Inference of three-dimensional hot-spot and shell morphology in inertial confinement fusion experiments using a convolutional neural network

K. M. Woo [1,*] J. Buck [2] K. Churnetski,[1] R. Betti,[1] C. A. Thomas,[1] C. Stoeckl [1] P. V. Heuer [1] C. Kanan [3]
P. B. Radha [1] J. Carroll-Nellenback [1] K. S. Anderson [1] T. J. B. Collins [1] L. Ceurvorst,[1]
V. Gopalaswamy [1] and S. T. Ivancic [1]

[1]*Laboratory for Laser Energetics, Rochester, New York 14623, USA*
[2]*Department of Physics and Astronomy, University of Rochester, Rochester, New York 14623, USA*
[3]*Department of Computer Science, University of Rochester, Rochester, New York 14623, USA*

(Received 11 March 2024; accepted 9 May 2025; published 17 June 2025)

The performance of inertial confinement fusion (ICF) implosions is sensitive to the three-dimensional (3D) morphology of the hot-spot and shell configurations. The ability to infer shell-mass uniformity and reconstruct 3D hot spots is crucial for quantifying the degradation of ignition criteria and improving symmetry in ICF implosion experiments. In this work, we present a deep-learning convolutional neural network (CNN) for reconstructing 3D hot-spot and shell structures for ICF capsules. The 3D geometry of the hot spot is reconstructed from x-ray images measured from multiple lines of sight on OMEGA. The shell configuration is inferred indirectly through machine learning using a convolutional neural network extensively trained on a DEC3D simulation database. This simulation-dependent approach yields consistent agreement between reconstructed 3D shell densities and machine-learning optimized DEC3D simulation results. This work demonstrates a CNN framework that successfully reconstructs 3D capsule structures from two-dimensional images in ICF implosions.

DOI: 10.1103/krtw-bkc3

I. INTRODUCTION

Inertial confinement fusion (ICF) [1] represents a clean energy approach where deuterium (D) and tritium (T) nuclei fuse under extreme compression, releasing 14.1-MeV neutrons and 3.5-MeV alpha particles without generating long-lived radioactive waste. Recent laser indirect-drive experiments at the National Ignition Facility (NIF) [2] have demonstrated significant progress [3] in achieving the burning plasma condition [1]. In a plasma regime the alpha particle heating triggers a self-sustaining fusion burn wave that propagates through the compressed DT fuel. This advance brings us closer to future fusion power plants that will convert this clean energy into electricity for commercial use.

The OMEGA laser facility [4] focuses on achieving high-performance laser direct-drive ICF implosions through direct-laser illumination of spherical targets using 60 laser beams. One key mission [5] is to demonstrate the hydrodynamic-equivalent ignition that scales to NIF conditions through increased capsule dimensions and laser energy, from OMEGA's 30 kJ to NIF's 2 MJ. However, the presence of three-dimensional (3D) implosion asymmetries remains a key factor degrading OMEGA's performance [6–8].

Understanding 3D implosion asymmetries enables target designers to evaluate nonuniformity effects on implosion performance and identify the underlying degradation mechanisms. This effort requires a 3D reconstruction platform to integrate OMEGA's diagnostic measurements [9–15] such as hot-spot flows, ion temperatures, areal densities, and x-ray images. Recent studies [16–23] have tackled this challenge of reconstructing 3D x-ray emission and neutron source profiles from a limited set of two-dimensional (2D) line-of-sight (LOS) measurements. These 3D reconstruction approaches include spherical harmonic function optimization [16–18] for plasma emissivity and neutron source functions, algebraic reconstruction [19] for 3D electron temperature distributions, and neural networks [20] for extracting 3D morphological information from particle measurements and the analytic multirocket piston model [21–23] for 3D shell asymmetry analysis.

The 3D hot-spot reconstructions can be used to quantify residual kinetic energies [24–26]—a direct indicator of how low-mode asymmetries impact fusion yields. It can be used to guide the beam-to-beam laser power imbalance adjustment to improve ICF implosion symmetry. This approach [27,28] has been demonstrated on OMEGA to mitigate low-mode $\ell = 1, 2$ asymmetries through real-time laser power imbalance adjustments.

In this work, an advance 3D hot-spot reconstruction technique is presented using a deep-learning approach. We developed a convolutional neural network (CNN) model to reconstruct 3D hot-spot and shell structures from multiangle x-ray imaging data obtained during laser direct-drive ICF implosions. Measured x-ray images are used to constrain the 3D

*Contact author: kwoo@lle.rochester.edu

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

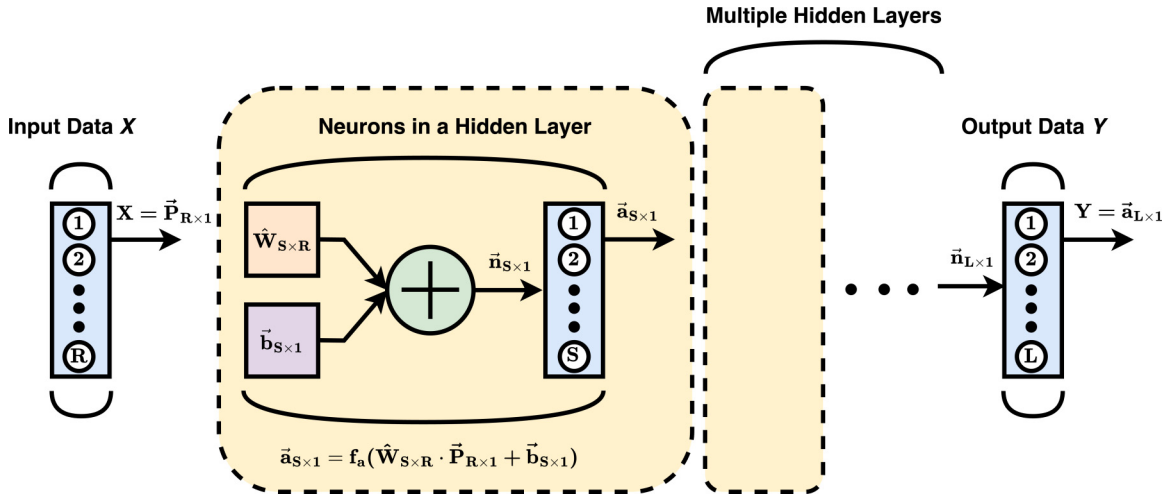


FIG. 1. A schematic representation of a neural network architecture as formalized by Hagan [29]. The input vector $\vec{X}$, with R features, is fed into a hidden layer of S neurons. Each neuron (depicted by a green circle marked with $+$) returns a nonlinear transformation of the weighted sum of inputs plus a bias term. The first hidden layer produces a nonlinear response vector $\vec{a} = f_a(\hat{W} \cdot \vec{P} + \vec{b})$, which then progresses through multiple subsequent layers, each performing increasingly abstract transformations, ultimately mapping to the output vector $\vec{Y}$.

hot-spot shape geometry, while the 3D shell profile is inferred indirectly from the CNN, trained with the DEC3D [25] single-mode simulation database. While hard x-ray signals (>5 keV) from OMEGA diagnostics weakly correlate [30–32] with the shell morphology, a physics-informed simulation-data-driven machine-learning approach is used to unfold this correlation. By training a CNN model on extensive DEC3D simulation data that incorporate 3D implosion physics, the correlation between x-ray images and the shell structure is established. The CNN model is trained on DEC3D simulations to learn the physical relationship between x-ray image features and the shell morphology. The hydrodynamic simulation database teaches the neural network how to map 2D x-ray image data to a 3D distribution such as the plasma emissivity, mass density, or electron temperature. CNN's performance was compared with Refs. [17,18], demonstrating fast and accurate 3D reconstructions of the plasma emissivity. The reconstructed 3D shell asymmetry for OMEGA shot 94017 agrees well with the DEC3D machine-learning (ML) reconstruction and HYDRA direct numerical simulation. The CNN model was applied to assess two key implosion metrics: the 4pi-averaged areal density and the hot-spot flow velocity. For shot 94017, the CNN-inferred values are consistent with DEC3D-ML reconstruction of 179 mg/cm^2 and 115 km/s within typical OMEGA experimental uncertainties.

This paper is organized as follows: Sec. II explains the neural network design, the generation of synthetic x-ray images, and the training method. Section III demonstrates the application of the CNN model to reconstruct OMEGA shot 94017, which was perturbed with a large mode $\ell = 1$. Reconstructed 3D core structures are compared with DEC3D-ML reconstruction results. Section IV summarizes the conclusions.

II. DEEP CONVOLUTIONAL ARCHITECTURES

Artificial neural networks (ANNs) provide powerful computational frameworks to create complex nonlinear mappings between input and output data. Here, ANNs are briefly explained. As illustrated in Fig. 1, the input data, denoted by the vector $\vec{X}$, undergo nonlinear transformation as it propagates through a hidden layer, where each neuron computes an output $a = f(\hat{W} \cdot \vec{X} + b)$, by applying a nonlinear activation function f to the weighted sum of inputs plus a bias term b, with $\hat{W}$ representing the weight matrix. The output vector $\vec{a}$ from each hidden layer serves as the input to the subsequent layer, creating an increasingly abstract transformation to map the output data, denoted by the vector $\vec{Y}$. The neural network learns from data by adjusting weight and bias parameters during the stochastic gradient decent optimization process to minimize prediction errors.

In ICF, multilayer neural networks have been applied to predict experimental observables using transfer learning [33]. In this approach, the network is initially trained on a large simulation database and subsequently fine tuned using a limited set of experimental data. It has been demonstrated that the neural network can effectively predict experimental observables.

