

Vision-Based Process Optimization for Hybrid Metal 3D Printing

Hybrid metal 3D printing combines Gas Metal Arc Welding (GMAW)-based deposition with CNC machining on a single machine, producing finished metal parts without separate post-processing equipment. This research focuses on using Wire Arc Additive Manufacturing (WAAM) for building metal parts layer by layer, offering high deposition rates and material efficiency. However, variations in weld-bead geometry due to process-parameter interactions result in uneven layer surfaces and increased machining requirements. The variation in height directly determines how much material machining must remove. Smaller height variation means tighter tolerances, shorter machining time, and less tool wear. Minimizing height variation requires careful process parameter selection, but existing measurement methods are either too complex or not suited for in-process use. This paper presents an algorithm that measures bead height variation from side-profile camera images and recommends parameter optimization to minimize height deviations. A U-Net segmentation model pretrained on ImageNet extracts the bead surface profile from each photograph. Training on a small set of manually labeled images was satisfactory to achieve strong segmentation performance in bead height detection. Bead height values extracted from the images were validated against Coordinate Measuring Machine (CMM) measurements and showed close agreement on peak height (H_{max}) across the tested beads. These validated measurements were then used to train a Gaussian Process Regression (GPR) model, which identifies a candidate parameter combination predicted to minimize height variation and reports a confidence interval for each prediction. Unlike deterministic models, GPR communicates the reliability of each prediction, helping operators decide whether a recommended setting is ready for production. The algorithm runs automatically with only a fixed camera, requiring no additional specialized metrology equipment.

Keywords: *hybrid metal 3D printing, bead height variation, Gaussian Process Regression, U-Net segmentation, process parameter optimization*

Introduction

The WAAM method deposits metal wire layer by layer using an electric arc. This method allows building near-net-shape parts at deposition rates and material costs that powder-bed additive manufacturing cannot match for medium-to-large components in titanium, steel, and aluminum (Williams et al. 2016). That cost advantage has driven growing adoption across aerospace, defense, automotive, and heavy industry.

The hybrid 3D metal printing system at Western Michigan University advances WAAM by integrating CNC machining on the same platform. After each bead layer is deposited by the GMAW welding head, a cutting tool mills the surface before the next layer is deposited (Ikonov 2020). This alternating

1 cycle produces dimensionally accurate, fully finished parts without relocating
2 the workpiece.

3 Within this hybrid cycle, the properties of the welding bead affect the final
4 part quality. When bead height varies, the CNC machining step must remove
5 more material to reach the required surface finish. That extra step increases
6 machining time and accelerates tool wear. The key metric is the difference
7 between the maximum and minimum bead height within a single deposit, $\Delta H =$
8 $H_{\max} - H_{\min}$. Minimizing ΔH reduces machining requirements, which is one of
9 the central economic drivers of the hybrid process.

10 The deposited bead height can be measured using metrology tools such as
11 Coordinate Measuring Machines (CMMs), structured-light scanners, and laser
12 profilometers. These tools are accurate but slow and expensive, and interrupt the
13 deposition workflow as each measurement requires the machine to stop, the part
14 to cool, and in some cases, the workpiece to be repositioned or transferred to a
15 separate measurement station before deposition can resume. This research takes
16 a different approach, to in-process measurement, using a fixed camera to capture
17 side-profile images of each bead immediately after deposition, without stopping
18 the machine or moving the workpiece. Next, a trained neural network extracts
19 the height profile directly from the image. The substrate plate, the flat metal base
20 on which the bead is deposited, is visible in every frame and serves as an in-
21 image calibration reference.

22 The welding bead height depends on process parameters like machine
23 deposition speed (feed rate), welding wire feed rate, nozzle standoff height from
24 the substrate, and welding power trim (voltage and amperage control). In this
25 study, four parameters were varied across defined ranges of feed rate from 203
26 to 406 mm/min (8 to 16 in/min), wire feed rate at 5,588 and 6,350 mm/min (220
27 and 250 in/min), nozzle standoff height at three levels between 9.525 and 12.70
28 mm (0.375 and 0.500 in), and trim at three levels between 0.9 and 1.1. These
29 ranges yield 162 possible parameter combinations in total. Testing all 162
30 experimentally is impractical, so 69 beads were deposited across representative
31 combinations, and a model trained on those measurements can predict ΔH for
32 any combination and find the setting that minimizes it.

33 To identify the optimal parameter combination, Gaussian Process
34 Regression (GPR) was chosen as the model because it returns a calibrated
35 confidence interval alongside every prediction. Other methods like Random
36 Forest or neural network approaches, do not provide this directly for the given
37 sample sizes. In practice with GPR, a prediction paired with a quantified interval
38 is far more useful to an operator than a bare number. It tells the operator directly
39 whether the recommended setting is reliable enough to run in production.
40

1 Literature Review

3 *Wire Arc Additive Manufacturing and Hybrid Processing*

5 WAAM has attracted significant research attention since the early 2000s as
6 a lower-cost alternative to powder-bed fusion for medium and large metal
7 components. Gas Metal Arc Welding (GMAW), Gas Tungsten Arc Welding
8 (GTAW), and plasma arc process variants have been demonstrated across
9 titanium alloys, aluminum, and steel (Williams et al., 2016). Plasma arc
10 deposition in particular has shown consistently good results for titanium alloys,
11 producing Ti-6Al-4V components with mechanical properties that meet
12 aerospace structural requirements (Martina et al., 2012).

13 Wire-feed processes when compared with other additive manufacturing
14 methods like powder-bed fusion and subtractive manufacturing use material
15 more efficiently and are better suited for building medium-to-large components
16 where minimizing waste and machining time matters (Ding et al., 2015; Guo and
17 Leu, 2013). Despite these advantages, controlling the geometry of each
18 deposited layer remains a persistent challenge. Thermal gradients between
19 successive passes induce distortion and residual stress that accumulate
20 throughout the build and ultimately affect the dimensional accuracy of the
21 finished part (Denlinger et al., 2015; Li et al., 2018). The same thermal history
22 also drives microstructure changes within the deposit, which is why monitoring
23 what happens during and immediately after each pass is important rather than
24 waiting until the part is complete (Xiong et al., 2020).

25 The hybrid 3D metal printing machine used in this study operates on a five-
26 axis CNC platform that alternates between GMAW deposition and milling
27 within each layer (Ikononov, 2020). Subsequent work from the same laboratory
28 demonstrated that forward-reverse and zigzag deposition path strategies reduce
29 height inconsistencies at layer transitions, achieving build heights above 8 inches
30 on the Z-axis (Ikononov, 2025). With deposition path strategies established, the
31 remaining lever for controlling bead geometry is the welding parameter
32 combination within each pass, which is the focus of this research. Even with
33 deposition path strategies in place, maintaining part quality and dimensional
34 accuracy across a full build remains sensitive to how well process parameters
35 are selected and controlled at each pass (Jafari et al., 2021).

