

# Machine Learning Enabled Sensor Fusion for In Situ Defect Detection in Laser Powder Bed Fusion



Zackary Snow  
Luke Scime  
Amir Koushyar Ziabari  
Brian Fisher  
Vincent Paquit

**CRADA Final Report**

**January 2024**

**NFE-21-08807**



## DOCUMENT AVAILABILITY

**Online Access:** US Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free via <https://www.osti.gov>.

The public may also search the National Technical Information Service's [National Technical Reports Library \(NTRL\)](#) for reports not available in digital format.

DOE and DOE contractors should contact DOE's Office of Scientific and Technical Information (OSTI) for reports not currently available in digital format:

US Department of Energy  
Office of Scientific and Technical Information  
PO Box 62  
Oak Ridge, TN 37831-0062  
**Telephone:** (865) 576-8401  
**Fax:** (865) 576-5728  
**Email:** [reports@osti.gov](mailto:reports@osti.gov)  
**Website:** [www.osti.gov](http://www.osti.gov)

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.

Manufacturing Science Division

**MACHINE LEARNING ENABLED SENSOR FUSION FOR IN SITU DEFECT  
DETECTION IN LASER POWDER BED FUSION**

Zackary Snow  
Luke Scime  
Amir Koushyar Ziabari  
Brian Fisher  
Vincent Paquit

December 2023

CRADA Report No. NFE-21-08807

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





## CONTENTS

CONTENTS .....	iii
LIST OF FIGURES .....	iv
LIST OF TABLES .....	v
ABBREVIATIONS .....	vi
ABSTRACT.....	1
1. INTRODUCTION .....	1
2. RESULTS .....	4
2.1 PHASE I—TASK 1: PRINTING TEST ARTIFACTS .....	4
2.2 PHASE I—TASK 2: X-RAY COMPUTED TOMOGRAPHY OF TEST ARTIFACTS.....	6
2.3 PHASE I—TASK 3: SENSOR FUSION AND FLAW DETECTION .....	7
2.4 PHASE II—TASK 1: INDUSTRY-RELEVANT PERFORMANCE.....	14
2.5 PHASE II—TASK 2: FACILITATING HUMAN INTERPRETABILITY .....	21
2.6 X-RAY COMPUTED TOMOGRAPHY DEVELOPMENTS UNDER PHASE II.....	22
3. CONCLUSION.....	26
4. REFERENCES .....	27

## LIST OF FIGURES

Figure 1. A representative in situ sensor data stack from an L-PBF printer consisting of information from three different sensing modalities and three different layers. ....	2
Figure 2. Post-build image of RTRC Build #0849. ....	4
Figure 3. Summary of process monitoring data collected during the current study. ....	5
Figure 4. Composite image of registered XCT scan data for RTRC Build #0849 showing the spatial distribution of internal flaws across the build plate.....	7
Figure 5. Pixel classes available for training in Peregrine. ....	8
Figure 6. DSCNN segmentation results for Layer #95 for the five coupons just below the spatter generators (P43-47) from RTRC Build #0849.....	9
Figure 7. Estimate of the Phase I segmentation model’s PoD for a single coupon (P43) from RTRC Build #0849. ....	12
Figure 8. PoD curve for the INDE system, with successful detections shown by green circles and unsuccessful detections depicted with red crosses.....	13
Figure 9. Representative sensor data collected from Layer #556 of the build performed at RTRC.....	15
Figure 10. XCT Supercomposite of layer 515-610, summarizing the porosity structure throughout the gauge regions of the fatigue coupon blanks. ....	17
Figure 11. Segmentation results for Layer #556 from the Phase II build using (a) the Phase I segmentation model and (b) the DMSCNN trained for Phase II. ....	18
Figure 12. In situ detection hit rate as a function of XCT flaw size for the Phase II network applied to fatigue coupon gauge sections only. ....	19
Figure 13. Example of (a) an internal flaw detected in the in situ data and (b) its corresponding indication in the XCT data. ....	20
Figure 14. Point cloud representations of the intended geometry (orange) of a subset of the IRC and the registered XCT data (purple) of the final printed component. ....	20
Figure 15. Receiver-operator curves and precision-recall curves for the validation dataset. ....	21
Figure 16. Example XCT slices using various reconstruction algorithms: model-based iterative reconstruction (MBIR), Feldkamp–Davis–Kress (FDK), and the CAD-DLMBIR method used here. ....	22
Figure 17. Reconstructions of the pentagon geometry using CAD-DLMBIR trained on a different system with other geometries and materials. ....	23
Figure 18. Qualitative demonstration of Simurgh results on 10× reduction in the number of views for a sample that was more than 30% thicker than the samples used during training. ....	23
Figure 19. A single slice from the reconstructed volume using standard and BHCN-MBIR reconstructions. ....	24
Figure 20. A sample slice from the reconstruction of a cut section of the IRC, along with its segmentation and an expanded view of a region of interest within the slice.....	25

## LIST OF TABLES

Table 1. Build strategies employed to produce coupons suitable for training, validation, and testing of the DL network in Peregrine .....	5
Table 2. Breakdown of the number of XCT coupons scanned from each of the four builds .....	6
Table 3. Testing pixel-wise confusion matrix by predicted classes for the five XCT coupons just below the spatter generators (P43-47) .....	10
Table 4. Testing pixel-wise confusion matrix by ground-truth annotations for the five XCT coupons just below the spatter generators (P43-47) .....	11
Table 5. Automated segmentation results generated for the validation data of the Phase II build.....	21

## ABBREVIATIONS

AM	additive manufacturing
BHCN-MBIR	beam hardening correction network–model-based iterative reconstruction
CAD-DLMBIR	computer-aided design–deep learning model-based iterative reconstruction
CED	circular equivalent diameter
DL	deep learning
DSCNN	dynamic segmentation convolutional neural network
DMSCNN	dynamic multilabel segmentation convolutional neural network
FDK	Feldkamp–Davis–Kress
INDE	in situ nondestructive evaluation
IRC	industrially relevant component
L-PBF	laser powder bed fusion
MBIR	model-based iterative reconstruction
MDF	Manufacturing Demonstration Facility
MPM	melt pool monitoring
NIR	near-infrared
ORNL	Oak Ridge National Laboratory
OT	optical tomography
PoD	probability of detection
RTRC	RTX Technology Research Center
SED	spherical equivalent diameter
XCT	x-ray computed tomography

## ABSTRACT

Laser powder bed fusion (L-PBF) additive manufacturing is among the metal 3D printing technologies most broadly adopted by the manufacturing industry. The current industry qualification paradigm for critical-application L-PBF parts relies heavily on expensive nondestructive inspection techniques such as x-ray computed tomography (XCT), which significantly limits the use cases of L-PBF. In situ monitoring of the process promises a less expensive alternative to ex situ testing, but existing sensor technologies and data analysis techniques struggle to detect subsurface flaws (e.g., porosity and cracking) on production-scale L-PBF printers. RTX Technology Research Center has licensed Oak Ridge National Laboratory's Peregrine software package—a printer- and camera-agnostic data analytics tool designed specifically for detecting process anomalies using in situ data collected during powder bed printing. The goal of this project was to feed temporally rich, multimodal sensor data, including visible light, integrated near-infrared, and spatially mapped coaxial melt pool thermal emission data, into Peregrine to enable the detection of subsurface flaws. XCT data was used as ground-truth training data to allow Peregrine's deep learning algorithms to recognize anomalies in these complex data streams in test artifacts and industrially relevant geometries. Completion of this program has seen the successful implementation of multimodal, multilayer sensor data footprints for training machine learning models in Peregrine. Flaws detected in XCT data have been successfully detected directly from this in situ data footprint, and initial analyses of the in situ probability of detection has been conducted, showing performance levels commensurate with traditional nondestructive evaluation methods. The in situ monitoring methodology was then applied to an industrially relevant component that was using post-build nondestructive evaluation, highlighting the utility of the proposed method for hard-to-inspect additive manufacturing components. As a direct result of this program, two journal manuscripts [1], [2] have been published in *Additive Manufacturing*, with additional manuscripts planned following program completion.

