# Validating IFE Concepts with Machine Learning Driven Design Optimization

V. Gopalaswamy, R. Betti, J. Knauer, C. Thomas, P. Heuer, and E. M. Campbell
*Laboratory for Laser Energetics, Rochester, USA, 14623*

D. A. Mariscal
*Lawrence Livermore National Laboratory, Livermore, USA 94550*

(Dated: January 28, 2022)

## Executive Summary

Optimization and validation of inertial fusion energy (IFE) concepts will be necessary for use of these concepts in commercial IFE. Optimization schemes have primarily been considered in the context of low repetition rates (LRR) and long experimental lead times. On an HRR facility, these considerations are altered, and open up avenues for applying advanced machine learning (ML) optimization schemes. This paper outlines some of the possibilities for rapid design optimization in an HRR facility, as well as some operational considerations for best utilization of such schemes. While they have the most utility at an HRR facility, these approaches can be developed and tested at existing LRR facilities, and could inform the optimal configuration of future HRR facilities.

### I. INTRODUCTION

Given the recent successes in indirect and direct drive ICF, the question of inertial fusion energy is receiving renewed interest. Many candidate designs have been suggested for use for IFE, whether informed by analytic theory, simulations, experiments, or some combination thereof. Some of these designs resemble those currently performed on the National Ignition Facility[1] (NIF) and OMEGA Laser Facility[2], while others are radical departures from conventional designs.

Lessons learned from experiments on the NIF and OMEGA suggest that the designs that eventually prove optimal are unlikely to be the initially suggested design, whether due to unexpected physics, unforeseen engineering issues, or both. Before a concept could be used for commercial IFE, it would therefore need to be validated, whether on existing LRR or future HRR research facilities (see white paper by P. Heuer et al.).

On LRR facilities, designs are typically only altered incrementally, and thus progress towards an optimal design occurs over a multi-year timescale. The slow rate of progress is due to the low repetition rate, but in its absence will also be driven by the long lead times necessary to manufacture acceptable targets and qualify pulse shapes,assimilate results from recently executed experiments,and iterate on existing designs.

The low repetition rate of the facility also results in data starvation, which further exacerbates the difficulty of predicting the result of non-incremental changes to the design and forces a single-shot paradigm that increases sensitivity to rare events[3]. The relative lack of flexibility on LRR facilities also results in low adaptivity on experimental campaigns, which are typically frozen well in advance of the experiments, and thus have difficult adapting to unexpected outcomes during the course of a campaign. A combination of low predictive capability, incremental progress, low adaptivity and data starvation not only delays the optimization of existing schemes, but also prevents the trial and validation of novel concepts.

An HRR IFE research facility could avoid most of these issues, and enable the use of advanced ML techniques for design optimization. What follows will discuss how machine learning (ML) can be used to design ICF implosions in an HRR environment, as well as the changes necessary in how implosions are fielded to enable such a scheme. Though these techniques will have the greatest utility at an HRR facility, we propose that they be developed at existing facilities to motivate and develop ML-augmented ICF design for a facility that would require it in the future.

## II. ACCELERATING ICF DESIGN WITH ML

The experimental loop for both LRR and HRR facilities contain the following steps:

1. Using prior data or knowledge, propose a candidate design,
2. Acquire appropriate targets, and qualify pulse shapes on the facility,
3. Execute experiment,
4. Assimilate experimental data and augment the existing dataset

A laser-driven ICF design is specified by the properties of the target (e.g. layer thicknesses and material compositions) and the properties of the laser (e.g. wavelength, beam incidence angles, beam pulse shapes). The resulting parameter space is vast and high ($N >> 10$) dimensional. Historically, using purely empirical methods to traverse such a space has not been considered tractable, and physics-based models have been used instead.

The traditional approach to ICF design optimization has been to utilize the highest fidelity radiation-hydrodynamic (RH) simulations[4–10] as the basis for performance expectations. The process of using the RH simulations typically involves providing a subspace of the full parameter space over which simulations are run, and subsequent identification of high performing regions. Such a labor and compute intensive approach is well suited to an LRR facility, but is not easily adapted to an HRR facility.

