An ultra-fast method for generating synthetic down-scattered neutron data for inertial confinement fusion implosions

J. H. Kunimune,[1, a)] D. T. Casey,[2] and J. A. Frenje[1]
[1)] *Plasma Science and Fusion Center, Massachusetts Institute of Technology. 167 Albany St., Cambridge, Massachusetts, USA.*
[2)] *Lawrence Livermore National Laboratory, 7000 East Ave., Livermore, California, USA.*

(Dated: 7 March 2026)

In inertial confinement fusion experiments at the National Ignition Facility, asymmetries are probed by a variety of neutron diagnostics including neutron imaging systems, real-time neutron activation diagnostics (RTNADs), and neutron spectrometers. It is often useful to generate synthetic data based on these diagnostics to validate and tune models. However, current methods of doing so using Monte Carlo particle tracing are slow. In this paper, an ultra-fast method is presented for generating synthetic neutron images, RTNAD data, and spectrometry data using line integrals and 3D convolutions. While it does not contain as much physics as particle tracing codes, it is thousands of times faster and produces nearly identical data. This enables analysis techniques that depend on generating large amounts of synthetic data, and will thus prove very useful to the study of asymmetries going forward.

I. INTRODUCTION

In inertial confinement fusion (ICF) experiments at the National Ignition Facility (NIF), spherical capsules of deuterium–tritium (DT) fuel are compressed to high pressures using x-ray radiation. At the time of peak compression, the fuel forms a central hot-spot with a temperature of about 10 keV and a density of about 50 g/cm^3, surrounded by a dense fuel-shell with a density of about 500 g/cm^3 and a temperature of about 200 eV, as shown schematically in Figure 1. The conditions in the hot-spot are sufficient for thermonuclear ignition to occur, and the assembly's inertia holds it in place long enough for the thermonuclear burn wave to propagate into the dense fuel (about 100 ps).[1]

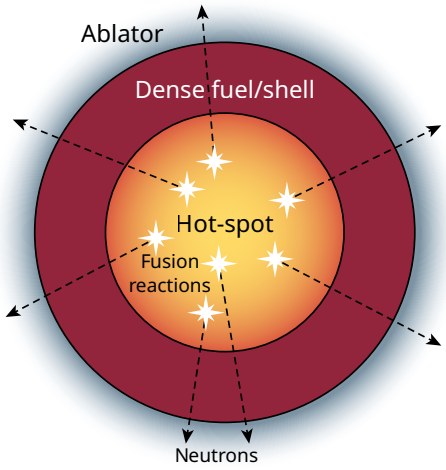

FIG. 1. A schematic of an ICF implosion at the time of peak compression. Fusion reactions occur in the central hot-spot, which is surrounded by a relatively cold dense fuel-shell.

Critical to the performance of such experiments is the symmetry with which the fuel is imploded. Ideally, the implosion would be perfectly spherical, resulting in a maximal transfer of kinetic energy from the radiation drive to the hot-spot thermal energy. In practice, asymmetries in the drive and capsule manufacturing result in non-spherical shapes. These asymmetries waste energy and reduce the fusion yield.[2,3] Thus, diagnosing them is a key goal of ICF research.

In this context, neutron imaging is an important diagnostic. At the NIF, three neutron imagers are used to measure neutron sources from three lines of sight.[4,5] Two of them are capable of distinguishing neutrons in two energy ranges. Primary neutrons, with energies between 13 and 17 MeV, are emitted directly from the hot-spot; primary neutron images are images of the hot-spot. Down-scattered neutrons, with energies between 6 and 12 MeV, elastically scatter off of fuel atoms before exiting the implosion; down-scattered neutron images are combined with primary neutron images to provide approximate images of the fuel density, in a process known as fluence compensation.[6]

Another diagnostic used for measurement of asymmetry is the real-time neutron activation detectors (RTNADs). These comprise up to 48 zirconium caps placed around the target chamber, each measuring the primary neutron yield. In a spherical implosion, the primary neutron emission is isotropic and all detectors record the same value. When asymmetries are present, though, variations in the dense fuel's areal density result in variations among the detectors.[7]

Another is the suite of five neutron time-of-flight spectrometers, which measure the down-scatter ratio (DSR) on multiple lines of sight. The DSR is the yield ratio of down-scattered neutrons between 10 and 12 MeV to pri-

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

mary neutrons between 13 and 15 MeV. The angularly-averaged DSR is a measure of compression. In addition, it is measured along multiple lines of sight, providing a measure of asymmetry similar to that provided by the RTNADs.[8]

Diagnostics like these can be used to validate and tune models of ICF implosions. An important aspect of this is the generation of synthetic data – in order to judge whether a given model is consistent with the data, synthetic neutron images and RTNAD data must be created based on the modeled profiles and first principles such that they can be compared to the experimental data. Unfortunately, generating synthetic data can be slow. Down-scattered neutron images, RTNAD data, and DSR data obtained along multiple lines of sight are complicated to calculate, and as a result are often done using time-intensive Monte Carlo particle tracing.

In this paper, a simple method for generating down-scattered neutron images and RTNAD data using convolutions is presented. Synthetic data is generated and presented for an example morphology generated by a 3D multi-rocket-piston simulation. A simpler example with more numerical details is also given in Appendix A This method makes several simplifying assumptions, but produces synthetic data that are practically indistinguishable from more exact data produced by Monte Carlo particle tracing – and in one quadrillionth the time.

II. CALCULATION OF PRIMARY NEUTRON IMAGES

As a test case, a multi-rocket-piston model[9] is used to obtain realistic 3D profiles of ion temperature $T_{\text{ion}}(\vec{\mathbf{r}})$, DT density ρ_{DT}, and carbon density ρ_{C}. Figure 2 shows the ion temperature profile in the hot-spot and the extent of the dense DT fuel surrounding it, and Figure 3 shows a lineout of the density profiles. This model generates an isobaric burning hot-spot surrounded by layers of DT and carbon with zero temperature. However, any combination of profiles is possible, including ones where the carbon is mixed into the DT, or where there is significant burn in the shell. It is assumed that these profiles are stationary in time.

