

Domain decomposition for data-driven reduced modeling of large-scale systems

Ionuț-Gabriel Farcaș *

The University of Texas at Austin, Austin, TX, 78712

Rayomand P. Gundeveia †

Jacobs Engineering Group, Inc., Edwards Air Force Base, CA 93524

Ramakanth Munipalli ‡

Air Force Research Laboratory, Edwards Air Force Base, CA 93524

Karen E. Willcox §

The University of Texas at Austin, Austin, TX, 78712

This paper focuses on the construction of accurate and predictive data-driven reduced models of large-scale numerical simulations with complex dynamics and sparse training datasets. In these settings, standard, single-domain approaches may be too inaccurate or may overfit and hence generalize poorly. Moreover, processing large-scale datasets typically requires significant memory and computing resources which can render single-domain approaches computationally prohibitive. To address these challenges, we introduce a domain decomposition formulation into the construction of a data-driven reduced model. In doing so, the basis functions used in the reduced model approximation become localized in space, which can increase the accuracy of the domain-decomposed approximation of the complex dynamics. The decomposition furthermore reduces the memory and computing requirements to process the underlying large-scale training dataset. We demonstrate the effectiveness and scalability of our approach in a large-scale three-dimensional unsteady rotating detonation rocket engine simulation scenario with over 75 million degrees of freedom and a sparse training dataset. Our results show that compared to the single-domain approach, the domain-decomposed version reduces both the training and prediction errors for pressure by up to 13% and up to 5% for other key quantities, such as temperature, and fuel and oxidizer mass fractions. Lastly, our approach decreases the memory requirements for processing by almost a factor of four, which in turn reduces the computing requirements as well.

I. Introduction

One of the main goals of computational science and engineering is to enable predictive simulations that inform design and decisions in engineering, physics, chemistry or medicine. To this end, the recent advances in algorithms, computer architectures and high-performance computing enable detailed and accurate such simulations at scale. However, making quantitative predictions over long time intervals or performing many-query applications such as design optimization, uncertainty quantification, and control remains challenging even on large supercomputers.

In this paper, we focus on the construction of physics-based data-driven reduced-order models (ROMs) of large-scale numerical simulations characterized by complex dynamics and sparse training datasets, with the goal of replacing the computationally expensive model with a cheap but sufficiently accurate ROM. The complexity of the underlying dynamics and the sparsity of the training datasets render the construction of accurate and predictive data-driven ROMs challenging via standard, global approaches. An additional challenge is that processing large-scale training datasets

*Postdoctoral fellow, Oden Institute for Computational Engineering and Sciences, ionut.farcas@austin.utexas.edu, AIAA Member.

†Aerospace Research Engineer, rayomand.gundeveia.ctr@afrl.af.mil, AIAA Member.

‡Senior Aerospace Research Engineer, Combustion Devices, ramakanth.munipalli@us.af.mil, AIAA Senior Member.

§Director, Oden Institute for Computational Engineering and Sciences, kwillcox@oden.utexas.edu, AIAA Fellow.

usually requires significant memory and computing resources. In this work, we address these challenges by introducing a domain decomposition formulation into the construction of data-driven ROMs.

Domain decomposition methods [1, 2] are prevalent in scientific computing and especially in high-performance computing where they are used to solve large-scale numerical simulations by decomposing them into smaller pieces that fit, for example, into the memory of one node on a compute cluster or supercomputer. Domain decomposition has also been used in model reduction. A large number of approaches target so-called decomposable systems, that is, systems comprising a set of well-defined components. These approaches are typically based on offline-online decompositions. In the offline phase, individual reduced bases are constructed for each component. In the online phase, a ROM for the full system is constructed using domain-decomposition approaches that enforce solution compatibility across interfaces. The majority of such methods are intrusive and have been developed in the context of linear parametric problems. Prominent examples include the reduced-basis element (RBE) method proposed in Ref. [3] and the static condensation RBE (scRBE) method [4–7]. Reduced bases and domain decomposition were used in Ref. [8] in the context of networks and complex parametrized geometries. In recent years, component-based approaches were extended to nonlinear problems as well. For example, Ref. [9] proposed an algebraically non-overlapping decomposition based on the least-squares Petrov-Galerkin approach and Ref. [10] formulated a method for nonlinear elliptic partial differential equations (PDEs) based on the partition of unity and adaptive enrichment to update the local reduced spaces. Ref. [11] proposed an intrusive component-based approach for large-scale systems, where the components are either ROMs or reduced fidelity full-order models. Furthermore, the predictive capabilities of the components were enhanced via online basis adaptation [12, 13]. Ref. [14] developed a domain-decomposed ROM approach for coupled heterogeneous systems without assuming a-priori knowledge of the coupled model.

Domain decomposition can also be used to split the full domain into subdomains followed by constructing ROMs for each subdomain. In the context of computational fluid dynamics simulations, for example, Refs. [15, 16] formulated domain-decomposed non-intrusive ROM strategies by leveraging Gaussian process regression or radial basis function interpolation to construct a set of local approximation functions (hypersurfaces) for each subdomain. Furthermore, Ref. [17] used domain decomposition together with balanced truncation for a class of PDE-constrained optimization problems governed by linear time-dependent advection-diffusion equations. This work was extended in Ref. [18] to the optimization of linear Stokes problems. Ref. [19] used a sparse estimation procedure based on local spatial patches suitable for complex flows. Other approaches besides domain decomposition for reducing complex, nonlinear problems include reduction to spectral submanifolds [20], including a recently developed data-driven variant [21], and model reduction with online basis adaptation and sampling [13].

Despite these advances, constructing accurate, predictive, and scalable ROMs of large-scale systems with complex dynamics and sparse training datasets remains challenging. For instance, the underlying simulation codes are complex and often proprietary, which makes the implementation of intrusive approaches difficult. Moreover, scalability is a major challenge due to the size of the datasets that must be generated for training. Our goal in the present work is to show that integrating domain decomposition into physics-based data-driven reduced modeling allows a flexible and easy-to-implement construction of accurate and scalable ROMs of large-scale systems.

To construct physics-based data-driven ROMs, we use Operator Inference (OpInf) [22, 23], noting that our domain-decomposed approach is also applicable to other data-driven methods. As a representative example of a large-scale application that is challenging to reduce via standard single-domain approaches, we consider a three-dimensional unsteady rotating-detonation rocket engine (RDRE) combustion chamber with 75, 675, 600 degrees of freedom and a sparse training dataset. The corresponding high-fidelity simulation is computationally very expensive, requiring about 1 million CPU-hours on 16,060 cores on a supercomputer for 1 ms of simulated physical time. The large computational cost prevents the utilization of the high-fidelity model for state predictions over longer time intervals or relevant engineering many-query tasks such as design optimization. Moreover, the size of the resulting simulation datasets is usually large, typically on the order of a few hundred gigabytes of disk storage. Hence, constructing accurate and scalable ROMs for these complex applications is essential in practice. We note that due to the complexity of the underlying simulation codes, constructing intrusive ROMs for these problems remains challenging.

The remainder of this paper is organized as follows. Section II introduces the setup for large-scale numerical simulations and summarizes the OpInf approach for constructing physics-based data-driven ROMs. Section III presents the proposed approach for constructing domain-decomposed ROMs from data, which we instantiate for OpInf. Section IV demonstrates the accuracy gains and scalability of the formulated domain-decomposed approach in a large-scale unsteady 3D RDRE combustion chamber with a sparse training dataset. We draw conclusions in Section V.

II. Single-domain physics-based data-driven reduced modeling

Section II.A introduces the setup for large-scale high-fidelity simulations. Section II.B summarizes the OpInf approach to learning physics-based ROMs from data.

A. Setup for large-scale high-fidelity simulations

We consider the large-scale system of nonlinear equations defined on the time domain $[t_{\text{init}}, t_{\text{final}}]$, with t_{init} denoting the initial time and t_{final} the final time,

$$\dot{\mathbf{q}} = \mathbf{f}(t, \mathbf{q}), \quad \mathbf{q}(t_{\text{init}}) = \mathbf{q}_{\text{init}}, \quad (1)$$

where $\mathbf{q}(t) \in \mathbb{R}^n$ is the n -dimensional vector of state variables at time $t \in [t_{\text{init}}, t_{\text{final}}]$, \mathbf{q}_{init} is a specified initial condition, and $\mathbf{f} : [t_{\text{init}}, t_{\text{final}}] \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a nonlinear function that defines the time evolution of the system state. The governing equations (1) are written as an n -dimensional system of ordinary differential equations (ODEs) representing the large-scale system of discretized PDEs over the computational domain denoted by $\Omega \subset \mathbb{R}^d$, where typically $d = 2, 3$. The state dimension, $n \in \mathbb{N}$, scales with the (large) dimension of the spatial discretization, which we denote by $n_x \in \mathbb{N}$. Letting $n_s \geq 1$ denote the number of state variables, we have that $n = n_s \times n_x$.

We consider the setup where training data are available over the time interval $[t_{\text{init}}, t_{\text{train}}]$, where by $t_{\text{train}} < t_{\text{final}}$ we denote the end of the training time horizon. We collect n_t snapshots over the training horizon by solving the high-fidelity model (1) and recording the high-dimensional state solution at n_t time instants, where we typically have that $n \gg n_t$. The state solution at time t_k is referred to as the k th snapshot and is denoted \mathbf{q}_k . The n_t snapshots are collected into a matrix $\mathbf{Q} \in \mathbb{R}^{n \times n_t}$ with \mathbf{q}_k as its k th column:

$$\mathbf{Q} = \begin{bmatrix} | & | & & | \\ \mathbf{q}_1 & \mathbf{q}_2 & \dots & \mathbf{q}_{n_t} \\ | & | & & | \end{bmatrix}.$$

Given \mathbf{Q} , the goal is to construct accurate and predictive ROMs. In the following, we summarize physics-based data-driven reduced modeling using the OpInf approach.

B. Learning physics-based data-driven reduced models via Operator Inference

OpInf is a scientific machine learning approach for learning ROMs with polynomial structure from data, where the structure is specified by the governing equations [22, 23]. For more general types of nonlinearities, lifting transformations can be used to expose (sometimes approximate) polynomial structure in the lifted governing equations [24].

The first step in OpInf is to compute the (global) representation of the available snapshots in a low-dimensional subspace. This is achieved by first computing the thin singular value decomposition (SVD) of the snapshot matrix \mathbf{Q}^* ,

$$\mathbf{Q} = \mathbf{V}\mathbf{\Sigma}\mathbf{W}^\top,$$

where $\mathbf{V} \in \mathbb{R}^{n \times n_t}$ contains the left singular vectors, $\mathbf{\Sigma} \in \mathbb{R}^{n_t \times n_t}$ is a diagonal matrix containing the singular values of \mathbf{Q} in non-decreasing order $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{n_t}$, where σ_j denotes the j th singular value, and $\mathbf{W} \in \mathbb{R}^{n_t \times n_t}$ contains the right singular vectors. The first $r \ll n$ columns of \mathbf{V} , i.e., the left singular vectors corresponding to the r largest singular values form the POD basis $\mathbf{V}_r \in \mathbb{R}^{n \times r}$. The low-dimensional representation of the snapshots in the reduced-order linear subspace spanned by \mathbf{V}_r is then obtained by computing

$$\hat{\mathbf{Q}} = \mathbf{V}_r^\top \mathbf{Q} \in \mathbb{R}^{r \times n_t}. \quad (2)$$

The next step is to determine the reduced operators. For example, for a ROM with a quadratic state dependency

$$\dot{\hat{\mathbf{q}}} = \hat{\mathbf{A}}\hat{\mathbf{q}} + \hat{\mathbf{H}}(\hat{\mathbf{q}} \otimes \hat{\mathbf{q}}), \quad (3)$$

we must determine the linear and quadratic reduced operators $\hat{\mathbf{A}} \in \mathbb{R}^{r \times r}$ and $\hat{\mathbf{H}} \in \mathbb{R}^{r \times r^2}$. OpInf determines the reduced operators that best match the projected snapshot data in a minimum residual sense by solving the following linear least-squares minimization problem

$$\underset{\hat{\mathbf{A}}, \hat{\mathbf{H}}}{\operatorname{argmin}} \left\| \hat{\mathbf{Q}}^\top \hat{\mathbf{A}}^\top + (\hat{\mathbf{Q}} \otimes \hat{\mathbf{Q}})^\top \hat{\mathbf{H}}^\top - \dot{\hat{\mathbf{Q}}}^\top \right\|_F^2 + \lambda_\ell \|\hat{\mathbf{A}}\|_F^2 + \lambda_q \|\hat{\mathbf{H}}\|_F^2, \quad (4)$$

*Note that it is often needed to shift and scale \mathbf{Q} prior to computing the global representation of the available snapshots in a low-dimensional subspace, especially in problems with multiple state variables with different scales such as pressure and species mass fractions in reactive flows.

where F denotes the Frobenius norm and $\dot{\hat{\mathbf{Q}}} \in \mathbb{R}^{r \times n_t}$ denotes the time derivative of the reduced snapshot data. If the high-dimensional time derivatives $\dot{\mathbf{Q}} \in \mathbb{R}^{n \times n_t}$ are provided by the simulation code, $\dot{\hat{\mathbf{Q}}}$ is obtained analogously to (2) as $\dot{\hat{\mathbf{Q}}} = \mathbf{V}_r^T \dot{\mathbf{Q}}$. Otherwise, it must be estimated numerically. In the context of OpInf ROMs for single-element rocket combustors, the works in [25, 26] showed that using higher-order finite difference schemes yields accurate approximations when the training snapshots were not subjected to downsampling. When downsampling the training snapshots, finite difference approximations may become inaccurate and alternative methods such as time-spectral methods [27] can be utilized instead. However, in cases where the snapshots are significantly downsampled, as seen in the RDRE scenario considered in the present paper, accurately approximating $\dot{\hat{\mathbf{Q}}}$ becomes challenging. Given the sensitivity of least-squares minimizations to perturbations in the right-hand side [28], inaccuracies in approximating $\dot{\hat{\mathbf{Q}}}$ can lead to inaccurate ROM predictions. In such cases, we use the discrete-time version of (3) and a corresponding fully discrete version of (4) to derive the reduced operators of the discrete-time ROM. In this approach, snapshot time derivatives are no longer needed, and are replaced by time-shifted snapshot data, akin to the dynamic mode decomposition (DMD) technique for learning linear discrete-time systems [25, 29, 30]; see also [31, 32].

