SAND2022-7907C

Key challenges in adapting data-driven approaches to modeling of inertial fusion concepts

William E. Lewis,[1] Patrick F. Knapp,[1] and Kristian Beckwith[1]

Sandia National Laboratories, Albuquerque, New Mexico 87185 USA

I. EXECUTIVE SUMMARY

Data-driven approaches are increasingly being adopted to study magnetic confinement fusion (MCF), magneto-inertial confinement fusion (MIF), and inertial confinement fusion (ICF) schema. These methods have demonstrated applications in this space ranging from automated extraction of features from experimental data and dynamical systems control during the operation phase to intelligent optimization algorithms and multi-objective statistical inference. At their heart, data-driven applications rely on models which are computationally cheap to evaluate and/or greedy model evaluation that exploits the trade-off between exploration of unsampled model inputs and an approximate understanding of output structure based on previous sampled inputs. For the computationally expensive and complex multi-physics models required to understand fusion experiments this typically means applying more-or-less black-box machine learned models that are quick to evaluate, transferring the primary effort to an offline learning phase whereupon the algorithm discovers a mapping directly from physics model input space to (typically highly spatio-temporally integrated) experimentally observable quantities. At the same time, it is expected that there is still significant unknown physics important for understanding ICF/MIF targets en route to an engineering-relevant end-state for inertial fusion energy (IFE) designs. Improving the resolution of experimental diagnostics and physical accuracy of simulation tools is clearly a requirement on this path. Due to significant computational and experimental costs, it is critical that the fusion community continue to push towards the application of state-of-the-art reduced order modeling approaches capable of capturing the dynamics of the system under consideration with high spatio-temporal fidelity. In this white paper, we provide a brief description of the status of the application of machine learning to MIF/ICF and highlight exemplars of current state-of-the-art in the ideas machine-learning based ROM techniques. The adaptation of these methods to problems in ICF presents a significant open challenge that should be addressed by the community to accelerate physics discovery and aid in design of IFE relevant systems.

II. INTRODUCTION

Significant effort in high energy density physics and ICF related research has been placed on applying methods in data science and statistics to accelerate experiment design, physics discovery, and improve uncertainty quantification. In essentially all cases, these efforts involve the evaluation of one or more forward models of the form

$$f(X(\theta,t),\theta) \to \{g_i\}, \tag{1}$$

where f represents a composition of an often-complex set of simulations (e.g. multi-resolution or simulation restarts) with experiment design parameters θ, and synthetic di-

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.

