
Certified data-driven physics-informed greedy auto-encoder simulator

Xiaolong He

Department of Structural Engineering
University of California, San Diego
La Jolla, CA, 92093
xih251@ucsd.edu

Youngsoo Choi

Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
Livermore, CA, 94550
choi15@llnl.gov

William D. Fries

Applied Mathematics
School of Mathematical Sciences
University of Arizona
Tucson, AZ, 85721
frieswd@math.arizona.edu

Jonathan L. Belof

Materials Science Division
Physical and Life Science Directorate
Lawrence Livermore National Laboratory
Livermore, CA, 94550
belof1@llnl.gov

Jiun-Shyan Chen

Department of Structural Engineering
University of California, San Diego
La Jolla, CA, 92093
js-chen@ucsd.edu

Abstract

A parametric adaptive greedy Latent Space Dynamics Identification (gLaSDI) framework is developed for accurate, efficient, and certified data-driven physics-informed greedy auto-encoder simulators of high-dimensional nonlinear dynamical systems. In the proposed framework, an auto-encoder and dynamics identification models are trained interactively to discover intrinsic and simple latent-space dynamics. To effectively explore the parameter space for optimal model performance, an adaptive greedy sampling algorithm integrated with a physics-informed error indicator is introduced to search for optimal training samples on the fly, outperforming the conventional predefined uniform sampling. Further, an efficient k -nearest neighbor convex interpolation scheme is employed to exploit local latent-space dynamics for improved predictability. Numerical results demonstrate that the proposed method achieves 121 to $2,658 \times$ speed-up with 1 to 5% relative errors for radial advection and 2D Burgers dynamical problems.

1 Introduction

Physical simulations have played an increasingly significant role in the developments of engineering, science, and technology. However, high-fidelity forward physical simulations can be computationally intractable even with high-performance computing, prohibiting their applications to problems that require a large number of forward simulations [1–9]. In recent years, several reduced-order models (ROMs) have been integrated with latent-space learning algorithms [10–13]. However, the latent-space dynamics models of these methods are often complex with limited *interpretability*. Many methods have been developed for the identification of interpretable governing laws from data, including symbolic regression that searches both parameters and governing equations simultaneously

[14–17], parametric models that fit parameters to equations of a given form, such as the sparse identification of nonlinear dynamics (SINDy) [18, 19], and operator inference [20–22]. Champion, et al. [19] applied an auto-encoder for nonlinear projection and SINDy to identify simple ordinary differential equations (ODEs) that govern the latent-space dynamics. The auto-encoder and the SINDy model were trained interactively to achieve simple latent-space dynamics. However, the proposed SINDy-autoencoder method is not parameterized and generalizable. Many non-intrusive ROMs have been developed based on POD-based linear projection with latent-space dynamics captured by polynomials through operator inference [20–25]. Due to the limitation of the POD-based linear projection, these non-intrusive ROMs have difficulties with advection-dominated problems. To address this challenge, Fries et al. [26] proposed a parametric latent-space dynamics identification (LaSDI) framework, in which an auto-encoder was applied for nonlinear projection and a set of local dynamics identification (DI) models were introduced to identify local latent-space dynamics. The LaSDI framework can be viewed as a generalization of the aforementioned non-intrusive ROMs built upon latent-space dynamics identification since it allows linear or nonlinear projection and enables latent-space dynamics to be captured by flexible DI models based on general nonlinear functions. However, the lack of interaction between the auto-encoder and the DI models due to a sequential training procedure leads to a strong dependency on the complexity and quality of the latent-space dynamics on the auto-encoder architecture, which could pose challenges to the subsequent dynamics identification and thus affect the model performances. Most importantly, all these approaches rely on *predefined* training samples, such as uniform or Latin hypercube sampling, which may not be optimal for the best model performance in the prescribed parameter space.

