

Exploring Feature Extraction Strategies for In-Situ Fault Detection of a Metal Additive Manufacturing Process with Time-Series Detection Models

ALVIN CHEN, PETROS SPILIOPOULOS,
FOTIS KOPSAFTOPOULOS and SANDIPAN MISHRA

ABSTRACT

Metal additive manufacturing (AM), particularly through Laser Powder Bed Fusion (LPBF), enables the fabrication of complex geometries with applications in aerospace, automotive, and other high-performance industries. These components must often meet stringent requirements for durability, weight, and cost. However, the AM process remains susceptible to internal defects such as overmelting and spatter, which originate in the melt pool and adversely affect part quality. Real-time monitoring of the melt pool using imaging sensors has emerged as a critical approach for detecting such defects, though the high dimensionality of image data necessitates data compression to enable efficient, real-time anomaly detection. This work investigates and evaluates various image compression strategies for their effectiveness in supporting unsupervised, real-time anomaly detection within the LPBF process. Building upon a predictive model trained on a healthy baseline, the detection algorithm identifies faults by assessing deviations in incoming signals. The study reveals that certain compression methods may obscure these deviations, limiting detection sensitivity. Therefore, an appropriate compression technique is essential to maintain detection performance while reducing computational complexity. Both physically intuitive features such as melt pool dimensions and machine learning-based feature extraction methods are examined. The efficacy of each approach is evaluated across several defect types commonly encountered in metal AM builds. These results assess the anomaly distinguishability of each method, offering guidance for the development of robust in-situ fault detection for metal AM.

Alvin Chen, PhD Student, Email: chena17@rpi.edu, Intelligent Structural Systems Lab (ISSL), Department of Mechanical, Aerospace and Nuclear Engineering, Rensselaer Polytechnic Institute, Troy, NY, USA

Petros Spiliopoulos, PhD Student, Email: spilip@rpi.edu, Intelligent Structural Systems Lab (ISSL), Department of Mechanical, Aerospace and Nuclear Engineering, Rensselaer Polytechnic Institute, Troy, NY, USA

Fotis Kopsaftopoulos, Assistant Professor. Intelligent Structural Systems Lab (ISSL), Department of Mechanical, Aerospace and Nuclear Engineering, Rensselaer Polytechnic Institute, Troy, NY, USA

Sandipan Mishra, Professor. Intelligent Systems, Automation and Control (ISAAc), Department of Mechanical, Aerospace and Nuclear Engineering, Rensselaer Polytechnic Institute, Troy, NY, USA

INTRODUCTION

Metal 3D printing is a form of Additive Manufacturing (AM), a manufacturing technique where 3D components are created by adding material to create the final geometry. In contrast, traditional methods of manufacturing are based on removing extra material from an initial part. Consequently, AM significantly reduces material waste, making it a cost-effective approach for producing complex geometries. A metal-based AM technology is Laser Powder Bed Fusion (LPBF), in which a layer of metal powder is evenly spread over a controlled surface and selectively melted by a laser beam according to a predefined pattern. Upon solidification, this forms one layer of the desired component, and the process repeats to build a complete 3D structure.

However, LPBF is susceptible to defects such as overmelting and spatter, which negatively impact the mechanical properties and structural integrity of the printed parts [1]. These anomalies manifest in the melt pool dynamics and are typically captured using Near-Infrared (NIR) imaging systems [2]. The melt pool image allows for direct visual inspection of the process and can be used in applications of automated defect detection algorithms for in-situ monitoring [3]. While it is common to apply machine learning models directly to these image signals, doing so typically requires large, labeled datasets for effective training, demanding significant time and resources [4].

To reduce dependence on large datasets and avoid labor-intensive labeling, simpler detection algorithms can be employed [5]. These algorithms benefit from dimensionality reduction techniques that decrease computational complexity without compromising detection performance. Accordingly, the design and implementation of effective image compression strategies for melt pool data are essential. By emphasizing process-relevant features during compression, such strategies can enhance the speed and robustness of real-time defect detection in LPBF systems.

