LLM Agent-guided Stochastic Control of Multi-layer and Multi-track Laser Powder Bed Fusion Process

Apurba Sarker
apurba@vt.edu

Virginia Tech
Tajwar Razib
BUET
Anup Saha
BUET
Sourav Saha
Virginia Tech

Research Article

Keywords: Laser powder bed fusion, LLM agents, Bayesian calibration, Feedforward control, Agentic AI

Posted Date: April 7th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9315895/v1

Additional Declarations: The authors declare no competing interests.

LLM Agent-guided Stochastic Control of Multi-layer and Multi-track Laser Powder Bed Fusion Process

Apurba Sarker[1], Tajwar Razib[2], Anup Saha[2], Sourav Saha[1*]

[1*]Kevin T. Crofton Department of Aerospace and Ocean Engineering, Virginia Polytechnic and State University, 1600 Innovation Drive, Blacksburg, VA, 24060, United States.
[2]Department of Mechanical Engineering, Bangladesh University of Engineering and Technology, Dhaka, Bangladesh.

*Corresponding author(s). E-mail(s): souravsaha@vt.edu;

Abstract

Achieving uniform melt-pool geometry in laser powder bed fusion (LPBF) remains difficult due to nonlinear parameter interactions, stochastic powder-bed conditions, and progressive heat accumulation. These challenges are intensified by the lack of a sustainable digital manufacturing platform that can unify data, models, and control. To address this issue, this work introduces a unified LLM-agent-guided framework that integrates stochastic calibration and real-time control within a single pipeline. The work uses two co-ordinated agents. **CalibAgent** performs Bayesian Markov Chain Monte Carlo calibration of an analytical thermal model of a multi-layer, and multi-track LPBF process using NIST AM-Bench melt-pool data, inferring distributions over absorptivity and thermal diffusivity. **ControlAgent** then uses the calibrated model to compute spatially varying laser-power schedules via analytical inversion, requiring no iterative optimization and executing in milliseconds. Demonstrations on IN718 across four scan geometries (horizontal and 45° cross-hatch, concentric triangular hatch, and a multi-layer L-shaped build) show that the controller reduces steady-state depth fluctuations by 83%–99% and eliminates turnaround spikes of up to 41 μm. In the multi-layer case, it maintains near-uniform depth across four layers without retraining or parameter adjustment. The framework requires no task-specific training, reward engineering, or expert tuning, enabling immediate extension to new materials and scan strategies.

Keywords: Laser powder bed fusion, LLM agents, Bayesian calibration, Feedforward control, Agentic AI

1 Introduction

Metal additive manufacturing (AM) has fundamentally transformed the production capacity, time, and quality of high-value structural components [12]. Across aerospace, biomedical, and defense industries, AM enables the fabrication of geometrically complex parts with minimal material waste and shortened lead times [40]. Unlike conventional subtractive methods, AM builds parts layer by layer through the local deposition and consolidation of material [53]. Critically, final part quality encompassing density, microstructure, mechanical properties, surface finish, and dimensional accuracy is not determined at a single moment rather it is continuously shaped throughout every layer of the build [8, 19]. This layer-wise evolution makes AM uniquely sensitive to precise, continuous management of process conditions.

In laser-based metal AM processes, the primary controllable inputs laser power, scan speed, hatch spacing, layer thickness, and scan strategy do not influence part quality independently [7]. Instead, they interact nonlinearly through a cascade of coupled physical phenomena [42]. The laser first determines the local energy density, which drives absorption and heat generation; these, in turn, govern the temperature field, melt pool geometry, fluid dynamics, solidification rate, and ultimately the microstructure and residual stress [26, 28]. Over fifty distinct parameters have been identified as potentially impacting process outcomes in laser powder bed fusion (LPBF) alone [15, 31]. The consequence is that no single parameter can be adjusted in isolation. Identifying process windows that satisfy all quality constraints therefore entails exploring a high-dimensional, nonlinearly coupled parameter space that is expensive to sample experimentally and challenging to model computationally.

At the heart of these challenges lies the melt pool, the small volume of molten material created beneath the laser during each scan. Melt pool geometry, thermal history, and stability collectively determine the solidification microstructure, grain morphology, porosity content, and interlayer bonding that define the structural integrity of the final part [8, 26]. When the melt pool is too shallow, adjacent layers fail to bond sufficiently, producing lack-of-fusion voids [14, 52]. When it is too deep, the process transitions into keyhole mode, where vapor depression and recoil pressure collapse to form spherical subsurface pores [21, 58]. Between these extremes lies a narrow operating window; the essential control problem is to keep the melt pool within this window everywhere along the scan trajectory.

Achieving this level of control demands fast and reliable melt-pool prediction across different materials and process conditions. It needs to be fast enough to predict melt pool behavior rapidly enough to inform parameter scheduling during the scan itself. High-fidelity models can capture detailed physics, but they are too expensive for real-time use and large parameter studies [2, 32, 45]. By contrast, analytical heat-conduction models offer closed-form descriptions of a moving heat source that evaluate in milliseconds, are interpretable in terms of physical parameters, and provide sufficiently accurate representations of melt pool geometry in the conduction-dominated regimes typical of LPBF [2, 68]. More broadly, the preference for fast, interpretable, physics-based models aligns with recent work on mechanistic artificial intelligence in manufacturing, where physics and data-driven methods are combined for prediction rather than relying on black-box models [46, 50, 62].

Classical analytical heat-conduction models, beginning with Rosenthal's moving heat-source solution [48] and the broader conduction framework of Carslaw and Jaeger [6], form the basis of many reduced-order LPBF thermal models. Later developments introduced distributed heat-source forms, such as the Goldak double-ellipsoidal model, while retaining fast evaluation [13]. For the small, conduction-dominated melt pools typical of LPBF, such models remain physically justified [66]. Their accuracy has also been improved through temperature-dependent properties [23], Gaussian heat-source descriptions [10, 63], and volumetric formulations [41, 67]. As a result, analytical models continue to provide a fast, scalable, and physically interpretable basis for melt-pool prediction, process-map construction, and LPBF optimization [3, 61]. Despite their promise for LPBF control, existing analytical applications have been largely confined to idealized or geometrically limited scenarios. Most prior analytical control studies considered single-track or single-layer geometries, where the thermal history is relatively simple and the dominant source of melt pool variation—residual heat accumulation across multiple vectors and successive layers—is absent by construction [56, 65]. In practice, however, LPBF parts are built from hundreds of overlapping scan tracks across many layers, and it is precisely this multi-track, multi-layer regime where thermal nonuniformity is most severe. At track turnarounds, the combination of scan deceleration and residual preheat from adjacent vectors drives sharp melt pool depth spikes [11, 52]. Across layers, progressive substrate heating shifts the baseline thermal state upward, causing depth to increase with each successive layer even under constant power [4]. These are not edge cases; rather, they are the dominant operating conditions of real LPBF builds, yet they remain unaddressed by existing analytical control frameworks [27, 30, 57].

The utility of any analytical model further depends on the accuracy of its physical parameters. Deterministic calibration strategies, which seek a single best-fit parameter vector, are inadequate for two reasons. First, they treat calibration as a point-estimation problem and discard uncertainty information, directly relevant to the expected variability of melt-pool behavior. Second, the LPBF process is inherently stochastic due to powder-bed inhomogeneity, local thermal history, and measurement uncertainty [26, 29]. A controller calibrated on a deterministic model therefore operates with implicit overconfidence. What is needed instead is a probabilistic calibration strategy that infers a distribution over model parameters from experimental data, propagates that uncertainty through the model, and yields predictions that match not only mean melt-pool geometry but also the observed scatter. Bayesian calibration provides a principled framework for this purpose, enabling uncertainty quantification and statistically consistent predictive intervals [22, 24], thereby making the model honest about what it knows—and what it does not. This perspective aligns with recent efforts in computational intelligence for engineering systems, which emphasize building predictive models grounded in physical insight rather than relying solely on empirical fitting [37, 49].

Finally, even a physically well-calibrated analytical control system requires a practical execution layer that adapts to different materials, scan strategies, and quality objectives without demanding reconfiguration of code. Reinforcement learning (RL) approaches address this execution problem by learning a control policy through environment interaction [55], and have been explored in manufacturing and additive manufacturing for adaptive process control [25, 47]. However, RL methods typically require extensive training data or simulation rollouts, careful reward engineering, and substantial computational cost. Moreover, learned policies are often difficult to interpret and do not generalize reliably across materials, machines, or process conditions without retraining [9]. A more natural fit for orchestrating a well-defined computational pipeline in response to high-level user objectives is an LLM-based, tool-using agentic framework. Unlike a static script that executes a fixed sequence of function calls, an agentic system interprets structured numerical outputs at each step and branches its execution accordingly: it may accept, modify parameters, or rerun a computation depending on the returned diagnostics. This distinction is critical—the same agent handles cases that require different sampling budgets, different controller hyperparameters, and different numbers of refinement iterations, all determined at runtime from the data rather than from hard-coded logic. Such systems leverage the reasoning and tool-use capabilities of large instruction-tuned models to translate natural-language instructions into structured tool calls [51, 64]. Unlike RL-based controllers, this approach requires no task-specific training, no reward design, and no policy optimization. Instead, it builds on zero-shot and tool-augmented reasoning capabilities of foundation models [60], making the execution layer extensible to new materials or scan geometries by modifying tool definitions rather than retraining model components.

This work presents an LLM-agent-guided framework for stochastic calibration and feedforward control of multi-track, multi-layer LPBF. For this purpose, two coordinated agents are used. Bayesian (Markov Chain Monte Carlo)MCMC calibration [20] of a volumetric Gaussian analytical thermal model is performed by *CalibAgent* using melt-pool data from the NIST AM-Bench dataset [39]. Posterior distributions of absorptivity and thermal diffusivity are inferred. A spatially varying laser-power schedule is then computed by *ControlAgent* from the calibrated model so that uniform melt-pool depth is targeted. In both agents, decisions are made from tool outputs. Sampling effort and reruns are selected by CalibAgent based on dataset size and convergence diagnostics. The computed schedule is accepted or refined by ControlAgent based on the depth-uniformity results of the controlled run. The framework is demonstrated on IN718 for horizontal and 45-degree cross-hatch scans, concentric triangular paths, and a multi-layer L-shaped geometry. Across these cases, reductions of 83% to 99% in steady-state melt-pool depth standard deviation are achieved. Turnaround depth spikes of up to 41 μm under constant power are also removed. Both the mean and scatter of the experimental melt-pool dimensions are reproduced by the stochastic calibration. Predicted width and depth are obtained as 142.4 ± 4.6 μm and 123.2 ± 4.0 μm, compared with experimental values of 141.3 ± 2.7 μm and 122.4 ± 2.1 μm. The paper is organized as follows. In Section 2, the model, calibration method, and agent framework are described. In Section 3, the calibration and control results are presented. In Section 4, directions for future work are discussed. In Section 5, the main findings are summarized.

2 Methodology

The current study develops a framework that integrates two coordinated LLM agents into a unified calibration and control pipeline, as shown in Figure 1. **CalibAgent** infers full posterior distributions over the thermal model's physical parameters from experimental melt pool data. It is capable of capturing both mean response and process variability. **ControlAgent** then applies a feedforward controller to generate spatially varying laser power schedules that enforce uniform melt pool depth across complex scan geometries. Both agents operate through structured tool calls to computational backends. At each step, the agent receives structured outputs from the environment, interprets those outputs using explicit decision rules, and selects the next action accordingly. In the calibration workflow, the relevant returned quantities include case availability, number of experimental measurements, MCMC acceptance rate, and posterior uncertainty. In the control workflow, the relevant returned quantities include uncontrolled and controlled depth statistics, residual spike amplitudes, and the achieved reduction in depth variation. These returned quantities determine whether the current result is accepted, whether execution settings are modified, or whether one additional run is triggered. Detailed descriptions of each component are provided in the following subsections.

2.1 Analytical Thermal Model for Melt Pool Prediction

2.1.1 Model Description

The analytical approach in this study adopts a volumetric Gaussian heat source distribution for thermal modeling in order to reproduce the melt pool geometries [41]. The governing equation is presented in Eq. (1),

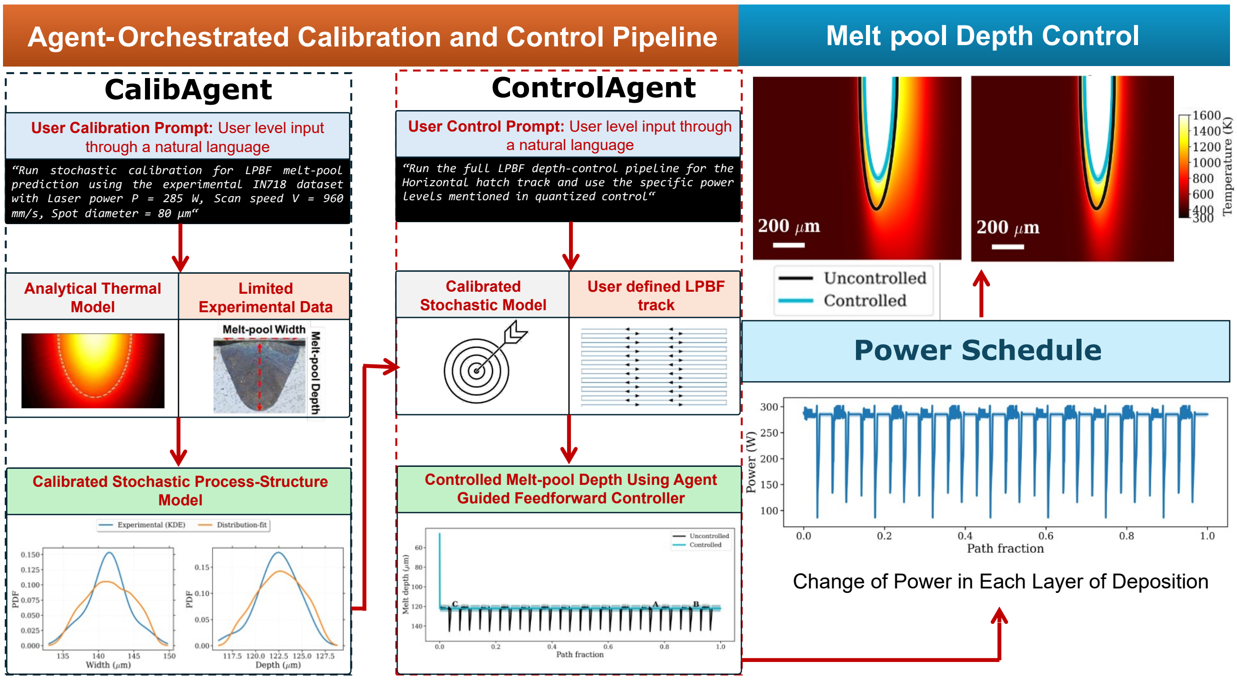

Fig. 1: Overview of the LLM-agent guided framework for LPBF melt pool control. CalibAgent performs Bayesian MCMC calibration of an analytical thermal model from experimental measurements, and ControlAgent uses the calibrated model to compute a laser power schedule targeting uniform melt pool depth across multi-track, multi-layer scan geometries.

where T_0 is the initial temperature (K), η is the absorptivity, and P is the laser power (W). The heat source is modeled as

$$T(x_i,t) = \frac{6\sqrt{3}\,\eta P}{\rho c_p \pi^{3/2}\sqrt{\lambda_1\lambda_2\lambda_3}} \exp\left(-\frac{3x_{ij1}^2}{\lambda_1^2} - \frac{3x_{ij2}^2}{\lambda_2^2} - \frac{3{x_{ij3}^*}^2}{\lambda_3^2}\right) + T_0, \qquad x_{ij3}^* = z_{\text{scale}}\, x_{ij3} \tag{1}$$

where, $\lambda_i = 12\alpha(t-\tau) + \sigma_i^2$, for $i = 1,2,3$. $(x_{ij1}, x_{ij2}, x_{ij3})$ represent the components of the relative position vector $(x_i - x_j)$ between the point of interest x_i and the heat source location x_j. Here, x_{ij1} and x_{ij2} denote the in-plane coordinates, while x_{ij3} corresponds to the build direction.

In the present implementation, a build-direction scaling factor z_{scale} is additionally introduced for the third spatial coordinate. This modifies the effective distance in the x_{ij3} direction only and therefore controls the thermal decay with depth without changing the in-plane (x, y) spread. The parameter serves as an effective anisotropy correction that compensates for depth-shape deviations not fully captured by the simplified analytical heat-source formulation, thereby improving agreement with the experimentally observed melt-pool aspect ratio. Accordingly, the third relative coordinate is replaced by the scaled quantity $x_{ij3}^* = z_{\text{scale}}\, x_{ij3}$ in the exponential term of Eq. (1).

To account for residual heat accumulation from earlier scan locations, the analytical model retains only a localized subset of prior heat-source points in the moving-source summation. For an evaluation point $\mathbf{x}_{\text{eval}}$ at time t_{eval}, only prior source points j satisfying

$$\|\mathbf{x}_{\text{eval}} - \mathbf{x}_j\| \le r_{\text{hist}}, \qquad t_{\text{eval}} - t_j \le \tau_{\text{hist}} \tag{2}$$

are included in the temperature evaluation. Here, r_{hist} is the spatial history radius and τ_{hist} is the temporal history window. Together, these parameters define the spatial and temporal extent of residual-heat carryover from previously deposited scan points. These values can be specified by the user in the track file provided to the control agent, allowing different settings to be assigned for different tracks and scanning strategies.

To allow the heat source model to operate, the following simplifying assumptions are made:

- The temperature field develops as a direct response to the concentrated energy input from the laser beam.
- Thermal energy propagates solely through heat conduction mechanisms. Radiation and other modes of heat transfer are considered negligible.