To address these issues, we present a gLaSDI framework for accurate and efficient physics-informed data-driven reduced-order modeling. For an optimal sampling of the parameter space, an adaptive greedy sampling algorithm integrated with a physics-informed residual-based error indicator is introduced to search for optimal and minimal training samples on the fly. The proposed gLaSDI framework contains an auto-encoder for nonlinear projection to discover intrinsic latent representations and a set of local DI models to capture local latent-space dynamics, which is further exploited by an efficient k -nearest neighbors (k-NN) convex interpolation scheme. The auto-encoder training and dynamics identification take place interactively to achieve an optimal identification of simple latent-space dynamics. Our numerical experiments demonstrate the effectiveness of the proposed gLaSDI framework and a considerable speed-up is achieved.

2 Greedy Latent Space Dynamics Identification (gLaSDI)

Let us consider a system of ODEs characterizing a parameterized nonlinear dynamical system, which can also be considered as a semi-discretized equation of a system of partial differential equations (PDEs) with a spatial domain Ω ,

$$\frac{d\mathbf{u}(t; \boldsymbol{\mu})}{dt} = \mathbf{f}(\mathbf{u}, t; \boldsymbol{\mu}), \quad t \in [0, T], \quad \text{with} \quad \mathbf{u}(0; \boldsymbol{\mu}) = \mathbf{u}_0(\boldsymbol{\mu}), \quad (1)$$

where T is the final time; $\boldsymbol{\mu}$ is the parameter in a parameter domain \mathcal{D} ; $\mathbf{u} \in \mathbb{R}^{N_u}$ is the parameterized time-dependent solution to the dynamical system with an initial state \mathbf{u}_0 ; \mathbf{f} denotes the velocity. With the implicit backward Euler time integrator, an approximate solution to Eq. (1) can be obtained by solving the nonlinear system of equations $\mathbf{u}_n = \mathbf{u}_{n-1} + \Delta t \mathbf{f}_n$, where $\mathbf{u}_n := \mathbf{u}(t^n; \boldsymbol{\mu})$, and $\mathbf{f}_n := \mathbf{f}(\mathbf{u}(t^n; \boldsymbol{\mu}), t^n; \boldsymbol{\mu})$. The residual function is defined as

$$\mathbf{r}(\mathbf{u}_n; \mathbf{u}_{n-1}, \boldsymbol{\mu}) = \mathbf{u}_n - \mathbf{u}_{n-1} - \Delta t \mathbf{f}_n. \quad (2)$$

The physics-informed adaptive greedy sampling is performed on a discrete parameter space ($\mathcal{D}^h \subseteq \mathcal{D}$) and $\mathbb{D} \subseteq \mathcal{D}^h$ denotes a set of N_μ selected training sample points. Let's consider the i -th training sample point $\boldsymbol{\mu}^{(i)} \in \mathcal{D}^h$ and $\mathbf{u}_n^{(i)}$ as the associated solution at the n -th time step of the dynamical system in Eq. (1). The solutions at all time steps are arranged in a snapshot matrix denoted as $\mathbf{U}^{(i)} = [\mathbf{u}_0^{(i)}, \dots, \mathbf{u}_{N_t}^{(i)}]$. The *reconstruction loss* of the auto-encoder is defined as $\mathcal{L}_{recon} := \|\mathbf{U} - \hat{\mathbf{U}}\|_{L_2}^2$, where $\mathbf{U} = \{\mathbf{U}^{(j)}\}_{j=1}^{N_\mu}$ and $\hat{\mathbf{U}} = \{\hat{\mathbf{U}}^{(j)}\}_{j=1}^{N_\mu}$, $\hat{\mathbf{U}}^{(i)} = [\hat{\mathbf{u}}_0^{(i)}, \dots, \hat{\mathbf{u}}_{N_t}^{(i)}]$, $\hat{\mathbf{u}}_n^{(i)} = \phi_d(\mathbf{z}_n^{(i)})$, and $\mathbf{z}_n^{(i)} = \phi_e(\mathbf{u}_n^{(i)})$; ϕ_e is an encoder and ϕ_d is a decoder; $\mathbf{z}_n^{(i)} \in \mathbb{R}^{N_z}$ denotes the latent variable with $N_z \ll N_u$.

