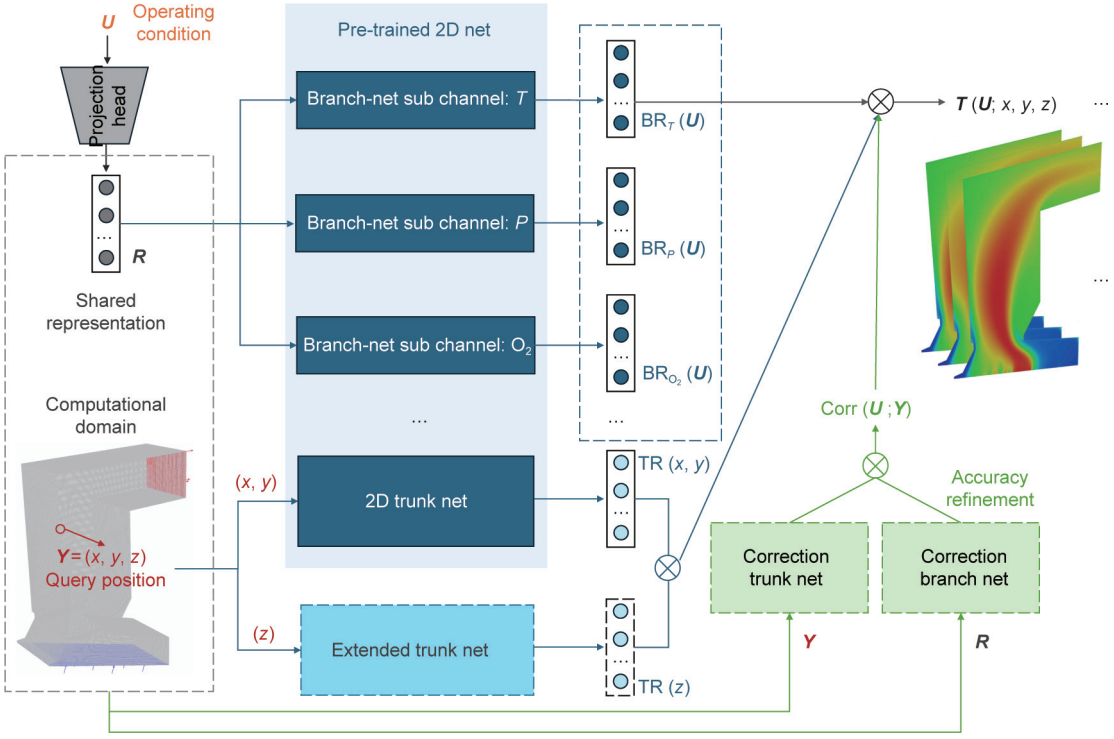

Fig. 3. The schematic diagram of the dimension expansion module. Based on the pre-trained reconstruction net on 2D (x and y) dataset, the additional trunk net on z-direction coordinates and accuracy refinement sub-net can be further trained for 3D reconstruction at complex conditions on the small high-fidelity dataset.

As illustrated in Fig. 4, the output of the branch networks from the pre-trained combustion model can be combined and processed using a set of attention-based adaptive merging layers $\{A_k\}$.

$$\boldsymbol{c}_k = A_k\left(\left[\mathbf{BR}_1(\boldsymbol{R}), \ldots, \mathbf{BR}_{N_f}(\boldsymbol{R})\right]\right) = A_k\left(\left[\boldsymbol{c}_1, \ldots, \boldsymbol{c}_{N_f}\right]\right),\ k = \mathrm{NO}, \mathrm{HCN}, \mathrm{NH_3} \tag{13}$$

These layers dynamically fuse the internal features of the pre-trained model through the self-attention mechanism, applying adaptive weighting to highlight the most relevant feature contributions for the target variables. Then, the additional branch nets for new variables will process the fused feature and predict the additional variables:

$$\phi_k(\boldsymbol{Y}) = (\boldsymbol{p} \cdot \boldsymbol{p}_z) \cdot \mathbf{BR}_k(\boldsymbol{c}_k) \tag{14}$$

In Section S2.1 in Appendix A, we demonstrated the details of the attention-based adaptive merging layers. This extension builds upon the pre-trained model, which has been trained on eight physical fields related to combustion, enabling accurate and efficient exploration of these additional variables under data-limited conditions.

2.2. Multi-objective optimization

2.2.1. Building response surface

The response surface model serves as a surrogate model for fitting specific input parameters to output responses, primarily employed to bypass computationally intensive tasks, such as high-fidelity CFD modeling [39]. For complex combustion systems, response surfaces enable rapid analysis of key design factors and operational parameters' effects on combustion efficiency and pollutant emissions, facilitating system control and optimization [40].

Various approaches exist for establishing response surfaces, ranging from fundamental methods like polynomials, kriging, and radial basis functions to neural networks, all of which have been previously applied across different research and application areas [41]. However, building a response surface is fundamentally a data-fitting problem. When dealing with multi-parameter conditions, if the high-fidelity dataset is insufficient, the response surface model may suffer from excessive prediction errors, rendering it impractical for application.

In this work, based on key scalars obtained from a limited 3D high-fidelity dataset, we employ neural networks to construct response surfaces targeting two critical output variables for prediction: the average mass fraction of CO at the boiler outlet M_{CO} and the normalized NO emission M_{NO} ($\mathrm{mg{\cdot}Nm^{-3}}$; standardized to a 6% oxygen concentration). We will demonstrate that building upon the pre-trained reconstruction model, the accuracy of the response surfaces using limited data can be largely improved.

Here, the input to the model is the conditional vector $\boldsymbol{U}$, and we explore two main pathways for prediction. The first pathway is direct prediction using a neural network. In this approach, we directly train a neural network M_{T_U} to predict the target values T_U based on the input vector.

$$M_{T_\mathrm{U}}(\boldsymbol{U}) = T_\mathrm{U} \tag{15}$$

The second pathway is feature extraction and fusion using pre-trained MFRNet. In this approach, the input vector $\boldsymbol{U}$ is first processed by the branch net modules of the pre-trained MFRNet, which yields intermediate features $\{\boldsymbol{c}_1, \ldots, \boldsymbol{c}_{N_f}\}$ used for reconstructing the N_f physical fields. These features are then fused using another attention-based feature fusion layer $\boldsymbol{A}_{T_\mathrm{U}}$ to obtain a combined feature $\boldsymbol{c}_{T_\mathrm{U}}$.

$$\boldsymbol{A}_{T_\mathrm{U}}\left(\left[\boldsymbol{c}_1, \ldots, \boldsymbol{c}_{N_f}\right]\right) = \boldsymbol{c}_{T_\mathrm{U}} \tag{16}$$

The fused feature $\boldsymbol{c}_{T_\mathrm{U}}$ is subsequently passed to an MLP M_{T_U} for target value prediction.

$$M_{T_\mathrm{U}}(\boldsymbol{c}_{T_\mathrm{U}}) = T_\mathrm{U} \tag{17}$$

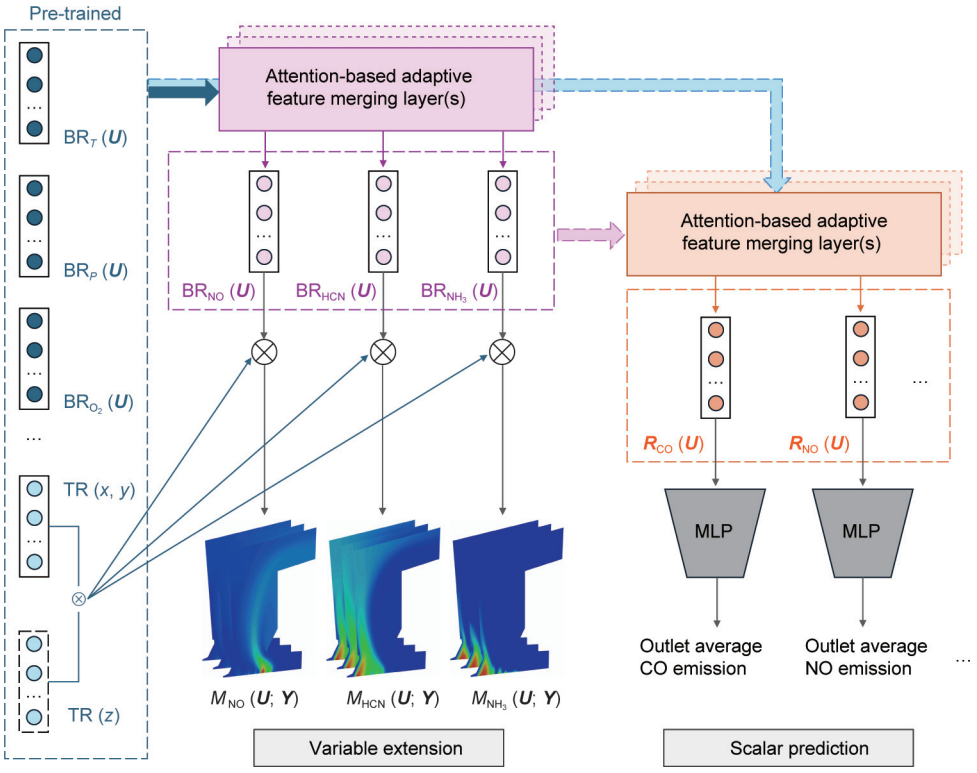

Fig. 4. The schematic diagram of variable extension and scalar prediction via attention-based feature merging layers.

More specifically, for CO prediction, the second pathway involves a two-step process. First, we establish a feature merge network and train the MLP on a 2D dataset to predict the target value from $\boldsymbol{U}$. Then, to further improve the prediction accuracy, an additional correction network is trained on a 3D dataset to refine the predicted values. For NO prediction, the intermediate features reconstructed from 3D multi-field data are directly integrated. The resulting fused features are used to train an MLP on the 3D dataset for target value prediction.

2.2.2. Pareto optimization

With response surfaces in place, we optimize combustion operation. As an illustrative example, we define a multi-objective optimization task to demonstrate the utility of the developed response surfaces. Given a fixed capacity, the task involves adjusting operational parameters, such as air distribution and grate movement, to maximize combustion efficiency while minimizing pollutant emissions.

This optimization task is highly relevant in real-world combustion systems, as it addresses the critical balance between efficiency and environmental impact [42]. In industrial settings, operators often seek to maximize combustion efficiency to reduce fuel consumption and costs while simultaneously minimizing pollutant emissions to comply with environmental regulations and mitigate harmful effects on human health and the environment. However, the determination of suitable combinations of operating parameters under fluctuating conditions, like capacity requirement and fuel properties, is not a trivial task, especially in real-time monitoring that requires immediate responses.

