APPLICATION OF MACHINE LEARNING TO HIGH-REPETITION-RATE LASER-PLASMA PHYSICS ON THE PATH TO INERTIAL FUSION ENERGY

Blagoje Đorđević*,
Peer-Timo Bremer
Gerald Jackson Williams,
Tammy Ma,
Derek A. Mariscal

Lawrence Livermore National Laboratory,
Livermore, United States

Abstract:

One of the grand challenges of the plasma physics community is mastering controlled nuclear fusion as an energy source, with one approach being inertial confinement fusion (ICF). ICF is an extremely complex scientific and engineering problem that spans many physical regimes and requires precise control of the system over many orders of magnitude in space and time. Recent scientific achievements have raised our confidence in the feasibility of this goal, but much work remains to make inertial fusion energy a reality. An important research thrust has been the implementation of machine learning on ICF and specifically on the high-repetition-rate laser systems needed to make fusion energy practical. With an eye to technology transfer, there has been work attempting to operate, understand, and control of HRRLs on smaller laser-plasma experiments and associated modeling efforts. Presented here will be a series of examples of how machine learning is applied to these topics at LLNL.

Keywords:

Inertial Confinement Fusion, Machine Learning, Plasma Physics, Lasers, Nuclear Energy.

INTRODUCTION

Nuclear fusion is the process that powers the stars and harnessing it for inertial confinement energy (IFE) has been a goal of the physics community for over 60 years [1]-[4]. Until recently, the goal of energy breakeven remained elusive for the field of inertial confinement fusion (ICF) and high-energy-density (HED) plasma physics [5], [6]. However, on December 4th, 2022, the National Ignition Facility (NIF) at Lawrence Livermore National Laboratory (LLNL) achieved breakeven, where 2.05 megajoules of laser energy was delivered onto the target and 3.15 megajoules of fusion energy output was generated. This breakthrough entailed a scientific gain of Q_S=1.5 and has heralded a renewal of interest in fusion energy. Here the lasers cause a capsule of deuterium and tritium (DT) to compress into a ball of burning plasma, converting mass into energy, visualized in Figure 1. However, there is much work to be done before this can become an energy source. A critical hurdle is that NIF typically is only able to fire once a day, while to generate enough energy to be useful for electricity we need ~10 Hz operations. So, to achieve a viable energy source, we need to develop and optimize high-repetition-rate laser (HRRL) systems.

Correspondence:
Blagoje Đorđević

e-mail:
bzdjordje@gmail.com

DOI: 10.15308/Sinteza-2023-2-8

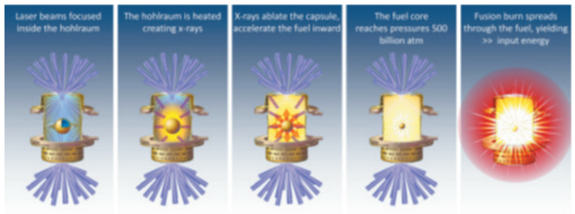

Figure 1 - The process of achieving nuclear fusion at NIF [4], [7], produced by J. Lett.

Managing NIF systems safely and precisely requires great expenditure and effort. This becomes exponentially more challenging when we wish to not only elevate the gain of our systems 100 times but also to increase the rate of operations 1,000,000-fold. To address this challenge, resources have been devoted to applying machine learning (ML) to fusion research. ML has experienced a renaissance during the past decade with the development of deep neural networks (NNs) and the advent of relatively cheap GPU technology. There has been great interest in applying ML to ICF in recent years, such as developing surrogate models of expensive simulations and experiments, identifying optical defects, and suggesting new configurations to optimize performance.

Since its inception, the ICF community has been deeply involved in advanced computational methods and technologies and LLNL and the general HED community have been taking advantage of advances in ML. Artificial Intelligence and ML are being pursued to help design targets, interpret diagnostics, analyze experiments, and improve simulation predictions. ML can accelerate the rate of learning in tandem with HRRL systems in ways not previously possible [8]. Surrogate models of sub-physics models and entire simulations allow for rapid investigation of wide parameter spaces for design. ML may also offer a new way to bridge physical regimes, as ICF covers many orders of magnitude. A common approach to ML at LLNL is to generate large ensembles of simulation results leveraging the supercomputing capabilities available [9], [10], [11], as NNs are data-intensive.

