

# Defect Detection and Monitoring in Metal Additive Manufactured Parts through Deep Learning of Spatially Resolved Acoustic Spectroscopy Signals.

Jacob Williams<sup>(a)</sup>, Paul Dryburgh<sup>(b)</sup>, Adam Clare<sup>(b)</sup>, Prahalada Rao<sup>(c)</sup>, Ashok Samal<sup>(a)</sup>.

- (a) Computer Science and Engineering Department, University of Nebraska-Lincoln, Lincoln, NE 68588, United States. Email: [jakew42@gmail.com](mailto:jakew42@gmail.com); [samal@cse.unl.edu](mailto:samal@cse.unl.edu)
- (b) Advanced Component Engineering Laboratory (ACEL), University of Nottingham, Nottingham, NG7 2RD, United Kingdom. Email: [Paul.Dryburgh@nottingham.ac.uk](mailto:Paul.Dryburgh@nottingham.ac.uk); [Adam.Clare@nottingham.ac.uk](mailto:Adam.Clare@nottingham.ac.uk)
- (c) Department of Mechanical and Materials Engineering, University of Nebraska-Lincoln, NE 68588, United States. Corresponding Author Email: [rao@unl.edu](mailto:rao@unl.edu)

## Abstract

Laser powder bed fusion (LPBF) is an additive manufacturing (AM) process that promises to herald a new age in manufacturing by removing many of the design and material-related constraints of traditional subtractive and formative manufacturing processes. However, the level and severity of defects observed in parts produced by the current class of LPBF systems will not be tolerated in safety-critical applications. Hence, there is a need to introduce information rich process monitoring to assess part integrity simultaneously with fabrication so that opportune corrective action can be taken to minimize part defects. Spatially Resolved Acoustic Spectroscopy (SRAS) is a uniquely positioned nondestructive acoustic microscopy sensing approach which has been successfully used to probe the mechanical properties and assess the presence of defects in LPBF parts. However, the technique is sensitive to extraneous phenomena, such as surface reflectivity which occur within the LPBF system, and may occlude identification of surface breaking and sub-surface defects. With a view to applying the SRAS technique for in-process monitoring in a production-scale LPBF environment, and to overcome the foregoing challenge, this study proposes the use of a *deep learning* convolutional neural network termed as Densely-connected Convolutional Block Architecture for Multimodal Image Regression (DCB-MIR) that invokes SRAS-derived acoustic velocity maps of the part

as input data and translates them to an output resembling an optical micrograph. Through this approach, we demonstrate that defects, such as porosity and surface imperfections in titanium- and nickel-alloy samples made using LPBF, which were not clearly discernable in the as-measured SRAS acoustic map and obscured by artifacts in the optical image, are accurately identified. To quantify the accuracy of the approach, the cosine similarity between the predicted output images and target optical images was used as the objective function of DCB-MIR. The cosine similarity between the acquired SRAS signatures and the corresponding as-measured optical micrographs of samples typically ranged between -0.15 to 0.15. In contrast, when the optical micrograph-like images derived from DCB-MIR proposed in this work were compared with the optical signatures, the cosine similarity improved in the range of 0.25 to 0.60.

## KEYWORDS

Additive Manufacturing, Process Monitoring, Defect Detection, Spatially Resolved Acoustic Spectroscopy, Neural Networks, Deep Learning.

## Introduction

The overarching goal of this work is to reliably detect defects in metal parts fabricated using additive manufacturing (AM) processes through data acquired from a novel sensing approach called Spatially Resolved Acoustic Spectroscopy (SRAS). This is an important research area, because, despite the enormous progress in AM technology, the wider use additively manufactured parts is hindered owing to the lack of process consistency and presence of defects [1]. For instance, case studies by the aerospace industry report that the *buy-to-fly* ratio with AM technology is as low as 7:1 compared to 20:1 with traditional subtractive machining [2, 3]. However, given propensity for defect formation, and nascent quality

certification procedures, industries, such as the aerospace and defense sector, are hesitant to use AM parts in mission-critical assemblies due to component integrity concerns [4-6].

Figure 1 shows a schematic of form of AM process known as laser powder bed fusion (LPBF). A high-power laser, typically rated at 300 W to 500 W is used to selectively irradiate a bed of metallic powder feedstock, fusing the regions of powder together. Following a layer of deposition, a recoater system is used to spread a virgin layer of powder across the bed, the process then repeats building up a part layer-by-layer [7]. LPBF is well suited to creating parts with intricate features and internal geometry, however the low build rates seen in LPBF lead to high specific costs. This is exacerbated by the propensity for LPBF systems to produce parts with defects [1, 8].

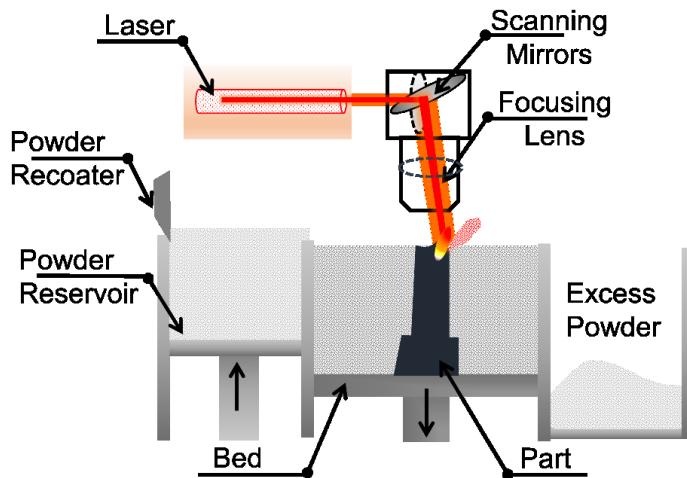


Figure 1: The schematic diagram of the laser-based powder bed fusion (LPBF) process [9].

Defects in LPBF are narrowed to four main causes [10]: (a) improper selection of process parameters, such as inordinately high or low laser power can both lead to porosity [11, 12]; (b) inherently poor part design [9]; (c) machine-related factors, such as the calibration of the laser, gap between the recoater and powder bed, optical aberrations from the focusing lens; and (d) feedstock powder material-related causes, such as presence of extraneous contaminants, characteristic distribution of the powder particle size, etc [13]. These factors can occur

concurrently, and lead to defects at multiple scales, ranging from the microstructure, macroscale (pinhole porosity and inclusions) to mesoscale (geometric distortion).

To further compound these challenges, traditional statistical optimization is incompatible with LPBF due to the multitude of adjustable processing parameters [4, 5]. For instance, O'Reagan et al. demarcate over fifty adjustable parameters in LPBF; it is not practicable to investigate each of these parameters and their interference using statistical experiments [14]. Additionally, the sample sizes in the AM process are small owing to the time, expense, and inherent slow production rate of the process. Hence making a large number of samples for subsequent destructive testing is not viable in AM [4, 5].

Everton et al. [15] investigated the development of in-situ monitoring technologies for LPBF and found that despite industry need for in-situ inspection, the majority of techniques implemented in practice have thus far focused on monitoring with the intention of understanding the process phenomena, rather than for in-situ identification of component defects. The authors concluded that development of new sensors that allow for inspection with a resolution consistent with commonly occurring material discontinuities is required. Additionally, variation in processing parameters has been shown to have a significant effect on the as-built microstructure and phase distribution within the component [16, 17]. Given the importance of microstructure for mechanical performance, a monitoring technique for LPBF should also have the capability to probe the microstructure [18].

The continuing objective of this work is to realize the integration of the Spatially Resolved Acoustic Spectroscopy (SRAS) instrumentation within a functional AM system to capture measurements during the build. The impetus for the use of SRAS in AM stems from the ability to detect and differentiate between surface and subsurface defects [19]; measurements of the

texture which can be correlated to build parameters; capability to measure rough surfaces; and inform targeted rework actions [20]. These prior results, uniquely position SRAS to tackle many of the monitoring and certification problems posed by LPBF.

Currently identification of surface breaking defects from SRAS data-sets relies upon user interpretation of the optical map obtained concomitantly with the acoustic velocity map. With a view to improving the reliability of defect identification and automating the process to apply this technique in a real-world manufacturing environment, this work utilizes the two parallel datasets generated by a SRAS inspection to construct a third artificial dataset, free from the artefacts, such as variations in surface reflectivity, which is not available when working with real samples and without repeated calibration trials. This is particularly relevant when considering the typical as-deposited surface finish of LPBF components, and the likelihood of loose powder resting on the sample [21]. Additionally, direct interpretation of the acoustic map is challenging due to a signal sensitivity arising from a variety of features such as subsurface and microstructure, alongside defects.

As a step towards realizing the objective of in-situ defect monitoring in LPBF with SRAS, we use the SRAS-acquired velocity map as an input and match it to the form of a binary defect map that is readily interpretable. This SRAS image to defect image transformation is done using a specially tailored deep learning convolutional neural network called Densely-connected Convolutional Block Architecture for Multimodal Image Regression (DCB-MIR). As a consequence of such mapping, the features that are not discernable in the original SRAS image are clearly discerned, whilst unintended variations in the optical dataset are disregarded.

## Methodology

### INTRODUCTION TO SPATIALLY RESOLVED ACOUSTIC SPECTROSCOPY (SRAS)

Both the technique and current instrumentation SRAS sensors have been reported in depth previously (c.f., Ref. [22]), thus, the following section provides only a brief overview. SRAS is a non-contact, non-destructive laser ultrasound technique, used for characterizing the physical properties, such as microstructure and grain orientation of a sample [22].

The operating principle of the SRAS apparatus used in this work is as follows; a schematic diagram is produced in Figure 2. An optical mask (grating pattern) is imaged on to the sample using a pulsed laser of known wavelength. These laser pulses initiate acoustic waves through thermoelastic expansion, this is done without ablation or damage to the sample. SRAS does not depend on time-of-flight measurements, but rather monitors the frequency of the wave packet.

Using a second probe laser, and a split-photodiode or knife edge detector, the perturbation caused by the surface acoustic wave can be measured and the characteristic frequency of the wave packet extracted using the Fourier transform. Thus, knowing the wavelength ( $\lambda$ ) and measuring the frequency ( $f$ ) the characteristic velocity can be calculated, using the relationship  $v = f\lambda$ . This SRAS-derived characteristic velocity can be further translated into image form, and is hence forth is referred as the SRAS velocity map. Additionally, the level out light returned to the photo-diode can be used to generate a second distinct dataset, analogous to an optical map.

