SANDIA REPORT
SAND2024-11571
Printed September 2024

Sandia National Laboratories

Learning Operators for Structure-Informed Surrogate Models

Anthony Gruber

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185
Livermore, California 94550

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology & Engineering Solutions of Sandia, LLC.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@osti.gov
Online ordering: http://www.osti.gov/scitech

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Road
Alexandria, VA 22312

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: https://classic.ntis.gov/help/order-methods

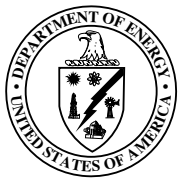

ABSTRACT

This report summarizes the work performed under the author's two-year John von Neumann LDRD project, which involves the non-intrusive surrogate modeling of dynamical systems with remarkable structural properties. After a brief introduction to the topic, technical accomplishments and project metrics are reviewed including peer-reviewed publications, software releases, external presentations and colloquia, as well as organized conference sessions and minisymposia. The report concludes with a summary of ongoing projects and collaborations which utilize the results of this work.

This page intentionally left blank.

CONTENTS

This page intentionally left blank.

LIST OF FIGURES

1. INTRODUCTION

At the time this John von Neumann fellowship LDRD was funded, the DOE's Modeling and Simulation (ModSim) infrastructure lacked cost-effective surrogates which incorporate known qualitative, quantitative, and topological structure present in the high-fidelity models governing many applications of interest. Conversely, it was recognized that having direct access to mathematical operators which represent the fundamental physical and mathematical correspondences underlying empirical measurements allows for lightweight models which are accurate, explainable, and generalizable to unseen scenarios. This motivated the present work, which aimed in the following directions:

- Development of a rigorous mathematical foundation for the surrogate modeling of systems with Hamiltonian, metriplectic, and Hodge-de Rham structure.
- Development of practical structure-informed surrogates which are (1) accurate, (2) stable, (3) convergent, (4) efficient, and (5) non-intrusive.
- Deployment of the mentioned surrogates in situations which demonstrate their inherent explainability and superior generalization to unseen timescales and parameter ranges.

My project succeeded in addressing these criteria in a way that substantially improved the state of the art in non-intrusive and structure-informed surrogate modeling. As this project was beginning, surrogate modeling suffered from the following deficiencies:

- An over-reliance on "general purpose" methods such as the data-driven training of multilayer perceptrons (MLPs), which lack both explainability and theoretical guarantees on out-of-distribution inputs.
- Poor-performing structure-informed techniques characterized by the weak enforcement of physical constraints or the requirement of intrusive access to the high-fidelity model, often leading to unpredictable and untrustworthy results in practice.
- Hard-constrained methods for incorporating Hamiltonian and metriplectic structure which are insufficiently general, require prohibitive code-intrusion, and are not guaranteed to approximate arbitrary systems within their structural class.

To address these troublesome difficulties, my project concentrated on three important mathematical structures of widespread interest frequently encountered in applications to climate modeling, electromagnetism, and thermomechanics: Hamiltonian structure, metriplectic structure, and Hodge-de Rham structure. Adapting techniques from varied areas including projection-based model reduction, machine learning, convex optimization, finite element analysis, exterior algebra, and tensor calculus, the project was able to:

- Develop the first method for Hamiltonian operator inference (OpInf) which is applicable to both canonical and noncanonical Hamiltonian systems and places no additional requirements on the form of the reduced basis, producing a non-intrusive and provably Hamiltonian surrogate from only snapshots of the state variable [1].
- Remedy the general lack of variational consistency in projection-based Hamiltonian surrogate modeling, leading to intrusive and non-intrusive models of canonical Hamiltonian systems which are provably bounded in error and orders of magnitude more performant than the current state-of-the-art at the same computational cost [2].
- Publish open-source software implementing the non-intrusive Hamiltonian surrogate models in [1, 2].
- Produce the first machine-learned metriplectic surrogate model which is provably universally approximating and admits a bound on generalization error [5].
- Extend the research on bracket-based surrogate models to graph-structured data, producing machine-learnable structure-preserving surrogates on graphs respecting Hamiltonian, gradient flow, double bracket, and metriplectic dynamics [7].
- Formulate and develop the first cohomology-aware reduced-order model based on proper orthogonal decomposition, enabling a Hodge-de Rham complex preserving discretization of standard- and mixed- form partial differential equations (PDEs) which correctly approximates harmonically dominant solutions [6].

Other notable research accomplishments include a structure-preserving Hamiltonian OpInf which automatically captures parametric dependence in the Hamiltonian functional [3], as well as a non-intrusive method for the hyper-reduction of nonlinear reduced models which respects variational consistency [8].

The next chapter provides a brief description of some technical highlights from this project. For further details, the reader is referred to the publications mentioned in Section 3.

2. TECHNICAL HIGHLIGHTS

This chapter briefly highlights some impactful research goals attained during this project.

2.1. Hamiltonian Operator Inference

A primary focus of this project was on the design of non-intrusive model reduction algorithms for Hamiltonian systems based on ideas from operator inference (OpInf) [9]. Recall that the Hamiltonian formalism provides a mechanical framework encompassing a wide variety of conservative dynamical systems which arise from a least-action variational principle. Given a phase space P of solutions, it reduces the problem of understanding a complicated dynamical system to the simpler problem of understanding a scalar-valued function $H : P \to \mathbb{R}$, called the Hamiltonian, along with a skew-symmetric Poisson bracket $\{\cdot,\cdot\} : P \times P \to P$ encoding the dynamical structure. More formally, given a state vector x with components in $P \subset \mathbb{R}^N$, it follows that any Hamiltonian system can be written in the form

$$\dot{x} = \{x, H(x)\} = L(x)\nabla H(x), \tag{2.1}$$

for some Hamiltonian $H : \mathbb{R}^N \to \mathbb{R}$ and some (potentially degenerate) Poisson matrix $L : \mathbb{R}^N \to \mathbb{R}^{N\times N}$, $L^\intercal = -L$ which is antisymmetric and satisfies the Jacobi identity[1]. Note that this expression immediately shows that the Hamiltonian is conserved along trajectories of the system, since its derivative in time $\dot{H} = \{H, H\} = 0$ vanishes by antisymmetry.

Hamiltonian systems are highly relevant as dynamical models and come in two variants, known respectively as *canonical* and *noncanonical*. The distinction lies fundamentally in their associated *Poisson structure*, which is symplectic in the case of canonical systems and degenerate otherwise. In particular, when the phase space $P \subset \mathbb{R}^{2M}$ separates cleanly into position and momentum variables (q, p), the matrix field $L = J = [0\ I; -I\ 0]$ is the usual (constant) symplectic matrix of dimension $2M$, and the system (2.1) becomes

$$\dot{x} = \begin{pmatrix}\dot{q}\\ \dot{p}\end{pmatrix} = \begin{pmatrix}0 & I\\ -I & 0\end{pmatrix}\begin{pmatrix}\partial_q H\\ \partial_p H\end{pmatrix} = J\nabla H(x),$$

which are the canonical Hamiltonian equations of motion. Conversely, it is often the case that the Poisson matrix field L contains functions in its kernel (e.g., mass, momentum), known as *Casimir invariants*, which disrupt this canonical structure. This has consequences for model reduction: not all techniques for building effective surrogates in the canonical case carry over to the noncanonical one. This project investigated both classes of systems, culminating in the works [1, 2] which are now highlighted.

[1] This is a technical condition involving L and its partial derivatives, which reads component-wise as $\sum_{\ell=1}^{N}\left(L_{il}L_{jk,l} + L_{jl}L_{ki,l} + L_{kl}L_{ij,l}\right) = 0,\ 1 \le i, j, k \le N$.

2.1.1. Noncanonical Hamiltonian operator inference

We developed the first non-intrusive OpInf model for noncanonical Hamiltonian systems in [1], which provably respects Hamiltonian structure and requires only snapshot data from the full-order model (FOM), often a system of semi-discrete partial differential equations (PDEs). Recall that OpInf is a projection-based technique for building reduced-order models (ROMs) in a solution space of coefficients corresponding to some reduced basis $U \in \mathbb{R}^{N\times n}$, typically generated by the method of proper orthogonal decomposition (POD), where $N \gg n$ denotes the dimension of the FOM solution. Assuming snapshot data $X \in \mathbb{R}^{N\times\tau}$ for the FOM system is given and a conserved quantity $H : \mathbb{R}^N \to \mathbb{R}$ has been identified (or postulated), the noncanonical Hamiltonian OpInf in [1] seeks a solution to the convex problem

$$\operatorname*{argmin}_{\bar{L}\in\mathbb{R}^{n\times n}} \left|\hat{X}_t - \bar{L}\hat{\nabla}H(X)\right|^2, \quad \text{s.t.} \quad \bar{L}^\intercal = -\bar{L},$$

where $\hat{X} = U^\intercal X$, $\hat{\nabla}H(X) = U^\intercal \nabla H(X)$, and $\hat{X}_t \approx \partial_t \hat{X}$ denotes an approximation to the temporal derivative of X. Key to this approach is strong enforcement of the antisymmetry condition $L^\intercal = -L$, which guarantees that the learned operator $\bar{L}$ defining the reduced Poisson structure leads to a valid Hamiltonian reduced-order model (ROM) for the coefficients of the approximate solution $\tilde{x} = U\hat{x}$ in the basis U, given as

$$\dot{\hat{x}} = \bar{L}\nabla\hat{H}(\hat{x}) = \bar{L}U^\intercal \nabla H(U\hat{x}).$$

It can be shown that the proposed approach is Hamiltonian-conserving and provably convergent to the corresponding intrusive Hamiltonian ROM based on Galerkin projection and the approximation $\hat{L} = U^\intercal LU$. An example in comparison to other methods is given in Figure 2-1, where accuracy and stability benefits can be observed. This work has paved the way for future investigations into non-intrusive ROMs for noncanonical Hamiltonian systems of larger scale, such as Maxwell's equations for plasma or the shallow water equations for oceanic modeling.

