ORNL/LTR-2017/285

Oak Ridge National Laboratory
Progress Update on Segmentation of Neutron Tomographic Images, 2017

Paul Hausladen
Deniz Aykac
Matthew Blackston
John Sparger

May 31, 2017

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

ORNL/LTR-2017/285

Nuclear Security and Isotope Technology Division

PROGRESS UPDATE ON SEGMENTATION OF NEUTRON TOMOGRAPHIC IMAGES, 2017

Paul Hausladen
Deniz Aykac
Matthew Blackston
John Sparger

Date Published: May 31, 2017

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

1. INTRODUCTION

The present report documents progress regarding segmentation of fast-neutron tomographic images of objects consisting of assemblies of machined parts with distinct boundaries. This report satisfies the fiscal year 2017 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. The second of these goals 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 addresses continued progress in the component of the project concerned with image segmentation.

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 avenue for benefiting from the high confidence of imaging methods without revealing the 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 segmented into discrete parts. The properties of a constituent part identified as being constructed of 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) 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 down 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 previous year, the project team reported progress on the first step. The present report describes progress on the second step, namely, operator-guided segmentation of 3D fast-neutron tomographic images. In particular, the team implemented a multiphase level set approach for image segmentation in three dimensions. The multiphase approach eliminated the issues associated with overlaps and gaps in the previous approach that segmented one material at a time.

In this report, the multiphase level sets approach to image segmentation will be described, first in general terms, and then in explicit mathematical detail. Then, initial results on example 3D data sets will be shown. Last of all, future directions for effort will be enumerated.

2. APPROACH

Suppose $g(x,y)$ is an image, such as an x ray, map, or photograph. In many instances, the image g consists of a number of identifiable objects where, within the objects, the image varies smoothly and/or slowly, and across most of the boundaries between the objects, the image varies discontinuously or rapidly. 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 [1],

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

This functional has three terms, each of whose meaning 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 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 [2] provide a framework to calculate the optimal segmentation using a "level-set" formulation. This formulation has the advantages that 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 give an update equation to solve for the optimal $\emptyset_i$ by gradient descent.

Prior to detailing the mathematical formulation, it is useful to consider an example image segmentation. In this case, a photograph of Dr. Donald Hornback is shown in Figure 1 (a) along with its segmentation (b). Again, the purpose of the segmentation is to provide a generally recognizable simplification (cartoon) of the original image, which it does.

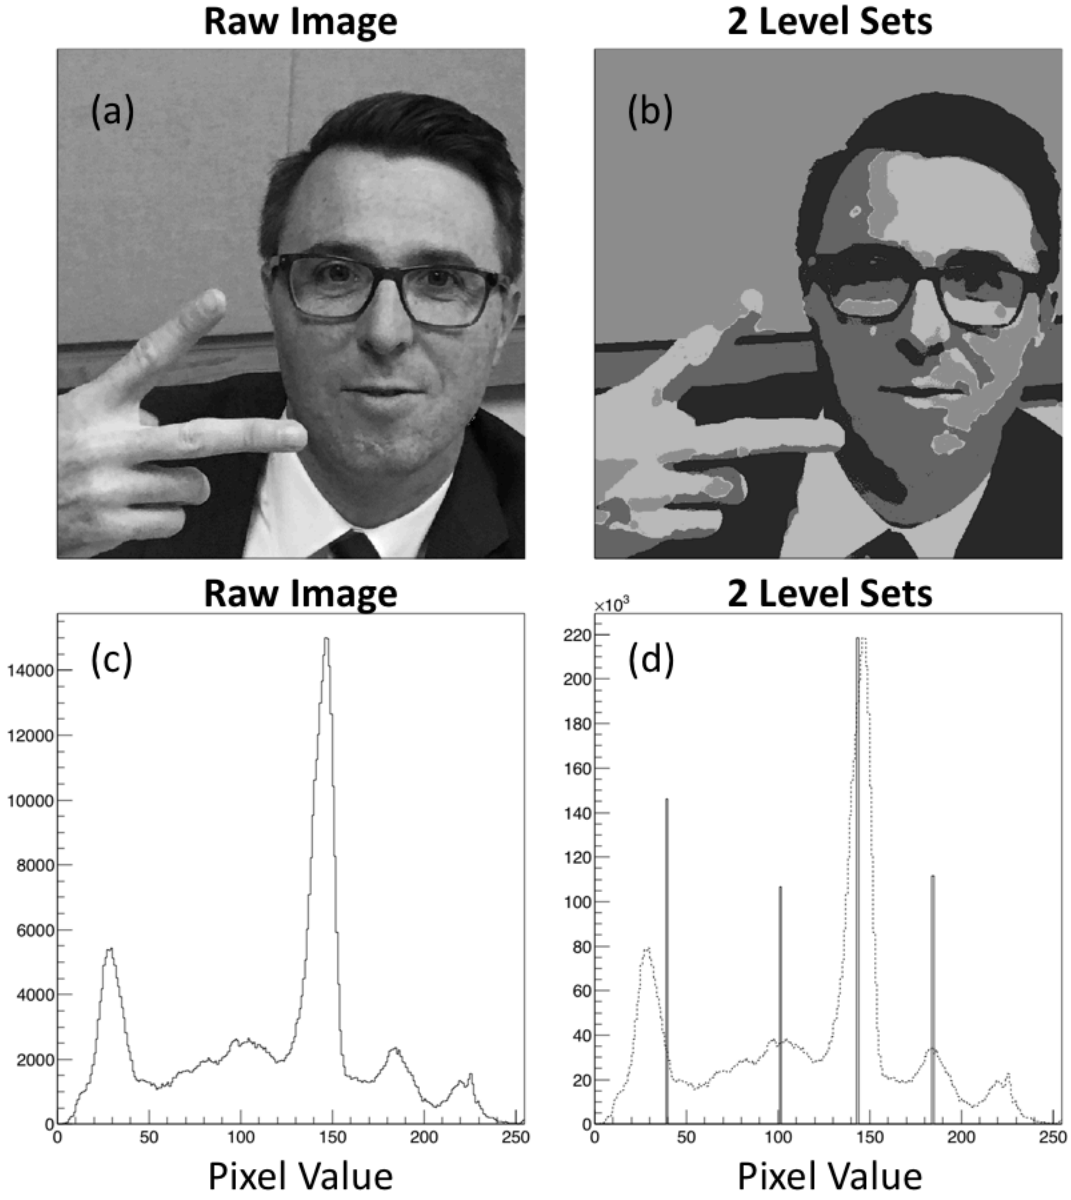

Figure 1. An (a) photograph and (b) image segmentation of Donald Hornback. The image segmentation reduces the distribution of gray levels from the photograph (c) to the four gray levels shown in (d).

A histogram of the frequency of occurrence of grayscale values is shown for the original image in Figure 1 (c) and for the segmented image in Figure 1 (d). In the segmented image, two level sets are used, which allows $2^2 = 4$ regions having different grayscales. In this simplification, some details are lost, such as the features of Dr. Hornback's face. Also, the segmentation is by gray level, so it does not identify the person vs. the background, but instead identifies constituent gray levels. The four regions, each having different gray levels, are enumerated in Figure 2. For fast neutron tomographic images, we will want a similar process to identify segmentations based on image values since distinct image values should identify distinct materials.

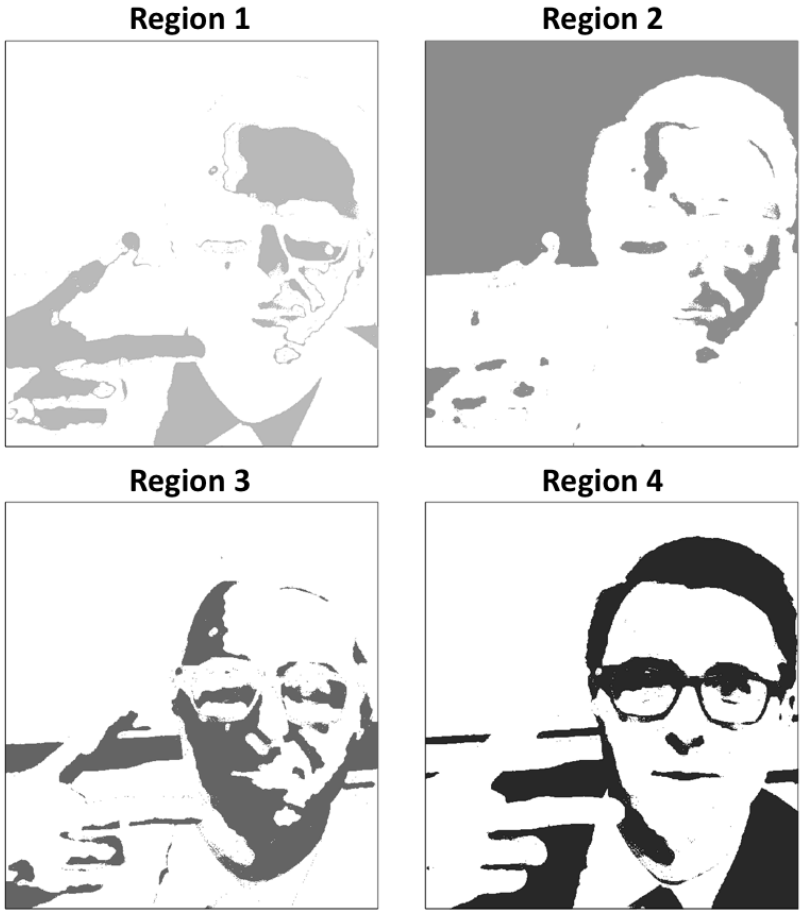

Figure 2. The four constituent regions, each having its own gray level, of the image segmentation shown in Figure 1 (b).

3. IMAGE SEGMENTATION AND THE MULTIPHASE LEVEL-SETS APPROACH

The intent of this section is to provide a compilation of the equations used to solve for image segmentation with one or more level sets in 2D and 3D. A single level set allows an image to be partitioned into two regions (an object and a background), but with m level sets, the image can be partitioned into 2^m regions.

In the remainder of this section, first, the equations associated with a single level set in two dimensions will be reviewed. Then, the equations associated with multiple level sets (multiphase) will be reviewed. Finally, the equations for multiple level sets in three dimensions will be enumerated.

3.1 SEGMENTATION WITH ONE LEVEL SET IN 2D

In the remainder of this document, the nomenclature will change to agree with Ref. [2] rather than Ref. [1]. For an image u_0 that can be divided into region Ω_1 having average image value c_1 and a complementary region Ω_2 having average image value c_2 (that is, the image is approximated as piecewise constant), the Mumford-Shah energy functional can be rewritten:

$$E(c_1, c_2, C) = \int_{\Omega_1}(u_0 - c_1)^2 dxdy + \int_{\Omega_2}(u_0 - c_2)^2 dxdy + \mu\, length(C)\,, \tag{1}$$

where the curve C is the boundary between regions Ω_1 and Ω_2. The optimal segmentation problem is to find the curve C (and thus the regions Ω_1 and Ω_2) that minimizes this energy. In the level-set formulation of Chan and Vese [4], the energy functional can be rewritten in terms of a Lipschitz continuous function $\emptyset$ whose zero-level set defines the boundary $C = \{(x, y)|\, \emptyset(x, y) = 0\}$. Then, region Ω_1 corresponds to the set of points where $\emptyset(x, y) > 0$ and region Ω_2 corresponds to the set of points where $\emptyset(x, y) < 0$. Under these circumstances, the energy term that we seek to minimize can be rewritten:

$$\mathrm{E}(\mathrm{c}_1, \mathrm{c}_2, \phi) = \int_{\Omega}(\mathrm{u}_0 - \mathrm{c}_1)^2 \mathrm{H}(\emptyset)\mathrm{dxdy} + \int_{\Omega}(\mathrm{u}_0 - \mathrm{c}_2)^2 (1 - \mathrm{H}(\emptyset))\mathrm{dxdy} + \mu \int_{\Omega} |\nabla \mathrm{H}(\emptyset_{\mathrm{i}})|\,,$$

where $H(\emptyset)$ is the Heaviside function that equals 1 where $\emptyset(x, y) > 0$ and 0 where $\emptyset(x, y) < 0$, and the length term has been rewritten as the integral of the magnitude of the gradient of the Heaviside function. The function $\emptyset(x, y)$ that minimizes this reformulated Mumford-Shah energy functional is the solution to the Euler-Lagrange equation. Parameterizing the gradient descent by the artificial time t, we get the partial differential equation (PDE)

$$\frac{\partial \emptyset}{\partial \mathrm{t}} = \delta_{\varepsilon}(\emptyset)\left[\mu\ \mathrm{div}\left(\frac{\nabla \emptyset}{|\nabla \emptyset|}\right) - (\mathrm{u}_0 - \mathrm{c}_1)^2 \mathrm{H}(\emptyset) + (\mathrm{u}_0 - \mathrm{c}_2)^2 (1 - \mathrm{H}(\emptyset))\right], \tag{2}$$

where t is an artificial time variable that parameterizes the descent iteration, and $H_\varepsilon(\emptyset)$ and $\delta_\varepsilon(\emptyset)$ are regularized Heaviside and Dirac functions (to ensure differentiability), respectively, given by

$$\mathrm{H}_{\varepsilon}(\emptyset) = \frac{1}{2}\left[1 + \left(\frac{2}{\pi}\right)\tan^{-1}\frac{\emptyset}{\varepsilon}\right], \quad \delta_{\varepsilon}(\emptyset) = \mathrm{H}'_{\varepsilon}(\emptyset) = \frac{1}{\pi}\left(\frac{\varepsilon}{\varepsilon^2 + \emptyset^2}\right). \tag{3}$$

