

Report on Progress of Correlation of In-Situ and Ex-Situ Data and the Use of Artificial Intelligence to Predict Defects



Luke Scime
James Haley
William Halsey
Alka Singh
Michael Sprayberry
Amir Ziabari
Vincent Paquit

September 2020

M2TC-20OR0403017

DOCUMENT AVAILABILITY

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

Website www.osti.gov

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

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

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

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

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

Transformational Challenge Reactor

**REPORT ON PROGRESS OF IN-SITU AND EX-SITU DATA AND THE USE OF
ARTIFICIAL INTELLIGENCE TO PREDICT DEFECTS**

Luke Scime
James Haley
William Halsey
Alka Singh
Michael Sprayberry
Amir Ziabari
Vincent Paquit

Date Published: September 2020

M2TC-20OR0403017

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

CONTENTS

LIST OF FIGURES	v
LIST OF TABLES	vi
ACRONYMS	vii
ABSTRACT	1
1. INTRODUCTION	1
2. DIRECT OBSERVATION OF DEFECTS	2
2.1 DIRECT ANOMALY DETECTION	2
2.2 LASER MISALIGNMENT	4
2.3 OXIDATION BANDING	6
2.4 REGISTRATION OF IN-SITU AND EX-SITU DATA	7
3. BULK MACHINE LEARNING CORRELATIONS	8
3.1 FEATURE VECTORS	8
3.2 MACHINE LEARNING APPROACH	8
3.3 PROCESS OUTCOMES ACROSS A L-PBF PRINT AREA	9
4. LOCALIZED MACHINE LEARNING CORRELATIONS	11
4.1 EX-SITU TESTING PROCEEDURE DEVELOPMENT	11
4.2 PRELIMINARY EX-SITU TESTING RESULTS	12
4.3 FEATURE VECTORS	17
4.4 MACHINE LEARNING APPROACH	17
5. A PATH FORWARD	20
5.1 TESTING SPECIMENS AT SCALE	20
5.2 DOWN-SELECTION OF TESTING PROCEDURES	21
5.3 IMPROVEMENTS TO L-PBF IN-SITU SENSING CAPABILITIES	21
5.4 IMPROVEMENTS TO THE MACHINE LEARNING MODELS	21
5.5 MACHINE LEARNING UNCERTAINTY QUANTIFICATION	21
6. CONCLUSION	22
APPENDIX A. TCR DIGITAL PLATFORM	A-1
A-1. THE DIGITAL PLATFORM	A-1
A-2. DATA STORAGE INFRASTRUCTURE AND DATABASE ARCHITECTURE	A-5
A-3. DIGITAL TOOL (WEB INTERFACE AND APIs)	A-5
A-4. PEREGRINE	A-8
A-5. RAVEN	A-13
A-6. SCOPS	A-13
A-7. PIGEON	A-14
A-8. SIMURGH	A-15
A-9. REFERENCES	A-19

LIST OF FIGURES

Figure 1. Powder coverage is evidently incomplete for parts P7, P8, and P12.	2
Figure 2. Plot of super-elevation detections throughout the height of the build.	3
Figure 3. The misprints are highlighted in orange by Peregrine in the left image and are visible ex-situ on the as-built part in the right image.	3
Figure 4. 3D reconstruction of the entire build showing the recoater streaking anomaly (blue) intersecting with the part identified as P3.	4
Figure 5. An image taken after the completion of the build; observe the witness line located approximately one third up the height of the build.	5
Figure 6. A reconstruction of the neural network pixel classifications, in which the witness line is also visible.	5
Figure 7. (left) Confirmation that the laser shift remains constant throughout the remainder of the build, (middle) the more pronounced shift toward the rear of the M2, and (right) the less pronounced shift toward the front.	6
Figure 8. A ConceptLaser M2 build with discoloration bands evident at various points in the build height.	7
Figure 9. Ex-situ XCT data registered to the in-situ printer coordinate system.	8
Figure 10. Post-build image of some of the pentagonal samples still attached to the build plate.	9
Figure 11. Clustering of in-situ feature vectors.	10
Figure 12. Clustering of the in-situ feature vectors.	10
Figure 13. Clustering of the in-situ feature vectors.	11
Figure 14. Diagram of a standard SS-J3 mechanical testing specimen (left) and a photo of a mechanical test apparatus (right).	12
Figure 15. (a) Side view of a build with a CAD model of samples overlaid and (b) an isometric view of the same CAD model.	12
Figure 16. Mechanical yield tensile strength (MPa) test data separated by tensile specimen build height and reported as a function of distance from the center of the as-built part.	14
Figure 17. Percent elongation test data separated by tensile specimen build height and reported as a function of distance from the center of the as-built part.	15
Figure 18. Ultimate tensile strength (MPa) test data separated by tensile specimen build height and reported as a function of distance from the center of the as-built part.	16
Figure 19. Each datapoint represents one of the 117 tensile specimens contained within the testing dataset.	18
Figure 20. These results were produced using variation #2 of the ML model shown in Figure 19.	19
Figure 21. These results were produced using variation #3 of the ML model shown in Figure 19.	19
Figure A-1. An overview of the Digital Platform.	A-1
Figure A-2. A representation of an example Digital Thread.	A-2
Figure A-3. A simplified Digital Workflow for the manufacture of a silicon carbide fuel element.	A-3
Figure A-4. Specimen numbering scheme following the ASTM 52900:2015(E) standard [1].	A-4
Figure A-5. Representation of the microservices-based architecture utilized by the Digital Tool.	A-6
Figure A-6. An example metadata entry form for an Operation.	A-7
Figure A-7. An example outcome of a search action.	A-8
Figure A-8. An example of the image viewer feature.	A-8
Figure A-9. Overview of Peregrine’s capabilities and operation.	A-10
Figure A-10. Peregrine’s internal data flow.	A-11
Figure A-11. (a) SCOPS imaging system, two stereo high-resolution visible cameras and an infrared camera, (b) visible trackable surface roughness of representative deposition, (c) infrared imaging showing thermal dissipation, (d) 3D map constructed from surface	

features, (e) tracked deformation induced from a combination of residual strain and thermal expansion from (f) 3D mapped infrared signal showing hot regions.	A-14
Figure A-12. Block diagram of the Simurgh framework.....	A-17
Figure A-13. Example test results on a synthetic data set.	A-18
Figure A-14. Example test results on real data set.	A-18

LIST OF TABLES

Table 1. Descriptive statistics derived from 353 tensile samples extracted from a single as-built L-PBF AM SS316L part.	17
--	----

ACRONYMS

AI	augmented intelligence
AM	additive manufacturing
API	application programming interface
BH	beam hardening
CAD	computer-aided design
CVI	chemical vapor infiltration
DED	directed energy deposition
DIC	digital image correlation
DL	deep learning
DSCNN	dynamic segmentation convolutional neural network
DSCNN-P	DSCNN-Perceptron
FDK	uncorrected standard analytical
FPGA	field programmable gate array
GPU	graphics processing unit
IoT	Internet of Things
LPBF	laser powder bed fusion
MBIR	model-based iterative reconstruction
MDF	Manufacturing Demonstration Facility
ML	machine learning
MP	megapixel
MWIR	mid-wave infrared
ORNL	Oak Ridge National Laboratory
QR	quick response [code]
SS316L	stainless steel 316 low-carbon
STL	standard triangle language
TCR	Transformational Challenge Reactor
TIP	Technology Innovation Program
UTS	ultimate tensile strength
XCT	x-ray computed tomography
YTS	yield tensile strength

ABSTRACT

The Transformational Challenge Reactor (TCR) program is leveraging additive manufacturing (AM) technologies to fabricate nuclear components which will be assembled into a fully functional microreactor core. Compared with traditional manufacturing technologies, AM technologies allow (1) real-time observation of the manufacturing process at a much higher resolution using in-situ monitoring technologies to capture the sensor signatures that scientifically describe each event occurring over time and space and (2) validation of the manufacturing process quality using domain-informed data analytics techniques as a potential qualification and certification methodology for the final component.

This report provides an update on the program work on in-situ and ex-situ data correlation and associated data analytics results. Examples are provided to illustrate progress with respect to laser powder bed fusion (L-PBF), binder jetting, computed tomography (CT) reconstruction, and mechanical testing. Elements of the Digital Thread and data management infrastructure are discussed in the main document, and an extensive supplemental appendix is provided detailing the Digital Platform, as well as its implementation and subcomponents. In conclusion, the path forward for the next fiscal year is also discussed.

1. INTRODUCTION

The ultimate goal of the TCR Digital Platform is to enable determination of meaningful correlations between in-situ sensing data and ex-situ reactor component properties, thereby demonstrating the potential of the “born qualified” and advanced manufacturing approaches in the nuclear application space. As a corollary, it is necessary for a large proportion of the Digital Platform infrastructure (hardware and software) to be in place before such correlations can be determined. This fiscal year (FY20) has focused on developing this infrastructure and prototyping certain high-risk aspects of the correlation workflow. Infrastructure developments are fully described in the Digital Platform supplementary material and other FY20 TCR reports. This report demonstrates the usage of various aspects of the Digital Platform infrastructure as they relate to these correlation tasks. At the end of this report, a roadmap is presented detailing the planned FY21 efforts in this space.

The TCR program recognizes that several different types or levels of correlation are possible. These correlation levels approximately correspond to the four types of augmented intelligence (AI) algorithms—descriptive, diagnostic, predictive, and prescriptive—as defined in the supplementary Digital Platform material. This report considers three correlation levels: direct observation of defects, bulk correlations driven by machine learning (ML), and localized correlations driven by ML.

Direct observation of defects can be performed using only descriptive and diagnostic AI algorithms. The TCR team can report several correlations based on direct observation for FY19 and FY20. Bulk correlations necessitate at least a simple predictive AI algorithm. These correlations summarize the in-situ data for an entire component and attempt to correlate those data with bulk properties measured ex-situ. Because this level of correlation does not consider component geometry, it is not strictly generalizable, and therefore it is not a primary target of the TCR program. Instead, the bulk correlations presented in this report should be considered as prototypes and proofs of concept for the localized correlation efforts. Localized correlations also require a predictive AI, and they attempt to correlate local in-situ data with some representation of local ex-situ component properties. As these correlations will address component geometry, they are more generalizable, so they remain the focus of the TCR data team in FY21.

In addition to infrastructure requirements, the existential challenge to implementing a localized correlation AI model is ex-situ data volume. The ML models being designed for this task require on the order of thousands of data points to be properly trained. For example, if the model is to be trained to

predict component yield strength based on the in-situ data, then thousands of tensile tests must be performed and tracked via Digital Threads within the Digital Platform. Furthermore, these tensile specimens must be diverse to accurately represent the variability expected in the full population of related reactor components. This report highlights the success of this approach in very specific cases, and it paves the way for next year's parameter space exploration campaign which is needed to develop the proposed generalizable approach to geometry-agnostic certification and qualification of additively manufactured nuclear components.

2. DIRECT OBSERVATION OF DEFECTS

2.1 DIRECT ANOMALY DETECTION

For reasons which are not currently entirely understood, the ConceptLaser M2 machine occasionally spreads a powder layer incorrectly. This can result in incomplete powder coverage of one or more parts in the build, as shown in Figure 1. This anomaly can be directly detected by Peregrine; for example, many instances of this event are detected in the build shown in Figure 2. Because these anomalies change the effective powder layer thickness in subsequent layers, the solidification behavior of the melt pools, and hence the corresponding microstructure and pore populations, are expected to be different than that seen under nominal conditions. Correlating these detected anomalies to specific microstructural features is a goal for FY21.




Figure 1. Powder coverage is evidently incomplete for parts P7, P8, and P12.




Figure 2. Plot of super-elevation detections throughout the height of the build. Each spike in super-elevation corresponds to an improper powder spread which intersects with part P7.

Misprints on the ConceptLaser M2 machine are directly detectable via Peregrine analysis. During a misprint, material is fused in a location not specified by the part designers nor indicated by the 3D model of the part. A direct comparison between Peregrine’s fused powder detections and the registered CAD information makes such a flaw highly evident, as shown in Figure 3. Detection of these defects was first demonstrated in FY19.




Figure 3. The misprints are highlighted in orange by Peregrine in the left image and are visible ex-situ on the as-built part in the right image.

Contamination of the recoating roller on the ExOne Innovent Binder Jet systems can cause recoater streaking anomalies. Because the powder layer is not spread properly in these regions, recoater streaking is expected to be detrimental to the density and mechanical properties of the silicon carbide components, although this has not yet been directly verified. Improvements in data collection and analysis in FY20 have made reliable detection of these potential defects possible; Figure 4 shows an example of a recoater streaking anomaly intersecting with a printed silicon carbide fuel element.




Figure 4. 3D reconstruction of the entire build showing the recoater streaking anomaly (blue) intersecting with the part identified as P3.

2.2 LASER MISALIGNMENT

Shifts in the laser alignment on the ConceptLaser M2 can be directly detected using Peregrine-driven analysis of the in-situ imaging data. Laser shifts may occur when a build is paused and restarted, as occurred for the tubes shown in Figure 5, after approximately 50 mm of build height. This “witness line” is eminently visible ex-situ and can also be observed in the reconstruction of the in-situ data (Figure 6). Detection of these defects was first demonstrated in FY19.




Figure 5. An image taken after the completion of the build; observe the witness line located approximately one third up the height of the build.

