

---

# Can KANs (re)discover predictive models for Direct-Drive Laser Fusion?

---

**Rahman Ejaz\***  
University of Rochester  
reja@lle.rochester.edu

**Varchas Gopalswamy**  
University of Rochester  
vgop@lle.rochester.edu

**Riccardo Betti**  
University of Rochester  
betti@lle.rochester.edu

**Aarne Lees**  
University of Rochester  
a.lee.lle.rochester.edu

**Christopher Kanan**  
University of Rochester  
ckanan@cs.rochester.edu

## Abstract

The domain of laser fusion presents a unique and challenging predictive modeling application landscape for machine learning methods due to high problem complexity and limited training data. Data-driven approaches utilizing prescribed functional forms, inductive biases and physics-informed learning (PIL) schemes have been successful in the past for achieving desired generalization ability and model interpretation that aligns with physics expectations. In complex multi-physics application domains, however, it is not always obvious how architectural biases or discriminative penalties can be formulated. In this work, focusing on nuclear fusion energy using high powered lasers, we present the use of Kolmogorov-Arnold Networks (KANs) as an alternative to PIL for developing a new type of data-driven predictive model which is able to achieve high prediction accuracy and physics interpretability. A KAN based model, a MLP with PIL, and a baseline MLP model are compared in generalization ability and interpretation with a domain expert-derived symbolic regression model. Through empirical studies in this high physics complexity domain, we show that KANs can potentially provide benefits when developing predictive models for data-starved physics applications.

## 1 Introduction

Inertial confinement fusion [1] (ICF) is an approach to nuclear fusion power and clean energy that uses high power lasers to achieve conditions similar to stellar interiors in the laboratory, and is governed by complicated and deeply nonlinear underlying physics. Sophisticated multi-physics solvers are traditionally used to understand the physics of ICF, but despite significant progress in high-fidelity integrated simulation capabilities for ICF [2, 3], simulations alone are not sufficient to predict outcomes of experiments *a priori* and their utility to design future experiments remains limited. This lack of predictive capability is a major impediment towards realizing nuclear fusion via inertial confinement as a viable clean energy source.

To bridge the gap between simulation predictions and experimental outcomes, data-driven approaches are utilized. Past approaches have relied on using physics intuition for prescribing functional forms (piece-wise power laws) for physically motivated models using quantities derived from simulations [4, 5, 6]. In order to relax the hypothesized functional form assumptions similar models to [4, 5, 6] were developed in [7], where multi-layer perceptrons (MLPs) were used due to their function approximation abilities indicated by the universal approximation theorem. However, [7] found that

---

\*Laboratory for Laser Energetics, University of Rochester

for their small-sized dataset the inclusion of applicable inductive biases was needed to improve prediction generalization and model interpretation to be in accordance with physics expectations.

Our work presented attempts to address physics-based regression problems that require minimal assumptions on the underlying functional terms to explain the data, but are not amenable to physics-informed or augmented solutions[8, 9, 10]. This can be either due to complexity in explicitly formulating desired model properties into discriminative penalties as is conventionally done in physics-informed learning (PIL) and inapplicability of model design methods such as physics-augmented learning [8]. In essence we showcase our work to an audience where limited data makes learning desired model properties solely from data alone challenging and it is not efficient or obvious on how to effectively construct a  $\mathcal{L}_{phys}$  term (defined in Eq. 1) which is augmented with the usual loss functions,

$$\mathcal{L}_{phys} = \sum_i^n \gamma_i \Psi_i(f(x_i; \theta)) \quad (1)$$

where  $n$  is the input feature dimension,  $\Psi$  is an operator which captures deviations from the desired property,  $f(\mathbf{x}; \theta)$  is the neural network and  $\gamma$  is the strength of the penalty. Constructing an effective  $\mathcal{L}_{phys}$  is a typical requirement for multiple methods that interleave physics with deep learning.