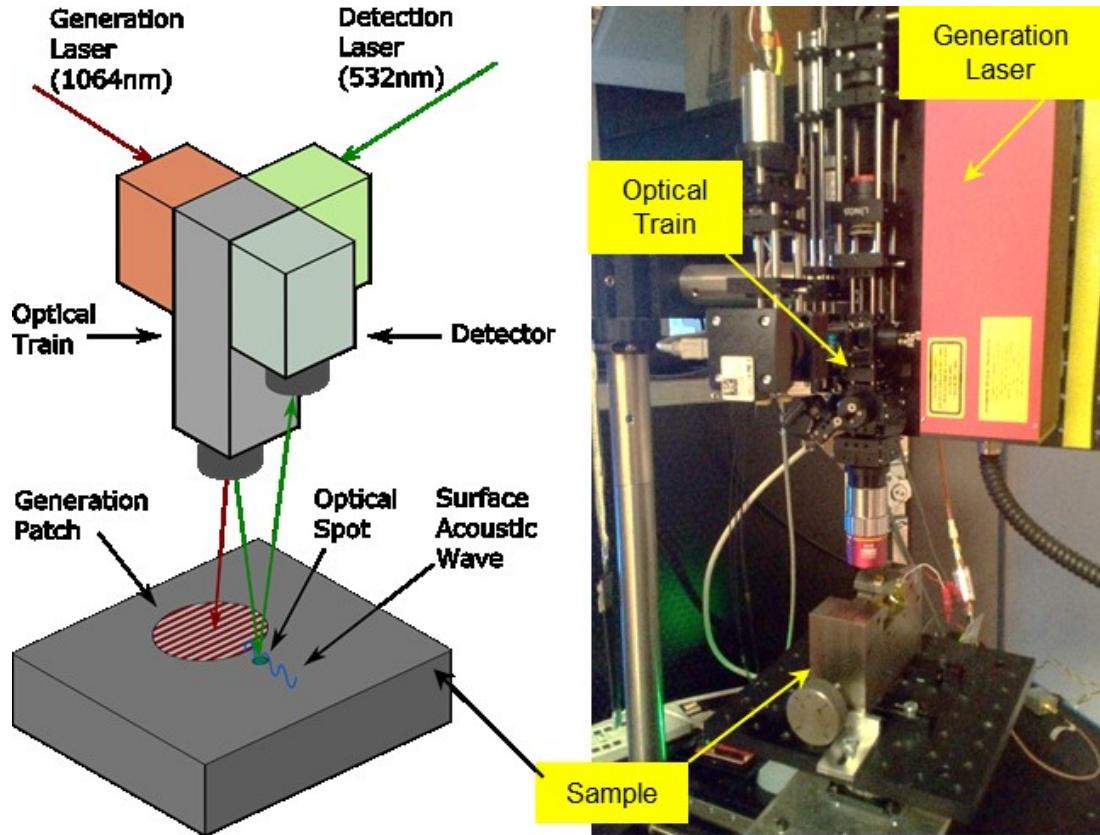


Figure 2: (left) Schematic of the SRAS instrument showing the surface acoustic wave generation laser, the laser used to detect the surface acoustic wave, and optical (patch) mask generated on the sample. (right) Photograph of third-generation SRAS instrumentation at Nottingham University.

The characteristic surface acoustic wave velocity is contingent on the orientation of the crystals and elastic properties of the sample. This information can then be used to obtain the grain size and orientation distribution. Indeed, as Li et al. have demonstrated that if multiple velocity images are taken with different acoustic wave propagation directions, then it is possible to recover the actual orientation of the grains [23].

The specifics of the SRAS sensor used in this work are as follows. The generation laser is a near-infrared (1064 nm) broadband Q-switched laser (AOT-YAG-10Q) which emits 1–2 ns pulses at a repetition frequency of 2 kHz. The detection laser is a 532 nm continuous wave laser (Laser Quantum Torus 532). The acoustic wavelength was  $\sim 24$   $\mu\text{m}$  and the generation patch size on the sample was  $\sim 200$   $\mu\text{m}$ . In this set-up the optical train is fixed and samples are raster scanned by linear stages (Physik Instrumente M-413 and M-414).

## EXPERIMENTAL PROCEDURE

A series of LPBF specimens were produced for inspection with the SRAS instrument system. The powder feedstock used was the common aerospace titanium (Ti-6Al-4V) and nickel alloys (CM247-LC) [24, 25]. The layer rotation was varied in the nickel sample set, whilst an additional titanium alloy sample set was manufactured with varying levels of tungsten by percentage weight, in order to study the effect of tungsten contamination on powder feedstock on LPBF builds. Specific build parameters for the three sample sets are given in Table 1. For reference purposes, the titanium alloy samples without contamination are labeled Ti(1); the titanium alloy samples with tungsten contaminants are labeled as Ti(2); and the nickel alloy samples are labeled Ni(1). Furthermore, both the sets of titanium samples are made using a rectilinear (meandering) laser scan hatch pattern, whilst the nickel samples are made using an island scan pattern, in which the hatches are made in small square sections. To further introduce heterogeneity in the processing conditions, the titanium samples are made in different LPBF systems. This purposely induced variation ensures that the approach presented is not specific to one type of alloy, process condition, and machine.

Lastly, all samples were mechanically polished prior to offline measurement by the SRAS system to ensure the surface provides a specular reflection, as necessitated by the knife edge detector. The 190 W sample was polished further in order to remove small layers of material to generate a quasi-volumetric dataset. The SRAS scan set-up was varied dependent on sample type; the scanning parameters, resolution and size of the data are also reported in Table 1.

Table 1: Processing parameters for all sample sets and scanning parameters.

Alloy	Titanium Alloy <i>Ti-6Al-4V</i>	Titanium Alloy <i>Ti-6Al-4V</i> contaminated with Tungsten	Nickel Alloy <i>CM247-LC</i>
Label	Ti(1)	Ti(2)	Ni(1)
Build System	Renishaw AM250	ReaLizer SLM50	ReaLizer SLM50
Size (xyz) [mm]	10 × 10 × 10	10 × 10 × 10	8 × 8 × 10
Number of samples	1	3	4
Laser Power [W]	190	82.5	100
Scan speed [mm s <sup>-1</sup> ]	600	500	400
Hatch spacing [μm]	75	90	50
Layer thickness [μm]	30	40	25
Hatch Rotation (°)	n/a	n/a	0, 15, 30 and 45
Tungsten Contamination (%)	n/a	1, 5 and 10	n/a
Scan Wavelength [μm]	24	24	24
Image Size (pixel)	2000 × 2000	1000 × 2000	1500 × 4000
Pixel size [μm]	5 × 5	10 × 10	2 × 5
Hatch pattern	Rectilinear	Rectilinear	Island
Approx. total data points	4,000,000	1,000,000	6,400,000

## DESIGN OF A CONVOLUTION NEURAL NETWORK FOR ANALYSIS OF SRAS SENSOR SIGNATURES.

### Primer to Neural Networks

Artificial neural networks (ANNs) are a form of machine learning originally inspired from neurophysiology [26]. ANNs are composed of basic computational units called neurons, which sum input from multiple sources (frequently the input data or other neurons) and apply a non-

linear transformation to generate an output. Neurons are connected to each other by scalar weights. These weights are generally learned through a process called backpropagation, in which the error at the end of the network (e.g. cross-entropy between the true and predicted values) is used to update the connections between neurons in all preceding layers. The simplest form of a neural network, called a multilayer perceptron (MLP), is composed of layers of neurons in which all neurons in one layer are connected to all neurons in the subsequent layer.

Typical MLPs, however, require a large amount of data to train and are prone to overfitting. This is because the MLP treats each data point as independent. In other words, the MLP does not inherently account for spatiotemporal correlations within the data set. This shortcoming is especially exacerbated when the data sets involve images. As the correlation between neighboring image pixels is not considered in MLP, the prediction is generally poor.

To overcome this lacuna with MLP, the proposed Densely-connected Convolutional Block Architecture for Multimodal Image Regression (DCB-MIR) uses a combination of convolutional layers and skip connections between layers, which will be clarified later. In the context of this work, the term *multimodal image regression* refers to mapping from one image input modality viz., SRAS velocity map, to another output image modality, namely the optical micrograph.

## Network Architecture

A diagram of the full architecture of DCB-MIR is shown in Figure 3, and can be described as a fully-convolutional block-densely-connected network. It consists of 6 blocks, plus input and output layers. Each block is composed of two convolutional layers, a batch-normalization layer, and a concatenation layer. The concatenation layer of each block both feeds directly into the next block and the concatenation layer of all subsequent blocks, which we refer to as the

block-densely-connected property. Similarly, the initial input feeds into both the first block and the concatenation layer of all subsequent blocks through a skip connection. This allows multiple scales of image processing to exist throughout the network. Hence, DCB-MIR preserves both the local, as well as wider area information of the image.

In the proposed network, the first block uses 12 2-dimensional convolutional filters. Each of these filters is composed of individual filters with convolutional kernels of size 7x7. The second block has 18 convolutional filters of size 5x5; and the remaining four blocks use 24 convolutional filters each of size 3x3. This configuration of filter sizes and number of filters aligns with previous work in the field, such as ResNet, and provides superior visual results than configurations using the same number or size of filters throughout the network [27]. In total, the final architecture has 102,273 trainable parameters. The first two blocks in the network are explained in depth later in this section, wherein the underlying mathematical operations will be clarified. The number of blocks and layers within a block were determined with systematic experimentation in a parameter space that satisfied the following criteria:

- 1) The receptive field must be large enough to account for the patterns of interest [28], and
- 2) The network must fit into memory of the available hardware while maintaining a reasonable batch size.

The limitations on data availability necessitated a smaller network size, thus limiting the impact of (2) on this study. As more data becomes available, the results may be improved with a deeper network trained on a multi-GPU set-up. In this work, the input to the network is a SRAS surface velocity signatures of 150 pixel  $\times$  150 pixel ( $\approx$  750  $\mu\text{m}$   $\times$  750  $\mu\text{m}$ ) in image form as shown in Figure 4(a), and its output is also image that resembles an optical micrograph (Figure 4 (c)); the output image is the same size as the input SRAS image.

This neural network is trained using supervised learning. The SRAS velocity map images is the input (Figure 4 (a)), and the corresponding as-measured optical micrographs are the output labels (Figure 4 (b)). This image-to-image mapping is the unique differentiating aspect of this work.

