Simulating Complex Physics at Lightning Speed

Karen E. Willcox
Multicore World
Christchurch, New Zealand
February 18, 2026

ODEN INSTITUTE
FOR COMPUTATIONAL ENGINEERING & SCIENCES

TEXAS
The University of Texas at Austin

ODEN

Addressing science, engineering & medicine grand challenges

through **new methods, scalable algorithms** & **high-performance computing**

Advancing the architecture–algorithm–application nexus

to transform **next-generation computational science**

Building the mathematical & statistical foundations

for **predictive science, data science, AI** & **machine learning**

The Oden Institute advances the field of **computational science and engineering** and its application to address the **grand challenge problems facing society**.

ODEN

1973

2026

50+ Years of Leadership

in Interdisciplinary Research & Education in Computational Engineering & Sciences

- Mathematical & Statistical Foundations
- High Performance Computing
- Scientific Machine Learning
- Computational Biology
- AI for Science & Engineering
- Autonomy
- Computational Geosciences
- Computational Physics
- Optimization, Inversion & UQ
- Computational Medicine
- Computational Chemistry
- Numerical Analysis & Numerical PDEs
- Computational Materials
- Computational Engineering

25	Research Centers/Groups
27	Principal Faculty
125	Affiliated & Core Faculty
90	Research Staff/Postdocs
35	Administrative & IT Staff
105	CSEM Students
176	GRAs (72 CSEM)
14	Visiting Faculty Fellows
$130M	Active Research

Decades of high-impact advances based on modeling & simulation for complex systems

Forward Simulations
Advancing scientific discovery
& engineering innovation

Forward simulation has been the backbone of engineering analysis for many decades

Decades of high-impact advances based on modeling & simulation for complex systems

Forward Simulations
Advancing scientific discovery & engineering innovation

Optimization & Inverse Problems
Advancing estimation, design & control

Uncertainty Quantification
Sensitivity analysis, design under uncertainty, risk assessment

Machine Learning & AI
Learning from data, accelerating workflows, …

Use of computing for complex systems spans well <u>beyond forward simulation</u>.

Decades of high-impact advances based on modeling & simulation for complex systems

Forward Simulations
Advancing scientific discovery & engineering innovation

Optimization & Inverse Problems
Advancing estimation, design & control

Uncertainty Quantification
Sensitivity analysis, design under uncertainty, risk assessment

Machine Learning & AI
Learning from data, accelerating workflows, …

Use of computing for complex systems spans well beyond forward simulation.

But: (1) UQ, ML & AI still in their infancy;
(2) not yet seen full impact across the entire lifecycle.

Physics-based models are powerful and bring predictive capabilities

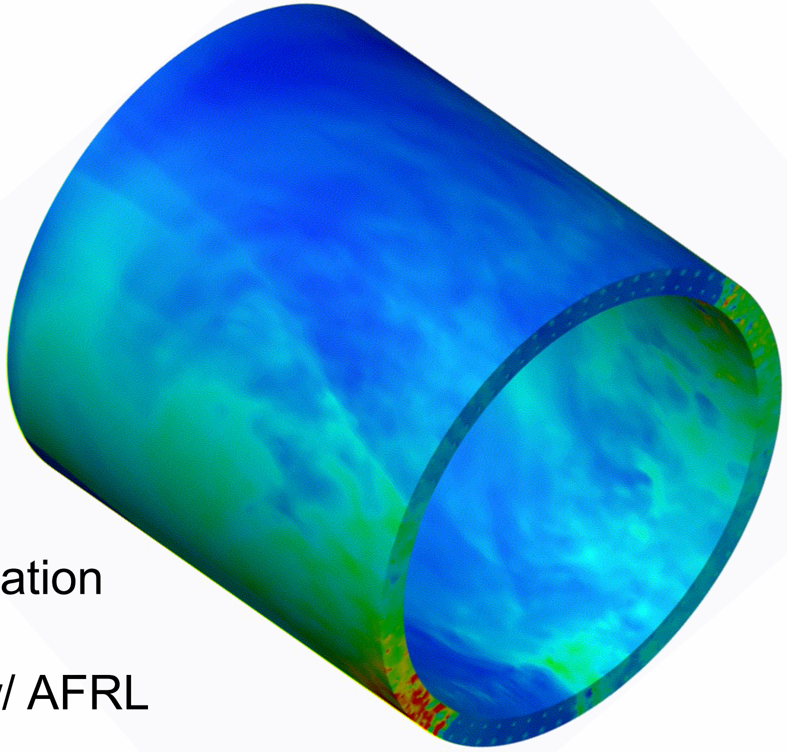

Rotating detonation
rocket engine
Ionut Farcas w/ AFRL

Large eddy simulation (LES) of
reactive Navier-Stokes equations
with 136M spatial dof and $\Delta t = 10^{-9}$

but they can be COMPUTATIONALLY PROHIBITIVE for design, control, UQ, or digital twins

Reduced-order models are critical enablers for engineering design, control, and UQ

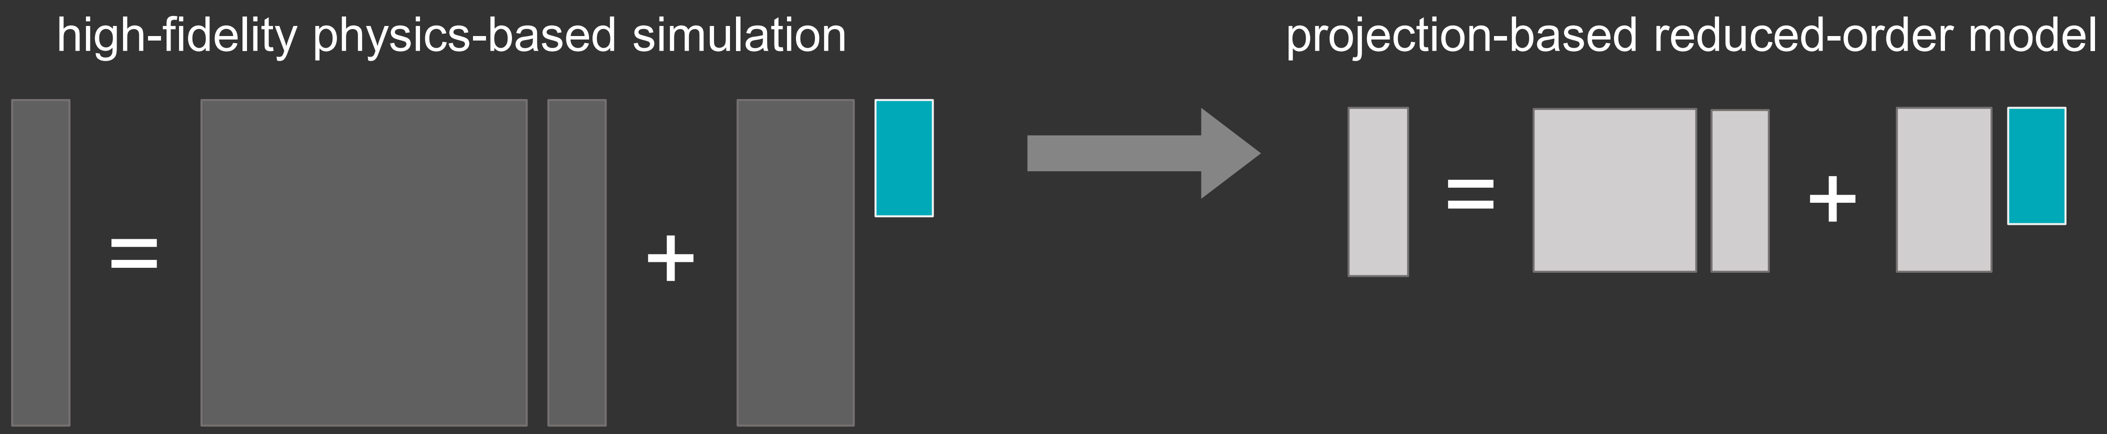

dimension $10^3 - 10^9$
solution time ~minutes / hours / days

dimension $10^1 - 10^3$
solution time ~seconds / minutes

1 **Train**: Solve PDEs to generate training data (snapshots)