2.1.2. Variationally consistent Hamiltonian operator inference

More techniques for projection-based model reduction are available in the canonical Hamiltonian case. Before our work in [2], state-of-the-art in this area relied on a specially designed basis called the proper symplectic decomposition [10] (PSD), whose defining property is that $U^\intercal J = JU^\intercal$, implying that canonical Hamiltonian systems project to canonical Hamiltonian systems. While this is a highly desirable property that guarantees stability and interpretability, the problem of computing a good PSD basis is nontrivial and only limited algorithms exist for accomplishing this. Moreover, it is frequently the case that these algorithms produce PSD bases which are not expressive enough to yield a good ROM, as can be seen from the projection errors in Figure 2-2.

In view of these issues, other methods for Hamiltonian model reduction have been developed which accommodate any desired reduced basis, but these suffer from a significant drawback: they are not variationally consistent. This means that the Galerkin projection of the true solution does not satisfy the discrete equations, which can dramatically affect the accuracy of these ROMs. The work

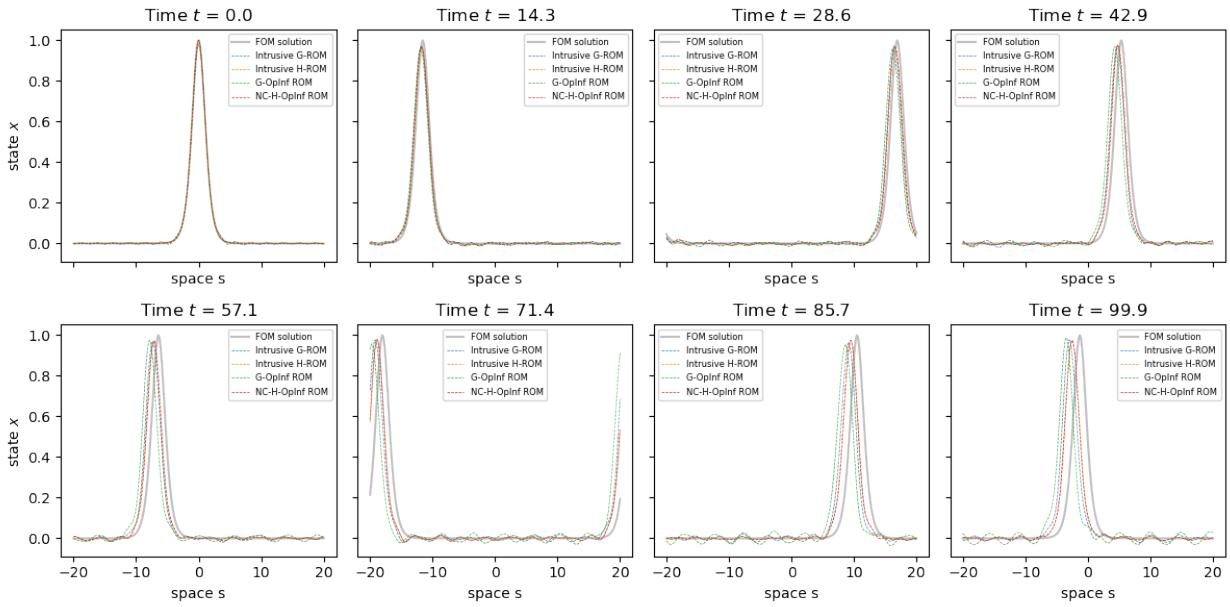

Figure 2-1. Snapshots in time corresponding to the FOM and ROM solutions to a predictive KdV equation example with $N = 1024$ and $n = 32$ modes. Note the lack of spurious oscillations in the H-OpInf ROMs from [1].

in [2] remedies this defect by proposing a new strategy based on Petrov-Galerkin projection (see Figure 2-3), which guarantees a variationally consistent, conservative, and symplectic Hamiltonian ROM. We show that this strategy can be applied in both intrusive and non-intrusive OpInf contexts, and prove an error estimate showing the trade-off between the expressivity of the reduced basis U and the deviation of the ROM from canonical Hamiltonian form. An example of the improvement gained from our variationally consistent approach is pictured in Figure 4-1 of Section 4 using a realistic example from Sandia's Albany code [11] with the material properties of steel. Remarkably, our variationally consistent approach is at present the only Hamiltonian method for projection-based ROM which can reliably predict the dynamics of an elastic bracket outside of the time interval where it has been trained.

2.1.3. Tensor parametric Hamiltonian operator inference

It is common for FOM systems of interest to rely on a number of parameters which control the behavior of their solutions. In this case, the operators which must be learned through OpInf carry around this parametric dependence, complicating the learning problem and preventing the application of standard methods for non-intrusive ROM. This has left a large gap in the research literature, as practitioners are left with only cumbersome or ad hoc methods for incorporating parametric dependence into OpInf surrogates.

To address this issue, our work in [3, 4] shows that, when the parametric dependence in the FOM is linear (or generalized linear), then a matrix operator $M(\mu) = T\mu \in \mathbb{R}^{N\times N}$ with dependence on the parameter vector $\mu \in \mathbb{R}^p$ can be written as the contraction of this vector against a *constant* order-3 tensor $T \in \mathbb{R}^{N\times N\times p}$. Therefore, if an effective (e.g., convex) learning problem can be formulated for

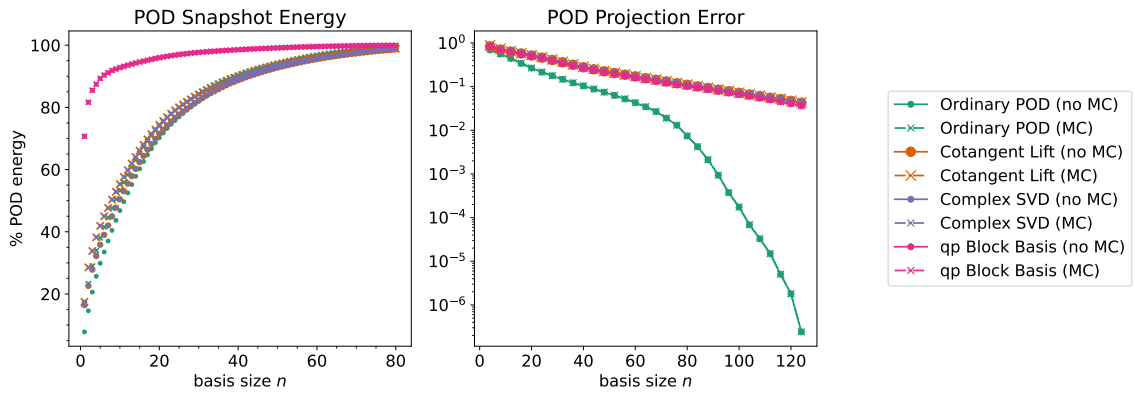

Figure 2-2. POD basis snapshot energies and projection errors corresponding to the plate example in [2]. Note that the PSD algorithms "Cotangent Lift" and "Complex SVD" are unable to produce an informative basis for the solution space.

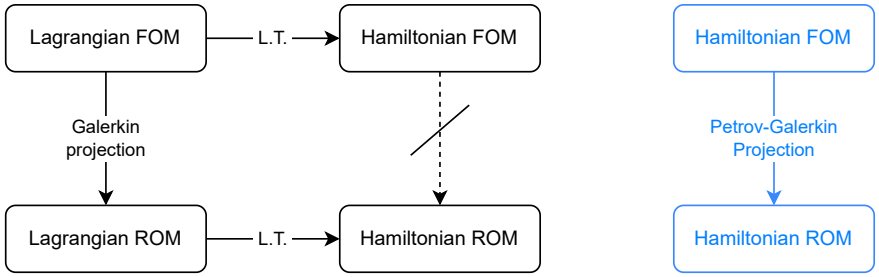

Figure 2-3. A diagram illustrating the behavior of Hamiltonian/Lagrangian FOMs under Galerkin projection. While (most) canonical Hamiltonian and Lagrangian systems are equivalent under the Legendre transformation (L.T.) at both the FOM and ROM levels, this equivalence does not commute with Galerkin projection. Conversely, the variationally consistent ROM presented in [2] projects Hamiltonian FOMs to Hamiltonian ROMs in every case.

the tensor T, then parametric dependence is captured automatically with standard OpInf principles. This represents a major step forward, as all previous non-intrusive ROM technology is enabled immediately and without additional modification.

