ORNL/SPR-2018/925

Oak Ridge National Laboratory Progress Update on Segmentation of Neutron Tomographic Images, 2018

Matthew Blackston
Deniz Aykac
Paul Hausladen

September 2018

Approved for public release.
Distribution is unlimited.

OAK RIDGE NATIONAL LABORATORY
MANAGED BY UT-BATTELLE FOR THE US DEPARTMENT OF ENERGY

DOCUMENT AVAILABILITY

Reports produced after January 1, 1996, are generally available free via US Department of Energy (DOE) SciTech Connect.

Website www.osti.gov

Reports produced before January 1, 1996, may be purchased by members of the public from the following source:

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone 703-605-6000 (1-800-553-6847)
TDD 703-487-4639
Fax 703-605-6900
E-mail info@ntis.gov
Website http://classic.ntis.gov/

Reports are available to DOE employees, DOE contractors, Energy Technology Data Exchange representatives, and International Nuclear Information System representatives from the following source:

Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831
Telephone 865-576-8401
Fax 865-576-5728
E-mail reports@osti.gov
Website http://www.osti.gov/contact.html

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

ORNL/SPR-2018/925

Nuclear Security and Isotope Technology Division

PROGRESS UPDATE ON SEGMENTATION OF NEUTRON TOMOGRAPHIC IMAGES, 2018

Matthew Blackston
Deniz Aykac
Paul Hausladen

Date Published: September 2018

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, TN 37831-6283
managed by
UT-BATTELLE, LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

This work is supported by the U.S. Department of Energy, Office of Defense Nuclear Nonproliferation R&D in the National Nuclear Security Administration (NA-22).

1. INTRODUCTION

The present report documents the techniques developed to segment fast-neutron tomographic images of objects consisting of assemblies of machined parts with distinct boundaries. This report satisfies the fiscal year 2018 technical deliverable for project OR16-3DTomography-PD3Jb, "3D Tomography and Image Processing Using Fast Neutrons," to report on the extension of geometric shape-finding algorithms to three dimensions. The project has two overall goals. The first of these goals is to extend associated-particle fast-neutron transmission and, particularly, induced-reaction tomographic imaging algorithms to three dimensions. This aspect of the project is beyond the scope of this report. The second goal is to automatically segment the resultant tomographic images into constituent parts and then extract information about the parts, such as the class of shape and potentially the shape parameters. This report describes progress for the image segmentation part of the project.

Imaging techniques have been developed as high-confidence methods for confirming the presence and configuration of special nuclear materials (SNM). This high confidence is achieved at the cost of revealing considerable information that may be undesirable to share with the operator of the equipment. One potential way to make use of the high-confidence of imaging methods without revealing imaging data is to employ automated analysis that can extract meaningful attributes of the SNM without showing imaging data to the operator. An essential step of this automated analysis is the segmentation of the image into its constituent parts. For example, a three-dimensional (3D) fast-neutron tomographic image of an assembly containing uranium could be reconstructed and then segmented into discrete parts. The properties of a constituent part identified as uranium would be of interest and could be extracted. For instance, the volume and density of the part would together determine its mass. Moreover, its shape might be expected to fall into a particular geometric class, such as a cube, cylinder, or sphere, and the parameters of the shape (such as side length, height, or diameter) would be expected to fall in a given range. In such an analysis, automated identification of the boundaries and properties of the uranium part is an essential step.

A key goal of the present work is to infer the boundaries of objects in fast-neutron tomographic images via the general method of image segmentation. Development of this capability can be broken into more manageable steps, including

1. operator-guided segmentation of two-dimensional (2D) fast-neutron tomographic images,
2. operator-guided segmentation of 3D fast-neutron tomographic images,
3. development of automated segmentation algorithms, and
4. extraction of shape parameters from constituent volumes in 3D.

In the first year, the project team reported progress on the first step [1]. Last year, progress on the second step, the extension into 3D, was reported with application to simulated data [2]. The present report describes additional progress on the second step, and progress toward the third step, automated segmentation algorithms. In particular, a multiphase level set approach for image segmentation in three dimensions has been applied to measured 3D data for the first time. Activities related to the fourth step will also be reported, but it will focus on estimation of 3D volumes rather than shape parameters.

In this report, the steps before image segmentation will be described, and the multiphase level sets approach to image segmentation will be briefly reviewed. This will be followed by a discussion of the issues encountered when applying the level set technique to measured 3D transmission images, and then a description of the inclusion of induced-reaction data will be provided. Finally, segmentation results for various measured target assemblies will be presented and the image segmentation portion of this project will be summarized and future directions will be discussed.

2. DATA PROCESSING STEPS BEFORE SEGMENTATION

Before the segmentation process is applied to data from a fast-neutron imager, the raw data must undergo several processing steps to produce the 3D image reconstruction, or 3D image, which is the input to the segmentation algorithms. If problems exist in the data processing steps before the segmentation process, the segmentation may have problems correctly segmenting the images. The data processing steps will be described briefly so that some of these issues can be more clearly understood as they arise during a discussion of the results.

Figure 1 shows the entire data processing flow from data collection to segmentation. The analysis chain consists of the following steps:

- Data are collected with the target object placed at multiple heights and with multiple rotations relative to the imaging system. Multiple measurement heights were needed because the coincidence cone of tagged neutrons was not wide enough to view each target in its entirety with a single view. The data were collected with the target object on a table of fixed height, and the height of the imager was adjusted manually using a crank attached to a leadscrew upon which the entire imager was translated vertically. For most of the data in this report, three imager heights were sufficient to capture the entire target, though some took more. Target rotation is needed to perform tomography, which requires multiple lines of response through each point in the target. In this case, the imager was fixed rotationally, and the targets were rotated using a motorized turntable that allowed the relative view through the target to be modified by issuing commands to the rotary motor.

- Once the data are collected at all heights and target rotations, they must be calibrated and analyzed to produce projection data. A projection is a single view (i.e., rotation) through a target object. To produce projection data for transmission radiographs, events with times of flight and detection locations consistent with transmission are selected. In addition, the contributions from random coincidences are subtracted and a method to estimate the amount of small angle scatter, which produces events that are indistinguishable from transmission, is performed to subtract the scattered component from the data. At present, the scatter subtraction method is a coarse approximation of the scattered component and its accuracy is highly dependent on the target geometry and the amount of shielding present. This variation in the accuracy of the scatter correction results in inhomogeneities in the transmission images that cause problems for the segmentation results reported below.

  Projection data are also produced for the induced-reaction analysis for identifying fissionable materials and scatter from hydrogenous material. The induced-fission analysis relies on counting single or multiple lower-energy neutrons that arrive at times consistent with fission energies. The hydrogen scatter analysis isolates events that are detected with the time and scatter angle correlation expected for single elastic scatters from hydrogen. Both analyses employ additional techniques to subtract contributions from random events that meet the selection criteria as well as background from detector scatter and scatter from other materials in the target. Any inaccuracies in the subtraction of these backgrounds are propagated to the 3D reconstructed images, which can also cause problems when the induced reconstructions are used to inform the segmentation process.

- Once the projection data are created, they are fed to the iterative reconstruction code, which constitutes the other major component of this project. The expectation that machined parts have homogeneous attenuation coefficients inspired the use of a total-variation constraint in the iterative solver that rewards solutions that produce more uniform regional attenuation coefficients. In principle, the total variation should result in reconstruction data that are more easily segmented than would result from an unconstrained reconstruction. The details of the iterative reconstruction have been reported in detail in Reference [3] and will not be covered here. However, it should be noted that the development of the iterative reconstruction has been a work in progress throughout the course of this project. This has sometimes resulted in the generation of artifacts in the 3D reconstructions that may also cause problems for the segmentation algorithm.

The induced reconstructions are performed in the same way as transmission reconstructions, with the addition of efficiency factors to account for the fact that scattered and fission neutrons will not travel along the incident path and will have reduced energy. In addition, the induced-reaction reconstructions rely on knowledge of the transmission reconstruction to correct for attenuation of fission or scattered neutrons leaving the target.

The reconstructions result in 3D voxelized maps containing imaged quantities that depend on the imaging type. For transmission, the attenuation coefficient for 14 MeV neutrons, μ_{trans}, is mapped; for hydrogen scatter, a quantity proportional to attenuation coefficient for elastic scatter from hydrogen, M_{hyd}, is mapped; and for induced fission, a quantity related to the attenuation coefficient from induced fission, M_{fiss}, is mapped. The reconstructed quantities for the induced reactions are symbolized using M rather than μ to indicate that they are not yet the attenuation coefficients for those interactions. These normalizations are still being studied.

The reconstructed images are the input to the 3D multiphase level set algorithm, which performs the 3D segmentation of the images.

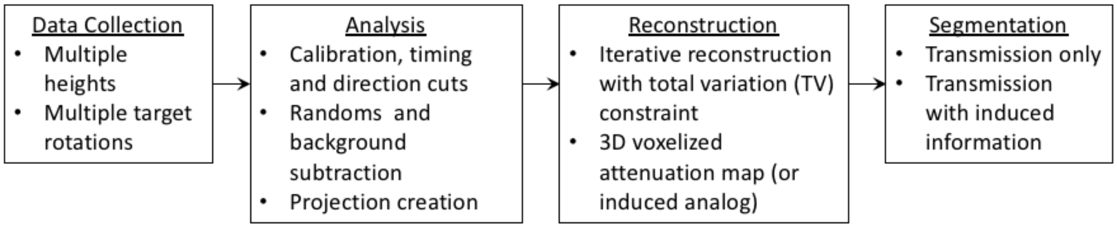

Figure 1. Diagram of the measurement and analysis flow. The image segmentation step comes after a significant amount of preprocessing.

3. REVIEW OF SEGMENTATION APPROACH

3.1 LEVEL SETS

In this section, the 3D multiphase level set formalism will be briefly reviewed. A description the 3D multiphase level set formalism was presented in last year's report [2]. Please refer to that report for full details of the implementation.

Suppose $g(x,y)$ is an image, such as an x-ray, map, or photograph. In many instances, the image g consists of several identifiable objects where, within the objects, the image varies smoothly or slowly, and the image varies discontinuously or rapidly across most of the boundaries between the objects. In these cases, the image is well modeled by a set of smooth functions, f_i, defined on a set of disjoint regions $R = \{R_i\}$ having boundaries $C = \{C_i\}$ that together cover the image. Then, the combination (f,C) is a generally recognizable simplification or cartoon of the original image that we refer to as a "segmentation" of the image. The optimal such segmentation is found by minimizing the Mumford–Shah energy functional [4],

$$E(f,C) = \mu^2 \iint_{R} (f-g)^2 dxdy + \iint_{R-C} \|\nabla f\|^2 dxdy + \nu|C|.$$

This functional has three terms, and the meaning of each term can be succinctly expressed. The first term requires that f approximates g. The second term requires that f does not vary much within each region R_i . The third term, sometimes referred to as the length term, requires that the boundaries C_i are as short as possible. The constants μ and ν are simply scale factors that govern the relative importance of the terms.

In the present work, we are chiefly interested in the case of tomographic images of machined parts. As a result, the approximation of constant f within regions R_i is realistic. With constant f, the second term in the functional is zero and can be eliminated entirely. In fact, constant f would exactly represent the "true" g, but the measured g will be a corrupted approximation of the true image that includes noise and the effects of limited resolution.

In general, Mumford and Shah do not provide a prescription for finding the optimal segmentation. However, Vese and Chan [5] provide a framework to calculate the optimal segmentation using a "level-set" formulation. This formulation is advantageous because it allows for automatic topology changes such as cusps and corners, operates on a grid, and automatically handles problems associated with gaps and overlaps between regions by exactly covering the domain by construction. In the Chan–Vese approach, the Mumford–Shah energy functional is rewritten in terms of a set of region boundaries C_i, where the C_i are represented as the zero-level sets of Lipschitz continuous functions $\emptyset_i$ such that $C_i = \{(x,y) | \emptyset_i(x,y) = 0\}$. The stationary points of this functional are recognized as being the solutions of the Euler–Lagrange equation, which identifies a set of partial differential equations whose solution gives the functions $\emptyset_i$ and the corresponding regions. Given an initial guess for the $\emptyset_i$, the partial differential equations provide an update equation to solve for the optimal $\emptyset_i$ by gradient descent.

4. SEGMENTATION OF MEASURED 3D TRANSMISSION DATA

The results from the application of the Chan–Vese multiphase level-set method with three level sets to a measured 3D image of a target consisting of a depleted uranium (DU) annulus, a high-density polyethylene (HDPE) block, and a hollow lead cube is shown in **Error! Reference source not found.**. In addition, tomographic slices through the bottom and top regions of the reconstruction data are shown in **Error! Reference source not found.** 3(a) and (c), with the segmentation results overlaid in (b) and (d).

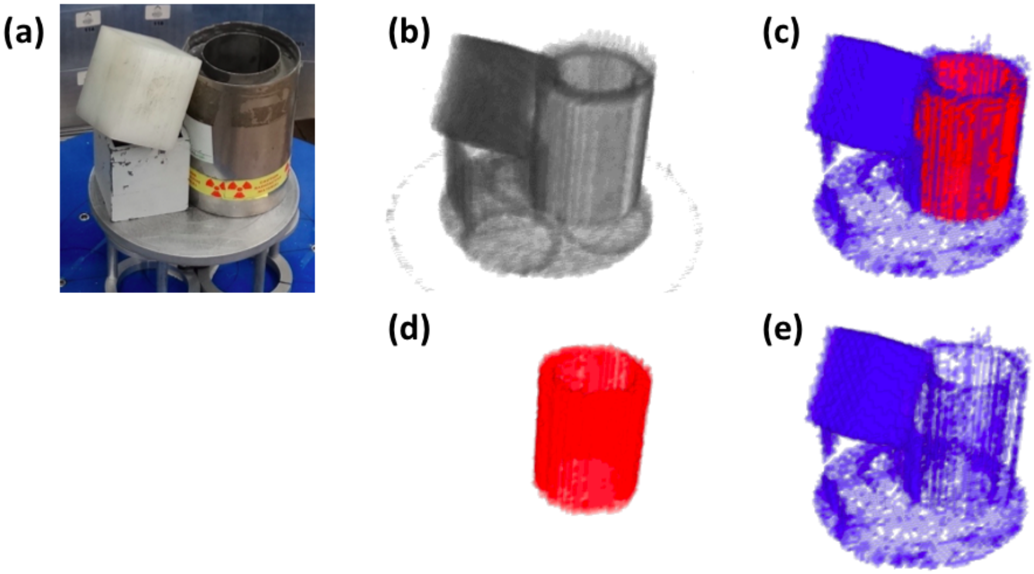

Figure 2. Results from the segmentation of the transmission image for a target consisting of a DU annulus, HDPE block, and hollow lead cube. (a) Photograph of the target object. The annulus is on the right, the HDPE block is white and sitting on the hollow lead cube, which is painted gray. (b) The reconstructed transmission image and (c) segmented result showing the two primary regions found. Each region is shown individually in (d) and (e).