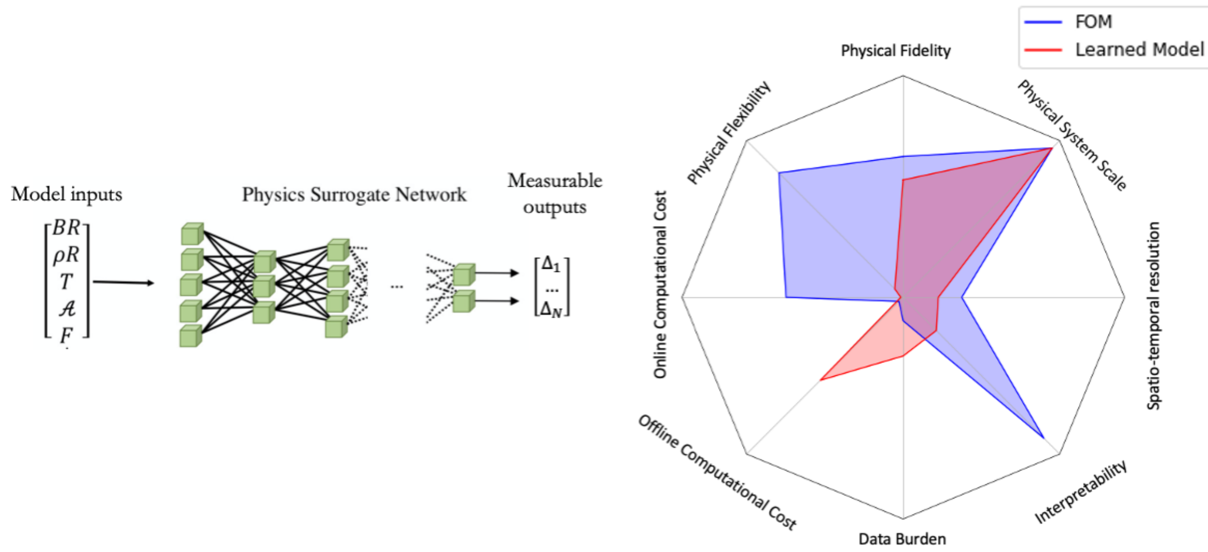

FIG. 1. Learned emulators or surrogate models generally trade online computation cost for offline cost, but often lose additional desirable capabilities of the original full order model.

agnostic calculations resulting in a set of (potentially time dependent or spatio-temporally averaged) states $X(\theta, t)$, which maps to a set of experimental observables $\{g_i\}$. To date, the program of machine learning as applied to ICF modeling has largely consisted of attempting to learn a function $\hat{f}(\theta) \rightarrow \{g_i\}$ that maps directly from the set of model inputs, θ, directly to the set of observables $\{g_i\}$ without reference to the time dependent physical system state. The function $\hat{f}$ is typically estimated by optimizing an objective function comparing the true function f to $\hat{f}$ on a set of example points

$$\mathcal{L}(f, \hat{f}, \mathbf{w}_i) \equiv \sum_{\mathbf{p}} d(f_{\mathbf{p}}, \hat{f}_{\mathbf{p}}(\mathbf{w}_i)) + R(*), \tag{2}$$

over any free parameters $\mathbf{w}_i$ prescribing $\hat{f}$. In the above, d is metric describing the discrepancy between the f and $\hat{f}$, while R is an additional function which may be used to regularize the problem or enforce constraints in the optimization.

We have selected as an example, the work of Refs. 1–3 for this purpose, but the qualitative aspects of this example transfer to other recent works (*e.g.* Refs. 4–7) that are more-or-less variations on the theme of optimizing Eq. 2 over parameters $\mathbf{w}_i$, with the outputs of f and $\hat{f}$ being highly integrated measurements. In Fig. 1 we present a qualitative assessment of the impact of this type of model emulation on a handful of non-exhaustive, correlated, and admittedly incompletely defined set of traits that could describe a modeling approach. The figure indicates qualitative judgements for relative relations between the different traits. It is indicated for example that application of the FOM, f, for an experimental analysis involves significant online computational costs. In this case, a single FOM evaluation takes order 10-100 CPU hours on an HPC system, with $O(10\text{k})$ simulations needed to conduct an analysis of a single experiment. Contrast this with the learned model, $\hat{f}$, for which the evaluation time is < 1 ms on a personal laptop. This computational cost gets shifted into a one-time offline computational cost to generate training data and optimize the machine learning model (provided the training data domain is sufficient to cover realistic experimental outcomes). At the same time, while training a model may require generation and storage (data burden) of somewhat more data than a single experimental analysis, this also presents a greedy approach, where new data generated only when needed to tune the model. Clearly, the learned model computational cost for analysis of experiments scales much more favorably with the number of experiments to be analyzed. We note however that this comes at the expense of physical flexibility. For example, in the case under consideration, a frozen in flux law was assumed for the magnetic field along with a static fusion fuel background in the FOM. If these assumptions are to be lifted, or, if one wishes to work in a spherical rather than cylindrical geometry, the emulator must be retrained. We also indicate a reduced "physical fidelity" of the surrogate model as the neural network-based approach does not fundamentally enforce physics constraints, but merely is shown to agree quantitatively with the physics model over the domain from which training data was drawn. Therefore, extrapolation presents a high risk for faulty analysis. In addition, as a black box map, the emulator may obscure the physics-based reasons for the model response.

III. BASIC RESEARCH NEEDS

The approach of directly surrogating inputs to outputs currently appears to be the most mature application of machine learning in the ICF community. The continued development of these approaches will certainly lead to additional discovery and therefore represents a continuing and crucial basic research need. However, due to the relative maturity of this paradigm, here we highlight areas in which basic research is required to push the types of traits indicated in Fig. 1 in new and favorable directions for ICF and IFE applications. Rather than being overly prescriptive, we will point out approaches that exist within the broader machine learning and reduced order modeling communities that have yet to see widespread adoption in the ICF community, indicating the potential impact that could be observed from adapting these methods to ICF applications.