Equation (2) then governs the evolution of the level set $\emptyset$ in order to solve the problem of two-phase segmentation of an image. Then, the values of c_1 and c_2 are given by

$$\mathrm{c}_1(\emptyset) = \mathrm{mean}(\mathrm{u}_0)\,, \mathrm{in}\,\{(\mathrm{x}, \mathrm{y}):\, \emptyset(\mathrm{t}, \mathrm{x}, \mathrm{y}) > 0\,\}, \mathrm{c}_2(\emptyset) = \mathrm{mean}(\mathrm{u}_0)\,, \mathrm{in}\,\{(\mathrm{x}, \mathrm{y}):\, \emptyset(\mathrm{t}, \mathrm{x}, \mathrm{y}) < 0\,\}.$$

To solve the PDE in Eq. (2), we apply a finite difference scheme where

$$\frac{\partial \emptyset}{\partial \mathrm{t}} \rightarrow \frac{\emptyset^{\mathrm{n}+1} - \emptyset^{\mathrm{n}}}{\Delta \mathrm{t}}.$$

The divergence terms are approximated using

$$\mathrm{div}\left(\frac{\nabla \emptyset}{|\nabla \emptyset|}\right) \rightarrow \hat{\Delta}_{\mathrm{x}}\emptyset + \hat{\Delta}_{\mathrm{y}}\emptyset.$$

More specifically, the $div\left(\frac{\nabla \emptyset}{|\nabla \emptyset|}\right)$ term in the x-direction can be expressed with a finite-difference scheme as follows:

$$\frac{\partial}{\partial x}\left(\frac{\frac{\delta \emptyset}{\delta x}}{|\nabla \emptyset|}\right) \rightarrow \quad \hat{\Delta}_x \emptyset = D_-^x \circ \left(\frac{D_+^x \emptyset_{i,j}^n}{\sqrt{(D_+^x \emptyset_{i,j}^n)^2 + (D_c^y \emptyset_{i,j}^n)^2}}\right)$$

$$= \left[\frac{D_+^x \emptyset_{i,j}^n}{\sqrt{(D_+^x \emptyset_{i,j}^n)^2 + (D_c^y \emptyset_{i,j}^n)^2}} - \frac{D_+^x \emptyset_{i-1,j}^n}{\sqrt{(D_+^x \emptyset_{i-1,j}^n)^2 + (D_c^y \emptyset_{i-1,j}^n)^2}}\right], \tag{4}$$

where D_-^x, D_c^x, D_+^x represent forward, central and backward difference operators in the x direction.

As described in Refs. [3, 4], we numerically approximate the iteration of Eq. (2) using the same finite difference scheme explained in Eq. (4) as

$$\frac{\emptyset_{i,j}^{n+1} - \emptyset_{i,j}^{n}}{\Delta t} = \delta_\varepsilon(\emptyset_{i,j}^n) \left\{ \begin{array}{c} \mu(\hat{\Delta}_x \emptyset_{i,j}^n + \hat{\Delta}_y \emptyset_{i,j}^n) - \left(u_{0,i,j} - c_1(\emptyset_{i,j}^n)\right)^2 H(\emptyset_{i,j}^n) \\ + \left(u_{0,i,j} - c_1(\emptyset_{i,j}^n)\right)^2 (1 - H(\emptyset_{i,j}^n)) \end{array} \right\}, \tag{5}$$

where Δt is the time step and $\phi^{n+1}{}_{i,j} = \phi\left(n\Delta t, x_i, y_j\right)$.

The initial $\phi^{0}{}_{i,j}$ is set to be the signed distance function to the initial curve, where the points inside the zero level-set are positive and those outside are negative. At the end of the above iterative process [Eq. (5)], when there is no change between successive iterations, the zero level-set of $\phi^{n}{}_{i,j}$ defines the border of the object of interest.

3.2 SEGMENTATION WITH MULTIPLE LEVEL SETS IN 2D

Using the method of Sect. 3.1, it is possible to segment an image one material (image value) at a time. However, this method is undesirable because of complications associated with having overlaps and gaps between the regions that are found. Instead, it is preferable to use the method of Vese and Chan [2] to use m level sets to simultaneously segment an image into 2^m regions that, by construction, have no overlaps or gaps.

For a single level set ϕ, the two possible regions are defined by the coordinates where $\phi > 0$ and $\phi < 0$. Similarly, for two level sets, there are four regions that can be defined. They are the coordinates where $(\phi_1 > 0, \phi_2 > 0)$, $(\phi_1 > 0, \phi_2 < 0)$, $(\phi_1 < 0, \phi_2 > 0)$, and $(\phi_1 < 0, \phi_2 > 0)$. These combinations are shown pictorially in Figure 3.

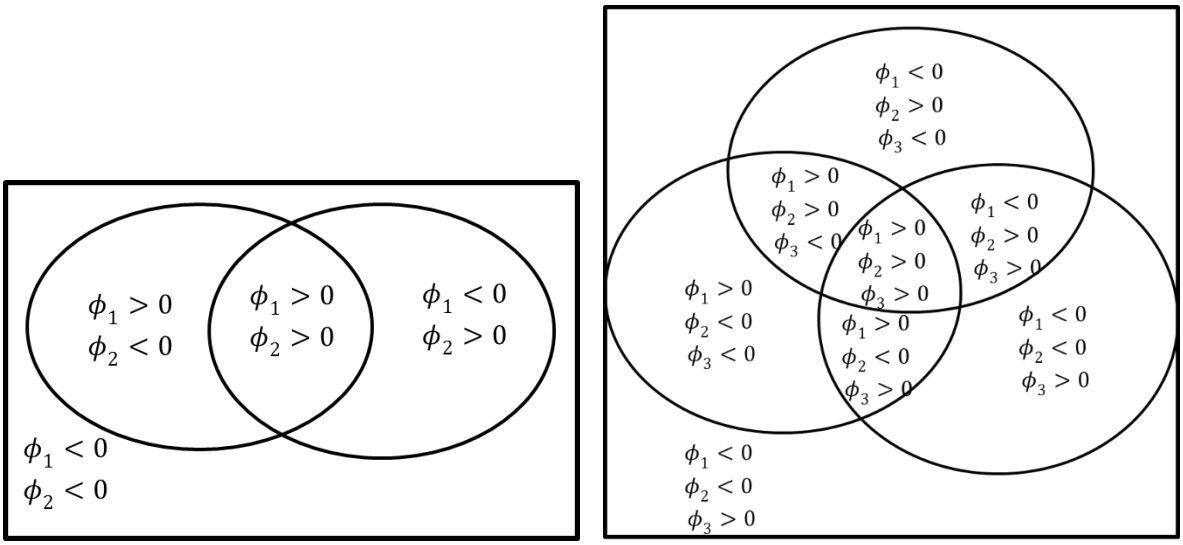

Figure 3. Left: Two level sets partition the domain into four regions ($\phi_1 > 0, \phi_2 > 0$), ($\phi_1 < 0, \phi_2 > 0$), ($\phi_1 > 0, \phi_2 < 0$), ($\phi_1 < 0, \phi_2 < 0$). Right: Three curves partitions the domain into 8 regions ($\phi_1 > 0, \phi_2 > 0, \phi_3 > 0$), ($\phi_1 < 0, \phi_2 > 0, \phi_3 > 0$), ($\phi_1 < 0, \phi_2 > 0, \phi_3 < 0$), ($\phi_1 < 0, \phi_2 < 0, \phi_3 < 0$), ($\phi_1 > 0, \phi_2 > 0, \phi_3 < 0$), ($\phi_1 > 0, \phi_2 < 0, \phi_3 > 0$), ($\phi_1 < 0, \phi_2 < 0, \phi_3 > 0$), ($\phi_1 > 0, \phi_2 < 0, \phi_3 < 0$).

Based on the earlier work, in Ref. [2], the multiphase level set energy functional can be written as

$$E_n(c, \phi) = \sum_{1 \le I \le n = 2^m} \int (u_0 - c_I)^2 \chi_I dxdy + \sum_{1 \le i \le m} \upsilon \int |\nabla H(\emptyset_i)| \,, \tag{6}$$

where $\Phi = (\phi_1, \phi_2, \dots, \phi_m)$ are the vector level set functions. Likewise, $H(\Phi) = \big(H(\phi_1), H(\phi_2), \dots, H(\phi_m)\big)$ are the vector Heaviside functions. The χ_I are the characteristic functions for each region (phase), that is, a combination of Heaviside functions equaling one for only a single region and zero for the remaining regions (phases). The c_I refer to the average image value of the I^{th} region.

In this generalized energy term, $E_n(c, \phi)$, m level set functions $\emptyset_i$ are used to segment the image. There are up to $n = 2^m$ possibilities for the vector values for $H(\Phi)$ in the image, allowing there to be up to $n = 2^m$ regions (phases) in the domain. Therefore, each pixel will belong to one phase by definition. The set of curves C_i is represented by the union of the zero level sets of the functions $\emptyset_i$. For illustration purposes, the above energy term can be written for $n = 4$ phases using $m = 2$ level set functions $\phi = (\emptyset_1, \emptyset_2)$. In particular,

$$\begin{aligned} E_4(c, \phi) = & \int_\Omega (u_0 - c_{11})^2 H(\emptyset_1) H(\emptyset_2) dxdy \\ & + \int_\Omega (u_0 - c_{10})^2 H(\emptyset_1)\big(1 - H(\emptyset_2)\big) dxdy \\ & + \int_\Omega (u_0 - c_{01})^2 (1 - H(\emptyset_1)) H(\emptyset_2) dxdy \\ & + \int_\Omega (u_0 - c_{00})^2 (1 - H(\emptyset_1))(1 - H(\emptyset_2)) dxdy \\ & + \mu \int_\Omega |\nabla H(\emptyset_1)| + \mu \int_\Omega |\nabla H(\emptyset_2)|, \end{aligned} \tag{7}$$

where the index I of the region average image value c_I has been replaced by the vector index of the level sets, which, for two level sets, can take the values $00, 01, 10$, and 11. Now, the analog of Eq. (2) can be written for the four-phase problem as

$$\frac{\partial \emptyset_1}{\partial t} = \delta_\varepsilon(\emptyset_1)\left[\mu\ div\frac{\nabla \emptyset_1}{|\nabla \emptyset_1|} - [((u_0 - c_{11})^2 - (u_0 - c_{01})^2)H(\emptyset_2) + ((u_0 - c_{10})^2 - (u_0 - c_{00})^2)(1 - H(\emptyset_2))]\right],$$

$$\frac{\partial \emptyset_2}{\partial t} = \delta_\varepsilon(\emptyset_2)\left[\mu\ div\frac{\nabla \emptyset_2}{|\nabla \emptyset_2|} - [((u_0 - c_{11})^2 - (u_0 - c_{10})^2)H(\emptyset_1) + ((u_0 - c_{01})^2 - (u_0 - c_{00})^2)(1 - H(\emptyset_1))]\right]. \quad (8)$$