Given the latent variable in the discrete form $\mathbf{Z}^{(i)} = [\mathbf{z}_0^{(i)}, \dots, \mathbf{z}_{N_t}^{(i)}]$ obtained from the encoder, the governing equation of its dynamics is approximated by a DI model, ψ_{DI} , based on a user-defined library of basis functions $\Theta(\mathbf{Z}^{(i)T})$, e.g., polynomial, trigonometric, and exponential functions, such that $\dot{\mathbf{Z}}^{(i)T} \approx \hat{\mathbf{Z}}^{(i)T} = \Theta(\mathbf{Z}^{(i)T})\Xi^{(i)}$, where $\Xi^{(i)}$ is a coefficient matrix. See [26] for more details.

To identify simple and smooth latent-space dynamics, two additional loss functions, $\mathcal{L}_{\dot{\mathbf{z}}}$ and $\mathcal{L}_{\dot{\mathbf{u}}}$, are constructed to enable interactive training of the auto-encoder and the DI model (*interactive Autoencoder-DI training*), as illustrated in Fig. 1, where $\mathcal{L}_{\dot{\mathbf{z}}} := \|\dot{\mathbf{Z}} - \hat{\dot{\mathbf{Z}}}\|_{L_2}^2$ ensures the consistency of the *latent-space dynamics gradients*, and $\mathcal{L}_{\dot{\mathbf{u}}} := \|\dot{\mathbf{U}} - \hat{\dot{\mathbf{U}}}\|_{L_2}^2$ ensures the consistency of the *physical dynamics gradients*. Note that $\dot{\mathbf{Z}}$ and $\dot{\mathbf{U}}$ can be computed through the chain rule, as illustrated in Fig. 1. The local DI models are considered to be point-wise (see more details in [26]), which means each local DI model is associated with a distinct sampling point in the parameter space. Hence, each sampling point has an associated DI coefficient matrix.




Figure 1: Schematics of the gLaSDI algorithm. The black square boxes in the right panel indicate the sampled parameter points in a 2D parameter space.

It is expected that the latent-space dynamics within a small neighborhood of a parameter point are similar. To exploit the local latent-space dynamics captured by the local DI models for enhanced parameterization and efficiency, a k -NN convexity-preserving partition-of-unity interpolation scheme based on Shepard function [27] or inverse distance weighting is employed. The associated DI coefficient matrix of a testing parameter $\mu \in \mathcal{D}$ is obtained by a convex interpolation of coefficient matrices of its k -nearest neighbors (existing sampling points), $\Xi_{interp} = \sum_{i \in \mathcal{N}_k(\mu)} \Psi^{(i)}(\mu) \Xi^{(i)}$. The k -nearest neighbors are selected based on the Mahalanobis distance between the testing and the training (existing) parameters. The interpolation functions $\Psi^{(i)}(\mu)$ with an inverse-distance kernel function is employed, preserving convexity and a partition of unity.

To address the issues of parameter dependency of local latent-space dynamics efficiently and effectively, a physics-informed adaptive greedy sampling procedure is introduced to construct an optimal training database $\mathcal{DB} = \{\mathbf{U}^{(i)}\}_{i=1}^{N_\mu}$ on the fly during off-line interactive autoencoder-DI training, as illustrated in Fig. 1, which corresponds to a set of parameters $\mathbb{D} = \{\mu^{(i)}\}_{i=1}^{N_\mu}$ sampled from \mathcal{D}^h with $N_\mu \ll |\mathcal{D}^h|$. The training database is first initialized with a small set of parameters located, e.g., at the corners or at the center of the parameter space. To enhance sampling reliability and quality, the model training is performed for a number of epochs before every greedy sampling, as illustrated