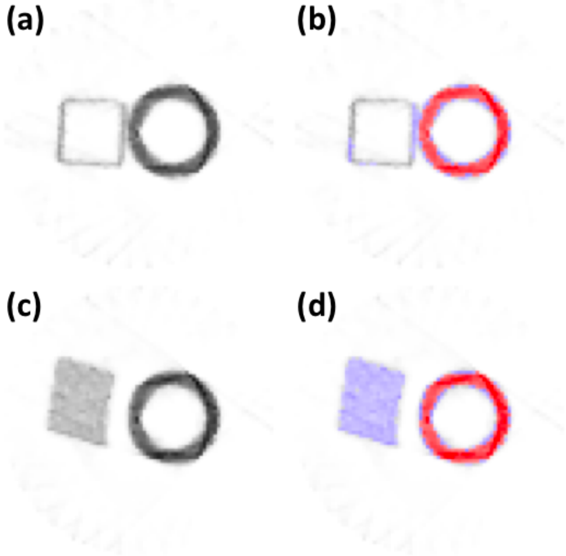

Figure 3. Tomographic slices from the transmission reconstruction overlaid with results from segmentation for the target consisting of a DU annulus, HDPE block, and hollow lead cube. (a) Tomographic slice of the transmission reconstruction for a slice through the DU annulus and the hollow lead cube and (b) overlay of the segmentation result on the transmission reconstruction for that slice. (c) Tomographic slice of the transmission reconstruction for a slice through the DU annulus and the HDPE block and (d) overlay of the segmentation result on the transmission reconstruction for that slice.

As seen from the figures, the segmentation roughly separates out the DU annulus from the HDPE block. The hollow lead cube was not found as a separate volume, but some parts of the wall were included in the

HDPE region. The transmission image that is used as input to the segmentation algorithm is somewhat noisy, as seen in **Error! Reference source not found.** 3(a) and (c), and this results in the segmentation algorithm attributing voxels that that should be part of the DU region to the HDPE region. This is especially the case near edges and boundaries of the materials, where the resolution of the imaging system produces an edge that is not sharp, leading to gradually decreasing attenuation values, some of which are consistent with the HDPE region and thus attributed there.

When the histogram of all transmission image values is plotted together with the histograms of values contained in each segmented region in Figure 4. Histograms of reconstructed transmission image values for the target consisting of a DU annulus, HDPE block, and hollow lead cube. All transmission image values are included in the black histogram. The red and blue histograms are the values included in the segmented regions depicted with the same colors in the figures above.

, it becomes clearer what this implementation of the Chan–Vese level set method is achieving. The segmentation is not significantly different from taking the transmission values and performing a course binning with a few large bins to determine regions. Because the transitional voxels at the boundaries have values that are similar to those within the other regions, those voxels are all grouped in a way that leaves poorly centralized regions. In principle, the inclusion of the constraint to minimize the length of the boundary in the energy function was intended to help alleviate some of the boundary issues, but this proved not to be the case.

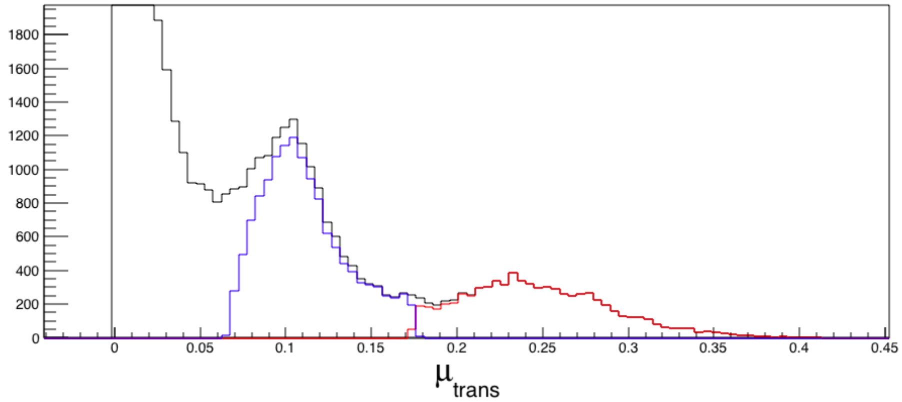

Figure 4. Histograms of reconstructed transmission image values for the target consisting of a DU annulus, HDPE block, and hollow lead cube. All transmission image values are included in the black histogram. The red and blue histograms are the values included in the segmented regions depicted with the same colors in the figures above.

4.1 CHALLENGES IN THE RECONSTRUCTION DATA

In addition to noise in the image, the input data had other problems that were challenging for the Chan–Vese implementation. In particular, the reconstructed images sometimes contained attenuation coefficients with two distinct groupings of values for the same material depending on the other materials present nearby. For instance, for one target that consisted of a DU annulus shielded by four different materials, the attenuation values for the half of the annulus shielded by steel and lead were distinct from the values for the half shielded by aluminum and HDPE. Figure 5. Histograms of transmission image values observed for two sides of a DU annulus shielded by different materials on each side. (a) Photograph of the target, showing the DU annulus in the center, shielded by steel (red), HDPE (white), aluminum (darker blue), and lead (lighter blue). (b) One tomographic slice showing two manually selected regions of the inner DU annulus (red and green) and one manually selected region denoting the outer steel shield (blue). This slice was taken from the upper half of the target, where it was shielded by HDPE and steel. (c) Histogram of transmission image values for all voxels (black) and regions outlined in (b) where the colors of each histogram correspond to the colors of the regions outlined in (b). The minimum of the x-axis has been selected to exclude the large peak at the bottom of the distribution containing small transmission image values.

shows a photograph of the target as well as a slice of the reconstruction that highlights three different regions. The regions outlined in red and green correspond to the same contiguous piece of DU, yet when the attenuation values are plotted for each half of the DU they peak at different values (see Figure 5. Histograms of transmission image values observed for two sides of a DU annulus shielded by different materials on each side. (a) Photograph of the target, showing the DU annulus in the center, shielded by steel (red), HDPE (white), aluminum (darker blue), and lead (lighter blue). (b) One tomographic slice showing two manually selected regions of the inner DU annulus (red and green) and one manually selected region denoting the outer steel shield (blue). This slice was taken from the upper half of the target, where it was shielded by HDPE and steel. (c) Histogram of transmission image values for all voxels (black) and regions outlined in (b) where the colors of each histogram correspond to the colors of the regions outlined in (b). The minimum of the x-axis has been selected to exclude the large peak at the bottom of the distribution containing small transmission image values.

[c]). The values from the red region are more consistent with the values observed from the steel region, which is outlined in blue in Figure 5. Histograms of transmission image values observed for two sides of a DU annulus shielded by different materials on each side. (a) Photograph of the target, showing the DU annulus in the center, shielded by steel (red), HDPE (white), aluminum (darker blue), and lead (lighter blue). (b) One tomographic slice showing two manually selected regions of the inner DU annulus (red and green) and one manually selected region denoting the outer steel shield (blue). This slice was taken from the upper half of the target, where it was shielded by HDPE and steel. (c) Histogram of transmission image values for all voxels (black) and regions outlined in (b) where the colors of each histogram correspond to the colors of the regions outlined in (b). The minimum of the x-axis has been selected to exclude the large peak at the bottom of the distribution containing small transmission image values.

(b).

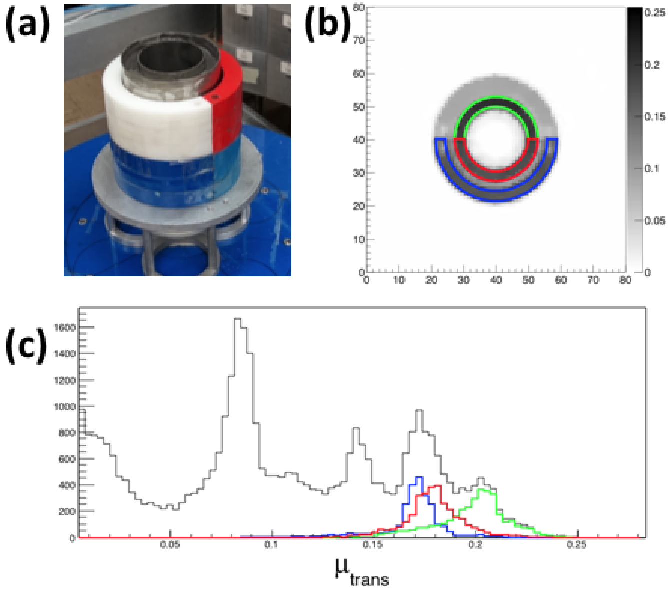

Figure 5. Histograms of transmission image values observed for two sides of a DU annulus shielded by different materials on each side. (a) Photograph of the target, showing the DU annulus in the center, shielded by steel (red), HDPE (white), aluminum (darker blue), and lead (lighter blue). (b) One tomographic slice showing two manually selected regions of the inner DU annulus (red and green) and one manually selected region denoting the outer steel shield (blue).

This slice was taken from the upper half of the target, where it was shielded by HDPE and steel. (c) Histogram of transmission image values for all voxels (black) and regions outlined in (b) where the colors of each histogram correspond to the colors of the regions outlined in (b). The minimum of the x-axis has been selected to exclude the large peak at the bottom of the distribution containing small transmission image values.

The difference in attenuation values can be attributed to differences in the accuracy with which the scatter correction method estimates the scattered component in the analysis step (see Section 2). The correction is more accurate for rays through the DU annulus that are shielded by lighter materials, such as aluminum and HDPE, than it is for parts shielded by steel and lead. As such, the inability of the segmentation algorithm to distinguish these materials is expected, and the fix involves modifications to the analysis steps rather than to the segmentation process.

In several cases, the ambiguity in transmission reconstruction values could be resolved with the additional information provided by induced-reaction reconstructions. For instance, DU and steel may be too close in value in the transmission image to distinguish with the currently implemented scatter correction, but only DU is capable of undergoing an induced fission reaction. The inclusion of induced-reaction images will be discussed more in a later section. First, a description of an attempt to account for these effects in the transmission images themselves will be discussed.

4.2 ATTEMPTS AT ACCOUNTING FOR PROBLEMS IN THE IMAGE DATA

The previous results have shown that the reconstructed data often suffer from problems related to inhomogeneities within single volumes and from noise in the reconstructed images. In response to this, the use of a method from Li [6] that modifies the Chan–Vese approach was investigated. The intent was to try to account for differences in the attenuation values within the same volume by treating the image as if it were subject to a slowly varying bias field.

The method implemented a multiplicative model of intensity inhomogeneity. An observed image can be modeled as

$$I = bJ + n,$$

where J is the true image, b is the component that accounts for the intensity inhomogeneity, and n is the additive noise. The true image J measures the physical properties of the object imaged, so it is assumed to be piecewise constant as in the Chan–Vese model.

We considered a circular neighborhood with a radius ρ centered defined by $O_y \triangleq \{\, x\colon |(y - x)| \leq \rho\}$, at each point $y \in \Omega$, where Ω is the entire domain. The bias field, b, is considered to be slowly varying so it can be approximated by a constant in neighborhood of each point in the domain. For a slowly varying bias field, $b(x)$ for all x in the circular neighborhood O_y are approximately equal to $b(y)$:

$$b(x) \approx b(y) \text{ for } x \in O_y.$$

The image model can be written as

$$I(x) \approx b(y)c_i + n(x) \text{ for } x \in O_y \cap \Omega_i,$$

where the assumption is that entire domain Ω is partitioned into $\{\, \Omega_i \}_{i=1}^{N}$. In multiphase level sets, N is the number of different intensities, c_i, in the image and $N = 2^n$, where n is the number of level sets used in multiphase level sets formulation. With a kernel function, $K(y-x)$, added to the formulation to work as a window function within the neighborhood, the local energy function ε_y is given by

$$\varepsilon_y = \sum_{i=1}^{N} \int_{\Omega_i} K(y-x)\, |I(x) - b(y)c_i|^2 \, dx,$$

where $K(y-x)$ is defined as,

$$K(y-x) = \begin{cases} e^{-|(y-x)|^2/2*\sigma^2} \\ 0 \end{cases}, \text{ for } |(y-x)| \le \rho.$$

The total energy term can then be defined as the integral of ε_y with respect to y over the image domain Ω:

$$\varepsilon \triangleq \int \varepsilon_y \, dy.$$

We can write the energy term and changing the orders of integration using multiphase level sets formulation,

$$\varepsilon(\phi, c, b) = \int \sum_{i=1}^{N} e_i(x) M_i(\phi(x)) dx\,,$$

where e_i and $M_i(\phi(x))$ are defined as

$$M_i(\emptyset_1(y), \ldots., \emptyset_k(y)) = \begin{cases} 1, y \in \Omega_i \\ 0, else \end{cases}$$

$$e_i(x) = \int K(y-x)|I(x) - b(y)c_i|^2 \, dy$$

The above energy term $\varepsilon(\phi, c, b)$ is used as "image term" in the energy formulation in multiphase level sets. The only difference with this method, Li's method, and the Chan–Vese's piecewise constant method is how the image term is defined. The length term from before ($\nu|C|$) is unchanged, so the total energy functional becomes

$$F(\phi, c, b) = \varepsilon(\phi, c, b) + \nu|C|\,.$$

By minimizing this equation, we obtain the result of image segmentation given by the level set functions and the estimation of bias field. The minimization problem can then be solved by using the gradient descent method.

$$\frac{\partial \emptyset}{\partial t} = -\frac{\partial F}{\partial \emptyset}$$

$$\frac{\partial \emptyset_1}{\partial t} = -\sum_{i=1}^{N} \frac{\partial M_i(\emptyset(x))}{\partial \emptyset_1} e_i + \mu\delta(\emptyset_1) div\left(\frac{\nabla \emptyset_1}{|\nabla \emptyset_1|}\right) \ldots$$

$$\frac{\partial \emptyset_k}{\partial t} = -\sum_{i=1}^{N} \frac{\partial M_i(\emptyset(x))}{\partial \emptyset_k} e_i + \mu\delta(\emptyset_k) div\left(\frac{\nabla \emptyset_k}{|\nabla \emptyset_k|}\right),$$

The e_is can be computed using following equivalent expression:

$$e_i(x) = I^2 - 2c_i I(b * K) + c_i^2 (b^2 * K),$$

where $*$ is the convolution operation.

