applied sciences

MDPI

Article

Variational Physics-Informed Neural Network for 3D Transient Melt Pool Thermal Modeling

Zhenghao Xu [1,2,3], Xin Wang [1,2,3,*], Yuan Meng [1,2,3,*], Mingwei Wang [1,2,3] and Xianglong Wang [1,2,3]

[1] State Key Laboratory of Advanced Forming Technology and Equipment, China Academy of Machinery Science and Technology Group, Beijing 100044, China
[2] Beijing National Innovation Institute of Lightweight Co., Ltd., Beijing 100083, China
[3] China Machinery Institute of Advanced Materials Co., Ltd., Zhengzhou 450001, China
* Correspondence: wangx05@tsinghua.org.cn (X.W.); mengyuan1383@163.com (Y.M.)

Abstract

Accurate prediction of transient melt pool thermal fields in Laser Powder Bed Fusion (LPBF) is essential for understanding melt pool geometry and defect formation mechanisms, yet conventional finite element methods (FEM) impose prohibitive computational costs for parametric process exploration. A variational physics-informed neural network (VPINN) framework is presented for 3D transient thermal modeling of a GH3536 single-track LPBF scan. The framework incorporates a continuously differentiable Goldak double-ellipsoid moving heat source, temperature-dependent thermophysical property surrogates, and an effective heat-capacity treatment of latent heat associated with solid–liquid phase change and vaporization. These components are embedded in a weak-form residual-minimization scheme with octree-adaptive domain decomposition, hierarchical Legendre test functions, and sequential sliding-window time marching. Effective absorptivity is inferred jointly with the network parameters, using sparse experimental melt pool profiles as supervision. Within a parametric study covering laser powers from 100 to 140 W and scan speeds from 1000 to 1500 mm/s, the predicted melt pool width, depth, and aspect ratio agree closely with FEM benchmarks and cross-sectional optical micrograph measurements across both supervised and held-out interpolation conditions, with total relative L_2 nodal temperature errors ranging from 3.23% to 6.75%. Following a one-time offline training investment of 15,323 s that simultaneously resolves the full parametric space, surrogate inference reduces per-condition query time from 3000–4000 s (FEM) to merely 4–5 s, delivering a speedup of two to three orders of magnitude and making the framework increasingly cost-effective for high-throughput parametric studies and digital-twin integration as the number of queried conditions grows.

Keywords: laser powder bed fusion; physics-informed neural network; variational formulation; melt pool thermal modeling; inverse calibration; additive manufacturing

Check for updates

Received: 6 March 2026
Revised: 27 March 2026
Accepted: 13 April 2026
Published: 14 April 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Laser Powder Bed Fusion (LPBF), also referred to as Selective Laser Melting (SLM), is characterized by highly localized, rapidly evolving thermal fields that dictate melt pool dimensions and thermal histories, thereby exerting first-order control over consolidation quality and defect propensity across process windows [1,2]. The melt pool width W and depth D provide compact geometric signatures of the underlying temperature gradients and cooling rates, making transient thermal prediction central to both mechanistic interpretation and practical process qualification in track-scale studies [3].

Computational heat transfer models for LPBF commonly rely on a moving volumetric heat source embedded in a nonlinear heat conduction formulation, because they provide a tractable mapping from process inputs to temperature fields while retaining compatibility with engineering-scale simulation workflows [4]. Among volumetric source representations, the Goldak double-ellipsoid model remains a de facto standard due to its flexible front–rear asymmetry and energy-conservative construction, which have motivated extensive adoption beyond welding and into laser-based manufacturing heat-transfer analysis [4]. Practical accuracy, however, is often limited by the need to calibrate effective laser–material energy coupling parameters against experimental observables, with effective absorptivity and related assumptions frequently acting as dominant uncertainty carriers that are often re-identified across scan conditions and material states [5].