## 1. INTRODUCTION

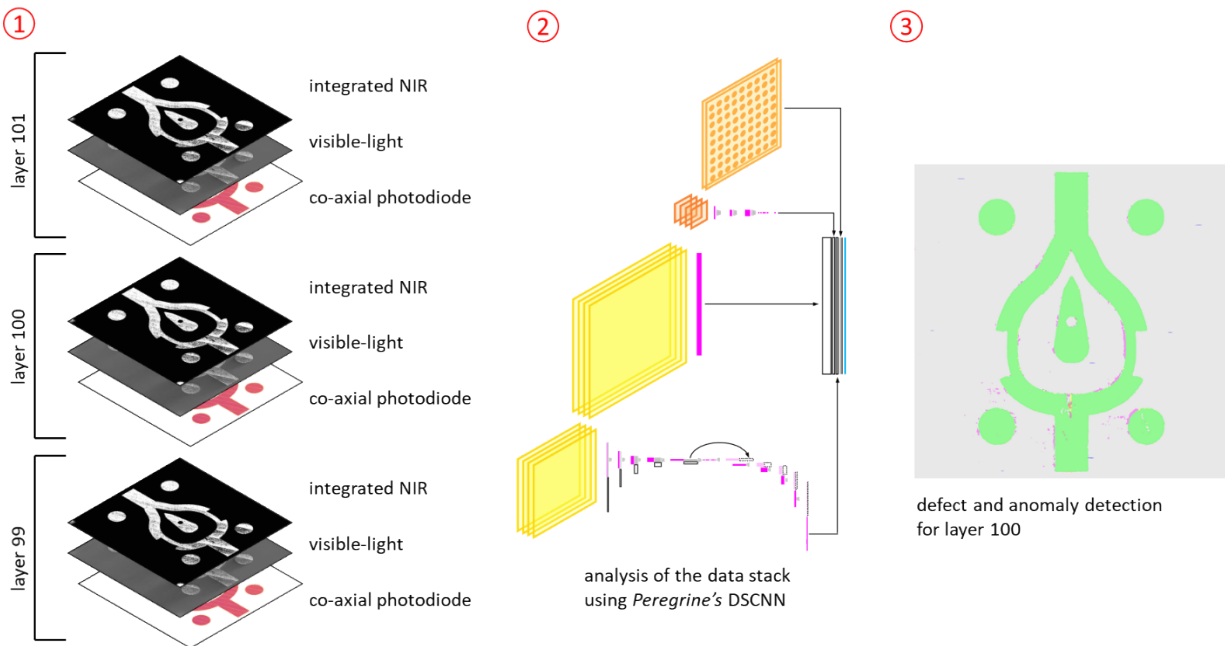
Laser powder bed fusion (L-PBF) additive manufacturing (AM) is among the metal 3D printing technologies most broadly adopted by the manufacturing industry. The current industry qualification paradigm for critical-application L-PBF parts relies heavily on expensive nondestructive inspection techniques such as x-ray computed tomography (XCT), which significantly limits the use cases of L-PBF. In situ monitoring of the printing process promises a less expensive alternative to ex situ testing, but existing sensor technologies and data analysis techniques struggle to detect subsurface flaws (e.g., porosity and cracking) on production-scale L-PBF printers. RTX Technology Research Center (RTRC) has licensed the Peregrine software package from Oak Ridge National Laboratory (ORNL). Peregrine is a printer- and camera-agnostic data analytics tool designed specifically for detecting process anomalies using in situ data collected during powder bed printing. In this work, temporally rich, multimodal sensor data, including visible light, integrated near-infrared (NIR), and spatially mapped coaxial melt pool thermal emission data, were fed into Peregrine to enable the detection of subsurface flaws. XCT data was used as ground-truth training data to allow Peregrine's deep learning (DL) algorithms to recognize anomalies in these complex data streams in test artifacts and an industrially relevant geometry.

Subsurface flaws such as pores and cracks can have significant effects on part performance, particularly under cyclic loading conditions. However, these flaws cannot typically be observed during the printing process simply by imaging the completed layer in the visible-light spectrum, which is the most common commercially available form of in situ monitoring for L-PBF systems. Other researchers have demonstrated that both off-axis and on-axis imaging in the infrared wavelengths can be used to observe these flaws [3], [4]. Previous techniques, however, suffer from either poor spatial resolutions or small fields of view, and challenges associated with human interpretation of the raw data make robust flaw detection difficult. Peregrine was used in a novel configuration to address both of these challenges.

Through its Dynamic Multilabel Segmentation Convolutional Neural Network (DMSCNN), Peregrine allows data from multiple sensor streams to be co-registered and fused together such that anomaly detections are made by leveraging the totality of the provided data [5]. RTRC has modified an EOS M290 L-PBF printer to capture three distinct image-like sensor streams during each layer of the print:

1. A high-resolution visible-light camera captures an image of the powder bed after laser fusion and after powder spreading. These data are readily human-interpretable but by themselves cannot be used to detect subsurface flaws.
2. A lower-resolution NIR imager captures integrated light from thermal emissions during layer fusion. Although thermal emissions are expected to correlate well to subsurface flaws, these data have a relatively low spatial resolution, can be challenging for a human to interpret, contain artifacts, and can be influenced by anomalies that formed in previous layers.
3. Two infrared photodiodes are used; one is coaxially aligned with the laser optics while the other observes the entire print area off-axis. This configuration allows the thermal emissions from the vapor plume, melt pool, and surrounding region to be measured. These data are captured in the time domain but can be spatially mapped. Although previous work has shown that these data can be very informative [6], they are generally difficult for a human to interpret.

The DMSCNN is also capable of incorporating temporal information by including sensor data from above and below the target print layer. Figure 1 shows a diagram of an example data stack. The primary goal of this project is to determine if such a data stack contains sufficient contextual and spatially resolved information to observe subsurface L-PBF flaws in the size range of 200–1,000  $\mu\text{m}$ .



**Figure 1. A representative in situ sensor data stack from an L-PBF printer consisting of information from three different sensing modalities and three different layers.** Peregrine's DMSCNN analyzes this data stack and outputs a segmented map of flaw and anomaly detections, which can be used to inform quality control decisions.

This data stack is extremely highly dimensional and not expected to be human-interpretable. Therefore, the DMSCNN was trained using XCT ground-truth data collected at the Zeiss Characterization Facility located at the Manufacturing Demonstration Facility (MDF), a US Department of Energy User Facility at


ORNL. Researchers at the MDF leveraged parallel work on XCT reconstruction to produce and coregister this high-quality training dataset. The XCT data were also be used to validate the performance of Peregrine's DMSCNN.

As a direct result of this program, two journal manuscripts [1], [2] have been published in *Additive Manufacturing*, with additional manuscripts planned following program completion.

## 2. RESULTS

### 2.1 PHASE I—TASK 1: PRINTING TEST ARTIFACTS

Four builds were completed at RTRC facilities on an EOS M290 using Inconel 718 powder. Figure 2 depicts an image of the completed build plate for one of the four builds. Four different coupon geometries were used in this study: (1) spatter generators, (2) XCT coupons, (3) powder capsules, and (4) overhanging structures. The spatter generators, denoted as the two large blocks in the upper left corner of the build plate, were intended to produce spatter particles that could land on the build plate and interfere with manufacturing the other components, thereby producing naturally occurring, stochastic lack-of-fusion porosity. The 46 XCT coupons were placed uniformly across the build plate to capture the spatial distribution of flaws and were XCT-scanned post-build to identify flaws. The powder capsules were used for post-build powder characterization, and the overhanging structures were held for future use and will not be discussed further in this report.



