

NON-INTRUSIVE REDUCED-ORDER MODELS FOR PARAMETRIC PARTIAL DIFFERENTIAL EQUATIONS VIA DATA-DRIVEN OPERATOR INFERENCE

SHANE A. MCQUARRIE*, PARISA KHODABAKHSHI†, AND KAREN E. WILLCOX*

Abstract. This work formulates a new approach to reduced modeling of parameterized, time-dependent partial differential equations (PDEs). The method employs Operator Inference, a scientific machine learning framework combining data-driven learning and physics-based modeling. The parametric structure of the governing equations is embedded directly into the reduced-order model, and parameterized reduced-order operators are learned via a data-driven linear regression problem. The result is a reduced-order model that can be solved rapidly to map parameter values to approximate PDE solutions. Such parameterized reduced-order models may be used as physics-based surrogates for uncertainty quantification and inverse problems that require many forward solves of parametric PDEs. Numerical issues such as well-posedness and the need for appropriate regularization in the learning problem are considered, and an algorithm for hyperparameter selection is presented. The method is illustrated for a parametric heat equation and demonstrated for the FitzHugh–Nagumo neuron model.

Key words. Parametric model reduction, operator inference, scientific machine learning, data-driven reduced model

MSC codes. 35B30, 35R30, 65F22

1. Introduction. Model reduction seeks to alleviate the computational burden of large-scale numerical simulations of dynamical systems by constructing reduced-order models (ROMs) that accurately capture the system dynamics, but which are much less expensive to solve than the high-fidelity models inherent in applications. The challenge is to generate ROMs from limited training data that respond well to changes in the scenario parameters that define the governing dynamics [8]. Such parametric ROMs are critical for enabling outer-loop applications such as design, inverse problems, optimization, and uncertainty quantification. Furthermore, as high-fidelity simulations become increasingly sophisticated and simulation data becomes more available, there is a growing need for non-intrusive model reduction methods, which aim to learn ROMs primarily from simulation data and/or outputs, as opposed to making a direct reduction of the underlying high-fidelity code that produced them [18]. Non-intrusive approaches combine data-driven learning with physics-based modeling in a way that enables both flexibility and robustness. This paper presents a framework for learning parametric ROMs in a non-intrusive fashion.

Adapting non-intrusive model reduction strategies to the parametric setting is an active area of research. One major model reduction strategy, Dynamic mode decomposition (DMD) [13, 44], learns a low-dimensional linear mapping based on state space data, approximating the eigenstates of the infinite-dimensional Koopman operator. The work in [49] targets parametric problems by incorporating DMD with an active subspace strategy to reduce the dimensionality of the parameter space. Methods based on the Loewner framework [2, 3], another common non-intrusive approach to model reduction, build a ROM based on input-output measurements and transfer functions. Loewner methods have been generalized to parameterized linear systems

*Oden Institute for Computational Engineering and Sciences, University of Texas at Austin, Austin, TX, USA (shanemcq@utexas.edu, kwillcox@oden.utexas.edu).

†Department of Mechanical Engineering and Mechanics, Lehigh University, Bethlehem, PA, USA (pak322@lehigh.edu).

by introducing additional degrees of freedom in the construction of the reduced-order transfer function to account for parametric dependencies [23]. Recent work in [10] blends ideas from classical model reduction with deep learning to construct a mapping from parametric inputs to state outputs. Deep learning approaches to model reduction aim to benefit from the flexibility of representing the state on a low-dimensional but nonlinear manifold [28].

Equation discovery methods, in which the governing equations of a dynamical system are learned from data, share some characteristics with non-intrusive model reduction methods. One class of equation discovery approaches uses sparse regression to identify the underlying partial differential equations (PDEs) from a set of data [43] or the key terms in a dynamical system within a library of potential nonlinear terms [12, 40, 42]. In a similar vein, the work in [4, 29] combines machine learning techniques with manifold learning algorithms to learn a macroscopic model for long-wavelength behavior corresponding to fine-scale measurement data. Each of these methods relies on an appropriate candidate library for the terms of the unknown equations and selects the best combination of terms based on data. In some cases where the original data is of high dimension, a dimensionality reduction technique is first applied so that the model selection occurs in a low-dimensional setting (see, e.g., [12, 16]). Model reduction approaches also utilize the paradigm of first seeking a low-dimensional representation for the system state which, for data-driven methods, leads to a learning problem posed in a reduced space.

Operator Inference (OpInf), introduced in [37], is a non-intrusive framework for model reduction of systems with polynomial nonlinearities. As with other non-intrusive approaches, the method does not require intrusive access to source code, instead inferring the ROM solely from initial conditions, simulation snapshots, and corresponding inputs. Known governing equations motivate the form of the ROM, and the operators defining the ROM are chosen by minimizing a data-driven residual in a reduced state space. The associated learning problem is linear, dense, relatively small, and has a closed-form solution. Since its introduction, the OpInf framework has been expanded in several ways: transforming variables (lifting) to induce the requisite polynomial structure [39, 48]; approximating nonpolynomial nonlinearities via the discrete empirical interpolation method (DEIM) [7]; regularizing the learning problem to enable performance on large-scale systems [24, 32]; re-projecting trajectories to exactly recover intrusive ROMs [36]; and accounting for algebraic equations arising from lifting transformations [26], to name a few. In these OpInf-based methods, parametric dependencies are addressed by learning separate ROMs for individual parameter samples, then interpolating either their reduced operators [36, 37] or their outputs [26]. However, interpolation in more than one or two parameter dimensions quickly becomes challenging due to the curse of dimensionality and Runge’s phenomenon.

In this paper, we show that the parametric structure of the governing equations of interest can be built directly into the OpInf regression problem—circumventing the need for interpolation—if the parametric dependencies have an affine form. Affine-parametric problems have been studied frequently in the context of reduced basis methods [21, 41, 54, 55], in which the preservation of the parametric structure by projection plays a key role [8]. The recent work [58] developed an OpInf framework for linear-quadratic affine-parametric problems in a semi-discrete setting, then applied the methodology successfully to the shallow water equations in non-traditional form with a single parameter. In this work, we pose the affine-parametric OpInf problem in the fully continuous setting and establish rigorous *a priori* conditions for determining the well-posedness of the inference problem. These conditions can be used to inform

the choice of parameter samples for training data. We extend the framework to systems of PDEs such that, different from [58], the system-level parametric structure is exactly preserved. Leveraging this structure preservation, we propose a robust regularization strategy which can be tailored to specific PDE operators. This step is key for obtaining stable and accurate ROMs. Our approach is applied to problems with two- and four-dimensional parameter spaces, demonstrating that our framework enables ROMs for parametric PDE systems with multiple parameters. In summary, the approach presented here includes the following key contributions: 1) formulating a general OpInf framework for affine-parametric systems of PDEs in a time-continuous setting, 2) analyzing the associated well-posedness conditions, and 3) implementing a flexible regularization strategy in the numerical solution of the inference problem.

The remainder of the paper is organized as follows. [Section 2](#) establishes the general methodology; [Section 3](#) extends the framework to systems of PDEs; [Section 4](#) details the computational aspects of solving the parametric OpInf problem; [Section 5](#) presents two numerical examples; and [Section 6](#) concludes the paper.

2. Non-intrusive Parametric Model Reduction. In [subsection 2.1](#), we show how the form of an appropriate ROM can be determined directly from the form of certain PDEs; [subsection 2.2](#) presents the OpInf approach for learning such ROMs from data and system structure. We also introduce a heat equation example for which numerical results are reported in [subsection 5.1](#).

2.1. Projection-based Reduced-order Models of Parametric PDEs. We target systems governed by parametric PDEs that are polynomial in state, which includes linear PDEs as well as a large class of nonlinear PDEs. Let $\Omega \subset \mathbb{R}^{d_x}$ be an open, bounded set with Lipschitz continuous boundary $\partial\Omega = \Gamma \cup (\partial\Omega \setminus \Gamma)$ and outward-pointing normal $\boldsymbol{\eta} \in \mathbb{R}^{d_x}$. For the time domain $[t_0, t_f] \subset \mathbb{R}$ and the parameter domain $\mathcal{P} \subset \mathbb{R}^{d_\mu}$, we consider the initial/boundary-value problem

$$(2.1a) \quad \frac{\partial u}{\partial t} = \mathcal{F}(u; \mu), \quad x \in \Omega, \quad t \in (t_0, t_f], \quad \mu \in \mathcal{P},$$

$$(2.1b) \quad u(x, t_0; \mu) = u_0(x; \mu), \quad x \in \Omega, \quad \mu \in \mathcal{P},$$

$$(2.1c) \quad u(x, t; \mu) = 0, \quad x \in \Gamma, \quad t \in [t_0, t_f], \quad \mu \in \mathcal{P},$$

$$(2.1d) \quad \boldsymbol{\eta} \cdot \nabla_x u(x, t; \mu) = 0, \quad x \in \partial\Omega \setminus \Gamma, \quad t \in [t_0, t_f], \quad \mu \in \mathcal{P},$$

where the unknown state variable $u(\cdot, t; \mu)$ is contained in a separable Hilbert space \mathcal{V} of real-valued functions satisfying the boundary conditions (2.1c)–(2.1d) with dual space \mathcal{V}^* , and $\mathcal{F} : \mathcal{V} \times \mathcal{P} \rightarrow \mathcal{V}^*$ is a spatial differential operator that depends on the free parameter $\mu \in \mathcal{P}$. We assume (2.1a)–(2.1d) has a unique solution u in at least the weak sense, meaning

$$(2.2) \quad \left\langle v, \frac{\partial u}{\partial t} \right\rangle = \langle v, \mathcal{F}(u; \mu) \rangle \quad \text{for all} \quad v \in \mathcal{V}, \quad t \in [t_0, t_f], \quad \mu \in \mathcal{P},$$

where $\langle \cdot, \cdot \rangle$ is the duality pairing of \mathcal{V} with \mathcal{V}^* .

We consider the setting in which \mathcal{F} has a polynomial structure with respect to the state u and its spatial derivatives. Many PDEs enjoy this structure or can be written in this form through a change of variables [39]. For brevity we consider a quadratic form, but higher-order (e.g., cubic) terms may also be included, as we will see later in [Example 3.1](#). Specifically, suppose

$$(2.3a) \quad \mathcal{F}(u; \mu) = \mathcal{C}(\mu) + \mathcal{A}(u; \mu) + \mathcal{H}(u, u; \mu),$$

where $\mathcal{C} : \mathcal{P} \rightarrow \mathcal{V}^*$, and where $\mathcal{A} : \mathcal{V} \times \mathcal{P} \rightarrow \mathcal{V}^*$ and $\mathcal{H} : \mathcal{V} \times \mathcal{V} \times \mathcal{P} \rightarrow \mathcal{V}^*$ are linear in each of their state arguments. Furthermore, assume that the operators \mathcal{C} , \mathcal{A} , and \mathcal{H} exhibit the following affine decompositions with respect to the parameter μ :

$$(2.3b) \quad \begin{aligned} \mathcal{C}(\mu) &= \sum_{p=1}^{q_c} \theta_c^{(p)}(\mu) \mathcal{C}^{(p)}, & \mathcal{A}(u; \mu) &= \sum_{p=1}^{q_A} \theta_A^{(p)}(\mu) \mathcal{A}^{(p)}(u), \\ \mathcal{H}(u, v; \mu) &= \sum_{p=1}^{q_H} \theta_H^{(p)}(\mu) \mathcal{H}^{(p)}(u, v), \end{aligned}$$

where $\mathcal{C}^{(p)} \in \mathcal{V}^*$, $\mathcal{A}^{(p)} : \mathcal{V} \rightarrow \mathcal{V}^*$, $\mathcal{H}^{(p)} : \mathcal{V} \times \mathcal{V} \rightarrow \mathcal{V}^*$, and the scalar-valued functions $\theta_c^{(p)}, \theta_A^{(p)}, \theta_H^{(p)} : \mathcal{P} \rightarrow \mathbb{R}$ are such that the sets $\{\theta_c^{(p)}\}_{p=1}^{q_c}$, $\{\theta_A^{(p)}\}_{p=1}^{q_A}$, and $\{\theta_H^{(p)}\}_{p=1}^{q_H}$ are each linearly independent. Note that the operators $\mathcal{C}^{(p)}$, $\mathcal{A}^{(p)}$, and $\mathcal{H}^{(p)}$ are independent of the parameter μ . This work considers the class of problems where the affine-parametric coefficient functions $(\theta_c^{(p)}, \theta_A^{(p)}, \theta_H^{(p)})$ are known. Such a structure may occur naturally in the known governing equations, be induced via approximation with the empirical interpolation method [6], or in some cases discovered from data via sparse regression [42, 43]. We leave for future work the possibility of simultaneously inferring the structure (2.3b) and constructing a ROM from data; see [30] for an example of simultaneous reduction in both the parameter and state space.

A projection-based ROM of (2.1a)–(2.1d) with \mathcal{F} as in (2.3a)–(2.3b) retains the affine-parametric polynomial structure of the system [8, 20]. Let $\{v_j\}_{j=1}^{\infty} \subset \mathcal{V}$ be an orthonormal set such that the solution u may be expressed with the expansion

$$(2.4) \quad u(x, t; \mu) = \sum_{j=1}^{\infty} \hat{u}_j(t; \mu) v_j(x).$$

Since $\langle v_i, v_j \rangle = \delta_{ij}$, the coefficients satisfy $\hat{u}_j(t; \mu) = \langle v_j, u(\cdot, t; \mu) \rangle$. A reduced model with $r \in \mathbb{N}$ degrees of freedom consists of time evolution equations for the coefficients $\hat{u}_1(t; \mu), \dots, \hat{u}_r(t; \mu)$; the approximate ROM solution \check{u} of (2.1a)–(2.1d) is then given by the sum (2.4), truncated to r terms:

$$(2.5) \quad \check{u}(x, t; \mu) = \sum_{j=1}^r \hat{u}_j(t; \mu) v_j(x).$$

Note that \check{u} is confined to the finite-dimensional subspace $\text{span}(\{v_1, \dots, v_r\}) \subset \mathcal{V}$. By substituting \check{u} for u in (2.2) with test function $v = v_i$, and using the form of \mathcal{F} from (2.3a)–(2.3b), we obtain

$$(2.6) \quad \begin{aligned} \frac{d\hat{u}_i}{dt} &= \sum_{p=1}^{q_c} \theta_c^{(p)}(\mu) \langle v_i, \mathcal{C}^{(p)} \rangle + \sum_{p=1}^{q_A} \theta_A^{(p)}(\mu) \sum_{j=1}^r \langle v_i, \mathcal{A}^{(p)}(v_j) \rangle \hat{u}_j \\ &\quad + \sum_{p=1}^{q_H} \theta_H^{(p)}(\mu) \sum_{j=1}^r \sum_{k=1}^r \langle v_i, \mathcal{H}^{(p)}(v_j, v_k) \rangle \hat{u}_j \hat{u}_k. \end{aligned}$$

Collecting (2.6) for $i = 1, \dots, r$ yields a uniquely-defined system of ordinary differential equations (ODEs) with state vector $\hat{\mathbf{u}}(t; \mu) = [\hat{u}_1(t; \mu) \ \dots \ \hat{u}_r(t; \mu)]^\top \in \mathbb{R}^r$,

$$(2.7a) \quad \frac{d}{dt} \hat{\mathbf{u}}(t; \mu) = \left(\sum_{p=1}^{q_c} \theta_c^{(p)}(\mu) \hat{\mathbf{c}}^{(p)} \right) + \left(\sum_{p=1}^{q_A} \theta_A^{(p)}(\mu) \hat{\mathbf{A}}^{(p)} \right) \hat{\mathbf{u}}(t; \mu) \\ + \left(\sum_{p=1}^{q_H} \theta_H^{(p)}(\mu) \hat{\mathbf{H}}^{(p)} \right) (\hat{\mathbf{u}}(t; \mu) \hat{\odot} \hat{\mathbf{u}}(t; \mu)),$$

$$(2.7b) \quad \hat{\mathbf{u}}(t_0; \mu) = \left[\langle v_1, u_0(\mu) \rangle \quad \cdots \quad \langle v_r, u_0(\mu) \rangle \right]^\top \in \mathbb{R}^r,$$

where $\hat{\mathbf{c}}^{(p)} \in \mathbb{R}^r$, $\hat{\mathbf{A}}^{(p)} \in \mathbb{R}^{r \times r}$, $\hat{\mathbf{H}}^{(p)} \in \mathbb{R}^{r \times \binom{r+1}{2}}$, and $\hat{\odot}$ denotes a compact Khatri-Rao product, i.e., $\mathbf{u} \hat{\odot} \mathbf{u}$ extracts the unique components of the Khatri-Rao product of \mathbf{u} with itself (see [Appendix A](#)). Note that (2.7a) and (2.3a)–(2.3b) are both quadratic in their respective state and affine-parametric with respect to μ , with the same functions $\theta_c^{(p)}$, $\theta_A^{(p)}$, and $\theta_H^{(p)}$ in the affine expansions.

The system (2.7a)–(2.7b) is a ROM for (2.1a)–(2.1d) in which the boundary conditions are embedded through the basis functions. The quality of the ROM as a surrogate for (2.1a)–(2.1d) depends heavily on the basis $\{v_j\}_{j=1}^r$, but the *form* of the equations is the same for any choice of basis. In [subsection 2.2](#), we leverage this property to develop a procedure for learning such a model without explicitly evaluating the terms in (2.6).