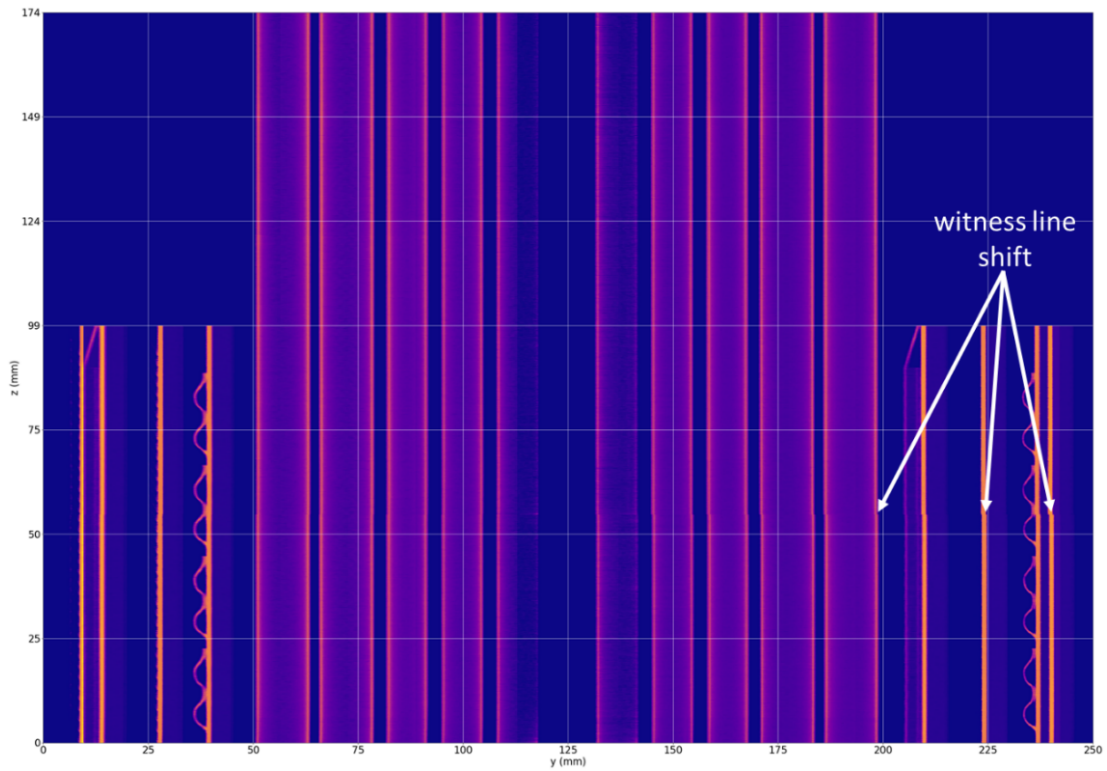


Figure 6. A reconstruction of the neural network pixel classifications, in which the witness line is also visible.

Closer inspection of the in-situ data (Figure 7) reveals that the laser shift is substantially more pronounced toward the rear of the M2 (laser module #1) than toward the front of the M2 (laser module #2). This contrast is also visually apparent ex-situ, where the front and rear of the M2 are left and right, respectively in Figure 5. Based on the in-situ data, the magnitude of the shift is approximately 500 μm in the rear and less than 100 μm in the front. It is also apparent that the shift is permanent, meaning that for the remainder of the build, the lasers remain offset from their original positioning at the beginning of the build.




Figure 7. (left) Confirmation that the laser shift remains constant throughout the remainder of the build, (middle) the more pronounced shift toward the rear of the M2, and (right) the less pronounced shift toward the front.

2.3 OXIDATION BANDING

Post-heat treatment, several bands of discoloration were visible on parts from a particular ConceptLaser M2 build (Figure 8). These bands directly correlate to fluctuations in the chamber oxygen content, as reported in the log file (Figure 8). During this build, the oxygen content fluctuated between 0.05 and 0.15%, which is atypical. Importantly, this correlation was identified by a technician using Peregrine to review the build data without direct support from the data team. Empowering the entire manufacturing team with the ability to identify and understand these types of correlations is absolutely critical for the success of the TCR program given the massive scale of the data to be produced during reactor fabrication.

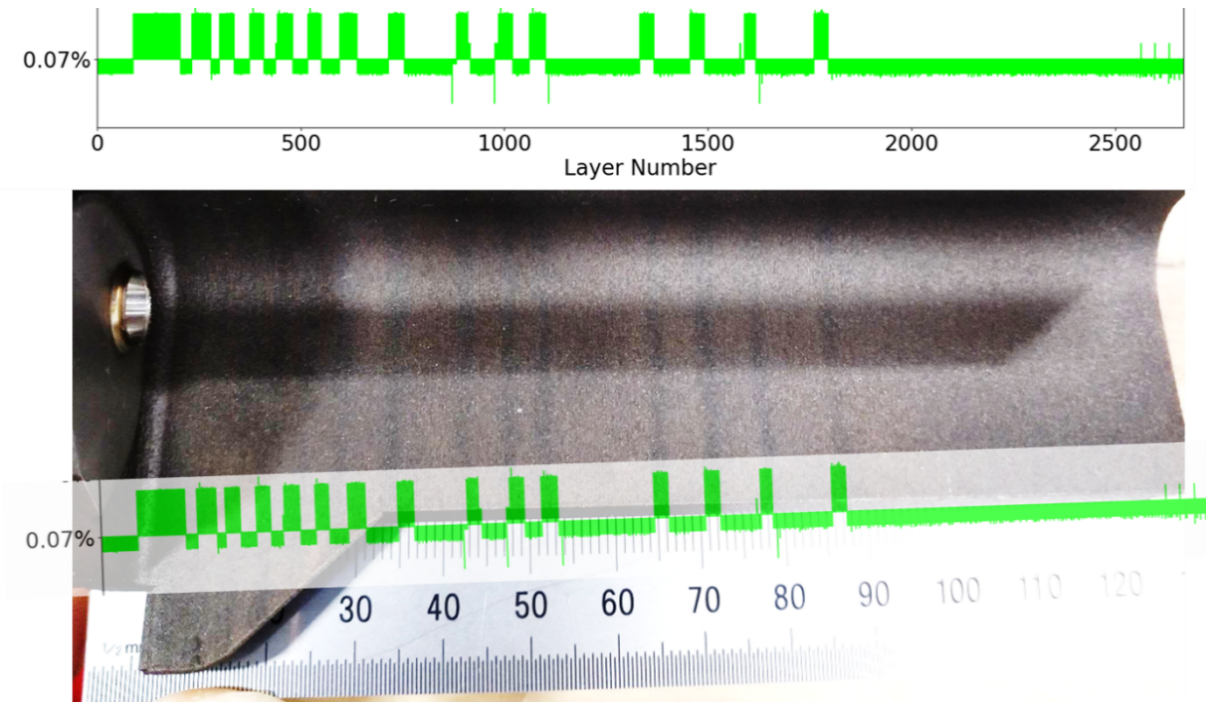


Figure 8. A ConceptLaser M2 build with discoloration bands evident at various points in the build height.
The chamber oxygen concentration data are manually overlaid on top of the image.

2.4 REGISTRATION OF IN-SITU AND EX-SITU DATA

In FY20, tools were added within Peregrine to allow for 2.5D co-registration between ex-situ x-ray computed tomography (XCT) data and the layer-wise imaging data and associated neural network anomaly detections. Figure 9 shows a demonstration of this co-registration technique applied to a build of test samples containing engineered porosity. This technique is expected to be fully leveraged in FY21. Furthermore, the use of Simurgh (Section A-8) to perform the XCT reconstruction using AI will dramatically improve the quality and utility of the XCT data. It is important to note that any correlations determined via this technique will be limited to features on the same size scale as the resolving power of the in-situ powder bed camera, regardless of the XCT scan resolution. This limitation is one of the motivations for increasing the imaging sensor resolution on the ConceptLaser M2 machine, as described in the FY20 TCR report entitled *Development of Monitoring Techniques for Laser Powder Bed Additive Manufacturing of Metal Structures*.




Figure 9. Ex-situ XCT data registered to the in-situ printer coordinate system. As a demonstration, yellow in-situ swelling detections are overlaid on top of the XCT data.

3. BULK MACHINE LEARNING CORRELATIONS

3.1 FEATURE VECTORS

A feature vector is a representative summary of a more complex data set. In the context of this report, feature vectors are designed to be ingested by predictive AI algorithms and used to predict ex-situ part properties. One way to summarize the in-situ Peregrine anomaly detections for a given part is to calculate the number of voxels (where the voxel size is defined by the camera resolution and the layer thickness) within that part which belong to each anomaly class. In this construction, a feature vector might look like $\{0.40, 0.35, 0.25\}$, indicating that 40, 35, and 25% of the voxels within that part belong to anomaly classes #1, #2, and #3, respectively.

This type of feature vector contains no information about the geometry of the part, yet it is well established that geometry-defect relationships play a critical role in determining part properties. Furthermore, such feature vectors ignore anomalies which are present near the part but do not lie directly on top of the part. For this reason, these types of feature vectors should only be used to compare and predict bulk properties for parts with identical geometries. This approach has been used as a prototype during development of the Digital Platform and the more advanced techniques described in the Section 4.3.

3.2 MACHINE LEARNING APPROACH

Feature vectors are *high dimensional*, meaning that they can only be fully represented in high dimensional space. Therefore, identifying correlations can be challenging. To address this issue, several different clustering algorithms are available which allow high dimensional vectors to be approximately represented in 2D space. The following subsections make use of the established t-distributed stochastic neighbor embedding (t-SNE) unsupervised ML algorithm for this task. Fundamentally, the t-SNE algorithm attempts to cluster bulk specimens with similar in-situ feature vectors together while separating samples with different feature vectors. It is important to note that all t-SNE plots are completely dimensionless, and the spacings between the clusters are highly nonlinear. Supervised ML techniques can also be applied to these feature vectors, but efforts in this work have focused on applying such supervised ML models only to the more localized feature vectors described in Section 4.3, as those relate directly to the ultimate goals of the TCR program.

3.3 PROCESS OUTCOMES ACROSS A L-PBF PRINT AREA

In FY20, identical pentagonal specimens (Figure 10) were fabricated on the ConceptLaser M2. Clustering of the bulk feature vectors demonstrates both the capabilities and limitations of this technique. The existence of clear clustering in Figure 11 indicates that the in-situ feature vectors can be used to detect variation in the parts correlated to their y position within the build chamber. Conversely, while minor variations in as-fabricated density were identified using ex-situ pycnometry, these variations do not appear to correspond to any of the in-situ feature clusters (Figure 12). These pentagonal samples were fabricated during two different builds. Comparison of the feature vectors from each build leads to several conclusions. Based on the clustering evident in Figure 13, it appears that the intra-build variation (in in-situ sensing signatures) is more significant than any inter-build variation—this result has positive implications for the laser powder bed additive manufacture of reactor components for TCR, as it suggests a relatively lower degree of variability between builds.



Figure 10. Post-build image of some of the pentagonal samples still attached to the build plate.




Figure 11. Clustering of in-situ feature vectors. Each dot corresponds to a single pentagonal specimen. Coloration indicates the y position of the pentagonal sample. In particular, samples with smaller y values (closer to the chamber door) have in-situ feature vectors which are distinct from the feature vectors of other specimens.




Figure 12. Clustering of the in-situ feature vectors. Each dot corresponds to a single pentagonal specimen. Coloration indicates the measured density of the as-fabricated specimen. The evident clusters do not appear to correlate with any trends in the ex-situ pycnometer data.




Figure 13. Clustering of the in-situ feature vectors. Each dot corresponds to a single pentagonal specimen. Coloration indicates the source build of the specimen.

4. LOCALIZED MACHINE LEARNING CORRELATIONS

4.1 EX-SITU TESTING PROCEEDURE DEVELOPMENT

In FY20, 586 AM SS316L samples from three different ConceptLaser M2 builds were mechanically tested. The lessons learned during these preliminary tests are being leveraged during planning of the full-scale testing campaign outlined in Section 5. As introduced earlier in this report, creating an AI model to estimate the mechanical properties of a part requires in-situ anomaly information, as well as large quantities of spatially correlated ex-situ testing data. Furthermore, because local part geometry can have a significant impact on the importance of a given defect, the part geometry must be represented computationally. To this end, it is critical that the spatial locations of each mechanical test specimen are carefully tracked within the Digital Platform. In FY20, standard SSJ3 tensile bars (Figure 14) were carefully extracted from builds using wire electrical discharge machining. The procedures necessary to maintain spatial tracking of such specimens were developed during this phase of the TCR program and will enable correlations between the measured mechanical properties and the voxel-wise in-situ sensing data. Figure 15 shows an efficient method for extracting several hundred SSJ3 tensile specimens from a single L-PBF build.




Figure 14. Diagram of a standard SS-J3 mechanical testing specimen (left) and a photo of a mechanical test apparatus (right).




Figure 15. (a) Side view of a build with a CAD model of samples overlaid and (b) an isometric view of the same CAD model.

Small ball punch testing was also performed on AM SS316L specimens during FY20. It was hoped that this form of mechanical testing would enable extremely rapid, localized measurements of part tensile properties. Unfortunately, initial results have not been promising for this use case. Specifically, AM parts with apparent major defects were tested, yet the mechanical properties measured by the small ball punch tests showed limited differences relative to nominal AM SS316L material. As a result, small ball punch testing may not be pursued in FY21; more details about the future ex-situ testing plan can be found in Section 5.

4.2 PRELIMINARY EX-SITU TESTING RESULTS

The tensile testing results for a subset of the specimens shown in Figure 15 are reported in Figures Figure 16–Figure 18 using standard box plots. Figure 16 reports the measured yield tensile strength (YTS), with the data subdivided by the specimen build height. The specimens analyzed in the top-left subplot were printed closest to the build plate, whereas the specimens in the bottom subplot were printed close to the

top of the as-built part. The distance of each test specimen from the centerline (as measured along the build chamber y -axis Figure A-4) of the as-built part is reported on the x -axis. Interestingly, an increased YTS is observed for the specimens extracted from nearer to the as-built part's edge, which is further away from the centerline. Note that the measured variation in the YTS is quite significant, with a variance of 100 MPa.

