

Deep-Learning-Enabled Bayesian Inference of Fuel Magnetization in Magnetized Liner Inertial Fusion

William E. Lewis,^{1, a)} Patrick F. Knapp,¹ Stephen A. Slutz,¹ Paul F. Schmit,¹ Gordon A. Chandler,¹ Matthew R. Gomez,¹ Adam J. Harvey-Thompson,¹ Michael A. Mangan,¹ David J. Ampleford,¹ and Kristian Beckwith¹
Sandia National Laboratories, Albuquerque, New Mexico 87185 USA

(Dated: 20 July 2021, DRAFT)

Fuel magnetization in magneto-inertial fusion (MIF) experiments improves charged burn product confinement, reducing requirements on fuel areal density and pressure to achieve self-heating. By elongating the path length of 1.01 MeV tritons produced in a pure deuterium fusion plasma, magnetization enhances the probability for deuterium-tritium reactions producing 11.8 – 17.1 MeV neutrons. Nuclear diagnostics thus enable a sensitive probe of magnetization. Characterization of magnetization, including uncertainty quantification, is crucial for understanding the physics governing target performance in MIF platforms, such as Magnetized Liner Inertial Fusion (MagLIF) experiments conducted at Sandia National Laboratories Z-facility. We demonstrate a deep-learned surrogate of a physics-based model of nuclear measurements. A single model evaluation is reduced from $\mathcal{O}(10 - 100)$ CPU hours on a high-performance computing cluster down to $\mathcal{O}(10)$ ms on a laptop. This enables a Bayesian inference of magnetization, rigorously accounting for uncertainties from surrogate modeling and noisy nuclear measurements. The approach is validated by testing on synthetic data and comparing with a previous study. We analyze a series of MagLIF experiments systematically varying preheat, resulting in the first ever systematic experimental study of magnetic confinement properties of the fuel plasma as a function of fundamental inputs on any neutron-producing MIF platform. We demonstrate that magnetization decreases from $BR \sim 0.5$ MG-cm to $BR \sim 0.2$ MG-cm as laser preheat energy deposited increases from $E_{\text{preheat}} \sim 460$ J to $E_{\text{preheat}} \sim 1.4$ kJ. This trend is consistent with 2D LASNEX simulations showing Nernst advection of the magnetic field out of the hot fuel and diffusion into the target liner.

I. INTRODUCTION

Magnetized Liner Inertial Fusion (MagLIF)¹ is a magneto-inertial fusion (MIF) concept being studied on the Z-machine^{2,3} at Sandia National Laboratories that has demonstrated thermonuclear neutron generation with temperatures of several keV and densities $\mathcal{O}(1 \text{ g/cm}^3)$ at stagnation.⁴ MagLIF relies on three stages to reach these conditions. First, an axial magnetic field ($10 - 20 \text{ T}$)⁵ is applied to an approximately 1 cm tall and 0.5 cm diameter cylindrical beryllium target or “liner” containing gaseous deuterium (DD) fuel at about $0.7 - 1.05 \text{ mg/cc}$. For the experiments discussed in this manuscript, liners with aspect ratio 6 (AR6, where $AR = r_o/\Delta r$, with r_o the outer liner radius, $\Delta r = r_o - r_i$, and r_i the inner liner radius) are used, although AR9 liners have also been regularly fielded. Second, a preheat pulse from a multi-kJ laser^{6,7} is shot into the target coupling $0.5 - 2 \text{ kJ}$ of energy to the fuel, serving to set the fuel on a high adiabat and ionize the deuterium to lock in the magnetic field. Finally, a current pulse ($\mathcal{O}(20 \text{ MA})$ peak current $\sim 100 \text{ ns}$ rise time) generates a $\mathbf{J} \times \mathbf{B}$ force that implodes the target, compressing the fuel to achieve fusion relevant conditions. The current-driven implosion heats the fuel quasi-adiabatically to multi-keV temperatures with convergence ratios (ratio of initial to final inner liner radius) reaching ~ 30 . In addition, the background magnetic field is compressed to high strengths, improv-

ing overall fuel magnetization and aiding in confinement of charged fusion products.

The MagLIF platform has continued to produce exciting results, such as the recently demonstrated peak performance of a primary deuterium-deuterium (DD) neutron yield exceeding 10^{13} (2 kJ equivalent were 50:50 DT fuel to be used).⁸ Furthermore, a recent computational study of the scaling of neutron yield as a function of input parameters indicates the potential to increase yield by another order of magnitude as experimental capabilities improve to allow increases in current delivered, laser preheat energy, and initial magnetic field.⁸ The extent to which performance can be enhanced by such improvements is sensitive to our understanding of thermal transport in the presence of the highly compressed background magnetic field in the target, which determines how well the fuel is thermally insulated. As a result, measurements of the magnetization during fusion burn is critical for advancing the MagLIF platform.

Understanding performance scaling and improving experimental design has increasingly relied on our ability to obtain a detailed characterization of the fuel plasma through various diagnostics. For example, self-emission x-ray images have been used to study morphology of the fuel plasma near stagnation^{9,10}, liner-radiograph images characterize the Magneto-Rayleigh-Taylor instability^{11–15}, x-ray spectra and yields illuminate fuel plasma composition¹⁶, temperature, liner areal density, and density¹⁷, and neutron diagnostics such as neutron time-of-flight (nToF) constrain ion temperature, liner areal density¹⁸, and magnetization^{19,20}. Unfortunately, routine analysis of experiments is a challenge. Expending

^{a)}Electronic mail: willewi@sandia.gov

resources required to obtain estimates of physical conditions from experiments using state-of-the-art models can become prohibitive, especially when conducting rigorous uncertainty quantification (UQ), which incurs significant additional computational costs.



These problems are ubiquitous to the broader Inertial Confinement Fusion (ICF) community. Motivated by these challenges, significant effort has been placed on developing machine-learning-based techniques to alleviate computational requirements, establish data-processing pipelines that reduce the need for significant user input that might bias results, and illuminate model inadequacies.^{9,21–23} In this work, we develop such a tool designed to better enable inference and UQ of the fuel magnetization during fusion burn in MagLIF experiments. Previous analysis of fuel magnetization has involved application of the physics based model developed and detailed in Refs. 19 and 20. Unfortunately, a single forward evaluation of that model requires a computation time of $\mathcal{O}(10\text{--}100)$ CPU hours on a high-performance computing cluster. Additional burden is incurred when attempting even a simple uncertainty calculation based on a model sensitivity analysis²⁰, while a rigorous treatment using a Bayesian statistics framework is not feasible for routine analysis. Due to this poor scaling, prior to the development of our method, an inferred value of BR had been published for only a single MagLIF experiment in Refs. 19 and 20.²⁴ Our approach relies on the use of an artificial neural network (ANN) to surrogate a physics based calculation of secondary neutron observables, resulting in a reduction of computational time for a single forward model evaluation down $\mathcal{O}(10)$ ms on a personal laptop. This enables fast inferences using the rigorous Bayesian statistics framework for uncertainty quantification.

Our paper is organized as follows. We begin by providing an overview of the physics relating neutron diagnostics to fuel magnetization. Next we detail our analysis pipeline. Particularly, logic behind the choices made in the analysis pipeline including data preprocessing, use of a deep-learning based surrogate model, and uncertainty quantification with Bayesian analysis will be covered. We apply the method to several synthetic data examples to validate the approach on cases with known solutions. Next, to confirm the efficacy of this new pipeline, we demonstrate consistency of our approach with an experiment analyzed previously with the model in Refs. 19 and 20. Finally, we analyze a series of MagLIF experiments where the total preheat energy was systematically varied, leading to measurable changes in the fuel magnetization during fusion burn. Particularly, by applying modern data science tools to experimental data and comparing with 2D LASNEX simulations^{25,26}, we demonstrate that the trend in fuel magnetization as a function of preheat is consistent with the Nernst effect. The Nernst effect is a thermo-electric effect whereby hot electrons transport magnetic flux along thermal gradients from hot to cold regions.²⁷ Due to the radial temperature gradients established by the laser preheat and the initially axially

oriented magnetic field, one might expect the Nernst effect to more effectively transport magnetic flux to the fuel-liner interface where it can resistively diffuse into the liner. Subsequent resistive diffusion of axial magnetic flux into the liner ultimately leads to the loss of some of the initial enclosed magnetic flux from the fuel. Our results indicate the importance of detailed understanding of these magnetic field transport mechanisms in determining fuel magnetization near stagnation, and is supported by recent computational investigation.²⁸ We conclude the paper with a discussion of possible improvements to the method and highlight additional physics we plan to investigate going forward.

II. OBSERVABLE EFFECTS OF FUEL MAGNETIZATION

Here we provide a summary of the details presented in Refs. 19 and 20. We emphasize the statement by the authors of those works that the large azimuthal magnetic field ($\gtrsim 50$ MG) driving the Z-pinch implosion in MagLIF experiments prohibits external probing of the fuel magnetization e.g. via proton deflectometry/radiography^{29–32}. Instead, we seek to understand how fuel magnetization alters the nuclear physics occurring near stagnation. To this end, consider the primary and secondary fusion reaction paths in the DD fuel given by



where the two branches of the primary DD reaction are roughly equally probable (energies reported in center-of-momentum frame). The upper channel for the primary reaction produces a ~ 2.45 MeV *primary* neutron while the lower path produces a 1.01 MeV triton that can subsequently interact with the deuterium background to generate a *secondary* DT neutron in with a kinematically allowed range of 11.8 – 17.1 MeV in the lab frame.³³

To demonstrate the impact of fuel magnetization on secondary neutron total yields and energy spectra, we utilize the model presented in Refs. 19 and 20. This model will henceforth be referred to as the full-order model (FOM) to distinguish it from the ANN surrogate model which we will construct to approximate the FOM in Sec. III B. The FOM is composed of a Landau-Fokker-Planck (LFP) model³⁴ used to track test-particle tritons through a background deuteron-electron plasma, accounting for electron-ion and ion-ion scattering, and the so-called MonteBurns model³⁵ to handle nuclear reactions with the background plasma. The LFP model assumes a charge neutralizing electron population in local thermal equilibrium (LTE), with equivalent electron and ion temperatures ($T_e = T_i$), so that a single temperature profile can be specified to describe the background fuel plasma (henceforth T). This assumption is reason-

able given the general consistency of experimentally observed T_e and T_i estimated from x-ray spectra and DD neutron time-of-flight (nToF) respectively.⁴ Furthermore, the experimentally estimated electron-ion collision time, $\tau_{e-i} \ll 1$ ps, is much faster than the stagnation duration ~ 2 ns, allowing plenty of time for electron-ion temperature equilibration during the quasi-adiabatic implosion.⁴ It bears mentioning that, from a physical standpoint, the electron temperature is more relevant. Particularly, fast tritons and α -particles slow down primarily on plasma electrons until they have lost the majority of their birth energy reaching an energy of 100 keV, where the DT reactivity begins to peak and ion-ion effects become more significant than at higher energy.

In keeping with the analysis of Ref. 20, the axially uniform, radially varying fuel profiles for a cylinder of radius R and finite height Z are given by

$$T(r) = T_c \left[1 - 0.9 \left(\frac{r}{R} \right)^2 \right], \quad (3)$$

$$\rho(r) = \rho_c \left[1 - 0.9 \left(\frac{r}{R} \right)^2 \right]^{-1}, \quad (4)$$

$$B_z(r) = B_{z,c} \left[1 - 0.9 \left(\frac{r}{R} \right)^2 \right]^{-1}. \quad (5)$$

In the above, $T(r)$ is the background fuel temperature, $\rho(r)$ the background density, and $B_z(r)$ the axially oriented magnetic field. Parameters not included in the above profiles are mix fraction, F (percent by atom) of fully ionized beryllium, and the fuel cylinder aspect ratio $\mathcal{A} = Z/R$. Note that the liner AR is not the same as \mathcal{A} describing the ratio of stagnation fuel plasma height to radius used in the FOM being discussed. Notice also that $B_z(r) \propto \rho(r)$, consistent with the frozen-in-flux law for highly conductive plasmas, but ignores flux redistribution effects associated with Nernst and resistive diffusion. We additionally follow the practice in Refs. 19 and 20 of fixing $R = 50 \mu\text{m}$. This is motivated by emission imaging and spectroscopy, which indicate that this is a reasonable choice, and has the added benefit of allowing us to make direct comparison to the analysis of those works. It is straightforward to modify the analysis we conduct to allow the fuel radius to vary, but is not expected to have a significant impact on the results presented in this work because the fundamental dependence is on BR and ρR , not B , ρ , or R .

Before proceeding to discuss the impact of magnetization, we discuss the inputs provided to the FOM and how they correspond to parameters in Eqs. 3-5. Input values of magnetic field, fuel density, temperature, aspect ratio, and mix fraction, denoted by B^{in} , ρ^{in} , T^{in} , \mathcal{A}^{in} , and F^{in} are provided to the FOM routine. Those values are transformed for use in the fuel profile equations according to the following prescription. First, T_c is found by requiring that the ratio of the DD neutron production rate to total fuel $(\rho R)^2$ be equal to that of a uniform fuel profile of temperature T^{in} , with both the radially varying and uniform profiles having the same radial extent. That is,

we solve

$$\frac{\int_0^R \frac{(\rho(r))^2}{2} [\langle \sigma_{\text{DD}} v \rangle (T(r))] r dr}{\left(\int_0^R \rho(r) dr \right)^2} = \frac{[\langle \sigma_{\text{DD}} v \rangle (T^{\text{in}})]}{8}, \quad (6)$$

for T_c , where $[\langle \sigma_{\text{DD}} v \rangle (T(r))]$ is the (locally) Maxwell-averaged DD-reactivity depending on the temperature at radius r .³⁶ Note that Eq. 6 is independent of ρ_c and ρ^{in} , and can be used to set T_c as a function of T^{in} and mix fraction. Next ρ_c is found by requiring ρR for the radially varying fuel profile to match that of a uniform profile of the same radius with density given by ρ^{in} and uniform mix concentration of beryllium given by F^{in} . Finally, $B_{z,c}$ is found by requiring the radially varying profile to contain the same flux as a uniform profile of the same radius with magnetic field given by B^{in} . F^{in} requires no transformation, and $\mathcal{A}^{\text{in}} = Z/R$ is used with $R = 50 \mu\text{m}$ to obtain the cylinder height Z . We emphasize that these prescriptions do not rely on any experimental information, but rather provide a simple mapping between our input variables and model parameters that may aid in interpretability via an “equivalent” uniform fuel profile.