REMARK 2.1. *As written above, (2.7a) is the most general form of a quadratic ODE with affine-parametric structure as in (2.3b). However, as we will see in examples, the constant, linear, and/or quadratic terms are not always present in practice, and when they are, the number of terms in the affine expansions (i.e., q_c , q_A , and q_H) tends to be small. In other words, projection-based ROMs of the form (2.7a) have the same number of nonzero terms as the original PDE (2.1a).*

EXAMPLE 2.2 (Heat Equation). *Consider the one-dimensional spatial domain $\Omega = (0, 1) \subset \mathbb{R}$ and the parameter domain $\mathcal{P} = [.01, 2.5] \times [0.1, 2.5] \subset \mathbb{R}^2$. For a fixed $\bar{x} \in \Omega$, let $\chi_{[0, \bar{x}]}$ and $\chi_{[\bar{x}, 1]}$ be indicator functions over Ω ,*

$$\chi_{[0, \bar{x}]}(x) = \begin{cases} 1, & 0 < x < \bar{x}, \\ 0, & \bar{x} \leq x < 1, \end{cases} \quad \chi_{[\bar{x}, 1]}(x) = \begin{cases} 0, & 0 < x < \bar{x}, \\ 1, & \bar{x} \leq x < 1. \end{cases}$$

The following equation models the diffusion of heat through a one-dimensional rod composed of two materials with independent thermal diffusivities $\mu = (\alpha, \beta) \in \mathcal{P}$, with the temperature prescribed at each end of the rod:

$$(2.8a) \quad \frac{\partial u}{\partial t} = (\alpha \chi_{[0, \bar{x}]} + \beta \chi_{[\bar{x}, 1]}) \frac{\partial^2 u}{\partial x^2} \quad x \in \Omega, t \in (t_0, t_f], \mu \in \mathcal{P},$$

$$(2.8b) \quad u(x, t_0; \mu) = u_0(x; \mu), \quad x \in \Omega, \mu \in \mathcal{P},$$

$$(2.8c) \quad u(0, t; \mu) = u(1, t; \mu) = 0, \quad t \in [t_0, t_f], \mu \in \mathcal{P}.$$

The underlying Hilbert space containing the state u is $\mathcal{V} = H^2(\Omega) \cap H_0^1(\Omega)$, the set of twice weakly differentiable functions satisfying the homogeneous Dirichlet boundary conditions (2.8c). In the language of (2.3a)–(2.3b), $\mathcal{C} \equiv \mathcal{H} \equiv 0$ and

$$\mathcal{A}(u; \mu) = \theta_A^{(1)}(\mu) \mathcal{A}^{(1)}(u) + \theta_A^{(2)}(\mu) \mathcal{A}^{(2)}(u)$$

with $\theta_A^{(1)}(\mu) = \alpha$, $\theta_A^{(2)}(\mu) = \beta$, $\mathcal{A}^{(1)}(u) = \chi_{[0,\bar{x}]} \frac{\partial^2 u}{\partial x^2}$, and $\mathcal{A}^{(2)}(u) = \chi_{[\bar{x},1]} \frac{\partial^2 u}{\partial x^2}$. A projection-based ROM for this problem as in (2.7a) has the form

$$(2.9) \quad \frac{d}{dt} \widehat{\mathbf{u}}(t; \mu) = \left(\alpha \widehat{\mathbf{A}}^{(1)} + \beta \widehat{\mathbf{A}}^{(2)} \right) \widehat{\mathbf{u}}(t; \mu),$$

with the initial condition as in (2.7b).

2.2. Affine Operator Inference for PDEs. Constructing $\widehat{\mathbf{c}}^{(p)}$, $\widehat{\mathbf{A}}^{(p)}$, and $\widehat{\mathbf{H}}^{(p)}$ in the ODE system (2.7a) by evaluating the terms of (2.6) is an inherently intrusive process, requiring explicit access to the differential operators $\mathcal{C}^{(p)}$, $\mathcal{A}^{(p)}$, and $\mathcal{H}^{(p)}$. In practice, this is not always favorable. For example, when the training data comes from commercial software, access to the details of the discretization might be limited. In addition, variable transformations used to induce the polynomial structure (2.3a) make these differential operators computationally inaccessible [26, 32, 39]. In this section, we pose an affine-parametric inference procedure to *learn* a ROM through a data-driven optimization problem given only samples of the solution u , a finite basis $\{v_j\}_{j=1}^r$ (which can be computed from the solution data, as we will discuss in Section 4), and knowledge of the governing equations (2.1a)–(2.1c) and the parametric structure (2.3a)–(2.3b). We do not rely on evaluations of the differential operators as in (2.6), but rather on the induced ODE structure (2.7a)–(2.7b). When restricted to linear-quadratic PDEs of a single variable, the formulation in this section is equivalent to the approach in [58]; Section 3 generalizes our method to systems of PDEs. Although we focus on the quadratic form motivated by (2.3a), the theory presented here applies to affine-parametric polynomial systems of arbitrary order.

Consider (2.7a) with fixed, known integers $q_c, q_A, q_H \geq 0$ and affine coefficient functions $\theta = \{\theta_c^{(1)}, \dots, \theta_H^{(q_H)}\}$, which define the polynomial and affine-parametric structure of the system. Define

$$(2.10a) \quad \begin{aligned} \frac{d}{dt} \widehat{\mathbf{u}}(t; \mu) = \mathbf{F}(\widehat{\mathbf{O}}; \widehat{\mathbf{u}}, t, \theta, \mu) = & \left(\sum_{p=1}^{q_c} \theta_c^{(p)}(\mu) \widehat{\mathbf{c}}^{(p)} \right) + \left(\sum_{p=1}^{q_A} \theta_A^{(p)}(\mu) \widehat{\mathbf{A}}^{(p)} \right) \widehat{\mathbf{u}}(t; \mu) \\ & + \left(\sum_{p=1}^{q_H} \theta_H^{(p)}(\mu) \widehat{\mathbf{H}}^{(p)} \right) (\widehat{\mathbf{u}}(t; \mu) \widehat{\odot} \widehat{\mathbf{u}}(t; \mu)), \end{aligned}$$

where the as yet unknown *operator matrix* $\widehat{\mathbf{O}}$ is the concatenation

$$(2.10b) \quad \widehat{\mathbf{O}} = \left[\widehat{\mathbf{c}}^{(1)} \ \dots \ \widehat{\mathbf{c}}^{(q_c)} \mid \widehat{\mathbf{A}}^{(1)} \ \dots \ \widehat{\mathbf{A}}^{(q_A)} \mid \widehat{\mathbf{H}}^{(1)} \ \dots \ \widehat{\mathbf{H}}^{(q_H)} \right] \in \mathbb{R}^{r \times q(r)},$$

with column dimension $q(r) = q_c + q_A r + q_H \binom{r+1}{2}$. Equipped with the initial condition (2.7b), (2.10a)–(2.10b) describe the family of ROMs with the same polynomial and parametric form as (2.7a), with particular realizations determined by the operator matrix $\widehat{\mathbf{O}}$. The goal is to choose $\widehat{\mathbf{O}}$ such that the corresponding ROM accurately captures the dynamics of the governing PDE (2.1a)–(2.1d) for all $\mu \in \mathcal{P}$.

We learn $\widehat{\mathbf{O}}$ by solving a data-driven least-squares regression problem. Suppose we can sample the solution of (2.1a)–(2.1d) at s parameter values $\{\mu_i\}_{i=1}^s \subset \mathcal{P}$ and K times $\{t_j\}_{j=1}^K \subset [t_0, t_f]$. We define the loss function $\mathcal{L} : \mathbb{R}^{r \times q(r)} \rightarrow \mathbb{R}$ associated with these solution samples to be the sum of the residuals of (2.10a),

$$(2.11) \quad \mathcal{L}(\widehat{\mathbf{O}}) = \sum_{i=1}^s \sum_{j=1}^K \left\| \mathbf{F}(\widehat{\mathbf{O}}; \widehat{\mathbf{u}}, t_j, \theta, \mu_i) - \frac{d}{dt} \widehat{\mathbf{u}}(t; \mu_i) \Big|_{t=t_j} \right\|_2^2,$$

where $\widehat{\mathbf{u}}(t; \mu) = [\langle v_1, u(\cdot, t; \mu) \rangle \cdots \langle v_r, u(\cdot, t; \mu) \rangle]^\top \in \mathbb{R}^r$ as before, and with \mathbf{F} and $\widehat{\mathbf{O}}$ as in (2.10a)–(2.10b). To write the minimization of (2.11) in a standard form, define the matrices

$$\begin{aligned} \widehat{\mathbf{U}}(\mu_i) &= \begin{bmatrix} \widehat{\mathbf{u}}(t_1; \mu_i) & \cdots & \widehat{\mathbf{u}}(t_K; \mu_i) \end{bmatrix} \in \mathbb{R}^{r \times K}, \\ \widehat{\dot{\mathbf{U}}}(\mu_i) &= \begin{bmatrix} \left. \frac{d}{dt} \widehat{\mathbf{u}}(t; \mu_i) \right|_{t=t_1} & \cdots & \left. \frac{d}{dt} \widehat{\mathbf{u}}(t; \mu_i) \right|_{t=t_K} \end{bmatrix} \in \mathbb{R}^{r \times K}, \end{aligned}$$

which group the projected solution and the associated time derivatives for each parameter sample, and define the row vectors

$$\begin{aligned} \boldsymbol{\theta}_c(\mu_i) &= \begin{bmatrix} \theta_c^{(1)}(\mu_i) & \cdots & \theta_c^{(q_c)}(\mu_i) \end{bmatrix} \in \mathbb{R}^{1 \times q_c}, \\ \boldsymbol{\theta}_A(\mu_i) &= \begin{bmatrix} \theta_A^{(1)}(\mu_i) & \cdots & \theta_A^{(q_A)}(\mu_i) \end{bmatrix} \in \mathbb{R}^{1 \times q_A}, \\ \boldsymbol{\theta}_H(\mu_i) &= \begin{bmatrix} \theta_H^{(1)}(\mu_i) & \cdots & \theta_H^{(q_H)}(\mu_i) \end{bmatrix} \in \mathbb{R}^{1 \times q_H}, \end{aligned}$$

which encode the affine-parametric dependencies of \mathcal{C} , \mathcal{A} , and \mathcal{H} , respectively. Finally, define the *data matrix*

$$(2.12) \quad \mathbf{D} = \left[\begin{array}{c|c|c} \mathbf{D}_c & \mathbf{D}_A & \mathbf{D}_H \\ \hline \boldsymbol{\theta}_c(\mu_1) \otimes \mathbf{1}_K & \boldsymbol{\theta}_A(\mu_1) \otimes \widehat{\mathbf{U}}(\mu_1)^\top & \boldsymbol{\theta}_H(\mu_1) \otimes \left(\widehat{\mathbf{U}}(\mu_1) \widehat{\odot} \widehat{\dot{\mathbf{U}}}(\mu_1) \right)^\top \\ \vdots & \vdots & \vdots \\ \boldsymbol{\theta}_c(\mu_s) \otimes \mathbf{1}_K & \boldsymbol{\theta}_A(\mu_s) \otimes \widehat{\mathbf{U}}(\mu_s)^\top & \boldsymbol{\theta}_H(\mu_s) \otimes \left(\widehat{\mathbf{U}}(\mu_s) \widehat{\odot} \widehat{\dot{\mathbf{U}}}(\mu_s) \right)^\top \end{array} \right],$$

where $\mathbf{1}_K \in \mathbb{R}^K$ is a column vector of unity of length K . We then have

$$(2.13) \quad \text{pOpInf:} \quad \min_{\widehat{\mathbf{O}}} \mathcal{L}(\widehat{\mathbf{O}}) = \min_{\widehat{\mathbf{O}}} \left\| \mathbf{D} \widehat{\mathbf{O}}^\top - \mathbf{R}^\top \right\|_F^2,$$

where $\mathbf{D} \in \mathbb{R}^{sK \times q(r)}$, $\widehat{\mathbf{O}} \in \mathbb{R}^{r \times q(r)}$, and $\mathbf{R} = [\widehat{\dot{\mathbf{U}}}(\mu_1) \cdots \widehat{\dot{\mathbf{U}}}(\mu_s)] \in \mathbb{R}^{r \times sK}$. We call (2.13) the *affine-parametric Operator Inference problem* (pOpInf). Note that (2.13) is a linear least-squares problem which decouples by the columns of $\widehat{\mathbf{O}}^\top$, meaning the dynamics for each \hat{u}_i are learned independently [37].

If the data matrix \mathbf{D} has full column rank, then (2.13) has a unique closed-form solution [11, 19, 51]. However, \mathbf{D} is susceptible to rank deficiencies due to its Kronecker block structure. Theorem 2.3 establishes necessary conditions for the well-posedness of the pOpInf problem (2.13) and sufficient conditions in the monomial case (e.g., $\mathcal{C} = \mathcal{A} = 0$ but $\mathcal{H} \neq 0$ so that $\mathbf{D} = \mathbf{D}_H$).

THEOREM 2.3. *Let $\Theta_c \in \mathbb{R}^{s \times q_c}$, $\Theta_A \in \mathbb{R}^{s \times q_A}$, and $\Theta_H \in \mathbb{R}^{s \times q_H}$ be the matrices with entries*

$$(2.14) \quad [\Theta_c]_{ij} = \theta_c^{(j)}(\mu_i), \quad [\Theta_A]_{ij} = \theta_A^{(j)}(\mu_i), \quad [\Theta_H]_{ij} = \theta_H^{(j)}(\mu_i),$$

that is, the i th row of Θ_c is $\boldsymbol{\theta}_c(\mu_i)$, the i th row of Θ_A is $\boldsymbol{\theta}_A(\mu_i)$, and the i th row of Θ_H is $\boldsymbol{\theta}_H(\mu_i)$. If any of Θ_c , Θ_A , or Θ_H do not have full column rank, then the data matrix $\mathbf{D} = [\mathbf{D}_c \ \mathbf{D}_A \ \mathbf{D}_H]$ defined in (2.12) is rank deficient. Furthermore, if either

of the block matrices

$$\widehat{\mathbf{U}}_A = \begin{bmatrix} \widehat{\mathbf{U}}(\mu_1)^\top \\ \vdots \\ \widehat{\mathbf{U}}(\mu_s)^\top \end{bmatrix}, \quad \widehat{\mathbf{U}}_H = \begin{bmatrix} \left(\widehat{\mathbf{U}}(\mu_1) \widehat{\odot} \widehat{\mathbf{U}}(\mu_1) \right)^\top \\ \vdots \\ \left(\widehat{\mathbf{U}}(\mu_s) \widehat{\odot} \widehat{\mathbf{U}}(\mu_s) \right)^\top \end{bmatrix}$$

do not have full column rank, then neither does \mathbf{D} . On the other hand, if Θ_A and each $\widehat{\mathbf{U}}(\mu_i)^\top$ ($i = 1, \dots, s$) have full column rank, then so does \mathbf{D}_A ; and if Θ_H and each $\left(\widehat{\mathbf{U}}(\mu_i) \widehat{\odot} \widehat{\mathbf{U}}(\mu_i) \right)^\top$ have full column rank, then so does \mathbf{D}_H .

Proof. Note that $\Theta_c = \mathbf{D}_c$, the leftmost block of \mathbf{D} . Therefore, if Θ_c does not have full column rank, then \mathbf{D} has linearly dependent columns. [Lemma B.1](#) gives the result for the remaining cases: if Θ_A or $\widehat{\mathbf{U}}_A$ do not have full column rank, apply [Lemma B.1](#) with $\mathbf{y}_i = \theta_A(\mu_i)^\top$, $\mathbf{Z}_i = \widehat{\mathbf{U}}(\mu_i)^\top$, and $\mathbf{W} = \mathbf{D}_A$ to show \mathbf{D}_A is rank deficient; a similar argument holds for Θ_H and $\widehat{\mathbf{U}}_H$, showing \mathbf{D}_H is rank deficient. The results to the converse also hold by [Lemma B.1](#). \square

Selecting appropriate parameter samples is an important issue for all parametric model reduction methods [8]. [Theorem 2.3](#) provides numerically relevant guidance for parameter selection in the context of (2.13): for a set of parameter values $\{\mu_i\}_{i=1}^s$, the (small) matrices Θ_c , Θ_A , and Θ_H can be explicitly formed and checked for rank deficiencies without any information about the solution trajectories. If any of the matrices are unsatisfactorily conditioned, different parameter values must be chosen at which to sample the solution. This is especially relevant for settings in which only a few solution trajectories can be afforded, i.e., s is relatively small. Note that if $s < \max\{q_c, q_A, q_H\}$, then the problem is guaranteed to be ill-posed. [Section 4](#) further addresses numerical issues that may arise from other sources of poor conditioning.

EXAMPLE 2.4 (Heat Equation). Recall the problem introduced in [Example 2.2](#). Given parameter samples $\{\mu_i = (\alpha_i, \beta_i)\}_{i=1}^s$, we learn a ROM of the form (2.9) by solving the pOpInf problem (2.13) with operator and data matrices

$$\widehat{\mathbf{O}} = \begin{bmatrix} \widehat{\mathbf{A}}^{(1)} & \widehat{\mathbf{A}}^{(2)} \end{bmatrix} \in \mathbb{R}^{r \times 2r}, \quad \mathbf{D} = \begin{bmatrix} \alpha_1 \widehat{\mathbf{U}}(\mu_1)^\top & \beta_1 \widehat{\mathbf{U}}(\mu_1)^\top \\ \vdots & \vdots \\ \alpha_s \widehat{\mathbf{U}}(\mu_s)^\top & \beta_s \widehat{\mathbf{U}}(\mu_s)^\top \end{bmatrix} \in \mathbb{R}^{sK \times 2r}.$$

Per [Theorem 2.3](#), the rank of the data matrix \mathbf{D} depends on each $\widehat{\mathbf{U}}(\mu_i)$ and the matrix

$$\Theta_A = \begin{bmatrix} \theta_A^{(1)}(\mu_1) & \theta_A^{(2)}(\mu_1) \\ \vdots & \vdots \\ \theta_A^{(1)}(\mu_s) & \theta_A^{(2)}(\mu_s) \end{bmatrix} = \begin{bmatrix} \alpha_1 & \beta_1 \\ \vdots & \vdots \\ \alpha_s & \beta_s \end{bmatrix} \in \mathbb{R}^{s \times 2}.$$

In particular, if Θ_A does not have full column rank, then the data matrix \mathbf{D} will be rank deficient. Our goal, then, is to choose $s \geq 2$ parameter samples such that Θ_A has a small condition number.