2. LASER TECHNOLOGY

Current ICF laser systems operate once a day and are constrained not only by drive sourcing but also by safety concerns. In one study a convolutional neural network (CNN) was trained on the input and output spectra of a 1-D physical model of the laser amplification process and classified the output spectra into safe/unsafe categories [12]. The CNN model achieved 98% accuracy on this binary classification problem. In a similar context, a Bayesian CNN model was used [13], which follows the architecture presented in Ref. [14], whereby a dropout layer is added before each layer and allows for probabilistic inference. An active learning training procedure was used, which incrementally selects the most informative data instances to be included in training. This allows for quicker convergence than training a regular CNN model. A common problem for HRRL systems is drift, where the laser pointing and performance change over time. To account for this the CNN model was exposed to multiple datasets in a sequential manner. One of the most desirable traits of NN-based surrogate models is their relative speed of deployment after they have been trained.

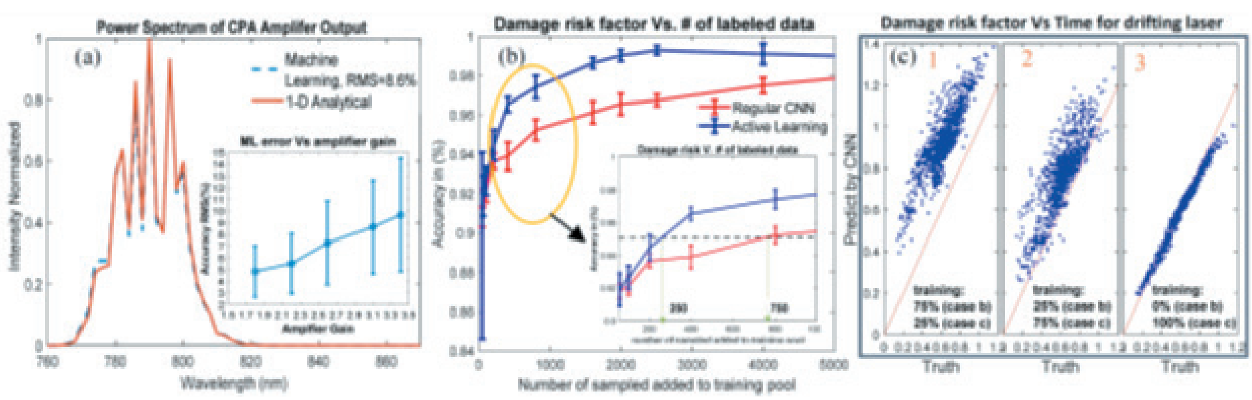

Figure 2 - (left) Sample of output power spectrum of an amplifier predicted by machine learning (blue) and 1D physical model (red). (center) Accuracy of damage risk as a function of the number of training data for regular (red) and active learning (blue). (right) Damage risk factor predicted by the model when trained during system drift. Reproduced from [12], with permission of Optica Publishing.

It is expected that in an IFE system, where diagnostic analysis and software would be encoded into surrogates generated via mL that allow for rapid analysis and control of the IFE operation loop. We are currently prototyping such techniques on smaller HRRL systems as described here.

Before every shot, NIF performs extensive maintenance of the optical systems that generate and deliver the lasers to the target capsule. High-fluence lasers slowly chip away at the lenses and mirrors in the laser chain. Damage sites are removed but sometimes microscopic sub-surface cracks remain that can easily grow when exposed to the high thermal stresses of the laser. In Ref. [15], a network was developed that was able to identify optical defects more rapidly and accurately than a human expert. In this case they took pre-trained, publicly available networks such as AlexNet [16] and ResNet [17] with more than 20,000 categories [18]. These networks were not trained on optical defects, but the intuition built into them allowed them to easily learn the features present in the NIF optics dataset of only 2,813 truth-labeled images.