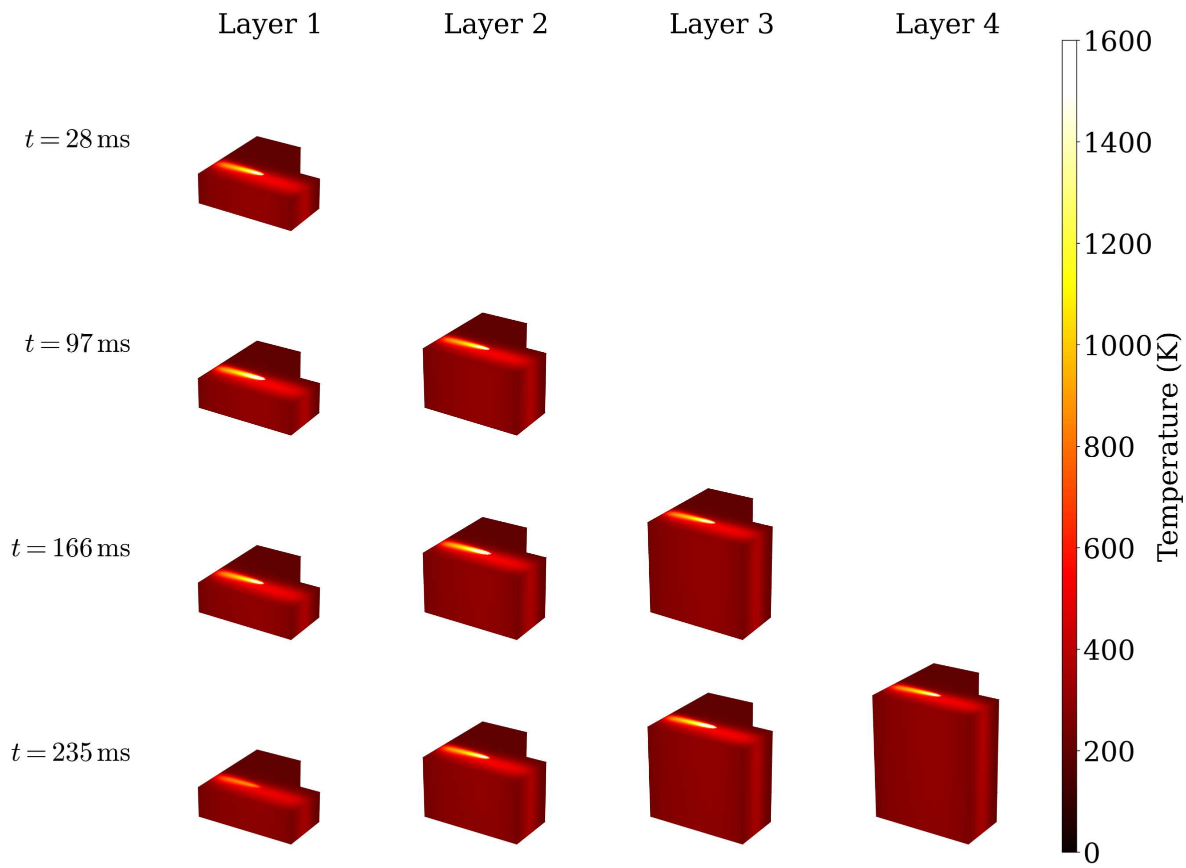

Fig. 2: Three-dimensional surface temperature field snapshots for the L-shaped multi-layer geometry, shown at four time points within each layer's scan (rows: t = 35, 104, 173, 242 ms) across all four build layers (columns: Layers 1–4). The bright concentrated hot spot marks the current laser position, while the surrounding red-to-dark gradient shows how heat diffuses into the substrate.

- The computational domain is regarded as semi-infinite [63] as vaporized alloying elements in the powder caused by the laser beam are neglected due to the small dimension of the melt pool.
- Process parameters such as laser power, scanning speed, powder size, and powder distribution influence the material behavior in metal additive manufacturing processes, as they may alter the predicted temperature field [36]. Temperature-dependent material properties during simulation are treated using the enrichment method [54]. This method employs an iterative scheme that updates the temperature-dependent material properties.

The multi-layer, multi-track thermal build case for an L-shaped geometry is demonstrated in Figure 2.

Table 1: Properties of IN718 used for simulation [1].

Symbol	**Meaning**	**Value**
ρ	Density	8193.3 [kg/m^3]
c_p	Specific Heat	435 [J/kg/K]
k	Thermal Conductivity	$0.016T + 6.39$ [W/m/K]
T_s	Solidus Temperature	1537 [K]
T_l	Liquidus Temperature	1609 [K]
T_v	Vaporization Temperature	3005 [K]
Q_f	Latent Heat of Fusion	295 [kJ/kg]

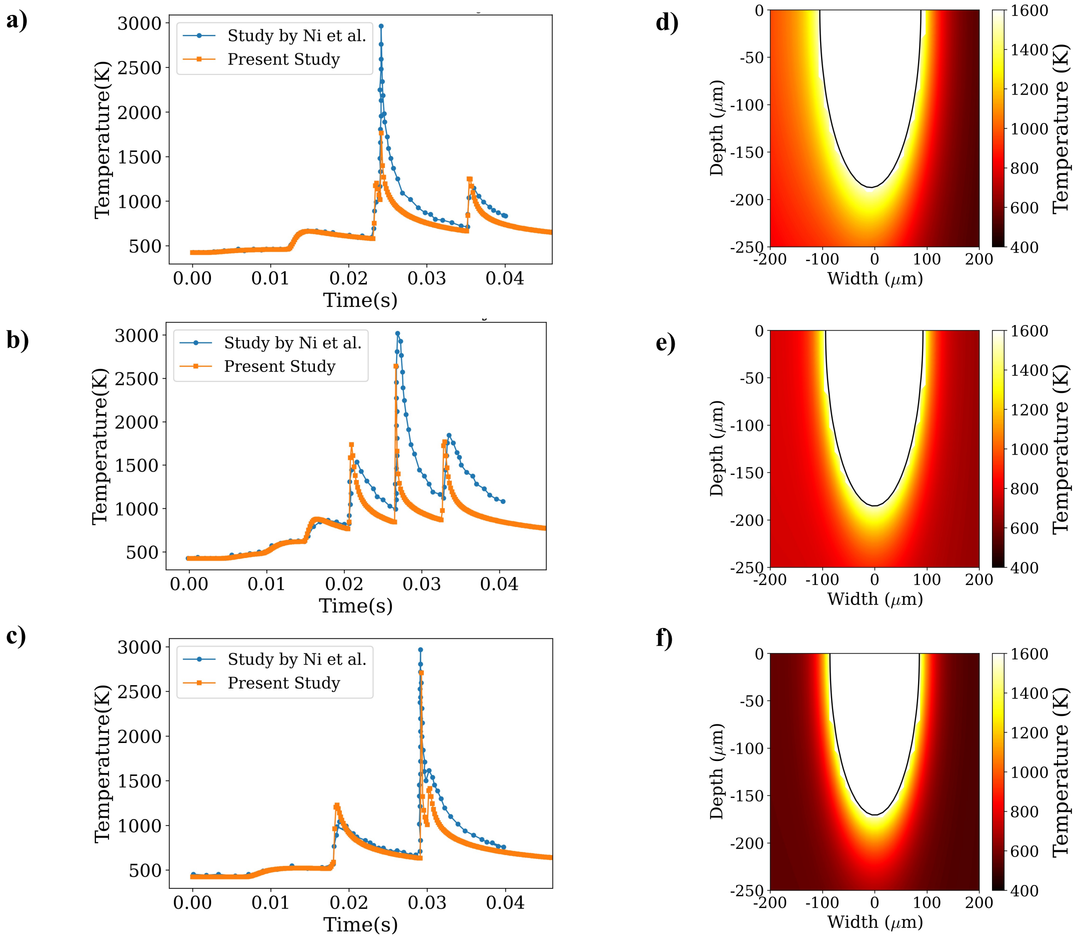

Fig. 3: Validation of the analytical thermal model against the study by Ni et al. [41] for the 5th scan of Layer 1. Temperature histories at three representative path locations (a) the first point, (b) the midpoint, and (c) the last point of the scan show that the present model closely reproduces the reference predictions, capturing both the sharp heating peak as the laser passes and the subsequent cooling decay. The corresponding melt pool cross-sections at each location (d, e, f) show the predicted temperature field with the white dashed contour marking the melt boundary. For visualization, the temperature colormap is clipped at the melting temperature to indicate the melt pool.

2.1.2 Model Validation

To assess the reliability of the developed analytical thermal model, the predicted temperature histories are validated against the numerical results reported by Ni et al. [41] for the 5th scan of Layer 1. The material considered in this study is Inconel 718, and its properties are summarized in Table 1. Three representative locations along the scan track are selected: the first point, the midpoint, and the last point of the scan. Figure 3 (a)–(c) present the temporal temperature evolution at these positions, comparing the present model with the reference study. At all three locations, the proposed model successfully captures the characteristic rapid temperature rise during laser irradiation followed by the gradual cooling phase after the laser passes. The peak temperatures predicted by the model are in close agreement with the reference data, with only minor deviations observed near the maximum values. The cooling trends and the overall shape of the thermal cycles also demonstrate strong consistency with the reference study results. Slight differences in peak magnitude and decay rate are observed due to the uncertainties in heat source calibration and numerical discretization.

In addition to temporal validation, spatial validation is performed by examining the predicted melt pool cross-sections at the corresponding locations, as shown in Figure 3 (d)–(f). The temperature contours exhibit the expected elongated melt pool geometry along the scanning direction. The white dashed contour indicates the melting isotherm, clearly defining the melt boundary. The predicted melt pool shape and depth are physically consistent with reported laser powder bed fusion characteristics, confirming that the model accurately reproduces both the transient thermal response and the melt pool morphology.

2.2 Agent-Guided Offline Stochastic Calibration Strategy

2.2.1 Description of Dataset Used For Calibration

Experimental melt pool measurements are taken from the NIST data publication [39]. The dataset serves as supplemental single-track experiments aligned with the NIST Additive Manufacturing Benchmark (AM-Bench) single-track challenges (2018 and 2022) and are associated with prior spot-size and fidelity studies. The dataset contains single-track laser scans produced using Yb-fiber (ytterbium-doped fiber) lasers on bare plates of Inconel 625 (IN625) and Inconel 718 (IN718) across three laser powder bed fusion (LPBF) systems. Laser power, scan speed, and laser spot diameter are varied, and each track is cross-sectioned and metallographically prepared. Optical micrographs of etched cross sections are provided together with corresponding melt pool depth and width measurements.

Melt pool depth and width are defined using a bounding-box procedure on the optical micrographs. Depth is the greatest vertical extent and width is the greatest lateral extent of the melt pool in the cross section while the top of the bounding box aligned to the plate surface such that any material above the surface is excluded. The annotated labels height and width on the images reflect the bounding-box dimensions and may not always correspond to depth/width depending on melt pool aspect ratio and the spreadsheet should be used as the authoritative source for numeric values.

For the stochastic calibration in this study, the IN718 single-track subset at the fixed process condition $P = 285\,\mathrm{W}$ and $V = 960\,\mathrm{mm/s}$ with spot diameter $80\,\mu\mathrm{m}$ is selected, and the observation set $\mathcal{D} = \{(W_i, D_i)\}_{i=1}^{N}$ is constructed directly from the corresponding rows. This choice provides repeated measurements under a single nominal (P, V, spot) setting, which is essential for inferring not only mean melt pool response but also the experimentally observed variability.

2.2.2 Stochastic Calibration Formulation

Traditional deterministic calibration seeks single best-fit parameter values that minimize prediction error. However, this approach cannot capture the inherent variability observed in experimental melt pool measurements, which arises from stochastic fluctuations in powder bed conditions, local thermal history, and measurement uncertainty. This study formulates calibration as a stochastic inverse problem in which model parameters are treated as random variables whose distributions are inferred from experimental data. Following classical Bayesian calibration frameworks for computer models [20, 24], uncertainty is explicitly represented in the calibration parameters rather than absorbed into deterministic tuning constants.

Two analytical-model parameters are treated as stochastic, thermal absorptivity η and the thermal diffusivity multiplier α. The effective absorptivity is dependent on local reflectivity, powder morphology, and surface condition. Likewise, the effective heat diffusion behavior depends on conduction into the substrate local powder packing, and prior thermal history, which are not explicitly captured by the analytical model. The laser beam radius σ_m is reported from the dataset spot diameter ($D_{4\sigma} = 80\,\mu\mathrm{m}$) via the relation $\sigma_m = D_{4\sigma}/4 = 20\,\mu\mathrm{m}$. The depth-scaling factor z_{scale} is fixed deterministically: it is obtained in a preliminary *mean-fit* step in which a Powell optimizer matches the analytical model to the experimental mean melt pool width and depth. This two-stage decomposition isolates the parameters that primarily govern variability (η, α) from those that govern mean geometry (z_{scale}), reducing the stochastic inference problem to a well-identified four-dimensional space.

The stochastic calibration is therefore written in terms of the hyperparameter vector

$$\boldsymbol{\theta} = \left[\mu_\eta,\ \mu_\alpha,\ \sigma_\eta,\ \sigma_\alpha\right], \tag{3}$$

where $\mu_{\{\cdot\}}$ and $\sigma_{\{\cdot\}}$ denote the mean and standard deviation of each stochastic physical parameter. For a given $\boldsymbol{\theta}$, an ensemble of $N_s = 30$ physical-parameter samples. Each sample represents one possible realization of the process:

$$\eta^{(j)} \sim \mathcal{N}(\mu_\eta, \sigma_\eta^2), \qquad \alpha^{(j)} \sim \mathcal{N}(\mu_\alpha, \sigma_\alpha^2), \qquad where\, j = 1, 2, \ldots, N_s, \tag{4}$$

with σ_m and z_{scale} held fixed. Each sample is propagated through the analytical forward model $\mathcal{F}$. This gives one predicted melt-pool width and one predicted melt-pool depth:

$$\left(W^{(j)}, D^{(j)}\right) = \mathcal{F}\left(\eta^{(j)}, \alpha^{(j)};\ z_{\mathrm{scale}}\right) \qquad where\, j = 1, 2, \ldots, N_s, \tag{5}$$

In this way, one hyperparameter vector $\boldsymbol{\theta}$ does not produce just one prediction. It produces a distribution of predicted width and depth values. That is important, because the experiments also give a distribution, not a single number. The calibration should therefore match the full spread of the data, not only the average.

The posterior distribution over $\boldsymbol{\theta}$ is

$$p(\boldsymbol{\theta} \mid \mathcal{D}) \propto p(\boldsymbol{\theta})\, p(\mathcal{D} \mid \boldsymbol{\theta}), \tag{6}$$

where $\mathcal{D} = \{(W_i, D_i)\}_{i=1}^{N}$ are the experimental observations. Uniform priors are assigned within physically admissible bounds Ω (Table 2), derived from prior NIST dataset studies [38, 59]. Rather than assuming a parametric noise model, a distribution-level discrepancy functional $J(\boldsymbol{\theta})$ is defined and converted to a likelihood [5]:

$$p(\mathcal{D} \mid \boldsymbol{\theta}) \propto \exp\bigl(-\beta\, J(\boldsymbol{\theta})\bigr), \quad \beta = 1, \tag{7}$$

where J combines moment matching and Kernel Density Estimation (KDE)-based Kullback–Leibler divergence (KLD):

$$J(\boldsymbol{\theta}) = w_m\, L_{\text{mean}} + w_s\, L_{\text{std}} + w_{kl}(\text{KL}_W + \text{KL}_D) + w_c\, L_{\text{corr}} + \lambda_r \|\boldsymbol{\sigma}\|^2, \tag{8}$$

Here the mean term checks whether the predicted average width and depth are correct. The standard-deviation term checks whether the predicted scatter is correct. The KL-divergence terms compare the predicted and experimental distributions to refine the fit. The correlation term keeps the link between width and depth. This matters because, in the experiments, width and depth do not vary independently. The regularization term prevents the inferred standard deviations from shrinking to zero or growing unstable.

with weights $w_m = 2.0$, $w_s = 5.0$, $w_{kl} = 1.1$, $w_c = 5$, and regularization $\lambda_r = 10^{-4}$. The larger weight on the variance-related part reflects the main goal of this calibration that is to reproduce the experimental scatter, not just the mean trend.

The posterior is explored with a Metropolis–Hastings random-walk MCMC sampler [18, 35]. Starting from the mean-fit estimate, the chain runs for 30,000 steps with a 12,000-step burn-in. The sampler converges to an acceptance rate near 0.5, consistent with efficient exploration of this low-dimensional posterior. The maximum a posteriori (MAP) estimate is used for downstream predictions. The entire algorithm is summarized in Algorithm 1 of C.

Table 2: Admissible domain Ω for the four stochastic calibration hyperparameters. The beam radius $\sigma_m = 20\,\mu$m and depth-scaling factor z_{scale} are fixed deterministically and are not subject to MCMC inference.

Parameter	Description	Lower bound	Upper bound
μ_η	Mean absorptivity	0.45	0.9
μ_α	Mean diffusivity multiplier	0.6	1.2
σ_η	Std. dev. of absorptivity	0.005	0.06
σ_α	Std. dev. of diffusivity multiplier	0.015	0.10

2.2.3 CalibAgent: Agent-Based Calibration Control

CalibAgent acts as the LLM-based decision layer on top of the stochastic calibration strategy described above. For each user request, the agent first determines whether the requested (P, V, spot) condition exists in the experimental dataset and then inspects the metadata associated with that case, particularly the number of available experimental measurements n_{points}. This quantity is used to choose the initial MCMC budget, since smaller datasets generally require longer chains to obtain stable posterior summaries. The burn-in and ensemble settings are then chosen consistently with this budget.

After the calibration tool returns, the agent evaluates the returned calibration summary using the MCMC acceptance rate, the number of data points, and the relative posterior spread of the inferred parameters, such as σ_η/μ_η and σ_α/μ_α. These quantities are used to categorize the result into three quality levels *acceptable*, *marginal*, or *poor*. Acceptable runs are retained directly. Marginal or poor runs trigger a stronger rerun in which the sampling budget is increased. The rerun settings are computed from the previous run rather than chosen arbitrarily, so the agent adapts the calibration effort to the difficulty of the case. To avoid unbounded retry behavior, the agent performs at most two total calibration attempts, including the initial run. It then

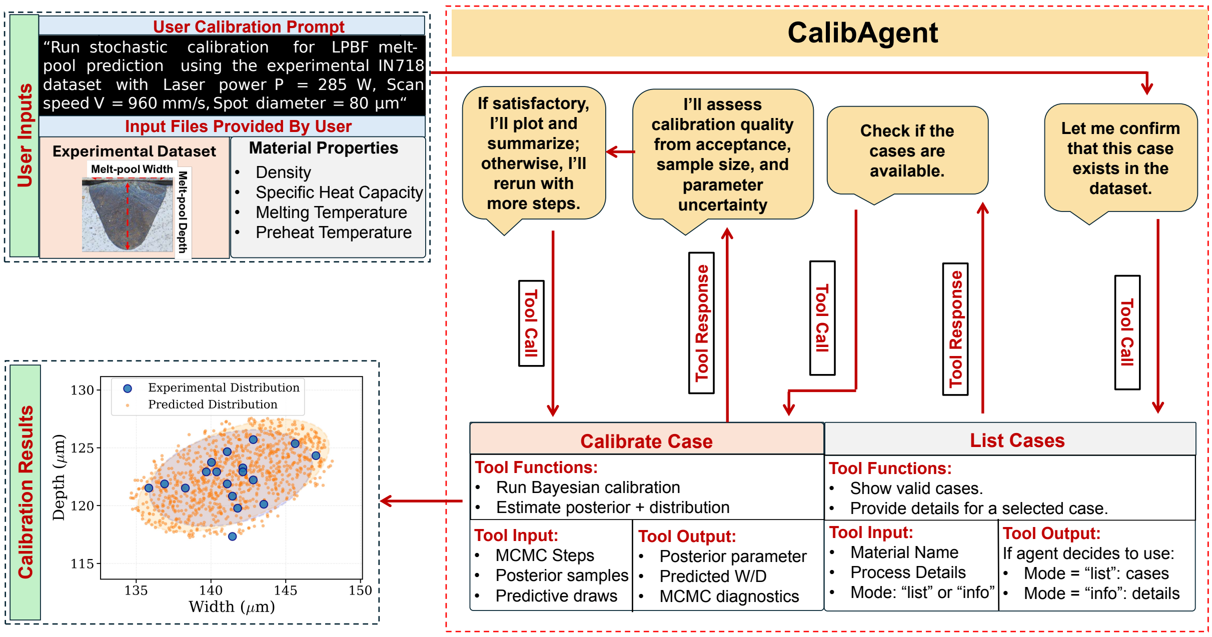

Fig. 4: CalibAgent pipeline demonstrating how inputs (material and process parameters) flow through dataset querying and case selection into a calibration run. Finally in calibration tool Bayesian inference is performed and reports and posterior summaries are generated.

returns the best available result together with the associated diagnostics. In this work, the reported runs use the tool-calling LLM identifier **llama3-groq-tool-use:8b**. The methodology is explained in Figure 4.

