

Variationally Consistent Hamiltonian Model Reduction*

Anthony Gruber[†] and Irina Tezaur[‡]

Abstract. Although ubiquitous as first-principles models for conservative phenomena, Hamiltonian systems present numerous challenges for model reduction even in relatively simple, linear cases. Here, we present a method for the projection-based model reduction of canonical Hamiltonian systems that is variationally consistent for any choice of linear reduced basis: Hamiltonian models project to Hamiltonian models. Applicable in both intrusive and nonintrusive settings, the proposed method is energy-conserving and symplectic, with error provably decomposable into a data projection term and a term measuring deviation from canonical form. Examples from linear elasticity with realistic material parameters are used to demonstrate the advantages of a variationally consistent approach, highlighting the steady convergence exhibited by consistent models where previous methods reliant on inconsistent techniques or specially designed bases exhibit unacceptably large errors.

Key words. variational consistency, Hamiltonian systems, projection-based model order reduction (PMOR), operator inference, structure preservation

MSC codes. 65P10, 74H15, 70H05

DOI. 10.1137/24M1652490

1. Introduction. Hamiltonian systems represent a simple description of purely conservative dynamics and are often viewed as idealized limits of realistic physical systems, e.g., the Euler equations as the zero-viscosity limit of the Navier–Stokes equation. Despite their relative simplicity, the subclass of canonical Hamiltonian systems still exhibits many challenges for model reduction. This is especially the case for projection-based reduced-order models (ROMs), where the governing equations are projected onto a data-driven basis for full-order model (FOM) solutions. It is straightforward to show that a naïve basis projection will not preserve Hamiltonian structure in the ROM (cf. [29]), meaning that the reduced ODE system, which updates the solution coefficients, will not be energy-conserving. This means that approximations built from a naïve projection-based ROM will not preserve critical invariants of the FOM, leading to artificial instabilities and other unphysical behavior in time.

To describe this more concretely, consider a canonical Hamiltonian system in the state $\mathbf{x} = (\mathbf{q} \ \mathbf{p})^\top \in \mathbb{R}^N$, where $N = 2M \in \mathbb{R}$, comprised of position and momentum variables,

$$(1.1) \quad \dot{\mathbf{x}} = \mathbf{J} \nabla H(\mathbf{x}), \quad \mathbf{x}_0 = \mathbf{x}(0),$$

*Received by the editors April 8, 2024; accepted for publication (in revised form) by G. Gottwald August 24, 2024; published electronically February 4, 2025.

<https://doi.org/10.1137/24M1652490>

Funding: This work was supported by the John von Neumann fellowship at Sandia and the U.S. Department of Energy, Office of Advanced Scientific Computing Research, under the Multifaceted Mathematics for Predictive Digital Twins (M2dt) project.

[†]Center for Computing Research, Sandia National Laboratories, Albuquerque, NM 87185 USA (adgrube@sandia.gov).

[‡]Quantitative Modeling and Software Engineering, Sandia National Laboratories, Livermore, CA 94551 USA (ikalash@sandia.gov).

where $\mathbf{J} = -\mathbf{J}^\top \in \mathbb{R}^{N \times N}$ is the full-rank and skew-symmetric symplectic matrix $\mathbf{J} = [\mathbf{0} \ \mathbf{I}; -\mathbf{I} \ \mathbf{0}]$ and $H : \mathbb{R}^n \rightarrow \mathbb{R}$ is a Hamiltonian functional [22]. Denote the flow map of (1.1) by $\varphi(t, \mathbf{x}_0) = \varphi_t(\mathbf{x}_0)$ so that $\varphi_t(\mathbf{x}_0)$ represents the position along the integral curve of (1.1) beginning at \mathbf{x}_0 after time t . The standard intrusive Galerkin ROM based on proper orthogonal decomposition (POD) [39, 20] utilizes a matrix of snapshots $\mathbf{X} = \{\varphi(t_i, \mathbf{x}_0)\}_{i=0}^K \in \mathbb{R}^{N \times K}$, collected offline, to construct the POD basis matrix $\mathbf{U} \in \mathbb{R}^{N \times n}$, $n = 2m$, through the truncated singular value decomposition (SVD) $\mathbf{X} \approx \mathbf{U} \mathbf{\Sigma} \mathbf{V}^\top$. If $\tilde{\mathbf{x}} = \mathbf{U} \hat{\mathbf{x}}$ is an approximate state given by Galerkin projection onto the span of \mathbf{U} , it follows that $\mathbf{U}^\top \dot{\varphi}_t(\mathbf{U} \hat{\mathbf{x}}_0) = \dot{\hat{\mathbf{x}}}$ is the projected vector field, and the differential equation

$$(1.2) \quad \dot{\hat{\mathbf{x}}} = \mathbf{U}^\top \mathbf{J} \nabla H(\mathbf{U} \hat{\mathbf{x}}), \quad \hat{\mathbf{x}}_0 = \mathbf{U}^\top \mathbf{x}_0,$$

defines the natural Galerkin ROM corresponding to (1.1). While this projection accomplishes the primary goal of restricting the dynamics to a low-dimensional subspace of the FOM data, it is easy to check that the Hamiltonian structure guaranteeing Lyapunov stability of the system is not preserved through this procedure since $(\mathbf{U}^\top \mathbf{J})^\top = -\mathbf{J} \mathbf{U} \neq -\mathbf{U}^\top \mathbf{J}$. This has obvious negative consequences: Besides the loss of long-time stability, there is no guarantee of energy conservation or symplectic structure anymore. To see this, let $\hat{H} = H \circ \tilde{\mathbf{x}}$ be the reduced Hamiltonian. Then $\nabla \hat{H} = \mathbf{U}^\top \nabla H$, and so the instantaneous energy satisfies

$$\dot{\hat{H}} = \dot{\hat{\mathbf{x}}} \cdot \nabla \hat{H} = \mathbf{U}^\top \mathbf{J} \nabla H \cdot \mathbf{U}^\top \nabla H = \mathbf{P}_U \nabla H \cdot \mathbf{J} \nabla H,$$

which likely does not vanish in view of the projection $\mathbf{P}_U := \mathbf{U} \mathbf{U}^\top$ onto the column span of \mathbf{U} . Additionally, the reduced dynamics are no longer symplectic: there is no closed and nondegenerate differential two-form preserved along the flow, meaning that volumes in phase space are unnaturally distorted by this evolution. It has been shown that this lack of Hamiltonian structure in the straightforward Galerkin ROM leads to low-order systems which are unacceptably sensitive and cannot forecast beyond their training range, even in the case of periodic dynamics [29, 15, 17].

This unfortunate deficiency of projection-based ROMs in the Hamiltonian case has been well studied, and several solutions are available at present, each with their own benefits and drawbacks. One popular approach involves utilizing a symplecticity-preserving approach such as the proper symplectic decomposition (PSD) to compute the reduced basis \mathbf{U} , which includes basis algorithms such as the cotangent lift and complex SVD [29, 19]. In each case, the basis $\mathbf{U} \in \mathbb{R}^{N \times n}$ is constructed to satisfy the equivariance condition $\mathbf{U}^\top \mathbf{J} = \mathbf{J}_n \mathbf{U}^\top$, where $\mathbf{J}_n \in \mathbb{R}^{n \times n}$ is the canonical symplectic matrix in n dimensions [22]. This admits the pleasing interpretation of projecting the canonical Hamiltonian FOM (1.1) directly to a canonical Hamiltonian ROM, which is obviously property preserving. It also enables the direct construction of this Hamiltonian ROM through straightforward Galerkin projection since the evolution of the reduced state becomes

$$(1.3) \quad \dot{\hat{\mathbf{x}}} = \mathbf{U}^\top \dot{\tilde{\mathbf{x}}} = \mathbf{U}^\top \mathbf{J} \nabla H(\tilde{\mathbf{x}}) = \mathbf{J}_n \mathbf{U}^\top \nabla H(\mathbf{U} \hat{\mathbf{x}}) = \mathbf{J}_n \nabla \hat{H}(\hat{\mathbf{x}}),$$

which is a low-order Hamiltonian system in canonical form. Sometimes called the symplectic lift [29], this technique effectively converts the problem of preserving Hamiltonian structure in projection-based ROM with constructing an informative \mathbf{J} -equivariant reduced basis

$\mathbf{U} \in \mathbb{R}^{N \times n}$ for the dynamics. However, as useful as the symplectic lift is for constructing Hamiltonian ROMs, it breaks down when the bases used are not informative enough to accurately represent system dynamics. This is particularly noticeable in realistic mechanics simulations: when the values \mathbf{q}, \mathbf{p} of position and momentum are separated by orders of magnitude, the basis \mathbf{U} resulting from the above procedures is often not informative enough to produce a useful Hamiltonian ROM of the form (1.3), and this effect persists even when these scales are equalized in the training data (cf. section 5.2). Moreover, the symplectic lift is only feasible for canonical Hamiltonian systems where the symplectic matrix $\mathbf{J} \neq \mathbf{J}(\mathbf{x})$ is a known constant, which limits its applicability to many systems of practical interest [17].

The issues with constructing the equivariant basis necessary for symplectic lift described above led to another line of work involving ROMs which preserve the Hamiltonian of structure of (1.1) [15]. In this case, the goal was to allow the use of an arbitrary reduced basis \mathbf{U} so that the potentially high approximation error of the symplectic lift could be sidestepped and Hamiltonian ROMs could be extended to the noncanonical case. This led to a nice strategy in [15] based on the observation that the equivariance condition $\mathbf{U}^\top \mathbf{J} = \mathbf{J}_n \mathbf{U}^\top$ is sufficient but not necessary for a Hamiltonian ROM. Treating this condition instead as an overdetermined system in the reduced symplectic matrix $\hat{\mathbf{J}}$, the skew-symmetric least-squares solution $\hat{\mathbf{J}} = \mathbf{U}^\top \mathbf{J} \mathbf{U}$ leads directly to the ROM

$$(1.4) \quad \dot{\hat{\mathbf{x}}} = \mathbf{U}^\top \mathbf{J} \mathbf{U} \mathbf{U}^\top \nabla H(\mathbf{U} \hat{\mathbf{x}}) = \hat{\mathbf{J}} \nabla \hat{H}(\hat{\mathbf{x}}), \quad \hat{\mathbf{x}}_0 = \mathbf{U}^\top \mathbf{x}_0.$$

Since it is both energy-conserving and symplectic, this has the advantage of being a valid, fully determined Hamiltonian ROM for almost all bases \mathbf{U} (cf. Proposition 2.2). These properties follow since the instantaneous energy satisfies

$$\dot{\hat{H}} = \dot{\hat{\mathbf{x}}} \cdot \nabla \hat{H} = \hat{\mathbf{J}} \nabla \hat{H} \cdot \nabla \hat{H} = -\nabla \hat{H} \cdot \hat{\mathbf{J}} \nabla \hat{H} = 0,$$

and the skew-symmetric matrix $\hat{\mathbf{J}}$ is constant in \mathbf{x} , meaning it satisfies the Jacobi identity equivalent to closure of the governing symplectic two-form. The latter condition ensures that volume densities are instantaneously conserved in phase space and that no artificial dissipation can be introduced into the scheme (cf. [7]). While the ROM (1.4) is not canonical unless \mathbf{U} is \mathbf{J} -equivariant, it is in some sense the reduced-order Hamiltonian system nearest to (1.2) for a given \mathbf{U} , making it advantageous as a structure-preserving alternative. Moreover, the least-squares technique generating (1.4) extends directly to noncanonical Hamiltonian systems since it does not require the construction of a special reduced basis. On the other hand, there is clearly an additional column space projection $\mathbf{P}_U = \mathbf{U} \mathbf{U}^\top$ in (1.4) which arises “extraneously” via the least-squares projection of \mathbf{J} on \mathbf{U} . It turns out that this additional projection in the reduced vector field $\dot{\hat{\mathbf{x}}}$ is enough to significantly harm ROM accuracy in some cases of interest, which will be seen in section 5.2. Fundamentally, this occurs because the ROM (1.4) fails to be variationally consistent with the original Hamiltonian FOM: the Hamiltonian variational principle obeyed by the reduced state $\hat{\mathbf{x}}$ is not the Galerkin projection of the Hamiltonian variational principle obeyed by the original state \mathbf{x} (see Figure 1), and therefore the Galerkin projection of solutions to (1.1) fails to satisfy the reduced equations (1.4). This necessitates a different model, developed in sections 2 and 3, which does not require inconsistent modifications to produce a reduced-order Hamiltonian vector field.

1.1. Related work. Before discussing the contributions of this work in detail, it is prudent to mention several other approaches recently proposed for maintaining Hamiltonian structure in intrusive as well as nonintrusive projection-based ROMs. Beginning with the intrusive case, Sockwell develops a Hamiltonian structure-preserving approach in [40] based on the work of [15], which targets noncanonical systems and includes a useful centering procedure for preserving linear Casimir invariants, such as mass and momentum. Buchfink, Glas, and Haasdonk derive a globally optimal symplectic reduced basis for “canonizable” Hamiltonian systems with a periodic solution in the sense of PSD [6], and further extensions of PSD are offered by Afkham and Hesthaven in [1] and Hesthaven, Pagliantini, and Ripamonti in [19]. More specifically, the authors of [1] propose a greedy extension of symplectic basis generation, which promises to be more cost-effective than traditional POD and exhibits an exponential convergence rate. The related work [19] proposes a dynamical symplectic reduced basis method for low-rank Hamiltonian systems based on evolving the reduced space along a trajectory locally constrained to lie on the tangent space of the high-dimensional dynamics. Beyond purely Hamiltonian dynamics, extensions of structure-preserving model reduction methods to problems with dissipation have been the subject of several recent works, including [2] for the case of Hamiltonian systems with dissipation and [16] for the case of metriplectic systems which combine energy conservation with entropy generation.

In addition to intrusive model reduction approaches, recent years have seen the emergence of data-driven structure-preserving operator inference (OpInf) ROMs. The premise of data-driven OpInf, first introduced by Peherstorfer and Willcox in [28], is that reduced operators in a projection-based ROM can be learned nonintrusively (i.e., without access to the FOM operators or code) given a set of FOM snapshots by solving an optimization problem. In this direction, several recent papers have devised OpInf algorithms which preserve Hamiltonian structure. The first of these works [38], termed H-OpInf, preserves canonical Hamiltonian structure in OpInf ROMs built using a “cotangent lift” (i.e., PSD) reduced basis. This approach is extended to quadratic reduced bases in [36] toward improving accuracy and efficiency for problems exhibiting a slowly decaying Kolmogorov n -width [30]. Again focusing on canonical Hamiltonian systems, this second approach features a quadratic manifold cotangent lift algorithm for reduced basis construction as well as a blockwise-quadratic approximation of each of the states, \mathbf{q} and \mathbf{p} . This line of work is extended in [17] to accommodate both canonical and the more general case of noncanonical Hamiltonian systems, enabling linear Casimir-preserving Hamiltonian ROMs built with any choice of linear reduced basis. Further extensions of Hamiltonian structure-preserving nonintrusive model reduction methods to nonlinear manifold ROMs using weakly symplectic autoencoders are developed in [5], and an approach for applying OpInf to nonlinear Hamiltonian problems with polynomial nonlinearities is developed in [4]. A more general approach for building structure-preserving ROMs for problems with symmetric negative-definite operators and linear or quadratic nonlinearities is presented in [33], where a Tikhonov regularization-inspired approach is introduced to improve the stability of the underlying optimization problem. Further extensions to nonlinear Hamiltonian systems are explored in [41], which hypothesizes that general Hamiltonian systems can be transformed to higher-dimensional nonlinear systems with cubic Hamiltonians through a lifting procedure [31]. Finally, a generalization of [38] to systems that are either conservative or dissipative is proposed in [14], which focuses on the preservation of gradient structure in

OpInf ROMs. Notably, the authors of this work also derive an a priori error estimator for their OpInf ROMs which inspired Theorem 4.4 presented here.

Before concluding this literature overview, it is worth noting that (under mild regularity conditions) the canonical Hamiltonian formulation (1.1) has an analogous Lagrangian formulation in which the first-order-in-time system (1.1) for the state $\mathbf{x} = (\mathbf{q} \ \mathbf{p})^\top \in \mathbb{R}^{2M}$ consisting of positions and momenta is replaced with an equivalent second-order-in-time system of so-called Euler–Lagrange equations for the unknown positions $\mathbf{q} \in \mathbb{R}^M$. For linear systems with no damping, the Euler–Lagrange equations take the form

$$(1.5) \quad M\ddot{\mathbf{q}} + K\mathbf{q} = \mathbf{f},$$

where $M, K \in \mathbb{R}^{M \times M}$ are the symmetric positive definite mass and stiffness matrices respectively, and $\mathbf{f} \in \mathbb{R}^M$ is an external force vector. The Lagrangian formulation is preferred in some applications, e.g., solid dynamics, since it contains half as many degrees of freedom relative to the analogous Hamiltonian formulation; hence, a number of researchers have also developed model reduction methods that preserve Lagrangian structure. Remarkably, Lall, Krysl, and Marsden showed in [21] that performing a Galerkin projection of the second-order-in-time Euler–Lagrange equations preserves Lagrangian structure for linear and nonlinear problems provided that no hyperreduction is employed, making structure preservation simpler in this case.¹ Approaches for preserving Lagrangian structure in OpInf models, termed L-OpInf ROMs, have been considered recently by Sharma and Kramer in [35] and Sharma et al. in [37]. While the basic premise of [35] is to formulate and solve a constrained optimization problem for the reduced operators \hat{M} and \hat{K} , [37] proposes a two-step approach that first learns the best-fit linear Lagrangian ROM via Lagrangian operator inference and subsequently applies a structure-preserving machine learning method to learn nonlinearities in the reduced space.

Remark 1.1. While continuous models predicated on Hamiltonian/Lagrangian variational principles are automatically conservative, practical FOMs/ROMs need to be advanced with a symplectic time integrator in order to be energy-conserving along their trajectories. This work employs the Newmark–Beta integrator for Lagrangian systems (1.5) and the implicit midpoint method (trapezoid rule) for Hamiltonian systems (1.1), both of which ensure that energy is conserved at the fully discrete level.

1.2. Contributions. The present work aims to combine the variational consistency of the \mathbf{J} -equivariant PSD approach (1.3) with the flexibility of the least-squares approach (1.4) to yield a Hamiltonian ROM which outperforms previous methods on moderate-to-large scale Hamiltonian problems of interest. In particular, the goal is to develop a Hamiltonian ROM which is variationally consistent with the governing Hamiltonian FOM (1.1), which allows for the use of any desired trial space basis \mathbf{U} , and which can be deployed in both intrusive and nonintrusive settings. This necessitates a new technique, presented here and applicable to

¹For maintaining Lagrangian and Hamiltonian structure in ROMs for nonlinear Lagrangian systems, several structure-preserving hyperreduction approaches have been developed, e.g., reduced basis sparsification and matrix gappy POD [8], energy-conserving sampling and weighting [13], and the symplectic discrete empirical interpolation method [29]; as this work focuses on linear problems, these methods are not discussed at length here.

ROMs constructed using both intrusive and nonintrusive (OpInf) methods. Its core idea is simple: while the trial space basis should be treated as fixed for optimal dimension reduction, there is flexibility in the space of test functions used to generate the ROM. It is shown that an inspired choice of test space yields a Hamiltonian ROM which is valid for any canonical Hamiltonian FOM and requires nothing more than Petrov–Galerkin orthogonality. Specifically, this work contributes the following:

- A procedure to generate variationally consistent Hamiltonian ROMs from canonical Hamiltonian FOMs with minimal assumptions on the trial space basis, applicable in both intrusive and nonintrusive settings.
- Interpretable error estimates showing the decomposition of ROM error into two distinct parts: basis projection error and deviation-from-canonicity error.
- Seamless incorporation of existing ROM technology, such as a centering procedure which preserves linear invariant quantities as well as reprojection which ensures that nonintrusive OpInf ROMs preasymptotically recover the performance of their intrusive counterparts.
- Three-dimensional solid mechanical examples from linear elasticity showing improvements over existing state-of-the-art methods for both intrusive and nonintrusive OpInf ROMs.

The remainder of this paper is organized as follows. Section 2 describes the core framework, leading to a novel variationally consistent intrusive Hamiltonian ROM which accommodates any desired linear reduced basis. Section 3 extends this approach to nonintrusive OpInf ROMs, culminating in a straightforward algorithm for inferring the relevant operators from snapshot data. An error analysis of the proposed intrusive and nonintrusive ROMs is performed in section 4, where interpretable error estimates for these approaches are derived and discussed. The numerical performance of the proposed ROMs is further examined in section 5, focused on realistic examples from linear elasticity. Finally, a concluding summary and a discussion of future work is presented in section 6.