The objective of this study is to evaluate and compare various image compression techniques for their suitability in real-time, unsupervised anomaly detection during metal AM. This investigation extends prior work that demonstrated the feasibility of fault detection using sparse training data [6, 7], with the goal of improving detection performance and scope. The proposed detection framework operates by comparing observed melt pool signals to predictive models trained exclusively on nominal (defect-free) data, using statistical hypothesis testing. However, in some cases, certain anomalies, when compressed, may exhibit similar variance as healthy data, thereby reducing detection reliability. As such, the choice of compression strategy is critical to enhancing both the sensitivity and responsiveness of the anomaly detection process. In this work, we analyze melt pool features such as melt pool dimensions that are relevant to process outcomes and compare these domain-specific features to machine-learned representations obtained through data-driven compression methods.

EXPERIMENTAL SETUP

Melt pool images are acquired in-situ using a coaxially mounted NIR camera, which captures 64×64 pixel grayscale images (0-255) at a sampling rate of 2 kHz. The sensor configuration, illustrated in Figure 1, follows the setup described in prior work [6]. Due to the coaxial arrangement, the camera captures melt pool reflections via the scanning

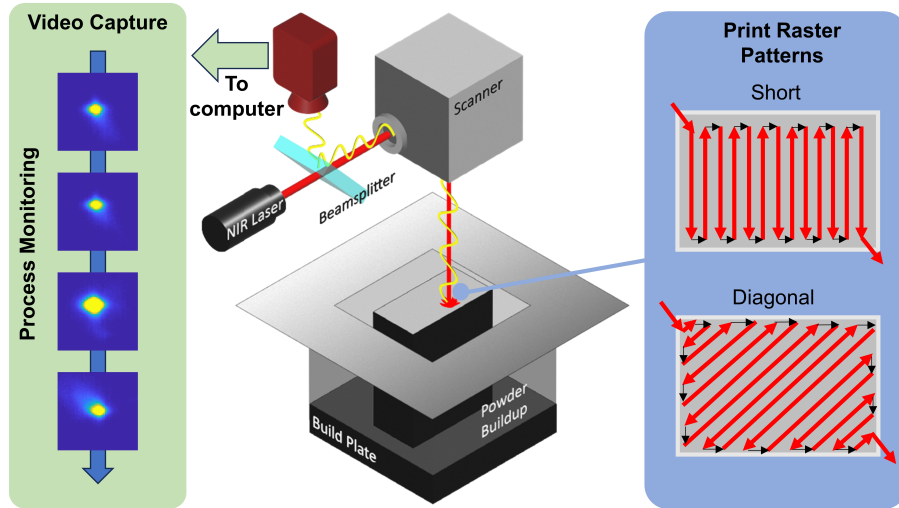


Figure 1. Sensor configuration (center), raster patterns (right), and sample melt pool images (left). A coaxial near-infrared (NIR) camera is used to monitor the melt pool in situ during the printing process. Two raster orientations are employed to fabricate a 2×4 cm rectangular prism. The raster patterns consist of alternating scan lines, creating "turnaround" points along the borders of each layer. These regions experience elevated energy deposition, resulting in periodic increases in melt pool size that correlate with the raster interval.

system, enabling consistent imaging regardless of the scanner's position across the plane.

The dataset comprises 150 total printed layers across three parts, each forming a 2×4 cm rectangular prism. Each layer is sintered in one of two raster paths, either a 90° or 45° orientation, corresponding to short and diagonal scan directions, respectively, as shown in Figure 1. The laser takes approximately 5.5 seconds to travel across the entire layer, yielding about 11000 images per layer. Anomalies are naturally and artificially occurring, with defects such as spatter or overmelting highlighted in the dataset.

Fluctuations in the melt pool length arise from the geometric periodicity. At locations of dense rasterization (such as at spots where the laser path turns back on itself), the local energy density will increase, consequently increasing the melt pool size. Additionally, the loose, unmelted powder will limit heat transfer along the edges of the solidified part. These geometric effects are deterministic, depending on the chosen raster pattern and influence the melt pool dynamics. Under healthy operation, the melt pool size will increase when at a local heating point, affecting the monitored signal.

The compression strategies presented in this work focus on interpretable characteristics of the melt pool, such as the melt pool area. Consequently, larger melt pools during healthy operation may result in false alarms when a compressed signal is examined, due to their similarity to an anomalous melt pool. To address that, an autoregressive (AR) model is trained on one layer (11000 samples) to compensate for the expected hotspot locations. The resulting error (model residuals) is expected to be normally distributed.