36 *Bead Geometry Prediction: Classical and Machine Learning (ML) Approaches*

37 Predicting bead geometry from welding parameters has been studied across
38 several modeling approaches. Response Surface Methodology (RSM) is a
39 classical method, fitting a polynomial to data collected at a designed set of
40 parameter combinations and optimizing the fitted surface to identify optimal
41 settings. RSM works well for simple cases but breaks down when parameter
42 interactions are nonlinear or when more than a handful of inputs are involved,
43 both of which apply in the present study (Geng et al., 2017). Neural network
44 models applied to GMAW bead geometry showed early on that data-driven
45 approaches could capture these nonlinear relationships without assuming a fixed
46 functional form, though getting reliable predictions from small experimental
47 datasets was still a challenge (Xiong et al., 2014).

1 Artificial neural networks/ML handle more complex parameter-to-
2 geometry relationships than RSM without assuming a fixed functional form.
3 This makes them better suited for cases where multiple welding parameters
4 interact in ways a polynomial cannot fit. However, they overfit easily on small
5 datasets and return only a point estimate with no indication of how reliable that
6 prediction is (Wang et al., 2021).

7 Deep neural networks push accuracy further by incorporating additional
8 input features such as prior-layer geometry measurements, but this requires
9 stopping to measure each layer before the next deposit, which is impractical in
10 most WAAM setups and negates the speed advantage of the process. Combining
11 support vector machines with response surface methods has also been explored
12 as a way to identify optimal process windows more reliably (Shin, 2024; Xiao et
13 al., 2022). When the parameter space is continuous or too large for exhaustive
14 testing, Bayesian optimization offers a more principled way to find good settings
15 by learning from each experiment and focusing subsequent trials on the most
16 promising regions (Snoek et al., 2012). More broadly, the range of machine
17 learning tools available for manufacturing process control has grown
18 considerably, making it increasingly practical to adapt general-purpose
19 frameworks for specific deposition tasks (Müller and Guido, 2016).

20 Gaussian Process Regression (GPR) provides a distinct approach to this ML
21 space. It is non-parametric, making no assumption about the functional form of
22 the parameter-to-height deviation (ΔH) relationship, and returns a full predictive
23 distribution rather than a point estimate. This probabilistic output is particularly
24 valuable at low sample counts, which are common in WAAM experiments, and
25 has been demonstrated to perform well in metal AM contexts (Rasmussen and
26 Williams, 2006; Tapia et al., 2016). The present work applies the same GPR
27 method to WAAM bead height variation, with the key distinction that the
28 training targets are derived from an image rather than contact measurement.

29 30 *Vision-Based Measurement of Welding Bead Geometry*

31
32 Non-contact measurement of WAAM bead geometry covers a wide range
33 of accuracy, cost, and complexity. Among the more established techniques,
34 structured-light scanning projects a pattern of light onto the bead surface and
35 captures the distortion of that pattern with a camera to reconstruct a full 3D
36 surface map, while laser line scanning sweeps a laser line across the surface and
37 records the reflected profile to build up a point-by-point height map. Both
38 techniques produce detailed geometry data with high precision and have been
39 used in several research systems for in-process monitoring. However, that
40 capability comes at a cost with dedicated hardware, calibration fixtures, and
41 mandatory pauses between layers all adding overhead that slows down the build
42 cycle and makes these methods difficult to deploy in a production WAAM
43 environment.

44 A two-dimensional (2D) side-profile captured by a camera provides a
45 simpler and far more rapid measurement approach. The bead image is captured
46 with a high-resolution camera positioned perpendicular to the deposition

1 direction. Because the camera angle and distance are fixed, each pixel in the
2 image corresponds to a known physical dimension. A single calibration step,
3 measuring a reference object of known size visible in the frame, in this case, the
4 substrate plate. The thickness of the substrate plate establishes the pixel-to-
5 millimeter calibration factor that applies to every image in the dataset. The
6 practical difficulty is the feature extraction. Classical approaches like adaptive
7 thresholding work well under stable lighting but breaks down when contrast at
8 the bead edge is low or when a background object bleeds into the detection zone.

9 Deep learning segmentation handles these edge cases more gracefully. A
10 practical solution to this annotation problem came from transfer learning, where
11 a network already trained on a large general-purpose image dataset is adapted to
12 a new domain with far fewer labeled examples. The pretrained network brings
13 general visual knowledge that transfers well even when the new domain looks
14 very different from the original training data (Pan and Yang, 2010). One such
15 model, U-Net, originally developed by Ronneberger, Fischer, and Brox (2015)
16 for biomedical image segmentation, has become the standard architecture for
17 pixel-level tasks with limited labeled data. Its encoder-decoder structure with
18 skip connections preserves spatial resolution all the way to the output. That
19 matters here because the bead surface boundary is what the measurement
20 actually depends on.

21 Using a ResNet34 encoder pretrained on ImageNet reduces the annotation
22 burden further. ResNet architectures were built around the idea of residual
23 connections, which allow very deep networks to train stably by letting gradients
24 flow directly through shortcut paths rather than only through the learned layers.
25 That depth is what makes them effective feature extractors across a wide range
26 of visual tasks (He and Sun, 2015). The encoder already captures general edge
27 and texture features, so only the decoder needs domain-specific training. In this
28 work, 22 labeled bead images were enough to reach 93% pixel-level agreement
29 between the predicted bead region and the hand-drawn ground truth. Other
30 segmentation architectures were considered during development. DeepLabV3+
31 uses a different approach called Atrous Spatial Pyramid Pooling to capture
32 context at multiple scales simultaneously, which works well for natural images
33 but can struggle to preserve spatial resolution at small batch sizes or non-
34 standard image dimensions (Chen et al., 2018).

35 36 37 **Methodology**

38 39 *Experimental Setup and Dataset*

40
41 The welding beads in this study were deposited onto substrate plates using
42 a hybrid 3D metal printing machine. The GMAW welding head is mounted on a
43 five-axis CNC machine and deposits each bead along a straight path at a
44 controlled machine feed rate, with the wire feed rate and nozzle height standoff
45 from the substrate set for each bead. Four parameters were varied across the
46 experiment. Table 1 lists the parameter ranges. A total of 69 beads were

1 deposited across 12 plates, identified by the naming convention PlateID_BeadID
2 (e.g., P01_BC01 through P12_BC06).