2. Variationally consistent intrusive Hamiltonian ROM. The present approach is motivated by the standard coordinate-free formulation of Hamiltonian systems, of which more details can be found in, e.g., [22]. Recall that a canonical Hamiltonian system can be alternatively described as a symplectic gradient flow, defined in terms of the symplectic structure on the even-dimensional \mathbb{R}^N ($N = 2M$) induced by a canonical closed and nondegenerate two-form $\omega \in \bigwedge^2 \mathbb{R}^N$. Particularly, (1.1) is equivalent to $\dot{\mathbf{x}} = \mathbf{X}_H$, where \mathbf{X}_H is the symplectic gradient of H defined through the differential of the Hamiltonian function, $dH = \mathbf{X}_H \lrcorner \omega := \omega(\mathbf{X}_H, -)$, which is a mapping from tangent vector fields on \mathbb{R}^N to tangent vector fields on \mathbb{R} . To see why this is an equivalence, notice that the negative of the canonical Kähler structure on \mathbb{R}^N induces a smooth matrix field $\mathbf{J} \in \mathbb{R}^{N \times N}$ satisfying $\omega(\mathbf{v}, \mathbf{w}) = \mathbf{v} \cdot \mathbf{J} \mathbf{w}$ for any $\mathbf{v}, \mathbf{w} \in T_{\mathbf{x}} \mathbb{R}^N \simeq \mathbb{R}^N$ at any point $\mathbf{x} \in \mathbb{R}^N$, which coincides with the constant symplectic matrix \mathbf{J} in (1.1). This implies that for any $\mathbf{v} \in \mathbb{R}^N$, the differential of the Hamiltonian satisfies

$$dH(\mathbf{v}) = \mathbf{v} \cdot \nabla H = \mathbf{X}_H \cdot \mathbf{J} \mathbf{v} = \mathbf{v} \cdot \mathbf{J}^\top \mathbf{X}_H,$$

and therefore the symplectic gradient \mathbf{X}_H is given by

$$\mathbf{X}_H = \mathbf{J}^{-\top} \nabla H = \mathbf{J} \nabla H,$$

confirming that the evolution of the state is given by $\dot{\mathbf{x}} = \mathbf{X}_H = \mathbf{J}\nabla H$ in agreement with (1.1).

Remark 2.1. Note that here, the implied definiteness condition $\omega(\mathbf{v}, \mathbf{J}\mathbf{v}) = -|\mathbf{v}|^2 < 0$ for vectors $\mathbf{v} \neq \mathbf{0}$ is opposite to what is required for the canonical Kähler form on \mathbb{R}^N . Since either choice leads to the same mathematical structure, this choice is made so that signs are consistent with the FOM (1.1).

While the expression $\dot{\mathbf{x}} = \mathbf{X}_H$ in terms of the symplectic gradient \mathbf{X}_H may look like a straightforward rewrite of the FOM (1.1), the advantage of this expression lies in its ability to convey a simple requirement for any Hamiltonian ROM derived from a canonical Hamiltonian FOM. Particularly, it follows that a reduced Hamiltonian ROM in terms of the approximate state $\tilde{\mathbf{x}} = \mathbf{U}\hat{\mathbf{x}}$ must obey the low-order evolution equation

$$\dot{\tilde{\mathbf{x}}} = \mathbf{X}_{\hat{H}} = \hat{\mathbf{J}}^{-\top} \nabla \hat{H},$$

involving the reduced symplectic gradient $\mathbf{X}_{\hat{H}} \in \mathbb{R}^n$, $n = 2m \ll N$, corresponding to some reduced Hamiltonian function $\hat{H} : \mathbb{R}^n \rightarrow \mathbb{R}$ and some closed and nondegenerate symplectic two-form $\hat{\omega} \in \bigwedge^2 \mathbb{R}^n$ defining the reduced symplectic matrix $\hat{\mathbf{J}} = \mathbf{U}^\top \mathbf{J} \mathbf{U} \in \mathbb{R}^{n \times n}$. Importantly, notice that this reduced symplectic form $\hat{\omega}$ need not coincide with the canonical symplectic form on \mathbb{R}^n so that $\hat{\mathbf{J}} = \mathbf{B}^\top \mathbf{J}_n \mathbf{B}$ can be conjugate to the canonical symplectic matrix $\mathbf{J}_n \in \mathbb{R}^{n \times n}$ by an *unknown* invertible matrix $\mathbf{B} \in \mathbb{R}^{n \times n}$. In particular, this means that the transpose of the reduced symplectic matrix $\hat{\mathbf{J}}^\top = -\hat{\mathbf{J}}$ may not be the inverse of $\hat{\mathbf{J}}$ anymore, i.e., $\hat{\mathbf{J}}^\top \neq \hat{\mathbf{J}}^{-1}$. This represents the fact that the reduced symplectic form $\hat{\omega}$ and hence the reduced symplectic matrix $\hat{\mathbf{J}}$ are allowed to be noncanonical on the reduced space \mathbb{R}^n . This contrasts significantly with the case of \mathbf{J} -equivariant reduced bases \mathbf{U} , where $\hat{\mathbf{J}} = \mathbf{J}_n$, and so the canonicity condition $\hat{\mathbf{J}}^{-1} = \hat{\mathbf{J}}^\top$ is satisfied by construction.

In view of this discussion, there are two main ingredients to the variationally consistent Hamiltonian ROM presented here. The first and most essential is a way to manufacture a reduced symplectic gradient $\mathbf{X}_{\hat{H}} \in \mathbb{R}^n$ from the full-order symplectic gradient $\mathbf{X}_H \in \mathbb{R}^N$ using only projection onto orthonormal reduced bases. This will be accomplished via an appropriately chosen Petrov–Galerkin projection, which effectively replaces the symplectic matrix \mathbf{J} on the right-hand side of (1.1) with a reduced inverse-transpose $\hat{\mathbf{J}}^{-\top}$, making the resulting ROM variationally consistent with the FOM by construction. This is remarkably different from the case of standard Galerkin projection onto the trial space spanned by the columns of \mathbf{U} , where the resulting ROM will not be Hamiltonian unless the reduced basis \mathbf{U} is \mathbf{J} -equivariant (cf. Figure 1). The second ingredient is a useful but nonessential centering procedure in the trial space basis \mathbf{U} , which is commonly used in POD ROMs and ensures that the value of the Hamiltonian function H is correct at the beginning of the ROM trajectory and hence will remain correct throughout its evolution. While this particular centering has previously been used only for intrusive ROMs, its applicability is extended to the nonintrusive OpInf case here (see section 3.3), which may be of independent interest. Combined with the idea of reprojection [27], this enables a nonintrusive Hamiltonian OpInf ROM that preasymptotically recovers the intrusive ROM, enabling its use in the case that only snapshot data are available.

2.1. Variationally consistent Petrov–Galerkin projection. The particular basis projection employed here exploits the fact that the symplectic structure of canonical Hamiltonian




Figure 1. A diagram illustrating the behavior of Hamiltonian/Lagrangian FOMs under Galerkin projection. While (canonical) Hamiltonian and Lagrangian systems which are “regular” are equivalent under the Legendre transformation (L.T.) at both the FOM and the ROM levels, this equivalence does not commute with Galerkin projection onto a general reduced basis $\mathbf{U} \in \mathbb{R}^{N \times n}$. Conversely, the variationally consistent ROM presented here projects Hamiltonian FOMs to Hamiltonian ROMs in every case.

systems is nondegenerate everywhere, meaning that the symplectic matrix \mathbf{J} has full matrix rank. Therefore, given any trial space basis $\mathbf{U} \in \mathbb{R}^{N \times n}$, the columns of the “rotated” matrix \mathbf{JU} form a valid test space basis of rank n , meaning that the Petrov–Galerkin projection $(\mathbf{JU})^\top \dot{\boldsymbol{\varphi}}_t(\mathbf{U}\hat{\mathbf{x}}) = \mathbf{U}^\top \mathbf{J}^\top \mathbf{U} \dot{\hat{\mathbf{x}}} = \hat{\mathbf{J}}^\top \dot{\hat{\mathbf{x}}}$ leads to a valid Hamiltonian ROM provided that the reduced symplectic matrix $\hat{\mathbf{J}} = \mathbf{U}^\top \mathbf{JU}$ remains full rank. Noting that $\mathbf{J}^\top \mathbf{J} = -\mathbf{J}^2 = \mathbf{I}$, this ROM is expressible as

$$(2.1) \quad \hat{\mathbf{J}}^\top \dot{\hat{\mathbf{x}}} = \mathbf{U}^\top \nabla H(\mathbf{U}\hat{\mathbf{x}}) = \nabla \hat{H}(\hat{\mathbf{x}}), \quad \hat{\mathbf{x}}_0 = \mathbf{U}^\top \mathbf{x}_0,$$

which is equivalent, at least formally, to the expression $\dot{\hat{\mathbf{x}}} = \hat{\mathbf{J}}^{-\top} \nabla \hat{H}(\hat{\mathbf{x}})$ from before. However, it remains to be sure that the differential equation (2.1) uniquely specifies the evolution of the reduced-order state $\hat{\mathbf{x}}$ since the full rank condition on $\hat{\mathbf{J}}$ is not obvious. The following result shows that this condition is indeed satisfied for almost every orthonormal trial space basis \mathbf{U} of even dimension that may be constructed.

Proposition 2.2. *The reduced-order symplectic matrix $\hat{\mathbf{J}} = \mathbf{U}^\top \mathbf{JU} \in \mathbb{R}^{2m \times 2m}$ is invertible for almost every column-orthonormal reduced basis $\mathbf{U} \in \mathbb{R}^{2M \times 2m}$, $\mathbf{U}^\top \mathbf{U} = \mathbf{I}$.*

Proof. Let $N = 2M$ and $n = 2m$. Denote by S the subset of the Stiefel manifold $V_n(\mathbb{R}^N) = \{\mathbf{U} \in \mathbb{R}^{N \times n} \mid \mathbf{U}^\top \mathbf{U} = \mathbf{I}\}$ of orthonormal n -frames in \mathbb{R}^N satisfying $\det \hat{\mathbf{J}} = \det(\mathbf{U}^\top \mathbf{JU}) \neq 0$ for all $\mathbf{U} \in S$, i.e., $S = \{\mathbf{U} \in V_n(\mathbb{R}^N) \mid \det \hat{\mathbf{J}} \neq 0\}$. This subset is Zariski-open since $\det \hat{\mathbf{J}}$ is a polynomial in the entries of \mathbf{U} and nonempty since $\mathbf{U} = (\mathbf{e}_1 \ \dots \ \mathbf{e}_m \ \mathbf{e}_{M+1} \ \dots \ \mathbf{e}_{M+m})$ satisfies $\det \hat{\mathbf{J}} = \det \mathbf{J}_n = 1 \neq 0$, where \mathbf{e}_i denotes the i th standard basis vector in \mathbb{R}^{2M} . It follows that S is dense in the Stiefel manifold $V_n(\mathbb{R}^N)$, and since polynomials are real-analytic and hence cannot vanish on a set of positive measure without being identically zero, it must be that S has full Lebesgue measure in this space. Therefore, almost every reduced basis $\mathbf{U} \in V_n(\mathbb{R}^N)$ lies in S , as claimed. ■

Remark 2.3. Bases \mathbf{U} for vector subspaces on which the reduced symplectic form $\hat{\omega}$ pulls back to be identically zero, e.g., $\mathbf{U} = (\mathbf{e}_1 \ \dots \ \mathbf{e}_m)$, are called isotropic and can be characterized by the vanishing of the reduced symplectic matrix $\hat{\mathbf{J}} \equiv \mathbf{0}$. More generally, bases \mathbf{U} for vector subspaces on which the determinant $\det \hat{\mathbf{J}} = 0$ is zero cannot be symplectic themselves since this implies that $\hat{\mathbf{J}}$ has a nontrivial kernel. Therefore, these subspaces must

be excluded from consideration when building any canonical Hamiltonian ROM. On the other hand, Proposition 2.2 shows that, given a reduced basis \mathbf{U} , the reduced symplectic matrix $\hat{\mathbf{J}}$ will be invertible with probability one, which is likely why this technical point is not considered in previous work on Hamiltonian ROMs.

Notice that the proposed ROM (2.1) is both energy-conserving and symplectic: The reduced symplectic matrix $\hat{\mathbf{J}}$ is obviously constant in $\hat{\mathbf{x}}$, implying closure in the reduced symplectic 2-form $\hat{\omega}(\hat{\mathbf{v}}, \hat{\mathbf{w}}) = \hat{\mathbf{v}} \cdot \hat{\mathbf{J}} \hat{\mathbf{w}}$, and the skew symmetry in $\hat{\mathbf{J}}^{-\top} = -\hat{\mathbf{J}}^{-1}$ implies that the instantaneous Hamiltonian satisfies

$$\dot{\hat{\mathbf{H}}} = \dot{\hat{\mathbf{x}}} \cdot \nabla \hat{\mathbf{H}} = \hat{\mathbf{J}}^{-\top} \nabla \hat{\mathbf{H}} \cdot \nabla \hat{\mathbf{H}} = -\nabla \hat{\mathbf{H}} \cdot \hat{\mathbf{J}}^{-\top} \nabla \hat{\mathbf{H}} = 0.$$

Remarkably, this variationally consistent ROM is also noncanonical unless $\mathbf{U}^\top \mathbf{J} = \hat{\mathbf{J}} \mathbf{U}^\top$, in which case $\mathbf{U}^\top \mathbf{J} \mathbf{U} = \mathbf{J}_n$ and (2.1) reduces to the same canonical Hamiltonian ROM (1.3) produced by a PSD or symplectic lift. This use of a noncanonical reduced symplectic matrix $\hat{\mathbf{J}} = \mathbf{U}^\top \mathbf{J} \mathbf{U}$ resulting from an arbitrary basis \mathbf{U} is similar to the approach in [15] leading to the ROM (1.4), although now it is apparent that the direct use of the symplectic matrix $\hat{\mathbf{J}}$ in defining the reduced vector field $\dot{\hat{\mathbf{x}}}$ is kinematically inconsistent with the FOM; what is instead needed for consistency is its inverse-transpose $\hat{\mathbf{J}}^{-\top}$, which only reduces to $\hat{\mathbf{J}}$ in the special case of a symplectic lift. The form of the ROM (2.1) also highlights the remarkable fact that a canonical Hamiltonian ROM is not always appropriate even when reducing a canonical Hamiltonian FOM, meaning that the flexibility $\hat{\mathbf{J}} = \mathbf{B}^\top \mathbf{J}_n \mathbf{B}$ of conjugation by an invertible matrix \mathbf{B} afforded in the reduced symplectic matrix may actually increase expressiveness in the resulting Hamiltonian ROM.

Remark 2.4. When the FOM (1.1) is equivalently described as $\mathbf{J}^\top \dot{\mathbf{x}} = \nabla H(\mathbf{x})$, again using the nondegeneracy of the symplectic structure, there is an interesting dual interpretation of the ROM (2.1) as the natural Galerkin ROM with trial/test space given by the image of \mathbf{U} and denoted by $\text{im } \mathbf{U}$. This again displays its variational consistency since the Hamiltonian ROM (2.1) is derived directly from a Hamiltonian FOM through a straightforward basis projection.

2.2. Intrusive centering. At this point, a variationally consistent intrusive Hamiltonian ROM (2.1) has been presented which is valid for almost any orthonormal reduced basis \mathbf{U} and is guaranteed to conserve the energy level $\hat{H}(\hat{\mathbf{x}}_0) = H(\mathbf{P}_U \mathbf{x}_0)$ exactly. Since $H(\mathbf{P}_U \mathbf{x}_0)$ is not necessarily equal to the true energy level $H(\mathbf{x}_0)$ of the corresponding FOM trajectory, it can be beneficial to introduce a well-known centering procedure which ensures that this discrepancy is eliminated. While this centering will be formulated in terms of the initial state \mathbf{x}_0 , note that all benefits of the centering discussed here remain valid when $\bar{\mathbf{x}}$ is a linear combination of state snapshots $\{\mathbf{x}_i\}$, say, the mean initial condition along several trajectories. This can be beneficial when perturbations from the mean are easier to resolve than perturbations from any particular trajectory.

In the intrusive case, the technique of using a centered POD basis for ROM development is well known and has many benefits. In [18, 3, 10, 9], it is used to enforce inhomogeneous Dirichlet boundary conditions in Galerkin ROMs, while in [40, 15, 17], it is used for the purpose of preserving linear invariants in Hamiltonian ROMs. The main idea at work is the inclusion of an affine shift in the basis \mathbf{U} so that it becomes “centered around zero” in whatever way is

appropriate for a given application. This is most easily accomplished by choosing \mathbf{U} through the truncated SVD of a centered snapshot matrix $\mathbf{X} - \bar{\mathbf{X}} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$, where $\bar{\mathbf{X}}$ is a matrix each column of which is a judiciously chosen constant vector $\bar{\mathbf{x}}$. In the case of Hamiltonian systems, it is advantageous to choose this centering vector $\bar{\mathbf{x}} = \mathbf{x}_0$ to be the initial condition for the Hamiltonian flow, as this ensures that the ROM approximation $\hat{\mathbf{x}} = \mathbf{x}_0 + \mathbf{U}\hat{\mathbf{x}}$ satisfies the zero initial condition $\hat{\mathbf{x}}_0 = \mathbf{0}$ and maintains the true initial energy value $\hat{H}(\hat{\mathbf{x}}_0) = H(\mathbf{x}_0)$ along the entire ROM trajectory.

Before moving to the discussion of OpInf ROMs, it is worth mentioning another benefit of centered POD important for noncanonical Hamiltonian systems $\dot{\mathbf{x}} = \mathbf{J}(\mathbf{x})\nabla H(\mathbf{x})$, where $\mathbf{J}(\mathbf{x})$ may have a nontrivial kernel. In [40], it was observed that centering of the reduced basis \mathbf{U} by the initial condition \mathbf{x}_0 has the added benefit of preserving linear invariant quantities of a Hamiltonian FOM system, known as linear Casimir invariants (or Casimirs), in a corresponding Hamiltonian ROM. To see why this is the case, first notice that a linear invariant can be written as $C(\mathbf{x}) = \mathbf{c}^\top \mathbf{x}$, where $\mathbf{c} = \nabla C \in \mathbb{R}^N$ is a constant vector. It follows that the Casimir

$$C(\mathbf{x}_0 + \mathbf{U}\hat{\mathbf{x}}) = \mathbf{c}^\top \mathbf{x}_0 + \mathbf{c}^\top \mathbf{U}\hat{\mathbf{x}}$$

will be conserved at the value $\mathbf{c}^\top \mathbf{x}_0$ along the ROM trajectory $\hat{\mathbf{x}}$ provided that the POD basis \mathbf{U} is constructed so that $\mathbf{c}^\top \mathbf{U}\hat{\mathbf{x}} = 0$. Moreover, this will automatically be true when \mathbf{U} is defined through the truncated SVD $\mathbf{X} - \mathbf{X}_0 \approx \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$: Since each column \mathbf{u}_i of \mathbf{U} takes the form $\mathbf{u}_i = (\mathbf{X} - \mathbf{X}_0)\sigma_i^{-1}\mathbf{v}_i$ in terms of a corresponding singular value σ_i and column \mathbf{v}_i of \mathbf{V} , it follows that

$$\mathbf{c}^\top \mathbf{U}\hat{\mathbf{x}} = \sum_{i=1}^n \hat{x}^i \mathbf{c}^\top \mathbf{u}_i = \sum_{i=1}^n \hat{x}^i \mathbf{c}^\top (\mathbf{X} - \mathbf{X}_0) \sigma_i^{-1} \mathbf{v}_i = 0,$$

where \hat{x}^i denotes the i th component of the reduced state $\hat{\mathbf{x}}$ and the final equality uses the fact that $\mathbf{c}^\top (\mathbf{X} - \mathbf{X}_0) = \mathbf{0}$ since all columns of the snapshot matrix \mathbf{X} lie on the same Casimir-preserving trajectory as \mathbf{x}_0 . While not immediately relevant for the ROMs formulated here, this useful fact enables Casimir-preserving schemes which are essential for realistic physics in more complicated Hamiltonian systems, the variationally consistent model reduction of which is the subject of future work.

Remark 2.5. To improve expressivity of the basis \mathbf{U} , it is occasionally useful to include additional snapshots in the matrix \mathbf{X} , such as snapshots of ∇H . When linear Casimir preservation is important, care has to be taken to ensure that these added snapshots still satisfy the conservation condition $\mathbf{c}^\top \mathbf{U}\hat{\mathbf{x}} = 0$. One simple way to accomplish this is to project out the component of ∇H in the direction of \mathbf{c} so that $\nabla^\perp H = \mathbf{P}_c^\perp \nabla H = \left(\mathbf{I} - \frac{\mathbf{c}\mathbf{c}^\top}{|\mathbf{c}|^2}\right) \nabla H$ is included in the snapshot matrix instead of ∇H (cf. [40]). This is equivalent to using the natural (i.e., Riemannian) gradient of the Hamiltonian H , which is tangent to the linear subspace $\{\mathbf{x} \mid \mathbf{c}^\top \mathbf{x} = 0\} \subset \mathbb{R}^N$.

3. Variationally consistent Hamiltonian operator inference. Recall that operator inference (OpInf) [28] is a useful procedure by which the operators governing a POD ROM are learned directly from snapshot data for the FOM system, ensuring that the resulting ROM is

solvable without intrusion into the FOM code. To apply the proposed variationally consistent Hamiltonian ROM (2.1) in nonintrusive OpInf settings, it is necessary to understand which challenges with OpInf for Hamiltonian systems have already been addressed and which remain to be overcome. Here, some previous work on Hamiltonian OpInf is recalled, and an algorithm is presented which enables the use of the variationally consistent approach presented in section 2 in cases where only snapshot data are available.