This model is the 8B-parameter Groq Tool-Use variant built on Meta's Llama 3 family [16, 34] and tuned to produce structured tool/function calls using tool name and JSON arguments, which is well-suited for reliably triggering the **list_cases** and **calibrate_case** actions used in this study [16, 17, 43]. Ollama provides the tool-calling interface by supplying tool schemas to the model and feeding tool results back into the context, enabling a closed-loop LLM–environment interaction [44].

CalibAgent interacts with the calibration code through two tools:

- **list_cases**: Returns a compact list of available dataset cases as JSON records containing Material, Laser Power (W), Laser Scan Speed (mm/s), and Spot Diameter (μm). This tool is used to verify that a requested case exists and to resolve missing case information before calibration is attempted.
- **calibrate_case**: Executes a complete calibration for one selected case with specified material, P, V, and spot diameter. Inputs include material name, power in Watts, scan speed in mm/s, spot diameter in μm, and the calibration mode, which in the present study combines moment matching and KL divergence. The tool returns a structured JSON summary containing fitted parameter statistics, MCMC acceptance diagnostics, predictive melt-pool statistics, and file paths to saved plots and result files. These returned quantities are then used by the agent to decide whether the calibration should be accepted or rerun with a stronger sampling budget.

CalibAgent runs as a short loop in which the tool-calling model **llama3-groq-tool-use:8b** reads the current context and either requests a tool call or stops when enough information is available. Every interaction is recorded in a machine-readable JSONL log so the full decision path is transparent and reproducible. A short description of the recorded events are demonstrated below to discuss the workflow of the agent.

- **llm_request**: the exact message list provided to the LLM (system prompt + user prompt + any previous tool results), along with the model identifier and a timestamp. Prompts are discussed further in B.1.
- **llm_response**: the LLM output. If the LLM decides to act, the response contains a tool request with a tool name and a JSON argument dictionary.
- **tool_call**: the tool that is executed and the exact arguments used (this is the concrete “action” taken in the environment).
- **tool_result**: the tool output returned to the agent (typically a JSON summary from the calibration backend), which becomes part of the context for the next LLM step.

In the typical flow, the agent first queries the dataset to verify that the requested process condition exists and to retrieve the associated case metadata. It then selects an initial calibration budget from the dataset size and launches the calibration tool. After each run, the returned diagnostics are evaluated in a fixed sequence, the acceptance rate is checked to verify that the Markov chain is exploring the posterior effectively, the posterior spreads are checked to determine whether the inferred uncertainty is excessively broad, and the predicted melt-pool width and depth statistics are compared against the experimental summary. If these indicators are satisfactory, the result is accepted immediately. If they indicate a marginal or poor run, the agent increases the calibration budget and reruns the calibration. After this the best available calibration summary is returned.

It is worth noting what the agent layer contributes beyond what a deterministic script would provide. The computational tools–MCMC sampler, forward model, diagnostic checks–would function identically if called from a fixed Python script. However, the agent mediates between these tools through structured interpretation of their outputs: it selects sampling budgets from dataset metadata, evaluates convergence diagnostics to decide between acceptance and rerun, and computes rerun parameters from effective sample size rather than from predetermined values. Adding a new evaluation criterion (e.g., a Gelman–Rubin convergence check across multiple chains) requires only adding a field to the tool's return value and a line in the system prompt; the agent's decision loop does not need to be rewritten. This extensibility is a direct consequence of mediating tool selection through natural-language interpretation of tool descriptions rather than through hard-coded control flow.

2.3 Agent-Guided Open-Loop Feedforward Control Framework

2.3.1 Build Cases and Scan Strategies

Four scan geometries of increasing complexity are used to evaluate the feedforward control framework. The horizontal and 45° cross-hatch patterns cover a 3×3 mm square domain with uniform hatch spacing, and together isolate the effects of track-end turnaround spikes and inter-track residual heat accumulation under straight and angled scan directions respectively. The concentric triangular pattern introduces sharp 60° corner events at three locations per contour with progressively increasing preheat toward the centroid, providing a non-rectilinear test case. The L-shaped multi-layer geometry extends the evaluation to a realistic part cross-section with geometric discontinuities and layer-to-layer substrate preheating, directly targeting the conditions most relevant to industrial LPBF builds. The tracks are described in more detail in A.

2.3.2 Open-Loop Feedforward Control Problem Formulation

The control objective is to maintain a spatially uniform melt pool depth z^* throughout the entirety of the scan trajectory. The scan trajectory is parameterized by normalized path fraction f. It provides a dimensionless measure of the laser position along the complete scan path. For any intermediate trajectory point i, the path fraction is defined as the distance traveled by the laser from the starting point up to point i, divided by the total length of the full scan path. Thus, the normalized path fraction corresponding to trajectory point i is given by,

$$f_i = \frac{L_i}{L_{\text{total}}}, \tag{9}$$

where L_i is the distance traveled by the laser from the start of the scan up to trajectory point i, and L_{total} is the total scan-path length. This normalization is useful because it allows melt-pool depth, laser power, and thermal-history quantities to be compared consistently across different scan geometries, even when the total path lengths are different.

Since the melt boundary is defined as the T_{melt} isotherm of the analytical thermal model (Eq. (1)), depth uniformity is enforced by requiring $T(\mathbf{x}_i, z^*, t_i) = T_{\text{melt}}$ at every evaluation point i along the path. The target depth z^* is set to the 55th percentile of the uncontrolled depth distribution, which represents a physically realizable and centrally representative value. At each evaluation index i, the trajectory is partitioned into a fixed historical contribution T_{fixed}, arising from all source locations outside a stride window, $\mathcal{W}_i$ of $s = 10$ points, and a controllable contribution, T_{ctrl}, that is linear in the unknown power P_{new} to be applied within $\mathcal{W}_i$. Exploiting this linearity and enforcing the melt boundary condition yields the closed-form inverse

$$P_{\text{new}} = P_{\text{ref}} \cdot \frac{T_{\text{melt}} - T_{\text{fixed}}(\mathbf{x}_i, z^*)}{T_{\text{ctrl}}(\mathbf{x}_i, z^*)}, \tag{10}$$

which is evaluated in a single step without iterative optimization. The computed power is rate-limited by a maximum allowable inter-step change $\Delta P_{\max}$ and clipped to the admissible range $[P_{\min}, P_{\max}]$ to ensure physical realizability. A secondary correction pass is subsequently applied to resolve residual undershoot locations that arise from geometric transitions such as track ends and corners, where the forward prediction accuracy of the first pass is reduced. The full derivation, decomposition, and algorithmic details are provided in E and Algorithm 2.

2.3.3 ControlAgent: LLM-Orchestrated Control Policy

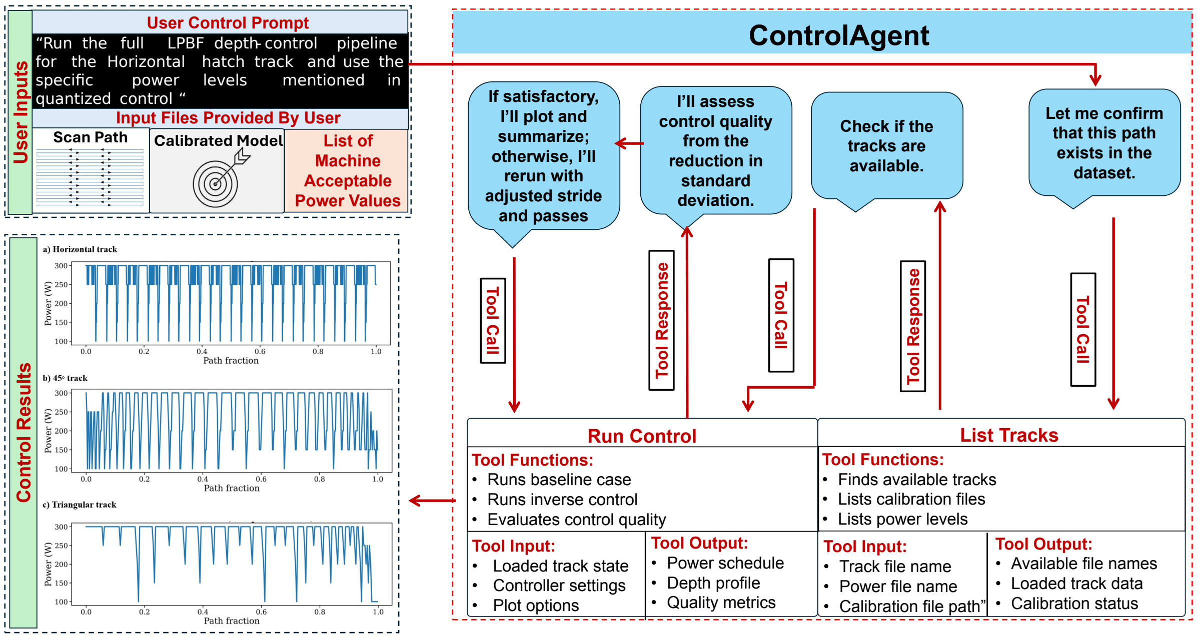

Fig. 5: ControlAgent pipeline demonstrating how the user-specified scan track flows through available track querying and configuration validation into a feedforward control run. The analytical inverse controller computes a spatially varying power schedule, and a control report with melt pool depth statistics returned.

The analytical inverse controller described above is orchestrated through an LLM agent, referred to as the **ControlAgent**. It serves as the decision-making and execution layer between the user and the computational control pipeline. For each requested scan geometry, the agent first confirms that a valid track configuration is available and then launches the control workflow using the calibrated analytical model. After the control run completes, the returned control summary is evaluated using several quantitative indicators: the uncontrolled and controlled steady-state depth standard deviations, the residual spike amplitudes near turnarounds or corners, the percentage reduction in depth variation, and the presence of any remaining localized excursions after the correction pass.

These indicators are used to classify the control result into three levels: *excellent*, *good*, and *poor*. Excellent and good runs are accepted directly. Poor runs trigger one adaptive refinement step in which controller hyperparameters are modified before the control tool is called again. In the present implementation, the adjustable settings include the evaluation stride, the number of correction passes, and the spike-fix window used in the second pass. The ControlAgent is therefore not limited to launching the controller once; it also evaluates the quality of the generated schedule and, when needed, modifies the controller settings in response to the observed depth profile. To keep the workflow bounded and reproducible, at most two total runs, including the initial control execution, are allowed. The ControlAgent is implemented using a locally deployed instruction-tuned model via the Ollama inference framework [44] and is configured with a tool-use fine-tuned variant (**llama3-groq-tool-use:8b**) [17, 43] to ensure reliable structured tool invocation. The agent is provided with a system prompt that defines its operational role, the available tool set, and the expected response format. The full system prompt and a sample user prompt are provided in B.2.

The ControlAgent interacts with the control pipeline through two tools:

- **list_tracks**: Returns a compact list of available scan-pattern configuration files as JSON records containing track type, geometry parameters, and hatch spacing. This tool is used to verify that a requested scan strategy exists and to resolve the intended configuration before control is attempted.
- **run_control**: Executes the complete feedforward control pipeline for a specified scan pattern. Inputs include the track name and controller settings. The tool runs the uncontrolled baseline at constant power, applies the analytical inverse controller, and saves all output figures to the designated directory. A structured JSON summary is returned containing uncontrolled and controlled depth statistics, spike-related metrics, percentage reduction in depth standard deviation, and file paths to all saved plots. These returned quantities are then used by the agent to determine whether the generated schedule is acceptable or whether controller settings should be refined and the control tool called once more.

Upon receiving a natural-language instruction specifying a scan track type, the ControlAgent first verifies the requested configuration and then executes the control tool with the current controller settings. The returned control summary is subsequently evaluated to determine whether the achieved depth-uniformity metrics are satisfactory. If the run is classified as excellent or good, the result is reported directly. If the run is classified as poor, the agent modifies selected controller settings and performs one additional control run before returning the final plots and statistics. The ControlAgent workflow, from natural-language input to saved plots and depth statistics, is illustrated in Figure 5.

As with CalibAgent, the ControlAgent's value lies not in the numerical computation itself but in the runtime decisions that connect the tools. The same agent handles single-run deterministic control, stochastic ensemble analysis, and combined calibration-to-control pipelines, dispatching to different tool sequences based on a natural-language interpretation of the user's objective. A conventional script would require a separate entry point or command-line flag for each workflow variant. When a control run is classified as needing improvement, the agent does not simply repeat the computation; it reads the suggested adjustments from the evaluation output and applies them before rerunning. Every tool call, returned diagnostic, and branching decision is recorded in a machine-readable interaction trace (D).

3 Results and Discussion

3.1 Stochastic Calibration Outcomes

Table 3 reports summary statistics for melt pool width W and depth D. The predicted mean values remain close to the experimental means, confirming that the calibration preserves the central tendency of the measurements. The predicted standard deviations are slightly larger than the experimental values because they include both the inferred parameter uncertainty and the randomness introduced during predictive sampling.

Figure 6 d) compares the experimental and predicted distributions in the joint W–D plane. The predicted cloud envelopes the experimental measurements and reproduces the observed elliptical variability pattern, confirming that the calibrated model captures the coupled fluctuations between width and depth. Figure 6 a) and b) shows the corresponding distributions for W and D separately. The KDE-smoothed experimental distributions are closely matched in location by the calibrated predictive distributions, with slight over-dispersion in the predicted spread consistent with the additional uncertainty carried from the inferred parameter posteriors. Together, Figures 6 a), b) and d) and Table 3 confirm that stochastic calibration produces a predictive distribution that matches the experimental data at both the summary-statistics level and the joint-distribution level.

Table 3: Experimental vs. predicted melt pool statistics summary.

Distribution	**Width, W (μm)**	**Depth, D (μm)**
Experimental	141.3 ± 2.7	122.4 ± 2.1
Predicted	142.4 ± 4.6	123.2 ± 4.0

Figure 6 c) validates the calibrated analytical model against an experimental cross-section. The left panel shows the predicted temperature field in the y–z plane at the nominal process condition, with the red contour marking the T_{melt} isotherm that defines the predicted melt pool boundary. The right panel shows the corresponding optical micrograph of the etched single-track cross-section from the NIST dataset. The predicted melt pool shape and depth are in close visual agreement with the experimental fusion zone,

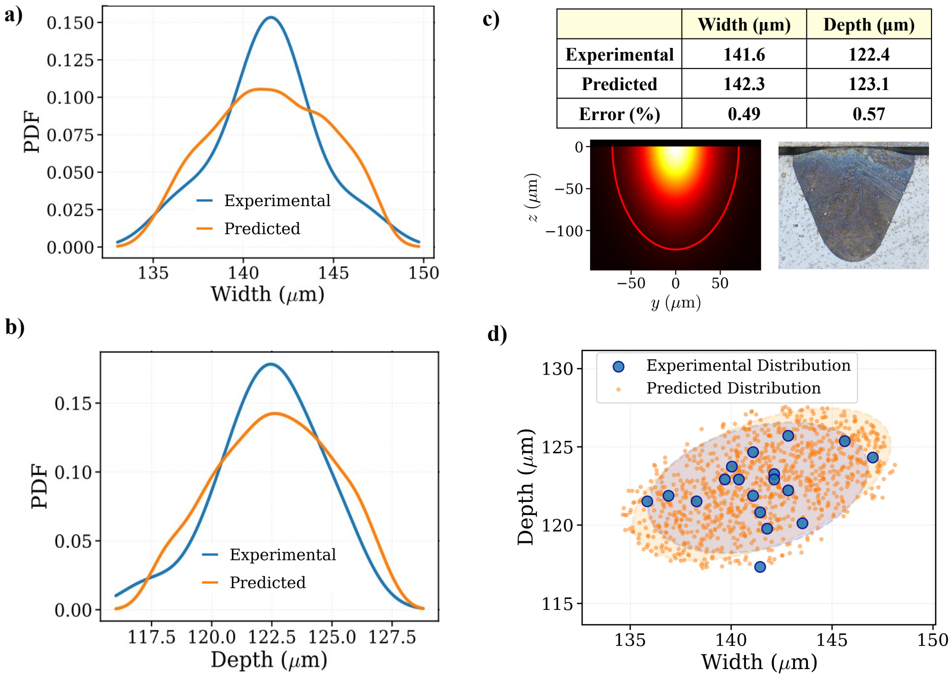

	Width (μm)	Depth (μm)
Experimental	141.6	122.4
Predicted	142.3	123.1
Error (%)	0.49	0.57

Fig. 6: Comparison of experimental and analytically predicted melt-pool geometry for the calibrated IN718 case: (a) width distribution, (b) depth distribution, (c) representative predicted thermal cross-section and experimental melt-pool cross-section with summary of width, depth, and relative error, and (d) joint width–depth distribution.

confirming that the calibrated volumetric Gaussian model captures the essential melt pool geometry at this process condition. More detailed statistics are discussed in C.

3.2 Control Outcomes for Single Layer Multi-track Cases

For the single-layer multi-track cases, uncertainty propagation is demonstrated by performing 10 stochastic control runs for each scan strategy using samples drawn from the calibrated parameter distributions. The depth and power-schedule figures report this distribution where the central curve represents the mean response and the surrounding spread indicates the variability across realizations.

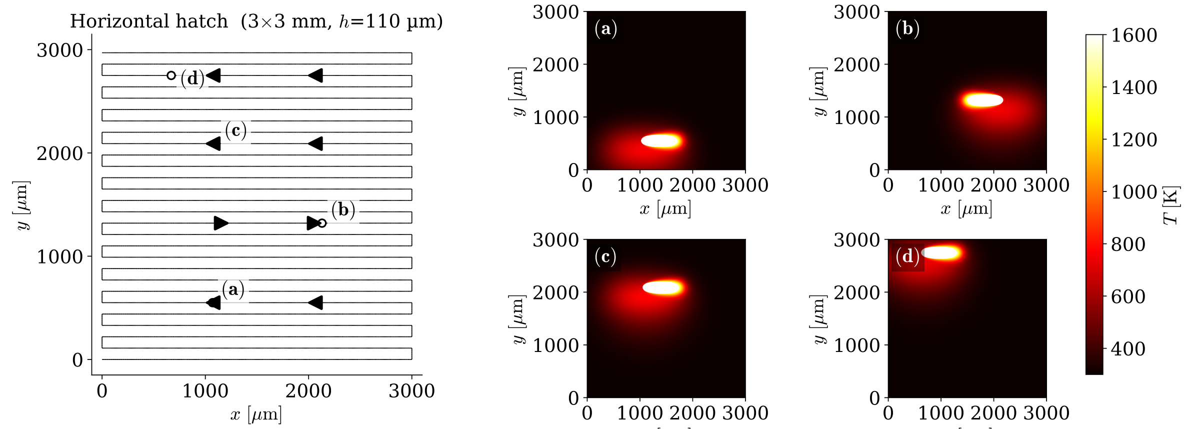

Fig. 7: Horizontal hatch scan strategy over a 3 × 3 mm region with hatch spacing $h = 110$ μm, showing the multi-track path layout and representative temperature snapshots along the scan. For visualization, the temperature colormap is clipped at the melting temperature to indicate the melt pool.

