

ASME TSD Webinar

Oct 14th, 2026



SciML and Frontier AI for Advancing Computational Modeling and Virtual Design of Combustion Engines




PINAKI PAL


Principal Research Scientist

Advanced Propulsion and Power Department


SIMULATION-DRIVEN ENGINE DESIGN OPTIMIZATION



CHALLENGES OF SIMULATION-DRIVEN ENGINE DESIGN



Competing objectives (efficiency, soot, NOx, noise, etc.)



*High-dimensional design spaces (10-20+ control variables)
Highly multi-modal & non-convex response surfaces*




CFD simulations are time-consuming & computationally expensive

ChemNODE & Phy-ChemNODE: Neural
Ordinary Differential Equation Frameworks
for Accelerating Stiff Chemical Kinetic
Computations in Combustion Simulations




DATA-DRIVEN MODELING OF STIFF CHEMICAL KINETICS


- Solving for detailed & stiff chemistry presents a major bottleneck in the application of combustion CFD for comprehensive parametric analysis and results in prolonged design cycles for combustion systems



IC engine




Gas Turbine Combustor



Rotating Detonation Combustor

A chemically reacting system (with no diffusion or convective transport) is given by:

$$\frac{d\Psi}{dt} = \dot{\omega}_{\Psi}(\Psi), \quad \Psi = [T, H_2, O_2, \dots]^T$$



We can replace the computation of $\dot{\omega}_{\Psi}$ using a neural network, $\mathcal{N}(\Psi; \theta)$, which learns to predict the non-linear source terms as functions of the thermochemical state of the system → **Amenable to variable time-stepping**

Conventional Data-driven “a-priori” Learning Approach

Train a neural network to minimize the difference between the predicted and actual source terms:

$$\mathcal{L} = \|\dot{\omega}_{\Psi} - \mathcal{N}(\Psi; \theta)\|_2^2 \rightarrow \text{Prone to unstable solution during deployment}$$

ChemNODE: NEURAL ODEs FOR CHEMICAL KINETICS

Original system

$$\Psi(t_f) = \Psi(t_0) + \int_{t_0}^{t_f} \mathbf{F}(\Psi) \Delta t_{\text{chem}}$$

Chemical
source terms

OH + O = H + O₂
OH + H = H₂ + O
OH + OH = H₂O + O
OH + H₂ = H₂O + H
H + H + M = H₂ + M
H + O + M = OH + M
O + O + M = O₂ + M
H + OH + M = H₂O + M
OH + CO = CO₂ + H
CO + O + M = CO₂ + M
Cl + H₂ = HCl + H
Cl + H₂O = HCl + OH
HCl + O = OH + Cl
H + Cl₂ = HCl + Cl

$$\frac{d\Psi}{dt} = \mathbf{F}(\Psi), \quad \Psi = [T, H_2, OH, \dots]^T$$

Advantages:



- ✓ Robust to autoregressive error accumulation over long-time horizons
- ✓ Can provide stiffness reduction
- ✓ Compatible with variable time-stepping

Training procedure:

Learn source term that minimizes trajectory error (*a-posteriori* learning)

Neural ODE

$$\Psi(t_f) = \Psi(t_0) + \int_{t_0}^{t_f} \mathcal{A}(\Psi) \Delta t_{\text{nn}}$$



PHYSICS-CONSTRAINED ChemNODE (Phy-ChemNODE)


Learning Dynamics in Latent Space using Autoencoder-NeuralODE

- Combining dimensionality reduction with latent space dynamics learning
 - Autoencoder** for mapping to and from low-dim latent space
 - NeuralODE** to capture dynamics in the reduced latent space


- Trained by combining the prediction loss, mapping loss, latent loss, & element conservation loss

$$L = ||\hat{y} - y||_1 + ||\psi(\phi(y)) - y||_1 + ||\phi(y) - z||_1$$

$$L_{total} = L + \lambda_1 L_{ele-1} + \lambda_2 L_{ele-2} \dots + \lambda_3 L_{ele-n}$$



Accurate, physically-consistent thermochemical state evolution with temporal stiffness reduction



~ 10-1000X speedup for CH₄/O₂ 32-species stiff chemistry

Kumar et al.

NeurIPS ML4PS 2023
CTM 2025

Kumar et al.