**Figure 2. Post-build image of RTRC Build #0849.** This image was captured by RTRC.

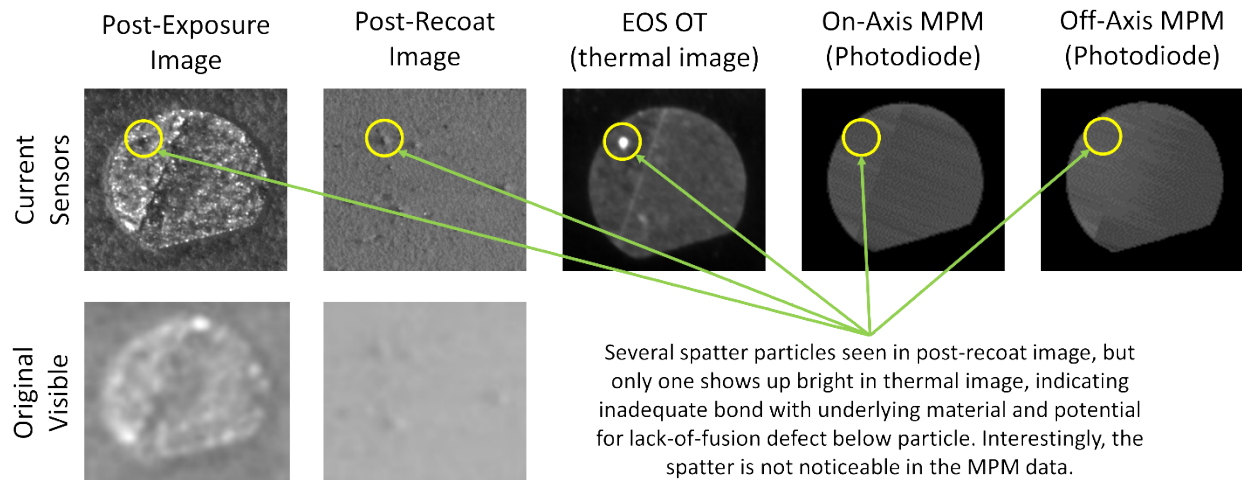
Although the build layout was kept constant across all four builds, each build used different combinations of processing parameters to generate test samples with flaw populations suitable for training, validating, and testing the segmentation model in Peregrine. The strategies employed for each build are outlined in Table 1. RTRC Build #0814 acted as a baseline build and used default process parameters on all parts aside from the two spatter generators. RTRC Build #0815 was identical to RTRC Build #0814, but the gas flow differential pressure was dropped from its default value of 0.79 mbar to 0.63 mbar. RTRC Build #0849 had the highest flaw concentration, employing numerous strategies for generating spatter. Notably, the process parameters used to build the spatter generators were chosen to generate severely keyholed melt pools. Finally, RTRC Build #0871 was intended to have an intermediate amount of spatter, returning all coupons to the default process parameters apart from the spatter generators, which still used the keyhole-prone parameter set.



**Table 1. Build strategies employed to produce coupons suitable for training, validation, and testing of the DL network in Peregrine.**

Build ID	Description
<b>RTRC 0814</b>	<ul style="list-style-type: none"> <li>• Default EOS parameters, except...</li> <li>• Spatter generators had 80 <math>\mu\text{m}</math> layer thickness; other parts had 40 <math>\mu\text{m}</math> layer thickness</li> <li>• Spatter generators printed first</li> </ul>
<b>RTRC 0815</b>	<ul style="list-style-type: none"> <li>• Same as #0814, except...</li> <li>• Gas flow was reduced to 0.63 mbar</li> </ul>
<b>RTRC 0849</b>	<ul style="list-style-type: none"> <li>• Same as #0815, except...</li> <li>• Spatter generators printed with known keyholing parameters</li> <li>• Layer thickness was 80 <math>\mu\text{m}</math> for all parts</li> <li>• Infill directions aligned with gas flow were allowed</li> </ul>
<b>RTRC 0871</b>	<ul style="list-style-type: none"> <li>• Same as #0815, except...</li> <li>• Spatter generators printed with known keyholing parameters</li> </ul>

Visible-light images, integrated NIR images, co-axial photodiode data, and off-axis photodiode data were collected for every layer of each of the four builds outlined in Table 1. As shown in Figure 3, the sensor loadout used in the present study demonstrated superior resolution and contrast compared with the industry standard sensors. The addition of the EOS optical tomography (OT) data and the two photodiode data streams was expected to improve Peregrine’s ability to detect flaws identified in XCT. For example, the spatter particle identified in Figure 3 is readily visible in the EOS OT data, but it is less apparent in the two industry-standard visible-light images. Notably, the spatter particle did not appear to be detectable in either of the two photodiode data modalities.



**Figure 3. Summary of process monitoring data collected during the current study.** Note that the visible-light camera originally installed on the machine was upgraded. Images from the upgraded camera demonstrated significantly higher resolution and contrast compared with the original camera’s visible-light images. The addition of the thermal and photodiode data was expected to increase the information content delivered to the segmentation model. MPM stands for melt pool monitoring.


## 2.2 PHASE I—TASK 2: X-RAY COMPUTED TOMOGRAPHY OF TEST ARTIFACTS

XCT scans of select coupons from the four builds were performed to interrogate the internal quality of the constructed components. Table 2 outlines the number of XCT coupons subjected to XCT scanning. Initial scans of XCT coupons from RTRC Builds #0814 and #0815 indicated low flaw counts, which would limit the training data available to Peregrine. In contrast, the specimens produced from RTRC Build #0849 contained numerous flaws that would provide samples of internal flaws to Peregrine for training. Consequently, all 46 of these XCT coupons were subjected to post-build XCT scanning.

**Table 2. Breakdown of the number of XCT coupons scanned from each of the four builds.**

<b>Build ID</b>	<b>Number of parts XCT-scanned.</b>
<b>RTRC 0814</b>	7
<b>RTRC 0815</b>	7
<b>RTRC 0849</b>	46
<b>RTRC 0871</b>	0

Figure 4 depicts a composite image of all 46 XCT scans conducted for RTRC Build #0849 spatially mapped to the build plate coordinate system defined in ISO/ASTM 52921 [7]. The composite image was constructed by projecting 100 layers' worth of XCT data from the XCT coupons onto a single image, taking the minimum value of each pixel across the 100 layers to generate the composite image. As shown in the figure, the spatial distribution of flaws across the build plate was heterogenous, with most flaws occurring just below the spatter generators (downstream from the gas flow). Notably, the XCT coupons that were not immediately downstream from the inert gas flow (see right side of the build plate in Figure 4) were relatively devoid of internal flaws. Flaws were identified in the XCT coupons algorithmically to record the flaw size and location on the build plate. In total, 5,915 flaws were identified in the 46 XCT coupons from RTRC Build #0849, with 1,689 having diameters greater than 200  $\mu\text{m}$  and the largest flaw having a diameter of 572  $\mu\text{m}$ . Considering the size and spatial distribution of flaws, these results indicate that many of the flaws larger than 200  $\mu\text{m}$  in diameter were generated via spatter-induced lack-of-fusion, which is known to be correlated with insufficient inert gas flow [8]. In keeping with this hypothesis, manual comparisons between the locations of large XCT flaws and the in situ sensor data indicated that spatter particles were frequently collocated with these flaws.



















**Figure 4. Composite image of registered XCT scan data for RTRC Build #0849 showing the spatial distribution of internal flaws across the build plate.** Notably, the most-porous XCT coupons were located just downstream from the spatter generators. The dashed green box denotes the approximate location of the spatter generators.

### 2.3 PHASE I—TASK 3: SENSOR FUSION AND FLAW DETECTION

Registered data streams from the process monitoring data (see Figure 3) were compared with the registered XCT volumes (see Figure 4) to generate ground-truth labels for supervised training Peregrine’s Dynamic Segmentation Convolutional Neural Network (DSCNN). A custom EOS M290 workspace was created in Peregrine with the available pixel classes shown in Figure 5. Note that during Phase I, a pixel could only belong to a single class, so class designations were mutually exclusive. During Phase II, the DSCNN was overhauled to create the multi-label DMSCNN, which can predict multiple classes simultaneously for each pixel. Note that although in total, 15 pixel classes were instantiated in the

Peregrine workspace, examples of only 8 pixel classes (denoted by the black boxes in Figure 5) were annotated in the training and testing data.