DATA COMPRESSION TECHNIQUES

In the following section the details for each compression strategy will be described. These strategies include both physical features of the melt pool and a data-driven machine learning technique. Each of the compression strategies is presented in Figure 2.

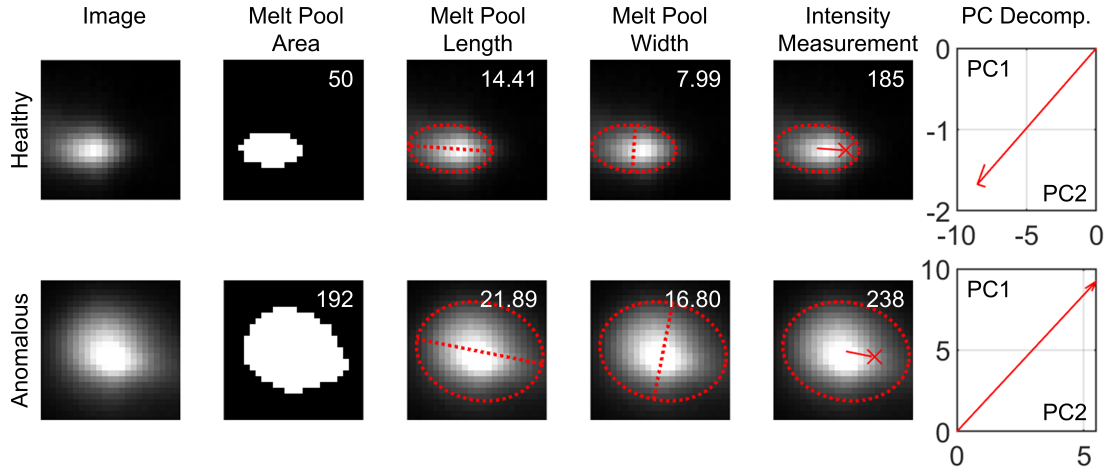


Figure 2. Sample healthy (top) and anomalous (bottom) melt pool images with visual representations of each compression strategy. Images are cropped to a 24×24 region of interest. Compression methods include melt pool area (pixel count above intensity threshold), major/minor axes of an ellipse fit, axial intensity along the major axis at a fixed offset from the centroid, and the first two principal components (PCs) projected onto a healthy basis. Compression values are shown in the top-right of each sub-image.

Melt Pool Area: The melt pool area is an interpretable metric that is proportional to the weld power and temperature [8]. This metric has been used in prior work [9] and serves as a baseline comparison. This is calculated by counting the number of pixels in each image that exceeds temperature threshold α (pixel intensity).

Melt Pool Length and Width: An ellipse is circumscribed around the (thresholded) melt pool image to estimate the properties of the melt pool. The major and minor axes of the ellipse correspond to the length and width of the melt pool. This ellipse is estimated with the “regionprops” function of the Matlab image processing toolbox [10] and a user-defined threshold α to create a binary melt pool shape to circumscribe the ellipse around. This allows for additional properties of the melt pool and captures the behavior of misshapen melt pools that may exhibit a similar area to healthy melt pools.

Liquidus Boundary: There are concentric *isothermal* rings around the centroid of the the melt pool [11]. The intensity of the pixel at a band near the healthy liquidus temperature region (liquid metal region boundary) from the melt pool center can then be used as a detection feature. Some distance d is chosen from the centroid of the melt pool along the major axis of a circumscribed ellipse. This intensity-based measurement complements geometric features by capturing thermal characteristics of the melt pool and can help differentiate anomalies with similar melt pool areas to healthy cases.

Principal Components: By applying principal component (PC) decomposition to the melt pool images, each image is transformed to a coordinate system such that the features with the greatest variance are ordered. By choosing the n th PC, each image can be compressed to a single value that carries the greatest variance of the data.

MODELING AND STATISTICAL TESTING

The following section provides an overview of the model used to compensate for the geometric effects on the melt pool dynamics and the tests used to evaluate the efficacy of each compression strategy. Refer to the authors’ prior work for additional detail.