in Fig. 1. To expedite the sampling procedure, a random-subset evaluation strategy is adopted. At the v -th sampling iteration, a small set of candidate parameters, $\mathbb{D}^{sub} \subseteq \mathcal{D}^h$ ($\mathbb{D}^{sub} \cap \mathbb{D}_{v-1} = \emptyset$), are randomly selected and the parameter $\boldsymbol{\mu}^*$ that maximizes a *residual-based error indicator*, $e^{res}(\boldsymbol{\mu})$, is selected to update $\mathbb{D}_v = \{\mathbb{D}_{v-1}, \boldsymbol{\mu}^*\}$ and $\mathcal{DB}_v = \{\mathcal{DB}_{v-1}, \mathbf{U}^*\}$. The residual error indicator is defined as

$$e^{res} = \frac{1}{N_{ts} + 1} \sum_{n=0}^{N_{ts}} \|\mathbf{r}(\hat{\mathbf{u}}_n; \hat{\mathbf{u}}_{n-1}, \boldsymbol{\mu})\|_{L_2}, \quad (3)$$

where the residual $\mathbf{r}(\hat{\mathbf{u}}_n; \hat{\mathbf{u}}_{n-1}, \boldsymbol{\mu})$ is physics-informed and obtained from Eq. (2) that relies on discretized governing equations; $N_{ts} \approx 0.1N_t$ to enhance the computational efficiency. The residual error indicator is calculated based on only predictions and positively correlated with the maximum relative error, as demonstrated in Fig. 1. The maximum residual error (e_v^{res}) of the training samples at the v -th iteration can be used to estimate the maximum relative error (e_v^{max}) in the parameter space based on a linear least-square fit between $\mathbf{E}_v^{res} = \{e^{res}(\hat{\mathbf{U}}(\boldsymbol{\mu}))\}_{\boldsymbol{\mu} \in \mathbb{D}_v}$ and $\mathbf{E}_v^{max} = \{e^{max}(\mathbf{U}(\boldsymbol{\mu}), \hat{\mathbf{U}}(\boldsymbol{\mu}))\}_{\boldsymbol{\mu} \in \mathbb{D}_v}$. The greedy sampling procedure continues until a prescribed number of sampled points or a target error tolerance is reached. The trained gLaSDI model can then be applied to efficiently predict dynamical solutions given a testing parameter in the prescribed parameter space.

3 Results and Discussion

The performance of gLaSDI is demonstrated by solving a two-dimensional parameterized radial advection problem and compared with that of LaSDI [26] (without autoencoder-DI interactive training and adaptive greedy sampling):

$$\frac{\partial u}{\partial t} + \mathbf{v} \cdot \nabla u = 0, \quad \Omega = [-1, 1] \times [-1, 1], \quad t \in [0, 3], \quad (4)$$

with a boundary condition $u(\mathbf{x}, t; \boldsymbol{\mu}) = 0$ on $\partial\Omega$ and an initial condition $u(\mathbf{x}, 0; \boldsymbol{\mu}) = \sin(w_1 x_1) \sin(w_2 x_2)$ parameterized by $\boldsymbol{\mu} = \{w_1, w_2\}$, where $\mathbf{v} = \frac{\pi}{2} d [x_2, -x_1]^T$ denotes the fluid velocity with $d = (1 - x_1^2)^2 (1 - x_2^2)^2$. The spatial domain is discretized by first-order periodic square finite elements constructed on a uniform grid of 96×96 discrete points. The fourth-order Runge-Kutta explicit time integrator with a uniform time step of $\Delta t = 0.01$ is employed. A discrete parameter space \mathcal{D}^h is constituted by the parameters of the initial condition, including $w_1 \in [1.5, 1.8]$ and $w_2 \in [2.0, 2.3]$, each with 21 evenly distributed discrete points in the respective parameter range. The distributions of the solution field and its gradient at a few time steps for the parameter case ($w_1 = 1.5, w_2 = 2.0$) are shown in Fig. 1. The gLaSDI model consists of *Linear* DI models and an auto-encoder that has an architecture of 9,216-100-3-100-9,216, with the numbers denoting the number of neurons in each hidden layer, and the latent dimension as 3. The training and testing are performed on a NVIDIA V100 (Volta) GPU with 3,168 NVIDIA CUDA Cores and 64 GB GDDR5 GPU Memory. The open-source TensorFlow library [28] and the Adam optimizer [29] are employed for model training. The gLaSDI training is performed until the total number of sampled parameter points reaches 25. A LaSDI model with the same architecture of the auto-encoder and DI models is trained using 25 predefined training points uniformly distributed in a 5×5 grid in the parameter space.

