ARL-TR-10284 ● Feb 2026

Application of Deep Learning Segmentation to Similar Contrast Materials in X-Ray CT Imaging

by Keaton Klaff and Jennifer Sietins

NOTICES

Disclaimers

ARL-TR-10284 ● FEB 2026

Application of Deep Learning Segmentation to Similar Contrast Materials in X-Ray CT Imaging

Keaton Klaff
SURVICE Engineering Company

Jennifer Sietins
DEVCOM Army Research Laboratory

REPORT DOCUMENTATION PAGE

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED	
		START DATE	END DATE
February 2026	Technical Report	1/1/2023	1/31/25

4. TITLE AND SUBTITLE
Application of Deep Learning Segmentation to Similar Contrast Materials in X-Ray CT Imaging

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
W15P7T-19-D-0126		
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**

6. AUTHOR(S)
Keaton Klaff and Jennifer Sietins

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
DEVCOM Army Research Laboratory ATTN: FCDD-RLW-MB Aberdeen Proving Ground, MD 21005	ARL-TR-10284

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

13. SUPPLEMENTARY NOTES
ORCIDs: Keaton Klaff, 0000-0003-0423-1503; Jennifer Sietins, 0000-0003-2547-1653

14. ABSTRACT
Accurate segmentation of X-ray computed tomography (CT) images is critical for quantitative analysis of composite materials. Segmenting samples with low-contrast material phases is often challenging and time-consuming. This study compares two segmentation approaches applied to CT data from an aluminum–silicon carbide (Al-SiC) metal matrix composite specimen. Dual-energy CT images were segmented using Zeiss' Dual Scan Contrast Visualizer (DSCoVr) software, and single-energy CT images were segmented using a 2D U-Net model implemented in Dragonfly Pro. The U-Net method was more precise in capturing phase boundaries and less prone to over-segmentation due to noise, at the expense of greater computational time and user effort. These results highlight a trade-off between accuracy and efficiency and offer guidance for selecting segmentation methods for low-contrast material phases in CT images.

15. SUBJECT TERMS
X-ray computed tomography, image segmentation, machine learning, dual-energy scanning, metal matrix composites

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT	b. ABSTRACT	C. THIS PAGE		
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UU	27

19a. NAME OF RESPONSIBLE PERSON	19b. PHONE NUMBER *(Include area code)*
Keaton Klaff	(520) 691-5252

STANDARD FORM 298 (REV. 5/2020)
Prescribed by ANSI Std. Z39.18

Contents

List of Figures iv

List of Tables iv

1. Introduction 1

2. Procedure 3

2.1 Material 3

2.2 CT Scanning 3

2.3 Bruker CTAnalyzer: Low-Energy Image Segmentation 4

2.4 Zeiss DSCoVr Software Segmentation 4

2.5 Dragonfly Pro ML/DL Segmentation 7

2.5.1 Data Preparation 7

2.5.2 Segmentation and Labeling for Segmentation Wizard 7

2.5.3 Training/Validation Data, Training via Deep Learning Tool 9

2.5.4 Application of DL Model to the Dual Energy–Combined Gray Scale Images 11

3. Results and Discussion 11

4. Conclusions 16

5. References 19

List of Symbols, Abbreviations, and Acronyms 20

Distribution List 21

List of Figures

Figure 1. 2D pixel intensity histogram of both low- and high-energy Al-SiC XCT scans. 5

Figure 2. Detailed view of Al and SiC cluster in the 2D pixel intensity histogram. 6

Figure 3. Zeiss' DSCoVr software interface with both Al-SiC XCT scans loaded. 6

Figure 4. Intensity calibration window in Dragonfly Pro software. 7

Figure 5. Example of paintable regions produced by the Smart-Grid tool. 8

Figure 6. Example of a key frame (top) with model predictions (bottom). 9

Figure 7. Training window of the Deep Learning Tool populated with parameters used to train the 2D U-Net segmentation model. 10

Figure 8. Training Parameters tab within the Deep Learning Tool. 11

Figure 9. Image matrix of the reconstruction results of the DSCoVr software and corresponding segmentation results from Bruker's CTAnalyzer, Zeiss' DSCoVr software, and the Dragonfly DL model. 12

Figure 10. Binary segmentation of SiC phase (red) created in CTAnalyzer is overlaid on corresponding gray scale image (Slice 500). 13

Figure 11. Binary segmentation of SiC phase (red) created in Zeiss' DSCoVr software is overlaid on corresponding gray scale image (Slice 500). 13

Figure 12. Binary segmentation of SiC phase (red) created in Dragonfly Pro is overlaid on corresponding gray scale image (Slice 500). Image on the left is from the low-energy image stack, and image on the right is from the dual-energy image stack. 14

Figure 13. Multi-ROI segmentation of low-energy dataset in Dragonfly Pro viewed in the ZX image plane: Al, SiC, and air phases in blue, orange, and green, respectively. 15

Figure 14. Pop-out window for the 3D Object Counter in ImageJ populated with settings used to analyze each segmented image slice. 15

List of Tables

Table 1. Parameters used during low- and high-energy XCT imaging of Al-SiC sample 3

Table 2. Numerical results of running the 3D Object Counter on segmented image slices from each segmentation method. 16

1. Introduction

Aluminum–silicon carbide (Al-SiC) is a metal matrix composite (MMC) used for various applications within the aerospace and defense industries from structural components to ballistic protection. MMC materials can be difficult to manufacture, and internal cracks, voids, or undesired phase distribution results in poor performance or premature failure of materials and parts. Nondestructive imaging techniques such as X-ray computed tomography (XCT) are a useful tool for identifying defects and inform engineers of manufacturing changes to improve quality and performance.[1,2]

Nondestructive testing via XCT allows researchers to assess material composition, phase distribution, and internal defects that cannot typically be analyzed without physically dissecting samples. Using standard laboratory XCT machines, researchers can image materials without using a synchrotron X-ray source. Lab-based X-ray systems are limited by their size and photon energy production, and therefore cannot achieve the same material penetration, resolution, and contrast of a synchrotron source.[3]

Multi-material samples can be difficult to image, especially when their material phases have similar X-ray attenuation coefficients, and such is the case for Al-SiC MMCs.[4,5] This means during scanning the materials absorb the X-ray photon's energy similarly, which translates to very little contrast between their measured pixel intensities in each radiograph. The lack of contrast is a limiting factor when segmenting the different phases during image processing, and standard methods will not typically yield accurate results.

