Reconstruction of residual kinetic energy in inertially confined fusion implosions at the National Ignition Facility using multiple heterogeneous data sources and neural networks

J. H. Kunimune,[1, a] D. T. Casey,[2] B. Kustowski,[2] M. Jones,[2] L. Divol,[2] T. M. Johnson,[1] S. G. Dannhoff,[1] A. DeVault,[1] and J. A. Frenje[1]
[1] *Plasma Science and Fusion Center, Massachusetts Institute of Technology. 167 Albany St., Cambridge, MA, USA.*
[2] *Lawrence Livermore National Laboratory. 7000 East Ave., Livermore, CA, USA.*

(Dated: 23 June 2026)

Three-dimensional (3D) asymmetries and their associated residual kinetic energy (RKE) are a major performance-degradation mechanism in inertial-confinement fusion (ICF) implosions at the National Ignition Facility (NIF). These asymmetries can be diagnosed with the three neutron imaging systems fielded with orthogonal views to the implosion, as well as a suite of real-time neutron activation diagnostics (RTNADs) and neutron spectrometers. Conventional tomographic reconstructions are typically used to reconstruct the 3D morphology of the hot-spot and surrounding dense fuel in an implosion using neutron images, but the reconstruction problem is ill-posed with only three imaging lines of sight. In this work, a 3D reconstruction technique has been developed and used to overcome these limitations by combining information from the neutron images, RTNAD data, and neutron spectra to infer RKE. Markov chain Monte Carlo is used along with a group of neural networks to fit a physics model to the observed data to robustly reconstruct the 3D morphology of the hot-spot and dense fuel in a NIF implosion. Bayesian statistics provides quantification of the uncertainty in the inference. Initial inferences find that the normalized RKE on several NIF experiments is between 2 and 5% of the total implosion energy, and is inversely related to the yield of an implosion, as expected. This technique is being developed to help guide designs toward minimizing asymmetries in ICF implosions and achieving higher performance.

I. INTRODUCTION

In an inertial confinement fusion (ICF) experiment at the National Ignition Facility (NIF), a 2 MJ laser irradiates the inside of a gold cylinder called a *hohlraum* to generate x-rays that spherically compress a capsule of deuterium–tritium (DT) fuel suspended inside it. The implosion of the capsule generates an extremely hot plasma in which fusion reactions take place. The implosion at the time of minimum volume can be divided into a central hot-spot and a surrounding dense fuel-shell. In an ideal implosion, the hot-spot reaches high enough temperatures and densities for a thermonuclear reaction to ignite and propagate into the surrounding dense fuel – this is called *ignition*.[1] In reality, ignition is often prevented by a number of degradation mechanisms.

One of the most significant degradations is 3D asymmetries. While an ideal ICF implosion is perfectly spherical, in reality, small perturbations are always present in the capsule and drive that grow into large 3D asymmetries and flows in the shell at the time of minimum volume. The kinetic energy associated with these flows is called *residual kinetic energy* (RKE) because it remains in the shell instead of being converted to heat energy in the hot-spot. RKE has a major impact on the performance of ICF implosions[2] but cannot be directly measured. Thus, new diagnostic techniques are needed to better understand asymmetries and RKE so that they can be addressed and implosion performance can be improved.

At the NIF, asymmetries are diagnosed by a variety of imagers, including three neutron imaging systems. Two of the neutron imaging systems are time gated to provide both an image of neutron emission above 13 MeV (the primary neutrons that exit the implosion at their birth energy without scattering) and an image of neutron emission between 6 MeV and 12 MeV (the down-scattered neutrons that lose energy colliding with fuel nuclei before exiting the implosion).[3,4] The primary neutron images are used to infer the morphology of the hot-spot, and the down-scattered neutron images are combined with the primary neutron images in a process known as *fluence-compensation* to infer the morphology of the surrounding dense fuel.[5]

Traditionally, these images alone are used to infer the 3D shape of the implosion using tomographic reconstruction.[6] However, because there are only three available primary neutron images and two available down-scattered neutron images, the amount of available information is limited. Tomographic reconstruction with only two or three lines of sight is severely ill-posed, and can produce implausible and inaccurate results unless the implosion is assumed to have some symmetry.[7] In addition, because neutron images contain no temporal information, tomographic reconstructions can only recon-

[a] Author to whom correspondence should be addressed: kunimune@mit.edu

struct the shape of the shell, not its velocity distribution or RKE.

Asymmetries are also diagnosed by a variety of instruments that are not imagers. One is the real-time neutron activation detector (RTNAD) array.[8] The RTNADs comprise 48 zirconium pucks around the NIF target chamber that each measure the primary neutron yield. Because the primary neutron yield is attenuated by the dense fuel, variations among the RTNADs are indicative of dense fuel asymmetries. By making assumptions about the dense fuel shape, low-mode (mode-1 and mode-2) asymmetries can be inferred from the RTNAD data.

Another diagnostic is the suite of neutron time-of-flight spectrometers.[9] Each spectrometer measures the mean energy of the primary neutrons, which is Doppler-shifted according to the mean velocity of the hot-spot relative to the spectrometer's line of sight. By combining measurements from seven spectrometers around the target chamber, the velocity vector $\vec{\mathbf{v}}_{\mathrm{HS}}$ can be reconstructed in 3D. This velocity vector is closely coupled to the mode-1 asymmetry in the implosion. In addition, each spectrometer measures the down-scatter ratio (DSR) – the ratio of neutron yield in the energy range 10–12 MeV to neutron yield in the energy range 13–15 MeV – which is proportional to areal density. This provides a measure of the implosion's overall areal density.

Each of these diagnostics probes the shape of the hot-spot and surrounding dense fuel, and each is routinely used in isolation to infer the mode and magnitude of 3D asymmetries in NIF implosions.[10,11] However, these inferences are often inconsistent with each other. This is because when analyzed individually, each only samples part of the implosion, and thus only sees a subset of the possible asymmetry modes.

This is compounded by the fact that all of these diagnostics only probe the shape of the implosion around bang-time (the time of peak neutron emission rate). This is problematic when trying to infer the amount of RKE at the time of minimum volume. Bang-time can occur hundreds of picoseconds after minimum-volume, and during that time the outward acceleration of the shell both amplifies asymmetries and increases the total amount of kinetic energy – some of this added energy is associated with asymmetric modes and some is not. This is illustrated by the simulation shown in Figure 1, which shows a simulation (using the model described in Section II) where the kinetic energy (as a percentage of the total implosion energy) rises from 1.9% at minimum-volume to 10.5% at bang-time. Thus, analyses that depend on measurements of the areal density distribution at bang-time may provide inaccurate results when inferring the RKE at minimum-volume. These errors are difficult to quantify without detailed modeling.

A methodology based on neural networks and a physics model was recently developed by Kunimune et al. to combine information from primary neutron images and the neutron spectrometers to provide a complete self-consistent 3D reconstruction of the implosion.[12] This methodology proved capable of accurately reconstructing hot-spot shape parameters from synthetic data. However, it did not work for experimental data. Furthermore, the reconstructions it produced did not have uncertainty quantification, making it impossible to know how unique or trustworthy the found solutions were. Uncertainty quantification is especially important for 3D reconstructions based on limited amounts of data, as the problem is heavily ill-posed. While the use of a physics model to constrain the results improves the situation, there may still be a wide range of morphologies that match the data.

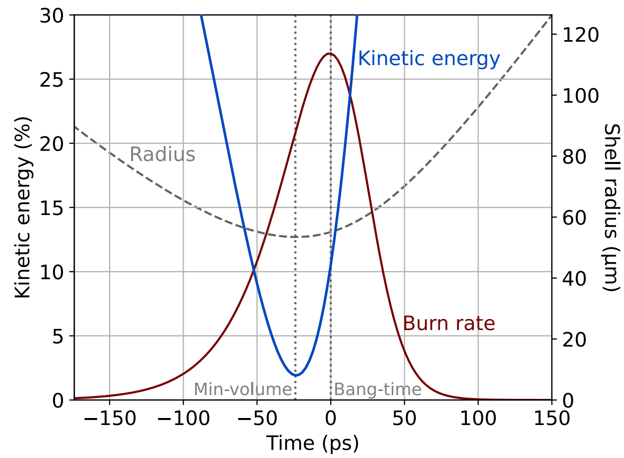

FIG. 1. The evolution of the shell kinetic energy in a simulated implosion with a yield of 9×10^{16} neutrons. RKE is defined as the minimum value of kinetic energy, which – for this simulation – occurs when the volume is at its smallest. Because of alpha-heating, the burn rate does not peak until tens of picoseconds later, at bang-time, at which point outward acceleration has increased the kinetic energy in the shell substantially. This simulation assumes the thickness of the shell is constant, and so does not account for the nonzero kinetic energy associated with compression and shock propagation within the shell.

A new methodology, building off of the work by Kunimune et al., has therefore been developed to overcome these challenges and robustly diagnose the 3D morphology and RKE of an implosion at the time of minimum volume. This methodology combines information available from the imagers, the RTNADs, and the neutron spectrometers, and connects it to the implosion dynamics using a simple physics model. It uses Bayesian statistics to perform uncertainty quantification, validating the uniqueness of the solutions and ensuring inferences based on this technique are robust even when the amount of available data is limited.

This methodology is based on a forward fit. Forward fitting is a technique commonly used to infer unknown parameters from data. The general process, depicted in Figure 2, is as follows. First some experimental data is collected (in this work, the data comprises neutron images, RTNAD data, and the spectrometer-measured

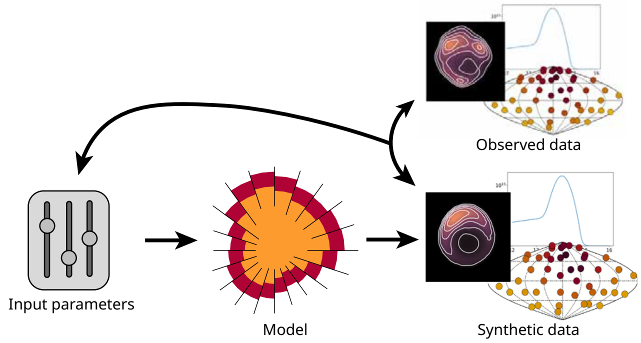

FIG. 2. The forward-fitting process, shown in broad strokes. It starts with the collection of some experimental data. Some input parameters are chosen by guessing, and fed into a physics model to obtain synthetic data. The synthetic data is compared to the observed data, and the differences between them inform an update to the input parameters. The input parameters are thus iteratively adjusted until the synthetic data matches the observed data. Neural networks are used to speed up the model evaluations.

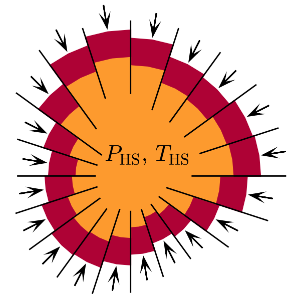

FIG. 3. The rocket-piston model of an implosion.[13] The hot-spot (orange) is assumed to be isobaric and have negligible mass. The pistons (red) have mass and are initially accelerated radially inward and after peak compression outward by the outer ablation pressure and the hot-spot pressure, respectively.

yield, DSR, and $\vec{\mathbf{v}}_{\mathrm{HS}}$). Then a physics model is used, with arbitrarily chosen input parameters, to generate synthetic data of the same type. The synthetic data is compared to the observed data, and the input parameters are iteratively tuned to improve the agreement between them until they match. The optimal input parameters thus obtained can then be used with the physics model to calculate any desired information about the best fit to the experiment.

This methodology is quite general and can be used to fit a morphology to any set of data, given an appropriate choice of model, an appropriate way to quantify the discrepancy between synthetic and observed data, and an appropriate method of tuning the parameters. Section II of this paper will discuss the model that has been chosen for this work, the so-called rocket-piston description of an implosion. Section III will discuss the application of Bayesian statistics to quantify the discrepancy and the use of Markov chain Monte Carlo to tune the input parameters. Section IV will discuss the use of a neural-network based surrogate model to speed up the calculation. Section V will demonstrate the methodology applied to synthetic and experimental NIF data. Finally, Section VI will discuss future work, and Section VII will conclude.