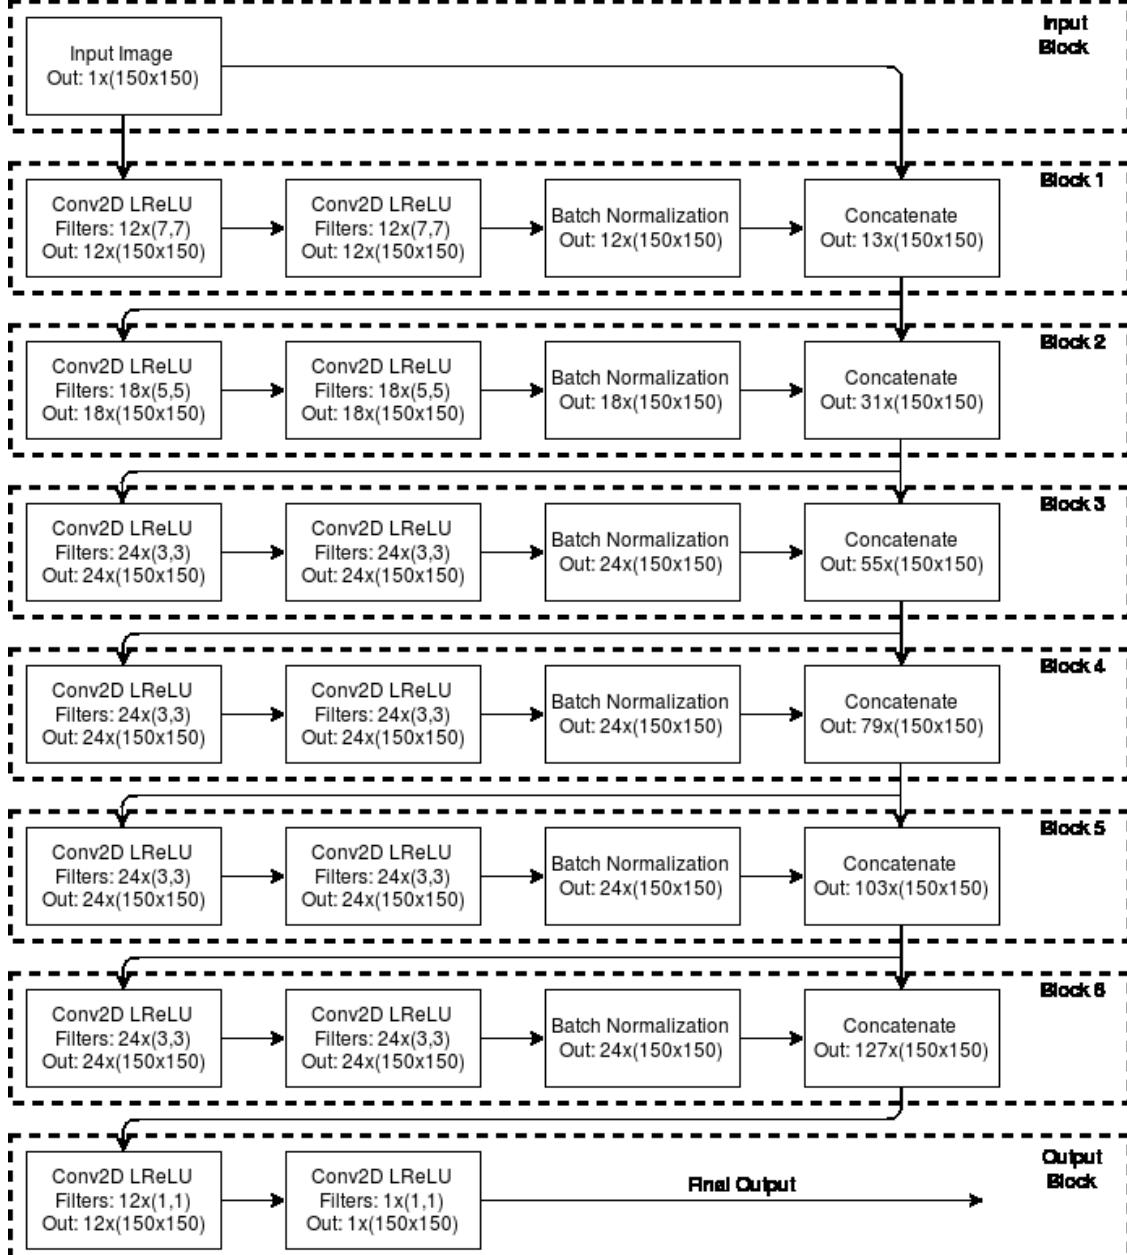


Figure 3: The architecture of the convolutional neural network proposed in this work. The network has six blocks, each with four layers. The layers labeled Conv2D LReLU is shorthand for a 2D-convolutional filter, whose output is passed through an activation function called the leaky rectified linear unit (LReLU). The number of convolutional filters in the layer are referenced in the next line, followed by the size of the filter. Lastly, the size of the output image is also specified.

DCB-MIR builds upon recent convolutional neural network architectures such as ResNet and DenseNet [27, 29]. This neural network architecture also takes inspiration from previous work in the areas of image super resolution [30, 31], the prediction of depth maps from RGB cameras [32], and approaches in image segmentation [27, 33].

Similar work has also been performed in biomedical image segmentation, demonstrating that convolutional network architectures with skip connections can achieve state of the art results in domains with limited data availability [34]. All of these techniques represent state of the art results in their respective fields, unmatched to this point by either traditional image processing or traditional machine learning techniques. Between these successes and the relatively few parameters in our model due to its fully convolutional nature, this approach can accurately model image to image regression. This is particularly important, since this domain lacks the data availability often required by deep learning to be successful.

Furthermore, to the best of our knowledge, traditional techniques in either machine learning or image processing have not been used in this type of image modality to image modality regression. The traditional techniques are ill suited to the complex, but strongly local, relationship exhibited by these modalities.

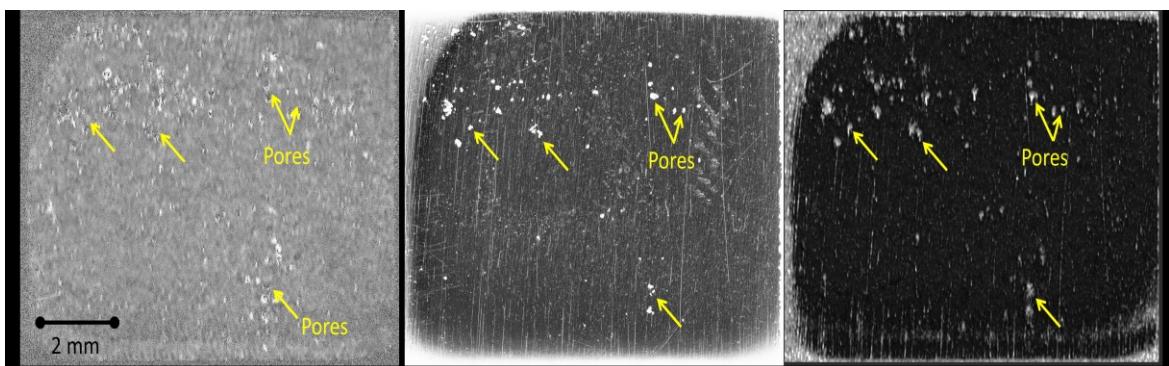


Figure 4: (a) SRAS signal velocity map for a titanium alloy sample, (b) the corresponding optical micrograph, and (c) the image predicted by the DCB-MIR neural network proposed in this work. Representative pores, that are clearly discernable in the optical micrograph are shown in (b), are largely occluded by noise in the SRAS signal (a). On processing the SRAS signal with a deep learning DCB-MIR neural network designed in this work, the pores are clearly discerned in image (c).

## Convolutional Layers

A convolutional neural network can be described as that which features one or more layers that learn certain number of convolutional filters. The objective of the convolutional layers is to learn filters that extract pertinent features of the SRAS image. These features can then be combined by future layers of the network to create more complicated features, eventually resulting in an approximation of the optical micrograph. More formally, let  $I$  be a 2-dimensional (2D) image of size  $m \times n$  and  $F$  be a 2D convolutional kernel of size  $k \times k$ . The result of a convolution of  $I$  with  $F$  at pixel  $i, j$  of image  $I$  is then defined as:

$$(I * F)(i, j) = a(i, j) = \sum_m \sum_n I(m, n) F(i - m, j - n) \quad (1)$$

The physical meaning of a convolutional filter is conveyed in Figure 5. The filter is akin to a mask or lens of certain size placed over the image. Here a  $4 \times 4$  convolutional filter is placed on the (1,1) pixel position of image  $I$ . This convolution operation results in a scalar value. The mask (filter) is then moved to the pixel position (1,2) and another scalar value is obtained. This operation is repeated pixel-by-pixel throughout the image (in practice the edges of the image  $I$  are padded with zero values), and ultimately results in a  $m \times n$  matrix, i.e., same size as that of the input image  $I$ .

Furthermore, signals evaluated by convolutional networks may have multiple channels or dimensions of data. To accommodate multidimensionality, if  $h$  is the number of incoming channels of data to a convolutional layer, then,  $h$  individual filters akin to the one shown in Figure 5 will be learned, i.e., one individual filter per channel. The results of each of these filters is subsequently summed and passed through an activation function  $f(\cdot)$ . Thus, the final result at pixel  $i, j$  can be defined as under ( $b$  is an additive bias term per convolutional filter, which we will omit including in further equations for simplicity of explanation):

$$A(i,j) = f \left( \sum_h (I_h * F_h)(i,j) \right) + b \quad (2)$$

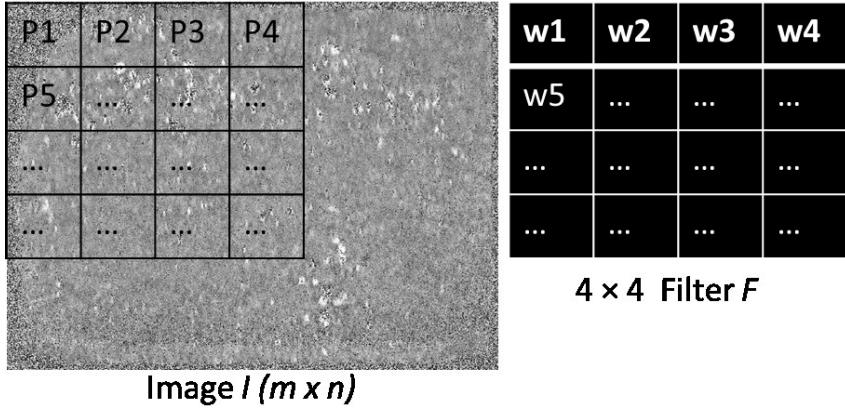
The leaky rectified linear unit (LReLU) activation function was used for all convolutional layers [35]. The LReLU activation function is a modification of the standard rectified linear unit (ReLU) function and takes the following form,

$$f(x) = \begin{cases} x & \text{if } x \geq 0 \\ ax & \text{if } x < 0 \end{cases} \quad (3)$$

The LReLU activation function was chosen to avoid the so-called *dying ReLU problem*, wherein the standard linear activation functions in a neural network are liable to map to 0. Experimentally, the LReLU led to more reliable convergence than the standard ReLU activation function. In this work  $a$  is set at the value of 0.3.