Given such profiles, the simplest datum to synthetically generate is the primary neutron image. Assuming that attenuation by the shell is uniform and that neutron emission is isotropic at all points, the primary neutron image can be computed as a line integral of the fuel density and fusion reactivity as follows:

$$I_{\text{prim}}(x,y) = \int_{-\infty}^{\infty} S_{\text{prim}}(x\hat{\mathbf{x}} + y\hat{\mathbf{y}} + z\hat{\mathbf{z}})\mathrm{d}z \tag{1}$$

$$S_{\text{prim}}(\vec{\mathbf{r}}) = \int_{-\infty}^{\infty} \frac{1}{4\pi}\left(\frac{\rho_{\text{DT}}(\vec{\mathbf{r}})}{m_{\text{d}} + m_{\text{t}}}\right)^2 \langle\sigma v\rangle(\vec{\mathbf{r}})\,\mathrm{d}t \tag{2}$$

Here, $\hat{\mathbf{z}}$ is the unit vector in the direction of the imager, $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ are unit vectors along the imager's x- and y-axes, m_{D} is the mass of a deuterium atom, m_{T} is the mass of a tritium atom, $\langle\sigma v\rangle$ is the DT reactivity, and $\int \mathrm{d}t$ is an integral over time.

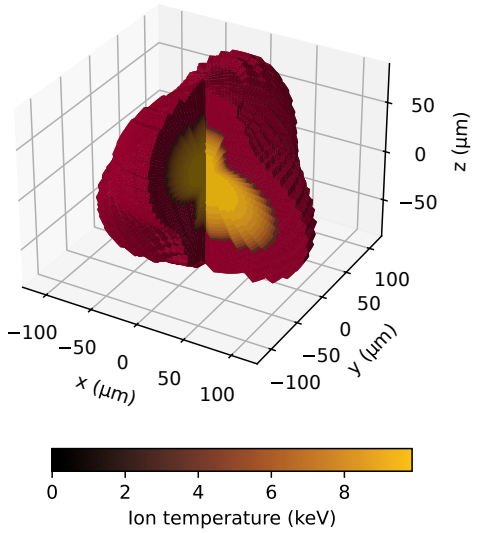

FIG. 2. The 3D morphology of an ICF implosion perturbed by a spherical harmonic with $l = 3$ and $m = 2$, generated by a multi-rocket-piston model. The dark red is the dense fuel, and the orange is the hot-spot, which is color-coded by ion temperature. Not shown is the remaining carbon ablator surrounding the dense fuel.

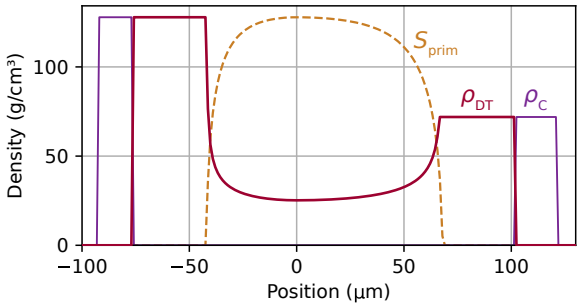

FIG. 3. A lineout of the morphology in Figure 2, showing the shape of the neutron source and the density profiles. The two sides differ in radius and density due to the shape asymmetry.

This line integral is evaluated for 10 000 pairs of x- and y-values, resulting in a 100×100 pixel grid.

In reality, asymmetries in the shell can cause variations in the attenuation of primary neutrons (indeed, it is precisely this variation that RTNADs measure). While it is possible to account for this without affecting this algorithm's time-complexity, the variations along a single imaging line of sight tend to be very small, and this does not affect the result significantly.

To demonstrate this, this method is compared to a more precise method using the Monte Carlo particle tracing code MCNP,[10] which is a widely used and thoroughly

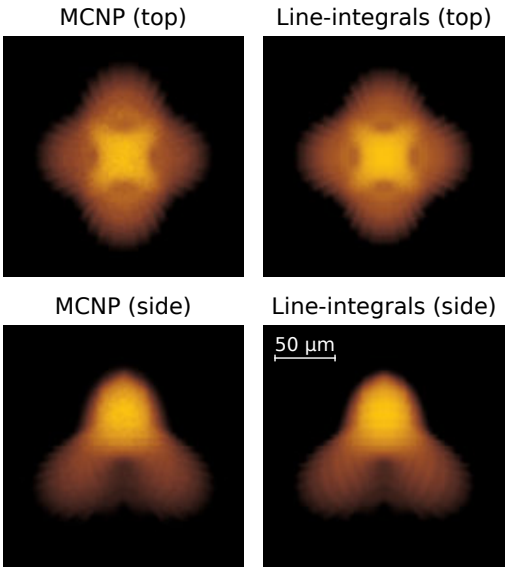

FIG. 4. Comparison of the primary neutron images of the implosion shown in Figure 2, as generated by MCNP (left) and the line integral method (right). The top images are taken from above while the bottom images are taken from the side of the implosion.

benchmarked code that allows the tallying of neutrons at a point by energy and direction of motion. The simulated morphology is built in MCNP out of geometric primitives, and neutrons are generated from each of the hot-spot cells in proportion to the reactivity and volume of each cell. Neutrons are subsequently allowed to scatter in the DT mass contained in each of the hot-spot and dense fuel cells, as well as the carbon mass contained in the ablator cells beyond them. Pinhole-image tallies that exclude neutrons below 13 MeV are used to obtain synthetic primary neutron images. As this is a Monte Carlo approach, these images contain statistical error. However, they thoroughly account for the neutron transport, including attenuation, absorption, inelastic scattering, neutrons that scatter multiple times, and down-scattered neutrons that retain enough energy to be recorded in the primary neutron image.

The line integral method and MCNP are both used to generate synthetic primary neutron images for the morphology shown in Figure 2. The comparison for images obtained along two lines of sight is shown in Figure 4. The line integral method has none of the statistical noise that is present in the MCNP image, but aside from that, there are no significant visible differences.

Running both on the same computing system, the line integral method took 0.0384 CPU-seconds to generate three images, while MCNP took 20 CPU-hours. MCNP can be made faster by running with a smaller number of particles, but this would increase the amount of statistical error in the image.