3
4 **Table 1.** *Experimental Parameter Ranges*

Parameter	Min	Max	Values / Levels
Feed Rate mm/min (in/min)	203 (8)	406 (16)	203, 229, 254, 279, 305, 330, 356, 381, 406 (8, 9, 10, 11, 12, 13, 14, 15, 16)
Wire Feed Rate mm/min (in/min)	5,588 (220)	6,350 (250)	5,588, 6,350 (220, 250)
Nozzle Height mm (in)	9.525 (0.375)	12.700 (0.500)	9.525, 11.113, 12.700 (0.375, 0.4375, 0.500)
Trim	0.9	1.1	0.9, 1.0, 1.1

5
6 *Image Acquisition and Calibration*

7
8 After each bead cooled to ambient temperature, it was photographed with a
9 high-resolution camera fixed perpendicular to the bead's long axis. The camera
10 position, lens focal length, and manual focus were held constant across the entire
11 dataset. Between captures, the substrate plate was slid laterally to center the
12 target bead in the frame. This ensures the pixel-to-millimeter calibration is
13 identical for every image in the dataset.

14 The calibration reference is the substrate plate visible at the bottom of every
15 frame. The plate top surface defines $Z = 0$, which is the base of the deposited
16 bead, and the plate bottom edge provides the second reference point. The plate
17 thickness was measured to be 6.15 mm. For each image, the algorithm detects
18 both plate edges and computes the pixel-to-millimeter scale using the known
19 plate thickness. This gives a physical resolution of 0.015 mm per pixel, which is
20 adequate for measuring ΔH across the dataset.

21 Figure 1 shows a representative raw side-profile photograph of a WAAM
22 bead on the substrate plate. The bright region at the top of the frame is the
23 background above the bead crown. The dark central body is the deposited bead
24 itself. The plate top edge, visible at the bottom of the bead, defines the $Z = 0$
25 reference from which all height measurements are taken.

Figure 1. *Raw side-profile photograph of a steel WAAM bead on substrate plate*



26
27
28 *Vision Algorithm*

29
30 Classical Detection Baseline (Versions 1–5)

31 Five versions of a classical computer vision algorithm were developed and
32 tested before deep learning was adopted. These results of the five algorithms are

1 discussed here, as they identify the failure modes that led to the adoption of the
2 CNN approach rather than relying on an off-the-shelf solution.

3 Version 1 of vision algorithm used GrabCut, a classical interactive
4 segmentation algorithm, to separate the bead from the background using a
5 bounding box approach. At the macro level it correctly identified the bead
6 region, but the detected top edge tracked the silhouette rather than the actual
7 surface contour. This produced heights 0.5–1.5 mm above the true bead crown.
8 Version 2 replaced GrabCut with adaptive thresholding. For each image column,
9 background brightness was sampled from a fixed reference row range, and the
10 detection threshold was set at 70% of that local background. The first pixel
11 darker than the threshold was taken as the bead top. This tracked the actual
12 surface rather than the silhouette and gave the first physically meaningful
13 measurements. A Savitzky-Golay smoothing filter (window = 61, polynomial
14 order = 3) was applied to reduce noise in the extracted profile.

15 Versions 3 through 5 of vision algorithm added to Version 1 median-filter
16 spike removal (Version 3), background-patch clamping via NaN interpolation
17 (Version 4), and dynamic right-trim detection based on the edge profile gradient
18 (Version 5). The background-patch clamping addressed a recurring artifact
19 where a bright region caused the detector to jump to the image top boundary. By
20 Version 5, the classical algorithm correctly processed 51 of 69 images (74%).
21 The remaining 18 failures shared three root causes. The plate top row varied
22 from row 279 to row 350 across images, making a fixed reference unusable.
23 Some beads had low contrast at the bead-plate junction, where the threshold
24 could not separate the two surfaces. Several images also had a background bead
25 to the right of the target, confusing the boundary detector. Heuristic fixes for any
26 one of these problems tend to break the others. This motivated the transition to
27 deep learning segmentation.

28 29 CNN Segmentation Model

30 For each image, the trained U-Net will produce a binary segmentation mask.
31 Rather than assuming a fixed plate row, the algorithm identifies the median row
32 of the bottom plate region in the mask, then scans upward from that point to find
33 the sharpest brightness transition. That transition corresponds to the plate top
34 surface. This image-specific plate row detection is the main reason the CNN
35 algorithm succeeds on the 18 images where the classical algorithm failed.

36 The segmentation model is a U-Net with a ResNet34 encoder pretrained on
37 ImageNet, implemented using the segmentation-models-pytorch library. The U-
38 Net architecture was chosen for three reasons. First, its encoder-decoder
39 structure with skip connections preserves spatial resolution at the output. The
40 bead surface edge is a one-pixel-wide boundary that must be precisely localized,
41 so full output resolution is essential. Second, U-Net generalizes well from small
42 training datasets when a pretrained encoder is used. The encoder captures
43 general visual features and only the decoder needs to learn domain-specific
44 patterns. Third, U-Net avoids the spatial size collapse encountered with
45 DeepLabV3+ at this image resolution and batch size of one. The Atrous Spatial


1 Pyramid Pooling module in DeepLabV3+ compressed spatial dimensions to a
2 1×1 map, destroying all positional information (Chen et al., 2018).

3 Twenty-two bead images were manually annotated using LabelMe (Wada,
4 2016), an open-source Python tool for adding polygon labels to images. A single
5 polygon was drawn tracing the bead outline in each image. The annotations were
6 converted to binary PNG masks and split into 18 training and 4 validation images
7 using a fixed random seed. Training augmentations included horizontal flip
8 (probability 0.5), random brightness and contrast adjustment (0.4), Gaussian
9 noise (0.3), and Contrast Limited Adaptive Histogram Equalization (CLAHE)
10 normalization.

11 Training ran for 50 epochs with an AdamW optimizer at learning rate
12 1×10^{-4} , cosine annealing schedule, and a combined Binary Cross-Entropy and
13 Dice loss weighted equally. The region-of-interest image size used during
14 training was 384×768 pixels, with a batch size of 4. The best validation
15 Intersection over Union (IoU) of 0.93, a measure of how well the predicted bead
16 region overlaps with the hand-labeled ground truth, where 1.0 means perfect
17 overlap, was reached at epoch 20 and held through epoch 50 with no sign of
18 overfitting. This indicates that 22 labeled images are adequate for this task when
19 starting from an ImageNet-pretrained encoder.