A. Develop and integrate surrogates of physics/algorithm submodules

One way to improve the physical fidelity and flexibility of surrogate models is to surrogate submodules of multiphysics codes. Recently, for example, Ref. 8 demonstrated the inclusion of a neural network model to emulate atomic physics computations of non-local thermodynamic equilibrium spectral opacities which was included inline with HYDRA calculations. In Ref. 9, the inviscid Euler equations are solved using a convolutional neural network in the pressure projection component of the velocity update step to enforce incompressibility. This replaces the typical approach of numerically solving the large sparse linear system corresponding to the associated Poisson equation, which represents the most computationally intensive component of the algorithm. In the context of magneto-hydrodynamics (MHD) equations, for example, some algorithms for magnetic diffusion and radiation transport involve numerically solving large sparse linear systems, so that a similar approach may approximate the solution or be used to precondition iterative solvers, potentially significantly reducing computational costs. We also note that there is a growing effort to apply ideas in randomized linear algebra to matrix decomposition and data-driven discovery of sparse representations[10–13]. Application of these methods to ICF relevant codes may prove fruitful for improving algorithm efficiency. Finally, we note recent work in the context of large eddy simulations (LES) to utilize neural networks to replace typically gradient-based subgrid scale (SGS) models[14–16]. Adapting these methods to ICF and HED problems with a more realistic physics is an open and potentially transformative application space. As for any algorithm development endeavor, assessing the computational cost – physical accuracy trade-off of these methods will be critical, requiring input from experts in hydrodynamic and MHD numerical methods as well as in the fields of data science and statistics.

B. Develop surrogate models for dynamical system evolution

A more holistic approach of directly surrogating the dynamical evolution of simulation field quantities, i.e. developing reduced order models (ROMs), is another active area of research whose adaptation to ICF problems presents a basic research challenge. Recent efforts along these lines have treated problems including 2D chemically reactive flows, shock propagation in 1D Burger's equation[17], nonlinear evolution of the Rayleigh-Taylor instability[18], and many others[19,20]. These fall into a wide range of different categories ranging from more "standard" linear projection based methods (e.g. POD) to those based on nonlinear methods, such as the work of Ref. 17, where a CNN autoencoder architecture is trained to nonlinearly project the data onto a low-dimensional manifold where the dynamics can be optimally, and kinematically consistently evolved. Furthermore, by incorporating various regression methods in the ML community, it may be possible to ensure that such models may maintain flexibility and fidelity over a range of possible initial conditions, such as variations in geometry or other parametrized boundary conditions (see e.g. Ref. 17). Of particular interest, many of these methods allow analytic access to gradients or enable computational speedups of high enough order so that sensitivity methods for uncertainty quantification are readily applicable to develop deeper physical insight. Importantly, many of the most mature approaches in projection based ROM methods fundamentally fail in the presence of traveling waves, transients, intermittent phenomena, continuous spectra, and strong nonlinearity. A number of solutions have been proposed to address these issues in specific cases[18,21,22]. However, the development of more generalized solutions appropriate to problems in ICF/MIF will be required to apply new developments in the ROM space.

C. Pursue data-compression ideas from machine learning and compressed sensing

A critical consideration when pursuing the application of machine learning to problems in ICF is the significant data burden represented by such an endeavor. Complex 3D simulations with potentially highly resolved spatio-temporal features in multiple self-consistent field quantities often result in the need for significant inline data reduction strategies. These approaches typically involve expert-based decisions on what data can be safely discarded while maintaining integrity of the fundamental physics insights derived. Clearly the development of methods outlined above require access to sufficiently large collections of data containing the quantities and fields of interest to be learned by ML models. Data compression ideas utilized in ROM approaches, such as more standard linear decomposition methods (e.g. SVD, Fourier) that promote sparsity, auto-encoding neural networks[6,17,23], mrDMD[21], and even ideas from the field of compressed sensing should be brought to bear for reducing the data burden presented by the training requirements for these methods. Expert driven decisions regarding such representations (e.g. spherical/cylindrical harmonic bases, wavelet decomposition, etc.) will likely also feed into improved understanding in the application of data driven surrogate modeling efforts outlined above.