III. CALCULATION OF DOWN-SCATTERED NEUTRON IMAGES

The down-scattered neutron images are complicated to compute, because down-scattered neutron emission is not isotropic. Rather, it depends on the angle θ between the trajectory of the incoming primary neutron and that of the outgoing down-scattered neutron.

By assuming that each down-scattered neutron experiences exactly one scattering event, and that primary neutron attenuation is negligible, this direction-dependent flux at each scattering location $\vec{\mathbf{r}}$ can be expressed as an integral over all possible birth locations $\vec{\mathbf{r}}'$. Since each down-scattered neutron imaging system only sees neutrons that end up going toward it, the final trajectory can be treated as fixed.

$$I_{\mathrm{DS}}(x,y) = \int_{-\infty}^{\infty} S_{\mathrm{DS}}(x\hat{\mathbf{x}} + y\hat{\mathbf{y}} + z\hat{\mathbf{z}})\mathrm{d}z \tag{3}$$

$$S_{\mathrm{DS}}(\vec{\mathbf{r}}) = \iiint \frac{S_{\mathrm{prim}}(\vec{\mathbf{r}}')}{(\vec{\mathbf{r}} - \vec{\mathbf{r}}')^2} \cdot \frac{\rho(\vec{\mathbf{r}})}{m_{\mathrm{d}} + m_{\mathrm{t}}} \cdot \tag{4}$$

$$\left(\frac{\mathrm{d}\sigma_{\mathrm{d}}}{\mathrm{d}\Omega}(\theta) + \frac{\mathrm{d}\sigma_{\mathrm{t}}}{\mathrm{d}\Omega}(\theta)\right)\mathrm{d}^3\vec{\mathbf{r}}' \tag{5}$$

$$\theta = \arccos\left(\hat{\mathbf{z}} \cdot \frac{\vec{\mathbf{r}} - \vec{\mathbf{r}}'}{|\vec{\mathbf{r}} - \vec{\mathbf{r}}'|}\right) \tag{6}$$

Here, $\frac{\mathrm{d}\sigma_{\mathrm{d}}}{\mathrm{d}\Omega}(\theta)$ is the differential scattering cross-section of a 14 MeV neutron on a deuterium nucleus (clipped to only be nonzero for angles that generate neutrons in the range 6–12 MeV), $\frac{\mathrm{d}\sigma_{\mathrm{t}}}{\mathrm{d}\Omega}(\theta)$ is the same for a tritium nucleus, m_{d} and m_{t} are the masses of deuterium and tritium, respectively, and $\rho(\vec{\mathbf{r}})$ is the mass density.

In practice, the shell of an ICF implosion is not just composed of DT, but also contains remnants of the carbon ablator. This can be accounted for by generating multiple down-scattered neutron images – one for each material – and summing them together. However, because the carbon only has a small contribution to the neutron scattering, it is easier and almost as accurate to do both materials at once by using an effective DT density: $\rho = \rho_{\mathrm{DT}} + \frac{7}{29}\rho_{\mathrm{C}}$. This ignores differences in the shapes of the differential scattering cross-sections of DT and carbon, approximating the carbon cross-section as 24% of the DT cross-section – a good approximation for the neutron energies involved.[11]

Calculating this quadruple integral numerically would be computationally expensive. However, S_{DS} can be written as a convolution followed by a point-wise multiplication, which can be performed relatively quickly using fast Fourier transforms:

$$S_{\mathrm{DS}} = (S_{\mathrm{prim}} * f) \odot \rho \tag{7}$$

$$f(\Delta\vec{\mathbf{r}}) = \frac{1}{|\Delta r|^2} \cdot \frac{1}{m_{\mathrm{d}} + m_{\mathrm{t}}} \cdot \left(\frac{\mathrm{d}\sigma_{\mathrm{d}}}{\mathrm{d}\Omega}(\theta) + \frac{\mathrm{d}\sigma_{\mathrm{t}}}{\mathrm{d}\Omega}(\theta)\right) \tag{8}$$

$$\theta(\Delta\vec{\mathbf{r}}) = \arccos\left(\hat{\mathbf{z}} \cdot \frac{\Delta\vec{\mathbf{r}}}{|\Delta r|}\right) \tag{9}$$

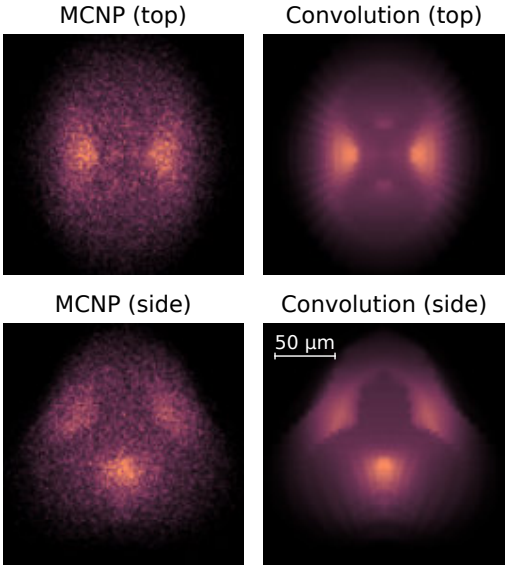

FIG. 5. Comparison of the down-scattered (6–12 MeV) neutron images of an implosion, without fluence compensation, as generated by MCNP (left) and the convolution method (right). The top images are taken from above while the bottom images are taken from the side of the implosion.

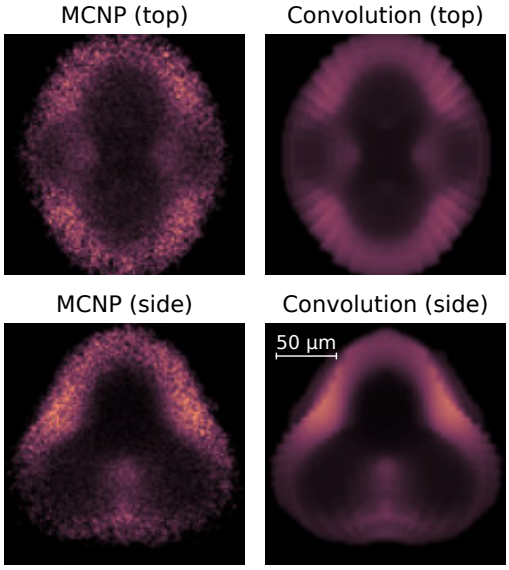

FIG. 6. Comparison of the fluence-compensated down-scattered (6–12 MeV) neutron images of an implosion as generated by MCNP (left) and the convolution method (right). The top images are taken from above while the bottom images are taken from the side of the implosion.