II. ROCKET-PISTON MODEL

A simplified multi-rocket-piston model is used for the forward fit. It represents the implosion as a set of massive pistons around a hot-spot, as depicted in Figure 3. The pistons are initially accelerated radially inward by the ablation pressure, and then after peak compression are accelerated radially outward by the hot-spot pressure.[13]

In this model, the radiation temperature and power inside the hohlraum is determined by the hohlraum model given by Callahan et al.,[14] and this is converted to a mass ablation rate using the empirical scalings found by Olsen et al.[15] A 16×32 spherical grid of pistons is initialized at the nominal capsule radius, and each is left free to move radially under the influence of the external ablation pressure and internal hot-spot pressure while the carbon mass in each piston decreases due to the ablation, as described by models by Hurricane et al.[10,11]. The hot-spot, defined as the volume contained within the pistons, has two state variables – pressure and central temperature, or equivalently, total energy and total mass – which evolve according to the dynamic power balance model developed by Springer et al.[16] with one simplification: rather than modeling the propagation of shocks through the shell, the shell in each piston is treated as rigid, and a constant multiplier is applied to the shell mass to account for the fact that, during deceleration, only part of it has been shocked and thus only part of it can affect the stagnation of the hot-spot.

Asymmetries are seeded by initializing each piston with a slightly different mass. Specifically, the mass of carbon and dense fuel in the ith piston is set to

$$m_{\mathrm{C},i} = \frac{1}{n} M_{\mathrm{C}} \left(1 + \sum_{l=1}^{5} \sum_{m=-l}^{l} a_{l,m} Y_l^m(\theta_i, \phi_i) \right) \tag{1}$$

$$m_{\mathrm{DT},i} = \frac{1}{n} M_{\mathrm{DT}} \left(1 + \sum_{l=1}^{5} \sum_{m=-l}^{l} a_{l,m} Y_l^m(\theta_i, \phi_i) \right) \tag{2}$$

where n is the number of pistons, M_{C} is the nominal total ablator mass, M_{DT} is the nominal total DT ice mass, $a_{l,m}$ is a perturbation coefficient for the $\langle l, m \rangle$ mode, and $Y_l^m(\theta, \phi)$ is a spherical harmonic function. Typically only modes with $l \leq 5$ are used, but theoretically infinitely many modes may be included, as long as the number of pistons is increased accordingly.

The rocket-piston model is simple and can thus be executed relatively quickly – a single execution takes only a few seconds. However, the simplicity of this model also means that it is not as accurate as full hydrodynamic simulations with codes like HYDRA.[17] Several adjustments must be made to capture some of the finer points of the implosion dynamics. For example, the radiation temperature must be adjusted to tune the amount of ablator that remains unablated at bang-time, because the ablation rates that one obtains by naively fitting the hohlraum model to observed x-ray spectra tend to underpredict the actual amount of ablation, and the unphysically high shell mass can make the power balance equations unphysically stiff. Similarly, a multiplier must be applied to the alpha-heating term in the power balance equations to mimic the collective effect of degradation mechanisms that are not modeled, like mix between the ablator and fuel.

In addition, an ad-hoc non-radial flow term must be added to mimic a phenomenon that is often observed in the implosions where, due to the swings in the drive asymmetry over the course of the implosion, shell mass gets concentrated at the poles of the implosion.[18] The non-radial flow term siphons fuel mass from equatorial pistons to polar pistons to mimic this; the rate at which this happens is treated as a tuning parameter.

The model outputs are the pistons' masses m_{DT} and m_{C}, the pistons' radii R, and the hot-spot's mass and energy (or, equivalently, central ion temperature $T_{\mathrm{i},0}$ and pressure p) at every time-step. With this model, synthetic data is generated and used for comparison to the experimental results.[13] Yield is derived from the hot-spot conditions integrated in time; hot-spot velocity is derived by differentiating the hot-spot centroid in time; and the neutron images, RTNAD skymap, and DSR skymap are calculated by converting the piston geometry to a 3D voxel grid of temperatures and densities and then performing line-integrals as described by Kunimune et al.[19]

The synthetic neutron images are of higher quality than is typically obtained from experimental data, and must be modified to be directly comparable to it. To account for the fact that the neutron imaging systems only register their images relative to each other (they don't absolutely register them to the target chamber center), the images on each line of sight are translated to put the primary neutron image's center of mass at the center of the image domain. Further, to mimic the finite resolution of the imaging systems, all synthetic images are degraded by an ideal spatial low-pass filter. This is equivalent to blurring each image with the convolution kernel $P(r) \propto J_1(kr)/r$,[20] where $J_1(x)$ is a Bessel function of the first kind. The scale factor k of the convolution depends on the neutron yield, and is treated as a tuning parameter.

It is important to note that while the neutron imaging hardware has a point-spread function associated with it, this low-pass filter is distinct from that. Unlike the instrument's point-spread function, which is quite complex and lacks translational symmetry,[21] the low-pass filter is only a slight blur that approximates the combined effect of the instrument's point-spread function and the deconvolution algorithm that recovers the neutron source from the measured image.

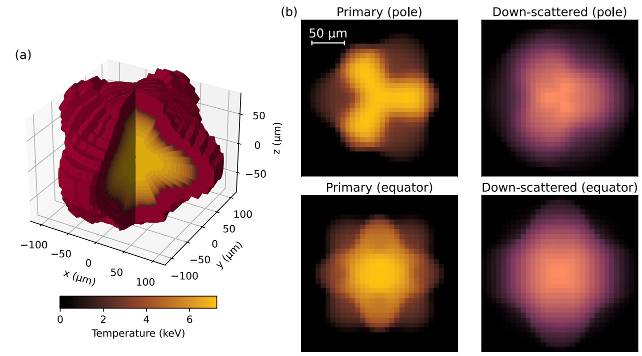

FIG. 4. Example 3D morphology and neutron images generated by the rocket-piston model. (a) The red regions are the shell, and the orange regions are the hot-spot, color-coded by ion temperature. (b) Four synthetic primary and down-scattered neutron images, generated by the implosion shown in a., generated along two different lines of sight.

Examples of the degraded synthetic neutron images generated by the rocket-piston model are shown in Figure 4.

III. FITTING WITH MARKOV CHAIN MONTE CARLO

Given the significant possibility that there are multiple morphologies that match the data, a rigorous approach such as Bayesian analysis must be applied. In Bayesian analysis, it is assumed that, because of noise and uncertainty in the measurements, the correct answer cannot be exactly determined. Therefore, instead of finding a single answer, we seek to find a *distribution.*

Bayes's theorem is used to quantify the *a posteriori* probability density $p(\mathbf{x}|\mathbf{d})$ that a given set of inputs is correct:

$$p(\mathbf{x}|\mathbf{d}) \propto p(\mathbf{x}) \cdot p(\mathbf{d}|\mathbf{x}) \tag{3}$$

In this context, $\mathbf{x}$ represents the input parameters to the rocket-piston model, and $\mathbf{d}$ represents the observed data. The likelihood $p(\mathbf{d}|\mathbf{x})$ can be thought of as a measure of fit quality, much like a χ^2 metric; while the *a priori* probability density $p(\mathbf{x})$ can be thought of as a regularization term.

The prior input distribution $p(\mathbf{x})$ is set to a multivariate truncated normal distribution.

$$p(\mathbf{x}) \propto \prod_j \begin{cases} \exp\left(-\frac{(x_j-\mu_j)^2}{2\sigma_j^2}\right) & \text{if } x_j^- \leq x_j \leq x_j^+ \\ 0 & \text{if } x_j < x_j^- \text{or } x_j > x_j^+ \end{cases} \tag{4}$$

Here, x_j is the jth input parameter, and μ_j, σ_j, x_j^-, and x_j^+ are the mean, standard deviation, lower bound,

TABLE I. The input parameters for the prior probability distribution used in the rocket-piston model. These parameters are used in Equation 4. Note that there are 23 different $\langle l,\, m\rangle$ pairs, resulting in 29 input parameters in total.

Input parameter	Parts	μ	σ	x^-	x^+
Initial pressure (kbar)	1	150	±50	0	∞
$\langle l,\, m\rangle$ mass perturbation	23	0	$\frac{\pm 0.02}{2l+1}$	$\frac{-0.10}{2l+1}$	$\frac{0.10}{2l+1}$
Alpha heating multiplier	1	1.00	±0.15	0	∞
Shocked mass fraction	1	0.60	±0.05	0.45	0.75
Non-radial flow term	1	0.16	±0.08	-1.00	1.00
Dense fuel adiabat	1	3.8	±0.6	1.5	10.0
Image blur width (µm)	1	12	±4	3	100

TABLE II. The composition of a typical 3085-vector output by the rocket-piston model, used to calculate the likelihood of the data. Some shots may have more output data, and others may have less.

Output datum	Parts
Logarithm of yield	1
Down-scatter ratio	1
Hot-spot velocity vector	3
RTNAD spherical harmonic fit	8
Primary neutron image 1	1024
Primary neutron image 2	1024
Down-scattered neutron image	1024

and upper bound of its prior distribution, respectively. It was found that 29 input parameters were needed to access the range of morphologies seen in experiments; the full list of variable parameters, along with their distribution parameters, is shown in Table I. The prior distributions, for the most part, are chosen based on physical intuition and are much broader than the posterior distributions, indicating that they do not significantly influence the result. The exception is the shocked mass fraction, which has strict limits based on physics constraints, and the high-mode-number mass perturbations, which must be limited to prevent the algorithm from straying into extreme shapes that the model cannot handle (the rocket-piston model becomes non-physical when asymmetries become so strong that the shell passes through the origin).

Note that the shape perturbations are restricted to 23 modes: all fifteen modes with $1 \leq l \leq 3$ and $-l \leq m \leq l$, plus $\langle 4,\, 0\rangle$, $\langle 4,\, \pm 4\rangle$, $\langle 5,\, 0\rangle$, $\langle 5,\, \pm 4\rangle$, and $\langle 5,\, \pm 5\rangle$. These are the modes that are expected to be seeded by features of the capsule and drive. All inputs not listed in Table I – including the capsule dimensions, the initial hot-spot density, and the drive strength multiplier – are assigned a reasonable arbitrary value and treated as fixed throughout the fitting process. Limiting the number of input dimensions like this makes the sampling faster and more effective.

To calculate the likelihood, the synthetic data outputs and observed data outputs must be arranged into two vectors that can be easily compared. The images are each cropped to a 240 µm × 240 µm square, sampled to a pixel size of 7.5 µm, and then flattened into 1024 individual pixel values. The array of RTNAD data is normalized to its mean and then fit to a linear combination of spherical harmonic functions with $l \leq 2$; those fit coefficients are put in the output vector. The total number of elements in the output vector varies from shot to shot, as some shots have more data associated with them than others, but a typical output vector has 3085 components, enumerated in Table II. The model itself is treated as a function $\mathbf{f}(\mathbf{x})$ that takes a vector in the 29-dimensional input space and generates a vector in the 3085-dimensional output space.

The likelihood of the data $p(\mathbf{d}|\mathbf{x})$ is then calculated as another multivariate normal distribution.

$$p(\mathbf{d}|\mathbf{x}) \propto \prod_k \exp\left(-\frac{(d_k - f_k(\mathbf{x}))^2}{2\epsilon_k^2}\right) \tag{5}$$

Here, d_k is the kth observed output, and $f_k(\mathbf{x})$ is the corresponding synthetic output based on the input vector $\mathbf{x}$. ϵ_k is the standard error on that measurement. For image pixels, which do not have errors reported with them, ϵ_k is set to the mean image value, as this was found to result in the images being weighted about the same as the scalars by the algorithm.

To find a plausible solution, the Adam algorithm,[22] a gradient-based optimizer, is first used to find the $\mathbf{x}$ that maximizes $p(\mathbf{x}|\mathbf{d})$. The Hamiltonian Monte Carlo algorithm,[23] a Markov chain Monte Carlo sampler, is subsequently used to sample the posterior probability distribution. The variation among the obtained samples is used to quantify uncertainty in the inferred values. Because $p(\mathbf{x}|\mathbf{d})$ is complicated and may have multiple local maxima, the replica exchange technique[24] is used to ensure that the sampling is thorough and accurate. This builds off of code originally written by Mike Jones and Jim Gaffney, which uses the implementations of these algorithms included in TensorFlow Probability.[25]

IV. ACCELERATION WITH MACHINE LEARNING

In practice, the sampling cannot be performed using the rocket-piston model directly. This is because obtaining a good characterization of the posterior probability distribution requires thousands of samples, and the Hamiltonian Monte Carlo algorithm requires tens of evaluations of $\nabla p(\mathbf{x}|\mathbf{d})$ for each sample. Evaluating $p(\mathbf{x}|\mathbf{d})$ requires the evaluation of $\mathbf{f}(\mathbf{x})$, which takes multiple seconds, and evaluating $\nabla p(\mathbf{x}|\mathbf{d})$ with the method of finite differences requires evaluating $p(\mathbf{x}|\mathbf{d})$ thirty times. As a result, obtaining a fit with uncertainty quantification for a single set of data would take months of computation. If $\mathbf{f}(\mathbf{x})$ can be automatically differentiated, then the method of finite differences can be foregone, but a single fit would still take over a week.