In this work, we quantify the efficiency of combustion using the average mass fraction of CO at the outlet, with lower values indicating higher efficiency. Additionally, we use the emissions of NO as the controlled pollutant indicator, as NO is a major contributor to the formation of ground-level ozone and acid rain, and its emissions are strictly regulated in many countries. By optimizing these two objectives simultaneously, we demonstrate how the MFRNet-based response surfaces can identify optimal operating conditions that balance efficiency and environmental concerns. This framework is also directly transferable for arbitrary target quantities of concern for the combustion system. Our multi-objective optimization task can be formulated as:

$$\begin{aligned}\underset{U}{\operatorname{argmin}} f(\boldsymbol{U}) &= (f_1(\boldsymbol{U}), f_2(\boldsymbol{U})) \\ &= (\text{NO emission},\ \text{CO emission})\end{aligned} \tag{18}$$

where $f_1(\boldsymbol{U})$ and $f_2(\boldsymbol{U})$ typically exhibit a conflicting relationship, necessitating a search for an optimal trade-off.

We employ the distributed evolutionary algorithms in Python (DEAP) framework to implement the non-dominated sorting genetic algorithm II (NSGA-II) to optimize the condition parameters for this multi-objective task [43]. Although gradient-based optimization via automatic differentiation is also feasible, the optimization landscape for industrial-scale combustion systems is usually highly non-convex and characterized by multiple local optima. In contrast, NSGA-II circumvents these challenges by efficiently exploring the design space without relying on precise gradient information. Furthermore, the multi-objective nature of our task would require consolidating multiple objectives into a single scalar function. Using NSGA-II aligns with common practices in digital twin studies, ensuring robustness and versatility in optimizing complex, real-world systems.

Initially, we generate 100 different initial conditions and derive their corresponding CO and NO missions using the response surfaces. The emissions data are then evaluated, assigning a higher priority to points with lower emissions. This selection mechanism ensures that high-priority points are propagated into subsequent generations. Additionally, to maintain diversity across generations and prevent clustering in certain regions, NSGA-II calculates the crowding distance for each solution [44]:

$$d_i = \sum_{n=1}^{N_{\text{obj}}} \left(f_n^{i+1} - f_n^{i-1}\right) / \left(f_n^{\max} - f_n^{\min}\right) \tag{19}$$

where d_i is the crowding distance for the ith individual, N_{obj} is the number of objectives, and $f^{\max}$ and $f^{\min}$ are the maximum and minimum values observed in the population for each objective. Solutions with larger crowding distances are given priority during the selection process. Through crossover and mutation processes, new generations are produced. Each generation retains elite solutions from both prior and current generations to preserve high-quality individuals within the population. Ultimately, we identify the optimized condition parameters, represented as points on the Pareto front [43]:

$$\left\{\boldsymbol{U}_{\text{pareto}_1}, \boldsymbol{U}_{\text{pareto}_2}, \ldots, \boldsymbol{U}_{\text{pareto}_{100}}\right\} \tag{20}$$

To finalize the optimization process, we employ the Euclidean distance approach to determine the optimal conditions for minimizing emissions. The solution that minimizes the Euclidean distance to the ideal point is selected:

$$d(\phi, \phi^*) = \sqrt{\sum_{n=1}^{N_{\text{obj}}} \left(\phi_n - \phi_n^*\right)^2} \tag{21}$$

where ϕ is a solution in the Pareto front and ϕ^* is the theoretical ideal point with the lowest possible emissions. The solution ϕ that minimizes this distance represents the best achievable trade-off between the conflicting objectives of minimizing NO and CO emissions.

2.3. Data generation and pre-processing

To construct the datasets used for training the digital twin, we employed a state-of-the-art comprehensive CFD model capable of capturing the multi-phase fuel bed combustion and the in-furnace as-phase combustion processes. The detailed settings of the CFD model and the data generation process have been demonstrated in our previous works [20,45–47], and thus here we only provide a brief demonstration. More details have also been summarized in Section S1 in Appendix A.

The numerical model consists of a thermally-thick particle treatment for the multiphase fuel bed and a detailed freeboard combustion model. The 2D and 3D temperature fields were generated as ground truth using this CFD framework. For the 2D simulations, a mesh consisting of 30 265 grid points was used. This mesh approximates the central cross-section of the furnace and neglects side-wall effects. For the 3D simulations, a mesh with 2.61 million grid points was created to model the entire furnace. Grid independence tests were performed to determine the optimal mesh sizes for both 2D and 3D computations.

The NO_x formation and reduction mechanisms were only implemented for the 3D high-fidelity simulations [45]. For the heterogeneous combustion process in the fuel bed, the model accounted for the release of fuel-N, as well as homogeneous and heterogeneous oxidation and reduction reactions. In the furnace freeboard, the model included the decomposition of tar-N, as well as the oxidation and reduction pathways for key precursors such as NH_3 and HCN.

To efficiently explore the high-dimensional condition space (denoted as $\boldsymbol{U}$), we adopted an incremental training strategy combined with LHS. In each round r, a group of 12 candidate cases ($\boldsymbol{G}_r = \boldsymbol{U}_1^r, \ldots, \boldsymbol{U}_{12}^r$) was selected using LHS from the condition space. For each case in $\boldsymbol{G}_r$, 2D CFD simulations were performed for the fuel

bed and freeboard combustion processes. From these 12 cases, two cases were randomly selected for 3D high-fidelity CFD simulations. After 24 rounds of sampling, a total of 288 2D cases and 48 3D cases were collected. This dual-dataset approach enabled the evaluation of temperature field reconstruction in both 2D and 3D spaces, as verified in our previous work [20]. It is worth noting that the computational time required for 2D CFD simulations is significantly lower than that needed to obtain 3D high-fidelity data. Using an advanced micro devices (AMD) 3970X 32-core workstation, the average completion time for a single 2D CFD simulation is around 1.5 h, while a 3D CFD simulation takes approximately 48–56 h to reach convergence. Therefore, utilizing 2D simulations to thoroughly explore the parameter space first proves to be a highly efficient strategy. All physical fields were pre-processed through global min–max normalization. For a given physical field ϕ_i, we will generate for training,

$$\widehat{\phi}_i = \frac{\phi_i - \phi_i^{\min}}{\phi_i^{\max} - \phi_i^{\min}} \tag{22}$$

where $\phi_i^{\min}$ and $\phi_i^{\max}$ denote the global minimum and maximum values of the field across all training data, respectively. This normalization ensures that the inputs and outputs of the neural network are scaled to comparable ranges, which facilitates stable and efficient training.

Finally, the configurations of all the neural network modules trained in this work, together with the hyper-parameters, training procedures, and evaluation metrics, are demonstrated in Section S2.2 in Appendix A. In the following sections, for visualization and quantitative comparison, we selected three cases from the test set of both 2D and 3D datasets. Their detailed settings of operating conditions are shown in Table 1.

3. Results and discussion

3.1. Multi-modal inputs and multi-field reconstruction

Real-world combustion systems provide various types of partial information, including operational parameters, temperature measurements, and flame images, each capturing internal operating conditions to varying degrees. To enhance digital twin performance, our framework integrates diverse inputs by mapping them into a unified latent space via feature matching. This approach enables robust multi-field reconstruction even when some inputs are missing, thereby improving model reliability in practical scenarios. Without loss of generality, we select operational parameters and temperature measurement data as two input modalities to evaluate multi-modal, multi-field reconstruction.

We introduce a contrastive learning strategy that employs two neural network-based projection heads. These projection heads transform the conditional vector $\boldsymbol{U}$ along with 25 randomly sampled temperature measurement points $\boldsymbol{S}$ into latent space. The transformation facilitates the subsequent multi-field reconstruction tasks by aligning the modalities within a shared manifold. To validate this alignment, we applied *K*-means clustering algorithm [48]. We comparatively analyzed the latent space structures with and without applying the contrastive learning loss supervision using *t*-distributed stochastic neighbor embedding (*t*-SNE) for visualization, as shown in Fig. 5. We expect information from the same case but different sources to map to the same point (i.e., to form overlapping clusters); thus we can verify that operational conditions and sparse sensor measurements align meaningfully for the same downstream tasks.

The clustering results reveal a significant distinction in the manifold alignment when contrastive learning is utilized. Without contrastive learning, the clusters corresponding to the two types of inputs showed minimal overlap, indicating distinct contributions to the multi-field reconstruction task from each input type. This scenario necessitates the presence of both input types to achieve effective reconstruction. Conversely, when trained with contrastive learning loss, the features from both input types exhibit nearly identical clustering in the latent space, with a high degree of overlap within clusters. This indicates a successful alignment of different input types, enabling the model to perform multi-field reconstruction from any singular type of input. Despite the fact that we specified four clusters for the classification of the latent features, the algorithm identified only two distinct clusters for features generated from conditional vector when contrastive loss is not applied, as shown in Fig. 5(a). This behavior can be attributed to the inherent structure, which may not exhibit sufficient variability or separable groupings to justify four distinct clusters. *K*-means relies on minimizing within-cluster variance, and if the data points in a dataset are relatively homogeneous or aligned along fewer natural groupings, the algorithm may converge to fewer effective clusters, even if more were requested [49].

Table 2 presents the test mean squared error (MSE) for multi-field reconstruction using both $\boldsymbol{U}$ and $\boldsymbol{S}$ as joint inputs, comparing scenarios with and without the application of contrastive learning loss. Additionally, we report the MSE for reconstructions performed using only $\boldsymbol{U}$ or $\boldsymbol{S}$ as inputs in models trained with contrastive learning loss. Further visualization of the multi-field reconstruction is provided in Fig. 6, showcasing the reconstruction capabilities from Case 1 to 3. We randomly selected these three operating conditions with significant differences in combustion modes from the test set, which effectively demonstrates the applicability of the digital twin model we established to various complex combustion conditions. Their temperature fields exhibit significantly different patterns, reflecting the variability of flame behavior within the furnace. Correspondingly, the distributions of flow fields as well as the concentrations of key species also show substantial differences. Changes in parameters within a high-dimensional condition space can profoundly alter the reactions occurring within the combustion system, highlighting the critical importance of digital twin systems.

By comparing the reconstruction results from $\boldsymbol{U}$ and $\boldsymbol{S}$, it is clear that introducing contrastive learning for feature fusion does not adversely affect the error rates of multi-field reconstruction, compared to directly merging information from both sources. Both strategies exhibit comparable reconstruction capabilities on the current dataset. With contrastive learning, the model acquires the ability to reconstruct fields from single modal inputs without significant loss in precision, highlighting the robustness of the feature fusion approach. Although our numerical comparisons indicate that a configuration without contrastive learning (Setting-2)

Table 1
Detailed operating conditions of three selected cases in the test set for model validation.