Following Ref. [33], Tikhonov regularization is added to (4) to reduce overfitting and to account for model misspecification or other sources of error, with scalar regularization hyperparameters $\lambda_\ell, \lambda_q \in \mathbb{R}$. The procedure proposed in [33] for finding the optimal values for λ_ℓ, λ_q involves solving (4) and ensuing the resulting ROM over a time horizon $[t_{\text{init}}, t_{\text{reg}}]$ with $t_{\text{reg}} \geq t_{\text{final}}$ over two nested loops: an outer-loop over candidate values for λ_ℓ and an inner-loop over candidate values for λ_q ; see also [26]. Nonetheless, this search is embarrassingly parallel and computationally cheap in general since (4) depends on the reduced dimension $r \ll n$ and can therefore be solved efficiently via standard least squares methods [28]. We note that especially in problems with complex dynamics such as the RDRE scenario considered in this work regularization is key to obtaining accurate OpInf ROMs.

We note a few important points about our OpInf formulation. Firstly, note that the quadratic reduced operator $\hat{\mathbf{H}} \in \mathbb{R}^{r \times r^2}$ in (3) is not uniquely defined as written due to the symmetry of quadratic products (i.e., $q_1 q_2 = q_2 q_1$), leading to redundant coefficients in the OpInf optimization problem statement (4). To ensure a unique solution, we eliminate the redundant degrees of freedom, which amounts to learning an operator of dimension $r \times r(r+1)/2$ as described in [23]. Secondly, we impose a restriction on the maximum reduced dimension of the OpInf ROM so as to ensure that the OpInf learning problem (4) is not under-determined. This maximum reduced dimension is set by ensuring that the number of operator coefficients to be learned does not exceed the number of training data points available. As discussed in Refs. [22] and [23], the least squares problem (4) decouples into r independent least squares problems; nonetheless, in scenarios with sparse training datasets, the sparsity of data can impose restrictions on the ROM dimension. Another approach might be to permit (4) to be under-determined and instead employ ℓ_1 -based sparsification instead of Tikhonov regularization, similarly to Ref. [34], for example. Lastly, we note that additional structure can be imposed on the OpInf ROM. For example, one might impose energy conservation by adding constraints to the entries of the reduced quadratic operator for a quadratic ROM [35], or impose Hamiltonian [36, 37] and Lagrangian [38] structure.

III. Domain-decomposed physics-based data-driven reduced modeling

This section presents our proposed approach, which integrates domain decomposition into physics-based data-driven reduced modeling of large-scale systems with sparse training datasets. Section III.A presents the domain-decomposed formulation of the large-scale governing equations. This formulation is extended to data-driven reduced modeling and instantiated for OpInf in Section III.B.

A. Domain-decomposed formulation of the governing equations

We begin by splitting the computational domain Ω into $k \in \mathbb{N}_{\geq 2}$ subdomains $\Omega_1, \Omega_2, \dots, \Omega_k \in \mathbb{R}^d$ such that $\Omega = \cup_{i=1}^k \Omega_i$ and $k \leq n_x$. In the limit when $k = n_x$, the computational domain is split into n_x subdomains with a single spatial degree of freedom per subdomain. We denote by $\mathbf{q}^{(i)}(t) \in \mathbb{R}^{n_i}$ the components of the full-domain vector of state variables $\mathbf{q}(t)$ at time t with size $\mathbb{N} \ni n_i < n$ corresponding to the i th subdomain.

For each subdomain Ω_i , let $I(i) = \{j : j \in \{1, 2, \dots, k\} \setminus \{i\}, \Omega_j \text{ is adjacent to } \Omega_i\}$ denote the set of indices of all

subdomains adjacent to it. We can then rewrite the governing equations (1) as a system of coupled ODEs:

$$\left\{ \begin{array}{l} \dot{\mathbf{q}}^{(1)} = \underbrace{\mathbf{f}_{1,1}(t, \mathbf{q}^{(1)})}_{\text{interior } \Omega_1} + \underbrace{\sum_{j \in I(1)} \mathbf{f}_{1,j}(t, \mathbf{q}^{(j)})}_{\text{coupling } \Omega_1 \rightarrow \Omega_j} \\ \dot{\mathbf{q}}^{(2)} = \underbrace{\mathbf{f}_{2,2}(t, \mathbf{q}^{(2)})}_{\text{interior } \Omega_2} + \underbrace{\sum_{j \in I(2)} \mathbf{f}_{2,j}(t, \mathbf{q}^{(j)})}_{\text{coupling } \Omega_2 \rightarrow \Omega_j} \\ \vdots \\ \dot{\mathbf{q}}^{(k)} = \underbrace{\mathbf{f}_{k,k}(t, \mathbf{q}^{(k)})}_{\text{interior } \Omega_k} + \underbrace{\sum_{j \in I(k)} \mathbf{f}_{k,j}(t, \mathbf{q}^{(j)})}_{\text{coupling } \Omega_k \rightarrow \Omega_j} \end{array} \right. \quad (5)$$

where $\mathbf{f}_{i,i} : [t_{\text{init}}, t_{\text{end}}] \times \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{n_i}$ models the time evolution of the state in the interior of the i th subdomain and $\mathbf{f}_{i,j} : [t_{\text{init}}, t_{\text{end}}] \times \mathbb{R}^{n_j} \rightarrow \mathbb{R}^{n_i}$ models the state evolution at the interface between two adjacent subdomains i and j . Note that we did not specify whether the decomposition is done with or without overlapping; we will do so in the next section.

B. Domain-decomposed Operator Inference

1. General considerations

Starting from the domain-decomposed formulation (5) of the governing equations, we seek the construction of accurate, predictive, and scalable physics-based domain-decomposed ROMs from data. The generic steps to construct such ROMs are to (i) split the training dataset into subsets corresponding to subdomains, (ii) use these subsets to construct reduced bases for each subdomain, and (iii) determine the corresponding interior and coupling reduced operators. Once constructed, the reduced system of coupled ODEs is integrated over the desired time horizon to obtain the reduced solutions for each subdomain, which are then used to determine the approximate global, full-domain solution. Constructing accurate data-driven domain-decomposed ROMs using a non-overlapping decomposition can be challenging in practice, especially in problems characterized by complex dynamics such as the RDRE scenario considered in Section IV. In such problems, it is likely that the corresponding reduced solutions will have sharp transitions or oscillations at the interfaces even when the reduced dimensions of the subdomain ROMs are large, which in turn will lead to inaccurate global solutions (see Refs. [11, 39], for example, for similar discussions). To mitigate this challenge, we formulate our domain-decomposed approach in terms of overlapping subdomains. That is, for two adjacent subdomains i and j , we have that $\Omega_i \cap \Omega_j = \Omega_{i \cap j}$ where $\Omega_{i \cap j}$ denotes the overlapping region in the physical space. Notice that overlapping implies that $\sum_{i=1}^k n_i > n$.

The decomposition can be influenced by several factors, including the domain topology, the anisotropy of the dynamics, and the size of the training dataset. For instance, if the computational domain consists of multiple similar components, it is natural to decompose the domain into subdomains corresponding to each component. An illustration of this approach can be found in Ref. [11], where a multi-injector rocket combustor was decomposed with each subdomain representing an injector element. Additionally, if the complex dynamics are concentrated in specific regions of the domain, it can be advantageous to incorporate this localization into the decomposition process. This approach was considered in Ref. [40], for example, where the computational domain was decomposed based on the shock location. Another crucial consideration is the size of the training dataset. When dealing with a large training dataset, it is essential to consider the available computational resources for managing the dataset effectively, which in turn may dictate a required level of partitioning.

In the following, we present our formulation for data-driven domain-decomposed ROMs with overlapping for OpInf, noting that it can be straightforwardly applied to other data-driven approaches such as DMD. We use the notation $DD-k$ to refer to the decomposition of Ω into k overlapping subdomains and $DD-k$ OpInf to refer to the corresponding domain-decomposed OpInf formulation.

2. Learning the domain-decomposed reduced operators via Operator Inference

Based on the overlapping decomposition of Ω into $\Omega_1, \Omega_2, \dots, \Omega_k$, we split the snapshot matrix \mathbf{Q} into k snapshot matrices $\mathbf{Q}^{(i)} \in \mathbb{R}^{n_i \times n_t}$, $i = 1, 2, \dots, k$. In the first step in DD- k OpInf, we compute the representation of these snapshots in low-dimensional subspaces for all k subdomains[†]. This is achieved by first computing the thin SVD of $\mathbf{Q}^{(i)}$,

$$\mathbf{Q}^{(i)} = \mathbf{V}^{(i)} \boldsymbol{\Sigma}^{(i)} \left(\mathbf{W}^{(i)} \right)^\top,$$

where $\mathbf{V}^{(i)} \in \mathbb{R}^{n_i \times n_t}$, $\boldsymbol{\Sigma}^{(i)} \in \mathbb{R}^{n_i \times n_t}$, and $\mathbf{W}^{(i)} \in \mathbb{R}^{n_t \times n_t}$. The left singular vectors corresponding to the $r_i \ll n_i$ largest singular values form the POD basis $\mathbf{V}_{r_i}^{(i)} \in \mathbb{R}^{n_i \times r_i}$ for the i th subdomain. The low-dimensional representation of the snapshots in the reduced-order subspace spanned by $\mathbf{V}_{r_i}^{(i)}$ is then obtained by computing

$$\hat{\mathbf{Q}}^{(i)} = \left(\mathbf{V}_{r_i}^{(i)} \right)^\top \mathbf{Q}^{(i)} \in \mathbb{R}^{r_i \times n_t}. \quad (6)$$

For the continuous-time OpInf formulation, we must also estimate $\dot{\hat{\mathbf{Q}}}^{(i)} \in \mathbb{R}^{r_i \times n_t}$.

Remark 1 *The size of the overlapping region has a direct effect on the corresponding POD bases in that it influences how well these bases capture the dynamics in the neighboring subdomains. For two adjacent subdomains with similar dynamics, a small overlapping region generally suffices. In contrast, a larger overlapping region is typically needed for neighboring subdomains with anisotropic dynamics to capture potentially sharp transitions, for example.*

We note that if the dynamics within the individual subdomains exhibit self-similarity, it may be feasible to compute POD bases for a single domain or a few subdomains and subsequently use these bases for the remaining subdomains. For example, Ref. [11] considered such an approach in the context of intrusive ROMs for rocket combustors with multiple injector elements. In this approach, one would need to accommodate potential dynamics variations from subdomain to subdomain. One strategy would be to parametrize the boundaries of the subdomain(s) for which the POD bases are computed; however, finding such a parametrization for problems with complex dynamics such as the RDRE scenario considered in our numerical experiments in Section IV is not straightforward. We therefore do not pursue this idea in this paper, but rather we compute a POD basis for each subdomain individually.