For image analysis, a specialized class of deep-learning architecture known as convolutional neural networks is needed. Convolutional neural networks are inspired by the local processing mechanisms of neurons in the cat's visual cortex discovered by Hubel and Wiesel [34], pioneered by Fukushima's neocognitron [35] and refined by Lecun's [36] gradient-based learning approach. CNNs have revolutionized pattern recognition by automatically learning feature representations directly from the raw input data. There are three key principles to enable them to achieve such remarkable performance. First, CNN neurons connect only to a small region of the input space known as local receptive fields. This allows the network to detect elementary visual features such as oriented edges, end points, and corners locally. The dimensions of these local receptive fields correspond to the filter size in convolutional layers, determining the spatial extent of feature detection. Second, CNN neurons share the same weight at each network level, which fundamentally reduces the number of parameters while ensuring the translation invariance. It

means that the same feature detector can identify patterns regardless of where they appear in the image. The weight sharing reduces the number of trainable free parameters and, therefore, improves the model's generalization from limited training data compared to traditional multilayer neural networks. Third, the CNN implements the spatial subsampling, which progressively reduces the spatial resolution of feature maps while simultaneously increasing the number of feature maps as the network deepens. This bipyramid architecture extracts a large degree of shift, scale, and distortion invariance to geometric transformation of the input image. These powerful techniques make CNNs a promising candidate for image analysis tasks [37]. Early applications of CNNs in medical image analysis include the use of a two-layer CNN for lung nodule detection [38].

A. CNN model

In this section, we describe our CNN architecture developed for reconstructing three-dimensional hot spots in ICF implosions, utilizing x-ray images measured from multiple viewing angles. Note that the weight-sharing and spatial subsampling characteristics of CNNs ensure an efficient transformation of the 2D input images into a 3D data structure.

X-ray images measured at different LOSs contain 3D morphology information of the mass density and electron temperature distributions. On the OMEGA laser, there are four x-ray imagers available for measuring x rays emitted from the high-temperature, low-density hot spot, surrounded by a spherical DT or a plastic shell. These imagers are the gated monochromatic x-ray microscope (GMXI) [12] at $(\theta = 96^\circ, \phi = 54^\circ)$, the Kirkpatrick-Baez x-ray microscope (KBframed or KBXI) [13] at $(\theta = 106^\circ, \phi = 342^\circ)$, the single line-of-sight time-resolved x-ray imager (SLOS-TRXI) [14] at $(\theta = 45^\circ, \phi = 234^\circ)$, and the spatially resolved electron temperature (SR-TE) [15] diagnostics at $(\theta = 101^\circ, \phi = 134^\circ)$. Measured x-ray signals are dominated by photons with energies $\approx$5 keV for GMXI, KBXI, and SLOS-TRXI, and $\approx$10 keV for SR-TE. The setup of these four x-ray imagers on OMEGA is displayed in Fig. 2.

Our deep-learning CNN architecture, shown in Fig. 3, implements Lecun's bipyramid architecture [36]. The convolutional layers progressively extract and compress spatial features from high to low dimensions in the encoder stage, before expanding them in the decoder stage for 3D reconstruction. The dense layer (a_2) with 512 neurons serves as a bottleneck in the latent space, creating a compressed representation that captures essential spatial relationships and introduces nonlinearities through rectified linear unit (ReLU) activation. Together with the subsequent dense layer (a_3), this architecture enables an efficient mapping between 2D image features and 3D reconstruction. The dimension of the latent space increases with the complexity of 3D core structures.

The four input x-ray images are rendered at a resolution of 20×20 pixels, processed by two convolutional layers. These layers apply a convolution filter of 3×3 pixels, acting on each pixel of the input images. After the convolution operation, each convolved pixel element p is fed into a neuron, undergoing a nonlinear transformation to give a neuron sum $n = a(wp + b)$. The weight w and bias b are fitting

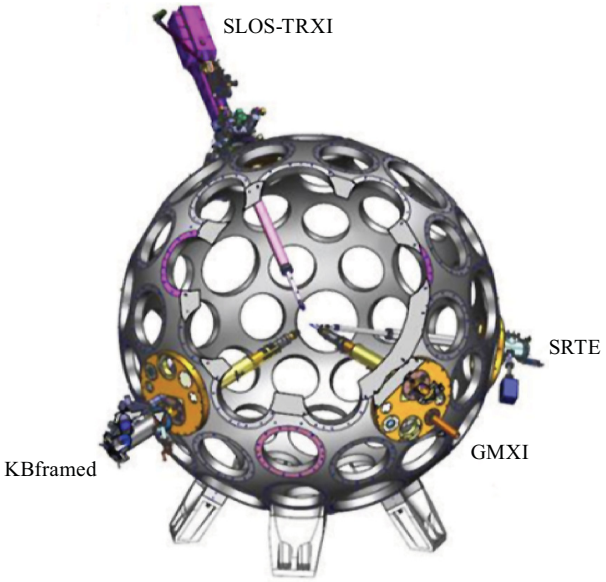


FIG. 2. Experimental setup of four x-ray imagers on OMEGA.

parameters to be optimized. The activation function a performs a nonlinear transformation of the input signal $wp + b$. Results of a are stored in the 128 and 256 feature maps. The dimensions of feature maps are reduced from 18×18 pixels in Conv2D_1 to 16×16 pixels in Conv2D_2, which extract lower-dimensional features. The information of feature maps is compressed into the one-dimensional (1D) latent space through a_1 and a_2 dense layers. The decoding process is obtained by connecting with two 1D dense layers, a_2 and a_3, to 3D target data with dimensions of $20 \times 20 \times 20$ pixels. The ReLU is used as the nonlinear activation function in Conv2D_1, Conv2D_2, and a_2 layers.

The two-layer Conv2D architecture demonstrates optimal performance for our 3200-sample dataset. The CNN architecture exhibits dataset-dependent scaling: for large datasets (>10 000 samples), additional Conv2D layers with doubled filter counts are required, while for smaller datasets (<1000 samples), fewer feature map counts in both Conv2D_1 and Conv2D_2 are needed. For example, our CNN model in a separate study [39] adds additional Conv2D layers for its 18 000-sample dataset.

To reconstruct all hydrodynamic primitive variables (flow velocities, mass density, and pressure), the linear activation is used instead of ReLU in the final dense layer a_3 to allow prediction across the full continuous range of positive and negative values of flow velocity. While ReLU activation in preceding layers captures nonlinearities, its zero-clipping property would limit prediction of bidirectional quantities such as flow velocities, making it unsuitable for the final dense layer. The ADAM optimizer minimizes the loss function $\mathcal{L}$, defined as the root-mean-squared error between true (y_{true}) and predicted (y_{pred}) target functions, as described in Eq. (4).

Our CNN architecture is designed with flexible inputs to accommodate one to four x-ray images, because the hardware performance limitations can reduce the number of available x-ray images on certain shot days. Section II C presents a comprehensive model validation test using the complete four x-ray

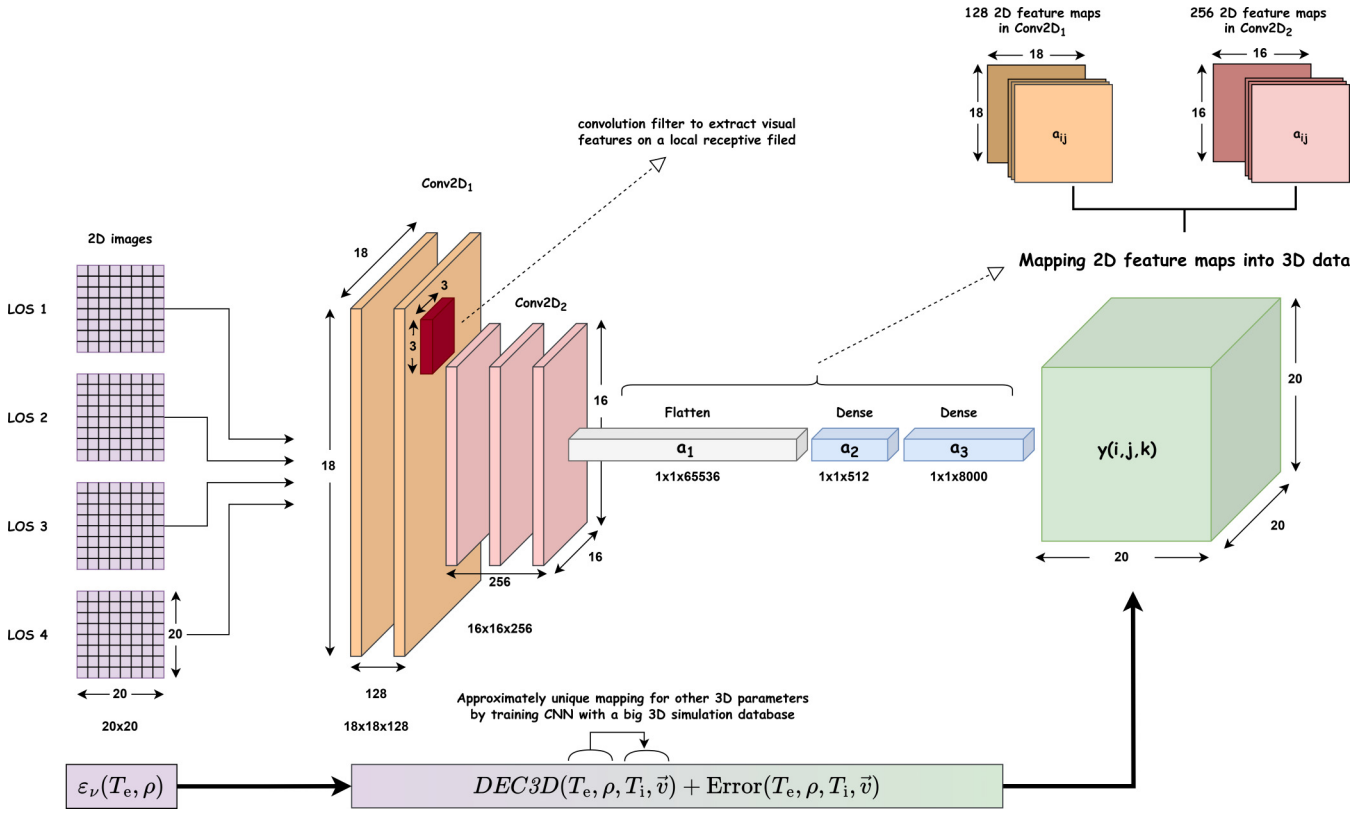


FIG. 3. The neural network design of the deep-learning convolutional neural network, which maps 2D x-ray images to 3D distributions of the plasma emissivity ε_ν, mass density ρ, or electron temperature T_e.