Figure 17 reports the measured elongation, with the data subdivided by the specimen build height. The x -axis indicates the distance of each specimen from the centerline of the as-built part. No clear trends with respect to build height or centerline distance are exhibited for elongation, and variation between the specimens is relatively minimal.

Figure 18 reports the measured ultimate tensile strength (UTS), with the data subdivided by the specimen build height. The x -axis indicates the distance of each specimen from the centerline of the as-built part. While there is no clear trend with respect to centerline distance, a significant drop in mean UTS is observed for the specimens printed closest to the top of the as-built part. Additional mechanical testing must be performed to verify and understand this observation.

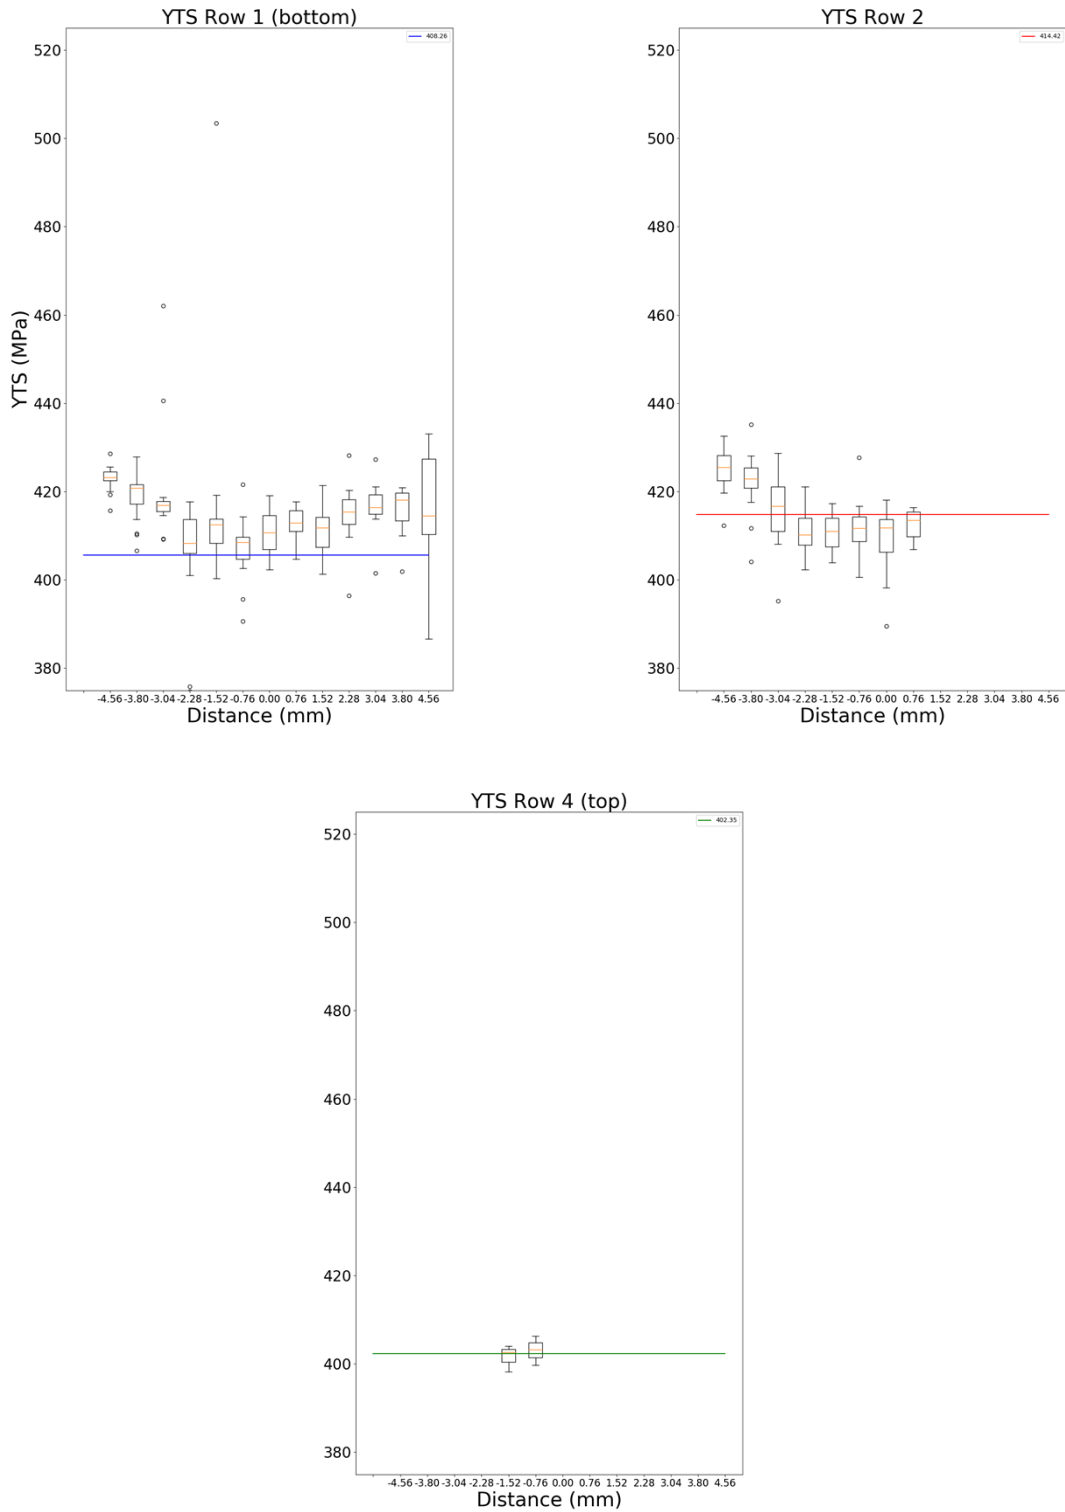


Figure 16. Mechanical yield tensile strength (MPa) test data separated by tensile specimen build height and reported as a function of distance from the center of the as-built part. The horizontal lines denote the mean of the specimens with planar positions common to all three build heights.

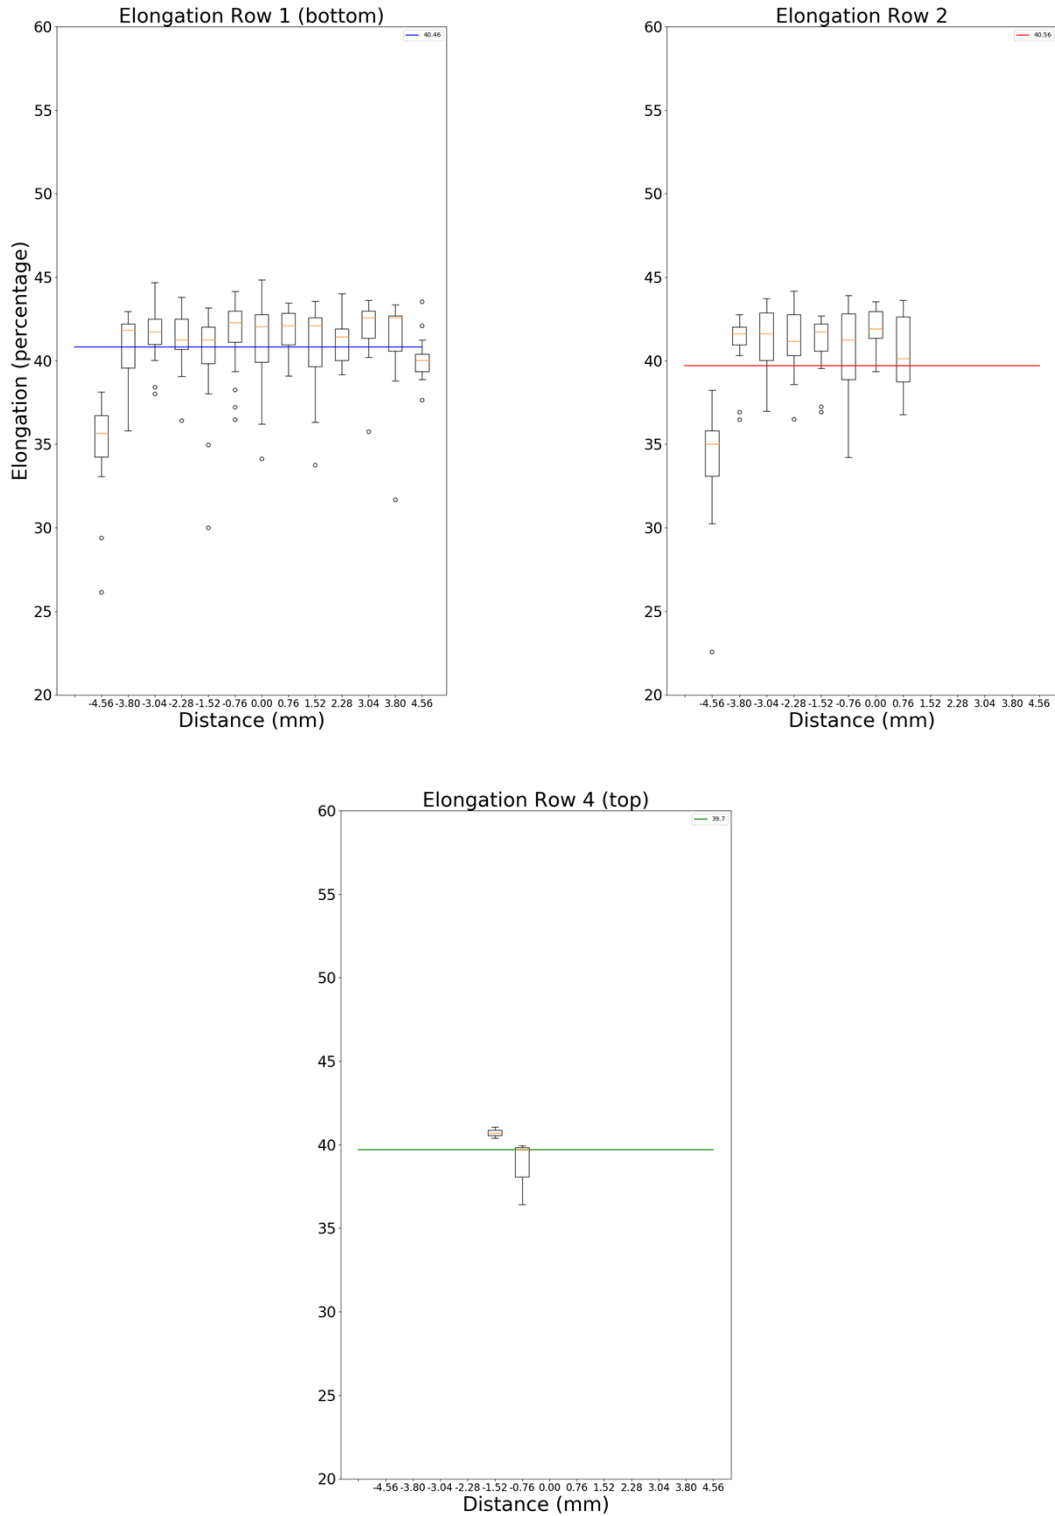


Figure 17. Percent elongation test data separated by tensile specimen build height and reported as a function of distance from the center of the as-built part. The horizontal lines denote the mean of the specimens with planar positions common to all three build heights.