The first two blocks of the proposed network are further detailed for clarity (Figure 6). Before we begin explaining the structure of the network we clarify the notation. The  $j^{\text{th}}$  convolutional filter in Layer  $a$ , and Block  $b$  is represented as  $\mathbf{F}_j^{a,b}$ . Each convolutional filter consists of a group of  $h$  individual filters, an individual filter  $F_{i,j}^{a,b}$ , within the convolutional filter  $\mathbf{F}_j^{a,b}$  is represented as  $F_{i,j}^{a,b}$ . To reiterate, the individual filter  $F_{i,j}^{a,b}$  is a  $k \times k$  matrix (kernel function) which acts on the input image per the convolution operation in Eqn. (1).

In Block 1, Layer 1 there are 12 convolutional filters  $\mathbf{F}_j^{1,1} j \in \{1 \dots 12\}F$  each with an individual filter (i.e.,  $h = 1$ ) of size  $7 \times 7$ . Hence, we have a total of 12 convolution operations performed, leading to 12 derived copies of the input image  $I$  labeled  $A_{i,j}^{1,1}, i = 1, j \in \{1 \dots 12\}$ . To reiterate, convolving an image  $I$  with an individual filter  $F_{1,j}^{1,1}$ , and passing through an activation function results in the output  $A_{1,j}^{1,1}$ . The size of each matrix  $A_{1,j}^{1,1}$  is identical to the input image  $I$ .



$$\begin{aligned}
 (I * F)(1,1) &= a(1,1) = p_1 \cdot w_1 + p_2 \cdot w_2 + p_3 \cdot w_3 \dots \\
 (I * F)(1,2) &= a(1,2) = p_2 \cdot w_1 + p_3 \cdot w_2 + p_4 \cdot w_3 \dots \\
 &\vdots \\
 (I * F)(m,n) &
 \end{aligned}$$

Figure 5: A pictorial explanation of a convolutional operation between an image  $I$  of size  $m \times n$ , and a filter  $F$  of size  $4 \times 4$  (consider each box as a single pixel for explanation purposes). Each pixel in the filter is progressively multiplied with the corresponding pixel in the image  $I$  over a  $4 \times 4$  area. In this picture a pixel for the Image  $I$  is expressed as  $P_i$ . For instance, beginning with the  $(1,1)$  location in the image, we obtain a sum  $a(1,1)$ . When the filter (or mask) is moved to the right by 1 pixel to be placed on the image position  $(1,2)$  the convolution operation results in the scalar output  $a(1,2)$ , and so on. In Block 1, Layer 2, the 12 derived images  $A_{i,j}^{1,1}$ ,  $i = 1, j \in \{1 \dots 12\}$  from the first layer are again passed through twelve convolutional filters,  $F_k^{2,1}, k \in \{1 \dots 12\}$ . The structure of the network predicates that the number of individual filters within a convolutional filter; the number of individual filters should be equal to the number of output images from layer 1. Since there are 12  $A_{i,j}^{1,1}$ 's from layer 1, the number of individual filters within each convolution filter in the second layer is also 12. Accordingly, the filters in the second layer are labeled  $F_{i,j}^{2,1}$   $i, j \in \{1 \dots 12\}$ . After the second convolution operation, there are a total of 144 derived copies of the image  $I$ , labeled  $A_{i,j}^{2,1}$ ,  $i \in \{1 \dots 12\}, j \in \{1 \dots 12\}$ .

The resulting outputs from a particular convolutional filter are summed together. This summing operation of outputs  $A_{i,j}^{2,1}$   $i \in \{1 \dots 12\}$ . The output from the  $j^{\text{th}}$  convolutional filter at the end of the second convolutional layer is represented as  $\mathbf{A}_j^1 = \sum_{i=1}^{i=12} A_{i,j}^{2,1}$ . Subsequently, following the summing operation in Block 1, each of the  $\mathbf{A}_j^1$ 's is exactly the size of the input image  $I$ . These  $\mathbf{A}_j^1$ 's are subjected to two more layers, namely, batch normalization and concatenation, the need for which is explained below.

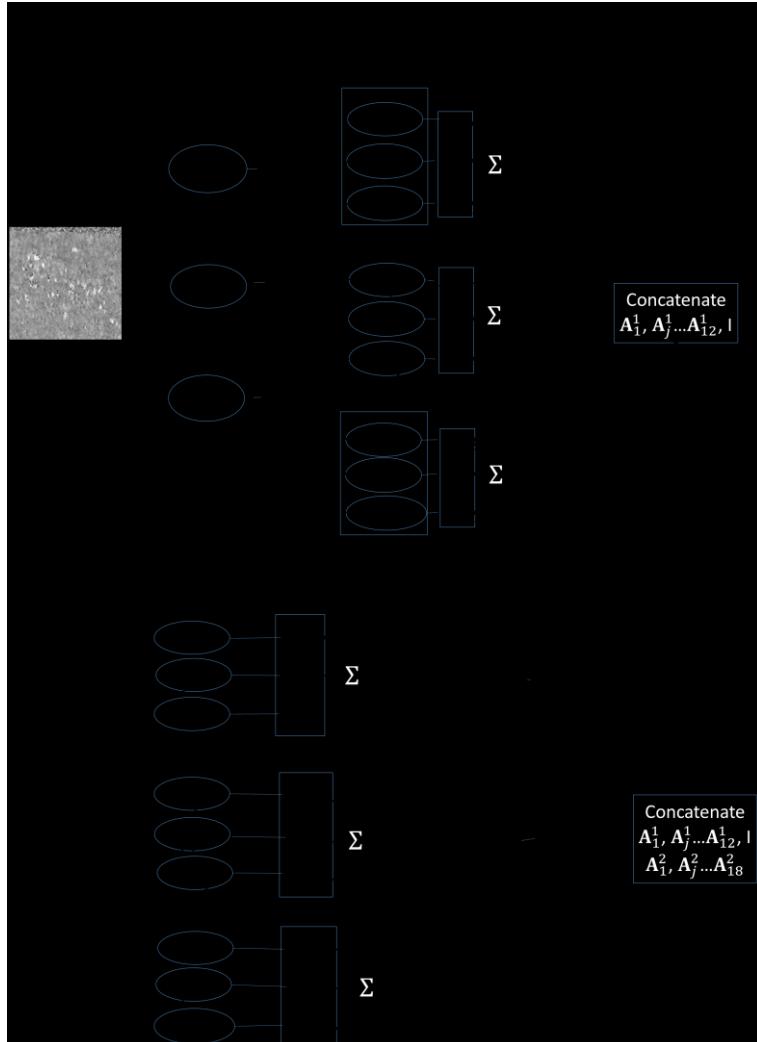


Figure 6: Schematic diagram of the first two blocks of the convolutional neural network. Each block has two convolutional layers, which contain a certain number of convolutional filters referenced as  $\mathbf{F}_j^{a,b}$ , where  $a$  is the block number, and  $b$  is the layer number. Each convolutional filter in turn has a fixed number of individual filters represented as  $F_{i,j}^{a,b}$  wherein the subscript  $i$  indexes an individual filter within convolutional filter  $\mathbf{F}_j^{a,b}$ .

## Batch Normalization Layers

Batch normalization is a technique developed to account for the shift in the covariance of an incoming signal that occurs as the signal passes through multiple layers of a neural network [36]. Batch normalization layers reduce the effect of this shift by scaling the output to a normal distribution of zero mean and unity variance  $\mathbf{N}(0,1)$ , and subsequently, learning an additive and multiplicative values to shift the mean and variance of the distribution. Batch normalization layers are included at the end of each block both to ensure that the network has the ability to maintain a similar scales for the outputs from each block, similar to their use in DenseNet [29]. The network used in this work was found to have poor error convergence without the use of batch normalization layers in each block.

Let  $B$  be a batch of  $m$  inputs,  $\{x_1, \dots, x_m\}$ . A batch normalization layer calculates the mean,  $\mu_B$ , and variance,  $\sigma_B^2$ , of  $B$  and then applies the following normalization to each  $x_i$ , a small value  $\varepsilon$  is added to ensure that the denominator does not become zero.

$$\hat{x}_i = \frac{x_i - \mu_B}{\sqrt{\sigma_B^2 + \varepsilon}} \quad (4)$$

The result is then scaled by a factor  $\gamma$  and summed with an offset  $\beta$  to achieve the output of the batch normalization layer,

$$y_i \leftarrow \gamma \hat{x}_i + \beta \quad (5)$$

During testing, however,  $\mu_B$  and  $\sigma_B^2$  are not available. Thus, running estimates of the population expected value,  $E[x]$ , and population variance,  $V[x]$  are updated at each mini-batch as follows:

$$\begin{aligned} E_{i+1}[x] &\leftarrow \alpha E_i[x] + (1 - \alpha) \mu_B \\ V_{i+1}[x] &\leftarrow \alpha V_i[x] + (1 - \alpha) \sigma_B^2 \end{aligned} \quad (6)$$

The final values of  $E[x]$  and  $V[x]$  are then used for normalization at test time, in place of  $\mu_B$  and  $\sigma_B^2$ .

## Concatenation and Output Layers

The final layer in each block is a concatenation layer which simply stores the output of the block in which it is contained with the output of all previous blocks and the original image. Accordingly, multiple scales of the filtered image are available at all layers of the network. Following concatenation layer in Block 1, we have the 12  $\mathbf{A}_j^1$ 's and the original image  $I$  stored in memory.

The concatenated images are then passed to Block 2 having with 18 convolutional filters each with 13 individual filters of size  $5 \times 5$  in Layer 1, and 18 convolutional filters each with 18 individual filters of size  $5 \times 5$  in Layer 2. At the end of the concatenation layer in block 2, there are a total of 31 images (12  $\mathbf{A}_j^1$ 's,  $j \in \{1 \dots 12\}$  from block 1, 18  $\mathbf{A}_j^2$ 's,  $j \in \{1 \dots 18\}$  from block 2, and the original image  $I$ ).