Dual-energy scanning is a XCT method that uses two different peak photon energy spectra to collect images, taking advantage of different photon interaction mechanisms. Interactions between the material and X-ray photons at low energies are predominately driven by the photoelectric effect, while high-energy interactions are governed by Compton scattering.[6] The given energy at which either interaction occurs depends on the density and atomic number.[7] The difference in photon interaction with the materials during XCT scanning enhances contrast between similarly attenuating materials, such as Al and SiC.

In this experiment, a Zeiss Xradia 520 system was used to collect dual-energy XCT images. Using the Zeiss system, the sample is scanned twice, once at high X-ray voltage and again at low X-ray voltage. Both datasets are reconstructed and imported to Zeiss' Dual Scan Contrast Visualizer (DSCoVr) software, where image data was combined based on a selected value of mixing.[8]

Machine learning (ML) and other artificial intelligence (AI)-based computational models simplify or automate various image processing tasks. ML models use algorithms, such as linear regression, decision trees, and support vector machines, to produce models that make predictions and classifications. This can be applied to tasks such as image segmentation, filtering, and image enhancement. The unique distinction between ML models and other statistical data algorithms is that they are not explicitly programmed to achieve the task they are being used for. An ML model can be trained with supplied data to shape and improve their predictions. Training is accomplished by analyzing ground truth data, which is used to iteratively adjust values/coefficients within its computational algorithm to minimize the difference between predicted and actual values. That difference is quantified by a cost function that calculates an error value. The algorithm seeks to minimize the error of the model and iteratively adjusts its coefficients until either a desired minimum threshold of error is achieved or the model reaches a predetermined stopping point.[9,10]

Deep learning (DL) is a subset of ML that is modeled after the way the human brain learns and interprets data. DL models are often implemented for semantic segmentation of large image datasets for their ability to learn complex patterns from large datasets.[11] This report focuses on the use of a U-Net DL convolutional neural network (CNN) for semantic segmentation of XCT images. The structure of a typical CNN typically consists of multiple computational layers, including convolutional layers for feature extraction, pooling layers for down-sampling and image size reduction, and fully connected layers that produce classifications based on learned features from the previous layers.[12]

The U-Net architecture has a symmetric U-shaped encoder–decoder structure with skip connections. The encoder path comprises several convolutional and pooling layers that extract image features and down-sample input dimensions. The decoder uses convolutional layers, concatenation, and activation functions to progressively up-sample the input from the encoder to its original resolution. Skip connections between the encoder and decoder paths are included to help preserve pixel-based spatial information from being lost during down-sampling, enabling higher-accuracy segmentation.[13]

This report seeks to demonstrate the application of Dragonfly Pro's ML and DL modules to the segmentation of similarly attenuating materials in X-ray CT image data. Results of this method will be compared to those produced by performing dual-energy XCT and segmentation using Zeiss' DSCoVr software.

Utilization of AI algorithms will continue to improve and optimize analyses within the realm of nondestructive testing. Fast and accurate segmentation of MMC materials, such as Al-SiC, provides useful qualitative and quantitative information that aids optimization of material processing parameters. Segmentation results can also be applied to the generation of high-fidelity material response models, which help provide a better understanding of a material's in-situ response.

2. Procedure

2.1 Material

The material used for this study was an MMC consisting of an Al matrix reinforced with spherical SiC particles. The CT scan settings and image data were obtained from an existing dataset. Details of the Al-SiC material's fabrication and processing history were unknown to the authors.

The objective of this work was to compare image segmentation methodologies rather than to analyze the influence of manufacturing processes on microstructure. Therefore, the absence of processing details does not affect the validity of the comparative analysis

2.2 CT Scanning

High- and low-voltage scans were performed on a Zeiss Xradia 520 micro-CT system. The system settings used for both scans can be found in Table 1. Each scan was optimized to ensure sufficient pixel intensity at the detector. Both scans had an 8192.76- × 8192.76-µm field of view.

Table 1. Parameters used during low- and high-energy XCT imaging of Al-SiC sample

XCT imaging parameters	Low-energy scan	High-energy scan
X-ray peak voltage (kV)	50	150
X-ray peak current (µA)	80	67
X-ray peak power (W)	4	10
Geometric magnification	8.503	8.503
Source filter name	Air	HE6
Exposure (s)	6	5
Objective lens	0.4×	0.4×
Detector binning	2	2
Frame averaging	1	1
Source-to-sample distance (mm)	25.025	25.025
Detector-to-sample distance (mm)	187.762	187.762
No. of projections	2501	1901
Pixel resolution (µm)	8.00	8.00

Data from each scan was reconstructed using Zeiss' Scout-and-Scan Reconstructor software. A pre-reconstruction smoothing constant of 0.7 was applied to both datasets. A beam-hardening constant of 0.1 was applied to the high-energy data, and 0.3 to the low-energy data. Both datasets were exported as .txm files and .tiff image stacks.

2.3 Bruker CTAnalyzer: Low-Energy Image Segmentation

To demonstrate the difficulty in segmenting the SiC phase from the Al in the XCT images, a standard global threshold segmentation was performed on the low-energy image stack using Bruker's CTAnalyzer (version 1.18.4).

After importing the stack, a rectangular region of interest (ROI) of 604 × 485 pixels was created around the sample and some surrounding background air. The images inside the ROI had a global threshold filter with lower- and upper-bound values of 205 and 255, respectively. These values were chosen based on visual determination of the best threshold window. An automatic Otsu threshold filter was originally applied but produced poorer segmentation results and was not ultimately used for this analysis.

2.4 Zeiss DSCoVr Software Segmentation

Prior to processing the data using the DSCoVr software (version 14.0.14829.38124), Zeiss recommends performing byte scaling and using a nonlocal means image filter on each dataset, but it is not critical for producing good results from the dual-energy scanning. This was applied to both image datasets.

Once reconstructed, the image sets were imported to the DSCoVr software into the respective high- and low-energy windows. Once imported, the software automatically registers and aligns the image data to one another, and intensity values from both stacks are combined to create a 2D histogram (Figure 1).