Case	CR	ER	Primary air ratios	Grate velocities ($m \cdot s^{-1}$)
1	1.2	1.4	[0.22, 0.32, 0.24, 0.14, 0.08]	[0.0072, 0.0072, 0.0072, 0.0072]
2	0.8	0.7	[0.10, 0.18, 0.28, 0.28, 0.18]	[0.0072, 0.0072, 0.0062, 0.0062]
3	0.5	1.3	[0.20, 0.32, 0.20, 0.18, 0.10]	[0.0032, 0.0032, 0.0032, 0.0032]

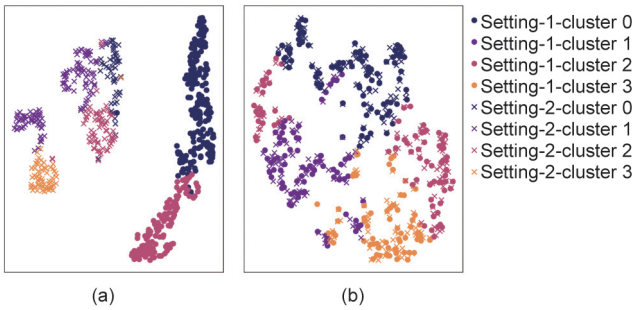

Fig. 5. Visualization of the *k*-means clustering (n_{clusters} = 4) of latent features generated by the projection heads: (a) without contrastive loss; (b) with contrastive loss. Setting-1: features generated from conditional vector ***U***; Setting-2: features generated from temperature measurement points ***S***. When trained with contrastive learning loss in (b), the features from both input types exhibit nearly identical clustering in the latent space, indicating a successful alignment of different input types.

achieves slightly lower MSE in some cases, it is important to note that the contrastive learning framework is designed primarily to improve the model's adaptability. This adaptability is crucial for scenarios where only single-modal inputs are available. For instance, settings using single-modal inputs without contrastive learning (Settings-5 and -6) yield reconstruction accuracies close to the multi-modal case but do not offer the same flexibility. In real-world digital twin systems, the ability to operate under varying data availability conditions is critical, ensuring that the model maintains stable performance in diverse operational scenarios.

The feature fusion framework demonstrated in this study has the potential for scalability. It can incorporate additional types of data, such as from more diverse measurement points or flame imagery, to construct a more versatile and flexible digital twin system. We also include a comparison with proper orthogonal decomposition (POD) on 2D data (Fig. S5 and Table S6 in Appendix A), where POD is shown to have significantly lower accuracy. The performance improvements observed with MFRNet underscore the advantages of the deep learning framework, particularly in effectively integrating diverse data modalities, which is challenging to implement using conventional methods. In subsequent chapters, we will explore the generalizability of our model, specifically the 2D MFRNet, using only the conditional vector ***U*** as input.

3.2. Dimension expansion

When constructing tailored digital twin systems for diverse combustion systems with varying structures, operational scenarios, and loads, the capability for dimension expansion becomes critically important. Dimension expansion allows models pre-trained on low-cost, low-precision 2D databases to achieve high-precision multi-field reconstructions on smaller 3D databases. This capability significantly reduces the reliance on high-precision data while effectively reconstructing system states in high-dimensional condition spaces. As mentioned earlier, our constructed 2D dataset comprises 288 operating conditions, while the 3D dataset contains 48 conditions. Considering that the computational time cost for a single 3D case CFD calculation is approximately 35 times that of a 2D simulation, it is evident that building the 2D dataset requires significantly fewer computational resources than the 3D dataset while covering more data points across the condition space.

Fig. 7 contrasts the training losses of the MFRNet model when directly trained on a 3D database against those obtained from a pre-trained 2D MFRNet followed by dimension expansion training. Specifically, Fig. 7(a) illustrates the learning curve when using the complete 3D dataset. In Fig. 7(b), we further demonstrate the range of test errors for both approaches using varying numbers of 3D cases in the training set. Table 3 further details the comparative reconstruction errors in the 3D multi-field tests for both training modalities.

During the training process, models that undergo dimension expansion on a pre-trained 2D MFRNet exhibit rapid convergence of both training and testing errors, maintaining very close levels between them. In contrast, direct training using the MFRNet architecture on a smaller 3D dataset results in a slower reduction in training loss, with a clear divergence between training and testing errors. This observation indicates that the dimension expansion training framework can more effectively capture additional features in small datasets and prevent overfitting, thereby significantly enhancing model accuracy. The results in Fig. 7(b) reveal that although the multi-field reconstruction test errors decrease with increasing dataset size for both approaches (with and without 2D pre-training), there is a substantial difference in their error ranges. Even with just 15 cases in the training set, the average test error achieved through 2D pre-training and dimension expansion reaches 0.01748, which significantly outperforms the reconstruction neural network directly trained with 40 3D cases. Notably, given that 3D data generation is approximately 35 times more expensive per case, generating the 2D dataset (288 cases) takes only roughly 30% of the computational cost of generating the 3D dataset, while the accuracy is greatly improved. This strongly demonstrates the effectiveness of the dimension expansion strategy in combining 2D and 3D datasets, enabling a substantial reduction in data requirements for training digital twin models while maintaining accuracy. From the perspective of reconstruction errors across various physical fields, it is evident that a small 3D dataset alone, with its extensive parameter space, is insufficient for training effective digital twin models capable of multi-field reconstruction. However, if MFRNet has already thoroughly explored the condition space on a 2D dataset and performed initial feature extraction, the 3D reconstruction task can be simplified to learning *z*-directional non-uniformities and overall accuracy refinement. For example, in the test dataset, models trained directly on the 3D data exhibit an average MSE of 4.28×10^{-3} for the temperature field, which is too high for downstream applica-

Table 2
The mean squared errors of multi-field reconstruction in the test set of the 2D dataset, with all quantities normalized to the range [0, 1].

Setting	Temperature	Pressure	*x*-velocity	*y*-velocity	M_{O_2}	M_{CO_2}	M_{H_2O}	M_{CO}
Setting-1	9.70×10^{-4}	7.24×10^{-5}	2.55×10^{-5}	2.92×10^{-5}	1.20×10^{-3}	1.34×10^{-3}	7.35×10^{-4}	1.21×10^{-3}
Setting-2	9.66×10^{-4}	6.28×10^{-5}	2.82×10^{-5}	2.81×10^{-5}	1.09×10^{-3}	1.15×10^{-3}	7.64×10^{-4}	1.33×10^{-3}
Setting-3	9.38×10^{-4}	7.37×10^{-5}	2.87×10^{-5}	3.89×10^{-5}	1.31×10^{-3}	1.44×10^{-3}	7.88×10^{-4}	1.40×10^{-3}
Setting-4	1.02×10^{-3}	7.89×10^{-5}	2.76×10^{-5}	3.49×10^{-5}	1.28×10^{-3}	1.29×10^{-3}	8.66×10^{-4}	1.43×10^{-3}
Setting-5	1.15×10^{-3}	7.01×10^{-5}	2.58×10^{-5}	3.98×10^{-5}	1.42×10^{-3}	1.59×10^{-3}	8.20×10^{-4}	1.62×10^{-3}
Setting-6	9.53×10^{-4}	6.72×10^{-5}	2.46×10^{-5}	3.10×10^{-5}	1.33×10^{-3}	1.51×10^{-3}	8.76×10^{-4}	1.20×10^{-3}

Setting-1: joint input of ***U*** and ***S*** with contrastive learning loss; Setting-2: joint input of ***U*** and ***S*** without contrastive learning loss; Setting-3: only input of ***U*** after trained with contrastive learning loss; Setting-4: only input of ***S*** after trained with contrastive learning loss; Setting-5: single-modal input trained on input of ***U***; Setting-6: single-modal input trained on input of ***S***; *M*: the mass fraction of species.