When numerically evaluating the kernel $f(\Delta\vec{\mathbf{r}})$, one must take care around the singularity where $\Delta\vec{\mathbf{r}} = 0$ and θ is undefined. If f is evaluated at the center of each voxel, then the central voxel will produce a value of *Not a Number*. The most rigorous way to avoid this is to calculate the average value of f in each voxel, which is finite for any finite voxel size. Alternatively, one can simply set $f(0)$ to zero, or evaluate $f(0)$ as if $\theta = 0$ and $|\Delta r| = L/2$ where L is the voxel size. If the number of voxels is sufficiently high, the difference between these treatments is small.

To compute a full down-scattered neutron image, S_{DS} only needs to be computed once. Multiple line integrals are subsequently performed to calculate I_{DS} in a 100×100-pixel grid, producing an image in the same manner as is done for the primary neutron images.

This neglects neutrons that scatter more than once, which are more isotropic than single-scatter neutrons. However, for current areal densities, the probability of a single scattering event is small (less than 10%). Because of this, and because doubly-downscattered neutrons tend to have low energies (below what is detected by the neutron imaging systems), neglecting them does not significantly affect the image. This approach also neglects the effect of attenuation, which makes the actual primary neutron fluence at each point less than the $\iiint S_{\mathrm{prim}}/\Delta r^2 \, \mathrm{d}^3\vec{\mathbf{r}}$ that's assumed here. However – again, because the scattering probability is small – this effect does not have a significant effect on the resulting image.

This method is again compared to MCNP. Altering the MCNP setup from before to record down-scattered neutron images rather than primary neutron images is as easy as changing the energy bounds of the synthetic detector from 13+ MeV to 6–12 MeV. The result for the morphology shown in Figure 2 is shown in Figure 5. To exaggerate the key features of the images for a more discerning comparison, the corresponding fluence-compensated neutron images are also shown in Figure 6. Again, the differences are small, aside from the large amount of statistical noise in the MCNP images. In this case, the convolution method took 4.20 CPU-seconds for three images, while the MCNP-generated images were produced at the same time as the primary images shown in Figure 4, thus requiring the same 20 CPU-hours.

IV. CALCULATION OF RTNAD AND DSR DATA

The unscattered yield measured by an RTNAD depends entirely on the average areal density seen by neutrons exiting the implosion in that direction, $\langle \rho L \rangle$. This can be expressed as the following line integral.

$$Y_{\mathrm{prim}} = Y_{\mathrm{total}} \exp\left(-\frac{\sigma_{\mathrm{d}} + \sigma_{\mathrm{t}}}{m_{\mathrm{d}} + m_{\mathrm{t}}} \langle \rho L \rangle \right) \tag{10}$$

$$\langle \rho L \rangle = \frac{\iiint S_{\mathrm{prim}}(\vec{\mathbf{r}}') \int_0^\infty \rho(\vec{\mathbf{r}}' + z\hat{\mathbf{z}}) \mathrm{d}z \mathrm{d}^3\vec{\mathbf{r}}'}{\iiint S_{\mathrm{prim}}(\vec{\mathbf{r}}') \mathrm{d}^3\vec{\mathbf{r}}'} \tag{11}$$

Here, $\hat{\mathbf{z}}$ is the unit vector in the direction of the RTNAD, and ρ is again the effective DT density accounting for car-

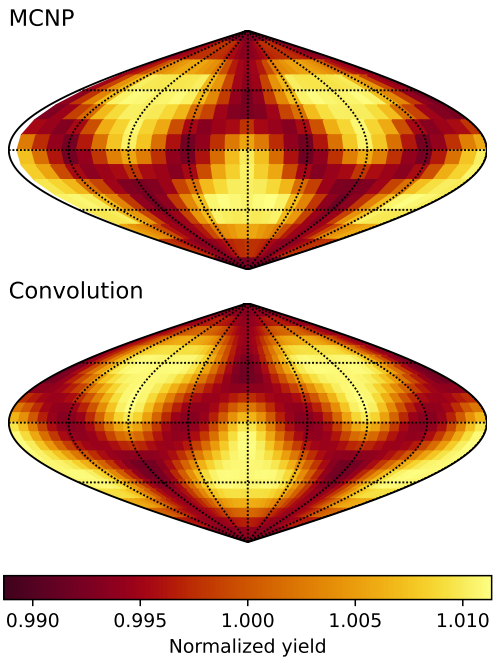

FIG. 7. Comparison of the RTNAD skymap of an implosion as generated by MCNP (top) and by the convolution method (bottom). For the sake of visualization, the normalized primary neutron yield is shown for every possible line of sight, rather than for the 48 discrete lines of sight the RTNADs actually probe.

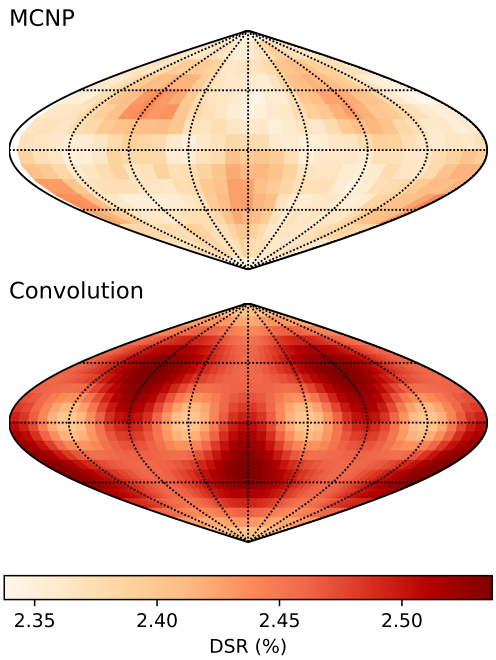

FIG. 8. Comparison of the DSR skymap of an implosion as generated by MCNP (top) and by the convolution method (bottom). For the sake of visualization, the DSR is shown for every possible line of sight, rather than for the five discrete lines of sight the neutron spectrometers actually probe. The same colormap is used for both, displaying a difference in magnitude between them.

bon. σ_d and σ_t are the total scattering cross-sections of 14 MeV neutrons on deuterium and tritium, respectively.