The histogram plots the intensity values of every pixel in both scans. Values from the low-energy scan are plotted on the x-axis, and high-energy on the y-axis, with the origin of the plot in the top left. The color values in the histogram represent the frequency in which a specific combination of intensity values from both scans occurs. The clusters seen on the plot correspond to materials present in the scan data.

The slider at the base of the plot changes the ratio used to combine the datasets in the final output. The ratio is determined based on the equation (cosx * low energy + sinx * high energy). The slider changes the angle of the crosshairs seen in the plot. When segmenting the combined dataset, this green line separates which intensity values are set to white in the image stack. The side with the gray half of the line includes all the intensity values that will appear black. All values on the side with the white line include the intensities that will appear white.

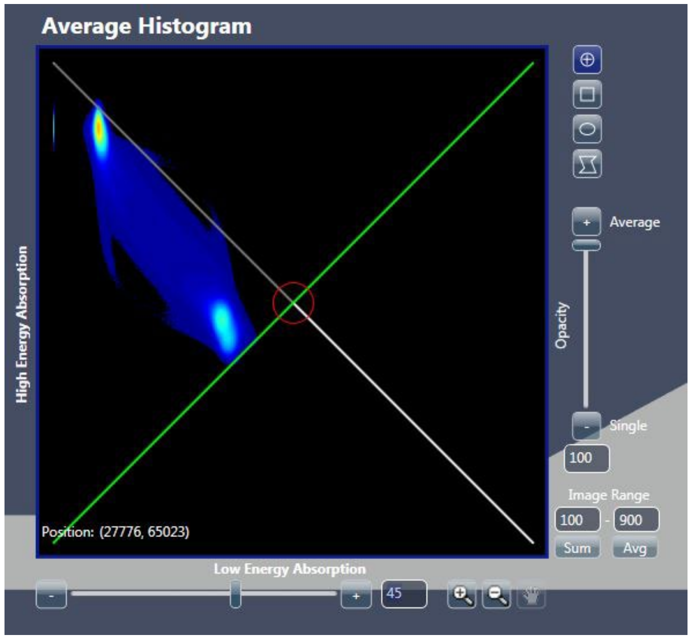

Figure 1. 2D pixel intensity histogram of both low- and high-energy Al-SiC XCT scans.

For the Al-SiC data, the histogram showed one cluster in the top-left corner for the air captured in the scan, and another cluster for the Al and SiC phases (Figure 2). The blue color values express the range of pixel intensities that fall in-between. The location of the crosshairs was set atop the Al-SiC cluster, splitting its two foci, with the intent to segment only the pixels corresponding to the SiC (Figure 3). This selection is subjective and based on multiple attempts to find a segmentation result that visually appears to best fit the data. The software does not provide automatic selection or clear indicators of the best choice, so it is up to the user to distinguish what values are optimal. The results of the placement were updated and shown in the Results window with the "Mask View" box checked. The angle of the crosshairs was left at 45°, which corresponds to equal mixing of the high- and low-energy scans. The resulting segmentation and gray scale version of the combined dataset were both exported as separate image stacks.

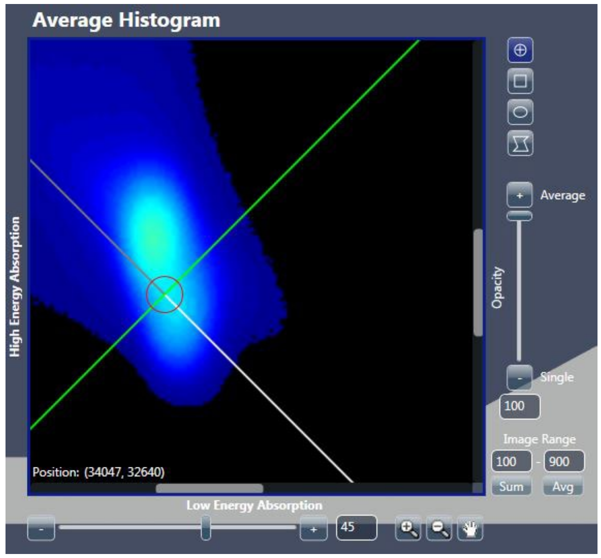

Figure 2. Detailed view of Al and SiC cluster in the 2D pixel intensity histogram.

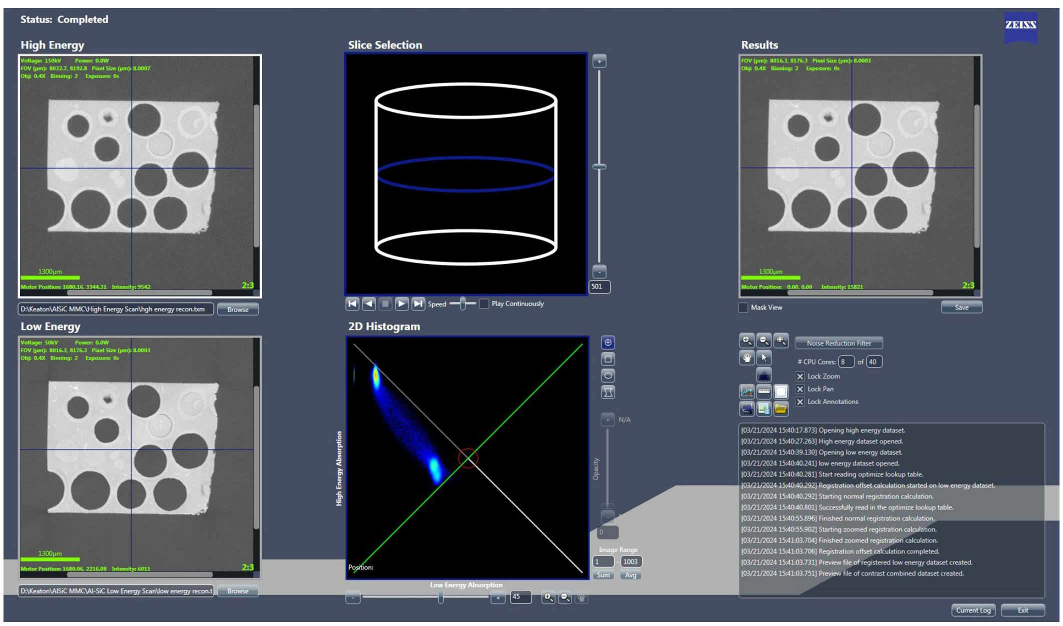

Figure 3. Zeiss' DSCoVr software interface with both Al-SiC XCT scans loaded.

