

FEDSM2026-182699

GRAPH NEURAL NETWORK REDUCED-ORDER MODELS FOR MULTIPHASE CFD: FLUIDIZED BED AND HYDRAULIC FRACTURING APPLICATIONS

Saurav Mitra
CPFD LLC
Houston, TX

Shashank K. Karra
CPFD LLC
Houston, TX

Keshav Ramchandran
CPFD LLC
Houston, TX

James Parker
CPFD LLC
Houston, TX

ABSTRACT

Multiphase Computational Fluid Dynamics (CFD) simulations are essential for analyzing industrial systems involving gas–solid flows. Consequently, high-fidelity CFD simulations are computationally intensive for design exploration and optimization and are generally impractical for real-time decision-making in large-scale systems. This work presents a Graph Neural Network (GNN)–based Reduced Order Modeling (ROM) framework that achieves approximately $80\text{--}105\times$ speedup for transient and time-averaged field predictions respectively, and up to three orders of magnitude speedup for isolated snapshot queries, while preserving predictive accuracy. The approach is demonstrated on two representative multiphase systems: a fluidized bed gasifier and proppant transport in hydraulic fractures.

The proposed framework is first evaluated on the proppant transport problem, where experimental measurements enable direct validation of ROM predictions. An Eulerian GNN-ROM is trained on CFD-generated Eulerian fields and validated against experimental measurements at 300, 630, and 1140 seconds, exhibiting strong spatial accuracy, with MAE consistently below 0.02 m.

The framework is then applied to a fluidized bed gasifier to assess its ability to generalize to unseen operating conditions using high-fidelity CFD as reference. GNN-ROMs are trained using Barracuda Virtual Reactor CFD data across multiple operating conditions defined by varying bottom air flow rates (0.691, 0.8, 0.9, and 1.0 kg/s). The model is evaluated at an intermediate, unseen flow rate (0.85 kg/s) and accurately predicts both time-averaged and transient fields for key quantities such as CH_4 , CO_2 , pressure, and particle volume fraction. Inference time is reduced from hours to under 10 minutes, with time-averaged relative errors below 7% and transient predictions preserving dominant spatial trends.

A key innovation of this work is the GNN architecture, which uniquely combines spatial graph convolutions on CFD meshes with explicit time-conditioning via Fourier feature embeddings.

This design enables high-resolution, temporally flexible predictions at arbitrary time instances without the constraints of iterative time-stepping, offering unprecedented flexibility and computational efficiency for complex multiphase flow systems.

This demonstrates the potential for generalization to unseen conditions and highlights GNN-ROMs as transformative tools for multiphase flow modeling, enabling real-time process control, rapid design optimization, and digital twin applications in energy and subsurface engineering.

Keywords: Reduced Order Modeling, Graph Neural Networks, Gasifier, Proppant Transport, Barracuda Virtual Reactor CFD

NOMENCLATURE

Roman Symbols

V	Graph nodes (CFD cells)
E	Graph edges (mesh connectivity)
G_k	Graph of the CFD snapshot at time index k
G^*	Graph used for inference
S_k	CFD snapshot at time index k
X_k	Node feature matrix for snapshot k
X^*	Node feature matrix for inference
X_{dyn}	Dynamic Eulerian node variables
Y_k	Predicted Eulerian field outputs at snapshot k
Y^*	Predicted Eulerian fields at inference time
u_k	Operating parameter at snapshot k
u^*	Operating parameter for inference
k	Snapshot index
N_s	Total number of training snapshots
t	Physical time
t^*	Query time for inference
u_i	CFD field value
$\hat{u}_i$	ROM field value

Greek Symbols

$\phi(t)$	Fourier feature embedding of time
$\phi(t_k)$	Time embedding at snapshot time t_k
ρ	Spatial correlation

Subscripts and Superscripts

$(\cdot)_k$	Snapshot index
$(\cdot)^*$	Inference quantity
$(\cdot)_i$	Node index

1. INTRODUCTION

Multiphase Computational Fluid Dynamics (CFD) simulations are essential for analyzing industrial systems involving gas–solid flows, including fluidized bed reactors, hydraulic fracture proppant transport, and energy conversion devices [1–4]. These systems exhibit strongly coupled nonlinear physics, multiscale transport phenomena, and complex geometries, requiring high spatial and temporal resolution to accurately resolve interphase momentum exchange, chemical reactions, and solids redistribution. As a result, high-fidelity CFD simulations often require hours to days of wall-clock time for a single operating condition, making design exploration and optimization computationally intensive and rendering real-time decision-making and digital-twin deployment impractical for large-scale systems [5–7].

Reduced-order models (ROMs) mitigate this computational burden by approximating the dynamics of high-fidelity simulations in a lower-dimensional representation. Classical projection-based ROMs, including Proper Orthogonal Decomposition (POD), Galerkin projection, and reduced-basis methods, have been successfully applied to linear or weakly nonlinear systems [8–11]. However, their applicability to industrial multiphase flows is often limited by strong nonlinearities, parametric variability, stability challenges, and the need for intrusive access to governing equations or problem-specific stabilization strategies [9–11].