AR Model

An autoregressive (AR) model [12] is trained on compressed, healthy melt pool data (one layer, 11,000 samples) to perform one-step predictions. Under nominal conditions, the model residuals are assumed to be independent, normally distributed errors. The AR model parameters are estimated via least squares, and model order is selected using the Bayesian Information Criterion (BIC) to best capture nominal signal behavior while filtering out geometry-induced variations.

Kullback-Leibler Divergence

To evaluate each compression method, the difference between compressed signal distributions is measured using Kullback-Leibler (KL) divergence [13], a standard metric for comparing probability distributions. For a test distribution $P(x)$ and baseline $Q(x)$, KL divergence is defined as:

$$D_{\text{KL}}(P \parallel Q) = \sum_{x \in \mathcal{X}} P(x) \log \frac{P(x)}{Q(x)} \quad (1)$$

As the objective of this work is a compression strategy that *best* characterizes the anomalous behavior of the system, the best performing strategy would yield the greatest difference between the KL divergence for a test healthy (D_{KL,H_0}) and a test anomalous (D_{KL,H_1}) case. Intuitively, a large KL divergence for an anomalous case implies that the distribution for signals differ from the (healthy) baseline. The choice of a strategy that maximizes the difference of KL divergences between a healthy and anomalous case highlights the effects of a fault and improves the detection capabilities of that signal.

Statistical Hypothesis Test

Detection is based on the following binary hypothesis test for the AR residuals $e[t]$:

$$\begin{aligned} H_0 : e[t] &\sim \text{iid} \mathcal{N}(0, \sigma_0^2) && \text{Null Hypothesis - Healthy Signal} \\ H_1 : e[t] &\not\sim \text{iid} \mathcal{N}(0, \sigma_0^2) && \text{Alternate Hypothesis - Anomalous Signal} \end{aligned} \quad (2)$$

For healthy performance, the AR model is expected to predict the signal well, accounting for the geometry-induced properties of the signal. When a fault occurs, this disturbs the prediction, leading to a large error (residual) and changes the statistical distribution of the residuals.

To assess the performance of each compression method, two statistical testing methods based off the above hypothesis methods were used. These tests are detailed in prior work [6, 9] and are briefly described.

The Residual variance test assumes the AR residuals have normally distributed residuals with a baseline variance σ_0^2 . A hypothesis test compares the variance of online residuals σ_u^2 to this baseline variance using an F -statistic (the quotient of residuals) that follows an f -distribution. If the ratio exceeds a threshold set by the false alarm rate α (type-I error), the signal is classified as anomalous.

The Sequential Probability Ratio Test (SPRT) [14] detects anomalies by evaluating whether the variance of model residuals deviates from a nominal baseline, using a cu-

mulative test statistic $\Lambda[t]$. This statistic is compared against upper and lower decision thresholds, determined by user-defined false alarm rate α and missed detection rate β . If the statistic exceeds these bounds, the signal is classified as either nominal or anomalous.

COMPRESSION STRATEGY EVALUATION

An AR model was trained for each compression strategy, with model order selected by minimizing the Bayesian Information Criterion (BIC), resulting in $n_a = 19$ across all cases. Each model was fit using a healthy baseline layer of 11,000 samples.

The resulting compressed signals are shown in Figure 3, which compares short time-series segments from a healthy case and an anomalous case with naturally occurring spatter faults. Due to the nature of these faults, it is difficult to differentiate the faults from the healthy case through observing the time series directly.

Detection parameter tuning is essential for reliable fault identification. For both the Sequential Probability Ratio Test (SPRT) and Residual Variance (RV) methods, thresholds were selected to balance sensitivity and false alarm (FA) rate as depicted in Figure 4. The FA rate is determined by checking detection results for a healthy segment (bottom, e-h), and the detection sensitivity is assessed by observing the anomaly detection rate for a test anomaly (top, a-d). For RV, the alarm rate is defined as the percentage of samples flagged as anomalous, while for SPRT, it is the proportion of anomalous decisions over all decisions (healthy and anomalous). Type-I error rate α was chosen as 5×10^{-2} , along the “elbow” of the alarm rate in the healthy dataset (Figure 4 e-f). Type-II error (missed fault β) is held constant for each plot at 0.01.