Even when the governing physics is restricted to thermal conduction with phase-change regularization, 3D transient simulations based on the FEM remain computationally expensive for parametric exploration, design-of-experiments loops, and inverse identification workflows that require repeated forward solves [6]. This cost barrier is amplified by the multiscale nature of LPBF, where steep near-source gradients coexist with comparatively smooth far-field diffusion, producing a persistent tension between numerical resolution demands and throughput requirements in process-window studies [7,8].

Physics-Informed Neural Networks (PINNs) provide a Scientific Machine Learning (SciML) alternative by embedding Partial Differential Equation (PDE) residuals and initial–boundary constraints into the learning objective, enabling forward solution and data-driven inverse discovery under sparse supervision through automatic differentiation [9,10]. Within the track-scale LPBF, PINNs have been applied at progressively increasing physical fidelity. Hosseini et al. [6] parametrically solved the single-track transient heat equation with a Gaussian heat source and constant properties in an unsupervised setting, achieving melt pool size errors below 5% relative to FEM across varied process conditions. Liao et al. [11] introduced semi-supervised calibration by fusing sparse infrared measurements with the heat equation, simultaneously recovering full-field temperatures and inferring unknown material parameters and laser absorptivity. Jiang et al. [12] extended data-physics coupling to predict melt pool dimensions across variable scan speeds from a single trained model using a small labeled dataset. Beyond conduction, Sharma et al. [13] embedded the coupled Navier–Stokes and energy equations to recover melt pool velocity and pressure fields from sparse temperature data alone, while Zhu et al. [14,15] reduced training cost through transfer learning and reality-augmented adaptive sampling strategies. Safari and Wessels [16] further extended the paradigm to multi-track scenarios via DeepONet operator learning. All of the above studies rely on strong-form pointwise residual minimization; the variational weak-form alternative has not been explored for 3D transient LPBF melt pool modeling.

VPINNs address these limitations by minimizing weak-form residuals over suitably chosen test spaces, reducing the differentiability burden on the trial network and improving robustness in problems with steep gradients and localized features [17]. Furthermore, hp-VPINNs enhance this approach through domain decomposition and high-order polynomial test functions, enabling localized resolution enrichment and more favorable accuracy–cost tradeoffs for multiscale transient phenomena [18]. Within the LPBF single-track context, this variational SciML perspective creates a direct pathway to amortized, parameterized thermal surrogates that can be queried rapidly while remaining anchored to the same governing equations used in established conduction-based benchmarks [14–16]. In this work, a variational physics-informed learning framework is formulated for 3D transient melt pool thermal modeling in single-track LPBF scanning of GH3536 under conduction-to-transition-mode conditions. The principal contributions are:

(1) The VPINN weak-form residual-minimization framework is applied for the first time to 3D transient LPBF melt pool thermal field prediction, reducing the differentiability requirement on the network from second to first order relative to strong-form PINNs.
(2) A heat-source-aware adaptive element weighting combined with causal sliding-window time marching is proposed to address the coupled optimization stiffness from the moving Goldak source and phase-change enthalpy peaks.
(3) Effective absorptivity η is treated as the sole trainable physical parameter, calibrated jointly with the network weights via sparse melt pool contour supervision, eliminating iterative outer-loop FEM calibration.

2. Methodology

2.1. Governing Equations

The transient thermal field $T(\mathbf{x}, t)$ in domain $\Omega \subset \mathbb{R}^3$ over $\mathcal{T} = [0, t_{end}]$ satisfies the nonlinear heat conduction equation governing LPBF single-track processing:

$$\rho(T)C_p^*(T)\frac{\partial T}{\partial t} = \nabla \cdot (k(T)\nabla T) + Q(\mathbf{x}, t),\ (\mathbf{x}, t) \in \Omega \times \mathcal{T} \tag{1}$$

where $C_p^*(T)$ denotes the effective specific heat incorporating latent heat contributions (Section 2.3) and $Q(\mathbf{x},t)$ denotes the Goldak volumetric source [19]. The domain is initialized at $T(\mathbf{x}, 0) = 298.15$ K, and all external boundaries $\partial\Omega$ satisfy a Robin condition combining natural convection and surface radiation:

