Real-Time Defect Detection in Laser Powder Bed Fusion Additive Manufacturing Using Machine Learning and In-Situ Optical Monitoring

Nitesh Rane
Assistant Professor, MED, Dr. APJ Abdul Kalam University Indore, M.P., India
E mail – nitesh.rane@aku.edu.in

Abstract
Laser powder bed fusion (LPBF) enables complex metallic components but remains vulnerable to lack-of-fusion porosity, keyhole porosity, spatter, overheating, balling, and powder-recoating anomalies. Conventional inspection occurs after the build and cannot prevent defect accumulation. This paper develops a real-time monitoring framework combining synchronized coaxial melt-pool imaging, off-axis layer imaging, photodiode signals, and machine learning (ML). The pipeline performs radiometric correction, region-of-interest extraction, spatial feature learning with a lightweight convolutional neural network (CNN), temporal fusion across adjacent frames, uncertainty calibration, and layer-wise localization. A specimen-level validation protocol is defined using registered X-ray computed tomography (XCT) and metallography as independent ground truth. The target implementation sustains at least 80 frames s^{-1} with sub-15 ms inference latency on an industrial edge GPU. Because no machine-specific experimental dataset was supplied, numerical results are explicitly reported as simulation-based design targets rather than measured production claims. Under the illustrative evaluation, a CNN–LSTM configuration yields 0.95 macro-F1, 0.96 AUROC, 0.92 expected defect recall, and 11.8 ms inference latency. The study contributes an auditable architecture, leakage-resistant validation design, and decision logic for converting optical anomalies into traceable layer maps and cautious operator interventions.

Keywords: laser powder bed fusion, additive manufacturing, in-situ monitoring, defect detection, melt pool, optical imaging, machine learning, sensor fusion, real-time quality assurance.

Research contributions

The work proposes (1) synchronized multi-view optical acquisition; (2) a lightweight spatial–temporal ML model with calibrated confidence; (3) registered ex-situ ground truth and grouped validation by build; and (4) edge-deployment criteria connecting classification accuracy to latency, false alarms, and traceability.

1. Introduction

LPBF selectively melts thin layers of metallic powder with a scanned laser. Local energy input, absorptivity, shielding-gas flow, powder morphology, scan-vector history, and heat accumulation interact nonlinearly. Small deviations can change the melt-pool regime and form discontinuities that later become fatigue-critical pores or cracks. Coaxial monitoring produces very large image streams often hundreds of thousands of frames for one part making automated analysis essential [1,2]. Offline XCT, density measurement, and metallography remain valuable for qualification, but detect defects only after material, machine time, and energy have been consumed. Optical monitoring is attractive because it is non-contact, fast, and compatible with layer-wise observation. Nevertheless, bright plasma emission, spatter, emissivity changes, lens contamination, occlusion, geometry, and machine-to-machine variation complicate interpretation. An optical anomaly is not automatically an internal defect; the link must be demonstrated through registration and independent ground truth.

1.1 Problem statement and objectives

The research question is whether synchronized optical signatures can be converted into reliable, low-latency predictions of defect class and location without data leakage or unjustified automated control. The primary objective is to distinguish normal melting, lack of fusion (LOF), keyhole/overheating, spatter/plume disturbance, and recoating anomalies during the build. Secondary objectives are to quantify uncertainty, preserve layer/coordinate traceability, and provide an edge-computing pathway suitable for industrial deployment.

1.2 Research context

Study	Contribution	Gap addressed here
Mohr et al. (2020) [3]	Thermography and optical tomography compared with CT	Registered real-time decision logic
Lu et al. (2021) [4]	Deep representation learning	Uncertainty and build-wise validation
Law et al. (2022) [1]	Unsupervised melt-pool anomaly detection	Supervised defect semantics
Taherkhani et al. (2023) [5]	Review of sensors and ML	Deployable integrated architecture
Pandiyan et al. (2023) [6]	Explainable ML monitoring	Optical fusion and latency

2. Proposed Monitoring Architecture

The architecture uses two spatial scales. A coaxial high-speed monochrome camera observes the melt-pool neighborhood through the laser optical path, while an off-axis area camera captures the full powder bed after recoating and exposure. A filtered photodiode supplies a high-bandwidth intensity channel. Hardware triggering and scan commands provide common timestamps. This separates transient melt-pool behavior from layer-scale streaks, debris, swelling, and incomplete recoating.