We show that in our particular use case of developing a new type of predictive model for direct-drive laser fusion the Kolmogorov-Arnold Networks (KANs) proposed by [11] has shown modest improvement for generalization on out-of-distribution (ood) data and comes closest in its interpretation with a domain expert derived model when compared with applicable PIL schemes. The PIL schemes are implemented through the  $\mathcal{L}_{phys}$  term and are constructed to include all the  $\Psi$  that can appropriately be considered prior knowledge. The ICF prediction model presented uses an input parameterization that makes it more general in directly relating the prediction quantity with experimental design features compared with the works of [4, 5], thus aiding in experimental design. We hypothesize that the learnable activation functions feature of KANs acts as a form of conformal/learnable inductive bias through learning of internal degrees of freedom (B-splines)[11]. The precise nature of the external and internal degrees of freedom and their role in making KANs successfully learn details of the physical system will be explored future works.

## 2 Dataset and Modeling Task

The empirical data used in this study come from cryogenic implosion experiments conducted on the OMEGA Laser Facility from 2016 to present that are relevant to fusion energy. The substantial difficulty in conducting these experiments limits the data to  $\sim 300$  samples, which precludes the use of many successful ML techniques. Using this database, the objective is to develop a model of the experimental fusion yield, which is the number of reactions and therefore is proportional to the energy released by nuclear fusion.

The fusion yield obtained in these experiments depends on the energy delivered to the fusion fuel target, details of energy delivery to the target via the laser power history[1], specifications of the fusion fuel payload[1], and 3-dimensional effects [5, 6]. Using a parameterization of the experimental fusion yield dependencies, the modeling task is to learn a mapping function which is given in general terms below,

$$Y^{\text{exp}} = f(E_L, R_{\text{out}}, \hat{M}, \hat{R}, \frac{R_b}{R_t}, \frac{\alpha}{\text{IFAR}}, \text{CR}, \hat{V}, \hat{Y}, T_{\text{min}}^{\text{max}}, \text{YOC}_{\text{He}}^{\text{sim}}) \quad (2)$$

where the inputs are described in more detail in Appendix A.1 and are physically motivated - however, the precise details of how the experimental yield depends on these inputs and their interactions is not yet well understood. We note that such fusion yield dependence parameterizations have been successful in the works of [4, 5, 6], though the parameterization in Eq 2 is more general than the previous mentioned works.

## 3 Methods

**Evaluation Criterion** The KAN based mapping model is compared in its prediction ability on ood experimental samples and in its agreement with the expert derived mapping function given in

Eq. 3. It should be noted that the domain expert derived model is not considered the "oracle" mapping relation as simplifying assumptions such as the absence of interactions between the input parameters is assumed in the formulation. The presence of complex interactions due to physics considerations is hypothesized [6, 7], however the nature and extent of the interactions is unknown. This motivates the use of automatic functional "finding" methods. However, when using these methods the inferred dependencies of  $Y^{\text{exp}}$  on input parameters (shown in Eq. 2) found need to be on average consistent with the physics of the model in Eq. 3. This is also used as an evaluation criterion and informs which model is preferable.

Due to design intent by physicists at OMEGA, the empirical database is curated into sets called "campaigns". However, it is not obvious how distinctions between campaigns correspond to distinctions in input parameter space. Therefore, for ood tests k-means clustering [12] is utilized to partition the dataset. The number of clusters is determined using a silhouette method [13] with the scan being restricted to the total number of unique campaigns in the dataset (29 campaigns). Given the clusters the dataset becomes  $\mathcal{D} = \{\mathcal{D}^1, \mathcal{D}^2, \dots, \mathcal{D}^N\}$  where  $N$  is the total number of clusters ( $N = 6$ ). A test dataset  $\mathcal{D}^{\text{test}} \in \mathcal{D}$  is chosen and correspondingly the train dataset becomes  $\mathcal{D}^{\text{train}} = \{\mathbf{X} | \mathbf{X} \in \mathcal{D} \wedge \mathbf{X} \notin \mathcal{D}^{\text{test}}\}$  where  $\mathbf{X}$  are the input-target pairs. The mean squared error (MSE),  $\frac{1}{N} \sum_{i=1}^N (\hat{y}_i - y_i)^2$ , where  $\hat{y}$  is the model prediction and  $y$  is the ground truth is compared across the models for all clusters used as  $\mathcal{D}^{\text{test}}$ . The prediction error  $\frac{\hat{y}_i - y_i}{y_i} \times 100$  ( $< > \pm \sigma$ ) on a new subset of experiments is also compared using these models as a means to probe prediction ability for future experiments.

