Accelerating Barracuda Virtual Reactor Workflow with AI-Powered ROMs

Saurav Mitra

CPFD Software

June 25, 2026

What CPFD Is Building: ML-ROM Capability Layers

1 **Eulerian Field ROM**

Fast prediction of fixed-mesh fields such as particle volume fraction and gas species.

2 **Lagrangian Field ROM**

Particle/parcels-aware ROM for transient solids motion, bed evolution, and regime behavior.

3 **Parameter-Conditioned ROM**

Interpolation across operating conditions such as velocity, air flow, injection rate, and material settings.

4 **Interactive Engineering Layer**

Snapshot queries, what-if studies, design exploration, calibration, and digital twin support.

MAIN TAKEAWAY

CPFD is not replacing Barracuda. CPFD is building acceleration layers around Barracuda-generated physics.

Why Acceleration Matters in Industrial Multiphase CFD

Acceleration Drivers

Design Space Exploration
Evaluate more geometries, operating conditions, and physics models

Optimization & Calibration
Systematic parameter tuning across high-dimensional spaces

Digital Twins
Real-time virtual replicas requiring continuous simulation updates

Turnaround Time
Compress weeks of analysis into hours for faster decision-making

AI is Shifting Toward

Transient Trajectory Learning
Capturing time-dependent behaviors across simulation runs

Parameter + Time Conditioning
Models that generalize across operating conditions and timesteps

Industrial-Scale Focus
Moving beyond academic benchmarks to real-world complexity

AI as Augmentation
Enhancing — not replacing — physics-based simulation tools

Why ML for CFD is Hard to Trust — and What it Must Satisfy

Sources of Skepticism

Conservation + Stability Risk
Small errors can blow up over long transients or stiff regimes

Out-of-Distribution Behavior
Models interpolate well but extrapolate poorly without clear warning

Validation Burden
Trust requires reproducibility across operating envelopes, not single-case agreement

Deployment Constraints
Industrial workflows demand large-mesh compatibility, solver integration, and traceability

Engineering Requirements

Solver-Consistent Target
Learn the behavior of validated CFD — not replace governing equations

Predictable Failure Modes
Known envelopes and guardrails for safe deployment boundaries

Clear Validation Pathway
Metrics tied to decision variables (e.g., lift, pressure drop, etc.)

Non-Intrusive Integration
Drop-in inference on existing meshes and fields, preserving solver workflows

The objective is not solver replacement — it is accelerating validated CFD with controlled, deployable surrogates.

Method Tradeoffs Under Industrial-Scale Constraints

	Main Strength	Industrial Limitation	Why CPFD chose Graph-ROM
POD	Fast	Limited generalization	Useful baseline
PINNs	Physics residuals	Expensive at large mesh scale	Not deployment-ready yet
Neural Operators	Promising	Active research for complex industrial geometry	Future direction
Graph-ROM	Mesh-aware and scalable	Needs high-quality training data	Current practical path

GNN-ROM: Continuous-Time Graph Operator

TRANSIENT TRAINING DATA

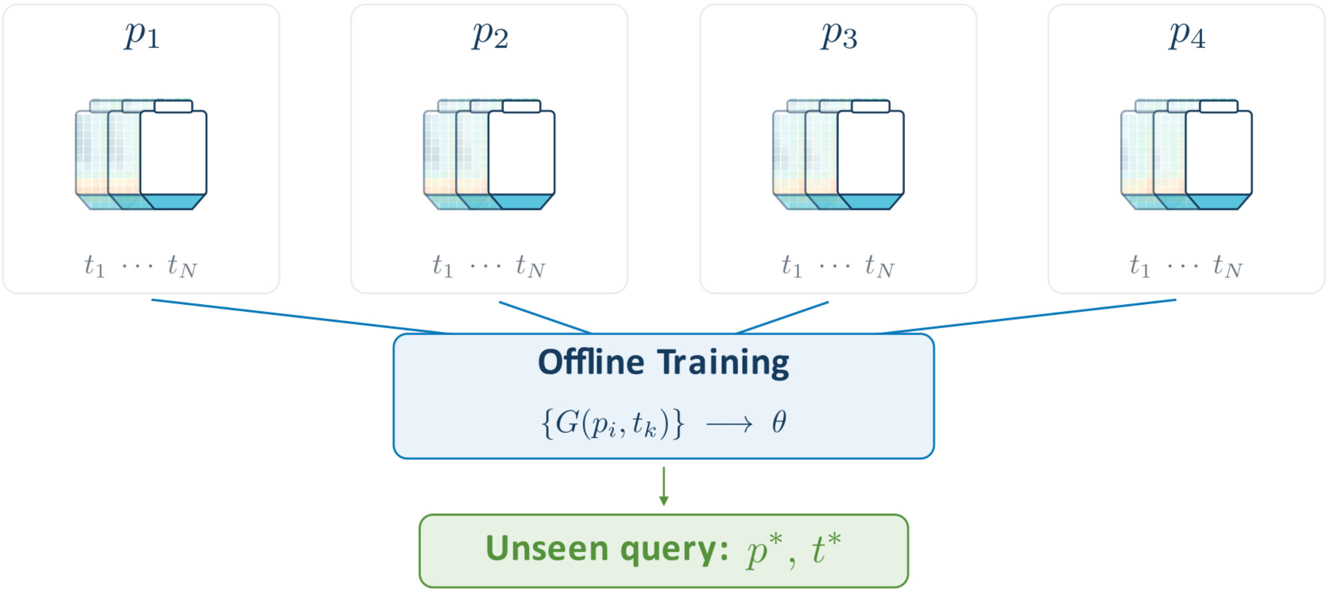

GNN-ROM: Fourier conditioned time

Application 1: Hydraulic Fracture (Validation)

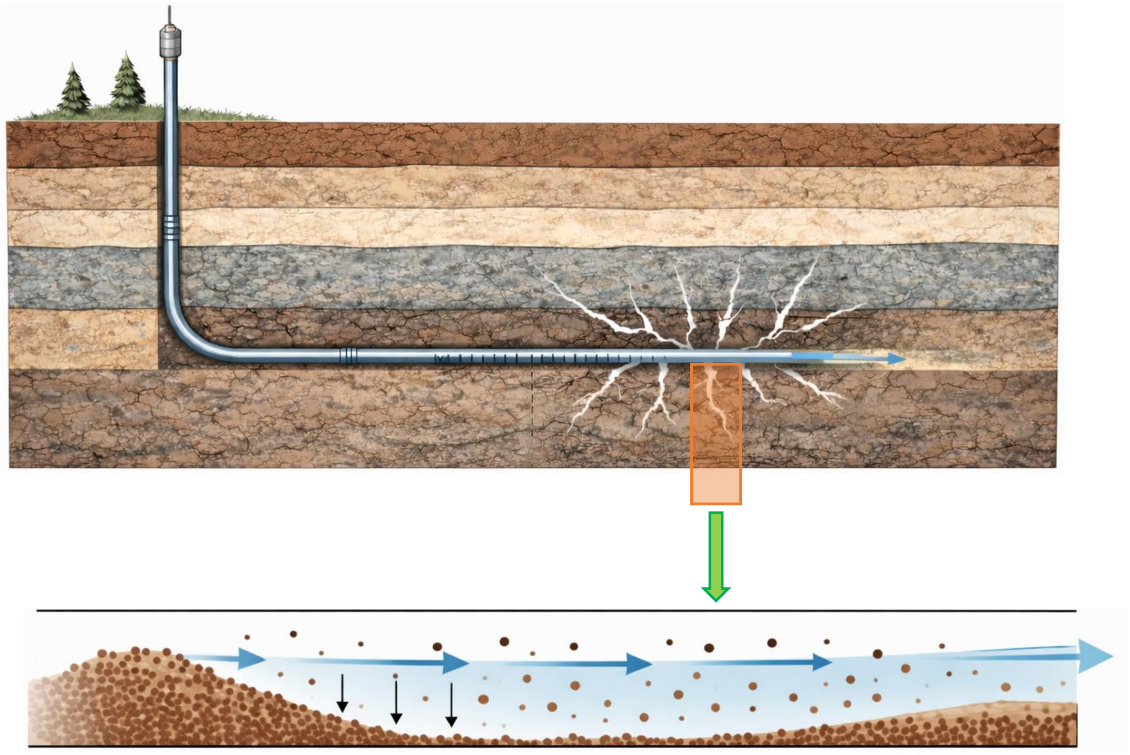

Gravity-driven settling + advection in a single fracture

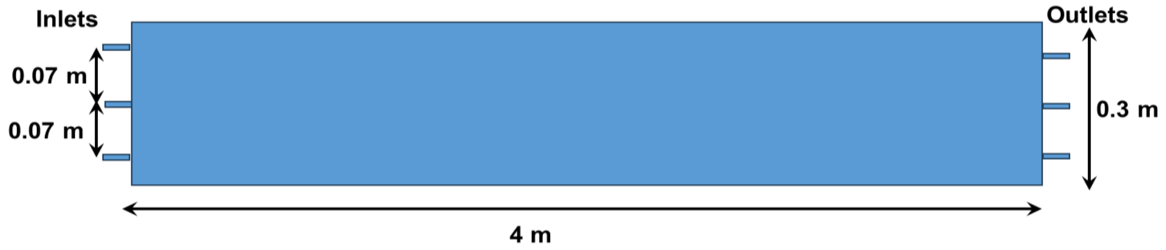

Laboratory Scale Experiment (Zhou et al., 2025)