Recent advances in machine learning have enabled data-driven ROMs capable of learning mappings between system states, parameters, and outputs directly from simulation data [12–16]. Deep autoencoders, convolutional neural networks, and recurrent architectures have demonstrated substantial computational speedups for fluid systems [13–15]. Nevertheless, many neural ROMs rely on autoregressive or time-marching formulations, which can accumulate error over long time horizons and restrict predictions to sequential temporal evaluation [16].

Graph Neural Networks (GNNs) provide a natural framework for learning spatially resolved dynamics directly on CFD meshes by representing control volumes as nodes and mesh adjacency as edges [17–19]. This mesh-aware representation preserves spatial locality and is well suited for transport-dominated multiphase flows on complex geometries. Prior studies have demonstrated the feasibility of GNN-based surrogate modeling for fluids and granular systems [17–19]. However, many existing GNN approaches adopt autoregressive time integration or latent-state evolution, which limit both temporal flexibility and introduce cumulative error during long-time prediction.

Alternative paradigms, including physics-informed neural networks (PINNs) and neural operator (NO) frameworks, have also been proposed for reduced-order modeling of parametric partial differential equations [20,21]. While these methods show promise in academic settings, their application to industrial-scale multiphase CFD problems involving complex geometries, stiff source terms, and detailed closure models remains limited.

In this work, we present a Graph Neural Network–based Reduced Order Modeling (GNN-ROM) framework that operates directly on Eulerian CFD meshes with fixed connectivity and incorporates explicit time conditioning via Fourier feature embeddings. Each CFD snapshot is represented as a graph whose nodes correspond to Eulerian control volumes and whose edges encode mesh adjacency, preserving the spatial structure of the finite-volume discretization. Temporal evolution and parametric dependence are introduced explicitly through node-level features augmented with operating parameters and time embeddings. This formulation enables single-shot prediction of full-field Eulerian quantities at arbitrary time instances without autoregressive time marching.

The proposed framework is demonstrated on two representative multiphase flow problems governed by Eulerian–Lagrangian formulations based on the Multiphase Particle-In-Cell (MP-PIC) method [22–24]. First, an Eulerian GNN-ROM is developed for proppant transport in hydraulic fractures and validated directly against laboratory-scale experimental measurements of proppant bed height [25]. Second, the framework is applied to a fluidized bed gasifier characterized by a complex, asymmetric geometry and a reacting gas–solid flow, consistent with the MP-PIC-based gasification model reported by Snider et al. [26]. In the latter case, high-fidelity simulations generated using Barracuda Virtual Reactor serve as reference data, and the GNN-ROM is evaluated for its ability to generalize to an unseen operating condition.

By combining mesh-aware graph representations, explicit time conditioning, experimental validation, and industrial-scale CFD benchmarks, this work advances the state of the art in non-intrusive reduced-order modeling for multiphase flow systems. The proposed formulation enables accurate, computationally efficient prediction of Eulerian flow fields across time and operating conditions, and supports applications in rapid design exploration, parametric studies, and digital-twin deployment.

2. METHODOLOGY

2.1 High-Fidelity CFD Framework

The high-fidelity multiphase CFD data used in this work are generated using Barracuda Virtual Reactor (BVR), a commercially validated solver for industrial gas–solid flow simulations. BVR employs a Multiphase Particle-In-Cell (MP-PIC) formulation [22–24], in which the gas phase is treated in an Eulerian framework while the particulate phase is represented by Lagrangian computational parcels. Momentum exchange between phases is modeled through interphase drag closures,

enabling efficient simulation of dense particulate systems with strong gas–solid coupling.

The MP-PIC approach is well suited for industrial multiphase applications characterized by large particle number densities and evolving particulate structures. In fluidized beds, BVR captures bed expansion, solids circulation, and bubble dynamics without resolving individual particle–particle collisions. In MP-PIC simulation of proppant transport in hydraulic fractures the Eulerian treatment of the carrier fluid combined with a parcel-based solids representation enables efficient simulation of bed formation and transport over engineering time scales. All simulations are performed on fixed computational meshes and produce time-resolved Eulerian field outputs, which serve as reference data for reduced-order modeling.

2.2 Problem Definition and ROM Scope

The objective of this work is to develop a non-intrusive Eulerian reduced-order model (ROM) that predicts key multiphase flow quantities across time and multiple operating conditions while retaining fidelity to high-resolution CFD simulations.

The proposed ROM is restricted to Eulerian field quantities defined on a fixed computational mesh. In the fluidized bed gasifier, the GNN-ROM predicts gas-phase species mass fractions, pressure, and particle volume fraction. For the hydraulic fracture application, the ROM predicts the Eulerian particle volume fraction field, from which derived metrics such as proppant bed height are computed. Although the underlying MP-PIC simulations employ Lagrangian parcels, the ROM operates exclusively on Eulerian outputs, ensuring consistent spatial discretization across all snapshots.

Parametric variability, such as inlet flow rate, is incorporated explicitly, and predictions are generated at arbitrary time instances within the temporal domain of interest. Unlike time-marching surrogate models, the proposed framework treats time as an explicit input and performs inference without sequential temporal integration, avoiding cumulative time-stepping error.

2.3 Graph Representation of CFD Fields

To enable learning directly on CFD discretizations while preserving spatial locality, the Eulerian field data are recast into a graph-based representation defined on the computational mesh. For a fixed mesh, the graph topology is constructed once and reused across all time instances and operating conditions.

