LaserTwin-AM: In-Situ Optical and 3D Point-Cloud Spatiotemporal Fusion for Closed-Loop Defect Mitigation in Laser Powder Bed Fusion

Arnulfo Salvo
Department of Computer Science, University of North Texas, Denton, TX, USA.
arnulfosalvo913@unt.edu

Abstract

Laser powder bed fusion has matured from a prototyping technique into a production pathway for geometrically complex metallic components, yet its industrial credibility remains constrained by stochastic defects that emerge at scales faster than conventional inspection can resolve. This paper presents LaserTwin-AM as a system-level architecture for closed-loop defect mitigation that fuses in-situ optical observations with layerwise three-dimensional point-cloud reconstructions inside a digital twin. The argument is not that sensing alone can eliminate porosity, lack of fusion, balling, spatter-induced irregularities, or delamination, but that a disciplined spatiotemporal fusion layer can convert heterogeneous signals into actionable process state estimates while preserving traceability, uncertainty, and operational safety. The proposed architecture distributes sensing, reconstruction, fusion, and actuation across edge and plant-level infrastructure, with explicit interfaces for model governance, calibration, data provenance, and human oversight. The paper examines structural trade-offs among latency, accuracy, coverage, bandwidth, and certification burden; compares the architecture with adjacent practices in cyber-physical manufacturing and multimodal inspection; and discusses fairness, liability, sustainability, and policy implications that arise when defect mitigation becomes partially autonomous. It concludes that LaserTwin-AM is best understood as a socio-technical infrastructure rather than a single model, and that its credibility depends as much on validation regimes, lifecycle governance, and deployment economics as on fusion accuracy alone.

Keywords

laser powder bed fusion; digital twin; multimodal fusion; point cloud; in-situ monitoring; closed-loop control; defect mitigation; additive manufacturing governance.

1. Introduction

Laser powder bed fusion occupies a distinctive position in advanced manufacturing because it couples extreme local physics with plant-scale production logic. The melt pool evolves under rapid heating and cooling, powder displacement, vapor recoil, and successive thermal cycles, producing defects whose causes may be distributed across scan strategy, material condition, gas flow, machine calibration, and part geometry [1], [2]. As production demands increase, the gap between what can be observed during build and what can be certified after build becomes a central systems problem. In-situ monitoring has advanced through optical imaging, thermography, acoustic sensing, and layerwise imaging, but these modalities are often deployed as passive records or isolated alarms rather than as closed-loop control inputs [3], [4], [5]. The result is a data-rich but action-poor environment in which defects are documented after they have formed.

LaserTwin-AM responds to this gap by treating the digital twin as an operational fusion environment rather than a visualization layer. The architecture combines in-situ optical observations with three-dimensional point-cloud reconstructions of each deposited layer so that transient melt-pool behavior can be interpreted against evolving geometry. Optical data provide high temporal resolution and sensitivity to thermal emission and spatter, while point clouds provide spatial context for layer topography, part boundaries, and deviations from nominal geometry. Neither modality is sufficient in isolation. Optical signals are difficult to register to part coordinates, while point clouds are comparatively slow and may miss sub-layer phenomena. The fusion problem is therefore not merely algorithmic; it is an infrastructure problem involving time synchronization, spatial calibration, data lifecycle management, uncertainty propagation, and actuation authority.

The contribution of this paper is a system-level account of how such fusion can support closed-loop defect mitigation without overstating the maturity of autonomous control. The discussion emphasizes architecture, trade-offs, governance, and deployment rather than narrow model performance. It considers how a digital twin can maintain a living state estimate across layers, how defect mitigation actions can be bounded by safety envelopes, and how validation and certification can accommodate systems that continuously adapt. It also examines fairness across materials and machines, sustainability implications of rework and scrap, and policy questions related to liability when an automated action alters a build. The paper positions LaserTwin-AM as a socio-technical infrastructure whose value depends on coupling sensing, decision logic, human authority, and auditable evidence into a coherent operational regime.