Based on a semi-implicit finite differences scheme, the $n+1^{st}$ approximation to the (i,j) pixel of $\emptyset_1$, $\emptyset_{1,i,j}^{n+1}$, can be obtained from the n^{th} approximation by numerically approximating Eq. (8).

$$\emptyset_{1,i,j}^{n+1} = \frac{1}{C}\left[\emptyset_{1,i,j}^n + \mu\,\Delta t\,\delta_\varepsilon\left(\emptyset_{1,i,j}^n\right)\left(C_1\emptyset_{1,i+1,j}^n + C_2\emptyset_{1,i-1,j}^n + C_3\emptyset_{1,i,j+1}^n + C_4\emptyset_{1,i,j-1}^n\right)\right] + \frac{\Delta t\,\delta_\varepsilon(\emptyset_{1,i,j}^n)}{C}\begin{bmatrix} -\left(u_{0,i,j} - c_{11}^n\right)^2 H_\varepsilon\left(\emptyset_{2,i,j}^n\right) - \left(u_{0,i,j} - c_{10}^n\right)^2\left(1 - H_\varepsilon\left(\emptyset_{2,i,j}^n\right)\right) \\ +\left(u_{0,i,j} - c_{01}^n\right)^2 H_\varepsilon\left(\emptyset_{2,i,j}^n\right) + \left(u_{0,i,j} - c_{00}^n\right)^2\left(1 - H_\varepsilon\left(\emptyset_{2,i,j}^n\right)\right) \end{bmatrix}, \quad (9)$$

where $C = 1 + \mu\Delta t\,\delta_\varepsilon(\emptyset_{1,i,j}^n)(C_1 + C_2 + C_3 + C_4)$, $c = (c_{11}, c_{10}, c_{01}, c_{00})$,

$$C_1 = \frac{1}{\sqrt{(\emptyset_{1,i+1,j}^n - \emptyset_{1,i,j}^n)^2 + \left[\frac{\emptyset_{1,i,j+1}^n - \emptyset_{1,i,j-1}^n}{2}\right]^2}},$$

$$C_2 = \frac{1}{\sqrt{(\emptyset_{1,i,j}^n - \emptyset_{1,i-1,j}^n)^2 + \left[\frac{\emptyset_{1,i-1,j+1}^n - \emptyset_{1,i-1,j-1}^n}{2}\right]^2}},$$

$$C_3 = \frac{1}{\sqrt{\left[\frac{\emptyset_{1,i+1,j}^n - \emptyset_{1,i-1,j}^n}{2}\right]^2 + (\emptyset_{1,i,j+1}^n - \emptyset_{1,i,j}^n)^2}},$$

$$C_4 = \frac{1}{\sqrt{\left[\frac{\emptyset_{1,i+1,j-1}^n - \emptyset_{1,i-1,j-1}^n}{2}\right]^2 + (\emptyset_{1,i,j}^n - \emptyset_{1,i,j-1}^n)^2}},$$

and

$$c_{11}(\phi) = mean(u_0)\,, in\,\{(x,y)\colon \emptyset_1(t,x,y) > 0\,, \emptyset_2(t,x,y) > 0\,\},$$
$$c_{10}(\phi) = mean(u_0)\,, in\,\{(x,y)\colon \emptyset_1(t,x,y) > 0\,, \emptyset_2(t,x,y) < 0\,\},$$
$$c_{01}(\phi) = mean(u_0)\,, in\,\{(x,y)\colon \emptyset_1(t,x,y) < 0\,, \emptyset_2(t,x,y) > 0\,\},$$
$$c_{00}(\phi) = mean(u_0)\,, in\,\{(x,y)\colon \emptyset_1(t,x,y) < 0\,, \emptyset_2(t,x,y) < 0\,\}.$$

3.3 SEGMENTATION WITH MULTIPLE LEVEL SETS IN 3D

The motivation for the present project has been to automatically identify a number of materials (that is, regions or phases) in fast neutron transmission images in 3D. In the following equations, we use three levels sets that, as before, yield a total of $2^3 = 8$ phases.

The previously presented 2D multiphase level set segmentation equations can be extended to 3D with minor modifications. Based on the semi-implicit finite differences scheme in 3D, $\emptyset_{1,i,j,k}^{n+1}$ can be obtained by solving Eq. (8) in 3D. In analogy to the 2D case, the length term can be identified as the portion multiplied by $1/D$:

$$\emptyset_{1,i,j,k}^{n+1} = \frac{1}{D}\Big[\emptyset_{1,i,j,k}^{n} + \mu\Delta t\,\delta_\varepsilon(\emptyset_{1,i,j,k}^{n})(D_1\emptyset_{1,i+1,j,k}^{n} + D_2\emptyset_{1,i-1,j,k}^{n} + D_3\emptyset_{1,i,j+1,k}^{n} + D_4\emptyset_{1,i,j-1,k}^{n} + D_5\emptyset_{1,i,j,k+1}^{n} + D_6\emptyset_{1,i,j,k-1}^{n})\Big] \tag{11}$$

$$+\frac{\Delta t\,\delta_\varepsilon(\emptyset_{1,i,j,k}^{n})}{D}\Big[+\big(u_{0,i,j,k} - c_{111}^{n}\big)^2 H_\varepsilon\big(\emptyset_{2,i,j,k}^{n}\big)H_\varepsilon\big(\emptyset_{3i,j,k}^{n}\big) + \big(u_{0,i,j,k} - c_{101}^{n}\big)^2\Big(1 - H_\varepsilon\big(\emptyset_{2,i,j,k}^{n}\big)\Big)H_\varepsilon\big(\emptyset_{3,i,j,k}^{n}\big) + \big(u_{0,i,j,k} - c_{110}^{n}\big)^2 H_\varepsilon\big(\emptyset_{2,i,j,k}^{n}\big)\big(1 - H_\varepsilon\big(\emptyset_{3,i,j,k}^{n}\big)\big) + \big(u_{0,i,j,k} - c_{100}^{n}\big)^2\Big(1 - H_\varepsilon\big(\emptyset_{2,i,j,k}^{n}\big)\Big)\Big(1 - H_\varepsilon\big(\emptyset_{3,i,j,k}^{n}\big)\Big) - \big(u_{0,i,j,k} - c_{011}^{n}\big)^2 H_\varepsilon\big(\emptyset_{2,i,j,k}^{n}\big)H_\varepsilon\big(\emptyset_{3,i,j,k}^{n}\big) - \big(u_{0,i,j,k} - c_{010}^{n}\big)^2 H_\varepsilon\big(\emptyset_{2,i,j,k}^{n}\big)\Big(1 - H_\varepsilon\big(\emptyset_{3,i,j,k}^{n}\big)\Big) - \big(u_{0,i,j,k} - c_{001}^{n}\big)^2\big(1 - H_\varepsilon\big(\emptyset_{2,i,j,k}^{n}\big)\big)H_\varepsilon\big(\emptyset_{3,i,j,k}^{n}\big) - \big(u_{0,i,j,k} - c_{000}^{n}\big)^2\big(1 - H_\varepsilon\big(\emptyset_{2,i,j,k}^{n}\big)\big)\big(1 - H_\varepsilon\big(\emptyset_{3,i,j,k}^{n}\big)\big)\Big],$$

where $D = 1 + \mu\Delta t\,\delta_\varepsilon(\emptyset_{1,i,j,k}^{n})(D_1 + D_2 + D_3 + D_4 + D_5 + D_6)$, $c = (c_{111}, c_{110}, c_{100}, c_{101}, c_{011}, c_{001}, c_{010}, c_{000})$, $\phi = (\emptyset_1, \emptyset_2, \emptyset_3)$.

$$D_1 = \frac{1}{\sqrt{(\emptyset_{1,i+1,j,k}^{n} - \emptyset_{1,i,j,k}^{n})^2 + \left|\frac{\emptyset_{1,i,j+1,k}^{n} - \emptyset_{1,i,j-1,k}^{n}}{2}\right|^2 + \left|\frac{\emptyset_{1,i,j,k+1}^{n} - \emptyset_{1,i,j,k-1}^{n}}{2}\right|^2}}$$

$$D_2 = \frac{1}{\sqrt{(\emptyset_{1,i,j,k}^{n} - \emptyset_{1,i-1,j,k}^{n})^2 + \left|\frac{\emptyset_{1,i-1,j+1,k}^{n} - \emptyset_{1,i-1,j-1,k}^{n}}{2}\right|^2 + \left|\frac{\emptyset_{1,i-1,j,k+1}^{n} - \emptyset_{1,i-1,j,k-1}^{n}}{2}\right|^2}}$$

$$D_3 = \frac{1}{\sqrt{\left|\frac{\emptyset_{1,i+1,j,k}^{n} - \emptyset_{1,i-1,j,k}^{n}}{2}\right|^2 + (\emptyset_{1,i,j+1,k}^{n} - \emptyset_{1,i,j,k}^{n})^2 + \left|\frac{\emptyset_{1,i,j,k+1}^{n} - \emptyset_{1,i,j,k-1}^{n}}{2}\right|^2}}$$

$$D_4 = \frac{1}{\sqrt{\left|\frac{\emptyset_{1,i+1,j-1,k}^{n} - \emptyset_{1,i-1,j-1,k}^{n}}{2}\right|^2 + (\emptyset_{1,i,j,k}^{n} - \emptyset_{1,i,j-1,k}^{n})^2 + \left|\frac{\emptyset_{1,i,j-1,k+1}^{n} - \emptyset_{1,i,j-1,k-1}^{n}}{2}\right|^2}}$$

$$D_5 = \frac{1}{\sqrt{\left[\frac{\emptyset^n_{1,i+1,j,k} - \emptyset^n_{1,i-1,j,k}}{2}\right]^2 + \left[\frac{\emptyset^n_{1,i,j+1,k} - \emptyset^n_{1,i,j-1,k}}{2}\right]^2 + (\emptyset^n_{1,i,j,k+1} - \emptyset^n_{1,i,j,k})^2}}$$

$$D_6 = \frac{1}{\sqrt{\left[\frac{\emptyset^n_{1,i+1,j,k-1} - \emptyset^n_{1,i-1,j,k-1}}{2}\right]^2 + \left[\frac{\emptyset^n_{1,i,j+1,k-1} - \emptyset^n_{1,i,j-1,k-1}}{2}\right]^2 + (\emptyset^n_{1,i,j,k} - \emptyset^n_{1,i,j,k-1})^2}}$$

and

$$c_{111}(\phi) = mean(u_0)\,, in\,\{(x,y): \emptyset_1(t,x,y) > 0\,, \emptyset_2(t,x,y) > 0, \emptyset_3(t,x,y) > 0\,\}$$
$$c_{110}(\phi) = mean(u_0)\,, in\,\{(x,y): \emptyset_1(t,x,y) > 0\,, \emptyset_2(t,x,y) > 0, \emptyset_3(t,x,y) < 0\,\}$$
$$c_{100}(\phi) = mean(u_0)\,, in\,\{(x,y): \emptyset_1(t,x,y) > 0\,, \emptyset_2(t,x,y) < 0, \emptyset_3(t,x,y) < 0\,\}$$
$$c_{101}(\phi) = mean(u_0)\,, in\,\{(x,y): \emptyset_1(t,x,y) > 0\,, \emptyset_2(t,x,y) < 0, \emptyset_3(t,x,y) > 0\}$$
$$c_{011}(\phi) = mean(u_0)\,, in\,\{(x,y): \emptyset_1(t,x,y) < 0\,, \emptyset_2(t,x,y) > 0, \emptyset_3(t,x,y) > 0\,\}$$
$$c_{001}(\phi) = mean(u_0)\,, in\,\{(x,y): \emptyset_1(t,x,y) < 0\,, \emptyset_2(t,x,y) < 0, \emptyset_3(t,x,y) > 0\,\}$$
$$c_{010}(\phi) = mean(u_0)\,, in\,\{(x,y): \emptyset_1(t,x,y) < 0\,, \emptyset_2(t,x,y) > 0, \emptyset_3(t,x,y) < 0\,\}$$
$$c_{000}(\phi) = mean(u_0)\,, in\,\{(x,y): \emptyset_1(t,x,y) < 0\,, \emptyset_2(t,x,y) < 0, \emptyset_3(t,x,y) < 0\,\}$$

4. EXAMPLE SEGMENTATION RESULTS

The project team expected to have measured 3D results available for inclusion in this report; however, delays in hardware availability precluded their inclusion. As a result, the example images used for segmentation are the result of image reconstructions of ray-trace simulations that only include the effects of finite counting statistics, detector size, and number of projections. In each case, the same neutron radiography phantom was used. It consists of two sets of concentric cylinders separated by 20 cm, as shown in Figure 4.