D. Develop and impose physics constrained models

Before concluding, we feel it important to emphasize that, as with any algorithm, there is potential for data-driven methods to produce unphysical results. Indeed, since many of these approaches lack insight into inherent physics constraints such as conservation laws or consistency of transport coefficients, or geometric constraints such as divergence-free magnetic fields, the risk is generally greater that ML based methods will fail to preserve important physical properties, even within the domain over which a given model was trained. Several approaches have been developed to produce results that approximately[24] or rigorously[25,26] obey physical laws. While it is not clear that these approaches significantly improve the fidelity of the results obtained, they may allow for better interpretability and improved extrapolatory power of the learned models. The continued development of such approaches is therefore critical to application spaces such as ICF where, for-example, performance cliffs and shifts in physics regimes may result in drastically different results between simulations and ML models that may go unnoticed. As a final note, we highlight that experimental data is also beginning to play a greater role in hunting for ways to improve the physical fidelity and quantification of uncertainty of simulation results (see *e.g.* Refs. 7 and 27). In the ICF and more broadly HED application spaces, such methods will additionally prove invaluable as they may reduce the data requirements for training machine learning models where the cost of simulations and experiment will continue to result in a data-starved environment.

IV. CONCLUSIONS

We have highlighted several areas in which basic research is needed to adapt existing methods in the data science literature to problems in ICF. We have chosen to focus on methods that reach well beyond recent successful work at adapting machine learning to address the problem of modeling the direct relation between simulation inputs and observables, however, we emphasize that there is still significant room to extend and apply such methods. while this discussion is incomplete, the overriding message that we wish to communicate is that extensive research is needed to adapt these promising approaches to the relevant physics problems. Navigating these ideas through the valley of death, such that the full promise of these methods can be realized in this application space should represent a high priority for basic research efforts on the near term. The cross-cutting nature of ML applications to theory and simulation will require continued collaboration across academia and national laboratories among experts in the development and application of algorithms to problems in ICF and practitioners of data science.

V. ACKNOWLEDGEMENTS

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 or the United States Government.

VI. REFERENCES

[1] P. F. Schmit, P. F. Knapp, S. B. Hansen, M. R. Gomez, K. D. Hahn, D. B. Sinars, K. J. Peterson, S. A. Slutz, A. B. Sefkow, T. J. Awe, E. Harding, C. A. Jennings, G. A. Chandler, G. W. Cooper, M. E. Cuneo, M. Geissel, A. J. Harvey-Thompson, M. C. Herrmann, M. H. Hess, O. Johns, D. C. Lamppa, M. R. Martin, R. D. McBride, J. L. Porter, G. K. Robertson, G. A. Rochau, D. C. Rovang, C. L. Ruiz, M. E. Savage, I. C. Smith, W. A. Stygar, and R. A. Vesey, Phys. Rev. Lett. **113**, 155004 (2014).

[2] P. F. Knapp, P. F. Schmit, S. B. Hansen, M. R. Gomez, K. D. Hahn, D. B. Sinars, K. J. Peterson, S. A. Slutz, A. B. Sefkow, T. J. Awe, E. Harding, C. A. Jennings, M. P. Desjarlais, G. A. Chandler, G. W. Cooper, M. E. Cuneo, M. Geissel, A. J. Harvey-Thompson, J. L. Porter, G. A. Rochau, D. C. Rovang, C. L. Ruiz, M. E. Savage, I. C. Smith, W. A. Stygar, and M. C. Herrmann, Physics of Plasmas **22**, 056312 (2015), http://dx.doi.org/10.1063/1.4920948.

[3] W. E. Lewis, P. F. Knapp, S. A. Slutz, P. F. Schmit, G. A. Chandler, M. R. Gomez, A. J. Harvey-Thompson, M. A. Mangan, D. J. Ampleford, and K. Beckwith, Physics of Plasmas **28**, 092701 (2021), https://doi.org/10.1063/5.0056749.