Continuing through the network results in 127 images of size 150 pixels  $\times$  150 pixels after the concatenation layer of Block 6. Layer 1 of the output block has 12 convolutional filters  $\mathbf{F}_j^{1, Output}$ ,  $j \in \{1 \dots 12\}$ , each with 127 individual filters of size  $1 \times 1$ ,  $F_{i,j}^{1, Output}$ ,  $i \in \{1 \dots 127\}$ ,  $j \in \{1 \dots 12\}$  resulting in 12 images labeled  $A_j^{1, output}$ ,  $j \in \{1 \dots 12\}$ . Finally, the twelve images ( $A_j^{1, output}$ ) are convolved through one last filter,  $\mathbf{F}_1^{2, Output}$  of size  $1 \times 1$  in Layer 2 of the output block, which combines (sums) them into a single output image of size 150 pixel  $\times$  150 pixel.

## Network Optimization

*Cosine Similarity:* The weights of the individual filters are obtained through the backpropagation algorithm [26]. Instead of the customary mean square error minimization, given the spatial (image-based) nature of the prediction, the network was trained to maximize the cosine similarity of the predicted optical image and the actual optical image. Cosine similarity of two vectors  $x$  and  $y$  is calculated as:

$$CS(x, y) = \frac{x^T y}{\|x\| \|y\|}$$

where  $T$  is the transpose operator and  $\|\cdot\|$  indicates the magnitude of the vector. The cosine similarity varies in the range [-1, 1]. Previous works in face recognition have seen benefit in using the cosine similarity measure over Euclidean distance or mean squared error [37].

While mean squared error is the standard measure in many similar applications, it was found to be less appropriate for this task. The tendency of mean squared error to over-exaggerate the effects of outliers lead to convergence instabilities, given the inherent noise in the SRAS signal. Further, in isolating defects, and particularly porosity, we posit that the pixel scale, i.e., exact value of the pixels, is of lesser importance, rather, the spatial pattern over a group of pixels determines the location of a defect. Since, the denominator in cosine similarity is a normalization term, the network is not tuned to learn the scale or value of each pixel in the output image. We found that this further simplifies the learning process and leads to faster and more accurate convergence when compared with mean squared error.

*$L_2$  Penalty Function:* In addition to cosine similarity, a small  $L_2$  penalty with  $\lambda = 0.001$  is applied, based on extensive trials, to all the individual filter weights in the network. This reduces the propensity of the network to overfit on training data by not allowing individual

connections to dominate the networks activity. Given a scale  $\lambda$  and weights  $w_j$ , the accrued L<sub>2</sub> loss is calculated as follows:

$$loss_{L_2} \leftarrow \lambda \sum_{j=1}^m w_j^2 \quad (7)$$

*Optimization Algorithm:* The Adam optimization algorithm was used to find the weights and bias value of the network. Adam is first-order gradient-based algorithm for the optimization of stochastic functions, informed by estimates of lower-order moments [38]. The optimizer is parameterized by four values:  $\beta_1$ , a decay value for the running average of the gradient;  $\beta_2$ , a decay value for the running average of the square of the gradient;  $\alpha$ , the learning rate; and  $\varepsilon$ , a value used to prevent division by zero. The values chosen for optimizing our network are as follows:  $\beta_1 = 0.9$ ;  $\beta_2 = 0.999$ ;  $\alpha = 0.1$ ; and  $\varepsilon = 0.1$ .  $\beta_2$  were chosen as the default values presented in . A significantly larger was chosen in concordance with and our empirical evaluation

## Preprocessing of the Data Set

A total of 10 pairs of SRAS and optical images were used in the training of this model. These images come from two different materials taken over three different samples. Three images (six total) were collected from each of two titanium alloy samples, referred to as Ti(1) and Ti(2), and four images were collected from a nickel sample labeled as Ni(1), as summarized previously in Table 1. The images used from the same samples were collected from different layers of printing with enough separation to prevent overlap in the optical images. Hence, the images from the same sample can be considered as heterogeneous. The samples are chosen such that they have distinctly different surface morphologies, therefore enabling the neural network to train over a more generalized space.

All optical and SRAS image pairs were first cropped to remove the background noise surrounding the sample. The final images used were approximately  $1\text{ cm} \times 1\text{ cm}$ . Next, an intensity scaling step was performed on each image to map the pixel intensity values to a range of  $[-1,1]$ , where neural networks are found to be most numerical stable. For each image, the mean value was subtracted and the image was then divided by the resulting maximum value. Since the network optimization metric used in this work, viz., cosine similarity, relies on the relative values in the image, rather than the absolute values, this form of preprocessing can be applied individually to each image, rather than using collected population statistics.

### **Training Procedure**

Referring to Table 1, out of the 10 images in the dataset (3 samples for Ti(1), 3 for Ti(2), and 4 for Ni(1)), 8 images were used for training, 1 for validation, and 1 for test in each run. The validation image was chosen out of the 9 non-test images with the stipulation that if an image from a type of material, such as Nickel (Ni(1)), was being tested, then the training set contained the rest of the 3 images for the Nickel. This exceedingly conservative testing procedure ensures that the training set contains images from different materials, thus imparting a degree of robustness to the network. In other words, the network trained in such a way is not constrained for mapping the SRAS output for only one type of material, but is generalizable across materials. We also note that, while validating on the same material would be ideal, given the limited number of images available, especially for the Nickel-alloy Ni(1) sample, training on more images from Ni(1) was found to produce visually superior results. The training procedure has the following steps.

*Step 1 – Dividing the images into windows.*

The SRAS and optical images held out for testing and validation were each divided into 150 pixel  $\times$  150 pixel regions (window size), and sampled with a *stride* (moving window shift) of 19 pixels, thus allowing for significant overlap between windows (131 pixels, window overlap). This sampling procedure generates approximately 10,000 sub-images for the SRAS-optical micrograph pair; we annotate these set of sub-images as  $(i)$ . The window shift or stride was chosen to have a large least common multiple with the window size to reduce situations where the beginning of one image frame aligns with the end of another, which can lead to a discernable gridding pattern when the image is reconstructed through averaging.

*Step 2 – Batching Images for Training*

A batch size of 40 sub-images was used for the stochastic gradient descent algorithm to satisfy computing hardware constraints. Each batch of 40 sub-images was comprised of five 150 pixel  $\times$  150 pixel sub-images sampled at random from each of the eight images held out for training. That is, a pair of pixel coordinates within the bounds of the original image were generated to serve as the upper left pixel of the sub-window used in training, for each of the 40 sub-images. As these images were generated dynamically during training, no prior splitting into sub-images was necessary, as in the train and test sets. It is also worth noting that due to the dimensions of the original images (1000 pixel  $\times$  1000 pixel at a minimum), the exact same sub-image was rarely presented twice to the network during training. A single update was then performed on the batch before discarding it and generating a new batch. Figure 7 demonstrates the selection of a batch.

### *Step 3 – Stopping the Training Process*

After a cycle of 100 updates is performed as described in Step 2, the performance of the network on the validation set was evaluated. If the average cosine similarity over all sub-images improved, the weights of the network were stored and training was continued. If the cosine similarity failed to improve for 20 cycles, training was halted and the weights of the network were reverted to the values stored at the maximal performance on the validation data.

### **Testing Procedure**

For each of the 150 pixel  $\times$  150 pixel sub-image regions ( $i$ ) of the SRAS signal from the test input-target set, the fully trained network was used to map to an optical micrograph image. Pixels in overlapping regions were averaged together to achieve the final predicted output image, with the following exception – a 5-pixel band at the outside of each of the 150 pixel  $\times$  150 pixel regions was discarded. This is reasonable, since, zero-padding was used in the convolutional layers to maintain image size throughout the network, the outermost pixels are found to be poor predictors of the original image.

The inclusion of these pixels led to a gridding pattern in the image, much like the issue with stride described in Step 1 of the training procedure. This final predicted image was then available for quantitative and qualitative comparison the original optical micrograph, as reported in the following section.

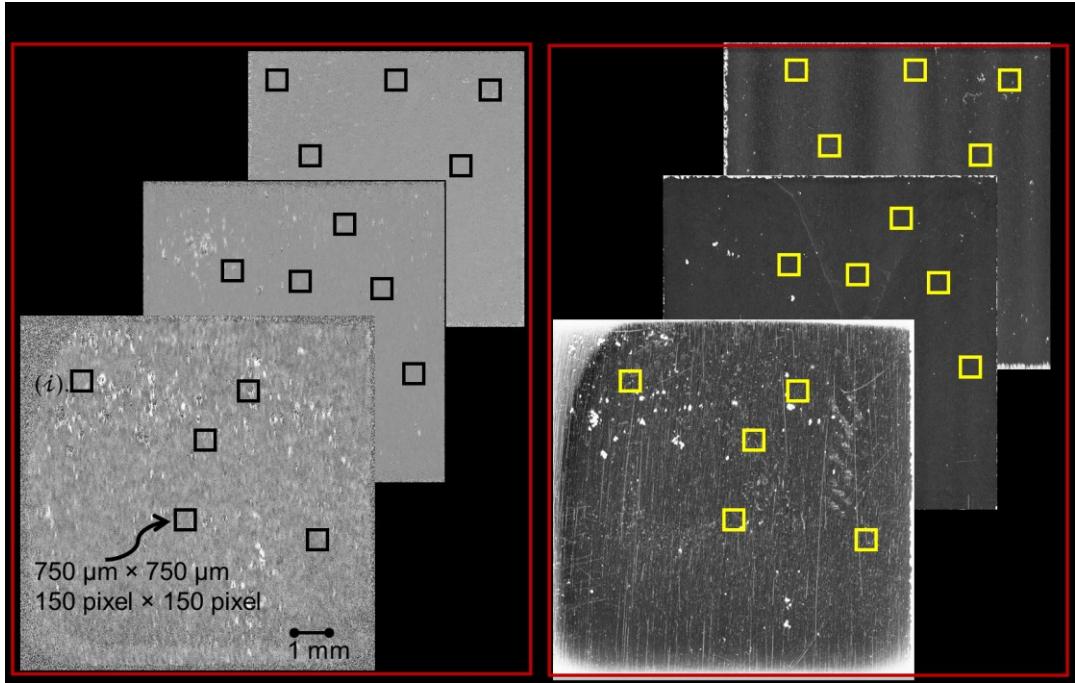


Figure 7: One batch in the training procedure includes selecting eight SRAS (input) and corresponding optical images (target). Next five areas are selected randomly for each input-target pair. Thus there are a total of 40 images per batch. With these 40 images the weights of the convolutional filters are updated continuously for 100 instances (updates). After 100 update epochs, the new set of 40 images is sampled (batch). The batch-and-update procedure is repeated 100 times.