2. Problem Domain and System Requirements

Laser powder bed fusion defects arise from coupled physical mechanisms that are difficult to isolate in production. Porosity may originate from gas entrapment or keyhole instability, lack of fusion may follow insufficient energy density or unfavorable powder spreading, and balling or spatter may alter subsequent layer continuity [1], [5]. These phenomena are not independent of geometry, since overhangs, thin walls, and support structures modify local heat transfer and powder deposition. A closed-loop system must therefore estimate not only whether an anomaly is present but also whether it is spatially consequential, temporally persistent, and correctable within the remaining build. This requirement pushes the architecture beyond single-frame anomaly detection toward stateful reasoning across time and space.

The sensing requirements are similarly demanding. Optical monitoring can capture melt-pool intensity, plume behavior, and spatter events at high frame rates, but it may be sensitive to emissivity changes, window contamination, and optical occlusion [3], [12]. Point-cloud reconstruction provides layerwise geometry and can reveal topographic discontinuities, but its acquisition rate, resolution, and registration accuracy vary with scanner configuration and calibration drift [10], [11]. Multimodal inspection research shows that fusing geometric and image data can improve defect detection when registration and feature selection are handled carefully, but such fusion also introduces dependencies on synchronization quality and cross-modal calibration [9]. A production-grade system must quantify these dependencies rather than assuming that more data necessarily produces better decisions.

From a system perspective, the requirements can be grouped around temporal alignment, spatial registration, uncertainty, actuation, and evidence. Temporal alignment must be precise enough that an optical event can be associated with the correct scan vector and layer region,

while spatial registration must map observations into a common part coordinate frame that remains stable across thermal distortion. Uncertainty must be represented because defect probabilities interact with production risk and because different actions carry different consequences. Actuation requirements include bounded authority, rollback or pause logic, and clear escalation to human operators. Evidence requirements include provenance, model versioning, calibration history, and the ability to reconstruct why a decision was made. These requirements are interdependent, and weakening one often shifts burden to another, making architecture a matter of deliberate trade-off rather than additive capability.

3. LaserTwin-AM Reference Architecture

The LaserTwin-AM architecture is organized as a layered system spanning edge sensing, geometric reconstruction, twin state estimation, fusion reasoning, decision support, actuation interfaces, and audit services. The edge layer hosts time-sensitive acquisition and preprocessing close to the machine, reducing latency and preserving raw data for later analysis when bandwidth or privacy constraints require selective transmission. The reconstruction layer converts layerwise scanner data into point-cloud representations and maintains calibration metadata that links sensor coordinates to the machine and part frames. The twin layer maintains a persistent digital state across the build, incorporating nominal design, actual layer geometry, process observations, and decision history. This separation is important because a twin that merely mirrors geometry cannot support mitigation, while a twin that ingests every sensor stream without structure becomes an unmanageable data lake.

The fusion core is the central analytical layer, but it is deliberately constrained by governance and safety envelopes. It receives temporally indexed optical features and spatially indexed point-cloud features, estimates defect likelihood and relevance, and proposes mitigation options with confidence and rationale. Decisions are not released directly to the machine without passing through an authority layer that enforces material-specific rules, machine limits, and operator permissions. This design reflects lessons from cyber-physical manufacturing architectures, where integration must balance autonomy with deterministic safety and maintainability [19]. It also reflects digital-twin research emphasizing that twins are most valuable when they connect models, data, and services across the lifecycle rather than when they remain isolated simulation artifacts [6], [7], [8].

A major architectural trade-off concerns centralization versus distribution. A plant-level twin can aggregate data across machines, support fleet learning, and provide consistent governance, but it may introduce communication latency and single points of failure. An edge-heavy design can react quickly and remain resilient when connectivity is intermittent, but it complicates model updates, cross-machine learning, and centralized audit. LaserTwin-AM therefore adopts a hybrid pattern in which latency-critical inference and safety checks run locally, while model training, calibration analytics, and long-term evidence management occur in controlled plant or cloud environments. This hybridity is not merely technical; it distributes accountability among shop-floor operators, process engineers, quality organizations, and platform providers, each of whom requires different access, observability, and override rights.

4. Spatiotemporal Fusion Methodology

The spatiotemporal fusion problem begins with the recognition that optical and point-cloud data live on different clocks and coordinate frames. Optical signals are continuous or high-frequency and are tied to scanner position, exposure settings, and sensor geometry, while

point clouds are discrete and associated with layer completion or periodic scanning. Fusion must align these streams into a common temporal window and spatial region so that a thermal anomaly can be associated with a topographic feature. This association is imperfect because thermal distortion changes geometry and because scanner timing may drift. The architecture therefore treats alignment as a probabilistic calibration problem supported by fiducial markers, machine logs, and periodic registration checks rather than as a one-time setup step.