$$-k(T)\frac{\partial T}{\partial \mathbf{n}} = h_c(T - T_\infty) + \epsilon\sigma_{SB}(T^4 - T_\infty^4) \tag{2}$$

where the ambient temperature is defined as $T_\infty = 298.15$ K, the convective heat transfer coefficient is uniformly set to $h_c = 15\,\mathrm{W/(m^2 \cdot K)}$, the surface emissivity is assigned a representative value of $\epsilon = 0.4$ for the nickel-based superalloy, and $\sigma_{SB} = 5.67 \times 10^{-8}\,\mathrm{W/(m^2 \cdot K^4)}$ denotes the Stefan-Boltzmann constant [20]. To avoid numerical scaling disparities, the temperature and time variables are non-dimensionalized as follows [21]:

$$u = \frac{T}{T_{\max}},\quad \widetilde{t} = \frac{t \cdot v}{L_{ref}} \tag{3}$$

where $T_{\max} = 3200$ K is the normalization temperature chosen to encompass the full ambient-to-near-vaporization range, $L_{ref} = 1.5$ mm is the reference length corresponding to the scan track length, and v is the physical scanning velocity. The spatial coordinates $\mathbf{x} = (x, y, z)$ retain their physical dimensions in millimeters and are not non-dimensionalized.

2.2. Moving Heat Source

To accurately capture the asymmetric thermal gradient of the melt pool while maintaining strict C^∞ differentiability for stable automatic differentiation within the VPINN framework, the laser–material interaction is modeled using a smoothly blended formulation of the Goldak double-ellipsoid moving heat source [19]. Conventional discrete quadrant switching introduces non-physical residual spikes during backpropagation; conversely, naively interpolating geometric semi-axes within the exponential argument compromises the analytical energy conservation of the domain integral. The present formulation adopts a field-level superposition strategy, constructing the total volumetric power density by blending two analytically conservative fields—a front ellipsoid Q_f and a rear ellipsoid Q_r—through a hyperbolic tangent smoothing weight $\omega(x, t)$:

$$Q(x,y,z,t) = \omega(x,t)Q_f(x,y,z,t) + (1-\omega(x,t))Q_r(x,y,z,t) \quad (4)$$

where individual power densities are defined strictly in their classical conservative forms:

$$Q_i(x,y,z,t) = \frac{6\sqrt{3}\, f_i\, \eta P}{abc_i\, \pi\sqrt{\pi}} \exp\left(-3\left[\frac{(x-vt)^2}{c_i^2} + \frac{y^2}{a^2} + \frac{(z-z_0)^2}{b^2}\right]\right),\ i \in \{f,r\} \quad (5)$$

The spatial transition weight acts as a continuously differentiable Heaviside function localized at the moving beam center:

$$\omega(x,t) = \frac{1}{2} + \frac{1}{2}\tanh\left(\frac{x-vt}{l_s}\right) \quad (6)$$

with smoothing length $l_s = 0.005$ mm—an order of magnitude below the transverse semi-axis a—ensuring that the field transition is confined within a sub-element domain without smearing the macroscopic Goldak geometry. Because Q_f and Q_r are each independently energy-conservative, the blending weight ω analytically drops out of the global energy integral. Consequently, the macroscopic temperature field is theoretically insensitive to the specific value of l_s provided it remains within the bounded range $0 \ll l_s < h_{\min}$, where the upper bound preserves sub-element spatial localization and the lower bound avoids floating-point saturation of the tanh argument. Furthermore, since the energy partitioning coefficients satisfy $f_f + f_r = 2$ by construction, interpolating scalar flux fields rather than geometric parameters strictly preserves the global integral $\iiint_\Omega Q d\Omega = \eta P$, guaranteeing energy conservation as a prerequisite for physically meaningful parameter inversion.