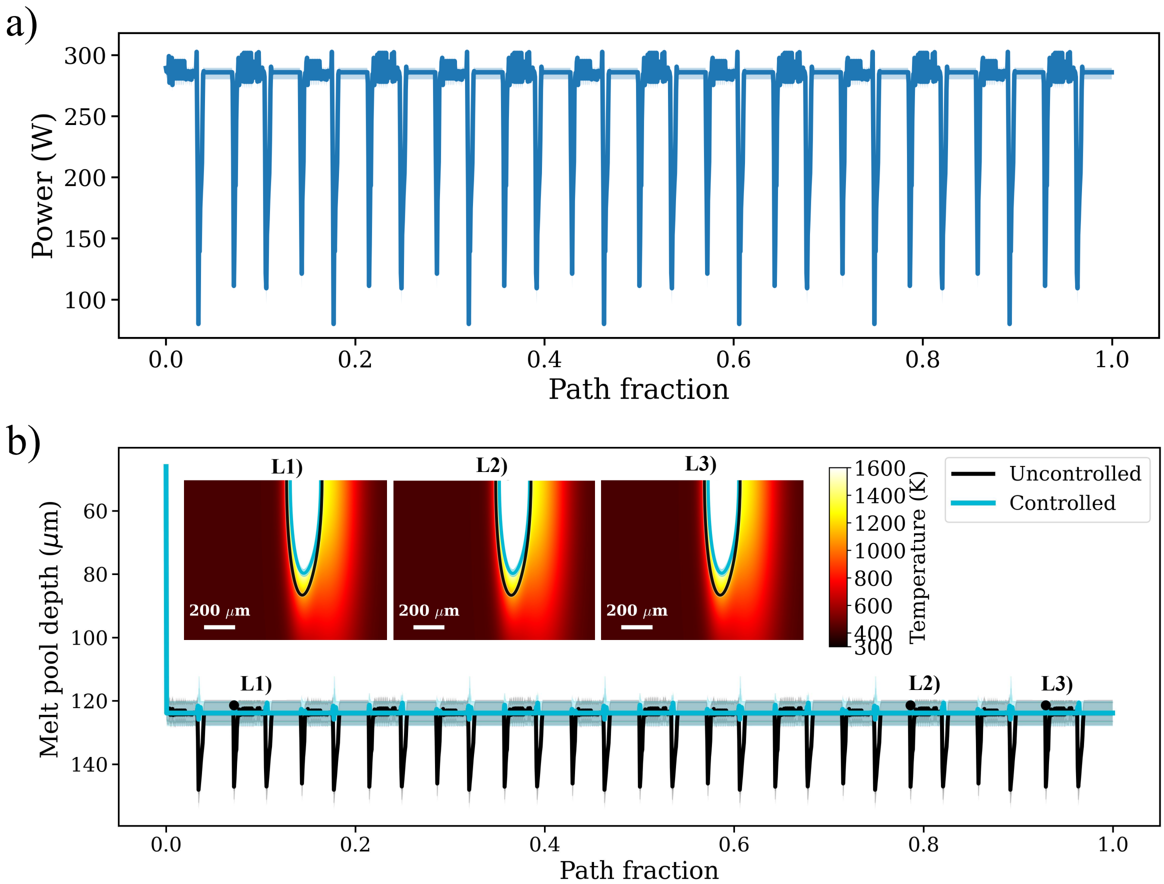

Fig. 8: Horizontal hatch control results: (a) laser power schedule computed by the feedforward controller, which remains near the nominal level along straight scan segments and applies sharp reductions at each track turnaround to offset local heat accumulation; (b) corresponding melt-pool depth response, where the uncontrolled case shows periodic depth spikes at repeated track transitions while the controlled case maintains nearly uniform depth along the full path. Points L1–L3 indicate representative cross-sections showing improved consistency of melt-pool geometry under control. For visualization, the temperature colormap is clipped at the melting temperature to indicate the melt pool.

3.2.1 Horizontal Cross-hatching scan path

Figure 7 shows the horizontal hatch scan path over the 3×3 mm domain, while Figure 8 b) compares the melt-pool depth histories for the uncontrolled and controlled cases. In the uncontrolled case, the laser power is kept constant at $P = 285$ W throughout the scan.

At the very beginning of the first track, the melt pool shows a short startup transient. This happens because the substrate is initially colder, so heat is conducted away more quickly, and the melt pool is temporarily shallower. After this brief initial drop, the response settles into a repeatable pattern. In the present case, the melt pool reaches this quasi-steady regime at a path fraction of about 0.011.

Once the startup transient has passed, the uncontrolled horizontal hatch exhibits a nearly constant baseline depth of 121.98 ± 3.17 μm across the stochastic runs. However, this baseline is interrupted repeatedly by sharp depth spikes at the track turnarounds. The mean spike peak depth is 144.79 ± 4.86 μm, which corresponds to an average overshoot of 22.81 ± 1.69 μm above the baseline depth. Physically, these spikes appear because each turnaround combines two effects: the laser slows down and reverses direction, and the nearby material has already been preheated by the previously scanned adjacent track. Together, these effects locally increase the effective energy input and drive the melt pool deeper. Similar behavior has been reported in prior studies, which show that residual heat accumulation and scan-direction changes near turn regions can produce pronounced local melt-pool deepening in LPBF [33, 52, 65].

The analytical inverse controller removes this behavior very effectively. As seen in Figure 8 b), the controlled depth remains centered around the same desired level, with a baseline depth of 121.995 ± 3.10 μm. In other words, the controller does not shift the intended melt depth; it mainly suppresses the unwanted fluctuations around it. The steady-state standard deviation is reduced from 5.24 ± 0.38 μm in the uncontrolled case to only 0.0296 ± 0.0669 μm under control, corresponding to a $99.49\% \pm 1.15\%$ reduction. At the same time, the turnaround spikes are fully removed in the detected statistics, with the mean spike overshoot reduced from 22.81 ± 1.69 μm to 0 μm. The statistics are summarized in Table 4.

Figure 8 a) explains how this is achieved. The controller keeps the power close to the nominal level over most of each straight track, but applies sharp reductions near the turnarounds, where overheating would otherwise occur. This behavior is physically consistent with the depth plot: the large uncontrolled spikes coincide with the locations where the controller demands the strongest local power reduction. Overall, the horizontal hatch case shows that the proposed control law can preserve the desired melt-pool depth while almost completely removing the repeated transition-induced depth excursions.

Table 4: Horizontal hatch depth statistics after the initial startup transient. Reported values are mean ± standard deviation over 10 stochastic runs.

Metric	**Uncontrolled**	**Controlled**
Baseline depth (μm)	121.98 ± 3.17	121.995 ± 3.10
Standard deviation (μm)	5.24 ± 0.38	0.03 ± 0.067
Mean spike peak depth (μm)	144.79 ± 4.86	121.99 ± 3.10
Mean spike overshoot (μm)	22.81 ± 1.69	0.00 ± 0.00
Steady-state standard deviation reduction: 99.49% ± 1.15%		

3.2.2 45° Cross-hatching scan path

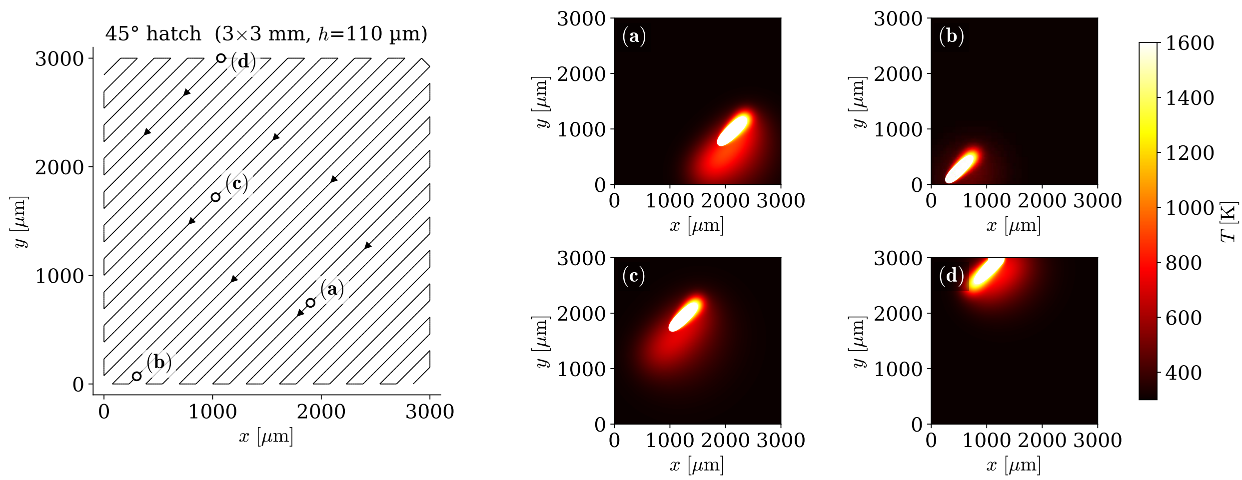

Fig. 9: 45° cross-hatch scan strategy over a 3×3 mm domain with hatch spacing $h = 110$ μm, showing the multi-track diagonal path layout and representative surface temperature snapshots at four path locations along the scan. For visualization, the temperature colormap is clipped at the melting temperature to indicate the melt pool.

Figure 9 shows the 45° hatch layout, and Figure 10 b) compares the uncontrolled and controlled melt-pool depth histories. Unlike the horizontal case, the diagonal tracks are clipped by the square boundary, so their lengths are no longer uniform. Tracks are shortest near the corners and longest near the middle of the domain. As a result the turnaround events are not evenly distributed along the scan history. This is why the uncontrolled depth curve show a non-uniform pattern of spikes, with a visibly denser cluster of excursions near the beginning and especially near the end of the scan.

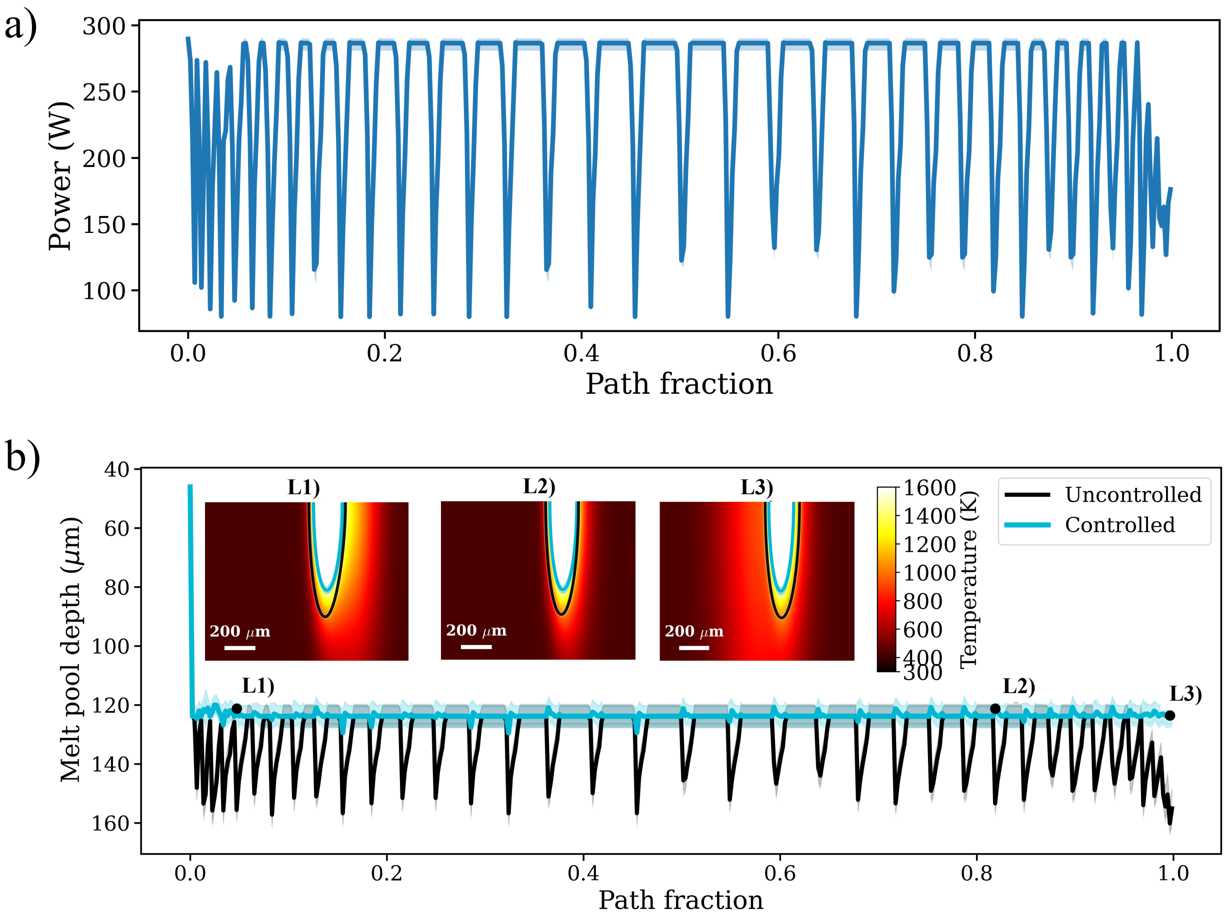

Fig. 10: 45° cross-hatch control results: (a) laser power schedule computed by the feedforward controller, with unevenly spaced reductions reflecting the non-uniform diagonal track lengths and stronger adjustments near clusters of short scan vectors; (b) corresponding melt-pool depth response, where the uncontrolled case exhibits irregular depth spikes due to variable track lengths and non-uniform residual heat accumulation, while the controlled case maintains near-uniform depth along the path. Points L1–L3 indicate representative cross-sections showing improved consistency of melt-pool geometry under control. For visualization, the temperature colormap is clipped at the melting temperature to indicate the melt pool.

After the startup transient, the uncontrolled case reaches a baseline depth of 121.92 ± 3.17 μm, but with much larger variation than in the horizontal hatch. The steady-state standard deviation is 9.40 ± 0.61 μm, and the mean spike overshoot is 33.37 ± 1.68 μm. Under control, the baseline depth is preserved at 121.99 ± 3.10 μm, while the standard deviation is reduced to 0.124 ± 0.225 μm. The mean spike overshoot is reduced to 0.73 ± 2.30 μm. This corresponds to a $98.77\% \pm 2.18\%$ reduction in steady-state standard deviation. A summary of these statistics is available in Table 5.

Figure 10 a) shows that the required power schedule is also less uniform than in the horizontal case. Larger local power reductions appear near the unevenly spaced turn regions, especially near the end of the scan where the shortest diagonal tracks are clustered. These results show that the controller remains effective even when the path geometry introduces variable track length and non-uniform residual heat buildup.

3.2.3 Concentric Triangular Scan Path

Figure 11 shows the concentric triangular scan layout, and Figure 12 b) compares the uncontrolled and controlled melt-pool depth histories. This case is more challenging than the hatch scans because the path moves continuously inward. Each new contour is shorter than the previous one and is deposited inside material that has already been heated by all outer contours. As a result, heat accumulates toward the triangle center, and the uncontrolled depth rises more strongly near the end of the scan.

Table 5: 45° hatch depth statistics after the initial startup transient. Reported values are mean ± standard deviation over 10 stochastic runs.

Metric	Uncontrolled	Controlled
Baseline depth (μm)	121.92 ± 3.17	121.99 ± 3.10
Standard deviation (μm)	9.40 ± 0.61	0.124 ± 0.225
Mean spike peak depth (μm)	155.28 ± 4.67	122.72 ± 3.78
Mean spike overshoot (μm)	33.37 ± 1.68	0.73 ± 2.30
Steady-state standard deviation reduction: 98.77% ± 2.18%		

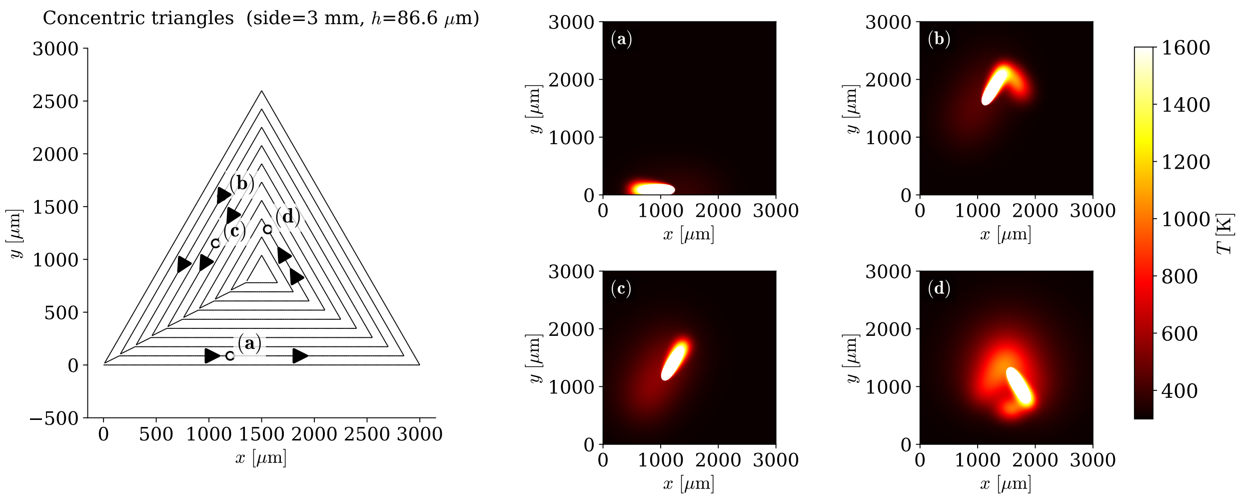

Fig. 11: Concentric equilateral triangle scan strategy with outer side length 3 mm and contour spacing $h = 86.6$ μm, showing the inward-spiraling path layout and representative surface temperature snapshots at four path locations along the scan. For visualization, the temperature colormap is clipped at the melting temperature to indicate the melt pool.

The triangle path also introduces repeated corner effects. Each contour contains three sharp 60° turns, which generate localized depth spikes in addition to the inward heat buildup. In the uncontrolled case, the steady-state standard deviation is 8.51 ± 0.54 μm, and the mean spike overshoot reaches 41.43 ± 2.48 μm, which is the largest among the single-layer cases considered here.

Under control, the baseline depth is preserved at 121.995 ± 3.10 μm, while the standard deviation is reduced to 1.43 ± 0.47 μm. The mean spike overshoot is reduced to 10.69 ± 3.19 μm. This corresponds to an $83.43\% \pm 4.39\%$ reduction in steady-state standard deviation. A summary of these statistics is available in Table 6. Figure 12 a) shows that the controller responds by reducing power at the repeated corner events and by lowering the overall power level as the path approaches the hotter inner contours. These results show that the controller remains effective even when both corner concentration and progressive inward heat accumulation are present.