To describe how this works more precisely, note that the intrusive Galerkin ROM with reduced basis $U \in \mathbb{R}^{N\times n}$ corresponding to a full-order model of the form $\dot{x} = M(\mu)x$ is given by

$$\dot{\hat{x}} = U^\intercal M(\mu) U\hat{x} = \left(\hat{T}\mu\right)\hat{x},$$

where the reduced tensor $\hat{T} \in \mathbb{R}^{n\times n\times p}$ is defined via $\hat{T}\mu = U^\intercal(T\mu)U$. In the OpInf context, this implies the goal of learning $\bar{T} \approx \hat{T}$ by solving the optimization problem

$$\underset{\bar{T}}{\operatorname{argmin}} \sum_{s=0}^{N_s} \left|\hat{X}_{s,t} - \left(\bar{T}\mu_s\right)\hat{X}_s\right|^2,$$

where the "hat" notation is as in Section 2.1.1 and the subscript s indexes the parameter snapshots. We show in [3, 4] that this problem is equivalent to a linear system in n^2p variables, which is readily solved with standard numerical techniques. It follows that the tensor parametric OpInf ROM

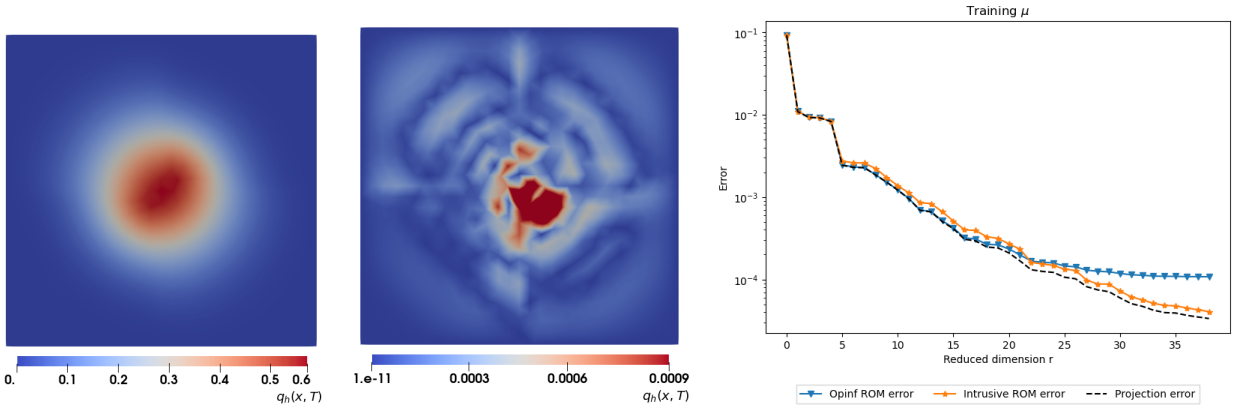

Figure 2-4. Solution to the heat equation (left) at a particular instant in time with different thermal conductivity in each quadrant. Absolute error plot (middle) of the solution at this time point and parameter value. A scaling study (right) showing decay of the error with the addition of basis modes.

$\dot{\hat{x}} = (\bar{T}\mu)\hat{x}$ for the coefficients $\hat{x}$ in the Galerkin approximation $\tilde{x} = U\hat{x}$ is a useful approximation to the intrusive ROM which requires no access to the FOM code and expresses the correct parametric dependence by construction.

The effectiveness of this procedure can be shown by considering the following initial boundary value problem for the heat equation on a domain $\Omega \subset \mathbb{R}^2$,

$$\begin{cases} \partial_t u(x,t) = c(\mu)\Delta u(x,t), & x \in \Omega \times (0,T], \\ u(x,0) = u_0(x), & x \in \Omega, \\ u(x,t) = 0, & x \in \partial\Omega \times (0,T], \end{cases} \tag{2.2}$$

where $c(\mu) : \mathbb{R}^p \to \mathbb{R}$ is the parameterized thermal conductivity

$$c(\mu) = \mu_1 1_{\Omega_1} + \mu_2 1_{\Omega_2} + \ldots + \mu_p 1_{\Omega_p}, \tag{2.3}$$

which is piecewise-constant on $\Omega = \bigcup_{i=1}^{p} \Omega_i$ since 1_{Ω_i} denotes the indicator function. In this case, the discrete FOM system generated by, e.g., the finite element method is of the form discussed above, and the tensor parametric technology applies straightforwardly. An example of the results in this case is shown in Figure 2-4, where it can be seen that the results are in good agreement with the intrusive parametric Galerkin ROM.

In addition, it is shown in [3, 4] that our tensor parametric OpInf can be adapted to preserve Hamiltonian structure as well. Following the techniques established in [1, 2], we show that it is enough to solve an inference problem of the form

$$\underset{\bar{T}}{\operatorname{argmin}} \frac{1}{2} \sum_{s=1}^{N_s} \left| \hat{C}_s - \hat{A}_s \left(\bar{T}\mu_s\right) \hat{B}_s \right|^2 \quad \text{s.t.} \quad \bar{T}\nu = \pm \hat{M}^{-1} \left(\bar{T}\nu\right)^{\intercal} \hat{M} \quad \forall \nu \in \mathbb{R}^p,$$

where $A_s \in \mathbb{R}^{n\times n}, B_s \in \mathbb{R}^{n\times\tau}, C_s \in \mathbb{R}^{n\times\tau}$ contain snapshot data and $M \in \mathbb{R}^{n\times n}$ is symmetric and positive semi-definite. Applying matrix calculus and the method of Lagrange multipliers, the

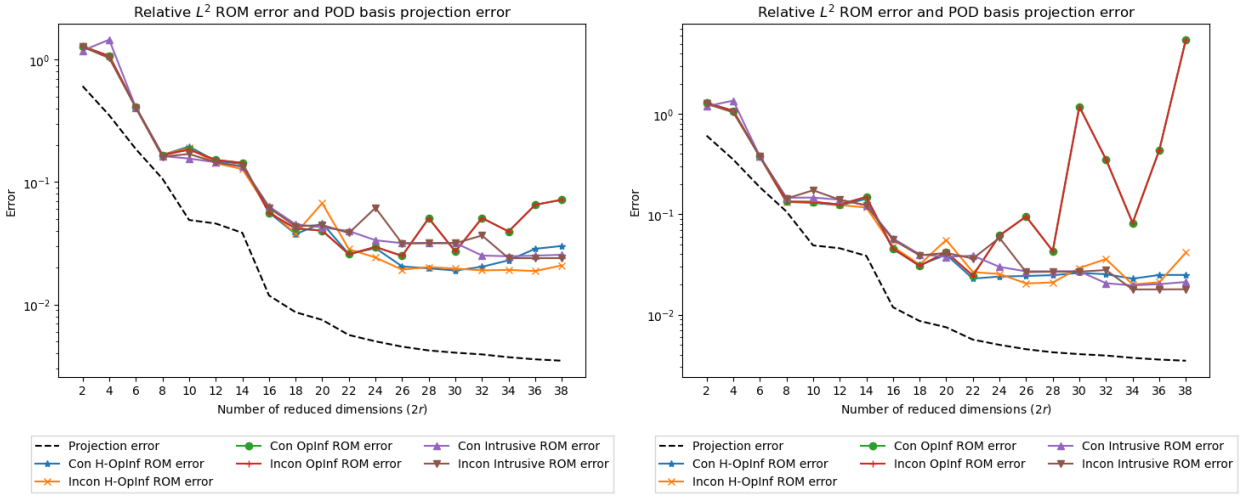

Figure 2-5. A comparison of ROMs for the parametric 1D wave equation with variable wave speed across four connected subdomains. Parameter-reproductive (left) and parameter-predictive (right) examples demonstrate that the Hamiltonian structure-preserving OpInf ROMs from [3, 4] (blue, orange) perform more stably than the structure-uninformed OpInf ROMs (green, red) with increasing basis dimension.

equivalence of this inference to a linear system can again be demonstrated, guaranteeing that the solution $\bar{T} \in \mathbb{R}^{n \times n \times p}$ is symmetric in the inner product defined by M, as required for Hamiltonian structure preservation. Applying this technique to a structure-preserving mixed finite element discretization of the 1D parametric wave equation

$$\begin{cases} \partial_{tt} q(x,t) = c(\mu)^2 \Delta q(x,t), & x \in \Omega \times (0,T], \\ q(x,t) = 0, & x \in \partial\Omega \times (0,T], \\ q(x,0) = q_0(x), & x \in \Omega, \\ \partial_t q(x,0) = 0, & x \in \Omega, \end{cases} \tag{2.4}$$

with parameterized wave speed $c(\mu)^2 : \mathbb{R}^p \to \mathbb{R}$,

$$c^2(\mu) = \mu_1^2 1_{\Omega_1} + \mu_2^2 1_{\Omega_2} + \ldots + \mu_p^2 1_{\Omega_p}, \tag{2.5}$$

then yields a Hamiltonian structure-preserving ROM given by $\dot{\hat{x}} = \hat{J}(\bar{T}\mu)\hat{x}$, which can be stepped forward in time[2] to generate an approximation $\tilde{x} = U\hat{x}$. The results of this are pictured in Figure 2-5, where it is shown that the proposed structure-preserving tensor parametric ROM offers stability advantages in predictive simulations. Future work will investigate both the unstructured and structured variants of this tensor parametric OpInf in the model reduction of large-scale parametric PDEs.

[2]Here, $\hat{J}$ is either $U^\intercal J U$ or $(U^\intercal J U)^{-1}$ depending on the consistency properties of the reduced basis U

2.2. Metriplectic Dynamics and Bracket-Based Surrogate Models

Metriplectic, also called GENERIC, systems [12, 13] provide a principled formalism for encoding dissipative dynamics in terms of thermodynamic variables that conserve energy and produce entropy. In particular, formal expression of the reversible and irreversible parts of state evolution with separate algebraic brackets leads to a general framework for maintaining essential conservation laws while simultaneously respecting dissipative effects. Said differently, the metriplectic formalism offers a principled extension to the Hamiltonian variational principle which captures dissipation in a complete thermodynamic way, making it explainable and useful for modeling nonconservative systems.