One alternative scheme would be to lower the fidelity of RH codes to fit the HRR shot cycle, followed by augmentation by ML models trained on historical data. The updated model can be applied to an existing large set of of RH simulations to draw candidates. This is the approach followed on OMEGA, and has proven effective in rapidly increasing performance[11, 12]. It is also relatively well suited for an HRR facility.

One issue with such a scheme is that it is limited to proposing designs from the parametrization of the candidate set. This can be made concrete by considering laser pulse shapes, which are broken up into components that have a specific function in the design, and are then parametrized. This limits the designs to a subset of all realizable pulse shapes. Though the existing parametrization of pulse shapes was physically motivated, it remains plausible that unknown physics or engineering effects could cause atypical pulse shapes to be high performance designs. Ideally, a method to non-parametrically represent and generate initial conditions is desired. Various ML techniques[13, 14] that may be well-suited for ICF exist, and should be investigated.

Another issue is the loss of physics fidelity, which becomes relevant for large extrapolations. To address this, one could consider replacing the RH codes in the experimental loop with a surrogate model[15–17] of the RH codes. Surrogate models require an enormous up-front cost in generating training data, but are subsequently trivial to evaluate. This decouples the RH code evaluation time from the fidelity level of the RH code itself, and could allow for extremely high fidelity RH codes to be used in the experimental loop of the scheme presented above. A subset of conducted experiments could be then simulated, and used to re-calibrated the surrogate model asynchronously from shot operations in order to maintain the quality of the surrogate in regions of interest.

Finally, with sufficiently high repetition rate it may be possible to *directly* predict the performance metrics from the initial conditions without resorting to intermediate models by constructing non-parametric 'black-box' ML models[18–20] trained solely on experimental inputs and outputs. The data-starved nature of LRR facilities limit the capability of such endeavours - However, an HHR facility could produce more data in two days than the last decade of OMEGA cryogenic implosions, and could potentially collect sufficient data to enable such approaches within a few months of operation. Though the extrapolative capability of a direct predictive model may be suspect, it benefits equally by being entirely unbiased by physics expectations that may not be valid in the regimes relevant for IFE, and could prove a valuable companion to the physics-augmented approach. We propose investigations into the volume of data necessary to use such a 'black-box' approach to confidently validate designs, which may inform the requirements for future facilities.

Once a predictive model is assembled, it can be used to search the parameter space for candidate designs[21, 22] in either a constrained or unconstrained manner. An example could be the investigation of a new ablator material with fewer defects, which may open up new optima near the stability boundary. Unlike in an LRR scheme, the search process would need to be relatively unsupervised to keep pace with the shot cycle. Metaheuristic search paradigms have been considered for ICF[23, 24], but have never been implemented due to the constraints of existing facilities. Multiple different approaches to design could also be combined via ensemble learning techniques[25]. For instance, multiple independent groups of researchers with their own predictive models could have blocks of investigation time allocated each year, followed by a joint investigation block where all the predictive models were ensemble averaged and compared. Testing these search paradigms, both in computer experiments as well as real implosion experiments on existing facilities would help prepare for rapid experimental design validation at an HRR facility.

## III. FACILITY CONSIDERATIONS

While the optimization strategies laid out in the previous section are exciting, some practical concerns must be addressed before such strategies could be implemented on an HRR facility while maintaining the rapid shot rate.

First, consider data assimilation. In the low-repetition rate environment of ICF, analysis of high value diagnostics is usually handled semi-manually, resulting in long lead times for key performance metrics. Many diagnostics are analog rather than digital, and it can take many hours or days for unprocessed data to be available. This has been historically justified given the low repetition rate, lack of adaptivity during an experimental campaign, and high value of each piece of diagnostic data in a data-starved environment. However, on an HRR facility that may have hundreds of diagnostics and $\sim$ minutes (or less) to complete data analysis, and where adaptive adjustment of designs are possible, it is imperative that digital replacements be developed for all high value analog diagnostics. It is also vital that the analysis of diagnostic data occur automatically and reliably, which may require the use of ML techniques[26], especially where image processing and tomography are concerned (see also white paper by D. A. Mariscal et al.)