imager diagnostic suite of GMXI, KB-Framed, SLOS-TRXI, and SRTE. When applying the CNN model to shot 94017, as discussed in Sec. III, we used data from only two x-ray imagers (GMXI and SLOS-TRXI), since SRTE diagnostic malfunction and signal-to-noise limitations in KB-Framed measurements prevented their inclusion.

Shot 94017 was selected for detailed analysis due to its large mode-1 implosion asymmetry, providing an ideal test case for the CNN model validation. While we focuses on this specific shot, results from applying our CNN model across 19 cryogenic OMEGA implosions have been published separately in a companion paper [39]. In that work, Fig. 3 presents a CNN 3D reconstruction for shot 109861 using all four x-ray image data.

B. Synthetic x-ray images

The training dataset is provided by DEC3D [25] simulations, a 3D radiation-hydrodynamic code that simulates deceleration-phase Rayleigh-Taylor (RT) instabilities in ICF implosions. Without laser energy deposition physics, DEC3D cannot model acceleration phase dynamics or ablator mixing effects that increase x-ray emission when higher-Z ablator material mixes into the hot spot [40]. In this work, the CNN model is trained on DEC3D single-mode simulations of $\ell = 1$ and 2 modes, which are dominant low-mode asymmetries in OMEGA implosions caused by target offset, laser power imbalance, and beam mispointing. Higher-mode training is limited by GMXI diagnostic signal-to-noise constraints that prevent reliable resolution of midmode x-ray features. Mode amplitudes ($\delta v \equiv \Delta v/v_0$) are defined as initial velocity perturbations Δv relative to the peak implosion velocity v_0 at deceleration onset, with mode $\ell = 1$ ranging from 1 to 20% and mode $\ell = 2$ from $\pm 20\%$. The negative and positive amplitude range for mode 2 captures both oblate ($\delta v > 0$) and prolate ($\delta v < 0$) configurations, while the 20% upper bound samples the typical range of hot-spot shapes observed in OMEGA experiments, including both pancakelike (oblate) and sausagelike (prolate) ICF hot spots. These perturbations are seeded on the inner-shell surface position R_{HS} using spherical harmonics $Y_\ell^m(\theta,\phi)$ with a radially decaying shape function $f(r;R_{\mathrm{HS}})$, giving a total perturbation of $\delta v Y_\ell^m(\theta,\phi) f(r;R_{\mathrm{HS}})$, where $f(r;R_{\mathrm{HS}})$ describes perturbations decaying both inward and outward from R_{HS}. Details on asymmetry initializations were described in Ref. [25].

While multimode simulations better match experimental conditions, the substantial memory requirements for storing comprehensive multimode training datasets currently restrict us to single modes. DEC3D can model time-dependent asymmetries throughout both deceleration and disassembly phases of ICF implosion dynamics. Future work will incorporate DRACO simulations to capture mixing effects in 3D reconstructions and expand computing resources to enable comprehensive single-mode and multimode spectra in the training dataset.

For each implosion design, a DEC3D simulation database is constructed by scanning through low modes ($\ell = 1$, 1–20%; $\ell = 2$, $\pm 20\%$). Using DEC3D hydrodynamic data for the mass density ρ and electron-temperature T_e profiles, synthetic x-ray images are produced along a straight-line transport toward a

given LOS:

$$\frac{\mathrm{d}I_\nu}{\mathrm{d}s} = \varepsilon_\nu - k_\nu I_\nu. \tag{1}$$

Here we have the spectral radiation intensity I_ν (W/m^2/Hz), the spectral emission ε_ν (W/m^3/Hz), the spectral absorption coefficient k_ν (1/m), and the distance s traveled by a photon. The subscript ν labels a quantity, measured per unit frequency from ν to $\nu + \mathrm{d}\nu$. In thermodynamic equilibrium, the spectral emissivity is given by the product of the total opacity (spontaneous and induced emission), $k'_\nu(1/\mathrm{m}) = k_\nu[1 - e^{-h\nu/(k_\mathrm{B}T_\mathrm{e})}]$, and the Planck spectrum $B_\nu(\mathrm{W/m^2}) = \frac{2h\nu^3}{c^2}[e^{h\nu/(k_\mathrm{B}T_\mathrm{e})} - 1]^{-1}$ according to Kirchhoff's law. Using Kramers's free-free opacity [41], $k_\nu = \frac{4}{3}\sqrt{\frac{2\pi}{3m_\mathrm{e}k_\mathrm{B}T_\mathrm{e}}}\frac{n_\mathrm{i}n_\mathrm{e}Z^2e^6}{hcm_\mathrm{e}\nu^3}$ (1/cm), the spectral plasma emissivity is

$$\varepsilon_\nu(n_\mathrm{i}, n_\mathrm{e}, T_\mathrm{e}) = \frac{8}{3}\sqrt{\frac{2\pi}{3m_\mathrm{e}}}\frac{Z^2e^6}{m_\mathrm{e}c^3}\frac{n_\mathrm{i}n_\mathrm{e}}{\sqrt{K_\mathrm{B}T_\mathrm{e}}}e^{-\frac{E_\nu}{K_\mathrm{B}T_\mathrm{e}}}. \tag{2}$$

Here we have the electron mass m_e, the plasma ionization state Z, the electron charge e, the speed of light c, the Planck constant h, and the photon energy $E_\nu = h\nu$ at frequency ν. In Eq. (2), an ionization state of $Z = 1$ is assumed, as the DT plasmas typically achieve a full ionization at stagnation. The Gaunt factor [41] is omitted because stagnating hot spots contain a nondegenerate plasma in the low-density and high-temperature environment. The absorption term $k_\nu I_\nu$ in Eq. (1) is neglected because the cold shell only absorbs soft x rays with energies $E_\nu \approx 1$ keV, which is below our measured x-ray photon energies. The normalized plasma emissivity, $\hat{\varepsilon}_\nu = \varepsilon_\nu/\max(\varepsilon_\nu)$, is used to form synthetic x-ray images because measured x-ray signals do not have the same calibration.

C. Training method

The CNN requires retraining for different implosion designs to account for varying RT growth factors that drive hot-spot distortions. The workflow includes DEC3D simulation data generation in ≈1 week, 4–5 h per simulation, and model training in ≈1–2 h. While this requires an initial computational cost, the trained CNN generates 3D reconstructions in seconds, compared to 1–2 weeks required for DEC3D ML reconstruction to be described in Sec. III B.

Each 3D dataset with dimensions of 20 × 20 × 20 pixels requires 2.5 MB of memory. A small finite number of DEC3D simulations were first produced. Each DEC3D simulation datum was rotated with all LOS angles and scaled with size factors to generate 3200 3D datasets to train the CNN model. In the DEC3D single-mode simulation database, RT instabilities were seeded along the z axis. Cartesian coordinates $\vec{X}_0$ from each DEC3D datum were rotated uniformly to all possible LOSs (θ_i, ϕ_j) and spatially scaled by a size factor $f \in [0.5, 4]$ to cover all possible hot-spot sizes and orientations in experiments:

$$\vec{X} = f \cdot \hat{M}_z(\phi_j) \cdot \hat{M}_y(\theta_i) \cdot \vec{X}_0. \tag{3}$$

The discretized polar angles are given by $\theta_i = i \times \pi/(N-1)$ and $\phi_j = j \times 2\pi/(N-1)$, where $N-1$ is the total number of bins in the polar and azimuthal angles. $\hat{M}_z(\xi)$ and $\hat{M}_y(\xi)$ are rotation matrices that rotate a vector counterclockwisely about the z and y axes, respectively, by an angle ξ.