3.1. OpInf for canonical Hamiltonian systems. Both of the Hamiltonian ROMs (1.4) and (2.1) can be effectively treated with OpInf techniques provided that the nonlinear part of ∇H is accessible online. In many examples of interest, this is not a large restriction since the relevant nonlinearities are vectorized over each degree of freedom separately (e.g., polynomials, $\sin(\mathbf{x})$) and can be easily simulated without understanding the underlying discretization or access to the FOM code. Conversely, the linear part of ∇H is heavily discretization-dependent and generally impossible to simulate without understanding of the full-order scheme or access to its code. To see how a Hamiltonian OpInf problem is set up in this case, consider writing the Hamiltonian in a way which separates out its quadratic part,

$$H(\mathbf{x}) = \frac{1}{2} \mathbf{x}^\top \mathbf{A} \mathbf{x} + f(\mathbf{x}),$$

where $\mathbf{A} \in \mathbb{R}^{N \times N}$ is constant in \mathbf{x} and $f(\mathbf{x}) : \mathbb{R}^N \rightarrow \mathbb{R}$ is a nonquadratic and likely nonlinear function of the state \mathbf{x} . In this form, the gradient of the Hamiltonian $\nabla H(\mathbf{x}) = \mathbf{A} \mathbf{x} + \nabla f(\mathbf{x})$ separates as well so that if the nonlinear part ∇f of ∇H is accessible, then the only quantity to be inferred is \mathbf{A} . Accomplishing this inference in a way which is compatible with the least-squares Hamiltonian ROM (1.4) has already been accomplished in [38, 17], so it remains to discuss this technique with the aim of adapting it to the case of the variationally consistent Petrov–Galerkin ROM (2.1) from section 2.

Recalling the canonical Hamiltonian operator inference (C-H-OpInf) procedure in [17], for an uncentered approximation $\tilde{\mathbf{x}} = \mathbf{U} \hat{\mathbf{x}}$ of the state, the goal is to infer the linear part $\hat{\mathbf{A}} = \mathbf{U}^\top \mathbf{A} \mathbf{U}$ of the reduced Hamiltonian gradient $\nabla \hat{H}(\mathbf{x}) = \mathbf{U}^\top \nabla H(\mathbf{U} \hat{\mathbf{x}})$ under the assumption that $\hat{\mathbf{A}} = \hat{\mathbf{A}}^\top$ is symmetric and that the nonlinear part $\nabla \hat{f}(\hat{\mathbf{x}}) = \mathbf{U}^\top \nabla f(\mathbf{U} \hat{\mathbf{x}})$ is either polynomial or can be simulated online. In this case, denoting the reduced state snapshot matrix by $\hat{\mathbf{X}} = \mathbf{U}^\top \mathbf{X}$, the reduced nonlinearity snapshot matrix by $\hat{\nabla} f(\mathbf{X}) = \mathbf{U}^\top \nabla f(\mathbf{X})$, and an approximation to the flow map velocity by $D_t(\mathbf{x}) \approx \dot{\varphi}_t(\mathbf{x}_0) = \dot{\varphi}_0(\mathbf{x})$, the problem to be solved is

$$\bar{\mathbf{A}} = \underset{\hat{\mathbf{A}} \in \mathbb{R}^{n \times n}}{\operatorname{argmin}} \left\| \mathbf{U}^\top D_t(\mathbf{X}) - \hat{\mathbf{J}} \left(\hat{\mathbf{A}} \hat{\mathbf{X}} + \hat{\nabla} f(\mathbf{X}) \right) \right\|_F^2 \quad \text{s.t.} \quad \hat{\mathbf{A}}^\top = \hat{\mathbf{A}},$$

which is convex and has a unique global minimum provided that $\hat{\mathbf{X}}$ has maximal rank (note that $\operatorname{rank} \hat{\mathbf{J}} = n$ by construction). Here, the velocity snapshot matrix $D_t(\mathbf{X})$ is an abuse of notation containing true or approximate time derivatives in each column so that $D_t(\mathbf{x}_k) \approx \dot{\varphi}_0(\mathbf{x}_k)$ at column k . In fact, letting $\hat{\mathbf{F}} = \hat{\nabla} f(\mathbf{X})$ denote the reduced nonlinearity snapshots and $\hat{\mathbf{X}}_t = \mathbf{U}^\top D_t(\mathbf{X})$ denote the approximate velocity field snapshots for brevity, straightforward differentiation from the obvious Lagrangian function shows that this minimization is solvable through the equivalent linear system

$$(3.1) \quad \left(\hat{\mathbf{J}}^\top \hat{\mathbf{J}} \oplus \hat{\mathbf{X}} \hat{\mathbf{X}}^\top \right) \text{vec } \bar{\mathbf{A}} = \text{vec} \left(\hat{\mathbf{J}}^\top \hat{\mathbf{X}}_t \hat{\mathbf{X}}_t^\top + \hat{\mathbf{X}} \hat{\mathbf{X}}_t^\top \hat{\mathbf{J}} - \hat{\mathbf{J}}^\top \hat{\mathbf{J}} \hat{\mathbf{F}} \hat{\mathbf{X}}^\top - \hat{\mathbf{X}} \hat{\mathbf{F}}^\top \hat{\mathbf{J}}^\top \hat{\mathbf{J}} \right),$$

where vec denotes the columnwise vectorization operator, \otimes denotes the Kronecker product of matrices, and $\mathbf{A} \oplus \mathbf{B} = \mathbf{A} \otimes \mathbf{B} + \mathbf{B} \otimes \mathbf{A}$. More details about the Kronecker product and vectorization can be found in, e.g., [34]. Moreover, note that many FOM systems of interest, including those in section 5.2, automatically compute the exact velocities $\dot{\varphi}_0(\mathbf{x})$, which can significantly improve the accuracy of the resulting OpInf ROM (cf. Theorem 4.4).

Remark 3.1. In the case that a symplectic lift basis \mathbf{U} is used, the inference in (3.1) reduces to the original H-OpInf procedure of [38].

Once $\bar{\mathbf{A}} \approx \hat{\mathbf{A}}$ has been inferred by solving (3.1), C-H-OpInf simulates the least-squares Hamiltonian ROM (1.4) online as

$$\dot{\hat{\mathbf{x}}} = \hat{\mathbf{J}} \left(\bar{\mathbf{A}} \hat{\mathbf{x}} + \nabla \hat{f}(\hat{\mathbf{x}}) \right), \quad \hat{\mathbf{x}}_0 = \mathbf{U}^\top \mathbf{x}_0.$$

While this minimally intrusive ROM has been shown in [17] to yield good results in several cases of interest, especially after making the approximation that $\hat{\mathbf{J}}^\top \hat{\mathbf{J}} \approx \mathbf{I}$, which is exact for \mathbf{J} -equivariant bases, it suffers from the same difficulties as the intrusive ROM (1.4), which uses the least-squares solution $\hat{\mathbf{J}} = \mathbf{U}^\top \mathbf{J} \mathbf{U}$, rendering it ineffective on moderate- to large-scale linear elastic problems frequently considered in practice (cf. section 5.2). Even on problems where the least-squares ROM (1.4) works well, the OpInf counterpart using the gradient approximation $\bar{\mathbf{A}} \approx \hat{\mathbf{A}}$ often plateaus in accuracy due to nontrivial optimization error [38, 17]. Thankfully, the work in [27] showed that there is a minimally intrusive way to remove this error almost entirely, which will now be discussed alongside the modifications necessary to develop an OpInf ROM for the variationally consistent ROM (2.1).

3.2. Re-projection. While OpInf is a useful technique for generating nonintrusive reduced models, the original procedure in [28] will generally never preasymptotically recover the intrusive Galerkin ROM. This is a straightforward consequence of how the FOM flow map $\varphi_t(\mathbf{x})$ interacts with Galerkin projection. More precisely, notice that when $\dot{\varphi}_t(\mathbf{x}_0) = \dot{\mathbf{x}}(t) := \mathbf{f}(\mathbf{x}(t))$ describes the FOM dynamics, both the reduced vector field $\mathbf{U}^\top \mathbf{f}(\mathbf{U} \hat{\mathbf{x}})$, corresponding to the (uncentered) intrusive ROM $\dot{\hat{\mathbf{x}}} = \mathbf{U}^\top \mathbf{f}(\mathbf{U} \hat{\mathbf{x}})$, and the \mathbf{U} -image of the reduced flow map $\mathbf{U} \hat{\varphi}_t(\hat{\mathbf{x}}_0) = \mathbf{P}_U \varphi_t(\mathbf{U} \hat{\mathbf{x}}_0)$ contain no component orthogonal to the image of the reduced basis \mathbf{U} . Conversely, integrating the FOM (1.1) for any length of time will almost certainly leave the image of \mathbf{U} even when the initial condition is contained there. This can be seen by letting \mathbf{V} be a basis for the orthogonal complement of \mathbf{U} (so that $\mathbf{U}^\top \mathbf{V} = \mathbf{0}$) and expressing the flow at time t with initial condition $\mathbf{P}_U \mathbf{x}_0$ as

$$\varphi_t(\mathbf{P}_U \mathbf{x}_0) = \mathbf{P}_U \varphi_t(\mathbf{P}_U \mathbf{x}_0) + \mathbf{P}_V \varphi_t(\mathbf{P}_U \mathbf{x}_0).$$

In particular, neither the state $\mathbf{x} = \varphi_t(\mathbf{P}_U \mathbf{x}_0)$ nor its velocity $\dot{\mathbf{x}} = \dot{\varphi}_t(\mathbf{P}_U \mathbf{x}_0)$ lies entirely in the image of \mathbf{U} when $\mathbf{P}_V \varphi_t(\mathbf{P}_U \mathbf{x}_0)$ is nonzero, and this has consequences for standard OpInf.

To see why this is the case, recall that standard OpInf makes use of (approximations to) the velocity data $\mathbf{U}^\top D_t(\mathbf{X}) \approx \mathbf{U}^\top \dot{\varphi}_0(\mathbf{X}) = \mathbf{U}^\top \mathbf{f}(\mathbf{X})$, each column of which looks like $\mathbf{U}^\top \mathbf{f}(\mathbf{x}_k)$ corresponding to the point $t = k\Delta t$ in time. While this quantity lies squarely in the

image of \mathbf{U} , it does not represent the velocity field of any reduced flow map. This can be seen by writing the state \mathbf{x} at any time t as

$$\mathbf{x} = \mathbf{P}_U \mathbf{x} + \mathbf{P}_V \mathbf{x} := \mathbf{U} \hat{\mathbf{x}} + \mathbf{V} \hat{\mathbf{y}}$$

in terms of coefficient vectors $\hat{\mathbf{x}} = \mathbf{U}^\top \mathbf{x}$ and $\hat{\mathbf{y}} = \mathbf{V}^\top \mathbf{x}$. It follows that the vector field $\mathbf{U}^\top \mathbf{f}(\mathbf{x})$ admits the expression

$$\mathbf{U}^\top \mathbf{f}(\mathbf{x}) = \mathbf{U}^\top \mathbf{f}(\mathbf{U} \hat{\mathbf{x}} + \mathbf{V} \hat{\mathbf{y}}) \neq \hat{\mathbf{f}}(\hat{\mathbf{x}}),$$

which is not equal to a low-dimensional vector field $\hat{\mathbf{f}}(\hat{\mathbf{x}})$ for any choice of $\hat{\mathbf{f}}$ due to the presence of $\mathbf{P}_V \mathbf{x}$, the portion of the state \mathbf{x} in the image of \mathbf{V} . This reflects the fact that $\mathbf{P}_U \boldsymbol{\varphi}_{\Delta t}(\mathbf{x}) \neq \mathbf{P}_U \boldsymbol{\varphi}_{\Delta t}(\mathbf{P}_U \mathbf{x})$ at the level of the FOM flow mapping since integration in time is not restricted to the image of \mathbf{U} and also means that a correction to the standard OpInf procedure is required to obtain preasymptotic recovery of the operators governing the intrusive ROM.

To remedy this issue, a minimally intrusive strategy was given in [27], termed “reprojection” and viewed through the lens of the Mori–Zwanzig formalism [26]. The idea behind this approach is to collect snapshot data for the FOM vector field \mathbf{f} at the reprojected states $\mathbf{P}_U \mathbf{X}$ instead of (or in addition to) the original states \mathbf{X} , in which case the “non-Markovian” snapshot matrix $\mathbf{U}^\top D_t(\mathbf{X}) \approx \mathbf{U}^\top \dot{\boldsymbol{\varphi}}_0(\mathbf{X}) = \mathbf{U}^\top \mathbf{f}(\mathbf{X})$ appearing in the OpInf minimization can be replaced with the “Markovian” matrix $\mathbf{U}^\top D_t(\mathbf{P}_U \mathbf{X}) \approx \mathbf{U}^\top \dot{\boldsymbol{\varphi}}_0(\mathbf{P}_U \mathbf{X})$, where we have abused notation slightly by vectorizing the velocity map $\dot{\boldsymbol{\varphi}}_t$ columnwise. Notice that this precisely addresses the issue seen before: Instead of drawing snapshots from the projected FOM vector field $\mathbf{U}^\top \mathbf{f}(\mathbf{U} \hat{\mathbf{x}} + \mathbf{V} \hat{\mathbf{y}})$ containing contributions from the image of \mathbf{V} , snapshots are drawn from the reduced vector field $\mathbf{U}^\top \mathbf{f}(\mathbf{P}_U \mathbf{x}) = \mathbf{U}^\top \mathbf{f}(\mathbf{U} \hat{\mathbf{x}})$, which depends only on the image of \mathbf{U} . This ensures that the time derivatives used in the OpInf procedure correspond directly to the targeted model form matching the intrusive ROM, allowing for lossless recovery of the intrusive operator independent of the reduced dimension n when these derivatives are exact, i.e., $D_t = \dot{\boldsymbol{\varphi}}_0$. The connection to statistical dynamics mentioned in [27] comes from considering the single-step flow map $\mathbf{x}_{k+1} = \boldsymbol{\varphi}_{\Delta t}(\mathbf{x}_k)$, where with similar reasoning to the above, it can be seen that the projection $\mathbf{P}_U \boldsymbol{\varphi}_{\Delta t}(\mathbf{x}_k)$ will generally depend on the entire time history of the state \mathbf{x}_k and therefore cannot be a low-dimensional flow map itself. The algorithm in [27] shows that reprojection successfully mitigates the optimization error in OpInf provided that the user is able to collect solution snapshots with prespecified initial conditions.

Remark 3.2. In the case that the derivative operator $D_t \approx \dot{\boldsymbol{\varphi}}_0$ is not exact but, e.g., approximated with finite differences, care must be taken to construct and differentiate a new trajectory using the single-step ROM flow map approximation $\mathbf{P}_U \boldsymbol{\varphi}_{\Delta t}(\mathbf{P}_U \mathbf{x}) \approx \mathbf{U} \hat{\boldsymbol{\varphi}}_{\Delta t}(\mathbf{U}^\top \mathbf{x})$. Then the approximate velocity is calculated as $D_t(\mathbf{x}) = \mathcal{FD}_{\Delta t}(\mathbf{P}_U \boldsymbol{\varphi}_{\Delta t}(\mathbf{P}_U \mathbf{x}))$, where $\mathcal{FD}_{\Delta t}$ denotes the finite difference approximation with increment Δt , which converges to $\dot{\boldsymbol{\varphi}}_0(\mathbf{x})$ as the increment size $\Delta t \rightarrow 0$ approaches zero. In particular, $D_t(\mathbf{x})$ is not linear in the state \mathbf{x} , and applying the differential operator $\mathcal{FD}_{\Delta t}$ requires samples along the reprojected trajectory defined by $\mathbf{P}_U \boldsymbol{\varphi}_{\Delta t} \mathbf{P}_U$.

3.3. Nonintrusive centering. It will be desirable to apply centering in the variationally consistent Hamiltonian OpInf ROM for the same reasons as discussed in section 2.2. However, this requires some extension of the ideas discussed there since it is commonly assumed in OpInf

that the approximate state $\tilde{\mathbf{x}}$ is uncentered and therefore defined by Galerkin projection $\tilde{\mathbf{x}} = \mathbf{U}\hat{\mathbf{x}}$ directly onto the span of the (uncentered) reduced basis \mathbf{U} . The reason for this is as follows: Using a linear FOM $\dot{\mathbf{x}} = \mathbf{M}\mathbf{x}$, where $\mathbf{M} \in \mathbb{R}^{N \times N}$ is a constant matrix operator for illustration, the usual intrusive Galerkin ROM with centered approximation $\tilde{\mathbf{x}} = \mathbf{x}_0 + \mathbf{U}\hat{\mathbf{x}}$ is given by

$$\dot{\tilde{\mathbf{x}}} = \mathbf{U}^\top \mathbf{M} \mathbf{x}_0 + \mathbf{U}^\top \mathbf{M} \mathbf{U} \hat{\mathbf{x}} = \mathbf{U}^\top \mathbf{M} \mathbf{x}_0 + \hat{\mathbf{M}} \hat{\mathbf{x}}, \quad \hat{\mathbf{x}}_0 = \mathbf{0}.$$

In the uncentered case, the reduced matrix operator $\hat{\mathbf{M}} = \mathbf{U}^\top \mathbf{M} \mathbf{U}$ is straightforward to infer without knowledge of \mathbf{M} following the standard procedure in [28]. On the other hand, when the approximate state $\tilde{\mathbf{x}}$ is centered, it is not immediately obvious how to nonintrusively access the term $\mathbf{U}^\top \mathbf{M} \mathbf{x}_0$ necessary for ROM simulation, which involves a product between the unknown matrix \mathbf{M} and the initial condition \mathbf{x}_0 . While it is possible to incorporate this product directly in the governing model form through a learnable constant term, e.g., [23] and [12], this may not yield the desired energy level preservation in the resulting ROM due to the presence of nontrivial optimization error. However, notice that the problematic term $\mathbf{U}^\top \mathbf{M} \mathbf{x}_0 = \mathbf{U}^\top \dot{\mathbf{x}}|_{t=0} =: \mathbf{U}^\top \mathbf{v}_0$ can be written in terms of the reduced initial velocity vector $\hat{\mathbf{v}}_0 = \mathbf{U}^\top \mathbf{v}_0$, which is available at least in approximate form. Therefore, as long as the time derivative of the snapshots can be computed or accurately estimated, which is necessary to carry out OpInf in the first place, it is feasible to carry out a centered inference for the reduced operator

$$\bar{\mathbf{M}} = \underset{\bar{\mathbf{M}}}{\operatorname{argmin}} \left\| \mathbf{U}^\top (D_t(\mathbf{X}) - \mathbf{V}_0) - \bar{\mathbf{M}} \mathbf{U}^\top (\mathbf{X} - \mathbf{X}_0) \right\|_F^2,$$

where \mathbf{V}_0 is a matrix each column of which is \mathbf{v}_0 . This leads to the centered OpInf POD ROM,

$$\dot{\tilde{\mathbf{x}}} = \hat{\mathbf{v}}_0 + \bar{\mathbf{M}} \hat{\mathbf{x}}, \quad \hat{\mathbf{x}}_0 = \mathbf{0},$$

which can be simulated without intrusion into the FOM code.

Remark 3.3. While expressing the inaccessible part of the centered ROM in terms of snapshots of the velocity field $\dot{\mathbf{x}}$ is a useful trick which is suitable for the present purposes, it is important to note that this strategy becomes more complicated for nonlinear FOM models or ROMs which utilize nonlinear reduced bases. For example, the simplest quadratic model $\dot{\mathbf{x}} = \mathbf{x}^2$ leads to a POD ROM requiring the term $(\mathbf{x}_0 + \mathbf{U}\hat{\mathbf{x}})^2 = \mathbf{x}_0^2 + 2 \operatorname{Diag}(\mathbf{x}_0) \mathbf{U}\hat{\mathbf{x}} + \operatorname{diag}_{1,3}(\mathbf{U} \otimes \mathbf{U})(\mathbf{x} \otimes \mathbf{x})$, where $\operatorname{Diag}(\mathbf{x}_0)$ is a diagonal matrix with entries \mathbf{x}_0 and the order-3 tensor $\operatorname{diag}_{1,3}(\mathbf{U} \otimes \mathbf{U})_{j k}^i = u_j^i u_k^i$. Applying the same trick in this case yields $\mathbf{v}_0 = \mathbf{x}_0^2$, which does not recover the entire contribution of \mathbf{x}_0 . On the other hand, given \mathbf{v}_0 as before, standard OpInf [23] can be carried out for the inaccessible operators $2 \operatorname{Diag}(\mathbf{x}_0) \mathbf{U}$ and $\operatorname{diag}_{1,3}(\mathbf{U} \otimes \mathbf{U})$ simultaneously, perhaps with some optimization error incurred. It is straightforward to check that this general recipe will work for all polynomial nonlinearities, and extending centering to nonintrusive ROMs with more complex nonlinearities is an interesting avenue for future work.