Owing to the interactive autoencoder-DI training, gLaSDI is able to capture simpler and smoother latent-space dynamics than LaSDI that has sequential and decoupled training of the auto-encoder and DI models, as shown in Figs. 2(a-b). A better agreement between the latent-space dynamics predicted by the encoder and the DI models is achieved by gLaSDI. Figs. 2(c-d) show that gLaSDI achieves a maximum relative error of 2.0% in the whole parameter space, lower than 5.4% of LaSDI.

When the parameter space is enlarged to $w_1 \in [1.5, 2.0]$ and $w_2 \in [2.0, 2.5]$, gLaSDI maintains a high accuracy with a maximum relative error of 3.3%, much lower than 24% of LaSDI, as shown in Figs. 3(c-d). The results demonstrate that given the same number of training parameters, gLaSDI with the physics-informed adaptive and sparse sampling can intelligently identify the optimal training parameter points to achieve higher accuracy than LaSDI based on predefined uniformly distributed training parameters. It is interesting to note that changing the parameter space affects the distribution of gLaSDI sampling. Compared with the high-fidelity simulation based on MFEM [30], the gLaSDI model achieves $121 \times$ speed-up.




Figure 2: Comparison between LaSDI and gLaSDI with the same model architecture for the radial advection problem with parameters $w_1 \in [1.5, 1.8]$ and $w_2 \in [2.0, 2.3]$. The number on each box denotes the maximum relative error of the associated parameter case. The black square boxes indicate the sampled training parameter points.




Figure 3: Comparison between LaSDI and gLaSDI with the same model architecture for the radial advection problem with parameters $w_1 \in [1.5, 2.0]$ and $w_2 \in [2.0, 2.5]$. The number on each box denotes the maximum relative error of the associated parameter case. The black square boxes indicate the sampled training parameter points.

A potential limitation is the residual computation, which might not be straight-forward, for example, if the source code of the full-order model is not available. Then, the physics information needs to be post-processed with the gLaSDI solution, which might not be accurate. Another limitation is that the training of the auto-encoder can be computationally expensive as the data size becomes large, which can be improved by applying the convolutional auto-encoder. The training efficiency can potentially be further improved by a combination of *pre-training* and *re-training*. The proposed gLaSDI framework is general and not restricted by the use of auto-encoders and DI models. Depending on applications, linear compression techniques with much lower training cost could be employed. Further, other system identification techniques or operator learning algorithms could be applied to identify latent-space dynamics. The auto-encoder architecture can be optimized to maximize generalization performance by integrating automatic neural architecture search into the proposed framework. The parameterization in this study only considers the parameters from the initial conditions of the problems. The proposed framework can be easily extended to account for other parameterization types, such as of material properties, which will be useful for inverse problems. All the codes to regenerate the results in this paper can be found in the GitHub repository (<https://github.com/LLNL/libROM/tree/gLaSDI/python/gLaSDI>). The following existing assets have been used: SindyAutoencoder with MIT license (<https://github.com/kpchamp/SindyAutoencoders>) and the MFEM library with BSD-3-Clause (<https://github.com/mfem/mfem>).