Thus, the viability of this technique hinges on a way to map the input parameters to synthetic data generated by the rocket-piston model quickly – ideally in a differentiable fashion. Such a way exists in neural networks. Specifically, neural networks can be used to build a surrogate model – a mathematical function that provides the same input-to-data mapping as the rocket-piston model but is faster by orders of magnitude.

To train a surrogate model, a training dataset – many sets of synthetic data generated by the rocket-piston model along with the corresponding sets of input parameters – is first needed. This dataset is built by executing the rocket-piston model two hundred thousand times with different combinations of input parameters. Each input parameter is drawn from a normal distribution for each simulation.

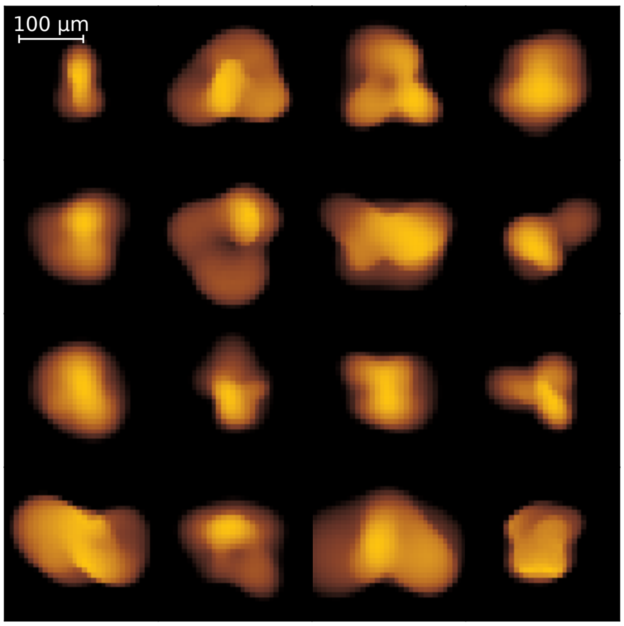

FIG. 5. Sixteen primary neutron images from an equatorial line of sight, randomly selected from the training dataset. Each is produced by the rocket-piston model run on a random set of input parameters. There are two hundred thousand such sets of images in total.

Figure 5 shows a few primary neutron images extracted from such a dataset. The spatial low-pass filter is not applied to these images, as that step is already fast and differentiable, and thus does not need to be folded into the surrogate model.

The surrogate model itself is based on the manifold- and cyclically-consistent architecture that was proposed by Anirudh et al.[26] and that was implemented in TensorFlow 2 with the extensibility to an arbitrary number of image modalities by Kustowski et al.[27] This architecture comprises two parts.

The first part is an autoencoder that maps all of the synthetic data into a 60-dimensional latent space. It comprises two neural networks: one that combines images and scalars into a latent space vector, and another that converts a latent space vector into images and scalars. These are trained together to maximize cyclic consistency between the two, and in the process defines a latent space that encodes the data.

Next, a pair of neural networks are trained to map input parameters to points in the latent space and back. These are trained to maximize both consistency with the rocket-piston-derived training data and cyclic consistency with each other.

A comparison between the data from the surrogate model and data from the rocket-piston model is shown in Figure 6. The surrogate model produces almost the same outputs as the rocket-piston model when given the same input parameters.

Deviations from a perfect match are caused by the limited size of the latent space, the limited size of the neural networks, the limited number of training data, and the limited number of training iterations. Increasing the size of the latent space and networks would improve the quality of the match during training, but increase the risk of over-fitting, resulting in the wrong answer for some inputs. Increasing the number of training data and training iterations would also increase the quality of the match but also increase the amount of computing resources needed to train the surrogate model.

As the surrogate model takes orders of magnitude less time to execute than the rocket-piston model and is automatically differentiable, the sampling process becomes tractable. Using the surrogate model to evaluate $\mathbf{f}(\mathbf{x})$, instead of the rocket-piston model, speeds up the evaluation of $p(\mathbf{x}|\mathbf{d})$ such that drawing 70 000 samples takes just 19 minutes on a single GPU. The first 20 000 of those samples are discarded to give the sampler time to reach the equilibrium distribution, leaving 50 000 samples from the posterior distribution. Because neural networks can be unreliable, thirty samples are then randomly drawn from those 50 000 and fed into the rocket-piston model to check the accuracy of the surrogate model.

In practice, naively training on every sample obtained with the above methodology would be inefficient, as the majority of simulations have either very low yields or very high yields. Training on this bimodal distribution would result in a surrogate model that is faithful to the rocket-piston model for sub-ignition shots and for super-ignition shots, but not for shots on the ignition cliff – right where the physics is the most complex and the most interesting. Thus, it is important to concentrate training samples in the areas of parameter space that are the most relevant. This is done in two ways.

The simplest way is to discard samples with a probability

$$p = 1 - \exp\left(-\left(\frac{17.75 - \log_{10} Y_\mathrm{n}}{0.625}\right)^2 - \left(\frac{2.75\% - \mathrm{DSR}}{0.875\%}\right)^2\right) \tag{6}$$

where Y_n is the simulated neutron yield and DSR is the

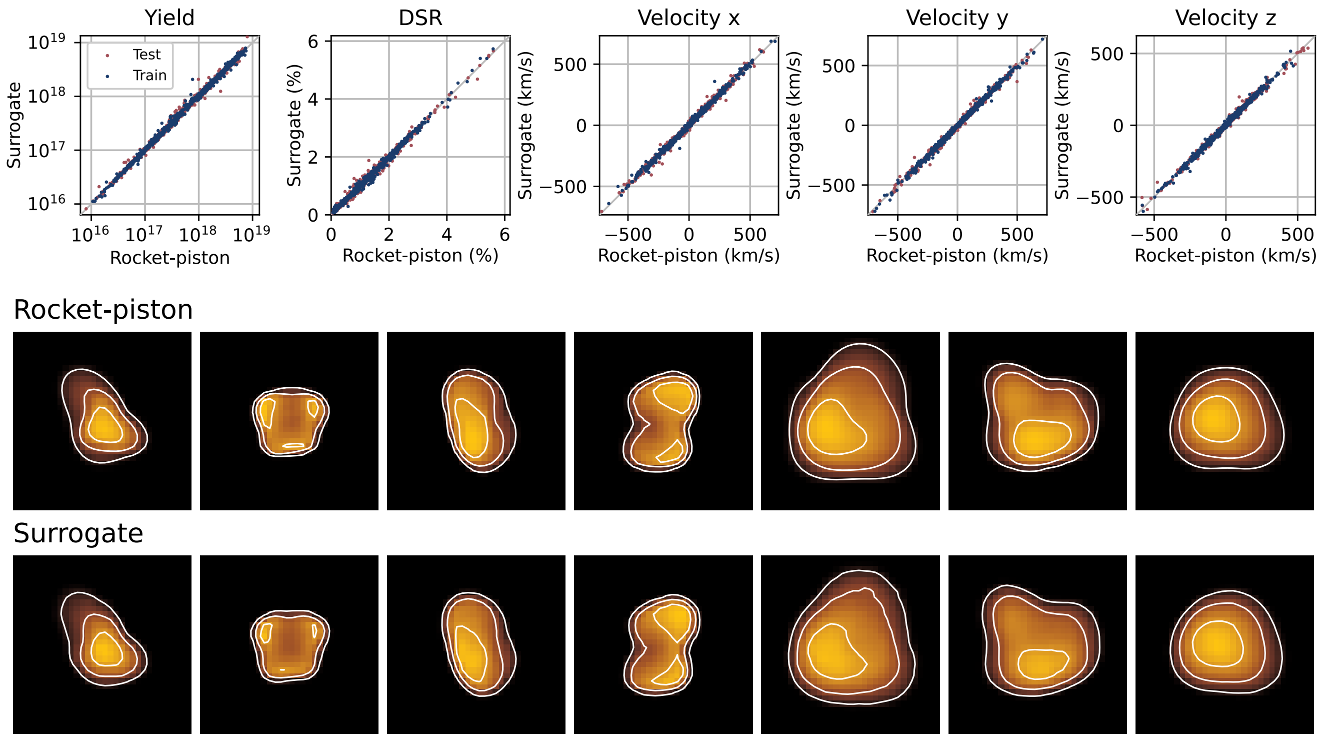

FIG. 6. Surrogate outputs versus the rocket-piston-model outputs for multiple sets of inputs. (top) Each plot represents a single scalar. Blue points were used to train the surrogate model while red points were set aside during the training process for testing. (bottom) The rocket-piston outputs in the form of equatorial primary neutron images (top row) contrasted with the surrogate outputs (bottom row). All of the shown samples were set aside during the training process for testing.

down-scatter ratio. This rejects more than half of the generated samples, concentrating the dataset on intermediate yields. If a surrogate model has already been trained, it can be used to calculate the rejection probability before running each simulation, thus rejecting uninteresting input parameter vectors before any time is wasted running them. Thus, the usual workflow for building a new surrogate model has four steps: generate a small dataset, use it to train a surrogate model that only predicts yield, generate a large dataset while using that surrogate model to filter out points with uninteresting yields, and then use the larger dataset to train a better surrogate model that predicts all outputs.

For particularly challenging reconstructions, an alternative sampling technique is to use a previous surrogate model and the fitting method described in Section III to choose new samples. Because Markov chain Monte Carlo is designed to find points that are close to the experimental data, the output represents points that are especially useful for training. Thus, this workflow has five steps: generate a dataset, use it to train a surrogate model, use Markov chain Monte Carlo to sample the posterior using that surrogate model, augment the dataset with points from the Markov chain, then use the larger dataset to train a better surrogate model. This technique is only useful if the original fit is flawed (see Section V C for examples of how this can happen).

V. RESULTS

A. Application to synthetic data

To illustrate how this methodology works, it is applied to some synthetic data. A random rocket-piston simulation was used to generate synthetic data, including yield, DSR, $\vec{\mathbf{v}}_{\mathrm{HS}}$, 48 RTNAD values, two primary neutron images, and one down-scattered neutron image. Arbitrary uncertainties were assigned to each datum based on their typical uncertainties in experimental data, and the algorithm described in Section III was used to infer the underlying input parameters of the model and morphology of the implosion. Since these synthetic data were generated by the very same rocket-piston model being used in the fit, a more or less perfect match should be obtained.