3.4. Variationally consistent Hamiltonian OpInf. It is now pertinent to discuss how OpInf can be effectively applied to the variationally consistent ROM (2.1) with or without a centered state approximation. Recalling the splitting $H(\mathbf{x}) = \frac{1}{2} \mathbf{x}^\top \mathbf{A} \mathbf{x} + f(\mathbf{x})$ of the Hamiltonian from section 3.1, the standard OpInf problem associated to the uncentered ROM (2.1) is given as

$$\bar{\mathbf{A}} = \operatorname{argmin}_{\hat{\mathbf{A}} \in \mathbb{R}^{n \times n}} \left| (J\mathbf{U})^\top D_t(\mathbf{X}) - \hat{\mathbf{A}}\hat{\mathbf{X}} - \hat{\nabla}f(\mathbf{X}) \right|_F^2 \quad \text{s.t.} \quad \hat{\mathbf{A}}^\top = \hat{\mathbf{A}}.$$

This problem has a unique minimum under the same rank constraints as in section 3.1 while requiring fewer matrix multiplications with the reduced-order Poisson matrix $\hat{\mathbf{J}} = \mathbf{U}^\top J\mathbf{U}$, which is an advantage of the Petrov–Galerkin formulation. On the other hand, it is clear from section 3.2 that the learned operator $\bar{\mathbf{A}}$ will not generally match the intrusive ROM operator $\hat{\mathbf{A}}$ due to closure error introduced by the velocity approximation $D_t(\mathbf{X})$ and the reduced nonlinearity approximation $\hat{\nabla}f(\mathbf{X})$. When snapshots of the projected state $\mathbf{P}_U \mathbf{x}$ can be collected, this error can be eliminated by carrying out the reprojected inference,

$$(3.2) \quad \bar{\mathbf{A}} = \operatorname{argmin}_{\hat{\mathbf{A}} \in \mathbb{R}^{n \times n}} \left| (J\mathbf{U})^\top D_t(\mathbf{P}_U \mathbf{X}) - \hat{\mathbf{A}}\hat{\mathbf{X}} - \hat{\nabla}f(\mathbf{P}_U \mathbf{X}) \right|_F^2 \quad \text{s.t.} \quad \hat{\mathbf{A}}^\top = \hat{\mathbf{A}},$$

where the approximate reprojected velocity field $D_t(\mathbf{P}_U \mathbf{X}) \approx \dot{\varphi}_0(\mathbf{P}_U \mathbf{X})$ is now used. Letting $\mathbf{X}_t = D_t(\mathbf{P}_U \mathbf{X})$ denote the reprojected velocity field and $\hat{\mathbf{F}} = \hat{\nabla}f(\mathbf{P}_U \mathbf{X})$ denote the reduced, reprojected gradient of the nonlinear term (these are different from the $\mathbf{X}_t, \hat{\mathbf{F}}$ defined before), it can be shown using [17, Theorem 3.1] that the minimization (3.2) is equivalent to the linear system

$$(3.3) \quad (\mathbf{I} \oplus \hat{\mathbf{X}}\hat{\mathbf{X}}^\top) \operatorname{vec} \bar{\mathbf{A}} = \operatorname{vec} \left(\mathbf{U}^\top J^\top \mathbf{X}_t \hat{\mathbf{X}}^\top + \hat{\mathbf{X}} \mathbf{X}_t^\top J\mathbf{U} - \hat{\mathbf{F}}\hat{\mathbf{X}}^\top - \hat{\mathbf{X}}\hat{\mathbf{F}}^\top \right).$$

Importantly, this inference yields a learned operator $\bar{\mathbf{A}} \approx \hat{\mathbf{A}}$ which matches the intrusive ROM operator $\hat{\mathbf{A}} = \mathbf{U}^\top \mathbf{A} \mathbf{U}$ up to numerical error from the linear solve. Of course, once $\bar{\mathbf{A}}$ has been determined from the above, a nonintrusive version of the variationally consistent ROM (2.1) can be simulated as

$$(3.4) \quad \hat{\mathbf{J}}^\top \dot{\hat{\mathbf{x}}} = \bar{\mathbf{A}}\hat{\mathbf{x}} + \nabla \hat{f}(\hat{\mathbf{x}}), \quad \hat{\mathbf{x}}_0 = \mathbf{U}^\top \mathbf{x}_0.$$

On the other hand, it is often desirable to include a shift $\tilde{\mathbf{x}} = \mathbf{x}_0 + \mathbf{U}\hat{\mathbf{x}}$ in the ROM state approximation for the reasons described in section 2.2, corresponding to the centered Hamiltonian ROM analogous to (2.1):

$$(3.5) \quad \hat{\mathbf{J}}^\top \dot{\hat{\mathbf{x}}} = \mathbf{U}^\top \nabla H(\mathbf{x}_0 + \mathbf{U}\hat{\mathbf{x}}) = \nabla \hat{H}(\hat{\mathbf{x}}), \quad \hat{\mathbf{x}}_0 = \mathbf{0}.$$

When $\nabla H = \mathbf{A}$ is linear, such as in the solid mechanical examples of section 5, or when the nonlinear part ∇f can be accessed intrusively, this centering carries over to OpInf. Applying the novel procedure described in section 3.3, it follows from (3.5) that

$$\begin{aligned} \hat{\mathbf{J}}^\top \dot{\hat{\mathbf{x}}} &= \mathbf{U}^\top \mathbf{A}(\mathbf{x}_0 + \mathbf{U}\hat{\mathbf{x}}) + \mathbf{U}^\top \nabla f(\mathbf{x}_0 + \mathbf{U}\hat{\mathbf{x}}) \\ &= (J\mathbf{U})^\top (\mathbf{v}_0 - J\nabla f(\mathbf{x}_0)) + \hat{\mathbf{A}}\hat{\mathbf{x}} + \mathbf{U}^\top \nabla f(\mathbf{x}_0 + \mathbf{U}\hat{\mathbf{x}}), \end{aligned}$$

where $\mathbf{v}_0 = \dot{\mathbf{x}}|_{t=0}$ denotes the FOM initial velocity. Therefore, to simulate the centered nonintrusive ROM corresponding to (3.5),

$$(3.6) \quad \hat{\mathbf{J}}^\top \dot{\hat{\mathbf{x}}} = (J\mathbf{U})^\top (\mathbf{v}_0 - J\nabla f(\mathbf{x}_0)) + \bar{\mathbf{A}}\hat{\mathbf{x}} + \nabla \hat{f}(\hat{\mathbf{x}}), \quad \hat{\mathbf{x}}_0 = \mathbf{0},$$

it is enough to solve the reprojected OpInf problem

$$(3.7) \quad \bar{\mathbf{A}} = \operatorname{argmin}_{\hat{\mathbf{A}} \in \mathbb{R}^{n \times n}} \left| \mathbf{C} - \hat{\mathbf{A}}\mathbf{U}^\top (\mathbf{X} - \mathbf{X}_0) - \mathbf{U}^\top \nabla f(\mathbf{X}_0 + \mathbf{P}_U (\mathbf{X} - \mathbf{X}_0)) \right|_F \quad \text{s.t.} \quad \hat{\mathbf{A}}^\top = \hat{\mathbf{A}},$$

where the matrix $\mathbf{C} = (\mathbf{J}\mathbf{U})^\top D_t(\mathbf{X}_0 + \mathbf{P}_U(\mathbf{X} - \mathbf{X}_0)) - D_t(\mathbf{X}_0) + \mathbf{J}\nabla f(\mathbf{X}_0)$ and each column of the matrix \mathbf{X}_0 is the initial condition \mathbf{x}_0 . There are a couple of things to point out here. First, reprojection involves both the state \mathbf{x} and the initial condition \mathbf{x}_0 in this case since if $\text{im } \mathbf{V} = (\text{im } \mathbf{U})^\perp$, then $\mathbf{x} = \mathbf{x}_0 + \mathbf{P}_U(\mathbf{x} - \mathbf{x}_0) + \mathbf{P}_V(\mathbf{x} - \mathbf{x}_0)$ is the appropriate decomposition in terms of the subspaces $\text{span } \mathbf{x}_0, \text{im } \mathbf{U}$ and $\text{im } \mathbf{V}$. Second, when the approximate time derivative operator D_t is linear, such as in the case that $D_t = \dot{\boldsymbol{\varphi}}_t = \mathbf{J}\mathbf{A}$ is exact, the term $(\mathbf{J}\mathbf{U})^\top D_t(\mathbf{X}_0 + \mathbf{P}_U(\mathbf{X} - \mathbf{X}_0)) - D_t(\mathbf{X}_0) = (\mathbf{J}\mathbf{U})^\top D_t(\mathbf{P}_U(\mathbf{X} - \mathbf{X}_0))$ conveniently simplifies. In this case, letting $\mathbf{X}_t = D_t(\mathbf{P}_U(\mathbf{X} - \mathbf{X}_0)) + \mathbf{J}\nabla f(\mathbf{X}_0)$ and $\hat{\mathbf{X}} = \mathbf{U}^\top(\mathbf{X} - \mathbf{X}_0)$, a reduced operator $\bar{\mathbf{A}} \approx \hat{\mathbf{A}}$ which approximates the intrusive operator up to numerical error is given by the solution to (3.3) with $\hat{\mathbf{F}} = \mathbf{0}$.

This gives all information necessary for carrying out analysis and experiments using the proposed Hamiltonian ROM in both intrusive and nonintrusive settings. For ease of use, the information presented here is summarized in Algorithm 3.1. Note in particular that all proposed variationally consistent Hamiltonian OpInf ROMs are efficient online, incurring exactly the same online cost as OpInf ROMs, which are structure-agnostic.

4. Analysis. The proposed variationally consistent Hamiltonian ROM has now been established in four forms: uncentered and intrusive (2.1), centered and intrusive (3.5), uncentered and nonintrusive (3.4), and centered and nonintrusive (3.6). It will now be shown that state error is bounded in each of these cases by easily interpretable quantities coming from the reduced basis projection and the deviation of these ROMs from being canonically Hamiltonian. To that end, denote the L^2 -norm of the solution $\mathbf{x}(t)$ on the interval $[0, T] \subset \mathbb{R}$ by

$$\|\mathbf{x}\|^2 = \int_0^T |\mathbf{x}(t)|^2 dt,$$

and denote projection onto the orthogonal complement of the span of \mathbf{U} by $\mathbf{P}_U^\perp = \mathbf{I} - \mathbf{U}\mathbf{U}^\top$. To develop an error estimate for the intrusive ROMs (2.1) and (3.5), first recall the following classical result, which bounds projection error in terms of the eigenvalues of the covariance matrix.

Proposition 4.1. *Let $\mathbf{U} \in \mathbb{R}^{N \times n}$ be the rank- n POD basis generated through snapshots of $\mathbf{x}(t)$ on the interval $[0, T] \subset \mathbb{R}$. Then the POD projection error satisfies*

$$\|\mathbf{P}_U^\perp \mathbf{x}\|^2 = \|\mathbf{x} - \mathbf{P}_U \mathbf{x}\|^2 = \sum_{i>n} \lambda_i,$$

where λ_i are the eigenvalues of the snapshot covariance matrix $\mathbf{C}(\mathbf{x}) = \int_0^T \mathbf{x}(t)\mathbf{x}(t)^\top dt$. In particular, if \mathbf{U} is instead built from centered snapshots $\mathbf{w}(t) = \mathbf{x}(t) - \bar{\mathbf{x}}$, then the projection error satisfies

$$\|\mathbf{P}_U^\perp \mathbf{w}\|^2 = \|\mathbf{P}_U^\perp (\mathbf{x} - \bar{\mathbf{x}})\|^2 = \sum_{i>n} \lambda_i,$$

where λ_i are now the eigenvalues of the centered covariance matrix

$$\mathbf{C}(\mathbf{x}) = \int_0^T (\mathbf{x} - \bar{\mathbf{x}})(\mathbf{x} - \bar{\mathbf{x}})^\top dt.$$

Algorithm 3.1 Variationally Consistent Hamiltonian Operator Inference (VC-H-OpInf).

Input: Snapshots $\mathbf{X}^{N \times (n_s+1)}$ of FOM solution corresponding to times $\mathcal{T} = \{t_0, \dots, t_{n_s}\}$; intrusive access to nonlinear term $\nabla f: \mathbb{R}^N \rightarrow \mathbb{R}^N$ in Hamiltonian gradient $\nabla H(\mathbf{x}) = \mathbf{A}\mathbf{x} + \nabla f(\mathbf{x})$; reduced dimension $n > 0$.

Output: Reduced operator $\bar{\mathbf{A}} \in \mathbb{R}^{n \times n}$ satisfying $\bar{\mathbf{A}}^\top = \bar{\mathbf{A}}$, which approximates the linear term $\hat{\mathbf{A}} = \mathbf{U}^\top \mathbf{A} \mathbf{U}$ in the reduced Hamiltonian $\hat{H}(\hat{\mathbf{x}}) = \frac{1}{2} \hat{\mathbf{x}}^\top \hat{\mathbf{A}} \hat{\mathbf{x}} + \hat{f}(\hat{\mathbf{x}})$.

- 1: Build a reduced basis $\mathbf{U} \in V_n(\mathbb{R}^N)$ from snapshot data \mathbf{X} .
- 2: **if** centering is employed **then**
- 3: **if** reprojection is employed **then**
- 4: Collect reprojected snapshots $\mathbf{x}_0 = \mathbf{x}_0$, $\mathbf{x}_k = \mathbf{P}_U \boldsymbol{\varphi}_{\Delta t}(\mathbf{x}_0 + \mathbf{P}_U(\mathbf{x}_{k-1} - \mathbf{x}_0))$ columnwise in matrix \mathbf{X} .

▷ Skip this if $\dot{\boldsymbol{\varphi}}_0$ is available.
- 5: Set reduced state snapshots $\hat{\mathbf{X}} = \mathbf{U}^\top(\mathbf{X} - \mathbf{X}_0)$.
- 6: Collect reprojected velocity snapshots
 $D_t(\mathbf{X}_0 + \mathbf{P}_U(\mathbf{X} - \mathbf{X}_0)) \approx \dot{\boldsymbol{\varphi}}_0(\mathbf{X}_0 + \mathbf{P}_U(\mathbf{X} - \mathbf{X}_0))$.
- 7: Set derivative matrix $\mathbf{X}_t = D_t(\mathbf{X}_0 + \mathbf{P}_U(\mathbf{X} - \mathbf{X}_0)) - D_t(\mathbf{X}_0) + \mathbf{J} \nabla f(\mathbf{X}_0)$.
- 8: Collect reprojected reduced nonlinearity snapshots
 $\hat{\mathbf{F}} = \mathbf{U}^\top \nabla f(\mathbf{X}_0 + \mathbf{U}(\mathbf{X} - \mathbf{X}_0))$.
- 9: **else**
- 10: Set reduced state snapshots $\hat{\mathbf{X}} = \mathbf{U}^\top(\mathbf{X} - \mathbf{X}_0)$.
- 11: Collect velocity snapshots $D_t(\mathbf{X}) \approx \dot{\boldsymbol{\varphi}}_0(\mathbf{X})$.
- 12: Set derivative matrix $\mathbf{X}_t = D_t(\mathbf{X}) - D_t(\mathbf{X}_0) + \mathbf{J} \nabla f(\mathbf{X}_0)$.
- 13: Collect reduced nonlinearity snapshots $\hat{\mathbf{F}} = \mathbf{U}^\top \nabla f(\mathbf{X})$.
- 14: **end if**
- 15: Solve (3.3) for the reduced operator $\bar{\mathbf{A}}$ used in the nonintrusive centered ROM (3.6).
- 16: **else**
- 17: **if** reprojection is employed **then**
- 18: Collect reprojected snapshots $\mathbf{x}_0 = \mathbf{P}_U \mathbf{x}_0$, $\mathbf{x}_k = \mathbf{P}_U \boldsymbol{\varphi}_{\Delta t}(\mathbf{P}_U \mathbf{x}_{k-1})$ columnwise in matrix \mathbf{X} .

▷ Skip this if $\dot{\boldsymbol{\varphi}}_0$ is available.
- 19: Set reduced state snapshots $\hat{\mathbf{X}} = \mathbf{U}^\top \mathbf{X}$.
- 20: Collect reprojected velocity snapshots $\mathbf{X}_t = D_t(\mathbf{P}_U \mathbf{X}) \approx \dot{\boldsymbol{\varphi}}_0(\mathbf{P}_U \mathbf{X})$.
- 21: Collect reprojected reduced nonlinearity snapshots $\hat{\mathbf{F}} = \mathbf{U}^\top \nabla f(\mathbf{P}_U \mathbf{X})$.
- 22: **else**
- 23: Set reduced state snapshots $\hat{\mathbf{X}} = \mathbf{U}^\top \mathbf{X}$.
- 24: Collect velocity snapshots $\mathbf{X}_t = D_t(\mathbf{X}) \approx \dot{\boldsymbol{\varphi}}_0(\mathbf{X})$.
- 25: Collect reduced nonlinearity snapshots $\hat{\mathbf{F}} = \mathbf{U}^\top \nabla f(\mathbf{X})$.
- 26: **end if**
- 27: Solve (3.3) for the reduced operator $\bar{\mathbf{A}}$ used in the nonintrusive uncentered ROM (3.4).
- 28: **end if**

Recall additionally that a function $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^p$ is Lipschitz continuous provided that there exists a finite constant $L_{\mathbf{f}} > 0$ such that for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$,

$$|\mathbf{f}(\mathbf{x}) - \mathbf{f}(\mathbf{y})| \leq L_{\mathbf{f}} |\mathbf{x} - \mathbf{y}|,$$

where the norms are Euclidean and defined in the relevant spaces. Letting $\lambda_{\max}(\cdot)$ denote selection of the maximum eigenvalue, the following error estimate is now proven for the variationally consistent Hamiltonian ROM in the intrusive setting.

Theorem 4.2. *Let $\mathbf{x}(t)$ denote the solution to the FOM (1.1) on the interval $[0, T]$, and let $\hat{\mathbf{x}}(t)$ denote the solution to the variationally consistent Hamiltonian ROM (3.5). If the Hamiltonian gradient ∇H is Lipschitz continuous on the \mathbf{x} -image of $[0, T]$, then the approximation error in the reproductive ROM satisfies*

$$\|\mathbf{x} - (\mathbf{x}_0 + \mathbf{U}\hat{\mathbf{x}})\| \leq c \sqrt{\sum_{i>n} \lambda_i} + \bar{c} \|\nabla H(\mathbf{x})\| \sqrt{\lambda_{\max}(\hat{\mathbf{J}}^{-\top} \hat{\mathbf{J}}^{-1} - \mathbf{I})}$$

for constants c, \bar{c} depending on $T, \mathbf{U}, \mathbf{J}, \nabla H$ and where λ_i are the eigenvalues of the centered covariance matrix. In particular, the ROM error is bounded and tends to zero as the reduced dimension n tends to the full dimension N .