### 3.1 Domain Expert Model

Intuition and physics knowledge is used to derive the domain expert model which is an evolution of the models in [4, 5, 6]. This model is represented as a piece-wise power law and is given below in Eq. 3.

$$Y^{\text{exp}} = E_L^{2.3} R_{\text{out}}^{-2.6} \hat{M}^{2.9} \hat{R}^{26.5} \left( \frac{R_b}{R_t} \right)^{1.6} \left( \frac{R_b}{R_t} \right)^{-3.0} \left( \frac{R_b}{R_t} \right)^0 \left( \frac{\alpha}{\text{IFAR}} \right)^{0.45} \left( \frac{\alpha}{\text{IFAR}} \right)^{-0.1} \hat{V}^{2.02} \hat{Y}^{0.78} T_{\text{min}}^{(-1.32)} (YOC_{\text{He}}^{\text{sim}})^{1.26}$$

$$\begin{matrix} R_b < 0.86 & 0.86 < R_b < 1 & R_b > 1 \\ R_t & & \end{matrix}$$
(3)

One major drawback of this parametrization strategy is the lack of interaction effects between the terms. Simple physics intuition, verified in some cases by multi-physics simulations[14] show that interaction effects can be important, but it is non-trivial to formulate them correctly in a closed-form parametrization. Capturing these effects is of high importance, not only to better predict the outcomes of future designs, but also because the details of interactions can point towards physical behaviors that are important in ICF experiments.

### 3.2 Physics-informed Model

For the MLP PIL model, a fully connected MLP,  $f_{\text{MLP}}(\mathbf{X}) = (\sigma \mathbf{W}_N \circ \sigma \mathbf{W}_{N-1} \circ \dots \circ \sigma \mathbf{W}_1) \mathbf{X}$  with a loss function  $\mathcal{L} = \mathcal{L}_{\text{MSE}} + \mathcal{L}_{\text{physics}}$  is used. The  $\mathcal{L}_{\text{physics}}$  is constructed such that it enforces properties that are definitively considered prior knowledge (explained in Appendix A.3).

$$\mathcal{L}_{\text{physics}} = \gamma_1 \left| \frac{\partial f(\mathbf{x}; \theta)}{\partial T_{\text{max}}^{\text{sim}}} \right| + \gamma_2 \left| \frac{\partial f(\mathbf{x}; \theta)}{\partial YOC_{\text{He}}^{\text{sim}}} \right|$$
(4)

$$\text{Where, } \frac{\partial f(\mathbf{x}; \theta)}{\partial T_{\text{max}}^{\text{sim}}} = \begin{cases} 0 & \text{if } \frac{\partial f(\mathbf{x}; \theta)}{\partial T_{\text{max}}^{\text{sim}}} \leq 0 \\ \frac{\partial f(\mathbf{x}; \theta)}{\partial T_{\text{max}}^{\text{sim}}} & \text{if } \frac{\partial f(\mathbf{x}; \theta)}{\partial T_{\text{max}}^{\text{sim}}} > 0 \end{cases} \text{ and } \frac{\partial f(\mathbf{x}; \theta)}{\partial YOC_{\text{He}}^{\text{sim}}} = \begin{cases} \frac{\partial f(\mathbf{x}; \theta)}{\partial YOC_{\text{He}}^{\text{sim}}} & \text{if } \frac{\partial f(\mathbf{x}; \theta)}{\partial YOC_{\text{He}}^{\text{sim}}} < 0 \\ 0 & \text{if } \frac{\partial f(\mathbf{x}; \theta)}{\partial YOC_{\text{He}}^{\text{sim}}} \geq 0 \end{cases}$$

For the baseline MLP model only the  $\mathcal{L}_{\text{MSE}}$  loss is used while the rest of the training configuration and model architecture is kept the same as the MLP with PIL model.