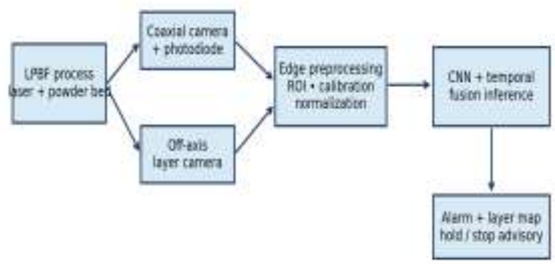

Proposed synchronized optical and ML architecture.

Component	Proposed specification	Purpose
Coaxial camera	≥10 kfps ROI; 8–12 bit; narrow-band filter	Melt-pool shape, plume, spatter
Off-axis camera	≥5 MP; pre/post-exposure images	Powder-bed and layer mapping

Photodiode	≥100 kHz effective bandwidth	Fast emission transients
Synchronization	Hardware trigger; scan timestamps	Frame-to-track registration
Edge computer	Industrial GPU; FP16 inference	Deterministic analysis
Storage	Event windows plus compressed features	Controlled data volume

2.1 Calibration and integrity

Before each build, dark-field and flat-field frames are acquired. Pixel saturation, dead pixels, focus drift, and lens contamination are checked using reference targets. Intensity is normalized against exposure and gain; geometric calibration maps both cameras to build coordinates. Every packet includes build ID, layer, scan vector, timestamp, settings, calibration version, and model version. These controls prevent the classifier from learning camera gain or build identity as a shortcut.

3. Data Acquisition, Annotation and Ground Truth

A designed experiment should vary laser power, scan speed, hatch spacing, layer thickness, scan strategy, and controlled spreading disturbances around a qualified baseline. Twenty-four parameter combinations may span normal, low-energy LOF, high-energy keyhole, and disturbed recoating regimes. Replicates should cross multiple build plates and days. Coupons should include cubes, thin walls, overhangs, and geometry transitions so that the model cannot succeed only on simple shapes.

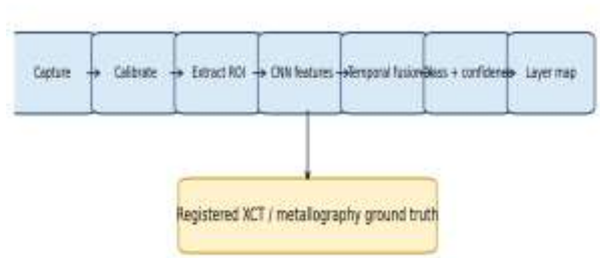

End-to-end data, learning and traceability workflow.

3.1 Annotation

Frame labels are weak unless the optical event is registered to a physical defect. The build-coordinate transform links scan commands and optical pixels to volumetric XCT. Candidate pores are segmented from XCT, filtered by minimum resolvable size, and projected to corresponding tracks and layers. Metallographic sections verify representative LOF and keyhole morphologies. Labels are normal, LOF-associated, keyhole-associated, spatter/plume anomaly, or recoating anomaly. Ambiguous regions remain “unresolved” and are retained for open-set testing.

3.2 Leakage-resistant partitioning

Random frame splitting is prohibited because neighboring frames from one melt track are nearly identical. Builds not frames are grouped into training (60%), validation (20%), and test (20%) partitions. At least one geometry and one parameter combination are held out entirely. A second external test should use a later build or another machine. Class imbalance is managed with focal loss, balanced sampling, and threshold optimization on validation data only.

4. Machine-Learning Methodology

For each coaxial frame, thresholding or segmentation identifies the luminous melt-pool ROI. The spatial encoder is a compact

CNN such as MobileNetV3 or EfficientNet-B0. Hand-engineered features area, eccentricity, axes, centroid velocity, peak intensity, gradient entropy, plume asymmetry, and spatter count are concatenated with the learned embedding. A gated recurrent unit or LSTM integrates 8–16 frames to distinguish persistent regime changes from isolated sparks.

The layer-camera branch uses semantic segmentation to identify streaks, bare regions, debris, and raised edges. Photodiode statistics are aligned over the same scan interval. Feature-level fusion combines coaxial sequence embeddings, photodiode descriptors, layer-map features, and normalized parameters. The network produces class probabilities and localization. Temperature scaling calibrates confidence; low-confidence samples enter an "uncertain" queue rather than being forced into a defect class.