2.5 Dragonfly Pro ML/DL Segmentation

2.5.1 Data Preparation

The low-energy XCT data was imported to the Dragonfly Pro software as a .txm file, which is the file format produced by Zeiss' reconstruction software. Before any segmentation or training began, the image data was converted to Century Units via intensity calibration. In the calibration window (Figure 4), intensity values for both the background and foreground (material) in the images can be selected by either clicking on the respective area in the XCT images or moving the background/material lines to the desired location on the intensity histogram. Raw intensity values of 4973.964 and 11775.433 were selected for the background and foreground, respectively. The dataset was calibrated to the new units, and the intensity histogram was normalized to values between 0 and 100.

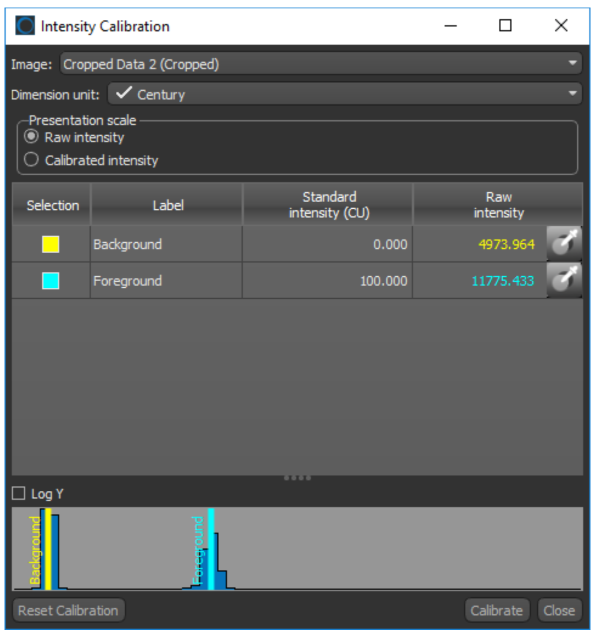

Figure 4. Intensity calibration window in Dragonfly Pro software.

2.5.2 Segmentation and Labeling for Segmentation Wizard

Labeled image data is required to train the DL model. Labeled images act as the ground truth data from which the DL algorithm learns how to accurately segment an unlabeled image stack. In Dragonfly, the Segmentation Wizard was used to quickly label images, generate several segmentation models, and perform training.

First, several slices containing key characteristic features of each material phase were selected as "Key Frames." In the "Classes" section, three classes were created: Al, SiC, and background. Within the key frames, pixels were labeled to each class using the ROI tools. The histogram, ROI painter, and Smart-Grid tools were primarily used for this step. The background was labeled using the histogram tool and its "Lower Otsu" threshold button. This calculates the threshold value that separates the material from the background based on the intensity histogram's pixel values. Since the gray scale values of the Al and SiC are very similar, there is no single threshold value that can be used to separate the phases using the histogram. Instead, the ROI painter and Smart-Grid tools were used. The ROI painter treats the cursor like a paint brush, labeling pixels in a circular- or square-shaped area. The Smart-Grid breaks the image into regions that are calculated based on local variations in pixel gray scale values, and its settings can alter the size of the predicted regions. Figure 5 demonstrates an example of paintable regions produced by the Smart-Grid tool. The Smart-Grid tool was used together with the ROI painter to quickly label regions marked on the image. When a region is selected with the painter tool, the area is then labeled as the selected class.

Figure 5. Example of paintable regions produced by the Smart-Grid tool.

Additionally, the histogram and ROI painter were used together to quickly label pixels. Thresholding values were set to highlight specific gray scale values on the key frame, then with the ROI painter, only the highlighted pixels could be labeled.

Accuracy of labeling is critical for producing an accurate model. Incorrectly labeled pixels supplied during training will diminish the accuracy of the segmentation models. For this reason, great care was taken while labeling each key frame.

Four key frames were sparsely labeled with pixels from each class. Training was initiated on three models using the "sparse data" preset in the Model Generation window. This preset generates two Random Forest ML models and one 2D U-Net model. Each model was trained using class data. Once training was finished, a labeled prediction of the most recent key frame was populated from each model (Figure 6). Model predictions can be applied to new key frames, limiting the need for additional manual labeling. Predictions from each model were not producing accurate segmentations at this stage.

Eight fully labeled key frames were produced by applying model predictions to each new key frame, then errors were manually corrected using the ROI tools. These frames were used to further train the models. Following training, the 2D U-Net model produced the most accurate predictions and was exported/published so it could be used in the Deep Learning Tool in the Dragonfly main window.

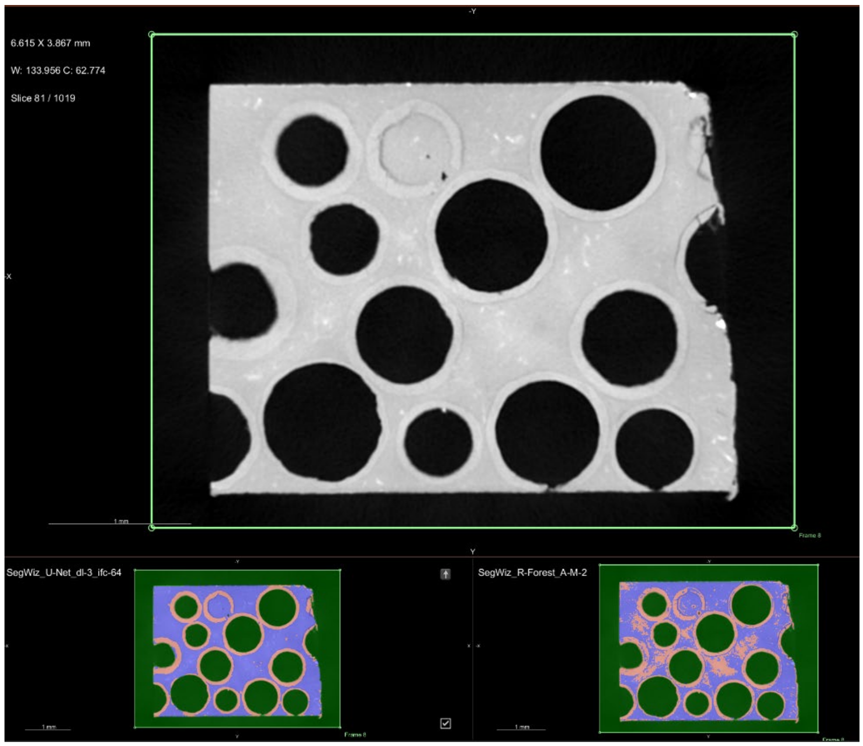

Figure 6. Example of a key frame (top) with model predictions (bottom).