Table 6: Concentric triangle depth statistics after the initial startup transient. Reported values are mean ± standard deviation over 10 stochastic runs.

Metric	Uncontrolled	Controlled
Baseline depth (μm)	121.92 ± 3.17	121.995 ± 3.10
Standard deviation (μm)	8.51 ± 0.54	1.43 ± 0.47
Mean spike peak depth (μm)	163.35 ± 5.26	132.69 ± 5.08
Mean spike overshoot (μm)	41.43 ± 2.48	10.69 ± 3.19
Steady-state standard deviation reduction: 83.43% ± 4.39%		

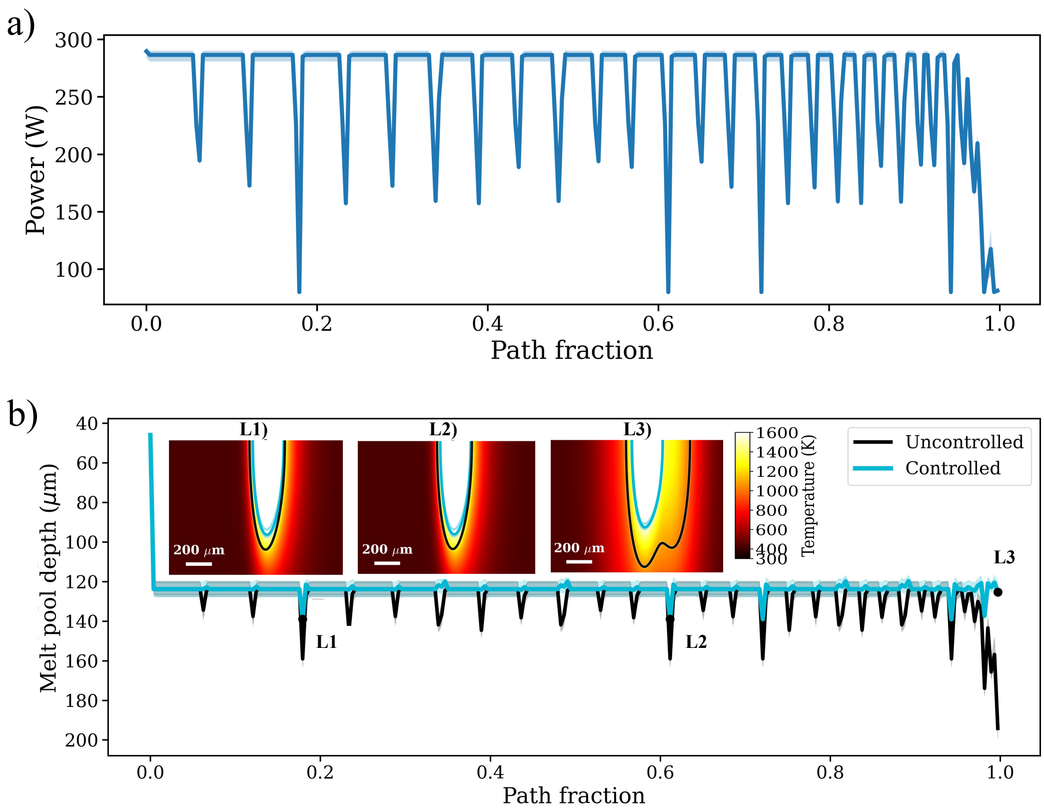

Fig. 12: Concentric triangle control results: (a) laser power schedule computed by the feedforward controller, with sharp reductions at repeated 60° corner events and a progressively lower overall power level as the scan moves inward toward the hotter central region; (b) corresponding melt-pool depth response, where the uncontrolled case exhibits both an inward rise in depth and repeated corner-driven spikes, while the controlled case maintains a much more uniform depth along the path. Points L1–L3 indicate representative cross-sections showing improved consistency of melt-pool geometry under control. For visualization, the temperature colormap is clipped at the melting temperature to indicate the melt pool.

3.3 L-shape Multi-layer Multi-track

The L-shape multi-layer case is the most demanding geometry considered in this study. It combines **three sources** of depth non-uniformity at the same time: turnaround spikes within each layer, track-length variation caused by the re-entrant corner, and inter-layer heat accumulation across successive layers. The dimension and track definitions are explained in Figure 13.

The multilayer thermal coupling is modeled directly in the analytical simulation. All scan points from all four layers are assembled into one 3D source history, with each layer placed at its own build height and time offset. As a result, melt-pool depth in a given layer is not determined only by the current scan vector. It is also affected by the residual heat deposited in the earlier layers below it. In this way, the lower layers thermally preheat the upper layers, and the full source history is carried into the depth evaluation. The layer-by-layer variation in depth is demonstrated in G.

Figure 14 shows the uncontrolled and controlled depth profiles for all four layers. In the uncontrolled case, periodic spikes are present throughout the scan, and the depth remains non-uniform along the path. These variations arise from local turn events, the change in track length across the L-shape, and the accumulated thermal history inherited from the previously deposited layers. The controlled case is much flatter. The target depth is maintained closely across the full path and across all four layers.

The improvement is summarized in Table 7. The baseline depth is preserved, while the standard deviation is reduced from 5.01 μm to 0.52 μm. The mean spike overshoot is reduced from 8.93 μm to 1.54 μm, and the spike count drops from 27 to 1. This corresponds to an 89.69% reduction in standard deviation and an 82.81% reduction in mean spike overshoot.

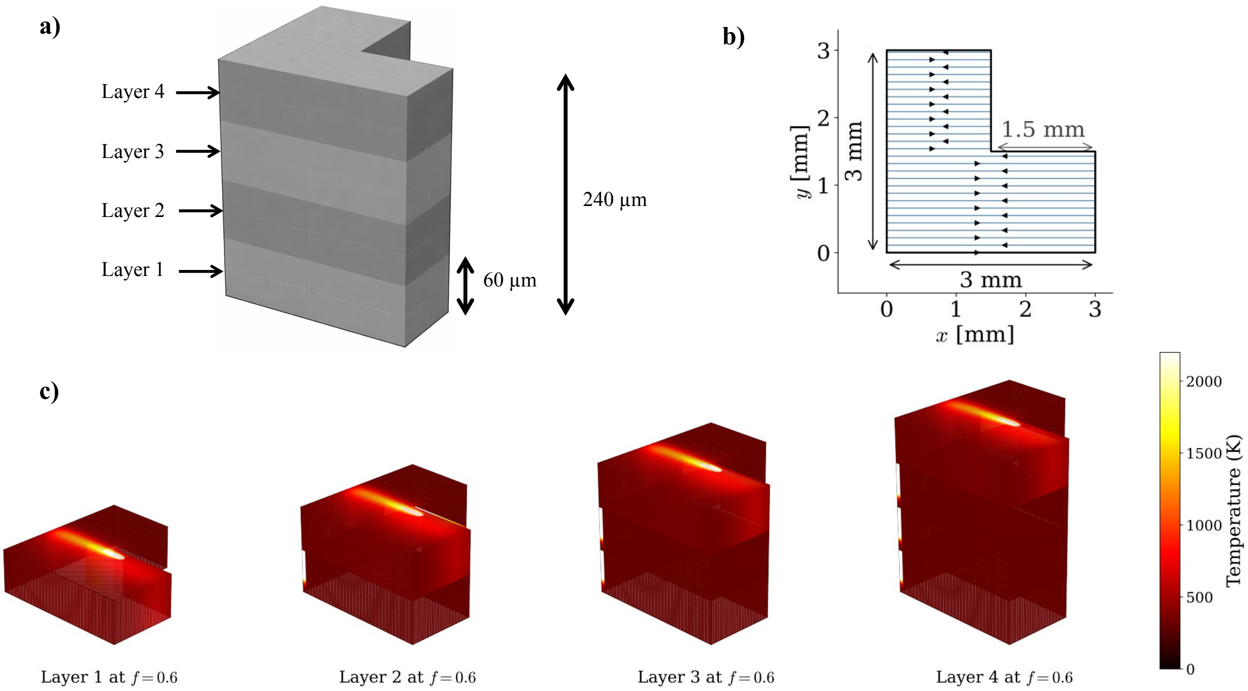

Fig. 13: L-shaped four-layer geometry and thermal snapshots used in the multi-layer study. (a) Schematic of the deposited L-shaped part with four layers, total height 240 μm, and layer thickness 60 μm. (b) L-shaped scan layout and laser travel direction over the 3×3 mm domain. (c) Diagonal temperature snapshots at path fraction $f = 0.6$ for Layers 1–4, illustrating the thermal state at the same path-fraction location in each layer.

These results show that the same analytically calibrated framework remains effective for a non-convex multi-layer geometry with thermal carryover between layers. No retraining is required between layers. Instead, the controller is applied sequentially, and each new layer is evaluated using the corrected thermal and power history produced by the earlier layers. Figure 15 shows the corresponding layer-wise power schedules, where the controller applies a spatially varying laser power along each L-shaped scan path to offset local thermal-history effects and maintain a nearly uniform melt-pool depth.

Table 7: L-shape multilayer depth statistics after the initial startup transient. Reported values are identical across all four controlled layers in the present simulation.

Metric	Uncontrolled	Controlled
Baseline depth (μm)	119.40	119.36
Standard deviation (μm)	5.01	0.52
Mean spike peak depth (μm)	128.34	120.90
Mean spike overshoot (μm)	8.93	1.54
Spike count	27	1
Steady-state standard deviation reduction: 89.69%		
Mean spike overshoot reduction: 82.81%		

3.4 Controller Performance Under Quantized Power Constraints

In practical LPBF systems, laser power cannot always be commanded as a fully continuous variable; instead, hardware and controller constraints may restrict the admissible power values to a discrete set. To assess controller robustness under such actuation limits, the present framework allows the user to specify a quantized power set. The controller then projects each requested power value onto the nearest admissible machine-realizable level. Figure 16 shows the resulting depth-control performance for the 45° hatch case when the

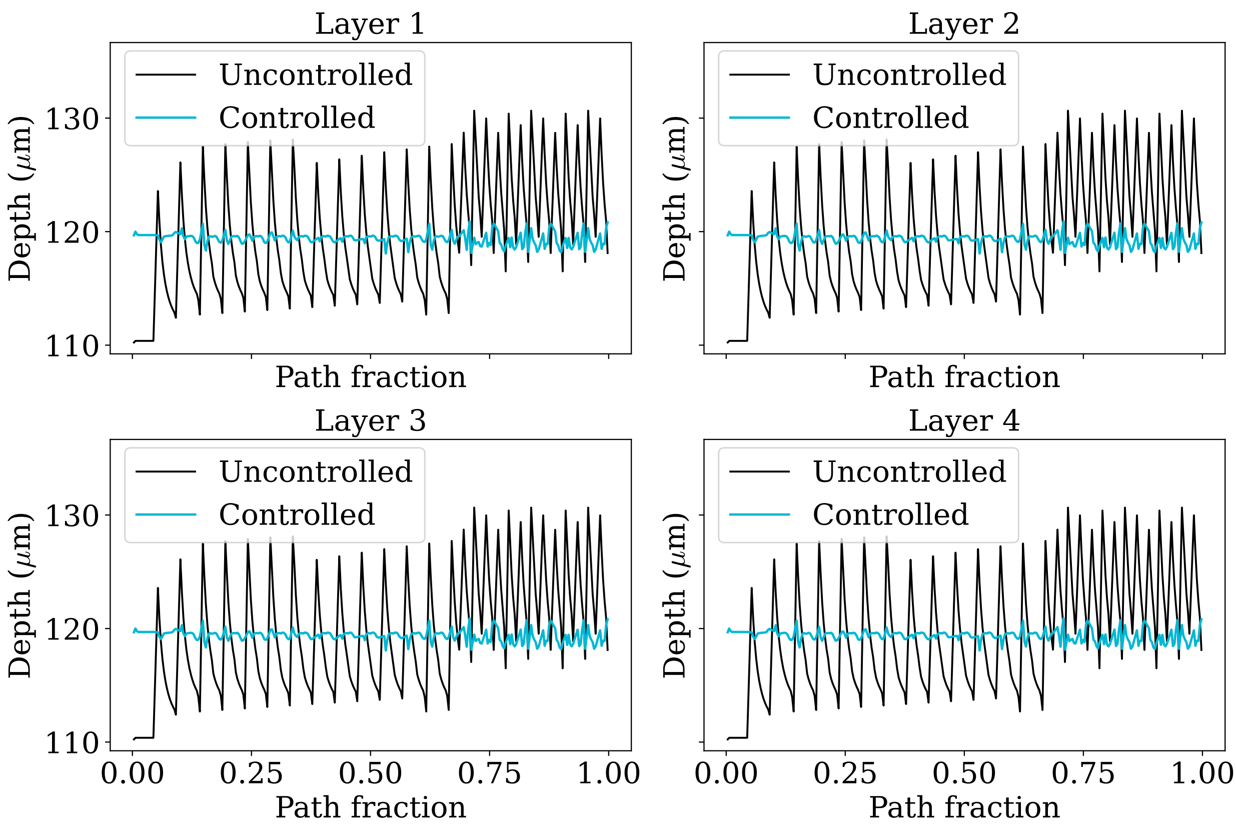

Fig. 14: Layer-by-layer melt pool depth control for the L-shaped multi-layer geometry. Uncontrolled and controlled depth profiles are shown for all four layers, with the target depth indicated by the dashed line.

allowable power values are restricted to 100–300 W in increments of 50 W. The corresponding quantized power schedule is provided in the F.

4 Limitations and Future Work

The present work develops a framework whose primary goal is to establish the computational pipeline for LLM-agent-guided stochastic calibration and feedforward control of LPBF processes, and to demonstrate its effectiveness within a physically grounded analytical modeling environment. Experimental validation of the predicted power schedules against in-situ melt pool measurements on a physical LPBF machine is therefore outside the scope of this paper and is identified as the natural and immediate next step. Future work will focus on closing this loop by implementing the computed power schedules on an instrumented LPBF system and comparing predicted and measured melt pool depth distributions across the scan geometries considered here. Extension to closed-loop feedback control using real-time pyrometry or melt pool imaging, and generalization to additional alloy systems and machine platforms, are also planned as the framework matures toward experimental deployment.

5 Conclusions

This work presents for the first time a unified LLM-agent-guided framework that jointly addresses stochastic calibration and feedforward melt pool depth control for multi-track, multi-layer LPBF processes. The main contributions are as follows:

- **Stochastic Bayesian calibration.** CalibAgent infers full posterior distributions over absorptivity and thermal diffusivity via Bayesian MCMC, reproducing both the mean and scatter of experimental IN718 melt pool dimensions from the NIST AM-Bench dataset (predicted: 142.4 ± 4.6 μm width, 123.2 ± 4.0 μm depth; experimental: 141.3 ± 2.7 μm, 122.4 ± 2.1 μm).

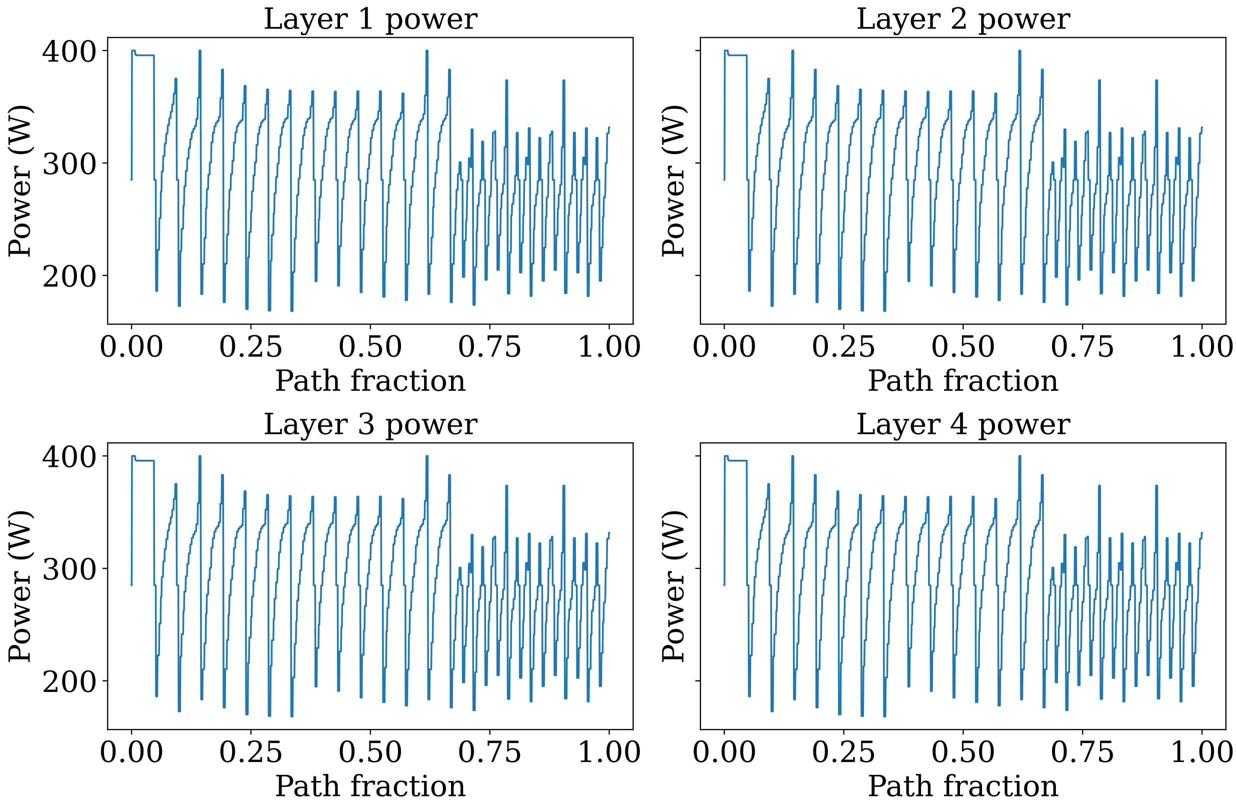

Fig. 15: Controlled laser power schedule for the four-layer L-shape print. The power is plotted against path fraction for each layer, showing the spatially varying input applied by the controller along the L-shaped scan trajectory to compensate for local thermal-history effects.

- **Analytical thermal solution for multi-layer and multi-track builds.** The article provides an analytical solution for multi-layer and multi-track laser manufacturing process. The solution can be extended to other laser-based processes for quick estimate of thermal field.
- **Limited data reliance.** The article shows a pathway of how to collaborate an LLM-agent with very sparse calibration data and produce a reliable computational method for repeatable prediction.
- **Analytical feedforward control.** ControlAgent computes spatially varying laser power schedules via single-step analytical inversion, reducing steady-state depth fluctuations by 83% to 99% and eliminating turnaround spikes of up to 41.4 μm across horizontal, 45°, and concentric triangular scan paths.
- **Multi-layer uniformity and hardware robustness.** For a four-layer L-shaped geometry, the controller maintains near-uniform depth across all layers without retraining and remains effective under discrete quantized power constraints of 50 W increments between 100 and 300 W.
- **Cross-Material generalizability:** The same agent workflow can be applied to different materials and process parameters without retraining, as supported by the IN625 calibration results in Table 9. More generally, the framework can be extended to new materials and process conditions without additional training by supplying the corresponding material-property file and experimental dataset.

