

Conditional generative adversarial networks for in-situ layerwise additive manufacturing data

Christian Gobert^a, Edel Arrieta^b, Adrian Belmontes^b, Ryan B. Wicker^b, Francisco Medina^b,
Brandon McWilliams^c

^a Oak Ridge Associated Universities, Belcamp, MD

^b The University of Texas at El Paso, El Paso, TX

^c CCDC Army Research Laboratory, Aberdeen Proving Ground, MD

Abstract

Conditional generative adversarial networks (CGANs) learn a mapping from conditional input to observed image and perform tasks in image generation, manipulation and translation. In-situ monitoring uses sensors to obtain real-time information of additive manufacturing (AM) processes that relate to process stability and part quality. Understanding the correlations between process inputs and in-situ process signatures through machine learning can enable experimental-driven predictions of future process inputs. In this research, in-situ data obtained during a metallic powder bed fusion AM process is mapped with a CGAN. A single build of two turbine blades is monitored using EOSTATE Exposure OT, a near-infrared optical tomography system of the EOS M290 system. Layerwise images generated from the in-situ monitoring system were paired with a conditional image that labeled the specimen cross-section, laser-scan stripe overlap and z-distance to part surfaces. A CGAN was trained using the turbine blade data set and employed to generate new in-situ layerwise images for unseen conditional inputs.

Keywords: machine learning, additive manufacturing, conditional generative adversarial networks, in-situ monitoring

Introduction

Metallic powder bed fusion (PBF) additive manufacturing (AM) is a process in which near-net shape parts are built directly from digital files in a layer-by-layer manner [1-4]. In-situ monitoring uses sensors to obtain real-time information during the build, aiming to build correlations between in-situ process signatures and post-process material quality [5-7]. The PBF AM EOS M290 systems employ a high-resolution camera at high frequency in the near-infrared range, capturing a layer-wise image representative of the material melting behavior across the build-plate [8]. The correlations between post-process material quality and laser-material interactions is understood to be impactful [1-7]. Building an understanding through machine learning tools between in-situ information and process inputs is therefore desirable.

Machine learning encompasses classes of computer algorithms that map inputs to desired outputs by optimizing parameters without explicitly being programmed [9]. Convolutional neural networks (CNNs) are multi-stage algorithms that can learn features through application of trainable convolutional filters [10-12]. Conditional generative adversarial networks (CGANs) attempt to learn a mapping from a conditional input to an observed images by pitting two neural networks, a generator and discriminator, against each other. The generator attempts to produce realistic-looking images from conditional inputs while the discriminator aims to distinguish

between generated images (i.e. fake) and observed images (i.e. real) [13-15]. The pix2pix [14] network is a demonstrated CGAN for images up to a resolution 256×256 pixels, using a U-Net [16] architecture as the generator and a patch-based CNN as the discriminator.

In this research the pix2pix network architecture was applied to EOSTATE Exposure OT in-situ data generated from two printed turbine blades. High-resolution near-infrared images were generated for each layer of the build. Identical sized conditional images were constructed which contained information on the STL cross-section, laser-scan stripe overlap, and z-distance to part surfaces. A subset of the data was used to train the pix2pix network, the remaining data was used for validation. Conditional inputs were generated from outside the dataset using geometric shapes to view the CGAN's generalization of new process inputs.

Material and Methods

Specimens and Process

Two turbine blades were printed side by side on an EOS M290 system. Each turbine blade occupied a bounding box of 80mm height, 62.5mm width and 25.5mm depth, a 3D view is depicted in Figure 1. Process parameters were set for Inconel 718; scan strategy was stripe-based with a width of 10mm, overlap of 0.12mm, and hatch distance of 0.11mm while the laser was set to 285W at 960mm/s; contours are offset by 0.012mm. The upskin process used 153W at 600mm/s with hatch distance of 0.09mm, and the downskin process used 95W at 785mm/s with hatch distance of 0.1mm. Both, upskin and downskin consisted of being 3 layers in thickness. Layer height was set to 0.04mm with 2000 build layers.




Figure 1. 3D view of single turbine blade.