2.5.3 Training/Validation Data, Training via Deep Learning Tool

Within the main Dragonfly window, two mask ROIs were generated to define the slices for a training dataset and a validation dataset. These were labeled "training mask" and "validation mask," respectively. Each set was populated with different images from the original stack.

These were defined by creating ROI masks over desired slices from the image stack. The training set covered 55 images spread out between slice numbers 454 to 588 (of 1004). The validation set included 20 images between slices 474 to 643. No images were shared between sets to mitigate the risk of overfitting the model to the training data.

AI image segmentation was used with Dragonfly's "Segment with AI" tool to segment both the slices marked as the training set and as the validation set. The previously trained U-Net model was applied, and a new multi-ROI was generated with the “Background,” “Al,” and “SiC” classes for each set. To clarify, one multi-ROI was made for the training set and another for the validation set.

The sparsely trained U-Net model did not correctly label all the images, so errors in the multi-ROIs were manually corrected using the same ROI tools as in the ML Segmentation Wizard. Once all labels in both datasets were corrected, they were ready to be supplied for training a new model within the Deep Learning Tool**.**

The Deep Learning Tool was opened to generate and begin training a new model. From the Deep Learning Tool window, a new 2D U-Net segmentation model with a depth level of 5 and an initial filter count of 64 was created. Within the training window, the full image stack was set as the “Input,” the training multi-ROI was set as the “Output,” and the training mask was set as the “Mask” (Figure 7). The “Data Augmentation” section shows the parameters used to modify the supplied images to supply the model with additional training data, which helps prevent over-fitting of the model to the training data. In this case, the standard augmentations were used (horizontal flip, vertical flip, rotations, shearing, scaling). The brightness augmentation was also selected with a range of 0.9 to 1.1. In the “Use Validation” section, the validation dataset and mask were selected.

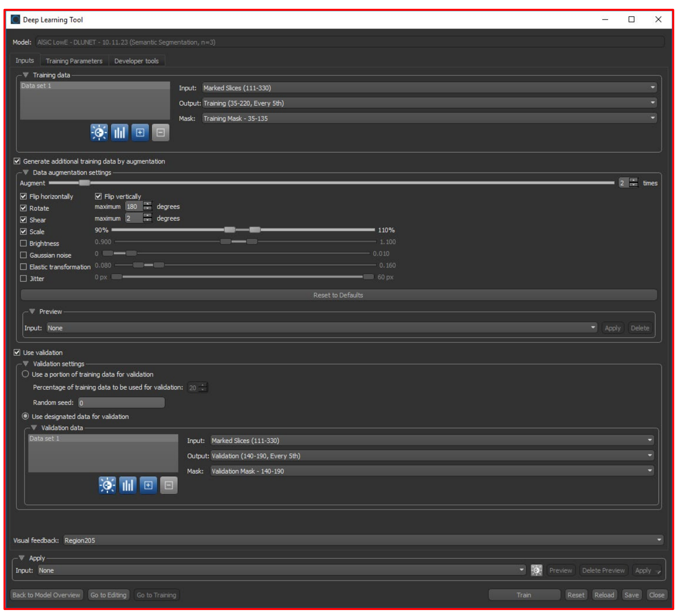

Figure 7. Training window of the Deep Learning Tool populated with parameters used to train the 2D U-Net segmentation model.

Prior to beginning the training, a rectangle was drawn on a slice of the original dataset. Within the DL training window, that region was selected in the Visual Feedback drop-down menu. During training, the image captured by the rectangle shows the model’s segmentation prediction for that area after each epoch.

Additional settings were available in the “Training Parameters” tab, but no changes were made to the standard inputs (Figure 8).

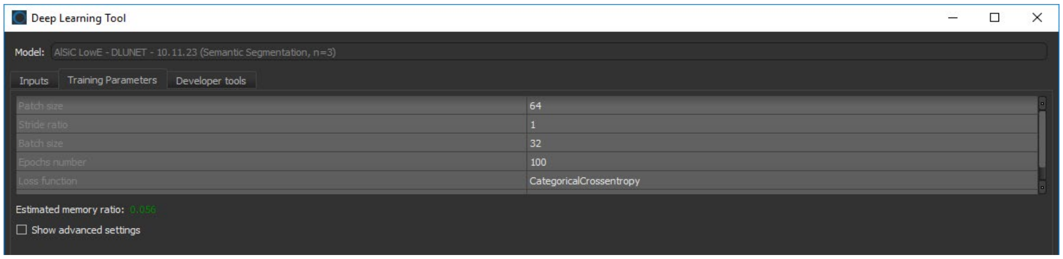

Figure 8. Training Parameters tab within the Deep Learning Tool.

Once training was complete, the segmentation was visually assessed by first applying the segmentation to a single preview slice and then to the full image stack. The segmentation was generated as a multi-ROI, from which the SiC class was copied to a standard ROI and exported as a binary image stack.

2.5.4 Application of DL Model to the Dual Energy–Combined Gray Scale Images

To further test the results of the DL model produced in Dragonfly and explore the reproducibility of segmentation results, the combined gray scale dataset created in DSCoVr was imported to a new session.

Before segmentation was applied, an intensity calibration was done on the image data to match the intensity scale of the low-energy scan that the model was trained on (from 0 to 1). Once calibrated, the Deep Learning Tool was used to load the previously trained U-Net model and perform the segmentation on the full image stack. No additional training was done on the model with this new dataset.

The SiC class of the newly generated multi-ROI was copied to an ROI and exported as a binary image stack.

3. Results and Discussion

Following application of the segmentation tools from Bruker, Zeiss, and Dragonfly Pro, five representative slices from the Al-SiC segmented image stacks were selected for visual comparison. The selected gray scale images from each dataset and their respective segmentations are shown in Figure 9.