Proof. First, suppose that $\mathbf{x}_0 = \mathbf{0}$ so that the ROM approximation is uncentered. Adding and subtracting the projection $\mathbf{P}_U \mathbf{x}$, it follows that

$$\mathbf{x} - \mathbf{U}\hat{\mathbf{x}} = (\mathbf{x} - \mathbf{P}_U \mathbf{x}) + (\mathbf{P}_U \mathbf{x} - \mathbf{U}\hat{\mathbf{x}}) =: \mathbf{P}_U^\perp \mathbf{x} + \mathbf{y}$$

so that the integrand of the ROM error decomposes into two distinct terms. Moreover, by adding and subtracting additional terms, it follows that the time derivative of \mathbf{y} can be expressed as

$$\begin{aligned} \dot{\mathbf{y}} &= \mathbf{P}_U \dot{\mathbf{x}} - \mathbf{U} \dot{\hat{\mathbf{x}}} = \mathbf{P}_U \mathbf{J} \nabla H(\mathbf{x}) - \mathbf{U} \hat{\mathbf{J}}^{-\top} \mathbf{U}^\top \nabla H(\mathbf{U}\hat{\mathbf{x}}) \\ &= \mathbf{U} \left(\mathbf{U}^\top \mathbf{J} - \hat{\mathbf{J}}^{-\top} \mathbf{U}^\top \right) \nabla H(\mathbf{x}) + \mathbf{U} \hat{\mathbf{J}}^{-\top} \mathbf{U}^\top (\nabla H(\mathbf{x}) - \nabla H(\mathbf{P}_U \mathbf{x})) \\ &\quad + \mathbf{U} \hat{\mathbf{J}}^{-\top} \mathbf{U}^\top (\nabla H(\mathbf{P}_U \mathbf{x}) - \nabla H(\mathbf{U}\hat{\mathbf{x}})). \end{aligned}$$

Noting that the spectral norm $|\mathbf{A}|^2 = \lambda_{\max}(\mathbf{A}^\top \mathbf{A}) = \lambda_{\max}(\mathbf{A} \mathbf{A}^\top)$ is invariant under unitary transformations then yields the norm bound $|\mathbf{U} \hat{\mathbf{J}}^{-\top} \mathbf{U}^\top| = |\hat{\mathbf{J}}^{-1}|$ along with the calculation

$$\begin{aligned} \left| \mathbf{U} \left(\mathbf{U}^\top \mathbf{J} - \hat{\mathbf{J}}^{-\top} \mathbf{U}^\top \right) \right|^2 &= \left| \mathbf{U}^\top \mathbf{J} - \hat{\mathbf{J}}^{-\top} \mathbf{U}^\top \right|^2 \\ &= \lambda_{\max} \left(\mathbf{U}^\top \mathbf{J} \mathbf{J}^\top \mathbf{U} - \hat{\mathbf{J}}^{-\top} \hat{\mathbf{J}}^\top - \hat{\mathbf{J}} \hat{\mathbf{J}}^{-1} + \hat{\mathbf{J}}^{-\top} \hat{\mathbf{J}}^{-1} \right) \\ &= \lambda_{\max} \left(\hat{\mathbf{J}}^{-\top} \hat{\mathbf{J}}^{-1} - \mathbf{I} \right). \end{aligned}$$

Since the gradient ∇H is assumed Lipschitz continuous on the \mathbf{x} -image of the compact set $[0, T]$, hence bounded by assumption, this implies that the norm of $\dot{\mathbf{y}}$ can be estimated as

$$\begin{aligned} |\dot{\mathbf{y}}| &\leq \left| \mathbf{U} \left(\mathbf{U}^\top \mathbf{J} - \hat{\mathbf{J}}^{-\top} \mathbf{U}^\top \right) \right| |\nabla H(\mathbf{x})| + \left| \mathbf{U} \hat{\mathbf{J}}^{-\top} \mathbf{U}^\top \right| L_{\nabla H} \left(\left| \mathbf{P}_U^\perp \mathbf{x} \right| + |\mathbf{y}| \right) \\ &= \sqrt{\lambda_{\max} \left(\hat{\mathbf{J}}^{-\top} \hat{\mathbf{J}}^{-1} - \mathbf{I} \right)} |\nabla H(\mathbf{x})| + \left| \hat{\mathbf{J}}^{-1} \right| L_{\nabla H} \left(\left| \mathbf{P}_U^\perp \mathbf{x} \right| + |\mathbf{y}| \right). \end{aligned}$$

Now using that the time derivative of the norm $\partial_t |\mathbf{y}| = (2|\mathbf{y}|)^{-1} \langle \dot{\mathbf{y}}, \mathbf{y} \rangle \leq |\dot{\mathbf{y}}|$ is less than the norm of the time derivative and defining the constant $c_2 = \left| \hat{\mathbf{J}}^{-1} \right|_{L_{\nabla H}}$ leads to a first-order differential equation for $|\mathbf{y}|$:

$$\partial_t |\mathbf{y}| - c_2 |\mathbf{y}| \leq |\nabla H(\mathbf{x})| \sqrt{\lambda_{\max}(\hat{\mathbf{J}}^{-\tau} \hat{\mathbf{J}}^{-1} - \mathbf{I})} + c_2 \left| \mathbf{P}_U^\perp \mathbf{x} \right|.$$

Therefore, multiplying by the integrating factor $e^{-c_2 t}$ and applying Gronwall's inequality along with $\mathbf{y}(0) = \mathbf{0}$ yields

$$\begin{aligned} |\mathbf{y}| &\leq \int_0^t e^{c_2(t-s)} \left(|\nabla H(\mathbf{x})| \sqrt{\lambda_{\max}(\hat{\mathbf{J}}^{-\tau} \hat{\mathbf{J}}^{-1} - \mathbf{I})} + c_2 \left| \mathbf{P}_U^\perp \mathbf{x} \right| \right) ds \\ &\leq \int_0^T e^{c_2(T-s)} \left(|\nabla H(\mathbf{x})| \sqrt{\lambda_{\max}(\hat{\mathbf{J}}^{-\tau} \hat{\mathbf{J}}^{-1} - \mathbf{I})} + c_2 \left| \mathbf{P}_U^\perp \mathbf{x} \right| \right) ds \\ &\leq C \left\| |\nabla H(\mathbf{x})| \sqrt{\lambda_{\max}(\hat{\mathbf{J}}^{-\tau} \hat{\mathbf{J}}^{-1} - \mathbf{I})} + c_2 \left| \mathbf{P}_U^\perp \mathbf{x} \right| \right\|, \end{aligned}$$

where $C = \|e^{2c_2(T-t)}\| = \sqrt{(e^{2c_2 T} - 1)/2c_2}$ and the last step applies the Cauchy-Schwarz inequality. It follows from this and the Minkowski inequality that

$$\begin{aligned} \|\mathbf{y}\| &\leq C\sqrt{T} \left\| |\nabla H(\mathbf{x})| \sqrt{\lambda_{\max}(\hat{\mathbf{J}}^{-\tau} \hat{\mathbf{J}}^{-1} - \mathbf{I})} + c_2 \left| \mathbf{P}_U^\perp \mathbf{x} \right| \right\| \\ &\leq C\sqrt{T} \left\| \nabla H(\mathbf{x}) \sqrt{\lambda_{\max}(\hat{\mathbf{J}}^{-\tau} \hat{\mathbf{J}}^{-1} - \mathbf{I})} \right\| + Cc_2\sqrt{T} \left\| \mathbf{P}_U^\perp \mathbf{x} \right\|, \end{aligned}$$

and therefore the state error is bounded by

$$\|\mathbf{x} - \mathbf{U}\hat{\mathbf{x}}\| \leq \left\| \mathbf{P}_U^\perp \mathbf{x} \right\| + \|\mathbf{y}\| \leq \left(1 + Cc_2\sqrt{T} \right) \left\| \mathbf{P}_U^\perp \mathbf{x} \right\| + C\sqrt{T} \left\| \nabla H(\mathbf{x}) \sqrt{\lambda_{\max}(\hat{\mathbf{J}}^{-\tau} \hat{\mathbf{J}}^{-1} - \mathbf{I})} \right\|,$$

yielding the conclusion when $\mathbf{x}_0 = \mathbf{0}$. Now given the ROM approximation $\tilde{\mathbf{x}} = \mathbf{x}_0 + \mathbf{U}\hat{\mathbf{x}}$ with nonzero centering $\mathbf{x}_0 \neq \mathbf{0}$, $\mathbf{w} = \mathbf{x} - \mathbf{x}_0$ satisfies $\dot{\mathbf{w}} = \dot{\mathbf{x}}$, and the ROM error integrand becomes

$$\mathbf{x} - (\mathbf{x}_0 + \mathbf{U}\hat{\mathbf{x}}) = (\mathbf{w} - \mathbf{P}_U \mathbf{w}) + (\mathbf{P}_U \mathbf{w} - \mathbf{U}\hat{\mathbf{x}}) =: \mathbf{P}_U^\perp \mathbf{w} + \mathbf{z}.$$

Since the projection error $\|\mathbf{P}_U^\perp \mathbf{w}\|$ can be bounded by Proposition 4.1, $\mathbf{w}_0 = \mathbf{0}$, and $\dot{\mathbf{w}} = \dot{\mathbf{x}}$, the proof of this case becomes identical to the proven case $\mathbf{x}_0 = \mathbf{0}$ with no centering. ■

Theorem 4.2 shows that the error in the variationally consistent ROM approximation can be separated into two distinct terms coming from the information lost during POD projection and the failure of the reduced symplectic matrix inverse $\hat{\mathbf{J}}^{-1}$ to square to minus the identity, respectively. Notice that this second contribution is zero when using a symplectic lift or PSD basis since the equivariance condition on the reduced symplectic matrix implies that $\hat{\mathbf{J}} = \mathbf{J}_n$ and the resulting ROM is canonically Hamiltonian. In this case, state error in

the ROM depends only on the POD projection error as consistent with standard Galerkin theory, illustrating why symplectic lift bases tend to lead to the best performance when they are expressive enough to reliably capture FOM state dynamics. Unfortunately, this is often not the case in realistic simulations, where the usual POD basis formed via the SVD of state snapshots outperforms known algorithms for symplectic lift bases by several orders of magnitude. In these instances, the balance of both terms in Theorem 4.2 is important, and it can happen that only (2.1) leads to usable results. Some examples to this effect can be found in section 5.2.

Remark 4.3. It is remarkable that every Hamiltonian ROM which does not rely on a specially designed basis introduces extra contributions to the ROM error in order to preserve Hamiltonian structure at the reduced level. In the case of the least-squares ROM (1.4), this term is essentially the projection $\mathbf{P}_U^\perp \nabla H$ of the gradient into $(\text{im } \mathbf{U})^\perp$ (cf. [15, Theorem 3.1]), while for the proposed variationally consistent ROM (2.1), it is the failure of the reduced symplectic matrix inverse $\hat{\mathbf{J}}^{-1}$ to be anti-involutive.

It is useful to verify that a similar error bound holds for the nonintrusive ROM (3.7). In fact, the same technique can be applied provided that care is taken with other sources of potential error. In particular, these are error in the time derivative approximation $D_t(\mathbf{X}) \approx \dot{\varphi}_0(\mathbf{X})$ and optimization error, which arises from solving the relevant OpInf problem. For each state \mathbf{x} , these quantities are defined as

$$\varepsilon_{\Delta t}(\mathbf{x}) = |\dot{\mathbf{x}} - D_t(\mathbf{x})|, \quad \varepsilon_A(\mathbf{x}) = |(\mathbf{J}\mathbf{U})^\top D_t(\mathbf{x}) - \mathbf{U}^\top \nabla \bar{H}(\mathbf{x})|,$$

where $\bar{H}(\mathbf{x}) = \frac{1}{2} \mathbf{x}^\top \mathbf{U} \bar{\mathbf{A}} \mathbf{U}^\top \mathbf{x} + f(\mathbf{x})$ denotes the approximate Hamiltonian defined by the OpInf solution $\bar{\mathbf{A}} \approx \hat{\mathbf{A}}$. In particular, note that $\varepsilon_A(\mathbf{P}_U \mathbf{x})$ is the relevant notion of approximation error when reprojection is employed and that this quantity vanishes when the time derivative approximation $D_t \approx \dot{\varphi}_0$ is exact. With these notions in place, there is the following analogue of Theorem 4.2.

Theorem 4.4. *Let $\mathbf{x}(t)$ denote the solution to the FOM (1.1) on the interval $[0, T]$, and let $\hat{\mathbf{x}}(t)$ denote the solution to the variationally consistent OpInf ROM (3.6). Under the same assumptions as Theorem 4.2 and provided that the time derivative approximation $D_t \approx \dot{\varphi}_t$ is Lipschitz continuous and the operator error $|\mathbf{A} - \mathbf{U} \bar{\mathbf{A}} \mathbf{U}^\top|$ is finite, the approximation error in the reproductive ROM satisfies*

$$\|\mathbf{x} - (\mathbf{x}_0 + \mathbf{U} \hat{\mathbf{x}})\| \leq c \sqrt{\sum_{i>n} \lambda_i} + \bar{c} \|D_t(\mathbf{x})\| \sqrt{\lambda_{\max}(\hat{\mathbf{J}}^{-\top} \hat{\mathbf{J}}^{-1} - \mathbf{I})} + \tilde{c} (\|\varepsilon_{\Delta t}(\mathbf{x})\| + \|\varepsilon_A(\mathbf{P}_U \mathbf{x})\|)$$

for constants c, \bar{c}, \tilde{c} depending only on $T, \mathbf{U}, \mathbf{J}, \nabla H$ and where λ_i are the eigenvalues of the centered covariance matrix. Moreover, when the time derivative operator $D_t = \dot{\varphi}_t$ is exact, this reduces to precisely the bound in Theorem 4.2.

Proof. First, note that Lipschitz continuity in the Hamiltonian gradient ∇H implies the same in the approximate Hamiltonian gradient $\nabla \bar{H}$ since by adding and subtracting $\mathbf{A}(\mathbf{x} - \mathbf{y})$ and applying the triangle inequality, it follows that

$$\begin{aligned} |\nabla \bar{H}(\mathbf{x}) - \nabla \bar{H}(\mathbf{y})| &= (\mathbf{U} \bar{\mathbf{A}} \mathbf{U}^\top (\mathbf{x} - \mathbf{y}) + \nabla f(\mathbf{x}) - \nabla f(\mathbf{y})) \\ &\leq (|\mathbf{A} - \mathbf{U} \bar{\mathbf{A}} \mathbf{U}^\top| + L_{\nabla H}) |\mathbf{x} - \mathbf{y}|, \end{aligned}$$

showing that $\nabla \bar{H}$ is Lipschitz continuous with (finite) constant $L_{\nabla \bar{H}} = L_{\nabla H} + |\mathbf{A} - \mathbf{U} \bar{\mathbf{A}} \mathbf{U}^\top|$. Now suppose that $\mathbf{x}_0 = \mathbf{0}$ so that the ROM approximation is uncentered. Adding and subtracting the projection $\mathbf{P}_U \mathbf{x}$ as before, it follows that

$$\mathbf{x} - \mathbf{U} \hat{\mathbf{x}} = (\mathbf{x} - \mathbf{P}_U \mathbf{x}) + (\mathbf{P}_U \mathbf{x} - \mathbf{U} \hat{\mathbf{x}}) =: \mathbf{P}_U^\perp \mathbf{x} + \mathbf{y}$$

so that the error decomposes into two distinct terms. Adding and subtracting more terms, it follows that the time derivative of \mathbf{y} can be expressed as

$$\begin{aligned} \dot{\mathbf{y}} &= \mathbf{P}_U \dot{\mathbf{x}} - \mathbf{U} \dot{\hat{\mathbf{x}}} = \mathbf{P}_U (\dot{\mathbf{x}} - D_t(\mathbf{x})) + \mathbf{P}_U (\mathbf{I} - \mathbf{U} \hat{\mathbf{J}}^{-\top} \mathbf{U}^\top \mathbf{J}^\top) D_t(\mathbf{x}) \\ &\quad + \mathbf{U} \hat{\mathbf{J}}^{-\top} \mathbf{U}^\top \mathbf{J}^\top (D_t(\mathbf{x}) - D_t(\mathbf{P}_U \mathbf{x})) \\ &\quad + \mathbf{U} \hat{\mathbf{J}}^{-\top} ((\mathbf{J} \mathbf{U})^\top D_t(\mathbf{P}_U \mathbf{x}) - \mathbf{U}^\top \nabla \bar{H}(\mathbf{P}_U \mathbf{x})) \\ &\quad + \mathbf{U} \hat{\mathbf{J}}^{-\top} \mathbf{U}^\top (\nabla \bar{H}(\mathbf{P}_U \mathbf{x}) - \nabla \bar{H}(\mathbf{U} \hat{\mathbf{x}})). \end{aligned}$$

Using again the invariance of the spectral norm under unitary transformations along with the calculation

$$\begin{aligned} |\mathbf{P}_U - \mathbf{U} \hat{\mathbf{J}}^{-\top} \mathbf{U}^\top \mathbf{J}^\top| &= |\mathbf{U} (\mathbf{U}^\top \mathbf{J} - \hat{\mathbf{J}}^{-\top} \mathbf{U}^\top) \mathbf{J}^\top| \\ &= |\mathbf{U}^\top \mathbf{J} - \hat{\mathbf{J}}^{-\top} \mathbf{U}^\top| = \lambda_{\max} (\hat{\mathbf{J}}^{-\top} \hat{\mathbf{J}}^{-1} - \mathbf{I})^{1/2} \end{aligned}$$

leads to the following norm estimate of $\dot{\mathbf{y}}$:

$$\begin{aligned} |\dot{\mathbf{y}}| &\leq |\dot{\mathbf{x}} - D_t(\mathbf{x})| + \lambda_{\max} (\hat{\mathbf{J}}^{-\top} \hat{\mathbf{J}}^{-1} - \mathbf{I})^{1/2} |D_t(\mathbf{x})| + |\hat{\mathbf{J}}^{-1}| |D_t(\mathbf{x}) - D_t(\mathbf{P}_U \mathbf{x})| \\ &\quad + |\hat{\mathbf{J}}^{-1}| |(\mathbf{J} \mathbf{U})^\top D_t(\mathbf{P}_U \mathbf{x}) - \mathbf{U}^\top \nabla \bar{H}(\mathbf{P}_U \mathbf{x})| + |\hat{\mathbf{J}}^{-1}| |\nabla \bar{H}(\mathbf{P}_U \mathbf{x}) - \nabla \bar{H}(\mathbf{U} \hat{\mathbf{x}})| \\ &\leq \varepsilon_{\Delta t}(\mathbf{x}) + \lambda_{\max} (\hat{\mathbf{J}}^{-\top} \hat{\mathbf{J}}^{-1} - \mathbf{I})^{1/2} |D_t(\mathbf{x})| + |\hat{\mathbf{J}}^{-1}| L_{D_t} |\mathbf{P}_U^\perp \mathbf{x}| + |\hat{\mathbf{J}}^{-1}| \varepsilon_A(\mathbf{P}_U \mathbf{x}) \\ &\quad + |\hat{\mathbf{J}}^{-1}| L_{\nabla \bar{H}} |\mathbf{y}|. \end{aligned}$$

Defining constants $c_0 = |\hat{\mathbf{J}}^{-1}|$, $c_2 = |\hat{\mathbf{J}}^{-1}| L_{\nabla \bar{H}}$, and $c_3 = |\hat{\mathbf{J}}^{-1}| L_{D_t}$ and applying $\partial_t |\mathbf{y}| \leq |\dot{\mathbf{y}}|$ again leads to a first-order differential equation,

$$\partial_t |\mathbf{y}| - c_2 |\mathbf{y}| \leq \varepsilon_{\Delta t}(\mathbf{x}) + c_0 \varepsilon_A(\mathbf{P}_U \mathbf{x}) + |D_t(\mathbf{x})| \sqrt{\lambda_{\max} (\hat{\mathbf{J}}^{-\top} \hat{\mathbf{J}}^{-1} - \mathbf{I})} + c_3 |\mathbf{P}_U^\perp \mathbf{x}|,$$

and therefore $\partial_t (e^{-c_2 t} |\mathbf{y}|)$ can be integrated in time and bounded in terms of the constant $C = \sqrt{(e^{2c_2 T} - 1)/2c_2}$ to yield the pointwise-in-time bound

$$|\mathbf{y}| \leq C (\|\varepsilon_{\Delta t}(\mathbf{x})\| + c_0 \|\varepsilon_A(\mathbf{P}_U \mathbf{x})\|) + C c_3 \|\mathbf{P}_U^\perp \mathbf{x}\| + C \|D_t(\mathbf{x})\| \sqrt{\lambda_{\max} (\hat{\mathbf{J}}^{-\top} \hat{\mathbf{J}}^{-1} - \mathbf{I})}.$$

Defining the constant $c_1 = \max\{1, c_0\}$ and integrating again in time then yields the desired estimate

$$\begin{aligned} \|\mathbf{x} - U\hat{\mathbf{x}}\| &\leq \left\| \mathbf{P}_U^\perp \mathbf{x} \right\| + \|\mathbf{y}\| \\ &\leq \left(1 + C_{c_3}\sqrt{T}\right) \left\| \mathbf{P}_U^\perp \mathbf{x} \right\| + C_{c_1}\sqrt{T} (\|\varepsilon_{\Delta t}(\mathbf{x})\| + \|\varepsilon_A(\mathbf{P}_U \mathbf{x})\|) \\ &\quad + C\sqrt{T} \|D_t(\mathbf{x})\| \sqrt{\lambda_{\max}(\hat{\mathbf{J}}^{-\top} \hat{\mathbf{J}}^{-1} - \mathbf{I})}. \end{aligned}$$

The general case $\mathbf{x}_0 \neq \mathbf{0}$ now follows similarly as in the proof of Theorem 4.2.

To finish the argument, note that the velocity approximation error $\varepsilon_{\Delta t} \equiv 0$ is identically zero when the derivative operator $D_t = \dot{\boldsymbol{\varphi}}_t$ is exact. It follows that the reprojected OpInf approximation error $\varepsilon_A \equiv 0$ is identically zero as well since this enables exact recovery of the intrusive operator $\bar{\mathbf{A}} = \hat{\mathbf{A}}$. The fact that the L^2 -norm of the FOM velocity can be expressed as $\|\dot{\mathbf{x}}\| = \|\mathbf{J}\nabla H(\mathbf{x})\| = \|\nabla H(\mathbf{x})\|$ then shows that the OpInf ROM recovers exactly the same error bound as given in Theorem 4.2. ■

Remark 4.5. While the error bound in Theorem 4.4 does not assume that reprojection is employed, it has been expressed in a way which is useful in this case. A slightly different variant of Theorem 4.4 with the same key terms can be established in the case that “vanilla” OpInf is assumed; cf. [14, Theorem 1] for a proof in the case of the least-squares Hamiltonian ROM (1.4). However, without reprojection, there can be troublesome optimization error which pollutes the learned solution $\bar{\mathbf{A}}$, leading to a preasymptotic plateau in accuracy as more basis modes are added (cf. [14, Figure 10]).