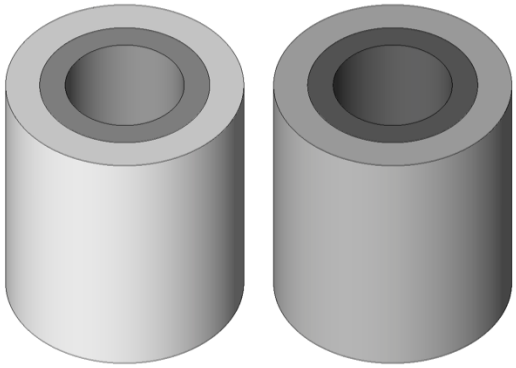

Figure 4. Radiography phantom used for image reconstruction and segmentation.

For each set, the inner cylinder has an inner diameter of 7.62 cm and an outer diameter of 12.7 cm. The outer cylinder has an inner diameter equaling the outer diameter of the smaller cylinder and an outer diameter of 17.78 cm. The attenuation coefficients of the cylinders are 0.28 cm^{-1} and 0.18 cm^{-1} for the pair of cylinders on the right and 0.22 cm^{-1} and 0.12 cm^{-1} for the pair of cylinders on the left. The values of the attenuation coefficients were chosen to correspond to materials for which ORNL has cylinders that can be measured once the imaging hardware is operational. For instance, 0.18 cm^{-1} corresponds to lead,

0.22 cm^{-1} corresponds to steel, and 0.28 cm^{-1} corresponds to depleted uranium. In addition, the set of values were chosen to have some (e.g., lead and steel) that are similar enough to potentially be difficult to distinguish from each other in noisy measurements.

Ray-tracing projection data were generated for 100 projections of the object, each projection rotated by 3.6 degrees compared to the previous. For the purpose of the ray tracing, the object and detectors were placed at distances of 57 cm and 143 cm from the source position, respectively. Images were reconstructed from the ray-tracing projection data. One set of images, henceforth referred to as the "ideal" images, was reconstructed from idealized ray-trace projection data having no statistical error. In these images, the error originates only from the finite sampling of the objects in the projection data. Likewise, a second set of images was generated. For these images, Poisson-distributed statistical noise was added to the projection data according to the expected number of detected neutrons for a source flux of 4×10^7 neutrons s^{-1} for 120 s. In the remainder of this section, these images will be referred to as the "noisy" images. Last of all, a third set of images was generated. The noisy projection data were also used to generate a third set of images, but the images were reconstructed using a total variation constraint. The total variation (TV) constraint is a constraint on the sum of the magnitude of the gradient of a reconstructed image and selects images that are largely piecewise uniform with sharp edges. It is an appropriate constraint for reconstructing objects consisting of machined parts such as those simulated for the present exercise. In the remainder of this section, these images will be referred to as the "TV constrained" images. Example (a) ideal, (c) noisy, and (d) TV-constrained noisy images can be seen in Figure 5. For comparison, a TV-constrained ideal image is also shown, although this image was not used to generate image segmentations.

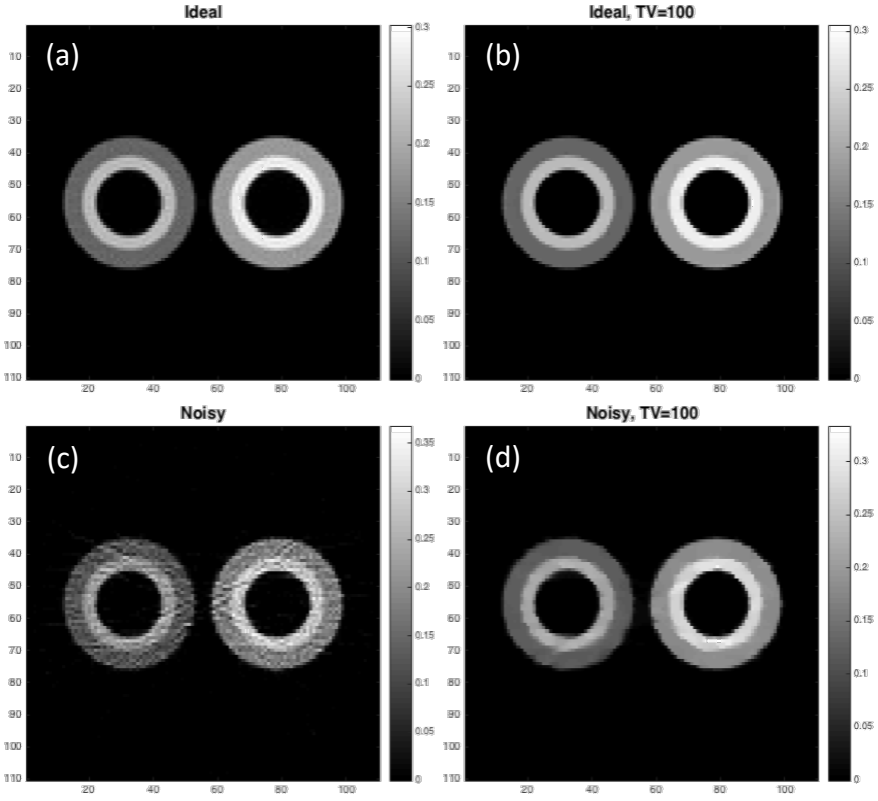

Figure 5. Example reconstructed images of (a) ideal and (c) noisy data. Example total-variation-constrained reconstructions for the same data are shown in (b) and (d).

4.1 SEGMENTATION RESULTS FOR THE IDEAL IMAGE

In this section, we report on segmentation results for the ideal image. First, initialization of the level sets are discussed. Then, the level sets will be shown as a function of iteration number in 2D and 3D. Finally, the resultant segmented volumes will be shown in 2D and 3D.

In order to perform the segmentation, we need an initial curve (or curves) that corresponds to the zero-level set of the function $\emptyset^0$. For simplicity, we use a sphere (or spheres) in 3D. Then, $\emptyset$ is the signed distance function in which zero is assigned to the zero-level set of the function. The points inside the zero-level set of the function are positive, and those outside are negative and increase in magnitude in proportion to the distance from the zero-level set. A number of initial conditions were investigated for the segmentation process and are shown in Figure 6 as a 2D slice (upper row) and for a single level set in 3D (lower row).

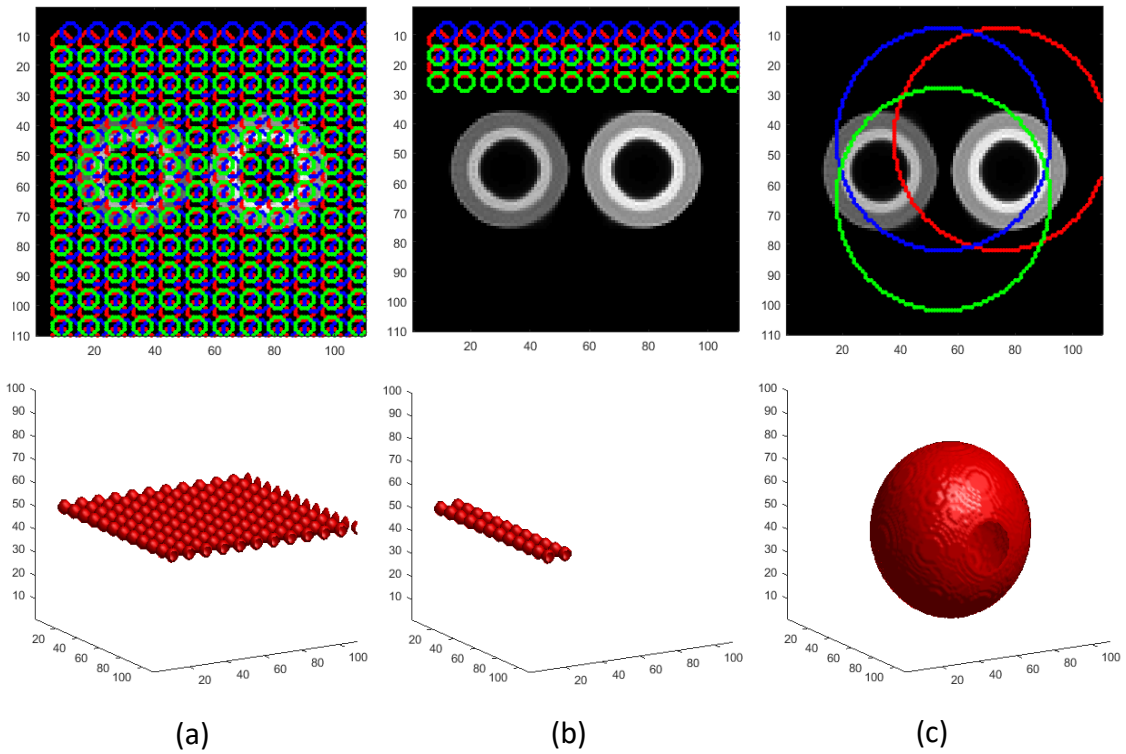

Figure 6. A number of initial level set configurations sliced on the middle of the z-axis (upper row) and rendered in 3D for one such curve (lower row). The configurations (a) cover the whole image with small spheres, (b) cover a background area of the image with small spheres, and (c) cover much of the image with a single sphere.

The separate initial configurations included (a) one covering the plane at the midpoint in z with small spheres, (b) one covering a background area of the plane at the midpoint in z with small spheres, and (c) one centering large spheres on the plane at the midpoint in z. During the process of exercising the code, we found that the final solution depended on the initial conditions of the level set. For gradient descent methods, it is often possible to find a local rather than a global minimum, so initial conditions can be important. In our experience, better performance was achieved with multiple spheres such as (a) and (b) in Figure 6 than with a single sphere, shown in (c).

For the ideal data set, there is no statistical error. As a result, the 3D image data ends up being a stack of copies of the single slice shown in Figure 5 (a), and the stack was padded on either end with empty images so that the cylinders have an extent smaller than the height of the data set and the bounding surfaces of the regions can close around the cylinders. Level-set-evolution results for the ideal data are shown in Figure 7 as a function of iteration number for a 2D plane that bisects the midpoint of the object. The corresponding 3D rendering of the level sets as a function of iteration number is shown in Figure 8, and several views of the terminal level sets (iteration 600) are shown in Figure 9. The thin vertical features are present because each of the images is identical, so edge values that are consistent with another material that are present in one layer will be present in all the layers.

The three level sets are then combined to form the eight regions. The regions are shown in Figure 10 in 2D, for a plane that bisects the object at the midpoint in z. The narrow rings that are apparent in some of the regions (e.g., the right side of region 2) are the result of selecting image pixels along a boundary gradient between two objects that inadvertently have the appropriate image value. As a postprocessing step, these small pieces are attributed to their most numerous neighbors. Then, the four regions containing objects (1–4) are shown rendered in 3D in Figure 11.