3. INERTIAL CONFINEMENT FUSION

An early example of ML and ICF involved approximately 4,000 2D HYDRA simulations of the capsule implosion selected with Latin hyper-cube sampling [19]. This study focused on the consequences of drive asymmetry, which is a perennial problem in ICF. Drive asymmetries lead to less efficient compression and may cause ignition to fail, due to hydrodynamic instabilities which pierce and cool the central hotspot. Using this data, a Gaussian Process (GP) surrogate model was trained on capsule properties, total neutron yield, and yield-over-clean, a ratio comparing a perturbed capsule implosion to a perfectly symmetrical one. Training studies found that 2,500 simulations or more resulted in a $R^2>0.95$. GP models are advantageous because they can give us statistically meaningful uncertainties, but do not scale as well with large datasets and dimensionality as other techniques.

Another study used a supervised ML algorithm trained on petabytes of ICF simulation data, with 60,000 2D capsule simulations, to identify a class of ICF implosions that are more robust to perturbations [20]. As opposed to a GP model, this study used a random forest model, which can handle large quantities of data and readily incorporate nonlinearities, which is important as we are looking for cliffs in the parameter space. To rapidly find the optimum in the simulation parameter space the Nelder-Mead simplex-based optimization routine was used [21]. The optimization algorithm, using rapid surrogate calls in lieu of direct simulations, identified an ovoid shape for the capsule to be more robust to perturbations of the total drive fluence. The discovered ovoid shape and the output yield of a baseline, spherical capsule versus an ovoid are shown in Figure 3. While the baseline, spherical capsule reaches high yields for a wider range of drive fluences, the ovoid is more robust to P1 mode asymmetry perturbations. It was concluded that coherent flows in the hot spot can stabilize shell deformations that may arise during stagnation. Such surrogate modeling work has continued with the application of NNs, e.g., transfer learning between different types of fidelity of ICF data [22] and uncertainty quantification [23].

Also, of great interest in ICF is how data can be condensed and represented in lower dimensional spaces. In Ref. [24] was introduced the concept of a manifold and cyclically consistent (MaCC) surrogate that uses a multi-modal and self-consistent NN that outperforms many other state-of-the-art models. The autoencoder structure approximates multimodal data and finds the optimal representation of it in the lowest-dimensional layer of the autoencoder, which we call the latent space.

Figure 3 - (left) Velocity (arrows), density (gray), ion temperature (left color contours), and pressure (right color contours) at the time of peak energy production for an ovoid implosion. (right) The surrogate's estimate of yield under changing total drive fluence for the baseline (round) and optimal (ovoid) cases. Reproduced from [20], with the permission of AIP publishing.

An inverse network trains in parallel with the surrogate model and cyclical consistency between the two acts as a regularization factor in an unsupervised fashion. The architecture was deployed on ICF datasets, specifically, scalars of note such as the yield and images of artificial diagnostic images of the neutron and X-ray emissions of the capsule implosion. ICF is a challenging subject to model and learn in ML contexts, data is sparsely sampled and highly nonlinear, and the development of advanced techniques such as MaCC will help the development of high-performing surrogate models.

4. LASER-PLASMA PHYSICS

Much effort and thought are dedicated directly to ICF by laser-plasma physicists, but there are several related research topics that synergize with the scientific and technological needs of ICF. A primary example is short-pulse, laser-plasma physics, where physical concepts and engineering constraints are shared. This has been the primary scientific focus of high-power, HRRL systems at LLNL [8]. Laser-plasma experiments, particularly laser-solid interactions, are inherently challenging. The advent of HRRL systems means that experiments often outrun our ability to not only analyze results but also to control and guide them. ML has been deemed as a necessary tool which allows our goals to align with the technical capabilities at hand.

Analogous to work in ICF, work was done applying ML to ensembles of particle-in-cell (PIC) simulations modeling laser-solid interactions. Initial efforts involved generating 1,400 1D PIC simulations which were used as the training set for NNs [25]. A fully connected NN was trained on scalar quantities of interest from the PIC-generated dataset, particularly the particle distributions as seen in Figure 4. In addition to hot electrons, a useful quantity of interest is the maximum ion energy, which following a self-similar model takes the form of

$$E_{max} = 2T_h \log\left[t_p + \left(t_p^2 + 1\right)^{1/2}\right]^2,$$

Equation 1 – Maximum ion energy from self-similar plasma expansion into a vacuum [26].

where T_h is the hot electron temperature, and t_p is the normalized time with respect to the plasma [26].

$$T_h \gg m_e c^2 \left(\sqrt{1 + 7.3\times10^{-19}\left(\lambda[\mu\mathrm{m}]\right)^2 I_0\left[\mathrm{W/cm^2}\right]} - 1\right)$$

Equation 2 – Hot electron temperature from ponderomotive scaling of laser-plasma interaction [27].

is quasi-empirical but roughly shows the dependence of hot electrons on intensity and wavelength [27].