Physical Configuration

Geometry
- 4.0 m x 0.3 m x 0.01 m
- Rectangular slot fracture

Fluid
- Slick water (~2.5 mPa-s)
- Shear-thinning behavior

Proppant
- 40/70 mesh ceramsite
- Initial particle vol frac = 6%
- Gaussian size distribution

Barracuda Virtual Reactor (CFD Model)
- ~170,000 Eulerian cells (470 x 36 x 10)
- 9 million Lagrangian parcels
- MP-PIC formulation
- Modified Beetstra drag (Chhabra adaptation)

ROM Scope
- Non-intrusive Eulerian ROM
- Input/Output: Particle volume fraction | Fixed mesh connectivity

Training Envelope

1.2 | 1.3 | 1.5 | 1.6

Inlet velocities (m/s)

Unseen Test Condition

1.4 m/s

Not used in training

1.2 — 1.3 — (1.4?) — 1.5 — 1.6

Application 1: Hydraulic Fracture (Validation)

CFD – Reference

GNN-ROM – Prediction

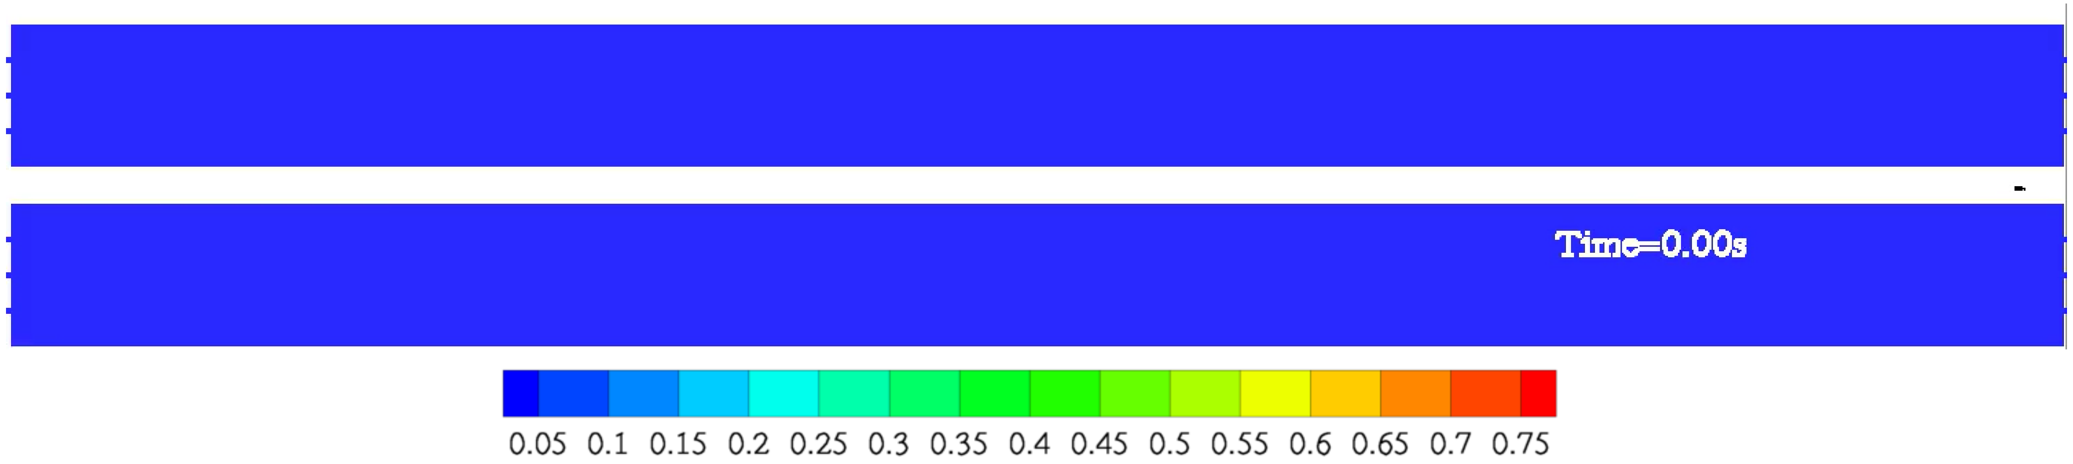

Particle volume fraction

$MAE_{bed\_height} < 0.02$ m

ROM tracks spatial deposition profile

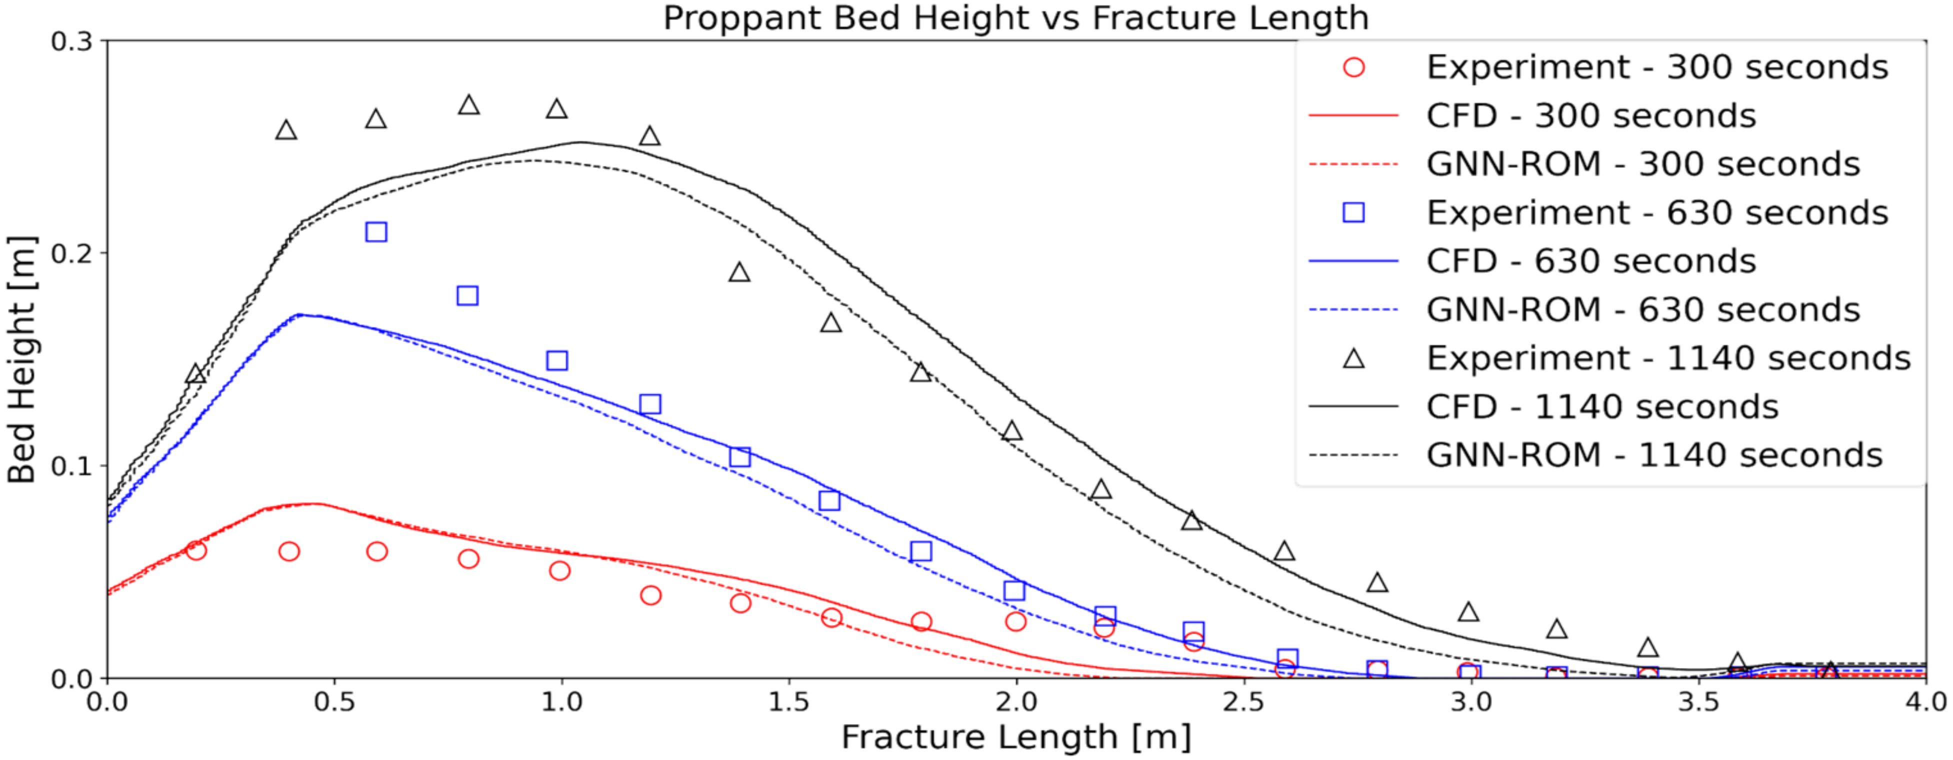

Geometry-Aware ROM: Multiple-Fracture Proppant Transport

Built on the PhysicsNeMo framework — joint research between CPFD and NVIDIA

ROM Capability Demonstrated

✓ Geometry-aware ROM prediction for multi-fracture proppant transport