Reduced-order models learned through pOpInf and those constructed via intrusive projection, i.e., by explicitly evaluating (2.6), are related in the following sense.

THEOREM 2.5. Let u be the unique solution of (2.1a)–(2.1d), where the differential operator \mathcal{F} has the quadratic, affine-parametric form described in (2.3a)–(2.3b).

Suppose there exists an orthonormal set $\{v_j\}_{j=1}^r \subset \mathcal{V}$, and let $\{\mu_i\}_{i=1}^s \subset \mathcal{P}$ be a finite set of parameter samples. Define the loss function

$$(2.15) \quad \mathcal{L}(\widehat{\mathbf{O}}) = \sum_{i=1}^s \int_{t_0}^{t_f} \left\| \mathbf{F} \left(\widehat{\mathbf{O}}; \widehat{\mathbf{u}}, t, \theta, \mu_i \right) - \frac{d}{dt} \widehat{\mathbf{u}}(t; \mu_i) \right\|_2^2 dt,$$

where $\widehat{\mathbf{u}}(t; \mu) = [\langle v_1, u(\cdot, t; \mu) \rangle \cdots \langle v_r, u(\cdot, t; \mu) \rangle]^\top \in \mathbb{R}^r$, and with \mathbf{F} and $\widehat{\mathbf{O}}$ given by (2.10a)–(2.10b). Consider the following conditions.

1. There exist functions $\hat{u}_j : [t_0, t_f] \times \mathcal{P} \rightarrow \mathbb{R}$, $j = 1, \dots, r$, such that the finite sum representation (2.5) is exact for all $x \in \Omega$, $t \in [t_0, t_f]$, and $\mu \in \{\mu_i\}_{i=1}^s$, i.e., there is no truncation error at the sample parameter values.
2. For $i = 1, \dots, s$, there exist times $\{\tau_{i,j}\}_{j=1}^K \subset [t_0, t_f]$, $K = 1 + r + \binom{r+1}{2}$, such that the matrix

$$\widetilde{\mathbf{D}}(\mu_i) = \begin{bmatrix} \mathbf{1}_K & \widetilde{\mathbf{U}}(\mu_i)^\top & \left(\widetilde{\mathbf{U}}(\mu_i) \widehat{\mathbf{O}} \widetilde{\mathbf{U}}(\mu_i) \right)^\top \end{bmatrix} \in \mathbb{R}^{K \times K}$$

is invertible, where $\widetilde{\mathbf{U}}(\mu_i) = [\widehat{\mathbf{u}}(\tau_{i,1}; \mu_i) \cdots \widehat{\mathbf{u}}(\tau_{i,K}; \mu_i)] \in \mathbb{R}^{r \times K}$.

3. The matrices $\Theta_c \in \mathbb{R}^{s \times q_c}$, $\Theta_A \in \mathbb{R}^{s \times q_A}$, and $\Theta_H \in \mathbb{R}^{s \times q_H}$ of (2.14) have full column rank.

If condition 1 holds, then the loss function \mathcal{L} has a global minimizer $\widehat{\mathbf{O}}$ satisfying $\mathcal{L}(\widehat{\mathbf{O}}) = 0$. If conditions 2 and 3 also hold, then that minimizer is unique.

Proof. Assume condition 1 (no truncation error at the parameter samples). Then the system of ODEs (2.7a) with operators $\widehat{\mathbf{c}}^{(p)}$, $\widehat{\mathbf{A}}^{(p)}$, and $\widehat{\mathbf{H}}^{(p)}$ derived from (2.6) holds exactly for all $x \in \Omega$, $t \in [t_0, t_f]$, and $\mu \in \{\mu_i\}_{i=1}^s$, i.e., it is equivalent to (2.1a)–(2.1d) at the parameter samples. Constructing $\widehat{\mathbf{O}}$ from these operators, we have $\frac{d}{dt} \widehat{\mathbf{u}}(t; \mu) = \mathbf{F}(\widehat{\mathbf{O}}; \widehat{\mathbf{u}}, t, \theta, \mu)$ exactly for each $\mu \in \{\mu_i\}_{i=1}^s$. By construction, $\mathcal{L}(\widehat{\mathbf{O}}) = 0$, which—since \mathcal{L} is non-negative—shows that $\widehat{\mathbf{O}}$ is a global minimizer of \mathcal{L} .

To prove uniqueness, assume conditions 2 and 3 and suppose $\widehat{\mathbf{O}}^{(1)}$ and $\widehat{\mathbf{O}}^{(2)}$ both minimize \mathcal{L} . By the previous argument, $\mathcal{L}(\widehat{\mathbf{O}}^{(1)}) = \mathcal{L}(\widehat{\mathbf{O}}^{(2)}) = 0$, which implies

$$(2.16) \quad \mathbf{0}_r = \mathbf{F}(\widehat{\mathbf{O}}^{(1)}; \widehat{\mathbf{u}}, \tau_{i,j}, \theta, \mu_i) - \mathbf{F}(\widehat{\mathbf{O}}^{(2)}; \widehat{\mathbf{u}}, \tau_{i,j}, \theta, \mu_i) = \mathbf{F}(\widetilde{\mathbf{O}}; \widehat{\mathbf{u}}, \tau_{i,j}, \theta, \mu_i)$$

for $i = 1, \dots, s$ and $j = 1, \dots, K$, where $\mathbf{0}_r \in \mathbb{R}^r$ is a column vector of zeros and

$$\widetilde{\mathbf{O}} = \widehat{\mathbf{O}}^{(1)} - \widehat{\mathbf{O}}^{(2)} = \begin{bmatrix} \widetilde{\mathbf{c}}^{(1)} & \dots & \widetilde{\mathbf{c}}^{(q_c)} & | & \widetilde{\mathbf{A}}^{(1)} & \dots & \widetilde{\mathbf{A}}^{(q_A)} & | & \widetilde{\mathbf{H}}^{(1)} & \dots & \widetilde{\mathbf{H}}^{(q_H)} \end{bmatrix} \in \mathbb{R}^{r \times q(r)}.$$

For fixed i , (2.16) with $j = 1, \dots, K$ can be written as the linear system

$$\widetilde{\mathbf{D}}(\mu_i) \begin{bmatrix} \sum_{p=1}^{q_c} \theta_c^{(p)}(\mu_i) \widetilde{\mathbf{c}}^{(p)} & | & \sum_{p=1}^{q_A} \theta_A^{(p)}(\mu_i) \widetilde{\mathbf{A}}^{(p)} & | & \sum_{p=1}^{q_H} \theta_H^{(p)}(\mu_i) \widetilde{\mathbf{H}}^{(p)} \end{bmatrix}^\top = \mathbf{0}_{K \times r},$$

where $\mathbf{0}_{K \times r} \in \mathbb{R}^{K \times r}$ is the zero matrix. As each $\widetilde{\mathbf{D}}(\mu_i)$ is invertible (condition 2), for $i = 1, \dots, s$ we have

$$\sum_{p=1}^{q_c} \theta_c^{(p)}(\mu_i) \widetilde{\mathbf{c}}^{(p)} = \mathbf{0}_r, \quad \sum_{p=1}^{q_A} \theta_A^{(p)}(\mu_i) \widetilde{\mathbf{A}}^{(p)} = \mathbf{0}_{r \times r}, \quad \sum_{p=1}^{q_H} \theta_H^{(p)}(\mu_i) \widetilde{\mathbf{H}}^{(p)} = \mathbf{0}_{r \times \binom{r+1}{2}}.$$

Collecting these equations component-wise for $i = 1, \dots, s$ yields the equations

$$\Theta_c \begin{bmatrix} [\widetilde{\mathbf{c}}^{(1)}]_j \\ \vdots \\ [\widetilde{\mathbf{c}}^{(q_c)}]_j \end{bmatrix} = \Theta_A \begin{bmatrix} [\widetilde{\mathbf{A}}^{(1)}]_{jk} \\ \vdots \\ [\widetilde{\mathbf{A}}^{(q_A)}]_{jk} \end{bmatrix} = \Theta_H \begin{bmatrix} [\widetilde{\mathbf{H}}^{(1)}]_{jk} \\ \vdots \\ [\widetilde{\mathbf{H}}^{(q_H)}]_{jk} \end{bmatrix} = \mathbf{0}_s$$

for every index pair j, k . But Θ_c , Θ_{A_j} , and Θ_H each have full column rank (condition 3), implying $[\tilde{\mathbf{c}}^{(p)}]_j = [\tilde{\mathbf{A}}^{(p)}]_{jk} = [\tilde{\mathbf{H}}^{(p)}]_{jk} = 0$ for all j, k, p . Thus, $\tilde{\mathbf{O}}$ is the zero matrix, hence $\hat{\mathbf{O}}^{(1)} = \hat{\mathbf{O}}^{(2)}$. \square

Theorem 2.5 easily extends to systems of arbitrary polynomial order. Hence, for finite-dimensional PDEs (condition 1) with a polynomial, affine-parametric structure, the operators defined via (2.6) uniquely minimize the loss function \mathcal{L} defined in (2.15), provided that the dynamics are sufficiently diverse (as quantified by conditions 2 and 3). As the time step between the data samples defining the pOpInf loss \mathcal{L} of (2.11) decreases, i.e., $t_{j+1} - t_j \rightarrow 0$ for all j , then \mathcal{L} converges to \mathcal{L} (up to a constant), in which case the pOpInf ROM and the intrusive projection-based ROM will be the same. For infinite-dimensional PDEs, condition 1 is realized as the number of basis functions in the approximation increases, i.e., $r \rightarrow \infty$. Thus, there are two factors explaining why ROMs learned through pOpInf are not guaranteed to agree with (2.6): 1) the truncation error of the representation (2.5), and 2) the finite sampling of the solution in time. See [36, 52] for additional work connecting the intrusive and data-driven ROMs.

3. Affine Operator Inference for Systems of PDEs. Systems of partial differential equations with affine polynomial structure admit low-dimensional representations similar to (2.7a)–(2.7b) [18]. Consider the system of d partial differential equations where each equation can be written as in (2.1a)–(2.1d):

$$(3.1a) \quad \frac{\partial u_\ell}{\partial t} = \mathcal{F}_\ell(u_1, \dots, u_d; \mu), \quad \ell = 1, \dots, d,$$

where each state variable u_ℓ is contained in a Hilbert space \mathcal{V}_ℓ of real-valued functions satisfying appropriate homogeneous boundary conditions. Suppose each spatial differential operator $\mathcal{F}_\ell : \mathcal{V}_1 \times \dots \times \mathcal{V}_d \times \mathcal{P} \rightarrow \mathcal{V}_\ell^*$ is polynomial in the variables u_1, \dots, u_d and their spatial derivatives,

$$(3.1b) \quad \mathcal{F}_\ell(u_1, \dots, u_d; \mu) = \mathcal{C}_\ell(\mu) + \sum_{m=1}^d \mathcal{A}_{\ell,m}(u_m; \mu) + \sum_{m=1}^d \sum_{n=m}^d \mathcal{H}_{\ell,mn}(u_m, u_n; \mu),$$

where $\mathcal{C}_\ell : \mathcal{P} \rightarrow \mathcal{V}_\ell^*$, $\mathcal{A}_{\ell,m} : \mathcal{V}_m \times \mathcal{P} \rightarrow \mathcal{V}_\ell^*$, $\mathcal{H}_{\ell,mn} : \mathcal{V}_m \times \mathcal{V}_n \times \mathcal{P} \rightarrow \mathcal{V}_\ell^*$, and each $\mathcal{A}_{\ell,m}$ and $\mathcal{H}_{\ell,mn}$ is linear in each state argument. This is the natural multivariate extension of (2.3a). We further assume that each operator \mathcal{C}_ℓ , $\mathcal{A}_{\ell,m}$, and $\mathcal{H}_{\ell,mn}$ has an affine-parametric expansion, that is,

$$(3.1c) \quad \begin{aligned} \mathcal{C}_\ell(\mu) &= \sum_{p=1}^{q_{\mathcal{C}_\ell}} \theta_{\mathcal{C}_\ell}^{(p)}(\mu) \mathcal{C}_\ell^{(p)}. & \mathcal{A}_{\ell,m}(u; \mu) &= \sum_{p=1}^{q_{\mathcal{A}_{\ell,m}}} \theta_{\mathcal{A}_{\ell,m}}^{(p)}(\mu) \mathcal{A}_{\ell,m}^{(p)}(u). \\ \mathcal{H}_{\ell,mn}(u, v; \mu) &= \sum_{p=1}^{q_{\mathcal{H}_{\ell,mn}}} \theta_{\mathcal{H}_{\ell,mn}}^{(p)}(\mu) \mathcal{H}_{\ell,mn}^{(p)}(u, v). \end{aligned}$$

For each $\ell = 1, \dots, d$, let $\{v_{\ell j}\}_{j=1}^{r_\ell} \subset \mathcal{V}_\ell$ be an orthonormal set. We assume a reduced representation \check{u}_ℓ of u_ℓ , confined to $\text{span}(\{v_{\ell 1}, \dots, v_{\ell r_\ell}\}) \subset \mathcal{V}_\ell$, given by

$$(3.2) \quad \check{u}_\ell(x, t; \mu) = \sum_{j=1}^{r_\ell} \hat{u}_{\ell j}(t; \mu) v_{\ell j}(x), \quad \hat{u}_{\ell j}(t; \mu) = \langle v_{\ell j}, u_\ell(\cdot, t; \mu) \rangle_{\mathcal{V}_\ell},$$

where $\langle \cdot, \cdot \rangle_{\mathcal{V}_\ell}$ is the duality pairing of \mathcal{V}_ℓ with its dual. Inserting \check{u}_ℓ into (3.1a) for each ℓ and proceeding as before, we obtain d ODEs that serve as a ROM for (3.1a)–(3.1c):

$$\begin{aligned}
 \frac{d}{dt} \hat{\mathbf{u}}_\ell(t; \mu) &= \mathbf{F}_\ell(\hat{\mathbf{O}}_\ell; \hat{\mathbf{u}}_1, \dots, \hat{\mathbf{u}}_d, t, \theta, \mu) \\
 &= \left(\sum_{p=1}^{q_{c_\ell}} \theta_{c_\ell}^{(p)}(\mu) \hat{\mathbf{c}}_\ell^{(p)} \right) + \sum_{m=1}^d \left(\sum_{p=1}^{q_{A_{\ell,m}}} \theta_{A_{\ell,m}}^{(p)}(\mu) \hat{\mathbf{A}}_{\ell,m}^{(p)} \right) \hat{\mathbf{u}}_m(t; \mu) \\
 (3.3a) \quad &+ \sum_{m=1}^d \left(\sum_{p=1}^{q_{H_{\ell,mm}}} \theta_{H_{\ell,mm}}^{(p)}(\mu) \hat{\mathbf{H}}_{\ell,mm}^{(p)} \right) (\hat{\mathbf{u}}_m(t; \mu) \hat{\odot} \hat{\mathbf{u}}_m(t; \mu)) \\
 &+ \sum_{m=1}^d \sum_{n=m+1}^d \left(\sum_{p=1}^{q_{H_{\ell,mn}}} \theta_{H_{\ell,mn}}^{(p)}(\mu) \hat{\mathbf{H}}_{\ell,mn}^{(p)} \right) (\hat{\mathbf{u}}_m(t; \mu) \odot \hat{\mathbf{u}}_n(t; \mu)),
 \end{aligned}$$

$$(3.3b) \quad \hat{\mathbf{u}}_\ell(t; \mu) = \left[\langle v_{\ell,1}, u_\ell(\cdot, t; \mu) \rangle \quad \cdots \quad \langle v_{\ell,r_\ell}, u_\ell(\cdot, t; \mu) \rangle \right]^\top \in \mathbb{R}^{r_\ell},$$

$$(3.3c) \quad \hat{\mathbf{O}}_\ell = \left[\hat{\mathbf{c}}_\ell^{(1)} \quad \cdots \quad \hat{\mathbf{c}}_\ell^{(q_{c_\ell})} \mid \hat{\mathbf{A}}_{\ell,1}^{(1)} \quad \cdots \quad \hat{\mathbf{A}}_{\ell,d}^{(q_{A_{\ell,d}})} \mid \hat{\mathbf{H}}_{\ell,11}^{(1)} \quad \cdots \quad \hat{\mathbf{H}}_{\ell,dd}^{(q_{H_{\ell,dd}})} \right],$$

where $\hat{\mathbf{c}}_\ell^{(p)} \in \mathbb{R}^{r_\ell}$, $\hat{\mathbf{A}}_{\ell,m}^{(p)} \in \mathbb{R}^{r_\ell \times r_m}$, $\hat{\mathbf{H}}_{\ell,mm}^{(p)} \in \mathbb{R}^{r_\ell \times \binom{r_m+1}{2}}$, and $\hat{\mathbf{H}}_{\ell,mn}^{(p)} \in \mathbb{R}^{r_\ell \times r_m r_n}$ ($m \neq n$). Here \odot denotes the Khatri-Rao product, whereas $\hat{\odot}$ is the compact Khatri-Rao product that omits redundant terms (see Appendix A). Higher-order operators (e.g., cubic terms) can be similarly accounted for.