2 **Identify structure**: Identify a low-dimensional manifold

3 **Reduce**: Project PDE model onto the low-dimensional manifold

Model reduction meets machine learning

Machine learning

“Machine learning (ML) is a field of study in artificial intelligence concerned with the development and study of statistical algorithms that can learn from data and generalize to unseen data and thus perform tasks without explicit instructions.” [Wikipedia]

Reduced-order modeling

“Model order reduction (MOR) is a technique for reducing the computational complexity of mathematical models in numerical simulations. As such it is closely related to the concept of metamodeling, with applications in all areas of mathematical modelling.” [Wikipedia]

ML surrogates vs. model reduction

Machine learning

"…statistical algorithms that can learn from data and generalize to unseen data and thus perform tasks without explicit instructions." [Wikipedia]

Reduced-order modeling

"Model order reduction (MOR) is a technique for reducing the computational complexity of mathematical models in numerical simulations." [Wikipedia]

Model reduction methods have grown from Computational Science, with focus on *reducing* high-dimensional models that arise from physics-based modeling, whereas machine learning has grown from Computer Science, with focus on *learning* models from black-box data streams.

[Swischuk et al., *Computers & Fluids*, 2019]

ML surrogates vs. model reduction

Machine learning

- Models may or may not generalize
- Large training data requirements
- Non-intrusive, portable & flexible
- Accessible & available
- Massive uptake outside the expert community

Reduced-order modeling

- Deep body of theory & methods, (stability, structure preservation, error estimators, …)
- Highly expert community
- Methods are inaccessible & intrusive
- Limited uptake outside the expert community

We aim to blend the predictive power of physics-based methods & the speed of ML

Define the **structure of the reduced model**

$$\dot{\hat{\mathbf{x}}} = \widehat{\mathbf{A}}\hat{\mathbf{x}} + \widehat{\mathbf{B}}\mathbf{u} + \widehat{\mathbf{H}}(\hat{\mathbf{x}} \otimes \hat{\mathbf{x}})$$

Non-intrusive learning by inferring reduced operators from simulation data [Peherstorfer & W., 2016]

ÔpInf

Reduced-order models are solved in <1sec, bringing **physics predictive power** off the designer's supercomputer and **into the operational world.**

Available online at www.sciencedirect.com
ScienceDirect

Comput. Methods Appl. Mech. Engrg. 306 (2016) 196–215

Computer methods in applied mechanics and engineering
www.elsevier.com/locate/cma

Data-driven operator inference for nonintrusive projection-based model reduction

Benjamin Peherstorfer*, Karen Willcox

Department of Aeronautics & Astronautics, MIT, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

Received 7 September 2015; received in revised form 2 February 2016
Available online 6 April 2016

Abstract

This work presents a nonintrusive projection-based model reduction approach for full models based on time-dependent partial differential equations. Projection-based model reduction constructs the operators of a reduced model by projecting the equations of the full model onto a reduced space. Traditionally, this projection is intrusive, which means that the full-model operators are required either explicitly in an assembled form or implicitly through a routine that returns the action of the operators on a given vector; however, in many situations the full model is given as a black box that computes trajectories of the full-model states and outputs for given initial conditions and inputs, but does not provide the full-model operators. Our nonintrusive operator inference approach infers approximations of the reduced operators from the initial conditions, inputs, trajectories of the states, and outputs of the full model, without requiring the full-model operators. Our operator inference is applicable to full models that are linear in the state or have a low-order polynomial nonlinear term. The inferred operators are the solution of a least-squares problem and converge, with sufficient state trajectory data, in the Frobenius norm to the reduced operators that would be obtained via an intrusive projection of the full-model operators. Our numerical results demonstrate operator inference on a linear climate model and on a tubular reactor model with a polynomial nonlinear term of third order.

Keywords: Nonintrusive model reduction; Data-driven model reduction; Black-box full model; Inference

1. Introduction

Model reduction seeks to construct reduced models that provide accurate approximations of the full model solutions with orders of magnitude reduction in computational complexity. We consider here projection-based model reduction

The Operator Inference problem

Given (1) a physical/natural system with known governing equations, and (2) a set of data in the form of state snapshots (experimental or simulation)

Infer a reduced-order model that recovers the given data and provides a predictive capability to rapidly simulate unseen conditions

$$\min_{\mathbf{O}} ||\mathbf{DO} - \mathbf{R}||$$

$\mathbf{O}$: low-dimensional operators that define the reduced model

$\mathbf{D}, \mathbf{R}$: data matrix / forcing from simulation and/or experimental data

We use:

- the **physics** to define the structured form of the model we seek
- **projection-based model reduction** to cast the inference in a reduced coordinate space and to provide error estimates in some settings
- **inverse theory** to analyze the structure of the resulting problem and treat it numerically
- **numerical linear algebra** to achieve efficient scalable algorithms

The Operator Inference ROM form is defined by projection-based reduction theory

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}$$

Approximate
$\mathbf{x} \approx \mathbf{V}\hat{\mathbf{x}}$
$\mathbf{V} \in \mathbb{R}^{N \times r}$

Residual: N eqs $\gg r$ dof

$$\mathbf{r} = \mathbf{V}\dot{\hat{\mathbf{x}}} - \mathbf{A}\mathbf{V}\hat{\mathbf{x}} - \mathbf{B}\mathbf{u}$$

Project
$\mathbf{W}^\top \mathbf{r} = 0$
(Galerkin: $\mathbf{W} = \mathbf{V}$)

$$\dot{\hat{\mathbf{x}}} = \hat{\mathbf{A}}\hat{\mathbf{x}} + \hat{\mathbf{B}}\mathbf{u}$$

Full-order model (FOM)
state $\mathbf{x} \in \mathbb{R}^N$
input $\mathbf{u} \in \mathbb{R}^{N_i}$

Inside-Out View

Reduced-order model (ROM)
state $\hat{\mathbf{x}} \in \mathbb{R}^r$, $r \ll N$

In classical projection approaches:

$$\hat{\mathbf{A}} = \mathbf{V}^\top \mathbf{A}\mathbf{V}$$
$$\hat{\mathbf{B}} = \mathbf{V}^\top \mathbf{B}$$

OPERATOR INFERENCE reflects the structure of a physics-based model to learn reduced models from data

full-order model (FOM):

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}$$

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{H}(\mathbf{x} \otimes \mathbf{x}) + \mathbf{B}\mathbf{u}$$

represent high-dimensional state $\mathbf{x} \in \mathbb{R}^N$ in a low dimensional basis $\mathbf{V} \in \mathbb{R}^{N \times r}$: $\mathbf{x} \approx \mathbf{V}\mathbf{x}_r$

reduced-order model (ROM):

$$\dot{\hat{\mathbf{x}}} = \hat{\mathbf{A}}\hat{\mathbf{x}} + \hat{\mathbf{B}}\mathbf{u}$$

$$\dot{\hat{\mathbf{x}}} = \hat{\mathbf{A}}\hat{\mathbf{x}} + \hat{\mathbf{H}}(\hat{\mathbf{x}} \otimes \hat{\mathbf{x}}) + \hat{\mathbf{B}}\mathbf{u}$$