✓ Trained across multiple fracture spacings: **3–30 m**

✓ Generalization across fracture spacing: **10 m and 18 m**

✓ CFD-like particle volume fraction fields

✓ Inference time: **< 1 second per query**

✓ Supports fast screening of alternative geometry layouts and flow conditions

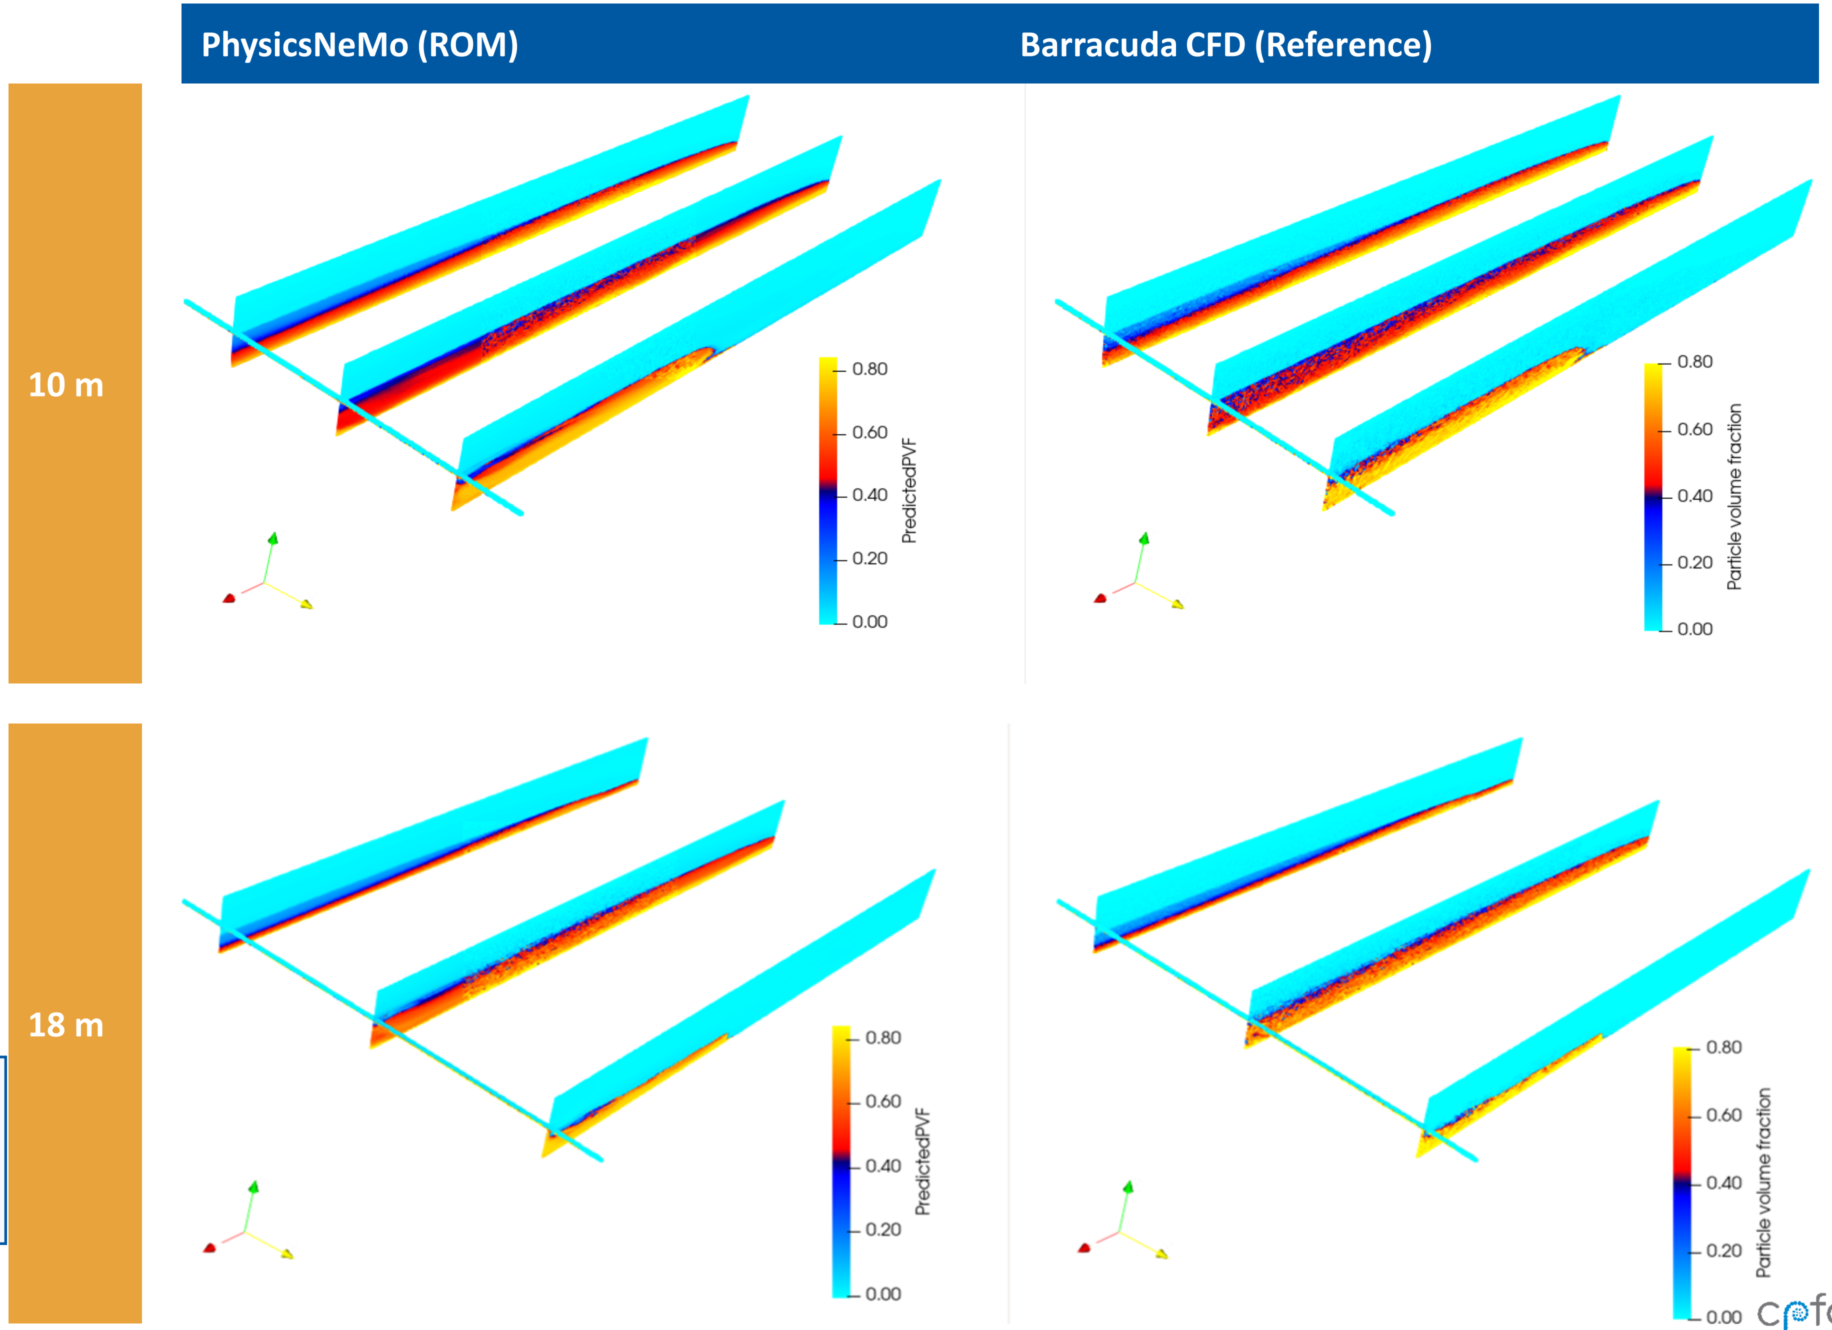

Collaboration with NVIDIA

PhysicsNeMo is NVIDIA's open-source physics-ML framework

Application 2: Fluidized-Bed Gasifier (Generalization)

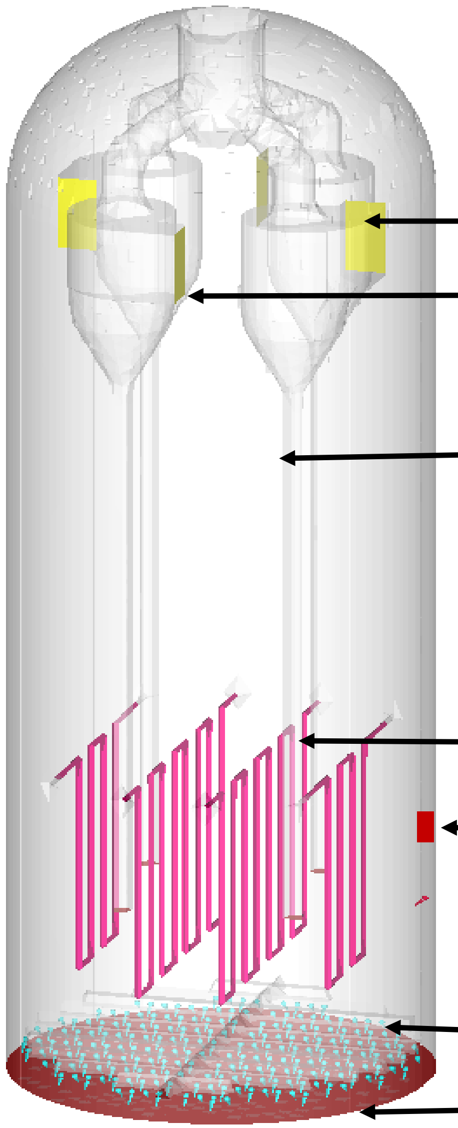

Baseline Operating Conditions	
Fluidizing Air Flow	0.691 kg/s
Air Temperature	700 K
Air Composition	77% N_2 / 23% O_2
Steam Flow	0.296 kg/s
Steam Temperature	900 K
Steam Injection	204 nozzles
Wall Temperature	~1300 K
Outlet Pressure	200 kPa

Barracuda Virtual Reactor
80,010 Eulerian control volumes
1.1 million Lagrangian parcels
MP-PIC formulation
Tang drag correlation
Volume-averaged chemistry

Chemistry: Devolatilization, Char oxidation, Steam gasification

Strongly coupled hydrodynamics + reactions + transport

Training Envelope
0.691 kg/s
0.8 kg/s
0.9 kg/s
1.0 kg/s

Fluidizing air mass flow rates

Unseen Test Condition

0.85 kg/s

Application 2: Fluidized-Bed Gasifier (Time-Average Fields)