Figure 7 summarizes the results. Note that, while the surrogate model was used to identify the most probable set of input parameters, all reconstructed quantities shown here are based on the rocket-piston model. Figure 7a shows the simulated morphology used to generate the data beside the reconstructed morphology. Figure 7b shows a comparison of the original scalar data that were used as the observation for the fit contrasted with the reconstructed scalar data. Some discrepancies – for example, the fact that the reconstructed x-component of

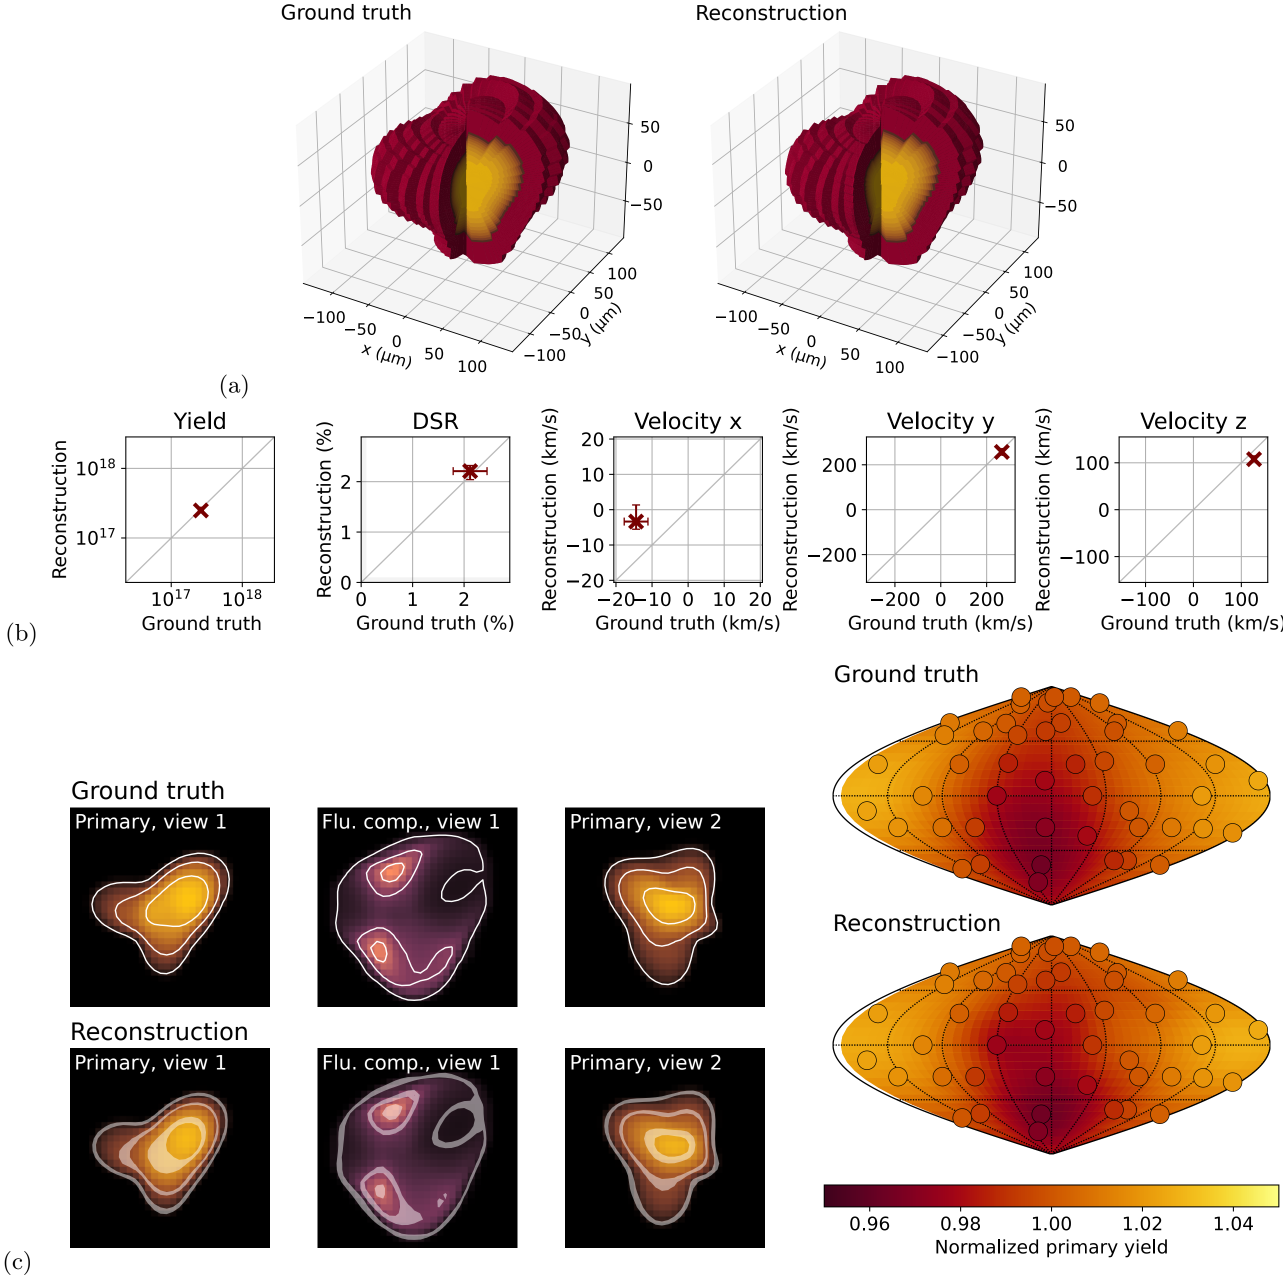

FIG. 7. Comparison of a rocket-piston simulation and synthetic data with the reconstructed simulation and data. (a) The true versus reconstructed 3D morphology of the hot-spot and dense fuel. They are indistinguishable to the eye. (b) Comparisons of true and reconstructed scalar data. All match closely except the x-component of velocity. (c) Comparisons of true (top) and reconstructed (bottom) images and RTNAD sky-maps. Each image displays its 17%, 50%, and 83% contours (the width of the reconstructed contour illustrates the fit's uncertainty). They also agree very well.

$\vec{\mathbf{v}}_{\mathrm{HS}}$ is too small – arise due to the imperfect nature of the surrogate model. However, the key parameters of the implosion – the yield and DSR – are nearly exactly reproduced by the rocket-piston model. Figure 7c shows comparisons of the equatorial primary neutron image, the equatorial fluence-compensated down-scattered neutron image, the polar primary neutron image, and the normalized RTNAD sky-map between the original synthetic data and the reconstructed synthetic data. Again, the fit between the synthetic measurement and the reconstruction is very good, as expected, indicating that this methodology is working correctly.

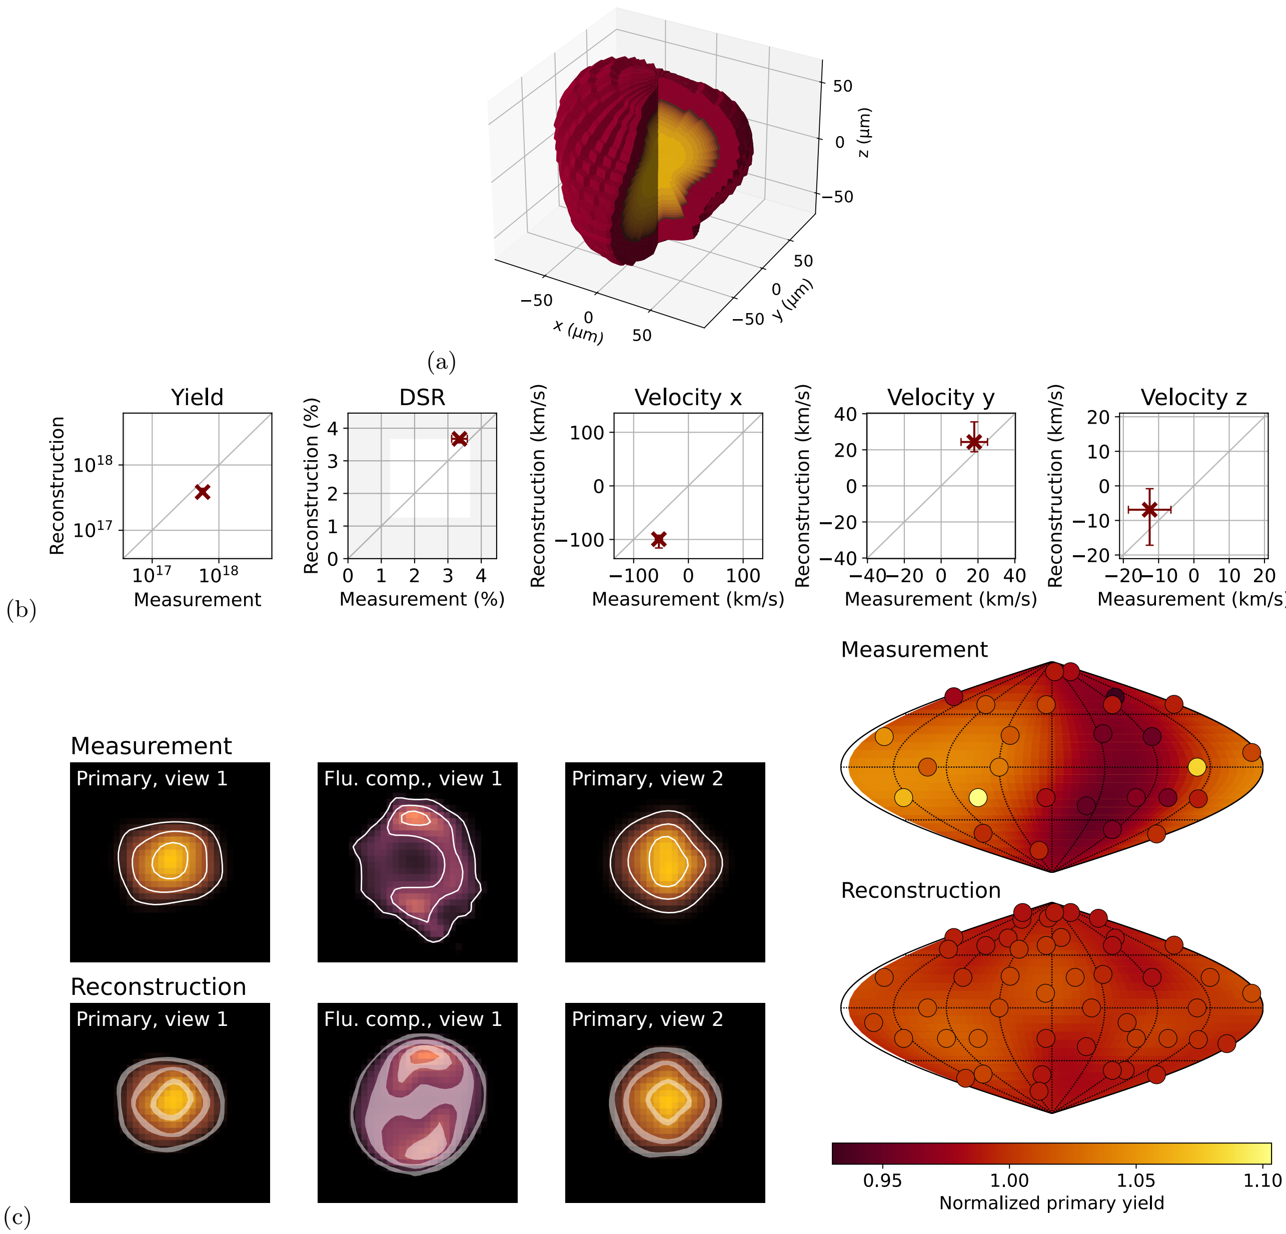

FIG. 8. Comparison of experimental data from June 28, 2024, with the reconstructed morphology and data. (a) The reconstructed 3D morphology of the hot-spot and dense fuel. (b) Comparisons of measured and reconstructed scalar data. The X marks the solution that maximizes $p(\mathbf{x}|\mathbf{d})$, the horizontal error bars are the reported error of the measurement, and the vertical error bars represent the root-mean-square spread in the posterior in each direction. All are decently matched except the x-component of $\vec{\mathbf{v}}_{\mathrm{HS}}$. (c) Comparisons of measured (top) and reconstructed (bottom) images and RTNAD sky-maps. Each image displays its 17%, 50%, and 83% contours (the width of the reconstructed contour illustrates the fit's uncertainty). The reconstructions are somewhat more symmetric than the data, but the overall structure of the images is well reconstructed.

B. Application to experimental data

Let's explore a more scientifically interesting case – real experimental data. The shot in question is the DT implosion carried out at NIF on June 28, 2024. This implosion produced 6×10^{17} neutrons (1.8 MJ), and had a target gain of about 0.8. The available data, shown in Figure 8, includes the yield, the DSR, $\vec{\mathbf{v}}_{\mathrm{HS}}$, 27 RTNAD values, two primary neutron images, and one down-scattered neutron image.

In this case, a perfect fit to the data is not expected, and thus the uncertainty quantification is more important. Figure 8a shows the highest-posterior morphology for the implosion given all the observed data. The corresponding synthetic scalar data are compared to the observations in Figure 8b, and the synthetic images and RTNAD sky-map are compared to the observations in Figure 8c.