In Figure 4c (RV), a sliding-window size of $M = 50$ was selected, as alarm rates stabilized beyond this value. For the SPRT, the upper threshold parameter q was set to 1.25 to avoid both missed faults ($q \leq 1.1$, Figure 4d) and high FAs ($q \geq 1.3$, Figure 4h) in the baseline C80 compression strategy. These settings ensure both methods maintain high detection performance with minimal false alarms.

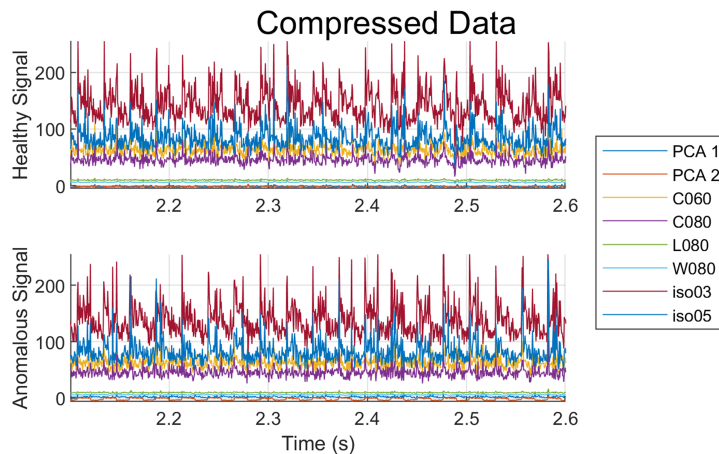


Figure 3. Compressed melt pool time series comparing a healthy signal (top) to an anomalous one (bottom) with naturally occurring spatter anomalies. The impact of anomalies is difficult to discern in the time series. Shown compression strategies include PCA components (1, 2), melt pool area at two thresholds (C60, C80), ellipse axes (L80, W80), and pixel intensities 3 and 5 pixels from the center along the major axis (iso03, iso05).

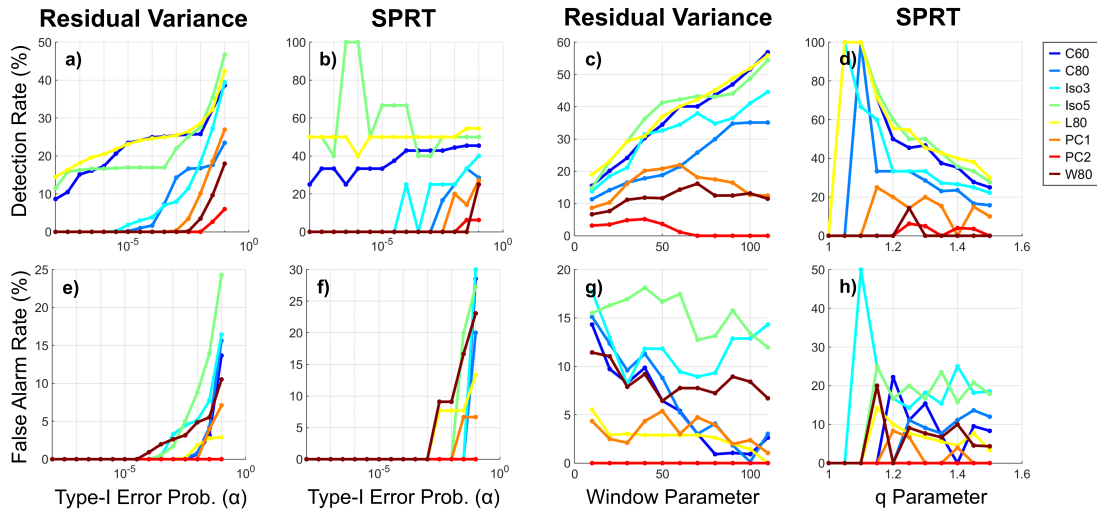


Figure 4. Parameter search plots for two test cases Spatter (top, a-d) and Healthy (bottom, e-h) for two different statistical detection methods and 8 compression strategies. For each detection method, two parameters are modified, the type-I error (false alarm rate, a, b, e, f), and the variance estimation window (residual variance, c,d) or upper threshold factor q (SPRT, d,h). For SPRT, the lower threshold β (type-II error) is constant at 0.01. The final selection of each parameter is chosen by minimizing false alarm rate and maintaining test anomaly detections.