For fixed $\emptyset$ and b, the optimal c that minimizes the energy $\varepsilon(\phi, c, b)$ is given by

$$\hat{c} = \frac{\int (b * K) I M_i(\emptyset(y))}{\int (b * K) M_i(\emptyset(y))},$$

and for fixed $\emptyset$ and c, the optimal b that minimizes the energy $\varepsilon(\phi, c, b)$ is given by

$$\hat{b} = \frac{(IJ^1) * K}{J^2 * K},$$

where $J^1 = \sum_{i=1}^{N} c_i M_i$, and $J^2 = \sum_{i=1}^{N} c_i^2 M_i$.

Level set parameters of $\mu = 0.0001 \times 255^2$ and $\Delta t = 0.1$ were used and the convolution Gaussian kernel, *K*, is constructed as $w \times w$ mask with $w = 4 * \sigma + 1$ with $\sigma = 2$. An example of the Li method applied to an image to successfully remove bias is shown for an image taken from reference [6] in Figure 6. Example image demonstrating the intended effect of implementing Li's method for accounting for bias in an image during segmentation. (a) The original image, (b) bias in the original image, (c) bias corrected image, (d) histogram of original image values, and (e) histogram of bias corrected image values.

.

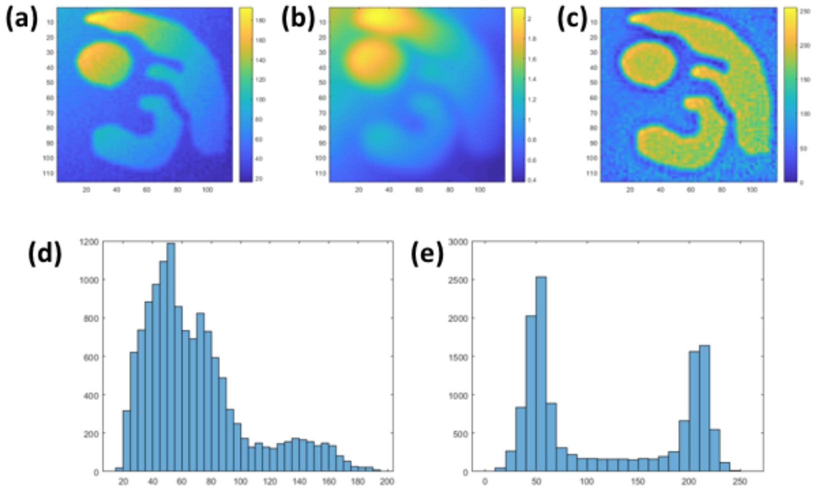

Figure 6. Example image demonstrating the intended effect of implementing Li's method for accounting for bias in an image during segmentation. (a) The original image, (b) bias in the original image, (c) bias corrected image, (d) histogram of original image values, and (e) histogram of bias corrected image values.

An example of Li's method applied to our data is shown in Figure 7. Example of applying Li's method to measured data to account for bias in an image during segmentation. (a) The original transmission image for an object with an inner annulus of DU, shielded by steel and HDPE, (b) the calculated bias in the original image, (c) the bias corrected image, (d) the histogram of original image values, and (e) the histogram of bias corrected image values.

for the target with the DU annulus and composite shield shown earlier. For this case, the DU region becomes more homogeneous (Figure 7. Example of applying Li’s method to measured data to account for bias in an image during segmentation. (a) The original transmission image for an object with an inner annulus of DU, shielded by steel and HDPE, (b) the calculated bias in the original image, (c) the bias corrected image, (d) the histogram of original image values, and (e) the histogram of bias corrected image values.

[c]), but still has similar image values as the steel. This is likely the best that can be achieved at this point because the inhomogeneity is caused by an underlying problem in the analysis. The final histogram shows that there is still essentially a single peak of image values that come from the DU and steel regions. It is possible that this or a similar method could be modified or developed that would better handle some of these homogeneities and image noise, but it requires further study. Additionally, the application of such a method would have be done so that the underlying physical information carried by the attenuation lengths was not lost in the process.

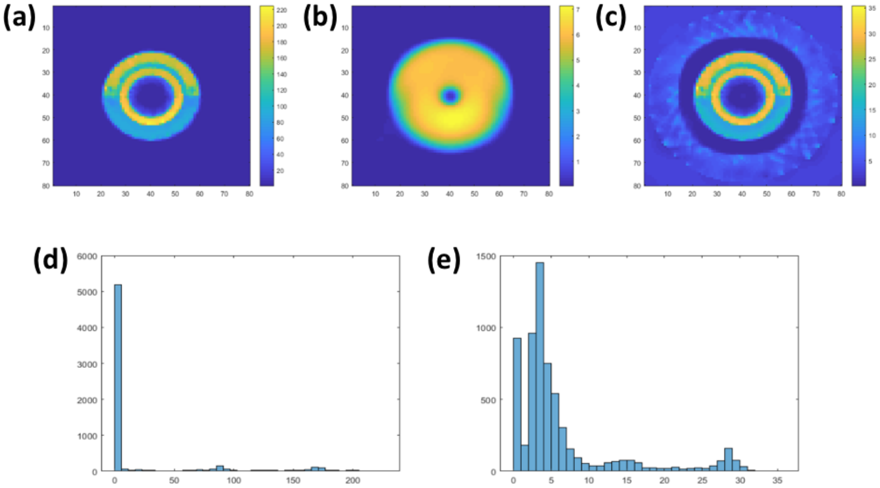

Figure 7. Example of applying Li’s method to measured data to account for bias in an image during segmentation. (a) The original transmission image for an object with an inner annulus of DU, shielded by steel and HDPE, (b) the calculated bias in the original image, (c) the bias corrected image, (d) the histogram of original image values, and (e) the histogram of bias corrected image values.

5. INCLUSION OF INDUCED-REACTION RECONSTRUCTIONS

The image reconstructions from induced reactions provides an additional means of discriminating between different materials that reconstruct to the same or similar attenuation coefficients in the transmission images. The induced-fission reconstructions help to distinguish DU from other heavy materials present in the target, and the hydrogen-scatter reconstructions distinguish hydrogenous materials such as HDPE from other lighter materials, such as aluminum, in the reconstruction. A description of how each induced reconstruction is incorporated into the segmentation process will be described in the following subsections.

5.1 INDUCED-FISSION IMAGING DATA

The induced-fission imaging reconstructions provide a map of where material exists within a target from which multiple neutrons get created and detected at times consistent with fission. There are several factors that make the induced-fission reconstructions of poorer resolution than the transmission reconstructions. First, the induced-fission reconstructions rely on lines of response defined by the initial neutron direction, which is not known to high precision. Second, induced-fission imaging is less precise because neutrons initially directed toward non-fissionable materials can scatter neutrons into nearby regions of the target that do contain fissionable materials and cause a fission to occur. Since the induced-fission analysis correlates such events to the initial neutron direction, it will appear as if the nearby materials are also producing multiple neutrons. Finally, there are other interactions that produce multiple neutrons, such as the (n, 2n) reaction. These interactions become particularly relevant for materials such as lead, which has a non-negligible (n, 2n) cross section at 14 MeV.

In Figure 5. Histograms of transmission image values observed for two sides of a DU annulus shielded by different materials on each side. (a) Photograph of the target, showing the DU annulus in the center, shielded by steel (red), HDPE (white), aluminum (darker blue), and lead (lighter blue). (b) One tomographic slice showing two manually selected regions of the inner DU annulus (red and green) and one manually selected region denoting the outer steel shield (blue). This slice was taken from the upper half of the target, where it was shielded by HDPE and steel. (c) Histogram of transmission image values for all voxels (black) and regions outlined in (b) where the colors of each histogram correspond to the colors of the regions outlined in (b). The minimum of the x-axis has been selected to exclude the large peak at the bottom of the distribution containing small transmission image values.

, the histogram of transmission values for half of a DU annulus were shown to be nearly identical to the values from the steel that surrounded it. In principle, this is a case where the induced-fission reconstruction can help break the degeneracy for those two regions. Figure 8 shows plots of the induced-fission image values vs. transmission image values for manually defined regions for that target, which are highlighted in different colors in Figure 8(a). In Figure 8(b), the induced-fission and transmission image values for voxels from all the highlighted regions are plotted. When values for the steel region (Figure 8[c]), the DU region shielded by HDPE and aluminum (Figure 8[d]), and the DU region shielded by lead and steel (Figure 8[e]) are displayed individually, it becomes clear that data from each material occupies slightly different regions of the plot. However, the regions still overlap significantly, making it difficult to create an algorithm that will properly separate the steel from the DU, as will be seen in Section 6.3.

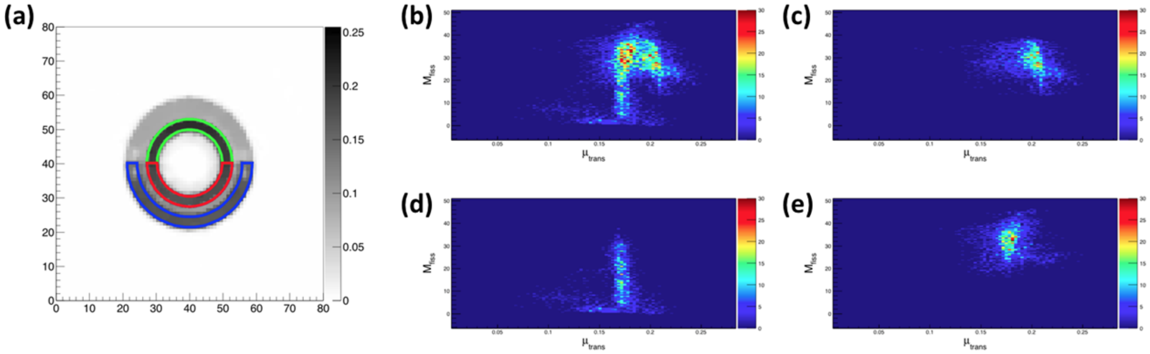

Figure 8. Induced-fission image values vs. transmission image values for selected regions in the target consisting of a DU annulus and composite shield. (a) One tomographic slice of the transmission reconstruction showing manually selected regions of the inner DU annulus (red and green) and the outer steel shield (blue). (b) Induced-fission vs. transmission values for all steel and DU regions of the target (red, blue, and green regions). The remaining panels show these results separated out by individual region: (c) the DU region shielded by HDPE and aluminum (green region), (d) the steel region (blue region), and (e) the DU region shielded by steel and lead (red region).

A method was developed to use the induced-fission reconstructions after the 3D transmission segmentation is performed to separate the fissionable DU regions from non-fissionable regions with similar transmission reconstruction values. The method consists of the following steps:

1. Perform the 3D segmentation of the transmission imaging data using Chan–Vese's multiphase level sets method as described above. The result of this is 2^n regions, Ω_i, where n is the number of level sets used, and i is the index for a given region. For all segmentations of measured data shown in this report, $n = 3$. An example of what these regions look like can be found in Figure 2. Results from the segmentation of the transmission image for a target consisting of a DU annulus, HDPE block, and hollow lead cube. (a) Photograph of the target object. The annulus is on the right, the HDPE block is white and sitting on the hollow lead cube, which is painted gray. (b) The reconstructed transmission image and (c) segmented result showing the two primary regions found. Each region is shown individually in (d) and (e).

2. . Call a 2D slice of a region $\Omega_{i,z}$.

3. Prepare the 3D induced-fission reconstruction by thresholding it to discard all values less than 10% of the maximum and then erode with a sphere with a radius of two voxels. Call the resulting induced-fission reconstruction data M_f, and call a 2D slice of this result $M_{f,z}$.

4. For each region, $\Omega_{i,z}$, determine the fraction of voxels in the region that contain non-zero induced-fission values using the thresholded and eroded induced-fission result, $M_{f,z}$. If the fraction is larger than 1/3, provisionally label $\Omega_{i,z}$ as a fissionable region. The value of 1/3 was determined through experimentation. This step is performed independently for each 2D slice.

5. For a region $\Omega_{i,z}$ that has been assumed to be fissionable, discard any voxels that have transmission values larger than 2 standard deviations (2σ) from the mean transmission values of that region, where the mean and standard deviation is calculated in $\Omega_{i,z}$ for that 2D slice. Call the 3D collection of voxels that remain from all 2D slices Ω'_i.

6. With the voxels in a 3D region Ω'_i that has been provisionally labeled as fissionable, a ratio is calculated from the means of the induced-fission and transmission image values:

$$R_i^F = \frac{\overline{M}_{fiss}(\Omega'_i)}{\overline{\mu}_{trans}(\Omega'_i)}$$

To continue to be labeled as a fissionable region, R_i^F must be > 2, and $\overline{\mu}_{trans}(\Omega'_i)$ must be > 0.16 for Ω'_i. These thresholds were determined through experimentation.

7. For a region Ω'_i that is still labeled as fissionable, a final check is performed. A ratio of the average transmission values for the following combinations of intersections with the collection of 2D labeled slices and full 3D segmentation results is calculated:

$$R_i^T = \frac{\overline{\mu}_{trans}(\Omega'_i \text{ AND } \Omega_i)}{\overline{\mu}_{trans}(\sum_{j \neq i}(\Omega'_j))}$$

In this expression "$\sum_{j,fiss}(\Omega'_j)$ AND NOT Ω_i" indicates that any voxel from any region in Ω' that has been provisionally labeled as fissionable, excluding those that contained the 3D segmentation of the current region (Ω_i) being tested, are used in the average. Through experimentation, it was found that regions where $R_i^T > 1.2$ should not be labeled as fissionable.

The ratio R_i^F in step 5 is an empirical attempt at imposing a physically meaningful threshold on the fraction of the total cross section that comes from fission to lower the chance that non-fissionable material that generates a reduced amount of doubles is labeled as fissionable. An example of how such a ratio would be able to distinguish DU from steel can be seen by examining the transmission and induced-fission reconstruction values plotted in Figure 8. At present the threshold is simply an empirically determined value because the overall normalization of the induced-fission reconstruction values, which is done in a combination of the analysis and reconstruction steps, is not yet accurate. If this is corrected in future work, it is possible that theoretical expectations for the ratio of cross sections can be used to discriminate non-fissionable materials. For now, the value was experimentally selected to optimally distinguish the DU from the specific collection of materials used for measurements in the project. This was found to be successful for distinguishing DU from steel, provided that they were separated from each other by a gap or different material. If they were contiguous, this method was not effective.