	Powder
	Printed
	Printed (non-optimal)
	Stripe Boundary
	Misprint
	Edge Swelling
	Super-Elevation
	Hopping Recoater
	Streaking Recoater
	Mounding Powder
	Incomplete Spreading
	Debris
	Spatter on Part (non-flaw)
	Spatter on Powder
	XCT Flaw (small)
	XCT Flaw (large)


**Figure 5. Pixel classes available for training in Peregrine.** Although in total, 15 pixel classes were instantiated in the Peregrine workspace, examples of only 8 pixel classes (denoted by the black boxes) were annotated.

Pixels of process monitoring data corresponding to the locations of XCT flaws were labeled as being either “*small*” or “*large*”. Flaws with spherical equivalent diameters (SED) approximately 150  $\mu\text{m}$  and larger were considered “*large*”. Once the XCT flaw pixels were semiautomatically annotated, the remaining pixels were annotated by a trained operator to identify regions of successfully printed material (*printed*), intentionally unmelted powder (*powder*), spatter particles that did not correspond to XCT flaws [*Spatter on Part (non-flaw)*], and more. For Phase I, a total of 16 layers’ worth of in situ sensor data (156,466,091 pixels) were annotated, 12 of which were used for training, and the remaining 4 layers were used for testing the trained network. A key contribution of this work was the extension of Peregrine to accept multiple layers of data to annotate and train a single layer of data, as described in Figure 1. For each layer, the layer above and the layer below were used as input, allowing the DSCNN to gain context from a 3D snapshot of data instead of only a single layer of 2D data.

Figure 6 shows preliminary segmentation results of the trained DSCNN within Peregrine for Layer #95 of the testing data for the five coupons just below the spatter generators from RTRC Build #0849. The top row of data shows the XCT flaw segmentation results for these five coupons, and the second and third rows show the corresponding EOS OT integrated data and ground-truth annotations, respectively. The final row shows Peregrine’s predicted classes for each pixel. As shown in the figure, the DSCNN appears to readily distinguish successfully printed material and unmelted powder particles in the powder bed. Stripe boundaries were also detectable, although there appears to be some confusion between stripe boundaries and printed material.

Of particular interest are the spatter particles on printed material that did not correspond flaws [*Spatter on Part (non-flaw)*] and large XCT flaws [*XCT Flaws (large)*]. Despite having very similar signatures in the process monitoring data, the DSCNN appears to be able to distinguish the difference between these two classes—albeit with a considerable amount of confusion between the two classes. It is hypothesized that the DSCNN is distinguishing between these two classes by looking at a highly dimensional feature space resulting from the multiple layers of multimodal sensor data. This highly dimensional feature space is not easily interpreted by human operators, making machine learning models an ideal candidate for such a task. Flaws with diameters of approximately 300  $\mu\text{m}$  and above were frequently detected successfully, and smaller flaws were sometimes missed, being predicted as either successfully printed material or

spatter particles that did not result in a flaw. Notably, the resolution of the sensor data was approximately 100  $\mu\text{m}$ , so flaws smaller than 2–3 pixels ( $\sim 300 \mu\text{m}$ ) are not expected to be easily detected by any algorithm.



**Figure 6. DSCNN segmentation results for Layer #95 for the five coupons just below the spatter generators (P43-47) from RTRC Build #0849.**

The pixel-wise confusion matrices by predicted classes and ground-truth labels are given in Table 3 and Table 4, respectively. Specifically, the confusion matrix presented in Table 3 shows the percentage of pixels predicted to be a certain class that were actually another class according to the ground-truth testing labels. For example, of all the pixels predicted by the trained DSCNN network to be successfully printed material, 98.3% actually were printed material. Of the remaining 1.7% of pixels predicted to be printed material, 1.1% were labeled as stripe boundaries in the ground-truth data, and the remaining 0.6% pixels predicted to be printed material were a combination of the remaining six pixel classes used. Several features are notable. First, the quantitative pixel-wise statistics presented in Table 3 support the qualitative findings presented in Figure 6. The trained network is particularly good at detecting powder and printed material, with 98.9% and 98.3% of predicted powder and printed pixels, respectively, correctly belonging to those classes in the ground-truth data. Similarly, 89.0% of predicted stripe boundaries were actually stripe boundaries, although 10.4% of predicted stripe boundaries were instead misclassified as printed material. Additionally, none of the pixels corresponding to either recoater streaking or small XCT flaws actually belonged to those classes. The poor performance of recoater streaking predictions can be attributed to the lack of recoater streaking in the region of interest defined in Figure 6; in the case of the small XCT flaws, the poor prediction performance can be attributed to the limited spatial resolution of the sensors relative to the size of the flaws belonging to this class.

As noted previously, there was considerable confusion between pixels corresponding to large XCT flaws and spatter particles on the parts that did not correspond to an internal flaw. This issue can be seen quantitatively in Table 3. Of all pixels predicted to be a large XCT flaw, only 46.3% of them were actually labeled as such in the ground-truth data. Of the remaining 53.7% of predicted large XCT pixels, 9.2% were actually spatter particles that did not result in a flaw, and 43.5% were printed material. Similarly, 39.9% of the pixels predicted to be spatter particles that did not result in a flaw were labeled as such, and an additional 14.3% were actually large XCT flaws. As noted previously, some confusion

between these two classes was expected given the similar nature of the process monitoring signatures in these locations. Finally, predictions of spatter particles that landed on the powder bed were overly aggressive, with only 13.4% of predicted spatter on powder being labeled as such in the ground-truth testing data. In contrast, 85.9% of predicted spatter on powder pixels were actually powder.

**Table 3. Testing pixel-wise confusion matrix by predicted classes for the five XCT coupons just below the spatter generators (P43-47).**

	Powder	Printed	Stripe Boundary	Recoater Streaking	Spatter on Part (non-flaw)	Spatter on Powder	XCT Flaw (small)	XCT Flaw (large)	Ground Truth Classes
Powder	98.9	0.4	0.2	100.0	0.0	85.9	0.0	0.0	
Printed	0.0	98.3	10.4	0.0	39.9	0.1	93.2	43.5	
Stripe Boundary	0.0	1.1	89.0	0.0	5.9	0.0	6.8	1.0	
Recoater Streaking	1.0	0.0	0.0	0.0	0.0	0.6	0.0	0.0	
Spatter on Part (non-flaw)	0.0	0.1	0.1	0.0	39.9	0.0	0.0	9.2	
Spatter on Powder	0.0	0.0	0.0	0.0	0.0	13.4	0.0	0.0	
XCT Flaw (small)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
XCT Flaw (large)	0.0	0.1	0.2	0.0	14.3	0.0	0.0	46.3	
Predicted Classes									

Conversely, the confusion matrix presented in Table 4 shows the percentage of pixels belonging to a certain class in the ground-truth labels that were predicted to be from another classes. For example, the trained DSCNN was able to successfully identify 97.4% of printed material, and the remaining 2.6% of pixels labeled as printed in the ground-truth data were predicted to be another class. The results presented in Table 4 support those presented in Table 3. The network was successfully able to detect powder and printed pixels but struggled to distinguish the difference between large XCT flaws and spatter particles that did not result in a flaw. Of the pixels labeled as large XCT flaws in the ground-truth data, only 39.4% were labeled as such. The remaining large XCT flaw pixels were largely labeled as either spatter particles on the part that did not result in a flaw (14.8%) or printed material (43.7%).