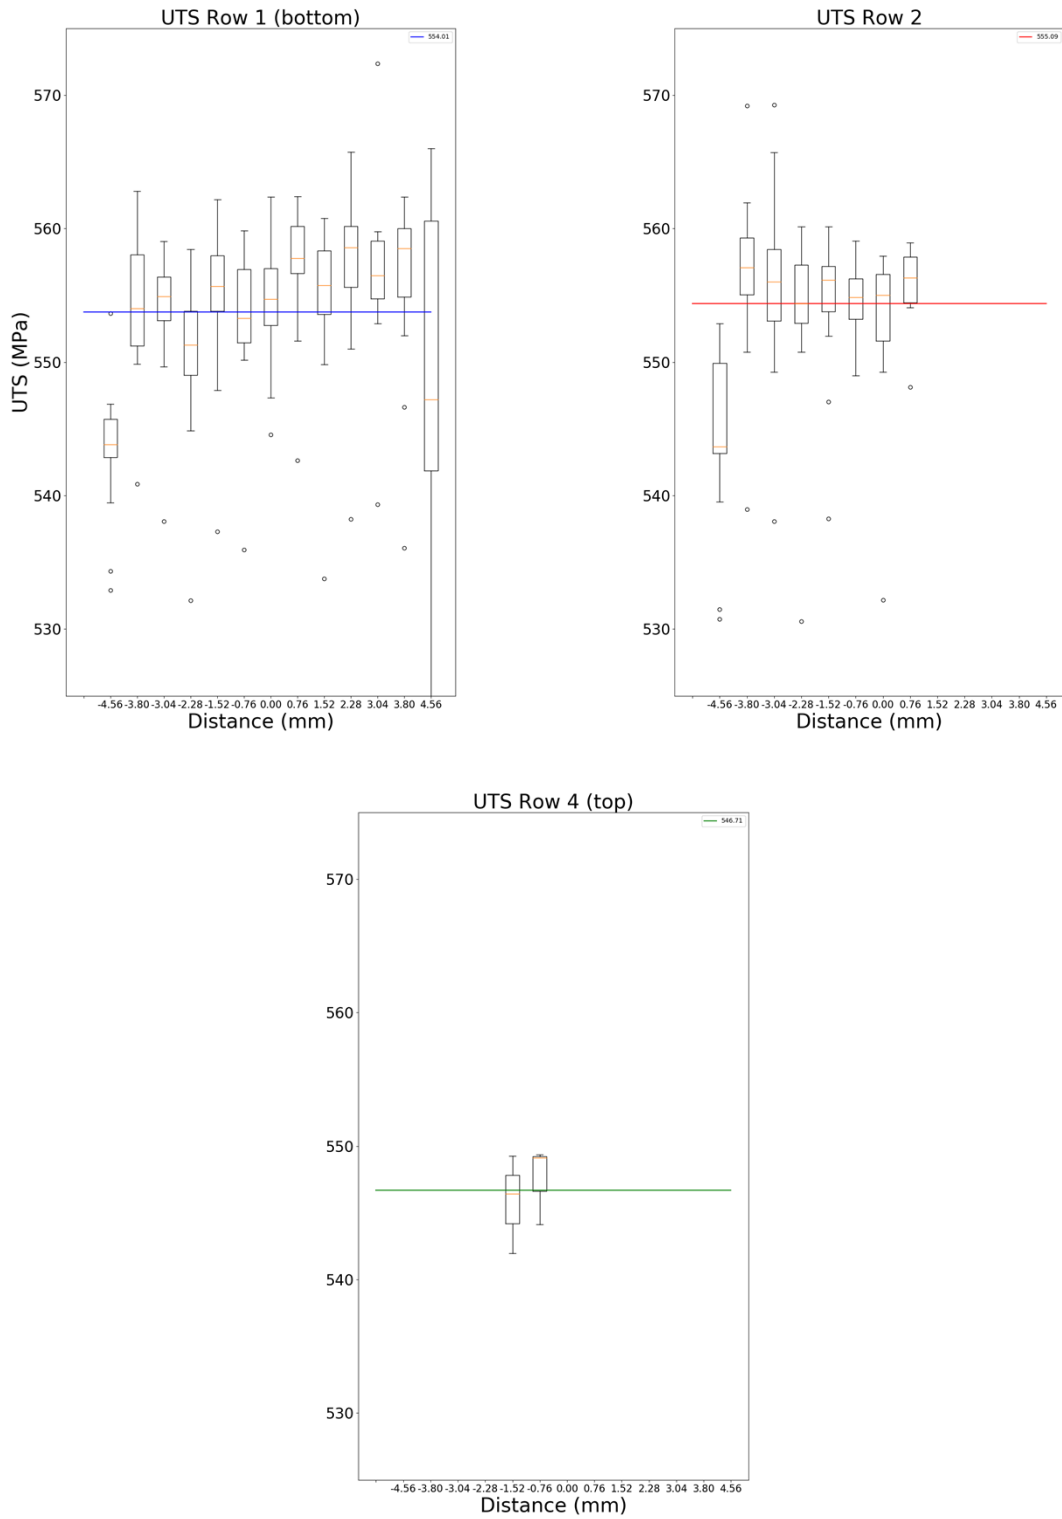


Figure 18. Ultimate tensile strength (MPa) test data separated by tensile specimen build height and reported as a function of distance from the center of the as-built part. The horizontal lines denote the mean of the specimens with planar positions common to all three build heights.

The results from the preliminary tensile testing campaign are summarized in Table 1. It is worth reiterating that at this time these trends are only valid for this particular build and these specific processing conditions. The testing plan detailed in Section 5.1 will determine if these trends are generalizable across as-built L-PBF AM SS316L components.

Table 1. Descriptive statistics derived from 353 tensile samples extracted from a single as-built L-PBF AM SS316L part.

Property	Mean	Variance
Yield tensile strength (MPa)	414	100
Percent elongation (%)	40.4	9.14
Ultimate tensile strength (MPa)	553	52.8

4.3 FEATURE VECTORS

The bulk feature vectors introduced in Section 3.1 have several limitations, namely that they do not encode (i.e. describe) the local part geometry, so they cannot generalize to diverse part geometries or provide localized part property predictions. As a result, a new feature vector structure was developed in FY20 to address these limitations. This feature vector design is a prototype and is expected to be iterated throughout FY21.

Instead of summarizing the Peregrine data for an entire part, these feature vectors use a Gaussian kernel to compute a weighted average of Peregrine anomaly detections in the region surrounding each voxel. Therefore, each voxel within each part will have a unique feature vector instead of the entire part being represented by a single feature vector. This has the added advantage of encoding information about the region beyond the nominal part boundary: for example, spatter can only be reliably detected in the powder surrounding a part, but it may be highly correlated to the energy density used to fabricate that part. Information about the local part geometry is also encoded by the feature vectors. Specifically, the distances of each voxel from the nearest part edge, as well as from the top and bottom part surfaces, are captured. These geometry metrics were chosen based on the TCR team’s physical understanding of the powder bed additive manufacturing processes.

The TCR data team has also developed methods for processing and spatially mapping temporal data from the machines’ embedded sensors. These sensors report fluctuations in environmental conditions and the status of machine subsystems. Additional detail describing the spatial mapping process are included in the FY20 TCR report entitled *Development of Monitoring Techniques for Laser Powder Bed Additive Manufacturing of Metal Structures*. Such fluctuations across layers have already been shown to affect physical attributes of the parts (as seen in Figure 8), and the ability to spatially map these fluctuations within a layer will allow these data to be included in the local feature vectors.

Because these feature vectors are designed by humans, they may not be optimized as they would be in a deep learning (DL) algorithm. However, by designing the feature vectors a priori, the number of ex-situ tests which must be performed to train the ML model is reduced by many orders of magnitude. A current limitation of this design is a rather large computer memory requirement during analysis of a full-scale build. Efforts in FY21 will include reducing or accommodating the memory footprint and iterating on the design of the feature vectors themselves.

4.4 MACHINE LEARNING APPROACH

In FY 20, a preliminary ML model was developed to predict mechanical part properties based on the localized feature vectors described in Section 4.3. This model topology is a prototype and is expected to

be iterated throughout FY21 as increasingly large amounts of training data become available. This prototype approach is a general linear model using a Gaussian family distribution.

The prototype model was trained on data from 235 of the tensile specimens introduced in Section 4.1. Once trained, the model was used to predict the YTS values of the remaining 117 tensile specimens. Figure 19 compares the predicted YTS values to the measured YTS values for these 117 specimens. If the model had made perfect predictions, then all of the datapoints would lie on the overlaid line. Figures Figure 20–Figure 21 show similar results given reasonable variations of the prototype model. The variation in the predictions amongst the presented models is one example of the uncertainty common to ML techniques. Uncertainty quantification is discussed in more detail in Section 5.5. In this particular situation, the in-situ Peregrine-detected classes which were most correlated with YTS include *powder*, *printed area*, *recoater streaking*, and *swelling*. It is worth reiterating that neither this ML model nor any of the resultant conclusions are generalizable beyond this specific as-built part in this specific AM build at this time.




Figure 19. Each datapoint represents one of the 117 tensile specimens contained within the testing dataset. The ML-predicted YTS values are reported on the x -axis, whereas the true measured YTS values are reported on the y -axis. If the ML model were perfect, then all of the datapoints would lie on top of the overlaid line.




Figure 20. These results were produced using variation #2 of the ML model shown in Figure 19.




Figure 21. These results were produced using variation #3 of the ML model shown in Figure 19.

5. A PATH FORWARD

With the Digital Platform infrastructure and testing procedures developed in FY20, the TCR digital team will focus on scaling up the ex-situ testing campaign in FY21. There will be a particular emphasis on designing build and specimen geometries to effectively induce, observe, and characterize AM-relevant in-process defects. As a corollary, the reactor-relevant part property metrics will be down-selected to a handful of scalable and appropriate ex-situ tests. Changes to the in-situ sensing systems on the L-PBF AM machines will be implemented and the in-situ feature vectors and ML models presented in Section 4 will be optimized. Importantly, uncertainty quantification of the AI techniques will also be considered. Each of these challenges is further discussed in the following subsections.

5.1 TESTING SPECIMENS AT SCALE

As discussed throughout this report, training a relay of AI models to robustly predict part properties based on in-situ data requires on the order of thousands of test specimens. Each of these specimens must be carefully tracked in both the physical and digital realms. Much of the efforts in FY20 revolved around developing the Digital Platform infrastructure necessary for managing thousands of these Digital Threads. Procedures for data registration across Operations were also developed. In FY21, more automated specimen tracking (e.g., printed quick response [QR] codes) and data registration techniques must be developed and implemented to handle the anticipated volume of test specimens.

Perhaps even more important than collecting the requisite number of test specimens is collecting a sufficient variety of specimens. Specifically, the specimens used for training must be representative of the true population. This is particularly challenging for AM, because process variability is extremely high, often stochastic, and many of the underlying mechanisms are not fully understood. The first priority of the TCR digital team in FY21 will be to create a specimen manufacturing plan designed to capture the widest range of reasonable process variability possible within the scope of the TCR program. It is expected that capturing a fully representative training dataset will not be possible, so additional algorithmic solutions are being considered (Section 5.5). A nonexclusive list of process variability considerations for L-PBF AM is presented below:

1. Layer times
2. Build heights
3. Location within the build chamber
4. Printing orientation
5. Laser module used for melting
6. Processing parameters
7. Time since machine maintenance and calibration
8. Feedstock characteristics and utilization history

Finally, it is well known that local part geometry can have a significant effect on AM in-process defect populations and their effects on final part properties. In particular, there are significant differences between “bulk” and “thin-wall” geometries as well as “down-skins” and “up-skins.” It can be particularly difficult to probe non-bulk regions of an AM build via mechanical testing. Doing so will require careful design of the as-printed parts, and it may require the use of nonstandard test specimen geometries. Another potentially viable approach is to directly test geometries analogous to those which will be used within the reactor core. Testing such non-bulk geometries will be another focus of FY21.

5.2 DOWN-SELECTION OF TESTING PROCEDURES

Because an extremely large number of samples must be fabricated and tested to train each ML model, the part properties of interest for the TCR program must be down-selected. For example, while YTS, fracture toughness, fatigue life, thermal conductivity, and corrosion resistance may all be of interest, only a subset of these properties can be reasonably investigated using AI. Because the selected properties will influence test specimen geometries and require differing amounts of testing time per specimen, this down-selection process must occur in tandem with the experimental design process outlined in Section 5.1.

For the purposes of developing the TCR Digital Platform, it is also important that the selected properties show significant variation over the expected AM processing window and that their underlying mechanisms can be directly or indirectly observed by the available in-situ sensing capabilities. To a first degree, this suggests that the selected properties should be defect sensitive. As an example, YTS is not expected to be particularly defect sensitive, but fracture toughness may be sensitive to various AM porosities. Microstructure sensitivity may also be viable, but this will depend strongly upon the post-build heat treatment chosen for the TCR AM SS316L components.

5.3 IMPROVEMENTS TO L-PBF IN-SITU SENSING CAPABILITIES

Improvements to the in-situ sensing suite on the ConceptLaser M2 AM machine will focus primarily on increasing imaging resolution and enabling additional direct defect detection capabilities. These planned changes are detailed in the FY20 TCR report entitled *Development of Monitoring Techniques for Laser Powder Bed Additive Manufacturing of Metal Structures*.

5.4 IMPROVEMENTS TO THE MACHINE LEARNING MODELS

The TCR digital team considers optimization of the localized property prediction ML models to be a relatively low-risk endeavor. Indeed, these types of ML models are well understood and relatively trivial to iteratively improve. However, proper optimization cannot begin until the rigorous testing campaign is underway, producing substantial quantities of high-quality and representative data.

5.5 MACHINE LEARNING UNCERTAINTY QUANTIFICATION

There are several potential sources of uncertainty in the predictions produced by ML algorithms. First, the hyperparameters, model topology, and learned weights can and are expected to influence the predicted outcomes. This form of uncertainty is relatively easily quantified and controlled using contemporary hyperparameter optimization techniques and Monte Carlo simulations. The second form of uncertainty is the result of how a human observer may interpret the output of the ML model. That is, if the model is a “black box,” then how does the human know whether or not to trust its predictions? While this is still an open research topic in the AI field, the TCR team plans to address this issue by selecting model topologies which allow humans to plainly observe which features within the feature vectors are most responsible for making a given prediction.

Finally, ML models are known to perform poorly in extrapolation (vs. interpolation) scenarios. Therefore, it is absolutely critical that the training datasets are representative of the true data populations. Unfortunately, as discussed in Section 5.1, ensuring that this condition is fully met is likely not possible within the scope of the TCR program. To ameliorate this issue, the TCR team proposes to implement a heuristic “circuit-breaker” within the model to first confirm that an interpolation is being performed and not an extrapolation. This pre-check will be implemented by comparing the feature vector of the new data for which a value is being predicted to all of the feature vectors contained within the training database. If this feature vector is substantially similar to previously observed feature vectors, then we expect to trust

the output of the ML model. If, however, the new feature vector is very different than the training data, the model will output “I do not know the answer” instead of outputting an extrapolated prediction. Determining a robust threshold for this heuristic will require experimentation and is expected to be a focus of the latter portion of FY21, continuing into FY22.

6. CONCLUSION

This report validates the TCR team’s conceptual data-driven methodology to correlate in-situ and ex-situ measurements for the purpose of ensuring quality control and providing certification for additively manufactured nuclear components. Multiple experiments and associated tests have been analyzed and interpreted to confirm this approach and to highlight gaps which must be addressed to attain a generalizable approach to certification. To address the later, a scientifically driven path forward has been established which relies on a series of design of experiments to produce high quality datasets. This will allow for exploration of the parameter spaces of the relevant additive manufacturing systems. At the end of the period of performance, a robust Digital Platform will be provided that will serve as the scientific foundation for certification of modern nuclear components produced via advanced manufacturing processes.