Eq. (1) defines each CFD snapshot S_k ,

$$G_k = (V, E, X_k), \quad (1)$$

where V and E denote the fixed sets of nodes and edges derived from the CFD mesh, and X_k is the node feature matrix associated with snapshot k . Nodes (V) correspond to Eulerian control volumes, and edges (E) encode mesh connectivity based on shared faces, yielding a sparse graph structure consistent with the

underlying finite-volume discretization. For the structured meshes considered here, this connectivity is equivalent to a six-neighbor (i, j, k) face-sharing stencil and is constructed once and reused across all snapshots. All node-level attributes, including static geometric quantities such as cell centroids and volumes, dynamic Eulerian field variables, operating parameters, and time embeddings are contained exclusively in X_k .




FIGURE 1 MESH-TO-GRAPH REPRESENTATION OF EULERIAN FIELD, COLOR REPRESENTS THE MAGNITUDE OF THE FIELD.

The graph topology remains invariant across snapshots, reflecting the use of a fixed computational mesh. Temporal evolution and parametric variability are therefore captured exclusively through changes in node-level features rather than graph connectivity, enabling efficient reuse of the mesh topology during training and inference.

Node features are assembled as shown in Eq. (2),

$$X_k = [X_{dyn}(S_k), X_{geom}, u_k, \phi(t_k)] \quad (2)$$

where dynamic Eulerian fields (e.g., particle volume fraction, species mass fractions) extracted from the CFD solution are augmented with operating parameters (u_k) and an explicit time embedding ($\phi(t_k)$). Here, X_{geom} denotes static geometric attributes (e.g. cell centroids and volumes), invariant across snapshots.

Edges are used solely to define neighborhoods for message passing, and no explicit edge features are introduced. Boundary cells are treated identically to interior cells; cells adjacent to physical boundaries naturally have fewer neighbors due to mesh connectivity. Boundary effects are therefore encoded implicitly through node-level features, consistent with the non-intrusive nature of the ROM.

Under this formulation, the reduced-order model learns a nonlinear mapping as in Eq. (3),

$$F_\theta : G_k \mapsto Y_k \quad (3)$$

where predictions are produced directly on the original CFD mesh. The temporal information t_k and operating conditions u_k are implicitly encoded within the node feature representation X_k , and are therefore not included as separate inputs. At inference, the same graph topology is reused to construct G^* , enabling

single-shot predictions of full Eulerian fields at arbitrary query times and operating conditions.

2.4 GNN-Based Reduced Order Model Architecture




FIGURE 2 PROPOSED GNN-ROM MODEL ARCHITECTURE

The reduced-order model is formulated as a graph neural network (GNN) that approximates the nonlinear mapping between Eulerian CFD fields, operating parameters, and time. The model operates directly on the graph representation defined in Section 2.3 and performs spatial learning through message passing over the fixed mesh connectivity.

Given a graph snapshot G_k , the GNN-ROM computes node-wise latent representations through a sequence of L message-passing layers. Message passing is implemented using graph attention layers, which allow adaptive weighting of neighboring control volumes, though the formulation is agnostic to the specific aggregation operator. At layer l , node features are updated according to Eq. (4),

$$h_i^{(\ell+1)} = U^{(\ell)}(h_i^{(\ell)}, A^{(\ell)}(\{h_j^{(\ell)} : (i, j) \in E\})), \quad (4)$$

where $h_i^{(0)} = X_{k,i}$ and $A^{(\ell)}(\cdot)$ and $U^{(\ell)}(\cdot)$ denote learnable aggregation and update operators, respectively. This formulation enables the model to learn localized spatial interactions induced by advection, diffusion, and interphase coupling while respecting mesh adjacency.

After L message-passing layers, the final node representations are mapped to Eulerian field predictions through a readout operator, given by Eq. (5),

$$Y_k = R_\theta(h^{(L)}) \quad (5)$$

The GNN-ROM learns the nonlinear mapping in Eq. (6),

$$Y_k = F_\theta(V, E, X_k), \quad (6)$$

where temporal dependence and parametric variability are encoded explicitly in the node features via time embeddings and operating parameters. The graph $G_k = (V, E, X_k)$, fully defines the

model input, ensuring consistent notation. Predictions are obtained in a single forward pass without sequential time integration.

In the present implementation, the GNN-ROM consists of $L = 4$ message-passing layers with a hidden feature dimension of 128 at each layer. Graph attention is employed for aggregation, followed by a nonlinear activation function (ReLU) and layer normalization to improve training stability. These choices balance model expressiveness and computational efficiency.

Each message-passing layer propagates information across one level of mesh adjacency; consequently, stacking L layers yields an effective receptive field spanning approximately L rings of neighboring control volumes. This behavior is analogous to increasing the spatial stencil size in a finite-volume discretization, enabling the GNN-ROM to capture spatially extended transport effects while maintaining locality.

The model is trained using the Adam optimizer with a fixed learning rate of 0.001, for 50 epochs, and a batch size of 32. Standard early stopping based on validation loss is employed to prevent overfitting.