The ratio R_i^T in step 6 is a final check that was found to be necessary to ensure that lead was not labeled as fissionable. As mentioned previously, lead creates doubles from reactions such as (n,2n) that are difficult to distinguish from fission. However, the application of the R_i^T threshold ensured that the transmission coefficients were consistent with those expected for fissionable materials.

See Section 6.6 for segmentation results demonstrating the use of the induced-fission reconstruction to distinguish DU from other heavy materials with similar attenuation coefficients. Several of these labeling criteria have an *ad hoc* nature at this point. More general purpose and physically meaningful procedures can be explored in the future when the reconstruction data are more quantitatively accurate.

5.2 HYDROGEN-SCATTER IMAGING DATA

For the data examined in the project, the need for the induced-fission reconstructions just discussed was to help distinguish between materials that had similar transmission reconstruction values due to inaccuracies in the analysis and reconstruction methods. There are other cases where, ignoring analysis inaccuracies, the transmission values should be the same for different materials, as is the case in instances where HDPE and aluminum are present in the same target. Each material has a nearly identical value of attenuation coefficient for 14 MeV neutrons and therefore *should* be segmented as the same material using transmission alone. This is a case where the ambiguity would be impossible to resolve without additional information. For the case of HDPE and aluminum, the fact that HDPE is hydrogenous means the use of the hydrogen elastic scatter reconstruction allows the two materials to be distinguished.

Figure 9 demonstrates how helpful the hydrogen scatter reconstruction can be for distinguishing HDPE from aluminum. The location of the regions containing HDPE and aluminum within the target is shown in the tomographic slice in Figure 9(a). Vertically, the HDPE is in the upper half of the target and the aluminum is in the lower half (a photograph of the target can be seen in Figure 5. Histograms of transmission image values observed for two sides of a DU annulus shielded by different materials on each side. (a) Photograph of the target, showing the DU annulus in the center, shielded by steel (red), HDPE (white), aluminum (darker blue), and lead (lighter blue). (b) One tomographic slice showing two manually selected regions of the inner DU annulus (red and green) and one manually selected region denoting the outer steel shield (blue). This slice was taken from the upper half of the target, where it was shielded by HDPE and steel. (c) Histogram of transmission image values for all voxels (black) and regions outlined in (b) where the colors of each histogram correspond to the colors of the regions outlined in (b). The minimum of the x-axis has been selected to exclude the large peak at the bottom of the distribution containing small transmission image values.

[a]). In Figure 9(b), the attenuation coefficients for aluminum and HDPE are shown in green and blue, respectively, and are seen to overlap significantly. However, if the hydrogen scatter image values are plotted versus the transmission image values, the two materials can be easily distinguished. Values from voxels containing both materials are shown in Figure 9(c), and they are shown individually for aluminum and HDPE in Figure 9(d) and (e), respectively.

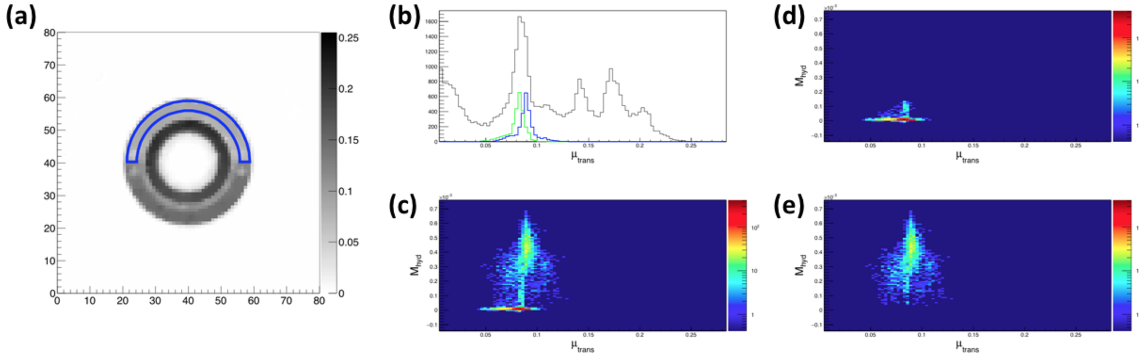

Figure 9. Hydrogen elastic scatter vs. transmission image values for selected regions in the target consisting of a DU annulus and composite shield. (a) A single tomographic slice from the transmission reconstruction with the manually defined region isolating the known location of HDPE and aluminum in the target outlined in blue. (b) Histogram of transmission voxel values for all voxels (black), the HDPE region (blue), and the aluminum region (green). The minimum of the x-axis was selected to exclude the large peak at the bottom of the distribution containing small transmission values. (c) Hydrogen scatter vs. transmission image values for both HDPE and aluminum regions of the target. The remaining two panels show hydrogen scatter vs. transmission values for the (d) aluminum and (e) HDPE regions individually.

The process for using the hydrogen scatter image data is nearly identical to the process used for induced-fission, with a few exceptions. The process consists of the following steps:

1. Perform the 3D segmentation of the transmission imaging data using Chan–Vese's multiphase level sets method as described above, resulting in 2^n 3D regions, Ω_i.

2. Prepare the 3D hydrogen-scatter reconstruction data by thresholding it to discard all values less than 10% of the maximum. Call the resulting hydrogen-scatter reconstruction data M_h, and call a 2D slice of this result $M_{h,z}$.

3. For each region, $\Omega_{i,z}$, determine the fraction of voxels in the region that contain non-zero hydrogen-scatter values using the thresholded hydrogen-scatter result, $M_{h,z}$, and provisionally label as hydrogenous if the fraction is larger than 1/3. This step is performed independently for each 2D slice.

4. For a region $\Omega_{i,z}$ that has been assumed to be hydrogenous, discard any voxels that have transmission values that are larger than 2 standard deviations (2σ) from the mean transmission values of that region, where the mean and standard deviation is calculated in $\Omega_{i,z}$ for that 2D slice. Call the 3D collection of voxels that remain from all 2D slices Ω'_i.

5. With the voxels from a 3D region Ω'_i that has been provisionally labeled as hydrogenous, a ratio is calculated from the means of the hydrogen-scatter and transmission values:

$$R_i^H = \frac{\overline{M}_{hyd}(\Omega'_i)}{\bar{\mu}_{trans}(\Omega'_i)}$$

   For Ω'_i to continue to be labeled as a hydrogenous region, two conditions are applied: $R_i^H > 0.006$ and $\bar{\mu}_{trans}(\Omega'_i) > 0.11$ must be true for Ω'_i to continue to be labeled as a hydrogenous region. These thresholds were determined through experimentation.

6. If desired, the user can view the results and perform a final sanity check for regions, Ω'_i, where HDPE and aluminum appear to be adjacent. Then the user can take the thresholded hydrogen scatter reconstruction, M_h, and erode it by a sphere with a radius of two pixels to create M_h^e. Label the voxels that include the intersection of both the 3D region, Ω'_i, and M_h^e as hydrogenous, and what is left of Ω'_i as a non-hydrogenous region.

See Section 6.5 for an example of a segmentation result that successfully distinguishes HDPE and aluminum using the hydrogen-scatter reconstruction as input.

As with the inclusion of the induced-fission imaged data, the inclusion of the hydrogen-scatter imaged data uses selection criteria that are somewhat *ad hoc* at this point. This situation will improve when the normalization of the hydrogen-scatter reconstruction values is confirmed to be correct, meaning that the values can be physically interpreted according to their relationship to the macroscopic cross section.

6. SEGMENTATION RESULTS

This section will review the results showing the application of the segmentation procedure described above to six target objects that contain several different materials in varying geometries. The only fissionable material available for these measurements was a DU annulus, and the hydrogenous materials used were HDPE, ordinary tap water, and the plastic scintillator within one of the neutron imaging detectors. The non-hydrogenous and non-fissionable materials used were lead, steel, and aluminum. The geometries vary from targets that are more easily segmented because the materials are well separated to targets that are more difficult to segment because there are multiple layers of materials placed directly adjacent to each other.

It should be emphasized that the complete segmentation algorithm that includes induced reactions is only able to identify fissionable regions, hydrogenous regions, and “other” regions. “Other” here simply denotes non-fissionable and non-hydrogenous, though there may be multiple “other” regions that distinguish different materials with different attenuation coefficients. This should not be confused with material identification because these methods are not assigning materials to each region. Rather, the segmentation is only attempting to isolate the *regions* containing distinct materials, even if the exact material type is unknown. However, regions will sometimes be referred to by their material names in the discussion below for reasons of clarity.

6.1 DEPLETED URANIUM ANNULUS, HIGH-DENSITY POLYETHYLENE BLOCK, AND HOLLOW LEAD CUBE

This target was described and the segmentation of the transmission reconstruction without the inclusion of induced-reaction reconstruction data was shown in Figure 2. Results from the segmentation of the transmission image for a target consisting of a DU annulus, HDPE block, and hollow lead cube. (a) Photograph of the target object. The annulus is on the right, the HDPE block is white and sitting on the hollow lead cube, which is painted gray. (b) The reconstructed transmission image and (c) segmented result showing the two primary regions found. Each region is shown individually in (d) and (e).

. The results after incorporating the induced-reaction reconstructions as described in Section 5 are shown in Figure 10, where the regions are much more distinctly defined than the original segmentation of the transmission data alone. The fissionable material is rendered in red, and the hydrogenous material is rendered in blue. The full procedure accurately finds the HDPE block to be hydrogenous, and the DU annulus to be fissionable. This target was a fairly easy case to segment because there was such large separation between the two layers of interest.

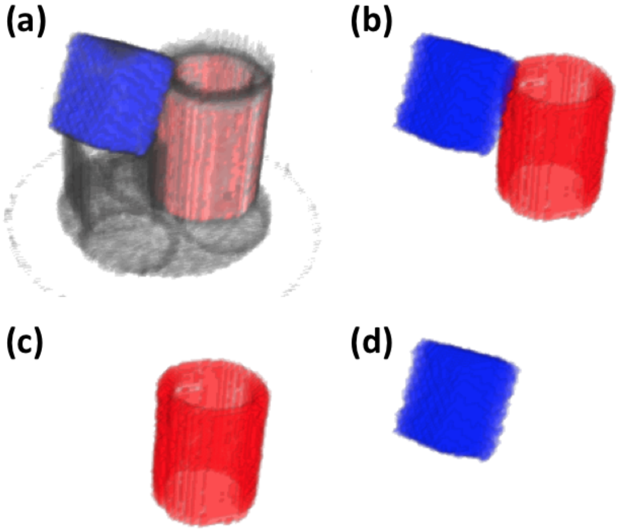

Figure 10. Results of image segmentation after inclusion of induced-reaction images for the target consisting of a DU annulus, HDPE block, and hollow lead cube. (a) The transmission image with final segmented result overlaid in blue for a hydrogenous region and red for a fissionable region, (b) the final segmented regions without the transmission reconstruction visible, and the (c) fissionable and (d) hydrogenous regions shown individually.

The hollow lead cube that held up the HDPE block did not segment to a separate region, likely due to the fact that the thickness of the walls was smaller than the reconstruction voxel size, resulting in a large variation in the attenuation values for that volume. This combined with noise in the reconstruction made it a difficult region to extract.

A comparison of results from the standard Chan–Vese level set algorithm applied only to transmission and the results after the additional processing with information from induced-reaction images are shown for tomographic slices at two heights of the target in Figure 11. From the segmentation of the transmission reconstruction in Figure 11(b) and (e) to the inclusion of the information from the induced-reaction in Figure 11(c) and (f), the regions have cleaned up considerably. Many of the edge pixels of the DU annulus, which were originally attributed to the HDPE region, no longer are, and the segmented regions do better at identifying the contiguous volumes corresponding to the actual objects in the target.

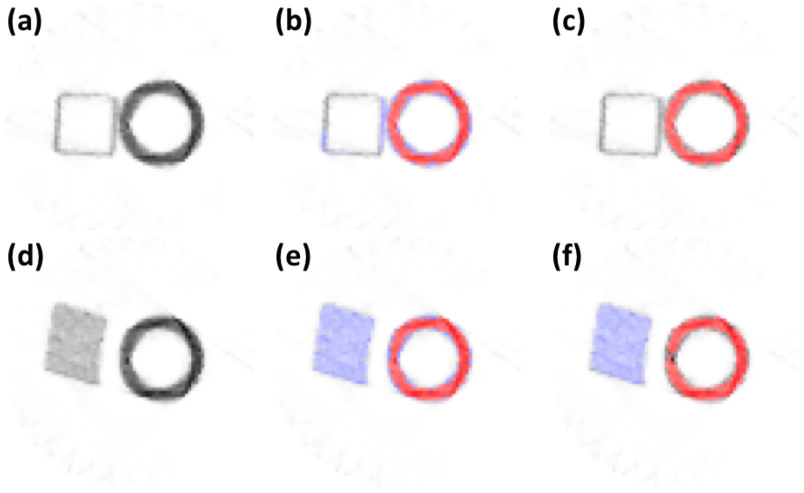

Figure 11. Tomographic slices through the transmission reconstruction with segmented results overlaid for the target consisting of a DU annulus, HDPE block, and hollow lead cube. For a slice containing the hollow lead cube and the DU annulus, the (a) transmission reconstruction, (b) overlay of the transmission-only segmentation, and (c) overlay of the segmentation that includes induced-reaction reconstructions. For a slice containing the hollow HDPE block and the DU annulus, the (d) transmission reconstruction, (e) overlay of the transmission-only segmentation, and (f) overlay of the segmentation that includes induced-reaction reconstructions. For the segmentations that include induced-reaction reconstructions, the red region corresponds to the fissionable region and the blue region corresponds to the hydrogenous region.

The histograms of attenuation values associated with each segmented region are shown in Figure 12. The region corresponding to the DU annulus did not change considerably after inclusion of the induced reactions, but the HDPE region contains considerably fewer voxels than before. The attenuation values associated with the final HPDE region appear to roughly correspond to the values expected if the peak observed in distribution of all values (the black histogram in Figure 12) was considered to be sitting on a continuum that was subtracted away.