4 Broader impact

This paper introduces a novel machine learning enhanced physics-guided data-driven reduced-order modeling strategy with adaptive sampling to accelerate high-dimensional physical simulations. The proposed gLaSDI framework is expected to have broad impacts on the computational science community and application potentials in a wide range of engineering and scientific domains. There is no negative consequence on ethics and society in this work.

Appendix

An additional experiment and results are presented to further demonstrate the effectiveness of the proposed framework. A two-dimensional (2D) parameterized inviscid Burgers equation is considered

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = \frac{1}{Re} \Delta \mathbf{u}, \quad \Omega = [-3, 3] \times [-3, 3], \quad t \in [0, 1], \quad (5)$$

with a boundary condition $\mathbf{u}(\mathbf{x}, t; \boldsymbol{\mu}) = \mathbf{0}$ on $\partial\Omega$, an initial condition $\mathbf{u}(\mathbf{x}, 0; \boldsymbol{\mu}) = ae^{-\frac{\|\mathbf{x}\|^2}{w^2}}$ parameterized by $\boldsymbol{\mu} = \{a, w\}$, and a Reynolds number $Re = 10,000$. A uniform spatial discretization with 60×60 discrete points is applied. The first order spatial derivative and the diffusion term are approximated by the backward difference scheme and the central difference scheme, respectively. The full-order model solutions are obtained by solving the semi-discretized system with the implicit backward Euler time integrator and a uniform time step of $\Delta t = 1/200$. A discrete parameter space \mathcal{D}^h is constituted by the parameters of the initial condition, including the width, $w \in [0.9, 1.1]$, and the amplitude, $a \in [0.7, 0.9]$, each with 21 evenly distributed discrete points in the respective parameter range. The solution fields of the first velocity component at different time steps for the parameter case ($a = 0.7, w = 0.9$) are shown in Fig. 4(a). The auto-encoder with an architecture of 7,200-100-5-100-7,200 (latent dimension as 5) and *quadratic* DI models are employed. The gLaSDI training is performed until the total number of sampled parameter points reaches 36. A LaSDI model with the same architecture of the auto-encoder and DI models is trained using 36 predefined training points uniformly distributed in a 6×6 grid for the parameter space.

Figs. 4(b-c) show that gLaSDI identifies much simpler and smoother latent-space dynamics than LaSDI, with a better agreement between the encoder and the DI predictions, which is attributed by the interactive autoencoder-DI training of gLaSDI, as illustrated in Fig. 1 and described in Section 2. Figs. 4(d-e) show that gLaSDI achieves the maximum relative error of 5% in the whole parameter space, much lower than 255% of LaSDI. The poor accuracy of LaSDI could be caused by the deviation between the DI predicted dynamics and the encoder predicted dynamics. It is also observed that gLaSDI tends to have denser sampling in the lower range of the parameter space. This demonstrates the importance of the physics-informed greedy sampling procedure. Compared with the high-fidelity simulation based on an in-house Python code, the gLaSDI model achieves $871 \times$ speed-up.

When the latent dimension is reduced from 5 to 3 and the polynomial order of the DI models is reduced from quadratic to linear, gLaSDI learns simpler latent-space dynamics, as shown in Fig. 5(b), and maintains a high accuracy with a maximum relative error of 4.6%, much lower than 22% of LaSDI, as shown in Figs. 5(c-d). Furthermore, simplifying the latent-space dynamics enhances the reduced-order modeling efficiency. Compared with the high-fidelity simulation, the gLaSDI model achieves $2,658 \times$ speed-up, which is 3.05 times the speed-up achieved by the gLaSDI model with a latent dimension of 5 and quadratic DI models.

Acknowledgements

This work was performed at Lawrence Livermore National Laboratory and partially funded by two LDRDs (21-FS-042 and 21-SI-006). Lawrence Livermore National Laboratory is operated by Lawrence Livermore National Security, LLC, for the U.S. Department of Energy, National Nuclear Security Administration under Contract DE-AC52-07NA27344.