The low-dimensional system (3.3a) retains the differential and parametric structure of the PDE system (3.1a)–(3.1b). The corresponding affine pOpInf problem decouples into d instances of (2.13) which can be solved *independently*:

$$(3.4) \quad \min_{\hat{\mathbf{O}}_\ell} \left\| \mathbf{D}_\ell \hat{\mathbf{O}}_\ell^\top - \mathbf{R}_\ell^\top \right\|_F^2, \quad \ell = 1, \dots, d,$$

where

$$\begin{aligned}
 \mathbf{R}_\ell &= \left[\hat{\mathbf{U}}_\ell(\mu_1) \quad \cdots \quad \hat{\mathbf{U}}_\ell(\mu_s) \right] \in \mathbb{R}^{r_\ell \times sK}, \\
 \hat{\mathbf{U}}_\ell(\mu_i) &= \left[\frac{d}{dt} \hat{\mathbf{u}}_\ell(t; \mu_i) \Big|_{t=t_1} \quad \cdots \quad \frac{d}{dt} \hat{\mathbf{u}}_\ell(t; \mu_i) \Big|_{t=t_K} \right] \in \mathbb{R}^{r_\ell \times K},
 \end{aligned}$$

and where each \mathbf{D}_ℓ is constructed from the matrices $\hat{\mathbf{U}}_1(\mu_1), \dots, \hat{\mathbf{U}}_d(\mu_s)$, where

$$\hat{\mathbf{U}}_\ell(\mu_i) = \left[\hat{\mathbf{u}}_\ell(t_1; \mu_i) \quad \cdots \quad \hat{\mathbf{u}}_\ell(t_K; \mu_i) \right] \in \mathbb{R}^{r_\ell \times K}.$$

See Appendix C for the general construction, and note that the number of terms in the PDE dictates the number of terms in the ROM and hence the size of the pOpInf problem. This is best illustrated by example.

EXAMPLE 3.1 (FitzHugh–Nagumo System). *Let $\Omega = (0, 1)$ and define the parameters $\mu = (\alpha, \beta, \gamma, \varepsilon) \in \mathbb{R}^4$. The following system of equations is a simplification of the Hodgkin–Huxley model for activation and deactivation in a spiking neuron [17, 33]:*

$$(3.5a) \quad \frac{\partial u_1}{\partial t} = \varepsilon \frac{\partial^2 u_1}{\partial x^2} + \frac{1}{\varepsilon} (-u_1^3 + 1.1u_1^2 - 0.1u_1 - u_2 + \alpha),$$

$$(3.5b) \quad \frac{\partial u_2}{\partial t} = \beta u_1 - \gamma u_2 + \alpha,$$

with initial conditions $u_1(x, t_0) = u_2(x, t_0) = 0$ and Neumann boundary conditions

$$(3.5c) \quad \left. \frac{\partial u_1}{\partial x} \right|_{x=0} = f(t) := -50000t^3 e^{-15t}, \quad \left. \frac{\partial u_1}{\partial x} \right|_{x=1} = 0.$$

We write (3.5a)–(3.5c) in the language of (3.1a)–(3.1c) as

$$\begin{aligned} \frac{\partial u_1}{\partial t} &= \theta_{c_1}^{(1)}(\mu) \mathcal{C}_1^{(1)} + \theta_{A_{1,1}}^{(1)}(\mu) \mathcal{A}_{1,1}^{(1)}(u_1) + \theta_{A_{1,1}}^{(2)}(\mu) \mathcal{A}_{1,1}^{(2)}(u_1) + \theta_{A_{1,2}}^{(1)}(\mu) \mathcal{A}_{1,2}^{(1)}(u_2) \\ &\quad + \theta_{H_{1,11}}^{(1)}(\mu) \mathcal{H}_{1,11}^{(1)}(u_1, u_1) + \theta_{G_{1,111}}^{(1)}(\mu) \mathcal{G}_{1,111}^{(1)}(u_1, u_1, u_1), \\ \frac{\partial u_2}{\partial t} &= \theta_{c_2}^{(1)}(\mu) \mathcal{C}_2^{(1)} + \theta_{A_{2,1}}^{(1)}(\mu) \mathcal{A}_{2,1}^{(1)}(u_1) + \theta_{A_{2,2}}^{(1)}(\mu) \mathcal{A}_{2,2}^{(1)}(u_2), \end{aligned}$$

with the operators and the affine coefficient functions given in Table 1. Here, $\mathcal{G}_{1,111}^{(1)}$ represents the cubic nonlinearity u_1^3 . This motivates a ROM of the form

$$\begin{aligned} \frac{d}{dt} \widehat{\mathbf{u}}_1(t; \mu) &= \theta_{c_1}^{(1)}(\mu) \widehat{\mathbf{c}}_1^{(1)} + \left(\theta_{A_{1,1}}^{(1)}(\mu) \widehat{\mathbf{A}}_{1,1}^{(1)} + \theta_{A_{1,1}}^{(2)}(\mu) \widehat{\mathbf{A}}_{1,1}^{(2)} \right) \widehat{\mathbf{u}}_1(t; \mu) \\ &\quad + \theta_{A_{1,2}}^{(1)}(\mu) \widehat{\mathbf{A}}_{1,2}^{(1)} \widehat{\mathbf{u}}_2(t; \mu) + \theta_{H_{1,11}}^{(1)}(\mu) \widehat{\mathbf{H}}_{1,11}^{(1)}(\widehat{\mathbf{u}}_1(t; \mu) \widehat{\odot} \widehat{\mathbf{u}}_1(t; \mu)) \\ &\quad + \theta_{G_{1,111}}^{(1)}(\mu) \widehat{\mathbf{G}}_{1,111}^{(1)}(\widehat{\mathbf{u}}_1(t; \mu) \widehat{\odot} \widehat{\mathbf{u}}_1(t; \mu) \widehat{\odot} \widehat{\mathbf{u}}_1(t; \mu)) + \theta_{B_1}^{(1)}(\mu) \widehat{\mathbf{B}}_1^{(1)} f(t), \\ \frac{d}{dt} \widehat{\mathbf{u}}_2(t; \mu) &= \theta_{c_2}^{(1)}(\mu) \widehat{\mathbf{c}}_2^{(1)} + \theta_{A_{2,1}}^{(1)}(\mu) \widehat{\mathbf{A}}_{2,1}^{(1)} \widehat{\mathbf{u}}_1(t; \mu) + \theta_{A_{2,2}}^{(1)}(\mu) \widehat{\mathbf{A}}_{2,2}^{(1)} \widehat{\mathbf{u}}_2(t; \mu), \end{aligned}$$

where the sizes of each discretized operators are listed in Table 1. The operator $\widehat{\mathbf{B}}_1^{(1)}$ accounts for the Neumann boundary condition on u_1 in (3.5c); it has the coefficient function $\theta_{B_1}^{(1)}(\mu) = \varepsilon$. The term $\widehat{\mathbf{G}}_{1,111}^{(1)}$ is the discretization of $\mathcal{G}_{1,111}^{(1)}$.

The corresponding pOpInf problem is (3.4) with $d = 2$ and

$$\begin{aligned} \widehat{\mathbf{O}}_1 &= \left[\widehat{\mathbf{c}}_1^{(1)} \quad \widehat{\mathbf{B}}_1^{(1)} \quad \widehat{\mathbf{A}}_{1,1}^{(1)} \quad \widehat{\mathbf{A}}_{1,1}^{(2)} \quad \widehat{\mathbf{A}}_{1,2}^{(1)} \quad \widehat{\mathbf{H}}_{1,11}^{(1)} \quad \widehat{\mathbf{G}}_{1,111}^{(1)} \right] \in \mathbb{R}^{r_1 \times q_1(r_1, r_2)}, \\ \mathbf{D}_1 &= \left[\mathbf{D}_{c_1} \mid \mathbf{D}_{A_1} \mid \mathbf{D}_{H_1} \mid \mathbf{D}_{G_1} \right] \in \mathbb{R}^{sK \times q_1(r_1, r_2)}, \\ \mathbf{D}_{c_1} &= \begin{bmatrix} \theta_{c_1}^{(1)}(\mu_1) \mathbf{1}_K & \theta_{B_1}^{(1)}(\mu_1) \mathbf{f} \\ \vdots & \vdots \\ \theta_{c_1}^{(1)}(\mu_s) \mathbf{1}_K & \theta_{B_1}^{(1)}(\mu_s) \mathbf{f} \end{bmatrix}, \\ \mathbf{D}_{A_1} &= \begin{bmatrix} \theta_{A_{1,1}}^{(1)}(\mu_1) \widehat{\mathbf{U}}_1(\mu_1)^\top & \theta_{A_{1,1}}^{(2)}(\mu_1) \widehat{\mathbf{U}}_1(\mu_1)^\top & \theta_{A_{1,2}}^{(1)}(\mu_1) \widehat{\mathbf{U}}_2(\mu_1)^\top \\ \vdots & \vdots & \vdots \\ \theta_{A_{1,1}}^{(1)}(\mu_s) \widehat{\mathbf{U}}_1(\mu_s)^\top & \theta_{A_{1,1}}^{(2)}(\mu_s) \widehat{\mathbf{U}}_1(\mu_s)^\top & \theta_{A_{1,2}}^{(1)}(\mu_s) \widehat{\mathbf{U}}_2(\mu_s)^\top \end{bmatrix}, \\ \mathbf{D}_{H_1} &= \begin{bmatrix} \theta_{H_{1,11}}^{(1)}(\mu_1) \left(\widehat{\mathbf{U}}_1(\mu_1) \widehat{\odot} \widehat{\mathbf{U}}_1(\mu_1) \right)^\top \\ \vdots \\ \theta_{H_{1,11}}^{(1)}(\mu_s) \left(\widehat{\mathbf{U}}_1(\mu_s) \widehat{\odot} \widehat{\mathbf{U}}_1(\mu_s) \right)^\top \end{bmatrix}, \end{aligned}$$

$$\begin{aligned}
 \mathbf{D}_{G_1} &= \begin{bmatrix} \theta_{G_{1,111}}^{(1)}(\mu_1) \left(\widehat{\mathbf{U}}_1(\mu_1) \widehat{\odot} \widehat{\mathbf{U}}_1(\mu_1) \widehat{\odot} \widehat{\mathbf{U}}_1(\mu_1) \right)^\top \\ \vdots \\ \theta_{G_{1,111}}^{(1)}(\mu_s) \left(\widehat{\mathbf{U}}_1(\mu_s) \widehat{\odot} \widehat{\mathbf{U}}_1(\mu_s) \widehat{\odot} \widehat{\mathbf{U}}_1(\mu_s) \right)^\top \end{bmatrix}, \\
 \widehat{\mathbf{O}}_2 &= \begin{bmatrix} \widehat{\mathbf{c}}_2^{(1)} & \widehat{\mathbf{A}}_{2,1}^{(1)} & \widehat{\mathbf{A}}_{2,2}^{(1)} \end{bmatrix} \in \mathbb{R}^{r_2 \times q_2(r_1, r_2)}, \\
 \mathbf{D}_2 &= \begin{bmatrix} \theta_{c_2}^{(1)}(\mu_1) \mathbf{1}_K & \theta_{A_{2,1}}^{(1)}(\mu_1) \widehat{\mathbf{U}}_1(\mu_1)^\top & \theta_{A_{2,2}}^{(1)}(\mu_1) \widehat{\mathbf{U}}_2(\mu_1)^\top \\ \vdots & \vdots & \vdots \\ \theta_{c_2}^{(1)}(\mu_s) \mathbf{1}_K & \theta_{A_{2,1}}^{(1)}(\mu_s) \widehat{\mathbf{U}}_1(\mu_s)^\top & \theta_{A_{2,2}}^{(1)}(\mu_s) \widehat{\mathbf{U}}_2(\mu_s)^\top \end{bmatrix} \in \mathbb{R}^{sK \times q_2(r_1, r_2)},
 \end{aligned}$$

where $\mathbf{f} = [f(t_1) \cdots f(t_K)]^\top \in \mathbb{R}^K$, $q_1(r_1, r_2) = 2 + 2r_1 + r_2 + \binom{r_1+1}{2} + \binom{r_1+2}{3}$, and $q_2(r_1, r_2) = 1 + r_1 + r_2$. The data matrix \mathbf{D}_1 and the operator matrix $\widehat{\mathbf{O}}_1$ have been modified from (3.3a)–(3.3c) to account for the cubic term present (3.5a).

REMARK 3.2. Even though since [14], the FitzHugh–Nagumo system has been widely used as a benchmark problem for the development of general nonlinear model reduction methods, a fully cubic intrusive model can be directly derived for this system (i.e., the nonlinear terms in the system are point-wise local and exactly quadratic and cubic), circumventing the need for the second layer of approximation introduced through hyper-reduction.

TABLE 1

The operators in the continuous setting, the associated affine coefficient functions, and the size of the reduced operators in discretized setting for the FitzHugh–Nagumo system of Example 3.1.

		continuous term	affine coefficient	discretized operator
Eq. (3.5a)	constant	$C_1^{(1)} = 1$	$\theta_{c_1}^{(1)}(\mu) = \alpha/\varepsilon$	$\widehat{\mathbf{c}}_1^{(1)} \in \mathbb{R}^{r_1}$
	input	$\frac{\partial u_1}{\partial x} \Big _{x=0} = f(t)$	$\theta_{B_1}^{(1)}(\mu) = \varepsilon$	$\widehat{\mathbf{B}}_1^{(1)} \in \mathbb{R}^{r_1}$
	linear	$\mathcal{A}_{1,1}^{(1)}(u) = \frac{\partial^2 u}{\partial x^2}$	$\theta_{A_{1,1}}^{(1)}(\mu) = \varepsilon$	$\widehat{\mathbf{A}}_{1,1}^{(1)} \in \mathbb{R}^{r_1 \times r_1}$
		$\mathcal{A}_{1,1}^{(2)}(u) = u$	$\theta_{A_{1,1}}^{(2)}(\mu) = -0.1/\varepsilon$	$\widehat{\mathbf{A}}_{1,1}^{(2)} \in \mathbb{R}^{r_1 \times r_1}$
	$\mathcal{A}_{1,2}^{(1)}(u) = u$	$\theta_{A_{1,2}}^{(1)}(\mu) = -1/\varepsilon$	$\widehat{\mathbf{A}}_{1,2}^{(1)} \in \mathbb{R}^{r_1 \times r_2}$	
quadratic	$\mathcal{H}_{1,11}^{(1)}(u, v) = uv$	$\theta_{H_{1,11}}^{(1)}(\mu) = 1.1/\varepsilon$	$\widehat{\mathbf{H}}_{1,11}^{(1)} \in \mathbb{R}^{r_1 \times \binom{r_1+1}{2}}$	
cubic	$\mathcal{G}_{1,111}^{(1)}(u, v, w) = uvw$	$\theta_{G_{1,111}}^{(1)}(\mu) = -1/\varepsilon$	$\widehat{\mathbf{G}}_{1,111}^{(1)} \in \mathbb{R}^{r_1 \times \binom{r_1+2}{3}}$	
Eq. (3.5b)	constant	$C_2^{(1)} = 1$	$\theta_{c_2}^{(1)}(\mu) = \alpha$	$\widehat{\mathbf{c}}_2^{(1)} \in \mathbb{R}^{r_2}$
	linear	$\mathcal{A}_{2,1}^{(1)}(u) = u$	$\theta_{A_{2,1}}^{(1)}(\mu) = \beta$	$\widehat{\mathbf{A}}_{2,1}^{(1)} \in \mathbb{R}^{r_2 \times r_1}$
		$\mathcal{A}_{2,2}^{(1)}(u) = u$	$\theta_{A_{2,2}}^{(1)}(\mu) = -\gamma$	$\widehat{\mathbf{A}}_{2,2}^{(1)} \in \mathbb{R}^{r_2 \times r_2}$

REMARK 3.3. An alternative approach to constructing a ROM for (3.1a)–(3.1c) is to consider the variables u_1, \dots, u_d as members of a product Hilbert space [38]. Let $\{\vec{v}_j\}_{j=1}^r \subset \mathcal{V}_\times = \mathcal{V}_1 \times \cdots \times \mathcal{V}_d$ be orthonormal with respect to the natural inner product

$$\langle \vec{u}, \vec{w} \rangle_{\mathcal{V}_\times} = \sum_{\ell=1}^d \langle u_\ell, w_\ell \rangle_{\mathcal{V}_\ell},$$

where $\vec{u} = (u_1, \dots, u_d)$ and $\vec{w} = (w_1, \dots, w_d)$. In this case, Galerkin projection yields a low-dimensional system with the structure of (2.7a)–(2.7b) where for each operator, the parametric coefficient functions are the union of the parametric coefficient functions for the operators of the same order from (3.1c). That is,

$$\{\theta_c^{(p)}\}_{p=1}^{q_c} = \bigcup_{\ell=1}^d \{\theta_{c_\ell}^{(p)}\}_{p=1}^{q_{c_\ell}}, \quad \{\theta_A^{(p)}\}_{p=1}^{q_A} = \bigcup_{\ell=1}^d \{\theta_{A_\ell}^{(p)}\}_{p=1}^{q_{A_\ell}}, \quad \{\theta_H^{(p)}\}_{p=1}^{q_H} = \bigcup_{\ell=1}^d \{\theta_{H_\ell}^{(p)}\}_{p=1}^{q_{H_\ell}}.$$

See [58] for details and an application to the shallow water equations. We call this approach monolithic because there is a single set of basis functions for all variables jointly, not individual basis functions corresponding to each variable as described in (3.2). The parametric dependencies in the corresponding ROM are determined by the governing equations, but the full system-level structure is not preserved.