The compact, bracket-based expression of metriplectic dynamics is crucial to its advantage: since conservative and dissipative contributions can be formally factored out, phenomenological models can be built by specifying each of these components separately. Precisely, as stated in [5], metriplectic dynamics on the finite or infinite dimensional phase space $\mathcal{P}$ are generated by a free energy function(al) $F : \mathcal{P} \to \mathbb{R}$, $F = E + S$ defined in terms of a pair $E, S : \mathcal{P} \to \mathbb{R}$ representing energy and entropy, respectively, along with two algebraic brackets $\{\cdot,\cdot\}, [\cdot,\cdot] : C^\infty(\mathcal{P}) \times C^\infty(\mathcal{P}) \to C^\infty(\mathcal{P})$ which are bilinear derivations on $C^\infty(\mathcal{P})$ with prescribed symmetries and degeneracies $\{S,\cdot\} = [E,\cdot] = 0$. Here $\{\cdot,\cdot\}$ is an antisymmetric Poisson bracket, which is a Lie algebra realization on functions, and $[\cdot,\cdot]$ is a degenerate metric bracket which is symmetric and positive semi-definite on the same function space. These brackets are often identified with symmetric matrix fields $L : \mathcal{P} \to \mathrm{Skew}_n(\mathbb{R}), M : \mathcal{P} \to \mathrm{Sym}_n(\mathbb{R})$ satisfying

$$\{F, G\} = \nabla F \cdot L \nabla G, \quad [F, G] = \nabla F \cdot M \nabla G,$$

for all functions $F, G \in C^\infty(\mathcal{P})$ and all states $x \in \mathcal{P}$. Using the degeneracy conditions along with $\nabla x = I$ and slightly abusing notation leads the standard equations for metriplectic dynamics,

$$\dot{x} = \{x, F\} + [x, F] = \{x, E\} + [x, S] = L\nabla E + M\nabla S,$$

which are provably energy conserving and entropy producing.

2.2.1. Machine-learned metriplectic surrogate models

The primary drawback of the metriplectic formalism lies in the fact that analytic expressions for the functionals E, S and the brackets defined by L, M are highly complicated, and no general algorithm exists to compute them for a given system of interest. To that end, the work in [5] focused on machine-learning surrogates for these systems from snapshot data in a way which guarantees their special mathematical structure is preserved. By exploiting ideas from exterior algebra, we showed that a general metriplectic system can be parameterized with hand-crafted learnable matrix fields that scale optimally with the number of degrees of freedom in the system and the rank of the dissipation field M. This lead to an architecture, diagrammed in [5], based on neural ordinary differential equations (NODEs) [14] which is provably universally approximating on metriplectic systems, admits a bound on generalization error, and achieves state-of-the-art performance on benchmark metriplectic learning tasks. Some results are pictured in Figure 2-6, which shows

the correct metriplectic evolution of two equilibrating gases. This lends confidence to the idea that the metriplectic learning strategy developed as part of this project can be applied to learn thermodynamically-consistent surrogates of real-world systems from observational data.

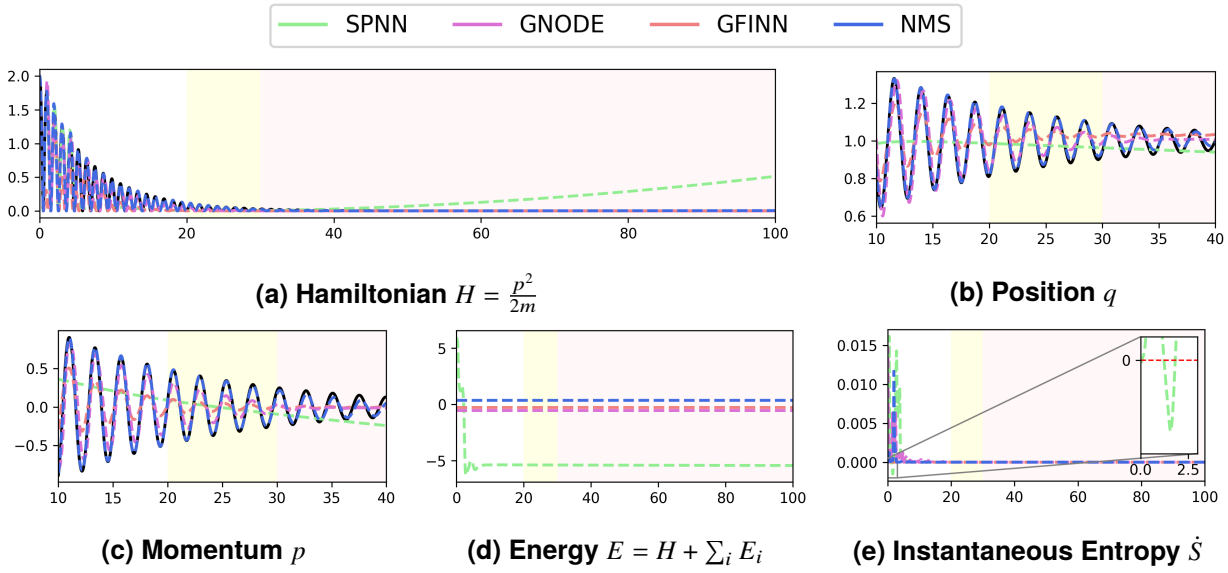

Figure 2-6. The ground-truth and predicted position, momentum, instantaneous entropy, and energies for a metriplectic benchmark involving the equilibration of gases in two containers. Our method, NMS [5], is provably metriplectic and most accurate in the training (white), validation (yellow), and testing (red) regimes when compared to previous metriplectic machine learning methods (SPNN, GNODE, GFINN).

2.2.2. Bracket-based graph neural networks

In addition to learning metriplectic systems from data, our work in [7] demonstrated that bracket-based dynamical systems can also be useful for designing graph neural network (GNN) architectures in a way that mitigates some of their inherent drawbacks. By applying ideas from the graph exterior calculus [15] to design GNN architectures reflecting the bracket-based structures in Table 2-1, we showed that such GNNs are able to mitigate the well-known *over-smoothing* problem [16] commonly encountered in GNN applications, whereby the feature representations at each node exhibit exponential homogenization with depth due to repeated Laplacian aggregations at each layer.

To understand how this works, note that graph-structured data comes equipped with a combinatorial notion of graph derivative defined by the incidence matrices of *k-cliques*–ordered subgraphs of $k+1$ nodes–on $(k+1)$-cliques. Denoting the set of k-cliques in the graph $\mathcal{G}$ by $\mathcal{G}_k$ and its dual space of linear functionals by Ω^k, these derivatives are commonly denoted as $d_k : \Omega_k \to \Omega_{k+1}$, and it can be shown that $d_{k+1} \circ d_k = 0$ for any k. Pictured in 2-7, the derivatives $d_0, d_0^\intercal$ and d_1 are simply the graph analogues of the usual vector calculus gradient, divergence, and curl.

Making use of the graph derivatives d_k and their adjoints $d_k^* : \Omega^{k+1} \to \Omega^k$ with respect to learnable inner products encoding graph attention (see [7] for details), bracket-based GNN architectures,

Formalism	Equation	Requirements	Completeness	Character
Hamiltonian	$\dot{x} = \{x, E\}$	$L^* = -L$, Jacobi's identity	complete	conservative
Gradient	$\dot{x} = -[x, E]$	$M^* = M$	incomplete	totally dissipative
Double Bracket	$\dot{x} = \{x, E\} + \{\{x, E\}\}$	$L^* = -L$	incomplete	partially dissipative
Metriplectic	$\dot{x} = \{x, E\} + [x, S]$	$L^* = -L, M^* = M,$ $L\nabla S = M\nabla E = 0$	complete	partially dissipative

Table 2-1. The abstract bracket formulations employed in [7]. "Conservative" indicates purely reversible motion, "totally dissipative" indicates purely irreversible motion, and "partially dissipative" indicates motion which either dissipates E (in the double bracket case) or generates S (in the metriplectic case).

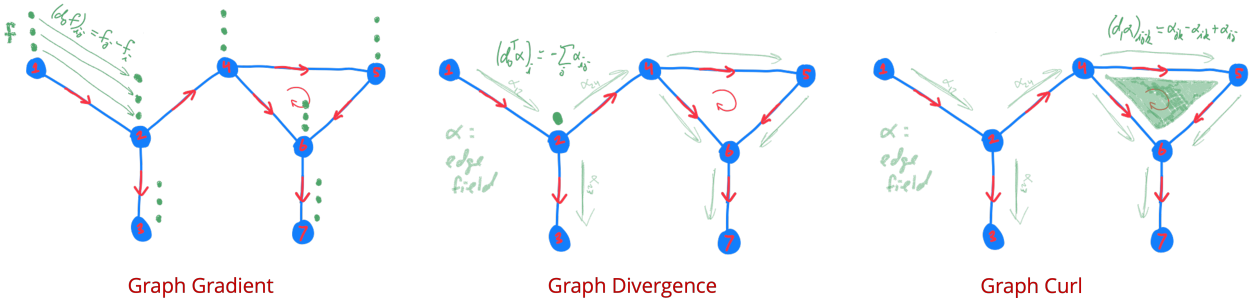

Figure 2-7. An illustration of the lowest-order graph derivative operators.

pictured in Figure 4-2 of Section 4, can be designed which are guaranteed to conserve or dissipate learned energy/entropy quantities. Although this idea is physics-based, its primary advantage in the graph setting lies in constraining the mechanism of the resulting GNNs in ways which are provably stable or otherwise advantageous for predictive accuracy. For example, our Hamiltonian bracket architecture designed in this way is incapable of over-smoothing in the traditional sense: since the network propagation is conservative by construction, there is no way for dissipative effects to homogenize network predictions as the network depth increases. We show in [7] that incorporating physics information into GNN design leads to improved accuracy in a variety of physics- and non-physics- based tasks including node classification, feature imputation, and trajectory prediction. It is particularly remarkable that the bracket-based structure-preserving character of our GNNs can improve predictive performance on large graph-structured datasets, such as citation networks, even when the task at hand has no physical interpretation.