The ROM form is inspired by classical intrusive physics-based model reduction, but the operators are learned directly from data

Our Operator Inference approach blends model reduction & machine learning

Define the **structure of the reduced model**

$$\dot{\hat{\mathbf{x}}} = \hat{\mathbf{A}}\hat{\mathbf{x}} + \hat{\mathbf{B}}\mathbf{u} + \hat{\mathbf{H}}(\hat{\mathbf{x}} \otimes \hat{\mathbf{x}})$$

Non-intrusive learning by inferring reduced operators from simulation data [Peherstorfer & W., 2016]

snapshots generated from projected simulation data

low-dimensional operators define the reduced model as a dynamical system

$$\min_{\hat{\mathbf{A}},\hat{\mathbf{B}},\hat{\mathbf{H}}} \left\| \hat{\mathbf{X}}^\top \hat{\mathbf{A}}^\top + \left(\hat{\mathbf{X}} \otimes \hat{\mathbf{X}}\right)^\top \hat{\mathbf{H}}^\top + \mathbf{U}^\top \hat{\mathbf{B}}^\top - \dot{\hat{\mathbf{X}}}^\top \right\|$$

minimum residual formulation leads to linear least squares

- For Markovian data, OpInf has preasymptotic **recovery of intrusive ROM** [Peherstorfer, 2020]
- Projection-based model reduction **preserves physics structure** by construction
- Projection-based reduced models are **interpretable** (evolution of modal coordinates)

Our Operator Inference approach blends model reduction & machine learning

Define the **structure of the reduced model**

$$\dot{\hat{\mathbf{x}}} = \hat{\mathbf{A}}\hat{\mathbf{x}} + \hat{\mathbf{B}}\mathbf{u} + \hat{\mathbf{H}}(\hat{\mathbf{x}} \otimes \hat{\mathbf{x}})$$

Non-intrusive learning by inferring reduced operators from simulation data [Peherstorfer & W., 2016]

- OpInf for non-polynomial nonlinear systems [Qian et al., 2020, Benner et al., 2020]
- Regularization is key [McQuarrie, Huang & W., 2021]
- OpInf formulation in PDE setting [Qian, Farcas & W., 2022]
- Parametric OpInf [McQuarrie, Khodabakhshi & W., 2023]
- Bayesian OpInf [Guo, McQuarrie & W., 2022]
- Quadratic Manifold OpInf [Geelen et al., 2023, 2024]
- Distributed OpInf [Farcas et al., 2024, 2025]
- Block-structured Opinf [Zastrow et al., 2025]
- Nested OpInf [Aretz & W., 2026]

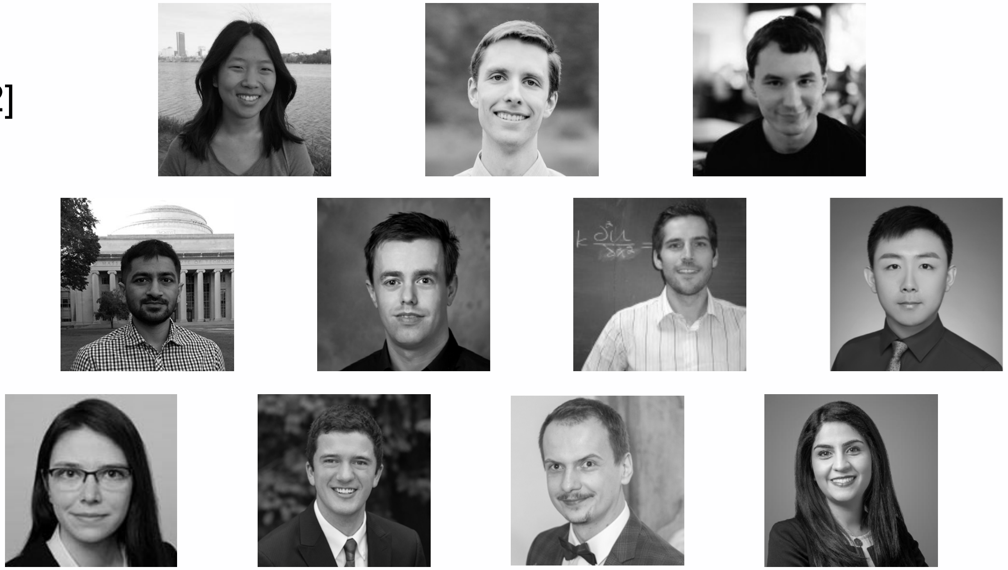

Modeling the combustion chamber of a rotating detonation rocket engine

- LES simulations of the reactive, viscous 3D Navier-Stokes equations
- Skeletal chemistry mechanism based on the Foundational Fuel Chemistry Model (FFCMy-12)
- Non-premixed fuel injection (gaseous methane and oxygen)

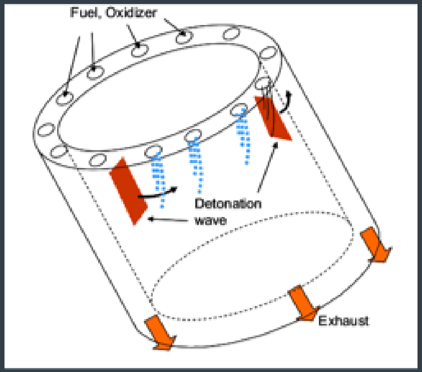

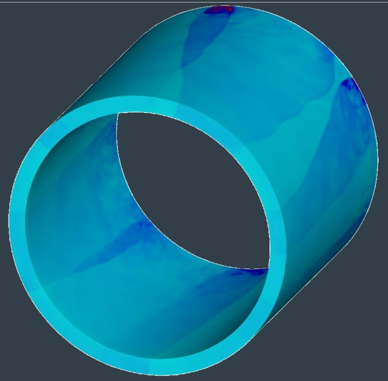

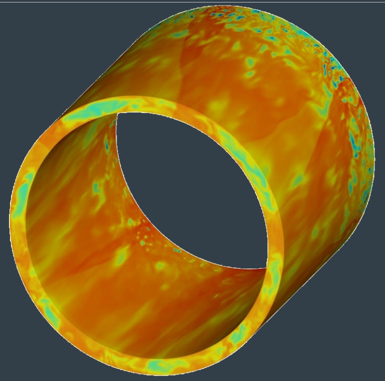

$$\frac{\partial}{\partial t}\begin{bmatrix}\rho\\ \rho v_x\\ \rho v_y\\ \rho E\\ \rho Y_1\\ \vdots\\ \rho Y_{n_{\mathrm{sp}}}\end{bmatrix}+\nabla\cdot\left(\begin{bmatrix}\rho v_x\\ \rho v_x^2+p\\ \rho v_x v_y\\ \rho v_x E+p v_x\\ \rho v_x Y_1\\ \vdots\\ \rho v_x Y_{n_{\mathrm{sp}}}\end{bmatrix}\vec{i}+\begin{bmatrix}\rho v_y\\ \rho v_x v_y\\ \rho v_y^2+p\\ \rho v_y E+p v_y\\ \rho v_y Y_1\\ \vdots\\ \rho v_y Y_{n_{\mathrm{sp}}}\end{bmatrix}\vec{j}-\begin{bmatrix}0\\ \tau_{xx}\\ \tau_{yx}\\ \tau_{xx}v_x+\tau_{yx}v_y-j_x^q\\ -j_{1,x}^m\\ \vdots\\ -j_{n_{\mathrm{sp}},x}^m\end{bmatrix}\vec{i}-\begin{bmatrix}0\\ \tau_{xy}\\ \tau_{yy}\\ \tau_{xy}v_x+\tau_{yy}v_y-j_y^q\\ -j_{1,y}^m\\ \vdots\\ -j_{n_{\mathrm{sp}},y}^m\end{bmatrix}\vec{j}\right)=\begin{bmatrix}0\\ 0\\ 0\\ 0\\ \dot{\omega}_1\\ \vdots\\ \dot{\omega}_{n_{\mathrm{sp}}}\end{bmatrix}$$