20 Figure 2 shows the training convergence curves. The left plot tracks training
21 loss and validation loss across all 50 epochs while both curves decrease steadily
22 and converge by epoch 30 with no divergence between them, confirming that the
23 model did not overfit to the training images. The right plot shows validation IoU
24 across epochs, which climbs quickly and stabilizes at 0.93 by epoch 10,
25 remaining consistent through the end of training.

Figure 2. *U-Net training convergence*



26 Bead Edge Extraction and Measurement

27
28 The bead top edge is extracted as the topmost white pixel per column of the
29 mask, above the detected plate row. The raw edge undergoes three-pass spike
30 removal, pass 1 uses a rolling median deviation tolerance of 20 pixels, pass 2
31 tightens to 12 pixels, and pass 3 to 7 pixels. After spike removal, a Savitzky-
32 Golay filter (window = 41, polynomial order = 3) smooths the profile. Per-bead
33 measurement boundaries are stored in trim_config.json as a left limit and a right
34 limit. These exclude the bead start ramp, where height builds up to steady state,
35 and the bead end zone, where spatter or wire retraction drops the height. The
36

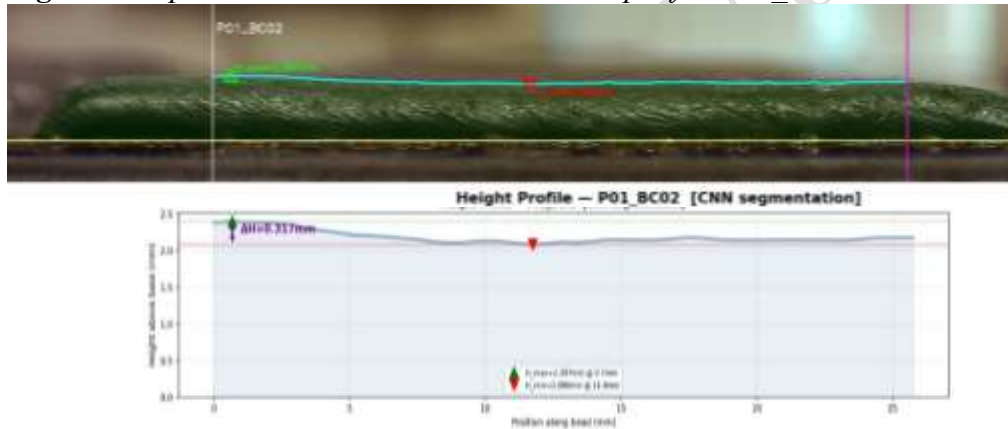
1 limits were set interactively for each of the 69 beads using an OpenCV tool. The
 2 operator views the detected edge overlaid on the image and clicks the correct
 3 boundaries. Figure 3 shows a representative annotated output.

4 From the trimmed, smoothed profile, H_{\max} is the maximum pixel height
 5 converted to millimeters, H_{\min} is the minimum, and $\Delta H = H_{\max} - H_{\min}$. The
 6 position of H_{\max} and H_{\min} along the bead axis is also recorded.

7 Figure 3 shows the measurement output for a representative bead,
 8 P01_BC02. The top panel displays the original side-profile photograph with the
 9 CNN segmentation overlay highlighted in color, showing the detected bead
 10 boundary and the measurement annotations for H_{\max} and H_{\min} marked at their
 11 respective positions along the bead. The bottom panel presents an enlarged
 12 graphical representation of the extracted height profile in millimeters, plotted
 13 along the bead length, with H_{\max} , H_{\min} , and the resulting ΔH clearly labeled. This
 14 enlarged view makes it easier to see how height varies across the bead and where
 15 the peak and minimum points fall relative to each other.

16

Figure 3. Representative bead measurement output for P01_BC02



17

18 *Cross-Batch Validation*

19

20 CMM measurements were taken directly on all 69 deposited steel beads.
 21 The probe was run along each bead to record H_{\max} and H_{\min} at consistent
 22 intervals, producing a physical reference dataset tied to the same beads
 23 photographed by the vision algorithm. Vision and CMM outputs were then
 24 compared bead by bead. Match quality was classified as high (H_{\max} error below
 25 0.05 mm), medium (0.05 to 0.183 mm), or low (above 0.183 mm). 62 of 69 beads
 26 (90%) fell in the high category, confirming that the vision algorithm produces
 27 H_{\max} values consistent with direct physical measurement across the full
 28 parameter range.

29

30

1 *Gaussian Process Regression Model*

3 Model Formulation

4 GPR models the unknown function $f(x)$ by mapping process parameters x
 5 = [feed rate, wire feed rate, nozzle height, trim] to ΔH as a sample from a
 6 Gaussian process with mean zero and covariance described by a kernel function
 7 $k(x, x')$. Given n training observations $\{X, y\}$, the predictive distribution at a new
 8 point x^* is Gaussian with mean and variance, as shown in equation 1 below:

$$10 \quad \mu(x^*) = k(x^*, X)[K(X, X) + \sigma_n^2 I]^{-1} y$$

$$12 \quad \sigma^2(x^*) = k(x^*, x^*) - k(x^*, X)[K(X, X) + \sigma_n^2 I]^{-1} k(X, x^*)$$

15 where $[K(X, X)]$ is the $n \times n$ training kernel matrix and σ_n^2 is the noise variance,
 16 see equation 2. Both equations have closed-form solutions, so prediction requires
 17 no iterative inference (Rasmussen and Williams 2006). The 95% confidence
 18 interval at x^* is $\mu(x^*) \pm 1.96 \sigma(x^*)$. This interval is the primary output reported
 19 for the optimal parameter recommendation.

21 Kernel and Hyperparameters

22 The GPR model uses a Matérn kernel with smoothness parameter $\nu = 2.5$.
 23 This controls how the model interpolates between observed data points. It was
 24 chosen over the squared exponential kernel because welding processes are not
 25 perfectly smooth. A small change in feed rate or nozzle height can produce a
 26 noticeable shift in bead geometry, and the Matérn kernel captures this more
 27 realistically. All four input parameters were scaled to the same range before
 28 fitting, so no single parameter dominates due to its units. The three kernel
 29 hyperparameters controlling prediction amplitude, length scale, and noise level
 30 were tuned automatically from the training data, using 15 independent starting
 31 points to avoid local optima.

32 Training and Parameter Search

33 The three beads classified as LOW confidence in the cross-batch validation,
 34 specifically P04_BC05, P05_BC05, and P05_BC06, all had H_{\max} match errors
 35 above 0.183 mm and were excluded from the training set, leaving 66 samples.
 36 Model performance was evaluated on the training set (MAE, R^2 , Pearson r) and
 37 by five-fold cross-validation (CV), in which the 66 samples were split into five
 38 roughly equal folds and each fold was predicted by a model trained on the
 39 remaining four. CV provides an estimate of out-of-sample prediction error that
 40 is independent of the training set fit.