The chosen normalizations of the radially varying fuel profiles aid in physical interpretation via “equivalent” uniform profile parameters, as opposed to simply setting the central axis values. It is worth pointing out the interpretation of the the normalization of $B_{z,c}$ in more detail. Without writing out the explicit profile we have

$$B^{\text{in}} R = \frac{\Phi_z}{\pi R} = \frac{2}{R} \int_0^R r B_z(r) dr, \quad (7)$$

with Φ_z being the axial flux for the uniform profile of radius R and axial field B^{in} . Notice that $B^{\text{in}} R$ is precisely the definition of *effective* BR given in Ref. 20, which was defined to maintain ties with existing MIF work while also directly corresponding with Φ_z . The significance of this definition is that Φ_z along with the fuel radial extent R determine the confinement properties of the fuel plasma. To see this, we appeal to conservation of the canonical angular momentum given by $\mathbf{L}^{\text{can}} = \mathbf{r} \times m\mathbf{v} - q\mathbf{A}$. Supposing an axially oriented field $\mathbf{B}(r) = B(r)\hat{z}$ for $r \leq R$, one has that $\mathbf{A}(r) = A(r)\hat{\phi}$, so that it is easy to see using $\mathbf{B} = \nabla \times \mathbf{A}$ and integrating over a circular cross section to radius r_0 one finds $A(r_0) = \Phi_{r_0}/(2\pi r_0)$, with Φ_{r_0} being the flux enclosed by radius r_0 . For a radially outward moving ion born at radius r_i , the initial canonical angular momentum is thus $L_0^{\text{can}} = -q\Phi_{r_i}$. If the ion is just confined to the fuel it will be orbiting at radius R , giving $L_f^{\text{can}} = mv_{\perp} R - q\Phi_R$. As a result, the condition for trapping is

$$\frac{\Phi_R - \Phi_{r_i}}{\pi R} \geq \sqrt{\frac{8mE_{\perp}}{q^2}}, \quad (8)$$

where $E_{\perp} = mv_{\perp}^2/2$. Notice that for $r_i = 0$ the left hand side reduces to the effective BR defined by Eq. 7. We also




FIG. 1. Top Left: Axial (solid) and Radial (dashed) secondary neutron spectra (normalized) with $\log_{10}(BR) = 4.0$, $\rho R \approx 2$ mg/cm^2 , $T = 3$ keV, $\mathcal{A} = 80$, and $F = 0$, showing severe anisotropy between the two views. Top Middle: Axial (solid) and Radial (dashed) secondary neutron spectra with $\log_{10}(BR) = 5.22$ and other parameters as in the $\log_{10}(BR) = 4.0$ case, showing the dip in the axially viewed spectrum beginning to fill in and the shoulders of the radial peak becoming less pronounced. Top Right: Axial (solid) and Radial (dashed) secondary neutron spectra with $\log_{10}(BR) = 6.5$ and other parameters as in the $\log_{10}(BR) = 4.0$ case, showing significant reduction in the anisotropy between views. Bottom: Impact of $\log_{10}(BR)$ on the ratio of DT to DD neutron yields.

pause to mention that, as discussed in Ref. 19, the birth gyroradii of DD tritons ($r_{L,t}$) and DT α -particles are nearly identical, with $r_{L,\alpha} \approx 1.07r_{L,t}$, while the stopping length for $E_\alpha = 3.5$ MeV α -particles is roughly half the stopping length of 1.01 MeV tritons. As a result this makes tritons a good surrogate for 3.5 MeV α -particles, so that measuring BR via secondary neutron signatures has direct relevance to our understanding of α -particle trapping.

Equipped with the understanding that BR is fundamentally a confinement parameter, we can move on to develop intuition for how this property impacts nuclear measurements. To do so, we will consider FOM calculations for a high aspect ratio $\mathcal{A} = 80$ cylinder, with radial areal density of deuterium $\rho R \approx 2$ mg/cm^2 , $T = 3$ keV, mix fraction $F = 0$. Note that we have dropped the superscript “in” for the input parameters for convenience. For the two cases of weakly magnetized $\log_{10}(BR) = 4.0$ ($R/r_{L,t} \approx 0.04$) and strongly magnetized $\log_{10}(BR) = 6.5$ ($R/r_{L,t} \approx 12.75$) respectively (where BR is in G-cm), one finds striking differences in the secondary neutron spectrum viewed along the axial and radial directions, as well as in the ratio of the secondary neutron yield to the primary neutron yield (\bar{Y}). Figure 1 shows that, in the weakly magnetized case, the radially and axially viewed spectra are highly anisotropic while for the highly magnetized case, the spectra, though

still anisotropic, are much less so. Furthermore, \bar{Y} has increased roughly two orders of magnitude.

These features can be explained by considering the secondary neutron spectrum produced by radially and axially moving tritons separately as seen along both the radial and axial directions. We first consider the weakly magnetized case $\log_{10}(BR) = 4.0$. For the axially viewed spectrum, the center-of-mass (CM) frame has a velocity component along the line of sight for tritons with a significant axial velocity component, causing an up- or down-shift in secondary neutron energy when transforming back to the lab frame. This Doppler-shift results in an accumulation of signal in the neutron spectra on either side of 14.1 MeV. On the other hand, tritons moving in a radial plane produce no such shift since there is no velocity component along the line of sight. However, since these tritons experience a much smaller path length before escaping the radially narrow fuel, there will be significantly fewer neutrons observed without an energy shift. For the radially viewed spectrum the situation is reversed, and axially moving tritons do not produce an energy shift away from 14.1 MeV when transforming from the CM to lab frame. Initially radially moving tritons do produce such a shift, but are observed with significantly lower frequency, again due to reduced path length in the deuterium background resulting from weak magnetic confinement.

The strong anisotropy between the radially and axially viewed spectrum characteristic of the weakly magnetized case begins to lift as fuel magnetization is increased. For $\log_{10}(BR) = 6.5$, initially radially moving tritons are effectively magnetically confined by the strong axial field, with average path length becoming $\sim Z$ instead of $\sim R$. As a result these primarily radially moving tritons begin to contribute substantially more to the measured signal for both views as compared to predominantly axially moving tritons. Some peak splitting is now observable in the radially viewed spectra due to transforming into the lab frame. This occurs because there is now a wider range radial velocity components along the line-of-sight that result in trapped tritons causing a Doppler-shift in energy when transforming to the lab-frame. Although the axially viewed spectrum technically still retains the double-peaked structure arising from axially moving tritons, the peak at 14.1 MeV from radially moving tritons which do not experience a significant doppler shift when transforming to the lab frame now dominates as many more radially moving tritons are magnetically confined to the fuel. Fig. 1 also includes an intermediate case with $\log_{10}(BR) = 5.22$ so that the interested reader may further investigate how these arguments apply.

Experimentally, nToF signals rather than spectra are measured along axial and radial lines of sight. The conversion between simulated neutron energy spectra and nToF is computed as

$$\frac{dN}{dt} = BH(t) * IRF(t) * LO(E) \frac{dN}{dE} \frac{dE}{dt}, \quad (9)$$

where $*$ indicates convolution, $BH(t)$ is the burn history of the fuel, which confines neutron production temporally, $IRF(t)$ is the temporal instrument response function, $LO(E)$ is the energy dependent light output response of the scintillator, dE/dt is a Jacobian computed from special relativity, and dN/dE is the energy spectrum computed from the FOM. The burn history is assumed to be Gaussian in time with a 2 ns full width half max (FWHM). While the burn history cannot be directly measured, this approximation is motivated by photo-conducting diamond (PCD) measurements of the x-ray output as a function of time, which indicate a 1.7–2 ns burn duration.⁴ The assumption of a ~ 2 ns Gaussian burn history is not expected to significantly bias the synthetic nToF due to the relatively broad energy spectrum of neutrons produced in secondary reactions. The IRF has been experimentally characterized as having ~ 4 ns FWHM with roughly Gaussian rise and exponential falloff. The LO response has been modeled using Monte Carlo N-Particle simulations, and is roughly linear over the relevant energy range. Although we use the more complex IRF and LO functions described above, we provide approximate analytic expressions (denoted by an over-line) for these so that it would be possible to

closely reproduce our results if desired:

$$\overline{IRF}(t) = \exp\left(-\frac{t-\mu}{\tau}\right) \exp\left(\frac{\sigma^2}{2\tau^2}\right) \times \left(1 + \operatorname{erf}\left(\frac{t-\mu-\frac{\sigma^2}{\tau}}{\sqrt{2}\sigma^2}\right)\right), \quad (10)$$

$$\overline{LO}(E) = mE + b. \quad (11)$$

The parameters for the above analytic expressions were found by fitting the analytic expressions to the values used in our analysis using a mean-square-error (MSE) minimization with the results $\mu = 11.5$ ns, $\tau = 1.9$ ns, $\sigma = 1.1$ ns, $m = 0.1$ MeV⁻¹, and $b = 1.67$. The amplitude of \overline{IRF} was not fit as this would constitute an overall normalization of the computed secondary neutron nToF signal which will be normalized to have peak value of unity anyways.

Given the linearity of the LO over the secondary neutron energy range and narrow FWHM for the temporal response and burn history compared to the several tens of nanoseconds difference in arrival time of the fastest and slowest secondary neutrons, the nToF signal is essentially equivalent to the spectrum and will exhibit similar behavior as a function of magnetic field. There is of course the factor of $\frac{dE}{dt}$, which will introduce some additional asymmetry due to the non-linearity of the transformation from the energy to time domain. The IRF is also not perfectly symmetric and will contribute to slightly different shapes of the nToF signals in comparison to the spectra.

In addition to reducing anisotropy between the radially and axially viewed secondary neutron spectra, fuel magnetization has a significant effect on the ratio of DT to DD neutrons as shown in Fig. 1. Refs. 19 and 20 discuss that this ratio \bar{Y} is an increasing function of $\langle\rho_D\ell\rangle$, with ρ_D the density of the deuterium background and ℓ the path length of the tritons in that background. Since we expect ℓ to be a strong function of fuel magnetization, this explains the observed behavior of \bar{Y} with BR . Particularly, Ref. 20 demonstrates that for small values of BR $\langle\ell\rangle \sim R$, while for large values of BR $\langle\ell\rangle \sim Z$. This accounts largely for the about 80 \times increase in \bar{Y} for the increase in BR shown in Fig. 1. Axial and radial secondary nToF measurements combined with primary and secondary neutron yields will be the signals of interest in inferring BR . For a more complete description of the sensitivities of the the secondary neutron spectra and yield to model parameters other than BR , we refer the reader to Refs. 19 and 20.

III. ANALYSIS FRAMEWORK

The basic ingredients of our analysis pipeline are indicated in Fig. 2. There are three key stages that will be discussed in this section. The first is a data preprocessing stage where neutron time-of-flight (nToF) signals are prepared for an automatic featurization, which includes an estimate of uncertainty in the extracted features. The

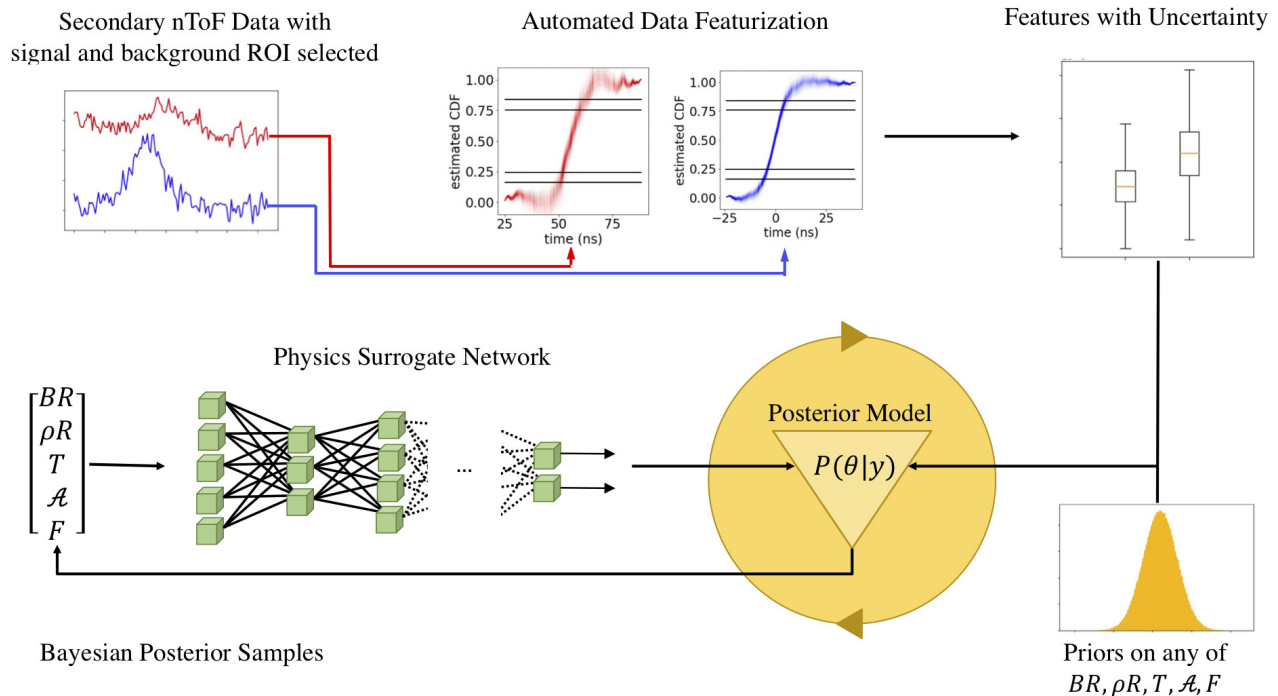


FIG. 2. Summary of the analysis pipeline. Beginning in the upper left and moving clockwise: A user provides cropped secondary nToF data to a deep convolutional network to compute the shape summary features with uncertainty using the algorithm described in the main text. These features, the primary and secondary neutron yields, nToF shape features, and priors on BR , ρR , T , \mathcal{A} , and mix fraction F are provided to the Bayesian posterior model. All prior distributions are assumed to be uniform except for T which is assumed normally distributed with mean value and uncertainty estimated from analysis of DD neutron nToF measurements. A neural network surrogate model based on the neutron physics model in Refs. 19 and 20 is used to compute nToF shape summaries and yields given model input parameters to compare to observed values. Posterior parameter samples are sampled using a gradient enhanced Markov Chain Monte Carlo algorithm known as No U-Turn Sampling³⁷.

second is the use of a deep neural network to surrogate a first principles physics calculation of the nToF and neutron yield features that can be collected experimentally and compared to the model. The third and final stage is a Bayesian inference to extract the posterior distribution of the fuel magnetization by comparing the physics surrogate model to experimental data under a statistical likelihood model along with appropriate prior information imposed on the system.