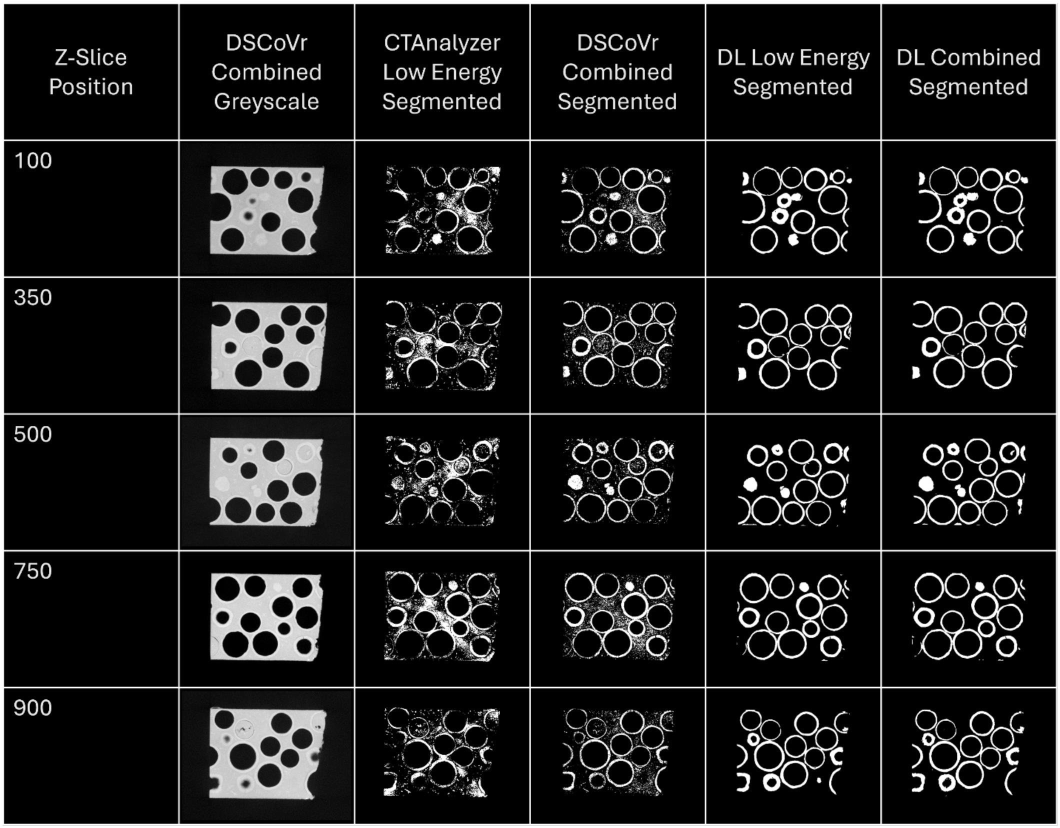

Figure 9. Image matrix of the reconstruction results of the DSCoVr software and corresponding segmentation results from Bruker's CTAnalyzer, Zeiss' DSCoVr software, and the Dragonfly DL model.

Visual inspection of outputs from each method reveals notable differences. Each segmentation identified the SiC phase with varied success, but no segmentation method produced perfect results.

The CTAnalyzer method shows how poorly an automatic thresholding method performs on low-contrast images. The SiC is under-segmented from the Al phase, and lots of noise particles are present throughout the image stack (Figure 10).

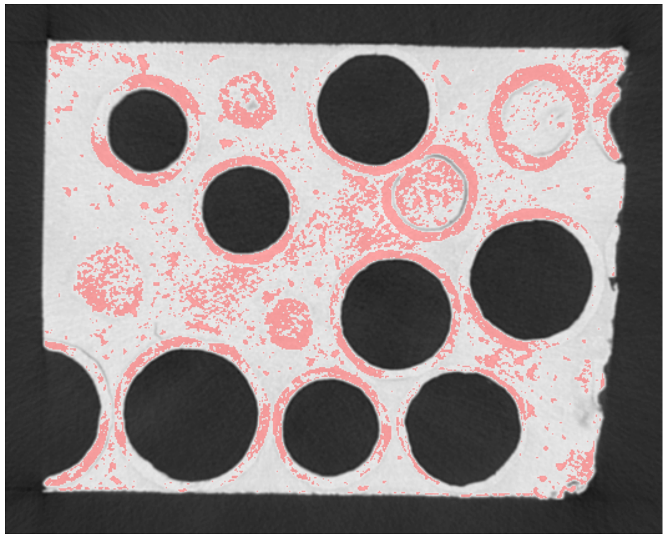

Figure 10. Binary segmentation of SiC phase (red) created in CTAnalyzer is overlaid on corresponding gray scale image (Slice 500).

The DSCoVr segmentation also captured a large amount of noise and over-segmented most of the SiC phase, which causes each spherical particle to appear as fragmented circles rather than solid material in 2D (Figure 11). There are also instances where individual particles of SiC were not captured at all.

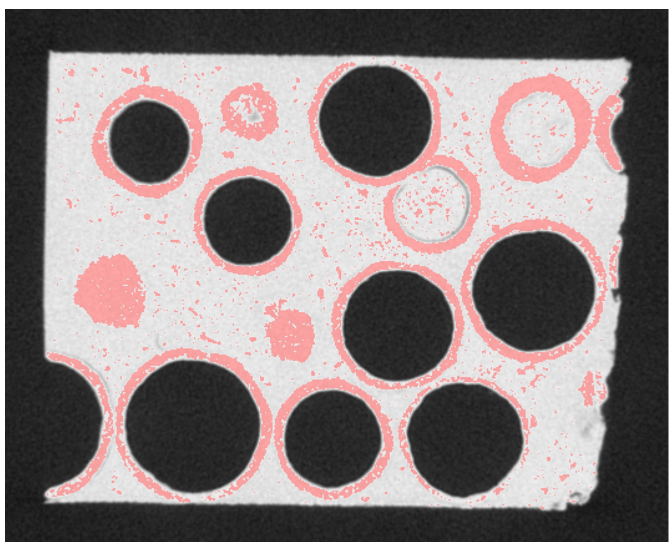

Figure 11. Binary segmentation of SiC phase (red) created in Zeiss' DSCoVr software is overlaid on corresponding gray scale image (Slice 500).