Modeling the combustion chamber of a rotating detonation rocket engine

1. Training data & data transformations

- LES simulation 136M spatial dof for full RDRE, timestep $\Delta t = 10^{-9}$s, ~1M CPU hours per 1ms on 16K cores
- Available data: 501 down-sampled combustion chamber snapshots over [2.50, 3.75] ms (~4 periods of two-wave system) interpolated onto structured grid with 4.2M dof
- 18 transformed state variables (+ scale & center): specific volume, pressure, 3D velocity, temperature, 12 species mass fractions (full chemistry data)
- Transformed training data: snapshots $\mathbf{X} \in \mathbb{R}^{76M \times 375}$

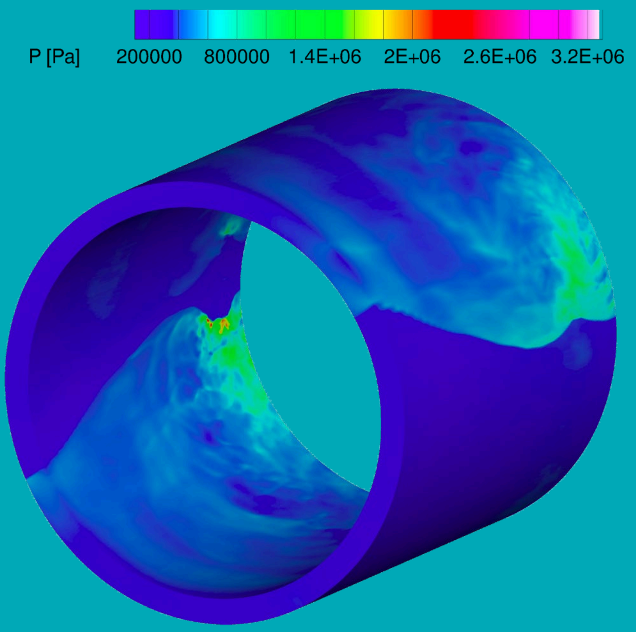

Two dominant co-rotating waves in the quasi-limit-cycle behavior of the flow

Modeling the combustion chamber of a rotating detonation rocket engine

2. Compute low-dimensional representation

- Snapshot matrix of transformed variables

$$\mathbf{X} \in \mathbb{R}^{76\mathrm{M}\times 375}$$

$$\hat{\mathbf{X}} = \mathbf{U}_r \boldsymbol{\Lambda}_r^{-1/2} \mathbf{C} \in \mathbb{R}^{r\times 375}$$

- Singular values guide the choice of r (low-data regime limits size of non-intrusive ROM)
- POD basis only computed if needed for reconstruction