**Table 4. Testing pixel-wise confusion matrix by ground-truth annotations for the five XCT coupons just below the spatter generators (P43-47).**

	Powder	Printed	Stripe Boundary	Recoater Streaking	Spatter on Part (non-flaw)	Spatter on Powder	XCT Flaw (small)	XCT Flaw (large)
Powder	96.8	0.4	0.0	1.2	0.0	1.7	0.0	0.0
Printed	0.4	97.4	1.3	0.3	0.3	0.0	0.1	0.2
Stripe Boundary	0.3	25.6	72.6	0.5	0.5	0.0	0.2	0.2
Recoater Streaking	18.9	15.6	0.2	64.9	0.1	0.3	0.0	0.0
Spatter on Part (non-flaw)	0.4	36.5	2.7	0.3	50.0	2.0	0.1	8.0
Spatter on Powder	7.4	0.4	0.0	0.0	0.0	92.0	0.0	0.0
XCT Flaw (small)	0.0	91.2	6.1	0.6	1.1	0.0	0.3	0.8
XCT Flaw (large)	0.0	43.7	2.0	0.0	14.8	0.1	0.1	39.4

Ground Truth Classes


Predicted Classes

Note that within the context of using process monitoring for in situ flaw detection, pixelwise annotation performance is not necessarily the most meaningful indicator of model performance because the goal of this work is not to accurately measure the flaw morphology but to identify individual instances of flaws. Therefore, the instance-wise performance of the network was assessed during Phase II of the present program. This work was accomplished by generating probability of detection (PoD) and false-positive rate curves based on instances of individual flaws.

Figure 7 depicts an estimate of Peregrine’s ability to detect XCT flaws directly from in situ monitoring data for a single component from RTRC Build #0849. The locations of flaws detected in the registered XCT data were compared with the locations of flaws detected by Peregrine in the registered in situ monitoring data. The size of flaws identified in the XCT and in situ data were estimated by their SED—the diameter of a sphere with a volume equivalent to the volume of the flaw in question. A histogram of XCT flaw sizes was generated, and the number of flaws within each size bin that were successfully detected by the DSCNN was recorded. Successful detections by Peregrine were determined using a voxel-to-voxel distance calculation between flaws detected in the XCT and in situ monitoring data, and instances of XCT flaws having DSCNN-detected flaws within 250  $\mu\text{m}$  were considered to have been successfully detected. This distance threshold was determined through estimates of the XCT and sensor registration error ( $\sim 100 \mu\text{m}$  in total) and considerations of interactions between spatter particles and the melt pool ( $\sim 150 \mu\text{m}$ ). The percentage of XCT flaws within each size bin that were successfully detected by Peregrine is plotted as a solid black line in Figure 7 and can be interpreted as an estimate of the PoD of



the Phase I DSCNN. The dashed black line denotes the percentage of flaws of a given size class and above that were successfully detected by Peregrine.



**Figure 7. Estimate of the Phase I segmentation model’s PoD for a single coupon (P43) from RTRC Build #0849.** The size of flaws detected in the XCT data is shown on the x-axis, and the number and percentage of flaws detected in each size class are plotted on the left and right y-axes, respectively. As shown, the percentage of XCT flaws successfully detected using Peregrine increased as the flaw size increased.


As shown in the figure, Peregrine was successfully able to detect all flaws with an SED of approximately 375  $\mu\text{m}$  or greater in part number P43 from RTRC Build #0849. However, the percentage of XCT flaws successfully detected dropped as the flaw size decreased, likely because of a lack of adequate spatial resolution and contrast. Although these preliminary results are promising, additional work was required to fully characterize Peregrine’s ability to detect XCT flaws from in situ data.

As part of Phase II activities, the PoD analysis was advanced to be consistent with MIL-HDBK-1823A [9], and the results were published in Snow et al. [2]. An excerpt from this manuscript is adapted below to describe the performance of the technique developed in this program, called the *in situ nondestructive evaluation* (INDE) system.

The flaw detection performance of the INDE system can be seen quantitatively in the PoD curve presented in Figure 8, which shows the PoD as a function of XCT flaw CED [circular equivalent diameter]. Following the guidelines presented in MIL-HDBK-1823 [9], XCT flaws that were successfully detected by the INDE system are shown as circles at the top of the figure, while missed detections are shown with cross-patterns at the



bottom of the figure. The solid black line shows the mean PoD curve, while the dashed line represents the lower 95% confidence interval. The solid red line corresponds to the  $a_{90/95}$ , and the gray region highlights XCT flaw sizes below the resolution limit of the in situ sensors. Finally, the histograms at the top of the figure show the flaw size distributions of XCT flaws that were and were not successfully detected, respectively. Note that the histograms are displayed on different y-axes to facilitate comparisons between the sizes of XCT flaws successfully detected by the INDE system (green) and XCT flaws that were not detected (red).



**Figure 8. PoD curve for the INDE system, with successful detections shown by green circles and unsuccessful detections depicted with red crosses.** As shown, the PoD improves significantly for larger flaws, with almost all flaws with CEDs greater than 400  $\mu\text{m}$  being successfully detected. Notably, many of the missed detections corresponded to confusion between the XCT flaw (large) and spatter-on-part (non-flaw) classes.

As shown by the cross-patterns at the bottom of the figure, almost all flaws with CEDs greater than 400  $\mu\text{m}$  were successfully detected by the INDE system, with only a single missed detection for an XCT flaw larger than 600  $\mu\text{m}$ . Further investigation of this missed flaw detection revealed that the anomaly was in fact detected by the INDE system, but it was misclassified as the *spatter-on-part (non-flaw)* class. Manual interrogation of the segmentation results revealed that many of the missed XCT flaw detections were the result of misclassifications as *spatter-on-part (non-flaw)*, reinforcing the confounding between *XCT flaw (large)* and *spatter-on-part (non-flaw)* pixels. For the current iteration of the INDE system, the  $a_{90/95}$  was found to be approximately 661  $\mu\text{m}$ . While the performance of the INDE system generally degraded for smaller flaws, particularly for flaws smaller than the resolution limit of the sensors, the INDE system was able to detect flaws smaller than the resolution limit, likely because the sensor signatures corresponding to those flaws were larger than the resolution limit. Close inspection of Figure 8] in [2] also shows that, in agreement with the findings of [10], many of the smaller XCT flaws appear to make up larger defect structures connected to a single spatter particle indication, which may also explain why successful detections were made for flaws that were below the resolution limit of the in situ sensor data. From a

mechanics-of-materials perspective, smaller flaws clustered together may be considered as a single flaw [11], so considering clustered flaws as a single flaw may both improve the PoD performance and be more relevant to the mechanical performance of the printed material. These combined effects likely affected the quality of the PoD curve since the functional form of the PoD curve assumes a clear separation between flaws that are detectable and those that are not. Since many small, detected flaws were likely tied to the same defect structure, the effect of flaw size on the PoD is muddled, resulting in an artificially poor PoD.


## 2.4 PHASE II—TASK 1: INDUSTRY-RELEVANT PERFORMANCE

Completion of the Phase I activities highlighted a significant limitation of the DSCNN network architecture when using XCT data as a label source for training. Although the use of XCT data during training allowed flaws visible in the ex situ characterization data to be annotated, any components or regions of components lacking registered XCT could not be annotated without compromising the veracity of the ground-truth labels. This issue is because class labels are mutually exclusive using the DSCNN model architecture (i.e., a single pixel could only be labeled as belonging to a single class). In reality, pixels in the sensor data represent regions of physical space in which multiple physical phenomena may be observed coincidentally. For example, a pixel could represent successfully printed material (i.e., belonging to the *printed* class) and a stripe boundary (i.e., the *stripe boundary* class), but, using the DSCNN architecture, a single label would have to be chosen. Since not every component on the build plate will have XCT data, the *XCT Flaw (large)* class cannot be labeled for these pixels, leaving only the remaining pixels' classes. However, *XCT Flaw (large)* pixels represents a spatially small pixel class, so any in situ data labeled as another class simply because they lacked XCT data to label would degrade the quality of these annotations to an unacceptable level. Therefore, components from the Phase I builds lacking XCT data were left unlabeled, resulting in poor performance for those part geometries and limiting the generalizability of the model and the scale at which this technology could be applied in an industrial setting.