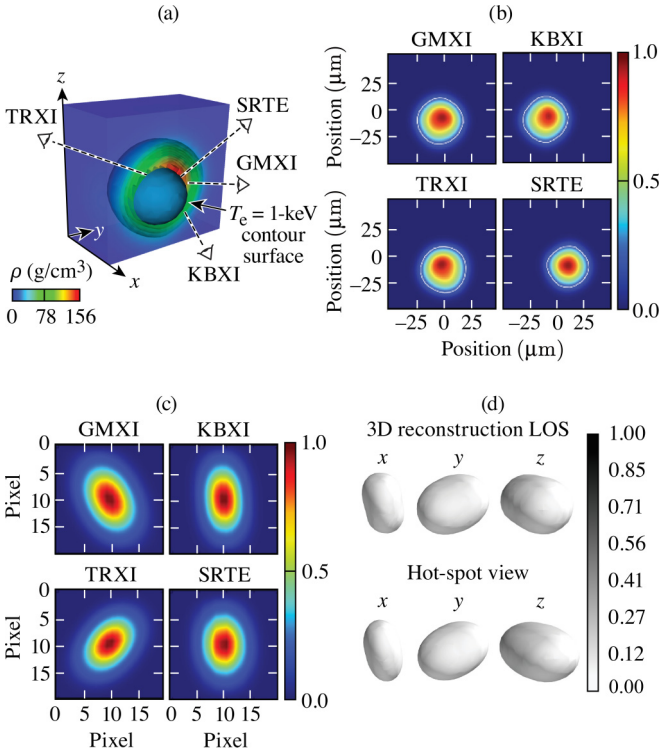


FIG. 4. (a) ρ and T_e profiles for DEC3D mode $\ell = 1$ simulation. (b) Corresponding synthetic x-ray images produced by this mode $\ell = 1$ hydrodynamic data in (a), measured at four LOSs. (c), (d) Comparison of 3D reconstructions of a rotated hot spot for mode $\ell = 2$. (c) The input synthetic x-ray images at four LOSs of this rotated mode $\ell = 2$ hot spot. (d) CNN (top) and SHG (bottom) reconstruction results for this mode $\ell = 2$ hot spot in (c).

The loss function $\mathcal{L}$ minimizes the root-mean-square error between predicted and true 3D reconstructions:

$$\mathcal{L} = \sqrt{\frac{1}{M^3}\sum_{x_i,y_j,z_k=1}^{M}[y_\mathrm{true}(x_i, y_j, z_k) - y_\mathrm{pred}(x_i, y_j, z_k)]^2}, \tag{4}$$

where $M = 20$ defines the dimensions of a $M \times M \times M$ cubic volume. The loss function computes the averaged root-mean-square difference between predicted (y_pred) and true (y_true) values at each voxel position (x_i, y_j, z_k). This metric quantifies the average prediction error across the entire reconstructed volume. Using this training approach, the CNN model achieves a strong generalization with about 2–3% prediction error, defined by Eq. (4), between normalized original and predicted 3D datasets across both training and test sets. All subsequent analyses in this article use trained models that satisfy this 2–3% accuracy range in both training and test predictions. Figure 5(b) demonstrates this performance to train a CNN model for reconstructing shot 94017.

Figure 4(a) shows the ρ and T_e profiles of a hot spot from a DEC3D simulation with a 7% mode $\ell = 1$ initial velocity perturbation. The velocity perturbation is introduced at the beginning of the deceleration phase (≈200 ps before

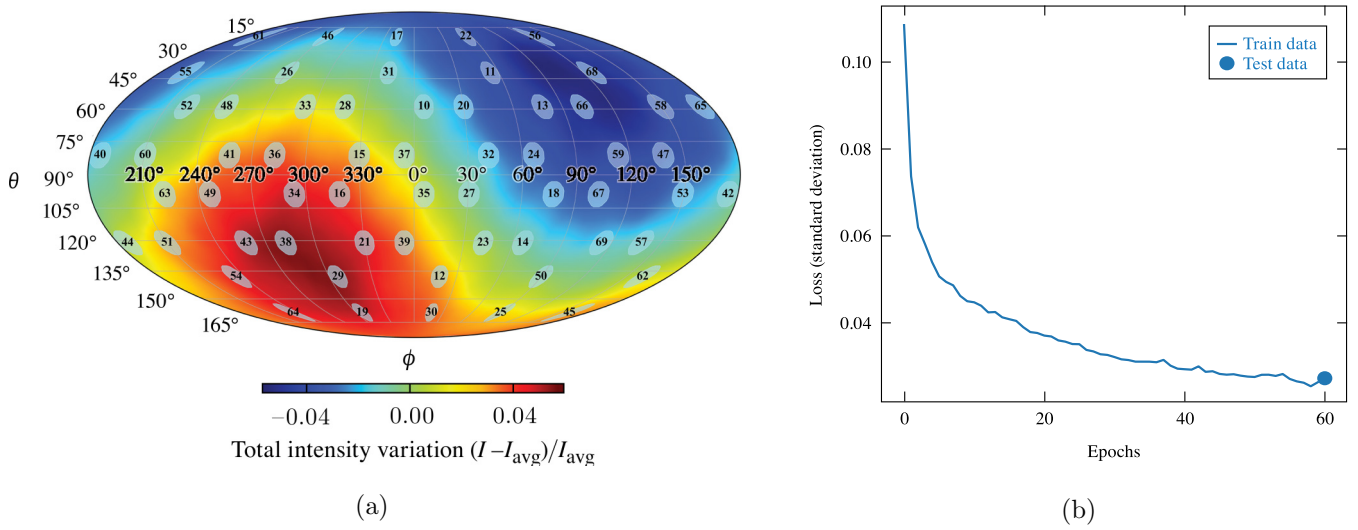


FIG. 5. (a) Hard-sphere illumination map showing on-target laser energy variation for shot 94017. Locations of OMEGA 60 laser beams are labeled by blue circles from port numbers 10 to 69. (b) Validation of the CNN model for shot 94017 using two x-ray images from GMXI and SLOS-TRXI lines of sight.

the bang time) and seeded along the polar axis, rotated to polar angles $\theta = 45^\circ$ and $\phi = 90^\circ$ with a unity size factor $f = 1$. The spectral radiation intensity emitted by a voxel with a plasma density ρ and an electron temperature T_e is calculated from Eq. (2), followed by tracking a straight-line ray trajectory to transport the radiation to the detector plane. Spectral filters were used to measure 5-keV photons by GMXI, KBXI, and TRXI, and 10-keV photons by SR-TE. Figure 4(b) shows the results of the rotated and spectrally filtered synthetic x-ray images. The SR-TE image of the hot spot is slightly smaller than others because detected photons are produced by a smaller high-temperature region inside the hot spot.

The model performance of the CNN is compared with our previous 3D hot-spot model [17,18] produced by the spherical harmonic (SHG) decomposition, which produces the normalized 3D plasma emissivity $\hat{\varepsilon}$:

$$\ln \hat{\varepsilon}(r,\theta,\phi) = \sum_{n=0}^{\infty} \sigma_n R^n \left(1 + \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} \sum_{k=0}^{\infty} A_{\ell mk} R^k Y_{\ell m}(\theta,\phi)\right)^n. \quad (5)$$

The SHG decomposition in Eq. (5) is given by a superposition of scalar spherical harmonics, encoded in two power series that expand the radial contour levels and the mode amplitude as a function of radius. Equation (5) was used to produce 2D projections of a rotated mode $\ell = 2$ at four LOSs. A training database was created using the second-order contour expansion ($\sigma_n = 0$ except for a nonzero σ_2). These 2D projections were used to train the CNN. Figures 4(c) and 4(d) show the four 2D input projections and compare $\hat{\varepsilon}_{SHG}$ produced by SHG (bottom) to the reconstructed profile $\hat{\varepsilon}_{CNN}$ by CNN (top). The standard deviation between $\hat{\varepsilon}_{SHG}$ and $\hat{\varepsilon}_{CNN}$ is 1.34%, relative to the peak of $\hat{\varepsilon}_{SHG}$. This result demonstrates an accurate 3D reconstructed hot spot by the CNN model. The training history of the CNN applied to OMEGA experimental data is detailed in Sec. III.

III. PERFORMANCE ON OMEGA

In this section, we present a detailed analysis of 3D reconstruction for shot 94017 using two independent methods: the CNN and DEC3D ML reconstructions. Additionally, Ref. [42] compared DEC3D results with HYDRA [43] direct numerical simulations for this shot. The quantitative agreement among three independent methods—physics-informed CNN reconstruction, DEC3D ML reconstruction, and HYDRA direct numerical simulation—provides a comprehensive validation of 3D core structures for shot 94017. The multishot analysis using the CNN model on OMEGA was reported separately in Ref. [39].

Shot 94017 has a large mode-1 implosion asymmetry, leading to significant ion temperature variations across different lines of sight, with approximately 2-keV difference between maximum and minimum apparent temperatures. Previous work [17] used a 3D ion-temperature reconstruction to determine the flow direction for shot 94017, identifying angular coordinates of $\theta = 81^\circ$ and $\phi = 112^\circ$. A direct measurement of the shell asymmetry was obtained by a magnetic recoil spectrometer (MRS) diagnostic [44], which measured the areal density (ρR). As reported in Table I, the measured ρR value was 214 mg/cm^2. The two primary objectives are to (1) assess the consistency among three independent properties—the inferred flow direction, the measured ρR, and the 3D shell asymmetry—and (2) validate CNN reconstruction results through comparison with both DEC3D ML reconstruction and HYDRA simulation.

A. CNN 3D reconstruction

Shot 94017 provides a critical test case for demonstrating the physics-informed CNN 3D reconstruction, as it features a well-characterized mode-1 asymmetry from beam

TABLE I. Three-dimensional reconstruction by DEC3D deceleration-phase simulations. The overall goodness of fit provided by DE3D simulation observables is quantified by the chi-squared statistic (χ^2/DOF), as illustrated in Fig. 8.