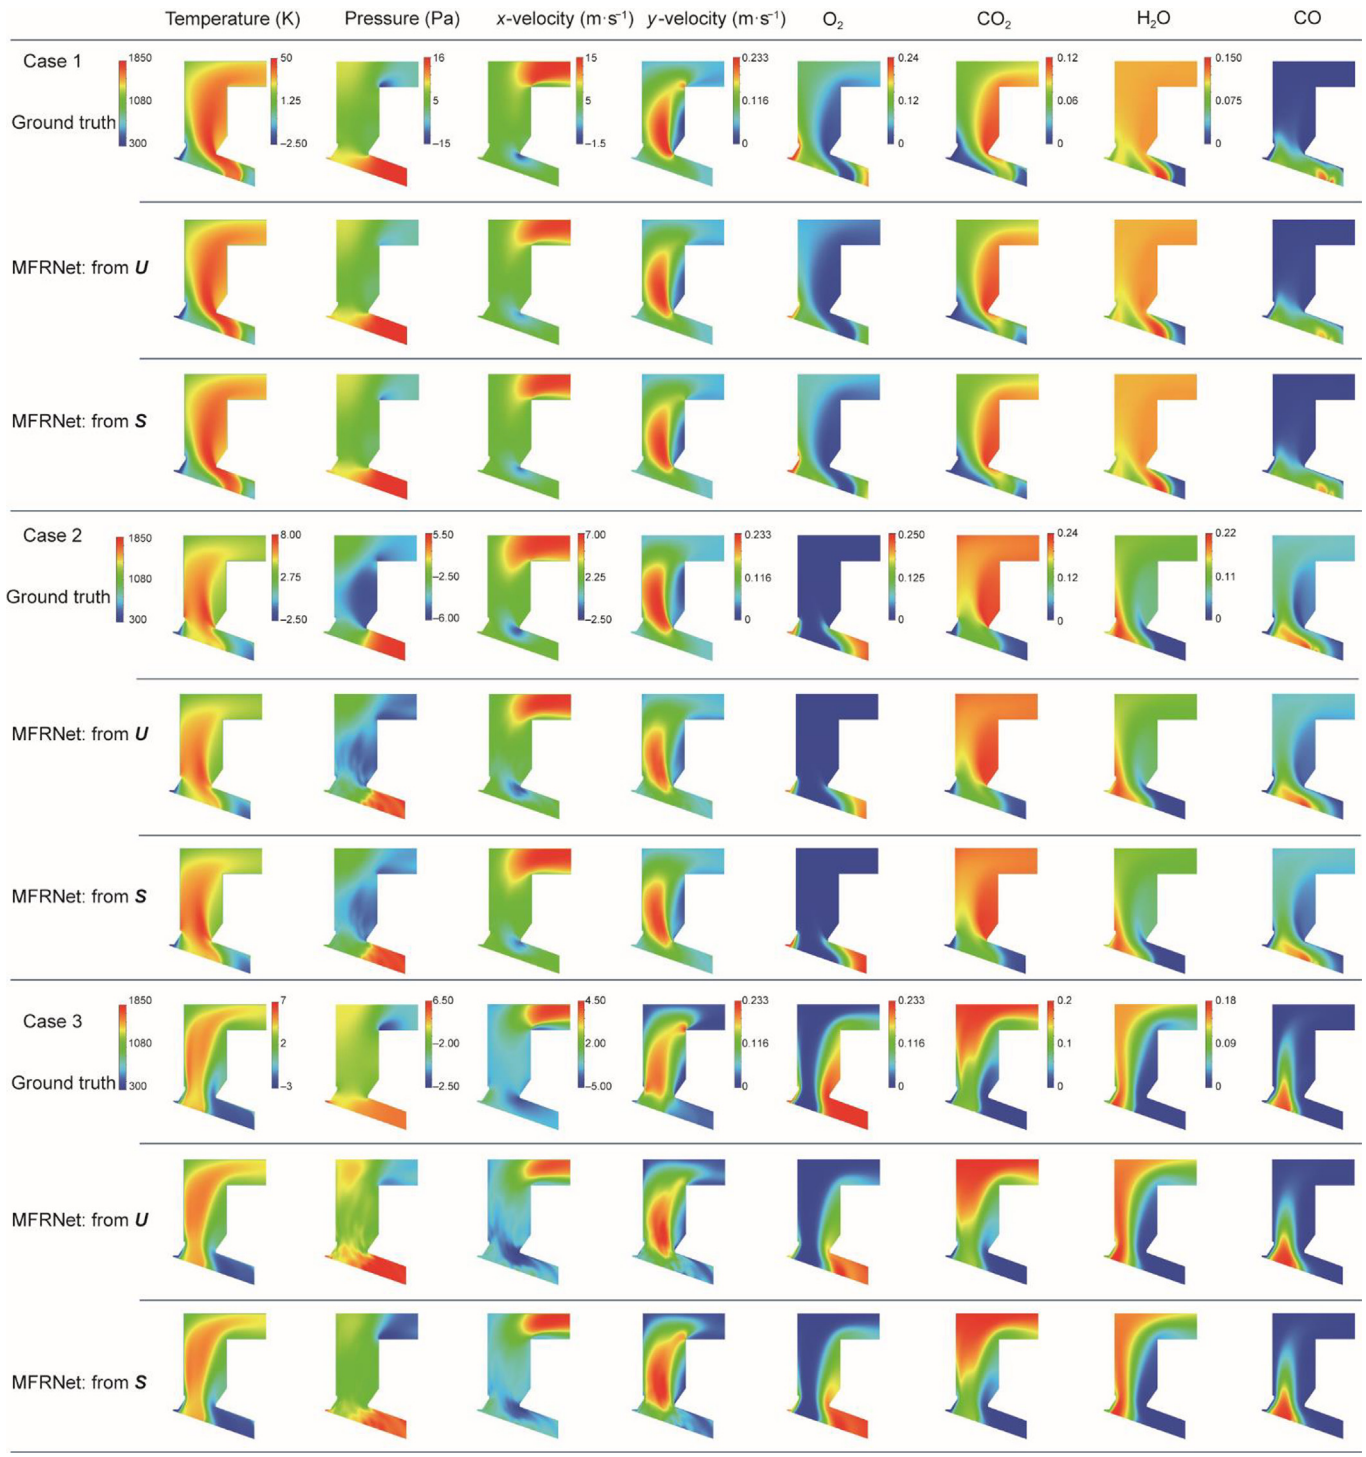

Fig. 6. Visualization of the 2D multi-field reconstruction trained with contrastive learning loss. The outputs from using only input of ***U*** and ***S*** are compared with the ground truth. The three cases show significant differences in combustion modes, reflecting the variability of flame behavior in high dimensional condition space. By introducing contrastive learning, the model achieves similar reconstruction accuracy with partial input information.

tions and may fail to capture essential flame distribution characteristics. On the other hand, models undergoing dimension expansion based on a 2D MFRNet achieve an average temperature field MSE of 1.14×10^{-3}, approaching the precision level of multi-field reconstructions on larger 2D datasets.

Given that the 3D dataset contains only approximately 15 % of the total cases in the 2D dataset, this result underscores the practical significance of dimension expansion in building effective and precise digital twin systems. It is acknowledged that comparisons between using a pre-trained network and a model trained from scratch involve trade-offs. In practical industrial applications, pre-training is typically conducted as a one-time investment, with its cost amortized over multiple uses. Although summing the number of epochs for pre-training and fine-tuning would shift the training curves, our results are presented to illustrate final performance improvement and computational efficiency. We emphasize

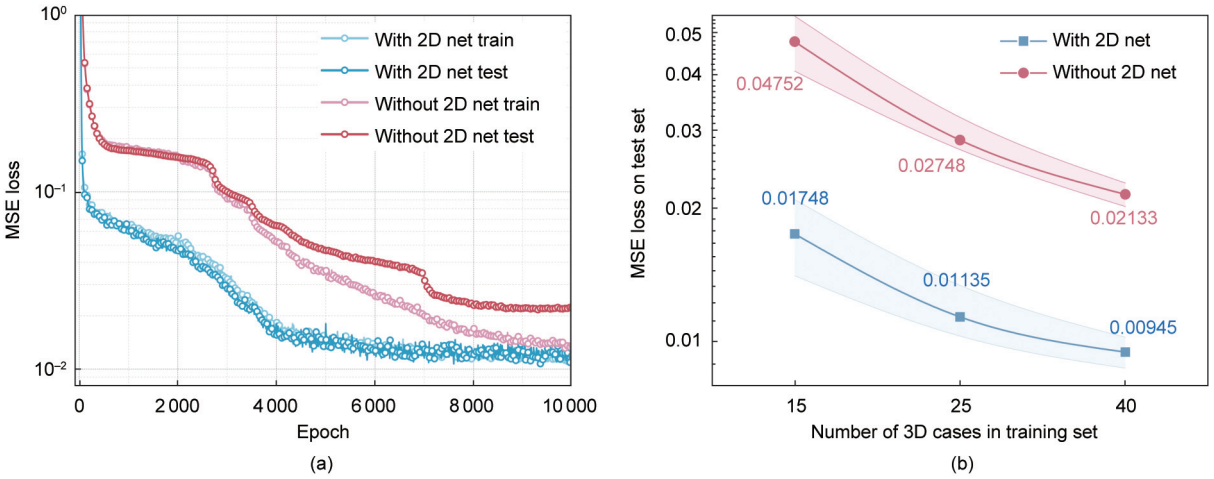

Fig. 7. Comparison of the losses in the 3D reconstruction training process between with and without pre-trained 2D reconstruction net. (a) The learning curves of reconstruction on the complete 3D dataset; (b) the total reconstruction MSE losses on the test set when trained on different numbers of 3D cases in the training set.

Table 3
The MSE of multi-field reconstruction in the test set of 3D dataset.

Setting	Temperature	Pressure	*x*-velocity	*y*-velocity	O_2	CO_2	H_2O	CO
Setting-7	1.14×10^{-3}	2.12×10^{-4}	7.28×10^{-5}	6.81×10^{-5}	2.87×10^{-3}	2.65×10^{-3}	8.65×10^{-4}	1.48×10^{-3}
Setting-8	4.28×10^{-3}	5.50×10^{-5}	1.50×10^{-4}	1.21×10^{-4}	5.99×10^{-3}	7.53×10^{-3}	1.63×10^{-3}	3.03×10^{-3}

Setting-7: dimension expansion on pre-trained 2D MFRNet; Setting-8: direct training MFRNet on 3D dataset.

that the comparisons are designed to reflect realistic conditions where data availability, network architectures, and training strategies vary.

Fig. 8 further showcases a side-view visualization of the 3D multi-field reconstruction from Case 1 to 3. In addition, Fig. 9 presents a top-view of three physical fields for Case 1, including temperature, CO, and O_2. The results for Cases 2 and 3 are provided in Figs. S1 and S2 in Appendix A. Compared to the 2D multi-field distributions of Cases 1 to 3 presented in the previous sections, the high-fidelity results obtained from 3D simulations show significant differences. On the one hand, in the *xy*-plane (i.e., the longitudinal cross-section), the 3D structure of the flow channel induces changes in the local flow field and mixing behavior. For instance, in the 2D simulation of Case 3, the mixing and burnout rate of CO were evidently overestimated, resulting in a relatively uniform temperature distribution. However, the 3D simulation reveals a more pronounced diffusion flame pattern. On the other hand, due to wall effects, noticeable non-uniformity is observed in the *z*-direction. These findings indicate that both dimension expansion and accuracy refinement must be considered when performing 3D multi-field reconstructions.

We compare the results of dimension expansion based on the 2D MFRNet with those obtained directly from training on the 3D dataset. In the side view presented in Fig. 8, the model trained with dimension expansion accurately reconstructs the key patterns in various physical fields. For instance, in the temperature field, critical features such as the ignition point, high temperature zones, and low temperature regions are faithfully replicated. From the top view in Fig. 9, the outputs of the dimension expansion training exhibit a faithful reproduction of the *z*-directional non-uniformity in temperature and component concentration distributions. The magnitudes and shapes of the patterns are both accurately and finely represented, indicating a successful reconstruction.

In contrast, the MFRNet trained directly on the 3D dataset does not achieve satisfactory reconstruction results, with significant quantitative errors across the physical fields. The visualization results highlight its inability to accurately reconstruct various physical quantities in 3D. For example, in the temperature field, there is a significant misalignment in the positioning of the flame, while the distributions and magnitudes of components like O_2 and CO show notable discrepancies.

This comparison underscores the effectiveness of dimension expansion from a 2D pre-trained model in enhancing the precision and accuracy of 3D reconstructions, particularly when dealing with smaller datasets that inherently lack extensive variability in high-dimensional spaces. Meanwhile, considering the computational cost, the generation of high-fidelity 3D CFD simulations requires approximately 48–56 h per case on a 64-core AMD Ryzen Threadripper 3970X central processing unit (CPU). In comparison, our MFRNet can be trained on a single RTX A6000 graphics processing unit (GPU) in about 2–3 h for the 2D reconstruction part and an additional 3–5 h for the 3D dimension and variable extension components. This represents a significant reduction in computational cost when contrasted with conventional CFD methods. Moreover, even compared with alternative machine learning approaches, our method yields substantial efficiency gains, particularly given the challenge posed by the irregular geometry of the computational domain, where traditional CNN-based architectures often struggle.

3.3. Variable extension

Variable extension is another critical capability for multi-physics systems. Its core concept lies in leveraging the intrinsic coupling and correlations within multi-physics systems to dynamically integrate learned features of existing physical fields. This enables the model to infer the distribution of new physical variables, even with smaller datasets. This capability has wide applications, especially considering the varying computational costs and requirements for different variables. For instance, when generating low-fidelity datasets using numerical methods, simplifying chemical reaction mechanisms or ignoring certain minor