2.5 Explicit Time Conditioning via Fourier Features

To enable flexible temporal inference without reliance on sequential time integration, the proposed GNN-ROM treats time as an explicit conditioning variable rather than a marching coordinate. Instead of advancing the solution through

incremental time stepping, the model learns a direct mapping from physical time and operating conditions to Eulerian field outputs.

Temporal dependence is introduced through a Fourier feature embedding of the scalar time variable, given by Eq. (7).

$$\phi(t) = [\sin(\omega_1 t), \cos(\omega_1 t), \dots, \sin(\omega_m t), \cos(\omega_m t)] \quad (7)$$

$\{\omega_j\}_{j=1}^m$ are fixed frequencies selected to span the temporal scales present in the training data. In the present implementation, 8–16 sine–cosine frequency pairs are used, logarithmically spaced from the inverse of the longest resolved transient to the inverse of the smallest snapshot interval.

For each snapshot k , the time embedding $\phi(t_k)$ is broadcast to all graph nodes and concatenated with the dynamic Eulerian variables and operating parameters when forming the node feature matrix X_k . During inference, the same procedure is applied using the query time t^* , allowing predictions to be generated at arbitrary time instances.

Because temporal information is introduced solely through node-level features, the underlying graph topology and spatial coupling defined by mesh connectivity remain unchanged. Temporal consistency is therefore evaluated *a posteriori* through transient error metrics and correlation analysis, rather than enforced through sequential integration, aligning the formulation with its intended use for rapid field evaluation.

2.6 Training and Inference Procedure

The GNN-ROM is trained in a supervised manner using time-resolved CFD snapshots generated across multiple operating conditions on a fixed computational mesh. Each training sample corresponds to a single CFD snapshot represented as a graph with explicit time conditioning.

Model training minimizes a mean-squared-error (MSE) loss computed between predicted and reference Eulerian field quantities over all mesh nodes. When multiple Eulerian variables are predicted, the loss is evaluated independently for each variable and aggregated.

All CFD snapshots from the training operating conditions are aggregated into a single dataset and divided into training and validation sets using an 80–20 split at the snapshot level. Snapshots were temporally subsampled such that adjacent time instances were not split across training and validation sets. Validation data are used exclusively for monitoring convergence and preventing overfitting. For the gasifier case, generalization is assessed using a separate test data set corresponding to an unseen operating condition.

All input and output variables are normalized using statistics computed from the training dataset only, and the same normalization is applied to validation and test data. Training is performed offline, and inference is conducted via a single forward pass without sequential time integration. Because predictions are independent across time instances, inference can

be efficiently batched across multiple query times or operating conditions.

Physical consistency of the GNN-ROM predictions was verified *a posteriori*. Predicted quantities are clipped within physically admissible bounds to prevent numerical artifacts; however, this post-processing step may mask unphysical excursions and should be interpreted with caution.

3. RESULTS AND DISCUSSION

This section presents the results of the proposed GNN-ROM framework, beginning with experimental validation for hydraulic fracture proppant transport, followed by an assessment of generalization to unseen operating conditions for a fluidized bed gasifier.

3.1 Hydraulic Fracture Proppant Transport

The hydraulic fracture configuration considered in this study is based on the laboratory-scale fracture simulation framework reported by Zhou et al. [25], which simplifies the geometry of field-scale hydraulic fractures while retaining the dominant mechanisms governing proppant transport, settling, and bed formation. Experimental measurements from Ref. [25] are used exclusively for validation and are not incorporated during model training.

The fracture is represented as a rectangular slot with dimensions of $4000 \times 300 \times 10 \text{ mm}^3$ (length $\times$ height $\times$ width), with flow directed along the fracture length and injection occurring at the inlet located on the left boundary, as shown in Figure 3. A multi-hole entrance configuration is employed experimentally to promote uniform inflow distribution across the fracture height. In the fracture, injected proppant settles under gravity and is advected downstream, forming a bed whose height is inferred from the Eulerian particle volume fraction field. The working fluid is slick water with an effective viscosity of $2.5 \text{ mPa}\cdot\text{s}$, and 40/70 mesh ceramsite proppants are injected at an initial particle volume fraction of 6%. The particle size distribution follows an approximately Gaussian profile consistent with the experimental setup in Ref. [25].




FIGURE 3 SCHEMATIC OF PROPPANT FRACTURE MODEL

High-fidelity CFD simulations are performed using Barracuda Virtual Reactor (BVR). The fracture domain is discretized using a structured mesh comprising approximately 170,000 Eulerian control volumes ($470 \times 36 \times 10$), while the particulate phase is represented using approximately 9 million Lagrangian parcels. Interphase momentum exchange accounts for the non-Newtonian shear-thinning behavior of slick water through a modified Beetstra drag formulation adapted by

Chhabra [27]. In the present work, only Eulerian field outputs are retained for reduced-order modeling.

Specifically, the GNN-ROM is trained using the Eulerian particle volume fraction field as the sole input and output variable for the hydraulic fracture case.

The GNN-ROM is trained using Eulerian CFD snapshots generated at pumping velocities of 1.2, 1.3, 1.5, and 1.6 m/s spanning the operational envelope of the experimental system. Model predictions are evaluated at an intermediate pumping velocity of 1.4 m/s, using CFD and experimental measurements




FIGURE 4 COMPARISON OF PARTICLE VOLUME FRACTION AT A) T=300s, B) T=630s, AND C) T=1140s

as references.