The Dragonfly DL segmentation applied to the low-energy dataset also produced some labeling errors. Little to no noise particles were segmented within the image stack, but some areas of the edge of the sample were incorrectly labeled as SiC. There are also some instances where the SiC particles were incorrectly labeled near their edges, leaving a small amount uncaptured by the segmentation model (Figure 12).

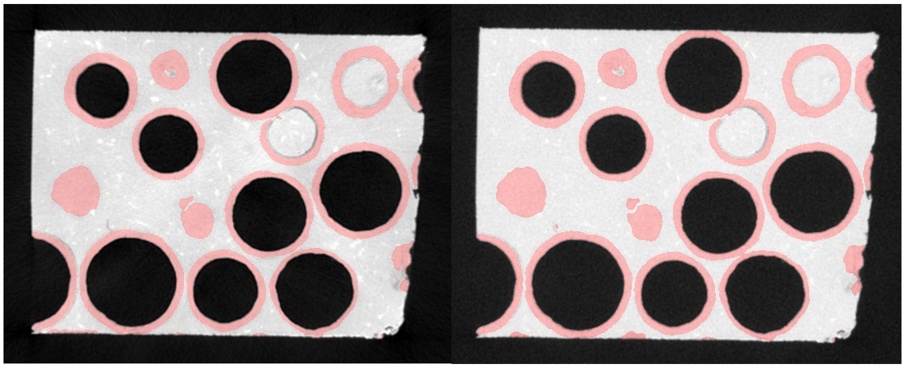

Figure 12. Binary segmentation of SiC phase (red) created in Dragonfly Pro is overlaid on corresponding gray scale image (Slice 500). Image on the left is from the low-energy image stack, and image on the right is from the dual-energy image stack.

Similar errors were seen when the same U-Net segmentation model in Dragonfly was applied to the dual-energy dataset. There are only slight differences between the segmentation results applied to the low- and dual-energy data. This is a positive result, since the segmentation model had not been trained on this "new" dataset. The geometric features in both datasets are identical, with the only differentiating factors being the gray scale intensity histogram and the enhanced contrast in the dual-energy results. Despite the lack of additional training, the model was able to segment the image stack similarly to the low-energy data. A pixel count of the white pixels in both the low- and dual-energy DL segmented stacks gives values of 27048153 and 27450035 pixels, respectively. This is a difference of 401882, or only 1.48% more pixels for the dual-energy stack.

When viewing the DL segmented images from a different plane, there are discontinuities that indicate segmentation in 3D is not completely accurate (Figure 13). This is likely because the model used was a 2D U-Net, which lacks sufficient spatial comprehension required to segment a 3D dataset without inaccuracies between image slices. A 2.5D or 3D model could have been used, but at the cost of higher computation times.

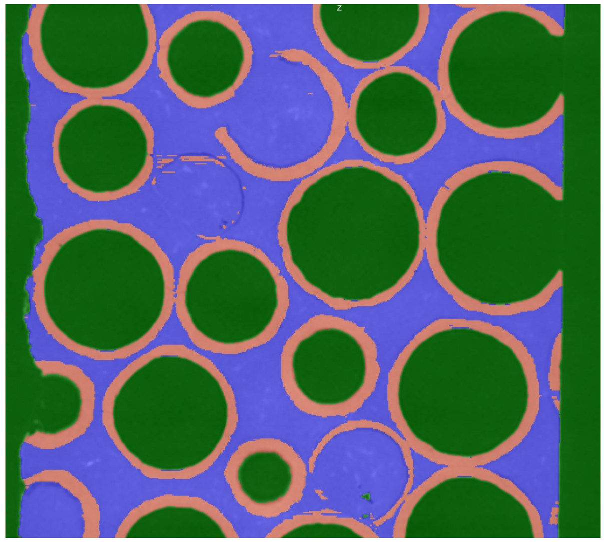

Figure 13. Multi-ROI segmentation of low-energy dataset in Dragonfly Pro viewed in the ZX image plane: Al, SiC, and air phases in blue, orange, and green, respectively.

As a simple method to quantitatively compare the segmentation results, the 3D Object Counter within FIJI/ImageJ was used to count the number of objects in each of the representative segmented images. The settings used in the 3D Object Counter pop-out window are shown in Figure 14. To avoid counting the smallest noise particles in each dataset, a minimum size filter of 10 pixels was used. The DSCoVr combined gray scale images were visually inspected, and each individual instance of SiC material was labeled with a number for each image slice. These values were used as a standard for comparison with the number objects in the segmented images. Table 2 contains the numerical results of the object counter analyses.

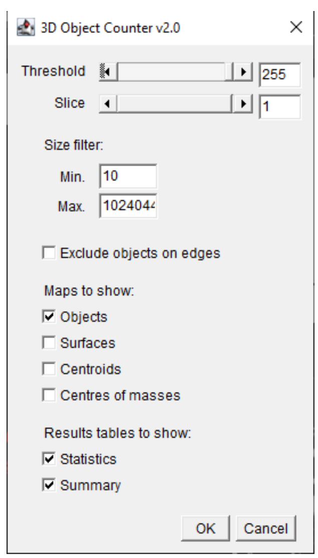

Figure 14. Pop-out window for the 3D Object Counter in ImageJ populated with settings used to analyze each segmented image slice.

Table 2. Numerical results of running the 3D Object Counter on segmented image slices from each segmentation method.

Image slice	Visual count	Low Energy - CTAn	Dual Energy - DSCoVr	Low Energy - DL	Dual Energy - DL
100	17	159	121	15	14
350	14	188	115	12	12
500	19	178	85	20	19
750	17	177	167	18	18
900	15	184	153	14	15

The results of the object counter express several key issues with each of the segmentation methods. First, in the case of the standard CTAnalyzer and DSCoVr segmentation, noise particles are a major source of error. The DSCoVr results have a slightly lower number of noise particles (reduction - average 26.4%, max 52.3%, min 5.65%). Calculation of SiC phase volume fraction or to count the total number of SiC particles from either method would yield inaccurate results. Efforts could be made to filter out these particles through morphological or despeckling operations, but since the regions of SiC are also over-segmented and fragmented, a portion of correctly labeled pixels would also be removed.