Figure 5. Kullback-Leibler (KL) divergence for various compression strategies across three healthy and three anomalous layers. The first principal component (PCA 1) shows the greatest separation, while the pixel intensity 5 pixels from the melt pool center (iso05) performs the worst with similar KL divergence values across all layers.

Using the outlined selection criteria, optimal performance for melt pool area and ellipse-based strategies was achieved with an intensity threshold of approximately 80 (on a 0–255 grayscale), which is believed to approximate the melt pool’s liquidus region despite the NIR camera’s lack of direct temperature measurement. The “iso” metric, based on pixel intensity at a fixed distance from the melt pool center, also performed best near this threshold. For PCA, the first two components were chosen as they capture most of the image variance. While not directly interpretable, they still highlight anomalies.

Figure 5 shows the KL divergence for each compression strategy following one-step AR prediction. The first principal component (PCA 1) yields the highest divergence between anomalous and healthy layers, outperforming the baseline melt pool area feature (C80). This indicates greater distributional separation, which may help reduce FAs.

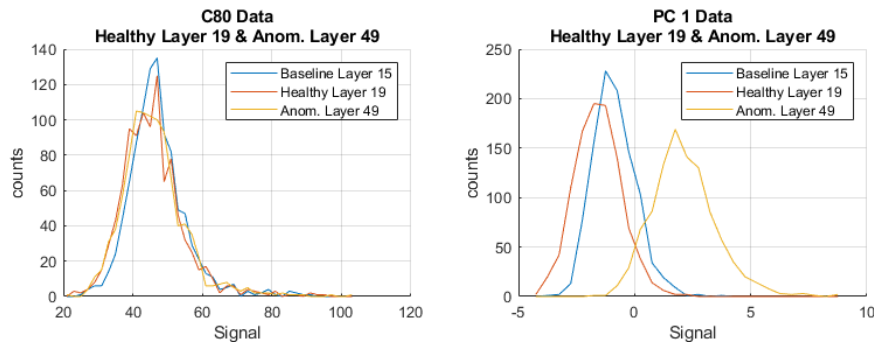


Figure 6. Compression signal distributions for baseline, test healthy, and test anomalous layers using the best (left, melt pool size) and worst (right, PCA 1) compression strategies. Melt pool length shows clear separation of distributions, while PCA 1 fails to distinguish anomalies.

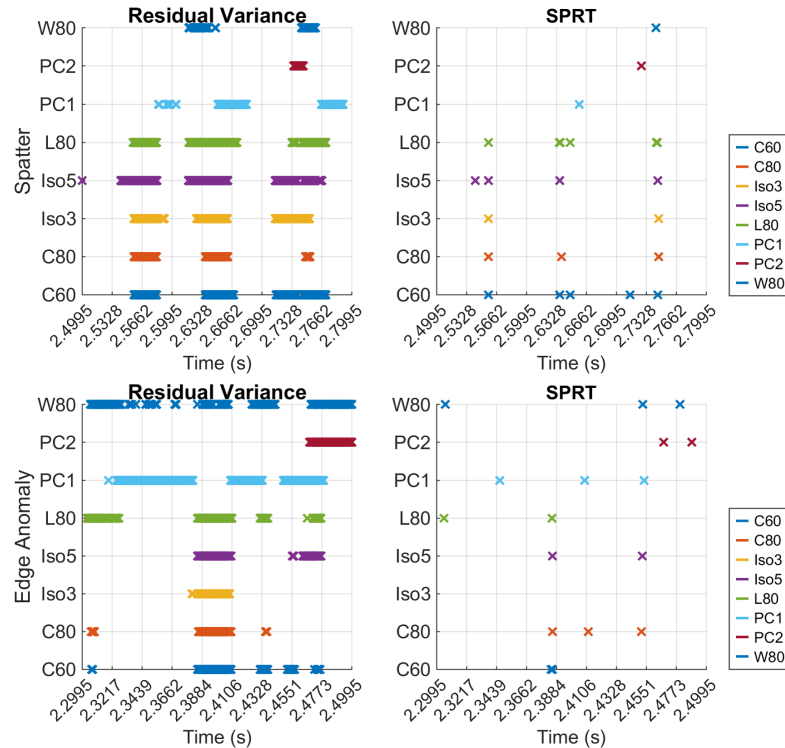


Figure 7. Detection results (marked with an x at detected location) for two statistical tests, a residual variance-based detection (left, parameters $\alpha = 0.05, M = 50$) and the sequential probability ratio test (right, parameters $\alpha = 0.05, \beta = 0.01, q = 1.25$). Depicted are detection results for several compression strategies for two test anomalous layers: a spatter dataset (top) and an overmelting anomaly (bottom).