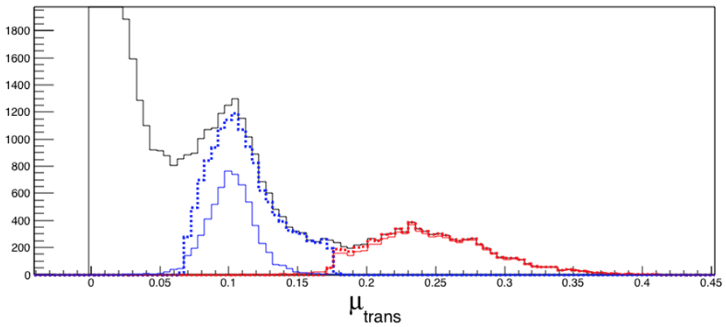

Figure 12. Histograms of transmission image values for all voxels and for the segmented regions for the target consisting of a DU annulus, HDPE block, and hollow lead cube. The histogram of transmission image values is shown for all voxels (black), the HDPE region after segmenting with transmission only (blue dotted) and after inclusion of induced-reaction reconstructions (blue), the DU region after segmenting with transmission only (red dotted) and after inclusion of induced-reaction reconstructions (red). The y-axis was chosen so that the large number of small reconstruction values are off scale to better see the distribution of values for the regions being segmented.

6.2 DEPLETED URANIUM ANNULUS, BLOCK DETECTOR, WATER BOTTLE IN STEEL PIPE

Another target with relatively well-separated materials is the one shown in Figure 13(a), which consists of a DU annulus, one of the neutron block detectors of the type used in this imaging system, and a steel tee pipe, within which sits a plastic water bottle partially filled with tap water. The detector is front-face down, meaning the approximately 10 × 10 × 5 cm^3 block of plastic scintillator is at the bottom of the detector housing. Although the volumes are relatively well separated, there are many small features in the transmission image (see Figure 13[b]) that are difficult to segment because they are so small. Like the previous target, the Chan–Vese level set method applied to the transmission data had difficulty successfully segmenting the target. Although the general features of the DU annulus and the plastic scintillator portion of the neutron block detector and water in the water bottle are segmented from each other, there are other regions of the target that are also included. Both regions contain large portions of the steel tee pipe, and the blue region (ostensibly associated with the plastic of the detector and water in the water bottle) also includes the edges of the DU annulus and some of the non-plastic interior parts of the block detector.

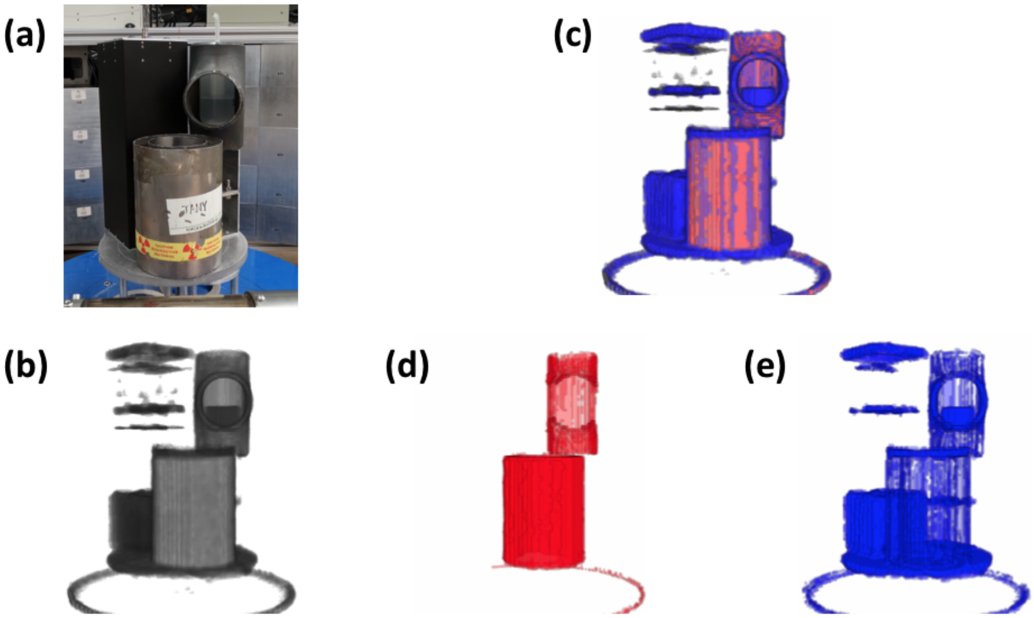

Figure 13. Results of image segmentation using only transmission imaging data for the target consisting of a DU annulus, a neutron block detector, and a water bottle in a steel tee pipe. (a) Photograph of the target showing the DU annulus in the front, a neutron block detector in the back left, and a partially filled water bottle within a steel tee pipe held up by an aluminum shelf in the back right. (b) Transmission reconstruction, (c) segmentation regions overlaid on transmission reconstruction, and each region shown separately in (d) and (e).

After inclusion of induced-reaction reconstruction data, the DU annulus is cleanly separated and identified as fissionable, but the region containing the plastic scintillator and water in the water bottle, while being correctly identified as hydrogenous, also have small portions of the steel pipe and regions near the bottom of the DU annulus included. The inclusion of some of the steel near the water is perhaps not surprising, given the fact that the hydrogen scatter imaged data has poorer resolution than transmission resulting in some of the scatter from hydrogen being attributed to nearby materials. It is unclear why some of the lower portion of the DU annulus is labeled as hydrogenous. The image values in that region of the reconstruction appear lower than what is observed for the rest of the DU annulus, so there might be a problem in the image data themselves.

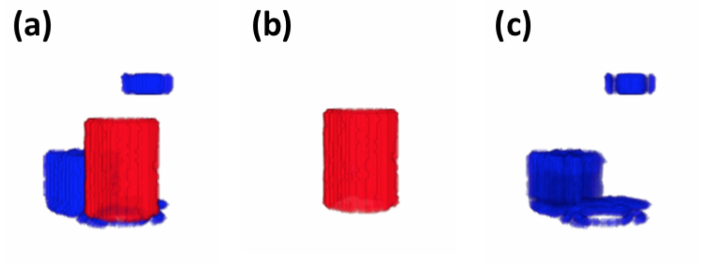

Figure 14. Results of image segmentation using both transmission and induced-reaction reconstruction data for the target consisting of a DU annulus, a neutron block detector, and a water bottle in a steel tee pipe. Segmentation regions shown (a) together and shown separately in (b) and (c). The red region corresponds to the fissionable region, and the blue region corresponds to the hydrogenous region.

Views of tomographic slices through the object at two heights are shown in Figure 15 for segmentations with and without induced-reaction data, where the degree of clean up, especially near the boundaries and in the steel pipe can be seen. Figure 16 shows the histograms of voxel attenuation values for all voxels and for the segmented regions before and after the inclusion of the induced-reaction reconstructions. Note that as before, the segmentation before induced-reaction reconstructions are used is similar to what would be expected if the reconstruction was simply histogrammed with course binning, but the full segmentation preserves more of the contiguousness of the volumes.

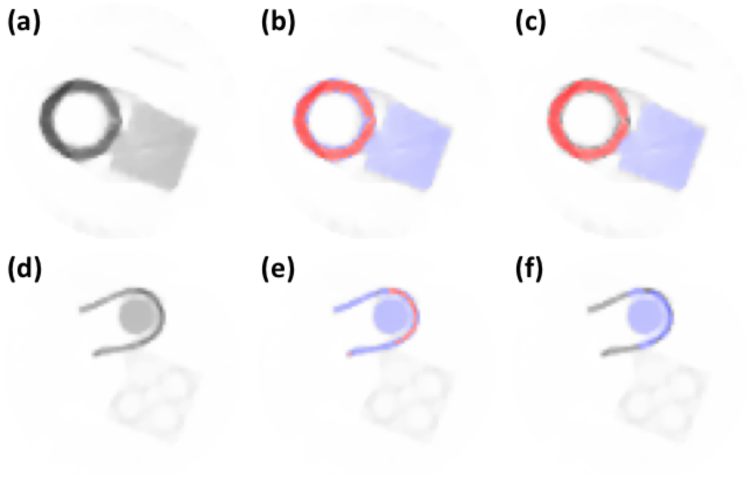

Figure 15. Tomographic slices through the transmission reconstruction with segmented results overlaid for the target consisting of a DU annulus, a neutron block detector, and a water bottle in a steel tee pipe. For a slice containing the DU annulus and the plastic part of the block detector, the (a) transmission reconstruction, (b) overlay of the transmission-only segmentation, and (c) overlay of the segmentation that includes induced-reaction reconstructions. For a slice containing the water-filled part of the water bottle in the steel tee pipe and the internal components of the block detector, the (d) transmission reconstruction, (e) overlay of the transmission-only segmentation, and (f) overlay of the segmentation that includes induced-reaction reconstructions. For the segmentations that include induced-reaction reconstructions, the red region corresponds to the fissionable region, and the blue region corresponds to the hydrogenous region.

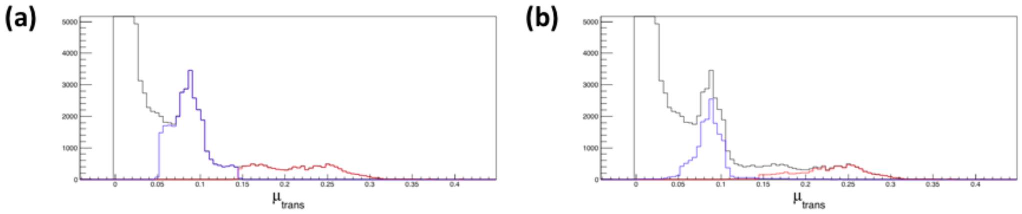

Figure 16. Histograms of transmission image values for all voxels and for the segmented regions for the target consisting of a DU annulus, a neutron block detector, and a water bottle in a steel tee pipe. The histogram of transmission reconstruction values is shown for all voxels (black) in both panels. (a) Histograms for each region after transmission-only segmentation and (b) with induced-reaction reconstructions included. The hydrogenous region is blue and the fissionable region is red. The y-axis was chosen so that the large number of small reconstruction values are off scale to better see the distribution of values for the regions being segmented.

6.3 DEPLETED URANIUM ANNULUS INSIDE COMPOSITE SHIELD

This target was introduced earlier when discussing problems in the reconstruction data that make segmentation difficult (see Figure 5). It consists of a DU annulus, surrounded by a composite shield. As the drawing in Figure 17 shows, the lower half of the shield is aluminum in the front and lead in the back, and the upper half of the shield is HDPE in the front and steel in the back. The challenge of this target is that the materials are closely packed, which means edge voxels will overlap, causing difficulty for the segmentation routine.

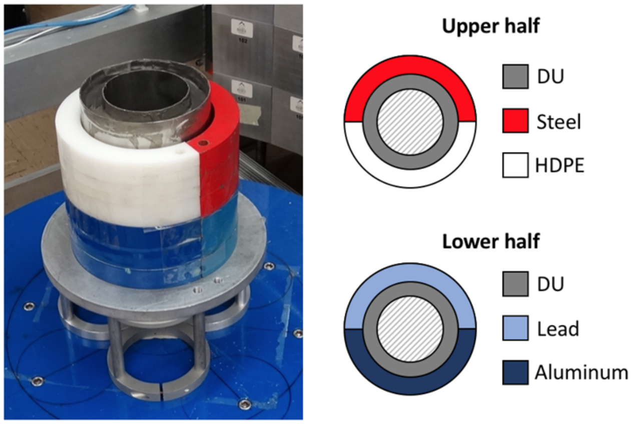

Figure 17. Photograph and material labeling for the target consisting of a DU annulus shielded by steel, HDPE, lead, and aluminum. The upper half of the target is shielded by steel and HDPE, and the lower half is shielded by lead and aluminum.

The segmentation of the transmission data without induced-reaction images results in three primary regions is shown in Figure 18. The first contains most of the DU annulus and the steel portion of the shield (Figure 18[d]), but the other two are somewhat mixed together. Nevertheless, a greater number of voxels for one of the regions (Figure 18[e]) corresponds primarily with the side of the shield containing HDPE and aluminum, and the other (Figure 18[f]) corresponds primarily with the portion containing lead. Surprisingly, because of what has already been seen regarding the difference in attenuation values for the two sides of the DU annulus, the DU annulus is segmented as the same volume, albeit with steel added. Because of the distribution of values observed in Figure 5, the steel might be expected to be grouped with only the adjacent DU with the other half of the DU identified as a distinct region, but this is not the case.

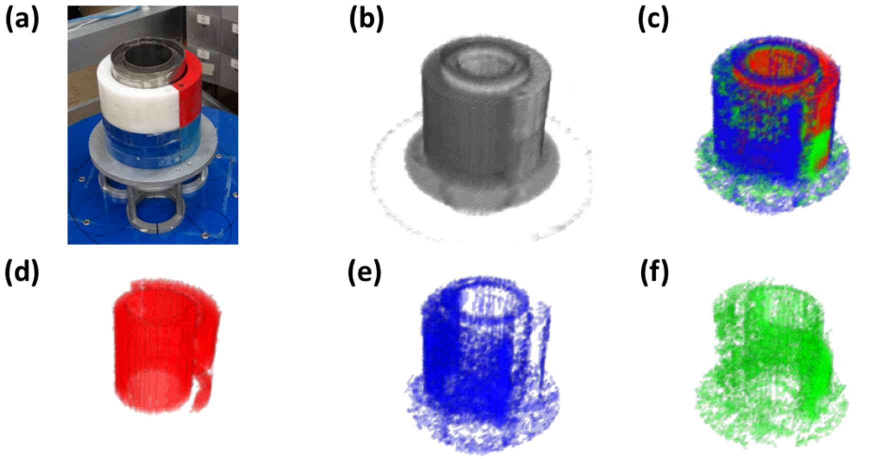

Figure 18. Results of image segmentation using only the transmission image for the target consisting of a DU annulus shielded by steel, HDPE, lead, and aluminum. (a) Photograph of the target, (b) transmission reconstruction, and (c) segmentation result. Each region is shown separately in (d), (e), and (f).

The segmentation results including the induced-fission data (Figure 19) confirm that the red region is fissionable, but the steel remains a part of the fissionable region. This is because the two volumes are adjacent to each other and were grouped in the initial transmission segmentation. The poor resolution of the induced-fission reconstruction makes it difficult to distinguish the steel as non-fissionable. This is most easily seen in Figure 21(c), where the induced fission reconstruction data is overlaid on the transmission reconstruction for a single tomographic slice in the part of the target containing steel and HDPE, showing the induced reconstruction overlapping the shielding layers outside the DU annulus. The same behavior was observed in the slice containing the lead and aluminum parts of the shield in Figure 20(c).