In-situ Data Stream

For each build layer, the EOS M290 records a near-infrared high-resolution optical image of the build surface recording the energy application in the melt-pool [8]. The sCMOS camera detects near-infrared images (900nm) by using a bandpass filter to block any emissions from plasma and reflected light from the laser. The Exposure OT system measures the intensity of the emissivity of the molten material, in a long-exposure photograph of the build plate at every layer. The amount

of emissivity can be interpreted as a fraction of the residual energy from the melting process; moreover, the law of thermal radiation formulated by Max Planck correlates emitted energy to temperature according to wavelength. In this context, image zones with marked intensity variations may be indicators for flaws and defects; it is important to mention, that a variation in intensity does not necessarily imply defects all the time. For example, too much emissivity could be detected in narrow geometries, given the saturation of energy due the short scan vectors, which may result in distortions or local metal vaporization. On the other hand, it is expected to have higher emissivity detected from the laser-scan stripe overlaps given that the laser spends more time along them. In this case, high emissivity from laser-scan stripe overlaps rarely represent defects. The resolution of the layerwise images is approximately $\sim 125 \mu\text{m}/\text{pixel}$ and are 8-bit RGB images where increasing image intensity corresponds to higher emissivity and increasing energy application. A sample image from the monitoring software is displayed in Figure 2, showing the captured image from a single layer of the turbine blade build.




Figure 2. Image of in-situ data stream from Exposure OT of single build layer of two turbine blades.

In figure 2 the stripes of red can clearly be seen that are a result of the laser-scan stripe overlaps. The dark red regions are a result of the upskin and downskin process parameters that are utilized for near-part surface regions. The increase in image intensity can be seen in the regions of acute angles created between the part contours and laser-scan stripe overlaps, most-likely due to the loitering of the laser in small regions during scanning.

Data Set

The data set for the CGAN network comprises of paired observed in-situ images and corresponding conditional images for each sample. A sample was generated for each turbine blade and build layer. The in-situ images were 8-bit RGB 512×512 pixels images cropped to the center of each turbine blade. The conditional images were 8-bit RGB 512×512 pixel images, centered on the turbine blade as well, with each image channel representing a process input. The red channel was used to indicate the locations of the part obtained from the STL, pixels with a red channel value equal to 255 indicated a part region. The green channel was used to indicate the location of laser-scan stripe overlaps, pixels with a green channel value equal to 255 indicated a region of laser-scan stripe overlap. The blue channel was used to indicate the distance in layers for an upskin or downskin process, pixels were assigned a standardized value that was mapped to the RGB space. Pixels were given an integer value from -3 to 3 standardized to $[0 \ 255]$ that indicated a pixels distance related to the upskin and downskin parameters. An input value of zero (standardized to

127 in RGB space) represented no upskin or downskin parameters. A pair of observed and conditional images is displayed in Figure 3.




Figure 3. Sample of observed in-situ image (left) and corresponding conditional image (right).

For network training, images were resized to 256×256 pixels due to the limitations in resolution for the pix2pix methodology and RGB values were mapped from $[0 \ 255]$ to $[0 \ 1]$. Generating the conditional images is time-consuming and only 582 samples were used, a sample encompasses a pair of observed and conditional images. In Isola et al. [14], demonstrations of the pix2pix network had been shown with data sets containing only 400 images.

Machine Learning Network

CGANs contain two neural networks, a generator (i.e. to generate realistic looking samples given a conditional input) and a discriminator (i.e. discriminate between observed and fake images). The generator input and output are the conditional image and the observed image, respectively. The generator is defined by down-sampling and up-sampling modules. A down-sampling module comprises of n number of 4×4 convolutions of stride 2, leaky ReLU activation [17] and batch normalization [18]. An up-sampling module comprises of a 2×2 up-sampling, n number of 4×4 convolutions of stride 1, leaky ReLU activation and batch normalization. The end of an up-sampling module is where the U-net skip connections are connected to the corresponding output of a down-sampling module. The generator comprises of five down- and up-sampling modules, from 256×256 pixels to 8×8 pixels. The number of filters used in each module increases with decreasing image size; being 32, 64, 128, 256, and 512 filters for each of the corresponding down- and up-sampling modules. The last layer of the generator is a 3-channel output to match the RGB observed images, using 32 4×4 convolutions and leaky ReLU activation followed by 3 4×4 convolutions with sigmoid activation [17].