This quadruple integral would be fairly computationally expensive, but can be written as a line integral of a convolution, which can be performed quickly using fast Fourier transforms:

$$\langle \rho L \rangle = \int_0^\infty \tilde{\rho}(z\hat{\mathbf{z}})\mathrm{d}z \tag{12}$$

$$\tilde{\rho} = \frac{S^*_\mathrm{prim} * \rho}{\iiint S_\mathrm{prim}(\vec{\mathbf{r}}')\mathrm{d}^3\vec{\mathbf{r}}'} \tag{13}$$

$$S^*_\mathrm{prim}(\vec{\mathbf{r}}) = S_\mathrm{prim}(-\vec{\mathbf{r}}) \tag{14}$$

As $\tilde{\rho}$ doesn't depend on $\hat{\mathbf{z}}$, the convolution only needs to be performed once, and Y_prim can be calculated along all 48 lines of sight using only inexpensive line integrals.

This is fairly exact; the only errors arise from the finite resolution of the arrays that are convolved and the approximation of the carbon differential scattering cross-section.

This method is once again compared to MCNP. To produce synthetic RTNAD data in MCNP, synthetic detectors are placed on a sphere around the implosion, measuring the 13–15 MeV neutron yield at each point. This can be done at the same time as computing the primary and down-scattered neutron images. The result for a continuum of lines of sight is shown in Figure 7. The convolution-generated skymap looks very similar to the MCNP-generated skymap. In this case, the convolution method took 0.343 CPU-seconds, compared to MCNP's 20 CPU-hours.

The DSR skymap can be inferred from the RTNAD skymap, shown in Figure 7, using a spherical convolution,[12] or calculated directly in MCNP by tallying 10 –12 MeV neutrons. The DSR skymaps from the convolution and from MCNP are compared in Figure 8. Like the RTNAD skymap, they have the same asymmetry, but unlike the RTNAD skymap, the magnitudes differ noticeably. The convolution method predicts an angularly averaged DSR of 0.0246, compared to the 0.0238 predicted by MCNP (a 3% difference). This may be due to differences between the cross-section data used by MCNP and the convolution algorithm.

V. COMPARISON TO EXPERIMENTAL DATA

While images generated by this method cannot be benchmarked directly to real experimental data – as the 3D density and temperature profiles are not generally known for an experiment – simulations can be tuned to approximate real experiments and thus provide rough comparisons. The NIF shot of June 5, 2021, is taken as an example. Based on the neutron images and RTNAD data from this shot, a 3D rocket-piston simulation

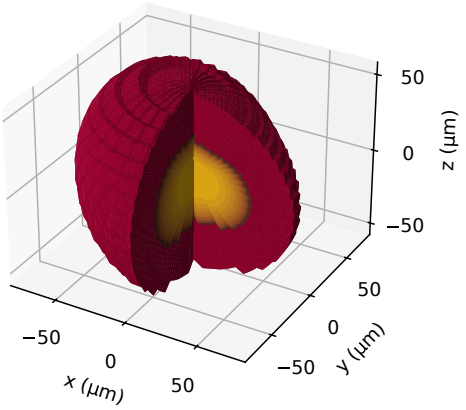

FIG. 9. The 3D morphology of a simulated ICF implosion tuned to approximate the NIF shot on June 5, 2021. The dark red is the dense fuel, and the orange is the hot-spot.

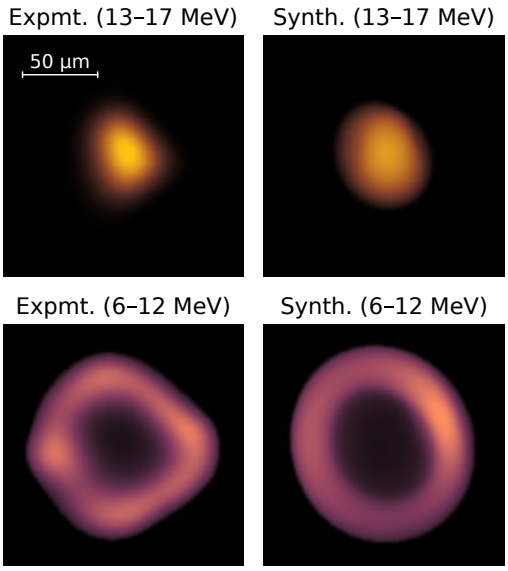

FIG. 10. Comparison of the neutron images of the NIF shot on June 5, 2021 (left), with the synthetic images generated by line integrals and 3D convolutions for the simulation shown in Figure 9 (right). The top row shows primary neutron images, while the bottom row shows fluence-compensated down-scattered images. All four images are from the same line of sight.

was constructed that matched the observed size, mode-1 asymmetries, and mode-2 asymmetries. Like the simulation shown in Figure 2, this simulation assumes an isobaric hot-spot with a centrally peaked temperature surrounded by layers of cold DT and carbon. The simulated morphology at bang-time is shown in Figure 9.

Neutron images are generated from the simulation and

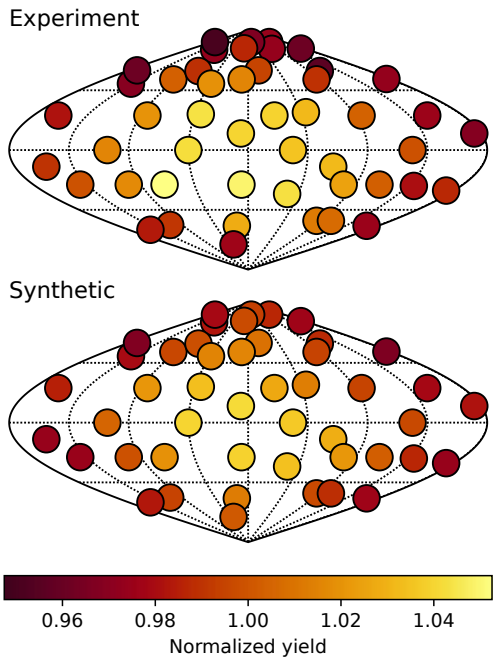

FIG. 11. Comparison of the RTNAD data from the NIF shot on June 5, 2021 (top), with the synthetic data generated by line integrals and 3D convolutions for the simulation shown in Figure 9 (bottom). Each circle represents one detector, and its location represents its position on the target chamber.