EXAMPLE 3.4 (Monolithic FitzHugh–Nagumo System). *With a monolithic basis as described in Remark 3.3, the FitzHugh–Nagumo system (3.5a)–(3.5c) motivates a ROM with the form*

$$\begin{aligned} \frac{d}{dt} \hat{\mathbf{u}}(t; \mu) &= \frac{\alpha}{\varepsilon} \hat{\mathbf{c}}^{(1)} + \alpha \hat{\mathbf{c}}^{(2)} + \left(\varepsilon \hat{\mathbf{A}}^{(1)} - \frac{1}{\varepsilon} \hat{\mathbf{A}}^{(2)} + \beta \hat{\mathbf{A}}^{(3)} - \gamma \hat{\mathbf{A}}^{(4)} \right) \hat{\mathbf{u}}(t; \mu) \\ &+ \frac{1}{\varepsilon} \hat{\mathbf{H}}^{(1)} (\hat{\mathbf{u}}(t; \mu) \hat{\odot} \hat{\mathbf{u}}(t; \mu)) - \frac{1}{\varepsilon} \hat{\mathbf{G}}^{(1)} (\hat{\mathbf{u}}(t; \mu) \hat{\odot} \hat{\mathbf{u}}(t; \mu) \hat{\odot} \hat{\mathbf{u}}(t; \mu)) + \varepsilon \hat{\mathbf{B}}^{(1)} f(t). \end{aligned}$$

Contrast this with the ROM prescribed in Example 3.1, which inherits the entire structure of the PDE system (3.5a)–(3.5b). In particular, the reduced-order operators in Example 3.1 correspond directly to individual PDE operators (see Table 1).

4. Computational Procedure. Solving the pOpInf problem (2.13) requires samples of the solution $u(x, t; \mu)$ and its time derivative at times $\{t_j\}_{j=1}^K$ for each selected parameter value $\{\mu_i\}_{i=1}^s$. The quality of the resulting ROM depends on how well the orthonormal basis functions $\{v_j\}_{j=1}^r$ represent the solution at each parameter value and throughout the spatial and temporal domains. We therefore adopt the widely used proper orthogonal decomposition (POD) [1, 9, 31, 46], defined by the set of orthonormal functions that minimize the mean squared projection error of the sample data, i.e., solving the problem

$$\min_{v_1, \dots, v_r \in \mathcal{V}} \sum_{i=1}^s \sum_{j=1}^K \left\| u(\cdot, t_j; \mu_i) - \sum_{\ell=1}^r \langle v_\ell, u(\cdot, t_j; \mu_i) \rangle v_\ell \right\|_{\mathcal{V}}^2 \quad \text{subject to} \quad \langle v_i, v_j \rangle = \delta_{ij},$$

where $\|v\|_{\mathcal{V}} = \sqrt{\langle v, v \rangle}$ is the natural norm on \mathcal{V} . This data-driven choice of basis optimally represents the solution at the sampled parameter values, although the resulting ROM does not share such guarantees [8].

The conditioning of the linear least-squares problem (2.13) depends on the data matrix \mathbf{D} . If the parameter samples are chosen so that the pitfalls described in Theorem 2.3 are avoided, then the condition number of \mathbf{D} depends on the nature of the solution at the quadrature points $\{t_j\}_{j=1}^K$. To improve the conditioning, we introduce a Tikhonov regularization [50] so that (2.13) becomes

$$(4.1) \quad \min_{\mathbf{O}} \left\| \mathbf{D} \hat{\mathbf{O}}^\top - \mathbf{R}^\top \right\|_F^2 + \left\| \mathbf{\Lambda} \hat{\mathbf{O}}^\top \right\|_F^2, \quad \mathbf{\Lambda} \in \mathbb{R}^{q(r) \times q(r)}.$$

The solution to this regularized problem satisfies the modified normal equations,

$$(4.2) \quad (\mathbf{D}^\top \mathbf{D} + \mathbf{\Lambda}^\top \mathbf{\Lambda}) \widehat{\mathbf{O}}^\top = \mathbf{D}^\top \mathbf{R}^\top.$$

The regularizer $\mathbf{\Lambda}$ can be parameterized in a number of ways [32, 38]. One choice that provides flexibility without introducing a large number of hyperparameters is the diagonal matrix $\mathbf{\Lambda} = \mathbf{\Lambda}(\lambda_1, \lambda_2)$ defined such that

$$(4.3) \quad \|\mathbf{\Lambda}(\lambda_1, \lambda_2) \widehat{\mathbf{O}}^\top\|_F^2 = \lambda_1^2 \left(\sum_{p=1}^{q_c} \|\widehat{\mathbf{c}}^{(p)}\|_2^2 + \sum_{p=1}^{q_A} \|\widehat{\mathbf{A}}^{(p)}\|_F^2 \right) + \lambda_2^2 \sum_{p=1}^{q_H} \|\widehat{\mathbf{H}}^{(p)}\|_F^2.$$

This regularization structure groups the operators defining the ROM according to their polynomial order and drives $\widehat{\mathbf{O}}$ toward the zero matrix as $\lambda_1, \lambda_2 \rightarrow \infty$. Therefore, the regularization drives the resulting ROM toward the globally stable zero system $\frac{d}{dt} \widehat{\mathbf{u}}(t; \mu) = \mathbf{0}$. As before, the extension of (4.1) to a system of PDEs is straightforward where each PDE is written as an independent pOpInf problem (see (3.4)).

We choose the regularization hyperparameters $\lambda_1, \lambda_2 \geq 0$ to minimize the mean squared training error

$$\frac{1}{s} \sum_{i=1}^s \sum_{j=1}^K \|\widehat{\mathbf{u}}(t_j; \mu_i) - \widetilde{\mathbf{u}}(t_j; \mu_i)\|_2^2,$$

where $\widehat{\mathbf{u}}(t_j; \mu_i)$ is the training data and $\widetilde{\mathbf{u}}(t_j; \mu_i)$ is the result of integrating the ODE $\frac{d}{dt} \widetilde{\mathbf{u}}(t) = \mathbf{F}(\widehat{\mathbf{O}}; \widetilde{\mathbf{u}}, t, \mu)$ defined by the solution $\widehat{\mathbf{O}}$ of (4.1) with regularization hyperparameters λ_1 and λ_2 . This is an optimization problem in the principal quadrant of \mathbb{R}^2 , which we carry out with a sparse grid search followed by a derivative-free search method [34]. The approach guarantees that the selected hyperparameters result in a ROM that is stable for all training parameter values μ_1, \dots, μ_s . To further promote ROM stability throughout the parameter domain \mathcal{P} , we may constrain the hyperparameter selection problem by introducing a set of \bar{s} parameters $\{\bar{\mu}_i\}_{i=1}^{\bar{s}} \subset \mathcal{P}$ at which we demand stability from the learned ROM. Specifically, we disqualify hyperparameter pairings (λ_1, λ_2) whenever the time integrator for the resulting ROM diverges for any $\mu \in \{\bar{\mu}_i\}_{i=1}^{\bar{s}}$. This approach is similar to the hyperparameter selection strategy of [32] for non-parametric problems wherein the objective is to predict beyond the temporal domain for which data are available; our strategy focuses on the prediction in the parametric space but could also be adapted to address prediction in time.

We now summarize the computational procedure for solving (2.13): 1) select parameter values $\{\mu_i\}_{i=1}^s$ such that the associated $\Theta_c, \dots, \Theta_G$ have full column rank; 2) sample the PDE solution $u(x, t, \mu)$ and its time derivative for $t \in \{t_j\}_{j=1}^K$ for each $\mu \in \{\mu_i\}_{i=1}^s$; 3) compute the POD basis associated with the sampled solution data; 4) use the POD basis to project the solution data, obtaining $\widehat{\mathbf{U}}(\mu_i)$ and $\dot{\widehat{\mathbf{U}}}(\mu_i)$; 5) form the data matrix \mathbf{D} and the time derivative matrix \mathbf{R} ; 6) choose optimal regularization hyperparameters and solve (4.1) with these hyperparameters. In the computational setting, we sample the solution in step 2 by obtaining approximate discretized *solution snapshots* via a high-fidelity solver. For example, let $\{\mathbf{x}_\ell\}_{\ell=1}^N \subset \Omega$ be a discretization of Ω , and define

$$\mathbf{U}(\mu_i) = \begin{bmatrix} u(\mathbf{x}_1, t_1; \mu_i) & \cdots & u(\mathbf{x}_1, t_K; \mu_i) \\ \vdots & & \vdots \\ u(\mathbf{x}_N, t_1; \mu_i) & \cdots & u(\mathbf{x}_N, t_K; \mu_i) \end{bmatrix} \in \mathbb{R}^{N \times K},$$

the *snapshot matrix* for parameter μ_i . The rank- r POD basis of [step 3](#) is comprised of the first r left singular vectors of the concatenated snapshot matrices, that is,

$$(4.4) \quad \Phi \Sigma \Psi^\top = \left[\mathbf{U}(\mu_1) \quad \cdots \quad \mathbf{U}(\mu_s) \right] \in \mathbb{R}^{N \times sK}, \quad \mathbf{V} = \Phi_{:,1:r} \in \mathbb{R}^{N \times r},$$

where $\Phi \Sigma \Psi^\top$ is the singular value decomposition (SVD). With this notation, the projection of [step 4](#) is given by $\hat{\mathbf{U}}(\mu_i) = \mathbf{V}^\top \mathbf{U}(\mu_i)$, $i = 1, \dots, s$. If the time derivatives of u are not provided by the high-fidelity solver, they may be estimated as finite differences of the solution snapshots. The time integration error of the high-fidelity model, as well as the approximation error accompanying finite differences for the time derivatives, are additional motivations for utilizing the regularization strategy described previously. [Algorithm 4.1](#) fully details the procedure. The algorithm is presented for the general case of a system of d partial differential equations, but it may be simplified to a case of a single PDE by setting $d = 1$.

REMARK 4.1. *There are several alternatives for addressing poor conditioning in (2.13). Notably, [58] uses a truncated QR decomposition to solve a column subset selection problem, in which the size of the truncation is an integer hyperparameter that is typically selected via the L-curve criteria [22]. In some settings, building physical constraints into (2.13) also acts as an implicit regularization (see, e.g., [45]). We choose the Tikhonov regularization defined in (4.3) because 1) it features a continuous hyperparameter search space, enabling a precise selection; and 2) the structure of the regularizer can be tailored to specific model structure, as we will see in [subsection 5.2](#).*

5. Numerical Examples. We now present numerical results for the heat equation and FitzHugh–Nagumo system introduced in [Section 2](#) and [Section 3](#), respectively. In each of these examples we learn ROMs using [Algorithm 4.1](#) and compare their performance to classical ROMs obtained through intrusive projection [8]. However, our pOpInf methodology applies to situations in which intrusive projection is infeasible due to, e.g., variable transformations to induce polynomial structure [39]. The code for these experiments can be found at <https://github.com/Willcox-Research-Group/affine-parametric-opinf>.

5.1. Heat Equation. We return to the heat equation of [Example 2.2](#) and [2.4](#), setting $\bar{x} = 2/3$ and $u_0(x; \mu) = 1 - (1 - x)^{50} - x^{50}$. For each of the $s = 5$ parameter samples shown in [Figure 1](#), we generate solutions by discretizing the spatial domain $\Omega = (0, 1)$ with a uniform grid of $N = 1000$ points and approximating the spatial derivative with second-order central finite differences. The resulting semi-discrete ODE, called the full-order model, is integrated in time with the first-order implicit Euler scheme on $K = 1500$ uniformly spaced time steps in $[t_0, t_f] = [0, 1.5]$. [Figure 1](#) also shows example snapshots for each parameter sample, demonstrating that variation in the parameters $\mu = (\alpha, \beta)$ determines the diffusion dynamics. To select the number of modes in the POD basis, define the cumulative energy

$$\mathcal{E}(r) = \sum_{j=1}^r \sigma_j^2 / \sum_{j=1}^N \sigma_j^2,$$

where σ_j is the j th singular value in the POD factorization (4.4). Note that $\mathcal{E}(r)$ is a nondecreasing function of r . We select r to be the smallest integer such that the

Algorithm 4.1 Regularized parametric Operator Inference for systems of PDEs

```

1: procedure POpINF(
    training parameter values  $\mu_1, \dots, \mu_s \in \mathcal{P}$ ,
    training snapshots  $\mathbf{U}_1(\mu_i), \dots, \mathbf{U}_d(\mu_i) \in \mathbb{R}^{N \times K}$  for  $i = 1, \dots, s$ ,
    affine coefficient functions  $\theta = \{\theta_{c_1}^{(1)}, \dots, \theta_{H_{d,dd}}^{(q_{H_{d,dd}})}\} : \mathcal{P} \rightarrow \mathbb{R}$ ,
    reduced dimensions  $r_1, \dots, r_d \in \mathbb{N}$ ,
    stability parameter values  $\bar{\mu}_1, \dots, \bar{\mu}_{\bar{s}} \in \mathcal{P}$  (optional, else  $\bar{s} = 0$ )
)
2: # Project training data to low-dimensional subspaces.
3: for  $\ell = 1, \dots, d$  do
4:      $\mathbf{V}_\ell \leftarrow \text{pod}([\mathbf{U}_\ell(\mu_1) \cdots \mathbf{U}_\ell(\mu_s)], r_\ell)$  # Rank- $r_\ell$  POD basis.
5:     for  $i = 1, \dots, s$  do
6:          $\hat{\mathbf{U}}_\ell(\mu_i) \leftarrow \mathbf{V}_\ell^\top \mathbf{U}_\ell(\mu_i)$  # Projected solution data.
7:          $\hat{\dot{\mathbf{U}}}_\ell(\mu_i) \leftarrow \frac{d}{dt} \hat{\mathbf{U}}_\ell(\mu_i)$  # Projected time derivatives.
8: # Construct pOpInf matrices.
9: for  $\ell = 1, \dots, d$  do
10:     $\mathbf{D}_\ell \leftarrow$  build the  $\ell$ th data matrix from  $\hat{\mathbf{U}}_1(\mu_1), \dots, \hat{\mathbf{U}}_d(\mu_s), \theta, \mu_1, \dots, \mu_s$ 
11:     $\mathbf{R}_\ell \leftarrow [\hat{\dot{\mathbf{U}}}_\ell(\mu_1) \cdots \hat{\dot{\mathbf{U}}}_\ell(\mu_s)]$ 
12: # Compute pOpInf solution with optimal hyperparameters.
13:  $\lambda_1^*, \lambda_2^* \leftarrow \text{argmin TRAININGERROR}(\lambda_1, \lambda_2)$ 
14: return REGOPINF( $\lambda_1^*, \lambda_2^*$ )

1: procedure TRAININGERROR( $\lambda_1, \lambda_2$ )
2:     $\hat{\mathbf{O}}_1, \dots, \hat{\mathbf{O}}_d \leftarrow \text{REGOPINF}(\lambda_1, \lambda_2)$ 
3:    if  $\bar{s} > 0$  then
4:        for  $i = 1, \dots, \bar{s}$  do # Check behavior at stability parameters.
5:            if integrating (3.3a),  $\ell = 1, \dots, d$ ,  $\mu = \bar{\mu}_i$ , over  $[t_0, t_f]$  diverges then
6:                return  $\infty$ 
7:        for  $i = 1, \dots, s$  do # Calculate error at training parameters.
8:             $\tilde{\mathbf{U}}_1(\mu_i), \dots, \tilde{\mathbf{U}}_d(\mu_i) \leftarrow$  integrate (3.3a),  $\ell = 1, \dots, d$ ,  $\mu = \mu_i$ , over  $[t_0, t_f]$ 
9:        return  $\frac{1}{sd} \sum_{\ell=1}^d \sum_{i=1}^s \|\hat{\mathbf{U}}_\ell(\mu_i) - \tilde{\mathbf{U}}_\ell(\mu_i)\|_F^2$ 

1: procedure REGOPINF( $\lambda_1, \lambda_2$ )
2:     $\mathbf{\Lambda}^2 \leftarrow \mathbf{\Lambda}(\lambda_1, \lambda_2)^\top \mathbf{\Lambda}(\lambda_1, \lambda_2)$  # Construct the regularizer.
3:    for  $\ell = 1, \dots, d$  do
4:         $\hat{\mathbf{O}}_\ell^\top \leftarrow (\mathbf{D}_\ell^\top \mathbf{D}_\ell + \mathbf{\Lambda}^2)^{-1} \mathbf{D}_\ell^\top \mathbf{R}_\ell^\top$  # Solve the  $\ell$ th pOpInf problem.
5:    return  $\hat{\mathbf{O}}_1, \dots, \hat{\mathbf{O}}_d$ 
    
```

residual energy $1 - \mathcal{E}(r)$, or the energy in the non-retained modes, lies below a fixed threshold $\epsilon > 0$. For this problem, setting $\epsilon = 10^{-7}$ results in $r = 12$, while setting $\epsilon = 10^{-10}$ results in $r = 19$ (see Figure 2).

We compute $\hat{\mathbf{O}}$ using Algorithm 4.1 with the $s = 5$ training parameter values and without any stability parameter values ($\bar{s} = 0$). Since (2.9) is a linear system, the regularization is parameterized by a single hyperparameter λ_1 , hence the minimiza-




FIG. 1. Experimental parameter samples (left) and associated snapshots at intermediate time $t = t_{75} = 0.075$ (right) for the heat equation problem (2.8a)–(2.8c). The vertical line $x = \bar{x}$ marks the point in the domain where the diffusion constant switches between α and β .