Feature representation must preserve both local and contextual information. Point-cloud learning methods demonstrate that raw geometric sets can be processed without imposing grid structure, which is attractive for irregular layer topography and complex part boundaries [10], [11]. Optical features, by contrast, may capture rapid melt-pool fluctuations that are not visible in post-layer geometry. A useful fusion strategy combines local descriptors near the scan location with a broader context describing part geometry, support structure proximity, and recent layer history. The twin can then reason about whether an anomaly is likely to be self-correcting, such as transient spatter that is later remelted, or persistent, such as a depression that propagates across layers. This stateful interpretation is where the digital twin adds value beyond conventional image classification.

Uncertainty must be propagated through the fusion process and expressed in operational terms. A high defect probability in a noncritical region may warrant recording and continued observation, while moderate probability near a structural feature may justify a pause or local remelting. These decisions require calibrated confidence and an understanding of model limitations across materials, geometries, and machine conditions. Research on process monitoring indicates that different sensing modalities respond differently to defect types, and that no single modality dominates across all conditions [3], [5], [13], [14]. Fusion should therefore expose modality-specific evidence rather than collapsing everything into an opaque score. The resulting twin state becomes an auditable composite of observations, inferred geometry, uncertainty, and prior decisions, enabling later review and continuous improvement.

5. Closed-Loop Defect Mitigation and Validation

Closed-loop mitigation in LaserTwin-AM is bounded by the reality that many corrective actions are irreversible or carry secondary risks. Reducing laser power may lower keyhole porosity but risk lack of fusion, increasing scan speed may reduce overheating but alter microstructure, and local remelting may heal topography while introducing residual stress. The architecture therefore distinguishes between observational actions, such as increased sampling or conservative pausing, and corrective actions, such as parameter modulation or remelting, which require stricter authority and validation. This graduated autonomy is consistent with measurement science frameworks for additive manufacturing, where sensing, control, and qualification must evolve together rather than as disconnected initiatives [18]. It also aligns with process-control research that emphasizes closed-loop stability and observability over isolated detection accuracy [4].

Defect-specific logic must account for spatial scale and temporal persistence. Spatter and balling may be detected optically and confirmed by subsequent layer geometry, while lack of fusion may appear as subtle topographic irregularity that becomes evident only across multiple layers. Delamination may require immediate pause because continued deposition can compound risk. The twin can support these distinctions by maintaining layerwise history and by comparing observed geometry against predicted distortion, drawing on work that links process conditions to distortion and residual stress [16], [17]. The aim is not to replace

metallurgical understanding with data-driven inference but to situate inference within a process model that constrains plausible actions and expected outcomes.

Validation must operate at multiple levels. At the sensor level, calibration and drift must be monitored. At the fusion level, defect detection must be tested against ground truth from destructive and nondestructive evaluation, including cases where defects are absent. At the system level, closed-loop interventions must be evaluated for their effect on build quality, cycle time, and operator workload. Digital twin validation is especially challenging because the twin is not a static model; it evolves with data and may be updated across machines. A credible validation regime therefore requires frozen model versions for certification, staged rollouts for updates, and continuous monitoring for performance drift. Comparisons with autonomous systems in other domains suggest that safe autonomy depends less on peak model accuracy than on fail-safe behavior, transparency, and the ability to recover from unanticipated conditions.

6. Governance, Fairness, Policy, and Sustainability

Governance is central to LaserTwin-AM because the system creates evidence that may be used for quality release, liability allocation, and regulatory review. Data ownership must be clarified when machine builders, material suppliers, part manufacturers, and end users contribute to the twin. Provenance must record sensor calibration, model version, operator intervention, and environmental conditions so that a decision can be reconstructed after the fact. Liability becomes particularly complex when an automated action alters a build and the resulting part later fails. Policy frameworks should therefore distinguish between decision support, supervised autonomy, and unsupervised control, with corresponding requirements for documentation, human oversight, and incident reporting. These concerns connect to broader debates about algorithmic accountability and the ethics of automated decision systems [20].