compared to the measured images. The primary neutron images and fluence-compensated down-scattered neutron images from the side of the implosion are shown in Figure 10. While the exact shape of the images differs because the simulation only accounts for mode-1 and mode-2 asymmetries, the general shape of the primary neutron image and the hollow structure of the fluence-compensated down-scattered neutron image are accurately reproduced by the simulated data. The experimental images (especially the primary neutron image) have softer edges than the simulated images. This is most likely because of the analysis technique used to extract the experimental images from the raw neutron imaging system data, which has imperfections due to counting statistics and metrological errors. A blur effect can be applied to the synthetic images to better approximate the experimental images. This may also be an indication that the true temperature profile was more peaked than the profile assumed in the simulation.

RTNAD data is also generated from the simulation and compared to the measured RTNAD data, and shown in Figure 11. There is a mode-1 asymmetry in the data which has been well replicated in the simulated data. Some discrepancies can be seen in individual detectors. This is most likely dominated by random error in the experimental measurements, which is not present in the simulation.

Comparisons like this can be used not only to evalu-

ate the accuracy of simulations data, but also to tune simulations to match experimental data. The speed of this synthetic data generation technique enables a great number of techniques that would be intractable with a slower Monte Carlo approach. For example, computationally intensive iterative techniques like gradient descent or Markov chain Monte Carlo can be used to find the simulation that best fits an experimental dataset. Alternatively, a large number of synthetic datasets can be generated and used to train a machine learning model such as a neural network to directly reconstruct data.[13] These techniques allow unique insights to be obtained from ICF implosion data.

VI. CONCLUSION

In ICF experiments at the NIF, x-ray radiation is used to implode a spherical capsule of fuel to high densities and temperatures for the purposes of thermonuclear fusion. Three important diagnostics used to diagnose asymmetry in ICF experiments are neutron imaging, RTNADs, and neutron spectrometry. Using these diagnostics to validate and tune models often requires that synthetic data be generated from first-principles. However, generating synthetic data for these diagnostics can be complicated and extremely time consuming.

In this work, an ultra-fast method is presented that uses line integrals and 3D convolutions, which can be evaluated quickly using fast Fourier transforms, to quickly produce high-quality synthetic neutron images, RTNAD data, and DSR data. This method makes several simplifying assumptions, but these assumptions are all well-founded and do not significantly affect the output, as is demonstrated by the comparison to MCNP-generated data.

While MCNP captures more physics than the line integral and convolution method – such as attenuation and multiple-scattered neutrons – and it allows an entire dataset to be generated all at once, the line integral and convolution outperforms it in two important ways. First, it is much faster, to such an extent that even when generating three primary neutron images, two down-scattered neutron images, and an RTNAD skymap sequentially, the line integral and convolution method is sixteen thousand times faster than using MCNP. Second, it generates data free of statistical error, which is prevalent in MCNP-generated data.

Overall, the great speed increase provided by this method over MCNP, as well as the fact that it is easy to implement, enables analysis techniques that depend on generating large amounts of synthetic data, such as forward-fitting a model to measured data, or generating a synthetic dataset to train a machine learning model.[13] It is also valuable for a wide variety of techniques and model studies that do not require large amounts of synthetic data, but benefit from shorter computation times regardless. This will thus prove very useful to the study of asymmetries in ICF implosions 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 acknowledges 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-2012975.

AUTHOR DECLARATIONS

Conflict of interest statement

The authors have no conflicts to disclose.

Author contributions

J. H. Kunimune: investigation (equal), writing (lead). D. T. Casey: investigation (equal), writing (supporting). J. A. Frenje: writing (supporting), supervision.

Data availability statement

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

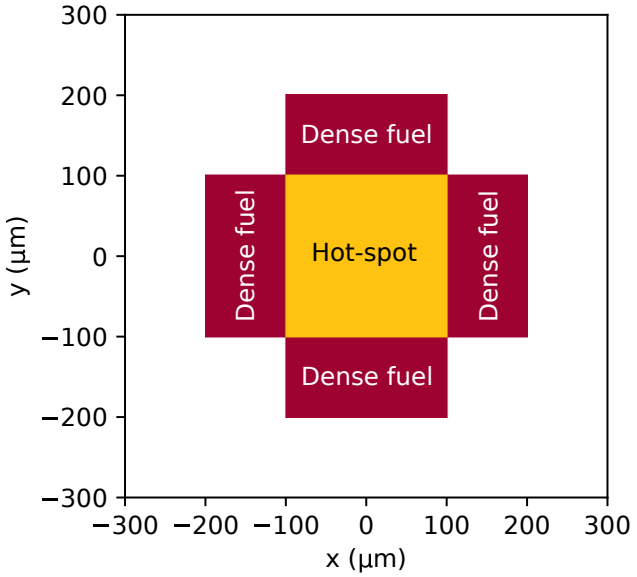

FIG. 12. A low-resolution two-dimensional morphology for an example calculation

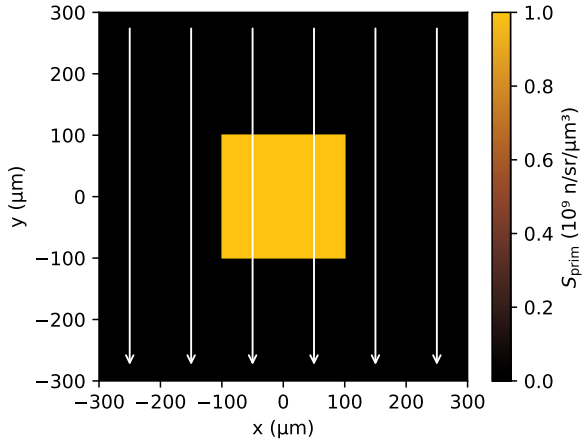

FIG. 13. The primary neutron emission, along with the lines along which it will be integrated to calculate the primary neutron image for an imager looking from below

Appendix A: An example calculation

To illustrate the equations in more specific terms, a calculation is performed here for a simple geometry with an extremely low resolutions. It comprises a 4×4×1 grid of voxels, each 100 μm in size. The central four voxels are designated as hot-spot, with an ion temperature of 9.20 keV and a DT density of 10 g/cm^2. The eight voxels directly adjacent to them are designated as dense shell, with an ion temperature of 0 keV and a DT density of 100 g/cm^2. This is illustrated in Figure 12. The hot-spot and shell are both confined to a single layer of voxels in the z-direction to enable easy visualization. For this example, there is no carbon.