NeurIPS ML4PS 2024
Front. Therm. Eng. 2025




ActivO: Active Machine Learning-based Optimizer for Accelerating Simulation- driven Engine Design




EXPLORATION vs EXPLOITATION OF DESIGN SPACE




- Which value of x (design variables) should we run the next batch of simulations?
- Goal: Reach global optimum in as few calls to $f(x)$ as possible



Wasteful; would take longer than needed to reach optimum



The optimizer may converge prematurely




Balanced! Converges to global optimum quickly


ACTIVE ML ALGORITHM FOR RAPID DESIGN OPTIMIZATION



Activo



- “Dynamic on-the-fly” exploration and exploitation based on the change in the response surface predicted by the weak learner



CFD-ActivO DRIVEN IC ENGINE OPTIMIZATION

Notation	Input Parameter	min	max	units
nNoz	Number of Nozzle holes	8	10	-
TNA	Total Nozzle Area	1	1.3	-
Pinj	Injection Pressure	1400	1800	bar
SOI	Start of injection timing	-11	-7	dATDC
Nang	Nozzle Inclusion Angle	73	83	deg
EGR	EGR fraction	0.35	0.5	-
Tivc	IVC temperature	323	373	K
Pivc	IVC pressure	2.0	2.3	bar
SR	Swirl Ratio	-2.4	-1	-


$$Merit = 100 * \left\{ \frac{160}{ISFC} - f(PMAX) - f(MPRR) - f(SOOT) - f(NO_x) \right\}$$

$$f(PMAX) = 100 * \left\{ \frac{PMAX}{220} - 1, \text{if } PMAX > 220 \right. \\ \left. 0, \text{if } PMAX \leq 220 \right\}$$


$$f(MPRR) = 10 * \left\{ \frac{MPRR}{15} - 1, \text{if } MPRR > 15 \right. \\ \left. 0, \text{if } MPRR \leq 15 \right\}$$

$$f(SOOT) = \left\{ \frac{SOOT}{0.0268} - 1, \text{if } SOOT > 0.0268 \right. \\ \left. 0, \text{if } SOOT \leq 0.0268 \right\}$$

$$f(NO_x) = \left\{ \frac{NO_x}{1.34} - 1, \text{if } NO_x > 1.34 \right. \\ \left. 0, \text{if } NO_x \leq 1.34 \right\}$$



Number of CFD simulations to achieve convergence



- ActivO shows 7-10X speedup (from months to days) and computational savings relative to current industry-standard optimization techniques
- The automated optimizer can be coupled with any CFD solver for rapid simulation-driven **virtual prototyping** of energy technologies

CFD-ActivO OPTIMIZATION OF MIXING EQUIPMENT




Multi-parameter optimization of turbulent mixer geometry

HPC4 ENERGY INNOVATION

2


OpenFOAM®

- Automatic mesh generation and CFD case setup
- Multiphase L-E CFD simulations



1

Turbulent jet mixer geometry



4



CASES

Automatic geometry morphing for the new design parameter sets obtained from ActivO

10-parameter design space

$$MF = \frac{\text{entrained \%}_{ref}}{\text{entrained \%}} + \frac{\text{mixing time}}{\text{mixing time}_{ref}} + f(\tau_{bub})$$

- Evaluate the performance of mixer designs (multi-parameter)
- Generate new promising sets of design parameters




ActivO

Active ML optimization tool

3

GENERATIVE AI FRAMEWORKS FOR FUEL DESIGN


- An end-to-end Gen AI framework based on variational autoencoder (VAE) learns molecular representations in a probabilistic latent space
- Leveraged ALCF for efficient training the AI model on large structural databases
- The Gen AI model, coupled with data-driven quantitative structure-property relationship (QSPR) models, generates novel fuel/additive candidates with desired properties



<https://doi.org/10.48550/arXiv.2504.12075>




Traditional ML




- Individual siloed models
- Require task-specific training
- Lots of human supervised training




Foundation models



- Massive multi-modal model
- Adapted with minimal training
- Pre-trained unsupervised learning



Reasoning models



- Deliberative multi-step logic
- Self-consistency checks
- Slower but more accurate inference

COMB-FLOW-AI: Frontier AI Platform Transforming Combustion Science and Technologies for Energy and National Security

Team Lead/POC: Sibendu Som (ANL)

Team Co-Leads: Pinaki Pal (ANL), Jackie Chen (SNL), Shashank Yellapantula (NLR), Ramanan Sankaran (ORNL)

Co-Investigators:

ANL: Chandrachur Bhattacharya, Kiran Yalamanchi, Yixuan Sun, Sandeep Madireddy, Navin Mahto

SNL: Martin Rieth

NLR: Kevin Griffin, Bruce Perry

ORNL: Pei Zhang, Paul Laiu

LBNL: Jean Sexton, Ann Almgren



U.S. DEPARTMENT OF
ENERGY

Office of Science | Advanced Scientific Computing Research



Genesis Mission



U.S. DEPARTMENT of ENERGY

Combustion & Flows Model Team Vision & Scope

Overarching Vision: *Drive rapid research, design, development, and deployment of next-gen fuel combustion-powered energy and propulsion technologies leveraging Frontier AI and DOE-NL expertise (computational, experimental, AI, HPC)*