A. Data Preprocessing

As discussed previously, the ratio of secondary to primary neutron yields as well as the widths of the axial and radial secondary neutron energy spectra are sensitive to fuel magnetization. In this section, we focus on detailing aspects that differ from the approach of Refs. 19 and 20. In particular, we utilize the widths and asymme-

tries of the axial and radial nToF measurements for the secondary neutrons directly rather than converting them to energy spectra. By utilizing the nToF measurements directly, we avoid uncertainties which can arise due to inverting the nToF signals into spectra. Particularly, the inversion requires an absolute timing fiducial which is not available. Additionally, division by small values from the relativistic Jacobian may amplify noise. Furthermore, deconvolution of the instrument response presents an ill-posed problem, further motivating a direct use of the nToF data. We note that, unless otherwise specified, all references to nToF signals refer to the secondary DT neutron peaks in those signals.

There are many reasonable approaches to characterizing the widths and asymmetries (generically, shape features) of the nToF signals, including full width at half max (FWHM), difference between left and right half width at half max (HWHM), and second- and third-moments (variance and skewness) about the mean of the nToF trace. Since the nToF signals can be rather noisy,

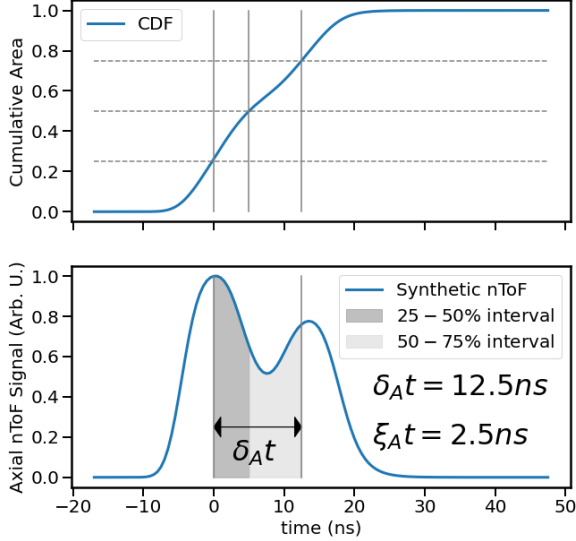


FIG. 3. Summary of the nToF width computation with the top panel showing the CDF computed from the spectrum in the lower panel. The CDF is used to find percentiles, which can then be used to compute shape features such as δ_{At} and ξ_{At} as described in the text.

finding the correct maximum peak from which to measure half-max points introduces significant uncertainty. Furthermore, computing statistical moments involves multiplying by a function of time that increases away from the nToF mean before integrating can amplify noise and make extracting uncertainty in the measured width difficult to formalize. In an effort to alleviate these issues, we use a percentile based definition of the nToF shape features as illustrated on a synthetic nToF signal in Fig. 3, which will tend to integrate out fluctuations produced by noise. As an example, in the ideal setting (no noise or background), the cumulative distribution function for the nToF signal

$$C_c(t) = \frac{\int_{-\infty}^t nToF_c(t') dt'}{\int_{-\infty}^{\infty} nToF_c(t') dt'}, \quad (12)$$

where $c = A, R$ indicates the axial or radial observation axis respectively, is computed. The two features we will define as δ_{At} and ξ_{At} can then be calculated as follows. First, find the crossings $C(t_{0.25}) = 0.25$, $C(t_{0.5}) = 0.5$, and $C(t_{0.75}) = 0.75$. After solving for $t_{0.25,0.5,0.75}$ we compute the features as

$$\delta_{At} = (t_{0.75} - t_{0.25}) \quad (13)$$

$$\xi_{At} = (t_{0.75} + t_{0.25} - 2 * t_{0.5}), \quad (14)$$

where the first equation is a measure of the central 50% width of the nToF signal, and the second measures the asymmetry of that interval about the 50%




FIG. 4. Example nToF shape computation on real data. Left panel shows CDFs computed after subtracting different fits to the background (inset). For each CDF we compute the shape features, giving a distribution of features (right panel) consistent with the experimental observation. The red curve in the right panel is a normal distribution with mean and standard deviation computed from the samples. The covariance matrix of these distributions of shape features is used as the estimate of uncertainty.

point. We additionally compute the same features, but using the central 68% interval (*i.e.* focusing on the wider 16% – 84% interval rather than the 25% – 75% interval). While the features using the wider interval do correlate fairly strongly with the corresponding features from the 25% – 75% interval, they contain non-redundant information (*i.e.* they are not perfectly linearly correlated). Intuitively, they capture information about the rate of change of width and asymmetry of the nToF, and the wider interval captures more information in the tails of the asymmetric nToF signal. However, despite the advantages of including the wider interval, we do not attempt to utilize an even larger interval (*e.g.* the 5% – 95% interval), as this unnecessarily introduces far too much uncertainty into the estimate when featurizing noisy signals as detailed in the next paragraphs. Furthermore, we expect that the late-time tail may be susceptible to neutron scattering effects that are as-of-yet unincorporated in the FOM, so we wish to avoid contaminating the features we extract with that effect to the greatest possible extent while maximizing useful information. We denote the features utilizing the 16% – 84% interval with the corresponding capitalized Greek symbols, and can perform the same procedure on a radial nToF signal, leading to a total of 8 shape features $\{\delta_{At}, \xi_{At}, \Delta_{At}, \Xi_{At}, \delta_{Rt}, \xi_{Rt}, \Delta_{Rt}, \Xi_{Rt}\}$. Notice that the computed shape features are invariant to translation of the nToF signal as long as a large enough region is included in the computation of the CDF. As a result, no

information about arrival time of the nToF peak is incorporated into the shape features.

We use the idealized procedure described above to generate training data for our physics surrogate model from simulated nToF signals as described in the next section. Unfortunately, the process is insufficient in treating real experimental nToF signals as these contain contributions from random noise, digitization of the signal, and an (assumed additive) background. Thus, experimental data requires additional steps to extract the nToF shape features with uncertainty. Note that the following process is applied to each of the axial and radial signals independently. As a first step, the analyst must provide the nToF signal with an estimate of the region containing the peak of the nToF signal. The rest of the signal will be treated as background data (denoted BG_{obs}). In order to obtain a good Bayesian fit to the background, a large region of background signal should be included with the nToF signal while avoiding any obvious pathologies at early or late times. To fit the axial (radial) nToF background, we use an exponential (linear) model of the form

$$BG_A(t, \gamma, \beta, \tau) = \gamma \exp^{-t/\tau} + \beta, \quad (15)$$

$$BG_R(t, m, b) = mt + b. \quad (16)$$

We point out here that in principle an exponential fit to the radial nToF background is possible, but we choose a linear model as there is no obvious curvature to the background in this case unlike the axial signal that typically shows an apparent exponential decay. This difference in background is consistent with the radial detector being further from the target than the axial detector, thus temporally separating the feature of interest from the scintillator decay caused by the earlier photon pulse.

For our Bayesian fitting procedure, we use a normal likelihood model, assuming that digitization noise does not invalidate the use of such a model. A Markov Chain Monte Carlo (MCMC) sampling method is used to generate samples from the posterior distribution for the background model parameters $P(\theta_c | BG_{\text{obs}})$ given the observed background data, where, again, $c = A, R$ indicates the axial or radial observation axis respectively, $\theta_A = \gamma, \beta, \tau$ and $\theta_R = m, b$. By generating a posterior for the background fit parameters, we can then estimate the uncertainty in the extracted nToF width as follows. Letting $\{BG_{c,i}\}_{i=1}^{N_{\text{MCMC}}}$ denote the background model associated with parameters $\{\theta_{c,i}\}_{i=1}^{N_{\text{MCMC}}}$ drawn from $P(\theta_c | BG_{c,\text{obs}})$, where $N_{\text{MCMC}} \sim \mathcal{O}(10^4)$ is the number of samples, we can compute a set of background subtracted signals $\{S_i\}_{i=1}^{N_{\text{MCMC}}}$ from which cumulative distribution functions can be calculated and processed as in the ideal case. The result is a set of nToF shape features $\{\delta_{ct_i}, \xi_{ct_i}, \Delta_{ct_i}, \Xi_{ct_i}\}_{i=1}^{N_{\text{MCMC}}}$ consistent with the observed nToF signal under the background. The mean and covariance matrix of this set of observations is used for the experimentally observed shape features and their (potentially correlated) uncertainty. For convenience, we note that the covariance matrix for a set of N observations of k quantities $\{\mathbf{X}_n\}_{n=1}^N$, with $\mathbf{X}_n \in \mathbb{R}^k$ is given by the

$k \times k$ matrix

$$\Lambda_{i,j} = \mathbb{E}_n[(X_i - E_n[X_i])(X_j - E_n[X_j])], \quad (17)$$

with $i = 1, 2, \dots, k$ and $j = 1, 2, \dots, k$ denoting a particular component of $\mathbf{X}_n \in \mathbb{R}^k$ and the \mathbb{E}_n denoting the expected value over the N observations. We note that this procedure amounts to assuming the observed shape features are described by a multivariate normal random variable. Though this assumption is imperfect, the normal approximation will conservatively overestimate the uncertainty in the experimentally extracted features. See Fig. 4 for a summary of this process applied to an experimental axial nToF signal. In addition to the nToF shape features, we utilize DD yield, Y_{DD} , as well as the ratio of DT to DD yield, \bar{Y} , as observations to constrain the parameters of the physics surrogate model in a Bayesian framework as described in Secs. III B and III C.

B. Physics Surrogate Model

Now that we have discussed which observations are going to be used to constrain our model input parameters, in particular the fuel magnetization, we discuss the construction of a fast deep-learning-based surrogate model. Unfortunately, the FOM presented in Refs. 19 and 20 takes $\mathcal{O}(10 - 100)$ core-hours to run for a single instance of its input parameters. While not slow on the scale of complex multi-physics simulations used to model fusion experiments, it does present a challenge when trying to perform UQ analysis requiring many tens- to hundreds-of-thousands of evaluations of the forward physics model for a single inference. To alleviate this issue, we build an artificial neural network (ANN) surrogate trained on data generated from the FOM. This ANN surrogate runs in $\mathcal{O}(10\text{ms})$, thus enabling experiment analysis with accurate UQ to be completed in a few minutes on a standard personal laptop. An additional benefit of surrogating the physics model with an ANN is the fact that gradients of the model outputs with respect to the inputs are analytically available and readily computed by many of the commonly used probabilistic programming packages used to create and train such ANNs. This will enable the use of certain tools in the Bayesian inference step of our analysis pipeline, and will be discussed in the next section. For now, we proceed to discuss the ANN training data and architecture, followed by a discussion of the performance of the surrogate model.

About 65k FOM evaluations are obtained using a Halton sampling scheme to generate approximately uniformly distributed random samples³⁸ over the hyper-rectangle defined by $\log_{10}(BR) \in [4.0, 6.5]$, $\log_{10}(\rho R) \in [-3.0, -1.5]$, $T \in [1.0, 6.0]$, $\log_{10}(\mathcal{A}) \in [0.9, 2.7]$, and $F \in [0.0, 0.2]$. The bounds of the sample domain are chosen to cover the ranges of fuel conditions that could be reasonably expected to occur in MagLIF experiments (example of typical conditions $BR \sim 0.3\text{MG-cm}$, $\rho R \sim 2\text{mg/cm}^2$, $T \sim 3\text{keV}$, $\mathcal{A} \sim 100$,

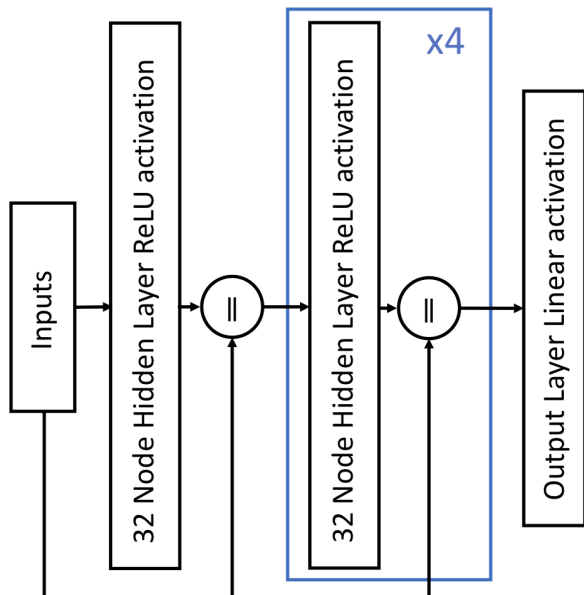


FIG. 5. Network Architecture where \parallel indicates concatenation. The blue plate is meant to indicate 4 layers with the same structure where the output of the previous layer is concatenated with the input before being fed to the next layer with the same structure but different trainable weights. Note that the output of the last layer is concatenated with the input before being fed to the linear output layer. We note that Resnet architectures were also found to work well.