To address this issue, a new model architecture was designed and implemented in Peregrine that could assign multiple labels to each pixel. The new architecture, called the DMSCNN, is a multiclass, multilabel network that effectively trains as a set of jointly learned binary classifiers in which pixels can be assigned to multiple classes or, crucially, can be ignored for specific classes during training. The overhaul of Peregrine's model architecture effectively allows pixels representing components lacking XCT data to be annotated for non-XCT-informed classes without inadvertently introducing inaccurate XCT ground-truth annotations into the training set. Additionally, it is now possible for the DMSCNN to represent different laser processing conditions because the anomaly classes are no longer mutually exclusive. This functionality is hypothesized to be a powerful tool when in situ imaging data are leveraged for process parameter development. Finally, the new architecture also allows the DMSCNN to report that a pixel is *unknown* (i.e., it cannot confidently predict its state given its current training). This work is a first step toward uncertainty quantification and providing justifications and confidence levels alongside AI-based predictions.

While the DMSCNN was being developed, a new build was completed by RTRC that included a sanitized geometry, called the *industrially relevant component* (IRC) in this report, representative of a structural component in a commercial turbofan engine. Logistical constraints with other programs at ORNL prevented the build from being replicated at the MDF, but a sufficient amount of training data was still generated by the build completed at RTRC. Figure 9 shows representative images of the sensor data collected during this build. Six different coupon geometries, printed out of Inconel 718 using the RTRC Build 0849 parameters described in Table 1, were used in this study: (1) spatter generators, (2) Zeiss XCT coupons, (3) pentagon geometries, (4) fatigue coupons, (5) the IRC, and (6) miniature versions of the

IRC. Similar to the Phase I activities, the spatter generators were designed to distribute spatter particles across the rest of the build plate while the Zeiss and pentagon geometries were printed to provide in situ data representing geometrically complex components that could also be XCT-scanned with relative ease. Fatigue coupons were also printed with the intention of assessing the fatigue performance of these components in the future and correlating the results with the in situ detections. Fatigue testing was not within the scope of this program, but it should be noted that the fiducial markers were added to the gauge sections of the fatigue coupons to facilitate the registration of XCT scans to in situ data for this critical region.




**Figure 9. Representative sensor data collected from Layer #556 of the build performed at RTRC.**

Finally, the IRC was printed to assess the generalizability of this in situ monitoring technique to difficult-to-inspect, industrially relevant geometries. The IRC was approximately 10 in. in diameter, made of a

radiographically dense material (Inconel 718), and had a complex internal structure that introduced numerous imaging artifacts into the XCT reconstructions. The dimensions of the IRC alone present a challenge for post-build nondestructive evaluation because commercial XCT systems use geometric magnification to fit objects of interest into the field-of-view of the detector. Therefore, scanning the entirety of the IRC would have resulted in an unacceptably low spatial resolution, been prone to significant XCT imaging artifacts because of the complex geometry, and demonstrated poor signal-to-noise ratios because of attenuation of the x-rays. These constraints on post-build nondestructive evaluation made the IRC an ideal candidate for assessing the in situ monitoring technique developed in this program, which should be less sensitive to these challenges than XCT. Miniaturized versions of the IRC were also printed to provide a mechanism for isolating the effect of complex geometries on post-build XCT, but they were not evaluated further in this work.

XCT scans were performed on the Zeiss geometries, pentagons, and the tops and gauge sections of the fatigue coupons, and a subset of this data was also registered to the in situ data and the Peregrine coordinate system. Given the XCT challenges of the IRC, the entirety of the component could not be XCT scanned nondestructively, but two subsections of the IRC were excised using wire electrical discharge machining, XCT-scanned, and registered to the in situ data. Notably, even though subsections of the IRC were able to be XCT scanned, their large size still resulted in larger voxels ( $77.96\text{ }\mu\text{m}$ ) than the XCT scans of the other components ( $17.28\text{ }\mu\text{m}$ ). Figure 10 shows a representative example of XCT data for the layers comprising the gauge section of the fatigue coupons, highlighting the nonuniform distribution of spatter-induced lack-of-fusion. Additional details regarding XCT reconstruction and analysis developments are presented in Section 2.6.




**Figure 10. XCT Supercomposite of layer 515-610, summarizing the porosity structure throughout the gauge regions of the fatigue coupon blanks.** The resolution of the fatigue coupon blanks was approximately  $17.28\ \mu\text{m}$ , and the resolution of the IRC subsection was only  $77.96\ \mu\text{m}$ . Notably, some features of the IRC subsections may not be accurately represented because, unlike the fatigue coupons, the slice geometry changed on a layer-wise basis.

Once the XCT data were registered, three additional layers' worth of in situ data from the Phase II build were annotated and added to the Peregrine training dataset. The previous layer-wise annotations from Phase I were also updated to make full use of the multiclass, multilabel capabilities of the DMSCNN. A new DMSCNN model was then trained and applied to the in situ data generated during the Phase II build performed by RTRC.

Figure 11 compares the segmentation performance of the trained Phase I and II segmentation models for Layer #556 of the Phase II build. Figure 11a shows the segmentation results using the model trained from Phase I, and Figure 11b shows the segmentation results of the same layer using the trained DMSCNN. As


shown in the figure, both models are able to detect numerous anomaly classes, including the *XCT Flaw (large)* class, in regions of the build lacking XCT data. However, the mutual exclusion of class labels resulted in poor segmentation performance of the Phase I model for the parts printed with the spatter generation parameter, which are denoted by the black dotted boxes in each image. Although XCT data were lacking for these components in both builds, the DMSCNN trained during Phase II was successfully able to identify these components as having pixels belonging to the *printed (non-optimal)* class, as well as accurate predictions of the *printed, stripe boundary*, *spatter-on-part (non-flaw)*, and *recoater streaking* classes, where appropriate. In contrast, the segmentation model trained during Phase I of this program was only able to provide accurate segmentation results for those components having XCT data and in situ data signatures that were similar to the training data. In particular, the spatter generator components displayed overpredictions of the *XCT Flaw (large)* class using the Phase I segmentation model. The ability of the DMSCNN to facilitate annotations for components lacking an ex situ data label source, such as XCT, loosens the restrictions on which pixels can be annotated for training, which should make the new network architecture more scalable in an industrial setting.



**Figure 11. Segmentation results for Layer #556 from the Phase II build using (a) the Phase I segmentation model and (b) the DMSCNN trained for Phase II.** As shown, regions of the build corresponding to parts that were not XCT-scanned can now be analyzed with reasonable results, whereas the Phase I segmentation model yielded poor segmentation results, particularly for the spatter generators.


Figure 12 shows the rate at which a flaw identified in the gauge regions of the fatigue coupons had a matching in situ *XCT Flaw (large)* detection. This in situ detection hit rate is given as a function of XCT flaw CED. Notably, only XCT flaws at least 1 mm from the part surface were considered for this analysis to avoid erroneous flaw detections in the XCT data commonly associated with the rough surfaces of AM components. Additionally, only those in situ *XCT Flaw (large)* instances that persisted for multiple layers were used in the hit rate analysis, which reduces the rate of false detections at the cost of a more conservative hit rate performance. As shown in the figure, all flaws with CEDs greater than approximately 625  $\mu\text{m}$  were successfully detected using the in situ monitoring technique developed in this program, showing comparable performance to the DSCNN developed in Phase I while now being generalizable to more complex geometries.






**Figure 12. In situ detection hit rate as a function of XCT flaw size for the Phase II network applied to fatigue coupon gauge sections only.**

Notably, a similar analysis was not conducted for the IRC subsections because of a limited number of internal flaws being evident in the XCT data. Although this lack of analysis presents challenges for verifying the performance of the in situ monitoring system's performance on the IRC, it also highlights the importance of the technique for inspecting radiographically dense materials and large and/or complex geometries. Despite this issue, the in situ monitoring method developed here was still able to successfully detect flaws inside of the IRC subsections, as demonstrated in Figure 13.



**Figure 13.** Example of (a) an internal flaw detected in the in situ data and (b) its corresponding indication in the XCT data.