## Results and Discussion

The SRAS velocity map, corresponding optical micrograph and predicted images for each of the samples are shown in Figure 8 through Figure 10. The first uncontaminated titanium samples (Ti(1)) polished and imaged at three layers are shown in Figure 8. The surface morphology of this sample is characterized by large acicular pores of diameter in the range of 100  $\mu\text{m}$ , characteristic of lack-of-fusion porosity. The neural network predicted image aptly captures the location, as well as the size of the pores.

More pertinently, the neural network was robust to artifacts that are seen in the optical images, such as the large number of scratches, Figure 8(b), and the solvent stains that invariably occur from sample polishing and vary the light level returned to the detector system but are not due indicators of defects. Similarly a series of vertical lines can be discerned in the optical image of the final scan, Figure 8(c), this corresponds to unexpected variations in the power-

output from the continuous wave-nature of the detection laser; as in previous cases, the predicted image successfully disregards these features to leave a map of only defects. Furthermore, it is known that the acoustic map is sensitive to subsurface features; in the acoustic datasets presented here there are several such indications of subsurface features [19]. These features are excluded from the predicted map, allowing segmentation of surface and subsurface defects.

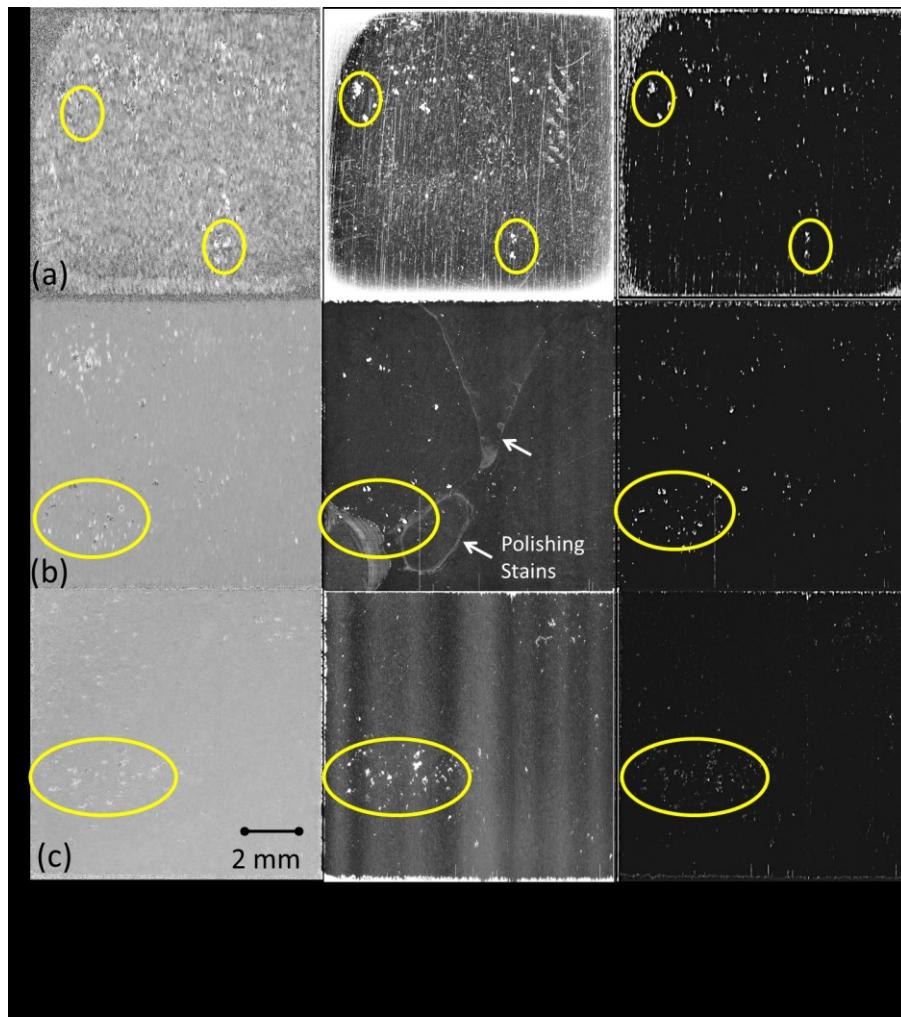


Figure 8: Results for the standard 190 W titanium sample polished to three different heights (a) nominal surface (b)  $\sim 100 \mu\text{m}$  below surface and (c)  $\sim 200 \mu\text{m}$  below surface. Representative areas of porosity are circled. Results predominantly show porosity in the offline optical micrograph, while not being demonstrably clear in the SRAS velocity map. The convolutional neural network maps these pores with clarity. Furthermore, the neural network is found to be robust to extraneous artifacts, such as polishing stains.

Figure 9 reports results for the three tungsten contaminated titanium samples (Ti(2)); these samples contain (a) 10%, (b) 5%, and (c) 1% tungsten contaminants by weight. Firstly, from the captured optical datasets it is discerned that the tungsten particles dispersed across the sample surface have a significantly different optical reflectance to the bulk material at 532 nm compared to the Ti(1) samples shown in Figure 8. Furthermore, samples suffer from remnant polishing marks which vary the light returned to the detector; in particular in Figure 9(b) for the 5% tungsten contaminated sample, the polishing stain found in the lower left has an optical response similar to the tungsten contamination. The neural network captures the position and size of the pores appreciably for samples shown in Figure 9(b) and (c), however, the neural network is not able to capture some of the prominent pores in image Figure 9(a). This under-prediction is likely caused by the nature of the training data: only ten samples were used, collected from two materials, under three different processing conditions. The combination of non-homogeneous data and limited sample size limits the ability of the network to generalize in some cases. Future work with a larger training dataset and more samples per material would likely increase sensitivity and alleviate under-prediction.

Figure 10 depicts the results for the nickel alloy sample, which have distinctive surface morphology. Whilst there are less discernable pores compared to the titanium alloy samples, the nickel maps are dominated by surface cracking, this is due to the hard to weld nature of CM247-LC [39]. Again, similar variations to the previous two sample sets can be seen in the optical datasets. The large conglomeration of artefacts in Figure 10(a) and (b) are from colloidal silica used for fine polishing. From observation of the predicted images, it is clear that the presence of tungsten in the Ti(2) samples is captured and translated in the DCB-MIR predicted images. The neural network captures a vast majority of the distinguishing characteristics for the nickel images. In Figure 10(b), as an example, samples have been manufactured in an island

scan pattern, with a 1 mm overlap. In both the acoustic and optical map this cross pattern is difficult to discern, but is clearly seen in the predicted image. This concurs with previous findings of higher cracking rates in border interfaces [40]. The shift in microstructural texture in these interface regions obfuscates the relevant features in the SRAS acoustic map.

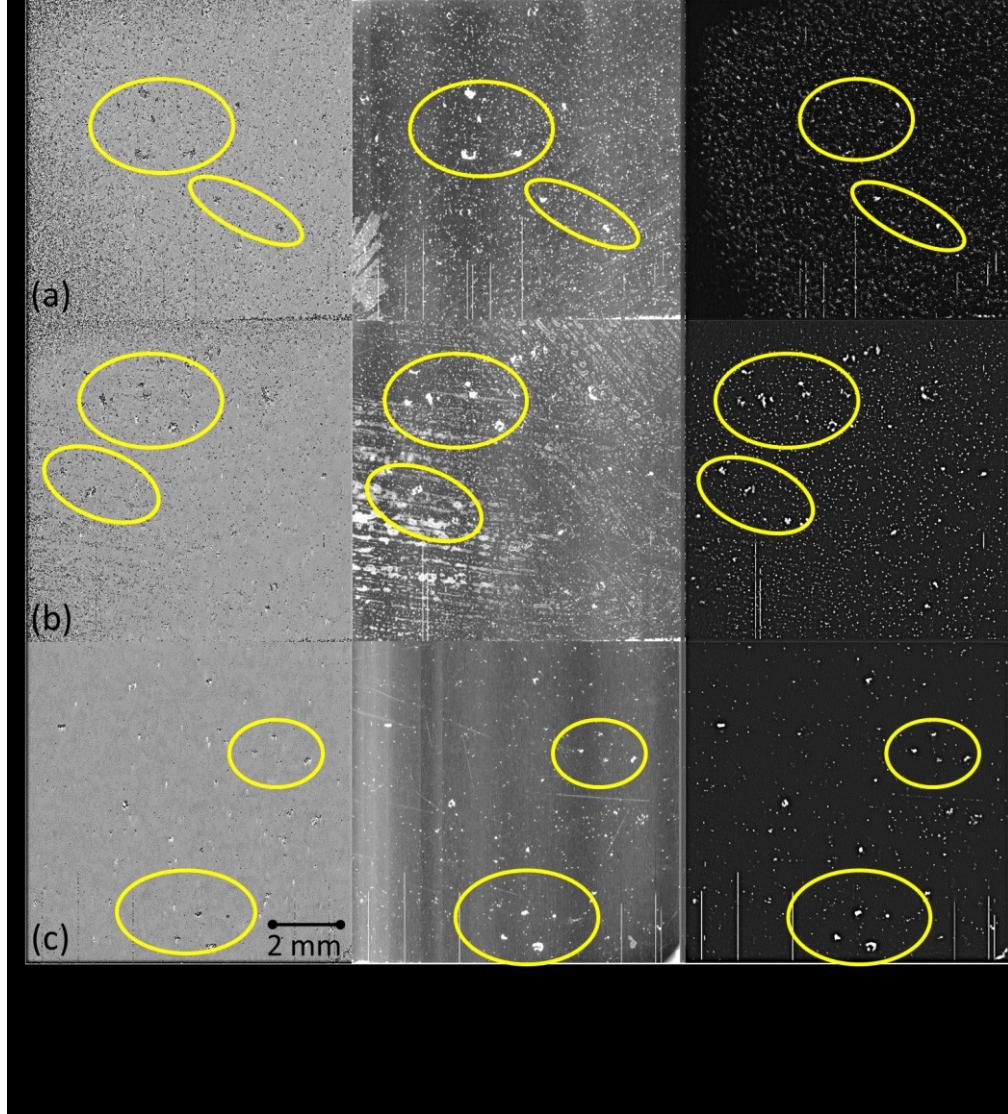


Figure 9: Results for the tungsten contaminated titanium alloy samples (Ti(2)) (a) 10%, (b) 5% and (c) 1% Tungsten. Representative areas of porosity are circled. Results depict porosity and surface defects, such as scratches in the offline optical micrograph, while not being demonstrably clear in the SRAS velocity map. The convolutional neural network captures these defects.