$F \sim 0.1$). We split this dataset into three parts, a training dataset of about 62k samples, a validation dataset of about 2.7k samples used to estimate out-of-sample (OOS) covariance (detailed later), and about 500 test samples not used to train or estimate OOS statistics for verification work in Sec. IV We transform these input values for use in the radially varying fuel profiles in Eqs. 3-5 as discussed in Sec. II. The FOM computes secondary DT neutron spectra, Y_{DD} , and \bar{Y} . The FOM outputs are processed to obtain the final set of features $\delta t_{A,R}$, $\xi t_{A,R}$, $\Delta t_{A,R}$, $\Xi t_{A,R}$, Y_{DD} and \bar{Y} that will be used in training our ANN. We remind the reader that the shape features are found through application of Eq. 9 and subsequent computation of shape features discussed in Sec. III A for clean synthetic data. Before providing the data to the ANN for training, we perform several additional scaling transformations. First, we utilize the (base 10) logarithm of Y_{DD} , and \bar{Y} . Second, we standard-normal scale the inputs and outputs by subtracting the mean and dividing by the standard deviation (both computed over the training data). Note that the experimental observations are also scaled according to the mean and standard deviation of the corresponding outputs in the training data for inference. Log-transforming ensures values spanning multiple orders of magnitude do not result in swamping of the training loss function by only the largest values.




FIG. 6. Mean Square Error as a function of training epoch for training and validation data showing that our gradient descent optimization of neural network weights has converged. The inset shows MSE on training sets of different sizes showing that we need $\mathcal{O}(10k)$ samples to cover the training domain well enough for our ANN model to obtain good generalization error.

Standard-normal scaling results in inputs and outputs spanning the same range and is also useful in ensuring no particular inputs or outputs dominate the loss function used to train the network.

In early iterations of this work, we found the neural network architectures that were used had difficulty fitting the training data despite having more than enough parameters to overfit by naive model complexity arguments. Inspired by the ResNet architecture that provides information from shallower layers of the network to deeper layers through skip connections, which can help to avoid vanishing gradients³⁹, we chose an architecture that feeds the input layer forward and concatenates it with the output of each successive hidden layer before feeding into the next layer. Our network has 5 hidden layers, each with 32 neurons and a ReLU activation function followed by a final linear output layer with 10 neurons for our 10 output features for a total of about 5.5k fit parameters. The model architecture is summarized in Fig. 5. The model is constructed in Python using the Keras package. The network weights are initialized with a Glorot uniform initialization⁴⁰ and the model is trained using a MSE loss function and an Adam optimizer⁴¹ for 1k epochs with a batch size of 128 for a total of about 500k gradient descent iterations. Figure 6 shows the MSE vs gradient descent epoch using the full 62k training data points, while the inset shows the final MSE after training versus training set size. The MSE evaluated on the training and validation datasets are nearly identical indicating that we are not overfitting, which would present as low MSE on the training data and high MSE on validation data as

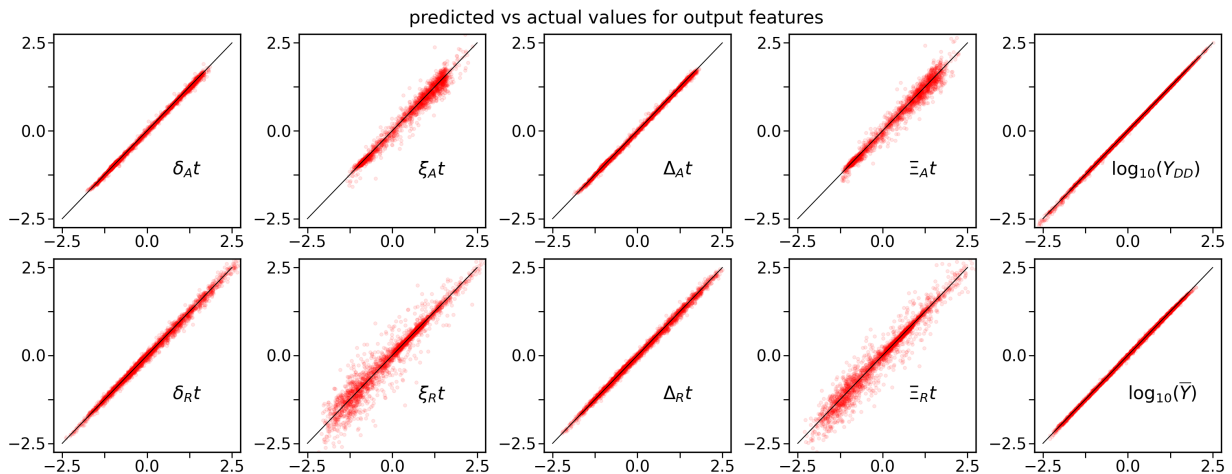


FIG. 7. Predicted (x-axis) vs actual (y-axis) features demonstrating that the model is doing a fair job at making predictions. The spread, as well as correlations in error among different outputs, are used to compute an error approximation used for the Bayesian inference as described in the text. Note that all values are roughly contained in the interval $[-3, 3]$ as all of the quantities have been shifted to have mean zero and are divided by the standard deviation of the training data, so that the interval corresponds to $\sim \pm 3\sigma$.

seen in the inset when using a small subset of training data to train the network. The good generalization of the network to unseen data as demonstrated by the close agreement between MSE on training and validation data also gives confidence that the training data sufficiently cover the sampling domain chosen.

When considering network performance, it is often useful to view a comparison of the network predictions to actual value. Figure 7 makes such a comparison on the validation data not used to train the network with predictions on the x-axis and actual values on the y-axis (all quantities standard-normal scaled) showing very good agreement. There is however a slight spread about the line $y = x$ that corresponds to perfect prediction. The remaining spread is likely due to a combination of statistical noise coming from variance in the numerical simulation as well as the surrogate model failing to capture meaningful variance in the training data due *e.g.* to under fitting. However, we did not find that increasing the model capacity improved the performance. In the next section on the Bayesian analysis procedure, we will model this spread by assuming that the discrepancy between the (vector) of predicted and actual values follows a multivariate zero-mean normal distribution with a covariance matrix that may be approximated from the validation data.

Figure 7 also demonstrates that there is some heteroskedasticity in the discrepancy (*i.e.* non-uniformity in the spread of predicted vs actual values for output features). It turns out that the heteroskedasticity in the asymmetry feature discrepancy is correlated with the input value of $\log_{10}(BR)$. Particularly, there is a greater spread in the discrepancy for lower values of $\log_{10}(BR)$. This is likely be a result of using the same maximum number of test particles in the FOM with weaker contri-

bution of test particles to the spectra and hence slightly lower statistics in the case of low $\log_{10}(BR)$. While it is in principle possible to attempt to model this heteroskedasticity, the zero-mean multivariate normal approximation using the covariance matrix calculated from the validation data should be a reasonable first approximation for the error in our predictions, and will overestimate the error in our ANN.

C. Bayesian Inference Model

All that remains in order to complete the description of our analysis pipeline is a discussion of the Bayesian inference model. It is our hope that this section will clearly indicate the various assumptions that have been made, making it obvious precisely which sources of uncertainty we capture, and under what conditions those uncertainties are accurately represented. We begin by detailing the full mathematical decomposition of the object of interest, the posterior distribution of our model parameters, into a likelihood and prior function (up to an unimportant normalization).

Bayesian inference gives a route to UQ of model parameters via providing a probability distribution, called the posterior distribution, which is weighted towards parameters that are the most consistent with observed data. If we let $\mathbf{x} = (\log_{10}(BR), \log_{10}(\rho R), T, \log_{10}(\mathcal{A}), F)$ denote the vector of model parameters and $\mathbf{y} = (\delta_{At}, \xi_{At}, \Delta_{At}, \Xi_{At}, \delta_{Rt}, \xi_{Rt}, \Delta_{Rt}, \Xi_{Rt})$ denote our observations, then the posterior distribution may be written as $p(\mathbf{x}|\mathbf{y})$. Normally one would then immediately rewrite this using Bayes' theorem as $p(\mathbf{x}|\mathbf{y}) = p(\mathbf{y}|\mathbf{x})p(\mathbf{x})/p(\mathbf{y})$ where the first term on the righthand side is known as the likelihood function, and

encodes assumptions about how the observed data are distributed about the model. The denominator $p(\mathbf{y})$ is known as the evidence and is often ignored, as it is constant with respect to parameters \mathbf{x} we wish to infer. The factor $p(\mathbf{x})$ is the prior distribution that encodes any prior knowledge or hard constraints with which our model parameters should be endowed. However, we do not start with such a decomposition. Instead, we first introduce an unobserved latent random variable which we shall denote \mathbf{z} , which will serve to propagate the variance of the first principles physics model about the predicted values from our ANN surrogate model into our final inference. Using standard rules of probability, we may thus write

$$p(\mathbf{x}|\mathbf{y}) = \int p(\mathbf{x}, \mathbf{z}|\mathbf{y})d\mathbf{z}. \quad (18)$$

In other words, we will work with the distribution $p(\mathbf{x}, \mathbf{z}|\mathbf{y})$ from which we can obtain $p(\mathbf{x}|\mathbf{y})$ by marginalizing out \mathbf{z} . Using Bayes' Theorem and standard manipulations of probability theory we then have

$$p(\mathbf{x}, \mathbf{z}|\mathbf{y}) \propto p(\mathbf{y}|\mathbf{x}, \mathbf{z})p(\mathbf{z}|\mathbf{x})p(\mathbf{x}) \quad (19)$$

where we have ignored the evidence term since it serves as an unimportant normalization when performing Markov Chain Monte Carlo sampling of the posterior. As the prior on our model parameters, the last term on the right hand side of Eq. 19 is taken to split into a product of terms, one for each of our model parameters.

The second term on the right hand side of Eq. 19 describes the distribution of the latent parameter (*i.e.* the first principles physics model) given our inputs. As discussed in the previous subsection, the discrepancy between the ANN and first principles model for our output features can be reasonably approximated as a zero-mean multivariate normal with the covariance matrix estimated from the data. In other words, we may write

$$p(\mathbf{z}|\mathbf{x}) \sim \mathcal{N}(f_{NN}(\mathbf{x}), \Lambda_{NN}), \quad (20)$$

where $f_{NN}(\mathbf{x})$ indicates evaluation of our neural network surrogate at the input parameters \mathbf{x} , and Λ_{NN} is the estimate of the covariance matrix using the validation data described in the previous subsection. We pause briefly to note that it would in principle be possible to treat the weights of the ANN surrogate in a fully Bayesian fashion in order to propagate the uncertainty from our surrogate model fit, however, such methods generally require significant additional computation time, and provided the high quality of our surrogate model are not likely to be a significant source of error in our model parameter inferences.

The first term in Eq. 19 describes how observed data are distributed about the latent variable used as a stand-in for our first principles physics model that, to reiterate, is never actually evaluated in the Bayesian inference, but is simply approximated as a normal random variable distributed about the ANN surrogate model prediction. We

assume that the data are distributed about this latent variable in the following fashion. First, we split each of the observations and latent variables into three separate pieces, one containing all of the nToF shape information, one containing the DD yield information, and one containing the DT to DD yield ratio information. That is, we represent \mathbf{y} by $\mathbf{y} = (\mathbf{y}_{\text{nToF}}, y_Y, y_{\overline{Y}})$ and similarly for \mathbf{z} . We then assume that the following holds

$$\begin{aligned} p(\mathbf{y}|\mathbf{x}, \mathbf{z}) &= p(\mathbf{y}|\mathbf{z}) \\ &= p(\mathbf{y}_{\text{nToF}}|\mathbf{z}_{\text{nToF}})p(y_Y|z_Y)p(y_{\overline{Y}}|z_{\overline{Y}}). \end{aligned} \quad (21)$$

The first line in the above equation simply expresses that once we have computed the model value (with added variance from using a surrogate), we no longer have any direct dependence of our observations on the model parameters. The second line of the equation is an expression of independence of the different measurements. Finally, we assume the following probabilistic models are appropriate

$$p(\mathbf{y}_{\text{nToF}}|\mathbf{z}_{\text{nToF}}) \sim \mathcal{N}(\mathbf{z}_{\text{nToF}}, \Lambda_{\text{nToF}}), \quad (22)$$

$$p(y_Y|z_Y) \sim \mathcal{N}(z_{\text{nToF}}, \Lambda_Y), \quad (23)$$

$$p(y_{\overline{Y}}|z_{\overline{Y}}) \sim \mathcal{N}(z_{\overline{Y}}, \Lambda_{\overline{Y}}), \quad (24)$$

where evaluation of covariance for the nToF features was discussed in Sec. III A, and the variance for the yield features are provided by standard analysis techniques¹⁸. Note that the uncertainty in nToF shape measurements is not assumed to be uncorrelated. The reason for this is that when we estimate those features from the data, we subtract different estimations of the background, which will introduce correlations into the measured values for a given view (axial or radial) as captured in the covariance matrix computed for those features discussed in the data preprocessing subsection. Notice however, that there is no reason to expect the errors in the radial and axial nToF features to be correlated as their backgrounds are fit independently. This results in an approximately block structure of the covariance matrix Λ_{nToF} extracted for the observed nToF shape features. In general, we observe extremely weak (near-zero) values in the components of the covariance matrix corresponding to correlations in error between the radial and axial channels.

IV. VALIDATION OF ANALYSIS FRAMEWORK

This section discusses results from applying the deep learning based Bayesian inference to both synthetic and experimental data. Application to synthetic data will allow us to confirm that the learned model and inferences conform to the physics imposed by the first principles model, at least for cases contained within the domain of the training data. Furthermore, by applying the method to synthetic data, which has noise added to mimic experimental observations, we gain confidence that the approach can be reasonably expected to generalize to experimental data. Consistency with a previous experimental

analysis further supports the use of our ANN surrogate to replace the FOM enabling Bayesian inference for UQ.