2.3. Cohomology-Aware Model Reduction

Cohomology theories are topological tools that allow practitioners to understand when a local solution to a differential equation on a subdomain extends to a global one on the entire domain, or, said differently, how drastically the Fundamental Theorem of Calculus fails to hold on a given domain. This is of crucial importance in ModSim, where activities like co-design require the geometric perturbation of solution domains along with the fast solution of PDEs on these domains. Enabling model reduction in these instances is difficult, as well-understood techniques such as POD do not respect the topological properties of the original domain, and so cannot automatically

approximate the appropriate space of solutions to the original problem when this topology is important.

In the ongoing work [6], co-funded by the ASCR MMICC-3 center M2dt, we have worked to remedy this defect by exploiting ideas from the finite element exterior calculus (FEEC) [17]. The key to our approach lies in the construction of an appropriate *cochain complex*, a sequence $(\tilde{\Omega}, d)$ of vector spaces $\tilde{\Omega}^k$ and linear maps d_k satisfying $d_k \circ d_{k-1} = 0$. Along with appropriate *chain maps* P_k between the FOM and ROM "rungs" of the complex (c.f. Figure 2-8) satisfying $d_k P_k = P_{k+1} d_k$, we can show that this construction induces well-defined endomorphisms $\bar{P}_k : H^k(\Omega) \to H^k(\tilde{\Omega})$ on cohomology. This implies that the cohomology of the reduced space will be inherited from the cohomology of the full space provided the projection mappings P_k are constructed in an appropriate way.

$$\begin{array}{ccccc}
\Omega_{k-1} & \underset{d^*_{k-1}}{\overset{d_{k-1}}{\rightleftarrows}} & \Omega_k & \underset{d^*_k}{\overset{d_k}{\rightleftarrows}} & \Omega_{k+1} \\
\downarrow{\scriptstyle P_{k-1}} & & \downarrow{\scriptstyle P_k} & & \downarrow{\scriptstyle P_{k+1}} \\
d^\dagger\tilde{\Omega}^k & \xrightarrow{d_{k-1}} & \tilde{\Omega}^k & \xrightarrow{d_k} & d\tilde{\Omega}^k
\end{array}$$

Figure 2-8. A commutative diagram illustrating the cochain complexes (Ω^k, d_k) and $(\tilde{\Omega}^k, d^k)$ which are used in building the cohomology-aware ROMs in [6], along with the data-driven chain maps P_k inducing morphisms on cohomology.

In constructing the projections $P_k : \Omega^k \to \tilde{\Omega}^k$, we make use of the *Hodge decomposition* $\Omega^k = \operatorname{im} d_k \oplus \operatorname{im} d_k^* \oplus \ker \Delta_k$ (see Figure 2-9), where d_k^* denotes the L^2-adjoint of d_k and $\Delta_k = d_{k-1} d_{k-1}^* + d_k^* d_k$ is the Hodge Laplacian. This is the differential k-form analogue of the usual Helmholtz-Hodge decomposition of vector fields, e.g., $v = \nabla\phi + \nabla \times \psi + h$. Importantly, the L^2-orthogonality of the different components of the Hodge decomposition allows for data-driven projections $P_k = \tilde{P}_d + \tilde{P}_{d^*} + \tilde{P}_h$ which are built from snapshots of the system state; each learned projection $\tilde{P}_l = U_l U_l^*$ is built according to the POD decomposition $X_l \approx U_l \Sigma_l V_l^\intercal$, where X_l denotes the projection of the snapshot set X onto the l^{th} component of the Hodge decomposition. It follows that P_k is orthogonal and maps each component of the Hodge decomposition into itself, enabling the commutative diagram shown in Figure 2-8 when $P_{k-1} = d_{k-1}^\dagger P_k d_{k-1}$ and $P_{k+1} = d_k P_k d_k^\dagger$ ($A^\dagger$ denotes the Moore-Penrose pseudoinverse of A).

A simple example of where this technique is necessary and improves greatly upon the current state of the art is given by Poisson's equation on a domain $[0, 1]^2$ with periodic boundary conditions:

$$\begin{aligned}
\Delta_0 u &= f && \text{in } (0,1)^2, \\
u(0, y) &= u(1, y) && \text{on } \{0, 1\} \times [0, 1], \\
u(x, 0) &= u(x, 1) && \text{on } [0, 1] \times \{0, 1\}.
\end{aligned} \tag{2.6}$$

This problem is ill-posed as written above since $\dim \ker \Delta_0 = \mathbb{R}$, reflecting the fact that $u' = u + c$ solves the system for any solution u and constant c. This means that the system (2.6) must be augmented with the *gauge condition* $(u, 1)_{L^2} = c$ for some constant c. This poses great challenges

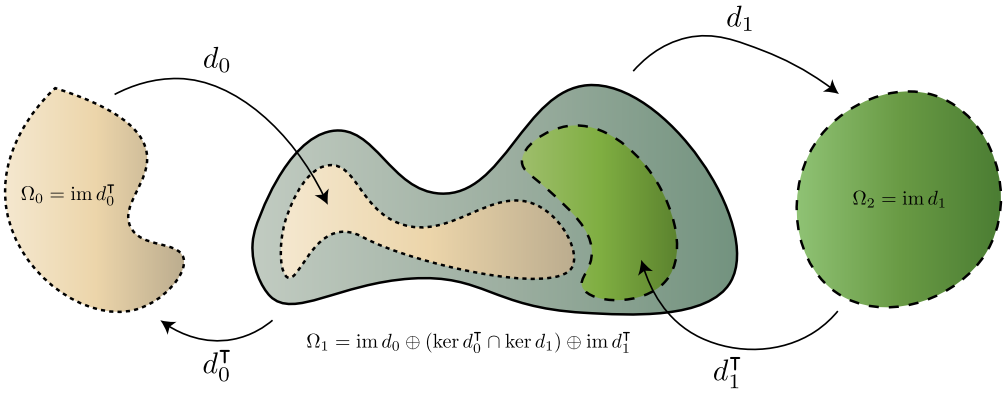

Figure 2-9. A cartoon illustration of the Hodge decomposition of one-forms (duals of vector fields) in the ℓ^2 inner product.

for model reduction, where snapshots of the system (2.6) have been collected into a matrix X, possibly with different gauges. In this case, a data-driven basis U constructed via standard POD will not induce a valid morphism on cohomology between functions on the problem domain M and the reduced domain span(U), and hence will not lead to a solvable reduced model in any desired gauge. To see this more precisely, note that a discrete analogue of (2.6) after, say, discretization by linear finite elements, is given by

$$\begin{aligned} \mathbf{K}\mathbf{u} &= \mathbf{M}_0\mathbf{f}, \\ \mathbf{1}^\intercal\mathbf{M}_0\mathbf{u} &= c\mathbf{1}^\intercal\mathbf{M}_0\mathbf{1}, \end{aligned} \tag{2.7}$$

where $\mathbf{M}_0, \mathbf{K}$ are the standard mass and stiffness matrices of the scheme and $\mathbf{u}, \mathbf{f}$ represent coefficients in the finite element basis. Now, if $\mathbf{U}$ is a data-driven basis constructed from snapshots of $\mathbf{u}$ satisfying (2.7), potentially for different c, it follows that the Galerkin-projected system to be solved for $\tilde{\mathbf{u}} = \mathbf{U}\hat{\mathbf{u}}$ is

$$\begin{aligned} \mathbf{U}^\intercal\mathbf{K}\mathbf{U}\hat{\mathbf{u}} &= \mathbf{U}^\intercal\mathbf{M}_0\mathbf{f}, \\ \mathbf{1}^\intercal\mathbf{M}_0\mathbf{U}\hat{\mathbf{u}} &= c\mathbf{1}^\intercal\mathbf{M}_0\mathbf{1}. \end{aligned}$$

However, this system is not solvable in general. First, note that it is highly unlikely that span $U \perp \ker \mathbf{K}$ since the harmonic part of each snapshot $\mathbf{u}$ contributes to the POD basis computation. Moreover, the gauge condition is fixed by the left-hand side involving the basis, and there are no degrees of freedom left in $\hat{\mathbf{u}}$ which can be used to satisfy the right-hand side for arbitrary c. This is a direct consequence of the lack of cohomology-preservation mentioned above, which has significant consequences on the accuracy of the ROM (c.f. Figure 2-10).