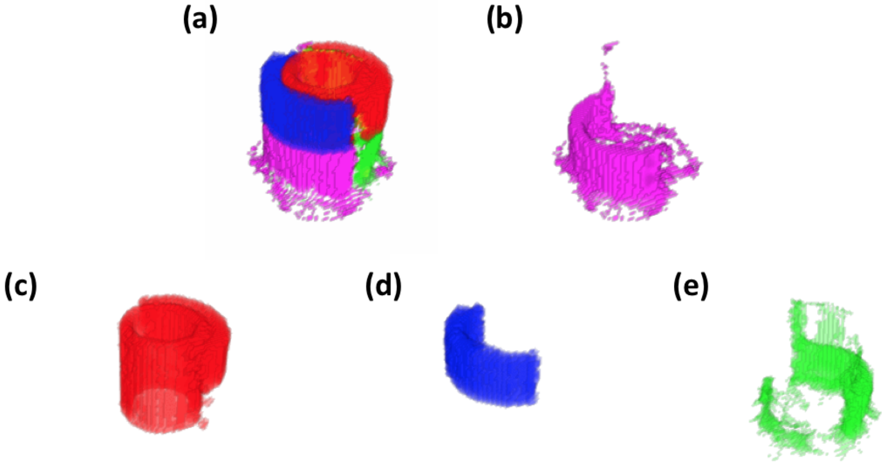

Figure 19. Results of image segmentation using both transmission and induced-reaction reconstruction data for the target consisting of a DU annulus shielded by steel, HDPE, lead, and aluminum. Segmentation regions are shown (a) together and shown separately in (b), (c), (d), and (e). The red region corresponds to the fissionable region and the blue region corresponds to the hydrogenous region.

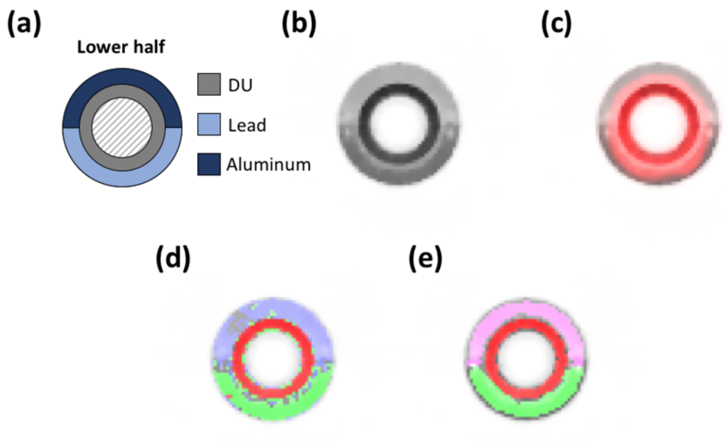

Figure 20. Tomographic slice through the transmission reconstruction with induced-fission reconstruction and segmented results overlaid for the part of the target consisting of a DU annulus shielded by lead and aluminum. For a slice containing the DU annulus shielded by lead and aluminum, (a) the drawing labeling material location, (b) the transmission reconstruction, (c) the induced-fission reconstruction overlaid on the transmission reconstruction, (d) overlay of the transmission-only segmentation on the transmission reconstruction, and (e) overlay of the segmentation that includes induced-reaction reconstructions on the transmission reconstruction.

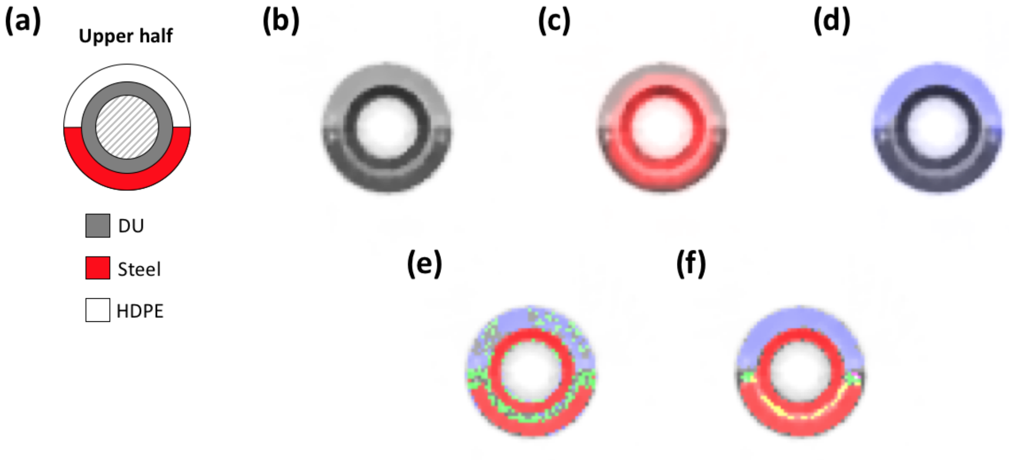

Figure 21. Tomographic slice through the transmission reconstruction with induced-reaction reconstructions and segmented results overlaid for the part of the target consisting of a DU annulus shielded by HDPE and steel. For a slice containing the DU annulus shielded by HDPE and steel, (a) the drawing labeling material location, (b) the transmission reconstruction, (c) the induced-fission reconstruction overlaid on the transmission reconstruction, (d) the hydrogen-scatter reconstruction overlaid on the transmission reconstruction, (e) overlay of the transmission-only segmentation on the transmission reconstruction, and (f) overlay of the segmentation that includes induced-reaction reconstructions on the transmission reconstruction. In (f), the red region corresponds to the what was labeled as the fissionable region and the blue region corresponds to the hydrogenous region.

The HDPE and aluminum parts of the shield are indistinguishable in transmission in theory and are shown to be so in practice in the transmission-only segmentation. However, results from this target were given earlier as an exemplar case for the use of hydrogen-scatter reconstructions to discriminate between the two materials. For this target, the HDPE was identified as hydrogenous (Figure 19[d]), and the remainder of the material (aluminum) on that side of the shield was found to be non-hydrogenous (Figure 19[e]). A tomographic slice through the part of the target shielded by HDPE and steel is shown with the hydrogen scatter reconstruction overlaid in Figure 21(d).

The segmentation results with and without induced-reaction reconstruction data are shown for tomographic slices in Figure 20 for the half of the DU annulus shielded by aluminum and lead and in Figure 21 for the half of the DU annulus shielded by steel and HDPE. The voxel reconstruction values for each region are shown for segmentation with and without induced-reaction reconstruction data in Figure 22. This figure nicely shows the aluminum and HDPE parts of the shields grouped together initially in Figure 22(a) in the blue histogram but distinguished in Figure 22(b) in the blue and magenta histograms. Despite the fact that the steel was not distinguished from the DU for this target, it is noteworthy that four of the primary five regions of this difficult target were able to be isolated.

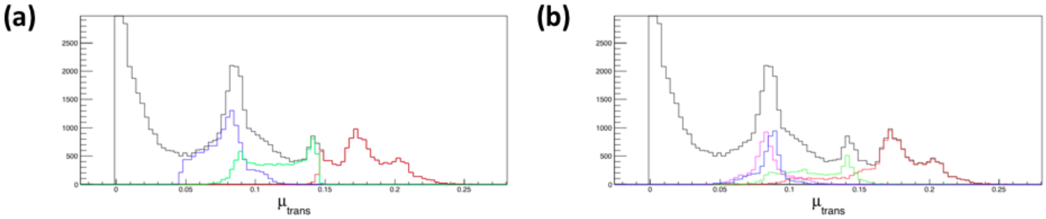

Figure 22. Histograms of transmission image values for all voxels and for the segmented regions for the target consisting of a DU annulus shielded by steel, HDPE, lead, and aluminum. The histogram of transmission image values is shown for all voxels in black in both panels. The segmentation results using transmission only are shown in (a), and the results after inclusion of induced-reaction reconstructions are shown in (b). The colors correspond to the colors used in the 2D and 3D plots in the figures above. The y-axis was chosen so that the large number of small reconstruction values are off scale to better see the distribution of values for the regions being segmented.

6.4 HIGH-DENSITY POLYETHYLENE ANNULUS INSIDE COMPOSITE SHIELD

This target is identical to the previous target with the exception of the replacement of the DU annulus in the center by a HDPE annulus. A drawing identifying the materials in the various regions can be found in Figure 23.

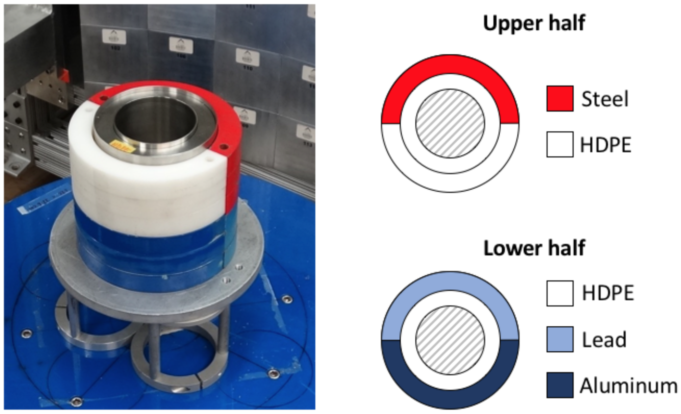

Figure 23. Photograph and material labeling for the target consisting of an HDPE annulus shielded by steel, HDPE, lead, and aluminum. The upper half of the target is shielded by steel and HDPE, and the lower half is shielded by lead and aluminum.

The segmentation of the transmission image produces only two regions for this target (Figure 24). One corresponds primarily to the steel and lead layers of the outer shield (Figure 24[d]), and the other to the

HDPE annulus and the aluminum and HDPE layers of the shield (Figure 24[e]). Note that previously, the steel portion of the shield was grouped with the inner DU annulus, whereas this time it is grouped with the lead. This is somewhat surprising given that the attenuation values for the lead and steel seem to produce two peaks in the distribution (Figure 28[a]).

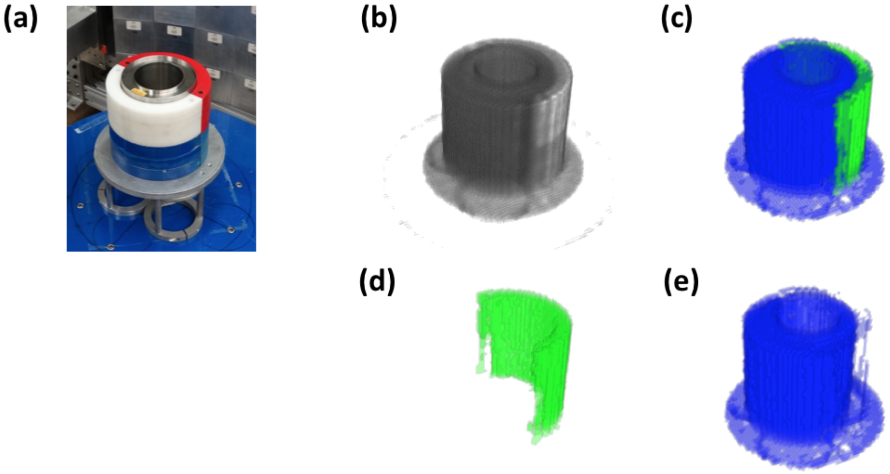

Figure 24. Results of image segmentation using only the transmission image for the target consisting of an HDPE annulus shielded by steel, HDPE, lead, and aluminum. (a) Photograph of the target, (b) transmission reconstruction, (c) segmentation result, and each region shown separately in (d) and (e).

As before, the use of hydrogen-scatter reconstruction data allows the region containing the HDPE outer shield and annulus to be identified as hydrogenous and separate from the aluminum region (Figure 25[d]). The aluminum region (Figure 25[c]) is noisy and contains voxels that belong to other materials, but it captures the primary aluminum region fairly well. Since the induced reactions do not help distinguish lead from steel, the region initially containing those two materials remains unchanged.

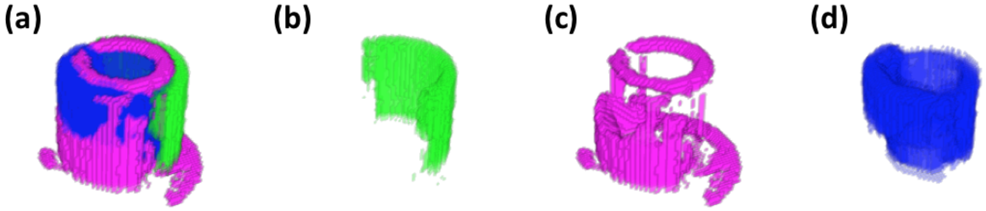

Figure 25. Results of image segmentation using both transmission and induced-reaction imaging data for the target consisting of an HDPE annulus shielded by steel, HDPE, lead, and aluminum. Segmentation regions are shown (a) together and shown separately in (b), (c), and (d). The blue region corresponds to the hydrogenous region.

Tomographic slices through the transmission reconstruction, with hydrogen-scatter reconstruction overlaid, and with segmentations with and without the induced-reaction reconstructions overlaid are shown for slices in the lower and upper halves of the target in Figure 26 and Figure 27, respectively.

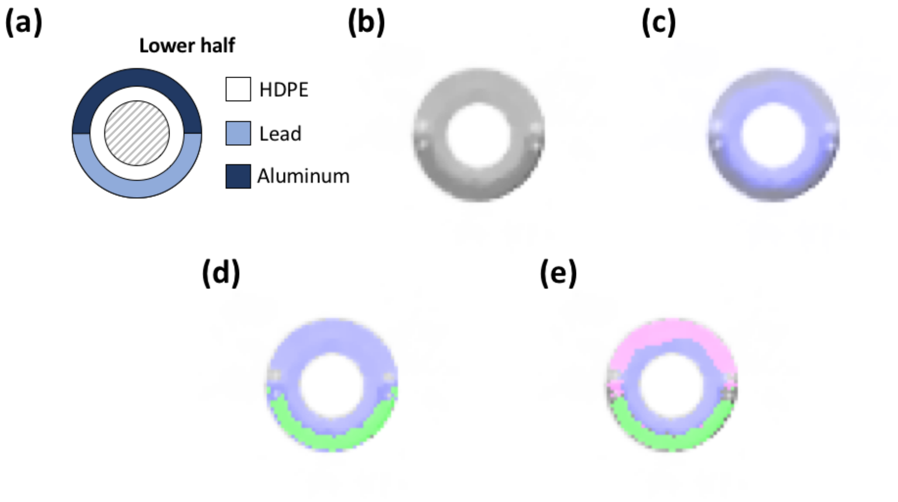

Figure 26. Tomographic slice through the transmission reconstruction with hydrogen-scatter reconstruction and segmented results overlaid for the part of the target consisting of an HDPE annulus shielded by lead and aluminum. For a slice containing the HDPE annulus shielded by lead and aluminum, (a) the drawing labeling material location, (b) the transmission reconstruction, (c) the hydrogen-scatter reconstruction overlaid on the transmission reconstruction, (d) overlay of the transmission-only segmentation on the transmission reconstruction, and (e) overlay of the segmentation that includes induced-reaction reconstructions on the transmission reconstruction. In (e), the blue region corresponds to the hydrogenous region.