APPENDIX A. TCR DIGITAL PLATFORM

A-1. THE DIGITAL PLATFORM

The Digital Platform encompasses the design intent information, networking, and computing infrastructure, in situ sensing hardware, metrology and characterization processes, data storage and management systems, and software and analytics tools necessary to create and effectively leverage a Digital Thread (aka a *digital twin*) for each part and test specimen created for the TCR program. Components of the Digital Platform that have TCR-specific definitions are capitalized throughout this report for clarity. The backbone of the Digital Platform is a searchable Database linking data from component design, in situ sensors, characterization, user entries, machine calibration Timelines, and feedstock utilization Timelines. This information is physically stored on various Data Servers and is accessible via the Digital Tool web interface and associated application programming interfaces (APIs). The Digital Tool and APIs are hosted on the Database Server. Data are processed and visualized on Compute Servers or local user computers (e.g., desktops and laptops) using various ORNL-developed Software Tools; the Software Tools access the Database using standardized APIs. Figure A-1 provides an overview of the Digital Platform.

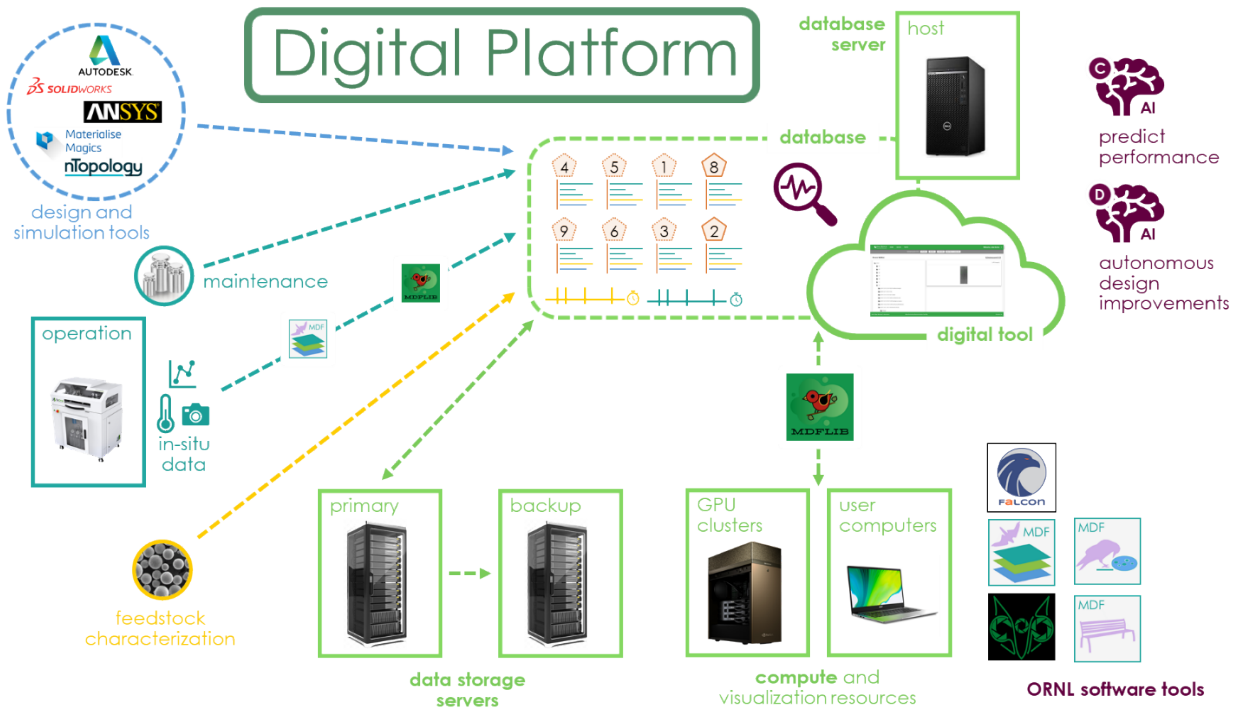


Figure A-1. An overview of the Digital Platform.

Each physical TCR component and test specimen will have a corresponding Digital Thread. Each Digital Thread is composed of building blocks referred to as Operations. Operations may be manufacturing processes, material characterization techniques, maintenance procedures, or purchases. Operations may act on a physical component, material feedstock (e.g., powder or wire), or a machine (e.g., a 3D printer). Example Operations include laser powder bed additive manufacturing (AM), annealing, wire electrical discharge machining, chemical vapor infiltration (CVI), scanning electron microscopy, tensile testing, purchasing a component from a supplier, and laser calibration. Some Operations may also instantiate new

parts or split an existing parent part into multiple child parts. The Digital Thread for a child part includes the thread for its parent.

Functionally, each Digital Thread is a list of Operations that have been performed on a given component and any corollary Operations that compose the machine maintenance and feedstock utilization Timelines. The Database stores these Digital Threads and allows users to access the data (e.g., in situ powder bed images) associated with each Operation. A combination of the Digital Tool web interface and various Software Tools is used to visualize and analyze the Operation data. Data from multiple Operations may be spatially co-registered and visualized using certain Software Tools. Figure A-2 shows a representation of an example Digital Thread.

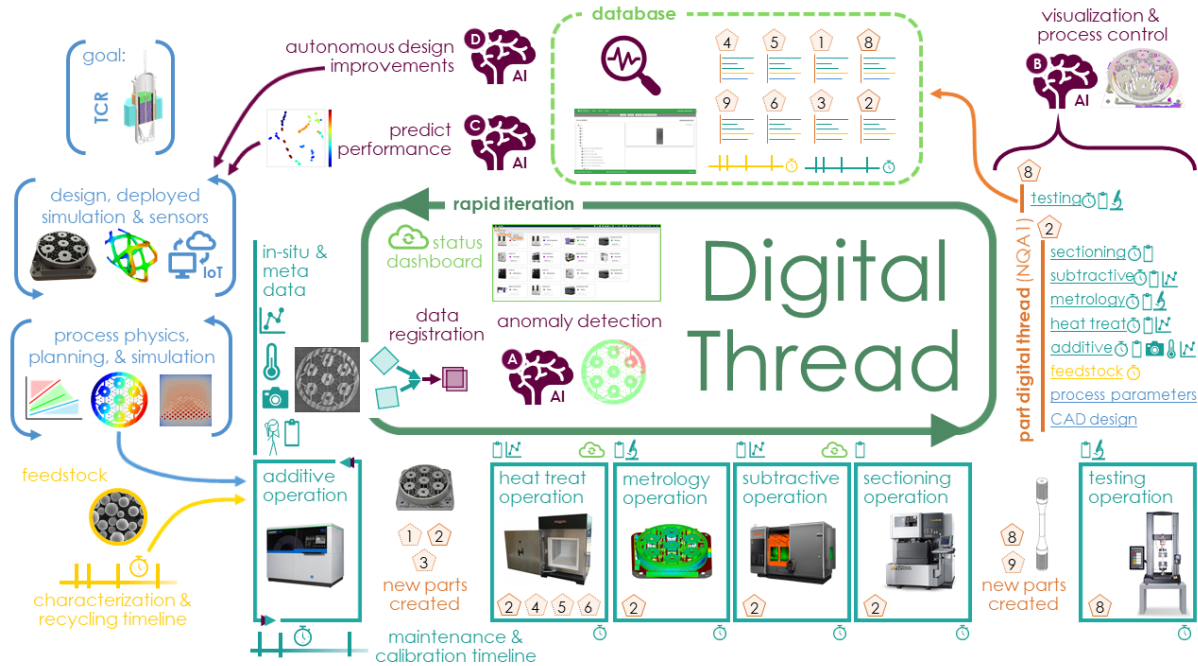


Figure A-2. A representation of an example Digital Thread. The process begins in the upper left-hand corner and follows the arrows counterclockwise.

At a more granular level, information is incorporated into a Digital Thread as part of the Digital Workflow. Many different types of users may interact with a given Digital Workflow, including reactor designers, advanced manufacturing engineers, machine technicians and operators, material scientists, data analysts, program managers, and public communications specialists. The composition of each Digital Thread, and therefore the corresponding Digital Workflow, varies dramatically, depending on the particular component. As an example, Figure A-3 shows a simplified Digital Workflow for a binder jet-printed silicon carbide reactor fuel element.

Digital Workflow

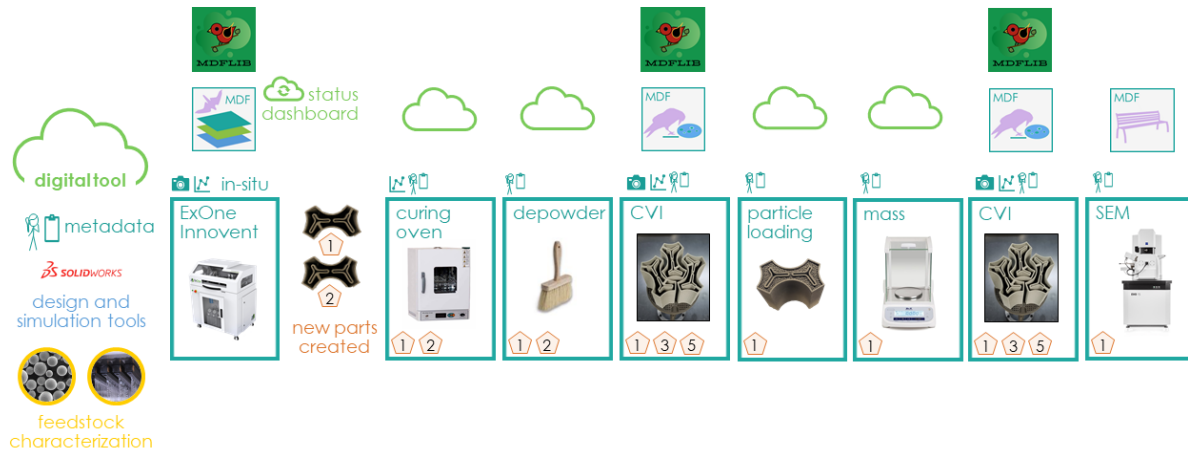


Figure A-3. A simplified Digital Workflow for the manufacture of a silicon carbide fuel element.

A primary goal of the TCR program is the prediction of reactor-relevant part properties based on in situ data collected during various advanced manufacturing processes. To this end, augmented intelligence (AI) algorithms will ingest collections of Digital Threads in order to identify correlations between in situ data and part properties. We deliberately refer to such algorithms as *augmented intelligence* instead of *artificial intelligence* to convey the necessity of keeping expert humans and their physics-based understanding of the world “in the loop” throughout these analyses. Critically, part property prediction will be accomplished through a relay of AI algorithms, with each hand-off leveraging human expertise to inspect the AI performance and to point the subsequent AI in the correct direction. This relay approach is essential to ingest the highly unstructured and extremely high dimensional in situ sensing data and to successfully link it to part property measurements without requiring a prohibitively large number of ex situ characterization experiments. To further clarify the terminology used throughout the TCR reports, *AI* is considered to be a broad umbrella of algorithms which includes heuristics, machine learning (ML), and deep learning (DL). The in situ/ex situ property correlation efforts for FY 2020 are described in a separate TCR report titled *Viability of Data Analytics to Ascertain Component Performance for Additive Manufacturing*. Consistent tracking of each reactor part and test specimen is a critical aspect of this approach. Development of a robust, integrated, and automatic specimen naming scheme was completed in FY 2020 and is illustrated in Figure A-4.





Figure A-4. Specimen numbering scheme following the ASTM 52900:2015(E) standard [1].

Importantly, AI is used for data analysis and visualization throughout the Digital Thread. AI usage for the TCR program can be approximately separated into four different application areas: *descriptive*, *diagnostic*, *predictive*, and *prescriptive*. Descriptive AI is used to analyze raw data, an example being segmentation of powder bed images and identification of process anomalies. Diagnostic AI summarizes large datasets and presents them to a human user or another algorithm. A diagnostic AI might flag certain layers of an AM Operation for closer inspection by a human, or it might decide on an autonomous process intervention to attempt to correct a detected defect. A predictive AI seeks to predict part properties (e.g., fracture toughness) based on in situ data, process parameter information, and part geometry. Demonstration of a predictive AI is an end goal of the TCR program. Finally, a prescriptive AI might autonomously modify a part design in order to improve the predicted performance; this class of AI is beyond the scope of the TCR program, but it is the ultimate goal for the Digital Platform.

In FY 2020, the Digital Platform transitioned from a conceptualization framework and a set of stand-alone tools into a highly flexible, scalable architecture with increasing integration between the Software Tools and the Digital Tool. The physical infrastructure needed to support the Digital Platform includes (1) a optical fiber-based dedicated network to connect all machines to the facility digital backbone, (2) dedicated Wi-Fi networks for Internet of Things (IoT) sensors, (3) redundant storage systems, and (4) multi-GPU (graphics processing unit) compute systems for AI model training and AI at the edge. FY 2020 also saw the standardization of the onboarding process for new Operations to be recognized by the Database. The remaining sections of this supplementary appendix describe the various components of the Digital Platform, including the ORNL-developed Software Tools. Changes and improvements made in FY 2020 are highlighted.

A-2. DATA STORAGE INFRASTRUCTURE AND DATABASE ARCHITECTURE

There is a diverse variety of data produced by a Digital Workflow during different Operations. The data produced vary with respect to the data source, data type, data format, and the scale of the data volume. The standardization of this wide variety of complex data is a challenging problem to solve. To address this complex data organization, the data are structured, organized, and stored on a dedicated Storage Server. This standardization facilitates users and services with improved data search, management, access, and retrieval functionality. The data are organized by the data source (i.e., Operation) and are further organized by the date of the Operation. The dedicated data Storage Server is used as the primary location for storing and accessing the data, whereas a secondary data Storage Server is used for monthly data backups. Some of the benefits that are achieved using a dedicated Storage Server for TCR data include the following.