Several features of the data are missing or distorted in the reconstruction. Notably, the reconstructed fluence-compensated down-scattered neutron image has a weaker left-right asymmetry than the observed image, the reconstructed primary neutron images have more traces of mid-mode asymmetry, the reconstructed RTNAD sky-map shows less overall asymmetry (the asymmetry in the data may be overestimated due to the large amount of random scatter in the RTNAD data on this shot), the reconstructed x-component of the hot-spot velocity is higher, and the observed z-component of $\vec{\mathbf{v}}_{\mathrm{HS}}$ is only recovered in some of the reconstructions (indicated by the vertical error bars in Figure 8.b). However, overall, the observations are reproduced quite well by the rocket-piston simulations. The reconstruction closely matches the yield and DSR, the direction of $\vec{\mathbf{v}}_{\mathrm{HS}}$, the low-mode asymmetries in the primary neutron images, and the bright caps in the fluence-compensated down-scattered neutron image. Overall, the fact that the reconstruction reproduces the broad features of the data gives confidence that this is a useful, though not perfect, representation of the actual implosion.

C. Quantifying goodness of fit

To more quantitatively assess the accuracy of the reconstruction, we calculate the normalized quadratic sum of residuals χ^2/n:

$$\chi^2/n = \frac{1}{n}\sum_{k=1}^{n}\frac{\left(d_k - f_k(\mathbf{x}^*)\right)^2}{\epsilon_k^2} \tag{7}$$

Here, $\mathbf{f}(\mathbf{x})$ represents the rocket-piston model, and $\mathbf{x}^*$ is highest-posterior input vector found by Markov chain Monte Carlo. The sum must only be carried out over the first n components of $\mathbf{f}(\mathbf{x})$ that relate to scalars, vectors, and spherical harmonic fits, excluding images (in the case of the typical vector described in Table II, $n = 13$). This is because the way the image uncertainty has been treated – as an independent Gaussian process on each pixel – is inaccurate. The errors of adjacent pixels are actually strongly positively correlated. This is not a problem in the fitting because the value of ϵ_k for each pixel can be tuned to compensate for this, such that the reconstructed image error is comparable to the reconstructed scalar error. However, this tuning does not preserve the statistical meaning of χ^2/n. Thus, without a more complex understanding of the image uncertainties, this statistical test is only valid for scalars and RTNAD modes.

Nevertheless, the 13 non-image outputs provide a sound basis for assessing the fit quality. If $\mathbf{x}^*$ accurately represents the implosion and Equation 5 accurately describes the likelihood of any given measurement error, then χ^2/n should be about 1. In reality, for the example shown in Figure 8, $\chi^2/n = 22.3$. This large value indicates statistically significant discrepancies between the best fit and the data. These can be attributed to two main sources.

TABLE III. χ^2/n statistics for reconstructions of three experiments: the simulated one shown in Figure 7, the 2024 one shown in Figure 8, and the 2022 one shown in Figure 9. The top row represents the actual amount of error, and the difference between the two rows estimates the amount of error in the surrogate model.

	Simulated	2024-06-28	2022-09-19
Rocket-piston χ^2/n	8.63	22.27	23.46
Surrogate $\tilde{\chi}^2/n$	0.54	9.69	3.67
Difference	8.09	12.58	19.79

The first source of error is deviation of the surrogate model from the rocket-piston model. While the surrogate model matches the rocket-piston model well over most of parameter space, large deviations in small parts of parameter space can cause large discrepancies between the reconstructed synthetic data and reality. This can be quantified by using the surrogate model to calculate the residuals instead of the rocket-piston model. For the example in Figure 8, this is $\tilde{\chi}^2/n = 9.7$. The difference between this and the true χ^2/n provides a good estimate of the amount of error in the surrogate model. For this shot, the difference is 12.6, which indicates that accounting for errors in the surrogate model increased the error by 12.6 (more than the amount of error present in the fit itself). Improving the surrogate model (by increasing the size of the networks, increasing the amount of training data, or increasing the number of training iterations) would reduce this difference, and thus possibly reduce the actual error.

The other source of error is limitations in the model. With only 29 input parameters and several strong physics assumptions, the rocket-piston model can only get so close to the true implosion morphology. This component of the error could be reduced by increasing the number of free input parameters, though this would also increase the amount of computational resources needed for the sampling and increase the surrogate model error.

These statistics provide a useful way to compare different implosions and assess the success of the algorithm. Table III shows χ^2/n for reconstructions of three experiments: the simulated one that is shown in Figure 7, and the one from June 28, 2024, which is shown in Figure 8, and one from September 4, 2022, which is shown in Figure 9. The latter reconstruction experienced a more severe failing of the surrogate model, producing synthetic images that look very different from the corresponding experimental images.

The χ^2/n statistics reflect this. The synthetic data re-

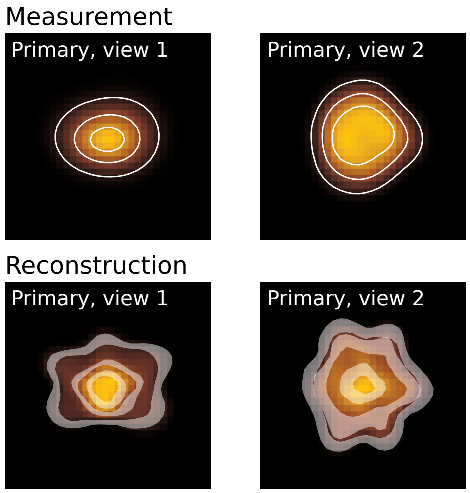

FIG. 9. Comparison of one of the images from the experiment of September 9, 2022, with the corresponding reconstructed image. Due to large discrepancies between the surrogate model and rocket-piston model, the reconstruction looks very different from the original data, indicating a failed fit.

construction has the lowest χ^2/n, as it is the most accurate fit. Error in the surrogate model brought $\tilde{\chi}^2/n$ down to 0.5 but raised the true χ^2/n to 8.6. The 2024 reconstruction, which was noticeably imperfect, has a higher χ^2/n at 22.3. Finally, the 2022 reconstruction, which can be seen by eye to be the worst match of the three, has the highest χ^2/n of 24.5. The failure of the surrogate model played a large role in the 2022 reconstruction, as is apparent from the fact that it has a bigger difference between χ^2/n and $\tilde{\chi}^2/n$ than either of the other two. This demonstrates that χ^2/n is an effective way to compare the quality of reconstructions, but also indicates that judging the synthetic data by eye is a good way to approximately determine when the algorithm has successfully found a good fit.

D. Inferred residual kinetic energy

At this point, the 3D density distribution of the shell can be calculated. Once one has a rocket-piston simulation that describes an implosion, it is straightforward to calculate the areal density ρR of each piston at the time of minimum volume. This allows the weighted harmonic mean of the areal density, ρR_{WHM}, to be calculated. When normalized to the arithmetic mean of the areal density, ρR_{mean}, this is a measure of 3D asymmetry and is related to RKE.[11] For the 2024 shot, the ratio $\rho R_{\mathrm{WHM}}/\rho R_{\mathrm{mean}}$ was found to be 95.06% $\pm$ 0.59%.

In addition to these averaged quantities, the areal density distribution can also be decomposed into spherical harmonic modes. The resulting modal decomposition of ρR is shown in Figure 10. This provides a quantitative way to assess the impact of individual asymmetry modes. In particular, it suggests that – in addition to mode $\langle 2, 0\rangle$, which is known to be prevalent and is always tuned for new experiment designs – mode $\langle 3, 2\rangle$ was especially impactful in this particular implosion.

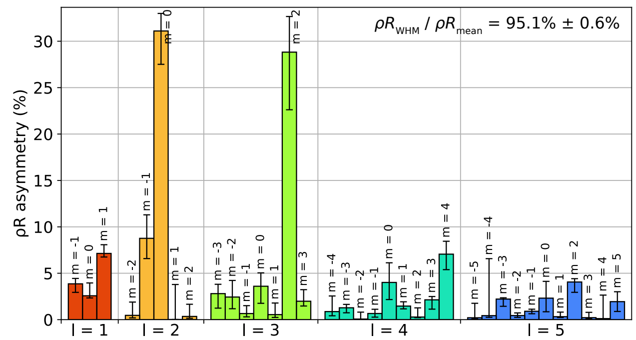

FIG. 10. The inferred perturbations from an isotropic areal density distribution in the experiment of June 28, 2024. The reciprocal of areal density in each piston is used to fit a sum of spherical harmonic functions; each bar shows the magnitude of one mode.

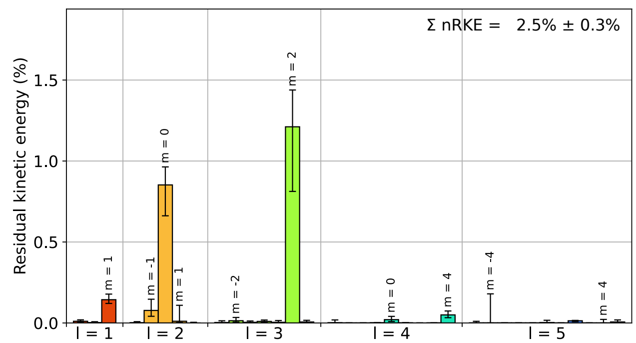

FIG. 11. The inferred distribution of normalized residual kinetic energy in the experiment of June 28, 2024. The velocity distribution of the shell is decomposed into spherical harmonic functions; each bar shows the energy stored in one mode.

By examining the piston velocities rather than their ρRs, one can also directly calculate RKE. The quantity of interest is the total RKE at the time of minimum volume. This quantity is divided by the peak kinetic energy during the implosion, yielding the normalized RKE (nRKE). On this shot, the nRKE was found to be 1.98% $\pm$ 0.45%.

Like ρR, this can be decomposed into individual modes, as shown in Figure 11. This presentation tends to emphasize the most significant modes – here, it especially highlights the importance of mode $\langle 3, 2\rangle$, which trapped about 1% of the total implosion energy in the shell. While this may seem small, rocket-piston simulations suggest that removing the $\langle 3, 2\rangle$ perturbation would have increased the yield by 53%. Work is in progress to validate the importance of this mode in this experiment, and determine what seeded this mode if so.

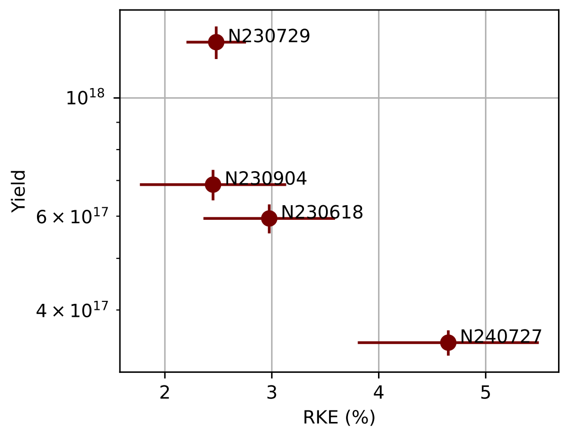

FIG. 12. The relationship between total inferred nRKE, as a percentage of the total implosion energy, and measured neutron yield for four NIF experiments with the same laser configuration.

This technique was applied to several of the high-yield experiments performed at the NIF. In Figure 12, results are illustrated for four experiments that used the same 2.05 MJ laser pulse. The inferred nRKE ranges from 2 to 5%. While the uncertainty in this inference varies greatly due to the fact that some shots have more data than others, there is a clear inverse correlation between the inferred nRKE in an implosion and the measured yield, as is expected from simulations and first principles. This gives us confidence that the method is working correctly. Future work will examine such trends in more detail to obtain new insights about implosion design.

It is worth noting that these numbers are comparable to those seen in 3D radiation–hydrodynamic simulations by Nora et al. Burn-off simulations of high-yield NIF implosions with experimentally typical mode-1 asymmetries exhibit nRKE up to about 2%, and simulations with typical mode-2 asymmetries exhibit nRKE up to about 8%.[28]

VI. FUTURE WORK

This technique is in some respects more limited than tomographic algorithms because the results are subject to the constraints imposed by the model. In other respects, it is far more general, as it can be used to combine all kinds of data rather than being limited to just images of a single type of radiation. And despite the constraints, it is able to find good matches between synthetic and experimental data. However, there are certain experiments where it fails to find good fits, so more work remains to further assess the method's robustness, and to apply it to a wider range of NIF experiments. This will enable it to extract new insights about the origins and role of 3D asymmetries in implosion performance.

The main limitations of this technique are currently the fidelity with which the rocket-piston model describes physical reality, and the fidelity with which the surrogate model imitates the rocket-piston model. Thus, modifications to this technique that address those limitations are promising areas of future study that could provide more accurate reconstructions for a wider diversity of shots.