For a single-domain ROM with a quadratic state dependency (cf. (3)) and linear coupling terms, the corresponding DD- k OpInf ROM reads

$$\begin{cases} \dot{\hat{\mathbf{q}}}^{(1)} = \underbrace{\hat{\mathbf{A}}_{1,1} \hat{\mathbf{q}}^{(1)} + \hat{\mathbf{H}}_{1,1} \left(\hat{\mathbf{q}}^{(1)} \otimes \hat{\mathbf{q}}^{(1)} \right)}_{\text{interior}} + \underbrace{\sum_{j \in I(1)} \hat{\mathbf{A}}_{1,j} \hat{\mathbf{q}}^{(j)}}_{\text{coupling}} \\ \dot{\hat{\mathbf{q}}}^{(2)} = \underbrace{\hat{\mathbf{A}}_{2,2} \hat{\mathbf{q}}^{(2)} + \hat{\mathbf{H}}_{2,2} \left(\hat{\mathbf{q}}^{(2)} \otimes \hat{\mathbf{q}}^{(2)} \right)}_{\text{interior}} + \underbrace{\sum_{j \in I(2)} \hat{\mathbf{A}}_{2,j} \hat{\mathbf{q}}^{(j)}}_{\text{coupling}} \\ \vdots \\ \dot{\hat{\mathbf{q}}}^{(k)} = \underbrace{\hat{\mathbf{A}}_{k,k} \hat{\mathbf{q}}^{(k)} + \hat{\mathbf{H}}_{k,k} \left(\hat{\mathbf{q}}^{(k)} \otimes \hat{\mathbf{q}}^{(k)} \right)}_{\text{interior}} + \underbrace{\sum_{j \in I(k)} \hat{\mathbf{A}}_{k,j} \hat{\mathbf{q}}^{(j)}}_{\text{coupling}}. \end{cases} \quad (7)$$

The linear coupling terms in (7) act like forcing/boundary condition terms between neighboring subdomains. We note that the coupling between adjacent subdomains may also include additional components such as quadratic terms. Extending (7) to such settings is trivial. Figure 1 provides a visual illustration of two example domains. Both domains are decomposed into three overlapping subdomains. The left plot shows a rectangular domain in which, from left to right, the first subdomain is coupled with the second subdomain, and the second subdomain is coupled with the third. Therefore, $I(1) = \{2\}$, $I(2) = \{1, 3\}$, and $I(3) = \{2\}$. On the right, we plot an annular, periodic domain in which all three sectors overlap, which means that $I(1) = \{2, 3\}$, $I(2) = \{1, 3\}$, and $I(3) = \{1, 2\}$.

[†]As noted in Section II.B, especially in problems with multiple state variables with different scales, we need to first shift and scale the given snapshot data prior to performing the overlapping decomposition and computing the POD basis in each subdomain.



Fig. 1 Two example domains decomposed into three overlapping subdomains. The overlapping regions are shown in red.

In the next step, we learn the reduced operators for each subdomain in (7) via OpInf by solving k independent linear least-squares minimizations. That is, in the i th least-squares minimization problem with $i = 1, 2, \dots, k$, we learn the interior linear and quadratic reduced operators $\hat{\mathbf{A}}_{i,i} \in \mathbb{R}^{r_i \times r_i}$ and $\hat{\mathbf{H}}_{i,i} \in \mathbb{R}^{r_i \times r_i^2}$ as well as the coupling operators $\hat{\mathbf{A}}_{i,j} \in \mathbb{R}^{r_i \times r_j}$ for all $j \in I(i)$ by solving

$$\begin{aligned} \underset{\hat{\mathbf{A}}_{i,i}, \hat{\mathbf{H}}_{i,i}, \hat{\mathbf{A}}_{i,j}, \quad j \in I(i)}{\operatorname{argmin}} \quad & \left\| \left(\hat{\mathbf{Q}}^{(i)} \right)^\top \hat{\mathbf{A}}_{i,i}^\top + \left(\hat{\mathbf{Q}}^{(i)} \otimes \hat{\mathbf{Q}}^{(i)} \right)^\top \hat{\mathbf{H}}_{i,i}^\top + \sum_{j \in I(i)} \left(\hat{\mathbf{Q}}^{(j)} \right)^\top \hat{\mathbf{A}}_{i,j}^\top - \left(\dot{\mathbf{Q}}^{(i)} \right)^\top \right\|_F^2 + \\ & \lambda_\ell^{(i)} \left(\left\| \hat{\mathbf{A}}_{i,i} \right\|_F^2 + \sum_{j \in I(i)} \left\| \hat{\mathbf{A}}_{i,j} \right\|_F^2 \right) + \lambda_q^{(i)} \left\| \hat{\mathbf{H}}_{i,i} \right\|_F^2. \quad (8) \end{aligned}$$

Notice that we use different regularization hyperparameters $\lambda_\ell^{(i)}, \lambda_q^{(i)} \in \mathbb{R}$ for each subdomain.

To find the optimal values for $\{\lambda_\ell^{(i)}, \lambda_q^{(i)}\}_{i=1}^k$, we extend the strategy formulated in Ref. [33] (cf. Section II.B) to the domain-decomposed setting. Since the reduced system of ODEs (7) is coupled, this involves solving (7) over $2k$ nested loops, that is, two loops for each subdomain for a ROM with quadratic structure. Therefore, the effort for finding separate hyperparameters for each subdomain scales exponentially with $2k$ for k overlapping subdomains. Even though this search is embarrassingly parallel, the overall computational cost can become prohibitive for a large number of subdomains. An alternative approach in such cases is to instead find a single pair of hyperparameters for all subdomains, keeping in mind that this strategy may impact the prediction accuracy of the ensuing DD- k OpInf ROM, especially when the dynamics in the underlying subdomains are heterogeneous. Once the values of the regularization hyperparameters are determined, the least-square minimization problems (8) are solved independently to find the interior and coupling reduced operators for all k subdomains.

In summary, the main ingredients needed in our DD- k OpInf approach are (i) knowledge of the structure of the governing equations, (ii) simulation data from the high-fidelity simulation code, and (iii) the coordinates of the underlying computational domain to perform the overlapping decomposition. From an implementation perspective, the proposed approach is flexible and easy to implement. After the decomposition is performed, the major required operations are the thin SVD (or equivalent approaches) of the subdomain snapshot matrices to determine the POD bases and linear least-squares minimization to learn the reduced interior and coupling operators. These are standard numerical linear algebra operations with efficient implementations provided by most scientific computing libraries.

3. Obtaining the full-domain solution using the domain-decomposed reduced solutions

Once the interior and coupling reduced operators are learned, the next step is to integrate the reduced system of coupled ODEs (7) over the time horizon $[t_{\text{init}}, t_{\text{final}}]$ and collect the k reduced solutions for all subdomains. The reduced initial conditions $\hat{\mathbf{q}}_{\text{init}}^{(i)} \in \mathbb{R}^{r_i}$ for each subdomain are determined by projecting their corresponding full-domain initial

conditions $\mathbf{q}_{\text{init}}^{(i)} \in \mathbb{R}^{n_i}$ onto the linear subspace spanned by the column vectors of $\mathbf{V}_{r_i}^{(i)}$ for all $i = 1, 2, \dots, k$:

$$\hat{\mathbf{q}}_{\text{init}}^{(i)} = \left(\mathbf{V}_{r_i}^{(i)} \right)^\top \mathbf{q}_{\text{init}}^{(i)}.$$

We denote by $\tilde{\mathbf{q}}^{(i)}(t) \in \mathbb{R}^{r_i}$ the reduced solution at time t corresponding to the i th subdomain. We then map $\tilde{\mathbf{q}}^{(i)}(t)$ back to the original space to obtain the approximate DD- k OpInf solutions in the original coordinates $\mathbf{q}_{\text{DD-}k}^{(i)}(t) \in \mathbb{R}^{n_i}$:

$$\mathbf{q}_{\text{DD-}k}^{(i)}(t) = \mathbf{V}_{r_i}^{(i)} \tilde{\mathbf{q}}^{(i)}(t) \in \mathbb{R}^{n_i}, \quad i = 1, 2, \dots, k.$$

To determine the full-domain approximate solution $\mathbf{q}_{\text{DD-}k}(t) \in \mathbb{R}^n$ using $\mathbf{q}_{\text{DD-}k}^{(1)}(t), \mathbf{q}_{\text{DD-}k}^{(2)}(t), \dots, \mathbf{q}_{\text{DD-}k}^{(k)}(t)$, we proceed as follows. The interior degrees of freedom in each subdomain solution are uniquely mapped to their corresponding indices in the full-domain solution. The overlapping degrees of freedom, however, cannot be uniquely mapped since they are computed twice for two adjacent subdomains. To obtain the full-domain solution, we smoothly combine the overlapping degrees of freedom. We illustrate this process in the following example.

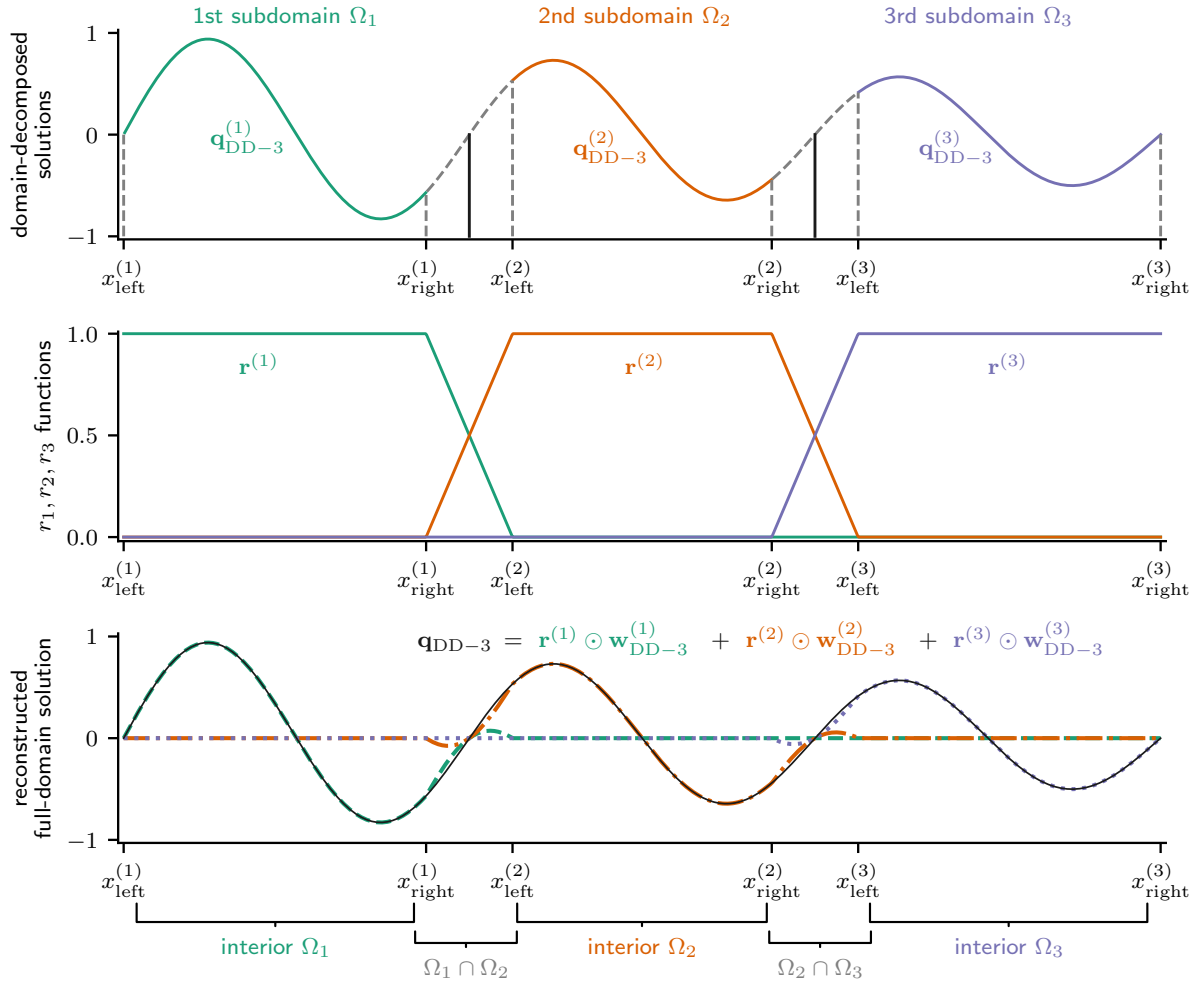


Fig. 2 Visual illustration of the process of determining the full-domain solution using domain-decomposed solutions using an example with three overlapping subdomains.

Consider a stationary damped sinusoidal wave function defined on a one-dimensional spatial domain $\Omega \subset \mathbb{R}$. We split Ω into three overlapping subdomains Ω_1, Ω_2 , and Ω_3 such that Ω_1 overlaps with Ω_2 and the Ω_2 overlaps with Ω_3 .

This is illustrated in the top part in Figure 2. The subdomains are separated by the black solid vertical lines. Moreover, let $\mathbf{q} \in \mathbb{R}^n$ denote the function evaluated at n spatial points and let $\mathbf{q}_{\text{DD-3}}^{(i)} \in \mathbb{R}^{n_i}$ denote the respective subdomain evaluations for $i = 1, 2, 3$. In addition, we denote by $\{x_{\text{left}}^{(i)}, x_{\text{right}}^{(i)}\}_{i=1}^3$ the spatial locations delimiting the interior of the three subdomains (shown on the x -axes).