Scientific Problems Addressed:

- Near-real-time forecasting and control of combustion device operation
 - physics-based simulations and experiments are time-consuming and costly
- Prediction and causal understanding of safety-critical extreme/rare events [*combustion instabilities*]
- Rapid co-design of fuels and combustion systems [*efficiency, emissions, fuel flexibility, operability, durability*]
 - current industry design cycles (clean-slate concept to full production) remain extremely long, e.g., 5-10 years for piston engines and 10-15 years for aviation gas turbines
 - true complexity of high-dimensional design spaces is not navigated; design improvements are incremental


Capability Leap with AI — What Becomes Possible:

- Truly unprecedented full co-optimization of fuel-engine systems for clean-slate design at scale
 - slash industry design cycles by ~10-100X
- Robust control strategies for reliable and safe operation of fuel-flexible engines
 - cut maintenance, repair, and overhaul (MRO) costs by ~5X

Focus Area 1 – Multimodal AI Foundation Modeling

Develop scientific AI foundation modeling frameworks to seamlessly integrate complex multimodal data from multi-fidelity simulations & experiments for multiscale-multiphysics modeling of combustion systems


- Benefits: Scientific Foundation Modeling (SciFM) offers significant advantages over traditional AI/ML approaches:
 - multimodal architecture and cross-modal understanding
 - generalizability
 - multi-task adaptability
- Capability target: SciFM architectures for Combustion and Flows leveraging large-scale heterogeneous (0D-4D) scientific datasets (sim/exp); demonstrate for downstream tasks



Focus Area 2 – Multi-Agent AI Frameworks for Workflow Automation

Develop scalable & end-to-end frontier AI-driven multi-agent workflows for combustion science and engineering that integrate hypothesis generation, multiscale modeling, experimental design, and data analysis



- Benefits: Agentic AI frameworks can massively accelerate novel hypothesis generation and evaluation process for design conceptualization, exploration, and optimization of fuel-engine combustion recipes
- Capability target: Human-in-the-loop automated multi-agent AI pipelines for decision support:
 - drive physics-based simulation and SciFM runs
 - generate scientific insights from data
 - orchestrate active learning and design exploration/optimization loops, etc.



Scientific AI Foundation Modeling

Spatiotemporal Forecasting of Multiscale Turbulent Reacting Flows

- A pre-trained vision transformer (ViT)-based deterministic PDE foundation model (FM), **Walrus**, was advanced by coupling it with a lightweight latent flow matching head to enable **probabilistic forecasting with UQ (WalrusFlow)**
- Fine-tuning of WalrusFlow was demonstrated on combustion CFD datasets for **spatiotemporal forecasting of complex turbulent reacting flow dynamics**
- Ongoing work focuses on improving stability and accuracy of autoregressive rollouts over long-time horizons




Sun et al.
NeurIPS FMTS Workshop 2026

Multi-agent AI Frameworks for Workflow Automation

AI Scientific Assistant Core (AISAC)


The general-purpose **AI Scientific Assistant Core (AISAC)** developed at Argonne for orchestration of multi-agent frameworks and for deploying domain-adapted scientific co-scientists

- **Bounded multi-agent orchestration** — driver/helper role separation; recursion, step, and tool-round limits; per-session interrupt controls; context management
- **Scientific workflow runtime** — in-built multi-purpose tools, shared artifacts, stateful sessions, validated code execution
- **Layered persistent memory** — conversation history, semantic recall, agent notes, learned skills
- **Live observability and control** — real-time trace inspection, replay, Stop / Redirect / Clarify for Human-in-the-Loop
- **Institutional interoperability** — vetted endpoints, Globus auth, MCP partner agents, HPC deployment, offline mode
- **User-friendly by default** — native desktop UI, project wizard, multi-project manager; domain scientists can build and deploy agents without framework expertise




Multi-agent AI Frameworks for Workflow Automation

Cross-facility Multi-Agent workflow for end-to-end orchestration of combustion simulations



~ 20X
acceleration
over manual simulation
workflows



The agentic AI workflow was showcased for the use case of a Jet-in-Crossflow (JICF) system

Combustion AI Co-Scientist: Chandrachur Bhattacharya, Pinaki Pal, Sibendu Som (ANL)
ModSim AI Agent: Jean Sexton (LBNL); Bruce Perry, Kevin Griffin (NLR)
Post-processing AI Agent: Martin Rieth (SNL)

THANK YOU

pal@anl.gov