However, as shown in Figure 6, the healthy and baseline distributions for the PCA signal (right) do not align well. In contrast, the melt pool area feature (C80, left) shows strong agreement between the healthy and baseline distributions. Although both methods successfully separate anomalous data, PCA 1 is less ideal due to the distribution mismatch, which can increase the false alarm rate and reduce detection reliability. The impact of this difference is further demonstrated in the detection results.

Detection results for both anomalous and healthy cases are shown in Figure 7, with a spatter dataset (top) and an overmelting anomaly (bottom). In the spatter case, all three

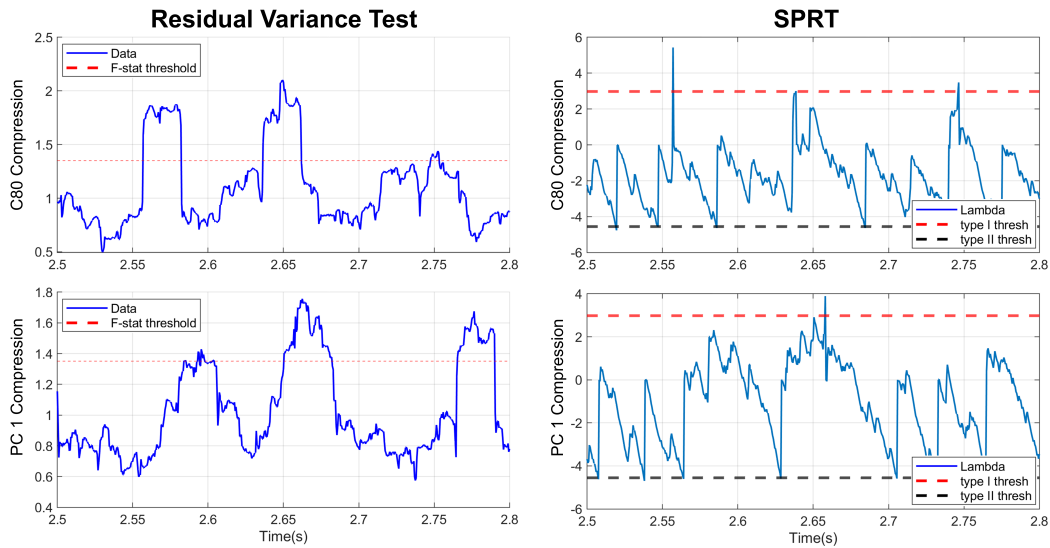


Figure 8. Sample detections for spatter dataset. Residual variance (left) and sequential probability ratio test (right) for two compression strategies: melt pool area (C80, top) and first PC (PC 1, bottom). Anomalies are declared when the detection metric (blue) exceeds the statistical threshold (red dashed line). Three spatter events are present in the segment, which are detected by all tests except for the SPRT on PC 1.

spatter events are detected by both algorithms when using melt pool dimension-based compression strategies (area and length), except for width. In contrast, the first principal component (PC 1) detects only partial anomalies with the RV test and misses events entirely with SPRT. The second PC fails to detect any spatter instances in either method. This is likely due to a mismatch in mean between the healthy and baseline PC signals, which the AR model cannot compensate for, resulting in diminished sensitivity.

Figure 8 illustrates this behavior in detail for C80 and PC 1. While all spatter events are detected using C80 in both tests, PC 1 misses the third event at 2.75s with RV and both the first and third events with SPRT. For the overmelting anomaly at 2.38s, neither PC captures the fault, whereas the melt pool dimension features successfully identify it in both detection tests.

The distribution of the compressed data and the results for the statistical tests indicate that the PCs highlight a feature where the characteristics of a spatter and overmelting anomaly do not appear clearly. As such, the melt pool dimension compression strategies (such as the melt pool length and size) are more ideal with these detection algorithms.