A. Testing on Synthetic Data

We begin by considering an inversion of the clean synthetic test data (no added noise) produced as discussed in Sec. III B. We emphasize that this test data was not used either to train the ANN surrogate nor to estimate covariance between surrogate and the first principles predictions as discussed in the previous section. We focus on highlighting consistency with the FOM. Unless otherwise noted, all priors are taken to be uniform over the sampling domain defined in Sec. III B to ensure results are not biased by providing additional information to the inversion procedure. The upper panel of Fig. 8 shows that the surrogate prediction with uncertainty for \bar{Y} as a function of $\log_{10}(BR)$ agrees well with the first principles calculation for two distinct values of fuel areal density, $\rho R \approx 32 \text{ mg/cm}^2$, at the upper ρR boundary of our training data, as well as $\rho R \approx 2 \text{ mg/cm}^2$ near the lower limit of our training data at $\rho R \approx 1 \text{ mg/cm}^2$. Other input parameters were held fixed at $T = 3 \text{ keV}$, $\mathcal{A} = 80$, and $F = 0.0$. Note that the surrogate model is able to well reproduce the trends of the FOM for these two cases, with the higher value of ρR demonstrating significantly less sensitivity of \bar{Y} to BR as expected¹⁹.

In the lower panel of Fig. 8, we see the posterior distributions of $\log_{10}(BR)$ inferred from the clean synthetic test for 5 different true values of $\log_{10}(BR)$. For the inferences, output features are computed from clean synthetic FOM nToF signals with known inputs, so that a well-defined truth value of the input parameters exists. The likelihood function, however, requires a finite variance to fit the model, so we utilize a small fractional uncertainty of 1% on the FOM outputs, equivalent to assuming a 1% uncertainty on the output features (in standard-normal scaled units). Figure 8 shows that as one approaches the two extremes, that the value of \bar{Y} becomes insensitive to BR and as a result, the posterior distribution widens (greater spread along the vertical axis). In particular, for low BR , the low- BR tail of the posterior is heavier than the high- BR tail of the posterior distribution for the high BR case, where other features have not yet saturated to the degree that \bar{Y} has. More generally, for very low or high BR we expect the inference to provide only an upper or lower bound respectively. This is consistent with the discussion of sensitivity of yield and shape features of the FOM to BR discussed in Ref. 20. Physically, the range over which \bar{Y} and nToF signals are sensitive to BR is the same range over which $R/r_{L,t} \gtrsim 1$, where the fraction of magnetically trapped tritons increases and eventually saturates, thereby losing sensitivity to BR .

Now we consider the inversion of a single synthetic data set ($\log_{10}(BR) \approx 5.34$, $\log_{10}(\rho R) \approx -1.79$, $T = 4.9 \text{ keV}$, $\log_{10}(\mathcal{A}) = 1.9$, and $F \approx 0.1$) treated in two different ways to assess the impact of noisy nToF signals




FIG. 8. Upper Panel: The surrogate model agrees quantitatively with the first principles calculations under the normal approximation of the surrogate model uncertainty (gray bands), with the FOM output \bar{Y} (black curves) being given as an example. The two curves are for fixed values $T = 3 \text{ keV}$, $\mathcal{A} = 80$, and $F = 0.0$, with the upper curve $\rho R \approx 32 \text{ mg/cm}^2$, which is the upper bound on areal density in our training data, and the lower curve with $\rho R \approx 2 \text{ mg/cm}^2$ lying fairly close the lower bound of our training data of $\rho R = 1 \text{ mg/cm}^2$. We can see the surrogate reproduces the expected effects of increasing the areal density of the fuel. Lower Panel: Inferred BR posterior distributions for values of BR shown in corresponding colored dots in the upper panel with $\rho R \approx 2 \text{ mg/cm}^2$, $T = 3 \text{ keV}$, $\mathcal{A} = 80$, and $F = 0.0$. For this violin plot, the width as a function of the vertical axis is proportional to the probability density associated with the vertical axis value. In this case, $\log_{10}(BR)$. The distributions are ordered along the x-axis according to the known true value of $\log(BR)$ so that the posterior distribution should overlap strongly along the vertical axis with this line. We show the 50% percentile point of each distribution in white, demonstrating good agreement with the true value. Also notice that the middle three cases also have narrow distributions, while the first and last case show wider distributions indicating a loss of sensitivity. Note that a relatively small number of outliers from the MCMC sampling cause the narrow tail towards low values of $\log_{10}(BR)$ for the posterior with true value of $\log_{10}(BR) = 5.75$, and that tail is not meaningful.

on our inference. Uniform priors were used for all the input parameters except T , which uses a standard normal prior with mean value centered on the true value and assumed 20% error. We use this prior as it is also how we analyze real experimental data. For experimental data, the temperature and its uncertainty used in this prior are extracted through an independent analysis

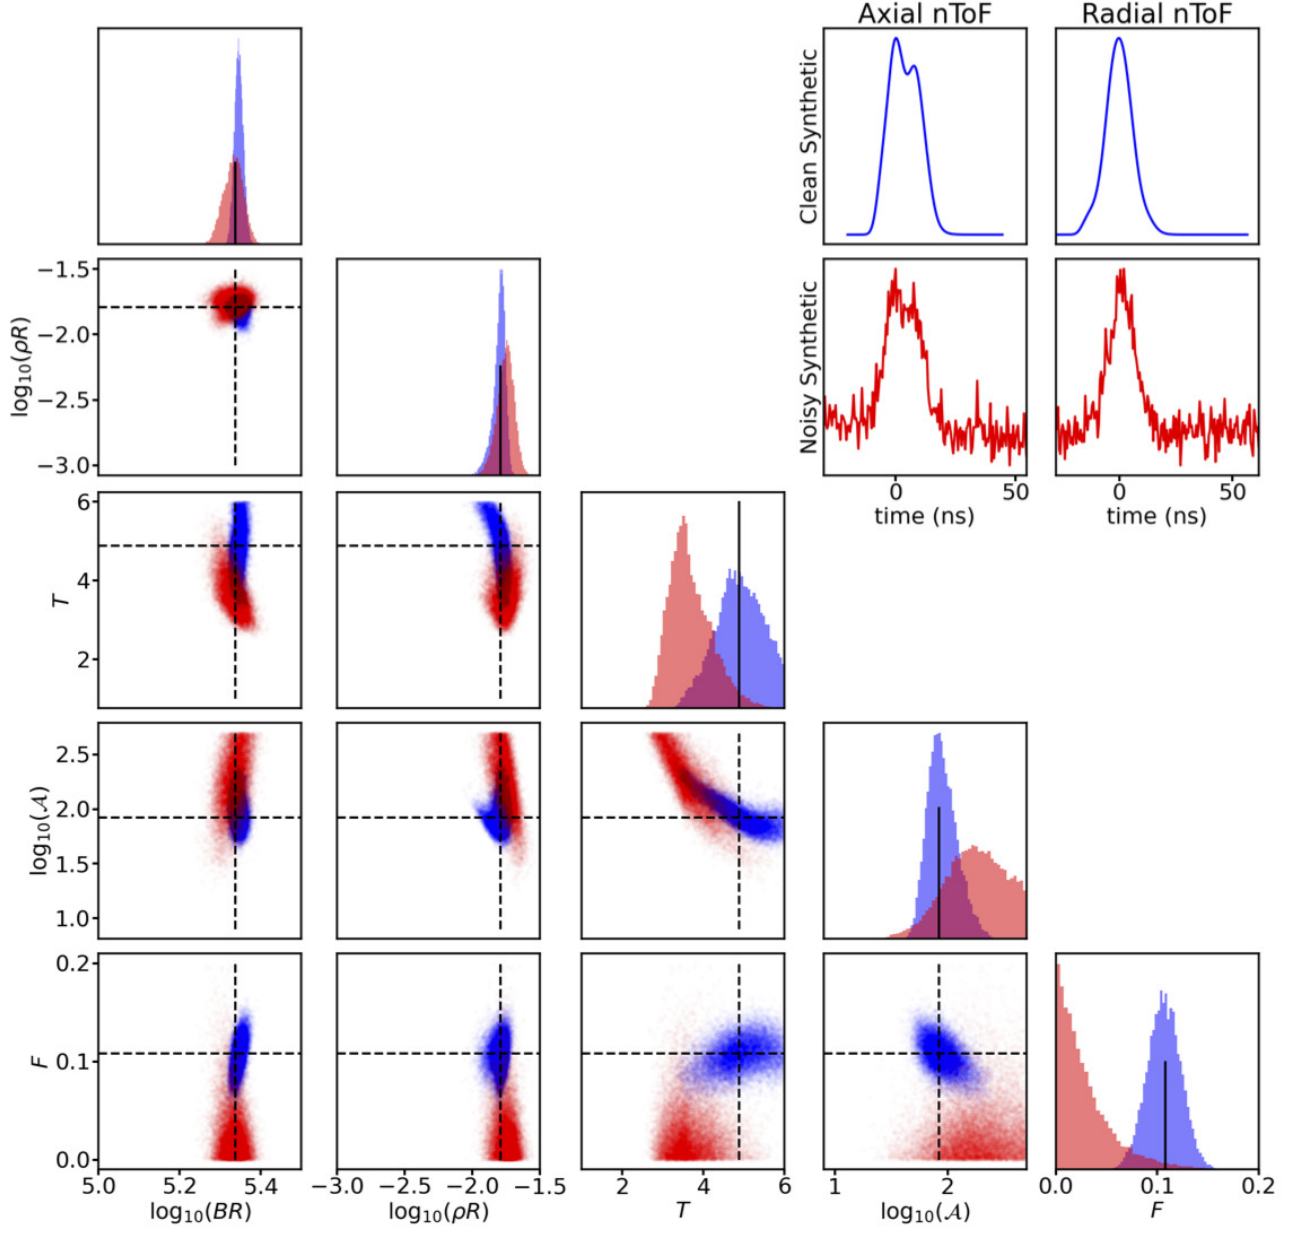


FIG. 9. Plot showing pairs of parameters from MCMC samples on clean (blue) and noisy (red) synthetic test data described in the text. Notice that the clean data set allows the inference of all parameters with reasonable accuracy, while adding noise allows only for accurate inference of the parameters that have the greatest impact on the observables. Histograms on the diagonal are the single parameter marginal posterior distributions corresponding to the parameter labeled on the bottom row beneath the histogram. The inset shows the clean and noisy versions of the synthetic nToF signals used in the analysis.

of the primary DD nToF peak. In practice, however, we find that this choice of prior has little influence on the inference of $\log_{10}(BR)$ as compared to a uniform prior. Since the posterior distribution dimension is too large to visualize, we present a corner- or pair-plot showing the pairwise joint distributions of MCMC samples. These

plots integrate along all but two of the parameter axes, showing how the remaining two parameters can vary together while still achieving a good fit to the observed data within uncertainties. Figure 9 the MCMC samples using the clean synthetic outputs from the FOM as observations (with 1% fractional uncertainty) in blue demon-

strating that the posterior distributions are well centered on the true physical parameters (black lines) in this case. For the second, noisy synthetic case, MCMC samples are shown in red in Fig. 9. To create noisy synthetics, we first add random noise and background to the synthetic nToF signals, characteristic of experiment. Particularly, we add an exponential (linear) background and additive gaussian noise to the synthetic axial (radial) synthetics, with parameters of these noise models selected at random so that the noisy synthetic is representative of experimental data. For the outputs Y_{DD} and \bar{Y} , we randomly bias the clean synthetic values by a fractional amount chose at random from $\mathcal{N}(0, 0.01)$ for each, and assume a statistical uncertainty about that value of 20%, typical of our experimental measurements. We then apply our full analysis pipeline and observe that the inference gives good agreement of $\log_{10}(BR)$ and $\log_{10}(\rho R)$ with their true values, while the other three parameters show significant bias. This makes physical sense as Ref. 20 shows that the dominant sensitivities of features such as \bar{Y} and the width features are to $\log_{10}(BR)$ and $\log_{10}(\rho R)$, so that one might expect noise to obfuscate the information that informs the inference of inputs resulting in only a weak variance in the observations. It is also important to point out that we have observed a handful of cases for which the inferred value for $\log_{10}(\rho R)$ was significantly biased towards larger values for the noisy synthetic data. This appears to be more significant for small values of $\log_{10}(BR)$, though additional analysis would be required to better understand this effect. As a result, we do not try to infer $\log_{10}(\rho R)$ using our approach. Finally, we note that upon investigating of a handful of examples, the case shown in Fig. 9 is representative of performance for inferring $\log_{10}(BR)$, except for when $\log_{10}(BR) \lesssim 4.75$ where in both the clean and noisy synthetic cases, one expects only to obtain an upper bound on $\log_{10}(BR)$. In that scenario, for the noisy case, the inference is biased further towards smaller values of $\log_{10}(BR)$. We take this as a good indication that we can safely treat parameters besides BR as nuisance parameters, whose values may be biased away from their correct physical value. Particularly, we expect that, integrating out all other parameters except BR will not significantly impact our inference of BR . This is important, as extracting BR will be the purpose of our analysis in Sec. V.

B. Comparison to Previous Work

We now turn to applying the method to data taken on MagLIF experiments. We will begin with a reanalysis of shot z2591 that was studied in Refs. 19 and 20. This early MagLIF shot was one of the first to give significant DD yields, $Y_{DD} \sim 2 \times 10^{12}$ along with a reasonable secondary DT nToF signal. Note again that for all experiment analysis, we take the fuel temperature to follow a normal distribution according to temperature measurements and uncertainties obtained using primary

DD neutron nToF signals. To minimize the possibility of introducing inconsistencies with the FOM, we do not apply any other informative priors. As previously indicated, however, we found the informative temperature prior had little impact on the overall results compared to a uniform prior. Figure 10 shows that the two separate analyses are consistent to within error bars, with this work demonstrating significantly tighter constraints on several of the inferred quantities. By using Bayesian inversion, we obtain full distributions for the input parameters, allowing us to rigorously define the 2.5 – 97.5% (2σ -like) confidence interval rather than direct propagation of errors on input parameters obtained via independent analysis as done in Refs. 19 and 20. Furthermore, we fit multiple shape features of the full nToF signal rather than a single ratio of the HWHMs of the high energy portion of the measured radial and axial energy spectra (computed by inverting the nToF signal) and include the primary DD neutron yield as an output fit by the model.