Figure 4 compares instantaneous particle volume fraction fields predicted by the high-fidelity CFD simulations and the GNN-ROM at representative times of 300 s, 630 s, and 1140 s. The ROM accurately reproduces the global morphology of the proppant bed, including the formation of a densely packed region near the inlet and the downstream tapering of the bed along the fracture length. The predicted bed front location and slope remain in close agreement with the CFD reference across all times, indicating that the graph-based representation preserves the dominant spatial transport characteristics despite the significant reduction in computational complexity.

To enable direct comparison with experimental measurements, proppant bed height profiles are extracted from Eulerian particle volume fraction fields using a consistent threshold-based criterion applied uniformly to both CFD and GNN-ROM predictions. Figure 5 presents comparisons of bed height profiles against experimental measurements at the same time instances. The GNN-ROM reproduces both the peak bed height near the inlet and the downstream decay trend with high fidelity, closely matching the CFD reference and experimental data. Minor discrepancies are primarily observed in regions of steep bed gradients, which also correspond to increased experimental variability.

Quantitative accuracy is assessed using the mean absolute error (MAE), defined as the spatial average of the absolute difference between predicted and experimentally measured bed height along the fracture length at each time instance. The mean absolute error (MAE) is computed as in Eq. (8).

$$MAE(t) = \frac{1}{N_x} \sum_{i=1}^{N_x} |h_{pred}(x_i, t) - h_{exp}(x_i, t)|, \quad (8)$$

where h_{pred} denotes the bed height predicted by CFD or the GNN-ROM, h_{exp} denotes the experimentally measured bed height, and x_i are the discrete sampling locations along the fracture length. Table 1 summarizes the MAE values for both the

CFD and GNN-ROM predictions. Across all evaluated times, the GNN-ROM maintains MAE values below 0.02 m and remains comparable to the CFD reference. Slightly higher errors at later times reflect increased bed complexity as proppant accumulation progresses, but the errors remain within acceptable engineering tolerances.

Computational performance is summarized in Table 2. All training and inference computations are performed on an NVIDIA DGX-1 system equipped with Tesla V100-SXM2 GPUs (32 GB). While GNN-ROM training is performed offline,

inference achieves an approximately 82× reduction in wall-clock time for transient field prediction compared to full CFD simulations, while maintaining close agreement with experimental measurements.

Overall, these results demonstrate that the proposed Eulerian GNN-ROM accurately captures both the spatial structure and temporal evolution of proppant transport and bed formation in hydraulic fractures. The ability to generate full-field predictions in a single forward pass, combined with strong agreement with experimental data, highlights the potential of the framework for rapid evaluation of proppant transport behavior in design and operational studies.




FIGURE 5 COMPARISONS OF BED-HEIGHT VALUES WITH EXPERIMENTAL VALUES

TABLE 1 MEAN ABSOLUTE ERROR (MAE) OF PROPPANT BED HEIGHT PREDICTIONS FOR CFD AND GNN-ROM AT DIFFERENT TIMES

Time (s)	CFD MAE (m)	GNN-ROM MAE (m)
300	0.0082	0.0088
630	0.0080	0.0107
1140	0.0175	0.0195

TABLE 2 COMPUTATIONAL COST COMPARISON: CFD vs GNN-ROM


Task	Number of snapshots	Wall-Clock Time (s)
CFD Simulation	600	244,800
CFD Data Generation (training set)	2400	979,200
GNN-ROM training	2400	6051
GNN-ROM inference	600	3000

The CFD data generation cost for the training dataset is a one-time offline expense, not included in the GNN-ROM training and inference, and is amortized over repeated ROM deployment.

3.2 Fluidized Bed Gasifier: Generalization to Unseen Operating Conditions

The fluidized bed gasifier case is used to assess the ability of the proposed GNN-ROM framework to generalize across operating conditions in a complex, industrially relevant geometry. Unlike the hydraulic fracture configuration in Section 3.1, experimental measurements are not available for this case; therefore, high-fidelity CFD simulations generated using Barracuda Virtual Reactor (BVR) are treated as the reference solution. The underlying configuration and multiphase reacting-flow formulation are consistent with the Eulerian–Lagrangian gasification model reported by Snider et al. [26]. Although the computational mesh is structured, geometric asymmetry present within the reactor breaks translational invariance and induces strongly nonuniform flow, solids circulation, and reaction patterns, providing a stringent test for reduced-order modeling and motivating a graph-based formulation that readily extends to unstructured meshes.

The gasifier configuration is illustrated in Figure 6. Fluidizing air is introduced through a bottom distributor, while steam is injected through a multi-nozzle sparger located above the distributor. Coal is fed laterally into the reactor, and product gases exit through a cyclone assembly at the top. The baseline operating condition corresponds to a fluidizing air mass flow rate of 0.691 kg/s at 700 K with an inlet composition of 77% N₂ and 23% O₂ by mass. Steam is injected at 0.296 kg/s and 900 K through 204 nozzles. Heating coils maintain the reactor wall temperature at approximately 1300 K, and the outlet pressure is fixed at 200 kPa.