Shot 94017	Yield (10^{13})	$\rho R^{\mathrm{MRS}}_{\mathrm{mg/cm^2}}$	$T^{\mathrm{H10}}_{\mathrm{i,keV}}$	$T^{\mathrm{10.4m}}_{\mathrm{i,keV}}$	$T^{\mathrm{12m}}_{\mathrm{i,keV}}$	$T^{\mathrm{15.8m}}_{\mathrm{i,keV}}$	$T^{\mathrm{15.9m}}_{\mathrm{i,keV}}$	$\Delta t_{\mathrm{w,ps}}$	$t_{\mathrm{b,ns}}$
Experimental observable	1.20	214	5.21	3.05	3.90	3.23	2.92	67	2.51
Simulation observable	1.00	232	4.01	2.49	3.06	3.26	2.49	62	2.57
Experimental error	0.08	22	0.36	0.21	0.27	0.23	0.20	5	0.05

mispointing errors. Figure 5(a) shows the on-target laser energy variation map for this shot, revealing a substantial mode-1 asymmetry of 7.4% oriented at $\theta = 46^\circ$ and $\phi = 129^\circ$, where variations are calculated relative to the mean energy across all 60 beams. The total laser energy variation pattern was calculated using a hard-sphere illumination model that accounts for combined effects of target offset, laser imbalance, beam mispointing, and beam-to-beam overlap. This known mode-1 asymmetry direction provides a validation reference for the CNN-reconstructed mass density and plasma emissivity distributions.

Among four x-ray imagers, only GMXI and SLOS-TRXI provided analyzable data, as KB-Framed and SRTE diagnostics did not yield viable images for 3D reconstruction. Figure 5(b) shows the validation performance of the CNN model with two-imager inputs, achieving both high training accuracy and strong generalization in reconstructing the 3D hot-spot plasma spectral emissivity $\varepsilon_{5\mathrm{keV}}$ at the characteristic 5-keV x-ray emission energy. The prediction accuracy is quantified by a low fitting error between test and predicted 3D distributions, with a standard deviation of only 2.7% relative to the peak normalized intensity. The model achieved consistent performance across both training and test predictions after approximately 60 epochs.

Reconstructed ρ and T_e distributions shown in Fig. 6(b) were obtained from the CNN model using two normalized x-ray images, measured by GMXI and TRXI as shown in Fig. 6(a). Standard deviation between modal projections and measurements, for GMXI and TRXI (4.4 and 5.2%, respectively) was lower than that of SHG. Asymmetries in projections shown in Fig. 6(a) of the normalized electron temperature $\hat{T}_e$, and the normalized mass density $\hat{\rho}$, along these two LOSs are physically consistent with the signature of skewed x-ray images for a mode $\ell = 1$ perturbed hot spot. Converging velocity fields below the dense RT spike convert the fluid kinetic energy into the fluid internal energy through PdV work [25], resulting in a higher electron temperature below the RT spike.

For this shot, the areal density (ρR) was obtained solely from the MRS detector, as other neutron time-of-flight (nTOF) diagnostics did not provide valid ρR measurements. The substantial measured value of $\rho R = 214$ mg/cm^2 (Table I) along the MRS LOS ($\theta = 117^\circ, \phi = 306^\circ$) indicates a signature of mode-1 driven shell compression, with a shell mass accumulation in this direction. The enhanced ρR measurement is a characteristic of mode $\ell = 1$ implosion asymmetries, where RT instability drives material flow and creates regions of significantly increased areal density along the direction of the RT spike development.

Using the CNN 3D reconstruction model, it confirms a predicted enhancement in the shell thickness along the MRS LOS. The corresponding 3D reconstructed shell morphology, visualized in Figure 6(b), shows an agreement with this asymmetric density compression, with the region of an enhanced shell density aligned with the MRS diagnostic line of sight in the $x > 0, y < 0, z < 0$ quadrant. This agreement of spatial correlation between the reconstructed density distribution and experimental ρR measurement validates the CNN model's ability to capture directional shell asymmetries. The inferred flow direction ($\theta = 81^\circ$ and $\phi = 112^\circ$) aligns with the expected RT instability dynamics, where the hot-spot flow drives from regions of thicker to thinner shell wall, as observed in the $x < 0, y > 0, z > 0$ quadrant in Fig. 6(b).

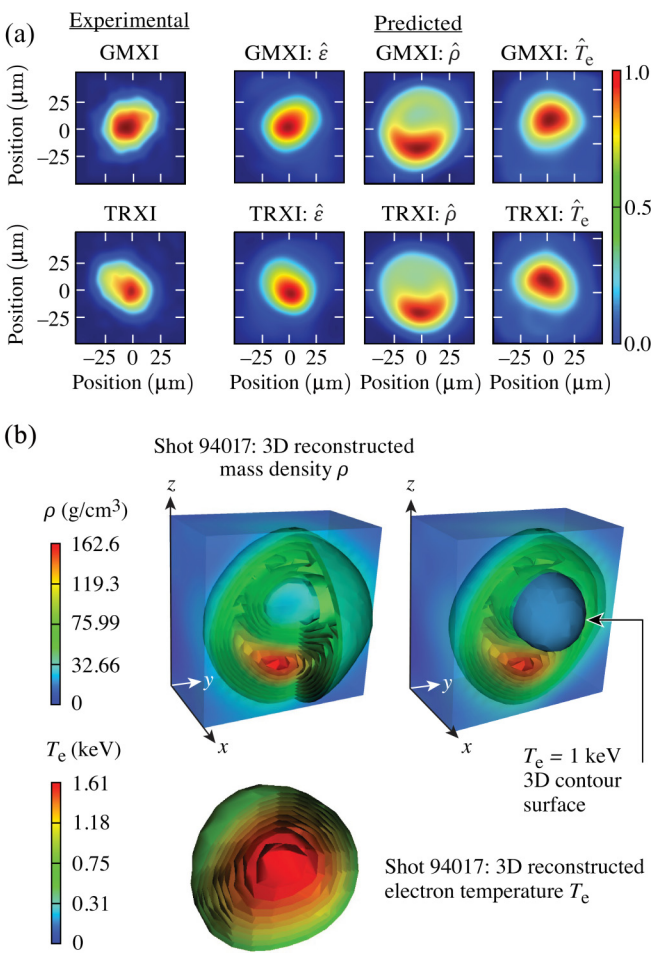


FIG. 6. 3D reconstruction for OMEGA shot 94017 using CNN. (a) X-ray images measured by GMXI and TRXI, and their predicted normalized plasma emissivity $\hat{\varepsilon}$, mass density $\hat{\rho}$, and electron temperature $\hat{T}_e$ projections at the same LOS. (b) 3D distributions of reconstructed shell mass density and hot-spot electron temperature.

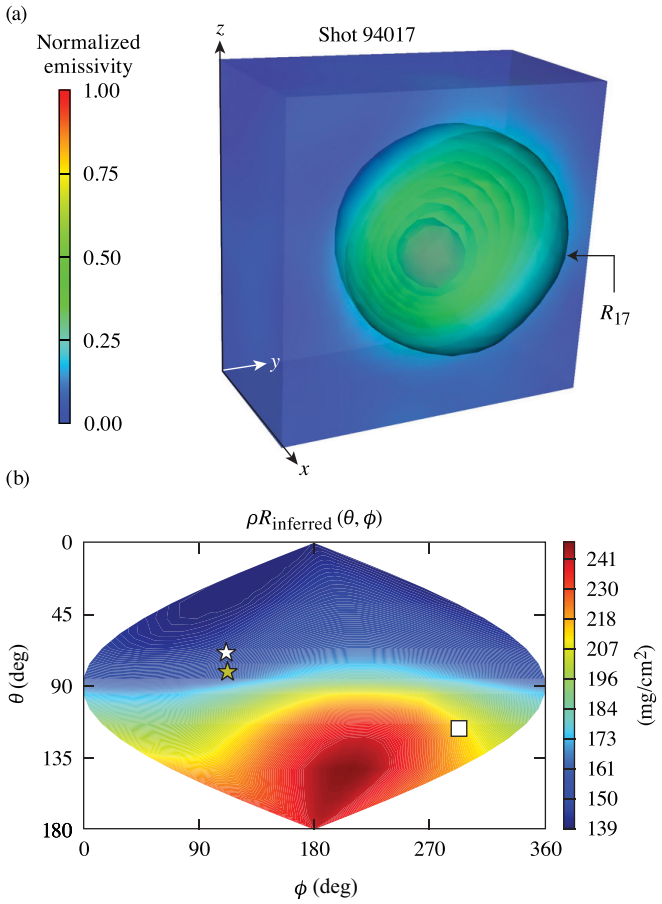


FIG. 7. 3D reconstruction for OMEGA shot 94017. (a) CNN-reconstructed normalized plasma emissivity showing the asymmetric hot-spot structure. (b) Angular distribution of the inferred areal density, with the yellow/white stars indicating the experimentally/DEC3D-ML inferred hot-spot flow directions and a white square indicating the line of sight of MRS areal density diagnostic.

Figure 7(a) shows the CNN-reconstructed normalized plasma emissivity for shot 94017, while Fig. 7(b) shows the angular distribution of the CNN-inferred areal density $\rho R_{\text{inferred}}(\theta, \phi)$, which corresponds to MRS areal-density measurements using forward-scattered neutrons. The experimentally inferred flow direction, marked by a yellow star in Fig. 7(b), aligns consistently with the low inferred ρR region, while the MRS diagnostics, marked by a white square, locates at the high ρR region.