This situation can be effectively remedied using the cohomology-aware technology introduced before. In this case, the coexact and harmonic components[3] X_{d^*}, X_h are extracted from the snapshot data and used to form the Galerkin approximation $\tilde{\mathbf{u}} = \mathbf{U}_{d^*}\hat{\mathbf{u}}_{d^*} + \mathbf{U}_h\hat{\mathbf{u}}_h$. Projection of the governing equations then leads to the solvable ROM:

$$\begin{aligned} \mathbf{U}_{d^*}^\intercal\mathbf{K}\mathbf{U}_{d^*}\hat{\mathbf{u}}_{d^*} &= \mathbf{U}_{d^*}^\intercal\mathbf{M}_0\mathbf{f}, \\ \mathbf{1}^\intercal\mathbf{M}_0\mathbf{U}\hat{\mathbf{u}}_h &= c\mathbf{1}^\intercal\mathbf{M}_0\mathbf{1}. \end{aligned}$$

[3] that there is no exact component X_d follows from a simple degree argument

Notice that this resembles a mixed-form finite element method due to the use of separate bases $\mathbf{U}_{d^*}, \mathbf{U}_h$ for the discrete Hodge components of the approximate solution $\tilde{\mathbf{u}}$. The impact of this is shown in Figure 2-10, where a reconstruction based on our cohomology-aware approach provides an accuracy gain of nine orders of magnitude over the naive POD-ROM on a simple manufactured problem.

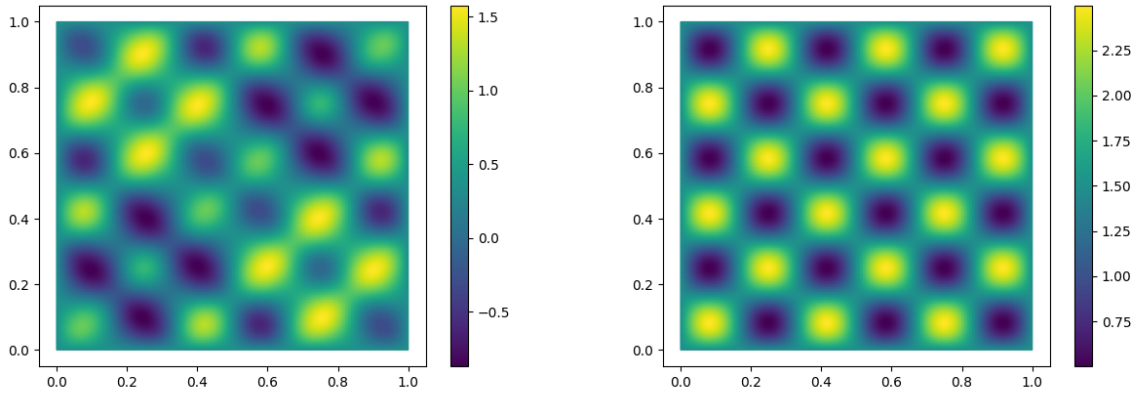

Figure 2-10. Poisson ROM reconstruction of a harmonically dominant function. Left: naive POD with $n = 4$ modes, relative error 75%. Right: HdR POD with $n = 3$ coexact and $n = 1$ harmonic mode, relative error 1.8×10^{-7}%.

Note that the advantages of this approach do not stop at Poisson's equation; the technique behind our cohomology-aware model reduction is generally applicable to any system of PDEs which take place on a topologically nontrivial domain. For example, the solution to Maxwell's equations inside a domain with obstacles (e.g., a circuit box) will depend crucially on the number and placement of these obstacles, and naive model reduction approaches such as standard Galerkin POD will not respect this. This can have dramatic impacts on important physical aspects of the system, such as charge conservation and magnetic flux, which are likely to produce un-physical and untrustworthy predictions at the ROM level. Therefore, it is essential that strategies such as that in [6] are developed in order to accelerate the rapid design and prototyping of electrical components necessary for carrying out Sandia's mission.

3. PROJECT METRICS

Throughout the duration of the project, I prioritized outreach activities and the dissemination of research results through publications, external presentations, and open-source software. This section summarizes the publications generated during my project, the conference and colloquium talks I gave, the students I mentored, the proposals I (co-)wrote, and the special sessions I (co-)organized.

3.1. Publications and Presentations

Publications. By September 2024, the project had four associated journal articles and peer-reviewed conference proceedings: two published [1, 7], one accepted for publication [2], and one more under review [5]. Three additional journal articles [3, 6, 8] and two CSRI Summer Proceedings papers [4, 18] are under active development and expected to be submitted by the middle of FY2025. Details of the completed papers are as follows:

- A. Gruber, I. Tezaur, "Canonical and noncanonical Hamiltonian operator inference", Comput. Methods Appl. Mech. Eng., 2023.
- A. Gruber, K. Lee, N. Trask, "Reversible and irreversible bracket-based dynamics for deep graph neural networks", Neural Information Processing Systems (NeurIPS) 2023.
- A. Gruber, I. Tezaur, "Variationally consistent Hamiltonian model reduction", SIAM J. Appl. Dyn. Sys. (to appear).
- A. Gruber, K. Lee, H. Lim, N. Park, N. Trask, "Efficiently parameterized neural metriplectic systems", arXiv:2405.16305 (under review for NeurIPS 2024).

Open-source Software. Source code for carrying out the Hamiltonian operator inference procedures in [1, 2] was released under SCR#:2917.0 and hosted at the Sandia Github site `https://github.com/sandialabs/HamiltonianOpInf`.

Presentations. I presented a total of 17 talks and posters (11 invited, 6 contributed) associated with the work completed during this project, with several more scheduled for later in the calendar year. The titles, dates, and locations of these events are given below.

Invited External Presentations

- "Learning metriplectic systems and other bracket-based dynamics", Mathematics Seminar series (virtual), University of Vienna, Vienna, Austria. (June 19, 2024).

- "Learning metriplectic systems and other bracket-based dynamics", Minisymposium on Mathematics of Machine Learning, CMS summer meeting, University of Saskatchewan, Saskatoon, Canada. (June 2, 2024).
- "Learning bracket-based dynamical systems for property preserving model reduction", guest lecture, UPenn, Philadelphia, PA. (Apr. 23, 2024).
- "Property-preserving model reduction in bracket-based dynamical systems", UNM Applied Mathematics Seminar Series, Albuquerque, NM. (March 25, 2024).
- "Learning Operators for Structure-Informed Surrogate Models", DOE ASCR PI Meeting, Albuquerque, NM. (poster; Jan. 8, 2024)
- "Data-driven dynamical systems with structural guarantees", S. Scott Collis Advanced Modeling & Simulations seminar series (virtual), El Paso, TX. (Nov. 10, 2023).
- "Property-preserving model reduction for conservative and dissipative systems", Numerical Analysis of Galerkin ROMs seminar series (virtual), INRIA, Bordeaux, France. (Oct. 10, 2023).
- "Data-driven surrogate models for bracket-based dynamical systems", Minisymposium on Data-driven Methods for Circuits and Devices, 2nd IACM MMLDE-CSET, El Paso, TX. (Sep. 27, 2023).
- "Mathematics in different settings" (virtual), Hong Duc University, Thanh Hóa, Vietnam. (May 20, 2023).
- "Energetically consistent model reduction for Hamiltonian and metriplectic systems", CRUNCH webinar (virtual), Brown University, Providence, RI. (Dec. 9, 2022).
- "Convolutional neural networks for data compression and reduced-order modeling", Minisymposium on Machine Learning for Large-Scale Scientific Data Analytics, SIAM MDS, (Sep. 28, 2022).

Contributed External Presentations

- "Flexible and variationally consistent Hamiltonian model reduction", Model Reduction and Surrogate Modeling (MORe2024), La Jolla, CA. (Sep. 10, 2024)
- "Learning metriplectic systems from full and partial state information", 16^{th} World Congress on Computational Mechanics, Vancouver, Canada. (July 23, 2024)
- "Reversible and irreversible bracket-based dynamics for deep graph neural networks" (poster), Advances in Neural Information Processing Systems (NeurIPS), New Orleans, LA. (Dec. 12, 2023)
- "Variational consistency in model reduction for conservative and dissipative systems", Minisymposium on Data-Driven Methods – Solids, A Conference Celebrating the 80^{th} Birthday of Thomas J.R. Hughes, Austin, TX. (Oct. 23, 2023)

- "Canonical and Noncanonical Hamiltonian Operator Inference", Minisymposium on Geometric Mechanics Formulations and Structure-Preserving Discretizations, 17^{th} U.S. National Congress on Computational Mechanics, Albuquerque, NM. (July 26, 2023)
- "Canonical and Noncanonical Hamiltonian Model Reduction", Workshop on Establishing Benchmarks for Data-Driven Modeling of Physical Systems, University of Southern California, Los Angeles, CA. (April 6, 2023)

Organization. Five special sessions at major international conferences were organized to promote outreach and make contact with state-of-the-art research in structure-informed surrogate modeling. The organizational activities undertaken during this project are as follows.

- Co-organizer, "Structure-preserving model reduction for large-scale systems", minisymposium, SIAM Computational Science and Engineering, Mar. 3-7, 2025.
- Co-organizer, "Geometric Mechanics Formulations and Structure-Preserving Discretizations for Models of Physical Systems", minisymposium, SIAM Joint Mathematics Meetings, Jan. 8-11, 2025.
- Co-organizer, "Advances in machine learning on graphs for physical sciences and data analysis", minisymposium, SIAM Mathematics of Data Science, Oct. 21-25, 2024.
- Co-organizer. "Geometric Mechanics formulations and structure-preserving discretizations for continuum mechanics and kinetics models", minisymposium, 16th World Congress on Computational Mechanics, July 21-26, 2024.
- Co-organizer. "Geometric Mechanics Formulations and Structure-Preserving Discretizations", minisymposium, 17th U.S. National Congress on Computational Mechanics, July 26-30, 2023.