Taken together, these contributions demonstrate a coherent and extensible agent-driven framework that unifies stochastic calibration with analytically grounded melt-pool control for complex LPBF builds. By combining Bayesian inference, closed-form thermal modeling, and single-step power-schedule inversion, the approach offers both scientific interpretability and practical robustness across geometries, scan strategies, and materials. Beyond the specific results shown here, the framework establishes a generalizable pathway for integrating LLM agents with physics-based models to accelerate process understanding, reduce experimental burden, and enable adaptive, data-efficient control in metal additive manufacturing

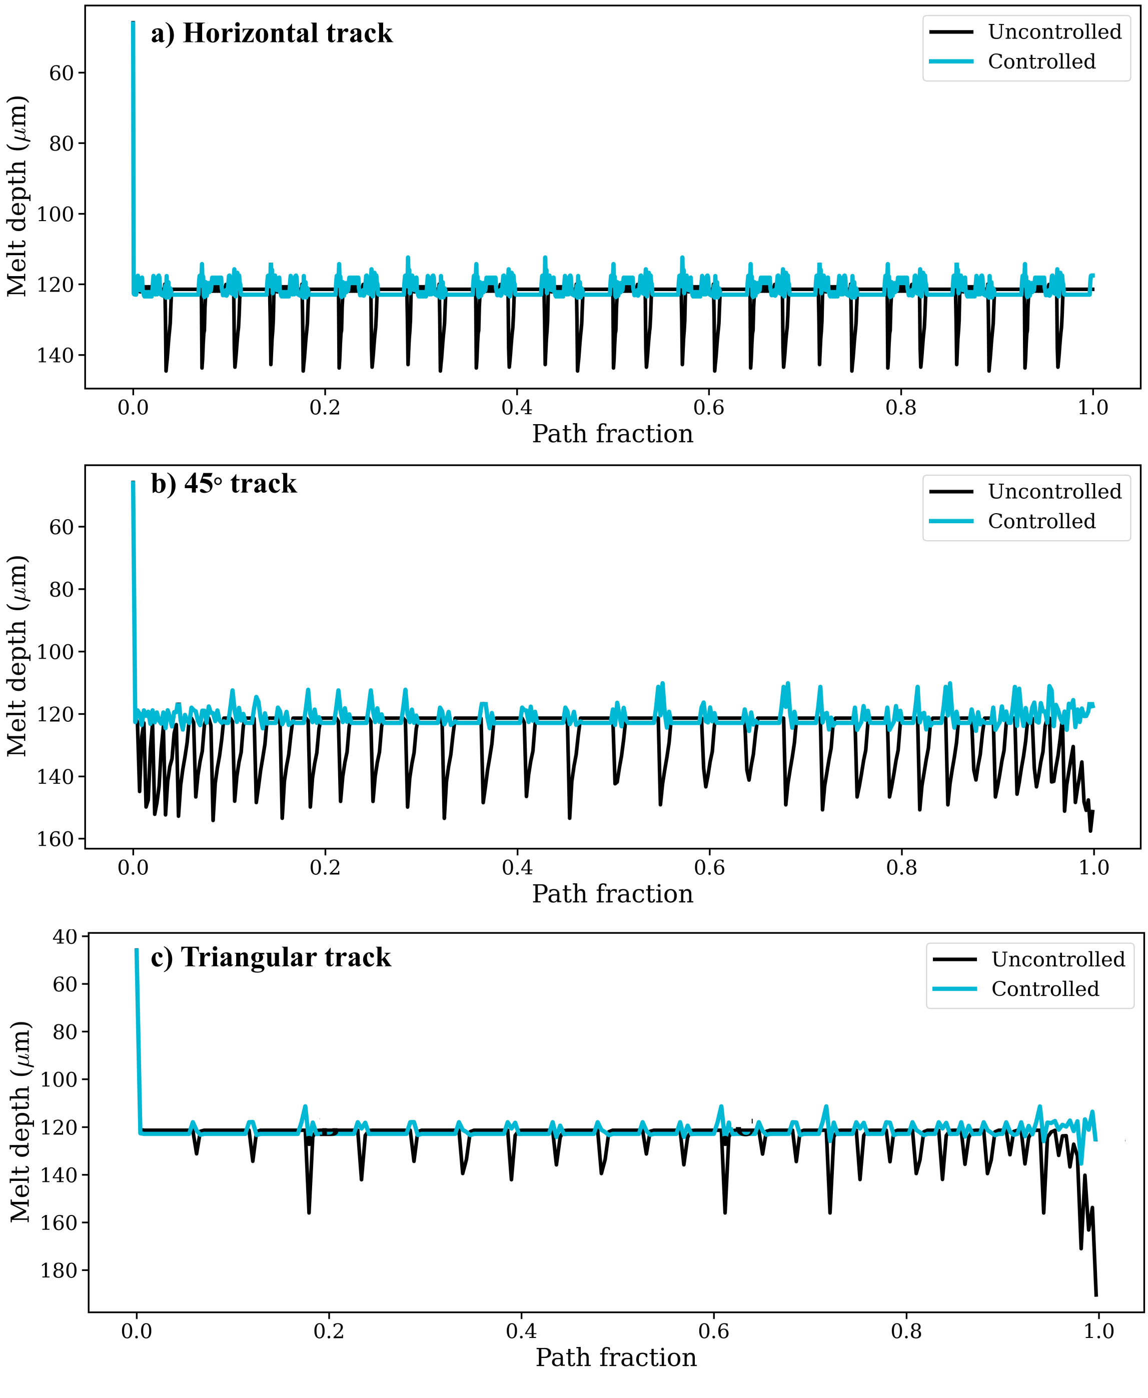

Fig. 16: Melt pool depth histories under quantized power constraints for the (a) horizontal, (b) 45°, and (c) concentric triangular scan paths. Rather than commanding a continuous power value, the controller is restricted to machine-realizable discrete power levels between 100 W and 300 W in increments of 50 W. Despite this constraint, the controller continues to suppress melt pool depth excursions effectively across all three geometries.

Code Availability

The code accompanying this paper is publicly available at github.com/Apurba-Sarker/agentic-stochastic-lpbf-control.

6 Acknowledgments

S. Saha gratefully acknowledges start-up funding by Kevin T. Crofton Department of Aerospace and Ocean Engineering, Virginia Tech, and Professor Ella M Atkins Faculty Fellowship.

7 Appendix

A Description of Build Cases and Scan Strategies

Four scan geometries of increasing complexity are used to evaluate the controller, all applied to IN718 at P = 285 W and V = 960 mm/s.

The Horizontal hatch is a 3×3 mm square domain scanned with 27 bidirectional parallel tracks at h = 110 μm. Tracks alternate direction in a serpentine pattern, producing 54 turnaround events where residual preheat and scan deceleration combine to drive periodic depth spikes. The mid-track interior reaches a quasi-steady depth, making this the baseline geometry for evaluating turnaround suppression.

The 45° cross-hatch has an identical domain and hatch spacing as the horizontal case, but tracks are oriented at 45° to the x-axis and clipped to the square boundary. Oblique track terminations produce asymmetric turnaround geometries and varying track lengths across the domain, testing the controller under non-axis-aligned thermal gradients.

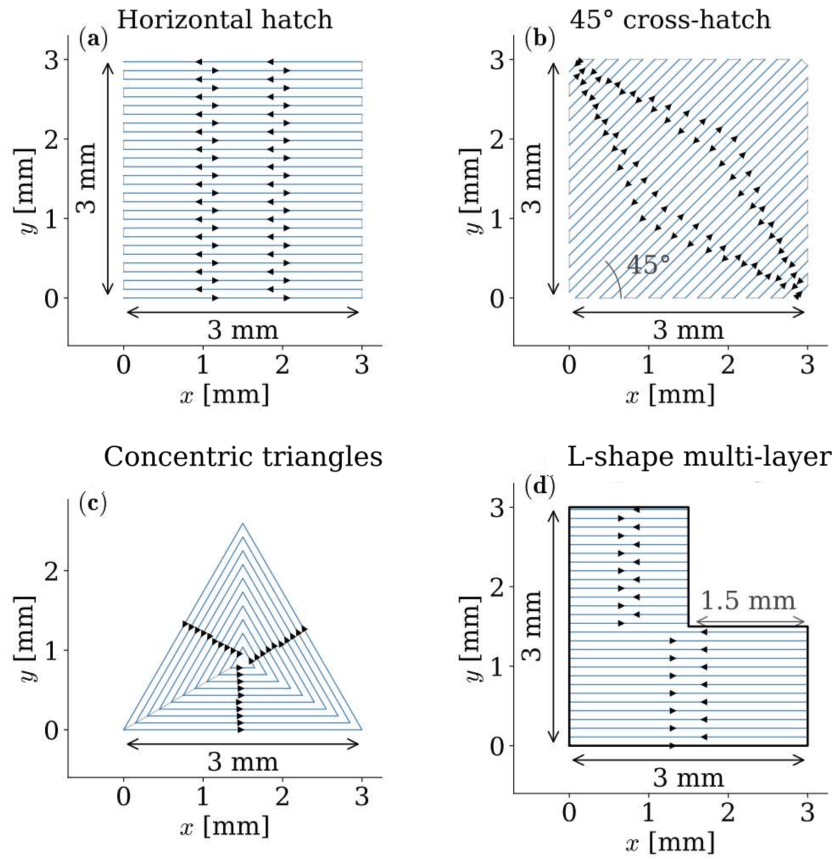

Fig. 17: Scan strategies used to demonstrate multi-track and multi-layer melt pool depth control. (a) Horizontal hatch: 27 bidirectional parallel tracks over a 3×3 mm domain with hatch spacing h = 110 µm. (b) 45° cross-hatch: diagonal parallel tracks oriented at 45° to the x-axis over the same domain and hatch spacing. (c) Concentric triangles: inward-spiraling equilateral triangle contours over a 3 mm side-length domain with h = 86.6 µm. (d) L-shaped multi-layer geometry: horizontal hatch scan repeated over successive build layers across an L-shaped cross-section (3×3 mm with a 1.5×1.5 mm upper-right cutout, h = 110 µm). Arrows indicate laser scan direction.

The concentric triangle is a inward-spiraling equilateral triangle contours starting from a 3 mm side-length outer triangle, with contour spacing h = 86.6 μm. Each contour introduces three sharp 60° corners with localized preheat accumulation, and the cumulative thermal history from outer contours progressively elevates the baseline temperature inward, producing a monotonically increasing depth trend under constant power.

L-shaped multi-layer. A non-convex L-shaped cross-section (3×3 mm with a 1.5×1.5 mm upper-right cutout) scanned with a horizontal hatch at h = 110 μm, repeated over four layers at 60 μm layer thickness. Track lengths vary abruptly at the re-entrant corner, and inter-layer preheating causes melt pool depth to grow with each successive layer under constant power. This is the most demanding case, combining geometric discontinuity, track-length variation, and cumulative multi-layer heating.

B System Prompts and User Prompts

B.1 CalibAgent

CalibAgent Prompt Structure

Role. CalibAgent is defined as a Bayesian calibration assistant for LPBF thermal models. The prompt instructs it to interpret a user request, inspect the dataset, launch stochastic calibration, read the returned diagnostics, and decide whether the result should be accepted or rerun.

Tool interface.

- **`list_cases`**: dataset-query tool with two functions:
  - list available calibration cases
  - return metadata for a selected case, including dataset size and recommended MCMC settings
- **`calibrate_case`**: calibration tool with two functions:
  - run Bayesian MCMC calibration for the selected case
  - evaluate the returned calibration diagnostics and determine whether to accept, warn, or rerun

Workflow constraints.

1. Inspect the dataset first if the requested case is incomplete.
2. Inspect the selected case before calibration.
3. Use the returned MCMC settings for the calibration run.
4. Always evaluate the returned diagnostics after calibration.
5. Allow at most one stronger rerun if quality is insufficient.

Decision logic. The prompt instructs the agent to use returned quantities such as dataset size, acceptance rate, and relative posterior spread to classify the result as `accept`, `warn`, or `rerun_more_steps`.

Strict rules.

- Never invent dataset cases or calibration outputs.
- Always inspect the case before calibration.
- Always evaluate the calibration after it completes.
- Keep the workflow bounded to one adaptive rerun.

Final report. The response must summarize the selected case, number of data points, chosen MCMC settings, fitted stochastic parameters, acceptance rate, and final calibration-quality verdict.

Example User Request

Run stochastic calibration for IN718 with:

- laser power $P = 285$ W
- scan speed $V = 960$ mm/s
- spot diameter = 80 μm

Inspect the case first, use the recommended MCMC settings, evaluate the returned calibration quality, and rerun only if stronger settings are needed.

B.2 ControlAgent

ControlAgent Prompt Structure

Role. ControlAgent is defined as an LPBF process-control assistant. The prompt instructs it to interpret a requested scan strategy, load the corresponding track configuration, run the analytical inverse controller, interpret the returned control-quality metrics, and decide whether the result should be accepted directly or improved through one bounded adaptive rerun.

Tool interface.

- **`list_tracks`**: track-query tool with two functions:
  - list available scan-track configurations
  - return the metadata needed to identify and validate the requested track
- **`run_control`**: control-execution tool with two functions:
  - run the analytical control workflow for the selected track, including baseline and controlled depth evaluation
  - evaluate the returned control diagnostics and determine whether the result should be accepted or rerun with adjusted controller settings

Workflow constraints.

1. Inspect the available scan tracks first if the requested geometry is ambiguous.
2. Validate the selected track before launching control.
3. Run the control workflow only after track validation has been completed.
4. Always evaluate the returned control diagnostics after execution.
5. Allow at most one adaptive rerun if the returned quality indicates that improvement is needed.

Decision logic. The prompt instructs the agent to use returned quantities such as uncontrolled and controlled depth standard deviation, improvement percentage, residual spike behavior, and suggested controller-setting adjustments to classify the result. The final action is selected from prompt-defined outcomes such as `accept`, `report_warning`, or `adjust_and_rerun`.

Strict rules.

- Never invent scan geometries, controller settings, or output files.
- Always validate the requested track before running control.
- Always evaluate the returned control quality after execution.
- Keep the workflow bounded to one adaptive rerun.

Final report. The response must summarize the selected scan track, uncontrolled and controlled melt-pool depth statistics, percentage improvement, and the final control-quality verdict. When uncertainty-aware or stochastic control is requested, the response must also summarize the variation across realizations.

Example User Request

Run the LPBF depth-control workflow for the **45-degree hatch** track with cross-sections enabled. Validate the track first, run the controller, evaluate the returned control quality, and rerun only if the returned diagnostics indicate that controller settings should be adjusted.

C Calibration Accuracy

Calibration accuracy is evaluated at the single-track condition (IN718, $P = 285$ W, $V = 960$ mm/s, $D_{4\sigma} = 80$ μm) by verifying MCMC convergence and by examining the inferred posterior distributions over the stochastic hyperparameters $\theta = (\mu_\eta, \mu_\alpha, \sigma_\eta, \sigma_\alpha)$. Post–burn-in traces (steps $12,000$–$30,000$) are shown in Figure 18. Bounded, stationary fluctuations are observed after burn-in, and persistent drift is not observed, which indicates that sampling is performed in a statistically stable regime. Opposing motion between μ_η and μ_α is observed and is interpreted as the expected compensatory relationship in conduction-dominated models: increased absorptivity is offset by increased effective heat diffusion to preserve similar melt-pool dimensions, producing the negative correlation reflected in the posterior.

Marginal posterior distributions are summarized in Figure 19. Concentration within physically admissible ranges is observed (e.g., $\sigma_\eta > 0$, $\sigma_\alpha > 0$), and heavy tails or boundary sticking are not observed. The

Table 8: Multi-chain MCMC convergence diagnostics for the single-track IN718 condition (P = 285 W, V = 960 mm/s, $D_{4\sigma}$ = 80 μm). Four independent chains were run for 30,000 steps each with 12,000 burn-in steps, leaving 23,000 post-burn-in samples per chain. Convergence is assessed using the Gelman–Rubin statistic $\hat{R}$ and effective sample size (ESS).

Parameter	$\hat{\mathbf{R}}$	**ESS**
μ_η	1.0970	103
μ_α	1.0194	148
σ_η	1.0117	177
σ_α	1.0122	191

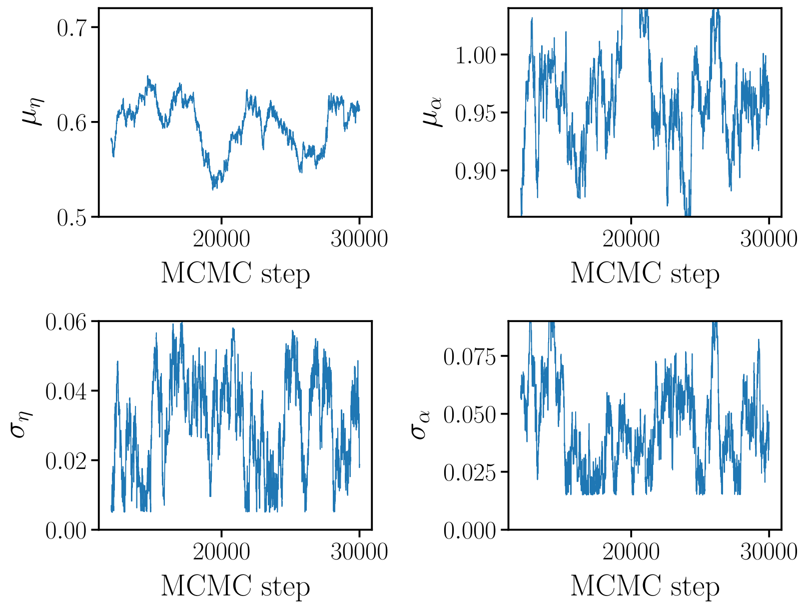

Fig. 18: Post–burn-in MCMC traces for the inferred hyperparameters $(\mu_\eta, \mu_\alpha, \sigma_\eta, \sigma_\alpha)$. Bounded, stationary fluctuations are observed after burn-in, indicating convergence. Opposing motion of μ_η and μ_α is consistent with the absorptivity–diffusivity tradeoff in conduction-based heat-source models.

inferred uncertainty levels $(\sigma_\eta, \sigma_\alpha)$ are interpreted as the variability required to reproduce the experimentally observed scatter in melt-pool width and depth, while the means (μ_η, μ_α) represent the central tendency of the calibrated parameterization. Collectively, these diagnostics indicate that a stable and interpretable posterior is obtained and is suitable for downstream posterior predictive propagation and uncertainty-aware feedforward scheduling.

The calibration methodology has been tested for IN625 as well as shown in 9.