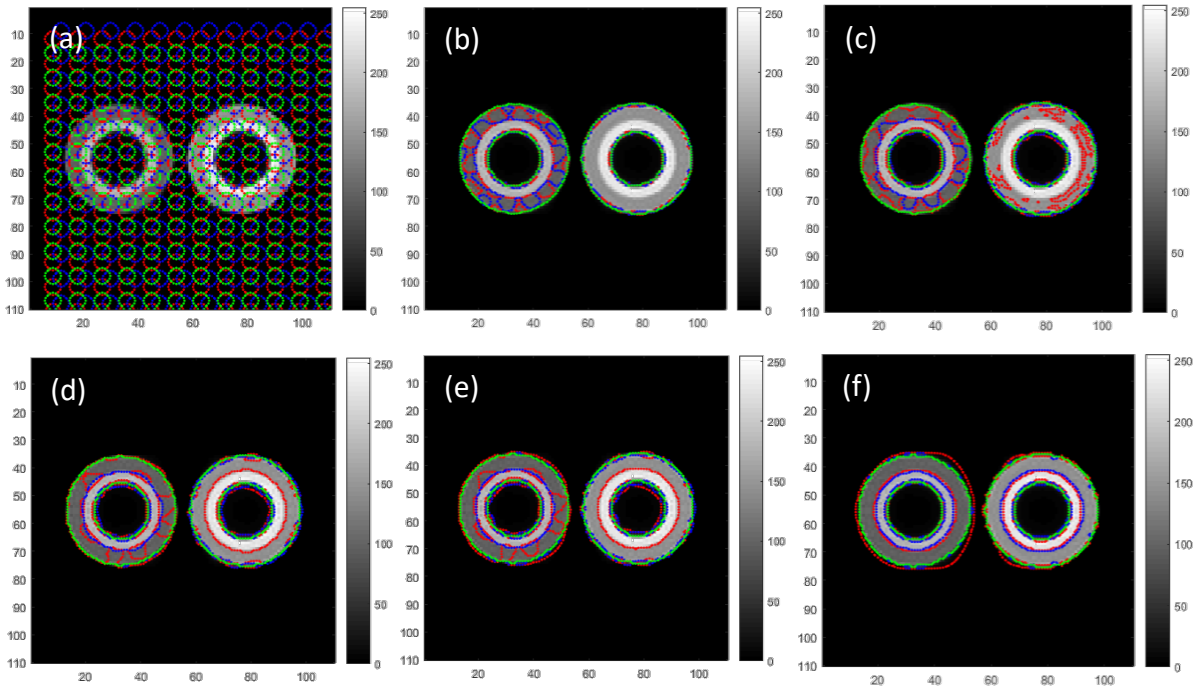

Figure 7. Segmentation as a function of iteration number for the ideal case, shown for the plane at the midpoint in z for iteration (a) 1, (b) 10, (c) 20, (d) 30, (e) 40, and (f) 600. For this segmentation, the scalar multiplier of the length term μ= 0.008*255^2, and the time step Δt= 0.8.

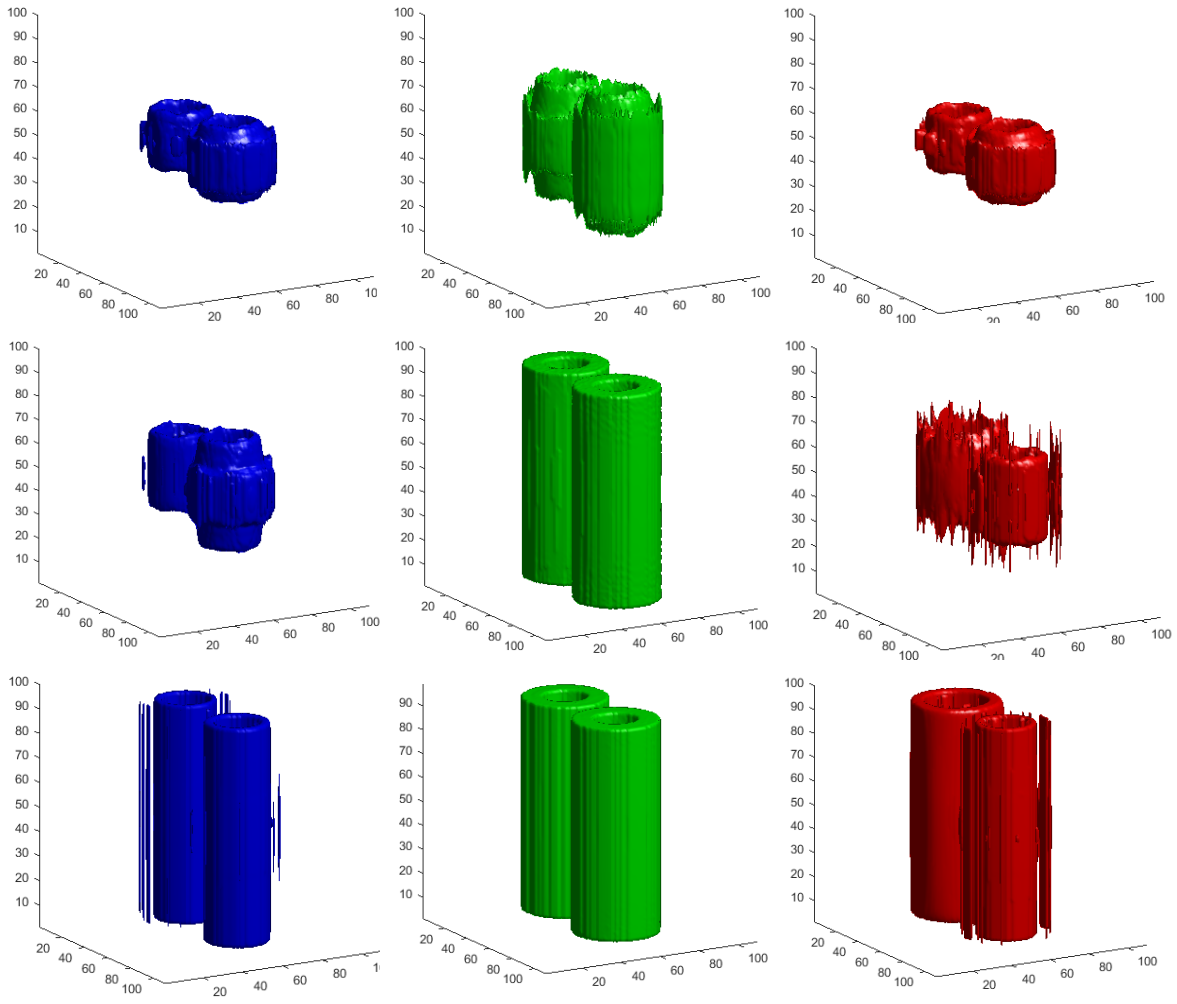

Figure 8. Segmentation results for the ideal case as a function of iteration number. Top row: the three zero-level set surfaces at iteration=10. Middle row: the three zero-level set surfaces at iteration=40. Bottom row: the three zero-level set surfaces at iteration=600.

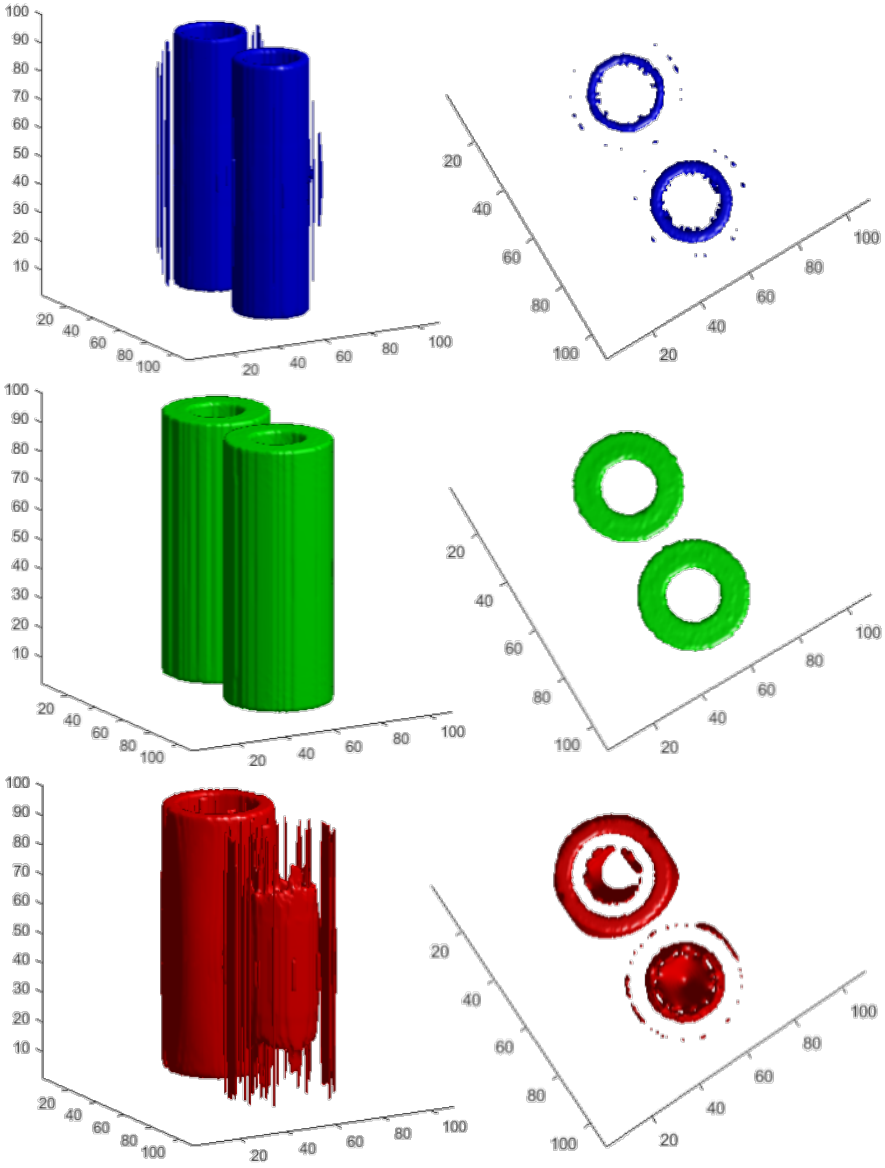

Figure 9. Two different views of the three zero-level set surfaces for iteration=300 for the ideal case.

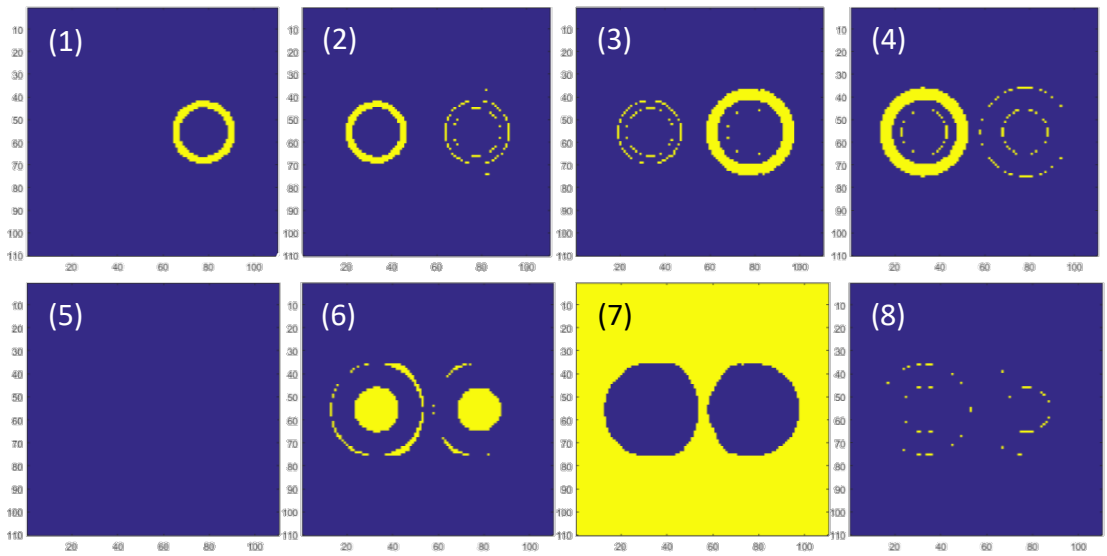

Figure 10. Final segmented regions for the ideal case (prior to postprocessing steps) shown along a plane bisecting the object at its midpoint. Regions 6 and 7 are both notionally empty regions, but the image reconstruction algorithm typically converges to a small nonzero value of the attenuation coefficient in interior voids.