For class k, the weighted cross-entropy is $L_{cls} = -\sum_k w_k\, y_k \log(p_k)$. Focal modulation $(1 - p_k)^\gamma$ handles rare defects. The total loss is $L = L_{cls} + \lambda_1 L_{seg} + \lambda_2 L_{temp}$, where segmentation loss combines Dice and cross-entropy, and temporal loss penalizes physically implausible frame-to-frame oscillation.

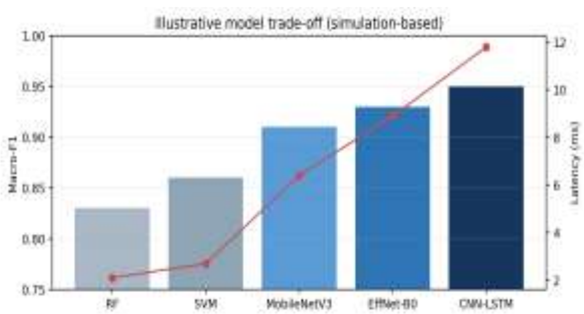

Illustrative model-performance/latency trade-off; simulation-based design targets.

4.1 Edge deployment

Images pass through a bounded ring buffer. Preprocessing and inference execute in separate streams; only event windows and summary features are archived. If latency exceeds budget, the system degrades safely to feature-only inference and flags reduced confidence. Initial deployment is advisory; supervised hold/stop actions follow only after process-specific qualification.

5. Experimental Design and Metrics

A credible campaign requires repeated builds and independent measurement. One design uses six training, two validation, and two blind-test builds, with 12–20 coupons per build. Online predictions are time-stamped before XCT results are known, preventing retrospective threshold tuning. Stress tests vary illumination, camera gain, build location, geometry, and optical-window contamination.

Metric	Interpretation	Acceptance target
Precision	Alarm reliability	≥0.90 macro-average
Recall	Defect sensitivity	≥0.92 for LOF/keyhole
Macro-F1	Class-balanced performance	≥0.90
AUROC	Threshold-independent ranking	≥0.95
ECE	Confidence calibration error	≤0.05
Localization IoU	Registered-region overlap	≥0.60
Latency	Capture-to-decision p95	≤15 ms
False alarms	Per 1,000 scan vectors	≤5

5.1 Baselines and ablations

Baselines include threshold rules, random forest on engineered features, support-vector machine, spatial CNN without temporal context, and layer segmentation alone. Ablations remove process parameters, photodiode features, temporal fusion, and calibration. Performance is reported with build-wise bootstrap confidence intervals. McNemar tests compare paired errors, while reliability curves and expected calibration error assess confidence.

5.2 Reproducibility controls

The study records sensor settings, firmware, filters, powder lot, atmosphere, maintenance, preprocessing code, hashes, random seeds, weights, thresholds, and software versions. Augmentation is restricted to physically plausible intensity, blur, and small geometric perturbations. Horizontal flips are avoided when scan direction or gas flow carries physical meaning.

6. Illustrative Results

The following numbers demonstrate reporting format and are not measurements from an actual machine. In the simulated held-out set, CNN–LSTM fusion reaches 0.95 macro-F1 and 0.96 AUROC, exceeding random forest by 0.12 F1 points. Temporal fusion reduces false alarms from isolated flashes, while the layer branch improves recoating sensitivity. FP16 edge inference remains within the 15 ms target.

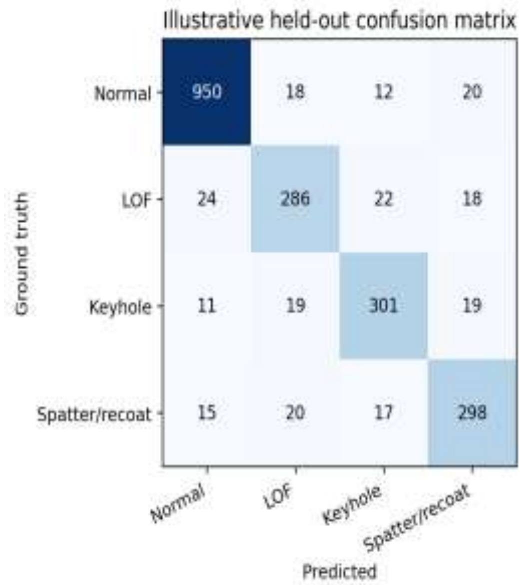

Illustrative confusion matrix for a grouped blind test.