Finally, difficulties in spatially registering XCT data to the in situ data were found to possibly result in complications during data annotation. For example, Figure 14 shows point cloud representations of the intended geometry (in orange) overlaid with the spatially registered XCT data (purple) of that same component. Because of residual stresses in the as-printed component (even after a stress relieving heat treatment was applied), the IRC subsection distorted upon excision, resulting in a macroscopic geometric difference between the XCT data volume and the intended geometry, potentially introducing annotation errors for the *XCT Flaw (large)* class. For the present work, regions of the IRC subsections that distorted were avoided during annotation, but it is noted that similar annotation errors may result from poor spatial registration of ex situ data to in situ data and should be considered carefully during annotation.



**Figure 14.** Point cloud representations of the intended geometry (orange) of a subset of the IRC and the registered XCT data (purple) of the final printed component. As shown, the printed component underwent significant distortion after being excised from the component, resulting in poor spatial alignment of the XCT and in situ data.




## 2.5 PHASE II—TASK 2: FACILITATING HUMAN INTERPRETABILITY

For this program, a suite of automated analysis tools for evaluating the performance of models training in Peregrine were developed to facilitate model comparisons between builds, machines, sensing suites, vendors, and users. For example, automatically generated tables of common machine learning performance metrics (e.g., accuracy, weighted accuracy, precision, recall) are reported for each class—both on the validation and testing data. Settings for the automated model evaluation are also recorded in a text file for record keeping, and layer-wise analysis timings are automatically recorded for each layer in the testing dataset. Reported timings include the time required to load sensor data, load template images, load annotations, calculate the class interfaces, apply the trained model, perform heuristics, and assign global labels to the layer. An example of the automated machine learning performance metrics is reported in Table 5.

**Table 5. Automated segmentation results generated for the validation data of the Phase II build**

class	accuracy	weighted accuracy	recall	precision	true negative rate	false positive rate	F1	ROC AUC	PR AUC
Powder	99.5%	99.3%	99.7%	99.7%	98.9%	1.1%	99.7%	100.0%	100.0%
Printed	99.5%	99.3%	99.0%	99.1%	99.7%	0.3%	99.0%	100.0%	99.9%
Printed (non-optimal)	99.7%	98.9%	98.0%	97.3%	99.8%	0.2%	97.7%	99.9%	98.9%
Stripe Boundary	99.2%	92.4%	85.5%	67.7%	99.4%	0.6%	75.6%	99.6%	83.9%
Streaking Recoater	99.0%	90.1%	80.9%	66.3%	99.3%	0.7%	72.9%	98.9%	79.3%
Spatter on Part (non-flaw)	99.6%	78.3%	56.9%	20.9%	99.7%	0.3%	30.5%	99.2%	32.5%
Spatter on Powder	98.9%	69.4%	38.8%	95.8%	100.0%	0.0%	55.2%	98.5%	76.6%
XCT Flaw (small)	99.9%	62.1%	24.3%	3.2%	99.9%	0.1%	5.7%	94.2%	2.8%
XCT Flaw (large)	99.8%	84.9%	69.9%	10.4%	99.8%	0.2%	18.1%	99.1%	36.1%

In addition to the tables of automated machine learning results metrics, receiver-operator curves and precision-recall curves are also automatically generated for both the validation and testing datasets (see Figure 15).




**Figure 15. Receiver-operator curves and precision-recall curves for the validation dataset.** These curves are automatically generated for both the validation and testing datasets to facilitate performance between models.


Comparative tools such as these could be used for an ablation study, which could show the sensor modalities that provide the most information to the DMSCNN. The time required to develop and implement the DMSCNN, as well as complications related to analyzing the XCT data, prevented ablation studies from being completed during the program’s period of performance, but such analyses are still of interest to the program participants and will be the subject of future collaborations.

## 2.6 X-RAY COMPUTED TOMOGRAPHY DEVELOPMENTS UNDER PHASE II

Several samples of three separate geometries were XCT-scanned throughout Phase II to provide training data for the DMSCNN. To improve the quality of the XCT data and associated training labels, DL-based reconstruction methods previously used for Phase I activities were further refined. Although these samples were of a different size, geometry, and material, the same trained model was able to successfully generate the necessary ground truth for the in situ monitoring data with minimal additional training. The computer-aided design–deep learning model-based iterative reconstruction (CAD-DLMBIR) [12] model was trained on three XCT reconstruction volumes of the Zeiss geometry (made out of stainless steel 316L) scanned on a Zeiss Metrotom system. Notably, the test samples had different shapes and geometry, with diameters ranging from 15 to 20 mm, printed with Inconel 718, and scanned on a Nikon H225 system. Despite having such out-of-distribution data, the retrained CAD-DLMBIR model was able to robustly perform reconstruction for these samples even when only small number of projections are available (see Figure 16 and Figure 17).




**Figure 16. Example XCT slices using various reconstruction algorithms: model-based iterative reconstruction (MBIR), Feldkamp–Davis–Kress (FDK), and the CAD-DLMBIR method used here.** As shown, reducing the number of projections by a factor of 10 significantly degrades the quality of the FDK reconstructions (industry standard), and the CAD-DLMBIR reconstruction shows image quality comparable with the much slower MBIR reconstruction with a fraction of the required views.



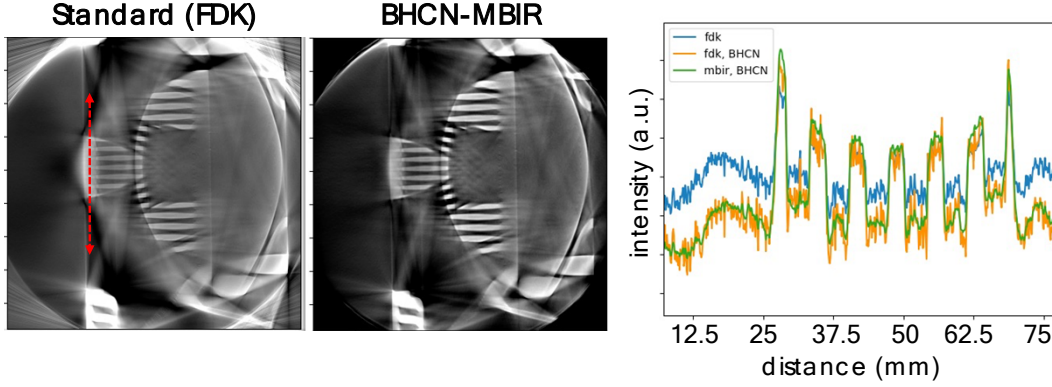
**Figure 17. Reconstructions of the pentagon geometry using CAD-DLMBIR trained on a different system with other geometries and materials.** The image quality of the reconstructed volumes on these out-of-distribution data, even with a sparse number of projections, showcases the generalizability of this reconstruction algorithm.

Additional exploratory reconstructions were performed for some of the samples using the Simurgh algorithm [13], and preliminary inspection of the Simurgh reconstructions showed promise for generating high-quality reconstructions (see Figure 18).



**Figure 18. Qualitative demonstration of Simurgh results on 10× reduction in the number of views for a sample that was more than 30% thicker than the samples used during training.**


Finally, the ORNL-developed DL reconstructions were evaluated for the IRC printed during Phase II. The IRC was more than 200 mm wide, limiting the resolution of the XCT scan to approximately 125  $\mu\text{m}$ . Instead of relying on the CAD-DLMBIR or Simurgh algorithms, which were not trained on this voxel size and such a complex geometry, the beam hardening correction network–model-based iterative reconstruction (BHCN-MBIR) method [14] was used (see Figure 19). However, even that technique had difficulty with reconstructing the full-sized sample, and more tuning was needed to improve the image quality.



**Figure 19. A single slice from the reconstructed volume using standard and BHCN-MBIR reconstructions.** A line profile along the arrow is plotted. Blue and green line profiles correspond to the standard and BHCN-MBIR reconstructions, respectively, showing higher contrast near the part edge and lower noise within the bulk. Although BHCN-MBIR fit slightly better, it still is far from perfect and needs further improvement before it can be used in a production environment. FDK stands for Feldkamp–Davis–Kress.