Fairness in this context concerns equitable performance across materials, geometries, machines, and facility conditions. A fusion model trained predominantly on one alloy or machine platform may perform poorly on another, not because the physics is entirely different but because sensor characteristics, powder morphology, and calibration practices differ. If closed-loop mitigation is available only to well-resourced manufacturers with extensive training data, smaller suppliers may face a capability gap that influences market access and part qualification. Governance should encourage shared benchmarks, calibrated data schemas, and transparent reporting of performance envelopes. It should also avoid treating fairness as a purely statistical property, since operational fairness depends on access to maintenance, training, and the ability to override automated recommendations.

Sustainability implications are similarly systemic. Defect mitigation can reduce scrap and rework, but it can also increase energy use through additional scanning, sensing, and data processing. The net benefit depends on defect severity, part value, and the carbon intensity of the production system. A twin that records rework and scrap can support lifecycle assessment and process optimization, but only if data management itself is resource-aware. Policy incentives that reward scrap reduction should be coupled with transparency about additional energy and data infrastructure burdens. In this sense, LaserTwin-AM is not automatically sustainable; it is a capability whose environmental value must be demonstrated through measured outcomes and disciplined deployment.

7. Deployment and Operational Trade-offs

Deployment requires integration with existing machine controllers, manufacturing execution systems, quality workflows, and cybersecurity boundaries. Many production machines are closed platforms with limited real-time interfaces, so the architecture must accommodate read-only monitoring, supervised advisory modes, and controller-mediated actuation. Edge devices must be hardened against shop-floor conditions, and network segmentation must protect both intellectual property and safety-critical functions. Maintenance is often underestimated: optical windows foul, scanners drift, point-cloud calibration degrades, and models become stale as materials or geometries change. A sustainable deployment plan treats calibration and model maintenance as recurring operational costs rather than one-time engineering tasks.

The economic case depends on the distribution of defect costs. For high-value aerospace or medical parts, early pause and rework may justify substantial sensing and compute investment. For lower-value parts, the same infrastructure may be uneconomic unless shared across a fleet or subsidized through quality assurance benefits. Latency versus accuracy is another recurring trade-off: local inference can act quickly but may lack fleet context, while plant-level analytics can improve accuracy but may not meet real-time constraints. Human factors also matter, because operators may mistrust systems that produce frequent alarms without clear explanations, and they may overtrust systems that appear authoritative. Interface design must therefore communicate uncertainty, rationale, and recommended actions without overwhelming users.

Scalability introduces further tensions. A fleet-level twin can enable cross-machine learning and consistent governance, but it also concentrates risk and raises questions about data sharing among competitors. Federated or privacy-preserving approaches may help, but they add complexity and can reduce interpretability. Resilience requires fallback modes that preserve production when fusion services are unavailable, and obsolescence planning requires that data schemas and model artifacts remain readable across equipment upgrades. These operational realities suggest that LaserTwin-AM should be deployed incrementally, beginning with traceable monitoring and advisory functions before moving toward bounded corrective actions. Such staging allows organizations to build calibration discipline, operator trust, and validation evidence before granting greater autonomy.

8. Conclusion

LaserTwin-AM frames closed-loop defect mitigation in laser powder bed fusion as an infrastructural challenge rather than a single sensing or modeling breakthrough. By fusing in-situ optical observations with layerwise point-cloud geometry inside a governed digital twin, the architecture can maintain a spatiotemporal state estimate that supports defect interpretation, bounded intervention, and auditable learning. The paper has argued that the value of such fusion depends on temporal alignment, spatial registration, uncertainty representation, actuation authority, and evidence management, all of which are system-level concerns. It has also emphasized that trade-offs among latency, accuracy, centralization, resilience, and maintainability shape what can be responsibly deployed in production.

The forward-looking agenda should prioritize validated uncertainty, staged autonomy, interoperable data schemas, and lifecycle governance. Material-specific validation remains essential because defects and corrective actions are not generic across alloys and geometries. Human oversight must be designed as an active component rather than a residual safety net. Policy frameworks should clarify liability and documentation expectations for supervised autonomy, while sustainability assessment should measure both scrap reduction and the

resource intensity of the twin infrastructure. If these conditions are met, LaserTwin-AM can serve as a credible pathway toward defect mitigation that is technically effective, operationally resilient, and socially accountable.