FIG. 2. Residual energy decay (left) and the six dominant POD basis functions (right) for the snapshot set generated at the parameter samples in Figure 1. For $r \geq 12$, we have $1 - \mathcal{E}(r) < 10^{-7}$; if $r \geq 19$, then $1 - \mathcal{E}(r) < 10^{-10}$.

tion (4.1) for this problem is given by

$$\min_{\hat{\mathbf{O}}} \left\| \mathbf{D}\hat{\mathbf{O}}^\top - \mathbf{R}^\top \right\|_F^2 + \lambda_1^2 \left(\|\hat{\mathbf{A}}^{(1)}\|_F^2 + \|\hat{\mathbf{A}}^{(2)}\|_F^2 \right).$$

Here, the time derivative data \mathbf{R} is estimated by first-order backward differences of the training states. The resulting ROM is integrated using the same implicit Euler scheme as the full-order model. To evaluate the performance of the ROM in terms of the parameters, we discretize \mathcal{P} in a 40×40 uniform grid. For each μ in the grid, we compute full-order solutions $\mathbf{u}(t_1; \mu), \dots, \mathbf{u}(t_K; \mu) \in \mathbb{R}^N$ and integrate the ROM to obtain reduced states $\tilde{\mathbf{u}}(t_1; \mu), \dots, \tilde{\mathbf{u}}(t_K; \mu) \in \mathbb{R}^r$, then compute the relative L^2 -norm error

$$(5.1) \quad \frac{\|\mathbf{V}\tilde{\mathbf{u}}(\cdot; \mu) - \mathbf{u}(\cdot; \mu)\|_{L^2([t_0, t_f])}}{\|\mathbf{u}(\cdot; \mu)\|_{L^2([t_0, t_f])}}, \quad \|\mathbf{w}(\cdot)\|_{L^2([t_0, t_f])} = \left(\int_{t_0}^{t_f} \|\mathbf{w}(t)\|_2^2 dt \right)^{1/2},$$

estimating the time integrals via the trapezoidal rule. Figure 3 shows the results for $r = 12$ and $r = 19$ and compares them to the relative projection error induced by the corresponding bases, given by

$$(5.2) \quad \frac{\|\mathbf{u}(\cdot; \mu) - \mathbf{V}\mathbf{V}^\top \mathbf{u}(\cdot; \mu)\|_{L^2([t_0, t_f])}}{\|\mathbf{u}(\cdot; \mu)\|_{L^2([t_0, t_f])}}.$$




FIG. 3. Relative L^2 projection errors (left), pOpInf ROM errors (center), and intrusive ROM errors (right) over the parameter domain \mathcal{P} for the heat problem (2.8a)–(2.8c). Both ROMs use the same basis with $r = 12$ (top) or $r = 19$ (bottom) basis functions. The parameter samples used to generate the training set are marked as stars.

Throughout \mathcal{P} , the relative projection error is on the order of 10^{-4} for $r = 12$ and 10^{-5} for $r = 19$, which indicates that the training snapshots have sufficient information to represent the solution well for any $\mu \in \mathcal{P}$. The results in Figure 3 highlight the ability of pOpInf ROMs to generalize beyond the training data; i.e., the ROMs perform well for (α, β) pairs away from the arc $\alpha^2 + \beta^2 = 4$ where the parameter samples lie.

Figure 3 also shows the error of the ROM based on intrusive Galerkin projection,

$$\frac{d}{dt} \hat{\mathbf{u}}(t; \mu) = \left(\alpha \bar{\mathbf{A}}^{(1)} + \beta \bar{\mathbf{A}}^{(2)} \right) \hat{\mathbf{u}}(t; \mu),$$

where the entries of $\bar{\mathbf{A}}^{(1)}$ and $\bar{\mathbf{A}}^{(2)}$ are computed as $[\bar{\mathbf{A}}^{(p)}]_{ij} = \langle v_i, \mathcal{A}^{(p)}(v_j) \rangle$, $p = 1, 2$, as in (2.6). Note that this requires explicit access to the operators $\mathcal{A}^{(1)}$ and $\mathcal{A}^{(2)}$. The performance of the pOpInf ROMs—which are constructed without access to $\mathcal{A}^{(1)}$ and $\mathcal{A}^{(2)}$ —is highly comparable to the performance of the intrusive ROMs. In particular, each ROM performs well near the line $\alpha = \beta$ and worse near the parameter domain boundaries $\alpha = 0.1$ and $\beta = 0.1$, with errors less than 0.6% for $r = 12$ and 0.03% for $r = 19$ throughout \mathcal{P} (see Table 2). Using 100 random samples in \mathcal{P} , the average CPU time to integrate the full-order model is ~ 0.087 s, while the ROMs of size $r = 12$ integrate in ~ 0.013 s. Hence the computational speedup factor is $0.087/0.013 \approx 6.7$ times. The full-order solver takes advantage of sparsity in a large system, while the ROM utilizes a small but dense system.

5.2. FitzHugh–Nagumo System. The neuron model (3.5a)–(3.5b) introduced in Example 3.1 features a four-dimensional parameter space. We generate training data at each of the $504 = 6 \times 6 \times 2 \times 7$ unique parameter realizations $\mu = (\alpha, \beta, \gamma, \varepsilon)$

TABLE 2
 Statistics for the relative errors displayed in [Figure 3](#).

		Projection	pOpInf ROM	Intrusive ROM
$r = 12$	maximum	4.716×10^{-4}	2.161×10^{-3}	5.520×10^{-3}
	median	2.297×10^{-4}	7.415×10^{-4}	3.862×10^{-4}
	minimum	1.610×10^{-4}	3.683×10^{-4}	1.904×10^{-4}
$r = 19$	maximum	2.659×10^{-5}	2.232×10^{-4}	2.767×10^{-4}
	median	1.034×10^{-5}	2.915×10^{-5}	4.386×10^{-5}
	minimum	5.895×10^{-6}	1.212×10^{-5}	7.643×10^{-6}

for

$$\begin{aligned}
 (\text{training set}) \quad & \alpha \in \{0.025, 0.035, \dots, 0.075\}, \quad \beta \in \{0.25, 0.35, \dots, 0.75\}, \\
 & \gamma \in \{2.0, 2.5\}, \quad \varepsilon \in \{0.010, 0.015, \dots, 0.040\}.
 \end{aligned}$$

For each parameter realization μ in the training set, we solve [\(3.5a\)–\(3.5b\)](#) by discretizing the domain $\Omega = (0, 1)$ with $N_x = 512$ spatial points and approximating the differential term with central finite differences. The total spatial dimension of this full-order model is thus $N = 2N_x = 1024$. Because the FitzHugh–Nagumo system is stiff, the spatially discretized equations benefit from a time integrator with an adaptive time step [\[15, 47\]](#). We make use of the implicit fifth-order Runge–Kutta Radau IIA method with adaptive time stepping [\[57\]](#), available in Python as `scipy.integrate.Radau` [\[56\]](#), with absolute and relative tolerances set to 10^{-6} . The solution is computed at equally spaced times over the time domain $[t_0, t_f] = [0, 4]$ with spacing $\delta t = 10^{-3}$, after which the time derivatives of the states are estimated with sixth-order finite differences. The results are recorded for every tenth snapshot-derivative pair, resulting in $K = 400$ snapshots per training parameter realization with spacing $\delta t = 10^{-2}$. [Figure 4](#) shows the phase plot of u_1 against u_2 at multiple spatial coordinate values for three of the training parameter samples, demonstrating that the system exhibits a diverse range of dynamical behaviors as the parameters are varied. To evaluate ROM performance with respect to the parameters, we also solve the full-order model at the $10,749 = 11 \times 11 \times 3 \times 31 - 504$ additional parameter realizations $\mu = (\alpha, \beta, \gamma, \varepsilon)$ with

$$\begin{aligned}
 (\text{testing set}) \quad & \alpha \in \{0.025, 0.030, \dots, 0.075\}, \quad \beta \in \{0.25, 0.30, \dots, 0.75\}, \\
 & \gamma \in \{2.00, 2.25, 2.50\}, \quad \varepsilon \in \{0.010, 0.011, \dots, 0.040\}.
 \end{aligned}$$

Parameter realizations in the training set are not included in the testing set.

The FitzHugh–Nagumo system [\(3.5a\)–\(3.5b\)](#) exhibits a rapid singular Hopf bifurcation with respect to ε wherein the limit cycle collapses to a stable fixed point as ε increases [\[5\]](#). Near this transition, the full-order model and ROM solutions are highly sensitive to ε and the time integration scheme (see [Figure 6](#)). Including such parameter realizations in the training and testing sets makes it difficult to assess ROM accuracy. For each parameter realization $\mu = (\alpha, \beta, \gamma, \varepsilon)$ in the training and testing sets, we compare the corresponding full-order model solution to the solutions at $\mu_{\varepsilon-} = (\alpha, \beta, \gamma, \varepsilon - 0.001)$ and $\mu_{\varepsilon+} = (\alpha, \beta, \gamma, \varepsilon + 0.001)$. If the relative difference between neighboring solutions exceeds 50%, that is, if

$$\max \left\{ \frac{\|\mathbf{u}(\cdot; \mu) - \mathbf{u}(\cdot; \mu_{\varepsilon-})\|_{L^2([t_0, t_f])}}{\|\mathbf{u}(\cdot; \mu)\|_{L^2([t_0, t_f])}}, \frac{\|\mathbf{u}(\cdot; \mu) - \mathbf{u}(\cdot; \mu_{\varepsilon+})\|_{L^2([t_0, t_f])}}{\|\mathbf{u}(\cdot; \mu)\|_{L^2([t_0, t_f])}} \right\} > 0.5,$$




FIG. 4. Phase portraits of training trajectories for the FitzHugh–Nagumo system (3.5a)–(3.5b), traced out at various points in the spatial domain. The center trajectory has a limit cycle, while the trajectories on the left and right converge to a single point. Small ε values de-emphasize the diffusion term and drive the system toward spatial homogeneity; larger ε values result in more variation across the spatial domain but which decreases with time.




FIG. 5. Decay of the residual energy in the training data for u_1 and u_2 in the FitzHugh–Nagumo problem (3.5a)–(3.5b) (left) and the corresponding selected basis sizes (right). Demanding that $1 - \mathcal{E}(r_\ell) < 10^{-7}$ requires $r_1 = 12$ and $r_2 = 9$ POD modes for u_1 and u_2 , respectively.

then we remove μ from the data. This prompts us to remove two parameter realizations from the training set ($\sim 0.4\%$ of the original training set) and eighty-nine parameter realizations from the testing set ($\sim 0.8\%$ of the original testing set). Hence the size of the final training set is $s = 504 - 2 = 502$ and the size of the final testing set is $10,749 - 89 = 10,660$, with size ratio $10,660/502 \approx 21.24$.

We select the number of POD modes r_1 and r_2 for the variables u_1 and u_2 , respectively, based on the residual energy $1 - \mathcal{E}(r_\ell)$ of the training data, $\ell = 1, 2$. For each of the resulting (r_1, r_2) pairs shown in Figure 5, we learn a pOpInf ROM via Algorithm 4.1 with $\bar{s} = 0$. To simplify the hyperparameter search, we regularize only the quadratic and cubic terms, so that the regularized pOpInf data residuals are

$$\min_{\hat{\mathbf{O}}_1} \left\| \mathbf{D}_1 \hat{\mathbf{O}}_1^\top - \mathbf{R}_1^\top \right\|_F^2 + \lambda_2 \left\| \hat{\mathbf{H}}_{1,11}^{(1)} \right\|_F^2 + \lambda_3 \left\| \hat{\mathbf{G}}_{1,111}^{(1)} \right\|_F^2, \quad \min_{\hat{\mathbf{O}}_2} \left\| \mathbf{D}_2 \hat{\mathbf{O}}_2^\top - \mathbf{R}_2^\top \right\|_F^2.$$

Regularization of the additional terms (λ_1 in (4.3)) was observed to have a marginal effect on the results. For each chosen basis size, we also compute a ROM based on intrusive projection for comparison with the pOpInf ROM. The intrusive ROM




FIG. 6. Phase portraits of trajectories for the FitzHugh–Nagumo system (3.5a)–(3.5b), traced out at various points in the spatial domain, at the parameter realizations $\alpha = 0.035$, $\beta = 0.75$, $\gamma = 2.5$, and three choices of ε . Near the Hopf bifurcation, small changes in ε result in rapid changes in the system behavior.

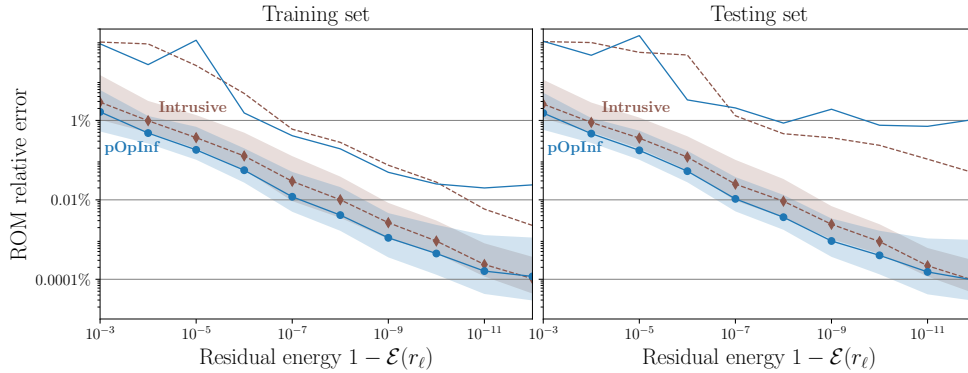


FIG. 7. Relative errors (calculated over space and time) of the learned pOpInf ROM and the intrusive ROM for the FitzHugh–Nagumo system. As residual energy decreases from left to right, the sizes of the underlying POD bases increase according to Figure 5. The shaded regions show the 10%–90% interdecile range of the error across all training samples (left) or testing samples (right), with the corresponding median and maximum errors denoted by the lines for the pOpInf ROMs and dashed lines for the intrusive ROMs.

is derived explicitly without any need for approximation of the nonlinear terms via DEIM (see Remark 3.2). The ROMs are integrated for each parameter quadruple $\mu = (\alpha, \beta, \gamma, \varepsilon)$ in the testing set using the same adaptive time integration scheme as the full-order model. Both ROMs are stable throughout the testing set with the exception of the pOpInf ROM with basis sizes $r_1 = 7$ and $r_2 = 5$, for which the time integrator diverges at three testing parameter realizations. For this case, we repeat Algorithm 4.1 with the same training data as before but now using these $\bar{s} = 3$ testing parameter realizations as the stability parameter inputs $\bar{\mu}_1, \bar{\mu}_2, \bar{\mu}_3$. The resulting ROM is then stable throughout the entire testing set.

The relative error compared to the full-order model is computed as in (5.1), yielding a single error value for each parameter realization in the training and testing sets. Figure 7 shows the 10% quantile, median, 90% quantile, and maximum of the relative errors for each ROM. The median relative error is similar in the training and testing sets and decreases steadily as the basis sizes increase. The pOpInf ROMs generally outperform the intrusive ROMs in terms of median error. This is unsurprising for the training set since the pOpInf ROMs minimize a residual driven by the training data,




FIG. 8. Phase portraits of testing trajectories for the FitzHugh–Nagumo problem (3.5a)–(3.5b), traced out at various points in the spatial domain. The solid lines are the full-order trajectories, and the dashed lines are the outputs of the learned pOpInf ROM with $r_1 = 12$ and $r_2 = 9$ POD modes. The total ROM relative errors in space and time are, from left to right, 0.028%, 0.190%, and 0.007%; the median relative error on the testing set for $r_1 = 12$ and $r_2 = 9$ is 0.011%.

whereas the intrusive ROMs only see the data through the construction of the POD bases. The tight interdecile range shows that the ROM error is mostly consistent throughout the parameter space; however, as the basis sizes increase the maximum pOpInf ROM error in both the training and testing sets eventually levels out while the maximum intrusive ROM error continues to decrease. For additional comparison, Figure 8 shows select trajectories of a single pOpInf ROM versus the full-order model. The reduced-order and full-order results are indistinguishable to the eye, which is representative of ROM performance throughout the testing set.

To probe the sensitivity of the method to the amount of training data, we repeat the numerical experiments by again estimating the time derivatives of the state with $\delta t = 10^{-3}$, then downsampling the snapshot data to only $K = 40$ snapshots per training parameter with uniform temporal spacing $\delta t = 10^{-1}$. This is equivalent to using only every tenth snapshot from the previous experiment. From this dataset we compute new POD bases and the corresponding intrusive ROMs. An initial application of Algorithm 4.1 with $\bar{s} = 0$ yields pOpInf ROMs that are stable except at three and ten testing parameter realizations for $(r_1, r_2) = (5, 4)$ and $(r_1, r_2) = (21, 17)$, respectively. For these models, reapplying Algorithm 4.1 with these testing parameter realizations as the stability parameter inputs $\bar{\mu}_1, \dots, \bar{\mu}_{\bar{s}}$ ($\bar{s} = 3$ and $\bar{s} = 10$, respectively) yields pOpInf ROMs that are stable over the entire testing set. Figure 9 shows the relative errors of the new intrusive and pOpInf ROMs over the testing set and compares the errors of this $K = 40$ experiment with the errors from the original $K = 400$ experiment. The $K = 40$ intrusive ROM error is similar to the $K = 400$ results shown in Figure 7, but the pOpInf ROM error levels out near $\mathcal{E}(r_\ell) = 10^{-10}$, after which the intrusive ROM outperforms the pOpInf ROM. Though not shown, the same trend occurs for the error over the training set. With $K = 40$, the median pOpInf ROM error is consistently slightly higher than the $K = 400$ pOpInf ROM error, and the gap between the two errors increase with the basis sizes. For this problem, Algorithm 4.1 produces more robust and accurate ROMs when more training data are available.