To obtain the full-domain solution $\mathbf{q}_{\text{DD-3}} \in \mathbb{R}^n$ using $\mathbf{q}_{\text{DD-3}}^{(1)}$, $\mathbf{q}_{\text{DD-3}}^{(2)}$, and $\mathbf{q}_{\text{DD-3}}^{(3)}$, we first construct three functions $r^{(i)} : \Omega \rightarrow [0, 1]$ – one for each subdomain – that keep the respective interior degrees of freedom unchanged and at the same time permit us to smoothly combine the corresponding overlapping degrees of freedom:

$$r^{(i)}(x) = \begin{cases} 1, & x_{\text{left}}^{(i)} \leq x \leq x_{\text{right}}^{(i)} \\ \left(x - x_{\text{left}}^{(i+1)}\right) / \left(x_{\text{right}}^{(i)} - x_{\text{left}}^{(i+1)}\right), & x_{\text{right}}^{(i)} < x < x_{\text{left}}^{(i+1)}, \quad i = 1, 2 \\ \left(x - x_{\text{right}}^{(i-1)}\right) / \left(x_{\text{left}}^{(i)} - x_{\text{right}}^{(i-1)}\right), & x_{\text{left}}^{(i)} < x < x_{\text{right}}^{(i-1)}, \quad i = 2, 3 \\ 0, & \text{otherwise.} \end{cases} \quad (9)$$

We then evaluate these functions at all spatial points in the full-domain to respectively obtain $\mathbf{r}^{(1)}$, $\mathbf{r}^{(2)}$, and $\mathbf{r}^{(3)} \in \mathbb{R}^n$, which are depicted in the center plot in Figure 2. We extend $\mathbf{q}_{\text{DD-3}}^{(1)}$, $\mathbf{q}_{\text{DD-3}}^{(2)}$, and $\mathbf{q}_{\text{DD-3}}^{(3)}$ to the full-domain (by adding zeros to all additional entries) to obtain $\mathbf{w}_{\text{DD-3}}^{(1)}$, $\mathbf{w}_{\text{DD-3}}^{(2)}$, $\mathbf{w}_{\text{DD-3}}^{(3)} \in \mathbb{R}^n$. The full-domain solution $\mathbf{q}_{\text{DD-3}}$ is computed as:

$$\mathbf{q}_{\text{DD-3}} = \sum_{i=1}^3 \mathbf{r}^{(i)} \odot \mathbf{w}_{\text{DD-3}}^{(i)}, \quad (10)$$

where \odot denotes the element-wise (Hadamard) product. This is depicted in the bottom plot in Figure 2. Computing the full-domain solution requires knowledge of the solution in each subdomain, but it does not require those subdomain solutions to take any specific functional form. It is trivial to extend this procedure to an arbitrary number of subdomains as well as to other types of domains such as periodic domains by suitably adjusting the definition of the functions in (9).

IV. Numerical demonstration in a large-scale rotating detonation rocket engine simulation

We demonstrate the capabilities of the presented domain-decomposed ROM approach in a large-scale three-dimensional unsteady RDRE scenario with 75, 675, 600 degrees of freedom and a sparse training dataset. RDREs are based on fuel injection in an axially symmetric chamber such as an annulus wherein, once ignited at suitable conditions, a system of spinning detonation waves is produced (see Ref. [41] for a detailed description). This concept, in addition to its mechanical simplicity, offers several other propulsive advantages compared to more conventional rocket engine designs, making RDREs an active area of research in recent years. However, the large computational cost of the corresponding simulations prohibits using the high-fidelity model for state predictions over long time intervals or for practically relevant engineering applications such as the design optimization of these devices. This motivates the development of accurate and predictive data-driven ROMs to ultimately enable these tasks at scale.

A. Problem setup for high-fidelity large-eddy simulations

The physics of the problem are modeled using the three-dimensional, reactive, viscous Navier-Stokes equations coupled with a skeletal chemistry-mechanism (FFCMy-12) based on the FFCM model [42]. The full RDRE simulation was performed using implicit large-eddy simulations (LES) via the AHFM (ALREST High-Fidelity Modeling) simulation code from the Air Force Research Laboratory (AFRL); for more details, please refer to Refs. [43, 44]. The computational mesh size for the full RDRE comprises around 136 million spatial cells. The time step, computed adaptively, is about one nanosecond in the quasi-steady state regime. The corresponding high-fidelity LES is computationally expensive, requiring about 1 million CPU-hours on 16,060 cores on a supercomputer to obtain 1 ms of simulated physical time. Due to the size of the resulting simulation datasets, the number of time instants for which the high-fidelity LES solutions can be saved to disk is limited, usually amounting to only a few hundred.

The scenario considered here comprises 72 discrete injector-pairs. The values of the input parameters characterizing the flow conditions, the mass flow-rate \dot{m} [$\text{kg} \cdot \text{s}^{-1}$] and equivalence ratio Φ , are $\dot{m} = 0.267 \text{ kg} \cdot \text{s}^{-1}$ and $\Phi = 1.16$. These values lead to the formation of two dominant co-rotating waves and no secondary waves in the quasi-limit cycle regime. We note that this setup, utilizing a sparser dataset, was also considered in Ref. [31] within the framework of data-driven parametric OpInf ROMs for RDRE combustion chambers.

B. Setup for data-driven reduced modeling

We are interested in constructing ROMs for the combustion chamber only. The ROM domain spans from 0.05 to 76.15 mm in the x direction, spans from -37.50 to 37.50 mm in both y and z directions, and has a fixed channel height of 4.44 mm throughout. The original simulation data has been interpolated onto a grid domain comprising $n_x = 4,204,200$ spatial degrees of freedom, with clustering of grid points at mid channel and closer toward the injector plane. Figure 3 plots a typical pressure field example (top) as well as the x - y and y - z extents of the domain (bottom). For more details about this scenario, including the geometry, we refer the reader to Ref. [45].




Fig. 3 Pressure field example (top), and the x - y and y - z extents of the computational domain (bottom). For more details, we refer the reader to Ref. [45].

Following the work in Ref. [46], we build the ROMs for the RDRE combustion chamber using the specific volume flow variables (specific volume, pressure and the three velocity components), chemical species mass fractions, and temperature, which makes most terms in the governing equations linear or quadratic. The data snapshots are thus transformed to represent the following $n_s = 18$ state variables:

$$q = [1/\rho \ p \ v_x \ v_y \ v_z \ w_{\text{CH}_4} \ w_{\text{O}_2} \ w_{\text{CO}_2} \ w_{\text{H}_2\text{O}} \ w_{\text{CO}} \ w_{\text{H}_2} \ w_{\text{OH}} \ w_{\text{CH}_2\text{O}} \ w_{\text{CH}_3} \ w_{\text{HO}_2} \ w_{\text{H}} \ w_{\text{O}} \ T]^T, \quad (11)$$

where ρ [$\text{kg} \cdot \text{m}^{-3}$] is density (and $1/\rho$ [$\text{m}^3 \cdot \text{kg}^{-1}$] is specific volume), p [Pa] is pressure, v_x, v_y, v_z [$\text{m} \cdot \text{s}^{-1}$] are the velocity components, w are the mass fractions of all chemical species (12 in total here), and T [K] is temperature. The total number of degrees of freedom is therefore $n = n_s \times n_x = 18 \times 4,204,200 = 75,675,600$. Since the $n_s = 18$ state variables in (11) have significantly different scales (for example, from the order of 10^4 to 10^6 Pa for pressure to between 0 and 1 for the species mass fractions), we center and scale the snapshot data variable-by-variable prior to constructing

ROMs. Each state variable is first centered around the mean field (over the training data) in that variable, and then scaled by its maximum absolute value, thus ensuring that the scaled variables do not exceed $[-1, 1]$.

Remark 2 *Centering and scaling of snapshots are essential for obtaining accurate OpInf ROMs, in order to prevent undesirable bias towards variables with larger values in the corresponding POD bases. Scaling is particularly critical in cases where the state comprises multiple physical quantities with differing physical scales. In this work, we perform centering around the mean value of each state variable over the training horizon. We explored two scaling parameters: maximum absolute value and standard deviation of centered variables. Our findings show that, for this particular problem, the choice of scaling parameters from these two options does not have a significant impact on OpInf ROM prediction quality. The key consideration is to maintain consistent scales for non-dimensionalized variables.*

Remark 3 *In some applications the choice of scaling parameters can have a large effect, and for our application there may be a better scaling choice that would lead to improved ROM performance. It remains an open and important challenge to identify optimal scaling strategies. A particularly interesting avenue of work would be to explore strategies that non-dimensionalize the snapshots in a physically consistent way. Non-dimensionalization and variable transformations are well known to play a key role in establishing desirable numerical properties in high-fidelity computational fluid dynamics simulations; their role in reduced-order modeling is much less explored. As an alternative to an explicit variable scaling, previous work has explored physically-motivated choices for inner products in reduced-order modeling, such as an energy-based inner product for compressible flows in Ref. [47] and an entropy normalization scheme that is related to the Chu energy norm in Ref. [48]. For further discussions on data normalization beyond non-intrusive data-driven ROMs, we refer to Refs. [11, 49].*

We have 501 down-sampled snapshots in $\mathbb{R}^{75,675,600}$ from the high-fidelity LES over the time interval $[2.50, 3.75]$ ms, which corresponds to roughly four full-cycle periods of the two-wave system. The time step between the down-sampled snapshots is about $2.5 \mu\text{s}$, which means that every 2500th time instant or so from the high-fidelity simulation was saved to disk. In our numerical experiments, we use the first $n_t = 375$ snapshots to train the ROMs, which correspond to the about the first three full cycles. The remaining fourth full cycle, comprising 126 snapshots, is used to assess the ROMs accuracy for predictions beyond the training horizon.

C. Results using domain-decomposed Operator Inference

In the following, we use the DD- k OpInf approach presented in Section III.B to construct domain-decomposed ROMs and compare it in terms of scalability and accuracy with the standard OpInf method applied to a single domain as summarized in Section II.B. Here, we refer to single-domain OpInf as *SD OpInf*. We note that due to the significant downsampling of the available snapshots, it is difficult to accurately estimate the time derivatives of the projected snapshots in both SD and DD- k OpInf. We therefore instead use the fully discrete version of OpInf where we learn a fully discrete ROM, rather than a semi-discrete ROM; see also [31, 32]. All ROM calculations were performed in standard double precision arithmetic using a shared-memory machine with 256 AMD EPYC 7702 CPUs and 2 TB of RAM, using the `numpy` and `scipy` scientific computing libraries in `python`.

We decompose the 360° annular combustion chamber into $k = 4$ overlapping sectors. Due to the periodicity of the underlying annular domain, we overlap the first and fourth sectors as well. This is plotted in Figure 4. Our experiments using two angle extents for the overlapping region, $2\pi/36 = \pi/18$ and $2\pi/18 = \pi/9$, indicated that the larger overlapping region, $\pi/9$, corresponding to the angle extent of four injectors, leads to the most accurate results. Therefore, in the following, we consider an overlapping region spanning $\pi/9$ radians.

Remark 4 *We have also performed numerical experiments in which we split the domain into $k = 8, 12$, and 24 overlapping sectors, using several values for the angle extent of the overlapping region as well. However, we did not obtain any gains in accuracy compared to the case with $k = 4$ overlapping sectors. This is because for these splits, we used a single, global set of hyperparameters for all k least-squares minimization problems (8) since it was computationally impractical to find separate regularization hyperparameters (see the discussion in Section III.B.2). Since we observed that the DD- k OpInf ROM accuracy depends sensitively on the choice of regularization parameters, using the same regularization hyperparameters rather than individual ones had a detrimental effect on the ROM accuracy. We therefore consider only a decomposition with $k = 4$ overlapping sectors. We note, nevertheless, that for different configurations, flow conditions, or number of detonation waves, decomposing the annular combustion chamber into a larger number of sectors could increase the accuracy of the domain-decomposed ROM.*




Fig. 4 Combustion chamber decomposition into four overlapping sector with overlapping region between two sectors of $\pi/9$ radians.

In general, for an RDRE scenario with N dominant waves, it is natural to decompose the annular combustion chamber into $k = 2^\eta N$ overlapping sectors, where $\eta \in \mathbb{N}_0$. Moreover, for a configuration with M injector pairs, the extent of the overlapping region should be at least $4\pi/M$ to ensure that it covers at least one injector pair in each sector. Choosing η and the size of the overlapping region depends on the specific configuration, flow conditions, and number of detonation waves. Based on our experiments in RDRE scenarios with a fixed number of co-rotating waves, an overlapping region covering two or four injectors is sufficient to ensure accurate approximations at interfaces. However, in scenarios with more complex dynamics, such as those involving subdominant or counter-propagating waves, a larger overlapping region may be necessary. To determine the value of η for a given extent of the overlapping region, one approach is to analyze the decay of the POD singular values as η increases and assess whether this leads to potential accuracy gains, keeping in mind that the computational cost of determining the regularization hyperparameters in the DD- k OpInf learning problems grows exponentially with the number of sectors.