5. Numerical examples. It is now appropriate to study the numerical performance² of the variationally consistent ROM proposed in sections 2 and 3. The primary metric that will be used for this purpose is relative ℓ^2 error between true and approximate state snapshots $\mathbf{X}, \tilde{\mathbf{X}} \in \mathbb{R}^{N \times n_t}$, which provides a normalized indicator of ROM performance:

$$R\ell_2(\mathbf{X}, \tilde{\mathbf{X}}) = \frac{\|\mathbf{X} - \tilde{\mathbf{X}}\|_F}{\|\mathbf{X}\|_F}.$$

In addition to error in the system state, signed error in the Hamiltonian function(al) will be measured along FOM and ROM trajectories, providing a complementary metric which detects violation of the governing physics. Note that this deviation is significantly influenced by the choice of time integrator; the implicit midpoint method (trapezoid rule), which is a symplectic integrator and hence energy-conserving, is used for all ROMs in order to remove this effect. Moreover, to facilitate visual comparison, this error will be measured in each case relative to the Hamiltonian value at the start of the ROM trajectory, which agrees with the FOM value of the Hamiltonian when centering is applied (cf. section 2.2). Since it will be instructive to compare the properties of POD bases used in building each ROM, the snapshot energy will

²Code for reproducing the results in this section can be found at <https://github.com/sandialabs/HamiltonianOpInf>.

also be evaluated: Given a rank r snapshot matrix \mathbf{X} with singular values $\{\sigma_i\}_{i=1}^r$ and POD basis size $n \leq r$, the snapshot energy is

$$E_s(\mathbf{X}, n) = \frac{\sum_{k=1}^n \sigma_k}{\sum_{k=1}^r \sigma_k}.$$

This quantity indicates how much of the total variance in the snapshots \mathbf{X} is captured by the reduced basis $\mathbf{U} \in \mathbb{R}^{N \times n}$.

To isolate the effect of enforcing variational consistency in Hamiltonian ROMs, every example below will be linear in the state \mathbf{x} . In fact, this simplified scenario is already enough to produce use cases which are considerably difficult for ROMs (cf. section 5.2), and it will be shown that the modifications proposed here may be necessary for acceptable performance in these cases. To provide an informative comparison, several ROMs are used in each of the given examples. “Intrusive G-ROM” refers to the most obvious Galerkin ROM (1.2), and “OpInf G-ROM” refers to the nonintrusive version of this based on operator inference. “Intrusive H-ROM (inconsistent)” refers to (1.4), while “Intrusive H-ROM (consistent)” refers to the proposed variationally consistent ROM (2.1). Operator inference H-ROMs corresponding to the variationally consistent case are also computed as in Algorithm 3.1 and denoted by “OpInf H-ROM (original)” and “OpInf H-ROM (reprojected)” to denote the absence (resp., presence) of reprojection (cf. section 3.2). In most cases, the centering procedures from sections 2.2 and 3.3 are applied to conserve the correct energy level along ROM trajectories, and this is denoted with “MC” in plot legends. To further probe similarities and differences, each considered ROM is applied using three different POD basis constructions: “Ordinary POD” denotes the basis \mathbf{U} formed using a simple SVD on snapshots $\mathbf{X} \approx \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$; “Cotangent Lift” denotes the \mathbf{J} -equivariant algorithm mentioned in section 1, where a block diagonal basis of rank $2m$ is built as $\mathbf{U} = \text{diag}(\mathbf{U}_{qp}, \mathbf{U}_{qp})$ with \mathbf{U}_{qp} defined by the rank m SVD of a concatenated snapshot matrix $(\mathbf{Q}|\alpha\mathbf{P})$ formed from position snapshots \mathbf{Q} and momentum snapshots \mathbf{P} which are scaled by a constant factor $\alpha = |\mathbf{Q}|_F / |\mathbf{P}|_F$ to equalize their influence; and “Block (q, p) ” denotes the block diagonal mixed basis of rank $2m$ built from the separate rank m SVDs of position and momentum snapshots; i.e., the Block (q, p) basis is computed as $\mathbf{U} = \text{diag}(\mathbf{U}_{qq}, \mathbf{U}_{pp})$, where $\mathbf{Q} \approx \mathbf{U}_{qq}\mathbf{\Sigma}_{qq}\mathbf{V}_{qq}^\top$ and $\mathbf{P} \approx \mathbf{U}_{pp}\mathbf{\Sigma}_{pp}\mathbf{V}_{pp}^\top$ denote the relevant decompositions for snapshots \mathbf{Q}, \mathbf{P} of position (resp., momentum). For completeness, the (scale-equalized) \mathbf{J} -equivariant “Complex SVD” algorithm described [29] is also presented in basis comparisons, though it is not considered in ROM building due to its near-identical performance to the cotangent lift.

5.1. Toy example: One-dimensional wave. First, it is instructive to see a toy example where the issues with the variationally inconsistent ROM (1.3) do not arise. Consider a one-dimensional linear wave equation with constant speed $c \in \mathbb{R}$ and wavefunction $\phi : [0, s] \times [0, T] \rightarrow \mathbb{R}$,

$$(5.1) \quad \begin{aligned} \phi_{tt}(s, t) &= c^2 \phi_{ss}(s, t), & 0 \leq s \leq l, \\ \phi(s, 0) &= h(y(s)), & \phi_t(s, 0) = 0, \end{aligned}$$

where periodic boundary conditions are applied and where the initial condition $h : [0, \infty) \rightarrow \mathbb{R}$ is a cubic spline:

$$h(y) = \begin{cases} 1 - \frac{3}{2}y^2 + \frac{3}{4}y^3 & 0 \leq y \leq 1, \\ \frac{1}{4}(2-y)^3 & 1 < y \leq 2, \\ 0 & y > 2, \end{cases} \quad y(s) = \left| s - \frac{1}{2} \right|.$$

This system is canonically Hamiltonian in the state variable $\mathbf{x} = (q \ p)^\top =: (\phi \ \phi_t)^\top$ involving the wavefunction and its first time derivative, which can be verified by defining the Hamiltonian functional

$$H(\mathbf{x}) = \frac{1}{2} \int_0^l (p^2 + c^2 q_s^2) ds$$

and writing (5.1) in Hamiltonian form:

$$(5.2) \quad \dot{\mathbf{x}} = \begin{pmatrix} \dot{q} \\ \dot{p} \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \begin{pmatrix} -c^2 \partial_{ss} & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} q \\ p \end{pmatrix} = \mathbf{J} \nabla H(\mathbf{x}).$$

Note that the gradient ∇H has been overloaded to mean the usual (Frechet) derivative of the functional H . To apply Hamiltonian POD ROMs to this system, it is necessary to discretize it in a way which preserves its variational structure. One such way is to subdivide the spatial interval $[0, l]$ into $(M - 1)$ subintervals of equal length so that the position and momentum functions q, p can be discretized along the M grid points bounding these intervals. Abusing notation, this produces a semidiscrete state $\mathbf{x} = (q \ p) \in \mathbb{R}^N$ of dimension $N = 2M$, with a corresponding discrete Hamiltonian

$$H(\mathbf{x}) = \frac{1}{2} \sum_{i=1}^{N/2} \left((p^i)^2 + c^2 \frac{(q^{i+1} - q^i)^2 + (q^i - q^{i-1})^2}{4\Delta x^2} \right).$$

Applying the implicit midpoint rule for time integration then yields a fully discrete system for the state \mathbf{x}_k at each time $k\Delta t$,

$$(5.3) \quad \frac{\mathbf{x}_{k+1} - \mathbf{x}_k}{\Delta t} = \mathbf{J} \mathbf{A} \left(\frac{\mathbf{x}_{k+1} + \mathbf{x}_k}{2} \right) = \begin{pmatrix} \mathbf{0} & \mathbf{I} \\ -\mathbf{I} & \mathbf{0} \end{pmatrix} \begin{pmatrix} -c^2 \mathbf{D}_2 & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \left(\frac{\mathbf{x}_{k+1} + \mathbf{x}_k}{2} \right),$$

where \mathbf{D}_2 denotes the usual circulant matrix resulting from a three-point stencil finite difference discretization of the one-dimensional Laplace operator. It is straightforward to check that the discrete Hamiltonian is exactly conserved along this evolution so that (5.3) is the discrete Hamiltonian counterpart to (5.1). For the examples here, the wave speed is fixed to $c = 0.1$, the length is fixed to $l = 1$, and the spatial domain is divided into $M = 500$ equally sized intervals to yield a discrete state vector \mathbf{x} of dimension $N = 1000$.

5.1.1. Time-reproductive simulation. To probe the performance of the various ROMs in a reproductive setting, snapshot data are collected every $\Delta t = 0.02$ seconds on the interval $[0, 10]$ and placed into a matrix $\mathbf{X} \in \mathbb{R}^{1000 \times 501}$. Applying different POD basis constructions in centered and uncentered cases leads to the snapshot energies and projection errors in Figure 2,




Figure 2. *POD basis snapshot energies and projection errors corresponding to the toy one-dimensional wave example in section 5.1.*

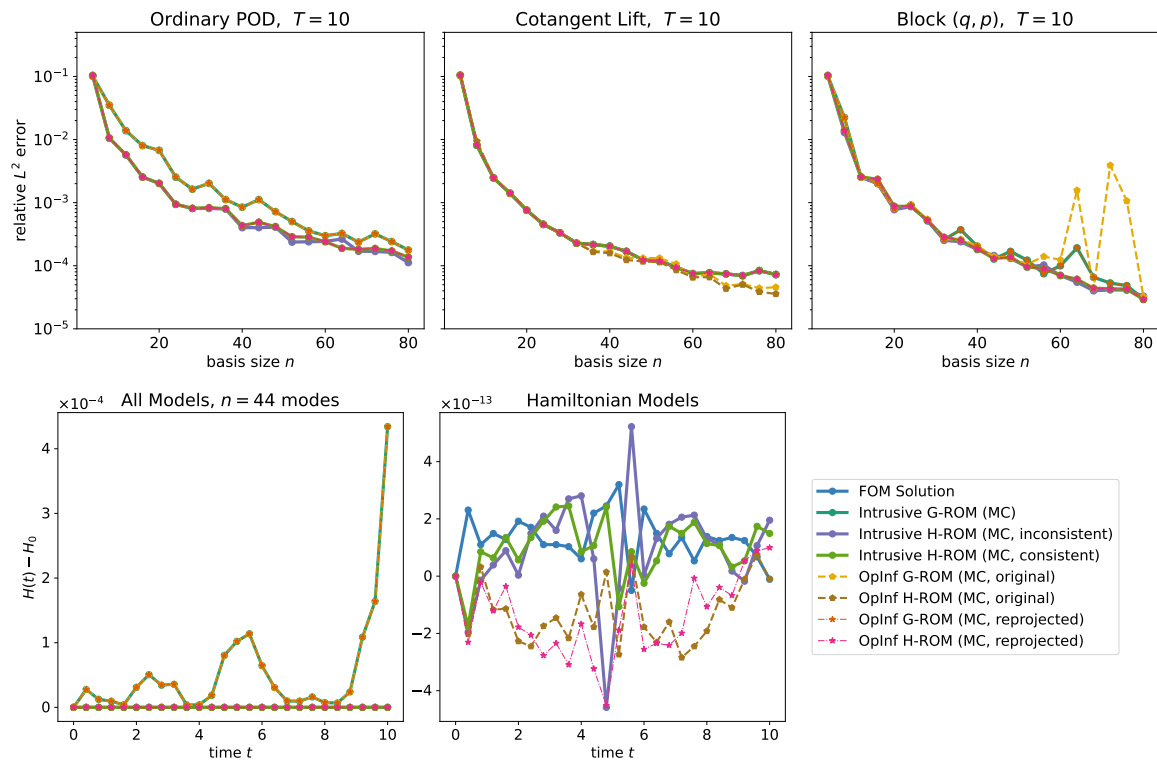


Figure 3. *ROM errors (top) and Hamiltonian errors (bottom) corresponding to the time-reproductive toy one-dimensional wave example in section 5.1. Hamiltonian error plots display ROMs using the Ordinary POD basis.*

where it is clear that adding basis modes to each method produces the slow but smooth error decay characteristic of problems with a large Kolmogorov n -width [30].

Results of applying each ROM to the POD-projected version of (5.3) under the centered approximation $\hat{\mathbf{x}} = \mathbf{x}_0 + \mathbf{U}\hat{\mathbf{x}}$ are displayed in Figure 3, along with the errors in the Hamiltonian along each evolution. Clearly, any algorithm is sufficient for producing an accurate ROM in

this case, although there is a modest benefit to applying a Hamiltonian ROM. Interestingly, the variational inconsistency introduced by (1.4) does not seem to harm ROM performance here, likely because this example is so simple. Notably, the Block (q, p) basis appears to lead to the most accurate ROMs, even though it did not have the lowest projection errors in Figure 2. On the OpInf side, it is clear that the reprojected OpInf ROMs match their intrusive counterparts exactly, as expected, while the vanilla OpInf ROMs are either very close to their intrusive counterparts or slightly unstable in the case of the OpInf G-ROM with the Block (q, p) basis. It is remarkable that both vanilla OpInf ROMs are slightly superconvergent in the Cotangent Lift case so that “non-Markovian” effects in the learned operators are actually helpful here.

The Hamiltonian errors in Figure 3 appear as expected. While all ROMs begin at the correct energy level due to the centering procedure employed, only the Hamiltonian ROMs are able to maintain this energy level to order 10^{-13} throughout their evolution. Again, it does not seem that variational inconsistency harms performance here, as the inconsistent H-ROM preserves the Hamiltonian to the same degree as the consistent ones. This is not exceptionally surprising: All that is required for energy conservation is skew symmetry of the reduced symplectic matrix $\hat{\mathbf{J}}$, which is obviously satisfied by the matrix $\hat{\mathbf{J}} = \mathbf{U}^\top \mathbf{J} \mathbf{U}$ used in the inconsistent case. Combined with the error plots in Figure 3, it is clear that there is remarkably little difference in overall performance between inconsistent and consistent H-ROMs in this example, confirming that it is not sophisticated enough to draw out such differences.

5.1.2. Time-predictive simulation. While each ROM performed similarly in the time-reproductive case, differences start to emerge when the domain of application is extended to $[0, 100]$ instead of the training domain $[0, 10]$. In this case, the same ROMs from above are tested on $[0, 100]$ in intervals of $\Delta t = 0.1$, a fivefold increase in time step and a tenfold increase in integration domain. It is easy to see from the ROM errors and Hamiltonian errors in Figure 4 that both the choice of POD basis and the choice of ROM algorithm are important to performance. Remarkably, the choice of centered approximation is also highly influential: Experiments in [17, Figure 6] show that OpInf G-ROMs will become unstable on this problem if centering is not applied.

It is interesting to observe that there is more separation between the errors of the G-ROMs and H-ROMs in the time-predictive case, which makes sense, as only the H-ROMs enforce the Hamiltonian structure necessary for correct dynamical evolution. Remarkably, although the Block (q, p) basis is again the most performant, it apparently leads to unstable G-ROMs as more modes are added. Similar to before, it seems that reprojection makes little difference on the OpInf ROMs using the Ordinary POD basis, although its effect is noticeable in the other bases. It is further interesting to note that the inconsistent H-ROM exhibits very slightly worse ROM error decay in the Ordinary POD case, perhaps displaying a consequence of its variational inconsistency. The Hamiltonian errors displayed in Figure 4 are remarkably similar to what is seen in the reproductive case, with all Hamiltonian ROMs conserving energy to the same order as the FOM throughout the entire length of time integration.

5.2. Three-dimensional linear elastic examples. The remaining examples are moderately sized problems from three-dimensional linear elasticity, which has the following equations of motion in terms of a displacement vector $\mathbf{q} \in \mathbb{R}^3$, material density $\rho > 0$, and Cauchy stress tensor $\boldsymbol{\sigma} \in \mathbb{R}^{3 \times 3}$:

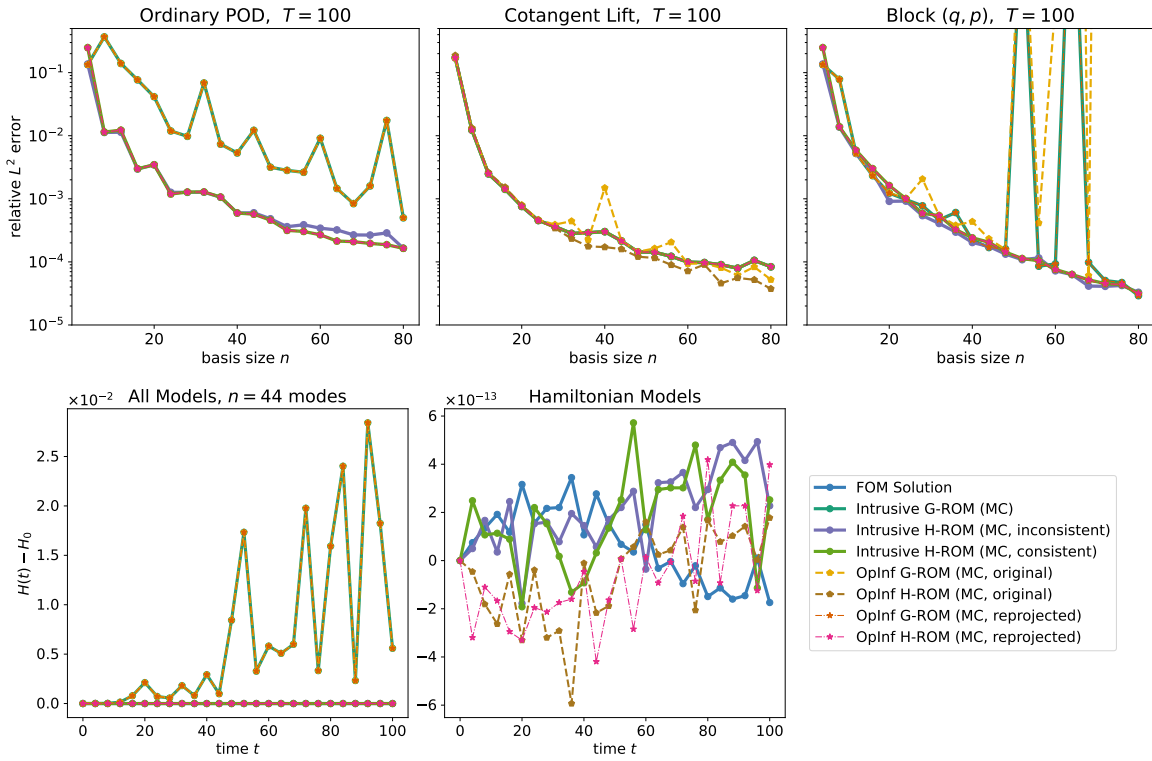


Figure 4. ROM errors (top) and Hamiltonian errors (bottom) corresponding to the time-predictive toy one-dimensional wave example in section 5.1. Hamiltonian error plots display ROMs using the Ordinary POD basis.

$$(5.4) \quad \rho \ddot{\mathbf{q}} = \nabla \cdot \boldsymbol{\sigma} \quad \text{on } \Omega \in \mathbb{R}^3.$$

It will be assumed that the material in question is elastic and follows Hooke's law so that the stress tensor $\boldsymbol{\sigma}$ has the explicit expression

$$\boldsymbol{\sigma} = \lambda \text{Tr}(\boldsymbol{\epsilon}) \mathbf{I} + 2\mu \boldsymbol{\epsilon}$$

in terms of Lamé coefficients $\lambda, \mu > 0$ and the symmetric strain tensor $\boldsymbol{\epsilon} := \frac{1}{2} [\nabla \mathbf{q} + (\nabla \mathbf{q})^\top]$. It can be shown (see [22]) that the Hamiltonian functional for (5.4) can be expressed concisely in terms of (noncanonical) position and velocity variables $(\mathbf{q}, \dot{\mathbf{q}})$:

$$(5.5) \quad H(\mathbf{q}, \dot{\mathbf{q}}) = \frac{1}{2} \int_{\Omega} \left(\rho |\dot{\mathbf{q}}|^2 + \lambda [\text{Tr}(\boldsymbol{\epsilon})]^2 + 2\mu |\boldsymbol{\epsilon}|^2 \right) dV.$$

It is interesting to note that, in one dimension, (5.5) reduces to the linear wave equation (5.1) considered in section 5.1. Considering three-dimensional problems introduces challenging solution features that are difficult for ROMs to capture, e.g., waves propagating in multiple spatial directions over potentially complex three-dimensional geometries (e.g., Figure 5(a) and (b)), which can complicate wave interaction. For physical realism, these model geometries are assumed to be made of steel, and the material parameters $E = 200$ GPa (Young's modulus), $\nu = 0.25$ (Poisson's ratio), and $\rho = 7800$ kg/m³ (density) lead to straightforward calculations




Figure 5. Geometries and meshes for the three-dimensional linear elastic examples, along with a one-dimensional cartoon illustrating the problem setup.

for the Lamé coefficients in the stress tensor σ via the formulas $\lambda = \frac{E\nu}{(1+\nu)(1-2\nu)}$ and $\mu = \frac{E}{2(1+\nu)}$. It is well known [11] that the CFL condition for (5.4) is given by

$$(5.6) \quad \Delta t_w \leq \frac{h}{c},$$

where h is the approximate average mesh resolution and $c = \sqrt{\frac{LM}{\rho}}$ is the speed of sound. Here, LM denotes the longitudinal elastic modulus, or “P-wave modulus,” as it is commonly referred to, which can be calculated from the elastic coefficients using the formula

$$(5.7) \quad LM := \frac{E(1-\nu)}{(1+\nu)(1-2\nu)}.$$

It is well known that time steps less than the one given by (5.6) are sufficient to resolve the highest-frequency waves that the medium can propagate. In view of this, the estimate (5.6) is used as a guide to inform the FOM time steps used in the numerical experiments presented below.