### 3.3 KAN Model

Differing from MLPs, KANs have no linear learnable weights  $\mathbf{f}_{\text{KAN}}(\mathbf{X}) = (\Phi_N \circ \Phi_{N-1} \circ \dots \circ \Phi_1) \mathbf{X}$ , each weight parameter in  $\Phi$  is replaced by a learnable univariate function constructed using B-spline basis functions [11]. The KAN based model is trained using solely the  $\mathcal{L}_{\text{MSE}}$  loss.

## 4 Results and Discussion

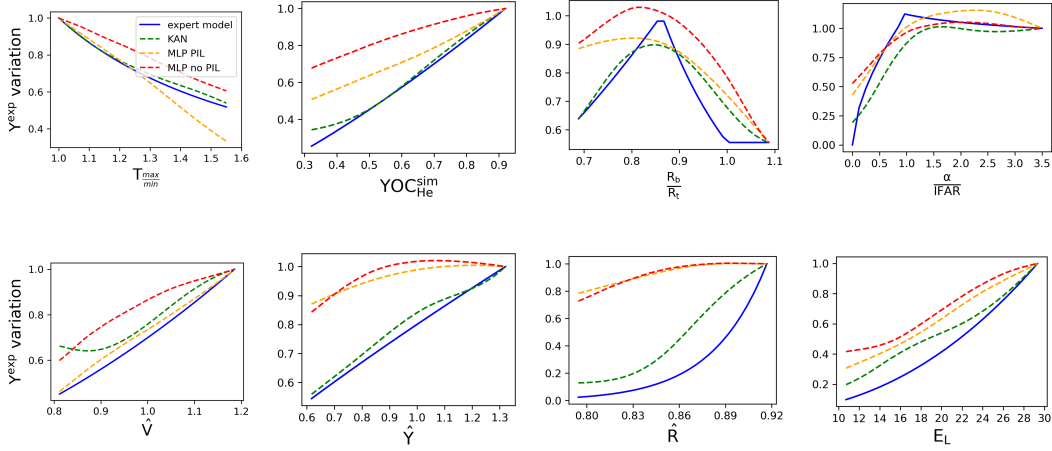


Figure 1: Inferred  $Y^{\text{exp}}$  variation across different models compared with the expert derived model, shown in arbitrary units (A.U)

**Main results** To interpret the learned behavior of the model, the partial dependence method [15] is utilized where  $f(x; \theta)$  is varied as a function of its input arguments. The average  $Y^{\text{exp}}$  variation on input parameters from the models is shown in Fig. 1 along with the domain expert derived model dependencies for direct comparison. All the models generally align with the expert derived model. However, the KAN model is in better agreement with the expert model as illustrated in Fig. 1. Important notable differences between the models are for the  $\frac{R_b}{R_t}$  and  $\frac{\alpha}{IFAR}$  terms. Due to multiple competing physics effects[6, 14] the optimum value of  $\frac{R_b}{R_t} \approx 0.85$ . While both models capture this nature, the strong dependence of  $Y^{\text{exp}}$  with  $\frac{R_b}{R_t}$  for  $\frac{R_b}{R_t} < 0.85$  is more accurately captured by the KAN model. The rapid fall off of for  $Y^{\text{exp}}$  below 0.85 is justified in the expert model based on detailed physics studies[14] and only the KAN model is able to correctly recover this behavior. We can make a similar case for the  $\frac{\alpha}{IFAR}$  term, though in this case the justification for the dependence inferred by the expert model is based largely on physical intuition. As  $\frac{\alpha}{IFAR}$  decreases,  $Y^{\text{exp}}$  is expected to undergo a transition from weak dependence to strong dependence, at which point  $Y^{\text{exp}}$  is expected to drop dramatically. KAN also recovers this behavior with more fidelity than the MLPs making it the preferable choice. Although, it can be argued that a suitable  $\mathcal{L}_{physics}$  to capture the above mentioned effects in the PIL model can be devised, these aspects of  $\frac{R_b}{R_t}$  and  $\frac{\alpha}{IFAR}$  have been uncovered through detailed and expensive high-fidelity physics simulations [14, 16] and dedicated experiments and as a result are included in the domain expert model. Thus, we postulate that information of this type isn't available *a priori* and shouldn't be included when investigating the merits of PIL schemes for this (re)discovery investigation. The functional approximation methods that can automatically uncover these properties from the data are preferable.