The NN can accurately reproduce the scalar quantities of interest and uncertainties were derived by taking a weighted average of a bundle of NNs, where the weighting was proportional to the average inverse loss of the NN at the end of its training [28]. An important quantity in laser-plasma physics is the preplasma, which is the exponential foot of the plasma generated by the low-intensity prepulse of the laser, depicted in Figure 4(left). Given the fast speed of a forward NN, we can use it for inverse modeling. Looking at a small collection of experimental data, we used our NN surrogate model, guided by a genetic algorithm, to estimate what the preplasma might be in those experiments. ML-trained surrogate models as such have been deployed in experiments during operations, allowing us to get live estimates of the experimental results using simulation and prior experimental data. To alleviate data sparsity, we focused on leveraging the transfer learning technique [29] to make our small, sparse datasets more robust. A preliminary effort used tens of thousands of analytical results to pretrain a composite NN architecture [30]. The NN was then retrained on the previous PIC dataset, resulting in greater performance.

Given our initial experience with ensemble simulations and ML, we generated a new dataset of PIC simulations that more closely approximated our experiments at the CSU ALEPH facility, consisting of 8,000 1D PIC simulations and approximately 500 2D simulations [31]. This work came at great computational cost, approximately three million process hours to generate the 1D dataset but two million for the 2D dataset. However, presupposing a hierarchical training framework [29], where we trained on different fidelities, we were able to keep costs within scope, and by transfer learning we were able to train surrogate models on 2D data with relatively high confidence. As an application of the higher-fidelity NN we did a parameter scan of laser properties over constant energy surfaces, i.e., $E \propto I_0\ \tau r_0^2$, varying from 1 joule to 4 Joules, depicted in Figure 5. The peak ion energy shifts from the long-pulse, low-intensity corner of parameter space for 1 joule to the short-pulse, high-intensity corner at 4 joules. More interestingly, at 2 joules we see a plateau in the long-pulse corner of parameter space. For HRRL applications, such as IFE, this is desirable as it represents a configuration more robust to perturbations, as laser alignment and pointing are not trivial matters. Laser delivery and control typically operate within error-bars of several percent at NIF, where an entire day or more may be spent by a team recalibrating the optics. At HRRL scales we have found such precision to be impossible with current technology and techniques and are looking to ML to assist with that.

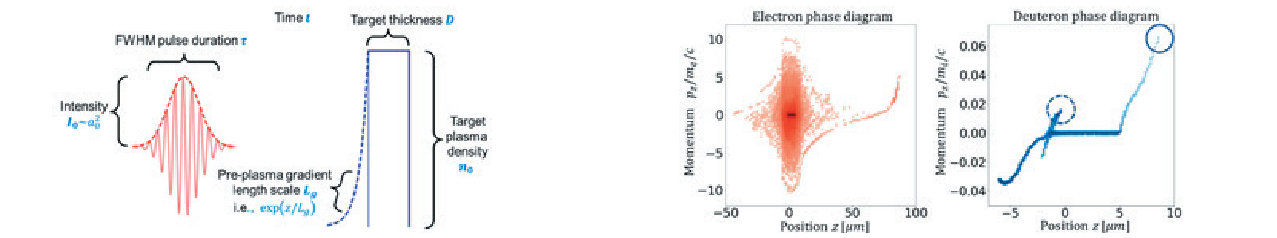

Figure 4 - (left) Setup of laser-solid interaction. (right) Phase space of electrons and ions after laser arrival. Reproduced from [25], with the permission of AIP publishing.