We can also gain insight into how well our method describes the experimental observations by comparing the range of features predicted by the model during the inference as shown in Fig 11. This figure demonstrates that the posterior parameters give rise to a narrower and less asymmetric shape of the axial nToF signal than is experimentally observed. This should not be too surprising when considering how significant the asymmetry of the axial spectrum for this shot is as seen in Fig. 18 of Ref. 20. The tendency for the model to bias towards lower than observed axial nToF width and asymmetry is typical to some degree of all the shots analyzed in this work, with generally better agreement occurring between the model and radial nToF shape features. This may arise from a variety of unmodeled physical effects. For example, scattering effects are expected to have a weaker impact on the radial detector due to the radial line of sight having better collimation and a better scattering environment, which may explain why the agreement with the axial shape features is generally not as good as for the radial nToF features. It is also possible that unaccounted for spatial inhomogeneities is fuel plasma conditions, or non-trivial magnetic field topology⁴² may contribute to deviation of the nToF signals from the ideal conditions laid out in the FOM. We hope to better address this in the future by incorporating some of these effects in the computation of synthetic nToF signals, but this goes beyond the scope of the current work. One will also note that in comparison to the analysis in Refs. 19 and 20, by including the DD yield as a quantity to fit, we are not able to obtain as close of an agreement with the yield ratio \bar{Y} (see Fig. 11 caption for a possible explanation). In addition, the uncertainty in the experimentally observed \bar{Y} is significantly larger (by about $2\times$) than is typical of our experiments. This is due to specifics with the way the activation sample was fielded, and may have also resulted in an unknown systematic bias in \bar{Y} . Note, however, that explicitly fitting to the DD yield helps to constrain ρR in a fashion consistent with the physics of the FOM.




FIG. 10. Demonstration that the analysis carried out in this work (red) is consistent with Ref. 20 (blue) for all parameters. Box plots for this work span 25 – 75% while whiskers (capped lines extending beyond the boxes) span 2.5 – 97.5% ($\sim \pm 2\sigma$), with the 50% shown as a red line. The uncertainty for the analysis in Ref. 20 (blue) does not correspond precisely to a 2σ definition, as they are extracted utilizing a different (non-Bayesian) procedure. We also note that Ref. 20 takes advantage of independent analysis methods to constrain all of the input parameters except BR , while we only use independent analysis of the DD neutron nToF peak to set a prior on the temperature T , making the overall agreement even more impressive.

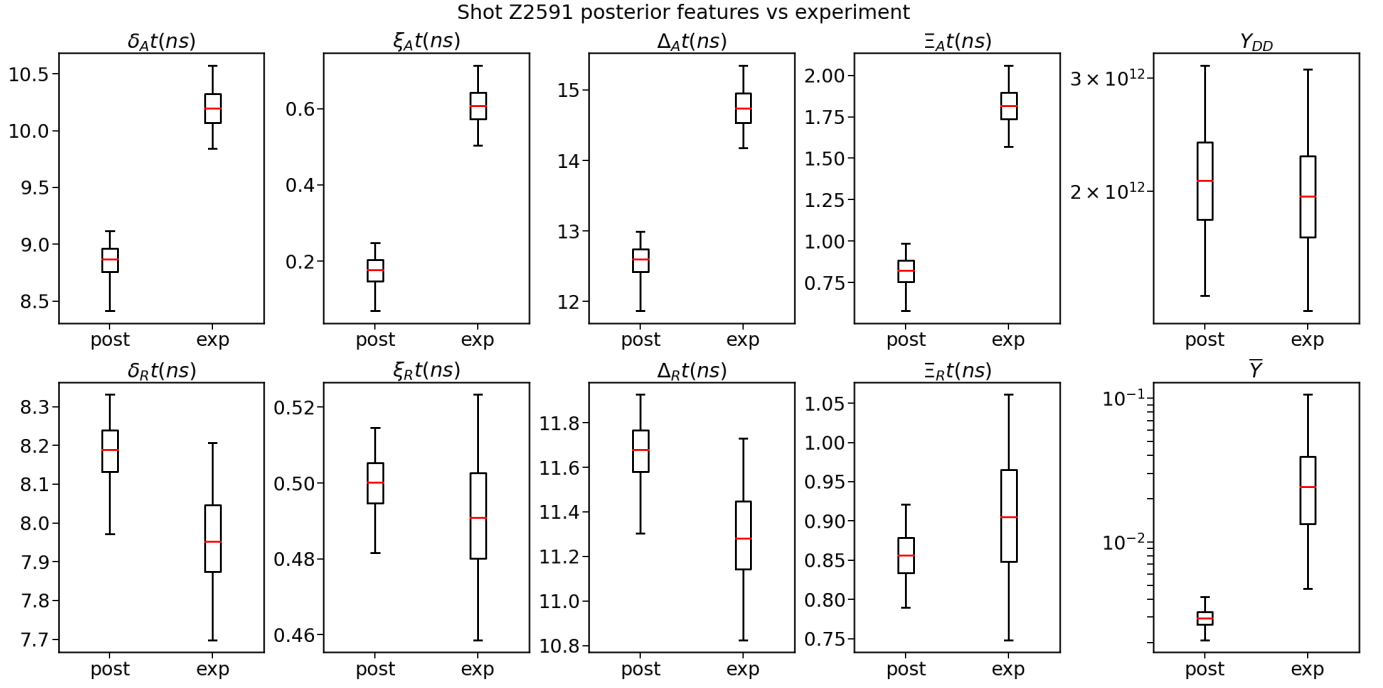


FIG. 11. A comparison of the posterior ANN outputs observations produced by our model that are most consistent with the experimental data. Notice that the axial nToF data are experimentally observed to have greater width and asymmetry than our model can predict when considering fitting additional data sources. This is consistent with fits of the energy spectra in Refs. 19 and 20. Note that the disagreement between the posterior and observed values of \bar{Y} , may be partially explained by appealing to the fact that due to specifics with the way the activation sample was fielded there was an about $2\times$ larger statistical uncertainty and potential systematic bias in \bar{Y} . Similarly fielded experiments have exhibited secondary DT neutron yields as low as $\sim 20\%$ of z2591, with values of Y_{DD} spanning $\sim 0.5 - 1.5\times$ the value for this shot. This indicates the possibility that the true value of \bar{Y} is significantly smaller than the best understood and reported value given here.

While the two approaches clearly show differences due to these discrepancies, the overall consistency of the two approaches, especially in regards to the inference of BR , encourages further application of our method to experiments in the search for interesting physical effects.

V. ASSESSING MAGNETIC CONFINEMENT IN MAGLIF PREHEAT EXPERIMENTS

We now present a new analysis, which constitutes the first time that the efficacy of magnetic confinement in a neutron-producing MIF platform has ever been systematically assessed as a function of fundamental inputs, in this case, laser preheat energy deposited in the fuel. We consider four experiments with all experimental condi-

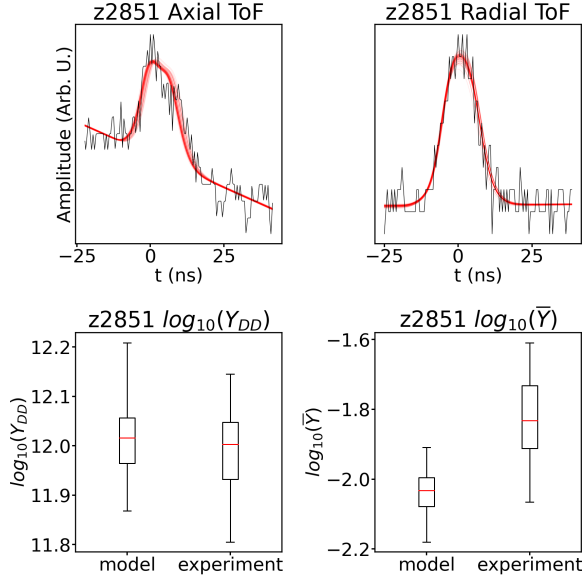


FIG. 12. 100 posterior samples are fed through the FOM to compute Y_{DD} , \bar{Y} , and (normalized) synthetic nToF signals. Each of the 100 nToF signals are fit to the observed signals with an amplitude, time shift, and background. Plotting all of the posterior FOM calculations shows reasonable agreement with experiment z2851. The axial nToF signal and \bar{Y} show the greatest discrepancy.

tions except total laser preheat energy held nominally identical. All the experiments used beryllium liners with an aspect ratio of 6 (AR6, where $AR = r_o/\Delta r$, with r_o the outer liner radius, $\Delta r = r_o - r_i$, and r_i the inner liner radius), initial deuterium fill density of 0.68 mg/cc, peak current of 16 MA, initial magnetic field of 10 T, and systematically varied preheat energies covering a range from about 0.5 – 1.5 kJ of energy coupled to the target. The reader is referred to Ref. 8 for additional details of experiment configuration.

Application of our ANN-based Bayesian inference obtains generally good agreement between experimental observations and features predicted by the model as seen for the lowest and highest coupled preheat in Figs. 12 and 13 respectively. For these figures, we collect 100 random MCMC samples of the input parameters from the posterior distribution to run through the FOM. We then compare the resulting nToF signal to the experiment by fitting background, amplitude, and a time shift. Our ability to obtain good agreement with the observations paired with the discussion of Sec. IV gives us confidence in our ability to obtain results for BR consistent with what would be obtained via direct application of the FOM.

Fig. 14 shows the systematic inferences of BR vs preheat for each of the four shots, demonstrating that the fuel magnetization decreases dramatically from the lowest to next highest preheat, after which the values level off to some extent. The preheat energy values and uncertain-




FIG. 13. 100 posterior samples are fed through the FOM to compute Y_{DD} , \bar{Y} , and (normalized) synthetic nToF signals. Each of the 100 nToF signals are fit to the observed signals with an amplitude, time shift, and background. Plotting all of the posterior FOM calculations shows excellent agreement with experiment z2839. This performance is also typical of z2977 and z3040 not shown to save space.

ties reported here are determined using energy measured at the output of the laser and deposited energy fractions inferred from surrogate off-line experiments utilizing similar target configurations and laser protocols. Energies for shots with an unconditioned laser beam were derived from⁴³, and energies for shots that used a distributed phase plate to smooth the laser beam profile were derived from⁴⁴. In order to understand the observed trend in BR , we ran detailed 2D LASNEX post-shot simulations of these experiments both with and without the Nernst effect to compare with the ANN-based inferences. The fuel magnetic confinement parameter is extracted from the simulations by directly computing $BR = \Phi_z/(\pi R)$ from the magnetic field profile at peak burn. The dashed curves and grey bands are obtained by taking the the mean and standard deviation of BR at a small number of simulations that vary simulation conditions to within experimental uncertainty at the different preheat values (initial B-field $\pm 5\%$, initial fill density $\pm 5\%$, and peak current delivered $\pm 10\%$). The uncertainty bands demonstrate that sensitivity of the simulation to the variation in initial conditions one might expect in experiments are not sufficient to swamp the impact that the Nernst effect has on determining the value of BR at peak burn. Furthermore, we note that the uncertainties in BR inferred from the experiments do not depend explicitly in any way on the uncertainty in the initial conditions of the experiment, but rather on uncertainties in the observed experimental outputs Y_{DD} , Y_{DT} , and the nToF shape fea-

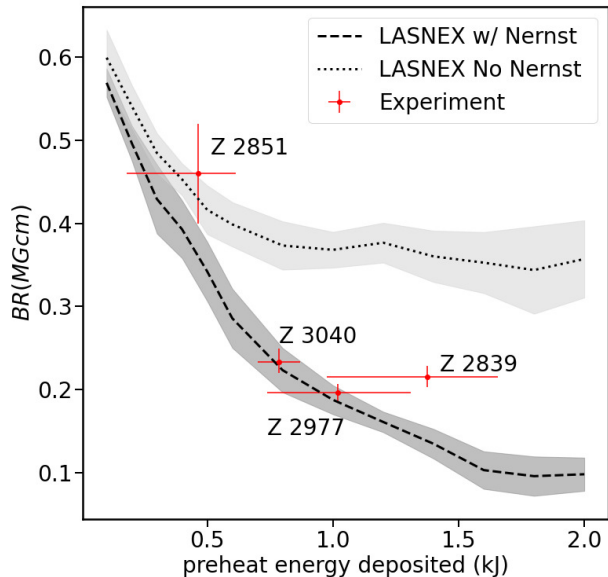


FIG. 14. Data points show the inferred value of BR vs preheat energy for the four shots in this series. The sharp decrease between the lowest and next-highest preheat values is consistent with expectations from Nernst physics and magnetic diffusion into the liner. The dashed lines and shaded bars show the mean and standard deviation of BR extracted from a small number of 2D LASNEX simulations that vary input parameters within experimental uncertainty both with and without the Nernst advection term as described in the text.

ture uncertainties extracted as described in Sec. III A. It is also interesting to note that the error bars on the inferred values of BR are significantly smaller for the three experiments with highest laser preheat energy than that for the lowest. We remind the reader that, as indicated by Figs. 8 and 9 and related discussion, the error bar scale will depend sensitively on both where in parameter space the inferred values lie, as well as the scale of signal to noise among the different observed features.