**Out-of-distribution tests** The KAN model on average ( $\langle \rangle \pm \sigma$ ) achieves an incremental advantage on MSE scores across the clusters ( $0.00394 \pm 0.00410$  vs next best  $0.00926 \pm 0.00506$ ). This translates to prediction errors of  $(0.6 \pm 5.9)\%$  using the KAN model,  $(1.2 \pm 9.5)\%$  using MLP PIL model and  $(9.6 \pm 4.2)\%$  using the MLP model when tasked to predict the outcomes of new experiments. The prediction error using the domain expert model is  $(12.8 \pm 10.4)\%$ .

**Discussion** The benefit of the KAN model in terms of generalization ability is not conclusive solely from these ood tests. The new experiments mentioned are not true ood samples as they have overlap with previous experiments. This is often the case in the domain of laser direct-drive fusion. Furthermore, the data comes from real-world experiments and thus has random variations which also complicates any benefit interpretation. To further elucidate any generalization benefit the use of variational autoencoders based segmentation[17] of the data would be investigated in future works. We note that for the  $T_{\text{min}}^{\text{max}}$  and  $YOC_{\text{He}}^{\text{sim}}$  terms the MLP PIL model can be brought in closer agreement with the expert model by tuning the hyperparameters  $\gamma_1$  and  $\gamma_2$ . An added advantage of KANs is

that for instances of nonphysical behavior such as the curl up behavior for low argument values in the  $YOC_{\text{He}}^{\text{sim}}$  and  $\hat{V}$  plots, the grid extension technique [11] could be performed post-training to fine-tune the spline hyperparameters which can mitigate such nonphysical behaviors. However, we leave this investigation on the choice of spline/grid hyperparameters pre-training and post-training for future works. In summary, the KAN model comes closest to reproducing the aggregate behavior of an expert derived model while also allowing for interaction effects between the inputs, which can possibly explain the better predictive ability on the new mentioned experiments. The physical interpretation and validity of these interaction effects will be investigated as the KAN is used to design new experiments, and will be discussed in subsequent publications.

## Acknowledgments and Disclosure of Funding

This material is based upon work supported by the Department of Energy under Award Nos. DE-SC0024381, DESC0022132, DE-SC0021072 and DE-SC0024456 as well as the Department of Energy (National Nuclear Security Administration) University of Rochester “National Inertial Confinement Program” under Award No. DE-NA0004144. This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

## References

- [1] John Nuckolls, Lowell Wood, Albert Thiessen, and George Zimmerman. Laser compression of matter to super-high densities: Thermonuclear (ctr) applications. *Nature*, 239(5368):139–142, 1972.
- [2] IV Igumenshchev, DH Edgell, VN Goncharov, JA Delettrez, AV Maximov, JF Myatt, W Seka, A Shvydky, S Skupsky, and C Stoeckl. Crossed-beam energy transfer in implosion experiments on omega. *Physics of Plasmas*, 17(12), 2010.
- [3] IV Igumenshchev, DT Michel, RC Shah, EM Campbell, R Epstein, CJ Forrest, V Yu Glebov, VN Goncharov, JP Knauer, FJ Marshall, et al. Three-dimensional hydrodynamic simulations of omega implosions. *Physics of Plasmas*, 24(5), 2017.
- [4] Varchas Gopalaswamy, Riccardo Betti, JP Knauer, N Luciani, D Patel, KM Woo, A Bose, IV Igumenshchev, EM Campbell, KS Anderson, et al. Tripled yield in direct-drive laser fusion through statistical modelling. *Nature*, 565(7741):581–586, 2019.
- [5] A Lees, R Betti, JP Knauer, V Gopalaswamy, D Patel, KM Woo, KS Anderson, EM Campbell, D Cao, J Carroll-Nellenback, et al. Experimentally inferred fusion yield dependencies of omega inertial confinement fusion implosions. *Physical review letters*, 127(10):105001, 2021.
- [6] Aarne Lees, Riccardo Betti, James P Knauer, Varchas Gopalaswamy, Dhruvir Patel, Ka Ming Woo, Ken S Anderson, E Michael Campbell, Duc Cao, Jonathan Carroll-Nellenback, et al. Understanding the fusion yield dependencies in omega dt-layered implosion experiments using a physics-based statistical mapping model. *Physics of Plasmas*, 30(1), 2023.
- [7] Rahman Ejaz, Varchas Gopalaswamy, A Lees, C Kanan, D Cao, and R Betti. Deep learning-based predictive models for laser direct drive at the omega laser facility. *Physics of Plasmas*, 31(5), 2024.
- [8] Ziming Liu, Yunyue Chen, Yuanqi Du, and Max Tegmark. Physics-augmented learning: A new paradigm beyond physics-informed learning. *arXiv preprint arXiv:2109.13901*, 2021.