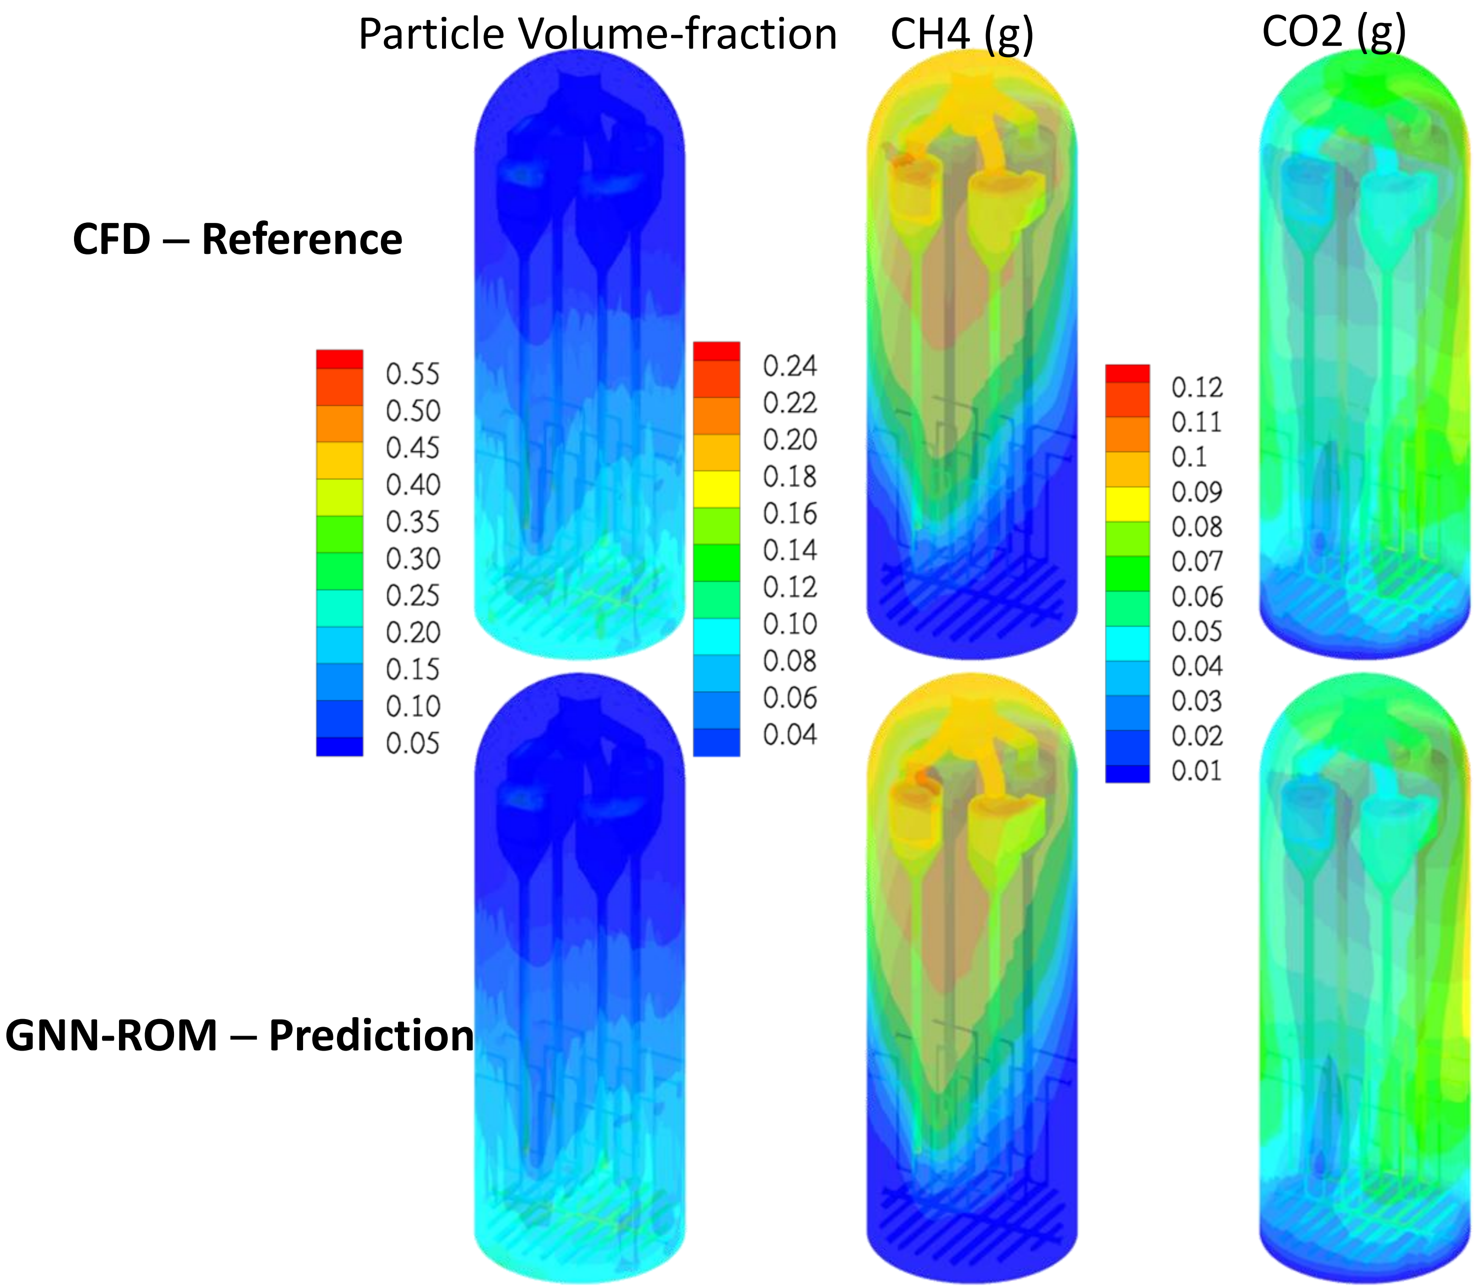

Unseen Operating Condition: Air Flow = 0.85 kg/s (Not Used in Training)

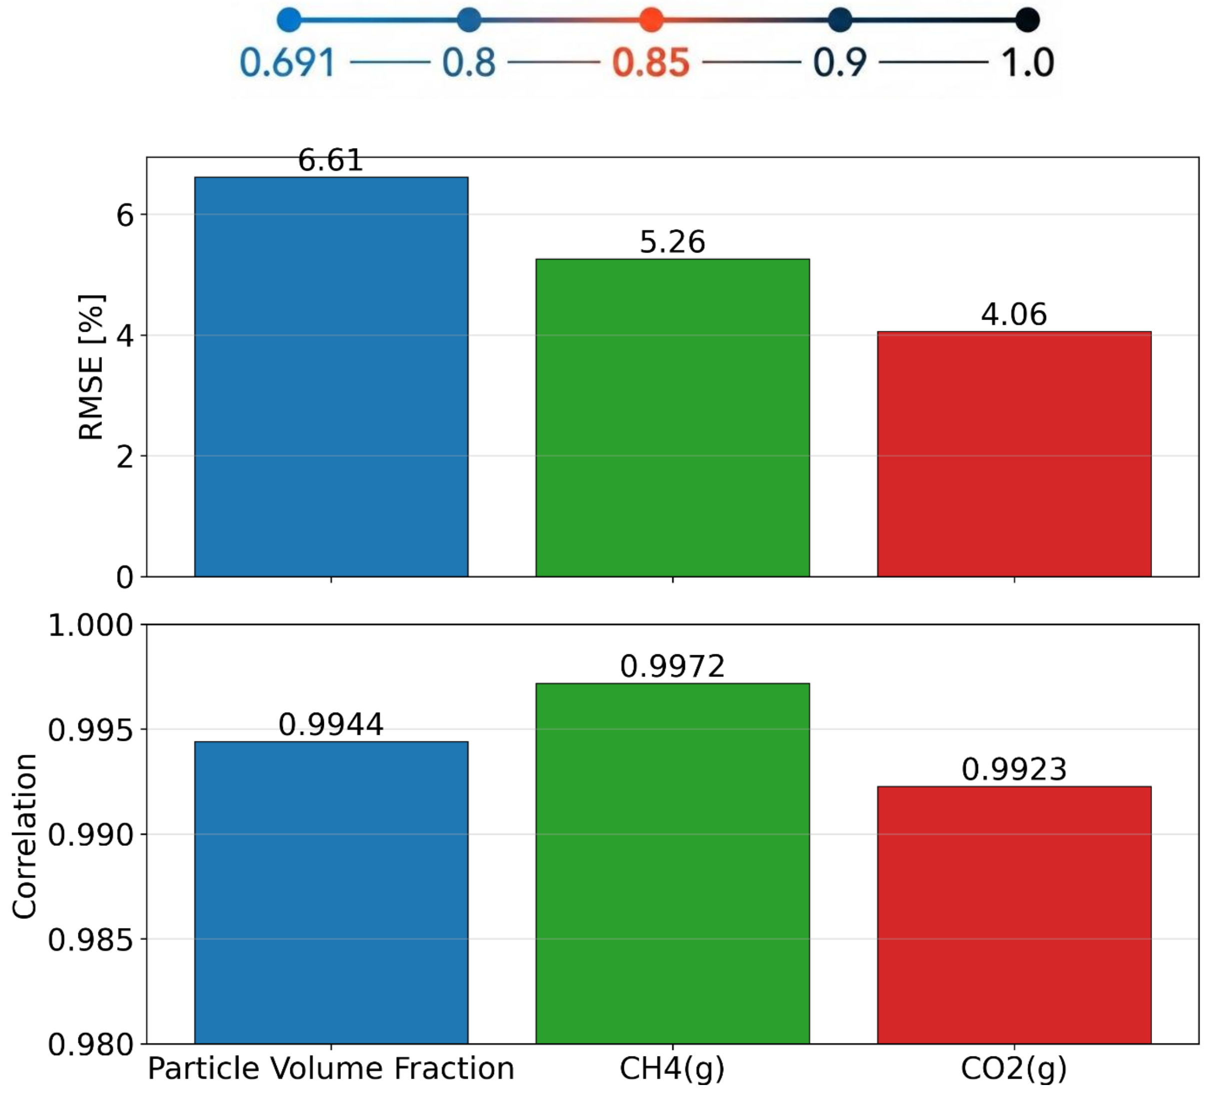

Generalizes to unseen operating condition with <7% RMSE and >0.99 spatial correlation