The discriminator input is a channel-wise concatenated array of the conditional input and observed image (or generated fake image depending on the sample). The output of the discriminator CNN is a patch multi-channel image of 8×8 pixels. For real and fake samples the output patches are set to a value of 1 and 0, respectively. The discriminator CNN includes five down-sampling modules with the image size reduced from 256×256 pixels to 8×8 pixels. Each down-sampling module

comprises of n number of 4×4 convolutions with stride of 2, leaky ReLU activation and batch normalization. The final layer of the discriminator CNN is an 8×8 output generated from a 3×3 convolutional filter with sigmoid activation. The number of filters used in each module increases with decreasing image size; being 32, 64, 128, 256, and 512 filters for each of the corresponding down-sampling modules.

Network Training, Validation and Analysis

The set of turbine blade samples are divided into a training and validation set at a ratio of 80/20, randomly sampled from the set. The training data was used to update network parameters. The validation data was only used to evaluate the network performance at the end of a training epoch and in the post-process analysis. Training a CGAN involves the training of the discriminator and generator.

During training the discriminator attempts to learn parameters that can enable it to discriminate between real and fake images by minimizing the loss function. The loss function for the discriminator was measured through the mean squared error of the patches and accuracy was used as a metric to gauge performance. For each iteration the discriminator is trained, the generator is also trained. The generator is trained to maximize the loss of the discriminator for samples that are fake (i.e. fool the discriminator). In this research, a batch size of one was used and the Adam optimizer [19] with learning rate of 0.0001, beta one of 0.9 and beta two of 0.999. The training was completed for 200 epochs, where each epoch included a number (equal to the number of training samples) of randomly sampled training iterations.

Results and Discussion

Network Training and Results

The CGAN was trained for 200 epochs, with the validation loss of the discriminator and generator being recorded for each epoch, displayed in Figure 4. Figure 5 displays the validation accuracy of the discriminator for each epoch.




Figure 4. Validation loss of the generator (orange) and discriminator (blue) from each training epoch.




Figure 5. Validation accuracy of the discriminator from each training epoch.

Examining Figure 5, the lowest accuracy was achieved around epoch 26. Figure 6 displays a set of samples generated from the generator network at epoch 26.




Figure 6. Validation samples and generated images from the epoch 26 network. The top row contains observed images, middle row contains conditional images, and the bottom row contains generated images, paired column-wise.

The generator tended to represent the high emissivity in the acute angles formed by the laser-scan strip overlap regions and STL contours. The generator seems to match the high emissivity on and around laser-scan stripe overlaps. There is an inability to represent the noise seen in ground truth images, due to the underlying architecture of the network.




Figure 7. Sample of turbine with observed image (left), conditional image (center) and generated image (right).

Figure 7 shows an example of how the network tended to overestimate the average emissivity intensity in some circumstances. The assumption for this phenomenon is first that the network does not have a complete understanding of the part geometry (i.e. layers below) and subsequent impact

that would have on energy saturation. Second the distribution of data was drawn from a section of the turbine blade with larger emissivity values, possibly biasing the generator. Third the network does not know the true path of the laser in the regions of the stripes and the sequence of neighboring stripes which would have an impact on energy saturation.

Generated Novel Samples

With a trained generator, novel inputs can be processed and studied. In Figure 8, three unique geometries with two laser-scan stripe overlap spacings are displayed, generated from the trained network of epoch 26.




Figure 8. Generate novel samples. There are three unique geometries (columns) and two unique laser-scan stripe overlap spacings (rows).

Figure 8 shows an example of how the average intensity of the specimen slice decreases with increasing separation of laser-scan stripe overlaps, due to less stripe overlaps over the entire layer. In Figure 9 two geometries are depicted with two instances of laser-scan stripes overlap regions.




Figure 9. Generated images for two unique geometries (rows) and two unique laser-scan stripe overlap regions (columns)

In Figure 9 the geometry of the circle is invariant to the angle of the laser-scan stripe overlap, unlike in the square, which impacts the average intensity of the image as a result of the acute angles created between scan-line stripe overlaps and part contours.