3. Infer reduced operators

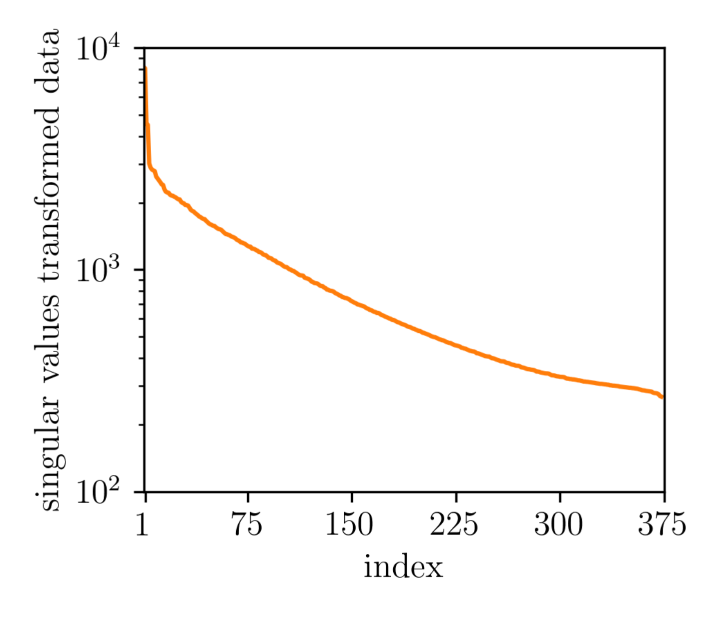

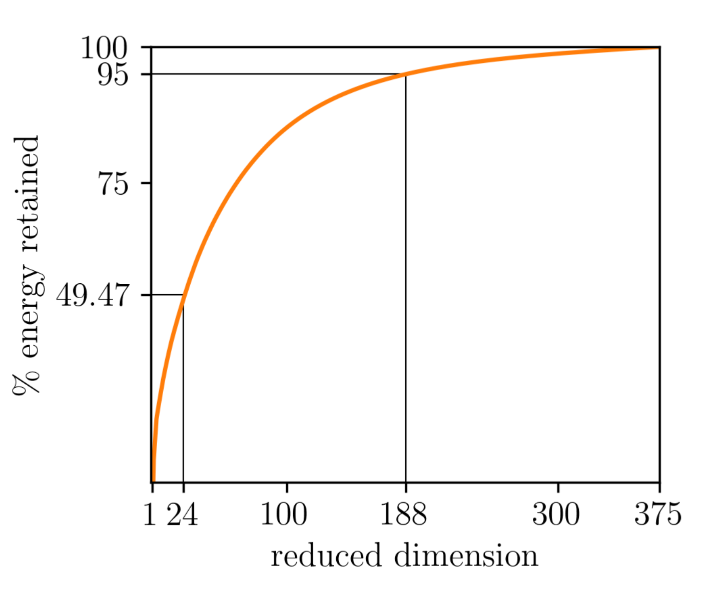

Parallel computations on Frontera (TACC) show excellent scalability

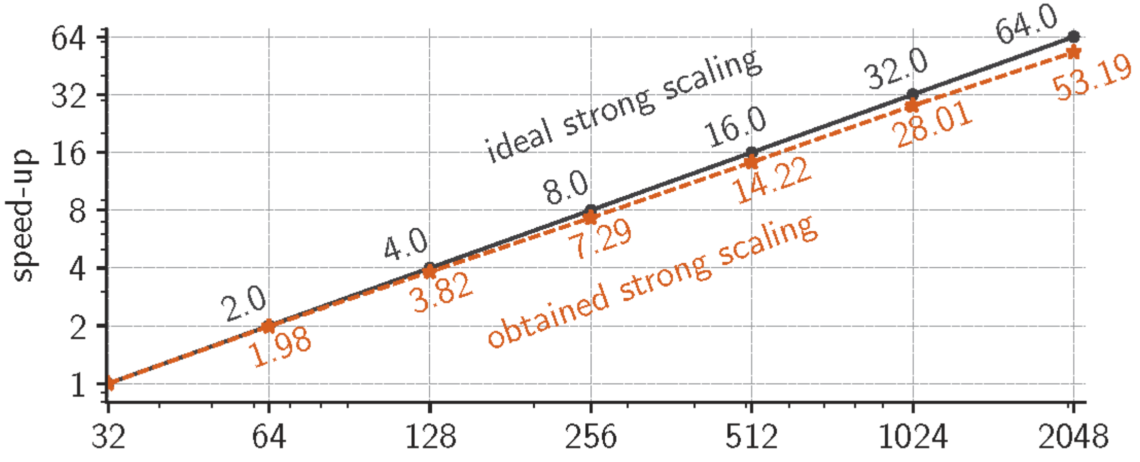

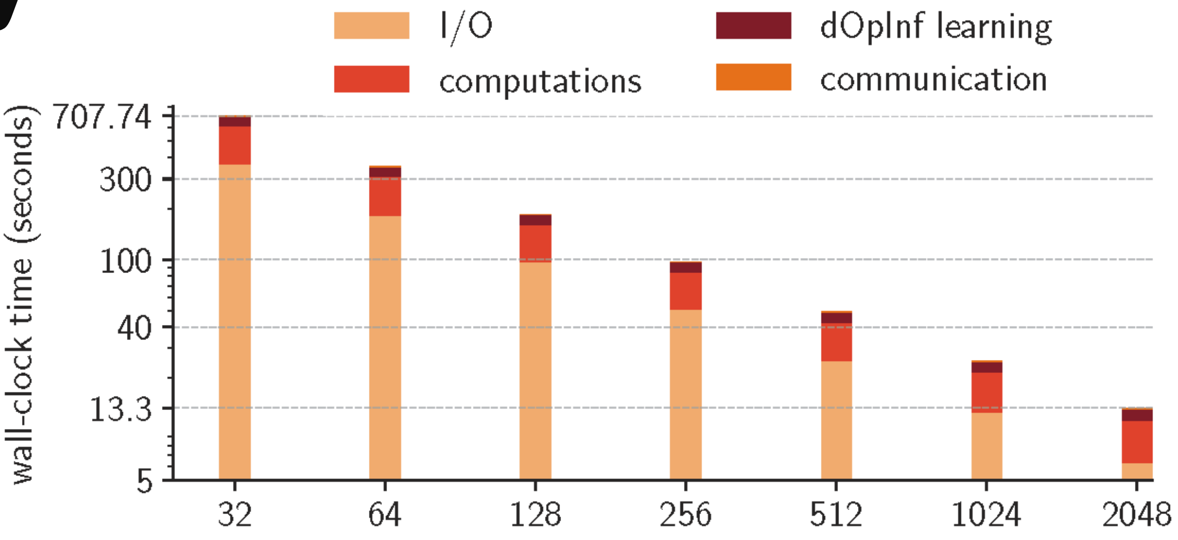

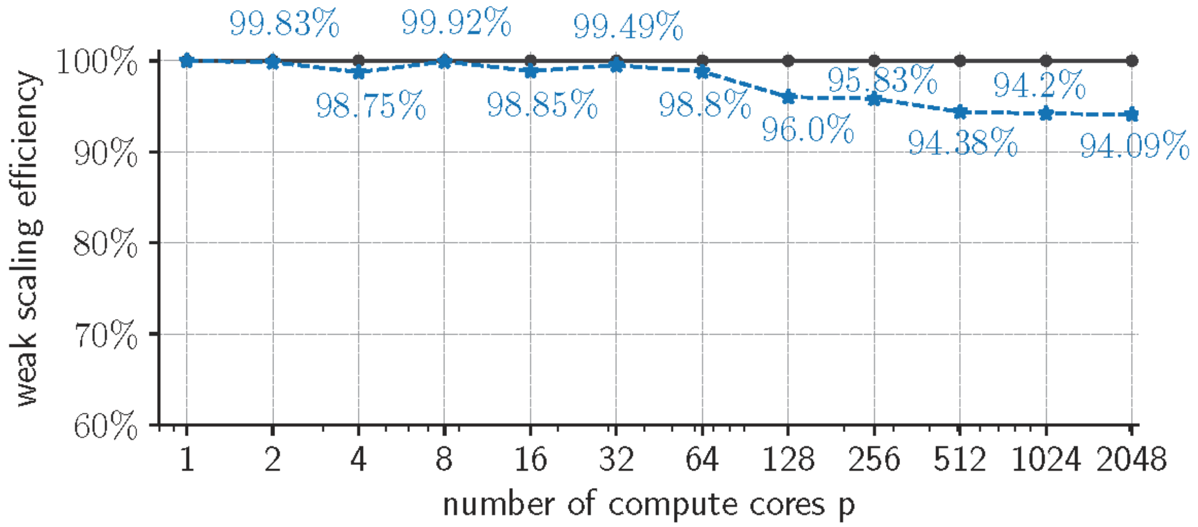