- [9] Ali Takbiri-Borujeni, Mohammad Kazemi, and Sam Takbiri. Physics-informed machine learning for fluid flow prediction in porous media. In *ICLR 2024 Workshop on AI4DifferentialEquations In Science*, 2024.
- [10] David Davini, Bhargav Samineni, Benjamin Thomas, Amelia Huong Tran, Cherlin Zhu, Kyung Ha, Ganesh Dasika, and Laurent White. Using physics-informed regularization to improve extrapolation capabilities of neural networks. In *Fourth Workshop on Machine Learning and the Physical Sciences (NeurIPS 2021)*, 2021.
- [11] Ziming Liu, Yixuan Wang, Sachin Vaidya, Fabian Ruehle, James Halverson, Marin Soljačić, Thomas Y Hou, and Max Tegmark. Kan: Kolmogorov-arnold networks. *arXiv preprint arXiv:2404.19756*, 2024.
- [12] Anil K Jain and Richard C Dubes. *Algorithms for clustering data*. Prentice-Hall, Inc., 1988.
- [13] Peter J Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.
- [14] V Gopalaswamy, R Betti, JP Knauer, A Lees, D Patel, AR Christopherson, IV Igumenshev, D Cao, KS Anderson, A Shvydky, et al. Using statistical modeling to predict and understand fusion experiments. *Physics of Plasmas*, 28(12), 2021.
- [15] Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232, 2001.
- [16] Valeri N Goncharov, TC Sangster, R Betti, TR Boehly, MJ Bonino, TJB Collins, RS Craxton, JA Delettrez, DH Edgell, R Epstein, et al. Improving the hot-spot pressure and demonstrating ignition hydrodynamic equivalence in cryogenic deuterium–tritium implosions on omega. *Physics of Plasmas*, 21(5), 2014.
- [17] Evan M Yu, Juan Eugenio Iglesias, Adrian V Dalca, and Mert R Sabuncu. An auto-encoder strategy for adaptive image segmentation. *arXiv preprint arXiv:2004.13903*, 2020.
- [18] Jia Wu, Xiu-Yun Chen, Hao Zhang, Li-Dong Xiong, Hang Lei, and Si-Hao Deng. Hyperparameter optimization for machine learning models based on bayesian optimization. *Journal of Electronic Science and Technology*, 17(1):26–40, 2019.
- [19] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [20] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [21] KM Woo, R Betti, OM Mannion, CJ Forrest, JP Knauer, VN Goncharov, PB Radha, D Patel, V Gopalaswamy, and V Yu Glebov. Inferring thermal ion temperature and residual kinetic energy from nuclear measurements in inertial confinement fusion implosions. *Physics of Plasmas*, 27(6), 2020.

## A Appendix / supplemental material

### A.1 Input Parameter Description

Table 1: Model input parameters

Parameter	Description
$E_L$	Total laser energy on fusion fuel target.
$R_{out}$	Target outer radius.
$\hat{M}$	$\frac{Mass_{fuel\ in\ ice\ form}}{Mass_{target\ total}}$
$\hat{R}$	$\frac{Radius_{inner}}{Radius_{outer}}$
$\frac{R_b}{R_t}$	Captures effects of $\sigma_{rms}$ illumination nonuniformity[6, 14].
$\frac{IFAR}{CR}$	Stability parameter[6, 16], captures effect of short-wavelength perturbations.
$\hat{V}$	Initial outer shell radius/inner shell radius at maximum compression.
$\hat{Y}$	Ad-hoc parameter capturing nonlinear effect of pulse shape on maximum achieved velocity.
$T_{min}^{max}$	Ad-hoc parameter capturing nonlinear effect of pulse shape on simulated yield.
$YOC_{He}^{sim}$	Captures degradation magnitude when ion temperature asymmetries are present.
	Captures the degrading effect of $^3He$ gas contamination due to fuel target’s age.