3.2. Mentoring and Training

Due to my proximity to and strategic alignment with the M2dt ASCR MMICC center and the fHNM CIS LDRD, I had the opportunity to formally and informally mentor graduate students affiliated with these programs. These students, my interactions with them, and the scope of their work are outlined below.

- **Arjun Vijaywargia, University of Notre Dame.** Arjun was a summer 2024 graduate intern funded by M2dt who continued as year round intern afterwards. Arjun's work led to the development of effective, structure-preserving tensor parametric techniques for operator inference, which will be submitted for journal publication in FY2025. Arjun's work is based on ideas generated during this project, and I served as his primary mentor.
- **Ian Moore, Virginia Tech.** Ian was a summer 2024 graduate intern funded by M2dt. Ian worked to develop methods for coupling non-intrusive operator inference models via the Schwarz alternating method. I served as a secondary technical mentor to Ian, regularly meeting and assisting him with method development and experimental design.

- **Rishi Pawar, University of Arizona.** Rishi was a 2023 summer graduate intern funded by fHNM who returned for summer 2024. His work investigated data-driven approaches based on operator inference and neural ordinary differential equations for the computation of flux conditions necessary for coupling systems across interfaces. I served as an informal technical mentor to Rishi during both summers, periodically assisting him in the design of effective reduced models based on these technologies.

- **Roxana Pohlmann, Vienna University of Technology.** Roxana was a self-funded visiting researcher at Sandia from September 2023 to January 2024, where she was advised by Nat Trask (1442) and Rekha Rao (1516) on graph neural network approaches for modeling problems in injection molding. I served as an informal technical mentor to Roxana, regularly meeting with her to offer guidance related to graph neural networks in general and our bracket-based architectures (c.f. [7]) in particular.

- **Edward Huynh, Arizona State University.** Edward was a 2023 summer graduate intern funded by fHNM who returned for summer 2024. His work involved maturing the theoretical understanding of term re-weighting in physics-informed neural network loss functions. I served as an informal technical mentor to Edward in summer 2023, periodically assisting him in theoretical arguments related to the analysis of neural network architectures. I also reviewed Edward's 2023 summer CSRI proceedings article entitled *A Preliminary Study for Obtaining Inverse Sobolev-Type Inequalities for Tanh Neural Networks*.

3.3. Strategic Partnerships

From project inception in FY22, a remarkable effort was made to cultivate new internal and external collaborations, as well as strengthen existing ones. Below the key collaborations fostered during this project are summarized.

- **Irina Tezaur, 8734.** Irina provided general mentorship and technical expertise related to solid mechanics and intrusive model reduction.

- **Nat Trask, University of Pennsylvania (formerly 1442).** Nat is an expert in scientific machine learning and compatible meshfree methods who provided general mentorship and technical guidance related to structure-preserving neural networks.

- **Pavel Bochev, 1400.** Pavel provided expertise and general guidance in numerical analysis and compatible finite elements, with particular emphasis on cohomology-aware methods.

- **Eric Cyr, 1442.** Eric provided technical guidance related to high-performance implementation and applications of structure-preserving methods to electromagnetics.

- **Max Gunzburger, Oden Institute at UT Austin.** Max provided technical guidance in intrusive model reduction particularly related to the implementation of strong boundary conditions.

- **Kookjin Lee, Arizona State University.** Kookjin is an expert in scientific machine learning and neural ordinary differential equations who provided technical support on projects involving structure-preserving machine learning.
- **Noseong Park, Korea Advanced Institute of Science and Technology.** Noseong is an expert in machine learning and large-scale data analytics who provided implementation assistance and student support for the learnable metriplectic models developed in [5].
- **Patrick Blonigan, 8738.** Patrick provided technical expertise related to the model reduction of large-scale Sanda-relevant systems, including the implementation of necessary hyper-reduction methods.
- **Alejandro Mota, 8363**. Alejandro provided technical expertise in solid mechanics and realistic material models.
- **Paul Kuberry, 1442.** Paul provided technical expertise in software development related to compatible mixed finite element schemes.
- **Chris Eldred, 1442.** Chris provided technical expertise in geometric mechanics and experimental design necessary for benchmarking structure-preserving methods.
- **Shane McQuarrie, 1441.** Shane provided technical expertise in non-intrusive model reduction and parametric operator inference.

These partnerships also led to two joint proposals with external parties and three internal LDRD proposals, listed below.

External Proposals

- Irina Tezaur and I collaborated with Nathan Urban (PI–BNL) and Youngsoo Choi (LLNL) on an ASCR AI for Science proposal *Few-Shot Optimal Learning of Surrogate Models for Uncertainty Quantification and Optimization of Exascale Simulations* (in progress).
- Irina Tezaur and I collaborated with Ramin Bostanabad (PI–UC Irvine) on an ASCR SciML proposal *Quantifying Epistemic and Aleatoric Uncertainties in Multi-fidelity Operator Learning* (not funded).

Internal Proposals (c.f. Section 4)

- I collaborated with Irina Tezaur (PI) on an FY25-FY28 ESRF LDRD proposal with acronym "AHeaD" (funded).
- I collaborated with Eric Cyr (PI) on a FY25-FY28 REHEDS LDRD proposal with acronym "DD-MAkER" (funded).
- I collaborated with Jonas Actor (PI, 1442) on a FY25-FY28 CIS LDRD proposal entitled "High-Dimensional Variational Gaussian Representation, Approximation, and Discretizations (H-DIVGRAD)" (not funded).

This page intentionally left blank.

4. PROJECT LEGACY

This LDRD project facilitated the development of new and improved methods for structure-informed surrogate modeling which are non-intrusive and guaranteed to satisfy key properties of their full-order counterparts. Due to resulting benefits in predictive accuracy, numerical stability, and physical realism, the products of this project may inspire further work in data-driven surrogate modeling with structural constraints, eventually bringing it to a level where it can be trusted for mission-relevant simulations. In particular, the Hamiltonian OpInf strategies highlighted in Chapter 2.1 have been demonstrated on moderate-to-large scale 3D linear elastic problems (c.f. Figure 4-1) and are ready to deploy to other large problems of interest whose solutions obey conservative dynamics. Similarly, the GNN-based networks from Chapter 2.2 have been applied to simulate reversible and irreversible dynamical behavior on graphs with 250K+ edges, and are ready for regression and classification tasks on other large graphs.

Overall, my project contributed to the CIS Investment Area goals and objectives in three principal ways:

- By producing mathematics and software in in support of Sandia structure-informed surrogate modeling efforts, specifically targeting Hamiltonian, metriplectic, and Hodge-de Rham systems, the project generated new capabilities for the non-intrusive surrogate modeling of systems with important mathematical structure.
- By disseminating this work through 4 peer-reviewed publications, 1 open-source software release, and 17 external talks, posters, and colloquia, as well as building community via the organization of 5 special sessions and minisymposia at leading international conferences, the project strengthened Sandia's reputation as a leader in structure-preserving modeling & simulation.
- By engaging with internal and external collaborators and institutions including 2 ASCR MMICC centers, mentoring 5 graduate students, and promoting numerous cross-center collaborations between 1400 and 8700, the project fostered ongoing research efforts and contributed to the development of Sandia's future workforce.

Remarkably, the project also produced several methods and technologies which have seen insertion both internally at Sandia and elsewhere. A brief summary of the current users of these products now follows.

M2dt MMICC center Multifaceted Mathematics for Predictive Digital Twins (M2dt) is a five-year (FY22-FY27), multi-institutional, ASCR-funded mathematics institute led out of UT Austin by Karen Willcox and Omar Ghattas. The Sandia PI is Irina Tezaur who manages a funding portfolio of $2.5M allocated over this time period. The M2dt MMICC-3 ASCR project is focused on the mathematical foundations of digital twins, including robust and structure-informed methods

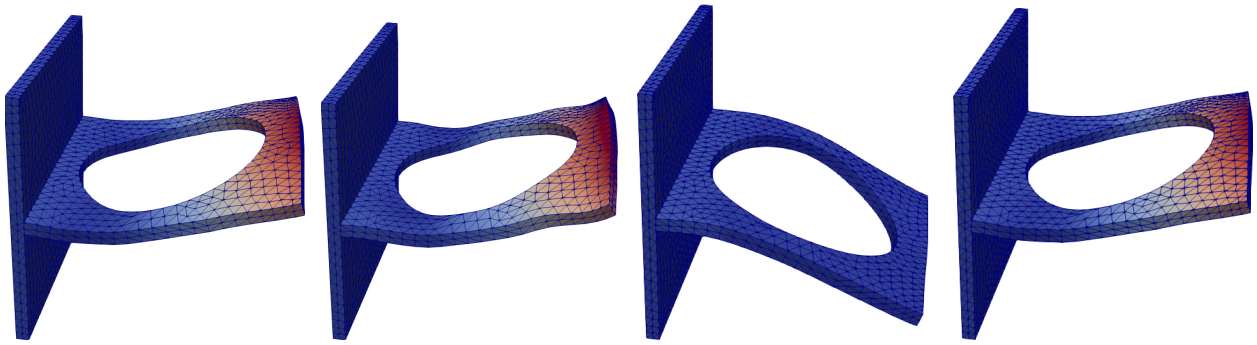

(a) FOM (b) Intrusive G-ROM (c) Intrusive H-ROM (inconsistent) (d) Intrusive H-ROM (consistent)

Figure 4-1. z–displacement solutions to the flexible bracket problem, scaled by a factor of 10 for plotting purposes, at the final time $t = 2 \times 10^{-2}$ s for the FOM (a) and various intrusive ROMs (b)–(d).

for data-driven model reduction to facilitate optimal control and optimization under uncertainty. This project provides M2dt with a suite of structure-preserving surrogate models which further the ultimate goal of trustworthy, real-time computational prediction and inference for informing a dynamically evolving physical system. The Hamiltonian models developed as part of this project have already shown promise for stably predicting the deformation of elastic brackets [2], see Figure 4-1, and development of cohomology-aware ROMs for models with complex topology based on the work in [6] will continue into FY25.