6. Conclusions. This paper has proposed a non-intrusive parametric model reduction method for parameterized PDEs based on the Operator Inference framework. The approach eliminates the need for interpolation by explicitly embedding the polynomial affine-parametric structure of the PDE system into the resulting ROM and

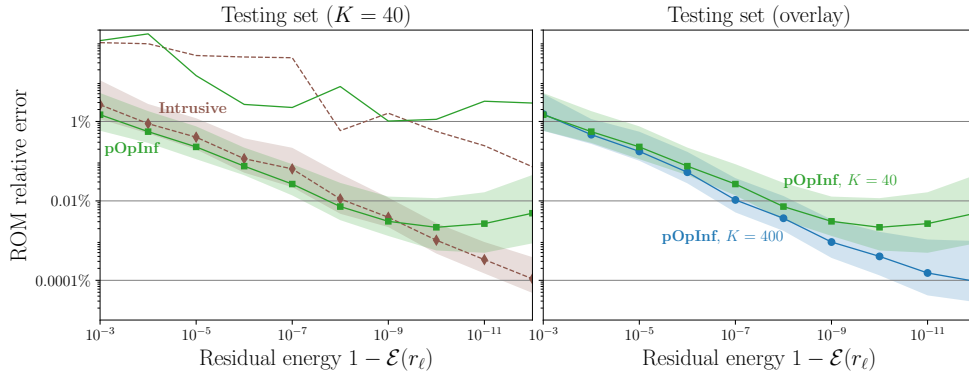


FIG. 9. Relative errors on the testing set (calculated over space and time) of the learned $pOpInf$ ROM for the FitzHugh–Nagumo system with only $K = 40$ snapshots per training parameter sample. The shaded regions show the 10%–90% interdecile range of the error; median and maximum errors are denoted by solid lines. On the left, a comparison to the corresponding intrusive ROM (dashed lines); on the right, a comparison to the $pOpInf$ ROM with $K = 400$ snapshots per training parameter from Figure 7.

uses an optimization-based regularization strategy to ensure well-posedness in the learning problem. The parametric ROMs can later be used in outer-loop applications to expedite model evaluations for any choice of the parameters. The efficacy of the method has been demonstrated for two numerical examples: a heat equation with a two-dimensional parametric space, and the FitzHugh–Nagumo system with a four-dimensional parametric space. The resulting ROMs are capable of capturing the behavior of the PDE for parameters outside of the training set and perform favorably on average when compared to the intrusive ROM with the same affine structure.

It was shown in the FitzHugh–Nagumo example that the learned ROMs successfully capture the inherently different behaviors of the system that come with changes in the parameters (see Figure 8). Yet, as with other data-driven approaches, the quality of inferred ROMs depends strongly on the training set, and one cannot expect a data-driven ROM to produce a particular dynamical behavior that differs wildly from the training data. An important future direction is therefore the automation of an efficient parameter sampling strategy for the offline stage. We further showed that the performance of $pOpInf$ degrades slightly as data becomes more sparse in time, although the results can still be quite good as long as the time derivatives are estimated accurately. Further research is needed to address problems for which data are sparse in time and cannot be upsampled to produce accurate time derivative estimates via finite differences. Finally, as the framework proposed here relies on an affine structure in the parametric dependence, future work should address adapting this framework to non-affine parametric problems or to problems for which the parametric structure is unknown.

Acknowledgements. This work has been supported in part by the US Department of Energy AEOLUS MMICC center under award DE-SC0019303, program manager W. Spotz; by the Air Force Center of Excellence on Multi-Fidelity Modeling of Rocket Combustor Dynamics under award FA9550-17-1-0195; and by the US Department of Energy National Nuclear Security Administration under award DE-NA0003969. The authors also wish to thank Elizabeth Qian and Vincent Martinez for their insightful comments.

Appendix A. Matricization of Tensors. Following the notation of [27], for matrices $\mathbf{W} \in \mathbb{R}^{r \times s}$ and $\mathbf{Z} \in \mathbb{R}^{m \times n}$, let $\mathbf{W} \otimes \mathbf{Z}$ denote the Kronecker product [27, 53]:

$$\mathbf{W} \otimes \mathbf{Z} = \begin{bmatrix} w_{11} & \cdots & w_{1s} \\ \vdots & \ddots & \vdots \\ w_{r1} & \cdots & w_{rs} \end{bmatrix} \otimes \mathbf{Z} := \begin{bmatrix} w_{11}\mathbf{Z} & \cdots & w_{1s}\mathbf{Z} \\ \vdots & \ddots & \vdots \\ w_{r1}\mathbf{Z} & \cdots & w_{rs}\mathbf{Z} \end{bmatrix} \in \mathbb{R}^{rm \times sn},$$

where w_{ij} is the component from the i th row and j th column of \mathbf{W} . The definition applies to vectors by setting $s = n = 1$. For $\mathbf{Z} \in \mathbb{R}^{m \times s}$ (i.e., $n = s$), define $\mathbf{W} \odot \mathbf{Z}$ to be the Khatri-Rao product, i.e., the column-wise Kronecker product [25, 27]:

$$\mathbf{W} \odot \mathbf{Z} := \left[\mathbf{w}_1 \otimes \mathbf{z}_1 \mid \mathbf{w}_2 \otimes \mathbf{z}_2 \mid \cdots \mid \mathbf{w}_s \otimes \mathbf{z}_s \right] \in \mathbb{R}^{rm \times s},$$

where \mathbf{w}_i and \mathbf{z}_i are the i th columns of \mathbf{W} and \mathbf{Z} , respectively.

Let $\mathbf{w} = [w_1 \cdots w_r]^\top \in \mathbb{R}^r$. As the product $\mathbf{w} \odot \mathbf{w} = \mathbf{w} \otimes \mathbf{w}$ has redundant terms (for instance, $w_1 w_2 = w_2 w_1$ appears twice), we introduce compact second- and third-order Khatri-Rao products, defined as

$$\mathbf{w} \hat{\odot} \mathbf{w} := \begin{bmatrix} w_1^2 \\ w_2(\mathbf{w}_{1:2}) \\ \vdots \\ w_r(\mathbf{w}_{1:r}) \end{bmatrix} \in \mathbb{R}^{\binom{r+1}{2}}, \quad \mathbf{w} \hat{\odot} \mathbf{w} \hat{\odot} \mathbf{w} := \begin{bmatrix} w_1^3 \\ w_2(\mathbf{w}_{1:2} \hat{\odot} \mathbf{w}_{1:2}) \\ \vdots \\ w_r(\mathbf{w}_{1:r} \hat{\odot} \mathbf{w}_{1:r}) \end{bmatrix} \in \mathbb{R}^{\binom{r+2}{3}},$$

where $\mathbf{w}_{1:i} = [w_1 \cdots w_i]^\top \in \mathbb{R}^i$ contains the first i entries of \mathbf{w} . Each entry of $\mathbf{w} \hat{\odot} \mathbf{w}$ is a product of 2 entries of \mathbf{w} , and each product appears exactly once; in other words, $\mathbf{w} \hat{\odot} \mathbf{w}$ contains the unique terms of $\mathbf{w}\mathbf{w}^\top$, the 2-fold tensorization of \mathbf{w} [27, 35]. Similarly, each entry of $\mathbf{w} \hat{\odot} \mathbf{w} \hat{\odot} \mathbf{w}$ is a unique product of 3 entries of \mathbf{w} . For matrices, the definition applies column-wise:

$$\left[\mathbf{w}_1 \mid \cdots \mid \mathbf{w}_s \right] \hat{\odot} \left[\mathbf{w}_1 \mid \cdots \mid \mathbf{w}_s \right] = \left[\mathbf{w}_1 \hat{\odot} \mathbf{w}_1 \mid \cdots \mid \mathbf{w}_s \hat{\odot} \mathbf{w}_s \right].$$

With this notation, any linear combination of products of two entries of \mathbf{w} can be represented by the product

$$\mathbf{z}^\top (\mathbf{w} \hat{\odot} \mathbf{w}) = z_1 w_1^2 + z_2 w_1 w_2 + z_3 w_2^2 + z_4 w_1 w_3 + \cdots, \quad \mathbf{z} = \begin{bmatrix} z_1 \\ z_2 \\ \vdots \end{bmatrix} \in \mathbb{R}^{\binom{r+1}{2}}.$$

This observation allows us to convert (2.6) to a compact matrix product representation. Specifically, we define $\hat{\mathbf{H}} \in \mathbb{R}^{r \times \binom{r+1}{2}}$ to be the matrix such that the i th component of the product $\hat{\mathbf{H}}(\hat{\mathbf{u}} \hat{\odot} \hat{\mathbf{u}})$ is given by

$$[\hat{\mathbf{H}}(\hat{\mathbf{u}} \hat{\odot} \hat{\mathbf{u}})]_i := \sum_{j=1}^r \sum_{l=1}^r \langle v_i, \mathcal{H}(v_j, v_l; \mu) \rangle \hat{u}_j \hat{u}_l.$$

Appendix B. Technical Lemma. The following lemma supports [Theorem 2.3](#). Here, $\mathbf{0}_k \in \mathbb{R}^k$ is the vector of k zeros and $\mathbf{1}_k \in \mathbb{R}^k$ denotes the vector of k ones.

LEMMA B.1. Let $\mathbf{y}_1, \dots, \mathbf{y}_s \in \mathbb{R}^q$ and $\mathbf{Z}_1, \dots, \mathbf{Z}_s \in \mathbb{R}^{k \times r}$. Consider the matrix

$$\mathbf{W} := \begin{bmatrix} \mathbf{y}_1^\top \otimes \mathbf{Z}_1 \\ \vdots \\ \mathbf{y}_s^\top \otimes \mathbf{Z}_s \end{bmatrix} \in \mathbb{R}^{sk \times qr}.$$

If either of the matrices

$$\mathbf{Y} := \begin{bmatrix} \mathbf{y}_1^\top \\ \vdots \\ \mathbf{y}_s^\top \end{bmatrix} \in \mathbb{R}^{s \times q}, \quad \mathbf{Z} := \begin{bmatrix} \mathbf{Z}_1 \\ \vdots \\ \mathbf{Z}_s \end{bmatrix} \in \mathbb{R}^{sk \times r}$$

do not have full column rank, then neither does \mathbf{W} . Conversely, if \mathbf{Y} and each $\mathbf{Z}_1, \dots, \mathbf{Z}_s$ have full column rank, then so does \mathbf{W} .

Proof. Assume that \mathbf{Y} does not have full column rank. Then there exists a nonzero vector $\boldsymbol{\alpha} \in \mathbb{R}^q$ such that $\mathbf{Y}\boldsymbol{\alpha} = \mathbf{0}_s$, that is, $\mathbf{y}_i^\top \boldsymbol{\alpha} = 0$ for $i = 1, \dots, s$. Using the mixed-product property of the Kronecker product,

$$(\mathbf{y}_i^\top \otimes \mathbf{Z}_i)(\boldsymbol{\alpha} \otimes \mathbf{1}_r) = (\mathbf{y}_i^\top \boldsymbol{\alpha}) \otimes (\mathbf{Z}_i \mathbf{1}_r) = 0 \otimes (\mathbf{Z}_i \mathbf{1}_r) = \mathbf{0}_k$$

for $i = 1, \dots, s$. Then $\mathbf{W}(\boldsymbol{\alpha} \otimes \mathbf{1}_r) = \mathbf{0}_{sk}$. But $\boldsymbol{\alpha} \otimes \mathbf{1}_r$ is a nonzero vector, which implies that the columns of \mathbf{W} are linearly dependent.

Next, suppose that \mathbf{Z} does not have full column rank. Then $\mathbf{Z}\boldsymbol{\beta} = \mathbf{0}_{sk}$ for some nonzero vector $\boldsymbol{\beta} \in \mathbb{R}^r$, implying $\mathbf{Z}_i\boldsymbol{\beta} = \mathbf{0}_k$ for $i = 1, \dots, s$. Then

$$(\mathbf{y}_i^\top \otimes \mathbf{Z}_i)(\mathbf{1}_q \otimes \boldsymbol{\beta}) = (\mathbf{y}_i^\top \mathbf{1}_q) \otimes (\mathbf{Z}_i\boldsymbol{\beta}) = (\mathbf{y}_i^\top \mathbf{1}_q) \otimes \mathbf{0}_k = \mathbf{0}_k$$

for $i = 1, \dots, s$, so that $\mathbf{W}(\mathbf{1}_q \otimes \boldsymbol{\beta}) = \mathbf{0}_{sk}$. Since $\mathbf{1}_q \otimes \boldsymbol{\beta}$ is a nonzero vector, \mathbf{W} does not have full column rank.

For the converse statement, assume that \mathbf{Y} and $\mathbf{Z}_1, \dots, \mathbf{Z}_s$ each have full column rank, and suppose $\boldsymbol{\gamma} = [\boldsymbol{\gamma}_1^\top \cdots \boldsymbol{\gamma}_q^\top]^\top \in \mathbb{R}^{qr}$, $\boldsymbol{\gamma}_j \in \mathbb{R}^r$, satisfies $\mathbf{W}\boldsymbol{\gamma} = \mathbf{0}_{sk}$, i.e., $(\mathbf{y}_i^\top \otimes \mathbf{Z}_i)\boldsymbol{\gamma} = \mathbf{0}_k$ for $i = 1, \dots, s$. Denoting $\boldsymbol{\Gamma} = [\boldsymbol{\gamma}_1 \cdots \boldsymbol{\gamma}_q] \in \mathbb{R}^{r \times q}$, we have

$$\mathbf{0}_k = (\mathbf{y}_i^\top \otimes \mathbf{Z}_i)\boldsymbol{\gamma} = \sum_{j=1}^q y_{ij} \mathbf{Z}_i \boldsymbol{\gamma}_j = \mathbf{Z}_i \left(\sum_{j=1}^q y_{ij} \boldsymbol{\gamma}_j \right) = \mathbf{Z}_i \boldsymbol{\Gamma} \mathbf{y}_i,$$

where y_{ij} is the j th entry of \mathbf{y}_i . Since each \mathbf{Z}_i has full column rank, it must be the case that $\boldsymbol{\Gamma} \mathbf{y}_i = \mathbf{0}_r$ for each $i = 1, \dots, s$, which in turn implies $\mathbf{Y}\boldsymbol{\Gamma}^\top = (\boldsymbol{\Gamma} \mathbf{Y}^\top)^\top = \mathbf{0}_{s \times r}$. But \mathbf{Y} having full column rank implies $\boldsymbol{\Gamma} = \mathbf{0}_{r \times q}$, hence $\boldsymbol{\gamma} = \mathbf{0}_{qr}$. Thus, the columns of \mathbf{W} are linearly independent, so \mathbf{W} has full column rank. \square

Appendix C. General Construction for PDE Systems. We provide here a general construction for the affine Operator Inference problem for systems of PDEs to learn reduced-order models of the form (3.3a)–(3.3b). The problem decouples into d instances of (2.13), that is,

$$\min_{\hat{\mathbf{O}}_\ell} \left\| \mathbf{D}_\ell \hat{\mathbf{O}}_\ell^\top - \mathbf{R}_\ell^\top \right\|_F^2, \quad \ell = 1, \dots, d,$$

where

$$\begin{aligned} \mathbf{D}_\ell &= \left[\mathbf{D}_{c_\ell} \mid \mathbf{D}_{A_{\ell,1}} \cdots \mathbf{D}_{A_{\ell,d}} \mid \mathbf{D}_{H_{\ell,11}} \quad \mathbf{D}_{H_{\ell,12}} \cdots \mathbf{D}_{H_{\ell,dd}} \right], \\ \hat{\mathbf{O}}_\ell &= \left[\hat{\mathbf{C}}_\ell \mid \hat{\mathbf{A}}_{\ell,1} \cdots \hat{\mathbf{A}}_{\ell,d} \mid \hat{\mathbf{H}}_{\ell,11} \quad \hat{\mathbf{H}}_{\ell,12} \cdots \hat{\mathbf{H}}_{\ell,dd} \right], \\ \mathbf{R}_\ell &= \left[\hat{\mathbf{U}}_\ell(\mu_1) \quad \cdots \quad \hat{\mathbf{U}}_\ell(\mu_s) \right] \in \mathbb{R}^{r_\ell \times sK}, \end{aligned}$$