As per common practice in the field of solid mechanics, (5.4) is discretized in space using the finite element method, yielding a semidiscrete system of the form (1.5) with $\mathbf{f} = \mathbf{0}$, where (again overloading notation) $\mathbf{q} \in \mathbb{R}^M$ is the discretized displacement field while $\mathbf{M} \in \mathbb{R}^{M \times M}$ and $\mathbf{K} \in \mathbb{R}^{M \times M}$ are the mass and stiffness matrices of the scheme, respectively. Letting $\mathbf{p} := \mathbf{M}\dot{\mathbf{q}}$ define the momentum leads to the discrete Hamiltonian state $\mathbf{x} := (\mathbf{q} \ \mathbf{p})^\top \in \mathbb{R}^N$ and allows (1.5) to be written as the following canonical Hamiltonian system:

$$(5.8) \quad \dot{\mathbf{x}} = \begin{pmatrix} \dot{\mathbf{q}} \\ \dot{\mathbf{p}} \end{pmatrix} = \begin{pmatrix} \mathbf{0} & \mathbf{I} \\ -\mathbf{I} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{K} & \mathbf{0} \\ \mathbf{0} & \mathbf{M}^{-1} \end{pmatrix} \begin{pmatrix} \mathbf{q} \\ \mathbf{p} \end{pmatrix} = \mathbf{J}\mathbf{A}\mathbf{x},$$

where $\mathbf{A} = \nabla H$ is the discrete Hamiltonian gradient and H is the quadratic discrete Hamiltonian:

$$H(\mathbf{x}) = \frac{1}{2} (\mathbf{q}^\top \mathbf{K} \mathbf{q} + \mathbf{p}^\top \mathbf{M}^{-1} \mathbf{p}).$$

5.2.1. Cantilever plate. The first test case considered presently is a classical solid mechanics benchmark, also considered in [17], involving a vibrating rectangular cantilever plate of size $0.2 \times 0.2 \times 0.03$ m, i.e., with spatial domain $\Omega = (0, 0.2) \times (0, 0.2) \times (0, 0.03) \in \mathbb{R}^3$ (see Figure 5(a)). Let $\mathbf{s} = (s^1 \ s^2 \ s^3)^\top \in \mathbb{R}^3$ denote the position vector. The left side of the plate is clamped in this benchmark, and a homogeneous Dirichlet boundary condition $\mathbf{q} = \mathbf{0}$ is imposed on the left boundary $\Gamma_l := \{s^2, s^3 \in \bar{\Omega} : s^1 = 0\}$. The remaining boundaries of the domain Ω satisfy homogeneous Neumann boundary conditions, indicating that these boundaries are free surfaces. An initial velocity of 100 m/s in the s^3 -direction is prescribed on the right boundary of the domain, $\Gamma_r := \{s^2, s^3 \in \bar{\Omega} : s^1 = 0.2\}$, so that $\dot{\mathbf{q}}(\mathbf{s}, 0) = (0, 0, 100)^\top$ for $\mathbf{s} \in \Gamma_r$. Figure 5(c) displays a one-dimensional cartoon illustrating the problem setup, whereby the initial velocity perturbation will cause the plate to vibrate and undergo a flapping motion. This vibration causes waves to form inside the plate, propagating in complex patterns across all three coordinate directions. In particular, this represents a challenging test case for Hamiltonian model reduction which previous methods are unable to reliably solve (cf. [17, Figure 29]).

Building the FOM corresponding to system (5.8) is accomplished using the open-source³ Albany-LCM multiphysics code base [32, 24, 25]. The domain Ω is discretized with a uniform mesh of $20 \times 20 \times 3$ hexahedral elements, as shown in Figure 5(a), and the FOM system (1.5) is advanced forward from time $t = 0$ to time $t = 2 \times 10^{-3}$ s to generate snapshots; here, a symplectic implicit Newmark time-stepping scheme with parameters $\beta = 0.25$ and $\gamma = 0.5$ and time step $\Delta t = 1.0 \times 10^{-7}$ s is used, and snapshots are collected every 1.0×10^{-5} s. It was verified that the time step used in the FOM integration is below the time step required to resolve the wave (1.8×10^{-6} s, from (5.6)). The transformation $\mathbf{p} = \mathbf{M}\dot{\mathbf{q}}$ is applied to generate momentum data, and the resulting 201 snapshots of the discrete state \mathbf{x} , each of length 10584, are used to build the various POD bases. Plots of the relevant snapshot energies and projection errors are shown in Figure 6. Notice in particular that the \mathbf{J} -equivariant constructions, i.e., the Cotangent Lift and Complex SVD, produce bases which are noticeably less expressive than the Ordinary POD basis, leading to projection errors which are several orders of magnitude higher after about 80 POD modes.




Figure 6. *POD basis snapshot energies and projection errors corresponding to the cantilever plate example in section 5.2.1.*

³Albany-LCM is available on Github at <https://github.com/sandialabs/LCM>.

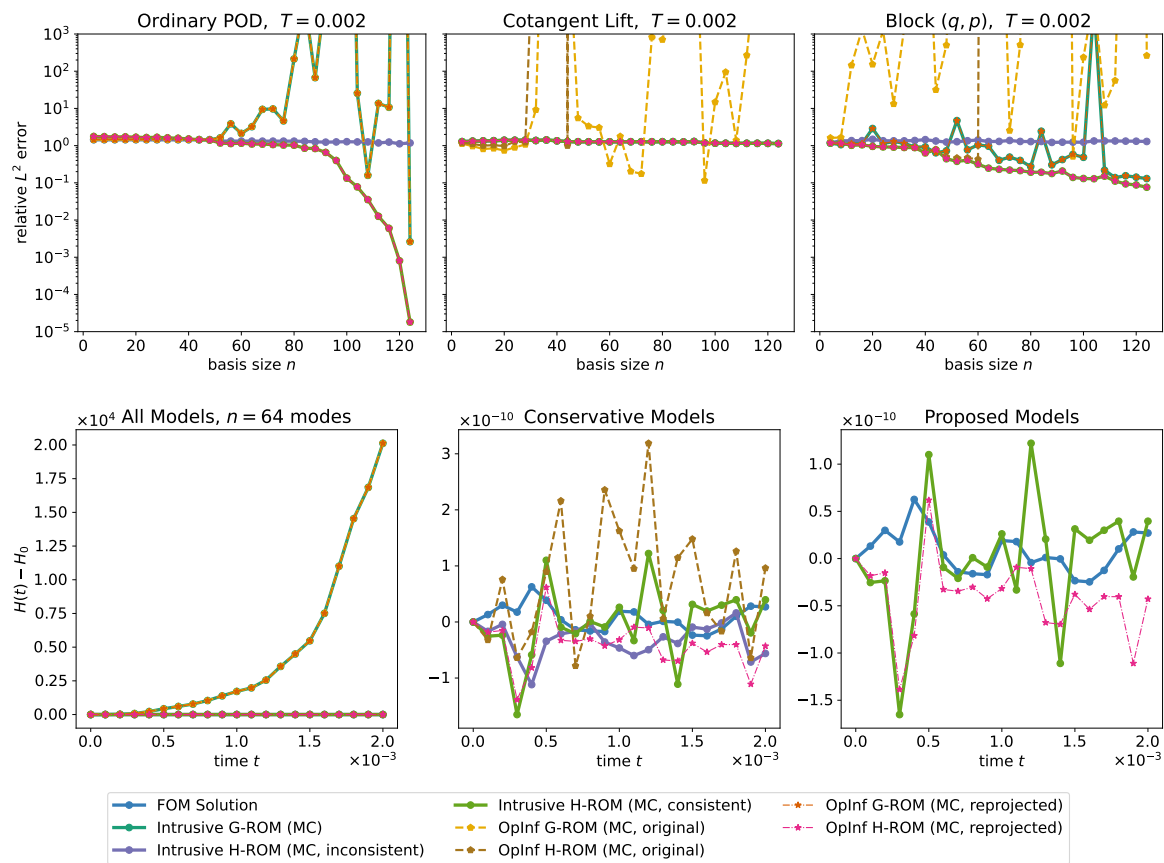


Figure 7. ROM state errors (top) and Hamiltonian errors (bottom) corresponding to the time-reproductive case of the cantilever plate example in section 5.2.1. Hamiltonian error plots display ROMs using the Ordinary POD basis.

Time-reproductive simulation. Again, the first test presented here involves a time-reproductive case, where the goal of each ROM is to reproduce the snapshots used for training. The results of this experiment are shown in Figure 7, where it is immediately clear that preserving a variationally consistent Hamiltonian structure is necessary for adequate performance. In each case, it is observed that the inconsistent H-ROM produces order-1 errors regardless of basis size, confirming that the least-squares projection used to enforce Hamiltonian structure is detrimental to ROM accuracy. Conversely, the variationally consistent H-ROM in both its intrusive and its nonintrusive versions converges smoothly at the rates expected based on the projection errors in Figure 6. Importantly, it can be seen that all ROMs built using the Cotangent Lift basis (which inherit variational consistency automatically) are unusable in this case, as this basis is not expressive enough to capture the variance in the solution snapshots.

Besides its implications on ROM accuracy, notice that preservation of Hamiltonian structure is also crucial for ROM stability: Each pictured G-ROM is unstable with the addition of basis modes even though some choices of reduced dimension n lead to relatively accurate solutions. This is an illustration of the commonly observed accuracy-versus-stability trade-off

seen in structure-preserving methods. While there is nothing which prevents a straightforward G-ROM from being as accurate as any H-ROM (although this is not observed here), Lyapunov stability is guaranteed only by the H-ROMs. The same goes for energy conservation: Figure 7 shows clearly that each Hamiltonian ROMs is conservative while the G-ROMs are not and also that the variationally inconsistent H-ROM (1.4) is almost as conservative as the variationally consistent one despite its very poor accuracy.

Time-predictive simulation. Beyond the reproductive case seen above, it is worthwhile to examine what happens when the ROMs describing the cantilever plate are simulated well beyond where they have been trained. Figure 9 shows the results of this in the case that the final time is $T = 5 \times 10^{-3}$ with the same time step $\Delta t = 1.0 \times 10^{-5}$, amounting to a test interval which is $2.5\times$ longer than the training interval. While this problem is periodic, note that these ranges are well below the period length of 0.02. This means that the dynamics on the unseen region $[0.002, 0.005]$ are substantially different than those seen by the ROMs during training, making this problem more difficult than the similar predictive problem in [17, section 5.5]. The error plots in Figure 9 confirm that only the Hamiltonian models are capable of stable results on this test case and that only the variationally consistent H-ROMs exhibit some semblance of convergence to a realistic solution. This is further confirmed by examining the trajectory of a single point displayed in Figure 8, which visibly highlights the artifacts in the inconsistent H-ROM solution. Again, it is clear from Figure 9 that the Cotangent Lift basis is unusable here, although it is remarkable that each of the nonreprojected (i.e., original) OpInf ROMs is highly unstable despite the \mathbf{J} -equivariance of the reduced basis. The plots

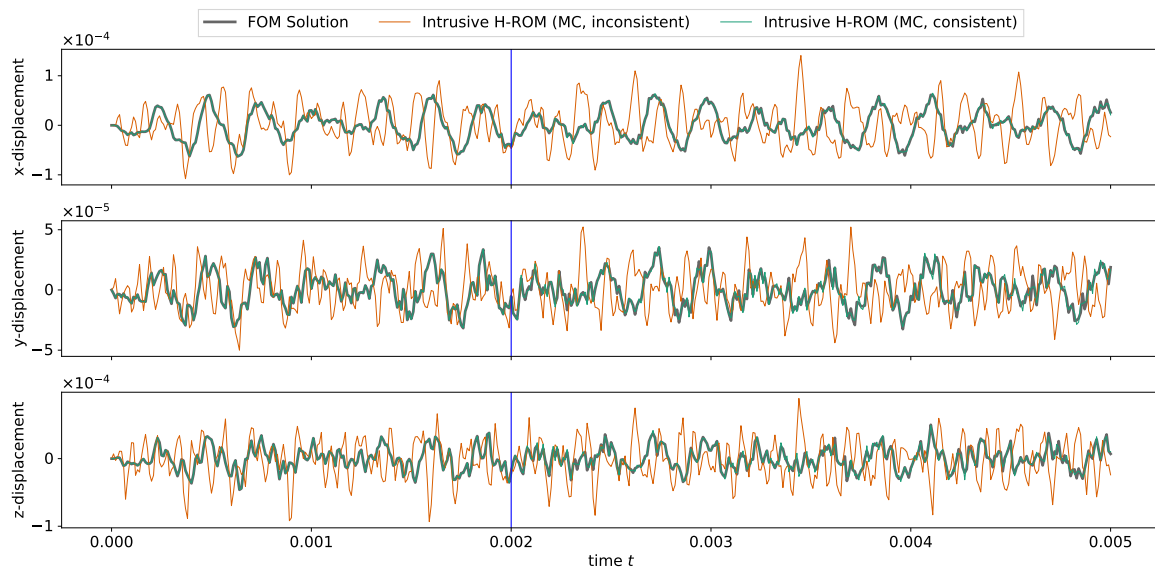


Figure 8. Time histories of the x -, y -, and z -displacement solutions to the cantilever plate example at a single node with global node ID 1000 and spatial coordinates $(x, y, z) = (0.12, 0.05, 0.02)$. The vertical blue line denotes the end of the training interval. All ROMs were computed with an Ordinary POD basis of $n = 100$ modes. Note that the trajectories of the variationally consistent OpInf ROMs (not pictured) lie directly on top of the variationally consistent intrusive H-ROM.

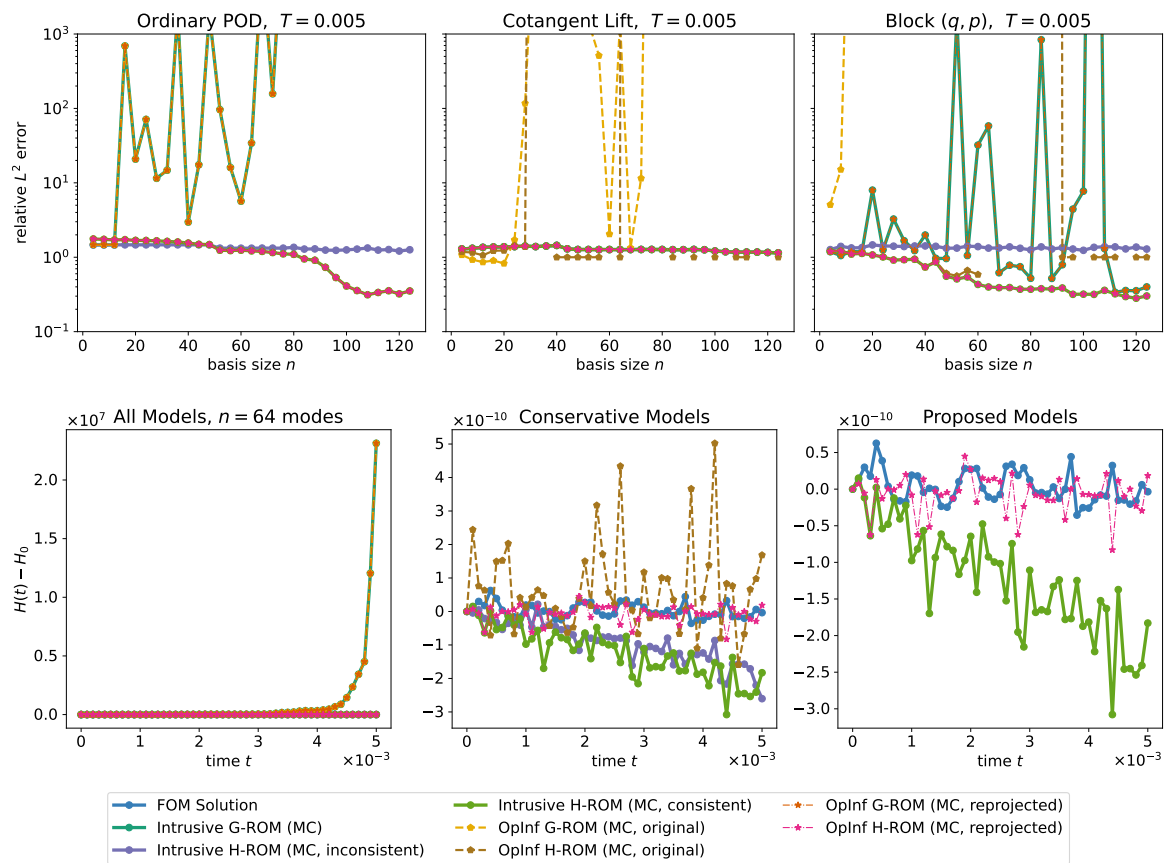


Figure 9. ROM state errors (top) and Hamiltonian errors (bottom) corresponding to the time-predictive case of the cantilever plate example in section 5.2.1. Hamiltonian error plots display ROMs using the Ordinary POD basis.

of Hamiltonian error tell a similar story: While each G-ROM artificially adds energy to the system over time, the H-ROMs are conservative to the same order as the FOM. Moreover, it is clear that variational consistency does not have much impact on conservation despite its large impact on accuracy.

5.2.2. Flexible bracket. The final example presented here involves a flexible bracket clamped at one end. Although the experimental setup is similar to the cantilever plate example, the problem is posed on a more complex geometry (see Figure 5(b)) featuring a circular hole whose presence weakens the bracket and complicates the propagation and interaction of three-dimensional waves propagating throughout the bracket. Additionally, both the final integration time and the time interval at which snapshots are collected are widened considerably. This has the effect of implicitly damping the oscillations in the metal, effectively making the dynamics easier to capture with a POD basis. It is interesting to see how this change, as well as the natural change caused by the inclusion of a more challenging mesh topology, affects ROM performance. To facilitate further comparison, the straightforward Lagrangian ROM (cf. [8]) arising from Galerkin projection onto a standard POD basis for the

displacement \mathbf{q} will be presented here alongside the Hamiltonian methods previously considered. Note that this method returns approximations to the displacement \mathbf{q} , velocity $\dot{\mathbf{q}}$, and acceleration $\ddot{\mathbf{q}}$, and therefore the approximate momentum $\tilde{\mathbf{p}} = \mathbf{M}\dot{\mathbf{q}}$ is postprocessed from the velocity approximation.

The domain on which the problem is posed is 20.5 cm wide in the x -direction and 20 cm wide in the y - and z -directions. The geometry is discretized using an unstructured tetrahedral mesh of 9640 elements, each with a mesh increment of approximately 5.8 mm (Figure 5(b)). The FOM system (1.5) is advanced forward from time $t = 0$ to time $t = 2 \times 10^{-2}$ s using $\Delta t = 1.0 \times 10^{-7}$ s, and the larger sampling time $\Delta t = 1.0 \times 10^{-4}$ s is used for snapshot collection. It was verified that the time step used to simulate the FOM was below the minimum time step required to resolve the wave (1.05×10^{-6} s, from (5.6)). As before, the transformation $\mathbf{p} = \mathbf{M}\dot{\mathbf{q}}$ is applied to generate momentum data, and the resulting 201 snapshots of the discrete state \mathbf{x} , each of length 15828, are used to build the various POD bases whose snapshot energies and projection errors are displayed in Figure 12. Remarkably, it is again the case here that the Ordinary POD basis is much more expressive than the other constructions, exhibiting projection errors which decay nearly to machine precision as basis modes are added. Moreover, even the Lagrangian construction suffers here, perhaps due to its utilization of only position information. For the associated ROMs, only reproduction of the snapshots in time is considered for this example, although all ROMs are simulated with and without centering to demonstrate the effect of this choice.

Results reproducing the motion of the flexible bracket are shown in Figure 13 for the ROMs without centering and in Figure 14 for the centered ROMs. Notice that the accuracy of each ROM is remarkably insensitive to this choice, with both sets of ROMs displaying similar performance across every choice of basis. It can also be seen that both the centered and the uncentered ROMs preserve the Hamiltonian to the same order, although only the centered ROMs preserve this quantity at the value prescribed by the FOM (not visible because the difference in values is plotted). Additionally, the results in both figures show that only the ROMs built with the Ordinary POD basis are capable of exhibiting convergence down to order 10^{-11} , while other bases are not expressive enough to lower the relative error beyond about 10^{-3} . The Lagrangian L-ROM in particular is quite stable but not highly performant. Moreover, as in the plate example, the variationally inconsistent H-ROM exhibits large errors at every basis size, while the consistent H-ROM developed here exhibits smooth convergence as expected by the projection errors in Figure 12.

Finally, several ROM solutions are plotted in in Figure 10. Subfigures (b)–(d) show the z -displacement of the flexible bracket problem at the final time $t = 2 \times 10^{-2}$ s, along with the FOM z -displacement solution in subfigure (a) and the absolute errors between the ROM and FOM z -displacement solutions in subfigures (e)–(g). The z -displacements in Figure 10 are scaled by a factor of 10 for visualization purposes, and all ROMs visualized included centering. The reader can observe some spurious oscillations in the intrusive G-ROM solution (Figure 10(b)) with respect to the FOM solution (Figure 10(a)). The situation is worse for the inconsistent H-ROM solution (Figure 10(c)), which has completely diverged from the FOM solution. In contrast, the consistent H-ROM solution (Figure 10(d)) matches the FOM solution remarkably well. These conclusions are supported by the absolute error plots shown in Figure 11.




Figure 10. 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, all calculated using centering (b)–(d).