We first highlight the benefits of domain decomposition in terms of cost reduction and improved scalability. Even though the training dataset comprises only 375 snapshots, loading these snapshots into memory alone requires about 211 GB of RAM. Processing them and computing the POD basis via the thin SVD of the snapshot matrix requires even more RAM (we needed about 840 GB in total). This means that access to either a large shared-memory machine (for a serial or multi-threaded implementation) or a distributed-memory computer such as a cluster or a supercomputer (for a distributed implementation) is necessary to construct SD OpInf ROMs. In contrast, by splitting the combustion chamber into four overlapping sectors, the full dataset can be processed in smaller and independent components in parallel. The sector with the largest number of degrees of freedom requires roughly 64 GB of RAM for loading the data into memory, which is a factor of $3.27\times$ reduction compared to the single-domain approach. This reduction decreases the overall memory and computing requirements and hence makes the construction of the domain-decomposed ROM more manageable. We note that domain decomposition leads to memory and compute savings even if other methods such as randomized SVD [50], streaming SVD [51], or the method of snapshots [52] are used to compute the POD basis.

We next present the gains in terms of accuracy due to using domain decomposition. The maximum reduced dimension for a quadratic SD OpInf ROM is $r = 24$. This value also represents the maximum reduced dimension for each sector in DD-4 OpInf with linear coupling terms (7). These reduced dimensions respectively set the maximum number of operator coefficients that can be inferred via the OpInf regression problems (4) and (8). We therefore use $r \leq 24$ in SD OpInf and $r_1 = r_2 = r_3 = r_4 \leq 24$ in DD-4 OpInf in all our experiments.

Remark 5 As noted above, the coupling terms in our DD-4 OpInf experiments are linear; cf. Eq. (7). Considering quadratic coupling terms as well would double the size of the DD-4 OpInf ROM which in turn would further lower the maximum reduced dimensions to $r_1 = r_2 = r_3 = r_4 = 16$ due to the sparsity of the training dataset. Nevertheless, since the decomposition is done using overlapping, the respective POD bases contain information about the coupling between neighboring subdomains. Therefore, as we will show in our results next, considering only linear coupling terms does not impact the accuracy of our DD-4 OpInf ROM.

Figure 5 plots the POD singular values (top left) and corresponding retained energy (top right) for the single-domain approach. It is not surprising that the POD singular values decay slowly, given the complexity of the underlying scenario and the large downsampling factor of the snapshots. For example, a reduced dimension of $r = 188$ is necessary to ensure that 95% of the total energy is retained. In contrast, the maximum reduced dimension for SD OpInf ($r = 24$) retains only 49.47% of the total energy. The POD singular values corresponding to the four sectors are plotted alongside the singular values for the single-domain approach in the bottom left figure. The decomposition decreases their magnitude but does not lead to any significant decrease in the decay rate. This is due to the symmetry of the domain and the (quasi-)periodicity of the corresponding solution fields. Lastly, the bottom right figure plots the corresponding total energy, which is consistent with what was observed in the singular value decay plot. We note that even though the maximum reduced dimension retains only a small percentage of the total energy, the ROMs issue accurate predictions for quantities such as pressure and temperature, as we will show next. This is because the POD basis is global, which means that it is constructed to represent all $n_s = 18$ transformed state variables (11). Since the dynamics corresponding to variables such as pressure are smoother than the dynamics corresponding to species mass fractions, which are highly oscillatory, the global POD basis will represent some variables more accurately. An alternative approach is to construct separate bases for individual variables or groups of variables, in which case the ROM dimension is given by the cumulative size of the respective bases. This approach, however, is infeasible here due to the sparsity of the training set.




Fig. 5 POD singular values (top left) and retained energy (top right) for SD OpInf. POD singular values (bottom left) and total energy (bottom right) for both DD-4 and SD OpInf.

Figure 6 plots the first eight POD modes for the centered (with respect to the mean over the training horizon) and scaled (with respect to the maximum absolute value of the centered data) pressure field; we refer the reader to Ref. [53],

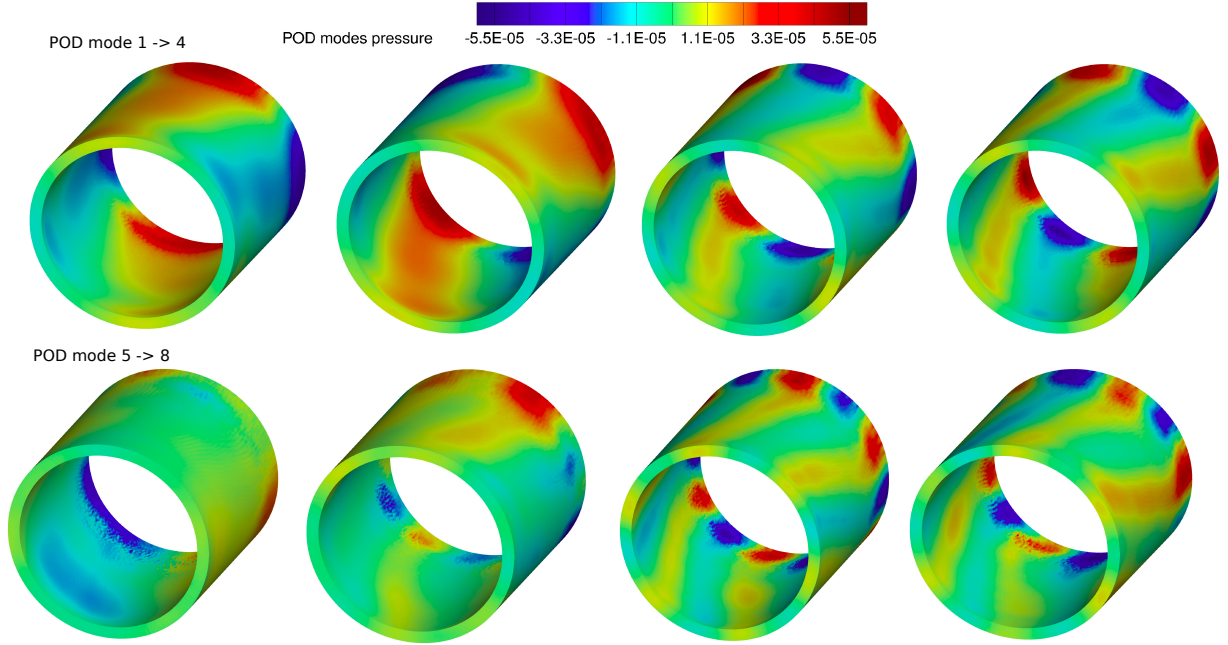


Fig. 6 First eight POD modes for the centered and scaled pressure field.




Fig. 7 First eight POD singular values for the centered and scaled pressure field.

for example, for a discussion about pressure POD modes in the context of reduced modeling of swirl injectors. We observe that modes 1-2, 3-4, and 7-8 come in pairs, which is explained by the (quasi)-periodicity of the underlying dynamics. In contrast, correlation between modes 5 and 6 seems poor, and this could be due to high amplitude fluctuations in coupled quantities such as entropy and enthalpy. In general, correlation in phase space indicates that dynamical information captured in the respective modes is adequate to describe the dynamics of larger scale features. Lack of correlation indicates coupling and other forms of nonlinearity. We furthermore plot, in Figure 7, the first eight POD singular values corresponding to the SD OpInf approach.

We also compare the amplitudes — over both training and prediction horizons — obtained by projecting the centered and scaled snapshots via (2) with those obtained via SD OpInf. Figure 8 plots two-dimensional phase portraits of the amplitudes corresponding to the five most energetic modes. The pairs involving the first four amplitudes show strong

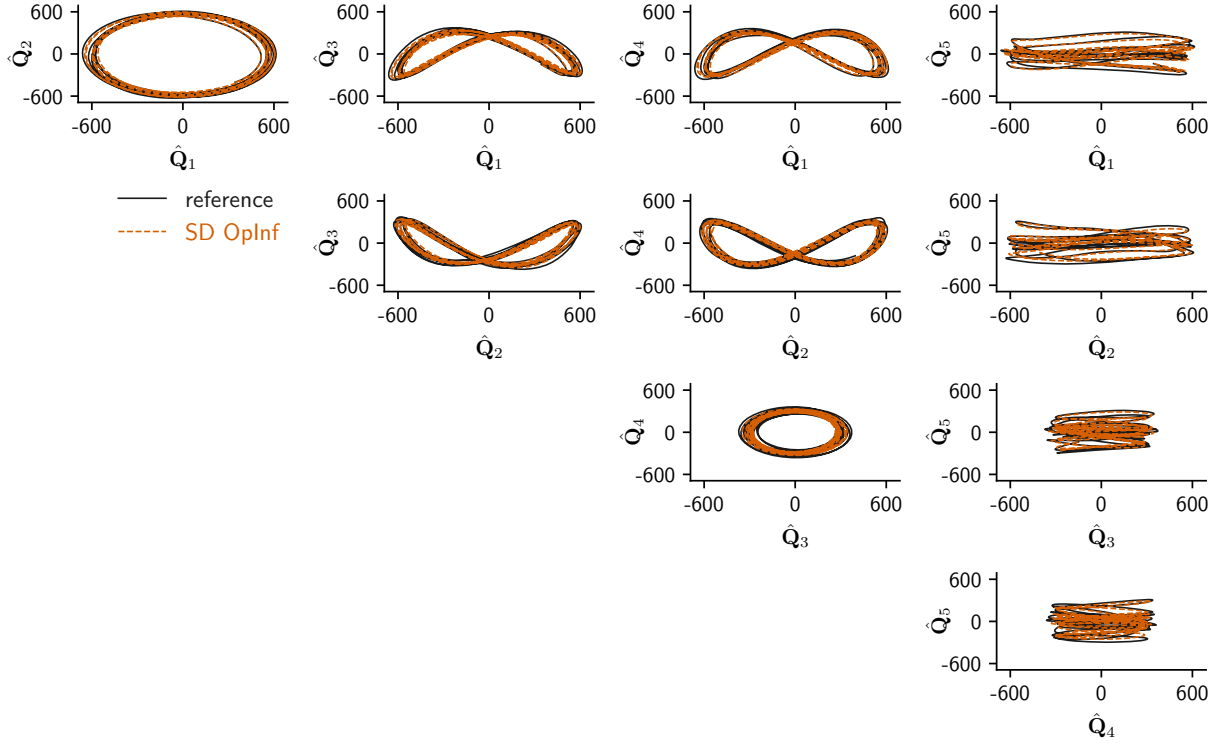


Fig. 8 Two-dimensional phase portraits of amplitudes over both training and prediction horizons corresponding to the five most energetic pressure POD modes.

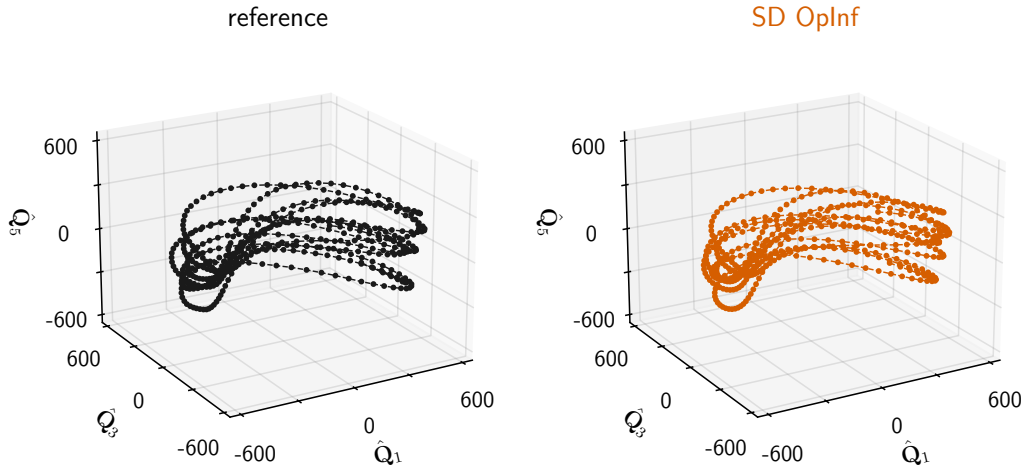


Fig. 9 Three-dimensional phase portrait of amplitudes for the first, third, and fifth most energetic pressure POD modes over both training and prediction horizons.

correlations whereas the pairs involving the fifth amplitude are much less correlated. This is consistent with the modes' structure observed in Figure 6. Furthermore, the SD OpInf approximate solutions closely match the reference results. While the 2D phase portraits involving the fifth amplitude do not exhibit significant correlations with other amplitudes, stronger correlations are observed in three-dimensional phase portraits. Figure 9 plots the 3D phase portrait depending

on the amplitudes corresponding to the first, third, and fifth most energetic modes.