For the DL segmented images, the object counter shows that there are some images where the number of objects fall both above and below the visual count. This discrepancy is likely due to a combination of touching of adjacent SiC particles in the binary images and the incorrect labeling of pixels at the edge of the sample.

4. Conclusions

In conclusion, two methods of segmenting Al and SiC phases were explored and the benefits and drawbacks of each technique are summarized below.

Dual-energy CT/ Zeiss' DSCoVr: In comparison to performing a standard XCT scan and reconstruction, the use of dual-energy scanning when imaging the Al-SiC sample reduced image noise and improved the contrast between the two phases. Using the DSCoVr software improved segmentation results over a standard CT scan and global threshold, but overall, the segmentation was still inaccurate with a significant number of erroneous noise particles. Segmentation of the SiC phase using the DSCoVr software is more subjective overall, since it is up to the user to manually determine the best result. Especially in this case, for the Al-SiC sample there was not a clear delineation between materials visible in the 2D pixel intensity histogram. It may be possible to achieve better segmentation results, but it is up to the user to guess and check selected dataset mixing and threshold values to achieve the best result.

This method required extra time up front to set up and perform two XCT scans and two reconstructions. Overall, less processing was needed in Zeiss' DSCoVr software compared to Dragonfly's DL module, but it resulted in noisier segmented images and over-segmentation of SiC particles. To perform quantitative analysis on the segmented results, additional image processing would likely be required. The segmentation results are subject to user bias, since there is some subjectivity in the selection of data mixing and thresholding values. Some workarounds are available to yield better results, including adding material fiducial markers (SiC, Al) to the XCT scan volume (as suggested by Zeiss). This can help improve results by clarifying the separation of Al and SiC foci in the 2D pixel intensity histogram, making it easier to properly select dataset mixing and threshold values.

Dragonfly DL: This method saves time up front by only requiring a single XCT scan and reconstruction, but significant time and manual input are needed to train the segmentation model. Results of the segmentation model are heavily dependent on the quality and quantity of labeled images used for training. In this study, a sufficiently accurate model was able to be trained using only a few labeled images, but inaccuracies in the segmentation still exist. A 2D U-Net model was used to save computational time over a 2.5D or 3D model architecture, but at the cost of spatial inaccuracies between image slices.

When the model was applied to the DE image stack without additional training, a similarly accurate segmentation was achieved, showing there is some level of generalization present in the model (images were identical geometrically, but with differing pixel intensities).

Dragonfly's DL model produced results with much less noise and fewer instances of over-segmentation than with the DSCoVr software. Some time and effort are required to supply a model with labeled images and to perform training, but as shown in this study, a model can produce accurate results with the input of only a few labeled images. With additional labeled data and more training, the accuracy of the DL model segmentation should gradually improve. Different training or model architecture parameters may result in better detail extraction and more accurate segmentation. 2D U-Net was used instead of 2.5D or 3D to prioritize speed of model training, but higher dimensional architecture would increase spatial accuracy. User bias can be introduced to the model through the process of labeling data for training but can be mitigated by using the automated segmentation tools supplied in the Dragonfly software (local Otsu ROI painting, automatic thresholding via intensity histogram, Smart-Grid tool, etc.). Additionally, the model's performance can be improved by supplying/training on XCT image data from new Al-SiC samples. In cases where multiple Al-SiC specimens need to be XCT-scanned, training the model with labeled data from multiple datasets will help prevent overfitting of the model to a single dataset. This could help improve segmentation accuracy when applying the model to new XCT datasets with minimal user input.

5. References

1. Mortensen A, Llorca J. Metal matrix composites. Annu Rev Mater Res. 2010;40:243–270.
2. Brunke O et al. Comparison between x-ray tube-based and synchrotron radiation-based μCT. Dev X-Ray Tomogr VI (Proc SPIE). 2008;7078: 70780U.
3. Fornaro J et al. Dual- and multi-energy CT: approach to functional imaging. Insights Imaging. 2011;2:149–159.
4. Hubbell JH, Seltzer SM. Tables of x-ray mass attenuation coefficients and mass energy-absorption coefficients (version 1.4). National Institute of Standards and Technology; 2004. NIST Standard Reference Database 126.
5. Große Hokamp N et al. Technical background of a novel detector-based approach to dual-energy computed tomography. Diagn Interv Radiol. 2020;26(1):68–71.
6. Huda W, Abrahams RB. Radiographic techniques, contrast, and noise in x-ray imaging. Am J Roentgenol. 2015;204(2):W126–W131.
7. Nachtrab F et al. Quantitative material analysis by dual-energy computed tomography for industrial NDT applications. Nucl Instrum Meth Phys Res A. 2011;633(Suppl):S159–S162.
8. Faouzi J, Colliot O. Classic machine learning methods. arXiv. 2023. arXiv:2310.11470.
9. Seo H et al. Machine learning techniques for biomedical image segmentation: an overview of technical aspects and introduction to state-of-art applications. Med Phys. 2020;47(5):e148–e167.
10. Ahmad IS et al. Deep learning models for CT image classification. Quant Imag Med Surg. 2024.
11. Voulodimos A et al. Deep learning for computer vision: a brief review. Comput Intell Neurosci. 2018;2018:7068349.
12. Ronneberger O et al. U-Net: convolutional networks for biomedical image segmentation. arXiv. 2015. arXiv:1505.04597.
13. Wang J et al. A comprehensive review of U-Net and its variants: advances and applications in medical image segmentation. arXiv. 2025. arXiv:2502.06895.

List of Symbols, Abbreviations, and Acronyms

2/3D	Two-/Three-Dimensional
AI	Artificial Intelligence
Al	Aluminum
ARL	Army Research Laboratory
CNN	Convolutional Neural Network
DEVCOM	U.S. Army Combat Capabilities Development Command
DL	Deep Learning
DSCoVr	Dual Scan Contrast Visualizer
HE	High Energy
LE	Low Energy
ML	Machine Learning
MMC	Metal Matrix Composite
ROI	Region of Interest
SiC	Silicon Carbide
XCT	X-Ray Computed Tomography

1 (PDF)	DEFENSE TECHNICAL INFORMATION CTR DTIC OCA
1 (PDF)	DEVCOM ARL FCDD RLB CI TECH LIB