References

- [1] S. McBane and Y. Choi, "Component-wise reduced order model lattice-type structure design," *Computer Methods in Applied Mechanics and Engineering*, vol. 381, p. 113 813, 2021.




Figure 4: Comparison between LaSDI and gLaSDI with an auto-encoder of 7,200-100-5-100-7,200 and *quadratic* DI models for the 2D Burgers problem. The number on each box denotes the maximum relative error of the associated parameter case. The black square boxes indicate the sampled training parameter points.




Figure 5: Comparison between LaSDI and gLaSDI with an auto-encoder of 7,200-100-3-100-7,200 and *linear* DI models for the 2D Burgers problem. The number on each box denotes the maximum relative error of the associated parameter case. The black square boxes indicate the sampled training parameter points.

- [2] D. Galbally, K. Fidkowski, K. Willcox, and O. Ghattas, “Non-linear model reduction for uncertainty quantification in large-scale inverse problems,” *International journal for numerical methods in engineering*, vol. 81, no. 12, pp. 1581–1608, 2010.
- [3] Y. Choi, G. Boncoraglio, S. Anderson, D. Amsallem, and C. Farhat, “Gradient-based constrained optimization using a database of linear reduced-order models,” *Journal of Computational Physics*, vol. 423, p. 109787, 2020.
- [4] S. McBane, Y. Choi, and K. Willcox, “Stress-constrained topology optimization of lattice-like structures using component-wise reduced order models,” *Computer Methods in Applied Mechanics and Engineering*, vol. 400, p. 115525, 2022.
- [5] S. Wang, E. d. Sturler, and G. H. Paulino, “Large-scale topology optimization using preconditioned krylov subspace methods with recycling,” *International journal for numerical methods in engineering*, vol. 69, no. 12, pp. 2441–2468, 2007.
- [6] D. A. White, Y. Choi, and J. Kudo, “A dual mesh method with adaptivity for stress-constrained topology optimization,” *Structural and Multidisciplinary Optimization*, vol. 61, no. 2, pp. 749–762, 2020.
- [7] Y. Choi, C. Farhat, W. Murray, and M. Saunders, “A practical factorization of a schur complement for pde-constrained distributed optimal control,” *Journal of Scientific Computing*, vol. 65, no. 2, pp. 576–597, 2015.
- [8] R. C. Smith, *Uncertainty quantification: theory, implementation, and applications*. Siam, 2013, vol. 12.