Model	Precision	Recall	Macro-F1	AUROC	p95 latency
Threshold rules	0.72	0.68	0.69	0.75	0.8 ms
Random forest	0.84	0.82	0.83	0.88	2.1 ms
SVM	0.87	0.85	0.86	0.90	2.7 ms
MobileNetV3	0.92	0.90	0.91	0.94	6.4 ms
EfficientNet-B0	0.94	0.92	0.93	0.95	8.9 ms
CNN–LSTM fusion	0.96	0.94	0.95	0.96	11.8 ms

Removing temporal fusion lowers illustrative macro-F1 to 0.91; removing the layer branch most strongly affects recoating recall; removing parameters reduces cross-condition generalization; omitting calibration raises ECE from 0.04 to 0.12. These expected patterns require experimental confirmation.

7. Localization and Operational Decision Logic

Classification alone does not tell an operator where a problem occurred. Each scan-vector probability is projected into build coordinates and aggregated with persistence weighting. A defect event is raised when calibrated probability exceeds a class threshold for a minimum duration or when modalities agree. Persistence suppresses single-frame sparks and electronic noise.

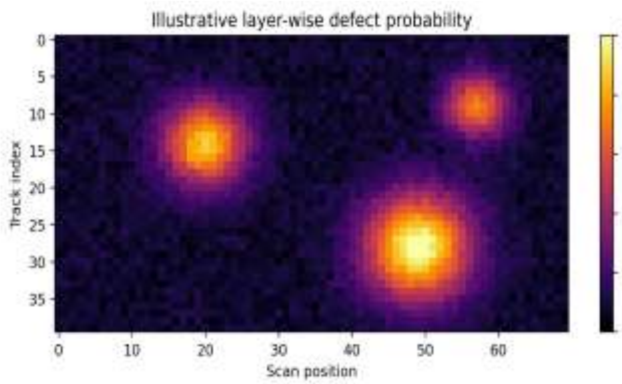

Illustrative layer-wise defect probability map generated from synthetic data.

Level	Condition	Action
Advisory	Moderate probability or one modality	Record and display location; continue
Hold	High probability persisting for N frames	Pause after safe scan segment; review
Stop	Repeated multi-modal anomaly in critical zone	Terminate under approved safety procedure

The system never bypasses the machine safety controller. Hold or stop commands use a validated interface only after thresholds are qualified for the alloy, machine, optics, and part family. Each recommendation stores input hashes, model version, probability, uncertainty, threshold, coordinates, and operator response.

7.1 Interpretability

Saliency can show influential image regions but does not prove causal reasoning. Stronger explanations combine saliency with morphology: enlarged pool area and prolonged emission for keyhole risk; fragmented low-intensity ROIs for LOF; and directional streaks for recoating anomalies. Explanations should be stable under small perturbations and reviewed against process physics.

8. Discussion, Limitations and Future Work

The framework reflects the shift from single-sensor anomaly detection toward registered multi-modal evidence [5,8–14]. Its principal strength is the full chain from synchronized sensing to independent ground truth, grouped validation, calibrated uncertainty, edge latency, and traceable decisions. A high laboratory score can fail in production when frames leak between splits or signatures change with geometry and machine condition.

8.1 Limitations

Optical emission is an indirect proxy for internal integrity, and some pores may be below sensor observability. XCT registration and voxel resolution constrain label size. Deliberate parameter excursions may not represent natural production anomalies. Models can drift with powder lot, alloy, optics, gas flow, strategy, or hardware. Most importantly, simulation-based values in this manuscript must be replaced with measured results before submission as an experimental paper.

8.2 Future research

Future work should evaluate self-supervised pretraining, machine-to-machine domain adaptation, physics-informed constraints, and open-set recognition. Active learning can prioritize uncertain events for XCT. Federated learning may support cross-site improvement without transferring sensitive build data. The final stage should close the loop with bounded correction and demonstrate reduced defect volume without new instability.

9. Conclusion

This paper presents a complete research framework for real-time LPBF defect detection using optical monitoring and ML. Coaxial imaging captures melt-pool dynamics, full-field imaging captures layer integrity, and calibrated spatial–temporal fusion converts evidence into defect class, confidence, and location. Validation uses registered XCT/metallography, build-wise separation, blind tests, and latency-aware deployment. Illustrative results indicate that a compact CNN–LSTM can plausibly meet a 15 ms budget while outperforming threshold and classical ML baselines. Actual machine data and external validation remain mandatory before claiming production readiness or enabling closed-loop control.