Application 2: Fluidized-Bed Gasifier (Transients)

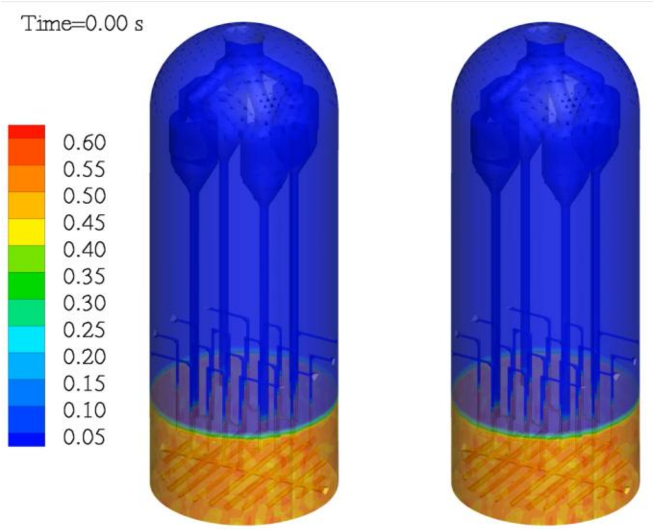

CFD – Reference GNN-ROM – Prediction

Particle Volume-fraction

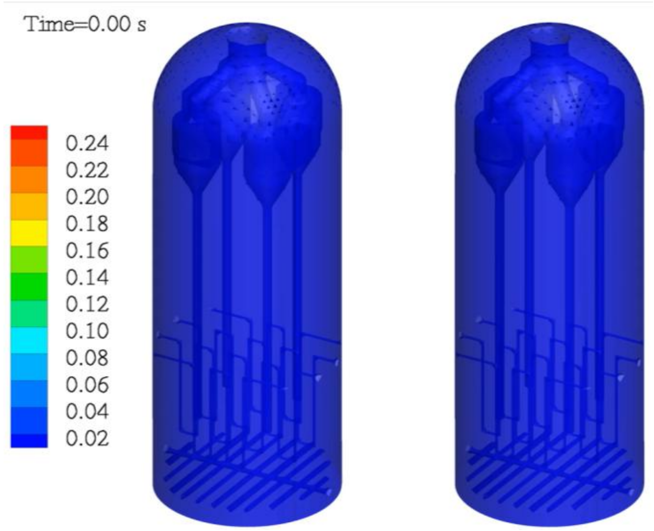

CFD – Reference GNN-ROM – Prediction

CH4 (g)