1. The Data Storage Server acts as a centralized data repository, and the data are available to all the computers within the same network.
2. Multiple users can access the same dataset, eliminating the need to make duplicate copies of data for different users.
3. Data integrity is maintained because the latest datasets are centralized and accessible.
4. The server provides better data management and implementation of security protocols, and it allows data access protocols to ensure that TCR data can only be accessed by authorized users.

A Database Server was installed, and a relational database was developed to store and organize the metadata produced from Operations and associated with the Digital Threads. PostgreSQL was selected as the standard database protocol; PostgreSQL is an open-source, object-relational database system. It is one of the most widely used database systems because of its strong community support, scalability, reliability, and consistent performance. PostgreSQL also provides a robust set of features to support complicated Digital Workflows.

The purpose of the PostgreSQL Database is to facilitate metadata collection, data organization, and storage, along with search and retrieval on the datasets stored on the Data Storage Server. To achieve these requirements, the Database was designed so that relevant data are stored in individual schemas. Some of the schemas present in the current database are as follows.

1. Users: contains information about authorized users of the Database
2. Roles: contains information about which users have access to each Operation and dataset
3. Projects: contains information about different projects at the Manufacturing Demonstration Facility (MDF)
4. Machines: contains information about different manufacturing equipment
5. Materials: contains information about the feedstock materials and their manufacturers
6. Operations: contains information about Operations supported by the Digital Workflows
7. Parts: contains information about all manufactured specimens and the Operations which compose their associated Digital Threads

In addition to the schemas listed above, there are additional AM machine-specific schemas to store information about build datasets and the files in each dataset. These schemas contain metadata that populate the Digital Thread of each manufactured component.

A-3. DIGITAL TOOL (WEB INTERFACE AND APIs)

The Digital Tool is a web-based application developed in Vue.js at ORNL. Vue.js was selected for the development of this platform because of its small size, high performance, and excellent flexibility. Vue.js

also has detailed documentation available, and it supports microservices architectures for developing web applications. Microservices architectures support incremental and modular growth of the application, thus making it a suitable framework for developing scalable infrastructure. Figure A-5 illustrates a microservices architecture.

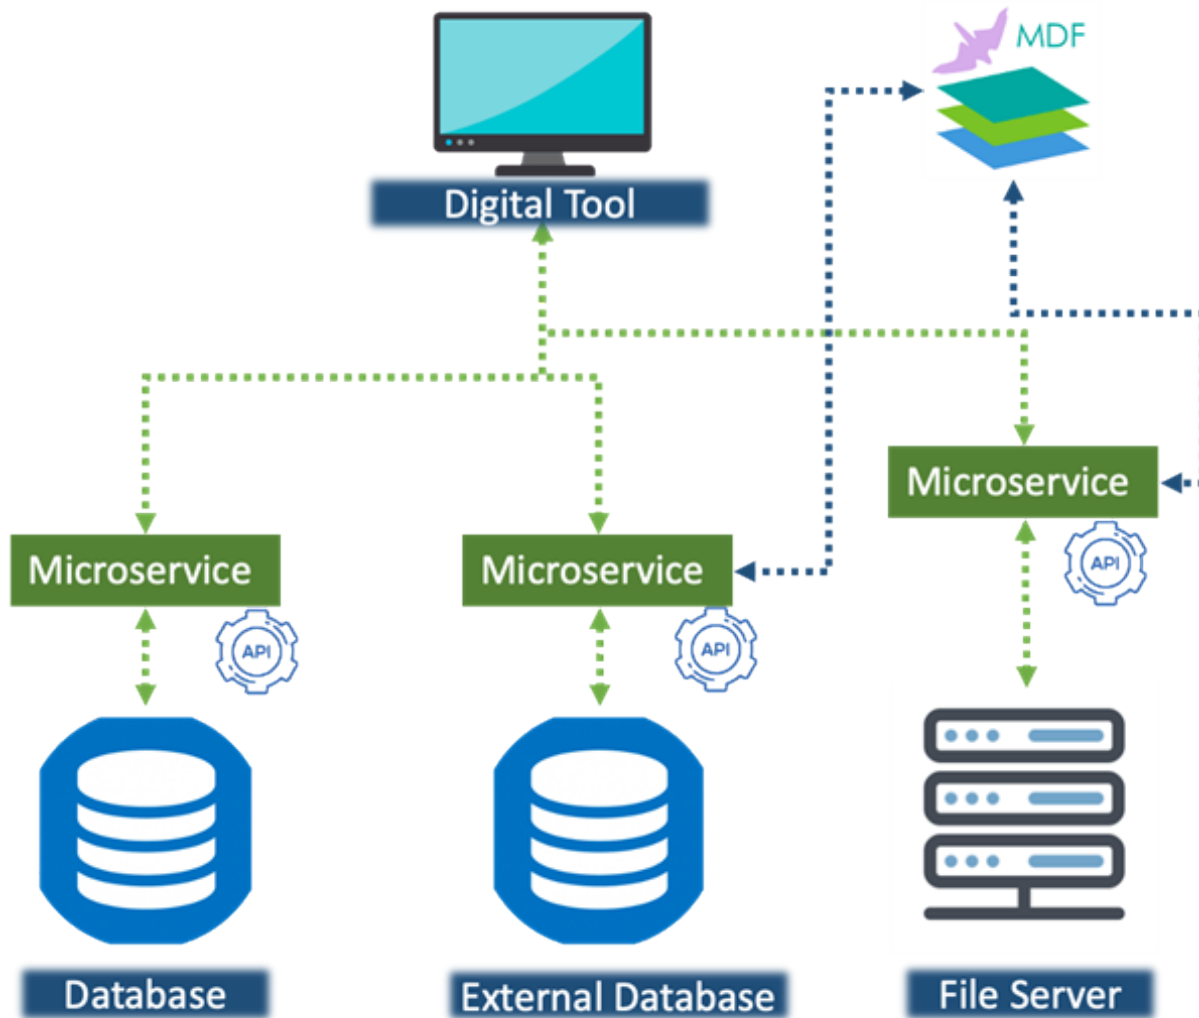


Figure A-5. Representation of the microservices-based architecture utilized by the Digital Tool.

The main components of the Digital Tool microservices architecture are as follows.

1. Digital Tool: the web application that is developed using JavaScript and html
2. Microservices: RESTful services developed in Python flask. These services run independently as a service and can communicate and pass information between data sources and the Digital Tool. Additionally, these services can be used by various Software Tools developed by the TCR data team. Currently, there are multiple independent microservices running which vary in functionality from querying databases, updating databases, and performing data transfer actions.

3. Data sources: microservices act as an interface between the applications and data sources. In this architecture, the data sources are the PostgreSQL Database and central Data Storage Server.

The Digital Tool provides online forms (Figure A-6) to help technicians record metadata before, during, or after an Operation. These metadata tracking forms were developed based on feedback from technicians and the ultimate data users. The forms are validated to ensure that the metadata collected are standardized, complete, and correct.

The screenshot shows a web application interface for recording build metadata. The header is green with the Oak Ridge National Laboratory logo and navigation links. The main content area is titled 'Build information:' and contains a form with three sections: 'Pre-build Info:', 'Post-build Info:', and 'Upload Files:'. The 'Pre-build Info:' section has several input fields: 'Build Name', 'Start Date' (with a date picker), 'Setup Technician', 'Startup Technician', and 'Powder Material Used'. There is also a toggle switch for 'Were template(s) used for this build?' set to 'No'. The 'Post-build Info:' and 'Upload Files:' sections are currently collapsed. A 'CREATE BUILD' button is located at the bottom of the form. The footer is green and contains the text 'Oak Ridge National Laboratory', 'Manufacturing Demonstration Facility', and 'Contact Us'.

Figure A-6. An example metadata entry form for an Operation.

The Digital Tool allows users to search (Figure A-7) the existing datasets based on the recorded metadata. This feature enables users to narrow down the desired datasets based on specific metadata, as well as less structured criteria such as the processing parameters used for an additive build.

The screenshot shows a web interface with a green header containing the Oak Ridge National Laboratory logo and navigation links (HOME, OPERATIONS, PARTS). A user is logged in as 'Alka Singh'. Below the header, there are buttons for 'UPLOAD', 'SEARCH', 'EXPLORE', 'WAITING TO ANALYZE', and 'STATS'. The main content area displays search results for 'Build(s): ConceptLaserM2-ORNL1'. A table lists various build actions with columns for Action, Name, Start Date, End Date, Status, Material, Setup Tech, Start Tech, and Was Test?. Below the table is a 'FILTER BY PARAMETER(S)' button and a pagination control showing 'Rows per page: 10' and '1-10 of 60'.

Action	Name	Start Date	End Date	Status	Material	Setup Tech.	Start Tech.	Was Test?
✓	Framatom Arch	2020-02-04	2020-02-04	Successful	316L/Praxair/27	Alka Singh	Alka Singh	Yes
✓	Airfoils & TCR Moderator Pieces	2020-02-07	2020-02-07	Successful	316L/Praxair/27	Alka Singh	Alka Singh	Yes
✓	Kairos Impeller	2020-02-12	2020-02-12	Successful	316L/Praxair/27	Alka Singh	Alka Singh	Yes
✓	MDF Framatome Fasteners 01	2020-02-26	2020-02-26	Successful	316L/Praxair/27	Alka Singh	Alka Singh	Yes
✓	Fastener Assembly	2020-02-06	2020-02-06	Successful	316L/Praxair/27	Alka Singh	Alka Singh	Yes
✓	Framatome Fastener Components	2020-02-14	2020-02-14	Successful	316L/Praxair/27	Alka Singh	Alka Singh	Yes
✓	TCR Moderator Pieces	2020-02-03	2020-02-03	Successful	316L/Praxair/27	Alka Singh	Alka Singh	Yes
✓	Framatom Middle Section	2020-02-05	2020-02-05	Successful	316L/Praxair/27	Alka Singh	Alka Singh	Yes
✓	Inner Mask Mold Bottom Section	2020-04-08	2020-04-08	Successful	316L/Praxair/27	Alka Singh	Alka Singh	Yes
✓	Theta Impeller and TCR Endcaps	2020-03-12	2020-03-12	Successful	316L/Praxair/27	Alka Singh	Alka Singh	Yes

Figure A-7. An example outcome of a search action.

The Digital Tool also allows users to explore the existing AM build datasets contained on the Data Storage Server. There are multiple data viewers (Figure A-8) to support exploring different data types including image viewers, STL viewers, PDF viewers, and log file viewers. This feature allows multiple users to access data simultaneously without requiring direct access to the stored data, thereby preventing accidental data modification.

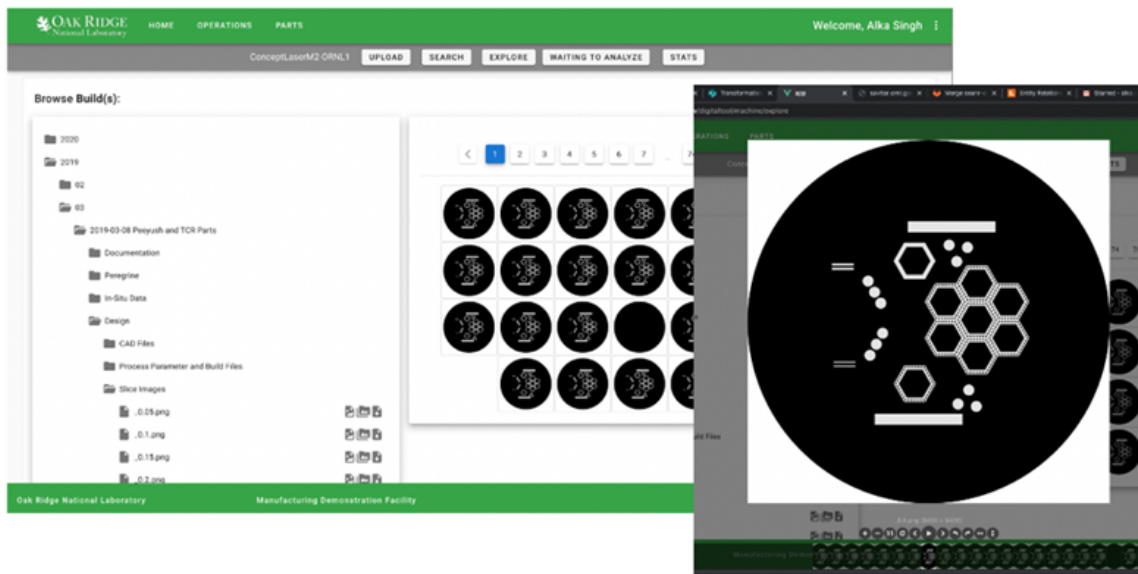


Figure A-8. An example of the image viewer feature.

A-4. PEREGRINE

Peregrine is a Software Tool developed in Python at ORNL. Peregrine is intended for use with powder bed AM machines (i.e., printers), including laser powder bed fusion (LPBF), electron beam powder bed

fusion (EB-PBF), and binder jetting technologies. In its fullest implementation, Peregrine autonomously collects and analyzes layer-wise imaging data, provides remote monitoring and process intervention capabilities, tracks metadata and part information, produces advanced visualizations of both the underlying data and the AI analysis results, and enables identification of correlations between in situ data and process parameters or ex situ measurements.

Peregrine's core capability is pixel-wise anomaly detection based on layer-wise images of the powder bed. This semantic segmentation of the data is achieved using a novel DL algorithm referred to as a *dynamic segmentation convolutional neural network* (DSCNN). The DSCNN architecture was designed specifically to overcome many of the challenges common to powder bed imaging data and is fully described in Scime et al. [2]. Beyond this core capability, Peregrine also empowers users with several advanced analytics and data tracking tools.