41 After training, the model was evaluated at all 162 combinations in the
 42 discrete parameter grid (9 feed rates \times 2 wire feeds \times 3 nozzle heights \times 3 trims).
 43 Each combination received a predicted mean ΔH and a 95% confidence interval.
 44 The combination with the lowest predicted mean is reported as the recommended
 45 parameter setting.

1 *How We Used Parameters to Predict Variation in Height (ΔH)*

2

3 Each bead in the dataset has four input parameters and one measured ΔH
4 value from the vision algorithm. We used these as input-output pairs to train the
5 GPR model. The four inputs are feed rate, wire feed rate, nozzle height, and trim.
6 The output is the ΔH measured from the bead photograph. So, the model learns
7 what combinations of those four parameters tend to produce flat beads and which
8 ones tend to produce variable ones.

9 Before fitting the model, all four parameters were normalized to a common
10 scale. Feed rate, wire feed rate, nozzle height, and trim each have different units
11 and ranges. Normalizing puts them on equal footing so the model does not treat
12 feed rate as more important just because its raw numbers are larger. Scikit-
13 learn's StandardScaler was used to do this, which shifts each input to zero mean
14 and unit variance.

15 The GPR model was then fit using GaussianProcessRegressor package from
16 scikit-learn (Pedregosa et al., 2011). The kernel was set to a constant multiplied
17 by a Matérn kernel with
18 $\nu = 2.5$, plus a White kernel for noise. Hyperparameters were tuned by
19 maximizing log marginal likelihood, with 15 restarts to avoid getting stuck in a
20 local optimum. Once the model is trained, predicting ΔH for any new parameter
21 combination is a single function call that returns both a mean prediction and a
22 standard deviation.

23 To find the best setting, we built a full grid of all 162 parameter
24 combinations from the experimental ranges in Table 1. We passed every
25 combination through the trained model and collected the predicted ΔH and the
26 95% confidence interval for each one. The combination with the lowest
27 predicted mean ΔH was selected as the recommendation. This grid search
28 approach is straightforward and works well here because the parameter space is
29 discrete and not too large. A more complex search strategy like Bayesian
30 optimization would be useful if the parameter space were continuous or much
31 larger (Snoek et al., 2012).

32

33

34 **Results**

35

36 *Vision Algorithm Measurement Outputs*

37

38 The algorithm successfully processed all 69 bead images. Table 2
39 summarizes the measurement statistics. Height variation (ΔH) ranged from
40 0.100 mm for the flattest bead (P10_BC04) to 1.067 mm for the most variable
41 (P05_BC06). The distribution is right-skewed, with most beads falling between
42 0.2 and 0.6 mm and a smaller group above 0.7 mm. The calibration scale of
43 66.34 px/mm was consistent across all images, confirming that the fixed camera
44 protocol maintained reproducibility across all 69 beads.

45

1 **Table 2.** *Vision Algorithm Measurement Summary*

Metric	Value	Remarks
Beads processed	69	Plates P01–P12
Calibration scale	66.34 px/mm	Plate thickness = 6.15 mm
H_{\max} range (mm)	1.991 – 3.878	
H_{\min} range (mm)	1.695 – 3.457	
ΔH range (mm)	0.100 – 1.067	
Mean ΔH (mm)	0.434	SD = 0.184 mm
Flattest bead	P10_BC04: $\Delta H = 0.100$ mm	
Most variable bead	P05_BC06: $\Delta H = 1.067$ mm	
CNN segmentation IoU	0.93	Best validation epoch 20/50
Classical algorithm success rate	51/69 (74%)	Before CNN adoption

2

3

4

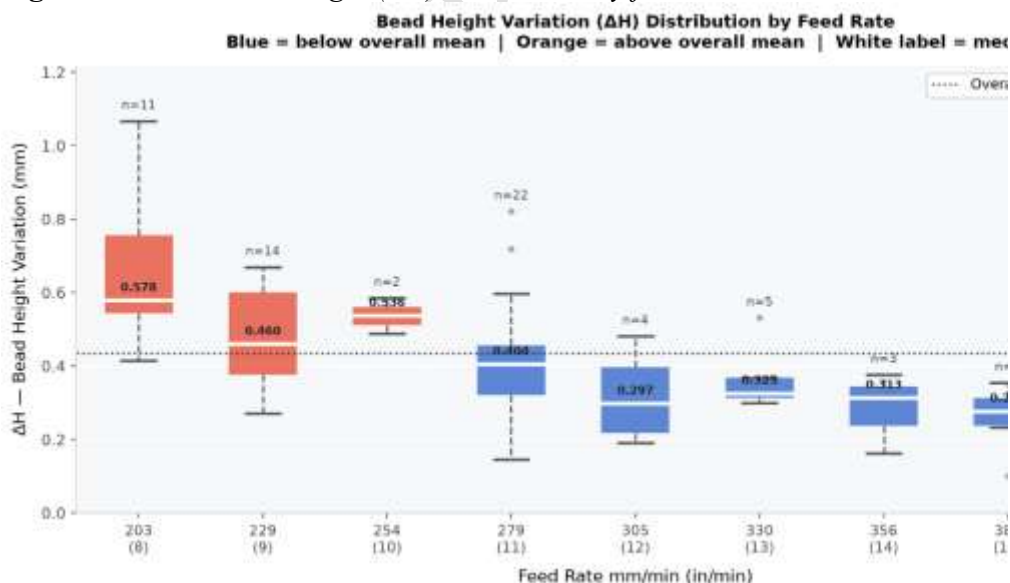
5

6

7

Figure 4 shows the ΔH distribution across all 69 beads grouped by feed rate. Median ΔH drops from 0.578 mm at 203 mm/min (8 in/min) to 0.348 mm at 406 mm/min (16 in/min). The spread also tightens at higher feed rates, meaning not just lower average variation but more consistent bead geometry. Feed rate is clearly the dominant driver of bead flatness.

Figure 4. *Variation in Height (ΔH) distribution by feed rate across all 69 beads*



8

9

Cross-Batch Validation Results

10

11

12

13

14

Table 3 presents the cross-batch validation metrics. The most operationally relevant result is H_{\max} agreement. 62 of 69 bead pairs (90%) fell in the HIGH confidence category with a mean match error of 0.017 mm against the physical reference data, which is well within the tolerance needed for machining

allowance decisions. The Pearson correlation between vision-derived ΔH and physical ΔH was $r = 0.361$ ($p = 0.0023$), indicating a statistically significant but moderate linear relationship between the two datasets. A correlation this low does not mean the algorithm is unreliable, but it reflects the fact that ΔH was measured across two separate deposition campaigns under slightly different conditions, introducing genuine variability in the bead profiles that neither dataset can fully account for. Both datasets agree closely on H_{\max} , but the full height profiles differ between campaigns, producing a $+0.212$ mm mean offset in ΔH . That offset reflects genuine deposition variability, not a systematic algorithm bias. For machining purposes, H_{\max} is the metric that matters, and on that measure the algorithm performs well.