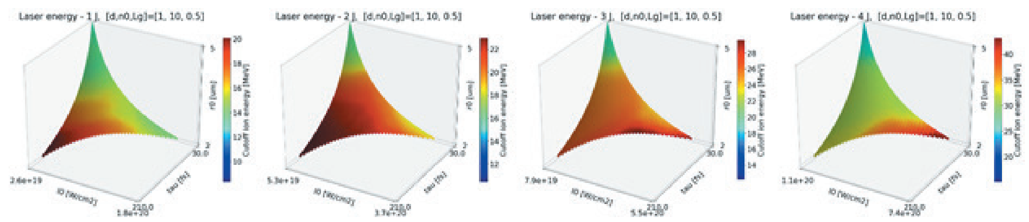

Figure 5 - Mapping of ion energy utilizing the transfer-learning-based 2D surrogate model. Each heatmap represents a constant energy surface within the parameter space. Reproduced from [29], with the permission of AIP publishing.

Of particular interest are the particle energy spectra, which are one of the primary observables in HRR experiments. A 1D CNN autoencoder was trained on energy spectra from the 1D dataset and then utilized in several different ways. First, the encoder was coupled to a fully connected NN (FCNN) to predict scalar quantities, such as the hot electron temperature from the energy spectrum. Second, the encoder was frozen and the autoencoder was retrained on the smaller 2D dataset, allowing us to effectively convert 1D predictions to 2D. Lastly, the decoder was spliced to a FCNN for inputs, but with two branches, where the whole network was trained on the 1D dataset and then again on the 2D dataset with 1D input quantities frozen. These are shown in Figure 6.

Similar work has been applied on the experimental side, again focusing on particle spectra. In one effort, a CNN was trained first on simulation spectra and correlated to a labelled hot electron temperature [32]. If the CNN was trained directly on the experimental the predictions showed a pronounced bias, plotted as open circles in Figure 7(a). However, if transfer learning was used first, then the NN was able to perform much better (closed circles). Synthetic data generation and subsequent training were also applied to ion diagnostics of greater complexity, where a 2D image simulating an ion beam passing through layered radiochromic film stacks allows us to infer the temporal and spatial distribution of the beam [33].

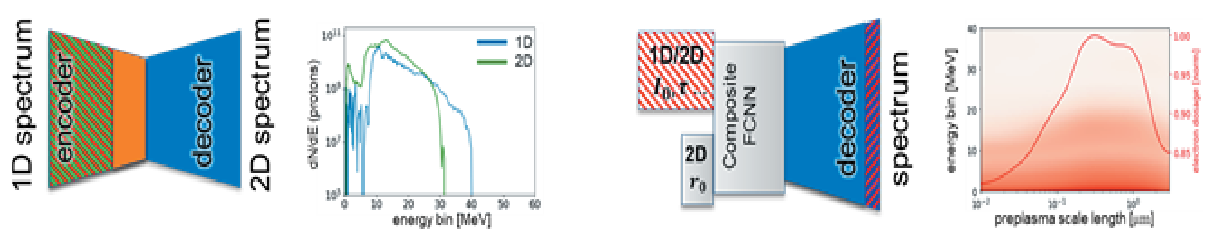

Figure 6 - Demonstration of how autoencoders can be applied to spectral data from PIC simulations. Reproduced from [29], with the permission of AIP publishing.

Figure 7 - (left) Calibration plot of the predicted electron temperatures using transfer learning. (right) CNN-based architecture used to correlate diagnostic images with scalar parameters of interest. Figures reproduced from [32] and [33], with permission of AIP publishing.

This image is passed through a CNN for data reduction and then given to a FCNN to predict the associated scalar values, as seen in Figure 8(b). Similar work using PCA for data reduction coupled to NNs has been applied to X-ray spectra relevant to ICF and laser-plasma physics [34]. Elevating the idealized results from simulations to our experimental observations is a primary goal in our community, although this is balanced by the significant uncertainties of experimental observables, which will need to be addressed if modeling, diagnostics, and control of IFE laser-plasma systems are to succeed.