The fidelity of the rocket-piston model can be improved by incorporating additional physics that are currently ignored, such as non-radial flows. Further work will be spent on determining what physics are needed to explain all features observed in the data, and refining the rocket-piston model to include them. Such refinement will likely need to be done continuously as experiments at the NIF achieve higher yields and push into new experimental regimes. More detailed comparisons of the model to codes like HYDRA will be performed to guide this effort.

The results of the surrogate model also occasionally diverge significantly from the results of the rocket-piston model. This can cause the sampling algorithm to think that it has found a set of input parameters that fit the data well, even though the synthetic data generated by the rocket-piston model with those same inputs do not match the data at all. This makes it very important that the rocket-piston model be used to check the validity of the solutions obtained with the surrogate model.

The fidelity of the surrogate model can be improved in several ways. One is to increase the amount of training data and the number of training iterations used for each neural network, though both would increase the amount of computational resources needed to generate the surrogate model. Another is increasing the numbers of nodes and layers in the surrogate model's neural networks, though this would also require more computational resources, and require more training data to avoid over-fitting.

One method that could improve the surrogate model's fidelity without increasing computation times would be to optimize the distribution of the training data in parameter space. The current method filters training samples by yield and DSR to concentrate points on the ignition cliff, but does not manage the distribution of points in any other dimensions. This can leave gaps in the training set, as seen in Figure 13. A more advanced sampling algorithm, such as the Delaunay-triangulation-based algorithm developed by Gaffney et al.,[30] would calculate the spacings of sampled points in all dimensions, identify gaps in the output space, and generate new training points in those gaps. Such an approach would yield training datasets with broader and more even coverage of the output space, which would yield surrogate models that are more reliably faithful to the rocket-piston model despite training on the same number of samples.

Another possibility is to remove the surrogate model altogether. If the rocket-piston model could be rewritten in a way that is fast and automatically differentiable, then it could be used in the sampling rather than the

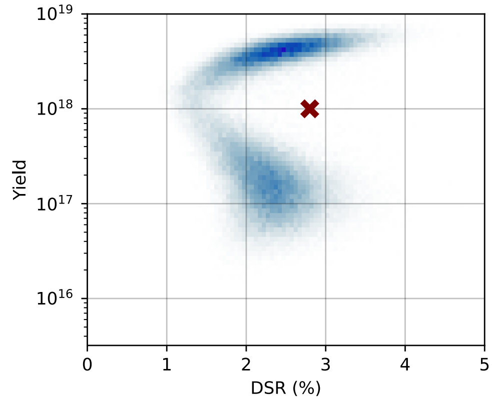

FIG. 13. Simulated yields versus DSR using the rocket-piston model that have been used to train the surrogate model. A red X marks the observed yield (10^{18} neutrons) and DSR (2.8%) measured on the milestone experiment on December 4, 2022.[29]

surrogate model. This would reduce error in the reconstruction and remove the need to spend time generating a synthetic dataset or training the surrogate model. However, whether the rocket-piston model can be sped up enough to make this approach viable is uncertain.

Future work will also include benchmarking the algorithm to ensure its accuracy. To do this, we will run the model on synthetic data generated by high-fidelity 3D hydrodynamic simulations. The simulations will provide a clear ground truth to compare the results to, to test for systematic errors in the sampling algorithm. Using a full hydrodynamic code to generate the data rather than the rocket-piston model will also enable us to test for systematic errors due to modeling assumptions.

As the limitations of this model are addressed and confidence is built in its results, it will become possible to fit the full range of DT experiments at NIF, which will make it possible to study RKE and asymmetries more holistically than has been done in the past. Eventually, this technique will be used to analyze every shot at the NIF to guide efforts to correct for and eliminate asymmetries.

This technique will also be generalized to other ICF applications. While in many cases, it will not provide benefits beyond what conventional analysis techniques can do, it has great potential in any situation where multiple diagnostics that probe related physics are used to obtain different types of data (scalars, vectors, images, or time-series) that cannot be combined using traditional techniques. For example, gamma-ray[3] and x-ray[31] diagnostics can be used instead of or in addition to neutron diagnostics. This should require little modification to what is presented here—the step in the model that converts the implosion morphology to synthetic data must be altered such that it generates x-ray images instead of neutron images. X-ray images would be generated in much the same way primary neutron images are generated now, but with a different emission profile. This framework will also be applied to future facilities, to determine how much information will be available from the planned suite of diagnostics and where those diagnostics should be placed to maximize their utility.

VII. CONCLUSION

3D asymmetries are a major performance-degradation mechanism in ICF implosions at the NIF. There is a large amount of information available about these asymmetries from various diagnostics on multiple lines of sight, but currently no method to integrate them into a single self-consistent reconstruction. In particular, the neutron imaging systems are used to reconstruct the 3D morphology of the hot-spot and dense fuel of implosions at the NIF. However, the reconstruction problem is ill-posed with only three imaging lines of sight. The RTNADs and neutron spectrometers are also used to reconstruct low-mode asymmetries in the hot-spot and dense fuel, but these reconstructions require assumptions about the implosion symmetry and are not sensitive to higher-mode asymmetries. Thus, it is desirable to combine information from all three diagnostics in a mathematically rigorous way, and in a way that quantifies uncertainty.

To that end, a new methodology has been developed that makes it possible to combine neutron images, RTNAD data, and neutron spectra to reconstruct the 3D morphology of the hot-spot and surrounding dense shell of a NIF implosion, and to infer the amount of RKE contained therein. This algorithm uses a multi-rocket-piston model as a basis to define the set of possible implosions, and uses Markov chain Monte Carlo to find the sets of model inputs that generate synthetic data consistent with the observed experimental data. A neural-network-based surrogate model is trained to mimic the rocket-piston model and used in its stead during most of this process to speed the sampling up by multiple orders of magnitude.

The result is a 3D model of the implosion with uncertainties quantified according to the amount and quality of the data. The model has been tested on both synthetic and experimental data, and provides good reconstructions in both cases. The nRKE is inferred from these reconstructions; it is fond to range from 2% to 5% in several NIF experiments, and it is found to vary inversely with the implosion yield, as expected. This methodology provides a framework to guide efforts to investigate and address mid-mode asymmetries at the NIF going forward.

ACKNOWLEDGMENTS

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. This work was also supported in part by the U.S. Department of Energy NNSA MIT Center-of-Excellence under Contract DE-NA0003868, and by the NNSA Laboratory Residency Graduate Fellowship under Contract DE-NA0003960. D. T. Casey and T. M. Johnson acknowledge support from the U.S. Department of Energy Early Career Research Program through the Office of Fusion Energy Sciences.

This document has been reviewed for release as document LLNL-JRNL-2008600.

AUTHOR DECLARATIONS

Conflicts of interest

The authors have no conflicts to disclose.

Author contributions

J. H. Kunimune: investigation (lead), writing (lead). D. T. Casey: investigation (supporting), writing (supporting), resources (equal), supervision (lead). B. Kustowski: investigation (supporting). M. Jones: investigation (supporting). L. Divol: investigation (supporting). T. M. Johnson: resources (equal). S. G. Dannhoff: resources (equal). A. DeVault: resources (equal). J. A. Frenje: supervision (supporting).

Data availability statement

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

REFERENCES

[1] S. Atzeni and J. M. ter Vehn, *The Physics of Inertial Fusion: Beam–Plasma Interaction, Hydrodynamics, Hot Dense Matter*, Oxford science publications (Clarendon Press, Oxford, 2004).

[2] A. L. Kritcher, A. B. Zylstra, C. R. Weber, O. A. Hurricane, D. A. Callahan, D. S. Clark, L. Divol, D. E. Hinkel, K. Humbird, O. Jones, J. D. Lindl, S. Maclaren, D. J. Strozzi, C. V. Young, A. Allen, B. Bachmann, K. L. Baker, T. Braun, G. Brunton, D. T. Casey, T. Chapman, C. Choate, E. Dewald, J.-M. G. Di Nicola, M. J. Edwards, S. Haan, T. Fehrenbach, M. Hohenberger, E. Kur, B. Kustowski, C. Kong, O. L. Landen, D. Larson, B. J. MacGowan, M. Marinak, M. Millot, A. Nikroo, R. Nora, A. Pak, P. K. Patel, J. E. Ralph, M. Ratledge, M. S. Rubery, D. J. Schlossberg, S. M. Sepke, M. Stadermann, T. I. Suratwala, R. Tommasini, R. Town, B. Woodworth, B. Van Wonterghem, and C. Wild, "Design of the first fusion experiment to achieve target energy gain $G > 1$," Phys. Rev. E **109**, 025204 (2024).

[3] F. E. Merrill, D. Bower, R. Buckles, D. D. Clark, C. R. Danly, O. B. Drury, J. M. Dzenitis, V. E. Fatherley, D. N. Fittinghoff, R. Gallegos, G. P. Grim, N. Guler, E. N. Loomis, S. Lutz, R. M. Malone, D. D. Martinson, D. Mares, D. J. Morley, G. L. Morgan, J. A. Oertel, I. L. Tregillis, P. L. Volegov, P. B. Weiss, C. H. Wilde, and D. C. Wilson, "The neutron imaging diagnostic at NIF," Rev. Sci. Instrum. **83**, 10D317 (2012).

[4] G. P. Grim, N. Guler, F. E. Merrill, G. L. Morgan, C. R. Danly, P. L. Volegov, C. H. Wilde, D. C. Wilson, D. S. Clark, D. E. Hinkel, O. S. Jones, K. S. Raman, N. Izumi, D. N. Fittinghoff, O. B. Drury, E. T. Alger, P. A. Arnold, R. C. Ashabranner, L. J. Atherton, M. A. Barrios, S. Batha, P. M. Bell, L. R. Benedetti, R. L. Berger, L. A. Bernstein, L. V. Berzins, R. Betti, S. D. Bhandarkar, R. M. Bionta, D. L. Bleuel, T. R. Boehly, E. J. Bond, M. W. Bowers, D. K. Bradley, G. K. Brunton, R. A. Buckles, S. C. Burkhart, R. F. Burr, J. A. Caggiano, D. A. Callahan, D. T. Casey, C. Castro, P. M. Celliers, C. J. Cerjan, G. A. Chandler, C. Choate, S. J. Cohen, G. W. Collins, G. W. Cooper, J. R. Cox, J. R. Cradick, P. S. Datte, E. L. Dewald, P. Di Nicola, J. M. Di Nicola, L. Divol, S. N. Dixit, R. Dylla-Spears, E. G. Dzenitis, M. J. Eckart, D. C. Eder, D. H. Edgell, M. J. Edwards, J. H. Eggert, R. B. Ehrlich, G. V. Erbert, J. Fair, D. R. Farley, B. Felker, R. J. Fortner, J. A. Frenje, G. Frieders, S. Friedrich, M. Gatu-Johnson, C. R. Gibson, E. Giraldez, V. Y. Glebov, S. M. Glenn, S. H. Glenzer, G. Gururangan, S. W. Haan, K. D. Hahn, B. A. Hammel, A. V. Hamza, E. P. Hartouni, R. Hatarik, S. P. Hatchett, C. Haynam, M. R. Hermann, H. W. Herrmann, D. G. Hicks, J. P. Holder, D. M. Holunga, J. B. Horner, W. W. Hsing, H. Huang, M. C. Jackson, K. S. Jancaitis, D. H. Kalantar, R. L. Kauffman, M. I. Kauffman, S. F. Khan, J. D. Kilkenny, J. R. Kimbrough, R. Kirkwood, J. L. Kline, J. P. Knauer, K. M. Knittel, J. A. Koch, T. R. Kohut, B. J. Kozioziemski, K. Krauter, G. W. Krauter, A. L. Kritcher, J. Kroll, G. A. Kyrala, K. N. L. Fortune, G. LaCaille, L. J. Lagin, T. A. Land, O. L. Landen, D. W. Larson, D. A. Latray, R. J. Leeper, T. L. Lewis, S. LePape, J. D. Lindl, R. R. Lowe-Webb, T. Ma, B. J. MacGowan, A. J. MacKinnon, A. G. MacPhee, R. M. Malone, T. N. Malsbury, E. Mapoles, C. D. Marshall, D. G. Mathisen, P. McKenty, J. M. McNaney, N. B. Meezan, P. Michel, J. L. Milovich, J. D. Moody, A. S. Moore, M. J. Moran, K. Moreno, E. I. Moses, D. H. Munro, B. R. Nathan, A. J. Nelson, A. Nikroo, R. E. Olson, C. Orth, A. E. Pak, E. S. Palma, T. G. Parham, P. K. Patel, R. W. Patterson, R. D. Petrasso, R. Prasad, J. E. Ralph, S. P. Regan, H. Rinderknecht, H. F. Robey, G. F. Ross, C. L. Ruiz, F. H. Séguin, J. D. Salmonson, T. C. Sangster, J. D. Sater, R. L. Saunders, M. B. Schneider, D. H. Schneider, M. J. Shaw, N. Simanovskaia, B. K. Spears, P. T. Springer, C. Stoeckl, W. Stoeffl, L. J. Suter, C. A. Thomas, R. Tommasini, R. P. Town, A. J. Traille, B. V. Wonterghem, R. J. Wallace, S. Weaver, S. V. Weber, P. J. Wegner, P. K. Whitman, K. Widmann, C. C. Widmayer, R. D. Wood, B. K. Young, R. A. Zacharias, and A. Zylstra, "Nuclear imaging of the fuel assembly in ignition experiments," Phys. Plasmas **20**, 056320 (2013).