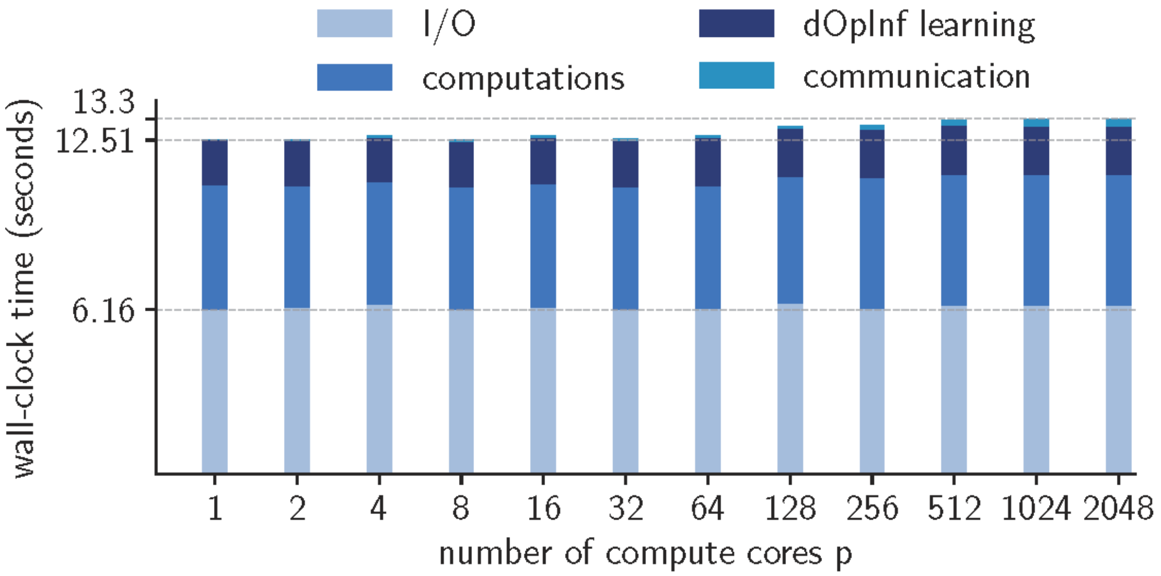

Farcas et al. *Computer Physics Communications*, 2025

Rotating detonation rocket engine simulation: weeks → milliseconds

76M → 24 dof

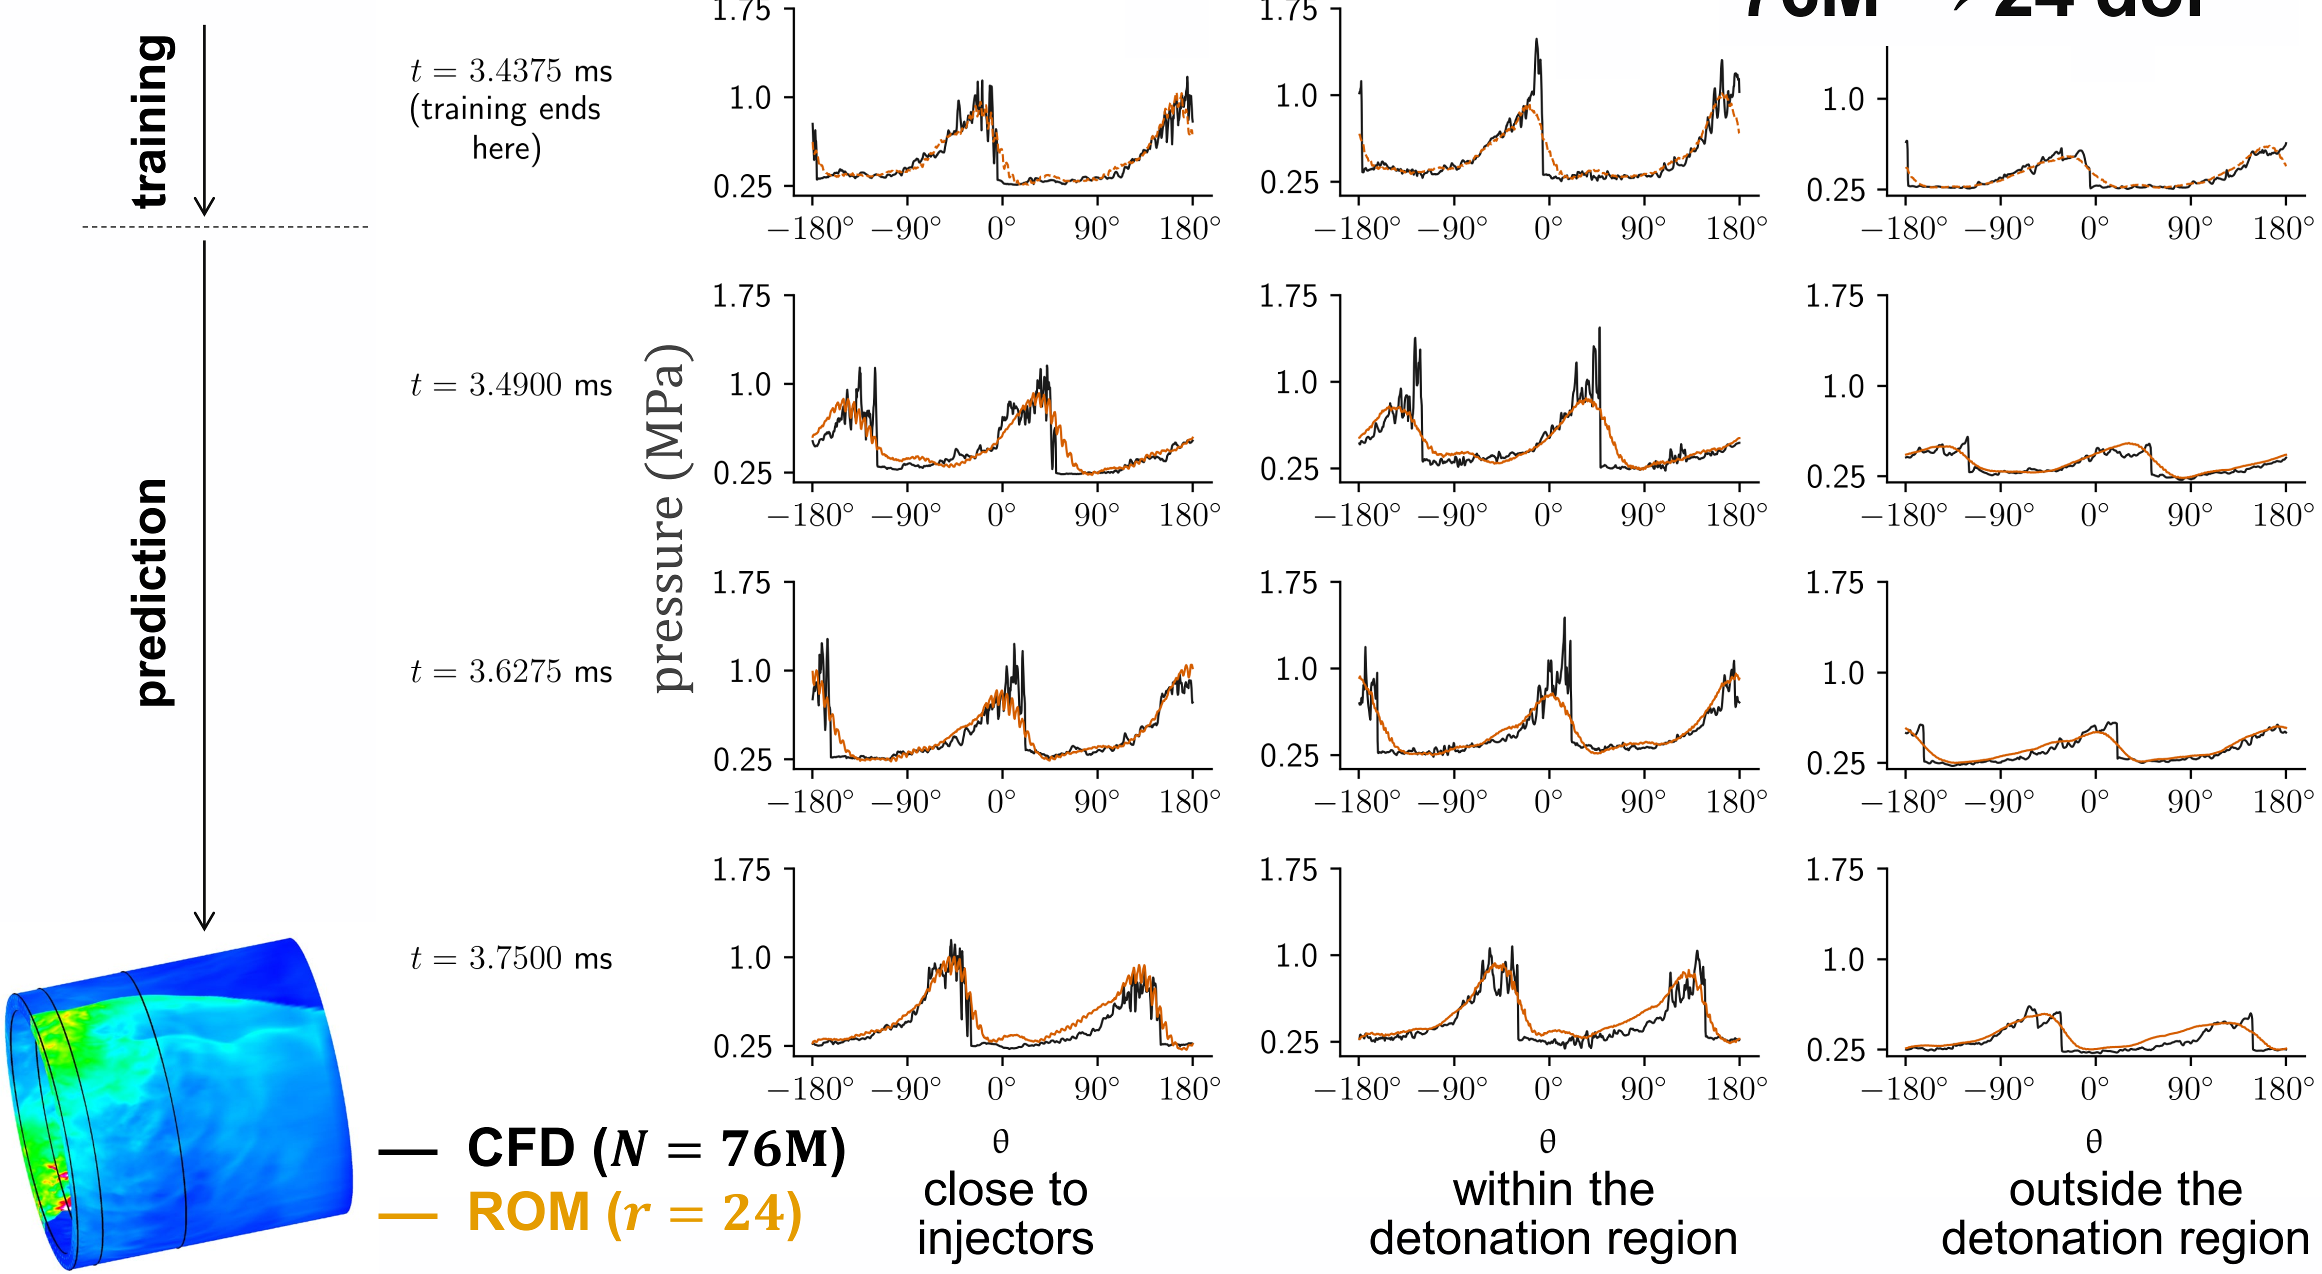

RDRE pressure contours: Reduced model captures coarse behavior but does not resolve all fine-scale dynamics