5. CONCLUSION

IFE is a challenging goal for the physics community that likely remains many years away. However, given recent successes in ICF, the community has been inspired to redouble its efforts and ML promises to help alleviate many issues. In this manuscript, we have reviewed how ML has been applied to IFE-relevant fields at LLNL, specifically laser technology, ICF, and basic laser-plasma physics. In fact, this is only a small sampling of how ML is being applied to scientific and engineering research at LLNL. With respect to the topics discussed, several upcoming projects are being pursued: adaptive, time-dependent laser system controls to address persistent issues in consistency; application of external structures to ICF capsules and numerical optimizing them to enhance implosion yield; and advanced networks architectures for synthesizing non-congruent datasets in the context of shaped, short-pulse laser-solid interactions. Much work is still needed to be done but ML has demonstrated its ability to accelerate the way we do science and will bring the reality of controlled nuclear fusion energy closer.

6. ACKNOWLEDGEMENTS

This work was funded under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract No. DE-AC5207NA27344 with funding support from Laboratory Directed Research and Development under tracking code 20-ERD-048, 22-ERD-022, 23-ERD-035, and DOE Office of Science SCW1772 and Early Career Research Program under SCW1651. Computational work was performed with Livermore Computing and used LLNL Grand Challenge allocations. This document was prepared as an account of work sponsored by an agency of the United States government.

7. REFERENCES

[1] J. Nuckolls, L. Wood, A. Thiessen and G. Zimmerman, "Laser Compression of Matter to Super-High Densities: Thermonuclear (CTR) Applications," *Nature*, vol. 239, no. 5368, pp. 139-142, 1972.

[2] J. Lindl, "Development of the indirect-drive approach to inertial confinement fusion and the target physics basis for ignition and gain," *Phys. Plasmas*, vol. 2, pp. 3933-4024, 1995.

[3] M. D. Rosen, "The physics issues that determine inertial confinement fusion target gain and driver requirements: A tutorial," *Phys. Plasmas*, vol. 6, pp. 1690-1699, 1999.

[4] T. Ma and R. Betti, "Inertial Fusion Energy Report of the 2022 Fusion Energy Sciences Basic Research Needs Workshop," 2022.

[5] M. D. Rosen, "The science applications of the high-energy density plasmas created on the Nova laser," *Physics of Plasmas*, vol. 3, no. 5, pp. 1803-1812, 1996.

[6] R. W. Lee, R. Petrasso and R. W. Falcone, "Science on High-Energy Lasers:," LLNL, Livermore, 1995.

[7] G. H. Miller, E. I. Moses and C. R. Wuest, "The National Ignition Facility: enabling fusion ignition for the 21st century," *Nucl. Fusion*, vol. 44, p. S228, 2004.

[8] T. Ma, D. Mariscal and R. Anirudh, "Accelerating the rate of discovery: Towards high-repetition-rate HED science," *Plasma Phys. Control. Fusion*, vol. 63, no. 10, p. 104003, 2021.

[9] B. Spears, S. Glenzer and M. J. Edward, "Performance metrics for inertial confinement fusion implosions: Aspects of the technical framework for measuring progress in the national ignition campaign," *Physics of Plasmas,* vol. 19, p. 056316, 2012.

[10] A. L. Kritcher, R. Town and D. Bradley, "Metrics for long wavelength asymmetries in inertial confinement fusion implosions on the national ignition facility," *Physics of Plasmas*, vol. 21, p. 042708, 2014.

[11] D. S. Clark, S. W. Haan and J. D. Salmonson, "Robustness studies of ignition targets for the national ignition facility in two dimensions," *Physics of Plasmas*, vol. 15, p. 056305, 2008.

[12] S. Herriot, T. Galvin and B. Ng, "Deep Learning for Real-Time Modeling of High Repetition Rate, Short Pulse CPA Laser Amplifier," *Conference on Lasers and Electro-Optics*, 2019.

[13] T. Galvin, S. Herriot and B. Ng, "Active learning with deep Bayesian neural network for laser control," *Optics and Photonics for Information Processing XII,* 2018.

[14] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning," *Proceedings of the 33rd International Conference on Machine Learning*, vol. 47, pp. 1050-1059, 2016.

[15] T. N. Mundhenk, L. M. Kegelmeyer and S. K. Trummer, "Deep learning for evaluating difficult-to-detect incomplete repairs of high fluence laser optics at the National Ignition Facility," *SPIE Proceedings,* 2017.

[16] A. Krizhevsky, I. Sutskever and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Communications of the ACM*, vol. 60, no. 6, pp. 84-90, 2017.