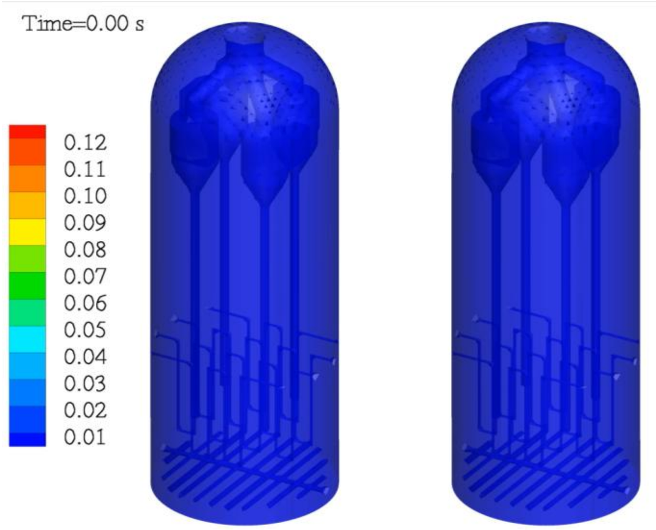

CFD – Reference GNN-ROM – Prediction

CO2 (g)

Unseen Operating Condition: Air Flow = 0.85 kg/s (Not Used in Training)

Key Observations:
Startup reorganization captured
Large-scale circulation preserved
Evolution of reaction fronts captured

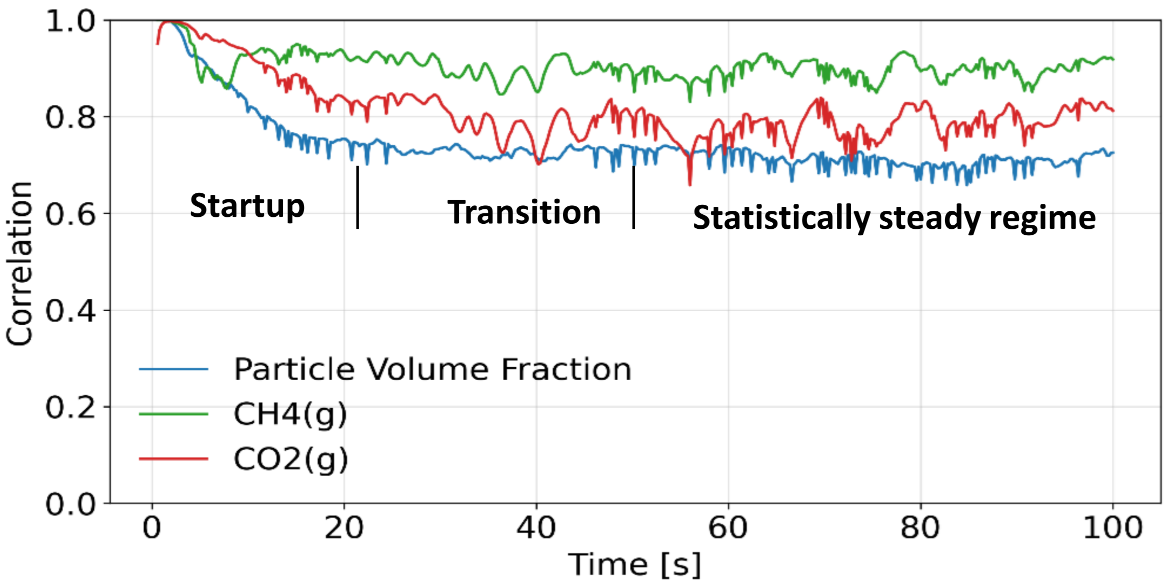

Why Lagrangian ROM Matters

Foundation: Barracuda already maps particles onto Eulerian cells → particle volume fraction & solids motion as field data

Eulerian GNN-ROM

Accelerates field prediction

✓ Speeds up Barracuda's **fixed-mesh field** predictions

✓ Returns **volume-fraction & velocity fields** on the Eulerian grid

✓ Familiar **mesh-aware** output format for CFD users

Lagrangian PointNet-GNN-ROM

Adds particle-cloud awareness — native to MP-PIC

✦ **Architecture:** PointNet captures global cloud state; GNN message passing captures local interactions

✦ **Unlocks engineering descriptors:** bed height & expansion, center of mass, solids spread, spout structure, segregation, regime-change indicators

✦ **Parameter-conditioned:** predicts cloud descriptors at unseen operating conditions

TAKEAWAY

Eulerian ROM accelerates field prediction. **Lagrangian ROM** turns particle-cloud evolution into fast engineering descriptors for fluidization, spouting & solids transport.

PointNet-GNN-ROM: Lagrangian Field

Application 3: Kuiper's Bed (Validation)