Next consider the process of acquiring a target. In principle, optimization schemes can search over the entire space of target specifications, but the requirement to fabricate and fill shells on-demand at sub-minute timescales is likely impractical. For shells, it may be possible to maintain a large stock of shell configurations that can sampled on demand. Wetted foam targets may also be filled in time, but solid layers likely cannot. Consequently, it may not be feasible to specify targets with a lead time under a few days, a constrained strategy may become necessary. Nevertheless, it is vital to minimize the fabrication time of a target to the extent possible to maximize the value of ML optimization schemes. ML techniques have proven useful in demand forecasting[27], and should be investigated in the context of maintaining sufficient inventory for adaptive design schemes.

Finally, consider the process of designing a pulse shape. ICF facilities typically require long lead times when qualifying pulse shapes in order to minimize damage to sensitive laser components. While such an approach to pulse shape design is appropriate for LRR facilities, it would severely limit the value of ML optimization schemes. For an HRR facility to be capable of using adaptive ML techniques, the ability to change pulse shapes significantly on-demand must be developed. Some safety concerns could be addressed via initial facility design (e.g. avoiding opposing beam ports), or could be assessed automatically upon review of the design (e.g. excessively high peak power), but other issues such as SBS backscatter are complex and design dependent. To mitigate the risk of such issues, the ML optimzation process could occur initially at low energies where they are expected to be below safety thresholds. As promising candidates are identified, they could then be stepped up gradually in energy while monitoring risk factors, with high risk candidates being removed from consideration as the energy increases. ML could also be used to accelerate pulse shaping on HRR facilities[28], especially in cases where ML optimization schemes request large variations in designs from shot to shot that could be challenging for existing pulse shaping technologies to keep pace with.


[1] E. M. Campbell and W. J. Hogan, The national ignition facility - applications for inertial fusion energy and high-energy-density science, Plasma Phys. Contr. Fusion **41**, 10.1088/0741-3335/41/12B/303 (1999).

[2] T. R. Boehly, Initial performance results of the omega laser system, Opt. Commun. **133**, 10.1016/S0030-4018(96)00325-2 (1997).

[3] J. Morio and M. Balesdent, *Estimation of rare event probabilities in complex aerospace and other systems: a practical approach* (Woodhead Publishing, 2015).

[4] M. M. Marinak, Three-dimensional hydra simulations of national ignition facility targets, Phys. Plasmas **8**, 10.1063/1.1356740 (2001).

[5] G. Zimmerman, D. Kershaw, D. Bailey, and J. Harte, The lasnex code for inertial confinement fusion, J. Opt. Soc. Am. **68** (1978).

[6] P. B. Radha, Two-dimensional simulations of plastic-shell, direct-drive implosions on omega, Phys. Plasmas **12**, 10.1063/1.1857530 (2005).

[7] J. Delettrez, R. Epstein, M. C. Richardson, P. A. Jaanimagi, and B. L. Henke, Effect of laser illumination nonuniformity on the analysis of time-resolved x-ray measurements in uv spherical transport experiments, Phys. Rev. A **36**, 10.1103/PhysRevA.36.3926 (1987).

[8] S. Atzeni, Fluid and kinetic simulation of inertial confinement fusion plasmas, Comput. Phys. Commun. **169**, 10.1016/j.cpc.2005.03.036 (2005).