The fidelity of the reconstructed shell morphology is fundamentally constrained by the physical accuracy and completeness of the DEC3D simulation database used for CNN training, which incorporates both 1D and 3D hydrodynamic physics to transform the information of x-ray image data into 3D reconstructed hot-spot and shell structures. The accuracy can be enhanced by integrating additional nuclear diagnostic measurements—areal density ρR, apparent ion temperature T_{i}, and flow velocity v—with existing x-ray images. These nuclear measurements provide complementary constraints on key 3D implosion characteristics: v directly quantifies hot-spot motions, T_{i} reveals the distribution of hot-spot residual kinetic energies, and ρR maps shell mass distribution patterns. Limitations of this simulation-dependent approach for shell reconstructions are discussed in Sec. III D.

B. Validation with DEC3D ML reconstruction

In this section, the DEC3D ML reconstruction is used as an independent validation method for shot 94017, providing a complementary analysis of the hot-spot and shell structures obtained from the CNN model. The objective is to assess the consistency of 3D reconstructed features. The methodology of DEC3D ML reconstruction proceeds as follows.

The initial seed of an *ad hoc* 3D velocity perturbation $\delta V_{3\text{D}}$ and 1D corrections at the beginning of the deceleration phase in DEC3D are optimized so as to simulate measured observables for shot 94017. DEC3D is a 3D radiation-hydrodynamic code [25], developed to simulate RT instabilities in the deceleration phase of ICF implosions. RT seeds are imposed as initial velocity perturbations on the shell's inner surface at the time t_{dec} of peak implosion velocity. Hydrodynamic data are postprocessed by the radiation-transport simulation software SPECT3D [45] and the neutron-transport simulation code IRIS3D [46]. These codes produce synthetic data to simulate the ion temperatures, hot-spot flows, areal densities, and x-ray images measured at different LOSs on OMEGA. The 1D correction factors $\delta Q_{1\text{D}} = (\delta V_{1\text{D}}, \delta \rho_{1\text{D}}, \delta P_{1\text{D}}, \delta P_{\text{e,1D}})$ are given by a set of multipliers applied to the 1D velocity $V_{1\text{D}}$, mass density $\rho_{1\text{D}}$, total pressure $P_{1\text{D}}$, and electron pressure $P_{\text{e,1D}}$ profiles at time t_{dec}. The parameter vector $\vec{P} = (\delta V_{3\text{D}}, \delta V_{1\text{D}}, \delta \rho_{1\text{D}}, \delta P_{1\text{D}}, \delta P_{\text{e,1D}})$ is optimized using the gradient descent method so that the loss function $\mathcal{L} = \frac{1}{N-D}\sum_{i=1}^{N} \chi_i^2$, defined as the χ^2 per degree of freedom which quantifies the goodness of fit between measurements and synthetic data, is minimized. Here, N is the total number of experimental observables, and D is the total number of fitting parameters. The minimum of $\mathcal{L}$ is obtained by an iterative solution of the linearized evolution $\frac{\mathrm{d}\mathcal{L}}{\mathrm{d}t} = \vec{\nabla}\mathcal{L} \cdot \frac{\mathrm{d}\vec{P}}{\mathrm{d}t}$, where t is the number of search steps. In each search step, the linear response $\vec{\nabla}\mathcal{L} = \sum_{i=1}^{D} (\partial \mathcal{L}/\mathcal{L}P_i)\hat{P}_i$, spanned by each unit vector $\hat{P}_i$ of the ith component, is calculated. The opposite direction of the gradient vector $\vec{\nabla}\mathcal{L}$ is used to perturb $\vec{P}$ along a direction of decreasing $\mathcal{L}$: $\vec{P}(t + \Delta t) = \vec{P}(t) - \varepsilon \cdot (\vec{\nabla}\mathcal{L}/|\vec{\nabla}\mathcal{L}|)\mathrm{d}\vec{P}$. The dimensionless learning rate, $0 \leqslant \varepsilon \leqslant 1$, determines the speed of convergence towards the local minimum of $\mathcal{L}$. The initial learning rate $\varepsilon = 0.2$ was used, and was reduced by half when the updated loss function is larger than the previous one: $\mathcal{L}(t + \Delta t) > \mathcal{L}(t)$.

To accelerate the optimization, a mode $\ell = 1$ perturbation $\delta V_{3\text{D}}^{\ell=1}(\theta = 99^\circ, \phi = 292^\circ)$, determined by ion temperature asymmetries [17], was seeded along the effective flow axis. A mode $\ell = 10$ perturbation $\delta V_{3\text{D}}^{\ell=10,m=5}(\theta = 0^\circ, \phi = 0^\circ)$ was seeded along the polar axis to simulate the midmode effect caused by the beam-to-beam overlapping nonuniformity on OMEGA. As shown in Fig. 8, $\mathcal{L}$ was dropped from 1478.51 to the minimum 5.98 after six search steps. The optimized $\vec{P}$ is $\delta V_{3\text{D}}^{\ell=1} = 43\%$, $\delta V_{3\text{D}}^{\ell=10} = 13\%$, $\delta V_{1\text{D}} = -7\%$, $\delta \rho_{1\text{D}} = -12\%$, $\delta P_{1\text{D}} = 0\%$, $\delta P_{\text{e,1D}} = 0\%$.

Table I compares DEC3D reconstruction results to the experimental measurements for the fusion yield, areal density

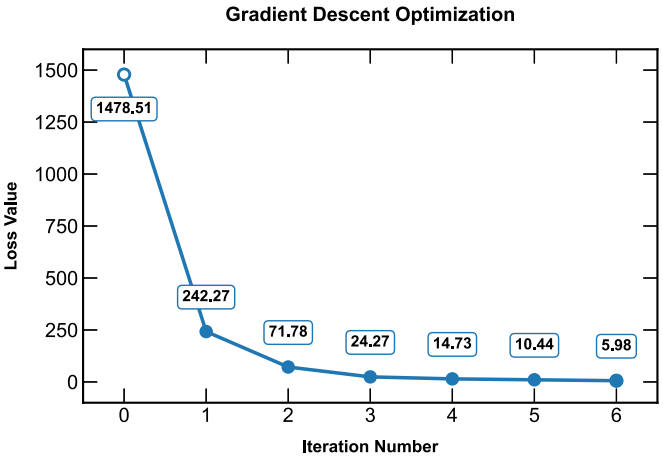


FIG. 8. History of gradient descent optimization to fit DEC3D simulation observables with experimental data for shot 94017. The loss value is given by the χ^2 per degree of freedom.

[ρR, measured by MRS], ion temperatures (T_i^X, measured by nTOF detectors), burn width (Δt_w), and bang time (t_b). The superscript "X" is the name of the nTOF detector, which is the beam-port name or the distance to the target chamber center. Since the true RT seeds were not applied to DEC3D, some discrepancies between simulated and measured observables were found within a few error bars. In Fig. 9, the simulated TRXI x-ray image shows a skewed hot-spot image similar to the measurement. The direction of skewness in the simulated x-ray image and the projection of the simulated hot-spot flow vector on TRXI LOS are parallel. Ion temperatures and areal densities inferred in $4-\pi$ LOSs show the expected cosine-square and cosine profiles, respectively, as discussed in Ref. [25]. The high down-scatter ratio region near the MRS LOS is consistent with the large ρR measurement. Reconstructed 3D ρ from CNN and DEC3D shows a good agreement in the alignment of the thicker shell wall in Figs. 6 and 9. In contrast to DEC3D simulation, midmode structures are not reconstructed because the CNN was trained without midmodes. The latter is because OMEGA x-ray detectors, measuring 5- to 10-keV photons, do not have the resolution required to capture midmode x-ray characteristics.The 4π-averaged areal density is a key parameter to quantify the ICF progress toward ignition [21,47] whereas hot-spot flows represent the loss of hot-spot compression caused by growing residual kinetic energies. These two parameters are estimated by CNN reconstructions.

TABLE II. Comparison of 4π-averaged areal density and hot-spot flow velocity between CNN and DEC3D-ML reconstructions for shot 94017. Overall goodness of fit provided by DE3D simulation observables is quantified collectively by the chi-squared statistic during the gradient descent optimization, as illustrated in Fig. 8.

	$\langle \rho R \rangle_{4\pi-\text{avg}}$ (mg/cm^2)	v_{flow} (km/s)
CNN	185 ± 6	100 ± 15
DEC3D-ML	179	115

Using the mass density profile from the CNN reconstruction shown in Fig. 6(b), areal densities $\rho R_{\text{CNN}}(\theta, \phi)$ across all solid angles can be calculated. Figure 7(b) shows the CNN-inferred ρR map, which corresponds to areal densities measured by MRS using forward-scattered neutrons. The CNN-inferred ρR map is scaled by a correction factor $\beta = \rho R^{\text{MRS}}_{\text{CNN}}/\rho R^{\text{MRS}}_{\text{exp}}$ such that $\rho R_{\text{inferred}}(\theta, \phi) = \rho R_{\text{CNN}}(\theta, \phi)/\beta$, where $\beta = 1.44$ for shot 94017 calibrates the inference to the measured value. From the ρR map shown in Fig. 7(b), the calculated 4π-averaged areal density is of 185 mg/cm^2. Using the analytic piston model [21,22], the mode-1 hot-spot flow velocity v_{flow} is estimated from the inferred shell asymmetry fraction f:

$$f = \frac{\rho R_{\text{max}} - \rho R_{\text{min}}}{\rho R_{\text{max}} + \rho R_{\text{min}}} = \frac{v_{\text{flow}}}{v_{\text{imp}}}, \tag{6}$$

where v_{imp} is the implosion velocity. Based on the 1D implosion velocity of 357 km/s and the shell asymmetry fraction $f = 0.28$ from Fig. 7(b), the estimated hot-spot flow velocity for shot 94017 is of 100 km/s. As shown in Table II, the DEC3D-ML reconstruction estimates a 4π-averaged areal density of 179 mg/cm^2 and flow velocity of 115 km/s, consistent with CNN-inferred values within experimental uncertainties, ± 20 mg/cm^2 for MRS ρR, and ± 30 km/s for flows.