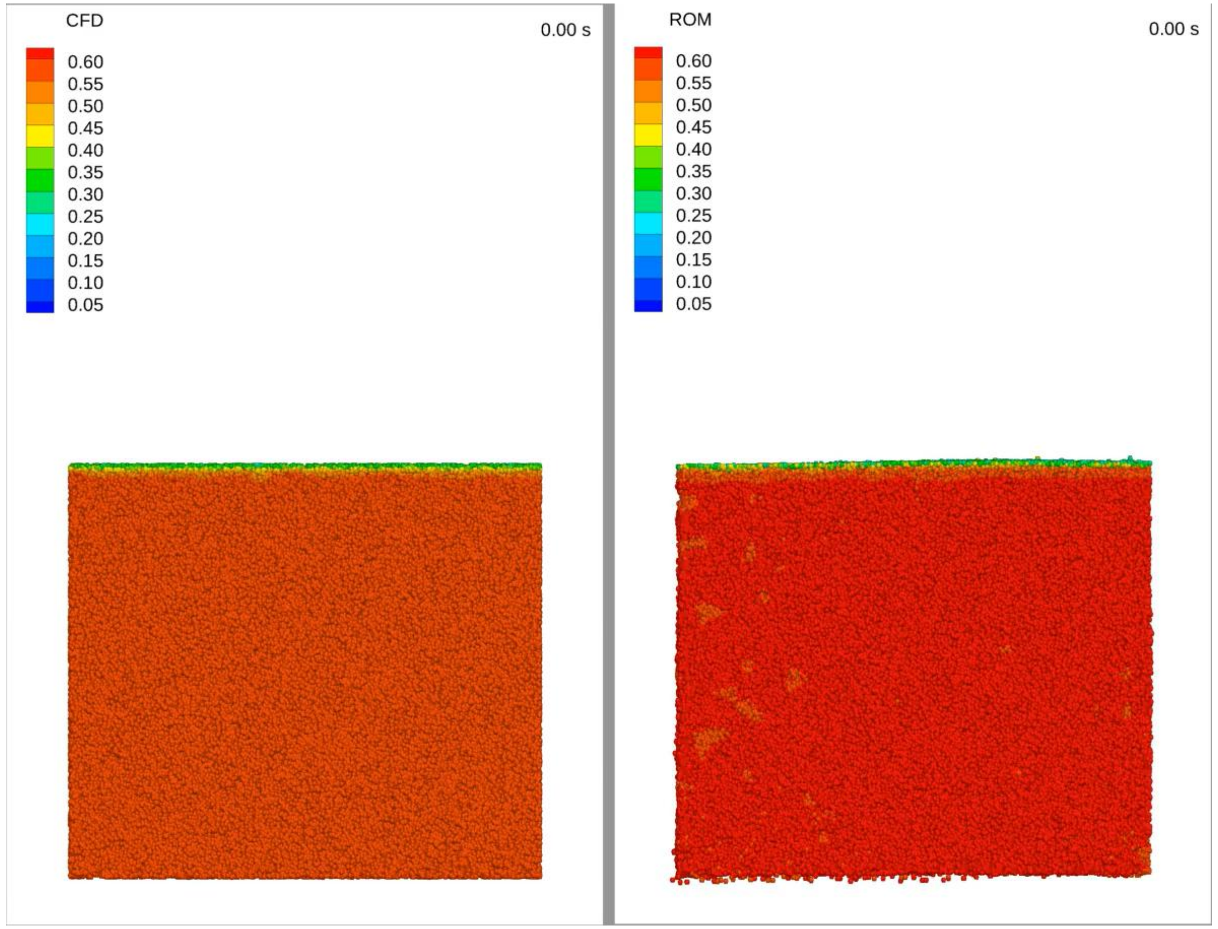

CFD – Reference **PointNet-GNN-ROM – Prediction**

Particle Volume-fraction

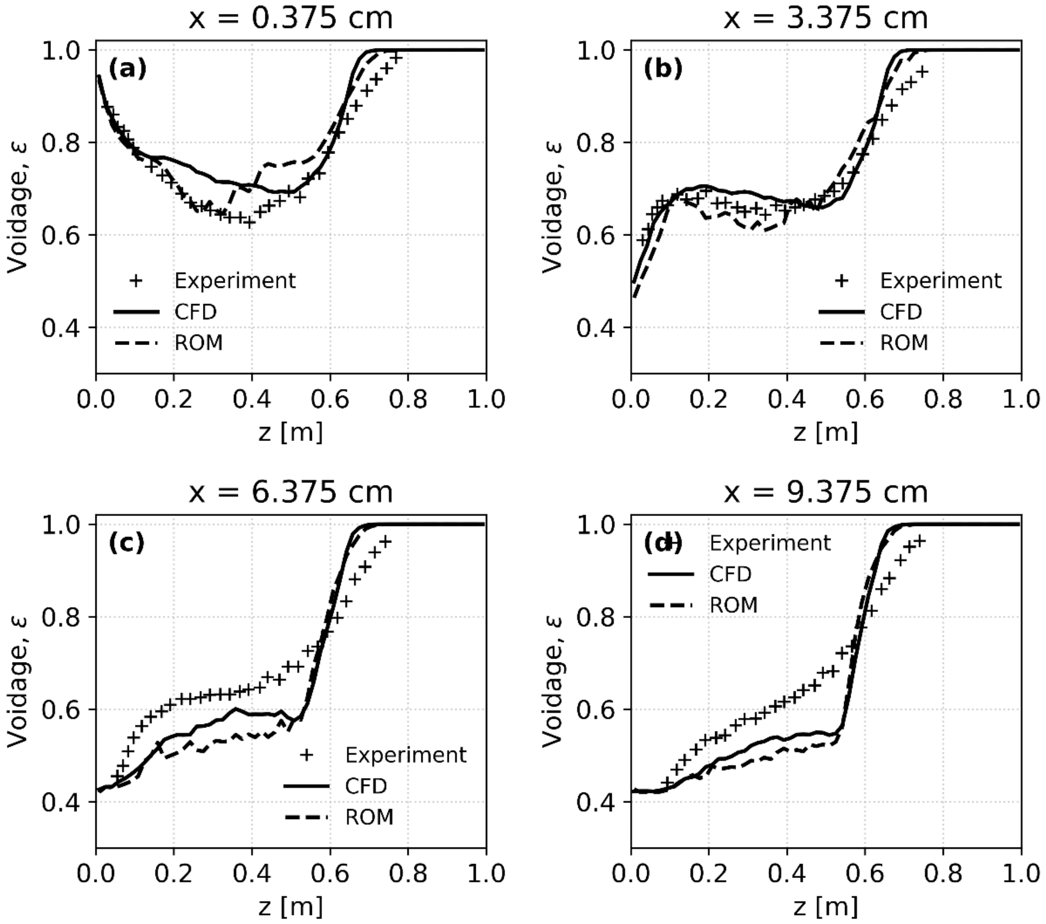

Training Envelope	Unseen Test Condition
8 m/s	10 m/s (Not used in training)
12 m/s	

Fluidizing air mass flow rates

PointNet-GNN-ROM: Lagrangian Field

Application 4: Spouted Bed (Testing Regime Change)

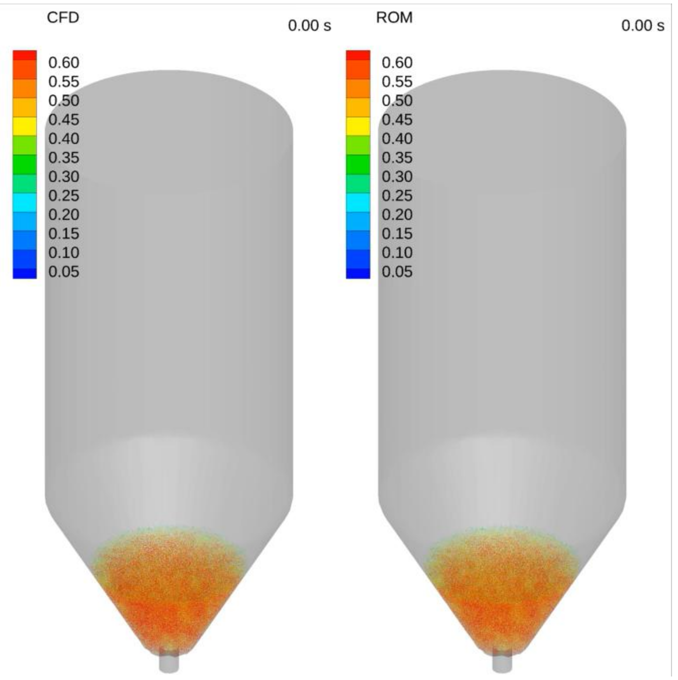

CFD – Reference PointNet-GNN-ROM – Prediction

Regime 1 – Center-jet velocity = 40 m/s

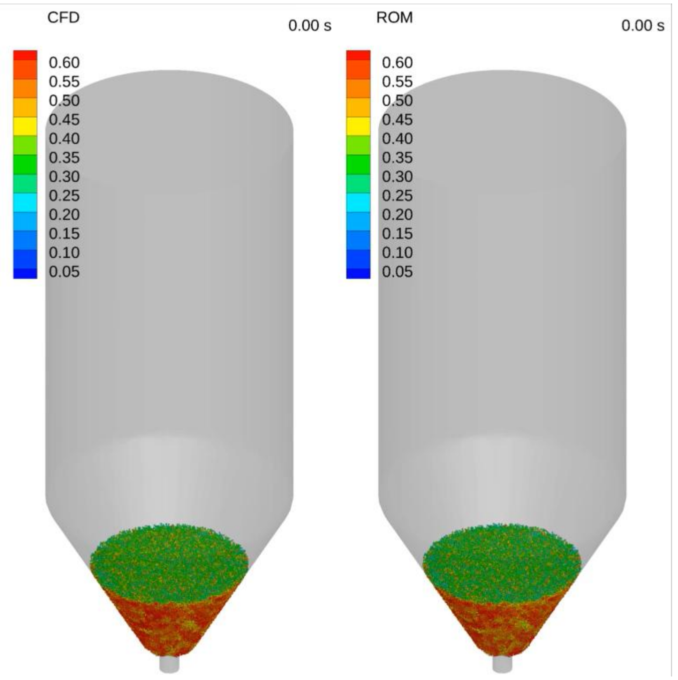

CFD – Reference PointNet-GNN-ROM – Prediction

Regime 2– Center-jet velocity = 52 m/s

Training Envelope	Unseen Test Condition
10 m/s	40 m/s (Not used in training)
30 m/s	52 m/s (Not used in training)
60 m/s	