SEA-CROGS MMICC center Scalable, Efficient and Accelerated Causal Reasoning Operators, Graphs and Spikes for Earth and Embedded Systems (SEA-CROGS) is another five-year, (FY22-FY27) multi-institutional, ASCR-funded mathematics institute led out of Brown University by George Karniadakis. The Sandia PI is Eric Cyr who manages a funding portfolio of $2.5M allocated over this time period. The SEA-CROGS MMICC-3 ASCR project aims to develop physics-informed machine intelligence to augment foundational understanding and optimal control of complex systems spanning extremely disparate scales, thereby greatly accelerating tasks like computational modeling, causal reasoning, etiology, and pathway discovery for Earth systems. In line with the "Graphs" thrust of SEA-CROGS, this project provides bracket-based structure-preserving graph neural networks which function as explainable prediction engines [7]. This technology is currently being inserted into work aligned with the "Operators" thrust which will continue into FY25: to accelerate the modeling of ice sheet dynamics, our structure-informed graph neural networks, see Figure 4-2, are actively being combined with deep operator networks (DeepONets) to facilitate the efficient and accurate prediction of field quantities output by ice sheet models.

AHeaD LDRD Accelerating the Analyst Workflow: Adaptive Hybrid modEls via DomAin Decomposition (AHeaD) is a new three-year LDRD (FY25-FY28) led by Irina Tezaur (8734) and funded at $2.2M total by the Engineering Sciences Investment Area. The project aims to expedite analyst workflows and significantly reduce mesh generation times by developing adaptive hybrid models through domain decomposition. AHeaD will adopt some of the operator inference (OpInf) technology developed as part of this project, particularly as it relates to parametric dependence and the enforcement of Dirichlet boundary conditions. Figure 4-3 shows a coupled FOM-ROM model

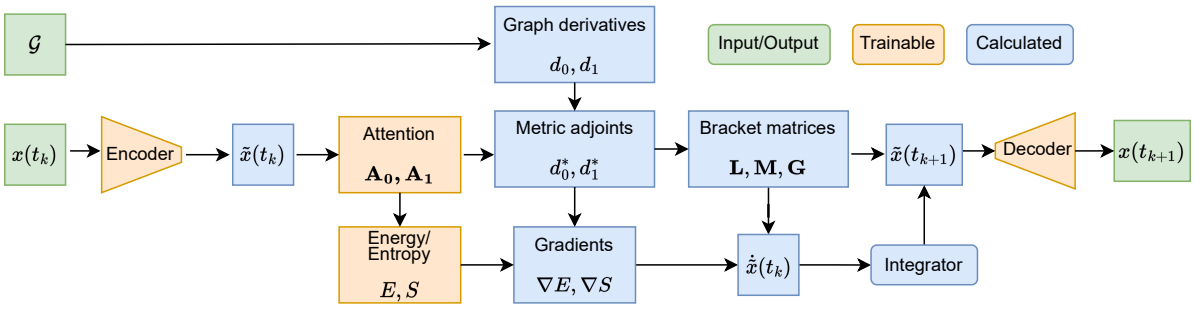

Figure 4-2. A schematic outlining the bracket-based GNN architecture from [7].

where the ROM is built with non-intrusive OpInf, and boundary conditions are enforced as in [6]. More details can be found in the forthcoming CSRI Proceedings article [18].

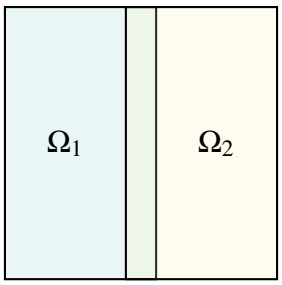

(a) Overlapping domain decomposition

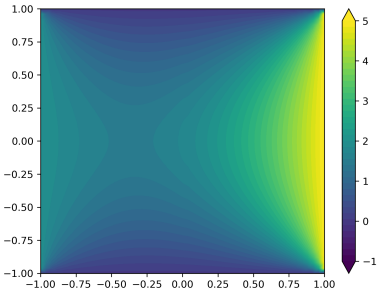

(b) Solution u at final time

Figure 4-3. OpInf-Schwarz coupling of a heat equation across two subdomains.

DD-MAkER LDRD Data-Driven Models Accelerating Cavity Electromagnetic Response Computations (DD-MAkER) is a new three-year LDRD (FY25-FY28) led by Eric Cyr (1442) and funded at $1.7M total by the REHEDS Inverstment Area. This project proposes a data-driven approach for modeling system generated electromagnetic pulse (SGEMP) that reduces analysis time, enabling rapid response to system design inquiries. DD-MAkER will develop methods informed by geometric and topological invariants for rapid predictions across parameter space, and so benefits from the Hamiltonian-preserving (e.g., [1, 2]) and cohomology-aware (e.g., [6]) methods developed as part of this project. As DD-MAkER also plans to consider machine-learned surrogates with structure, the project may also benefit from our methods for learning bracket-based dynamics developed in [7, 5].

Kookjin Lee's Lab The work on metriplectic systems [7, 5] done with my collaborator Kookjin Lee will be continued in his lab at Arizona State University. He is the recent recipient of the NSF CAREER award 2338909 "Accelerating Scientific Discovery via Deep Learning with Strong Physics Inductive Biases", which is focused on improved, structure-informed methods for data-driven system identification. Kookjin and his students will continue to expand on the foundation built during this LDRD, in collaboration with me as appropriate.

This page intentionally left blank.

REFERENCES

[1] A. Gruber and I. Tezaur, "Canonical and noncanonical Hamiltonian operator inference," *Computer Methods in Applied Mechanics and Engineering*, vol. 416, p. 116334, 2023.

[2] A. Gruber and I. Tezaur, "Variationally consistent Hamiltonian model reduction," *SIAM Journal on Applied Dynamical Systems*, (to appear).

[3] A. Vijaywargia, S. A. McQuarrie, and A. Gruber, "Tensor parametric operator inference with structure information," (in progress).

[4] A. Vijaywargia, S. A. McQuarrie, and A. Gruber, "Tensor parametric operator inference," *CSRI Summer 2024 Proceedings*, (to appear).

[5] A. Gruber, K. Lee, H. Lim, N. Park, and N. Trask, "Efficiently parameterized neural metriplectic systems," (under review).

[6] A. Gruber, P. Kuberry, C. Eldred, and P. Bochev, "Cohomology-aware model reduction," (in progress).

[7] A. Gruber, K. Lee, and N. Trask, "Reversible and irreversible bracket-based dynamics for deep graph neural networks," in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.

[8] P. Blonigan, A. Gruber, and I. Tezaur, "Non-intrusive hyper-reduction of nonlinear thermo-mechanical models," (in progress).

[9] B. Peherstorfer and K. Willcox, "Data-driven operator inference for nonintrusive projection-based model reduction," *Computer Methods in Applied Mechanics and Engineering*, vol. 306, pp. 196–215, 2016.

[10] L. Peng and K. Mohseni, "Symplectic model reduction of Hamiltonian systems," *SIAM Journal on Scientific Computing*, vol. 38, pp. A1–A27, 2022/11/05 2016.

[11] A. Salinger, R. Bartlett, A. Bradley, Q. Chen, I. Demeshko, X. Gao, G. Hansen, A. Mota, R. Muller, E. Nielsen, J. Ostien, R. Pawlowski, M. Perego, E. Phipps, W. Sun, and I. Tezaur, "Albany: Using Agile Components to Develop a Flexible, Generic Multiphysics Analysis Code," *Int. J. Multiscale Comput. Engng.*, vol. 14, pp. 415–438, 2016.

[12] P. J. Morrison, "A paradigm for joined hamiltonian and dissipative systems," *Physica D: Nonlinear Phenomena*, vol. 18, no. 1, pp. 410–419, 1986.

[13] M. Grmela and H. C. Öttinger, "Dynamics and thermodynamics of complex fluids. i. development of a general formalism," *Phys. Rev. E*, vol. 56, pp. 6620–6632, Dec 1997.

[14] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, "Neural ordinary differential equations," *Advances in neural information processing systems*, vol. 31, 2018.

[15] O. Knill, "The dirac operator of a graph," 2013.

[16] C. Cai and Y. Wang, "A note on over-smoothing for graph neural networks," *arXiv preprint arXiv:2006.13318*, 2020.

[17] D. N. Arnold, *Finite element exterior calculus*. SIAM, 2018.

[18] I. Moore, C. R. Wentland, A. Gruber, and I. Tezaur, "Domain decomposition-based coupling of operator inference reduced order models via the Schwarz alternating method," *CSRI Summer 2024 Proceedings*, (to appear).

DISTRIBUTION

Email—Internal

Name	Org.	Sandia Email Address
Technical Library	1911	sanddocs@sandia.gov

Hardcopy—Internal

Number of Copies	Name	Org.	Mailstop

Hardcopy—External

Number of Copies	Name(s)	Company Name and Company Mailing Address

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.