[5] D. T. Casey, P. L. Volegov, F. E. Merrill, D. H. Munro, G. P. Grim, O. L. Landen, B. K. Spears, D. N. Fittinghoff, J. E. Field, and V. A. Smalyuk, "Fluence-compensated down-scattered neutron imaging using the neutron imaging system at the National

Ignition Facility," Rev. Sci. Instrum. **87**, 11E715 (2016).

[6] P. L. Volegov, S. H. Batha, V. Geppert-Kleinrath, C. R. Danly, F. E. Merrill, C. H. Wilde, D. C. Wilson, D. T. Casey, D. Fittinghoff, B. Appelbe, J. P. Chittenden, A. J. Crilly, and K. McGlinchey, "Density determination of the thermonuclear fuel region in inertial confinement fusion implosions," J. Appl. Phys. **127**, 083301 (2020).

[7] P. L. Volegov, C. R. Danly, F. E. Merrill, R. Simpson, and C. H. Wilde, "On three-dimensional reconstruction of a neutron/x-ray source from very few two-dimensional projections," J. Appl. Phys. **118**, 205903 (2015).

[8] R. M. Bionta, G. P. Grim, K. D. Hahn, E. P. Hartouni, E. A. Henry, H. Y. Khater, A. S. Moore, and D. J. Schlossberg, "Real-time nuclear activation detectors for measuring neutron angular distributions at the National Ignition Facility," Rev. Sci. Instrum. **92**, 043527 (2021).

[9] A. S. Moore, D. J. Schlossberg, B. D. Appelbe, G. A. Chandler, A. J. Crilly, M. J. Eckart, C. J. Forrest, V. Y. Glebov, G. P. Grim, E. P. Hartouni, R. Hatarik, S. M. Kerr, J. Kilkenny, and J. P. Knauer, "Neutron time of flight (nToF) detectors for inertial fusion experiments," Rev. Sci. Instrum. **94**, 061102 (2023).

[10] O. A. Hurricane, D. T. Casey, O. Landen, A. L. Kritcher, R. Nora, P. K. Patel, J. A. Gaffney, K. D. Humbird, J. E. Field, M. K. G. Kruse, J. L. Peterson, and B. K. Spears, "An analytic asymmetric-piston model for the impact of mode-1 shell asymmetry on ICF implosions," Phys. Plasmas **27**, 062704 (2020).

[11] O. A. Hurricane, D. T. Casey, O. Landen, D. A. Callahan, R. Bionta, S. Haan, A. L. Kritcher, R. Nora, P. K. Patel, P. T. Springer, and A. Zylstra, "Extensions of a classical mechanics 'piston-model' for understanding the impact of asymmetry on ICF implosions: The cases of mode 2, mode 2/1 coupling, time-dependent asymmetry, and the relationship to coast-time," Phys. Plasmas **29**, 012703 (2022).

[12] J. H. Kunimune, D. T. Casey, B. Kustowski, V. Geppert-Kleinrath, L. Divol, D. N. Fittinghoff, P. L. Volegov, M. K. G. Kruse, J. A. Gaffney, R. C. Nora, and J. A. Frenje, "3D reconstruction of an inertial-confinement fusion implosion with neural networks using multiple heterogeneous data sources," Rev. Sci. Instrum. **95**, 073506 (2024).

[13] D. T. Casey, J. Kunimune, O. A. Hurricane, O. L. Landen, P. Springer, R. M. Bionta, C. V. Young, R. C. Nora, B. J. MacGowan, J. A. Gaffney, B. Kustowski, C. Weber, A. Kritcher, J. Milovich, S. Haan, M. G. Johnson, D. Schlossberg, S. Kerr, P. L. Volegov, D. N. Fittinghoff, V. Geppert-Kleinrath, C. H. Wilde, and M. Freeman, "A multi-rocket piston model to study three-dimensional asymmetries in implosions at the National Ignition Facility," High Energy Density Phys. **54**, 101172 (2025).

[14] D. A. Callahan, O. A. Hurricane, A. L. Kritcher, D. T. Casey, D. E. Hinkel, Y. P. Opachich, H. F. Robey, M. D. Rosen, J. S. Ross, M. S. Rubery, C. V. Young, and A. B. Zylstra, "A simple model to scope out parameter space for indirect drive designs on NIF," Phys. Plasmas **27**, 072704 (2020).

[15] R. E. Olson, G. A. Rochau, O. L. Landen, and R. J. Leeper, "X-ray ablation rates in inertial confinement fusion capsule materials," Phys. Plasmas **18**, 032706 (2011).

[16] P. T. Springer, O. A. Hurricane, J. H. Hammer, R. Betti, D. A. Callahan, E. M. Campbell, D. T. Casey, C. J. Cerjan, D. Cao, E. Dewald, L. Divol, T. Doeppner, M. J. Edwards, J. E. Field, C. Forrest, J. Frenje, J. A. Gaffney, M. Gatu-Johnson, V. Glebov, V. N. Goncharov, G. P. Grim, E. Hartouni, R. Hatarik, D. E. Hinkel, L. B. Hopkins, I. Igumenshchev, P. Knapp, J. P. Knauer, A. L. Kritcher, O. Landen, A. Pak, S. L. Pape, T. Ma, A. G. MacPhee, D. H. Munro, R. C. Nora, P. K. Patel, L. Peterson, P. B. Radha, S. P. Regan, H. Rinderknecht, C. Sangster, B. K. Spears, and C. Stoeckl, "A 3d dynamic model to assess the impacts of low-mode asymmetry, aneurysms and mix-induced radiative loss on capsule performance across inertial confinement fusion platforms," Nucl. Fusion **59**, 032009 (2018).

[17] S. H. Langer, I. Karlin, and M. M. Marinak, "Performance characteristics of HYDRA – a multi-physics simulation code from LLNL," in *High Perform. Comput. Comput. Sci.*, edited by M. Daydé, O. Marques, and K. Nakajima (Springer International Publishing, Cham, 2015) pp. 173–181.

[18] D. A. Callahan, N. B. Meezan, S. H. Glenzer, A. J. MacKinnon, L. R. Benedetti, D. K. Bradley, J. R. Celeste, P. M. Celliers, S. N. Dixit, T. Döppner, E. G. Dzentitis, S. Glenn, S. W. Haan, C. A. Haynam, D. G. Hicks, D. E. Hinkel, O. S. Jones, O. L. Landen, R. A. London, A. G. MacPhee, P. A. Michel, J. D. Moody, J. E. Ralph, H. F. Robey, M. D. Rosen, M. B. Schneider, D. J. Strozzi, L. J. Suter, R. P. J. Town, K. Widmann, E. A. Williams, M. J. Edwards, B. J. MacGowan, J. D. Lindl, L. J. Atherton, G. A. Kyrala, J. L. Kline, R. E. Olson, D. Edgell, S. P. Regan, A. Nikroo, H. Wilkins, J. D. Kilkenny, and A. S. Moore, "The velocity campaign for ignition on NIF," Phys. Plasmas **19**, 056305 (2012).

[19] J. H. Kunimune, D. T. Casey, and J. A. Frenje, "An ultra-fast method for generating synthetic down-scattered neutron data for inertial confinement fusion implosions," (2025), unpublished.

[20] J. S. Lim, *Two-dimensional signal and image processing* (Prentice Hall PTR, 1990) p. 29–30.

[21] P. Volegov, C. R. Danly, D. N. Fittinghoff, G. P. Grim, N. Guler, N. Izumi, T. Ma, F. E. Merrill, A. L. Warrick, C. H. Wilde, and D. C. Wilson, "Neutron source reconstruction from pinhole imaging at National Ignition Facility," Rev. Sci. Instrum. **85**, 023508 (2014).

[22] D. P. Kingma and J. Ba, "Adam: a method for stochastic optimization," (2017).

[23] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth, "Hybrid Monte Carlo," Phys. Lett. B **195**, 216–222 (1987).

[24] R. H. Swendsen and J.-S. Wang, "Replica Monte Carlo simulation of spin-glasses," Phys. Rev. Lett. **57**, 2607–2609 (1986).

[25] J. V. Dillon, I. Langmore, D. Tran, E. Brevdo, S. Vasudevan, D. Moore, B. Patton, A. Alemi, M. Hoffman, and R. A. Saurous, "TensorFlow Distributions," (2017), arXiv:1711.10604 [cs.LG].

[26] R. Anirudh, J. J. Thiagarajan, P.-T. Bremer, and B. K. Spears, "Improved surrogates in inertial confinement fusion with manifold and cycle consistencies," Proc. Natl. Acad. Sci. **117**, 9741–9746 (2020).

[27] B. Kustowski, J. A. Gaffney, B. K. Spears, G. J. Anderson, R. Anirudh, P.-T. Bremer, J. J. Thiagarajan, M. K. G. Kruse, and R. C. Nora, "Suppressing simulation bias in multi-modal data using transfer learning," Mach. Learn.: Sci. Technol. **3**, 015035–015052 (2022).

[28] R. C. Nora, N. Birge, D. Casey, C. Danly, E. L. Dewald, B. Z. Djordjevic, A. Do, M. Durocher, J. E. Field, D. Fittinghoff, M. S. Freeman, J. Gaffney, V. Geppert Kleinrath, S. Haan, K. Hahn, E. Hartouni, M. Hohenberger, S. Kerr, O. L. Landen, J. Milovich, A. S. Moore, P. Patel, D. J. Schlossberg, S. M. Sepke, B. K. Spears, P. L. Volegov, and C. Wilde, "Correlations between asymmetric compression, burn amplification, and hot-spot velocities in inertial confinement fusion implosions," Physics of Plasmas **30**, 092705 (2023), https://pubs.aip.org/aip/pop/article-pdf/doi/10.1063/5.0153421/18131669/092705_1_5.0153421.pdf.