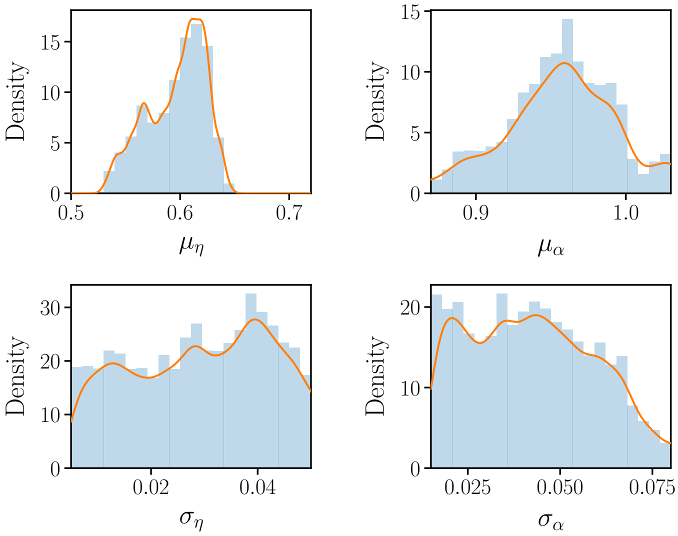

Fig. 19: Marginal posterior distributions of $(\mu_\eta, \mu_\alpha, \sigma_\eta, \sigma_\alpha)$. Blue bars denote posterior samples and the orange curve denotes a kernel density estimate. Unimodal, bounded distributions are observed, indicating well-behaved inference and quantifying calibration uncertainty for downstream posterior predictive propagation.

Table 9: Experimental vs. predicted melt pool statistics for IN625 at $P = 195$ W, $V = 800$ mm/s, and spot diameter 80 μm.

Distribution	**Width, W (μm)**	**Depth, D (μm)**
Experimental	135.19 ± 4.47	107.69 ± 7.22
Predicted	136.48 ± 6.46	108.91 ± 8.40

Algorithm 1 Stochastic Calibration of the Analytical LPBF Model

Require: Data $\mathcal{D} = \{(W_i, D_i)\}_{i=1}^{N}$, depth scaling factor z_{scale}, MCMC steps T, burn-in T_b, ensemble size N_s
Ensure: Posterior chain of $\boldsymbol{\theta}$ and MAP estimate $\boldsymbol{\theta}^*$
Initialize $\boldsymbol{\theta}^{(0)} = [\mu_\eta, \mu_\alpha, \sigma_\eta, \sigma_\alpha]$ from mean-fit calibration
for $t = 1$ to T **do**
  Propose $\boldsymbol{\theta}' = \boldsymbol{\theta}^{(t-1)} + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \Sigma_{\text{prop}})$
  if $\boldsymbol{\theta}'$ is outside prior bounds **then**
    Set $\boldsymbol{\theta}^{(t)} = \boldsymbol{\theta}^{(t-1)}$ and continue
  end if
  Draw N_s samples: $\eta^{(j)} \sim \mathcal{N}(\mu_\eta, \sigma_\eta^2)$ and $\alpha^{(j)} \sim \mathcal{N}(\mu_\alpha, \sigma_\alpha^2)$
  Evaluate $(W^{(j)}, D^{(j)}) = \mathcal{F}(\eta^{(j)}, \alpha^{(j)}; \sigma_m, z_{\text{scale}})$ for $j = 1, \ldots, N_s$
  Compute
  $J(\boldsymbol{\theta}') = w_m L_{\text{mean}} + w_s L_{\text{std}} + w_{kl}(KL_W + KL_D) + w_c L_{\text{corr}} + \lambda_r \|\boldsymbol{\sigma}\|^2$
  Compute $\log p(\boldsymbol{\theta}' \mid \mathcal{D}) = \log p(\boldsymbol{\theta}') - \beta J(\boldsymbol{\theta}')$
  Compute $\log r = \log p(\boldsymbol{\theta}' \mid \mathcal{D}) - \log p(\boldsymbol{\theta}^{(t-1)} \mid \mathcal{D})$
  Draw $u \sim \text{Uniform}(0, 1)$
  if $\log u < \log r$ **then**
    Accept: $\boldsymbol{\theta}^{(t)} = \boldsymbol{\theta}'$
  else
    Reject: $\boldsymbol{\theta}^{(t)} = \boldsymbol{\theta}^{(t-1)}$
  end if
end for
Discard the first T_b samples as burn-in
Estimate posterior summaries from the remaining samples
Set $\boldsymbol{\theta}^* \leftarrow \arg\max_{\boldsymbol{\theta}} \, p(\boldsymbol{\theta} \mid \mathcal{D})$
return posterior chain and $\boldsymbol{\theta}^*$

D Representative Agent Interaction Traces

D.1 Example interaction trace with CalibAgent

User
Calibrate IN718 $P = 285$ W, $V = 960$ mm/s, spot $= 80$ μm.

Agent thought
First verify that the requested case exists in the dataset. If found, retrieve case statistics and choose calibration settings accordingly.

Tool call
```
list_cases(material="IN718")
```

Tool result
Matching IN718 cases found in the dataset, including the requested operating condition.

Agent thought
The case appears available. Next, inspect the case metadata, especially the number of measurements, before selecting the MCMC budget.

Tool call
```
get_case_info(P_W=285, V_mmps=960, material="IN718", spot_um=80)
```

Tool result

```
{
"material": "IN718",
"P_W": 285.0,
"V_mmps": 960.0,
"spot_um": 80.0,
"n_points": 18,
"exp_W_mean_um": 141.34,
"exp_W_std_um": 2.72,
"exp_D_mean_um": 122.39,
"exp_D_std_um": 2.09
}
```


Agent thought
This case has 18 measurements, so use the intermediate calibration budget rather than the longest run. Proceed with Bayesian calibration using the selected settings.

Tool call
```
calibrate_case(P_W=285, V_mmps=960, burn_in=12000, material="IN718",
n_ensemble=30, n_steps=30000, spot_um=80)
```

Tool result
`Acceptance rate:` 0.488
`MAP estimate:` $\mu_\eta = 0.4865$, $\mu_\alpha = 1.1726$
`Posterior mean:` $\mu_\eta = 0.5705 \pm 0.0833$
`Posterior mean:` $\mu_\alpha = 1.0084 \pm 0.1287$
`Posterior spread:` $\sigma_\eta = 0.0307 \pm 0.0118$
`Posterior spread:` $\sigma_\alpha = 0.0367 \pm 0.0154$
`Predicted W:` 141.56 ± 4.60 μm
`Predicted D:` 122.64 ± 4.00 μm

Agent thought
The calibration completed successfully. Now verify whether the result quality is acceptable before returning the parameter set.

Tool call
```
evaluate_calibration()
```

Tool result

```
{
"accept_rate": 0.488,
"n_data_points": 18,
"sim_W_mean_um": 141.56,
"sim_D_mean_um": 122.64
}
```


Agent
The calibration is acceptable. I will retain this case with posterior parameters with predictive statistics consistent with the experimental width and depth measurements.

D.2 Example interaction trace with ControlAgent

User
Run the LPBF depth-control workflow for the 45-degree hatch pattern with cross-sections enabled.

Agent thought
First validate the requested scan track. Then run the control workflow, inspect the returned control-quality metrics, and determine whether the result should be accepted directly or improved through one adaptive rerun.

Tool call
```
list_tracks(mode="info", track_file="45deg")
```

Tool result

```
{
"track_file": "45deg",
"geometry": "45-degree cross-hatch",
"domain": "3 mm × 3 mm",
"status": "available"
}
```


Agent thought
The requested scan track is valid. Proceed with the control workflow for the selected geometry.

Tool call
```
run_control(track_file="45deg", include_xsections=true)
```

Tool result

```
{
"uncontrolled_std_um": 9.40,
"controlled_std_um": 0.124,
"improvement_pct": 98.77,
}
```

Agent
The control result is accepted. The validated 45-degree hatch track was used, the analytical inverse controller reduced the melt-pool depth standard deviation from 9.40 μm to 0.124 μm, and the corresponding reduction in steady-state depth variation is 98.77%.

E Analytical Inverse Controller for Melt Pool Depth Uniformity

The objective of the controller is to maintain a spatially uniform melt pool depth throughout the entire scan. Since the melt boundary is defined as the isotherm at which $T = T_{\mathrm{melt}}$, depth uniformity is equivalent to enforcing

$$T(\mathbf{x}_i,\, z^*,\, t_i) = T_{\mathrm{melt}} \quad \forall\, i \in \{1, 2, \ldots, N_{\mathrm{eval}}\} \tag{11}$$

at every evaluation point along the trajectory, where z^* is a fixed target depth and N_{eval} is the total number of evaluation points. The target depth z^* is selected as the q-th quantile ($q = 0.55$) of the uncontrolled depth distribution, ensuring that the controlled system targets a physically achievable and representative depth value rather than an extreme.

At evaluation index i, a control window $\mathcal{W}_i = [i,\, i + s)$ of stride s is defined, within which the laser power P_{new} is to be determined. All source contributions are partitioned into two groups:

$$T(\mathbf{x}_i, z^*, t_i) = \underbrace{\sum_{j \notin \mathcal{W}_i} G(\mathbf{x}_i, z^*, t_i, \mathbf{x}_j, \tau_j)\, P_j\, \Delta t_j}_{T_{\mathrm{fixed}}(\mathbf{x}_i, z^*)} + \underbrace{P_{\mathrm{new}} \sum_{j \in \mathcal{W}_i} G(\mathbf{x}_i, z^*, t_i, \mathbf{x}_j, \tau_j)\, \Delta t_j}_{P_{\mathrm{new}} \cdot \mathcal{G}_i} + T_0 \tag{12}$$

where $G(\cdot)$ denotes the Gaussian Green's function kernel from Eq. (1) evaluated with unit power, and $\mathcal{G}_i$ is the aggregated kernel sensitivity of the control window at the evaluation point. The term T_{fixed} represents the temperature contribution of all sources outside the control window and is independent of P_{new}. The term $P_{\mathrm{new}} \cdot \mathcal{G}_i$ is the controllable contribution, which is linear in P_{new}.

In practice, T_{fixed} is computed by evaluating the full thermal model with the power of sources in $\mathcal{W}_i$ set to zero, and $T_{\mathrm{ctrl}} = T_{\mathrm{total}} - T_{\mathrm{fixed}}$ is obtained by difference, where T_{total} is evaluated at the current power schedule.

Substituting the control objective $T = T_{\mathrm{melt}}$ into Eq. (12) and solving for P_{new} yields

$$P_{\mathrm{new}} = P_{\mathrm{ref}} \cdot \frac{T_{\mathrm{melt}} - T_{\mathrm{fixed}}(\mathbf{x}_i, z^*)}{T_{\mathrm{ctrl}}(\mathbf{x}_i, z^*)} \tag{13}$$

where P_{ref} is the reference power at which T_{ctrl} is evaluated. This expression constitutes a closed-form, single-step inversion of the thermal model — no iterative optimization or gradient computation is required. The result is clipped to the physically admissible range $[P_{\mathrm{min}},\, P_{\mathrm{max}}]$ and rate-limited by a maximum allowable inter-step power change ΔP_{max} to prevent abrupt transitions that could destabilize the melt pool:

$$P_{\mathrm{new}} \leftarrow \mathrm{clip}(P_{\mathrm{new}},\; P_{\mathrm{prev}} - \Delta P_{\mathrm{max}},\; P_{\mathrm{prev}} + \Delta P_{\mathrm{max}}) \tag{14}$$

The full controller operates in two sequential passes over the trajectory.

The trajectory is traversed from start to end. At each evaluation index i, the power P_{new} is computed via Eq. (13) and applied to the window $\mathcal{W}_i$. The controller then advances by s points and repeats. This pass produces a power schedule $\{P_j\}$ that approximately enforces the depth target at every stride point.

After Pass 1, the full depth evolution is recomputed over the controlled trajectory. Points at which the melt depth satisfies $d_i < z^* - \delta_{\mathrm{spike}}$, where δ_{spike} is a prescribed spike threshold, are identified as residual undershoot locations. For each such point, the inversion of Eq. (13) is applied locally with a narrower window centered on the spike. This correction is iterated up to N_{fix} times until no undershoot points remain or the iteration limit is reached.

F Quantized Power Schedule

Figure 20 shows the quantized power schedules corresponding to Figure 16. For the horizontal hatch (a), the schedule switches cleanly between 250–300 W mid-track and drops to 100–150 W at turnarounds. The 45° hatch (b) shows more frequent transitions due to irregular turn spacing. The triangular path (c) exhibits a progressive decrease in nominal power toward the interior, with corner events superimposed as brief drops.

Algorithm 2 Analytical Inverse Power Controller

Require: Scan trajectory, calibrated thermal model, nominal power P_0, control limits
Ensure: Controlled power schedule and controlled melt-pool depth profile

```
Initialize the power schedule with the nominal value P_0
Predict the uncontrolled melt-pool depth profile along the trajectory
Set the target depth from the uncontrolled profile

Pass 1: inverse power control
for each control location along the trajectory do
    Evaluate the local melt-pool depth error relative to the target
    Compute the required inverse power correction
    Apply power bounds and step-size limits
    Update the local power value
end for

Pass 2: spike mitigation
Recompute the depth profile using the updated power schedule
Detect locations where sharp depth spikes remain
for each detected spike region do
    Define a wider local correction window
    Recompute the local inverse power update over that window
    Replace the local power values with the corrected value
end for
Recompute the final controlled depth profile
return controlled power schedule and controlled depth profile
```

G Details on L-shaped Multi-layer Print

Figure 21 shows the depth increment of each layer relative to Layer 1 under constant power. The monotonic increase from Layer 2 to Layer 4 quantifies the cumulative substrate preheating that the controller in Section 3.3 compensates. Oscillations correspond to turnaround events whose amplitude grows with layer number as inter-track and inter-layer heat accumulation compound.

References

[1] Agazhanov AS, Samoshkin DA, Kozlovskii YM (2019) Thermophysical properties of inconel 718 alloy. Journal of Physics: Conference Series 1382:012175. https://doi.org/10.1088/1742-6596/1382/1/012175

[2] Agrawal AK, Rankouhi B, Thoma DJ (2022) Predictive process mapping for laser powder bed fusion: A review of existing analytical solutions. Current Opinion in Solid State and Materials Science 26(6):101024. https://doi.org/10.1016/j.cossms.2022.101024, URL https://doi.org/10.1016/j.cossms.2022.101024

[3] Agrawal AK, Rankouhi B, Thoma DJ (2022) Predictive process mapping for laser powder bed fusion: A review of existing analytical solutions. Current Opinion in Solid State and Materials Science 26(6):101024. https://doi.org/10.1016/j.cossms.2022.101024

[4] Ansari MJ, Nguyen DS, Park HW (2021) Numerical investigation of thermal behavior and melt pool morphology in multi-track multi-layer selective laser melting of the 316L steel. International Journal of Advanced Manufacturing Technology 116:1977–1991. https://doi.org/10.1007/s00170-020-06360-0

[5] Bissiri PG, Holmes CC, Walker SG (2016) A general framework for updating belief distributions. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 78(5):1103–1130

[6] Carslaw HS, Jaeger JC (1959) Conduction of Heat in Solids. Oxford University Press

[7] Chowdhury S, Nirsanametla Y, Prakash C, et al (2022) Laser powder bed fusion: a state-of-the-art review of the technology, materials, properties & defects, and numerical modelling. Journal of Materials Research and Technology 20:2109–2172. https://doi.org/10.1016/j.jmrt.2022.07.121, URL https://doi.org/10.1016/j.jmrt.2022.07.121

[8] DebRoy T, Wei HL, Zuback JS, et al (2018) Additive manufacturing of metallic components – process, structure and properties. Progress in Materials Science 92:112–224. https://doi.org/10.1016/j.pmatsci.2017.10.001, URL https://doi.org/10.1016/j.pmatsci.2017.10.001

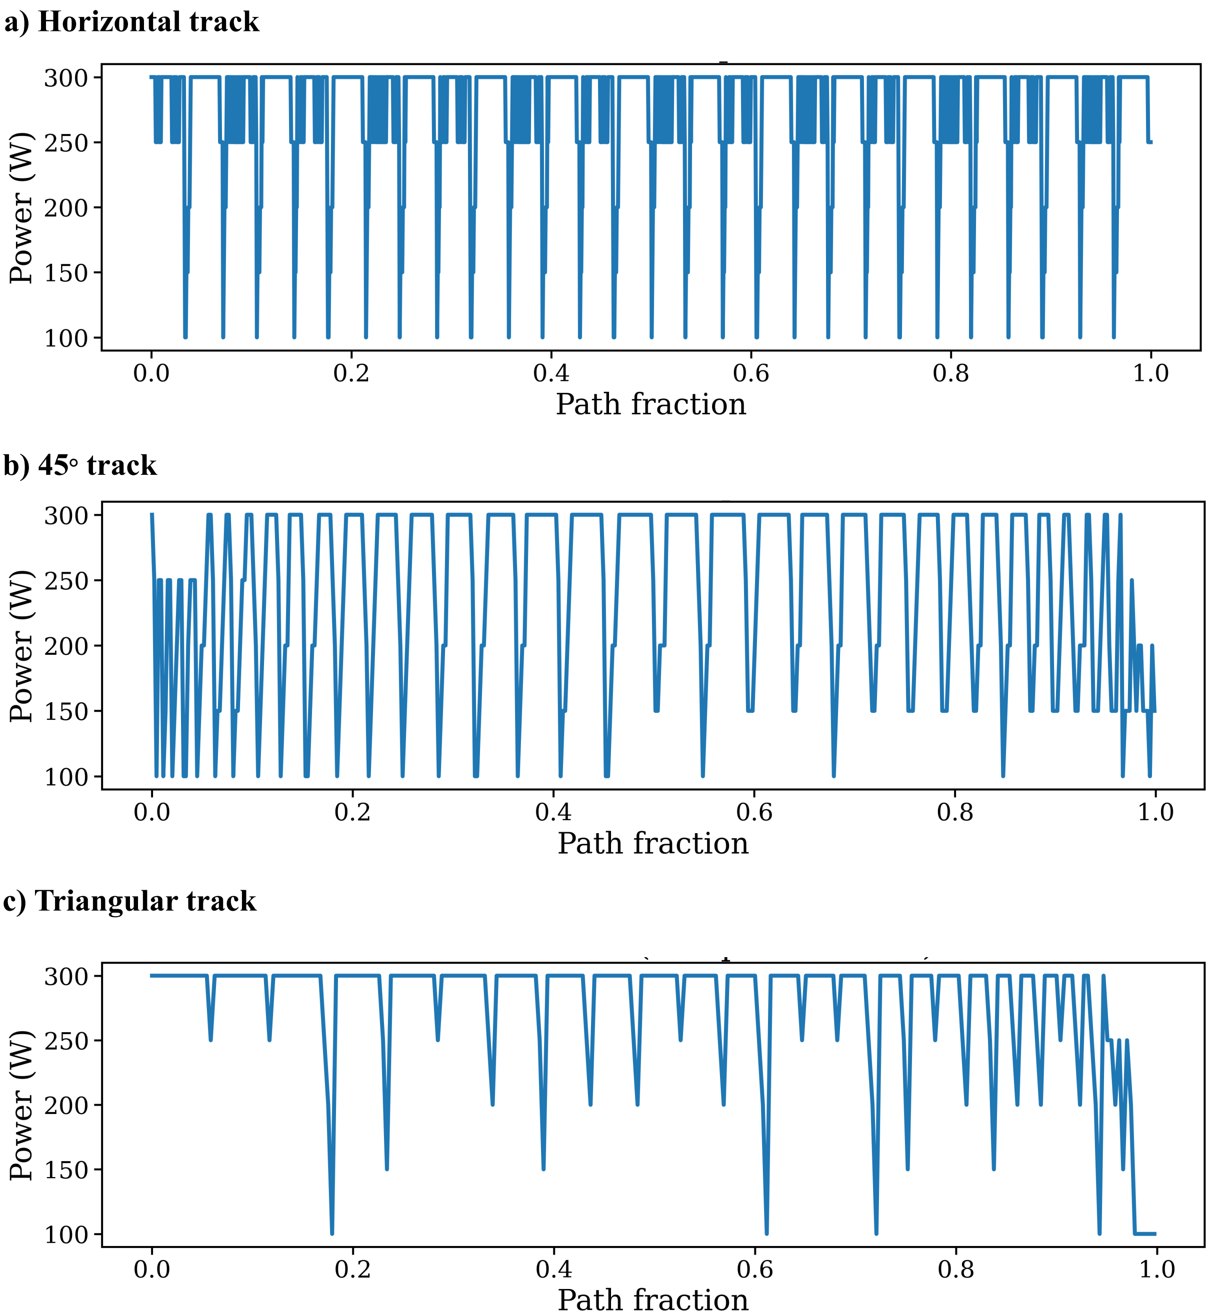

Fig. 20: Quantized feedforward laser power schedules for the three single-layer scan strategies: (a) horizontal hatch, (b) 45° hatch, and (c) concentric triangular path. The controller is restricted to a discrete set of machine-realizable power levels (100-300 W with increments of 50 W) rather than a continuous command range, causing the analytically computed schedule to snap to quantized values.