Figure 10 plots the squared L_2 relative errors of the SD and DD-4 OpInf ROMs for pressure (top left), temperature (top right), and fuel (bottom left) and oxidizer (bottom right) mass fractions. These four variables generally represent the most relevant quantities in RDRE simulations. We note that these errors were computed with respect to the original state variables. DD-4 OpInf leads to some accuracy gains for all four variables, for both training and prediction. The squared relative error for pressure is decreased by up to 12.60% for training and up to 12.80% for prediction. The gains for the other three variables are less significant, around 5% for both training and prediction. Nevertheless, these results indicate that domain decomposition can be used to construct accurate and scalable data-driven ROMs for complex and large-scale applications such as RDREs.

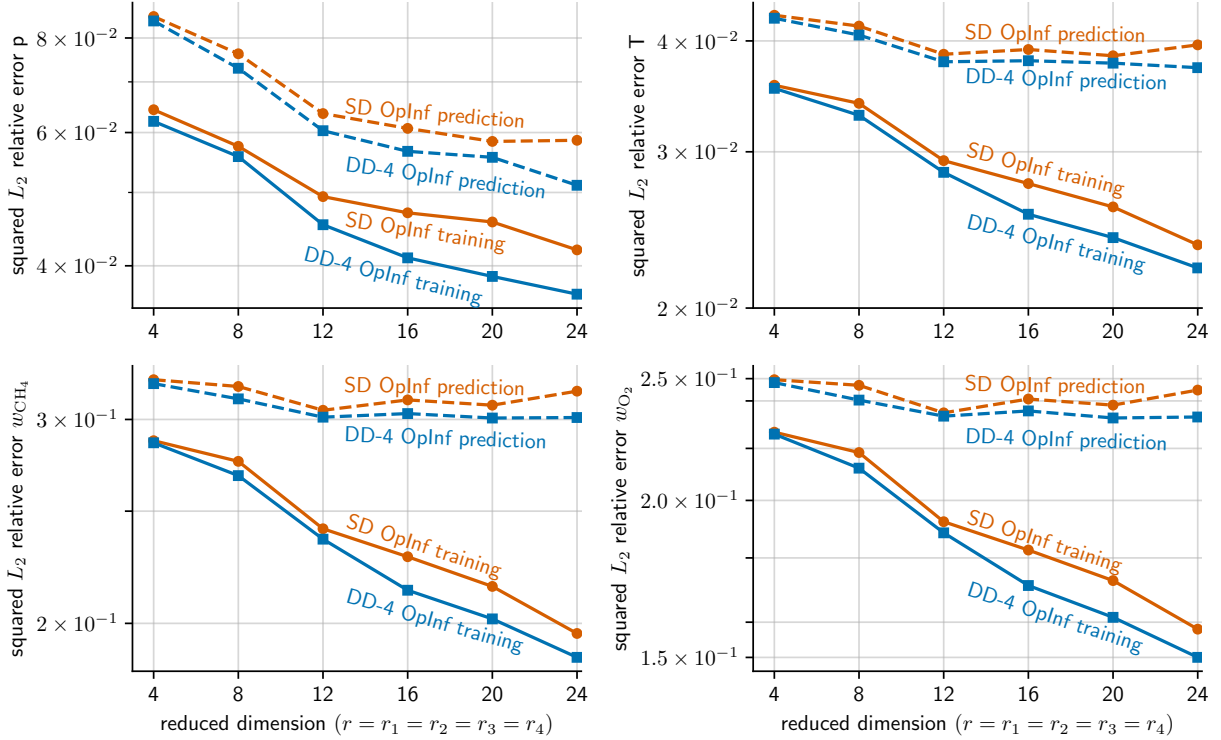


Fig. 10 Squared L_2 relative errors for pressure (top left), temperature (top right), and fuel (bottom left) and oxidizer (bottom right) mass fractions, for both training and prediction horizons.

For a more detailed understanding of the accuracy of the SD and DD-4 OpInf ROMs, we also compute the distribution of pointwise relative errors in the entire spatial domain, over both training and prediction horizons. Figure 11 compares the spatio-temporal distribution of relative errors for pressure in the original coordinates, denoted by $re(p)$. From bottom to top, we plot the percentages of spatial degrees of freedom that satisfy $re(p) \leq 5\%$ (bottom), $5\% < re(p) \leq 10\%$ (second plot from the bottom), $10 < re(p) \leq 20\%$ (second plot from the top), and $re(p) > 20\%$ (top). The percentages of spatial cells with large relative errors ($re(p) > 20\%$) are decreased by up to 6.60% (corresponding to 277,477 spatial cells in total) over the training horizon and up to 15.20% over the prediction horizon (corresponding to 639,038 spatial cells in total). This, in turn, leads to an increase in the percentages of cells with small errors ($re(p) \leq 10\%$).

We next use the ROM solutions obtained with the largest reduced dimensions ($r = r_1 = r_2 = r_3 = r_4 = 24$) to extract one-dimensional radial profiles at three representative locations close to the mid-channel. Figures 12–15 plot these profiles for pressure (Figure 12), temperature (Figure 13), fuel mass fraction (Figure 14), and oxidizer mass fraction (Figure 15). Axially, the first location is close to the injectors, the second location is further way from the injectors but still within the detonation region, and the third location is downstream of the detonation zone. The results corresponding to these three locations are respectively plotted in the first, second, and third column in each figure. The four rows in each figure plot the profiles at four time instants. The top row plots the profiles at the last time instant in the training horizon ($t = 3.5000$ ms). The remaining three rows plot the profiles at three time instants in the prediction horizon: the




Fig. 11 Distribution of pointwise relative errors for pressure over time interval [2.50, 3.75] ms. The dashed vertical lines mark the end of the training horizon.

40th ($t = 3.5375$ ms), the 100th ($t = 3.6875$ ms), and the 126th ($t = 3.7500$ ms) time instant. For pressure, the DD-4 OpInf accuracy gains observed in Figures 10 and 11 correspond to more accurate predictions for both the frequency and amplitude of the profiles. For the other three quantities, we do not observe a significant gain in accuracy due to DD-4 OpInf, which is consistent with the error plots in Figures 10. Overall, the ROMs perform well given the complexity of the underlying dynamics, which is clearly visible in the plotted profiles, especially at the location closer to the injectors.

Lastly, the runtimes of the SD and DD-4 OpInf ROMs, averaged over 1,000 runs, are $0.0296 \pm 5.012 \times 10^{-4}$ and $0.1256 \pm 4.4465 \times 10^{-4}$ seconds respectively. This translates into 12 (for SD OpInf) and 11 (for DD-4 OpInf) orders of magnitude reduction compared to the runtime of the high-fidelity model.

V. Conclusions

This paper introduced a domain decomposition formulation into the construction of a data-driven reduced model of large-scale systems. Our results for a large-scale simulation of a three-dimensional unsteady rotating detonation rocket engine combustion chamber with a sparse training dataset showed that domain decomposition improves the accuracy of the equivalent single-domain approach while decreasing the overall computational burden for preprocessing the training data. This reduction in computational costs is essential for enabling the construction of data-driven reduced models of large-scale applications. Overall, given the complexity of the scenario under consideration, our results show that domain decomposition can indeed provide a viable approach to constructing accurate, predictive, and scalable data-driven reduced models of large-scale systems. While these results illustrate the potential of the approach, they also highlight a number of open questions. Important future directions of work include alternative formulations of

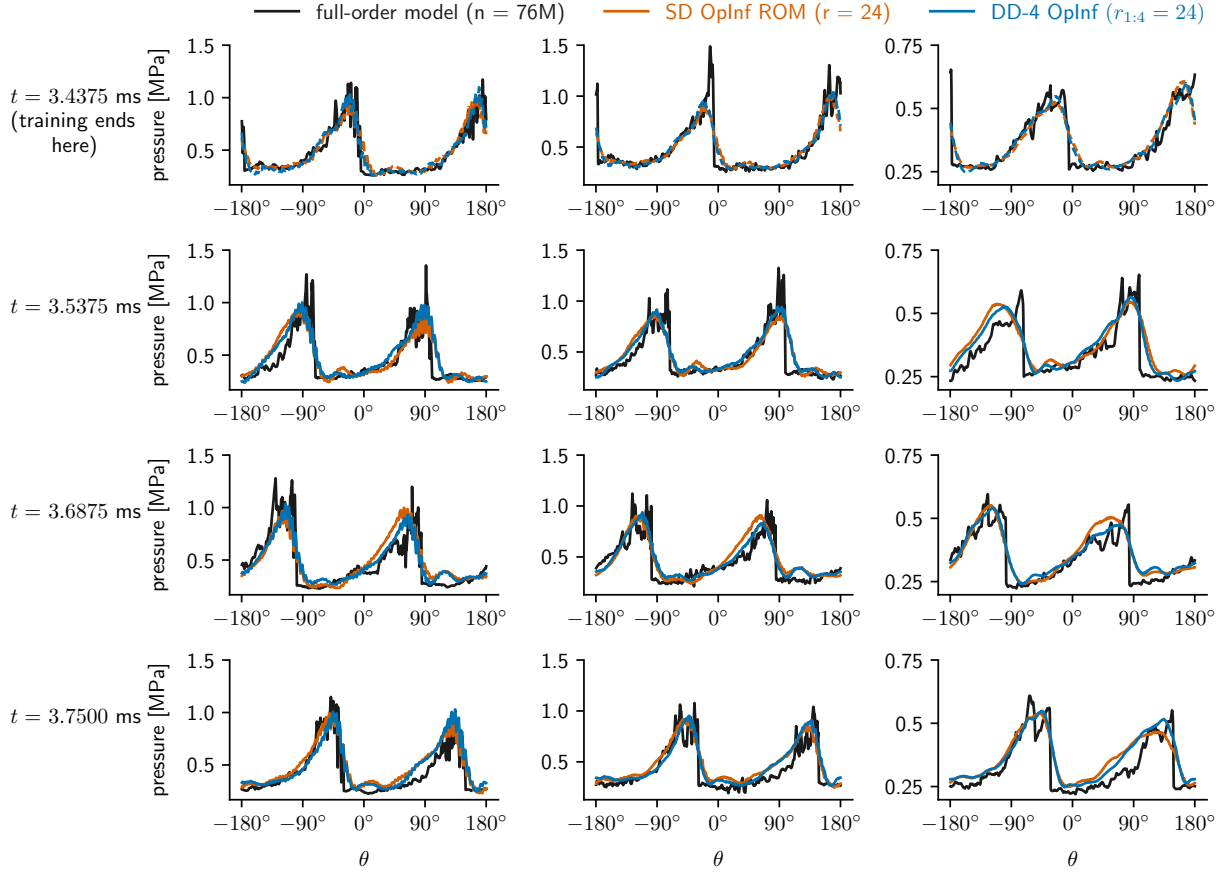


Fig. 12 One-dimensional circumferential pressure profile predictions. The columns plot the results at three representative locations close to the mid-channel. The rows plot the profiles at four representative time instants.

the domain-decomposed problem that explicitly enforce continuity across subdomains using constraints, strategies to optimally decompose the domain in a data-driven adaptive fashion, and an approach to identify optimal data scalings and transformations that may embed physical constraints in the reduced-order basis. In the broader context of aerospace propulsion, constructing reduced models that maintain the desired engineering accuracy while drastically decreasing the cost of computationally expensive high-fidelity simulations can have a strong impact and ultimately enable tasks such as design optimization which would be almost impossible to perform otherwise.

Acknowledgments

This work was supported in part by AFRL Grant FA9300-22-1-0001 and the Air Force Center of Excellence on Multifidelity Modeling of Rocket Combustor Dynamics under grant FA9550-17-1-0195. The views expressed are those of the author and do not necessarily reflect the official policy or position of the Department of the Air Force, the Department of Defense, or the U.S. government.