C. Discussion

The DEC3D ML reconstruction requires a 43% initial velocity perturbation at the peak implosion velocity to

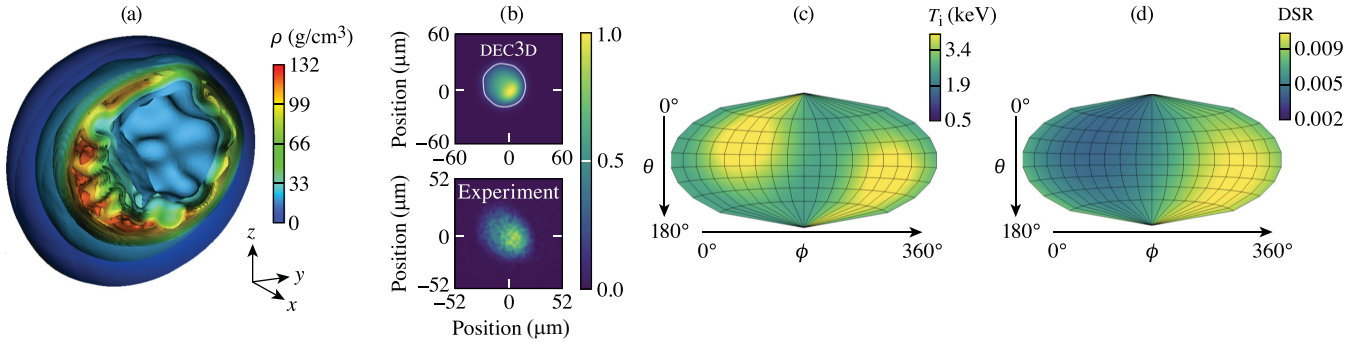


FIG. 9. Results of DEC3D reconstruction for OMEGA shot 94017: (a) the simulated 3D mass-density profile, (b) the simulated x-ray image at the TRXI LOS, (c) inferred ion temperatures in $4-\pi$ LOSs, and (d) inferred down-scattered neutron ratio in $4-\pi$ LOSs.

reproduce the observed mode-1 asymmetry. This substantial perturbation manifests experimentally as an approximately 2-keV difference between the maximum and the minimum apparent ion temperatures. Recent HYDRA 3D hydrodynamic simulations of shot 94017 were compared [42] with DEC3D ML reconstruction, showing strong agreement in both hot-spot and shell morphology and key performance metrics including fusion yield, bang time, and burn width. HYDRA's simulated flow velocity magnitude of 107 km/s and direction ($\theta = 68^\circ, \phi = 139^\circ$) closely match DEC3D ML's reconstructed flow velocity magnitude of 115 km/s and direction ($\theta = 69^\circ, \phi = 107^\circ$). Therefore, HYDRA simulated, DEC3D ML reconstructed, and experimentally inferred flow directions ($\theta = 68^\circ, \phi = 139^\circ$; $\theta = 69^\circ, \phi = 107^\circ$; and $\theta = 81^\circ$, $\phi = 112^\circ$, respectively) all align with the low ρR region from the CNN reconstruction in Fig. 7(b).

The CNN model has been applied to two correlation studies [26,39]: analyzing the relationship between CNN-reconstructed hot-spot residual kinetic energies $\mathrm{RKE}_{\mathrm{hs}}^{\mathrm{CNN}}$ and flow velocities, and between CNN-reconstructed mode-2 amplitudes and fusion yield degradation across 19 OMEGA shots. The first study [26] demonstrated strong correlation between hot-spot flow velocities and $\mathrm{RKE}_{\mathrm{hs}}^{\mathrm{CNN}}$, consistent with previous findings [25,47]. This analysis included two representative cases, identified through x-ray imaging, shot 94017 dominated by mode-1 asymmetry and shot 96808 dominated by mode-2 asymmetry. In the second study [39], a weak correlation was observed between the fusion yield degradation and mode-2 amplitudes derived from CNN-reconstructed hot-spot surface radius variations. This trend agrees with DRACO 2D simulations for typical OMEGA implosion conditions, where mode-2 laser power imbalance rarely exceeds 1%, resulting in simulated yield degradations ($Y_{\mathrm{2D}}/Y_{\mathrm{1D}}$) of less than 3%.

The CNN model generates 3D reconstructions in seconds compared to 1–2 weeks for DEC3D ML optimization. Each DEC3D iteration requires six simulations to compute parameter perturbations for mode-1, mode-10, and four 1D corrections (velocity, density, total/electron pressures), taking 4–5 h per simulation. With five to ten optimization iterations needed, the total reconstruction time reaches 1–2 weeks. While DEC3D provides detailed physics modeling, the CNN's rapid analysis capability enables practical shot-day applications such as real-time laser power balance adjustments for improving implosion symmetry [27,28].

D. Limitations and future improvements

Our current assessment of the shell morphology remains limited, and relies on computational simulation data. Particles emitting from or dynamically interacting with the shell provide direct insights into the capsule's 3D structural configuration. To address this limitation, we plan to train the CNN model using five additional knock-on deuteron images (KODI) [48] to provide direct diagnostic mapping of the shell's structure.

As a fundamental challenge in 3D reconstruction, the mapping from 2D images to 3D structures suffers from nonuniqueness, because some distinct 3D structures can produce degenerate or topologically equivalent 2D projections. Provided 3D structures produce observable x-ray signatures; the reconstruction uncertainty from limited LOS coverage is quantified by the model's prediction errors as shown in Fig. 5(b), which capture the CNN's accuracy across all rotation angles for a fixed LOS configuration. However, 3D reconstruction can fail to resolve distinct 3D structures in two scenarios: (1) when features are smaller than the diagnostic resolution or (2) when structures have insufficient electron temperature to produce detectable x-ray emission. To mitigate this geometric ambiguity, we plan to add five x-ray images from pinhole cameras along KODI lines of sight into the CNN model to provide more geometric constraints on the 3D hot-spot reconstruction. Therefore, there will be a total of nine x-ray images and five KODI images to provide improved hot-spot and shell morphology.

IV. CONCLUSION

In summary, a deep-learning CNN was developed to reconstruct three-dimensional hot-spot and shell morphologies in ICF implosions. The model unfolds the hot-spot structure from multiangle x-ray images while indirectly inferring the shell configuration through a physics-informed CNN pretrained on an extensive DEC3D simulation database. By integrating machine learning with multiview x-ray diagnostics, this approach provides rapid evaluation of implosion asymmetries—crucial for understanding 3D core structures. The CNN model was used to assess two key implosion metrics: the 4pi-averaged areal density and the hot-spot flow velocity. These parameters quantify how 3D implosion asymmetries affects ICF ignition threshold. Multiple independent approaches were used to validate the CNN model's accuracy. For OMEGA shot 94017, CNN-reconstructed profiles closely matched both DEC3D machine-learning reconstruction and HYDRA direct numerical simulation. The CNN model inferred a 4pi-averaged areal density of 185 ± 6 mg/cm^2 and flow velocity of 100 ± 15 km/s, consistent with DEC3D-ML values of 179 mg/cm^2 and 115 km/s within experimental uncertainties. This physics-informed CNN approach generates 3D reconstructions in seconds, providing real-time analysis of implosion asymmetries and allowing researchers to dynamically optimize implosion performance during experiments through rapid symmetry control.

ACKNOWLEDGMENTS

This material is based upon work supported by the US Department of Energy Fusion Energy Sciences program under Grants No. DE-SC0024381, No. DE-SC0024456, the STARFIRE IFE Hub, and the U.S. Department of Energy (National Nuclear Security Administration) University of Rochester "National Inertial Confinement Fusion Program" under Grant No. DE-NA0004144.

DATA AVAILABILITY

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

[1] S. Atzeni and J. Meyer-ter-Vehn, *The Physics of Inertial Fusion: Beam Plasma Interaction, Hydrodynamics, Hot Dense Matter* (Oxford University, New York, 2004).

[2] E. Moses, The National Ignition Facility: An experimental platform for studying behavior of matter under extreme conditions, Astrophys. Space Sci. **336**, 3 (2011).

[3] A. L. Kritcher *et al.*, Design of inertial fusion implosions reaching the burning plasma regime, Nat. Phys. **18**, 251 (2022).

[4] T. R. Boehly *et al.*, Initial performance results of the OMEGA laser system, Opt. Commun. **133**, 495 (1997).

[5] R. L. McCrory *et al.*, OMEGA ICF experiments and preparation for direct drive ignition on NIF, Nucl. Fusion **41**, 1413 (2001).

[6] S. X. Hu, D. T. Michel, A. K. Davis, R. Betti, P. B. Radha, E. M. Campbell, D. H. Froula, and C. Stoeckl, Understanding the effects of laser imprint on plastic-target implosions on OMEGA, Phys. Plasmas **23**, 102701 (2016).

[7] K. S. Anderson *et al.*, Effect of cross-beam energy transfer on target-offset asymmetry in direct-drive inertial confinement fusion implosions, Phys. Plasmas **27**, 112713 (2020).