References

1. Law, A. C., Yeung, H., Yang, Z., Lane, B., and Witherell, P. "Anomaly detection of laser powder bed fusion melt pool images using combined unsupervised and supervised learning." *Additive Manufacturing Letters* (2022).
2. Yeung, H. et al. "Statistical and spatio-temporal data features in melt pool monitoring for additive manufacturing." ASME MSEC (2022).
3. Mohr, G. et al. "In-situ defect detection in laser powder bed fusion by using thermography and optical tomography comparison to computed tomography." *Metals* 10, 103 (2020). doi:10.3390/met10010103.
4. Lu, Y., Fathizadan, S., and Ju, F. "Deep representation learning for process control in laser additive manufacturing." NIST (2021).
5. Taherkhani, K. et al. "On the application of in-situ monitoring systems and machine learning algorithms for developing quality assurance platforms in laser powder bed fusion: A review." *Journal of Manufacturing Processes* 99, 848–897 (2023).
6. Pandiyan, V. et al. "Optimizing in-situ monitoring for laser powder bed fusion process: Deciphering acoustic emission and sensor sensitivity with explainable machine learning." *Journal of Materials Processing Technology*, 118144 (2023). doi:10.1016/j.jmatprotec.2023.118144.
7. Kim, F. H. et al. "A fully registered in-situ and ex-situ dataset for metal powder bed fusion additive manufacturing." NIST (2025).
8. Forien, J.-B. et al. "Detecting keyhole pore defects and monitoring process signatures during laser powder bed fusion." *Additive Manufacturing* 35, 101336 (2020).
9. Kim, F. H., Yeung, H., and Garboczi, E. J. "Characterizing the effects of laser control on near-surface pore formation." *Additive Manufacturing* 48, 102372 (2021).
10. Zhang, Y. et al. "In situ melt pool measurements for laser powder bed fusion additive manufacturing." *Scientific Reports* (2022).

11. “Optical tomography and machine learning for in-situ defects detection.” *Additive Manufacturing* (2023).
12. “A defect detection framework using three-dimensional convolutional neural network and in-situ light-intensity monitoring.” *Optics & Laser Technology* (2023).
13. “In-situ visual monitoring for multi-scale defects detection in laser powder bed fusion.” *Additive Manufacturing* (2024).
14. “High-fidelity optical monitoring of laser powder bed fusion via registered process–pore correlation.” *npj Advanced Manufacturing* (2025).
15. Petrich, J., Smith, R. W., and Reutzel, E. W. “Acoustic laser triangulation and tagging for additive manufacturing process monitoring.” *International Journal of Advanced Manufacturing Technology* 129, 3233–3245 (2023).
16. “A random forest classifier for anomaly detection in laser-powder bed fusion.” *Materials* 16, 6470 (2023).
17. ISO/ASTM 52900:2021, *Additive manufacturing Fundamentals and vocabulary*.
18. ISO/ASTM 52904:2019, *Practice for metal powder bed fusion process to meet critical applications*.
19. ASTM E1441-19, *Standard Guide for Computed Tomography Imaging*.
20. Rohit Pandey, and Hardik B. Ramani, “Recent Development and Applications of Additive Layer Manufacturing Processes for Metal Components Production,” Journal of Innovative Engineering and Research (JIER) Vol.- 4, Issue -2, October 2021, pp. 11-14.
21. Amit D. Tajne, Prashant Kumar Shrivastava, “Optimization of Tool Wear for Cobalt based Superalloy,” Journal of Innovative Engineering and Research (JIER) Vol.- 7, Issue -1, April 2024, pp. 8-11.
22. Md Misbahuddin, Prashant Kumar Shrivastava, “Optimum Longitudinal and Bending Response Analysis of SS304 Stainless Steel Corrugated Core,” Journal of Innovative Engineering and Research (JIER) Vol.- 6, Issue -1, April 2023, pp. 37-44.
23. V Raghavendra Rao, Prashant Kumar Shrivastava, “Design and Analysis of Herringbone Gear using Polymers for using off Road Vehicles,” Journal of Innovative Engineering and Research (JIER) Vol. - 6, Issue -2, October, 2023, pp. 11-15.