Distribution Statement A: Approved for Public Release; Distribution is Unlimited. PA# AFRL-2023-5401

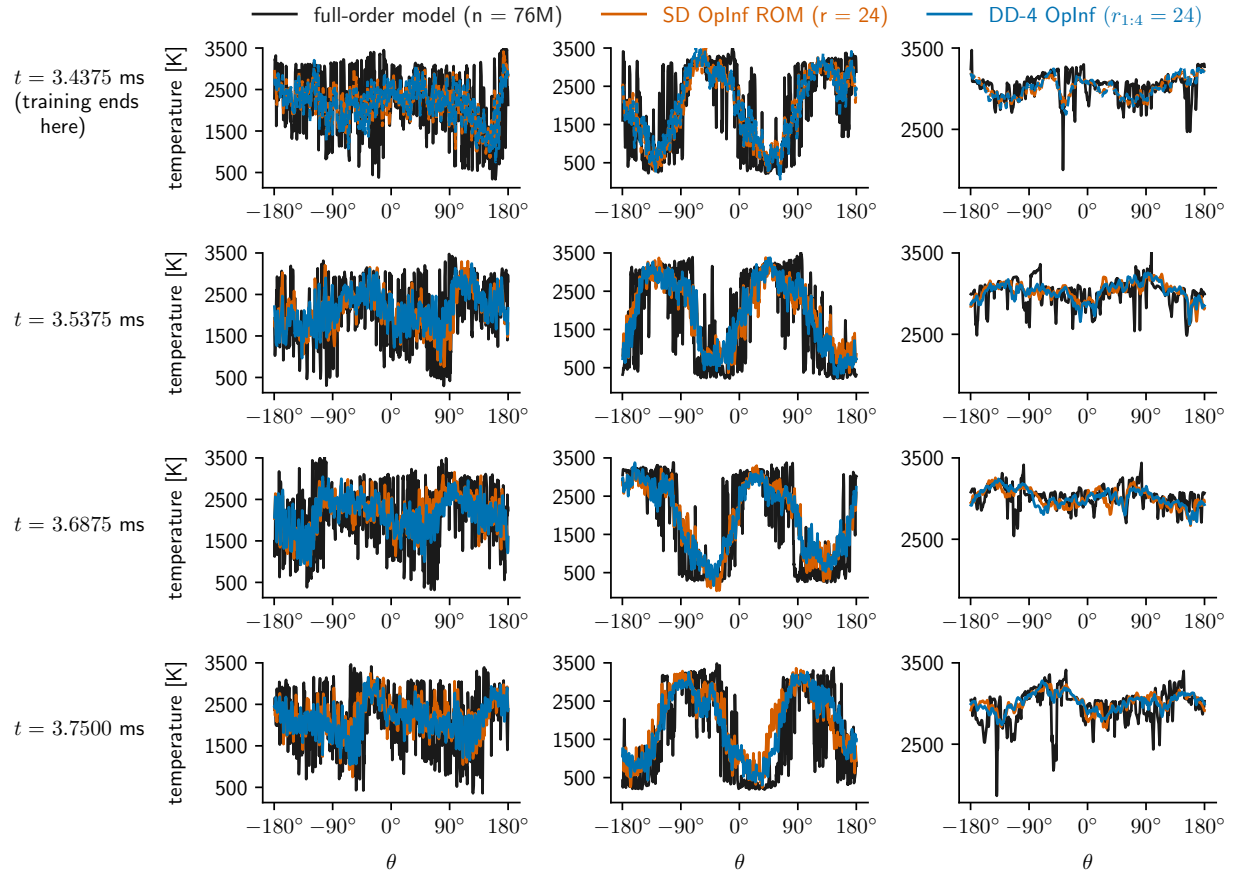


Fig. 13 One-dimensional circumferential temperature profile predictions. The columns plot the results at three representative locations close to the mid-channel. The rows plot the profiles at four representative time instants.

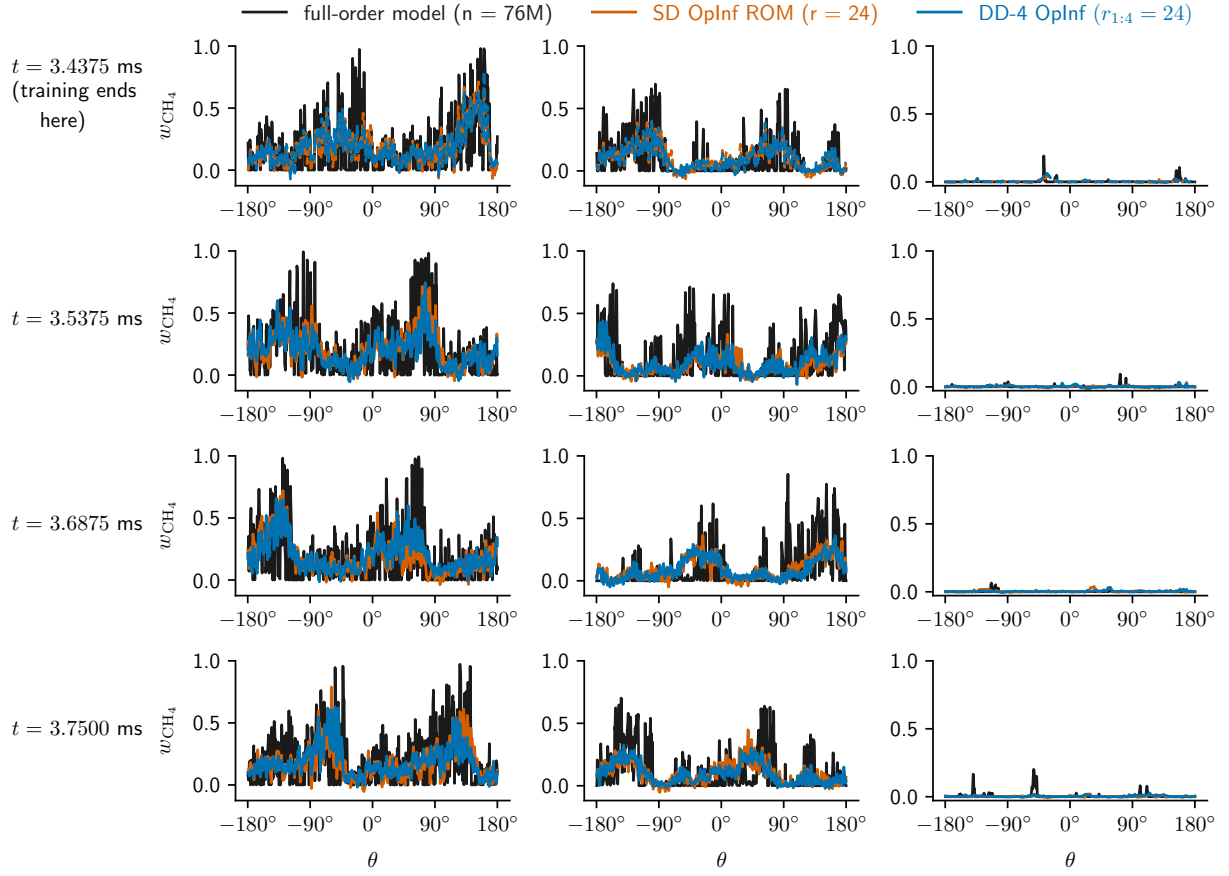


Fig. 14 One-dimensional circumferential w_{CH_4} profile predictions. The columns plot the results at three representative locations close to the mid-channel. The rows plot the profiles at four representative time instants.

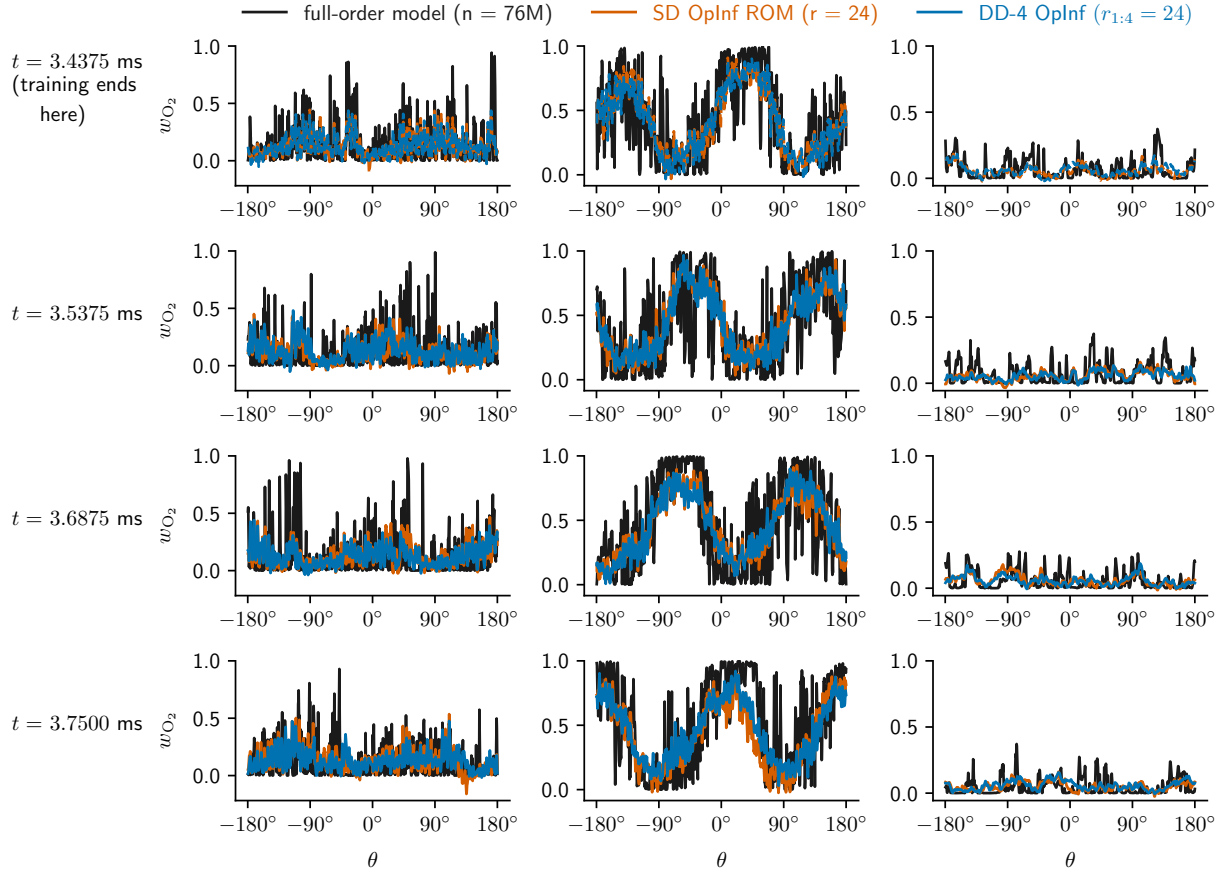


Fig. 15 One-dimensional circumferential w_{O_2} profile predictions. The columns plot the results at three representative locations close to the mid-channel. The rows plot the profiles at four representative time instants.

References

- [1] Dolean, V., Jolivet, P., and Nataf, F., *An Introduction to Domain Decomposition Methods*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2015. <https://doi.org/10.1137/1.9781611974065>.
- [2] Gropp, W. D., and Keyes, D. E., “Domain decomposition methods in computational fluid dynamics,” *International Journal for Numerical Methods in Fluids*, Vol. 14, No. 2, 1992, pp. 147–165. <https://doi.org/https://doi.org/10.1002/fld.1650140203>.
- [3] Maday, Y., and Rønquist, E. M., “A Reduced-Basis Element Method,” *Journal of Scientific Computing*, Vol. 17, No. 1, 2002, pp. 447–459. <https://doi.org/10.1023/A:1015197908587>.
- [4] Eftang, J. L., and Patera, A. T., “A port-reduced static condensation reduced basis element method for large component-synthesized structures: approximation and *A Posteriori* error estimation,” *Advanced Modeling and Simulation in Engineering Sciences*, Vol. 1, No. 1, 2014, p. 3. <https://doi.org/10.1186/2213-7467-1-3>.
- [5] Huynh, D., Knezevic, D., and Patera, A., “A static condensation reduced basis element method: Complex problems,” *Computer Methods in Applied Mechanics and Engineering*, Vol. 259, 2013, pp. 197–216. <https://doi.org/https://doi.org/10.1016/j.cma.2013.02.013>.
- [6] Phuong Huynh, D. B., Knezevic, D. J., and Patera, A. T., “A Static condensation Reduced Basis Element method : approximation and a *posteriori* error estimation,” *ESAIM: Mathematical Modelling and Numerical Analysis - Modélisation Mathématique et Analyse Numérique*, Vol. 47, No. 1, 2013, pp. 213–251. <https://doi.org/10.1051/m2an/2012022>.
- [7] Smetana, K., and Patera, A. T., “Optimal Local Approximation Spaces for Component-Based Static Condensation Procedures,” *SIAM Journal on Scientific Computing*, Vol. 38, No. 5, 2016, pp. A3318–A3356. <https://doi.org/10.1137/15M1009603>.
- [8] Iapichino, L., Quarteroni, A., and Rozza, G., “Reduced basis method and domain decomposition for elliptic problems in networks and complex parametrized geometries,” *Computers & Mathematics with Applications*, Vol. 71, No. 1, 2016, pp. 408–430.
- [9] Hoang, C., Choi, Y., and Carlberg, K., “Domain-decomposition least-squares Petrov–Galerkin (DD-LSPG) nonlinear model reduction,” *Computer Methods in Applied Mechanics and Engineering*, Vol. 384, 2021, p. 113997. <https://doi.org/https://doi.org/10.1016/j.cma.2021.113997>.
- [10] Smetana, K., and Taddei, T., “Localized Model Reduction for Nonlinear Elliptic Partial Differential Equations: Localized Training, Partition of Unity, and Adaptive Enrichment,” *SIAM Journal on Scientific Computing*, Vol. 45, No. 3, 2023, pp. A1300–A1331. <https://doi.org/10.1137/22M148402X>.
- [11] Huang, C., Duraisamy, K., and Merkle, C., “Component-Based Reduced Order Modeling of Large-Scale Complex Systems,” *Frontiers in Physics*, Vol. 10, 2022. <https://doi.org/10.3389/fphy.2022.900064>.
- [12] Huang, C., and Duraisamy, K., “Predictive reduced order modeling of chaotic multi-scale problems using adaptively sampled projections,” *Journal of Computational Physics*, Vol. 491, 2023, p. 112356. <https://doi.org/https://doi.org/10.1016/j.jcp.2023.112356>.
- [13] Peherstorfer, B., “Model Reduction for Transport-Dominated Problems via Online Adaptive Bases and Adaptive Sampling,” *SIAM Journal on Scientific Computing*, Vol. 42, No. 5, 2020, pp. A2803–A2836. <https://doi.org/10.1137/19M1257275>.
- [14] Discacciati, N., and Hesthaven, J. S., “Localized model order reduction and domain decomposition methods for coupled heterogeneous systems,” *International Journal for Numerical Methods in Engineering*, Vol. 124, No. 18, 2023, pp. 3964–3996. <https://doi.org/https://doi.org/10.1002/nme.7295>.
- [15] Xiao, D., Fang, F., Heaney, C., Navon, I., and Pain, C., “A domain decomposition method for the non-intrusive reduced order modelling of fluid flow,” *Computer Methods in Applied Mechanics and Engineering*, Vol. 354, 2019, pp. 307–330. <https://doi.org/https://doi.org/10.1016/j.cma.2019.05.039>.
- [16] Xiao, D., Heaney, C., Fang, F., Mottet, L., Hu, R., Bistrian, D., Aristodemou, E., Navon, I., and Pain, C., “A domain decomposition non-intrusive reduced order model for turbulent flows,” *Computers & Fluids*, Vol. 182, 2019, pp. 15–27. <https://doi.org/https://doi.org/10.1016/j.compfluid.2019.02.012>.
- [17] Antil, H., Heinkenschloss, M., Hoppe, R. H. W., and Sorensen, D. C., “Domain decomposition and model reduction for the numerical solution of PDE constrained optimization problems with localized optimization variables,” *Computing and Visualization in Science*, Vol. 13, No. 6, 2010, pp. 249–264. <https://doi.org/10.1007/s00791-010-0142-4>.