CONCLUDING REMARKS

This work presented a comparative evaluation of various melt pool image compression strategies for real-time, unsupervised fault detection for a metal AM LPBF process. Towards improving unsupervised statistical testing based on time-series AR modeling, each compression strategy was evaluated for its ability to isolate anomalies from a nominal process. Interpretable compression strategies such as melt pool length or size are better for detecting anomalies as principal component-based features are unable to capture an anomaly with the presented tests. This was evaluated by examining the distributions of each strategy and evaluating the corresponding detection results.

These results highlight the importance of selecting compression strategies that not only reduce data dimensionality but also preserve physical characteristics that are affected by process faults. Future work will apply the compression strategies developed in this work to reduce false alarms and improve the detection capabilities of the statistical detection algorithms presented.

ACKNOWLEDGMENT

This work was supported in part by the National Science Foundation under CMMI Award #2222250.

REFERENCES

1. Cai, Y., J. Xiong, H. Chen, and G. Zhang. 2023. "A review of in-situ monitoring and process control system in metal-based laser additive manufacturing," *Journal of Manufacturing Systems*, 70(August), doi:10.1016/j.jmsy.2023.07.018.
2. Li, X., H. Li, X. Chen, S. Shen, G. Zhang, H. Wei, Y. Hu, Z. Li, and L. Dai. 2025. "In-situ multi-eye monitoring of melt pool temperature field in laser additive manufacturing by light field camera," *Additive Manufacturing*, 102(July 2024), doi:10.1016/j.addma.2025.104747.
3. Shevchik, S., R. Wrobel, L. Quang T, V. Pandiyan, P. Hoffmann, C. Leinenbach, and K. Wasmer. 2024. "Unsupervised quality monitoring of metal additive manufacturing using Bayesian adaptive resonance," *Heliyon*, 10(12), doi:10.1016/j.heliyon.2024.e32656.
4. Mahmood, M. A., F. G. Alabtah, Y. Al-Hamidi, and M. Khraisheh. 2023. "On the laser additive manufacturing of high-entropy alloys: A critical assessment of in-situ monitoring techniques and their suitability," *Materials and Design*, 226, doi:10.1016/j.matdes.2023.111658.
5. Fu, Y., A. R. J. Downey, L. Yuan, T. Zhang, and A. Pratt. 2022. "Machine learning algorithms for defect detection in metal laser-based additive manufacturing : A review," *Journal of Manufacturing Processes*, 75(December 2021), doi:10.1016/j.jmapro.2021.12.061.
6. Chen, A., F. Kopsaftopoulos, and S. Mishra. 2023. "A data-sparse approach to in-situ fault detection and identification for metal additive manufacturing," *2023 International Workshop for Structural Health Monitoring*, doi:10.12783/shm2023/36939.
7. Chen, A., F. Kopsaftopoulos, and S. Mishra. 2025. "Online Fault Detection in Metal Additive Manufacturing Using Statistical Time Series Methods," *Structural Health Monitoring*, Under review.
8. Ren, Y. and Q. Wang. 2022. "Gaussian-process based modeling and optimal control of melt-pool geometry in laser powder bed fusion," *Journal of Intelligent Manufacturing*, 33(8):2239–2256, ISSN 15728145, doi:10.1007/s10845-021-01781-4.
9. Chen, A., F. Kopsaftopoulos, and S. Mishra. Presented Nov 2024. "Online Fault Detection for Metal Additive Manufacturing with Data-Driven Time Series Models," *2024 Dynamic Data Driven Applications Systems*.
10. MathWorks. 2024, "Image processing toolbox," .
11. DebRoy, T., H. L. Wei, J. S. Zuback, T. Mukherjee, J. W. Elmer, J. O. Milewski, A. M. Beese, A. Wilson-Heid, A. De, and W. Zhang. 2018. "Additive manufacturing of metallic components – Process, structure and properties," *Progress in Materials Science*, 92:112–224, ISSN 00796425, doi:10.1016/j.pmatsci.2017.10.001.
12. Ljung, L. 1999. *System Identification: Theory for the User*, Prentice Hall information and system sciences series, Prentice Hall PTR, ISBN 9780136566953.
13. Kullback, S. 1978. *Information Theory and Statistics*, Peter Smith, ISBN 9780844656250.
14. Wald, A. 1947. *Sequential Analysis*, John Wiley and Sons, 1st edn.