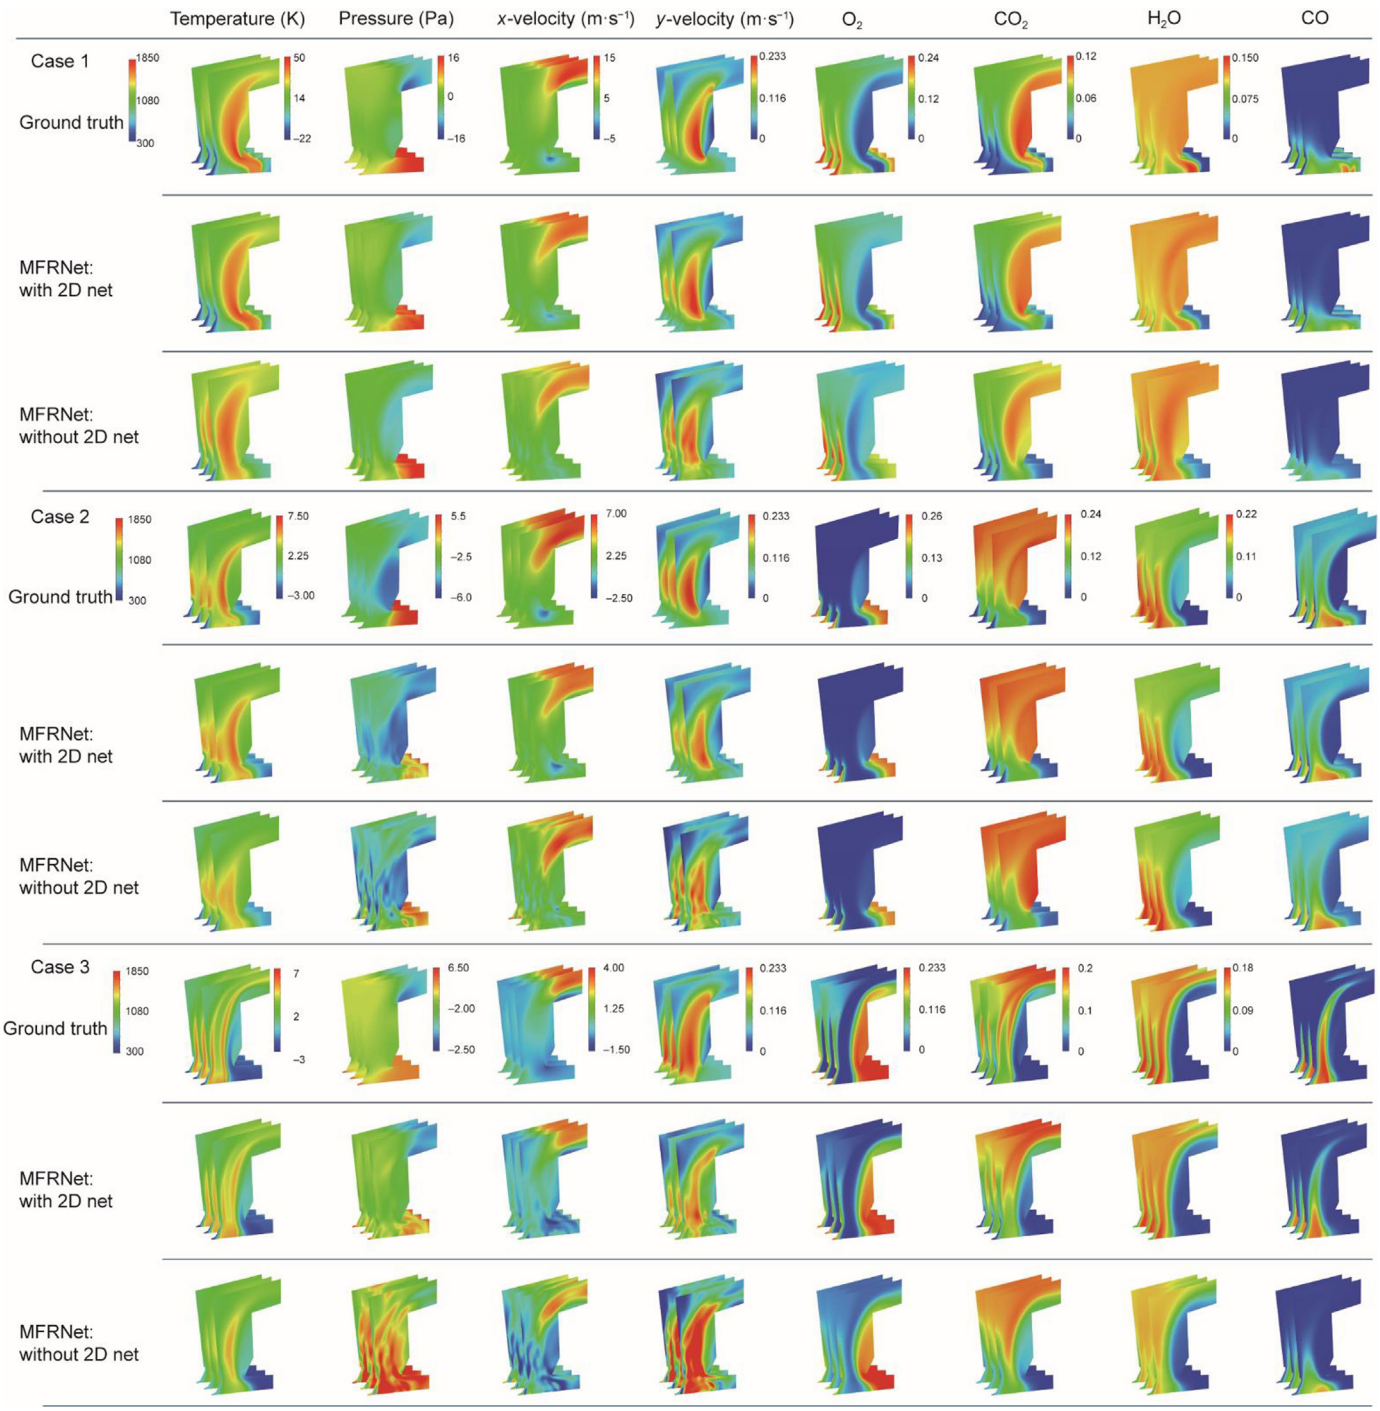

Fig. 8. Side-view visualization of the 3D multi-field reconstruction. The outputs from using dimension expansion with 2D net and direct training on 3D dataset without 2D net are compared.

species is often necessary to accelerate the computation of source terms [50].

In this study, we demonstrate the extension of a 2D combustion dataset to a 3D multi-field reconstruction, specifically focusing on the spatial distributions of NO_x-related species, including the mass fractions of NO, HCN, and NH_3.

Fig. 10 compares the total losses of NO_x prediction using two approaches: direct training on the 3D database with MFRNet and variable extension training based on a pre-trained 2D MFRNet. Additionally, Table 4 provides detailed reconstruction errors for NO_x-related variables in the 3D test set under both training modes. Notably, when applying variable extension after training on the combustion dataset, the initial loss during the training process is much lower, and the loss decreases significantly faster. In contrast, directly learning the distributions of NO_x-related variables from the smaller 3D dataset results in slower convergence and substantially higher test errors. Quantitatively, for NO, HCN, and NH_3, the test errors under the variable extension setting are nearly an order of magnitude lower than those from direct training. This highlights the importance of feature fusion in complex multi-field reconstruction tasks, particularly when considering the coupling relationships between different physical fields.

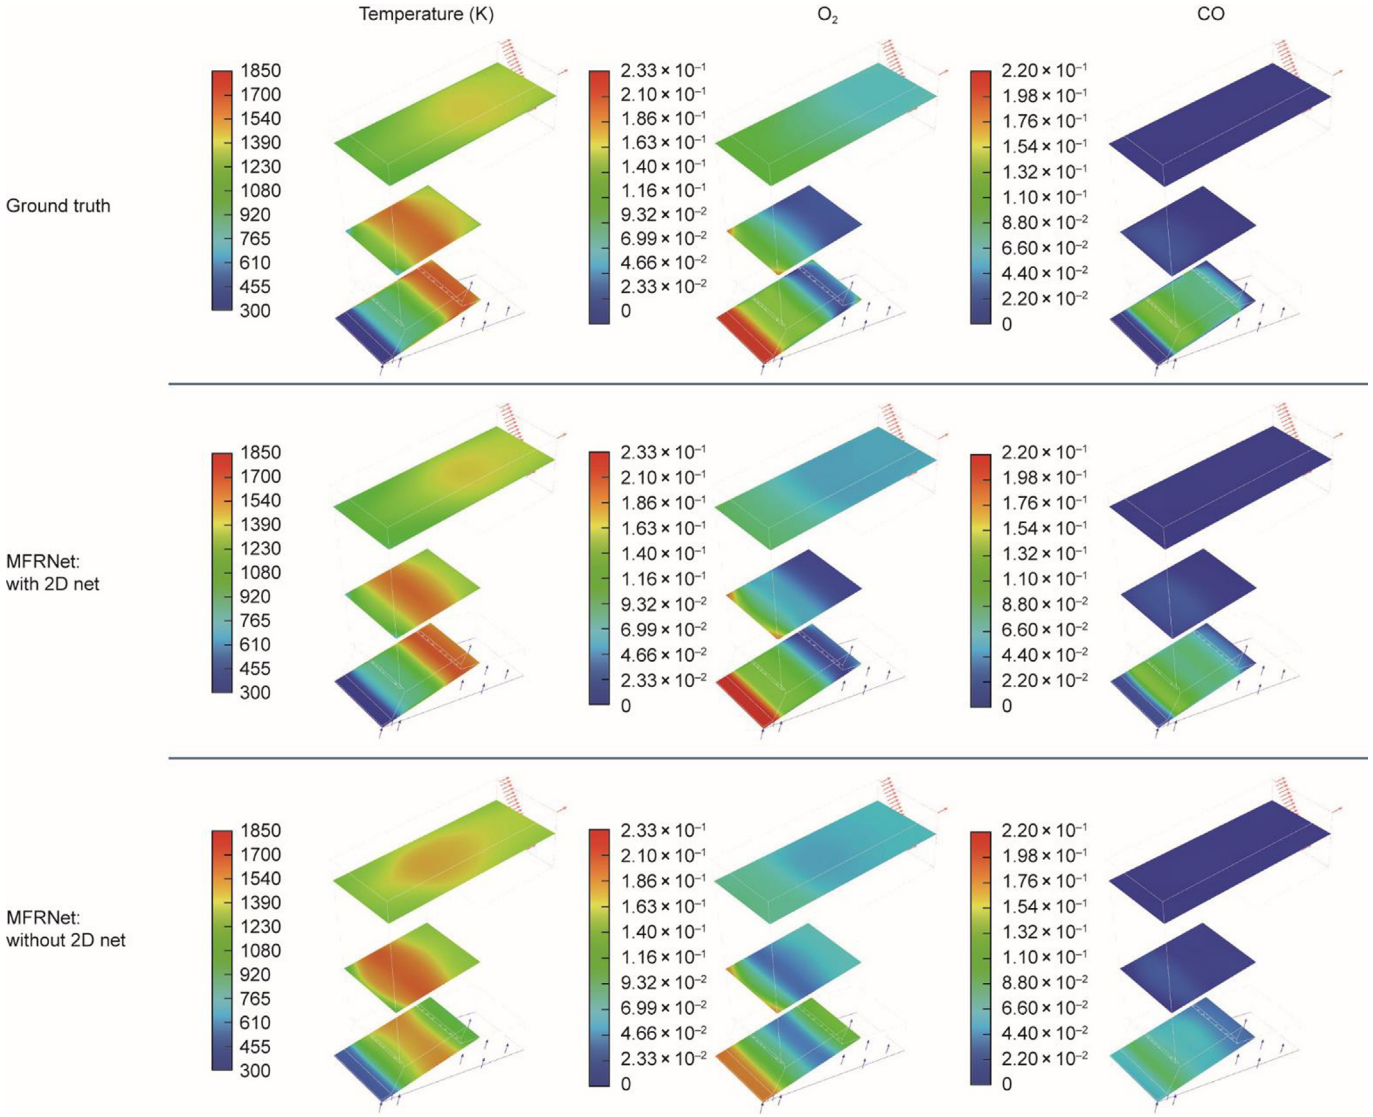

Fig. 9. Top-view visualization of the 3D multi-field reconstruction of Case 1. The outputs from using dimension expansion with pre-trained 2D reconstruction net and direct training on 3D dataset without 2D net are compared.