[9] Dulac-Arnold G, Levine N, Mankowitz D, et al (2021) Challenges of real-world reinforcement learning: Definitions, benchmarks and analysis. Machine Learning 110:2419–2468. https://doi.org/10.1007/s10994-021-05961-4

[10] Eagar TW, Tsai NS (1983) Temperature fields produced by traveling distributed heat sources. Welding Journal (Welding Research Supplement) 62(12):346s–355s

[11] Fox JC, Evans CJ, Weaver JS, et al (2023) Surface topography and melt pool behavior in rapid turnaround regions of laser powder bed fusion additive manufacturing of nickel superalloy 625. Wear

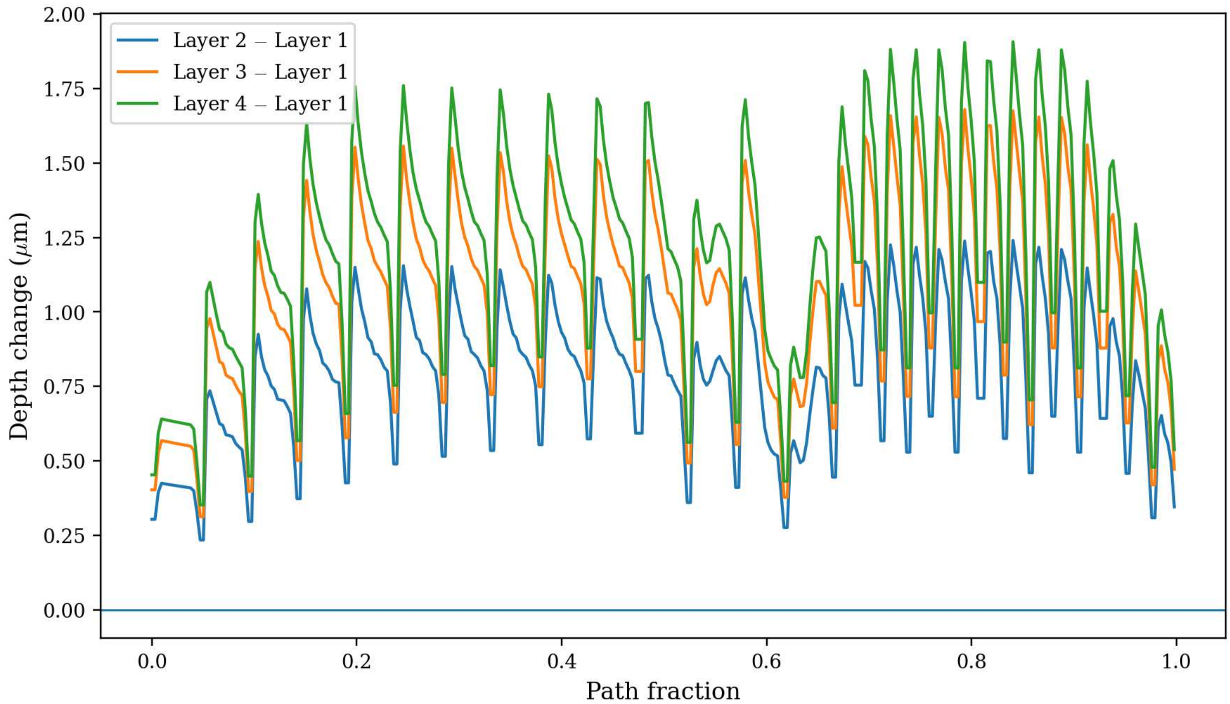

Fig. 21: Melt pool depth increment relative to Layer 1 as a function of path fraction for Layers 2, 3, and 4. The monotonically increasing layer-to-layer offset confirms that cumulative inter-layer substrate preheating progressively elevates the baseline melt pool depth under constant laser power, with the effect amplified at track turnarounds where residual heat from adjacent vectors compounds with the inter-layer thermal buildup.

523:204831. https://doi.org/10.1016/j.wear.2023.204831

[12] Frazier WE (2014) Metal additive manufacturing: A review. Journal of Materials Engineering and Performance 23(6):1917–1928. https://doi.org/10.1007/s11665-014-0958-z, URL https://doi.org/10.1007/s11665-014-0958-z

[13] Goldak J, Chakravarti A, Bibby M (1984) A new finite element model for welding heat sources. Metallurgical Transactions B 15(2):299–305. https://doi.org/10.1007/BF02667333

[14] Gordon JV, Narra SP, Cunningham RW, et al (2020) Defect structure process maps for laser powder bed fusion additive manufacturing. Additive Manufacturing 36:101552. https://doi.org/10.1016/j.addma.2020.101552, URL https://doi.org/10.1016/j.addma.2020.101552

[15] Gorgannejad S, Martin AA, Nicolino JW, et al (2023) Localized keyhole pore prediction during laser powder bed fusion via multimodal process monitoring and x-ray radiography. Additive Manufacturing 78(N/A). https://doi.org/10.1016/j.addma.2023.103810, URL https://doi.org/10.1016/j.addma.2023.103810

[16] Grattafiori A, et al (2024) The llama 3 herd of models. arXiv preprint arXiv:240721783 URL https://arxiv.org/abs/2407.21783

[17] Groq (2024) Groq llama-3-tool-use models: Top tool use & function calling. Blog post, URL https://groq.com/blog/introducing-llama-3-groq-tool-use-models

[18] Hastings WK (1970) Monte carlo sampling methods using markov chains and their applications. Biometrika 57(1):97–109. https://doi.org/10.1093/biomet/57.1.97

[19] Herzog D, Seyda V, Wycisk E, et al (2016) Additive manufacturing of metals. Acta Materialia 117:371–392. https://doi.org/10.1016/j.actamat.2016.07.019, URL https://doi.org/10.1016/j.actamat.2016.07.019

[20] Higdon D, Gattiker J, Williams B, et al (2008) Computer model calibration using high-dimensional output. Journal of the American Statistical Association 103(482):570–583

[21] Huang Y, Fleming TG, Clark SJ, et al (2022) Keyhole fluctuation and pore formation mechanisms during laser powder bed fusion additive manufacturing. Nature Communications 13:1170. https://doi.org/10.1038/s41467-022-28694-x, URL https://doi.org/10.1038/s41467-022-28694-x

[22] Kaipio J, Somersalo E (2005) Statistical and Computational Inverse Problems. Springer, https://doi.org/10.1007/b138659

[23] Keeley W, Turner R, Mitchell B, et al (2025) A development of the rosenthal equation for predicting thermal profiles during additive manufacturing. Thermo 5(2):16. https://doi.org/10.3390/thermo5020016

[24] Kennedy MC, O'Hagan A (2001) Bayesian calibration of computer models. Journal of the Royal Statistical Society: Series B 63(3):425–464

[25] Khadke A, Rai R (2020) Reinforcement learning-based control for additive manufacturing process optimization. Journal of Manufacturing Systems 57:404–417. https://doi.org/10.1016/j.jmsy.2020.09.009

[26] Khairallah SA, Anderson AT, Rubenchik A, et al (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. https://doi.org/10.1016/j.actamat.2016.02.014, URL https://doi.org/10.1016/j.actamat.2016.02.014

[27] Kirschbaum N, Wood N, Kim CE, et al (2025) Vector-level feedforward control of lpbf melt pool area using a physics-based thermal model. Additive Manufacturing 112:104981. https://doi.org/10.1016/j.addma.2025.104981, URL https://doi.org/10.1016/j.addma.2025.104981

[28] Kyogoku H, Ikeshoji TT (2020) A review of metal additive manufacturing technologies: Mechanism of defects formation and simulation of melting and solidification phenomena in laser powder bed fusion process. Mechanical Engineering Reviews 7(1):19–00182. https://doi.org/10.1299/mer.19-00182

[29] Leung CLA, Marussi S, Atwood RC, et al (2018) In situ x-ray imaging of defect and molten pool dynamics in laser additive manufacturing. Nature Communications 9:1355. https://doi.org/10.1038/s41467-018-03734-7

[30] Liao-McPherson D, Kolb M, Knapp GL, et al (2023) Layer-to-layer melt pool control in laser powder bed fusion. URL https://arxiv.org/abs/2311.10218, arXiv:2311.10218

[31] Lupi F, Pacini A, Lanzetta M (2023) Laser powder bed additive manufacturing: A review on the four drivers for an online control. Journal of Manufacturing Processes 103:413–429. https://doi.org/10.1016/j.jmapro.2023.08.022, URL https://doi.org/10.1016/j.jmapro.2023.08.022

[32] Mani M, Lane B, Donmez A, et al (2015) Measurement science needs for real-time control of additive manufacturing powder bed fusion processes. NIST Interagency/Internal Report (NISTIR) 8036 8036, National Institute of Standards and Technology (NIST), https://doi.org/10.6028/NIST.IR.8036, URL https://nvlpubs.nist.gov/nistpubs/ir/2015/NIST.IR.8036.pdf

[33] Martin AA, Calta NP, Khairallah SA, et al (2019) Dynamics of pore formation during laser powder bed fusion additive manufacturing. Nature Communications 10:1987. https://doi.org/10.1038/s41467-019-10009-2

[34] Meta AI (2024) Introducing meta llama 3: The most capable openly available llm to date. Blog post, URL https://ai.meta.com/blog/meta-llama-3/

[35] Metropolis N, Rosenbluth AW, Rosenbluth MN, et al (1953) Equation of state calculations by fast computing machines. The Journal of Chemical Physics 21(6):1087–1092. https://doi.org/10.1063/1.1699114

[36] Mirkoohi E, Malhotra R (2017) Effect of particle shape on neck growth and shrinkage of nanoparticles. In: Proceedings of the ASME 2017 12th International Manufacturing Science and Engineering Conference (MSEC2017), Additive Manufacturing; Materials, vol 2. ASME, Los Angeles, California, USA, https://doi.org/10.1115/MSEC2017-2811, URL https://doi.org/10.1115/MSEC2017-2811

[37] Mozaffar M, Liao S, Xie X, et al (2022) Mechanistic artificial intelligence (mechanistic-ai) for modeling, design, and control of advanced manufacturing processes: Current state and perspectives. Journal of Materials Processing Technology 302:117485

[38] Naderi M, Weaver JS, Deisenroth DC, et al (2023) On the fidelity of the scaling laws for melt pool depth analysis during laser powder bed fusion. Integrating Materials and Manufacturing Innovation 12(1):11–26. https://doi.org/10.1007/s40192-022-00289-w, URL https://doi.org/10.1007/s40192-022-00289-w

[39] National Institute of Standards and Technology (NIST) (2024) Single-track laser scan cross-sectional micrographs on in625 and in718 bare plates with melt pool depth and width measurements. NIST Public Data Repository, https://doi.org/10.18434/mds2-2923, URL https://doi.org/10.18434/mds2-2923, version 1.0.0

[40] Ngo TD, Kashani A, Imbalzano G, et al (2018) Additive manufacturing (3d printing): A review of materials, methods, applications and challenges. Composites Part B: Engineering 143:172–196. https://doi.org/https://doi.org/10.1016/j.compositesb.2018.02.012, URL https://www.sciencedirect.com/science/article/pii/S1359836817342944

[41] Ni J, Lammens N, Erdelyi H (2025) Thermal history prediction in selective laser melting using an improved analytical model. Progress in Additive Manufacturing https://doi.org/10.1007/s40964-025-01365-1

[42] Oliveira JP, LaLonde AD, Ma J (2020) Processing parameters in laser powder bed fusion metal additive manufacturing. Materials & Design 193(C):108762. https://doi.org/10.1016/j.matdes.2020.108762, URL https://doi.org/10.1016/j.matdes.2020.108762

[43] Ollama (2024) llama3-groq-tool-use:8b. Ollama model library entry, URL https://ollama.com/library/llama3-groq-tool-use%3A8b

[44] Ollama (2024) Tool calling. Documentation, URL https://docs.ollama.com/capabilities/tool-calling

[45] Pareek Y, Saadi AMAMA, Basak A, et al (2025) Real-time thermal state estimation and forecasting in laser powder bed fusion. arXiv preprint arXiv:2509.16114. URL https://arxiv.org/abs/2509.16114, arXiv:2509.16114

[46] Park C, Saha S, Guo J, et al (2025) Unifying machine learning and interpolation theory via interpolating neural networks. Nature Communications 16(1):8753

[47] Ren Y, Zhang Y, et al (2023) Deep reinforcement learning for adaptive process control in additive manufacturing. Additive Manufacturing 65:103350. https://doi.org/10.1016/j.addma.2023.103350

[48] Rosenthal D (1946) The theory of moving sources of heat and its application to metal treatments. Transactions of the ASME 68:849–866

[49] Saha S (2024) Mechanistic computational intelligence (mci) from theory to practice: Opening pathways for a new generation of engineering software. PhD thesis, Northwestern University

[50] Saha S, Gan Z, Cheng L, et al (2021) Hierarchical deep learning neural network (HiDeNN): an artificial intelligence (AI) framework for computational science and engineering. Computer Methods in Applied Mechanics and Engineering 373:113452

[51] Schick T, Dwivedi-Yu J, Dessì R, et al (2023) Toolformer: Language models can teach themselves to use tools. Advances in Neural Information Processing Systems 36

[52] Shrestha S, Chou K (2022) Residual heat effect on the melt pool geometry during the laser powder bed fusion process. Journal of Manufacturing and Materials Processing 6(6):153. https://doi.org/10.3390/jmmp6060153

[53] Slotwinski JA (2014) Additive manufacturing: Overview and NDE challenges. In: AIP Conference Proceedings, pp 1173–1177, https://doi.org/10.1063/1.4864953, URL https://doi.org/10.1063/1.4864953

[54] Steuben JC, Birnbaum AJ, Michopoulos JG, et al (2019) Enriched analytical solutions for additive manufacturing modeling and simulation. Additive Manufacturing 25:437–447. https://doi.org/10.1016/j.addma.2018.10.017, URL https://doi.org/10.1016/j.addma.2018.10.017

[55] Sutton RS, Barto AG (2018) Reinforcement Learning: An Introduction, 2nd edn. MIT Press

[56] Templeton WF, Hinnebusch S, Strayer S, et al (2023) Finding the limits of single-track deposition experiments: An experimental study of melt pool characterization in laser powder bed fusion. Materials & Design 231:112069. https://doi.org/10.1016/j.matdes.2023.112069, URL https://doi.org/10.1016/j.matdes.2023.112069

[57] Vallabh CKP, Hinnebusch S, To AC, et al (2024) Melt pool width measurement in a multi-track, multi-layer laser powder bed fusion print using single-camera two-wavelength imaging pyrometry. International Journal of Advanced Manufacturing Technology 132:2485–2499. https://doi.org/10.1007/s00170-024-13486-y

[58] Wang L, Zhang Y, Chia HY, et al (2022) Mechanism of keyhole pore formation in metal additive manufacturing. npj Computational Materials 8(22). https://doi.org/10.1038/s41524-022-00699-6, URL https://doi.org/10.1038/s41524-022-00699-6

[59] Weaver JS, Heigel JC, Lane BM (2022) Laser spot size and scaling laws for laser beam additive manufacturing. Journal of Manufacturing Processes 73:26–39. https://doi.org/10.1016/j.jmapro.2021.10.053, URL https://doi.org/10.1016/j.jmapro.2021.10.053

[60] Wei J, Wang X, Schuurmans D, et al (2022) Chain-of-thought prompting elicits reasoning in large language models. Advances in Neural Information Processing Systems 35

[61] Wolfer AJ, Aires J, Wheeler K, et al (2019) Fast solution strategy for transient heat conduction for arbitrary scan paths in additive manufacturing. Additive Manufacturing 30:100898. https://doi.org/10.1016/j.addma.2019.100898

[62] Xie X, Bennett J, Saha S, et al (2021) Mechanistic data-driven prediction of as-built mechanical properties in metal additive manufacturing. npj Computational Materials 7(1):86

[63] Yang YP, Jamshidinia M, Boulware P, et al (2018) Prediction of microstructure, residual stress, and deformation in laser powder bed fusion process. Computational Mechanics 61:599–615. https://doi.org/10.1007/s00466-017-1528-7, URL https://doi.org/10.1007/s00466-017-1528-7

[64] Yao S, Zhao J, Yu D, et al (2023) React: Synergizing reasoning and acting in language models. International Conference on Learning Representations (ICLR)

[65] Yeung H, Lane B (2020) A residual heat compensation based scan strategy for powder bed fusion additive manufacturing. Manufacturing Letters 25:10–14. https://doi.org/10.1016/j.mfglet.2020.07.005

[66] Zagade P, Gautham BP, De A, et al (2021) Analytical estimation of fusion zone dimensions and cooling rates in part scale laser powder bed fusion. Additive Manufacturing 46:102222. https://doi.org/10.1016/j.addma.2021.102222

[67] Zagade PR, Gautham BP, De A, et al (2024) Analytical modelling of scanning strategy effect on temperature field and melt track dimensions in laser powder bed fusion. Additive Manufacturing 82:104046. https://doi.org/10.1016/j.addma.2024.104046

[68] Zhao C, Shi B, Chen S, et al (2022) Laser melting modes in metal powder bed fusion additive manufacturing. Reviews of Modern Physics 94(4):045002. https://doi.org/10.1103/RevModPhys.94.045002, URL https://doi.org/10.1103/RevModPhys.94.045002