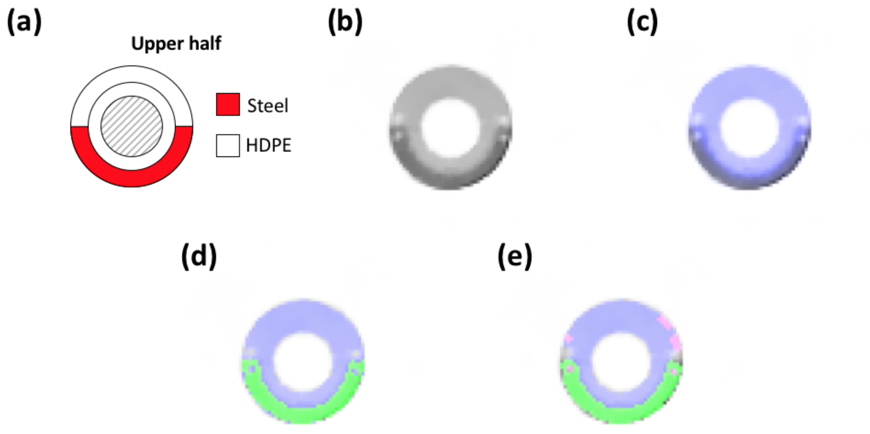

Figure 27. Tomographic slice through the transmission reconstruction with induced-reaction reconstructions and segmented results overlaid for the part of the target consisting of an HDPE annulus shielded by HDPE and steel. For a slice containing the HDPE annulus shielded by HDPE and steel, (a) the drawing labeling material location, (b) the transmission reconstruction, (c) the hydrogen-scatter reconstruction overlaid on the transmission reconstruction, (d) overlay of the transmission-only segmentation on the transmission reconstruction, and (e) overlay of the segmentation that includes induced-reaction reconstructions on the transmission reconstruction. In (e), the blue region corresponds to the hydrogenous region.

The histograms of transmission reconstruction values for each region can be seen in Figure 28. While the hydrogen-scatter data successfully distinguished the HDPE and aluminum, this segmentation was not a total success, given the fact that the steel and lead regions could not be distinguished despite the definite peaks in the distributions seen in Figure 28.

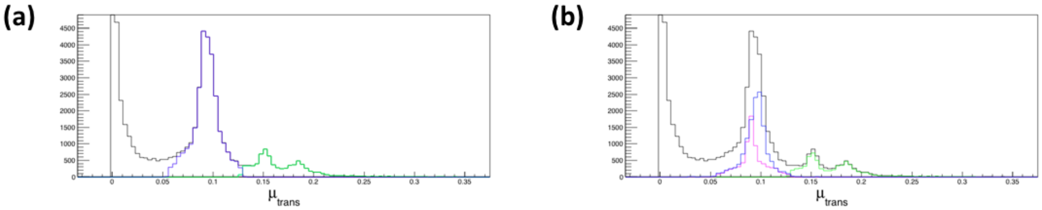

Figure 28. Histograms of transmission image values for all voxels and for the segmented regions for the target consisting of an HDPE annulus shielded by steel, HDPE, lead, and aluminum. The histogram of transmission reconstruction values is shown for all voxels in black in both panels. The segmentation results using transmission only are shown in (a), and the results after inclusion of induced-reaction reconstructions are shown in (b). The colors correspond to the colors used in the 2D and 3D plots in the figures above. The y-axis was chosen so that the large number of small reconstruction values are off scale to better see the distribution of values for the regions being segmented.

6.5 HIGH-DENSITY POLYETHYLENE, LEAD, AND ALUMINUM ANNULI

The target shown in Figure 29 contains an inner HDPE layer, surrounded by a middle layer of lead, and an outer layer of aluminum on the top half of the target. In the bottom half, the material ordering is reversed. This is the first target discussed that contains three layers of materials, although aluminum and HDPE are relatively light materials and therefore not overly challenging to penetrate. However, it provides a test of how dependent the segmentation is on material placement, since the HDPE is on the interior for half of the target, but the exterior for the other.

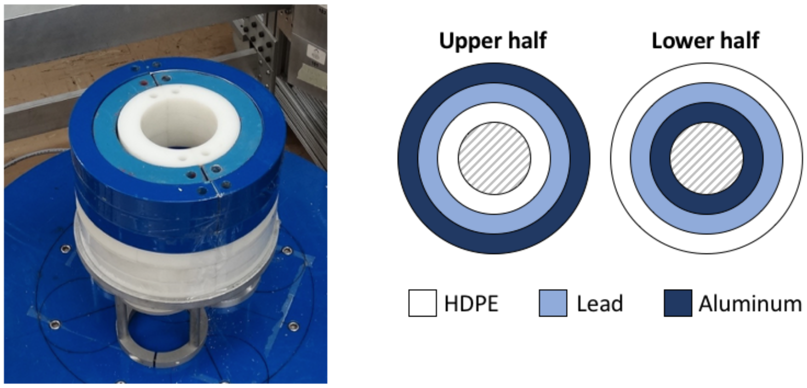

Figure 29. Photograph and material labeling for the target consisting of HDPE, lead, and aluminum annuli. The upper half of the target consists of the materials, from inner layers to outer layers, of HDPE, lead, and aluminum. For the lower half, the layers are reversed.

The segmentation results without use of the hydrogen scatter reconstruction is shown in Figure 30, where, not unexpectedly, two regions are found. The first is the middle layer, shown in Figure 30(d), and the other contains the inner and outer layers, each of which are made up of both HDPE and aluminum. The segmentation is behaving exactly as expected, since aluminum and HDPE should be indistinguishable using transmission alone. When the hydrogen-scatter reconstruction is included, the result shown in Figure 31 is obtained, where the hydrogenous HDPE is isolated from the non-hydrogenous part of the original region.

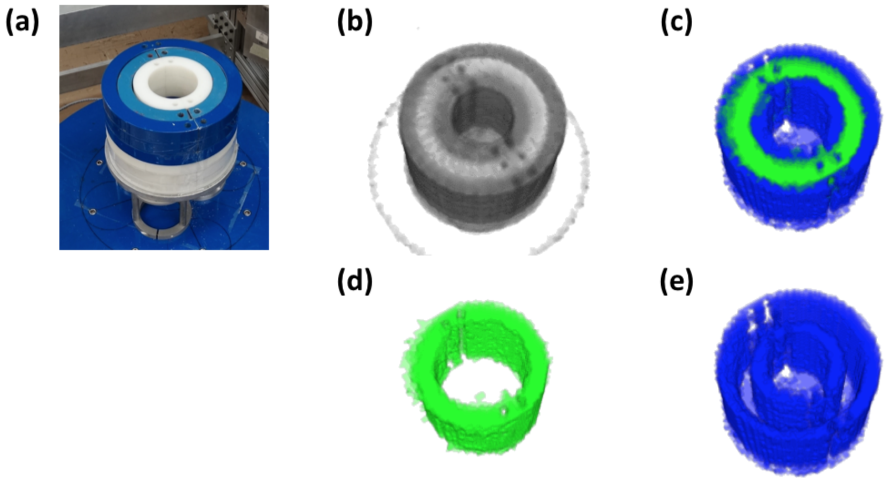

Figure 30. Results of image segmentation using only the transmission image for the target consisting of HDPE, lead, and aluminum annuli. (a) Photograph of the target, (b) transmission reconstruction, (c) segmentation result, and each region shown separately in (d) and (e).

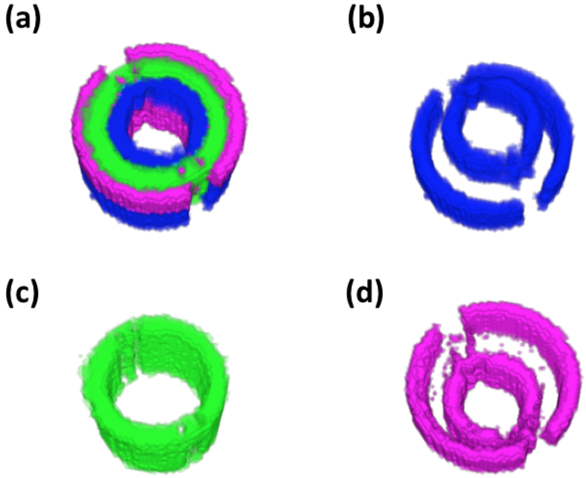

Figure 31. Results of image segmentation using both transmission and induced-reaction reconstruction data for the target consisting of HDPE, lead, and aluminum annuli. Segmentation regions are shown (a) together and shown separately in (b), (c), and (d). The blue region corresponds to the hydrogenous region.

Tomographic slices through the lower and upper halves of the target are shown in Figure 32 and Figure 33, respectively, with overlays from the hydrogen scatter reconstruction and the segmentations with and without hydrogen scatter included. Figure 34 shows the histogram of attenuation values for each segmented region, where the dividing of the initial blue region into hydrogenous and non-hydrogenous regions is clear when moving from transmission alone (Figure 34[a]) to transmission plus the inclusion of hydrogen scatter (Figure 34[b]).

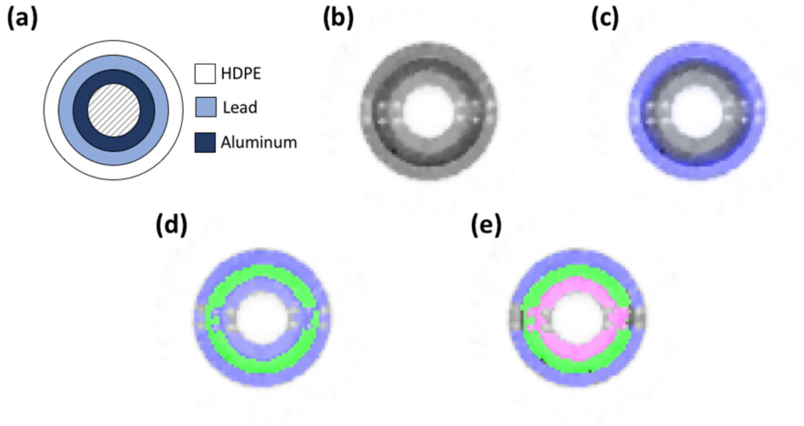

Figure 32. Tomographic slice through the transmission reconstruction with hydrogen-scatter reconstruction and segmented results overlaid for the lower half of the target consisting of HDPE, lead, and aluminum annuli. For a slice from the lower half of the target, (a) the drawing labeling material location, (b) the transmission reconstruction, (c) the hydrogen-scatter reconstruction overlaid on the transmission reconstruction, (d) overlay of the transmission-only segmentation on the transmission reconstruction, and (e) overlay of the segmentation that includes induced-reaction reconstructions on the transmission reconstruction. In (e), the blue region corresponds to the hydrogenous region.

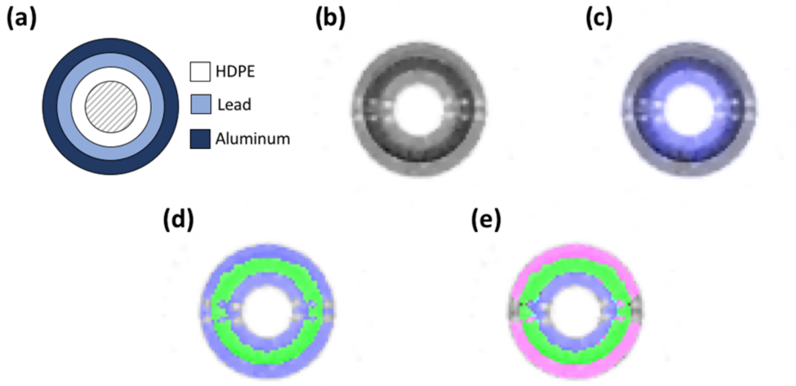

Figure 33. Tomographic slice through the transmission reconstruction with hydrogen-scatter reconstruction and segmented results overlaid for the upper half of the target consisting of HDPE, lead, and aluminum annuli. For a slice from the upper half of the target, (a) the drawing labeling material location, (b) the transmission reconstruction, (c) the hydrogen-scatter reconstruction overlaid on the transmission reconstruction, (d) overlay of the transmission-only segmentation on the transmission reconstruction, and (e) overlay of the segmentation that includes induced-reaction reconstructions on the transmission reconstruction. In (e), the blue region corresponds to the hydrogenous region.

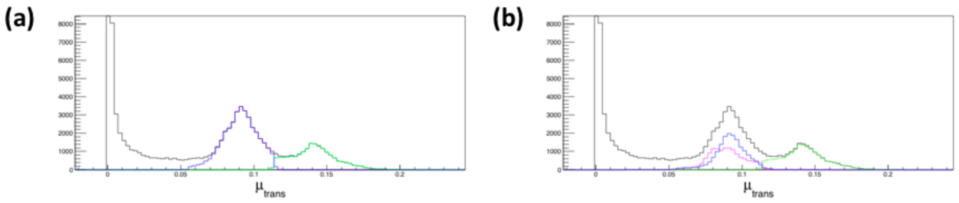

Figure 34. Histograms of transmission image values for all voxels and for the segmented regions for the target consisting of HDPE, lead, and aluminum annuli. The histogram of transmission image values is shown for all voxels in black in both panels. The segmentation results using transmission only are shown in (a), and the results after inclusion of induced-reaction reconstructions are shown in (b). The colors correspond to the colors used in the 2D and 3D plots in the figures above. The y-axis was chosen so that the large number of small reconstruction values are off scale to better see the distribution of values for the regions being segmented.

This segmentation was successful, but only because of the information available from the hydrogen-scatter image, demonstrating the power of the elastic scatter imaging for providing valuable information about regions of the target that would not be available through transmission alone.

6.6 DEPLETED URANIUM ANNULUS SHIELDED BY LEAD AND STEEL

The final target is by far the most challenging because it is the most difficult for the neutrons to penetrate. It contains three layers, from inner to outer: DU, lead, and steel. A photograph and drawing of the material locations are shown in Figure 35.