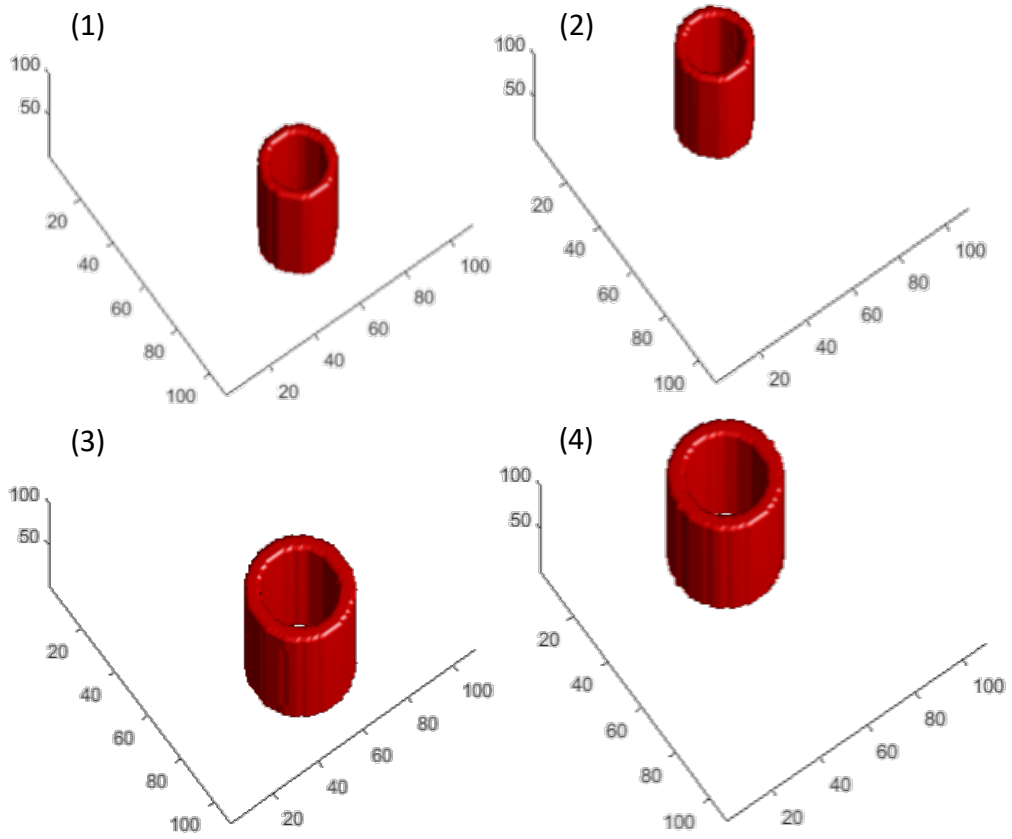

Figure 11. Final segmented surfaces of regions 1-4 for the ideal case (after postprocessing steps).

4.2 SEGMENTATION RESULTS FOR THE TV-CONSTRAINED NOISY IMAGE

For machined components, it is reasonable to assume that the resultant image should consist of smooth regions with rapid or discontinuous jumps between them. Such an image is sparse in gradient, so image reconstruction that constrains for total variation (the sum of the magnitude of the gradients over the image) will identify images that are more representative of the object that was imaged while producing images that are easier to segment (that is, the piecewise constant approximation is more accurate). As a result, image segmentation was performed on a reconstructed image of the same imaging phantom as was imaged in the previous section. However, this time, the data had statistical noise, and the imaging reconstruction used a TV constraint to partially counteract the statistical imperfections in the data. The segmentation using the TV-constrained image segmentation is nearly as good as that achieved with the ideal images.

The level-set-evolution results for the TV-constrained noisy data are similar to the results given in Sect. 4.1. The results are shown Figure 12 as a function of iteration number for a 2D plane that bisects the midpoint of the object. The corresponding 3D rendering of the level sets as a function of iteration number is shown in Figure 13, and several views of the terminal level sets (iteration 600) are shown in Figure 14.

The three level sets are then combined to form the eight regions. The regions are shown in Figure 15 in 2D for a plane that bisects the object at the midpoint in z. The narrow rings that are apparent in some of the regions (e.g., the right side of region 2) are the result of selecting image pixels along a boundary gradient between two objects that inadvertently have the appropriate image value. As a postprocessing step, these small pieces are attributed to their most numerous neighbors. Then, the four regions containing objects (1, 2, 7, and 8) are shown rendered in 3D in Figure 16.

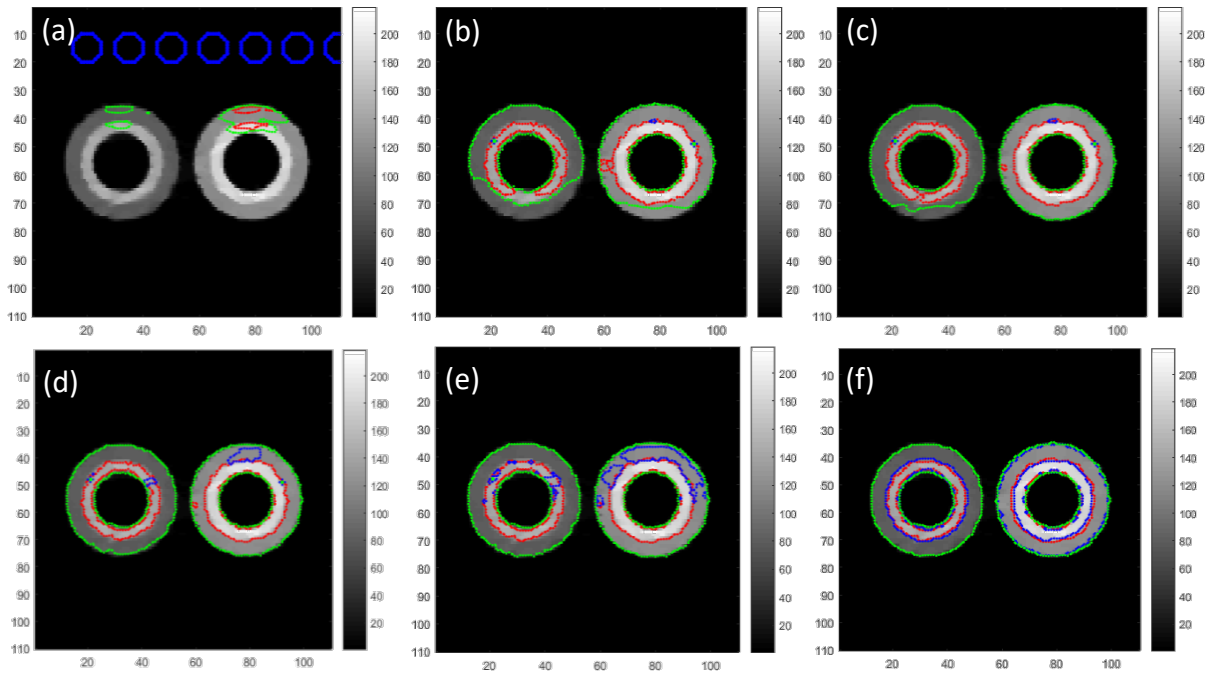

Figure 12. Segmentation as a function of iteration number for the TV-constrained noisy case, shown for the plane at the midpoint in z for iteration (a) 1, (b) 10, (c) 20, (d) 30, (e) 40, and (f) 600. For this segmentation, the scalar multiplier of the length term μ= 0.008*255^2, and the time step Δt= 0.8.

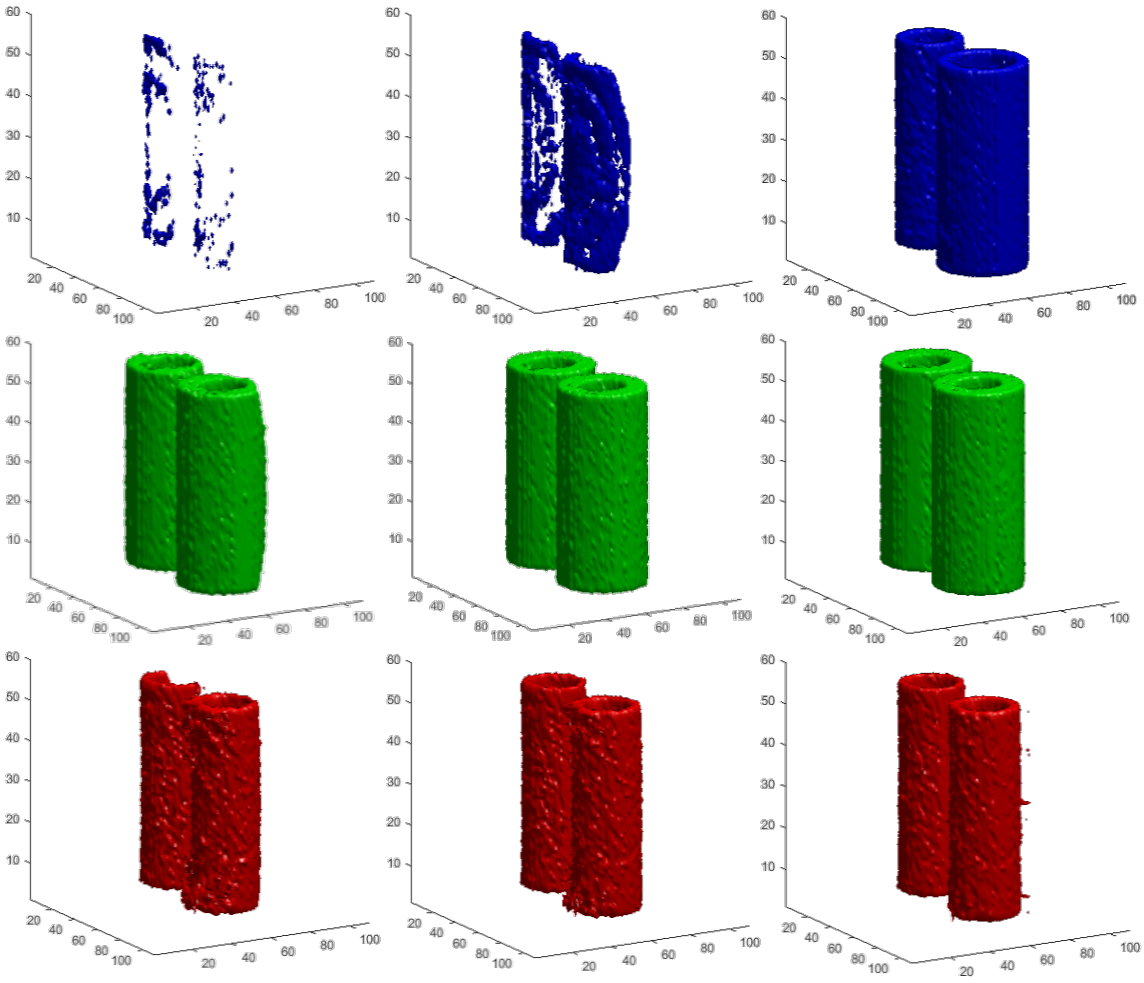

Figure 13. Segmentation results for the TV-constrained noisy data case as a function of iteration number. Top row: the three zero-level set surfaces at iteration=10. Middle row: the three zero-level set surfaces at iteration=40. Bottom row: the three zero-level set surfaces at iteration=600.