Fig. 11 visualizes the side-view reconstruction of 3D NO_x-related fields for Case 1. Results for Cases 2 and 3 are provided in Figs. S3 and S4 in Appendix A. The comparison includes results from variable extension based on 2D MFRNet and direct training on the 3D dataset. Unlike the 3D dimension expansion discussed in the previous section for combustion-related physical fields, the NO_x-related variables were not pre-learned in the 2D dataset. Thus, this task involves not only accuracy refinement in the z-direction but also learning the overall distribution characteristics of these variables from the small dataset. After applying variable extension based on the 2D MFRNet, the fundamental spatial distribution features of NO, HCN, and NH_3 are still well captured.

However, due to the limited size of the dataset and the higher demands of 3D reconstruction, certain local regions are not fully captured. For example, as shown in Fig. 11, the high-concentration distribution of NO at the rear of the grate and the localized peak of NH_3 at the middle of the grate are not perfectly reconstructed. Nevertheless, it is important to emphasize that the performance of variable extension remains significantly better than direct training on the 3D dataset. Without prior feature

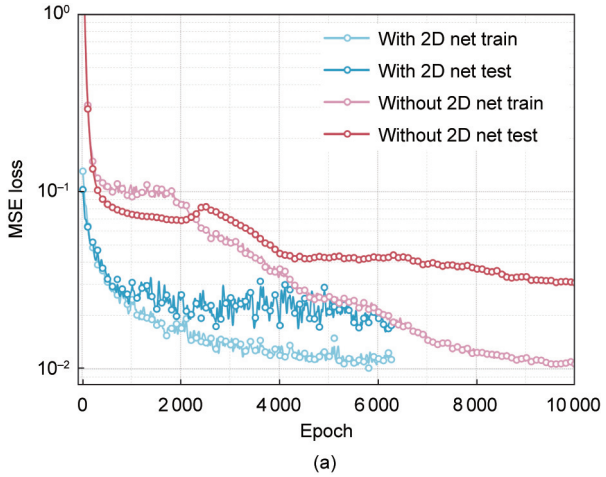

Fig. 10. Comparison of the losses in the 3D reconstruction training process of NO_x-related species between with and without pre-trained 2D reconstruction net.

Table 4
The mean squared errors of NO_x-related species reconstruction in the test set of 3D dataset. Setting-9: Variable extension on pre-trained 2D reconstruction net and 3D dataset; Setting-10: Direct training the reconstruction net on 3D dataset.

Setting	NO	HCN	NH_3
Setting-9	3.84×10^{-3}	2.04×10^{-3}	7.82×10^{-4}
Setting-10	1.27×10^{-2}	4.34×10^{-3}	9.98×10^{-3}

fusion, the distribution regions and numerical ranges of related variables are more prone to distortion during the learning process. It is also worth noting that generalization to extreme or unseen operational conditions is an important consideration in future works. Extrapolation to conditions far outside this range represents a challenging open question in machine learning. Our framework is designed to accommodate continual learning, allowing for progressive expansion of the training dataset and further improvements in generalization over time.

3.4. Scalar prediction

In the preceding sections, we demonstrated how dimension expansion and variable extension enable multi-field reconstruction on smaller high-fidelity 3D datasets. Beyond reconstruction tasks, scalar prediction is also a critical core module of a digital twin. These scalars represent global properties of combustion devices, such as combustion efficiency and pollutant emission indices. Establishing the response surface between input conditions and these scalars is a prerequisite for automating control and operational optimization in many industrial combustion systems [51]. However, the accuracy of scalar prediction is significantly influenced by dataset size. Reliable scalar values often require data derived from high-fidelity simulations, whereas those obtained from low-fidelity datasets tend to exhibit substantial biases that can hinder downstream tasks. In this section, we will further demonstrate that the detailed reconstruction of multiple physical fields in Section 3.1 not only validates the network's ability to capture high-fidelity spatial distributions but also extracts rich latent features that encapsulate the underlying combustion dynamics. These latent representations can be subsequently leveraged for direct prediction of key outlet scalars. This two-stage approach reduces the computational burden compared to full-field post-processing while still leveraging information from the reconstruction task.

In this study, we focus on predicting two scalars at the furnace outlet: the mass fraction of CO and the standardized NO concentration ($mg \cdot Nm^{-3}$, converted to values under 6% oxygen concentra-

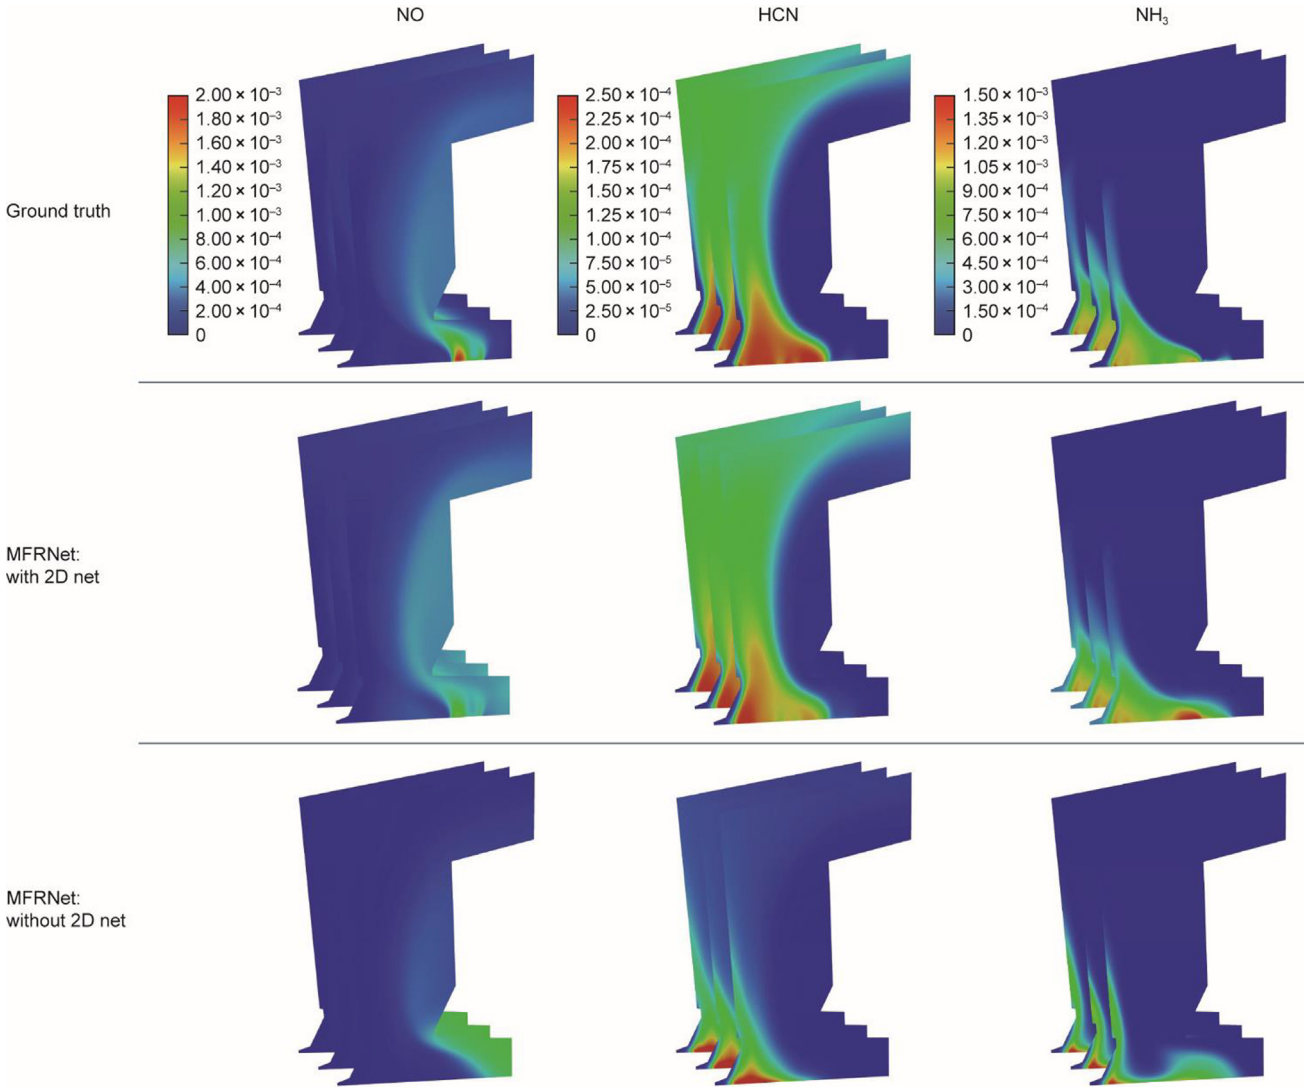

Fig. 11. Side-view visualization of NO_x-related species reconstruction of Case 1. The outputs from using variable extension with 2D reconstruction net and direct training on 3D dataset without 2D reconstruction net are compared.

tion). The CO mass fraction reflects the gas-phase burnout efficiency, while NO concentration is a key pollutant that needs to be tightly controlled. Notably, the 3D dataset provides limited data points for these scalars. For CO, we leverage 2D data for learning and subsequent correction, while the NO emission values are predicted based on the dynamic characteristics of various combustion fields. Fig. 12 illustrates the relationship between the predicted scalar values and the ground truth, including CO on the 2D dataset, CO on the 3D dataset, and NO on the 3D dataset. We compared the performance of scalar prediction using the MFRNet-based method with that obtained by directly training an MLP. Overall, by utilizing the intermediate features extracted from MFRNet trained on the multi-field reconstruction task, scalar prediction achieves significantly higher accuracy than direct scalar prediction using MLP under the same network structure and dataset size. The MFRNet further incorporates a trainable feature fusion layer for integrating latent features derived during the multi-field reconstruction phase. Although this results in a slightly larger model capacity compared to the direct MLP method, our experimental results demonstrate that this additional capacity is not the sole factor driving performance improvements. We performed additional testing with enlarged MLPs (Table S7 in Appendix A), confirmed that a mere increase in parameter count does not lead to appreciable gains in accuracy or generalizability. Instead, the effectiveness of the MFRNet is largely due to the improved feature fusion. This enhanced latent representation offers a compelling advantage in digital twin applications.

For the CO concentration at the outlet on both 2D and 3D datasets, the values vary by several orders of magnitude under different operating conditions. When using MLP for direct prediction, significant biases are observed, particularly in the low- and medium-concentration regions. In contrast, when leveraging models pre-trained on multi-field reconstruction tasks for 2D and 3D datasets and dynamically fusing their intermediate features via an attention mechanism, the predictions achieve excellent agreement with the ground truth, with R^2 values reaching 0.9955 on both datasets. For NO emission prediction, the error associated with MLP is much larger. When using directly the condition vector $\boldsymbol{U}$ as input for MLP to predict outlet NO concentration, the R^2 value is only 0.738, indicating low reliability and large deviations from the results. However, by combining the intermediate features of 3D MFRNet, the R^2 value for NO prediction improves significantly to 0.980. This highlights the interdependence between different physical fields and the necessity of multi-level data fusion for effective feature extraction.

In summary, this section reveals a crucial finding: reconstruction tasks and scalar prediction tasks can be integrated within a single framework through feature fusion strategies. By training on multi-field reconstruction across multiple physical fields, the resulting intermediate features can effectively enable high-accuracy scalar prediction on small datasets. This has significant implications for automated control and optimization tasks in industrial applications [52].