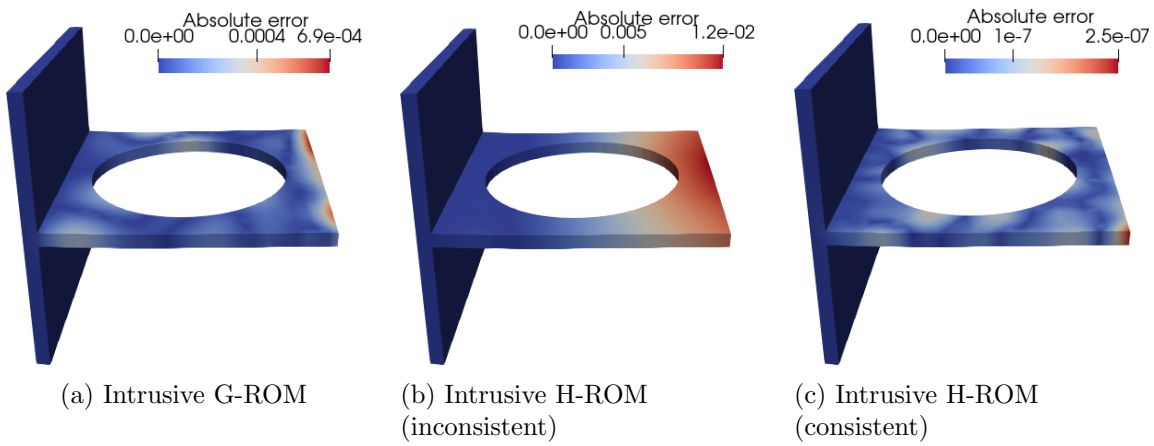


Figure 11. Spatial distribution of the absolute errors in the z -displacement for various ROM solutions to the flexible bracket problem, evaluated (with centering) at the final time $t = 2 \times 10^{-2}$ s.




Figure 12. POD basis snapshot energies and projection errors corresponding to the flexible bracket example in section 5.2.2.

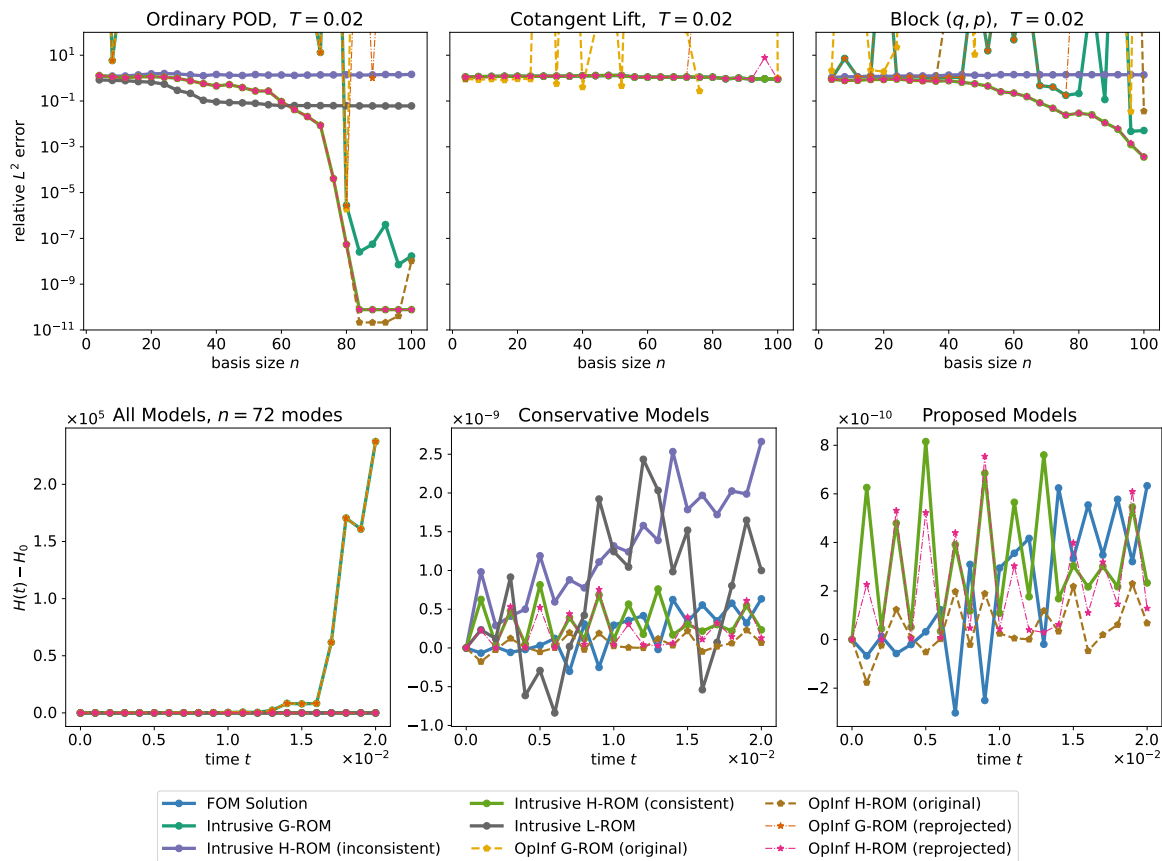


Figure 13. ROM state errors (top) and Hamiltonian errors (bottom) corresponding to the flexible bracket example in section 5.2.2 without ROM centering.

6. Conclusion. A variationally consistent method for the model reduction of canonical Hamiltonian systems has been proposed which is compatible with any linear orthonormal reduced basis and applicable in both intrusive and nonintrusive settings. In addition to exhibiting exact energy conservation and symplecticity, the proposed method does not require additional considerations beyond Petrov–Galerkin orthogonality and has been shown to produce accurate solutions in situations where previous methods cannot. Error estimates have been proven which interpret the error in the proposed method in terms of a balance between the usual basis projection error and an additional term arising from the failure of the reduced system to be canonically Hamiltonian. Since canonical Hamiltonian systems of appropriate regularity necessarily have a natural Lagrangian counterpart which is Legendre-dual, future work will compare and contrast the approach presented here with Lagrangian ROMs based on the second-order equations of motion. It is further interesting to consider the extension of these ideas to nonlinear problems, quadratic or nonlinear bases, and noncanonical Hamiltonian systems where there is no Lagrangian analogue. It is hoped that the work presented here will provide a useful starting point for researchers who wish to deploy ROMs on large-scale Hamiltonian systems with complex physics and realistic material parameters.

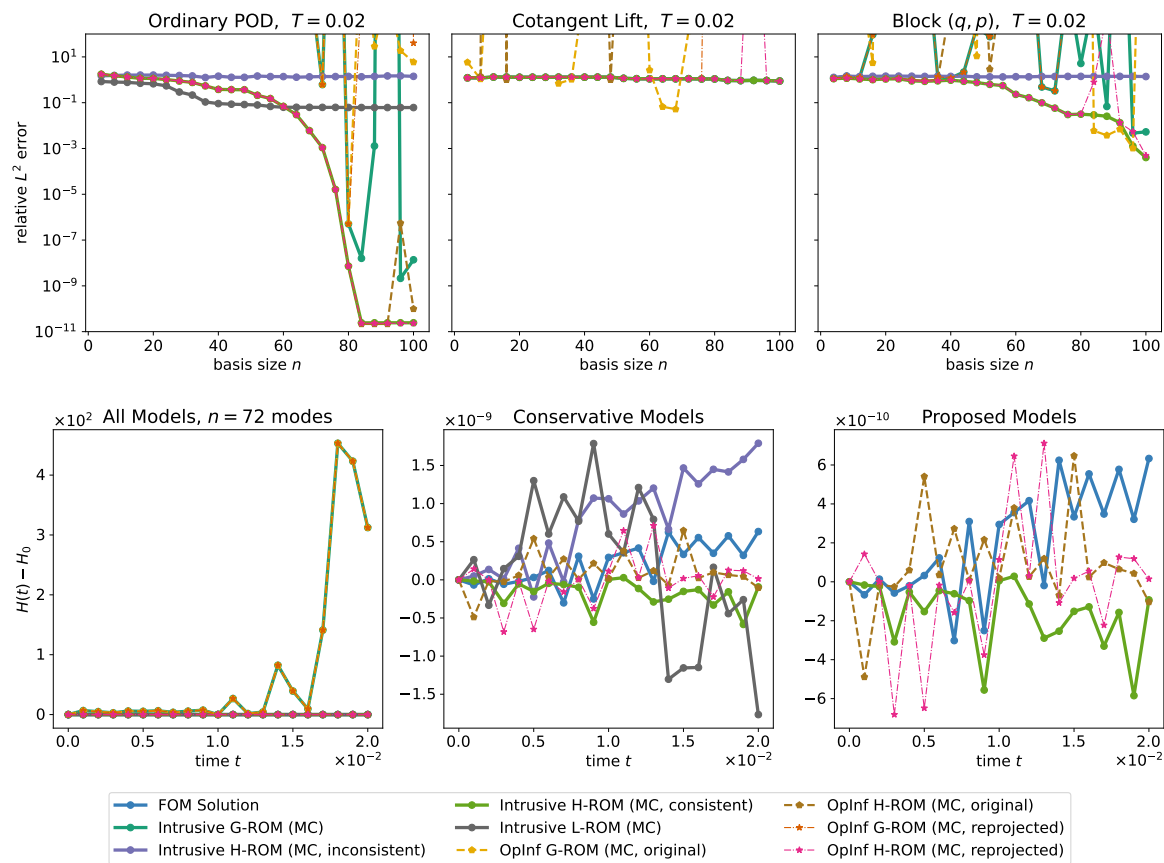


Figure 14. ROM state errors (top) and Hamiltonian errors (bottom) corresponding to the flexible bracket example in section 5.2.2 with ROM centering.

Acknowledgments. 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. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy (DOE) or the U.S. Government. This article has been coauthored by an employee of National Technology & Engineering Solutions of Sandia, LLC, under contract DE-NA0003525 with the DOE. The employee owns all right, title, and interest in and to the article and is solely responsible for its contents. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this article or allow others to do so for U.S. Government purposes. The DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<https://www.energy.gov/downloads/doe-public-access-plan>). The authors thank Elise Walker for careful proofreading and suggestions regarding the presentation of Proposition 2.2.

REFERENCES

- [1] B. M. AFKHAM AND J. S. HESTHAVEN, *Structure preserving model reduction of parametric Hamiltonian systems*, SIAM J. Sci. Comput., 39 (2017), pp. A2616–A2644, <https://doi.org/10.1137/17M1111991>.
- [2] M. AFKHAM AND J. HESTHAVEN, *Structure-preserving model-reduction of dissipative Hamiltonian systems*, J. Sci. Comput., 81 (2019), pp. 3–11, <https://doi.org/10.1007/s10915-018-0653-6>.
- [3] J. BARNETT, I. TEZAUR, AND A. MOTA, *The Schwarz Alternating Method for the Seamless Coupling of Nonlinear Reduced Order Models and Full Order Models*, Technical report, SAND2022-13334R; 710315, 2022, <https://www.osti.gov/biblio/1890056>.
- [4] 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*, Comput. Methods Appl. Mech. Engrg., 372 (2020), 113433, <https://doi.org/10.1016/j.cma.2020.113433>.
- [5] P. BUCHFINK, S. GLAS, AND B. HAASDONK, *Symplectic model reduction of Hamiltonian systems on nonlinear manifolds and approximation with weakly symplectic autoencoder*, SIAM J. Sci. Comput., 45 (2023), pp. A289–A311, <https://doi.org/10.1137/21M1466657>.
- [6] P. BUCHFINK, S. GLAS, AND B. HAASDONK, *Optimal bases for symplectic model order reduction of canonizable linear Hamiltonian systems*, IFAC-PapersOnLine, 55 (2022), pp. 463–468, <https://doi.org/10.1016/j.ifacol.2022.09.138>.
- [7] C. E. CALIGAN AND C. CHANDRE, *Conservative dissipation: How important is the Jacobi identity in the dynamics?*, Chaos, 26 (2016), 053101, <https://doi.org/10.1063/1.4948411>.
- [8] K. CARLBERG, R. TUMINARO, AND P. BOGGS, *Preserving Lagrangian structure in nonlinear model reduction with application to structural dynamics*, SIAM J. Sci. Comput., 37 (2015), pp. B153–B184, <https://doi.org/10.1137/140959602>.
- [9] A. DE CASTRO, P. BOCHEV, P. KUBERRY, AND I. TEZAUR, *Explicit synchronous partitioned scheme for coupled reduced order models based on composite reduced bases*, Comput. Methods Appl. Mech. Engrg., 417 (2023), 116398, <https://doi.org/10.1016/j.cma.2023.116398>.
- [10] A. DE CASTRO, P. KUBERRY, I. TEZAUR, AND P. BOCHEV, *A novel partitioned approach for reduced order model—Finite element model (ROM-FEM) and ROM-ROM coupling*, in Earth and Space 2022: Space Exploration, Utilization, Engineering, and Construction in Extreme Environments, C. B. Dreyer and J. Littell, eds., ASCE Press, Reston, VA, pp. 475–489, <https://doi.org/10.1061/9780784484470.044>.
- [11] R. DHARI, *Speed of Sound in Solids Calculator*, <https://www.omnicalculator.com/physics/speed-sound-solids>, 2024.
- [12] I.-G. FARCAS, R. P. GUNDEVIA, R. MUNIPALLI, AND K. E. WILLCOX, *Domain decomposition for data-driven reduced modeling of large-scale systems*, AIAA J., 62 (2024), <https://arc.aiaa.org/doi/10.2514/1.J063715>.
- [13] C. FARHAT, T. CHAPMAN, AND P. AVERY, *Structure-preserving, stability, and accuracy properties of the energy-conserving sampling and weighting method for the hyper reduction of nonlinear finite element dynamic models*, Internat. J. Numer. Methods Engrg., 102 (2015), pp. 1077–1110, <https://doi.org/10.1002/nme.4820>.
- [14] Y. GENG, J. SINGH, L. JU, B. KRAMER, AND Z. WANG, *Gradient preserving operator inference: Data-driven reduced-order models for equations with gradient structure*, Comput. Methods Appl. Mech. Engrg., 427 (2024), 117033, <https://doi.org/10.1016/j.cma.2024.117033>.
- [15] Y. GONG, Q. WANG, AND Z. WANG, *Structure-preserving Galerkin POD reduced-order modeling of Hamiltonian systems*, Comput. Methods Appl. Mech. Engrg., 315 (2017), pp. 780–798, <https://doi.org/10.1016/j.cma.2016.11.016>.
- [16] A. GRUBER, M. GUNZBURGER, L. JU, AND Z. WANG, *Energetically consistent model reduction for metriplectic systems*, Comput. Methods Appl. Mech. Engrg., 404 (2023), 115709, <https://doi.org/10.1016/j.cma.2022.115709>.
- [17] A. GRUBER AND I. TEZAUR, *Canonical and noncanonical Hamiltonian operator inference*, Comput. Methods Appl. Mech. Engrg., 416 (2023), 116334, <https://doi.org/10.1016/j.cma.2023.116334>.
- [18] M. D. GUNZBURGER, J. S. PETERSON, AND J. N. SHADID, *Reduced-order modeling of time-dependent PDEs with multiple parameters in the boundary data*, Comput. Methods Appl. Mech. Engrg., 196 (2007), pp. 1030–1047, <https://doi.org/10.1016/j.cma.2006.08.004>.

- [19] J. S. HESTHAVEN, C. PAGLIANTINI, AND N. RIPAMONTI, *Structure-preserving model order reduction of Hamiltonian systems*, Proc. Int. Cong. Math., 7 (2022), pp. 5072–5097.
- [20] P. HOLMES, J. LUMLEY, AND G. BERKOOZ, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, Cambridge University Press, Cambridge, 1996.
- [21] S. LALL, P. KRYSL, AND J. E. MARSDEN, *Structure-preserving model reduction for mechanical systems*, Phys. D, 184 (2003), pp. 304–318, [https://doi.org/10.1016/S0167-2789\(03\)00227-6](https://doi.org/10.1016/S0167-2789(03)00227-6).
- [22] J. E. MARSDEN AND T. RATIU, *Introduction to Mechanics and Symmetry: A Basic Exposition of Classical Mechanical Systems*, Springer-Verlag, Berlin, 1998.
- [23] S. A. MCQUARRIE, C. HUANG, AND K. E. WILLCOX, *Data-driven reduced-order models via regularised operator inference for a single-injector combustion process*, J. Roy. Soc. New Zealand, 51 (2021), pp. 194–211, <https://doi.org/10.1080/03036758.2020.1863237>.
- [24] A. MOTA, I. TEZAUR, AND C. ALLEMAN, *The Schwarz alternating method in solid mechanics*, Comput. Methods Appl. Mech. Engrg., 319 (2017), pp. 19–51, <https://doi.org/10.1016/j.cma.2017.02.006>.
- [25] A. MOTA, I. TEZAUR, AND G. PHILIPOT, *The Schwarz alternating method for transient solid dynamics*, Internat. J. Numer. Methods Engrg., 123 (2022), pp. 5036–5071, <https://doi.org/10.1002/nme.6982>.
- [26] E. J. PARISH AND K. DURAISAMY, *Non-Markovian closure models for large eddy simulations using the Mori-Zwanzig formalism*, Phys. Rev. Fluids, 2 (2017), 014604, <https://doi.org/10.1103/PhysRevFluids.2.014604>.
- [27] B. PEHERSTORFER, *Sampling low-dimensional Markovian dynamics for preasymptotically recovering reduced models from data with operator inference*, SIAM J. Sci. Comput., 42 (2020), pp. A3489–A3515, <https://doi.org/10.1137/19M1292448>.
- [28] B. PEHERSTORFER AND K. WILLCOX, *Data-driven operator inference for nonintrusive projection-based model reduction*, Comput. Methods Appl. Mech. Engrg., 306 (2016), pp. 196–215, <https://doi.org/10.1016/j.cma.2016.03.025>.
- [29] L. PENG AND K. MOHSENI, *Symplectic model reduction of Hamiltonian systems*, SIAM J. Sci. Comput., 38 (2016), pp. A1–A27, <https://doi.org/10.1137/140978922>.
- [30] A. PINKUS, *n-Widths in Approximation Theory*, Ergeb. Math. Grenzgeb. (3) 7, Springer-Verlag, Berlin, 1985.
- [31] E. QIAN, B. KRAMER, B. PEHERSTORFER, AND K. WILLCOX, *Lift & Learn: Physics-informed machine learning for large-scale nonlinear dynamical systems*, Phys. D, 406 (2020), 132401, <https://doi.org/10.1016/j.physd.2020.132401>.
- [32] A. SALINGER, R. BARTLETT, A. BRADLEY, Q. CHEN, I. DEMESKO, 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 component-based design to develop a flexible, generic multi-physics analysis code*, Int. J. Multiscale Comput. Engrg., 14 (2016), pp. 415–438, <https://doi.org/10.1615/IntJMultCompEng.2016017040>.
- [33] N. SAWANT, B. KRAMER, AND B. PEHERSTORFER, *Physics-informed regularization and structure preservation for learning stable reduced models from data with operator inference*, Comput. Methods Appl. Mech. Engrg., 404 (2023), 115836, <https://doi.org/10.1016/j.cma.2022.115836>.
- [34] K. SCHACKE, *On the Kronecker Product*, Master’s thesis, University of Waterloo.
- [35] H. SHARMA AND B. KRAMER, *Preserving Lagrangian structure in data-driven reduced-order modeling of large-scale dynamical systems*, Phys. D, 462 (2024), 134128, <https://doi.org/10.1016/j.physd.2024.134128>.
- [36] H. SHARMA, H. MU, P. BUCHFINK, R. GELEN, S. GLAS, AND B. KRAMER, *Symplectic model reduction of Hamiltonian systems using data-driven quadratic manifolds*, Comput. Methods Appl. Mech. Engrg., 417 (2023), 116402, <https://doi.org/10.1016/j.cma.2023.116402>.
- [37] H. SHARMA, D. A. NAJERA-FLORES, M. D. TODD, AND B. KRAMER, *Lagrangian operator inference enhanced with structure-preserving machine learning for nonintrusive model reduction of mechanical systems*, Comput. Methods Appl. Mech. Engrg., 423 (2024), 116865, <https://doi.org/10.1016/j.cma.2024.116865>.
- [38] H. SHARMA, Z. WANG, AND B. KRAMER, *Hamiltonian operator inference: Physics-preserving learning of reduced-order models for canonical Hamiltonian systems*, Phys. D, 431 (2022), 133122, <https://doi.org/10.1016/j.physd.2021.133122>.

- [39] L. SIROVICH, *Turbulence and the dynamics of coherent structures. I. Coherent structures*, Quart. Appl. Math., 45 (1987), pp. 561–571, <https://doi.org/10.1090/qam/910462>.
- [40] K. SOCKWELL, *Mass Conserving Hamiltonian-Structure Preserving Reduced Order Modeling for the Rotating Shallow Water Equations Discretized by a Mimetic Spatial Scheme*, Ph.D. thesis, Florida State University, 2019.
- [41] S. YILDIZ, P. GOYAL, T. BENDOKAT, AND P. BENNER, *Data-driven identification of quadratic representations for nonlinear Hamiltonian systems using weakly symplectic liftings*, J. Mach. Learn. Model. Comput., 5 (2024), pp. 45–71, <https://doi.org/10.1615/JMachLearnModelComput.2024052810>.