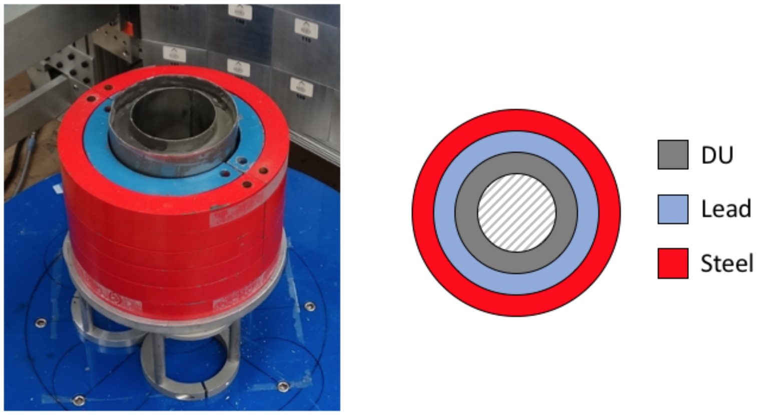

Figure 35. Photograph and material labeling for the target consisting of a DU annulus shielded by lead and steel.

The segmentation result for the transmission reconstruction data is shown in Figure 36, where two regions are identified. One is the middle layer of lead (Figure 36[e]) and the other is a combination of the inner DU and the outer steel (Figure 36[d]).

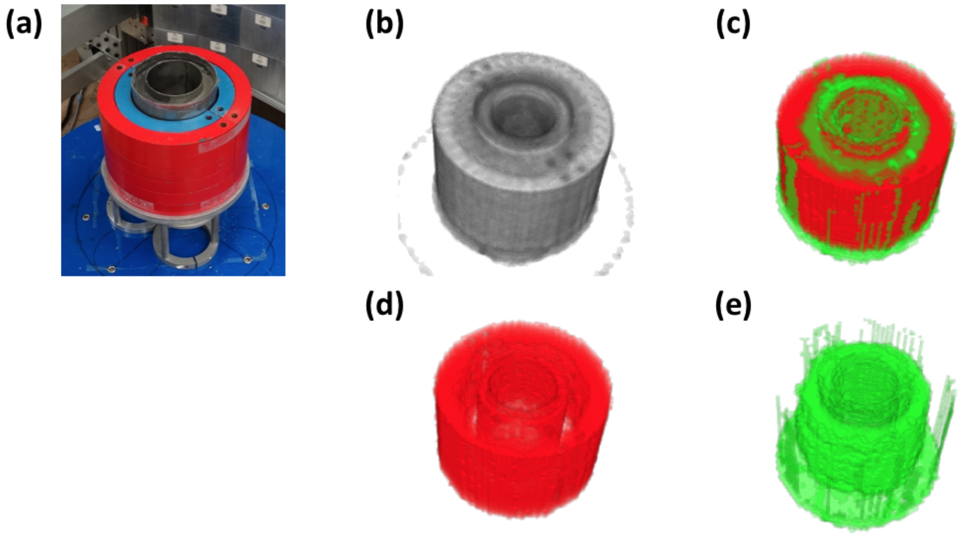

Figure 36. Results of image segmentation using only transmission reconstruction data for the target consisting of a DU annulus shielded by lead and steel. (a) Photograph of the target, (b) transmission reconstruction, (c) segmentation result, and each region shown separately in (d) and (e).

The steel and DU are nearly fully disentangled once the induced-fission reconstruction is included in the segmentation procedure, as seen in Figure 37. There are a few slices in the DU region (Figure 37[b]) that continue to pick out some of the steel, but for the most part the DU is isolated from the rest of the original region as a fissionable region and the layer containing steel is found to be an independent region.

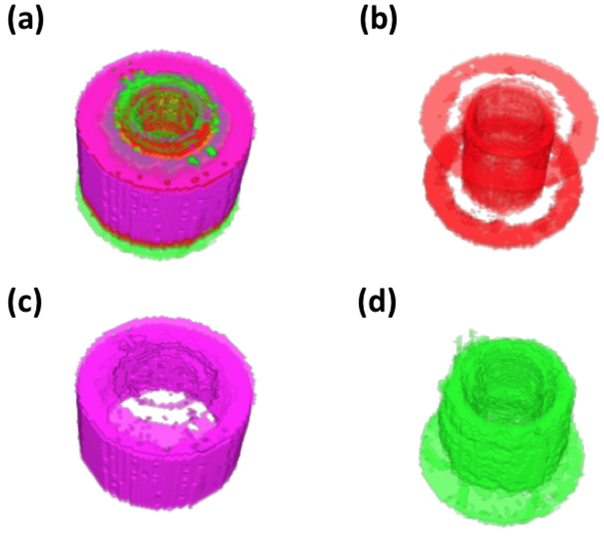

Figure 37. Results of image segmentation using both transmission and induced-reaction reconstruction data for the target consisting of a DU annulus shielded by lead and steel. Segmentation regions are shown (a) together and shown separately in (b), (c), and (d). The red region corresponds to the fissionable region.

A tomographic slice through the reconstruction with the induced-fission reconstruction and segmentation results overlaid are shown in Figure 38. Note that the inner DU layer was isolated from the other layers despite the fact that the induced-fission reconstruction blurs into all three layers (Figure 38[c]).

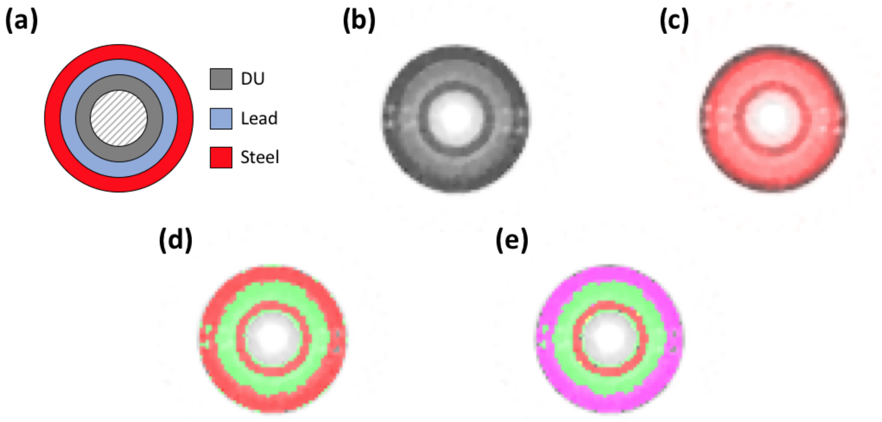

Figure 38. Tomographic slice through the transmission reconstruction with hydrogen-scatter reconstruction and segmented results overlaid for the target consisting of a DU annulus shielded by lead and steel. (a) A drawing labeling material location, (b) the transmission reconstruction, (c) the induced-fission reconstruction overlaid on the transmission reconstruction, (d) overlay of the transmission-only segmentation on the transmission reconstruction, and (e) overlay of the segmentation that includes the induced-fission reconstruction on the transmission reconstruction. In (e), the red region corresponds to the fissionable region.

From the histograms of transmission reconstruction values in Figure 39, note how the DU and steel, which were one region after segmentation using transmission only (Figure 39[a]), are distinguished after induced-fission is included (Figure 39[b]).

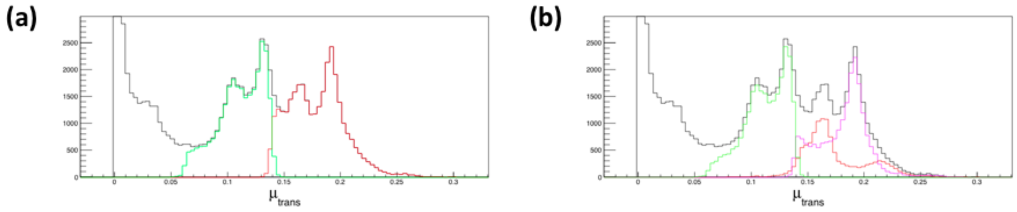

Figure 39. Histograms of transmission image values for all voxels and for the segmented regions for the target consisting of a DU annulus shielded by lead and steel. The histogram of transmission image values is shown for all voxels in black in both panels. The segmentation results using transmission only are shown in (a), and the results after inclusion of induced-reaction reconstructions are shown in (b). The colors correspond to the colors used in the 2D and 3D plots in the figures above. The y-axis was chosen so that the large number of small reconstruction values are off scale to better see the distribution of values for the regions being segmented.

7. VOLUME ESTIMATION

Once the reconstructed data are segmented into various regions, the total volume of each region can be estimated using the number of voxels inside the region and the known voxel volume. This was done for each segmentation result in this report, and the results are summarized in Table 1.

Table 1. Tabulated values of actual volumes of materials vs. volumes calculated from segmentation results.

Target	Fissionable (actual) [cm^3]	Fissionable (calculated) [cm^3]	Hydrogenous (actual) [cm^3]	Hydrogenous (calculated) [cm^3]	Other (actual) [cm^3]	Other (calculated) [cm^3]
DU annulus, HDPE block, hollow lead cube (Section 6.1)	984.6	682	1,333.2	598	—	—
DU annulus, block detector, water in steel (Section 6.2)	984.6	821	533.4	1,518	—	—
DU annulus, composite shield (Section 6.3)	984.6	1,046	463.9	446	1,391.7	1,527
HDPE annulus, composite shield (Section 6.4)	—	—	1,448.5	1,321	1,391.7	1,542
HDPE, lead, aluminum annuli (Section 6.5)	—	—	1,855.5	1,665	3,711.1	2,984
DU, lead, steel annuli (Section 6.6)	984.6	1,170	—	—	4,310.3	5,443

From these calculations, the volume calculations are found to be fairly close for some targets, but grossly off for others. For instance, the volume of fissionable material is found to be close to the true value for a couple of the cases but off by about 30% in the target with the DU annulus, HDPE block, and hollow lead cube. The volume of hydrogenous material is also close to the true values for three of the targets but differs by factors of nearly 2 and 3 for the first two targets listed. Finally, the volume of "other" material exhibits the same behavior—relatively close for some targets but off by as much as 30% for others.

8. SUMMARY AND DISCUSSION

This report presented results on the development of a method that segmented 3D images from a fast-neutron tomographic imaging system. The goal was to isolate the regions occupied by the individual materials in the images for any unknown configuration with a minimum of user input. The segmentation algorithms were applied to measured data for the first time with mixed success.

When the separation between relevant volumes was large, segmentation of the transmission reconstruction data alone was enough to isolate the distinct regions in some cases, such as for the target containing a DU annulus, HDPE block, and hollow lead cube. In the other cases, the induced-reaction reconstructions were necessary to isolate different materials that contained similar attenuation coefficients in the transmission data. Methods of postprocessing and cleaning up results to eliminate edge points attributed to the wrong region were explored and implemented. In addition, an attempt at modifying the energy functional to account for bias and noise in the imaged data was initiated but abandoned, possibly to be explored later.

The 3D multiphase level sets method provides a fast option to detect up to eight different materials, but the underlying problems in some of the input imaging data caused the segmentation routine to perform poorly depending on how the materials were configured in the target under investigation. For example, if steel surrounded DU, then the attenuation values for both materials was significantly lower than theoretical values because of a simplistic and inaccurate scatter correction. In these cases, the two materials would often reconstruct to very similar values, making segmentation more difficult. In addition, resolution effects at the boundaries of volumes, especially when materials were directly adjacent to each other, sometimes made it difficult for the segmentation routine to separate regions. Boundaries in these images are not sharp edges but rather cover several voxels, and an adequate way of dealing with this fact was not developed.

Like the segmentations themselves, attempts at calculating the volumes of the segmented regions encountered varying degrees of success. Some of the calculated volumes were near the actual volumes, but in some cases, they were off by significant amounts. Improvements to the input image data should improve the ability to perform segmentations, which in turn should improve volume calculations.

Future development of image segmentation methods for tomographic, fast-neutron imaging systems should address the following issues:

- Improvements to the data analysis and image reconstruction algorithms so that more accurate representations of the imaged data can be input to the segmentation routines. The highest priority among these improvements is the development of a more accurate scatter correction for transmission imaging.

- Better ways of handling fuzzy edges and noise in the images, potentially similar to what has been explored with Li's method, where the energy functional is modified to account for these effects in the imaged data.

- Additional study into how to best adjust the strength of the length term in the energy functional in such a way that segmentation solutions with stray islands of voxels are properly penalized so that larger bulk volumes are preferred. In addition, exploration into modifications of the image in the energy functional to achieve this result may also prove fruitful.

- Incorporation of transmission and induced-reaction reconstruction data as a one-step process, possibly by the inclusion of additional level sets for induced-fission, hydrogen-scatter, or a combination of both.

Overall, the results on these six data sets showed that the use of the combination of transmission, fission, and hydrogen scatter data to segment 3D images has enormous potential for identifying the location of distinct materials within unknown configurations. Future improvements in data analysis, image reconstruction, and segmentation algorithms should help realize this potential.

9. REFERENCES

[1] P. A. Hausladen and D. Aykac, "Progress Update on Segmentation of Neutron Tomographic Images," ORNL/LTR-2016/544, Oak Ridge National Laboratory, Oak Ridge, TN, 2016.

[2] P. A. Hausladen, D. Aykac, M. A. Blackston, and J. Sparger, "Progress Update on Segmentation of Neutron Tomographic Images, 2017," Report no. ORNL/LTR-2017/285, Oak Ridge National Laboratory, Oak Ridge, TN, 2017.

[3] P. A. Hausladen, M. A. Blackston, and J. Gregor, "Progress Update on Iterative Reconstruction of Neutron Tomographic Images, 2017," Report no. ORNL/SPR-2017/490, Oak Ridge National Laboratory, Oak Ridge, TN, 2017.

[4] D. Mumford and J. Shah, "Optimal Approximation by Piecewise Smooth Functions and Associated Variational Problems," *Comm. Pure Appl. Math.* 42:577–685, 1989.

[5] L. A. Vese and T. F. Chan, "A multiphase Level Set Framework for Image Segmentation using the Mumford and Shah Model," *International Journal of Computer Vision* 50(3), 271–293, 2002.

[6] C. Li, R. Huang, Z. Ding, J. C. Gatenby, D. N. Metaxas and J. C. Gore, "A Level Set Method for Image Segmentation in the Presence of Intensity Inhomogeneities with Application to MRI," *IEEE Trans. on Image Processing*, 20(3), July 2011.