**FIGURE 6 SCHEMATIC OF GASIFIER**

The computational domain consists of 80,010 Eulerian control volumes, with the particulate phase represented by approximately 1.1 million Lagrangian parcels. Interphase momentum exchange is modeled using the Tang drag correlation [28], and gas–solid reactions are treated using a volume-averaged chemistry formulation including coal devolatilization, char oxidation, and steam gasification. These settings represent a realistic industrial gasifier environment with tightly coupled hydrodynamics and chemistry.

The GNN-ROM is trained using CFD data generated at four fluidizing air mass flow rates (0.691, 0.8, 0.9, and 1.0 kg/s). Model performance is evaluated at an intermediate and previously unseen operating condition corresponding to a fluidizing air mass flow rate of 0.85 kg/s. This configuration is not included in the training dataset and therefore provides a direct test of parametric generalization. Predictions are assessed relative to CFD using qualitative field comparisons and quantitative error metrics.

Across all Eulerian quantities considered, the GNN-ROM preserves the dominant spatial structures and transport pathways observed in the CFD reference. While fine-scale temporal fluctuations are partially smoothed, this primarily affects high-frequency components that are under-resolved in the training snapshots. In particular, fluctuations occurring at timescales smaller than the CFD snapshot interval ($\Delta t \approx 0.2$ s) are not explicitly captured by the model, but do not significantly affect time-averaged or large-scale transport behavior. The model consistently captures large-scale circulation, solids holdup, and reacting regions of engineering relevance. Normalized error metrics can appear elevated during strongly transient phases or in regions of near-zero concentration; therefore, spatial correlation and qualitative field agreement are evaluated alongside RMSE to assess predictive fidelity.

Time-averaged flow fields are of primary interest for engineering analysis of gasifiers, as they govern overall solids holdup, species conversion, and reactor performance. Figure 7 compares CFD and GNN-ROM predictions for time-averaged particle volume fraction, CH₄ mass fraction, and CO₂ mass fraction at the unseen operating condition. Despite the asymmetric geometry and strong coupling between gas–solid transport and reactions, the GNN-ROM accurately reproduces the dominant spatial patterns observed in the CFD solution.

In particular, the model captures the dense solids region in the lower reactor, the large-scale circulation induced by fluidization, and vertical stratification of reacting species. Predicted CH₄ and CO₂ distributions closely match the CFD reference, resolving regions of enhanced reaction activity and transport-driven redistribution along the reactor height. The preservation of these large-scale spatial trends indicates that the mesh-based graph representation effectively encodes local transport interactions and geometry effects, even in the absence of explicit reaction modeling within the ROM.




FIGURE 7 COMPARISON OF TIME-AVERAGED RESULTS BETWEEN CFD AND GNN-ROM

Quantitative error metrics for the time-averaged fields are summarized in Figure 8. The normalized root-mean-square error (RMSE) remains below approximately 7% for all quantities considered, while spatial correlation coefficients exceed 0.99. The normalized root-mean-square error (RMSE%) is defined as in Eq. (9),

$$RMSE\%(t) = \frac{\sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{u}_i(t) - u_i(t))^2}}{\sqrt{\frac{1}{N} \sum_{i=1}^N u_i(t)^2}} \times 100 \quad (9)$$

where $u_i(t)$ and $\hat{u}_i(t)$ denote the CFD and GNN-ROM predictions at mesh node i , respectively, and N is the total number of Eulerian control volumes. The spatial correlation is quantified using the Pearson correlation coefficient, computed between flattened CFD and GNN-ROM field vectors at each time instance as shown in Eq. (10):

$$\rho(t) = \frac{\sum_{i=1}^N (u_i(t) - \bar{u}(t))(\hat{u}_i(t) - \bar{\hat{u}}(t))}{\sqrt{\sum_{i=1}^N (u_i(t) - \bar{u}(t))^2} \sqrt{\sum_{i=1}^N (\hat{u}_i(t) - \bar{\hat{u}}(t))^2}} \quad (10)$$




FIGURE 8 ERROR METRICS FOR TIME-AVERAGED FIELDS

The transient predictive capability of the GNN-ROM is evaluated by comparing instantaneous Eulerian fields and time-resolved error metrics against the CFD reference. Figure 9, Figure 10 and Figure 11 present transient comparisons of particle volume fraction, CH₄ mass fraction, and CO₂ mass fraction at representative time instances spanning startup and statistically unsteady operation. The GNN-ROM accurately captures the

evolution of dense solids regions, large-scale circulation structures, and reacting zones throughout the transient.

Time-resolved correlation coefficients and normalized RMSE metrics are reported in Figure 12 and Figure 13. While RMSE% values for transient species fields can appear elevated, this behavior is expected for reacting multiphase systems characterized by large regions of near-zero concentrations, sharp spatial gradients, and highly unsteady startup dynamics. In such regimes, small absolute deviations or slight phase shifts in reaction fronts can produce disproportionately large normalized errors and reduced instantaneous correlation. During strongly transient phases, the spatial correlation coefficient decreases to approximately 0.7, reflecting rapid spatial reorganization of the flow. Despite this, the GNN-ROM continues to capture the dominant spatial structures and transport pathways throughout the transient. Correlation values are not reported for snapshots with negligible spatial variance.