Reduced Order Models Boost Performance & Efficiency

Kuiper's Bed

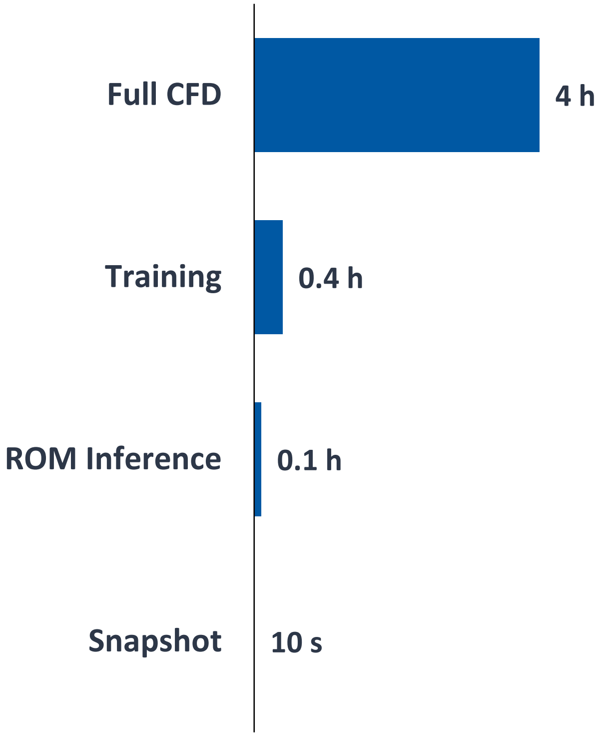

Spouted Beds

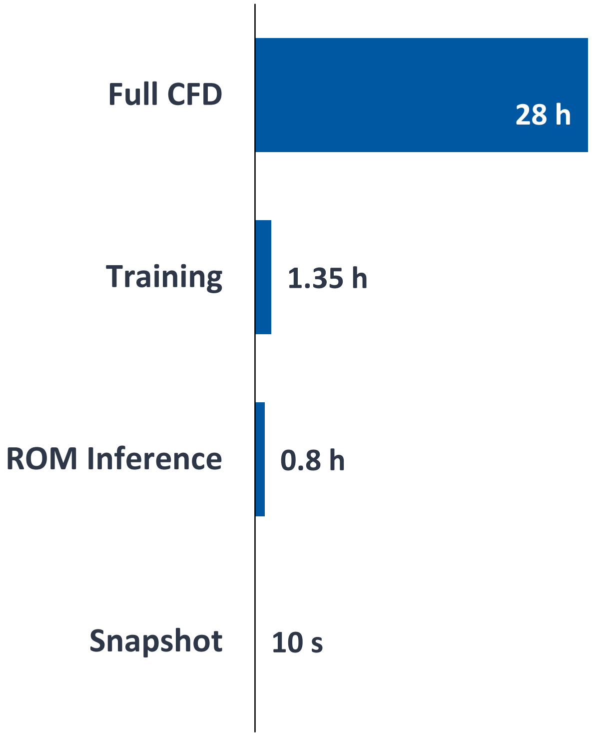

Hydraulic Fracture

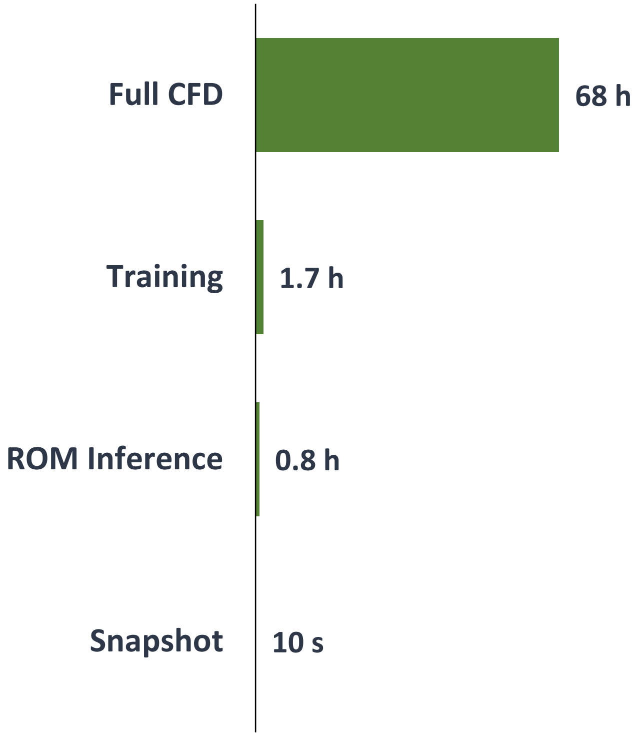

Fluidized-Bed Gasifier

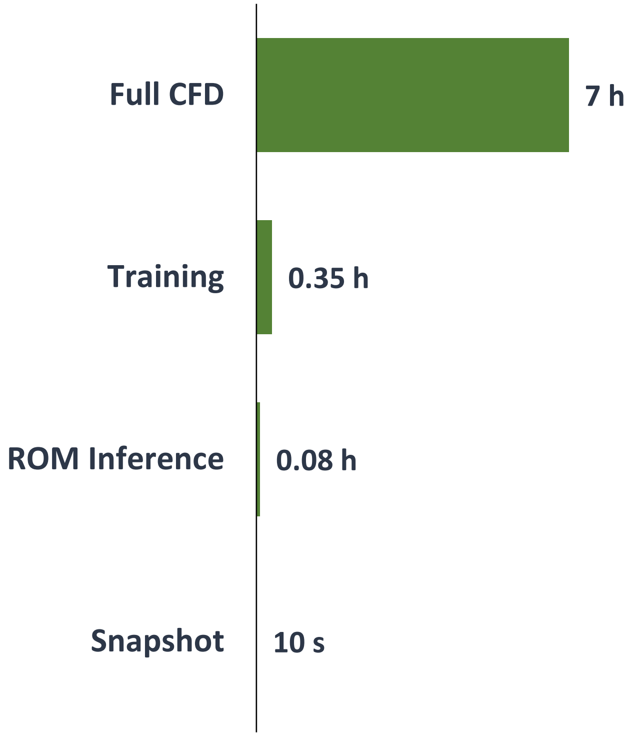

Key Findings

- ✓ **Eulerian ROM**: **80–90x Faster** than full CFD (Hydraulic Fracture, Fluidized-Bed Gasifier)
- ✓ **Lagrangian ROM**: **35–40x Faster** than full CFD (Kuiper's Bed, Spouted Beds)
- ✓**Real-time snapshot queries** (**<10 s**) on all geometries

GNN-ROM preserves MP-PIC physics while enabling deployment-scale transient prediction across diverse geometries.

Where CPFD's ROM Direction Differs

CPFD's differentiator: transient + multiphase + complex geometry ROMs built around validated Barracuda physics

Common AI-CFD demos	CPFD ROM direction
Slow training-data generation	**GPU-accelerated Barracuda data pipeline**
Broad generic training datasets	**Targeted ROMs with lower training-data burden**
Steady or time-averaged	**Transient evolution**
Single-phase or weak coupling	**Dense gas/liquid-solid multiphase**
Simple benchmark geometry	**Industrial reactor / fracture geometries**

BOTTOM LINE

This is not generic surrogate modeling. This is **ROM development for dense industrial multiphase CFD**.

From Validation to Deployment: What This Enables

Rapid Parametric Exploration

- Evaluate 10–100 operating points in hours, not weeks
- Interpolate across air flow / injection rates
- Identify optimal regimes

Model Calibration

- Match plant data via fast surrogate evaluation
- Tune drag, reaction, closure models
- Embed into solver workflows

Digital Twin Infrastructure

- Real-time field queries
- Snapshot <10 seconds
- 99% less energy per evaluation

Interactive Engineering

- What-if scenarios
- Startup transients
- Sensitivity studies

AI as acceleration layer — not solver replacement.

Looking Ahead

Release Timeline
Roadmap through Q4 2026

Features
- Eulerian field prediction: PVF, gas species, reacting-flow fields
- Lagrangian solids descriptors: bed height, COM, solids spread, spout structure
- Parameter-conditioned inference across operating envelopes
- Continuous-time snapshot queries
- ROM vs CFD validation reports
- Engineering metrics tied to design decisions
- What-if studies across flow conditions and geometries
- GUI-based ROM inference and visualization

Advanced Models
Collaboration with nVidia

- Geometry generalization
- More models via PhysicsNeMo
  - Transformers
  - Neural Operators
  - MeshGraphNet

NVIDIA®

Get In Touch
Connect with our team

Saurav Mitra
CPFD Software
cpfd-software.com

Request a Demo →