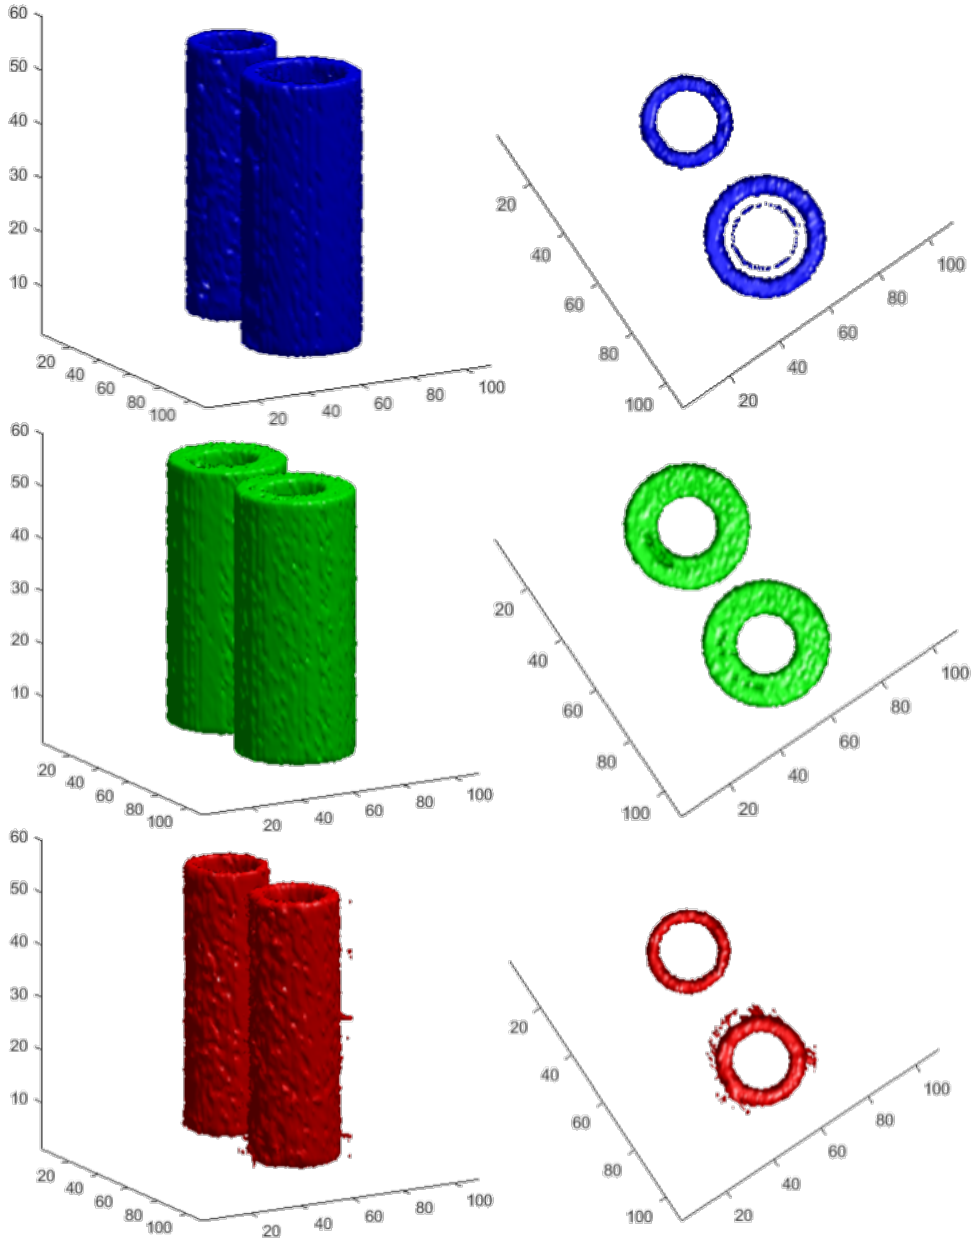

Figure 14. Two different views of the three zero-level set surfaces for iteration=300 for the TV-constrained noisy data case.

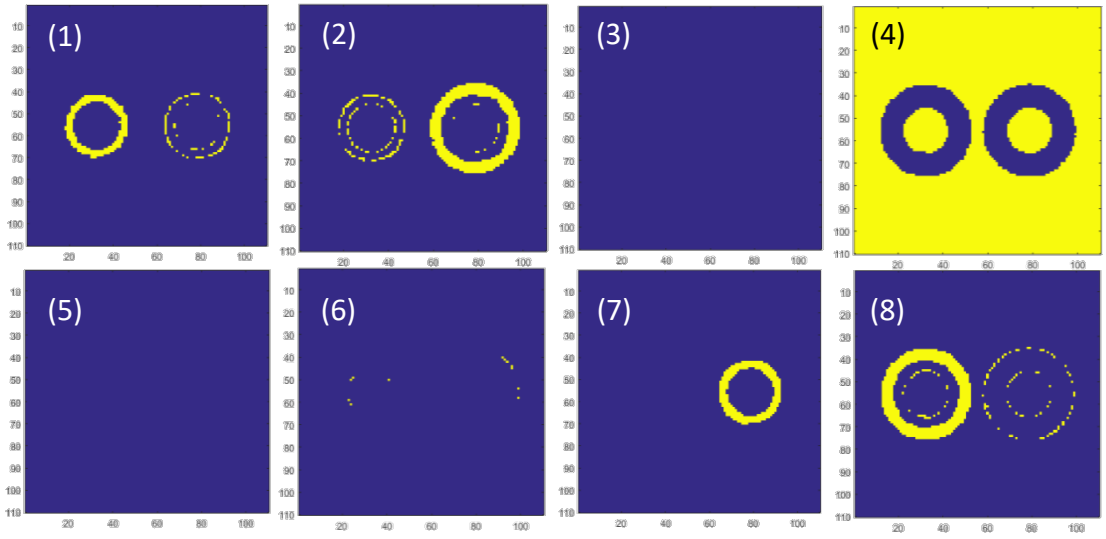

Figure 15. Final segmented regions for the TV-constrained noisy case (prior to post-processing steps) shown along a plane bisecting the object at its midpoint.

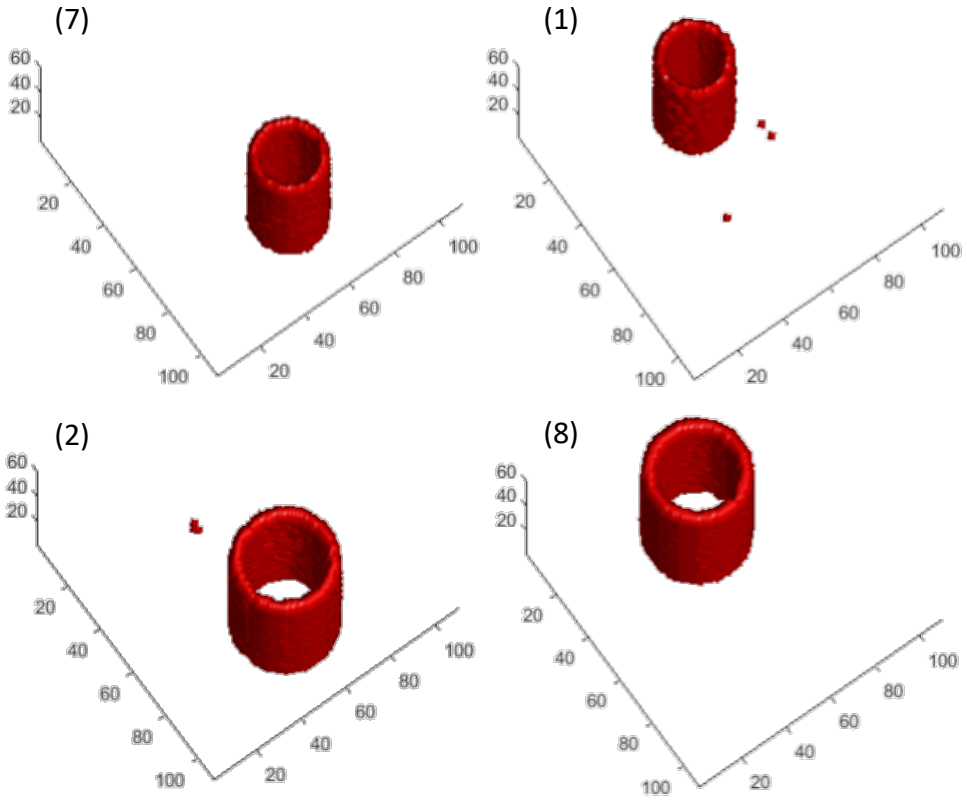

Figure 16. Final segmented surfaces of regions 1, 2, 7, and 8 for the TV-constrained noisy case (after postprocessing steps). The numbering is arbitrary and depends on the initial conditions of the gradient descent.

4.3 SEGMENTATION RESULTS FOR THE NOISY IMAGE

The magnitude of statistical noise in an image can cause significant difficulty for the image segmentation algorithm. To assess this issue for the present segmentation algorithm, image segmentation was performed for a reconstruction of the same data as was used in Sect. 4.2, but without the TV constraint. The resulting reconstructed image contained much more noise. As anticipated, the image segmentation achieved using this noisy data is markedly inferior to that achieved with the ideal image or TV-constrained image. Because of the magnitude of difficulty working with the noisy data, an amount of de-noising was necessary to achieve reasonable segmentation results. For the presented results, each 2D slice of the 3D data set was smoothed with a 3×3 median filter prior to segmentation.

Similar to the previous section, level-set-evolution results for the noisy data are shown in Figure 17 as a function of iteration number for a 2D plane that bisects the midpoint of the object. The corresponding 3D rendering of the level sets as a function of iteration number is shown in Figure 18, and several views of the terminal level sets (iteration 600) are shown in Figure 19.

The three level sets are then combined to form the eight regions. The regions are shown in Figure 20 in 2D, for a plane that bisects the object at the midpoint in z. The narrow rings that are apparent in some of the regions (e.g., the right side of region 2) are the result of selecting image pixels along a boundary gradient between two objects that inadvertently have the appropriate image value. As a postprocessing step, these small pieces are attributed to their most numerous neighbors. Then, the four regions containing objects were rendered in 3D (1, 5, 7, and 8) (shown Figure 21).

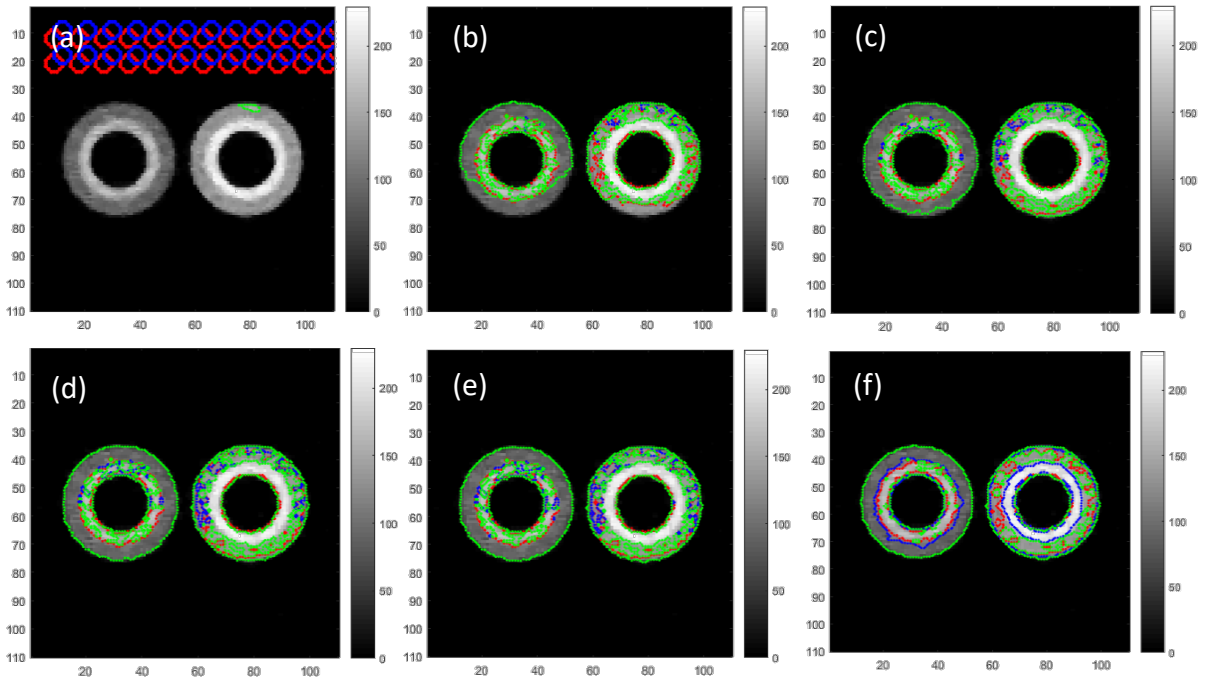

Figure 17. Segmentation as a function of iteration number for the noisy case, shown for the plane at the midpoint in z for iteration (a) 1, (b) 10, (c) 20, (d) 30, (e) 40, and (f) 600. For this segmentation, the scalar multiplier of the length term μ= 0.008*255^2, and the time step Δt= 0.8.