In addition to the operating condition shown, the trained GNN-ROM was evaluated qualitatively at several additional inlet air flow rates within the training envelope, including 0.75, 0.83, and 0.97 kg/s. For these cases, the ROM predictions remained physically consistent and reproduced the dominant spatial trends observed in the corresponding CFD solutions. Detailed quantitative metrics for these cases are omitted for brevity, as the focus of the present study is on demonstrating generalization to a representative unseen operating condition.




FIGURE 9 TRANSIENT COMPARISONS OF PARTICLE VOLUME FRACTION BETWEEN CFD AND GNN-ROM




FIGURE 10 TRANSIENT COMPARISONS OF CH₄(g) BETWEEN CFD AND GNN-ROM




FIGURE 11 TRANSIENT COMPARISONS OF $\text{CO}_2(\text{g})$ BETWEEN CFD AND GNN-ROM




FIGURE 12 CORRELATION METRIC FOR GNN-ROM AND CFD RESULTS




FIGURE 13 RMSE PERCENTAGE FOR GNN-ROM AND CFD RESULTS

TABLE 3 COMPUTATIONAL COST COMPARISONS: CFD VS GNN-ROM

Case Description	Number of snapshots	CFD Runtime (s)	GNN-ROM Training Time (s)	GNN-ROM Inference Time (s)
Time-averaged evaluation	501	25,200	N/A	241
Transient training	2004	N/A	1245	N/A
Transient inference	501	25,200	N/A	280.5

For the gasifier case, as seen in Table 3, the GNN-ROM achieves approximately $105\times$ speedup for time-averaged field prediction and $90\times$ speedup for transient field prediction relative to the corresponding CFD simulations, with all training performed offline.

Because temporal dependence is introduced explicitly through Fourier feature embeddings, the GNN-ROM can also predict isolated snapshots at arbitrary time instances using a single forward pass. In contrast, CFD requires sequential time marching from the initial condition to the query time. For late-time snapshot queries, this results in effective speedups approaching two to three orders of magnitude. Predicting a single snapshot at a late time using CFD requires advancing the solution through the full transient ($\approx 25,200$ s wall-clock time), whereas the GNN-ROM requires approximately 1–10 s per query. This results in an effective speedup approaching three orders of magnitude for isolated snapshot queries. This comparison reflects inference cost only, with all training performed offline.

4. CONCLUSION

This work presented a graph neural network-based reduced-order modeling (GNN-ROM) framework for high-fidelity multiphase CFD simulations that operates directly on Eulerian meshes and enables single-shot, time-conditioned prediction of full-field flow quantities. By recasting CFD snapshots into graph representations that preserve mesh connectivity and incorporating explicit time conditioning via Fourier feature embeddings, the proposed approach avoids intrusive model reduction and sequential time integration while maintaining spatial fidelity.

The framework was validated on proppant transport in hydraulic fractures, where GNN-ROM predictions of Eulerian particle volume fraction and derived proppant bed height showed strong agreement with laboratory-scale experimental measurements. Quantitative error metrics demonstrated accuracy comparable to high-fidelity CFD while reducing inference time by more than an order of magnitude for full transient prediction, with substantially larger effective speedups for isolated snapshot queries.

Generalization capability was assessed on a fluidized bed gasifier characterized by complex geometry and reacting gas–solid flow. Trained on multiple inlet air flow rates, the GNN-ROM accurately reproduced both time-averaged and transient Eulerian fields at an unseen operating condition using high-fidelity Barracuda Virtual Reactor simulations as reference. These results demonstrate the ability of the proposed formulation to generalize across operating conditions in industrial-scale multiphase systems without reliance on modal projections or latent-state evolution.

A systematic ablation and data-efficiency study is deferred to future work; however, the present paper already includes generalization to an unseen operating condition. Overall, the proposed GNN-ROM combines mesh-aware spatial learning, explicit time conditioning, and non-intrusive training in

a unified framework well suited for applications requiring rapid field evaluation, including design exploration, parametric studies, and digital-twin deployment. Future work will focus on mesh-invariant generalization, physics-aware regularization, and coupled Eulerian–Lagrangian reduced-order modeling.

ACKNOWLEDGEMENTS

The authors thank Sam Clark for his assistance with computational resource configuration and Tecplot-based post-processing support.