References

1. Khairallah, S. A., Anderson, A. T., Rubenchik, A., & King, W. E. (2016). Laser powder-bed fusion additive manufacturing: Physics of complex melt flow and formation mechanisms of pores, spatter, and denudation zones. Acta Materialia, 108, 36-45.

2. King, W. E., Anderson, A. T., Ferencz, R. M., Hodge, N. E., Kamath, C., Khairallah, S. A., & Rubenchik, A. M. (2015). Laser powder bed fusion additive manufacturing of metals; physics, computational, and materials challenges. Applied Physics Reviews, 2(4), 041304.

3. Everton, S. K., Hirsch, M., Stravroulakis, P., Leach, R. K., & Clare, A. T. (2016). Review of in-situ process monitoring and in-situ metrology for metal additive manufacturing. Materials & Design, 95, 431-445.

4. Tapia, G., & Elwany, A. (2014). A review on process monitoring and control in metal-based additive manufacturing. Journal of Manufacturing Science and Engineering, 136(6), 060801.

5. Grasso, M., & Colosimo, B. M. (2017). Process defects and in situ monitoring methods in metal powder bed fusion: A review. Measurement Science and Technology, 28(4), 044005.

6. Grieves, M., & Vickers, J. (2017). Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In F.-J. Kahlen, S. Flumerfelt, & A. Alves (Eds.), Transdisciplinary perspectives on complex systems (pp. 85-113). Springer.

7. Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. IEEE Transactions on Industrial Informatics, 15(4), 2405-2415.

8. Wang, Y., Zheng, P., & Xu, X. (2020). Digital twin-driven smart manufacturing: Connotation, reference model, applications and research issues. Robotics and Computer-Integrated Manufacturing, 61, 101837.

9. Li, X., He, H., Huang, J., Liu, R., Qian, T., & Hon, C. (2025). ASFM-AFD: multimodal fusion of AFD-optimized LiDAR and camera data for paper defect detection. IEEE Transactions on Instrumentation and Measurement, 74, 1-14.

10. Qi, C. R., Su, H., Mo, K., & Guibas, L. J. (2017). PointNet: Deep learning on point sets for 3D classification and segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 652-660). IEEE.

11. Qi, C. R., Yi, L., Su, H., & Guibas, L. J. (2017). PointNet++: Deep hierarchical feature learning on point sets in a metric space. Advances in Neural Information Processing Systems, 30, 5099-5108.

12. Lane, B., Whitenton, E., & Moylan, S. (2015). Thermographic measurements of the commercial laser powder bed fusion process at NIST. Rapid Prototyping Journal, 21(5), 574-584.

13. Ye, D., Hong, G. S., Zhang, Y., Zhu, K., & Fuh, J. Y. H. (2018). Defect detection in selective laser melting technology by acoustic signals with deep belief networks. International Journal of Advanced Manufacturing Technology, 96(5-8), 2791-2801.

14. Scime, L., & Beuth, J. (2018). Anomaly detection and classification in a laser powder bed additive manufacturing process using a trained computer vision algorithm. Additive Manufacturing, 19, 114-126.

15. Scime, L., & Beuth, J. (2019). A multi-scale convolutional neural network for autonomous anomaly detection and classification in a laser powder bed fusion additive manufacturing process. Additive Manufacturing, 24, 273-286.

16. Francis, J., & Bian, L. (2019). Deep learning for distortion prediction in laser-based additive manufacturing using big data. Manufacturing Letters, 20, 10-14.

17. DebRoy, T., Wei, H. L., Zuback, J. S., Mukherjee, T., Elmer, J. W., Milewski, J. O., Beese, A. M., Wilson-Heid, A., De, A., & Zhang, W. (2018). Additive manufacturing of metallic components – Process, structure and properties. Progress in Materials Science, 92, 112-224.

18. Mani, M., Lane, B., Donmez, A., Feng, S., Moylan, S., & Fesperman, R. (2015). Measurement science needs for real-time control of additive manufacturing powder bed fusion processes. National Institute of Standards and Technology.

19. Lee, J., Bagheri, B., & Kao, H. A. (2015). A cyber-physical systems architecture for Industry 4.0-based manufacturing systems. Manufacturing Letters, 3, 18-23.

20. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. Big Data & Society, 3(2), 2053951716679679.