[4] R. Nora, J. L. Peterson, B. K. Spears, J. E. Field, and S. Brandon, Statistical Analysis and Data Mining: The ASA Data Science Journal **10**, 230 (2017), https://onlinelibrary.wiley.com/doi/pdf/10.1002/sam.11344.

[5] J. A. Gaffney, S. T. Brandon, K. D. Humbird, M. K. G. Kruse, R. C. Nora, J. L. Peterson, and B. K. Spears, Physics of Plasmas **26**, 082704 (2019), https://doi.org/10.1063/1.5108667.

[6] R. Anirudh, J. J. Thiagarajan, P.-T. Bremer, and B. K. Spears, Proceedings of the National Academy of Sciences **117**, 9741 (2020), https://www.pnas.org/content/117/18/9741.full.pdf.

[7] K. D. Humbird, J. L. Peterson, J. Salmonson, and B. K. Spears, Physics of Plasmas **28**, 042709 (2021), https://doi.org/10.1063/5.0041907.

[8] G. Kluth, K. D. Humbird, B. K. Spears, J. L. Peterson, H. A. Scott, M. V. Patel, J. Koning, M. Marinak, L. Divol, and C. V. Young, Physics of Plasmas **27**, 052707 (2020), https://doi.org/10.1063/5.0006784.

[9] J. Tompson, K. Schlachter, P. Sprechmann, and K. Perlin, in *Proceedings of the 34th International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 70, edited by D. Precup and Y. W. Teh (PMLR, 2017) pp. 3424–3433.

[10] T. Sarlos, in *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)* (2006) pp. 143–152.

[11] V. Rokhlin, A. Szlam, and M. Tygert, SIAM Journal on Matrix Analysis and Applications **31**, 1100 (2010), https://doi.org/10.1137/080736417.

[12] P.-G. Martinsson, V. Rokhlin, and M. Tygert, Applied and Computational Harmonic Analysis **30**, 47 (2011).

[13] M. W. Mahoney, Found. Trends Mach. Learn. **3**, 123–224 (2011).

[14] F. Sarghini, G. de Felice, and S. Santini, Computers and Fluids **32**, 97 (2003).

[15] S. Rasp, M. S. Pritchard, and P. Gentine, Proceedings of the National Academy of Sciences **115**, 9684 (2018), https://www.pnas.org/doi/pdf/10.1073/pnas.1810286115.

[16]S. G. Rosofsky and E. A. Huerta, Phys. Rev. D **101**, 084024 (2020).

[17]K. Lee and K. T. Carlberg, J. Comput. Phys. **404** (2020), 10.1016/j.jcp.2019.108973.

[18]S. W. Cheung, Y. Choi, D. M. Copeland, and K. Huynh, "Local lagrangian reduced-order modeling for rayleigh-taylor instability by solution manifold decomposition," (2022).

[19]Y. Choi, P. Brown, W. Arrighi, R. Anderson, and K. Huynh, Journal of Computational Physics **424**, 109845 (2021).

[20]D. M. Copeland, S. W. Cheung, K. Huynh, and Y. Choi, Computer Methods in Applied Mechanics and Engineering **388**, 114259 (2022).

[21]J. N. Kutz, X. Fu, and S. L. Brunton, SIAM Journal on Applied Dynamical Systems **15**, 713 (2016), https://doi.org/10.1137/15M1023543.

[22]B. Lusch, J. N. Kutz, and S. L. Brunton, Nature Communications **9**, 4950 (2018).

[23]M. A. Kramer, AIChE Journal **37**, 233 (1991), https://aiche.onlinelibrary.wiley.com/doi/pdf/10.1002/aic.690370209.

[24]Y. Zhu, N. Zabaras, P.-S. Koutsourelakis, and P. Perdikaris, Journal of Computational Physics **394**, 56 (2019).

[25]M. Raissi, P. Perdikaris, and G. Karniadakis, Journal of Computational Physics **378**, 686 (2019).

[26]K. Carlberg, S. Guzzetti, M. Khalil, and K. Sargsyan, arXiv e-print (2019).

[27]J. Gaffney, D. Clark, V. Sonnad, and S. Libby, High Energy Density Physics **9**, 457 (2013).