Table 3. *Cross-Batch Validation Metrics*

Metric	Value	Interpretation
Total matched bead pairs	69	
HIGH confidence (H_{\max} error < 0.05 mm)	62 (90%)	Excellent H_{\max} agreement
MEDIUM confidence (0.05–0.183 mm)	4 (6%)	
LOW confidence (> 0.183 mm)	3 (4%)	Excluded from GPR training
Mean H_{\max} match error	0.017 mm	
Max H_{\max} match error	0.183 mm	3 bead pairs
Pearson r (ΔH : vision vs. physical)	0.361	$p = 0.0023$
Mean ΔH offset (vision – physical)	+0.212 mm	Reflects batch differences

GPR Model Performance

Table 4 presents the GPR model performance metrics. After hyperparameter optimization by log marginal likelihood maximization, the kernel settled on a Matérn length scale of 4.46 and a white noise level of 0.500, which is at the upper bound. That high noise level indicates the training labels contain variability that the four input parameters alone cannot explain. Some of that reflects real process scatter. Some reflects the cross-batch offset discussed above. Training MAE is 0.096 mm and R^2 is 0.375 on the training set. The five-fold cross validation MAE of 0.117 ± 0.029 mm is the more useful number, as it estimates out-of-sample error on unseen beads. The cross validation R^2 of -0.104 is discussed in Section 5.

Table 4. *GPR Model Performance Metrics*

Metric	Value	Notes
Training samples	66	HIGH + MEDIUM confidence
Excluded (LOW confidence)	3	P04_BC05, P05_BC05, P05_BC06
Kernel (optimized)	Constant \times Matérn($\nu=2.5$) + White	
Length scale ℓ	4.46 (normalized)	

Noise level σ^2_n	0.500	At upper bound
Log marginal likelihood	-91.59	
Training MAE	0.096 mm	
Training R^2	0.375	
Training Pearson r	0.625	$p < 0.001$
5-Fold CV MAE	0.117 ± 0.029 mm	
5-Fold CV R^2	-0.104 ± 0.133	See Section 5; negative CV R^2 expected when label noise \approx target variance
Mean 95% CI width	± 0.233 mm	

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

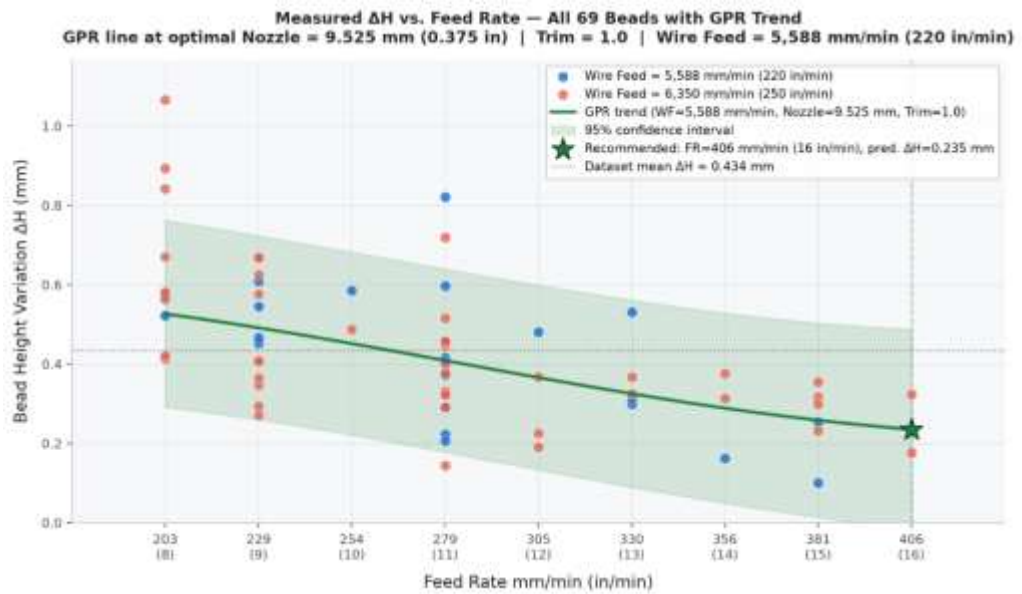
26

The five-fold cross-validation R^2 of -0.104 means that when the model is tested on beads it has not seen during training, it does not predict ΔH better than simply guessing the dataset average every time. This sounds like a poor result but it is actually expected given the nature of the data. The ΔH values used to train the model came from two separate deposition campaigns, the vision-derived measurements and the CMM reference measurements, and these two campaigns produced a mean offset of roughly 0.2 mm between them. That offset is not a measurement error but a reflection of real differences in how the beads were deposited across the two sessions. The problem is that this cross-batch variability is about the same size as the natural spread of ΔH values in the dataset, which has a standard deviation of 0.184 mm. When the noise in the training labels is as large as the variation the model is trying to predict, no model, regardless of how it is built, can produce a reliable cross-validation R^2 . The GPR model recognizes this situation on its own. During training, the kernel optimization pushed the noise parameter to its maximum allowed value, which is the model's way of signaling that the labels contain variability it cannot explain from the four input parameters alone. This is not a failure, it is the model being honest about the limits of the available data.

Figure 5 plots the measured ΔH for all 69 beads against feed rate, with the GPR predicted trend overlaid. The green line drops from a predicted 0.578 mm at 8 in/min to 0.235 mm at 16 in/min, confirming that higher feed rates produce more uniform beads. The shaded band shows the 95% confidence interval, which widens at the parameter space boundaries where training data is sparse.

1 *Recommended Parameter Combination*

2

Figure 5. ΔH vs. feed rate — measured beads and GPR predicted trend

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

Table 5. *Recommended Parameter Combination*

Parameter	Recommended Value	Unit
Feed Rate	406 (16)	mm/min (in/min)
Wire Feed Rate	5,588 (220)	mm/min (in/min)
Nozzle Height	9.525 (0.375)	mm (in)
Trim	1.0	dimensionless
Predicted ΔH (mean)	0.235	mm
95% CI	$[-0.017, 0.487]$	mm
Prediction uncertainty σ	0.129	mm
Reduction vs. dataset mean	46%	0.434 mm \rightarrow 0.235 mm

20

21

1 Discussion

2
3 The most practically significant result from the vision algorithm is the 0.017
4 mm mean H_{\max} error against the physical reference dataset. The machining step
5 on the hybrid system typically removes 0.5 to 1.5 mm per pass. Peak bead height
6 is known to within 0.017 mm, which is well within the tolerance needed for
7 machining allowance planning. That accuracy comes from a fixed high-
8 resolution camera and polygon annotation across 22 images.