- [18] Antil, H., Heinkenschloss, M., and Hoppe, R. H., “Domain decomposition and balanced truncation model reduction for shape optimization of the Stokes system,” *Optimization Methods and Software*, Vol. 26, No. 4-5, 2011, pp. 643–669. <https://doi.org/10.1080/10556781003767904>.
- [19] Callahan, J. L., Maeda, K., and Brunton, S. L., “Robust flow reconstruction from limited measurements via sparse representation,” *Physical Review Fluids*, Vol. 4, 2019, p. 103907. <https://doi.org/10.1103/PhysRevFluids.4.103907>.
- [20] Haller, G., and Ponsioen, S., “Nonlinear normal modes and spectral submanifolds: existence, uniqueness and use in model reduction,” *Nonlinear Dynamics*, Vol. 86, No. 3, 2016, pp. 1493–1534. <https://doi.org/10.1007/s11071-016-2974-z>.
- [21] Cenedese, M., Axås, J., Bäuerlein, B., Avila, K., and Haller, G., “Data-driven modeling and prediction of non-linearizable dynamics via spectral submanifolds,” *Nature Communications*, Vol. 13, No. 1, 2022, p. 872. <https://doi.org/10.1038/s41467-022-28518-y>.
- [22] Peherstorfer, B., and Willcox, K., “Data-driven operator inference for nonintrusive projection-based model reduction,” *Computer Methods in Applied Mechanics and Engineering*, Vol. 306, 2016, pp. 196–215. <https://doi.org/https://doi.org/10.1016/j.cma.2016.03.025>.
- [23] Kramer, B., Peherstorfer, B., and Willcox, K. E., “Learning Nonlinear Reduced Models from Data with Operator Inference,” *Annual Review of Fluid Mechanics*, Vol. 56, No. Volume 56, 2024, 2024, pp. 521–548. <https://doi.org/https://doi.org/10.1146/annurev-fluid-121021-025220>.
- [24] Qian, E., Kramer, B., Peherstorfer, B., and Willcox, K., “Lift & Learn: Physics-informed machine learning for large-scale nonlinear dynamical systems,” *Physica D: Nonlinear Phenomena*, Vol. 406, 2020, p. 132401. <https://doi.org/https://doi.org/10.1016/j.physd.2020.132401>.
- [25] Schmid, P. J., “Dynamic mode decomposition of numerical and experimental data,” *Journal of Fluid Mechanics*, Vol. 656, 2010, p. 5–28. <https://doi.org/10.1017/S0022112010001217>.
- [26] Qian, E., Farcaş, I.-G., and Willcox, K., “Reduced Operator Inference for Nonlinear Partial Differential Equations,” *SIAM Journal on Scientific Computing*, Vol. 44, No. 4, 2022, pp. A1934–A1959. <https://doi.org/10.1137/21M1393972>.
- [27] Scheffel, J., and Lindvall, K., “Temporal smoothing - A step forward for time-spectral methods,” *Computer Physics Communications*, Vol. 270, 2022, p. 108173. <https://doi.org/https://doi.org/10.1016/j.cpc.2021.108173>.
- [28] Golub, G. H., and Loan, C. F. V., *Matrix Computations*, The Johns Hopkins University Press, Baltimore, MD, 2013.
- [29] Kutz, J. N., Brunton, S. L., Brunton, B. W., and Proctor, J. L., *Dynamic Mode Decomposition*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2016. <https://doi.org/10.1137/1.9781611974508>.
- [30] Tu, J. H., Rowley, C. W., Luchtenburg, D. M., Brunton, S. L., and Kutz, J. N., “On dynamic mode decomposition: Theory and applications,” *Journal of Computational Dynamics*, Vol. 1, No. 2, 2014, pp. 391–421.
- [31] Farcas, I., Gundevia, R., Munipalli, R., and Willcox, K. E., “Parametric non-intrusive reduced-order models via operator inference for large-scale rotating detonation engine simulations,” *AIAA Scitech 2023 Forum*, 2023. <https://doi.org/10.2514/6.2023-0172>.
- [32] Farcas, I., Munipalli, R., and Willcox, K. E., “On filtering in non-intrusive data-driven reduced-order modeling,” *AIAA AVIATION 2022 Forum*, 2022. <https://doi.org/10.2514/6.2022-3487>.
- [33] McQuarrie, S. A., Huang, C., and Willcox, K. E., “Data-driven reduced-order models via regularised Operator Inference for a single-injector combustion process,” *Journal of the Royal Society of New Zealand*, Vol. 51, No. 2, 2021, pp. 194–211. <https://doi.org/10.1080/03036758.2020.1863237>.
- [34] Rubini, R., Lasagna, D., and Da Ronch, A., “The l1-based sparsification of energy interactions in unsteady lid-driven cavity flow,” *Journal of Fluid Mechanics*, Vol. 905, 2020, p. A15. <https://doi.org/10.1017/jfm.2020.707>.
- [35] Loiseau, J.-C., and Brunton, S. L., “Constrained sparse Galerkin regression,” *Journal of Fluid Mechanics*, Vol. 838, 2018, pp. 42–67. <https://doi.org/10.1017/jfm.2017.823>.
- [36] Sharma, H., Wang, Z., and Kramer, B., “Hamiltonian operator inference: Physics-preserving learning of reduced-order models for canonical Hamiltonian systems,” *Physica D: Nonlinear Phenomena*, Vol. 431, 2022, p. 133122. <https://doi.org/https://doi.org/10.1016/j.physd.2021.133122>.
- [37] Gruber, A., and Tezaur, I., “Canonical and noncanonical Hamiltonian operator inference,” *Computer Methods in Applied Mechanics and Engineering*, Vol. 416, 2023, p. 116334. <https://doi.org/https://doi.org/10.1016/j.cma.2023.116334>.

- [38] Sharma, H., and Kramer, B., “Preserving Lagrangian structure in data-driven reduced-order modeling of large-scale dynamical systems,” *Physica D: Nonlinear Phenomena*, Vol. 462, 2024, p. 134128. <https://doi.org/https://doi.org/10.1016/j.physd.2024.134128>.
- [39] Baiges, J., Codina, R., and Idelsohn, S., “A domain decomposition strategy for reduced order models. Application to the incompressible Navier–Stokes equations,” *Computer Methods in Applied Mechanics and Engineering*, Vol. 267, 2013, pp. 23–42. <https://doi.org/https://doi.org/10.1016/j.cma.2013.08.001>.
- [40] Lucia, D. J., King, P. I., and Beran, P. S., “Reduced order modeling of a two-dimensional flow with moving shocks,” *Computers & Fluids*, Vol. 32, No. 7, 2003, pp. 917–938. [https://doi.org/https://doi.org/10.1016/S0045-7930\(02\)00035-X](https://doi.org/https://doi.org/10.1016/S0045-7930(02)00035-X).
- [41] Bykovskii, F., Zhdan, S., and Vedernikov, E., “Continuous Spin Detonations,” *Journal of Propulsion and Power*, Vol. 22, 2006, pp. 1204–1216. <https://doi.org/10.2514/1.17656>.
- [42] Smith, G. P., Tao, Y., and Wang, H., “Foundational Fuel Chemistry Model Version 1.0 (FFCM-1),” , 2016.
- [43] Batista, A., Ross, M., Lietz, C., and Hargus, W. A., “Detonation Wave Interaction Classification in a Rotating Detonation Rocket Engine,” *AIAA Propulsion and Energy 2020 Forum*, 2020. <https://doi.org/10.2514/6.2020-3861>.
- [44] Lietz, C., Desai, Y., Munipalli, R., Schumaker, S. A., and Sankaran, V., “Flowfield analysis of a 3D simulation of a rotating detonation rocket engine,” *AIAA Scitech 2019 Forum*, 2019. <https://doi.org/10.2514/6.2019-1009>.
- [45] Bennewitz, J. W., Burr, J. R., Bigler, B. R., Burke, R. F., Lemcherfi, A., Mundt, T., Rezzag, T., Plaehn, E. W., Sosa, J., Walters, I. V., Schumaker, S. A., Ahmed, K. A., Slabaugh, C. D., Knowlen, C., and Hargus, W. A., “Experimental validation of rotating detonation for rocket propulsion,” *Scientific Reports*, Vol. 13, No. 1, 2023, p. 14204. <https://doi.org/10.1038/s41598-023-40156-y>.
- [46] Swischuk, R., Kramer, B., Huang, C., and Willcox, K., “Learning Physics-Based Reduced-Order Models for a Single-Injector Combustion Process,” *AIAA Journal*, Vol. 58, No. 6, 2020, pp. 2658–2672. <https://doi.org/10.2514/1.J058943>.
- [47] Rowley, C. W., Colonius, T., and Murray, R. M., “Model reduction for compressible flows using POD and Galerkin projection,” *Physica D: Nonlinear Phenomena*, Vol. 189, No. 1-2, 2004, pp. 115–129. <https://doi.org/https://doi.org/10.1016/j.physd.2003.03.001>.
- [48] Vogel, E. A., and Coder, J. G., “A novel entropy normalization scheme for characterization of highly compressible flows,” *Theoretical and Computational Fluid Dynamics*, Vol. 36, No. 4, 2022, pp. 641–670. <https://doi.org/10.1007/s00162-022-00617-y>.
- [49] Lumley, J. L., and Poje, A., “Low-dimensional models for flows with density fluctuations,” *Physics of Fluids*, Vol. 9, No. 7, 1997, pp. 2023–2031. <https://doi.org/10.1063/1.869321>.
- [50] Halko, N., Martinsson, P. G., and Tropp, J. A., “Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions,” *SIAM Review*, Vol. 53, No. 2, 2011, pp. 217–288. <https://doi.org/10.1137/090771806>.
- [51] Brand, M., “Fast low-rank modifications of the thin singular value decomposition,” *Linear Algebra and its Applications*, Vol. 415, No. 1, 2006, pp. 20–30. <https://doi.org/https://doi.org/10.1016/j.laa.2005.07.021>.
- [52] Berkooz, G., Holmes, P., and Lumley, J. L., “The Proper Orthogonal Decomposition in the Analysis of Turbulent Flows,” *Annual Review of Fluid Mechanics*, Vol. 25, No. 1, 1993, pp. 539–575. <https://doi.org/10.1146/annurev.fl.25.010193.002543>.
- [53] Chang, Y.-H., Wang, X., Zhang, L., Li, Y., Mak, S., Wu, C.-F. J., and Yang, V., “Reduced-Order Modeling for Complex Flow Emulation by Common Kernel-Smoothed Proper Orthogonal Decomposition,” *AIAA Journal*, Vol. 59, No. 9, 2021, pp. 3291–3303. <https://doi.org/10.2514/1.J060574>.