- [9] G. Biros *et al.*, *Large-scale inverse problems and quantification of uncertainty*. John Wiley & Sons, 2011.
- [10] B. Kim, V. C. Azevedo, N. Thuerey, T. Kim, M. Gross, and B. Solenthaler, “Deep fluids: A generative network for parameterized fluid simulations,” in *Computer Graphics Forum*, Wiley Online Library, vol. 38, 2019, pp. 59–70.
- [11] X. Xie, G. Zhang, and C. G. Webster, “Non-intrusive inference reduced order model for fluids using deep multistep neural network,” *Mathematics*, vol. 7, no. 8, p. 757, 2019.
- [12] C. Hoang, K. Chowdhary, K. Lee, and J. Ray, “Projection-based model reduction of dynamical systems using space–time subspace and machine learning,” *Computer Methods in Applied Mechanics and Engineering*, vol. 389, p. 114 341, 2022.
- [13] T. Kadeethum, F. Ballarin, Y. Choi, D. O’Malley, H. Yoon, and N. Bouklas, “Non-intrusive reduced order modeling of natural convection in porous media using convolutional autoencoders: Comparison with linear subspace techniques,” *Advances in Water Resources*, p. 104 098, 2022.
- [14] J. R. Koza, “Genetic programming as a means for programming computers by natural selection,” *Statistics and computing*, vol. 4, no. 2, pp. 87–112, 1994.
- [15] M. Schmidt and H. Lipson, “Distilling free-form natural laws from experimental data,” *science*, vol. 324, no. 5923, pp. 81–85, 2009.
- [16] M. Cranmer *et al.*, “Discovering symbolic models from deep learning with inductive biases,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 17 429–17 442, 2020.
- [17] M. Cranmer, *Pysr: Fast & parallelized symbolic regression in python/julia*, 2020.
- [18] S. L. Brunton, J. L. Proctor, and J. N. Kutz, “Discovering governing equations from data by sparse identification of nonlinear dynamical systems,” *Proceedings of the national academy of sciences*, vol. 113, no. 15, pp. 3932–3937, 2016.
- [19] K. Champion, B. Lusch, J. N. Kutz, and S. L. Brunton, “Data-driven discovery of coordinates and governing equations,” *Proceedings of the National Academy of Sciences*, vol. 116, no. 45, pp. 22 445–22 451, 2019.
- [20] B. Peherstorfer and K. Willcox, “Data-driven operator inference for nonintrusive projection-based model reduction,” *Computer Methods in Applied Mechanics and Engineering*, vol. 306, pp. 196–215, 2016.
- [21] E. Qian, B. Kramer, B. Peherstorfer, and K. Willcox, “Lift & learn: Physics-informed machine learning for large-scale nonlinear dynamical systems,” *Physica D: Nonlinear Phenomena*, vol. 406, p. 132 401, 2020.
- [22] P. Benner, P. Goyal, B. Kramer, B. Peherstorfer, and K. Willcox, “Operator inference for non-intrusive model reduction of systems with non-polynomial nonlinear terms,” *Computer Methods in Applied Mechanics and Engineering*, vol. 372, p. 113 433, 2020.
- [23] R. Geelen, S. Wright, and K. Willcox, “Operator inference for non-intrusive model reduction with nonlinear manifolds,” *arXiv preprint arXiv:2205.02304*, 2022.
- [24] M. Guo, S. McQuarrie, and K. Willcox, “Bayesian operator inference for data-driven reduced order modeling,” *arXiv preprint arXiv:2204.10829*, 2022.
- [25] S. A. McQuarrie, P. Khodabakhshi, and K. E. Willcox, “Non-intrusive reduced-order models for parametric partial differential equations via data-driven operator inference,” *arXiv preprint arXiv:2110.07653*, 2021.
- [26] W. D. Fries, X. He, and Y. Choi, “LaSDI: Parametric latent space dynamics identification,” *Computer Methods in Applied Mechanics and Engineering*, vol. 399, p. 115 436, 2022.
- [27] D. Shepard, “A two-dimensional interpolation function for irregularly-spaced data,” in *Proceedings of the 1968 23rd ACM national conference*, 1968, pp. 517–524.
- [28] M. Abadi *et al.*, “Tensorflow: A system for large-scale machine learning,” in *12th USENIX symposium on operating systems design and implementation (OSDI 16)*, 2016, pp. 265–283.
- [29] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [30] R. Anderson *et al.*, “MFEM: A modular finite element methods library,” *Computers & Mathematics with Applications*, vol. 81, pp. 42–74, 2021.

Checklist

1. For all authors...

- (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] See Section 3.
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A] The proposed work does not involve any societal aspects.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
- (a) Did you state the full set of assumptions of all theoretical results? [N/A]
 - (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See Section 3.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] see Section 3. More details can be found in the following GitHub repository, <https://github.com/LLNL/libROM/tree/gLaSDI/python/gLaSDI>.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Section 3.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- (a) If your work uses existing assets, did you cite the creators? [Yes] See Section 3.
 - (b) Did you mention the license of the assets? [Yes] See Section 3.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] See Section 3.
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [N/A] The license of the used asset allows our usage of the codes and distribution.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]