[8] O. M. Mannion *et al.*, Mitigation of mode-one asymmetry in laser-direct-drive inertial confinement fusion implosions, Phys. Plasmas **28**, 042701 (2021).

[9] O. M. Mannion, J. P. Knauer, V. Yu. Glebov, C. J. Forrest, A. Liu, Z. L. Mohamed, M. H. Romanofsky, T. C. Sangster, C. Stoeckl, and S. P. Regan, A suite of neutron time-of-flight detectors to measure hot-spot motion in direct-drive inertial confinement fusion experiments on OMEGA, Nucl. Instrum. Methods Phys. Res. Sect. A **964**, 163774 (2020).

[10] V. Y. Glebov, C. J. Forrest, K. L. Marshall, M. Romanofsky, T. C. Sangster, M. J. Shoup III, and C. Stoeckl, A new neutron time-of-flight detector for fuel-areal-density measurements on OMEGA, Rev. Sci. Instrum. **85**, 11E102 (2014).

[11] Z. L. Mohamed, O. M. Mannion, J. P. Knauer, C. J. Forrest, V. Yu. Glebov, C. Stoeckl, and M. H. Romanofsky, Application of an energy-dependent instrument response function to analysis of nTOF data from cryogenic DT experiments, Rev. Sci. Instrum. **92**, 043546 (2021).

[12] F. J. Marshall and J. A. Oertel, A framed monochromatic x-ray microscope for ICF, Rev. Sci. Instrum. **68**, 735 (1997).

[13] F. J. Marshall, R. E. Bahr, V. N. Goncharov, V. Yu. Glebov, B. Peng, S. P. Regan, T. C. Sangster, and C. Stoeckl, A framed, 16-image Kirkpatrick-Baez x-ray microscope, Rev. Sci. Instrum. **88**, 093702 (2017).

[14] W. Theobald *et al.*, The single-line-of-sight, time-resolved x-ray imager diagnostic on OMEGA, Rev. Sci. Instrum. **89**, 10G117 (2018).

[15] D. Cao, R. C. Shah, S. P. Regan, R. Epstein, I. V. Igumenshchev, V. Gopalaswamy, A. R. Christopherson, W. Theobald, P. B. Radha, and V. N. Goncharov, Interpreting the electron temperature inferred from x-ray continuum emission for direct-drive inertial confinement fusion implosions on OMEGA, Phys. Plasmas **26**, 082709 (2019).

[16] P. L. Volegov, C. R. Danly, D. Fittinghoff, V. Geppert-Kleinrath, G. Grim, F. E. Merrill, and C. H. Wilde, Three-dimensional reconstruction of neutron, gamma-ray, and x-ray sources using spherical harmonic decomposition, J. Appl. Phys. **122**, 175901 (2017).

[17] K. M. Woo *et al.*, Analysis of core asymmetries in inertial confinement fusion implosions using three-dimensional hot-spot reconstruction, Phys. Plasmas **29**, 082705 (2022).

[18] K. Churnetski *et al.*, Three-dimensional hot-spot x-ray emission tomography from cryogenic deuterium–tritium direct-drive implosions on OMEGA, Rev. Sci. Instrum. **93**, 093530 (2022).

[19] K. W. Wong and B. Bachmann, Three-dimensional electron temperature measurement of inertial confinement fusion hotspots using x-ray emission tomography, Rev. Sci. Instrum. **93**, 073501 (2022).

[20] L. Wu, S. Yoo, A. F. Suzana, T. A. Assefa, J. Diao, R. J. Harder, W. Cha, and I. K. Robinson, Three-dimensional coherent x-ray diffraction imaging via deep convolutional neural networks, npj Comput. Mater. **7**, 175 (2021).

[21] O. A. Hurricane *et al.*, An analytic asymmetric-piston model for the impact of mode-1 shell asymmetry on ICF implosions, Phys. Plasmas **27**, 062704 (2020).

[22] O. A. Hurricane *et al.*, 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).

[23] D. T. Casey *et al.*, A multi-rocket piston model to study three-dimensional asymmetries in implosions at the National Ignition Facility, High Energy Density Phys. **54**, 101172 (2025).

[24] A. L. Kritcher *et al.*, Metrics for long wavelength asymmetries in inertial confinement fusion implosions on the National Ignition Facility, Phys. Plasmas **21**, 042708 (2014).

[25] K. M. Woo *et al.*, Effects of residual kinetic energy on yield degradation and ion temperature asymmetries in inertial confinement fusion implosions, Phys. Plasmas **25**, 052704 (2018).

[26] K. Churnetski *et al.*, Three-dimensional reconstruction of inertial confinement fusion hot-spot plasma from x-ray and nuclear diagnostics on OMEGA, Bull. Am. Phys. Soc.PI02.00003 (2024).

[27] K. Churnetski *et al.*, Mitigation of low-mode perturbations with 3-D hot-spot x-ray emission tomography on laser-direct-drive implosions on OMEGA, Bull. Am. Phys. Soc.NO07.00004 (2023).

[28] K. M. Woo *et al.*, Experimental demonstration of 3D hot-spot shape symmetry control in laser direct-drive inertial confinement fusion implosions, Bull. Am. Phys. Soc.**T**O07.00010 (2024).

[29] M. T. Hagan, H. B. Demuth, M. H. Beale, and O. De Jesús, *Neural Network Design* (Hagan, Oklahoma, 2014).

[30] D. J. Schlossberg *et al.*, Observation of hydrodynamic flows in imploding fusion plasmas on the National Ignition Facility, Phys. Rev. Lett. **127**, 125001 (2021).

[31] D. T. Casey *et al.*, Evidence of three-dimensional asymmetries seeded by high-density carbon-ablator nonuniformity in experiments at the National Ignition Facility, Phys. Rev. Lett. **126**, 025002 (2021).

[32] D. T. Casey *et al.*, Diagnosing the origin and impact of low-mode asymmetries in ignition experiments at the National Ignition Facility, Phys. Rev. E **108**, L053203 (2023).

[33] K. D. Humbird, J. L. Peterson, J. Salmonson, and B. K. Spears, Cognitive simulation models for inertial confinement fusion: Combining simulation and experimental data, Phys. Plasmas **28**, 042709 (2021).

[34] D. H. Hubel and T. N. Wiesel, Receptive fields, binocular interaction and functional architecture in the cat's visual cortex, J. Physiol. **160**, 106 (1962).

[35] K. Fukushima, Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position, Biol. Cybernetics **36**, 193 (1980).

[36] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, Gradient-based learning applied to document recognition, Proc. IEEE **86**, 2278 (1998).

[37] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. W. M. van der Laak, B. van Ginneken, and C. I. Sánchez, A survey on deep learning in medical image analysis, Med. Image Anal. **42**, 60 (2017).

[38] S.-C. B. Lo, S.-L.A. Lou, J.-S. Lin, M. T. Freedman, M. V. Chien, and S. K. Mun, Artificial convolution neural network techniques and applications for lung nodule detection, IEEE Trans. Med. Imaging **14**, 711 (1995).

[39] K. Churnetski *et al.*, Three-dimensional reconstruction of laser-direct-drive inertial confinement fusion hot-spot plasma from x-ray diagnostics on the OMEGA laser facility, Rev. Sci. Instrum. **95**, 103521 (2024).

[40] T. Ma *et al.*, The role of hot spot mix in the low-foot and high-foot implosions on the NIF, Phys. Plasmas **24**, 056311 (2017).

[41] Y. B. Zel'dovichand and Yu. P. Raizer, *Physics of Shock Waves and High-Temperature Hydrodynamic Phenomena* (Dover, New York, 2002).

[42] K. S. Anderson, J. A. Marozas, S. C. Miller, T. J. Collins, K. M. Woo, V. N. Goncharov, A. Shvydky, S. M. Sepke, and M. M. Marinak, Progress in 3D modeling of direct-drive inertial confinement implosions using HYDRA, Bull. Am. Phys. Soc. **NO**07.00006 (2024).

[43] M. M. Marinak, G. D. Kerbel, N. A. Gentile, O. Jones, D. Munro, S. Pollaine, T. R. Dittrich, and S. W. Haan, Three-dimensional HYDRA simulations of National Ignition Facility targets, Phys. Plasmas **8**, 2275 (2001).

[44] J. A. Frenje *et al.*, First measurements of the absolute neutron spectrum using the magnetic recoil spectrometer at OMEGA, Rev. Sci. Instrum. **79**, 10E502 (2008).

[45] J. J. MacFarlane, I. E. Golovkin, P. Wang, P. R. Woodruff, N. A. Pereyra, SPECT3D: A multi-dimensional collisional-radiative code for generating diagnostic signatures based on hydrodynamics and PIC simulation output, High Energy Density Phys. **3**, 181 (2007).

[46] F. Weilacher, P. B. Radha, and C. Forrest, Three-dimensional modeling of the neutron spectrum to infer plasma conditions in cryogenic inertial confinement fusion implosions, Phys. Plasmas **25**, 042704 (2018).

[47] K. M. Woo, R. Betti, O. M. Mannion, C. J. Forrest, J. P. Knauer, V. N. Goncharov, P. B. Radha, D. Patel, V. Gopalaswamy, and V. Yu. Glebov, Inferring thermal ion temperature and residual kinetic energy from nuclear measurements in inertial confinement fusion implosions, Phys. Plasmas **27**, 062702 (2020).

[48] J. H. Kunimune *et al.*, Knock-on deuteron imaging for diagnosing the morphology of an ICF implosion at OMEGA, Phys. Plasmas **29**, 072711 (2022).