with

$$\begin{aligned} \mathbf{D}_{c_\ell} &= \begin{bmatrix} \boldsymbol{\theta}_{c_\ell}(\mu_1) \otimes \mathbf{1}_K \\ \vdots \\ \boldsymbol{\theta}_{c_\ell}(\mu_s) \otimes \mathbf{1}_K \end{bmatrix} \in \mathbb{R}^{K \times q_{c_\ell}}, \\ \mathbf{D}_{A_{\ell,m}} &= \begin{bmatrix} \boldsymbol{\theta}_{A_{\ell,m}}(\mu_1) \otimes \hat{\mathbf{U}}_m(\mu_1)^\top \\ \vdots \\ \boldsymbol{\theta}_{A_{\ell,m}}(\mu_s) \otimes \hat{\mathbf{U}}_m(\mu_s)^\top \end{bmatrix} \in \mathbb{R}^{K \times q_{A_{\ell,m}} r_m}, \\ \mathbf{D}_{H_{\ell,mm}} &= \begin{bmatrix} \boldsymbol{\theta}_{H_{\ell,mm}}(\mu_1) \otimes \left(\hat{\mathbf{U}}_m(\mu_1) \hat{\odot} \hat{\mathbf{U}}_m(\mu_1) \right)^\top \\ \vdots \\ \boldsymbol{\theta}_{H_{\ell,mm}}(\mu_s) \otimes \left(\hat{\mathbf{U}}_m(\mu_s) \hat{\odot} \hat{\mathbf{U}}_m(\mu_s) \right)^\top \end{bmatrix} \in \mathbb{R}^{K \times q_{H_{\ell,mm}} r_m (r_m + 1)/2}, \\ \mathbf{D}_{H_{\ell,mn}} &= \begin{bmatrix} \boldsymbol{\theta}_{H_{\ell,mn}}(\mu_1) \otimes \left(\hat{\mathbf{U}}_m(\mu_1) \odot \hat{\mathbf{U}}_n(\mu_1) \right)^\top \\ \vdots \\ \boldsymbol{\theta}_{H_{\ell,mn}}(\mu_s) \otimes \left(\hat{\mathbf{U}}_m(\mu_s) \odot \hat{\mathbf{U}}_n(\mu_s) \right)^\top \end{bmatrix} \in \mathbb{R}^{K \times q_{H_{\ell,mn}} r_m r_n}, \quad (n \neq m) \\ \hat{\mathbf{C}}_\ell &= \left[\hat{\mathbf{c}}_\ell^{(1)} \quad \cdots \quad \hat{\mathbf{c}}_\ell^{(q_{c_\ell})} \right] \in \mathbb{R}^{r_\ell \times q_{c_\ell}} \\ \hat{\mathbf{A}}_{\ell,m} &= \left[\hat{\mathbf{A}}_{\ell,m}^{(1)} \quad \cdots \quad \hat{\mathbf{A}}_{\ell,m}^{(q_{A_{\ell,m}})} \right] \in \mathbb{R}^{r_\ell \times q_{A_{\ell,m}} r_m} \\ \hat{\mathbf{H}}_{\ell,mm} &= \left[\hat{\mathbf{H}}_{\ell,mm}^{(1)} \quad \cdots \quad \hat{\mathbf{H}}_{\ell,mm}^{(q_{H_{\ell,mm}})} \right] \in \mathbb{R}^{r_\ell \times q_{H_{\ell,mm}} r_m (r_m + 1)/2} \\ \hat{\mathbf{H}}_{\ell,mn} &= \left[\hat{\mathbf{H}}_{\ell,mn}^{(1)} \quad \cdots \quad \hat{\mathbf{H}}_{\ell,mn}^{(q_{H_{\ell,mn}})} \right] \in \mathbb{R}^{r_\ell \times q_{H_{\ell,mn}} r_m r_n}, \quad (m \neq n) \\ \hat{\mathbf{U}}_\ell(\mu_i) &= \left[\hat{\mathbf{u}}_\ell(t_1; \mu_i) \quad \cdots \quad \hat{\mathbf{u}}_\ell(t_K; \mu_i) \right] \in \mathbb{R}^{r_\ell \times K}, \\ \dot{\hat{\mathbf{U}}}_\ell(\mu_i) &= \left[\left. \frac{d}{dt} \hat{\mathbf{u}}_\ell(t; \mu_i) \right|_{t=t_1} \quad \cdots \quad \left. \frac{d}{dt} \hat{\mathbf{u}}_\ell(t; \mu_i) \right|_{t=t_K} \right] \in \mathbb{R}^{r_\ell \times K}, \\ \boldsymbol{\theta}_{c_\ell}(\mu_i) &= \left[\theta_{c_\ell}^{(1)}(\mu_i) \quad \cdots \quad \theta_{c_\ell}^{(q_{c_\ell})}(\mu_i) \right] \in \mathbb{R}^{1 \times q_{c_\ell}}, \\ \boldsymbol{\theta}_{A_{\ell,m}}(\mu_i) &= \left[\theta_{A_{\ell,m}}^{(1)}(\mu_i) \quad \cdots \quad \theta_{A_{\ell,m}}^{(q_{A_{\ell,m}})}(\mu_i) \right] \in \mathbb{R}^{1 \times q_{A_{\ell,m}}}, \end{aligned}$$

$$\boldsymbol{\theta}_{H_{\ell,mn}}(\boldsymbol{\mu}_i) = \left[\theta_{H_{\ell,mn}}^{(1)}(\boldsymbol{\mu}_i) \quad \cdots \quad \theta_{H_{\ell,mn}}^{(q_{H_{\ell,mn}})}(\boldsymbol{\mu}_i) \right] \in \mathbb{R}^{1 \times q_{H_{\ell,mn}}}.$$

In practice, this construction is often sparse due to the limited number of terms in the governing PDE.

REFERENCES

- [1] V. ALGAZI AND D. SAKRISON, *On the optimality of the Karhunen-Loève expansion*, IEEE Transactions on Information Theory, 15 (1969), pp. 319–321.
- [2] A. C. ANTOUNAS AND B. D. Q. ANDERSON, *On the scalar rational interpolation problem*, IMA Journal of Mathematical Control and Information, 3 (1986), pp. 61–88.
- [3] A. C. ANTOUNAS, I. V. GOSEA, AND A. C. IONITA, *Model reduction of bilinear systems in the Loewner framework*, SIAM Journal on Scientific Computing, 38 (2016), pp. B889–B916.
- [4] H. ARBABI, J. E. BUNDER, G. SAMAEE, A. J. ROBERTS, AND I. G. KEVREKIDIS, *Linking machine learning with multiscale numerics: Data-driven discovery of homogenized equations*, JOM, 72 (2020), pp. 4444–4457.
- [5] S. M. BAER AND T. ERNEUX, *Singular Hopf bifurcation to relaxation oscillations*, SIAM Journal on Applied Mathematics, 46 (1986), pp. 721–739.
- [6] M. BARRAULT, Y. MADAY, N. C. NGUYEN, AND A. T. PATERA, *An ‘empirical interpolation’ method: application to efficient reduced-basis discretization of partial differential equations*, Comptes Rendus Mathématique, 339 (2004), pp. 667–672.
- [7] P. BENNER, P. GOYAL, B. KRAMER, B. PEHERSTORFER, AND K. WILLCOX, *Operator inference for non-intrusive model reduction of systems with non-polynomial nonlinear terms*, Computer Methods in Applied Mechanics and Engineering, 372 (2020), p. 113433.
- [8] P. BENNER, S. GUGERCIN, AND K. WILLCOX, *A survey of projection-based model reduction methods for parametric dynamical systems*, SIAM Review, 57 (2015), pp. 483–531.
- [9] G. BERKOOZ, P. HOLMES, AND J. LUMLEY, *The proper orthogonal decomposition in the analysis of turbulent flows*, Annual Review of Fluid Mechanics, 25 (1993), pp. 539–575.
- [10] K. BHATTACHARYA, B. HOSSEINI, N. B. KOVACHKI, AND A. M. STUART, *Model reduction and neural networks for parametric PDEs*, The SMAI Journal of Computational Mathematics, 7 (2021), pp. 121–157.
- [11] Å. BJÖRCK, *Numerical Methods for Least Squares Problems*, SIAM, Philadelphia, PA, 1996.
- [12] S. L. BRUNTON, J. L. PROCTOR, AND J. N. KUTZ, *Discovering governing equations from data by sparse identification of nonlinear dynamical systems*, Proceedings of the National Academy of Sciences, 113 (2016), pp. 3932–3937.
- [13] S. L. BRUNTON, J. L. PROCTOR, J. H. TU, AND J. N. KUTZ, *Compressed sensing and dynamic mode decomposition*, Journal of Computational Dynamics, 2 (2015), p. 165.
- [14] S. CHATURANTABUT AND D. C. SORENSEN, *Nonlinear model reduction via discrete empirical interpolation*, SIAM Journal on Scientific Computing, 32 (2010), pp. 2737–2764.
- [15] C. F. CURTISS AND J. O. HIRSCHFELDER, *Integration of stiff equations*, Proceedings of the National Academy of Sciences, 38 (1952), pp. 235–243.
- [16] U. FASEL, J. N. KUTZ, B. W. BRUNTON, AND S. L. BRUNTON, *Ensemble-sindy: Robust sparse model discovery in the low-data, high-noise limit, with active learning and control*, Proceedings of the Royal Society A, 478 (2022), p. 20210904.
- [17] R. FITZHUGH, *Impulses and physiological states in theoretical models of nerve membrane*, Biophysical Journal, 1 (1961), pp. 445–466.
- [18] O. GHATTAS AND K. WILLCOX, *Learning physics-based models from data: perspectives from inverse problems and model reduction*, Acta Numerica, 30 (2021), pp. 445–554.
- [19] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, vol. 3, John Hopkins University Press, Baltimore, MD, 2013.
- [20] W. R. GRAHAM, J. PERAIRE, AND K. Y. TANG, *Optimal control of vortex shedding using low-order models. Part I—Open-loop model development*, International Journal for Numerical Methods in Engineering, 44 (1999), pp. 945–972.
- [21] M. GREPL AND A. PATERA, *A posteriori error bounds for reduced-basis approximations of parametrized parabolic partial differential equations*, ESAIM: Mathematical Modelling and Numerical Analysis, 39 (2005), pp. 157–181.
- [22] P. C. HANSEN, *Analysis of discrete ill-posed problems by means of the L-curve*, SIAM Review, 34 (1992), pp. 561–580.
- [23] A. C. IONITA AND A. C. ANTOUNAS, *Data-driven parametrized model reduction in the Loewner framework*, SIAM Journal on Scientific Computing, 36 (2014), pp. A984–A1007.
- [24] P. JAIN, S. MCQUARRIE, AND B. KRAMER, *Performance comparison of data-driven reduced*

- models for a single-injector combustion process*, in AIAA Propulsion and Energy 2021 Forum, Virtual Event, 2021. Paper AIAA-2021-3633.
- [25] C. KHATRI AND C. R. RAO, *Solutions to some functional equations and their applications to characterization of probability distributions*, Sankhyā: The Indian Journal of Statistics, Series A, (1968), pp. 167–180.
- [26] P. KHODABAKHSHI AND K. E. WILLCOX, *Non-intrusive data-driven model reduction for differential algebraic equations derived from lifting transformations*, Oden Institute Report 21-08, University of Texas at Austin, 2021.
- [27] T. G. KOLDA AND B. W. BADER, *Tensor decompositions and applications*, SIAM Review, 51 (2009), pp. 455–500.
- [28] K. LEE AND K. T. CARLBERG, *Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders*, Journal of Computational Physics, 404 (2020), p. 108973.
- [29] S. LEE, M. KOOSHBAGHI, K. SPILLOTIS, C. I. SIETTOS, AND I. G. KEVREKIDIS, *Coarse-scale PDEs from fine-scale observations via machine learning*, Chaos: An Interdisciplinary Journal of Nonlinear Science, 30 (2020), p. 013141.
- [30] C. LIEBERMAN, K. WILLCOX, AND O. GHATTAS, *Parameter and state model reduction for large-scale statistical inverse problems*, SIAM Journal on Scientific Computing, 32 (2010), pp. 2523–2542.
- [31] J. LUMLEY, *The structures of inhomogeneous turbulent flow*, Atmospheric Turbulence and Radio Wave Propagation, (1967), pp. 166–178.
- [32] S. A. MCQUARRIE, C. HUANG, AND K. E. WILLCOX, *Data-driven reduced-order models via regularised operator inference for a single-injector combustion process*, Journal of the Royal Society of New Zealand, 51 (2021), pp. 194–211.
- [33] J. NAGUMO, S. ARIMOTO, AND S. YOSHIKAWA, *An active pulse transmission line simulating nerve axon*, Proceedings of the IRE, 50 (1962), pp. 2061–2070.
- [34] J. A. NELDER AND R. MEAD, *A simplex method for function minimization*, The Computer Journal, 7 (1965), pp. 308–313.
- [35] G. ONGIE, D. PIMENTEL-ALARCÓN, L. BALZANO, R. WILLET, AND R. D. NOWAK, *Tensor methods for nonlinear matrix completion*, SIAM Journal on Mathematics of Data Science, 3 (2021), pp. 253–279.
- [36] B. PEHERSTORFER, *Sampling low-dimensional Markovian dynamics for pre-asymptotically recovering reduced models from data with operator inference*, SIAM Journal on Scientific Computing, 42 (2020), pp. A3489–A3515.
- [37] B. PEHERSTORFER AND K. WILLCOX, *Data-driven operator inference for nonintrusive projection-based model reduction*, Computer Methods in Applied Mechanics and Engineering, 306 (2016), pp. 196–215.
- [38] E. QIAN, I.-G. FARCAŞ, AND K. WILLCOX, *Reduced operator inference for nonlinear partial differential equations*, SIAM Journal on Scientific Computing, 44 (2022), pp. A1934–A1959.
- [39] E. QIAN, B. KRAMER, B. PEHERSTORFER, AND K. WILLCOX, *Lift & Learn: Physics-informed machine learning for large-scale nonlinear dynamical systems.*, Physica D: Nonlinear Phenomena, 406 (2020), p. 132401.
- [40] M. QUADE, M. ABEL, J. NATHAN KUTZ, AND S. L. BRUNTON, *Sparse identification of nonlinear dynamics for rapid model recovery*, Chaos: An Interdisciplinary Journal of Nonlinear Science, 28 (2018), p. 063116.
- [41] G. ROZZA, D. HUYNH, AND A. PATERA, *Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations: application to transport and continuum mechanics*, Archives of Computational Methods in Engineering, 15 (2008), pp. 229–275.
- [42] S. RUDY, A. ALLA, S. L. BRUNTON, AND J. N. KUTZ, *Data-driven identification of parametric partial differential equations*, SIAM Journal on Applied Dynamical Systems, 18 (2019), pp. 643–660.
- [43] H. SCHAEFFER, *Learning partial differential equations via data discovery and sparse optimization*, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences, 473 (2017), p. 20160446.
- [44] P. J. SCHMID, *Dynamic mode decomposition of numerical and experimental data*, Journal of Fluid Mechanics, 656 (2010), pp. 5–28.
- [45] H. SHARMA, Z. WANG, AND B. KRAMER, *Hamiltonian operator inference: Physics-preserving learning of reduced-order models for canonical Hamiltonian systems*, Physica D: Nonlinear Phenomena, 431 (2022), p. 133122.
- [46] L. SIROVICH, *Turbulence and the dynamics of coherent structures. I. Coherent structures*, Quarterly of Applied Mathematics, 45 (1987), pp. 561–571.

- [47] G. SÖDERLIND, *Automatic control and adaptive time-stepping*, Numerical Algorithms, 31 (2002), pp. 281–310.
- [48] R. SWISCHUK, B. KRAMER, C. HUANG, AND K. WILLCOX, *Learning physics-based reduced-order models for a single-injector combustion process*, AIAA Journal, 58 (2020), pp. 2658–2672.
- [49] M. TEZZELE, N. DEMO, G. STABILE, A. MOLA, AND G. ROZZA, *Enhancing CFD predictions in shape design problems by model and parameter space reduction*, Advanced Modeling and Simulation in Engineering Sciences, 7 (2020), pp. 1–19.
- [50] A. N. TIKHONOV AND V. Y. ARSEININ, *Solutions of Ill-Posed Inverse Problems*, Wiley, New York, NY, 1977.
- [51] L. N. TREFETHEN AND D. BAU III, *Numerical Linear Algebra*, SIAM, Philadelphia, PA, 1997.
- [52] W. I. T. UY AND B. PEHERSTORFER, *Probabilistic error estimation for non-intrusive reduced models learned from data of systems governed by linear parabolic partial differential equations*, ESAIM: Mathematical Modelling and Numerical Analysis, 55 (2021), pp. 735–761.
- [53] C. F. VAN LOAN, *The ubiquitous Kronecker product*, Journal of Computational and Applied Mathematics, 123 (2000), pp. 85–100.
- [54] K. VEROY AND A. PATERA, *Certified real-time solution of the parametrized steady incompressible Navier-Stokes equations: Rigorous reduced-basis a posteriori error bounds*, International Journal for Numerical Methods in Fluids, 47 (2005), pp. 773–788.
- [55] K. VEROY, C. PRUD'HOMME, D. ROVAS, AND A. PATERA, *A posteriori error bounds for reduced-basis approximation of parametrized noncoercive and nonlinear elliptic partial differential equations*, in Proceedings of the 16th AIAA Computational Fluid Dynamics Conference, Orlando, FL, 2003. Paper AIAA-2003-3847.
- [56] P. VIRTANEN, R. GOMMERS, T. E. OLIPHANT, M. HABERLAND, T. REDDY, D. COURNAPEAU, E. BUROVSKI, P. PETERSON, W. WECKESSER, J. BRIGHT, S. J. VAN DER WALT, M. BRETT, J. WILSON, K. JARROD MILLMAN, N. MAYOROV, A. R. J. NELSON, E. JONES, R. KERN, E. LARSON, C. CAREY, Í. POLAT, Y. FENG, E. W. MOORE, J. VAND ERPLAS, D. LAXALDE, J. PERKTOLD, R. CIMRMAN, I. HENRIKSEN, E. A. QUINTERO, C. R. HARRIS, A. M. ARCHIBALD, A. H. RIBEIRO, F. PEDREGOSA, P. VAN MULBREGT, AND S. . CONTRIBUTORS, *SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python*, Nature Methods, 17 (2020), pp. 261–272.
- [57] G. WANNER AND E. HAIRER, *Solving ordinary differential equations II: Stiff and differential-algebraic problems*, vol. 375, Springer Berlin Heidelberg New York, 1996.
- [58] S. YILDIZ, P. GOYAL, P. BENNER, AND B. KARASÖZEN, *Learning reduced-order dynamics for parametrized shallow water equations from data*, International Journal for Numerical Methods in Fluids, 93 (2021), pp. 2803–2821.