[29] H. Abu-Shawareb, R. Acree, P. Adams, J. Adams, B. Addis, R. Aden, P. Adrian, B. B. Afeyan, M. Aggleton, L. Aghaian, A. Aguirre, D. Aikens, J. Akre, F. Albert, M. Albrecht, B. J. Albright, J. Albritton, J. Alcala, C. Alday, D. A. Alessi, N. Alexander, J. Alfonso, N. Alfonso, E. Alger, S. J. Ali, Z. A. Ali, A. Allen, W. E. Alley, P. Amala, P. A. Amendt, P. Amick, S. Ammula, C. Amorin, D. J. Ampleford, R. W. Anderson, T. Anklam, N. Antipa, B. Appelbe, C. Aracne-Ruddle, E. Araya, T. N. Archuleta, M. Arend, P. Arnold, T. Arnold, A. Arsenlis, J. Asay, L. J. Atherton, D. Atkinson, R. Atkinson, J. M. Auerbach, B. Austin, L. Auyang, A. A. S. Awwal, N. Aybar, J. Ayers, S. Ayers, T. Ayers, S. Azevedo, B. Bachmann, C. A. Back, J. Bae, D. S. Bailey, J. Bailey, T. Baisden, K. L. Baker, H. Baldis, D. Barber, M. Barberis, D. Barker, A. Barnes, C. W. Barnes, M. A. Barrios, C. Barty, I. Bass, S. H. Batha, S. H. Baxamusa, G. Bazan, J. K. Beagle, R. Beale, B. R. Beck, J. B. Beck, M. Bedzyk, R. G. Beeler, R. G. Beeler, W. Behrendt, L. Belk, P. Bell, M. Belyaev,

J. F. Benage, G. Bennett, L. R. Benedetti, L. X. Benedict, R. L. Berger, T. Bernat, L. A. Bernstein, B. Berry, L. Bertolini, G. Besenbruch, J. Betcher, R. Bettenhausen, R. Betti, B. Bezzerides, S. D. Bhandarkar, R. Bickel, J. Biener, T. Biesiada, K. Bigelow, J. Bigelow-Granillo, V. Bigman, R. M. Bionta, N. W. Birge, M. Bitter, A. C. Black, R. Bleile, D. L. Bleuel, E. Bliss, E. Bliss, B. Blue, T. Boehly, K. Boehm, C. D. Boley, R. Bonanno, E. J. Bond, T. Bond, M. J. Bonino, M. Borden, J.-L. Bourgade, J. Bousquet, J. Bowers, M. Bowers, R. Boyd, D. Boyle, A. Bozek, D. K. Bradley, K. S. Bradley, P. A. Bradley, L. Bradley, L. Brannon, P. S. Brantley, D. Braun, T. Braun, K. Brienza-Larsen, R. Briggs, T. M. Briggs, J. Britten, E. D. Brooks, D. Browning, M. W. Bruhn, T. A. Brunner, H. Bruns, G. Brunton, B. Bryant, T. Buczek, J. Bude, L. Buitano, S. Burkhart, J. Burmark, A. Burnham, R. Burr, L. E. Busby, B. Butlin, R. Cabeltis, M. Cable, W. H. Cabot, B. Cagadas, J. Caggiano, R. Cahayag, S. E. Caldwell, S. Calkins, D. A. Callahan, J. Calleja-Aguirre, L. Camara, D. Camp, E. M. Campbell, J. H. Campbell, B. Carey, R. Carey, K. Carlisle, L. Carlson, L. Carman, J. Carmichael, A. Carpenter, C. Carr, J. A. Carrera, D. Casavant, A. Casey, D. T. Casey, A. Castillo, E. Castillo, J. I. Castor, C. Castro, W. Caughey, R. Cavitt, J. Celeste, P. M. Celliers, C. Cerjan, G. Chandler, B. Chang, C. Chang, J. Chang, L. Chang, R. Chapman, T. D. Chapman, L. Chase, H. Chen, H. Chen, K. Chen, L.-Y. Chen, B. Cheng, J. Chittenden, C. Choate, J. Chou, R. E. Chrien, M. Chrisp, K. Christensen, M. Christensen, N. S. Christiansen, A. R. Christopherson, M. Chung, J. A. Church, A. Clark, D. S. Clark, K. Clark, R. Clark, L. Claus, B. Cline, J. A. Cline, J. A. Cobble, K. Cochrane, B. Cohen, S. Cohen, M. R. Collette, G. W. Collins, L. A. Collins, T. J. B. Collins, A. Conder, B. Conrad, M. Conyers, A. W. Cook, D. Cook, R. Cook, J. C. Cooley, G. Cooper, T. Cope, S. R. Copeland, F. Coppari, J. Cortez, J. Cox, D. H. Crandall, J. Crane, R. S. Craxton, M. Cray, A. Crilly, J. W. Crippen, D. Cross, M. Cuneo, G. Cuotts, C. E. Czajka, D. Czechowicz, T. Daly, P. Danforth, C. Danly, R. Darbee, B. Darlington, P. Datte, L. Dauffy, G. Davalos, S. Davidovits, P. Davis, J. Davis, S. Dawson, R. D. Day, T. H. Day, M. Dayton, C. Deck, C. Decker, C. Deeney, K. A. DeFriend, G. Deis, N. D. Delamater, J. A. Delettrez, R. Demaret, S. Demos, S. M. Dempsey, R. Desjardin, T. Desjardins, M. P. Desjarlais, E. L. Dewald, J. DeYoreo, S. Diaz, G. Dimonte, T. R. Dittrich, L. Divol, S. N. Dixit, J. Dixon, A. Do, E. S. Dodd, D. Dolan, A. Donovan, M. Donovan, T. Döppner, C. Dorrer, N. Dorsano, M. R. Douglas, D. Dow, J. Downie, E. Downing, M. Dozieres, V. Draggoo, D. Drake, R. P. Drake, T. Drake, G. Dreifuerst, O. Drury, D. F. DuBois, P. F. DuBois, G. Dunham, M. Durocher, R. Dylla-Spears, A. K. L. Dymoke-Bradshaw, B. Dzenitis, C. Ebbers, M. Eckart, S. Eddinger, D. Eder, D. Edgell, M. J. Edwards, P. Efthimion, J. H. Eggert, B. Ehrlich, P. Ehrmann, S. Elhadj, C. Ellerbee, N. S. Elliott, C. L. Ellison, F. Elsner, M. Emerich, K. Engelhorn, T. England, E. English, P. Epperson, R. Epstein, G. Erbert, M. A. Erickson, D. J. Erskine, A. Erlandson, R. J. Espinosa, C. Estes, K. G. Estabrook, S. Evans, A. Fabyan, J. Fair, R. Fallejo, N. Farmer, W. A. Farmer, M. Farrell, V. E. Fatherley, M. Fedorov, E. Feigenbaum, T. Fehrenbach, M. Feit, B. Felker, W. Ferguson, J. C. Fernandez, A. Fernandez-Panella, S. Fess, J. E. Field, C. V. Filip, J. R. Fincke, T. Finn, S. M. Finnegan, R. G. Finucane, M. Fischer, A. Fisher, J. Fisher, B. Fishler, D. Fittinghoff, P. Fitzsimmons, M. Flegel, K. A. Flippo, J. Florio, J. Folta, P. Folta, L. R. Foreman, C. Forrest, A. Forsman, J. Fooks, M. Foord, R. Fortner, K. Fournier, D. E. Fratanduono, N. Frazier, T. Frazier, C. Frederick, M. S. Freeman, J. Frenje, D. Frey, G. Frieders, S. Friedrich, D. H. Froula, J. Fry, T. Fuller, J. Gaffney, S. Gales, B. Le Galloudec, K. K. Le Galloudec, A. Gambhir, L. Gao, W. J. Garbett, A. Garcia, C. Gates, E. Gaut, P. Gauthier, Z. Gavin, J. Gaylord, C. G. R. Geddes, M. Geissel, F. Génin, J. Georgeson, H. Geppert-Kleinrath, V. Geppert-Kleinrath, N. Gharibyan, J. Gibson, C. Gibson, E. Giraldez, V. Glebov, S. G. Glendinning, S. Glenn, S. H. Glenzer, S. Goade, P. L. Gobby, S. R. Goldman, B. Golick, M. Gomez, V. Goncharov, D. Goodin, P. Grabowski, E. Grafil, P. Graham, J. Grandy, E. Grasz, F. R. Graziani, G. Greenman, J. A. Greenough, A. Greenwood, G. Gregori, T. Green, J. R. Griego, G. P. Grim, J. Grondalski, S. Gross, J. Guckian, N. Guler, B. Gunney, G. Guss, S. Haan, J. Hackbarth, L. Hackel, R. Hackel, C. Haefner, C. Hagmann, K. D. Hahn, S. Hahn, B. J. Haid, B. M. Haines, B. M. Hall, C. Hall, G. N. Hall, M. Hamamoto, S. Hamel, C. E. Hamilton, B. A. Hammel, J. H. Hammer, G. Hampton, A. Hamza, A. Handler, S. Hansen, D. Hanson, R. Haque, D. Harding, E. Harding, J. D. Hares, D. B. Harris, J. A. Harte, E. P. Hartouni, R. Hatarik, S. Hatchett, A. A. Hauer, M. Havre, R. Hawley, J. Hayes, J. Hayes, S. Hayes, A. Hayes-Sterbenz, C. A. Haynam, D. A. Haynes, D. Headley, A. Heal, J. E. Heebner, S. Heerey, G. M. Heestand, R. Heeter, N. Hein, C. Heinbockel, C. Hendricks, M. Henesian, J. Heninger, J. Henrikson, E. A. Henry, E. B. Herbold, M. R. Hermann, G. Hermes, J. E. Hernandez, V. J. Hernandez, M. C. Herrmann, H. W. Herrmann, O. D. Herrera, D. Hewett, R. Hibbard, D. G. Hicks, D. P. Higginson, D. Hill, K. Hill, T. Hilsabeck, D. E. Hinkel, D. D. Ho, V. K. Ho, J. K. Hoffer, N. M. Hoffman, M. Hohenberger, M. Hohensee, W. Hoke, D. Holdener, F. Holdener, J. P. Holder, B. Holko, D. Holunga, J. F. Holzrichter, J. Honig, D. Hoover, D. Hopkins, L. F. Berzak Hopkins, M. Hoppe, M. L. Hoppe, J. Horner, R. Hornung, C. J. Horsfield, J. Horvath, D. Hotaling, R. House, L. Howell, W. W. Hsing, S. X. Hu, H. Huang, J. Huckins, H. Hui, K. D. Humbird, J. Hund, J. Hunt, O. A. Hurricane, M. Hutton, K. H.-K. Huynh, L. Inandan, C. Iglesias, I. V. Igumenshchev, I. Ivanovich, N. Izumi, M. Jackson, J. Jackson, S. D. Jacobs, G. James, K. Jancaitis, J. Jarboe, L. C. Jarrott, D. Jasion, J. Jaquez, J. Jeet, A. E. Jenei, J. Jensen, J. Jimenez, R. Jimenez, D. Jobe, Z. Johal, H. M. Johns, D. Johnson, M. A. Johnson, M. Gatu Johnson, R. J. Johnson, S. Johnson, S. A. Johnson, T. Johnson, K. Jones, O. Jones, M. Jones, R. Jorge, H. J. Jorgenson, M. Julian, B. I. Jun, R. Jungquist, J. Kaae, N. Kabadi, D. Kaczala, D. Kalantar, K. Kangas, V. V. Karasiev, M. Karasik, V. Karpenko, A. Kasarky, K. Kasper, R. Kauffman, M. I. Kaufman, C. Keane, L. Keaty, L. Kegelmeyer, P. A. Keiter, P. A. Kellett, J. Kellogg, J. H. Kelly, S. Kemic, A. J. Kemp, G. E. Kemp, G. D. Kerbel, D. Kershaw, S. M. Kerr, T. J. Kessler, M. H. Key, S. F. Khan, H. Khater, C. Kiikka, J. Kilkenny, Y. Kim, Y.-J. Kim, J. Kimko, M. Kimmel, J. M. Kindel, J. King, R. K. Kirkwood, L. Klaus, D. Klem, J. L. Kline, J. Klingmann, G. Kluth, P. Knapp, J. Knauer, J. Knipping, M. Knudson, D. Kobs, J. Koch, T. Kohut, C. Kong, J. M. Koning, P. Koning, S. Konior, H. Kornblum, L. B. Kot, B. Kozioziemski, M. Kozlowski, P. M. Kozlowski, J. Krammen, N. S. Krasheninnikova, C. M. Krauland, B. Kraus, W. Krauser, J. D. Kress, A. L. Kritcher, *et al.* (The Indirect Drive ICF Collaboration), "Achievement of target gain larger than unity in an inertial fusion experiment," Phys. Rev. Lett. **132**, 065102 (2024).

[30]J. A. Gaffney, K. Humbird, M. Kruse, E. Kur, B. Kustowski, R. Nora, and B. Spears, "Iterative sampling of expensive simulations for faster deep surrogate training," Contributions to Plasma Physics **63**, e202200190 (2023), https://onlinelibrary.wiley.com/doi/pdf/10.1002/ctpp.202200190.

[31]B. Bachmann, A. L. Kritcher, L. R. Benedetti, R. W. Falcone, S. Glenn, J. Hawreliak, N. Izumi, D. Kraus, O. L. Landen, S. Le Pape, T. Ma, F. Pérez, D. Swift, and T. Döppner, "Using penumbral imaging to measure micrometer size plasma hot spots in Gbar equation of state experiments on the National Ignition Facility," Rev. Sci. Instrum. **85**, 11D614 (2014).