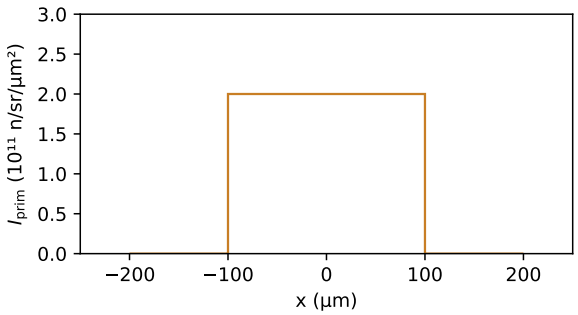

FIG. 14. The primary neutron image from the negative y-direction. Since the morphology is two-dimensional, the image is one-dimensional.

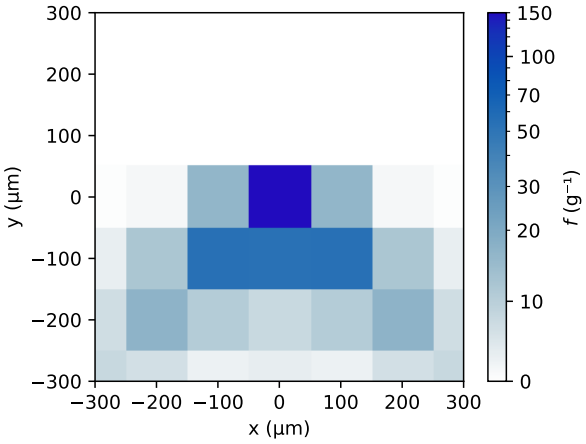

FIG. 15. The convolution kernel used to convert primary neutron emission to down-scattered neutron emission, for an imager looking from below

First, a primary neutron image is calculated. At 9.20 keV, the reactivity of DT is 8.78×10^{-17} cm^3/s. The neutron emission $S_{\text{prim}}(\vec{\mathbf{r}})$ as calculated by Equation 1 is therefore 1.00×10^{21} n/sr/cm^3. Its spatial extent is shown in Figure 13. To calculate the primary neutron image from the negative y-direction, one simply integrates this map over the shown vertical lines. The resulting primary image, shown in Figure 14, is a boxcar distribution 200 μm wide and 2.00×10^{19} n/sr/cm^2 tall.

To generate the down-scattered neutron image, we first calculate the scattering kernel $f(\Delta\vec{\mathbf{r}})$ as defined in Equation 8. This is shown in Figure 15 for an imager looking from the negative y-direction. The function f has been averaged throughout the volume of each voxel to ensure that it is finite everywhere. Convolving it with $S_{\text{prim}}(\vec{\mathbf{r}})$ and then multiplying element-wise by $\rho(\vec{\mathbf{r}})$, as per Equation 7, results in the down-scattered neutron emission map $S_{\text{DS}}(\vec{\mathbf{r}})$ shown in Figure 16. This is again integrated along the vertical lines, resulting in the image shown in Figure 17. The brightest pixels in the image are the central two, at 1.87×10^{17} n/sr/μm^2, primarily from the

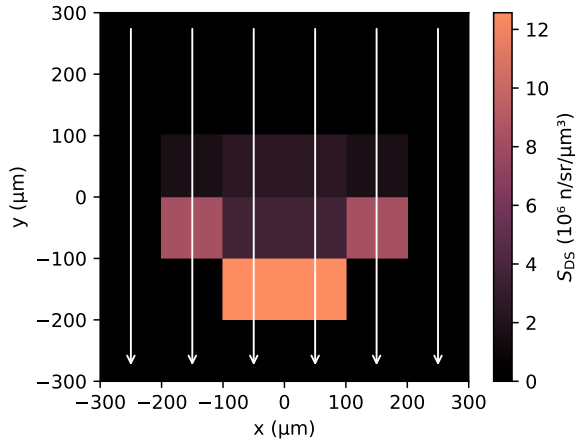

FIG. 16. The emission of 6–12 MeV neutrons traveling in the negative y-direction, along with the lines along which it will be integrated to calculate the down-scattered neutron image

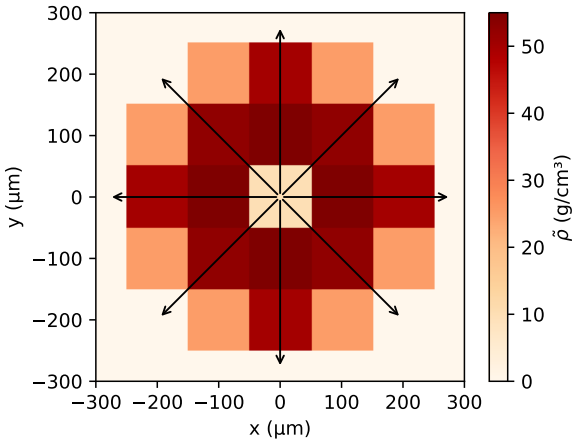

FIG. 18. The relative density distribution $\tilde{\rho}(\Delta\vec{\mathbf{r}})$, along with the radial lines along which it is integrated to obtain the average areal density

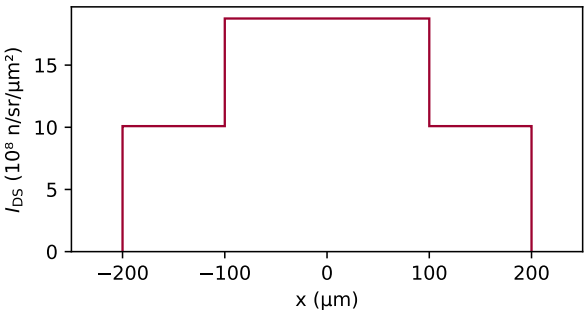

FIG. 17. The down-scattered neutron image from the negative y-direction. Since the morphology is two-dimensional, the image is one-dimensional.