REFERENCES

- [1] Anderson, T. B., and Jackson, R., 1967, “A Fluid Mechanical Description of Fluidized Beds,” *Ind. Eng. Chem. Fundam.*, **6**(4), pp. 527–539.
- [2] Gidaspow, D., 1994, *Multiphase Flow and Fluidization: Continuum and Kinetic Theory Descriptions*, Academic Press, San Diego, CA.
- [3] Patankar, N. A., and Joseph, D. D., 2001, “Modeling and Numerical Simulation of Particulate Flows by the Eulerian–Lagrangian Approach,” *Int. J. Multiphase Flow*, **27**(10), pp. 1659–1684.
- [4] Kunii, D. and Levenspiel, O. (1991) *Fluidization Engineering*. 2nd Edition, *Butterworth-Heinemann*.
- [5] Benner, P., Gugercin, S., and Willcox, K., 2015, “A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems,” *SIAM Rev.*, **57**(4), pp. 483–531.
- [6] Quarteroni, A., Manzoni, A., and Negri, F., 2016, *Reduced Basis Methods for Partial Differential Equations*, Springer, Cham, Switzerland.
- [7] Willcox, K., Ghattas, O., Heimbach, P., and Arnst, M., 2021, “The Role of Reduced-Order Models in Data Assimilation and Optimization,” *SIAM Rev.*, **63**(4), pp. 657–688.
- [8] Holmes, P., Lumley, J. L., and Berkooz, G., 1996, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, Cambridge University Press, Cambridge, UK.
- [9] Rowley, C. W., and Dawson, S. T. M., 2017, “Model Reduction for Flow Analysis and Control,” *Annu. Rev. Fluid Mech.*, **49**, pp. 387–417.
- [10] Noack, B. R., Morzyński, M., and Tadmor, G., 2011, *Reduced-Order Modelling for Flow Control*, Springer, Vienna, Austria.
- [11] Carlberg, K., Bou-Mosleh, C., and Farhat, C., 2011, “Efficient Nonlinear Model Reduction via a Least-Squares Petrov–Galerkin Projection,” *Int. J. Numer. Methods Eng.*, **86**(2), pp. 155–181.
- [12] Hesthaven, J. S., and Ubbiali, S., 2018, “Non-Intrusive Reduced Order Modeling of Nonlinear Problems Using Neural Networks,” *J. Comput. Phys.*, **363**, pp. 55–78.
- [13] Lee, K., and Carlberg, K. T., 2020, “Model Reduction of Dynamical Systems on Nonlinear Manifolds Using Deep Convolutional Autoencoders,” *J. Comput. Phys.*, **404**, p. 108973.
- [14] San, O., and Maulik, R., 2018, “Neural Network Closures for Turbulence Modeling,” *J. Comput. Phys.*, **366**, pp. 1–16.
- [15] Guo, X., Li, W., and Iorio, F., 2016, “Convolutional Neural Networks for Steady Flow Approximation,” *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.*, pp. 481–490.
- [16] Pathak, J., Wikner, A., Fussell, R., et al., 2018, “Hybrid Forecasting of Chaotic Processes: Using Machine Learning in Conjunction with a Knowledge-Based Model,” *Chaos*, **28**, 041101.
- [17] Sanchez-Gonzalez, A., Godwin, J., Pfaff, T., et al., 2020, “Learning to Simulate Complex Physics with Graph Networks,” *Proc. 37th Int. Conf. Mach. Learn. (ICML)*.
- [18] Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., and Battaglia, P. W., 2021, “Learning Mesh-Based Simulation with Graph Networks,” *9th International Conference on Learning Representations, ICLR 2021*.
- [19] Lino, M., Pfaff, T., Sanchez-Gonzalez, A., and Battaglia, P. W., 2022, “Learning Time-Dependent Partial Differential Equations with Graph Neural Networks,” *J. Comput. Phys.*, **450**, p. 110838.
- [20] Nikola Kovachki et al., 2023, “Neural Operator: Learning Maps Between Function Spaces with Applications to PDEs,” *JMLR* 24, no. 1: 1–97.
- [21] Raissi, M., Perdikaris, P., and Karniadakis, G. E., 2019, “Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations,” *J. Comput. Phys.*, **378**, pp. 686–707.
- [22] Andrews, M. J., and O’Rourke, P. J., 1996, “The Multiphase Particle-In-Cell (MP-PIC) Method for Dense Particulate Flows,” *Int. J. Multiphase Flow*, **22**(2), pp. 379–402.
- [23] Snider, D. M., 2001, “An Incompressible Three-Dimensional Multiphase Particle-in-Cell Model for Dense Particle Flows,” *J. Comput. Phys.*, **170**(2), pp. 523–549.
- [24] Snider, D. M., O’Rourke, P. J., and Andrews, M. J., 1998, “Sediment Flow in Inclined Vessels Calculated Using a Multiphase Particle-in-Cell Model,” *Int. J. Multiphase Flow*, **24**(8), pp. 1359–1382.
- [25] Zhou, H., Guo, J., Zhang, T., Li, M., Tang, T., and Gou, H., 2023, “Eulerian Multifluid Simulations of Proppant Transport with Different Sizes,” *Physics of Fluids*, **35**(4), 043314.
- [26] Snider, D. M., Clark, S. M., and O’Rourke, P. J., 2011, “Eulerian–Lagrangian Method for Three-Dimensional Thermal Reacting Flow with Application to Coal Gasifiers,” *Chemical Engineering Science*, **66**(6), pp. 1285–1295.
- [27] Chhabra, R. P., 1990, “Motion of Spheres in Power-Law (Viscoelastic) Fluids at Intermediate Reynolds Numbers: A Unified Approach,” *Canadian Journal of Chemical Engineering*, **68**(3), pp. 470–476.
- [28] Y. Tang, E. A. J. F. Peters, J. A. M. Kuipers, S. H. L. Kriebitzsch, and M. A. van der Hoef. “A new drag correlation from fully resolved simulations of flow past monodisperse static arrays of spheres”. *AIChE Journal*, **61**(2):688–698, 2015