9 The classical algorithm's 74% success rate means the basic imaging
10 geometry works in three out of four cases with no machine learning at all. The
11 remaining 26% is not random noise, but failures cluster in beads caused by
12 varied plate row position, a low-contrast bead, or a background bead bleeding
13 into the detection zone. Heuristic patches for any one of those failure modes tend
14 to break the others. A CNN trained on a diverse image set handles all three
15 without separate fixes. To apply this method on a new WAAM process, it is
16 recommended to start with the classical algorithm, then train the CNN once 20
17 or more labeled images are annotated.

18 An IoU of 0.93 from 22 training images is a strong result for a domain-
19 specific segmentation task. It means 93% of each bead's pixel area is correctly
20 classified on held-out validation images. Misclassification mostly affects
21 isolated boundary pixels, which are averaged out by the Savitzky-Golay
22 smoothing step. The key enabler is ImageNet pretraining on the ResNet34
23 encoder. General edge and texture knowledge from that pretraining transfers
24 well enough that only 22 domain-specific annotations were needed.

25 The response surface shows physical consistency and practical usefulness.
26 We found that feed rate has the dominant effect amongst all. Higher travel speed
27 means less energy deposited per unit length, which produces a flatter and more
28 uniform bead profile. That trend runs across the full parameter range and is
29 robust to label noise because it is a global pattern rather than a local one. An
30 operator can read the response surface and immediately know to stay in the high-
31 feed-rate zone without parsing any statistical output.

32 Unlike deterministic models such as neural networks or Random Forest,
33 GPR reports not just a predicted ΔH value but also how confident that prediction
34 is at every point in the parameter space (Rasmussen and Williams, 2006). This
35 confidence information is useful because it tells the operator where the model's
36 recommendations can be trusted and where more experimental data is needed.
37 Confidence is highest near feed rates of 254 to 305 mm/min (10 to 12 in/min),
38 where the training dataset has the most beads and the model has seen enough
39 examples to make reliable predictions. The recommended optimum sits at 406
40 mm/min (16 in/min), which is at the edge of that well-sampled zone. Because
41 fewer training beads exist at that setting, the model is less certain there and
42 reports a wider confidence interval — a prediction uncertainty of $\sigma = 0.129$ mm
43 and a 95% CI width of 0.254 mm. In practical terms this means the
44 recommended setting is promising but not yet fully confirmed by the available
45 data. Depositing three to five validation beads at that setting and comparing the
46 measured ΔH against the predicted range would directly confirm whether the

1 recommendation holds in practice and would also narrow the confidence interval
2 for that region of the parameter space.

3 The recommended setting of Feed Rate = 406 mm/min (16 in/min), Wire
4 Feed Rate = 5,588 mm/min (220 in/min), Nozzle Height = 9.525 mm (0.375 in),
5 and Trim = 1.0 pairs the highest feed rate tested in this study, sitting at the
6 boundary of the experimental range, with a lower wire feed and smallest nozzle
7 standoff. Higher feed rate means the torch covers more path per unit time, which
8 reduces heat input per unit length and physically explains the flatness gain. As a
9 secondary benefit, the higher travel speed also improves throughput. The trained
10 GPR model supports interactive use beyond the grid search. Given any
11 parameter combination, it returns a predicted ΔH with a 95% confidence interval
12 in seconds, which is useful when material constraints or joint geometry force a
13 setting away from the recommended optimum and the operator needs to know
14 the expected flatness penalty before committing to a build.

15 Future work should include the deposition of additional beads at high-
16 uncertainty parameter combinations, particularly at low feed rates and the
17 extremes of nozzle height, which would substantially reduce confidence interval
18 width across the response surface. Also, moving the system toward the real-time
19 closed-loop control by capturing the bead image immediately after deposition
20 while the machine repositions for the next pass within a build.

23 Conclusions

24
25 A U-Net segmentation model with a ResNet34 encoder trained on only 22
26 manually labeled images successfully achieves a segmentation IoU of 0.93,
27 confirming that ImageNet-pretrained encoders enable remarkably effective
28 manufacturing segmentation with limited annotation effort. This is a strong
29 result compared to classical segmentation approaches, which typically require
30 far larger labeled datasets to reach comparable accuracy and still struggle with
31 variable imaging conditions. The vision algorithm derived from this model
32 reliably measures bead peak height with a consistently low mean error against
33 physical reference data across 90% of matched bead pairs, achieving a level of
34 measurement accuracy that is well suited for machining allowance planning in a
35 production WAAM environment.

36 Classical computer vision methods achieve 74% success on the 69-bead
37 dataset but fail systematically on images with variable plate position, low
38 contrast, or background interference. The CNN handles all three of these
39 effectively and without separate fixes because it learned directly from labeled
40 bead images rather than relying on fixed rules, demonstrating clear superiority
41 over classical approaches in challenging imaging conditions.

42 The GPR model results are consistent with higher feed rates producing
43 lower ΔH , which lines up well with the physics of reduced heat input per unit
44 length. The model also clearly shows that the high-feed-rate region has the
45 widest prediction intervals, meaning that depositing a few more beads there
46 would yield the most reliable and confident parameter recommendations for that
47 zone.

1 The recommended parameter combination of Feed Rate = 406 mm/min (16
2 in/min), Wire Feed Rate = 5,588 mm/min (220 in/min), Nozzle Height = 9.525
3 mm (0.375 in), and Trim = 1.0 predicts $\Delta H = 0.235$ mm, a considerable 46%
4 reduction from the dataset mean, pending experimental validation. The
5 accompanying 95% confidence interval communicates this prediction
6 uncertainty directly and transparently to the operator rather than presenting a
7 bare point estimate.

8 The complete algorithm runs autonomously without manual intervention
9 and requires no specialized metrology equipment beyond a high-resolution fixed
10 camera, making it readily and practically deployable in WAAM environments
11 where structured-light or CMM-based measurement is impractical.

12 This work considerably advances the hybrid 3D metal printing research
13 program at Western Michigan University (Ikononov, 2020; Ikononov, 2025)
14 with a robust, data-driven measurement and optimization layer that brings the
15 process closer to fully automated unattended operation. Expansion of the
16 training dataset to boundary parameter regions is identified as a promising next
17 step toward achieving comprehensive in-process optimization.