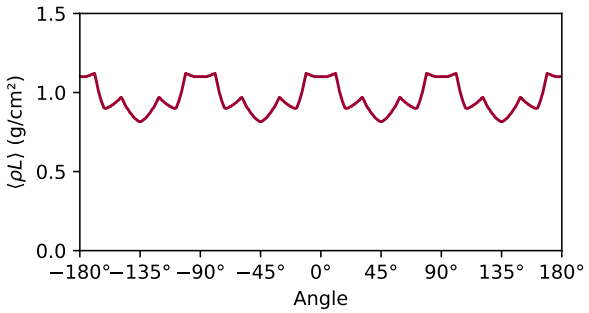

FIG. 19. The average areal density measured in all directions

dense fuel voxels that are directly backlit by the hot-spot. The two pixels on the sides also have some brightness due to side-scatter from the hot-spot.

To generate the $\langle\rho L\rangle$ skymap, which can then be used to extract RTNAD and DSR data, we first calculate the relative density distribution $\tilde{\rho}(\Delta\vec{\mathbf{r}})$ as defined in Equation 13. This distribution, shown in Figure 18, represents the average density distribution seen by a random primary neutron relative to its birth location. One obtains $\langle\rho L\rangle$ by integrating this along the shown radial lines. The result for a continuum of angles is shown in Figure 19. At multiples of 90°, it equals 1.10 g/cm^2, because every hot-spot voxel sees 1 g/cm^2 of dense fuel plus either 0.05 or 0.15 g/cm^2 of hot-spot.

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]M. Gatu Johnson, B. D. Appelbe, J. P. Chittenden, J. Delettrez, C. Forrest, J. A. Frenje, V. Y. Glebov, W. Grimble, B. M. Haines, I. Igumenshchev, R. Janezic, J. P. Knauer, B. Lahmann, F. J. Marshall, T. Michel, F. H. Séguin, C. Stoeckl, C. Walsh, A. B. Zylstra, and R. D. Petrasso, "Impact of asymmetries on fuel performance in inertial confinement fusion," Phys. Rev. E **98**, 051201 (2018).

[3]D. T. Casey, B. J. MacGowan, J. D. Sater, A. B. Zylstra, O. L. Landen, J. Milovich, O. A. Hurricane, A. L. Kritcher, M. Hohenberger, K. Baker, S. Le Pape, T. Döppner, C. Weber, H. Huang, C. Kong, J. Biener, C. V. Young, S. Haan, R. C. Nora, S. Ross, H. Robey, M. Stadermann, A. Nikroo, D. A. Callahan, R. M. Bionta, K. D. Hahn, A. S. Moore, D. Schlossberg, M. Bruhn, K. Sequoia, N. Rice, M. Farrell, and C. Wild, "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).

[4]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).

[5]V. E. Fatherley, S. H. Batha, C. R. Danly, L. A. Goodwin, H. W. Herrmann, H. J. Jorgenson, J. I. Martinez, F. E. Merrill, J. A.

Oertel, D. W. Schmidt, P. L. Volegov, C. A. Wilde, D. N. Fittinghoff, M. J. Ayers, D. A. Barker, G. P. Grim, R. L. Hibbard, N. D. Shingleton, and M. A. Vitalich, "System design of the NIF neutron imaging system north pole," in *Target Diagnostics Physics and Engineering for Inertial Confinement Fusion VI*, Vol. 10390, edited by J. A. Koch and G. P. Grim, International Society for Optics and Photonics (SPIE, 2017) p. 103900F.

[6]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), https://aip.scitation.org/doi/pdf/10.1063/1.4960065.

[7]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).

[8]A. S. Moore, E. P. Hartouni, D. Schlossberg, S. Kerr, M. Eckart, J. Carrera, L. Ma, C. Waltz, D. Barker, J. Gjemso, E. Mariscal, G. Grim, and J. Kilkenny, "The five line-of-sight neutron time-of-flight (nToF) suite on the National Ignition Facility (NIF)," Rev. Sci. Instrum. **92**, 023516 (2021).

[9]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).

[10]J. A. Kulesza, T. R. Adams, J. C. Armstrong, S. R. Bolding, F. B. Brown, J. S. Bull, T. P. Burke, A. R. Clark, R. A. Forster, III, J. F. Giron, A. S. Grieve, C. J. Josey, R. L. Martz, G. W. McKinney, E. J. Pearson, M. E. Rising, C. J. Solomon, Jr., S. Swaminarayan, T. J. Trahan, C. A. Weaver, S. C. Wilson, and A. J. Zukaitis, "MCNP code version 6.3.1 theory & user manual," Tech. Rep. LA-UR-24-24602, Rev. 1 (Los Alamos National Laboratory, Los Alamos, NM, MSM, 2024).

[11]J. A. Frenje, "Nuclear diagnostics for inertial confinement fusion (ICF) plasmas," Plasma Phys. Control. Fusion **62**, 023001 (2020).

[12]D. T. Casey, O. L. Landen, E. Hartouni, R. M. Bionta, K. D. Hahn, P. L. Volegov, D. N. Fittinghoff, V. Geppert-Kleinrath, C. H. Wilde, J. L. Milovich, V. A. Smalyuk, J. E. Field, O. A. Hurricane, A. B. Zylstra, A. L. Kritcher, D. S. Clark, C. V. Young, R. C. Nora, D. A. Callahan, B. J. MacGowan, D. H. Munro, B. K. Spears, J. L. Peterson, J. A. Gaffney, K. D. Humbird, M. K. G. Kruse, A. S. Moore, D. J. Schlossberg, M. Gatu-Johnson, and J. A. Frenje, "Three dimensional low-mode areal-density non-uniformities in indirect-drive implosions at the National Ignition Facility," Phys. Plasmas **28**, 042708 (2021).

[13]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), https://pubs.aip.org/aip/rsi/article-pdf/doi/10.1063/5.0205656/20028063/073506_1_5.0205656.pdf.