3.5. Multi-objective optimization

After establishing the response surface, we can proceed to multi-objective optimization, which is a crucial functionality for the practical application of the digital twin model. The multi-objective optimization task involves identifying a set of operational parameters $\boldsymbol{U}$ for a given capacity ratio that simultaneously control both CO and NO emission levels. This task has general applicability to various combustion systems and addresses a pressing issue in real-world applications: how to rapidly adjust combustion modes under fluctuating load conditions to achieve optimal combustion states [53]. Fig. 13 presents the Pareto front graphs obtained through the optimization process for different CR values. We selected three representative capacity ratios, including 0.75, 1.00, and 1.25, which correspond to low-load, design-load, and overload operating conditions, respectively. From the results, it is evident that a conflict exists between the objective functions at each capacity ratio.

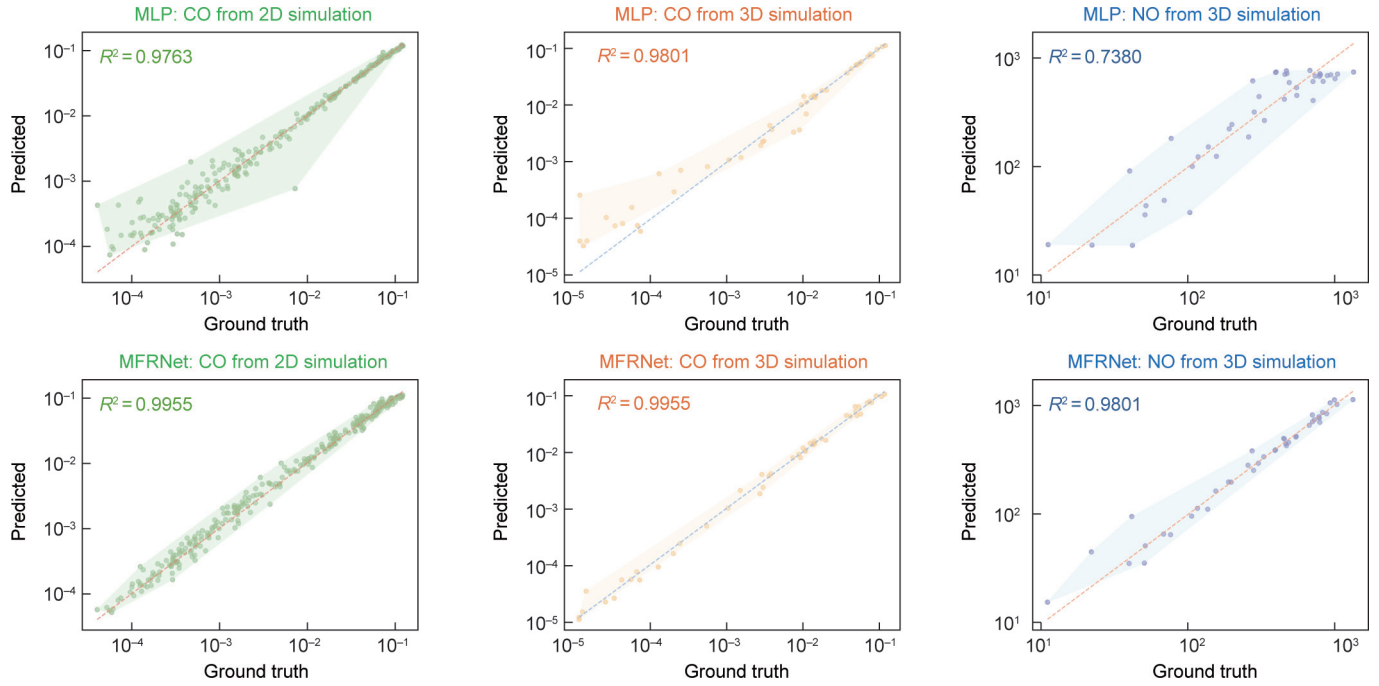

Fig. 12. Results of scalar prediction, including the mass fraction of CO (both from 2D and 3D datasets) and NO emission level (mg·Nm^{-3}) at the furnace outlet. The accuracy of prediction is compared between leveraging the pre-trained reconstruction modules of MFRNet, and direct MLP training.

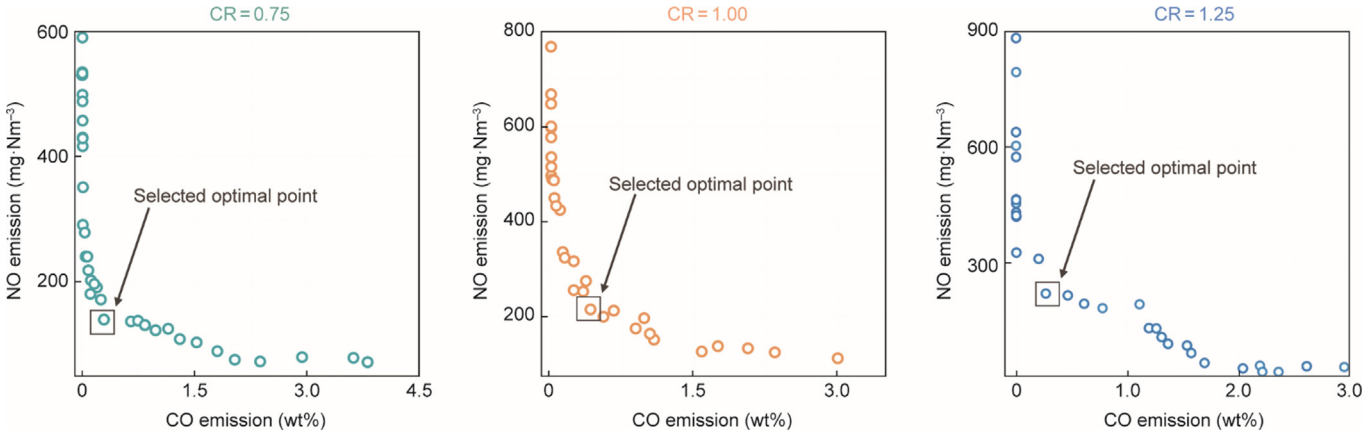

Fig. 13. Pareto points graph by the optimization process of different CR values.

Specifically, achieving optimal burnout efficiency requires sufficient oxygen supply and effective mixing. However, this often leads to local over-oxygenation and overheating, significantly increasing NO emissions. Conversely, suppressing NO formation may lead to incomplete combustion, resulting in higher CO emissions and reduced combustion efficiency.

Table 5 further summarizes the optimal operating conditions obtained through Pareto optimization for the three capacity ratios, along with the predicted CO and NO emission levels based on the MFRNet-derived response surface. In addition, we provide predictions from directly trained MLP-based response surfaces as well as validation results obtained through CFD simulations. The performance of directly trained MLP-based response surfaces, as presented in Section 3.4, represent the capability of data-driven system control solely relying on the small high-fidelity dataset. In multi-parameter optimization problems, the predictive accuracy of the response surface model is paramount. If the model cannot reliably characterize the burnout and emission profiles of the system under specific operating conditions, meaningful optimization becomes fundamentally unfeasible.

Taking CR = 1.0 as an example, the optimized operating condition suggests an equivalence ratio of 1.0, a delayed air distribution strategy, and a relatively low grate speed. This method aligns well with findings from recent studies [45], which indicate that for grate-fired boilers, combining delayed air distribution with a thick fuel bed can effectively control NO_x emissions while maintaining burnout efficiency. Remarkably, this operating strategy was identified solely through the data-driven response surface in this study.

Furthermore, the MFRNet-based response surface demonstrates high accuracy in predicting CO and NO emissions, particularly when compared to directly trained MLP-based response surfaces. For instance, at CR = 1.0, the relative errors of the MFRNet-based predictions for CO and NO emissions are 2.82 % and 8.89 %, respectively, while the directly trained MLP-based response surfaces yield relative errors of 34.22 % and 31.44 %, respectively. These results demonstrate that the MFRNet model accurately predicts combustion and pollutant distributions under various optimized conditions, offering valuable insights into the implementation details of each target operating condition.

It is worth noting that this study uses CO and NO as representative indicators to demonstrate how multi-objective optimization can be performed based on the MFRNet-derived response surface. While these metrics provide valuable insights, other performance measures, such as thermal efficiency and unburned carbon loss, are also important for a complete assessment of industrial combustion systems. Due to current data limitations and the focus of our study, these additional metrics were not incorporated into our optimization framework. Nonetheless, they represent promising directions for future research aimed at broadening the industrial relevance of digital twin applications. In practical applications, the optimization objectives can be further expanded to include these metrics. The key takeaway is that the MFRNet-based framework provides a robust foundation for flexible and extensible optimization tailored to specific needs.

4. Conclusions

In this study, we developed a digital twin framework based on the MFRNet, which effectively integrates data of different fidelities for hierarchical feature extraction and accuracy refinement. The framework combines reconstruction and optimization tasks into a unified workflow, providing a novel paradigm for data-driven intelligent operation of complex combustion devices. The key conclusions are as follows.

(1) When facing a complex condition space, pre-learning on a low-dimensional approximate dataset and subsequently expanding dimensions enables multi-field reconstruction on minimal 3D high-fidelity data, resulting in significant accuracy improvements compared to direct training on a smaller 3D dataset.

Table 5
Optimized condition vector and the predicted results by Pareto optimization using MFRNet-based response surfaces under different CR.

CR	Optimized condition vector	CFD simulation		MFRNet-based response surfaces		Direct training MLP	
		CO	NO	CO	NO	CO	NO
0.75	[0.75, 1.05,0.137, 0.141, 0.317, 0.299, 0.106, 0.0045, 0.0038, 0.0038, 0.0041]	0.6275	104.36	0.4982	112.87	0.7980	64.85
1.00	[1.0, 1.105, 0.085, 0.128, 0.288, 0.300, 0.199, 0.0041, 0.0041, 0.0039, 0.0039]	0.4237	207.27	0.4117	178.85	0.6687	122.09
1.25	[1.25, 0.975, 0.107, 0.086, 0.274, 0.291, 0.242, 0.0037, 0.0044, 0.0037, 0.0038]	0.8211	212.85	0.5827	228.92	1.155	145.43

(2) Variables not simulated in the 2D low-fidelity dataset can be learned with considerable accuracy on a smaller 3D dataset by dynamically integrating features obtained from the reconstruction process of learned variables. The MSE is an order of magnitude lower than with direct training, reflecting the inherent coupling of flow, heat transfer, and reactions in complex combustion systems.

(3) Reconstruction and scalar prediction tasks can be integrated into a unified framework using feature fusion. By dynamically integrating features corresponding to different operating conditions from pre-trained reconstruction models through an attention mechanism, a response surface with remarkable accuracy and generalization performance can be obtained using a small high-fidelity dataset.

(4) By establishing response surfaces for multiple key objective values, such as burnout efficiency and NO_x emissions, the model can perform multi-objective optimization, quickly identifying optimal operating conditions in high-dimensional condition spaces. Moreover, the digital twin reconstruction module can rapidly reconstruct multiple fields, providing a comprehensive evaluation of the combustor under target conditions.