[9] D. S. Clark, Radiation hydrodynamics modeling of the highest compression inertial confinement fusion ignition experiment from the national ignition campaign, Phys. Plasmas **22**, 10.1063/1.4906897 (2015).
[10] I. V. Igumenshchev, Three-dimensional modeling of direct-drive cryogenic implosions on omega, Phys. Plasmas **23**, 10.1063/1.4948418 (2016).
[11] V. Gopalaswamy, R. Betti, J. Knauer, N. Luciani, D. Patel, K. Woo, A. Bose, I. Igumenshchev, E. Campbell, K. Anderson, *et al.*, Tripled yield in direct-drive laser fusion through statistical modelling, Nature **565**, 581 (2019).
[12] A. Lees, R. Betti, J. Knauer, V. Gopalaswamy, D. Patel, K. Woo, K. Anderson, E. Campbell, D. Cao, J. Carroll-Nellenback, *et al.*, Experimentally inferred fusion yield dependencies of omega inertial confinement fusion implosions, Physical Review Letters **127**, 105001 (2021).
[13] C. Doersch, Tutorial on variational autoencoders, arXiv preprint arXiv:1606.05908 (2016).
[14] A. Creswell, T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta, and A. A. Bharath, Generative adversarial networks: An overview, IEEE Signal Processing Magazine **35**, 53 (2018).
[15] J. A. Gaffney, S. T. Brandon, K. D. Humbird, M. K. Kruse, R. C. Nora, J. L. Peterson, and B. K. Spears, Making inertial confinement fusion models more predictive, Physics of Plasmas **26**, 082704 (2019).
[16] K. D. Humbird, J. L. Peterson, B. Spears, and R. McClarren, Transfer learning to model inertial confinement fusion experiments, IEEE Transactions on Plasma Science **48**, 61 (2019).
[17] K. D. Humbird, J. L. Peterson, J. Salmonson, and B. K. Spears, Cognitive simulation models for inertial confinement fusion: Combining simulation and experimental data, Physics of Plasmas **28**, 042709 (2021).
[18] W. Samek, G. Montavon, S. Lapuschkin, C. J. Anders, and K.-R. Müller, Explaining deep neural networks and beyond: A review of methods and applications, Proceedings of the IEEE **109**, 247 (2021).
[19] E. Schulz, M. Speekenbrink, and A. Krause, A tutorial on gaussian process regression with a focus on exploration-exploitation scenarios, BioRxiv , 095190 (2017).
[20] A.-L. Boulesteix, S. Janitza, J. Kruppa, and I. R. König, Overview of random forest methodology and practical guidance with emphasis on computational biology and bioinformatics, Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery **2**, 493 (2012).
[21] L. M. Rios and N. V. Sahinidis, Derivative-free optimization: a review of algorithms and comparison of software implementations, Journal of Global Optimization **56**, 1247 (2013).
[22] W. H. Miner Jr, P. M. Valanju, S. P. Hirshman, A. Brooks, and N. Pomphrey, Use of a genetic algorithm for compact stellarator coil design, Nuclear fusion **41**, 1185 (2001).
[23] P. Hatfield, S. Rose, and R. Scott, The blind implosion-maker: Automated inertial confinement fusion experiment design, Physics of Plasmas **26**, 062706 (2019).
[24] T. Ma, D. Mariscal, R. Anirudh, T. Bremer, B. Djordjevic, T. Galvin, E. Grace, S. Herriot, S. Jacobs, B. Kailkhura, *et al.*, Accelerating the rate of discovery: toward high-repetition-rate hed science, Plasma Physics and Controlled Fusion **63**, 104003 (2021).
[25] O. Sagi and L. Rokach, Ensemble learning: A survey, Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery **8**, e1249 (2018).
[26] R. Simpson, D. Mariscal, G. Williams, G. Scott, E. Grace, and T. Ma, Development of a deep learning based automated data analysis for step-filter x-ray spectrometers in support of high-repetition rate short-pulse laser-driven acceleration experiments, Review of Scientific Instruments **92**, 075101 (2021).
[27] R. Carbonneau, K. Laframboise, and R. Vahidov, Application of machine learning techniques for supply chain demand forecasting, European Journal of Operational Research **184**, 1140 (2008).
[28] G. Genty, L. Salmela, J. M. Dudley, D. Brunner, A. Kokhanovskiy, S. Kobtsev, and S. K. Turitsyn, Machine learning and applications in ultrafast photonics, Nature Photonics **15**, 91 (2021).