The Nernst effect is believed to be one of the strongest physical mechanisms that exacerbate flux losses from the preheated fuel beyond simple resistive MHD predictions. Physically, the Nernst effect arises from the suprathermal, heat-conducting electrons transporting magnetic field along temperature gradients perpendicular to the background magnetic field orientation. The characteristic velocity at which this occurs is related to electron temperature gradients and perpendicular magnetic fields by $v_N = \beta_\lambda \nabla_\perp T_e / (eB)$, where β_λ is the so-called Nernst–Ettingshausen coefficient²⁷, which will itself have some dependence on B , T_e , and the electron and ion densities. In our experiments, axial orientation of the applied magnetic field and the predominantly radial temperature gradients established by the wall confinement of the cylindrical, preheated plasma leads to the expect-

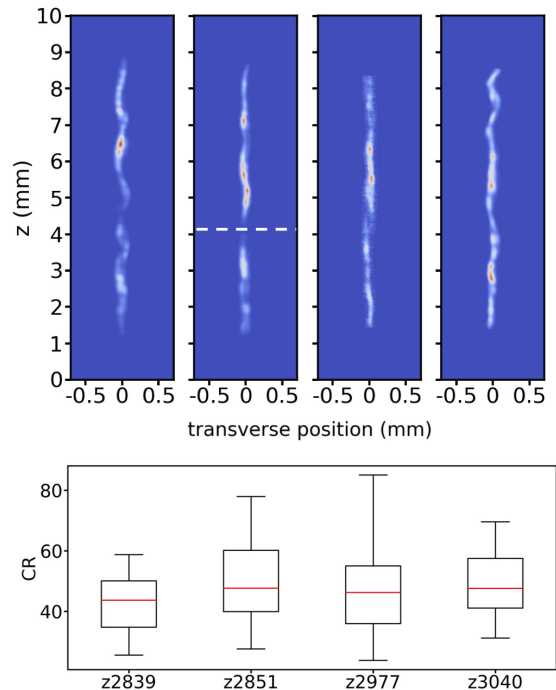


FIG. 15. Top: stagnation x-ray self-emission profiles of the four experiments, with a partially obscuring fiducial field on z2851 indicated by white dashed line. Bottom: Estimates of convergence ratio from the measured fuel self-emission profiles, with statistics provided by the axial variation in the estimate.

tation that the axial magnetic field will be advected out of the hot fuel by electrons opposite the direction of the temperature gradient towards the target wall, where it may then resistively diffuse into the liner. As the preheat energy delivered to the target is increased, these radial temperature gradients will become stronger. Naively then, one would expect an amplification of the magnitude of Nernst flux advection toward the cold walls surrounding the fuel as preheat energy is increased. This simple picture is consistent with the trend in both the experimental observations and LASNEX simulations including the Nernst effect.

Interestingly, the 2D LASNEX simulations give convergence ratios ($CR = r_{i,0}/r_{i,f}$) that are nearly independent of preheat and are approximately 42 ± 2 for both the Nernst and no-Nernst cases. As a result, the simulated values of BR are very nearly proportional to the ratio of final to initial flux

$$\frac{\Phi_{z,f}}{\Phi_{z,0}} = CR^{-1} \frac{(BR)_f}{(BR)_0} \approx \text{const.} \times (BR)_f. \quad (25)$$

For integrated MagLIF experiments, x-ray backlit liner radiography typically used to infer liner convergence ratios is unavailable due to the fact that the Z Beamlet Laser can be used to drive the backlighter foils⁴⁵ or the MagLIF laser preheat, but not both at the same time.

Shot Number	Laser Energy Coupled (kJ)	Y_{DD} ($\pm 20\%$)	Y_{DT} ($\pm 20\%$)	T_i ($\pm 20\%$) (keV)	BR (MG-cm)	CR
z2839	$1.375^{+0.280}_{-0.400}$	3.2×10^{12}	1.2×10^{10}	2.21	$0.215^{+0.013}_{-0.012}$	44^{+8}_{-9}
z2851	$0.465^{+0.150}_{-0.285}$	1.0×10^{12}	1.5×10^{10}	1.63	$0.460^{+0.060}_{-0.060}$	48^{+24}_{-13}
z2977	$1.02^{+0.290}_{-0.280}$	3.0×10^{12}	1.4×10^{10}	2.63	$0.196^{+0.011}_{-0.010}$	47^{+17}_{-12}
z3040	$0.785^{+0.085}_{-0.085}$	4.1×10^{12}	2.2×10^{10}	2.56	$0.233^{+0.016}_{-0.013}$	47^{+17}_{-13}

TABLE I. Summary of the laser preheat energy deposited, DD and DT yields, ion temperature, inferred value of BR , and estimated convergence ratios for the four shots in this experiment series. Note that uncertainties for all quantities except BR and CR are provided by experimental analyses not discussed in this manuscript, while error bars on BR and CR are obtained by determining the 16 – 84% ($\sim 1\sigma$) interval as described in the text.

However, we are able to obtain estimates of the fuel profile radius as a function of height from stagnation images, which can be used to estimate CR as shown in Fig. 15. We note that the profiles exhibit values of CR that are both consistent with the LASNEX simulation and appear to be independent of preheat. Unfortunately, the variation in the estimate of CR over the axial length of the fuel column is too high to use in a computation of $\Phi_{z,f}/\Phi_{z,0}$, so that we refrain from concluding anything about the magnetic flux loss specifically. For convenience, we have included Tab. I summarizing the relevant experimental observations, BR inference, and CR estimate according to the above procedure. As a concluding remark, we point out that, although resistive diffusion in LASNEX simulations without Nernst is sufficient to result in a slight decay in BR as preheat begins to increase, the fuel magnetization quickly plateaus at a significantly larger value than experimentally observed. The agreement of our inferences with the LASNEX simulations including Nernst suggests a significant role for this mechanism in determining fuel magnetization as well as magnetic flux compression and hence for the physics of α -particle confinement in future 50-50 DT MagLIF experiments.

VI. CONCLUSION AND FUTURE WORK

We have demonstrated a deep learning based surrogate model coupled with a Bayesian inference to replace a significantly slower physics model of nuclear diagnostics and provide uncertainty quantification for BR inferences. This framework enabled a landmark study of fuel magnetization trends in the neutron-producing magneto-inertial fusion platform MagLIF. Particularly, we found evidence that fuel magnetization decreases from $BR \sim 0.5$ MG-cm to $BR \sim 0.2$ MG-cm as laser preheat energy deposited in the fuel increases from $E_{\text{preheat}} \sim 460$ J to $E_{\text{preheat}} \sim 1.4$ kJ for otherwise nominally identical experimental conditions. Interestingly, our results are consistent with detailed 2D LASNEX post-shot simulations demonstrating flux loss due to the Nernst effect. Unfortunately, poorly constrained convergence ratios in the experiments preclude more direct statements about flux loss in experiments. Instead, we rely on inferences

of BR to support this claim.

While our results encourage the interpretation of Nernst as a critical mechanism for flux loss, it will be important to revisit the assumptions of our FOM and the simulations presented in this manuscript as better models become available. For convenience, we summarize those assumptions here, reminding the reader of their justification or need for future work where appropriate:

- **Thermally equilibrated electrons and ions:** The FOM assumes $T_e = T_i$. Refs. 4 and 17 estimate these values from x-ray spectra and DD nToF peak shape respectively, showing only minor discrepancies when accounting for experimentally measured uncertainties. In addition, with an estimated electron-ion collision time $\tau_{e-i} \ll 1$ ps that is much faster than the stagnation duration ~ 2 ns, the fuel plasma is collisional at the few keV temperature and solid density relevant to our experiments. We therefore find this assumption to be reasonably consistent with experimental observations.
- **2 ns Gaussian burn history:** We assume a Gaussian burn history with 2 ns full width half max (FWHM). Although we cannot directly measure the burn history, photo-conducting diamond (PCD) measurements of the x-ray output as a function of time indicate a 1.7 – 2 ns burn duration providing support for this assumption.⁴ Furthermore, due to the relatively broad energy spectrum of neutrons produced in secondary DT reactions, minor deviations from this assumption are not expected to significantly impact our results.
- **Dimensionality effects:** Our analysis assumes a purely radially varying fuel profile. X-ray diagnostics do not have the required spatial resolution to observe radial dependence in the $\sim 50\mu\text{m}$ radius stagnation column. As a result, the assumed functional form for radial profile cannot be further assessed by comparing to experiment at this time. However, the detailed study performed in Ref. 20 demonstrates that the FOM is relatively insensitive to the exact details of the radial profile. On the other hand, it is not difficult to infer from Fig. 15

that there is some extent to which the fuel conditions are axially varying. For example, previous analysis indicates variations in electron temperature on the order of $T_e = 3.0 - 3.2$ keV in the brighter (colder) regions and up to $T_e = 3.8$ keV in the dimmer regions for one particular experiment.¹⁷ Depending on the exact experimental configuration, axial mix profiles have also been observed to range from relatively uniform, with variations of 3 – 7% over the column, to cases with a few percent near the bottom of the target up to $\sim 15\%$ near the laser entrance window. Additional effects such as three-dimensional instability structure/stagnation morphology are also of significant interest. While incorporation of these effects is not currently possible, it may eventually become feasible to study the method on synthetic diagnostics produced from more complex two- or three-dimensional simulations that would help to reveal any systematic bias introduced by these assumptions in the FOM.

- **Magnetic field topology:** We assume a purely axial magnetic field. While this could technically fall under the previous topic, we find it worthwhile to address separately. Ref. 42 indicates that alternate magnetic field topologies can affect the asymmetry of the secondary neutron spectra, and may at least partially explain some of the residual discrepancy in the inferred axial nToF shape features in Fig. 11. Addressing this topic is of significant interest. While out of scope for this work, in principle it should be possible to apply our surrogate modeling approach to the model in Ref. 42, and will be addressed in followup work.
- **Use of nToF shape features:** Our surrogate model differs from previous application of the FOM in Refs. 19 and 20 by working directly with shape features of the measured nToF signals rather than shape features of spectra obtained by inverting that nToF data. In doing so, we avoid, among other things, the need to solve an ill-posed instrument response deconvolution problem and circumvent the need for an absolute timing fiducial for the nToF signal. Furthermore, we found that utilizing two distinct nested intervals characterizing the nToF width and asymmetry provided additional information by incorporating lower- and higher-energy neutrons contributing to the complex asymmetric spectra produced by the FOM (see discussion in Sec. III A). We found these features to be constraining when inverting synthetic datasets in Sec. IV A, even when noise was added to the synthetically produced observations. In future work including additional physical features, *e.g.* non-axial magnetic field topology, we will need to ensure that this featurization continues to provide sufficient information to constrain the inferred value of BR ,

and make adjustments as needed.

- **Comparison to two-dimensional simulations:** In Sec. V, we compared our BR inference to 2-D LASNEX simulations both with and without the Nernst advection term. While we do not currently have the capability able to assess BR with 3-D simulations, there has been recent progress in understanding magnetic field transport in 3-D (see Ref. 28). It would be interesting as our capabilities develop further to explore the impact of 3-D effects on fuel magnetization relative to the 2-D simulations presented here.

We are currently working on aggregating additional experimental data for analysis. Of particular interest are data sets where conditions such as the initial gas-fill density, current delivery, and initial magnetic field strength are systematically varied. The development and improvement of automated and data-driven analysis tools such as that demonstrated here will prove critical in providing fast and accurate assessment of high energy density experiment data with UQ. By enabling such tools, data science and machine learning methods will continue to aid in leveraging experimental data and simulation tools to the greatest possible extent, thereby contributing significantly to improving our understanding of the underlying physics in these systems. As we develop modeling capabilities and build surrogates from more complex models we will be able to lift some of the assumptions required in this work by capturing important physics such as fuel plasma inhomogeneities as well as scattering and geometry effects inherent in our experiments that may affect nToF signals. However, we believe that our detailed study of the deep-learned surrogate model and Bayesian inference procedure demonstrates the viability of this tool and encourages its use in assessing trends in magnetization/flux compression in other systematic experimental studies.

ACKNOWLEDGMENTS

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC (NTESS), a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy’s National Nuclear Security Administration (NNSA) under Contract No. 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.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon reasonable request.