In summary, this work addresses the data barrier for customizing data-driven digital twins for diverse industrial-scale combustion facilities. By integrating low- and high-fidelity datasets, we can achieve high-precision 3D multi-field reconstruction, scalar prediction, and multi-objective optimization with minimal cost.

CRediT authorship contribution statement

Linzheng Wang: Writing – original draft, Visualization, Investigation, Conceptualization. **Yaojun Li:** Visualization, Software, Methodology, Investigation. **Sili Deng:** Writing – review & editing, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

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

Acknowledgments

Linzheng Wang acknowledges the financial support from the International Post-doctoral Fellowship Program of the Biomass Energy Research Center, Shanghai Jiao Tong University. Sili Deng acknowledges the financial support from ExxonMobil through the Massachusetts Institute of Technology Energy Initiative as a Founding Member.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.eng.2025.08.020.

References

[1] Liu Z, Li M, Virguez E, Xie X. Low-carbon transition pathways of power systems for Guangdong–Hong Kong–Macao region in China. Energy Environ Sci 2024;17:307–22.
[2] Padovani D, Dimitriou P, Minav T. Challenges and solutions for designing energy-efficient and low-pollutant machines in off-road hydraulics. Energy Conv Manag X 2024;21:100526.
[3] Li X, Wang X, Wang H, He F. Low-temperature-solid combustion technology of biomass for pollution reduction: potentials and necessary fundamentals. ACS Omega 2023;8:43433–41.
[4] Larki I, Zahedi A, Asadi M, Forootan MM, Farajollahi M, Ahmadi R, et al. Mitigation approaches and techniques for combustion power plants flue gas emissions: a comprehensive review. Sci Total Environ 2023;903:166108.
[5] Li J, Zhou Q, He X, Chen W, Xu H. Data-driven enabling technologies in soft sensors of modern internal combustion engines: perspectives. Energy 2023;272:127067.
[6] Aliramezani M, Koch CR, Shahbakhti M. Modeling, diagnostics, optimization, and control of internal combustion engines via modern machine learning techniques: a review and future directions. Prog Energy Combust Sci 2022;88:100967.
[7] Zheng Z, Shafique M, Luo X, Wang S. A systematic review towards integrative energy management of smart grids and urban energy systems. Renew Sustain Energy Rev 2024;189:114023.
[8] Aversano G, Ferrarotti M, Parente A. Digital twin of a combustion furnace operating in flameless conditions: reduced-order model development from CFD simulations. Proc Combust Instit 2021;38:5373–81.
[9] Tao F, Zhang H, Zhang C. Advancements and challenges of digital twins in industry. Nat Comput Sci 2024;4:169–77.
[10] Liu X, Jiang D, Tao B, Xiang F, Jiang G, Sun Y, et al. A systematic review of digital twin about physical entities, virtual models, twin data, and applications. Adv Eng Inf 2023;55:101876.
[11] Acharya S, Khan AA, Päivärinta T. Interoperability levels and challenges of digital twins in cyber–physical systems. J Ind Inform Integr 2024;42:100714.
[12] Spinti JP, Smith PJ, Smith ST. Atikokan digital twin: machine learning in a biomass energy system. Appl Energy 2022;310:118436.
[13] Zhou L, Song Y, Ji W, Wei H. Machine learning for combustion. Energy AI 2022;7:100128.
[14] Parente A, Swaminathan N. Data-driven models and digital twins for sustainable combustion technologies. iScience 2024;27(4):109349.
[15] Hafeez MA, Procacci A, Coussement A, Parente A. Challenges and opportunities for the application of digital twins in hard-to-abate industries: a review. Resour Conserv Recycl 2024;209:107796.
[16] Ihme M, Chung WT, Mishra AA. Combustion machine learning: principles, progress and prospects. Prog Energy Combust Sci 2022;91:101010.
[17] Nemitallah MA, Nabhan MA, Alowaifeer M, Haeruman A, Alzahrani F, Habib MA, et al. Artificial intelligence for control and optimization of boilers' performance and emissions: a review. J Clean Prod 2023;417:138109.
[18] Zheng Y, Li Q, Zhu P, Li X, Zhang G, Ma X, et al. Study on multi-field evolution and influencing factors of coal spontaneous combustion in goaf. Combust Sci Technol 2023;195(2):247–64.
[19] Meng X, Karniadakis GE. A composite neural network that learns from multi-fidelity data: application to function approximation and inverse PDE problems. J Comput Phys 2020;401:109020.
[20] Wang L, Deng R, Zhang R, Luo Y, Deng S. 3-D full-field reconstruction of chemically reacting flow towards high-dimension conditions through machine learning. Chem Eng J 2024;499:156435.
[21] Buster G, Benton BN, Glaws A, King RN. High-resolution meteorology with climate change impacts from global climate model data using generative machine learning. Nat Energy 2024;9:894–906.
[22] Jung KS, Soriano BS, Chen JH, Khalil M. A Hessian-based transfer learning approach for artificial neural networks based chemical kinetics with a sparse dataset. Proc Combust Inst 2024;40(1–4):105390.
[23] Cui Z, Xu J, Liu W, Zhao G, Ma S. Data-driven modeling-based digital twin of supercritical coal-fired boiler for metal temperature anomaly detection. Energy 2023;278:127959.
[24] Cui C, Zhang J, Shen J. System-level modeling, analysis and coordinated control design for the pressurized water reactor nuclear power system. Energy 2023;283:128472.
[25] Donato L, Galletti C, Parente A. Self-updating digital twin of a hydrogen-powered furnace using data assimilation. Appl Thermal Eng 2024;236:121431.
[26] Nie X, Zhang W, Dong X, Medwell PR, Nathan GJ, Sun Z. Reconstructing temperature fields from oh distribution and soot volume fraction in turbulent flames using an artificial neural network. Combust Flame 2024;259:113182.
[27] Savarese M, Procacci A, Iavarone S, Giuntini L, De Paepe W, Parente A. A sparse sensing and chemical reactor network based framework for the development of physics-based digital twins of combustion devices. Proc Combust Inst 2024;40(1–4):105536.
[28] Zhang Y, Li J. Research on digital twin modeling method for combustion process based on model reduction. Case Stud Thermal Eng 2024;65:105619.
[29] Yang S, Wan MP, Chen W, Ng BF, Dubey S. Model predictive control with adaptive machine-learning-based model for building energy efficiency and comfort optimization. Appl Energy 2020;271:115147.
[30] Huang ZF, Soh KY, Islam MR, Chua KJ. Digital twin driven life-cycle operation optimization for combined cooling heating and power-cold energy recovery (CCHP-CER) system. Appl Energy 2022;324:119774.
[31] Sungur B, Basar C, Kaleli A. Multi-objective optimisation of the emission parameters and efficiency of pellet stove at different supply airflow positions based on machine learning approach. Energy 2023;278:127896.
[32] Xu J, Cui Z, Ma S, Wang X, Zhang Z, Zhang G. Data based digital twin for operational performance optimization in CFB boilers. Energy 2024;306:132532.
[33] Ali F, Dawood A, Hussain A, Alnasir MH, Khan MA, Butt TM, et al. Fueling the future: biomass applications for green and sustainable energy. Discov Sustain 2024;5:156.

[34] Zhang J, Wang Z, Li S, Wei P. A digital twin approach for gas turbine performance based on deep multi-model fusion. Appl Thermal Eng 2024;246:122954.
[35] Zhang Y, Ding J, Li Y, Ren Z, Feng K. Multi-modal data cross-domain fusion network for gearbox fault diagnosis under variable operating conditions. Eng Appl Artif Intel 2024;133:108236.
[36] Baltrušaitis T, Ahuja C, Morency LP. Multimodal machine learning: a survey and taxonomy. IEEE Trans Pat Anal Mach Intell 2018;41:423–43.
[37] Le-Khac PH, Healy G, Smeaton AF. Contrastive representation learning: a framework and review. IEEE Access 2020;8:193907–34.
[38] Lu L, Jin P, Pang G, Zhang Z, Karniadakis GE. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. Nat Mach Intell 2021;3:218–29.
[39] Smith JD, Sreedharan V, Landon M, Smith ZP. Advanced design optimization of combustion equipment for biomass combustion. Renew Energy 2020;145:1597–607.
[40] Rahimipetroudi I, Rashid K, Yang JB, Dong SK. Use of response surface methodology to optimize NO_x emissions and efficiency of w-type regenerative radiant tube burner under plasma-assisted combustion. J Clean Prod 2020;244:118626.
[41] Zhang Y, Dong W, Vandewalle LA, Xu R, Smith GP, Wang H. Neural network approach to response surface development for reaction model optimization and uncertainty minimization. Combust Flame 2023;251:112679.
[42] Jiang Y, Lee BH, Oh DH, Jeon CH. Optimization of operating conditions to achieve combustion stability and reduce NO_x emission at half-load for a 550-MW tangentially fired pulverized coal boiler. Fuel 2021;306:121727.
[43] Gao S, Zhang Y, Zhang Z, Tan D, Li J, Yin Z, et al. Multi-objective optimization of the combustion chamber geometry for a highland diesel engine fueled with diesel/*n*-butanol/$PODE_n$ by ANN-NSGA III. Energy 2023;282:128793.
[44] Ma H, Zhang Y, Sun S, Liu T, Shan Y. A comprehensive survey on NSGA-II for multi-objective optimization and applications. Artif Intell Rev 2023;56:15217–70.
[45] Zhang R, Wang L, Deng R, Luo Y. Numerical analysis of NO_x reduction in large-scale MSW grate furnace through in-bed combustion optimization using multi-section fuel bed model with thermally thick treatment. Appl Thermal Eng 2024;257:124156.
[46] Deng R, Wang L, Zhang R, Luo Y. Numerical modeling of fixed-bed cocombustion processes through the multiple thermally thick particle model. ACS Omega 2022;7:39938–49.
[47] Wang L, Zhang R, Deng R, Liu Z, Luo Y. Comprehensive parametric study of fixed-bed co-gasification process through multiple thermally thick particle (MTTP) model. Appl Energy 2023;348:121525.
[48] Ikotun AM, Ezugwu AE, Abualigah L, Abuhaija B, Heming J. *K*-means clustering algorithms: a comprehensive review, variants analysis, and advances in the era of big data. Inform Sci 2023;622:178–210.
[49] Jain AK. Data clustering: 50 years beyond *k*-means. Pattern Recognit Let 2010;31:651–66.
[50] Cai T, Zhao D, Gutmark E. Overview of fundamental kinetic mechanisms and emission mitigation in ammonia combustion. Chem Eng J 2023;458:141391.
[51] Guan J, Li Y, Liu J, Duan X, Shen D, Jia D, et al. Experimental and numerical research on the performance characteristics of OPLVCR engine based on the NSGA II algorithm using digital twins. Energy Conv Manag 2021;236:114052.
[52] Zhou H, Zhang R, Wang L, Luo Y. Comprehensive assessment and optimization of a middle- arch dual-channel municipal solid waste incinerator using numerical simulation methods. ACS Omega 2024;9:42010–26.
[53] Yao Q, Zhang Y, Wang X, Tian Z, Hu G, Du W. Investigation of NO_x emission under different burner structures with the optimized combustion model. Neurocomputing 2022;482:224–35.