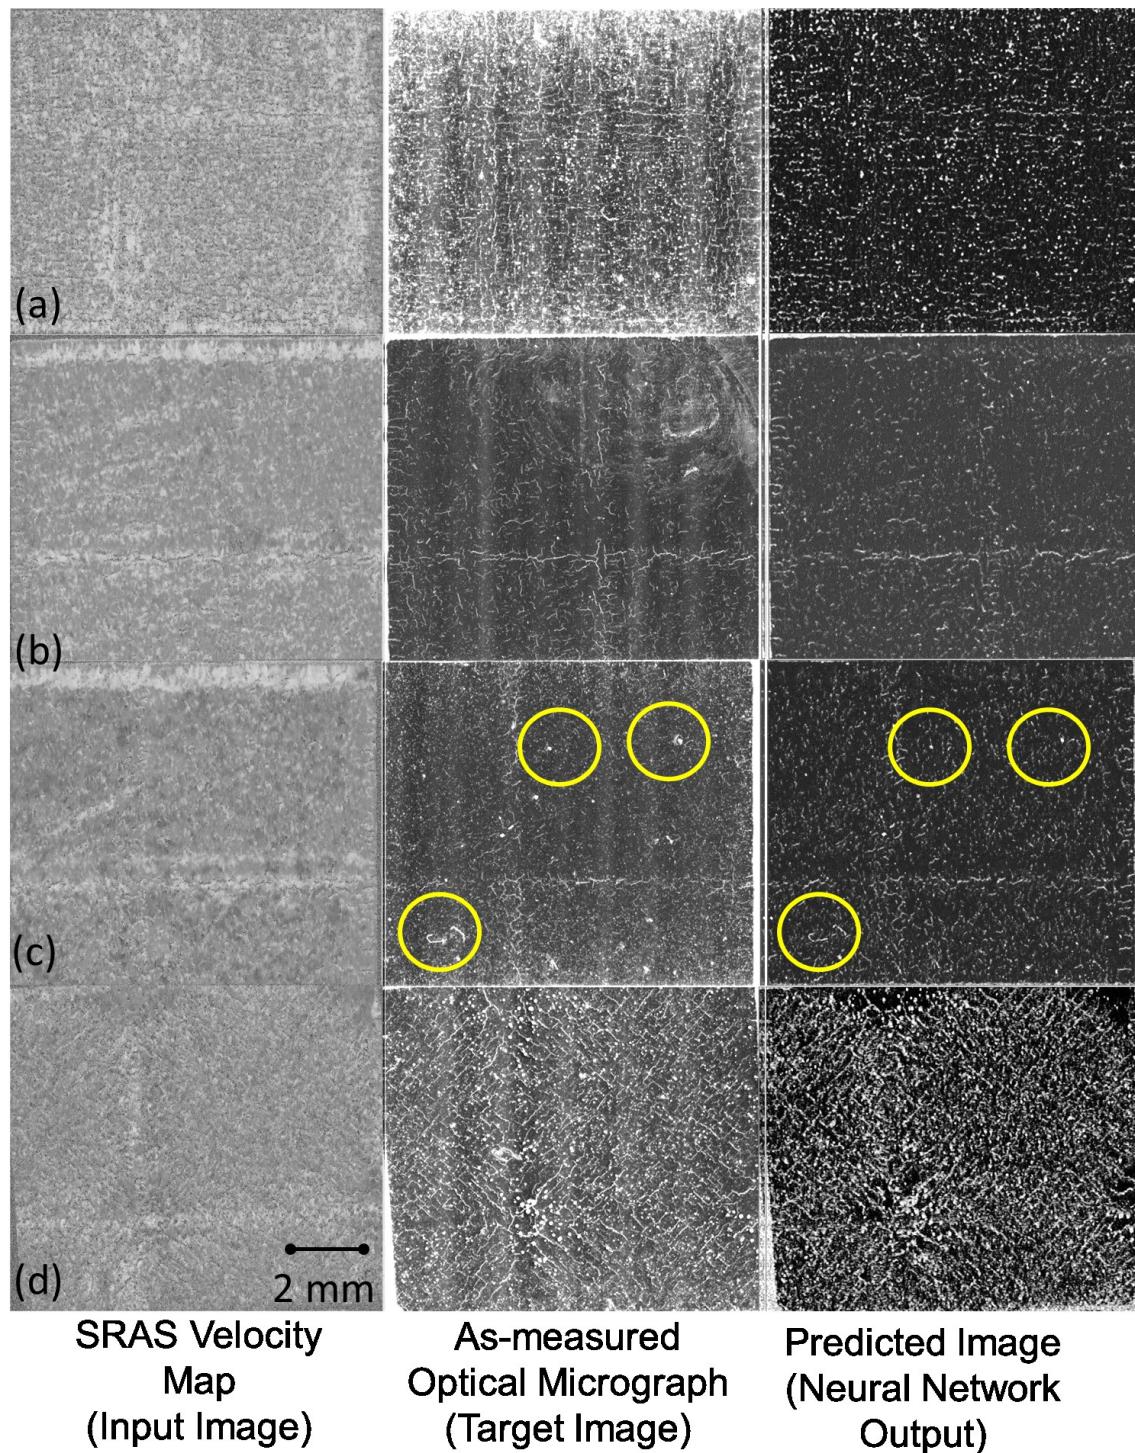


Figure 10: Results for the Nickel alloy samples (a) 0°, (b) 15°, (c) 30° and (d) 45°, that depicts a distinctive surface texture compared to the titanium-alloy samples. Representative areas of porosity are circled. The surface features are represented in the optical image recreated by DCB-MIR.

In all cases, the DCB-MIR predicted images have successfully captured the majority of defects accurately, in terms of size and location, whilst omitting features seen in the optical map from inconsequential variations in surface reflectance and scan artifacts. The predictive power of the approach is quantified in Table 2, which demonstrates a large improvement in cosine similarity across all sample images. For instance, the cosine similarity measured between the image mapped from the SRAS velocity map (predicted optical image) and corresponding optical micrograph for the titanium samples were in the range of 0.45 to 0.55, while in the case of the nickel samples the cosine similarity was in the range of 0.30 to 0.45.

As a quantitative comparison, the cosine similarity between the SRAS signal and optical micrograph was assessed to be in the range of -0.15 to 0.15. Accordingly, these results demonstrate that, using the convolutional neural network, the SRAS velocity maps are mapped to a form that captures the distinctive surface features, such as pores, which are hitherto occluded in the original SRAS velocity map signatures due to attendant noise. In essence, this approach has utilized the unique feature of SRAS where three distinct complementary datasets are generated under differing process conditions, and machines.

Table 2: Cosine similarities between original SRAS velocity maps and as-measured optical image pairs; and cosine similarity between the DCB-MIR predicted optical image and as-measured optical image pairs.

Image Label	Cosine Similarity between SRAS Acoustic Measurements and Measured Optical Image	Cosine Similarity between the optical image predicted by the DCB-MIR and as-measured optical image
<b>Ti(1) Sample 1</b>	-0.058	0.572
<b>Ti(1) Sample 2</b>	-0.104	0.450
<b>Ti(1) Sample 3</b>	0.153	0.241
<b>Ni 0°</b>	-0.126	0.526
<b>Ni 15°</b>	-0.071	0.391
<b>Ni 30°</b>	-0.059	0.278
<b>Ni 45°</b>	-0.099	0.471
<b>Ti(2) Sample 1</b>	-0.120	0.441
<b>Ti(2) Sample 2</b>	-0.133	0.542
<b>Ti(2) Sample 3</b>	-0.136	0.549

## Conclusions

In this work we developed and applied a deep learning convolutional neural network to map spatially resolved acoustic spectroscopy (SRAS) signatures of AM parts into a form resembling an optical micrograph, where only the defects are mapped on a binary scale. This form of information valuable because it is easily interpretable by the operator, and more pertinently, paves the way for a machine-learning module coupled to the SRAS instrument and embedded into a LPBF machine to automate defect detection.

Currently, detection of surface discontinuities by SRAS relies upon interpretation of the supplementary optical dataset. However, this dataset is affected by signal noise including, but not limited to, variations in surface reflection, surface artifacts, and variations in detection laser power. Looking ahead to the objective of using the SRAS instrument in-process during LPBF manufacture, the technique must be capable of reliably identifying defects, without repeated calibration studies and minimal user input.

Unlike traditional laboratory prepared samples, surfaces scanned during deposition will be rougher and be prone to loose powder pick-up. From the machine learning standpoint, this work advances the status quo in the analysis of image data obtained through contemporary process monitoring sensors, such as thermal and high speed optical cameras in AM, through a novel convolutional deep learning neural network-based image analysis approach.

Specific outcomes of this work are as follows:

1. Titanium (Ti-6Al-4V) and nickel (CM247-LC) alloy test samples were made using a laser powder bed fusion (LPBF) process under different process conditions and machines. The samples were polished, and SRAS and optical micrograph measurements were both taken offline. A deep learning neural network called Densely-

connected Convolutional Block Architecture for Multimodal Image Regression (DCB-MIR) was devised to translate the SRAS velocity signatures into a form resembling the corresponding optical micrograph. The network was trained with a data set that includes measurements from both the titanium- and nickel-alloy samples produced under differing LPBF processing conditions, as well as different LPBF systems, thus imparting generalizability. In other words, the network is robust to the type of alloy, process conditions, and LPBF machine.

2. The predicted defect map from the SRAS acoustic data were juxtaposed with the optical micrograph, and revealed a high degree of visual similarity. In particular, most of the large defects are predicted. The cosine similarity between the predicted micrograph images and as-measured optical micrograph averaged close to 50% (for titanium samples), which is a significant improvement over the cosine similarity between the as-measured optical micrograph and predicted micrograph (-15% ~ 15%).
3. Crucially, the predicted image reliably removed undesirable artifacts that obfuscate interpretation of the optical image. Furthermore, DCB-MIR is able to segment the acoustic map to reveal only surface breaking defects, ignoring other features such as variations in texture. Subsurface defects are not captured in the predicted map but in future these may be extracted by subtracting the predicted map from the original acoustic dataset.