20 References

- 21
22 Chen LC, Zhu Y, Papandreou G, Schroff F and Adam H (2018) Encoder-decoder with
23 atrous separable convolution for semantic image segmentation. In Proceedings of the
24 European Conference on Computer Vision (ECCV), Munich, Germany, 8–14
25 September 2018, pp. 801–818.
- 26 Denlinger ER, Heigel JC, Michaleris P and Palmer TA (2015) Effect of inter-layer dwell
27 time on distortion and residual stress in additive manufacturing of titanium and nickel
28 alloys. *Journal of Materials Processing Technology* 215: 123–131.
- 29 Ding D, Pan Z, Cuiuri D and Li H (2015) Wire-feed additive manufacturing of metal
30 components. *International Journal of Advanced Manufacturing Technology* 81(1–4):
31 465–481.
- 32 Geng H, Xiong J, Huang D, Lin X and Li J (2017) A prediction model of layer geometrical
33 size in wire and arc additive manufacture using response surface methodology.
34 *International Journal of Advanced Manufacturing Technology* 93(1–4): 175–186.
- 35 Guo N and Leu MC (2013) Additive manufacturing: technology, applications and research
36 needs. *Frontiers of Mechanical Engineering* 8(3): 215–243.
- 37 He K and Sun J (2015) Convolutional neural networks at constrained time cost. In
38 Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition
39 (CVPR), Boston, MA, USA, 7–12 June 2015.
- 40 Ikononov P (2020) Hybrid manufacturing system. U.S. Patent No. 10,688,581 B2.
41 Washington, DC: U.S. Patent and Trademark Office.
- 42 Ikononov P (2025) Hybrid 3D metal printing process optimization using ML and AI.
43 *Athens Journal of Technology and Engineering* 12(3): 181–194.
- 44 Jafari D, Vaneker THJ and Gibson I (2021) Wire and arc additive manufacturing:
45 opportunities and challenges to control the quality and accuracy of manufactured parts.
46 *Materials and Design* 202: 109471.
- 47 Kingma DP and Ba J (2015) Adam: a method for stochastic optimization. In Proceedings
48 of the International Conference on Learning Representations (ICLR), San Diego, CA,
49 USA, 7–9 May 2015.

- 1 Li Y, Han Q, Zhang G and Horváth I (2018) A layers-based approach for reducing distortion
2 in WAAM. *International Journal of Advanced Manufacturing Technology* 96(9–12):
3 3809–3820.
- 4 Loshchilov I and Hutter F (2019) Decoupled weight decay regularization. In *Proceedings*
5 *of the International Conference on Learning Representations (ICLR)*, New Orleans,
6 LA, USA, 6–9 May 2019.
- 7 Martina F, Mehnen J, Williams SW, Colegrove P and Wang F (2012) Investigation of the
8 benefits of plasma deposition for the additive layer manufacture of Ti-6Al-4V. *Journal*
9 *of Materials Processing Technology* 212(6): 1377–1386.
- 10 Müller AC and Guido S (2016) *Introduction to Machine Learning with Python*. Sebastopol,
11 CA: O'Reilly Media.
- 12 Pan SJ and Yang Q (2010) A survey on transfer learning. *IEEE Transactions on Knowledge*
13 *and Data Engineering* 22(10): 1345–1359.
- 14 Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M,
15 Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher
16 M, Perrot M and Duchesnay E (2011) Scikit-learn: machine learning in Python.
17 *Journal of Machine Learning Research* 12: 2825–2830.
- 18 Rasmussen CE and Williams CKI (2006) *Gaussian Processes for Machine Learning*.
19 Cambridge, MA: MIT Press.
- 20 Ribeiro MT, Singh S and Guestrin C (2016) Why should I trust you? Explaining the
21 predictions of any classifier. In *Proceedings of the ACM SIGKDD International*
22 *Conference on Knowledge Discovery and Data Mining*, San Francisco, CA, USA, 13–
23 17 August 2016, pp. 1135–1144.
- 24 Ronneberger O, Fischer P and Brox T (2015) U-Net: convolutional networks for biomedical
25 image segmentation. In *Proceedings of the Medical Image Computing and Computer-*
26 *Assisted Intervention (MICCAI)*, Munich, Germany, 5–9 October 2015, *Lecture Notes*
27 *in Computer Science*, vol. 9351. Springer, Cham.
- 28 Shin JH (2024) Prediction of metal additively manufactured bead geometry using deep
29 neural network. *Sensors* 24(19): 6250.
- 30 Snoek J, Larochelle H and Adams RP (2012) Practical Bayesian optimization of machine
31 learning algorithms. In *Advances in Neural Information Processing Systems (NIPS)*,
32 Lake Tahoe, NV, USA, 3–6 December 2012, vol. 25, pp. 2951–2959.
- 33 Tapia G, Elwany AH and Sang H (2016) Prediction of porosity in metal-based additive
34 manufacturing using spatial Gaussian process models. *Additive Manufacturing* 12:
35 282–290.
- 36 Wada K (2016) LabelMe: image annotation tool. GitHub repository. Available at:
37 <https://github.com/wkentaro/labelme> (accessed June 2026).
- 38 Wang Z, Zimmer-Chevret S, Léonard F and Abba G (2021) Prediction of bead geometry
39 with consideration of interlayer temperature effect for CMT-based wire-arc additive
40 manufacturing. *Welding in the World* 65(12): 2255–2266.
- 41 Williams SW, Martina F, Addison AC, Ding J, Pardal G and Colegrove P (2016) Wire +
42 arc additive manufacturing. *Materials Science and Technology* 32(7): 641–647.
- 43 Xiao X, Waddell C, Hamilton C and Xiao H (2022) Quality prediction and control in wire
44 arc additive manufacturing via novel machine learning framework. *Micromachines*
45 13(1): 137.
- 46 Xiong J, Lei Y, Chen H and Zhang G (2020) Thermal process and microstructure evolution
47 of thin-walled parts fabricated by WAAM with flux-cored wires. *Journal of*
48 *Manufacturing Processes* 52: 431–442.
- 49 Xiong J, Zhang G, Hu J and Wu L (2014) Bead geometry prediction for robotic GMAW-
50 based rapid manufacturing through a neural network and a second-order regression
51 analysis. *Journal of Intelligent Manufacturing* 25(1): 157–163.

- 1 Zuiderveld K (1994) Contrast limited adaptive histogram equalization. In Heckbert PS (ed)
- 2 Graphics Gems IV. San Diego, CA: Academic Press, pp. 474–485.

ONLY FOR REVIEW