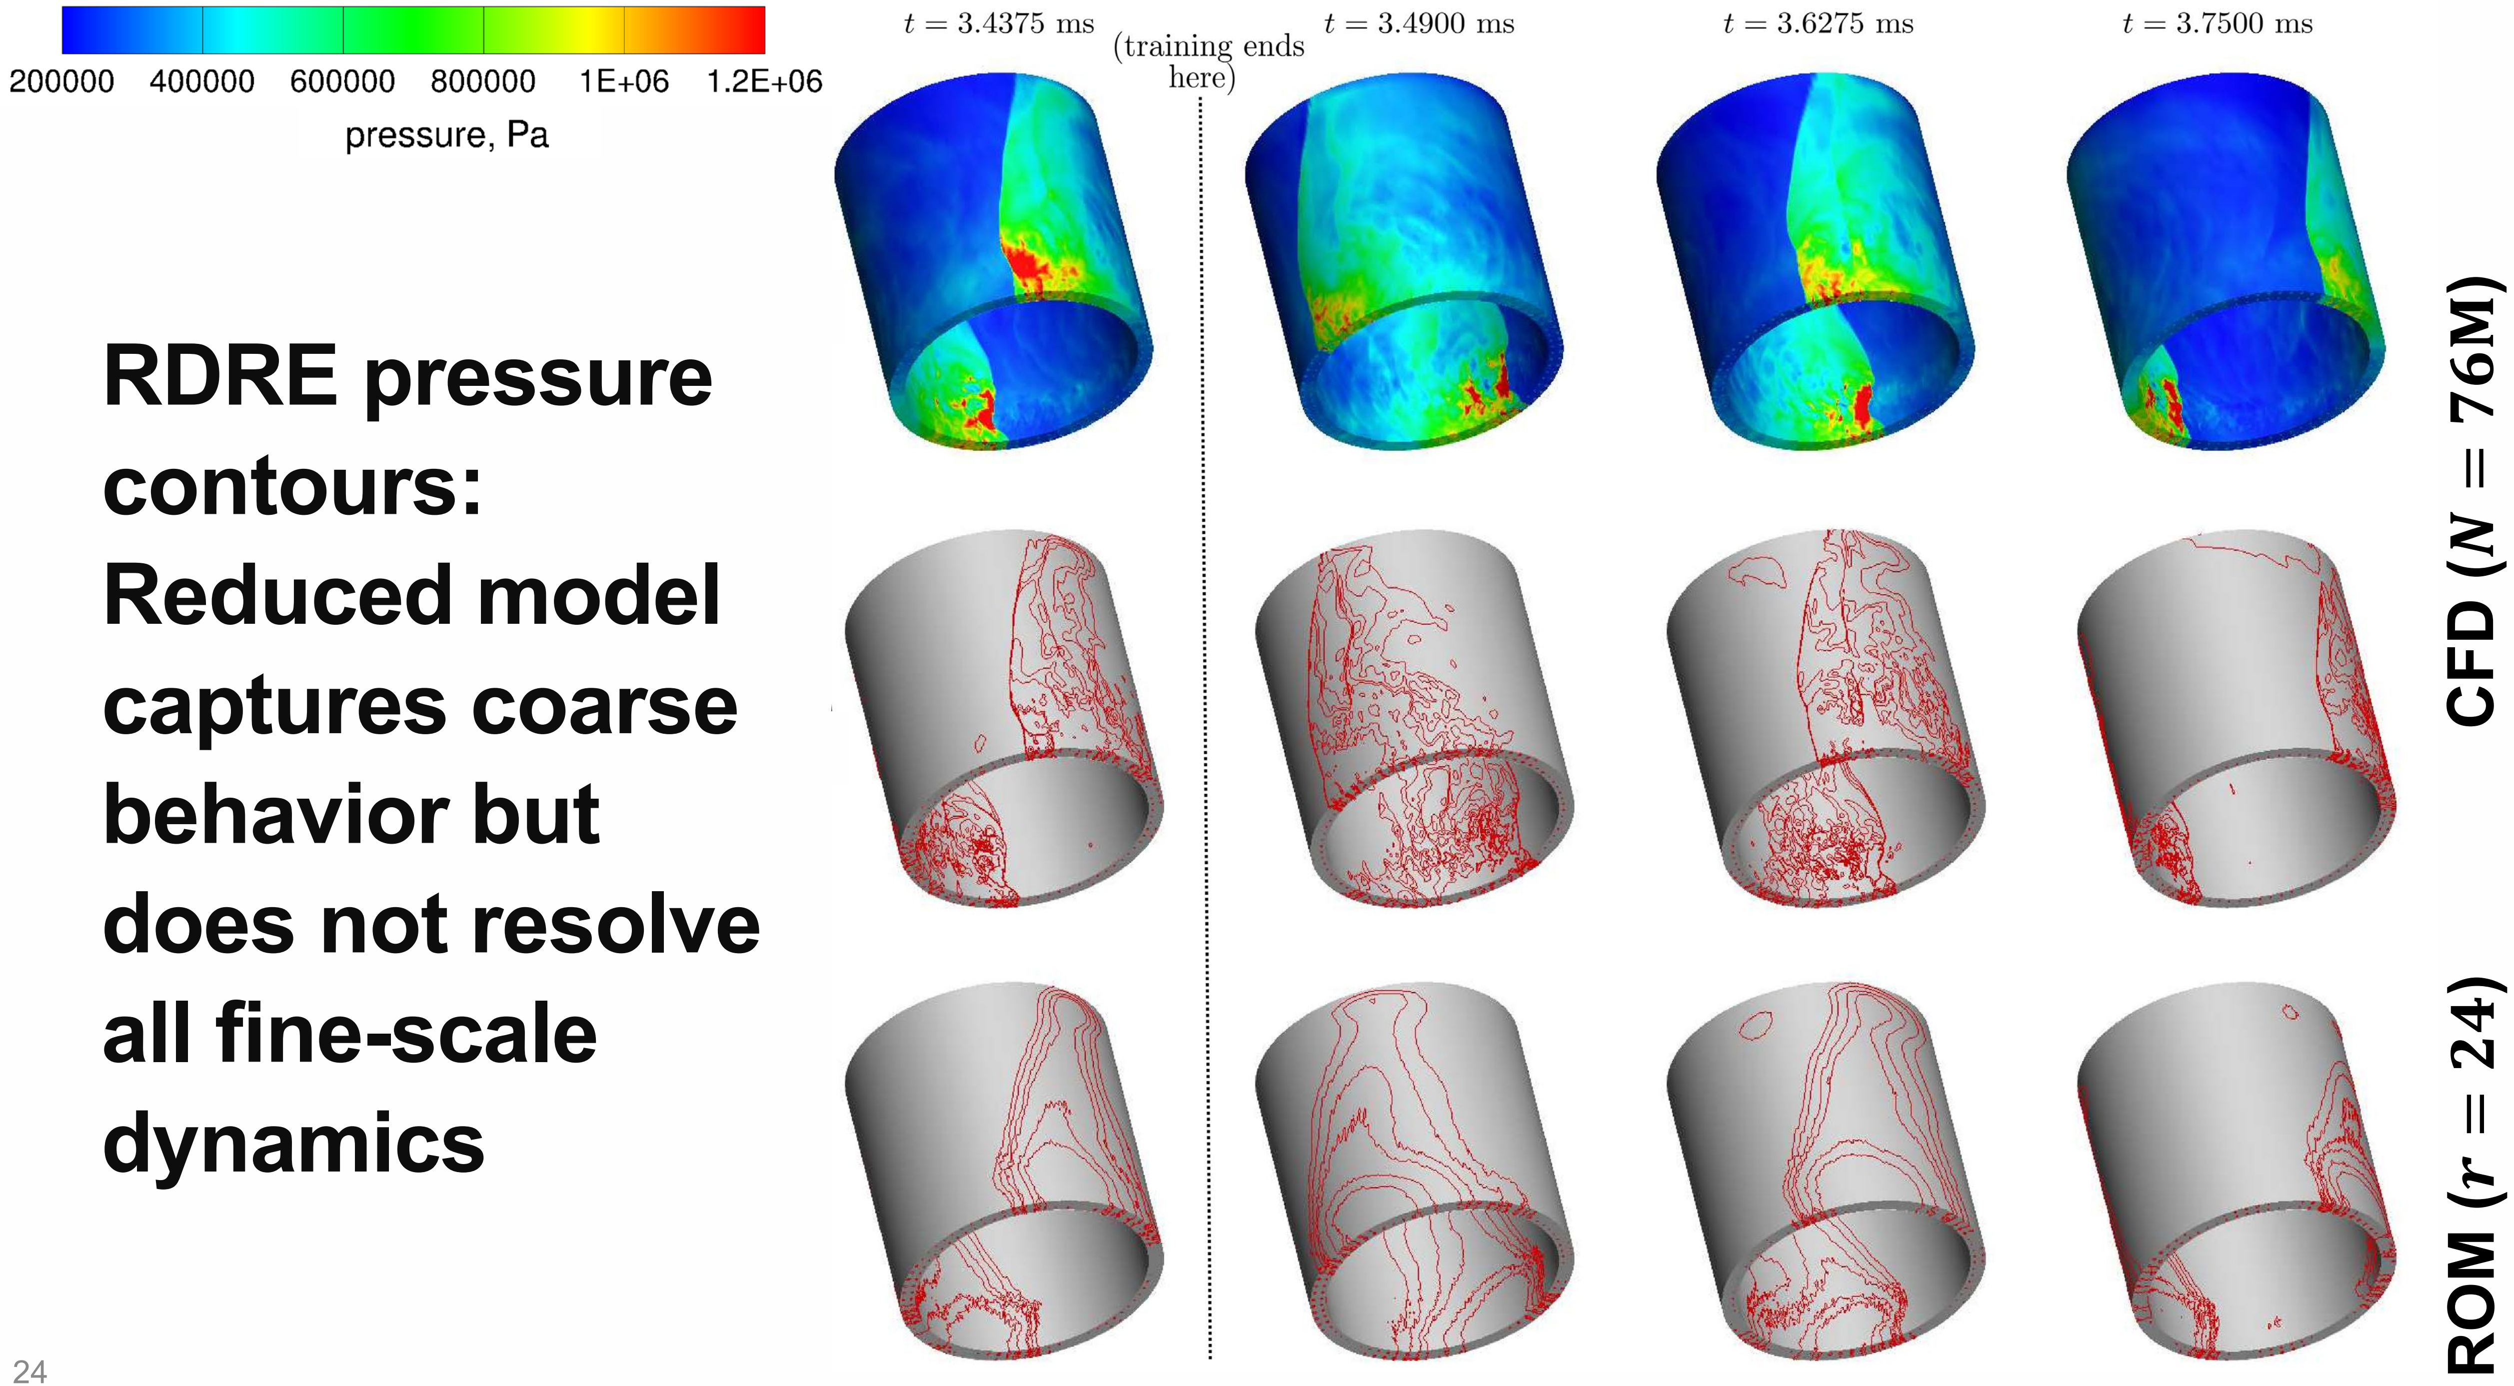

Looking to the Future

Decades of high-impact advances based on modeling & simulation for complex systems

Forward Simulations
Advancing scientific discovery & engineering innovation

Optimization & Inverse Problems
Advancing estimation, design & control

Uncertainty Quantification
Sensitivity analysis, design under uncertainty, risk assessment

Machine Learning & AI
Learning from data, accelerating workflows, …

Use of computing for complex systems spans well beyond forward simulation.

**But: (1) UQ, ML & AI still in their infancy;
(2) not yet seen full impact across the entire lifecycle.**

"A **digital twin** is a set of virtual information constructs that mimics the structure, context, and behavior of a natural, engineered, or social system (or system-of-systems), is **dynamically updated** with data from its physical twin, has a **predictive capability**, and informs **decisions** that realize value. **The bidirectional interaction between the virtual and the physical is central to the digital twin.**"