Conclusion

A CGAN was demonstrated for in-situ data generated using the EOSTATE Exposure OT system on the EOS M290. The trained generator was able to mimic characteristics of the in-situ images, but failed to represent the stochastic nature of the images and in instances overestimated the average emissivity. This work demonstrated an initial intuition of CGANs with in-situ data. Future research needs to explore the generalization capabilities of CGANs, with increasing scope. Ideally the conditional pair to an observed image would encompass the complete PBF AM process input space, i.e. geometry, material, process parameters, with the aim of simulating PBF AM builds through CGANs. Applying CGAN architectures that can accommodate larger resolution images would be utilized as well, as the limit to this approach was 256 pixels and ~250 microns/pixel. Once a CGAN is trained, generating images is relatively time inexpensive. Ideally, simulating builds with CGANS would complement the slow and computational expensive physics based models with data-driven predictions and aid in optimization of process inputs for PBF builds.

References

- [1] Frazier, W Metal additive manufacturing: a review, *Journal of Materials Engineering and Performance*, 23:1917–1928, April 2014
- [2] I. Gibson, D. Rosen, B. Stucker, *Additive manufacturing technologies: rapid prototyping to direct digital manufacturing*, Springer, New York, 2009.

- [3] Sames W J, List, F A, Pannala S, Dehoff R R, Babu S S 2016 The metallurgy and processing science of metal additive manufacturing, *International Materials Review*, 61:315-360 March 2016
- [4] Debroy T, Wei H L, Zuback J S, Mukherjee T, Elmer J W, Milewski J O, et al. Additive manufacturing of metallic components – process, structure and properties, *Progress in Materials Science*, 92:112-224, October 2017
- [5] Everton SK, Hirsch M, Stravroulakis P, Leach RK, Clare AT, Review of in-situ process monitoring and in-situ metrology for metal additive manufacturing, *Materials & Design*, 95:431-435, April 2016.
- [6] G. Tapia, A. Elwany, A review on process monitoring and control in metal-based additive manufacturing, *Journal of Manufacturing Science and Engineering*, 136(6):60801–60811, December 2014
- [7] Zhang B, Li Y, Bai Q, Defect formation mechanisms in selective laser melting: a review, *Chinese Journal of Mechanical Engineering* 30:515-527, March 2017
- [8] https://www.eos.info/press/case_studies/optical-tomography-enables-considerable-savings-in-quality-assurance
- [9] Bishop C M, *Pattern Recognition and Machine Learning*, Springer, New York, 2006
- [10] LeCun Y, Bengio Y, Hinton G, Deep learning, *Nature* 521:436-444, May 2015
- [11] LeCun Y, Kavukcuoglu K, Farabet C, Convolutional networks and applications in vision, *Proceedings of IEEE International Symposium on Circuits and Systems*, 253-256, 2010
- [12] Krizhevsky A, Sutskever I, Hinton G E, ImageNet classification with deep convolutional neural networks *Advances in Neural Information Processing Systems*, 25 , 2012
- [13] Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y, Generative adversarial networks, *Advances in Neural Information Processing Systems*, 27, 2014
- [14] Isola P, Zhu J, Zhou T, Efros A, Image-to image translation with conditional adversarial networks, *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1125-1134, 2017
- [15] Wang T, Liu M, Zhu J, Tao A, Kautz J, Cantanzaro B, High-resolution image synthesis and semantic manipulation with conditional GANS, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018
- [16] O. Ronneberger, P Fischer and T. Brox, U-Net: Convolutional Networks for Biomedical Image Segmentation, *International Conference on Medical image computing and computer-assisted intervention*, 234-241, October 2015
- [17] Nwankpa C, Ijomah W, Gachagan A, Marshall S, Activation functions: comparison of trends in practice and research for deep learning , November 2018
- [18] Ioffe S, Szegedy C, Batch normalization accelerating deep network training by reducing internal covariate shift, May 2015
- [19] Kingma D, Lei Ba J, Adam: a method for stochastic optimization, January 2017