[17] K. He, X. Zhang and S. Ren, "Deep Residual Learning for Image Recognition," *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[18] J. Deng, W. Dong and R. Socher, "ImageNet: A large-scale hierarchical image database," *2009 IEEE Conference on Computer Vision and Pattern Recognition,* 2009.

[19] R. Nora, J. L. Peterson and B. K. Spears, "Ensemble simulations of inertial confinement fusion implosions," *Statistical Analysis and Data Mining: The ASA Data Science Journal*, vol. 10, no. 4, pp. 230-237, 2017.

[20] J. L. Peterson, K. D. Humbird and J. E. Field, "Zonal flow generation in inertial confinement fusion implosions," *Physics of Plasmas*, vol. 24, no. 3, p. 032702, 2017.

[21] J. A. Nelder and R. Mead, "A Simplex Method for Function Minimization," *The Computer Journal*, vol. 7, no. 4, pp. 308-313, 1965.

[22] K. Humbird, L. Peterson and B. K. Spears, "Transfer Learning to Model Inertial Confinement Fusion Experiments," *IEEE Transactions on Plasma Science*, vol. 48, no. 1, pp. 61-70, 2020.

[23] J. A. Gaffney, S. T. Brandon and K. D. Humbid, "Making Inertial Confinement Fusion Models More Predictive," *Physics of Plasmas*, vol. 26, p. 082704, 2019.

[24] R. Anirudh, J. J. Thiagarajan, P.-T. Bremer and B. K. Spears, "Improved surrogates in inertial confinement fusion with manifold and cycle consistencies," *Proceedings of the National Academy of Sciences*, vol. 117, no. 18, pp. 9741-9746, 2020.

[25] B. Z. Djordjević, A. J. Kemp and J. Kim, "Modeling laser-driven ion acceleration with deep learning," *Physics of Plasmas*, vol. 4, no. 043105, p. 28, 2021.

[26] P. Mora, "Plasma Expansion into a Vacuum," *Phys. Rev. Lett.*, vol. 90, p. 185002, 2003.

[27] S. C. Wilks, A. B. Langdon and T. E. Cowan, "Energetic proton generation in ultra-intense laser–solid interactions," *Physics of Plasmas*, vol. 8, pp. 542-549, 2001.

[28] J. Thiagarajan, R. Anirudh and V. Narayanaswamy, "Single Model Uncertainty Estimation via Stochastic Data Centering," *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*, 2022.

[29] S. Bozinovski and A. Fulgosi, "The influence of pattern similarity and transfer learning upon the training of a base perceptron," *Proc. Symp. Informatica*, Vols. 3-121-5, 1976.

[30] B. Z. Djordjević, A. J. Kemp and J. Kim, "Characterizing the acceleration time of laser-driven ion acceleration with data-informed neural networks," *Plasma Physics and Controlled Fusion*, vol. 63, no. 9, p. 094005, 2021.

[31] B. Z. Djordjević, J. Kim and S. C. Wilks, "Transfer learning and multi-fidelity modeling of laser-driven particle acceleration," *Physics of Plasmas*, vol. 30, no. 4, p. 043111, 2023.

[32] K. Swanson, D. A. Mariscal and B. Z. Djordjevic, "Applications of machine learning to a compact magnetic spectrometer for high repetition rate, laser-driven particle acceleration," *Review of Scientific Instruments*, vol. 93, no. 10, p. 103547, 2022.

[33] D. A. Mariscal, B. Z. Djordjevć and R. Anirudh, "A flexible proton beam imaging energy spectrometer (PROBIES) for high repetition rate or single-shot high energy density (HED) experiments," *Review of Scientific Instruments,* vol. 94, no. 2, p. 023507, 2023.

[34] D. A. Mariscal, C. M. Krauland and G. G. Djordjević, "Enhanced analysis of experimental x-ray spectra through deep learning," *Physics of Plasmas*, vol. 29, no. 9, p. 093901, 2022.

[35] IEEE, *IEEE Editorial Style Manual for Authors, IEEE*, 2018.

[36] "Detection of Implicit Citations for Sentiment Detection," in *Proceedings of the Workshop on Detecting Structure in Scholarly Discourse*, Jeju Island, Korea, 2012.