- National Academies Study on Foundational Research Gaps and Future Opportunities for Digital Twins, 2024

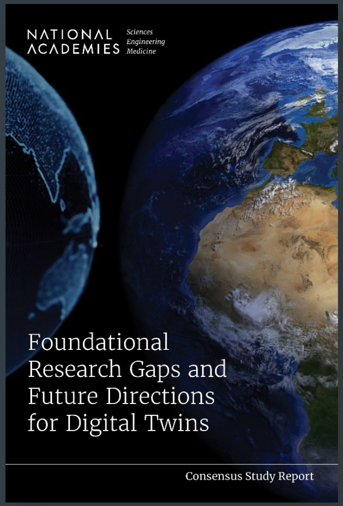

Reduced-order models are critical enablers for predictive digital twins

- There have been many advances in methods for parametric ROMs and nonlinear ROMs – these play an important role
- But ROM ability to handle complexity and wide range of operating conditions falls short for many digital twin applications
- Huge interest in medical digital twins, but relatively little ROM work for medical applications
- Cost of generating training data remains a barrier for many digital twin applications → this is a significant research need

Data-driven decisions

building the mathematical foundations and computational methods to enable design of the next generation of engineered systems

KIWI.ODEN.UTEXAS.EDU

ODEN INSTITUTE
FOR COMPUTATIONAL ENGINEERING & SCIENCES