- ¹S. A. Slutz, M. C. Herrmann, R. A. Vesey, A. B. Sefkow, D. B. Sinars, D. C. Rovang, K. J. Peterson, and M. E. Cuneo, *Phys. Plasmas* **17**, 056303 (2010).
- ²M. E. Savage, K. R. LeChien, M. R. Lopez, B. S. Stoltzfus, W. A. Stygar, D. S. Artery, J. A. Lott, and P. A. Corcoran, in *2011 IEEE Pulsed Power Conference* (2011) pp. 983–990.
- ³D. V. Rose, D. R. Welch, E. A. Madrid, C. L. Miller, R. E. Clark, W. A. Stygar, M. E. Savage, G. A. Rochau, J. E. Bailey, T. J. Nash, M. E. Sceiford, K. W. Struve, P. A. Corcoran, and B. A. Whitney, *Phys. Rev. ST Accel. Beams* **13**, 010402 (2010).
- ⁴M. R. Gomez, S. A. Slutz, A. B. Sefkow, D. B. Sinars, K. D. Hahn, S. B. Hansen, E. C. Harding, P. F. Knapp, P. F. Schmit, C. A. Jennings, T. J. Awe, M. Geissel, D. C. Rovang, G. A. Chandler, G. W. Cooper, M. E. Cuneo, A. J. Harvey-Thompson, M. C. Herrmann, M. H. Hess, O. Johns, D. C. Lamppa, M. R. Martin, R. D. McBride, K. J. Peterson, J. L. Porter, G. K. Robertson, G. A. Rochau, C. L. Ruiz, M. E. Savage, I. C. Smith, W. A. Stygar, and R. A. Vesey, *Phys. Rev. Lett.* **113**, 155003 (2014).
- ⁵D. C. Rovang, D. C. Lamppa, M. E. Cuneo, A. C. Owen, J. McKenney, D. W. Johnson, S. Radovich, R. J. Kaye, R. D. McBride, C. S. Alexander, T. J. Awe, S. A. Slutz, A. B. Sefkow, T. A. Hail, P. A. Jones, J. W. Argo, D. G. Dalton, G. K. Robertson, E. M. Waisman, D. B. Sinars, J. Meissner, M. Milhous, D. N. Nguyen, and C. H. Mielke, *Review of Scientific Instruments* **85**, 124701 (2014).
- ⁶P. K. Rambo, I. C. Smith, J. L. Porter, Jr., M. J. Hurst, C. S. Speas, R. G. Adams, A. J. Garcia, E. Dawson, B. D. Thurston, C. Wakefield, J. W. Kellogg, M. J. Slattey, H. C. Ives, III, R. S. Broyles, J. A. Caird, A. C. Erlandson, J. E. Murray, W. C. Behrendt, N. D. Neilsen, and J. M. Narduzzi, *Appl. Optics* **44**, 2421 (2005).
- ⁷P. Rambo, J. Schwarz, M. Schollmeier, M. Geissel, I. Smith, M. Kimmel, C. Speas, J. Shores, D. Armstrong, J. Bellum, E. Field, D. Kletecka, and J. Porter (2016) pp. 10014 – 10014 – 16.
- ⁸M. R. Gomez, S. A. Slutz, C. A. Jennings, D. J. Ampleford, M. R. Weis, C. E. Myers, D. A. Yager-Elorriaga, K. D. Hahn, S. B. Hansen, E. C. Harding, A. J. Harvey-Thompson, D. C. Lamppa, M. Mangan, P. F. Knapp, T. J. Awe, G. A. Chandler, G. W. Cooper, J. R. Fein, M. Geissel, M. E. Glinsky, W. E. Lewis, C. L. Ruiz, D. E. Ruiz, M. E. Savage, P. F. Schmit, I. C. Smith, J. D. Styron, J. L. Porter, B. Jones, T. R. Mattsson, K. J. Peterson, G. A. Rochau, and D. B. Sinars, *Phys. Rev. Lett.* **125**, 155002 (2020).
- ⁹M. E. Glinsky, T. W. Moore, W. E. Lewis, M. R. Weis, C. A. Jennings, D. J. Ampleford, P. F. Knapp, E. C. Harding, M. R. Gomez, and A. J. Harvey-Thompson, *Physics of Plasmas* **27**, 112703 (2020), <https://doi.org/10.1063/5.0010781>.
- ¹⁰D. J. Ampleford, C. A. Jennings, E. C. Harding, M. R. Gomez, T. Webb, P. F. Schmit, T. Awe, P. Knapp, H.-T. A. J., M. R. Weis, and et. al., In Preparation (2021).
- ¹¹D. B. Sinars, S. A. Slutz, M. C. Herrmann, R. D. McBride, M. E. Cuneo, K. J. Peterson, R. A. Vesey, C. Nakhleh, B. E. Blue, K. Killebrew, D. Schroen, K. Tomlinson, A. D. Edens, M. R. Lopez, I. C. Smith, J. Shores, V. Bigman, G. R. Bennett, B. W. Atherton, M. Savage, W. A. Stygar, G. T. Leifeste, and J. L. Porter, *Phys. Rev. Lett.* **105**, 185001 (2010).
- ¹²R. D. McBride, S. A. Slutz, C. A. Jennings, D. B. Sinars, M. E. Cuneo, M. C. Herrmann, R. W. Lemke, M. R. Martin, R. A. Vesey, K. J. Peterson, A. B. Sefkow, C. Nakhleh, B. E. Blue, K. Killebrew, D. Schroen, T. J. Rogers, A. Laspe, M. R. Lopez, I. C. Smith, B. W. Atherton, M. Savage, W. A. Stygar, and J. L. Porter, *Phys. Rev. Lett.* **109**, 135004 (2012).
- ¹³R. D. McBride, M. R. Martin, R. W. Lemke, J. B. Greenly, C. A. Jennings, D. C. Rovang, D. B. Sinars, M. E. Cuneo, M. C. Herrmann, S. A. Slutz, C. W. Nakhleh, D. D. Ryutov, J.-P. Davis, D. G. Flicker, B. E. Blue, K. Tomlinson, D. Schroen, R. M. Stamm, G. E. Smith, J. K. Moore, T. J. Rogers, G. K. Robertson, R. J. Kamm, I. C. Smith, M. Savage, W. A. Stygar, G. A. Rochau, M. Jones, M. R. Lopez, J. L. Porter, and M. K. Matzen, *Phys. Plasmas* **20**, 056309 (2013).
- ¹⁴T. J. Awe, R. D. McBride, C. A. Jennings, D. C. Lamppa, M. R. Martin, D. C. Rovang, S. A. Slutz, M. E. Cuneo, A. C. Owen, D. B. Sinars, K. Tomlinson, M. R. Gomez, S. B. Hansen, M. C. Herrmann, J. L. McKenney, C. Nakhleh, G. K. Robertson, G. A. Rochau, M. E. Savage, D. G. Schroen, and W. A. Stygar, *Phys. Rev. Lett.* **111**, 235005 (2013).
- ¹⁵T. J. Awe, C. A. Jennings, R. D. McBride, M. E. Cuneo, D. C. Lamppa, M. R. Martin, D. C. Rovang, D. B. Sinars, S. A. Slutz, A. C. Owen, K. Tomlinson, M. R. Gomez, S. B. Hansen, M. C. Herrmann, M. C. Jones, J. L. McKenney, G. K. Robertson, G. A. Rochau, M. E. Savage, D. G. Schroen, and W. A. Stygar, *Phys. Plasmas* **21**, 056303 (2014).
- ¹⁶P. F. Knapp, M. R. Gomez, S. B. Hansen, M. E. Glinsky, C. A. Jennings, S. A. Slutz, E. C. Harding, K. D. Hahn, M. R. Weis, M. Evans, M. R. Martin, A. J. Harvey-Thompson, M. Geissel, I. C. Smith, D. E. Ruiz, K. J. Peterson, B. M. Jones, J. Schwarz, G. A. Rochau, D. B. Sinars, R. D. McBride, and P.-A. Gourdain, *Physics of Plasmas* **26**, 012704 (2019), <https://doi.org/10.1063/1.5064548>.
- ¹⁷S. B. Hansen, M. R. Gomez, A. B. Sefkow, S. A. Slutz, D. B. Sinars, K. D. Hahn, E. C. Harding, P. F. Knapp, P. F. Schmit, T. J. Awe, R. D. McBride, C. A. Jennings, M. Geissel, A. J. Harvey-Thompson, K. J. Peterson, D. C. Rovang, G. A. Chandler, G. W. Cooper, M. E. Cuneo, M. C. Herrmann, M. H. Hess, O. Johns, D. C. Lamppa, M. R. Martin, J. L. Porter, G. K. Robertson, G. A. Rochau, C. L. Ruiz, M. E. Savage, I. C. Smith, W. A. Stygar, R. A. Vesey, B. E. Blue, D. Ryutov, D. G. Schroen, and K. Tomlinson, *Physics of Plasmas* **22**, 056313 (2015), <http://dx.doi.org/10.1063/1.4921217>.
- ¹⁸K. D. Hahn, G. A. Chandler, C. L. Ruiz, G. W. Cooper, M. R. Gomez, S. Slutz, A. B. Sefkow, D. B. Sinars, S. B. Hansen, P. F. Knapp, P. F. Schmit, E. Harding, C. A. Jennings, T. J. Awe, M. Geissel, D. C. Rovang, J. A. Torres, J. A. Bur, M. E. Cuneo, V. Y. Glebov, A. J. Harvey-Thompson, M. C. Herrman, M. H. Hess, O. Johns, B. Jones, D. C. Lamppa, J. S. Lash, M. R. Martin, R. D. McBride, K. J. Peterson, J. L. Porter, J. Reneker, G. K. Robertson, G. A. Rochau, M. E. Savage, I. C. Smith, J. D. Styron, and R. A. Vesey, *Journal of Physics: Conference Series* **717**, 012020 (2016).
- ¹⁹P. 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).
- ²⁰P. 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>.
- ²¹J. Gaffney, D. Clark, V. Sonnad, and S. Libby, *High Energy Density Physics* **9**, 457 (2013).
- ²²R. 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>.

- ²³R. 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>.
- ²⁴Note that the values of BR published in Ref. 8 were found using an early version of the surrogate model demonstrated in this manuscript. The inferred values have changed only slightly from those previously published values due to improvements in the surrogate model, and there is no impact on the conclusions therein.
- ²⁵G. B. Zimmerman and W. L. Kruer, Comments Plasma Phys. Controlled Fusion **2**, 51 (1975).
- ²⁶J. A. Harte, W. E. Alley, D. S. Bailey, J. L. Eddleman, and G. B. Zimmerman, UCRL- LR-105821-9-4, No. 4, page 150, Lawrence Livermore National Laboratory (1996).
- ²⁷S. I. Braginskii, *Reviews of Plasma Physics*, Vol. 1 (Consultants Bureau, New York, 1965).
- ²⁸M. R. Weis, A. J. Harvey-Thompson, and D. E. Ruiz, Physics of Plasmas **28**, 012705 (2021), <https://doi.org/10.1063/5.0029850>.
- ²⁹J. R. Rygg, F. H. Séguin, C. K. Li, J. A. Frenje, M. J.-E. Manuel, R. D. Petrasso, R. Betti, J. A. Delettrez, O. V. Gotchev, J. P. Knauer, D. D. Meyerhofer, F. J. Marshall, C. Stoeckl, and W. Theobald, Science **319**, 1223 (2008), <http://www.sciencemag.org/content/319/5867/1223.full.pdf>.
- ³⁰C. K. Li, F. H. Séguin, J. A. Frenje, M. Rosenberg, R. D. Petrasso, P. A. Amendt, J. A. Koch, O. L. Landen, H. S. Park, H. F. Robey, R. P. J. Town, A. Casner, F. Philippe, R. Betti, J. P. Knauer, D. D. Meyerhofer, C. A. Back, J. D. Kilkenny, and A. Nikroo, Science **327**, 1231 (2010), <http://www.sciencemag.org/content/327/5970/1231.full.pdf>.
- ³¹M. Hohenberger, P.-Y. Chang, G. Fiksel, J. P. Knauer, R. Betti, F. J. Marshall, D. D. Meyerhofer, F. H. Séguin, and R. D. Petrasso, Physics of Plasmas (1994-present) **19**, 056306 (2012).
- ³²C. Li, F. Séguin, J. Rygg, J. Frenje, M. Manuel, R. Petrasso, R. Betti, J. Delettrez, J. Knauer, F. Marshall, D. Meyerhofer, D. Shvarts, V. Smalyuk, C. Stoeckl, O. Landen, R. P. Town, C. Back, and J. Kilkenny, Phys. Rev. Lett. **100**, 225001 (2008).
- ³³M. D. Cable and S. P. Hatchett, Journal of Applied Physics **62**, 2233 (1987).
- ³⁴P. F. Schmit, K. Molvig, and C. W. Nakhleh, Physics of Plasmas (1994-present) **20**, 112705 (2013).
- ³⁵P. F. Knapp, D. B. Sinars, and K. D. Hahn, Physics of Plasmas (1994-present) **20**, 062701 (2013).
- ³⁶H.-S. Bosch and G. M. Hale, Nuclear Fusion **32**, 611 (1992).
- ³⁷M. D. Hoffman and A. Gelman, Journal of Machine Learning Research **15**, 1593 (2014).
- ³⁸J. H. Halton, Numerische Mathematik **2**, 84 (1960).
- ³⁹K. He, X. Zhang, S. Ren, and J. Sun, in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016) pp. 770–778.
- ⁴⁰X. Glorot and Y. Bengio, in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research, Vol. 9, edited by Y. W. Teh and M. Titterton (JMLR Workshop and Conference Proceedings, Chia Laguna Resort, Sardinia, Italy, 2010) pp. 249–256.
- ⁴¹D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” (2017), arXiv:1412.6980 [cs.LG].
- ⁴²B. Appelbe, J. Pecover, and J. Chittenden, High Energy Density Physics **22**, 27 (2017).
- ⁴³M. R. Gomez, S. A. Slutz, P. F. Knapp, K. D. Hahn, M. R. Weis, E. C. Harding, M. Geissel, J. R. Fein, M. E. Glinsky, S. B. Hansen, A. J. Harvey-Thompson, C. A. Jennings, I. C. Smith, D. Woodbury, D. J. Ampleford, T. J. Awe, G. A. Chandler, M. H. Hess, D. C. Lamppa, C. E. Myers, C. L. Ruiz, A. B. Sefkow, J. Schwarz, D. A. Yager-Elorriaga, B. Jones, J. L. Porter, K. J. Peterson, R. D. McBride, G. A. Rochau, and D. B. Sinars, IEEE Transactions on Plasma Science **47**, 2081 (2019).
- ⁴⁴A. J. Harvey-Thompson, M. Geissel, C. A. Jennings, M. R. Weis, M. R. Gomez, J. R. Fein, D. J. Ampleford, G. A. Chandler, M. E. Glinsky, K. D. Hahn, S. B. Hansen, E. C. Harding, P. F. Knapp, R. R. Paguio, L. Perea, K. J. Peterson, J. L. Porter, P. K. Rambo, G. K. Robertson, G. A. Rochau, C. L. Ruiz, J. Schwarz, J. E. Shores, D. B. Sinars, S. A. Slutz, G. E. Smith, I. C. Smith, C. S. Speas, K. Whitemore, and D. Woodbury, Physics of Plasmas **26**, 032707 (2019), <https://doi.org/10.1063/1.5086044>.
- ⁴⁵M. S. Schollmeier, P. F. Knapp, D. J. Ampleford, E. C. Harding, C. A. Jennings, D. C. Lamppa, G. P. Loisel, M. R. Martin, G. K. Robertson, J. E. Shores, I. C. Smith, C. S. Speas, M. R. Weis, J. L. Porter, and R. D. McBride, Review of Scientific Instruments **88**, 103503 (2017), <https://doi.org/10.1063/1.4994566>.