1. Each layer in a build can be flagged based on the DSCNN pixel segmentation results using either a set of user-defined heuristics (i.e., rules) or a learned ML model, referred to as the *DSCNN-Perceptron* (DSCNN-P).
2. Peregrine provides advanced data visualization, allowing the user to visualize the raw data, DSCNN, and DSCNN-P results in various formats. These visualizations include image overlays, plots of time series data, frequency analyses, specimen quality scoring, data projections in each orthogonal plane, 3D reconstructions, and time-lapse videos of the build.
3. Peregrine provides automatic specimen numbering and statistics. Up to two sets of computer-aided design (CAD) geometry information can be associated with each build, including the as-built geometries, which are referred to as *parts*, and the geometries of any *samples* removed from the components after the build is complete, such as a tensile bar. The term *specimen* is used by Peregrine to refer to either a *part* or a *sample*.
4. Metadata can be tracked for each build, along with reference images, log files (supported printers only), coaxial time series data (supported printers only), and registered ex situ data. Peregrine also includes 2.5D registration capability, allowing users to overlay ex situ data such as x-ray computed tomography (XCT) slice images on top of in situ data and DSCNN results.
5. Peregrine is capable of direct data collection and real-time analysis using most USB and ethernet cameras. During a *live* analysis, Peregrine can send alert emails to technicians, execute customizable macros to interface with a printer's user interface to effect process interventions, and provide remote viewing of the build status via a companion instance of Peregrine.
6. A limited number of statistics tools are also provided to enable the user to investigate potential correlations between Peregrine's in situ results and process parameters or ex situ measurements.

A high-level overview of Peregrine's capabilities is presented in Figure A-9, and Peregrine's internal data flow is summarized in Figure A-10. The figures illustrate bundling of all of the layer-wise imaging data into a single calibrated image stack, followed by a tiling operation before the data are fed into the DSCNN. The pixel classifications predicted by the DSCNN are then compared with the *template* image, and potentially the raw powder bed imaging data, using pixel heuristics. These comparisons create the final pixel classifications observed by the user. The output of the DSCNN is fed into the DSCNN-P to allow for *labeling* of the entire build layer, or *image stack*. Finally, the *labels* assigned by the DSCNN-P are combined with the global heuristics and presented to the user as layer flags.

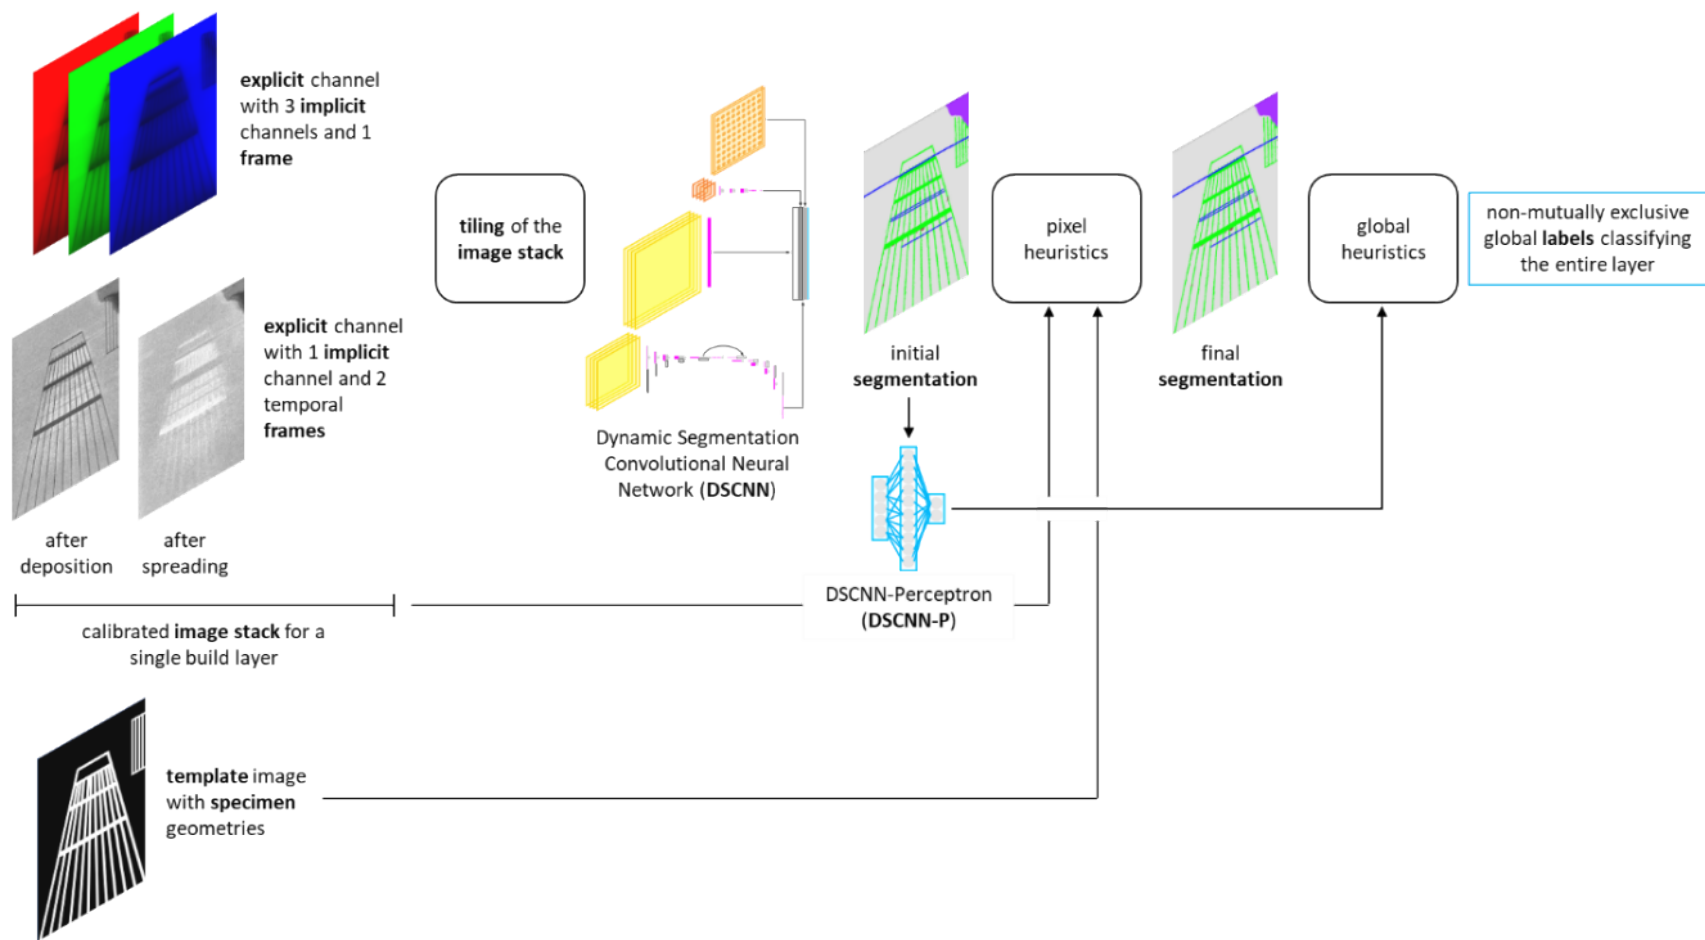


Figure A-10. Peregrine's internal data flow.

Peregrine has been under development since FY 2018. It is currently on major version 20 and is used by approximately 10 members of the TCR team. Users include data scientists, material scientists, and machine technicians. Over 100 pages of documentation and training materials have been written for Peregrine in FY 2020. In FY 2020, ORNL authored a journal paper describing the DSCNN that was published in *Additive Manufacturing*. Peregrine can now be licensed to companies, with over 10 companies already expressing interest. Below is a list of the major Peregrine improvements completed in FY 2020.

1. The software release process was formalized, including documentation, versioning, archiving, and bug reporting.
2. An immutable change log is now created for each build analysis to record all data edits and builds can now be marked as containing sensitive data.
3. Peregrine has been integrated into the Digital Platform, allowing data to be passed to and from the Database. Printer settings, calibrations, trained neural networks, and analyzed builds are hosted on the Data Storage Servers and are transparently accessible to all users at ORNL and those working remotely. Any trained user can directly update training data to improve AI model performance. Versioning improvements allow changes to be tracked as multiple users modify shared files.
4. A part numbering standard was introduced for the TCR program. Peregrine automatically numbers all *specimens* for each build. Automatic and consistent numbering is absolutely critical for tracking and maintaining the hundreds of thousands of Digital Threads to be stored on the Digital Platform.
5. The DSCNN architecture was upgraded to DSCNNv2. This architecture improves accuracy, segmentation of curvilinear structures, and segmentation of smaller structures, and it also reduces inference time and *tiling* artifacts. DSCNNv2 dramatically improves training flexibility and follows more industry standards for neural network design.
6. The entire data loading architecture was rewritten to allow for color images, 16-bit images, a truly arbitrary number of cameras, an arbitrary number of temporal *frames*, and more complex inter-layer lighting differences. These changes make Peregrine much more flexible as the imaging systems on several of the printers are improved in FY 2021.
7. XCT data can now be co-registered to the in situ imaging data and overlaid on top of the DSCNN predictions within Peregrine.
8. Data collected *live* are fully analyzed and accessible in real-time—this is critical for handling the massive TCR data burden.
9. The DSCNN-Perceptron can now be used to flag layers to bring them to the attention of a user or machine operator, and early results using this capability are promising. A training data collection and labeling campaign began in FY 2020 and will continue in FY 2021 to improve the utility of the DSCNN-P predictions.
10. The entire user interface was rebuilt based on over a year of user feedback. The primary goal was to enable faster, more efficient review of the analyzed build data and to enable crowdsourcing for collection of new training data.
11. The ConceptLaser X-Line 2000R machine type is now recognized by Peregrine.

12. Several artifacts can now be removed from the QM Meltpool data.

13. Log files can now be parsed for the ExOne Innovent and ExOne M-Flex machines.

A-5. RAVEN

Raven is a Software Tool developed in Python at ORNL. Raven is intended for use with the CVI process. Specifically, it allows the user to track the locations of the binder jet printed parts prior to insertion into the CVI furnace. CVI technicians will take a series of images of the parts as they are loaded onto the furnace platforms, or “grid plates.” These images are taken under controlled lighting conditions and background subtraction, along with other classical computer vision techniques. The images are used to determine the boundaries of each part. Each part is then assigned an identifier number produced by the Peregrine Software Tool during the instantiating printing process. The part locations and their corresponding identifiers are saved to a file which is uploaded to the Database via the Digital Tool web interface.

Development of Raven began in FY 2020, and it is currently on major version 1. Raven has not yet been deployed to users because COVID-19 restrictions delayed assembly and testing of the physical imaging system. Bringing Raven fully online is a priority for early FY 2021. Approximately 10 pages of documentation and training materials were written for Raven in FY 2020. The intended users of Raven will be the CVI machine technicians.

A-6. SCOPS

SCOPS is a set of imaging systems and an analysis method developed to analyze the complex thermomechanical history of parts printed with directed energy deposition (DED). The system uses stereo digital image correlation (DIC) to 3D map surface roughness features on parts as they are printed and tracks their motion over time. This allows for direct, in situ, noninvasive deformation and strain measurements over the entire part. This is combined with infrared thermal imaging to measure material expansion and contraction as a function of heat flow through the part. The SCOPS system is summarized graphically in Figure A-11.

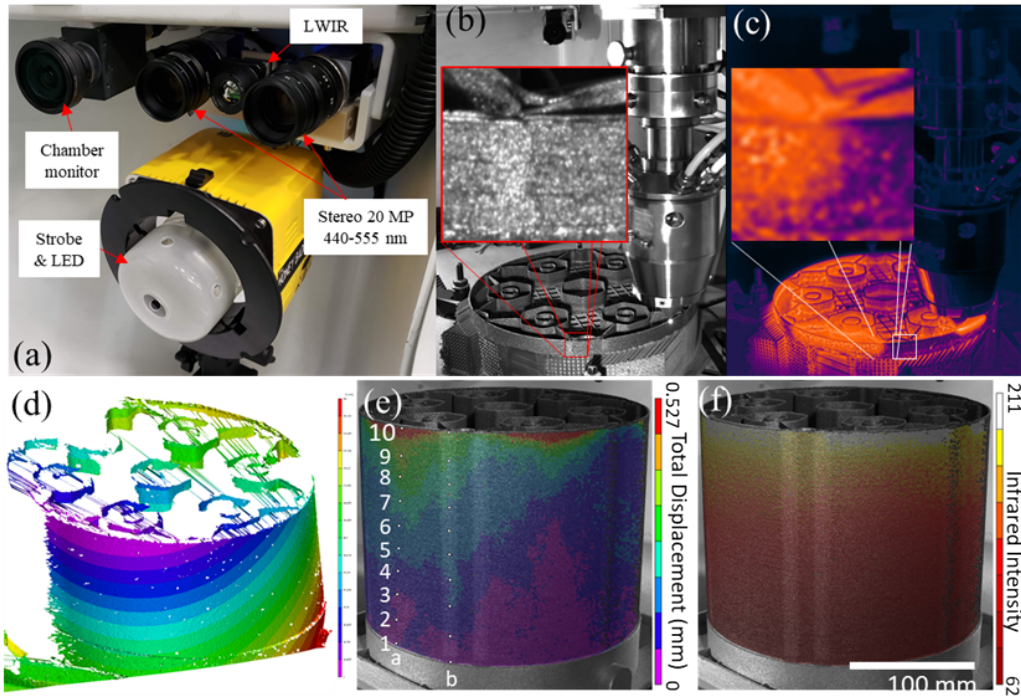


Figure A-11. (a) SCOPS imaging system, two stereo high-resolution visible cameras and an infrared camera, (b) visible trackable surface roughness of representative deposition, (c) infrared imaging showing thermal dissipation, (d) 3D map constructed from surface features, (e) tracked deformation induced from a combination of residual strain and thermal expansion from (f) 3D mapped infrared signal showing hot regions.