### A.2 Model architecture, hyperparameters and training details

**Architecture and hyperparameter optimization** An informal grid search was used to optimize the model architecture and hyperparameter configurations. In future works a more formal approach for hyperparameter optimization and architecture ablation would be used utilizing the bayesian hyperprop strategy [18].

**Compute** All models are trained on a desktop workstation with a single Nvidia RTX 6000 GPU. The KAN model takes  $\approx 30$  mins to train. The MLP and MLP PIL models take  $\approx 180$  mins to train on a single train-test split.

**Train-test split** Scikit-learn[19] KFold was using for 21 train-test partitions, all results apart from out-of-distribution results are presented using split number 8.

**K-means for out-of-distribution tests** Scikit-learn[19] silhouette\_score and KMeans were used to obtain clusters for out-of-distribution tests, six clusters were used.

Table 2: MLP PIL and MLP configuration

Hyperparameter	Value
Input layer width	16 (extra parameters for one-hot encoding of target composition)
Fully connected hidden layer widths	71-71-71
Learning rate	1e-3
Drop-out regularization	0.1
Optimizer	Adam [20]
Batch size	11
Epochs	3000
Early stopping on	train loss
Loss function MLP	$\mathcal{L}_{MSE}$
Loss function MLP PIL	$\mathcal{L}_{MSE} + \mathcal{L}_{phys}$

Table 3: KAN training configuration. The implementation is utilized from the following source, <https://github.com/Blealtan/efficient-kan.git>

Hyperparameter	Value
Input layer width	16 (extra parameters for one-hot encoding of target composition)
Number of layers	7
Layer widths	71
Grid size	5
Spline order	3
Learning rate	1e-3
Drop-out regularization	0.2
Optimizer	Adam [20]
Batch size	11
Epochs	500
Early stopping on	train loss
Loss function	$\mathcal{L}_{MSE}$

### A.3 $\mathcal{L}_{physics}$ explanation

The  $\mathcal{L}_{physics}$  in Eq. 4 for both terms can naively be applied as a penalty whenever deviations from the MLP PIL model is exhibited during training as for both terms simple physics consideration can be used to establish the validity of the expected trend.

The  $T_{min}^{max}$  can be trivially shown[21] to be related to the residual kinetic energy of the fusing plasma. Ideally, all the kinetic energy is converted to internal energy of the fusing plasma, which then results in fusion. Reducing the amount of kinetic energy  $\rightarrow$  internal energy monotonically reduces the reaction rate and therefore the number of reactions, thereby requiring that  $\frac{\partial Y^{exp}}{\partial T_{min}^{max}}$  must be negative.

The term  $YOC_{He}^{sim}$  expresses the degrading effect resulting due to the target's age which is the time between when target assembly begins and when the experiment is executed. As the target age increases  $YOC_{He}^{sim} \rightarrow 0$ . Due to natural  $\beta$  decay some of the fusion fuel is constantly converted in  ${}^3\text{He}$  which through established physics causes increased radiation losses and because this gas accumulates in the gas region of the fusion fuel target this leads to a reduction in the amount of fuel compression that can be achieved in the laser fusion scheme. It can simply be deduced that both of these effects monotonically reduce  $Y^{exp}$ .

In comparison to these two effects, the dependence of  $Y^{exp}$  on  $\frac{R_b}{R_t}$  or  $\frac{\alpha}{IFAR}$  is not trivial. Both these terms have competing effects that either reduce or increase  $Y^{exp}$ ,  $\frac{R_b}{R_t}$  going down means more degradation seed[14], but also means better laser inefficiency mitigation = better energy transfer. Not obvious which one is stronger, where the trade-off is maximized, if its uni-modal, monotonic, etc,  $\frac{\alpha}{IFAR}$  going down means better simulation performance but worse experiment performance. Again, not clear if its uni-modal, bimodal, monotonic, etc. It needs to be learned from the data.