One challenge with these data was that only 1,571 views were used to perform scanning on such a complex geometry, and at least 2,000 projections would likely be required given the size of the IRC and the XCT detector dimensions. For the large samples with complex geometries, it is critical to ensure that the XCT system and setup is suitable for the intended geometry so that the x-rays can penetrate the object sufficiently. Otherwise, the number of artifacts owing to beam blocking, beam hardening, and streaking because of dense complex geometry makes volume reconstruction very challenging.

To address these challenges, two subregions of the IRC, which were expected to represent different flaw populations, were excised from the IRC and rescanned. These individual sections were scanned with 3,142 views, and the final reconstructed volume voxel size was 78  $\mu\text{m}$ . BHCN-MBIR was again used to perform reconstruction, followed by a DL-based segmentation to extract internal pores. An example of the results is shown in Figure 20.



**Figure 20.** A sample slice from the reconstruction of a cut section of the IRC, along with its segmentation and an expanded view of a region of interest within the slice. Pixels are shown represented as (red) segmented flaws, (green) the component, and (purple) the background.

### 3. CONCLUSION

L-PBF is a promising technology for the aerospace and defense industries because of its ability to produce parts with fine feature size, complex internal geometries, and using a broad range of materials. However, part size, feature complexity, and the nature of the flaws found in L-PBF (e.g., crack-like voids) make the inspection of AM parts using conventional inspection techniques difficult or impossible at the necessary resolution. There exists an opportunity and unmet need to use commercially viable in situ sensors to supplement or offset conventional ex situ inspection in the certification of complex AM parts. However, those commercially viable in situ sensors are sensitive to a wide variety of process phenomena, which may not necessarily be pertinent to the detection of flaws in the AM process and have rendered traditional attempts to use in situ monitoring for flaw detection less robust than is necessary for part qualification.

The Peregrine software, coupled with XCT ground-truth data, was used to detect flaws identified in XCT scans directly from in situ process monitoring data with the goal of improving the probability of flaw detection to levels commensurate with traditional methods. The multilayer, multimodal input data stack allowed Peregrine to detect numerous subsurface flaws in the size range of 200–1,000  $\mu\text{m}$ . The present results represent a promising first step toward in situ flaw detection for commercial AM processes.

Additionally, a new model architecture, called the Dynamic Multilabel Segmentation Convolutional Neural Network (DMSCNN), was designed and implemented in Peregrine. The new architecture is a multiclass, multilabel network that allows pixels to be assigned to multiple classes and allows specific classes to be ignored in specific locations during training. Critically, the DMSCNN allows pixels representing components lacking XCT data to still be annotated, as well as allowing a pixel to be classified as *unknown* (i.e., it cannot confidently predict its state given its current training). This work is a first step toward uncertainty quantification and providing justifications and confidence levels alongside AI-based predictions.

Several tasks have been identified for future work. First, an ablation study should be performed, wherein different data streams will be selectively removed from the DMSCNN to assess the effect of each data modality on the segmentation results and PoD analyses. The segmentation performance would be compared with the model performance when all data modalities are included to determine the sensor data modalities that provide the most information to the network and how many print layers of sensor data are required to accurately detect flaws.

#### 4. REFERENCES

- [1] Z. Snow, L. Scime, A. Ziabari, B. Fisher, and V. Paquit, “Observation of Spatter-Induced Stochastic Lack-of-Fusion in Laser Powder Bed Fusion Using In Situ Process Monitoring,” *Addit. Manuf.*, vol. 61, p. 103298, Jan. 2023, doi: <https://doi.org/10.1016/j.addma.2022.103298>.
- [2] Z. Snow, L. Scime, A. Ziabari, B. Fisher, and V. Paquit, “Scalable in situ non-destructive evaluation of additively manufactured components using process monitoring, sensor fusion, and machine learning,” *Addit. Manuf.*, vol. 78, p. 103817, Sep. 2023, doi: [10.1016/j.addma.2023.103817](https://doi.org/10.1016/j.addma.2023.103817).
- [3] J.-B. Forien, N. P. Calta, P. J. DePond, G. M. Guss, T. T. Roehling, and M. J. Matthews, “Detecting keyhole pore defects and monitoring process signatures during laser powder bed fusion: A correlation between in situ pyrometry and ex situ X-ray radiography,” *Addit. Manuf.*, vol. 35, p. 101336, Oct. 2020, doi: [10.1016/j.addma.2020.101336](https://doi.org/10.1016/j.addma.2020.101336).
- [4] G. Mohr *et al.*, “In-Situ Defect Detection in Laser Powder Bed Fusion by Using Thermography and Optical Tomography—Comparison to Computed Tomography,” *Metals*, vol. 10, no. 1, p. 103, Jan. 2020, doi: [10.3390/met10010103](https://doi.org/10.3390/met10010103).
- [5] L. Scime, D. Siddel, S. Baird, and V. Paquit, “Layer-wise anomaly detection and classification for powder bed additive manufacturing processes: A machine-agnostic algorithm for real-time pixel-wise semantic segmentation,” *Addit. Manuf.*, vol. 36, p. 101453, Dec. 2020, doi: [10.1016/j.addma.2020.101453](https://doi.org/10.1016/j.addma.2020.101453).
- [6] B. Lane, E. Whinton, and S. Moylan, “Multiple sensor detection of process phenomena in laser powder bed fusion,” presented at the SPIE Commercial + Scientific Sensing and Imaging, J. N. Zalameda and P. Bison, Eds., Baltimore, Maryland, United States, May 2016, p. 986104. doi: [10.1117/12.2224390](https://doi.org/10.1117/12.2224390).
- [7] “Standard Terminology for Additive Manufacturing—Coordinate Systems and Test Methodologies.” ISO/ASTM International, 2019.
- [8] C. Schwerz, A. Raza, X. Lei, L. Nyborg, E. Hryha, and H. Wirdelius, “In-situ detection of redeposited spatter and its influence on the formation of internal flaws in laser powder bed fusion,” *Addit. Manuf.*, vol. 47, p. 102370, Nov. 2021, doi: [10.1016/j.addma.2021.102370](https://doi.org/10.1016/j.addma.2021.102370).
- [9] Department of Defense, “MIL-HDBK-1823A: Nondestructive Evaluation System Reliability Assessment.” Apr. 07, 2009.
- [10] M. Laleh *et al.*, “A critical insight into lack-of-fusion pore structures in additively manufactured stainless steel,” *Addit. Manuf.*, vol. 38, p. 101762, Feb. 2021, doi: [10.1016/j.addma.2020.101762](https://doi.org/10.1016/j.addma.2020.101762).
- [11] H. Masuo *et al.*, “Influence of defects, surface roughness and HIP on the fatigue strength of Ti-6Al-4V manufactured by additive manufacturing,” *Int. J. Fatigue*, vol. 117, pp. 163–179, 2018, doi: [10.1016/j.ijfatigue.2018.07.020](https://doi.org/10.1016/j.ijfatigue.2018.07.020).
- [12] A. Ziabari *et al.*, “Enabling rapid X-ray CT characterisation for additive manufacturing using CAD models and deep learning-based reconstruction,” *Npj Comput. Mater.*, vol. 9, no. 1, p. 91, May 2023.
- [13] A. Ziabari *et al.*, “Simurgh: A Framework for Cad-Driven Deep Learning Based X-Ray CT Reconstruction,” in *2022 IEEE International Conference on Image Processing (ICIP)*, Bordeaux, France: IEEE, Oct. 2022, pp. 3836–3867. doi: [10.1109/ICIP46576.2022.9898017](https://doi.org/10.1109/ICIP46576.2022.9898017).
- [14] O. Rahman *et al.*, “Neural Network-based Single-material Beam-hardening Correction for X-ray CT in Additive Manufacturing,” presented at the 17th International Meeting on Fully 3D Image Reconstruction in Radiology and Nuclear Medicine, Stony Brook, New York, 2023, pp. 324–327.