Apart from ongoing efforts to incorporate the SRAS setup with a LPBF apparatus, the following avenues will be explored in the future work within the machine learning context:

- Full image reconstruction currently relies on simple average of sub-images. Increased final image fidelity could be achieved through the use of more intelligent combination techniques. Alternatively, the need to stitch together some images could be avoided through the adaptation of auto-regressive neural network techniques, similar to those used in PixelCNN [41].
- The current models were trained under highly data-limited conditions. As more SRAS data becomes available, significant improvements in the model output should be possible.
- The features of defects vary significantly between materials. In this research, a multi-material model was necessary due to data constraints, but future models specific to individual materials should produce higher fidelity outputs.

## Acknowledgements

One of the authors (PKR) thanks the United States National Science Foundation (NSF) for funding his work through the following grants CMMI-1719388 CMMI-1739696 and CMMI-1752069 (CAREER). Specifically, the concept of applying using in-process sensing in conjunction with real-time, big data analytics analytics techniques, such a deep learning neural networks for defect detection and isolation in metal additive manufacturing applications was proposed and funded through CMMI-1752069 (CAREER).

## References

- [1] Huang, Y., Leu, M. C., Mazumder, J., and Donmez, A., 2015, "Additive Manufacturing: Current State, Future Potential, Gaps and Needs, and Recommendations," *Journal of Manufacturing Science and Engineering*, **137**(1), pp. 014001-014001-014010.
- [2] Thompson, M. K., Moroni, G., Vaneker, T., Fadel, G., Campbell, R. I., Gibson, I., Bernard, A., Schulz, J., Graf, P., Ahuja, B., and Martina, F., 2016, "Design for Additive Manufacturing: Trends, opportunities, considerations, and constraints," *CIRP Annals*, **65**(2), pp. 737-760.
- [3] Huang, R., Riddle, M., Graziano, D., Warren, J., Das, S., Nimbalkar, S., Cresko, J., and Masanet, E., 2016, "Energy and emissions saving potential of additive manufacturing: the case of lightweight aircraft components," *Journal of Cleaner Production*, **135**, pp. 1559-1570.
- [4] Gorelik, M., 2017, "Additive manufacturing in the context of structural integrity," *International Journal of Fatigue*, **94**(2), pp. 168-177.
- [5] Seifi, M., Gorelik, M., Waller, J., Hrabe, N., Shamsaei, N., Daniewicz, S., and Lewandowski, J. J., 2017, "Progress Towards Metal Additive Manufacturing Standardization to Support Qualification and Certification," *JOM*, **69**(3), pp. 439-455.
- [6] Wycisk, E., Solbach, A., Siddique, S., Herzog, D., Walther, F., and Emmelmann, C., 2014, "Effects of Defects in Laser Additive Manufactured Ti-6Al-4V on Fatigue Properties," *Physics Procedia*, **56**, pp. 371-378.
- [7] Gibson, I., Rosen, D. W., and Stucker, B., 2010, *Additive manufacturing technologies: rapid prototyping to direct digital manufacturing*, Springer,
- [8] Baumers, M., Dickens, P., Tuck, C., and Hague, R., 2016, "The cost of additive manufacturing: machine productivity, economies of scale and technology-push," *Technological Forecasting and Social Change*, **102**, pp. 193-201.
- [9] Montazeri, M., and Rao, P., 2018, "Sensor-Based Build Condition Monitoring in Laser Powder Bed Fusion Additive Manufacturing Process Using a Spectral Graph Theoretic Approach," *Journal of Manufacturing Science and Engineering*, **140**(9), pp. 091002-091002-091016.
- [10] Sames, W. J., List, F., Pannala, S., Dehoff, R. R., and Babu, S. S., 2016, "The metallurgy and processing science of metal additive manufacturing," *International Materials Reviews*, **61**(5), pp. 315-360.
- [11] Imani, F., Gaikwad, A., Montazeri, M., Rao, P., Yang, H., and Reutzel, E., 2018, "Process Mapping and In-Process Monitoring of Porosity in Laser Powder Bed Fusion Using Layerwise Optical Imaging" *Journal of Manufacturing Science and Engineering*, **140**(10), pp. 101009-101023.

[12] Maskery, I., Aboulkhair, N., Corfield, M., Tuck, C., Clare, A., Leach, R. K., Wildman, R. D., Ashcroft, I., and Hague, R. J., 2016, "Quantification and characterisation of porosity in selectively laser melted Al–Si10–Mg using X-ray computed tomography," *Materials Characterization*, **111**, pp. 193-204.

[13] Montazeri, M., Yavari, R., Rao, P., and Boulware, P., 2018, "In-process Monitoring of Material Cross-Contamination Defects in Laser Powder Bed Fusion," *Journal of Manufacturing Science and Engineering*, **140**(11), pp. 111001-111020.

[14] O'Regan, P., Prickett, P., Setchi, R., Hankins, G., and Jones, N., 2016, "Metal Based Additive Layer Manufacturing: Variations, Correlations and Process Control," *Procedia Computer Science*, **96**, pp. 216-224.

[15] Everton, S. K., Hirsch, M., Stravroulakis, P., Leach, R. K., and Clare, A. T., 2016, "Review of in-situ process monitoring and in-situ metrology for metal additive manufacturing," *Materials and Design*, **95**, pp. 431-445.

[16] Gong, X., Lydon, J., Cooper, K., and Chou, K., 2014, "Beam speed effects on Ti–6Al–4V microstructures in electron beam additive manufacturing," *Journal of Materials Research*, **29**(17), pp. 1951-1959.

[17] Thorsten, S., Vera, J., F., S. R., and Carolin, K., 2015, "Influence of the Scanning Strategy on the Microstructure and Mechanical Properties in Selective Electron Beam Melting of Ti–6Al–4V," *Advanced Engineering Materials*, **17**(11), pp. 1573-1578.

[18] Koester, L. W., Taheri, H., Bigelow, T. A., Collins, P. C., and Bond, L. J., 2018, "Nondestructive Testing for Metal Parts Fabricated Using Powder-Based Additive Manufacturing," *Materials Evaluation*, **76**(4), pp. 514-524.

[19] Smith, R. J., Hirsch, M., Patel, R., Li, W., Clare, A. T., and Sharples, S. D., 2016, "Spatially resolved acoustic spectroscopy for selective laser melting," *Journal of Materials Processing Technology*, **236**, pp. 93-102.

[20] Hirsch, M., Dryburgh, P., Catchpole-Smith, S., Patel, R., Parry, L., Sharples, S. D., Ashcroft, I. A., and Clare, A. T., 2018, "Targeted rework strategies for powder bed additive manufacture," *Additive Manufacturing*, **19**, pp. 127-133.

[21] Vayssette, B., Saintier, N., Brugger, C., Elmay, M., and Pessard, E., 2018, "Surface roughness of Ti-6Al-4V parts obtained by SLM and EBM: Effect on the High Cycle Fatigue life," *Procedia Engineering*, **213**, pp. 89-97.

[22] Smith, R. J., Li, W., Coulson, J., Clark, M., Somekh, M. G., and Sharples, S. D., 2014, "Spatially resolved acoustic spectroscopy for rapid imaging of material microstructure and grain orientation," *Measurement Science and Technology*, **25**(5), p. 055902.

[23] Li, W., Sharples, S. D., Smith, R. J., Clark, M., and Somekh, M. G., 2012, "Determination of crystallographic orientation of large grain metals with surface acoustic waves," *The Journal of the Acoustical Society of America*, **132**(2), pp. 738-745.

- [24] Donachie, M. J., 2000, Titanium: a technical guide, ASM international,
- [25] Harris, K., Erickson, G., and Schwer, R., 1984, "MAR M 247 Derivations—CM 247 LC DS Alloy, CMSX® Single Crystal Alloys, Properties and Performance," 5th Int. Symp, pp. 221-230.
- [26] Hagan, M. T., Demuth, H. B., Beale, M. H., and De Jesús, O., 1996, Neural network design, PWS Publishing, Boston, MA
- [27] He, K., Zhang, X., Ren, S., and Sun, J., 2016, "Deep residual learning for image recognition," Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 770-778.
- [28] Luo, W., Li, Y., Urtasun, R., and Zemel, R., 2016, "Understanding the effective receptive field in deep convolutional neural networks," Advances in neural information processing systems, pp. 4898-4906.
- [29] Huang, G., Liu, Z., Weinberger, K. Q., and van der Maaten, L., 2017, "Densely connected convolutional networks," Proceedings of the IEEE conference on computer vision and pattern recognition, p. 3.
- [30] Dong, C., Loy, C. C., He, K., and Tang, X., 2016, "Image Super-Resolution Using Deep Convolutional Networks," IEEE Transactions on Pattern Analysis and Machine Intelligence, **38**(2), pp. 295-307.
- [31] Tai, Y., Yang, J., and Liu, X., 2017, "Image super-resolution via deep recursive residual network," The IEEE Conference on Computer Vision and Pattern Recognition (CVPR).
- [32] Eigen, D., Puhrsch, C., and Fergus, R., 2014, "Depth map prediction from a single image using a multi-scale deep network," Advances in neural information processing systems, pp. 2366-2374.
- [33] Long, J., Shelhamer, E., and Darrell, T., 2015, "Fully convolutional networks for semantic segmentation," Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 3431-3440.
- [34] Drozdzal, M., Vorontsov, E., Chartrand, G., Kadoury, S., and Pal, C., 2016, "The Importance of Skip Connections in Biomedical Image Segmentation," Springer International Publishing, pp. 179-187.
- [35] Xu, B., Wang, N., Chen, T., and Li, M., 2015, "Empirical evaluation of rectified activations in convolutional network," arXiv preprint arXiv:1505.00853
- [36] Ioffe, S., and Szegedy, C., 2015, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," arXiv preprint arXiv:1502.03167
- [37] Nguyen, H. V., and Bai, L., 2011, "Cosine Similarity Metric Learning for Face Verification," Springer Berlin Heidelberg, pp. 709-720.

- [38] Kingma, D., and Adam, J. B., 2015, "A method for stochastic optimisation," International Conference for Learning Representations.
- [39] Henderson, M. B., Arrell, D., Larsson, R., Heobel, M., and Marchant, G., 2004, "Nickel based superalloy welding practices for industrial gas turbine applications," Science and Technology of Welding and Joining, **9**(1), pp. 13-21.
- [40] Carter, L. N., Martin, C., Withers, P. J., and Attallah, M. M., 2014, "The influence of the laser scan strategy on grain structure and cracking behaviour in SLM powder-bed fabricated nickel superalloy," Journal of Alloys and Compounds, **615**, pp. 338-347.
- [41] van den Oord, A., Kalchbrenner, N., Espeholt, L., Vinyals, O., and Graves, A., 2016, "Conditional image generation with pixelcnn decoders," Advances in Neural Information Processing Systems, pp. 4790-4798.