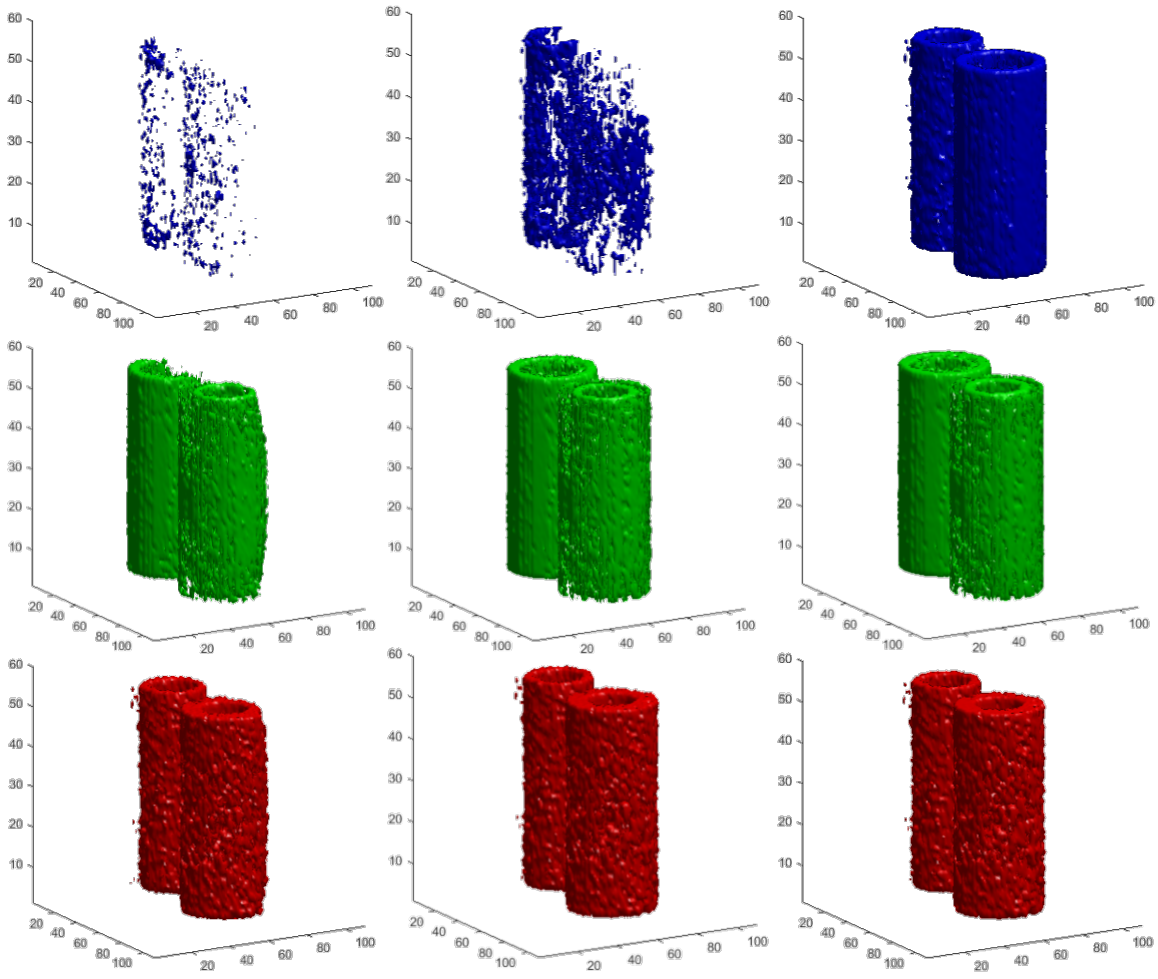

Figure 18. Segmentation results for the noisy data case as a function of iteration number. Top row: the three zero-level set surfaces at iteration=10. Middle row: the three zero-level set surfaces at iteration=40. Bottom row: the three zero-level set surfaces at iteration=600.

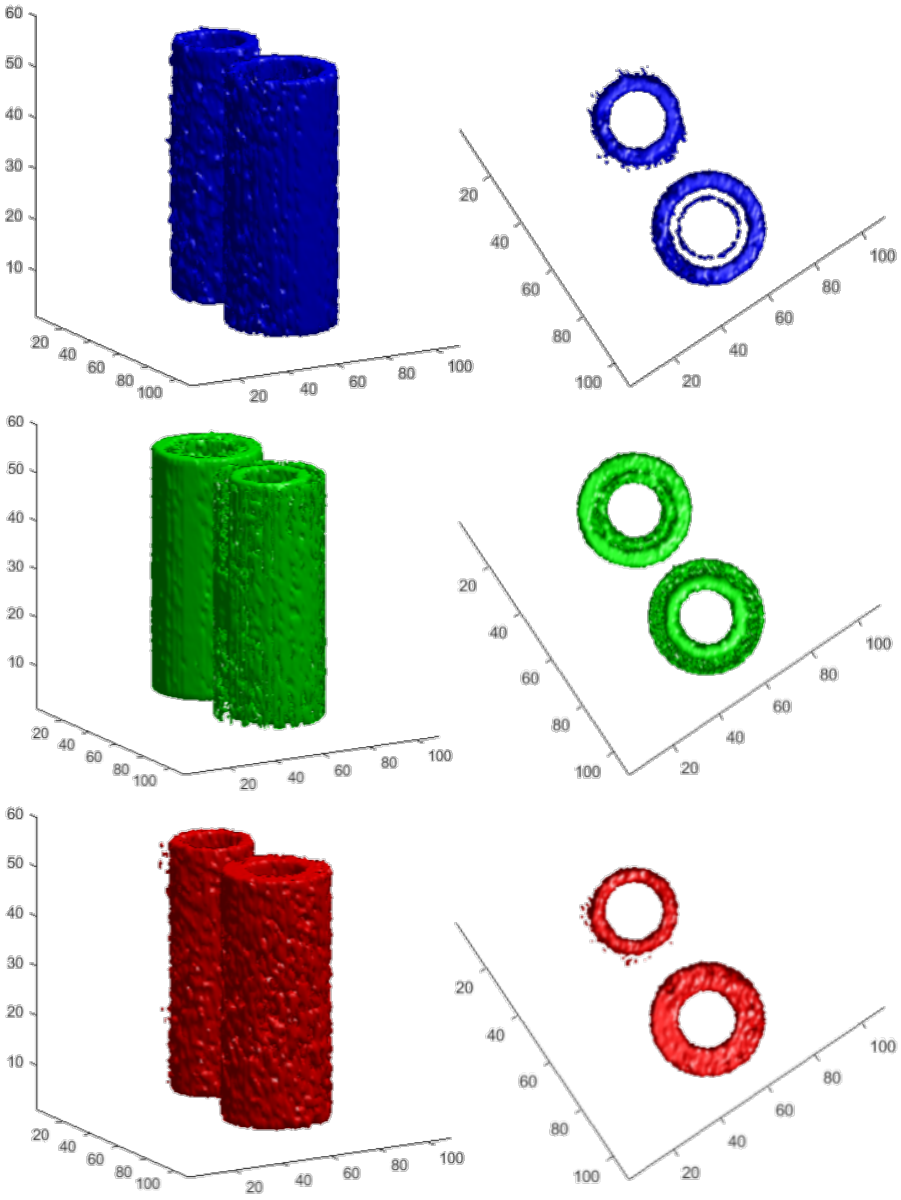

Figure 19. Two different views of the three zero-level set surfaces for iteration=300 for the noisy data case.

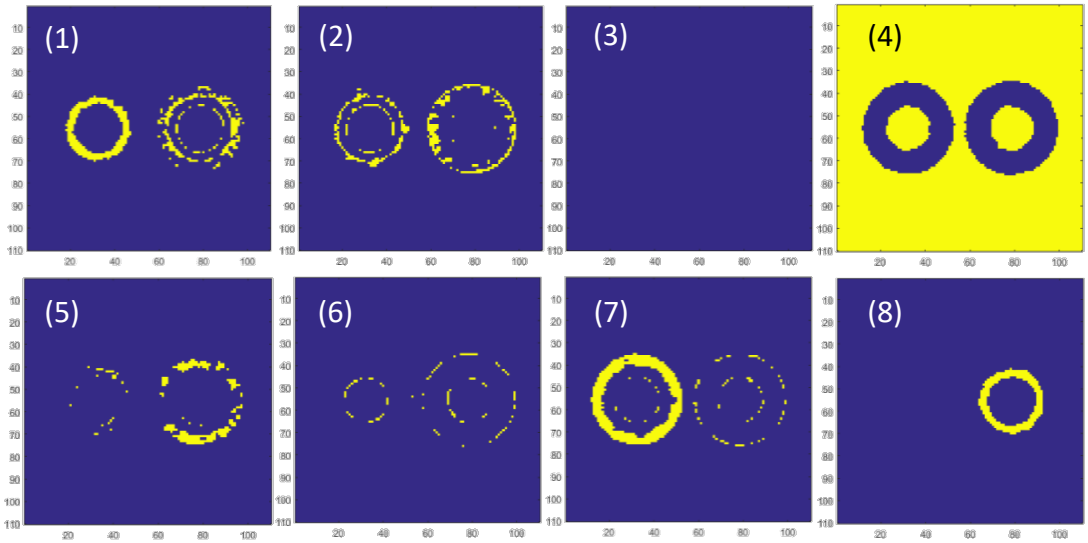

Figure 20. Final segmented regions for the noisy case (prior to post-processing steps) shown along a plane bisecting the object at its midpoint.

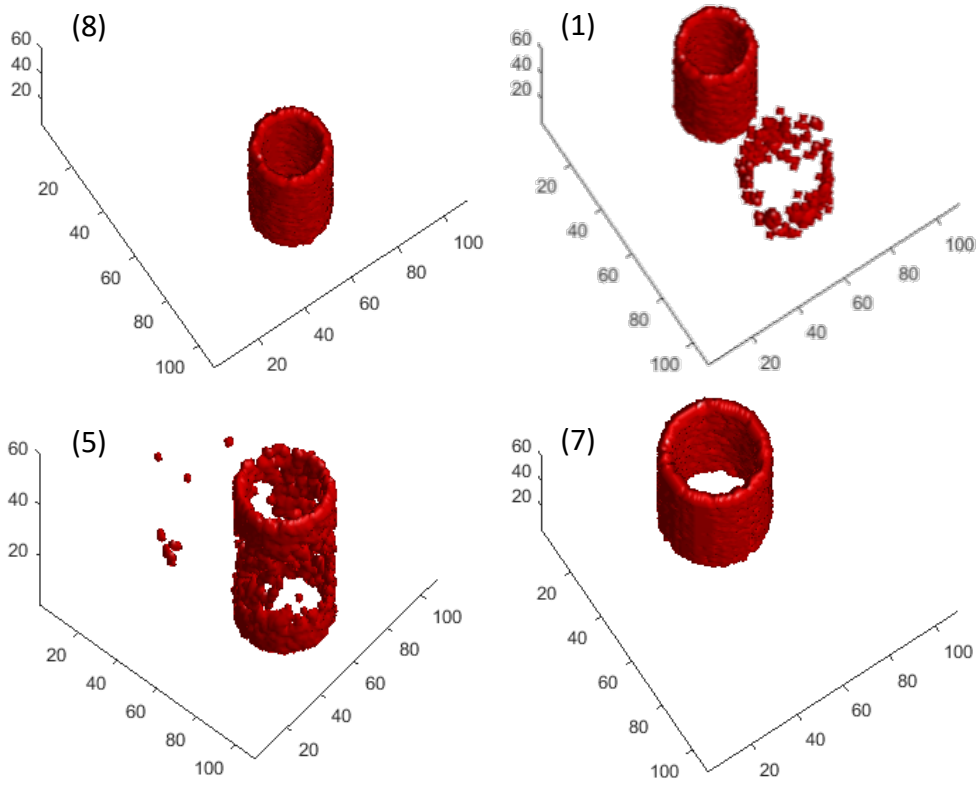

Figure 21. Final segmented surfaces of regions 1, 5, 7, and 8 for the TV-constrained noisy case (after postprocessing steps). The numbering is arbitrary and depends on the initial conditions of the gradient descent.

5. SUMMARY AND FUTURE WORK

The present document reports the implementation of multiphase image segmentation in three dimensions. While this milestone is notable, it is also useful to refocus on future effort. These future efforts will include

- analysis of 3D measured data as soon as it becomes available;
- exploration of strategies to identify the global rather than the local minimum;
- exploration of postprocessing or cleanup procedures to eliminate edge points that are attributed to the wrong region or potentially, exploration of modification of the energy functional and the associated gradient descent to more faithfully attribute boundary points; and
- addition of level sets to simultaneously segment for fission, hydrogen scatter, or other induced reaction images.

6. REFERENCES

[1] D. Mumford, and J. Shah, "Optimal approximation by piecewise smooth functions and associated variational problems," *Comm. Pure Appl. Math.* 42:577–685, 1989.

[2] L. A. Vese and T. F. Chan, "A multiphase Leve Set framework for image segmentation using the Mumford and Shah model," *International Journal of Computer Visio*n 50(3), 271–293, 2002.

[3] T. F. Chan and L. A. Vese "Active contours without edges," *IEEE Trans. on Image Processing,* 10(2): 266–277, 2001.

[4] G. Aubert and L. Vese, "A variational method in image recovery," *SIAM: J. Num. Anal*. 34(5), 1948–1979, 1997.