In FY 2020, two DED machines were outfitted with SCOPS: the BeAM modulo 400 and GKN cell 1. System hardware configuration was developed and optimized, along with supporting software for automated video capture and archival. Blackbody calibration was performed, and thermocouple validation of infrared imaging was automated. A validation experiment comparing DIC with neutron diffraction residual stress measurements was performed. Information from SCOPS provides an in situ measure of a previously inaccessible information stream, and it enables direct validation of finite element analysis (FEA) models, enhancing the accuracy of residual stress and microstructure simulation.

A-7. PIGEON

Pigeon was developed at ORNL and is both a Software Tool and a Python API. Pigeon is primarily an image segmentation software tool powered by a DSCNN. For more information regarding the DSCNN, users are referred to Scime et al. [2]. Pigeon's data flow is essentially identical to that represented in Figure A-10 for Peregrine. Some of Pigeon's key features include the following.

1. Pigeon accepts an arbitrary number of input *channels*, working seamlessly with both RGB and single-channel images. Pigeon also supports an arbitrary number of temporal *frames* associated with each *explicit channel*. For example, if two spatially registered color images, each with two temporal *frames*, are provided, then there are two *explicit channels* and 12 *implicit channels*.
2. Pigeon accepts images of arbitrary bit depth (up to 16 bits).

3. Pigeon accepts images of arbitrary resolution while ensuring that the output segmentation is always at the resolution of the input image.
4. The DSCNN is designed to work well with very high-resolution images of 20 MP or greater, even with limited GPU RAM available. It also provides sublinear analysis times, so doubling the image resolution will not double the inference time.
5. A low-resolution leg of the DSCNN always has a receptive field equal in size to the entire input image. This helps the DSCNN to learn to be robust to lighting variations and global scene changes.
6. The medium-resolution leg of the DSCNN has a receptive field which can be dynamically configured by the user. The size of this *tile* should be based on domain knowledge such that it includes the appropriate amount of contextual information.
7. The DSCNN can learn prior probability distributions that are a function of pixel location within the global image.
8. Class-wise imbalances are not a significant issue, because automatic rescaling of the loss function ensures that the DSCNN learns even extremely rare classes.
9. Full transfer learning from any DSCNN to any other DSCNN is supported.
10. The user can create multiple *workspaces* for working on different problems.
11. The DSCNN also outputs a deep feature vector which can be used by a perceptron to classify the entire *image stack*.

Development of Pigeon started in FY 2020. It is currently on major version 1 and is used by several members of the TCR team. Over 30 pages of documentation and training materials have been written for Pigeon in FY 2020. Users include data and material scientists. Pigeon is intended primarily as a development platform for creating different characterization tools. For example, Pigeon can serve as the framework for multiple Software Tools, such as one dedicated to characterizing powder feedstock and another focused on analyzing AM microstructures. Efforts are already under way to use Pigeon for a number of different aspects of the TCR Digital Platform, with a primary focus on automated material characterization.

A-8. SIMURGH

XCT is a process in which a 3D volume of an object of interest that represents attenuation coefficients of its comprising materials is reconstructed from 2D projections of the object acquired at different angles (views). XCT plays a critical role in nondestructive evaluation (NDE) and therefore certification/qualification of metal AM components. The quality of the 3D reconstruction algorithms used for dense metallic parts in complex systems and structures can be compromised by noise and artifacts such as streaks associated with metals, and a confounding effect called beam hardening (BH). Such artifacts and noise complicate the process of detection of salient defects (e.g., pores and cracks) in XCT images.

ORNL has developed an AI-based technique, *Simurgh*, that leverages CAD models of the parts, along with physics-based modeling, that shows significant improvement in XCT resolution by surpassing the state of the art in suppressing BH, metal artifacts, and noise in preliminary tests on both synthetic and real data. Figure A-12 shows a block diagram of the Simurgh technique. Figure A-13 and Figure A-14 show

example results on both synthetic and real data sets, respectively. For development and testing of the method, the data currently available at the MDF were used: XCT scans of printed airplane engine turbine blades. In FY 2021, these methods will be evaluated and developed further for application to TCR-relevant geometries.

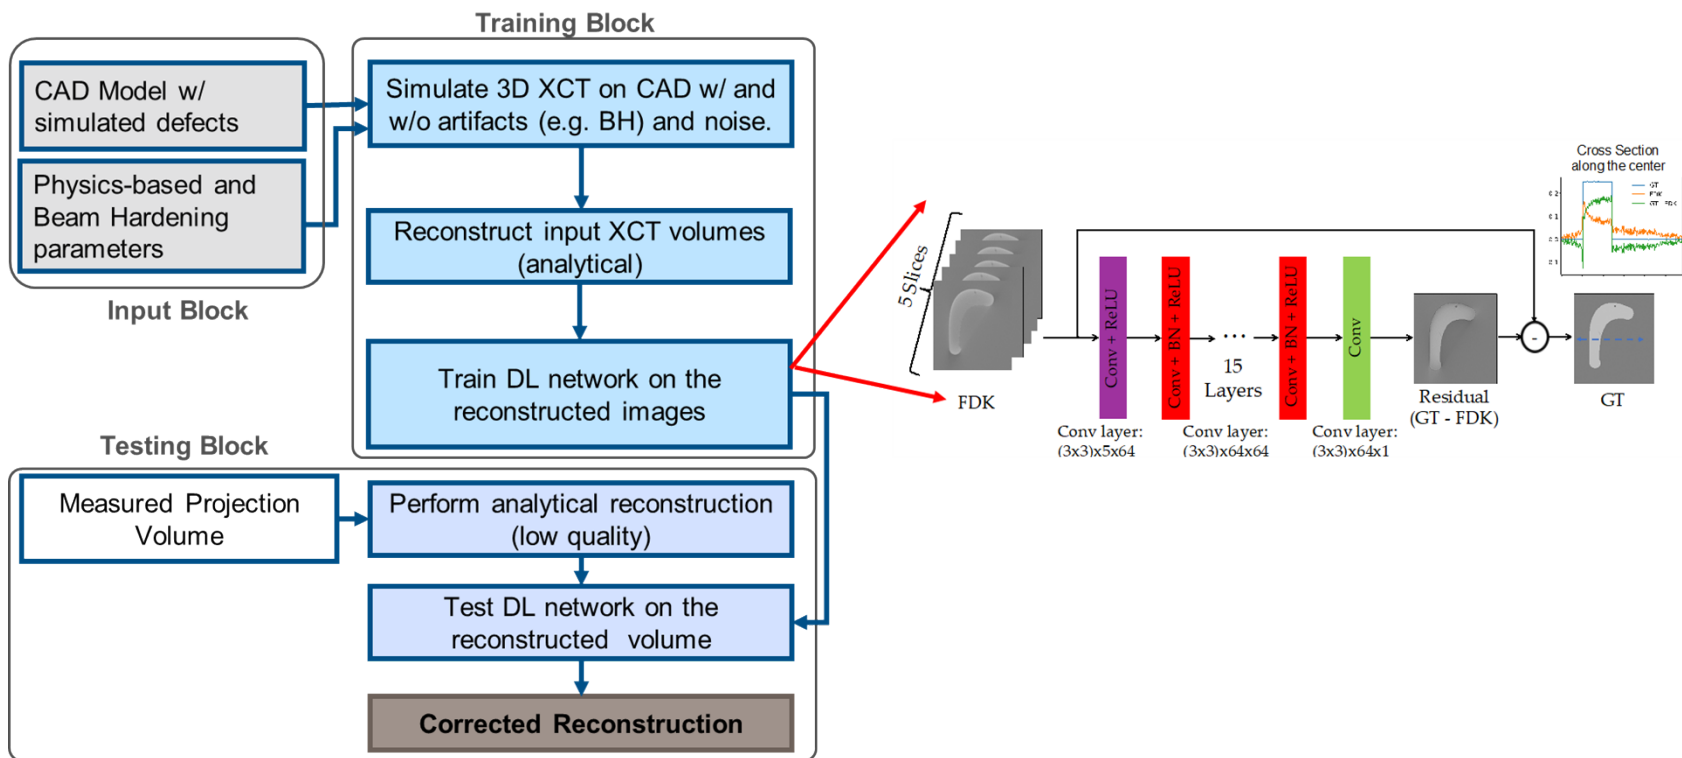


Figure A-12. Block diagram of the Simurgh framework. ORNL trains a convolutional neural network (AI-CT, top right) on the simulated (synthetic) data developed using CAD models and physics-based parameters.




Figure A-13. Example test results on a synthetic data set. Two different neural networks (MSD-Net [3] and AI-CT [4], [5]) to emphasize the modularity of the approach. The method is also compared against uncorrected standard analytical (FDK) approach [6] and state-of-the-art model-based iterative reconstruction (MBIR) [7]. Several areas of interest are highlighted to demonstrate the superiority of the work against the state of the art.

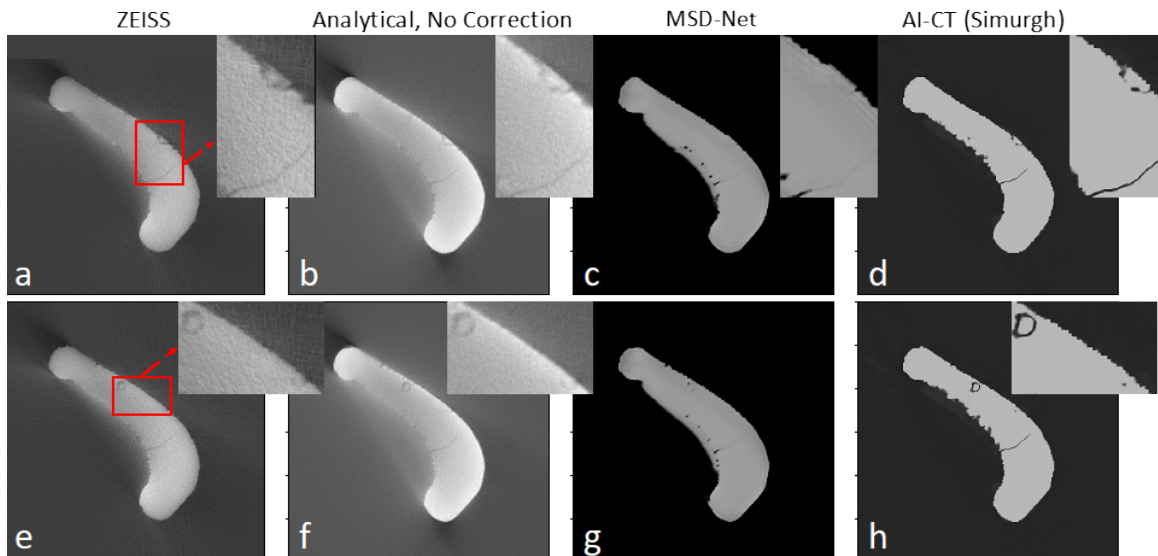


Figure A-14. Example test results on real data set. The proposed method is compared against the standard output from the ZEISS XCT system standard reconstruction, as well as the uncorrected standard analytical (FDK) approach and MSD-Net as a second neural network. Expanded views of areas of interest are highlighted in the inset images.

Functionally, a convolutional neural network (denoted as *AI-CT* in Figure A-12) is trained on the simulated (synthetic) data developed using CAD models and physics-based parameters. Then it is tested on synthetic and real data sets, as shown in Figure A-13 and Figure A-14, respectively. Two different neural networks—MSD-Net [3] and AI-CT [4], [5]—are tested to emphasize the modularity of this approach. The method is also compared against the standard output from the XCT system, as well as the uncorrected standard analytical (FDK) approach [6] and the state-of-the-art model-based iterative reconstruction (MBIR) approach [7].

ORNL has published a peer-reviewed conference paper on this topic, and a provisional patent has been filed by ORNL. A Technology Innovation Program (TIP) proposal was submitted to the Technology Transfer Office, and since its acceptance, the full proposal is in development. The Simurgh method is supported by ZEISS Quality Solutions, and they are considering licensing the technology from ORNL.

A-9. REFERENCES

- [1] ASTM. *Standard Terminology for Additive Manufacturing*. ISO / ASTM52900 - 15, 2015.
- [2] 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.* 36 (Dec. 2020):101453.
- [3] D. Pelt, K. Batenburg, and J. Sethian. “Improving Tomographic Reconstruction from Limited Data Using Mixed-Scale Dense Convolutional Neural Networks.” *J. Imaging* 4, no. 11 (Oct. 2018):128.
- [4] A. Ziabari, D. H. Ye, S. Srivastava, K. D. Sauer, J. B. Thibault, and C. A. Bouman. “2.5D Deep Learning for CT Image Reconstruction Using A Multi-GPU Implementation.” *Conf. Rec. - Asilomar Conf. Signals, Syst. Comput.* 2018-Octob (2019): 2044–2049.
- [5] A. Ziabari et al. “Beam Hardening Artifact Reduction in X-ray CT Reconstruction of 3D Printed Metal Parts Leveraging Deep Learning and CAD Models.” In *ASME IMECE 2020*, 2020.
- [6] L. A. Feldkamp, L. C. Davis, and J. W. Kress. “Practical Cone-Beam Algorithm.” *J. Opt. Soc. Am.* 1, issue 6 (1984) 612–619.
- [7] P. Jin, C. A. Bouman, and K. D. Sauer. “A Model-Based Image Reconstruction Algorithm with Simultaneous Beam Hardening Correction for X-Ray CT.” *IEEE Trans. Comput. Imaging* 1, no. 3 (Sept. 2015): 200–216.