Source geometry is assigned as physically motivated prior estimates consistent with the nominal laser spot radius and standard Goldak proportions: the transverse semi-axes $a = b = 0.05$ mm approximate the beam half-width typical of LPBF optics in the present power range, while the longitudinal semi-axes $c_f = 0.05$ mm and $c_r = 0.10$ mm preserve the classical front–rear ratio $c_r/c_f = 2$. Because η linearly scales the total absorbed power $P_{eff} = \eta P$ and thereby governs all isotherm extents monotonically, moderate perturbations in ellipsoid geometry can be largely compensated by a corresponding shift in the calibrated η as the sole trainable quantity—calibrated jointly with the solution network weights per Section 2.4—is therefore a deliberate regularization choice that avoids the parameter co-linearity inherent in jointly identifying shape and magnitude from sparse cross-sectional supervision, while concentrating all residual modeling uncertainty into a single physically interpretable scalar. The spatial distributions of the volumetric power density in the x-z, x-y, and y-z planes are depicted in Figure 1.

2.3. Thermophysical Properties and Phase Change Formulation

The thermophysical properties of GH3536 (thermal conductivity $k(T)$, density $\rho(T)$, and specific heat $C_p(T)$) are each represented by an independent Feedforward Neural Network (hereafter referred to as MaterialFNN), each comprising three hidden layers of 32 neurons with hyperbolic tangent activations. All three networks are trained exclusively on tabulated property data prior to any VPINN optimization. Specifically, the temperature-dependent reference data were extracted directly from the FLOW-3D material database. For each thermophysical quantity, the training dataset comprises approximately 30 discrete temperature-property pairs. These property surrogates were optimized via the Adam algorithm over a training cycle of 3000 epochs. Upon completion, their parameters $\theta_{mat} = \{\theta_k, \theta_\rho, \theta_{C_p}\}$ are frozen for the duration of physics-informed training, supplying $k(T)$, $\rho(T)$, and $C_p(T)$—together with their exact automatic-differentiation gradients—at arbitrary quadrature points as static material closure modules. This decoupled design pre-

vents PDE residual gradients from distorting physically calibrated property representations, a failure mode that joint co-optimization cannot reliably suppress even under anchoring regularization. The resulting property fits are depicted in Figure 2.

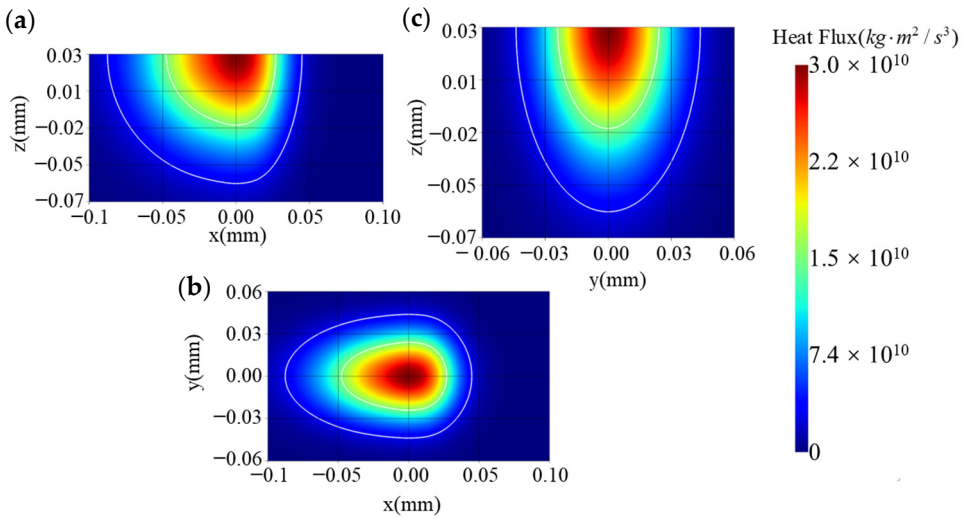


Figure 1. Power density distributions of the Goldak double-ellipsoid heat model: (**a**) *x*-*z* plane; (**b**) *x*-*y* plane; (**c**) *y*-*z* plane.

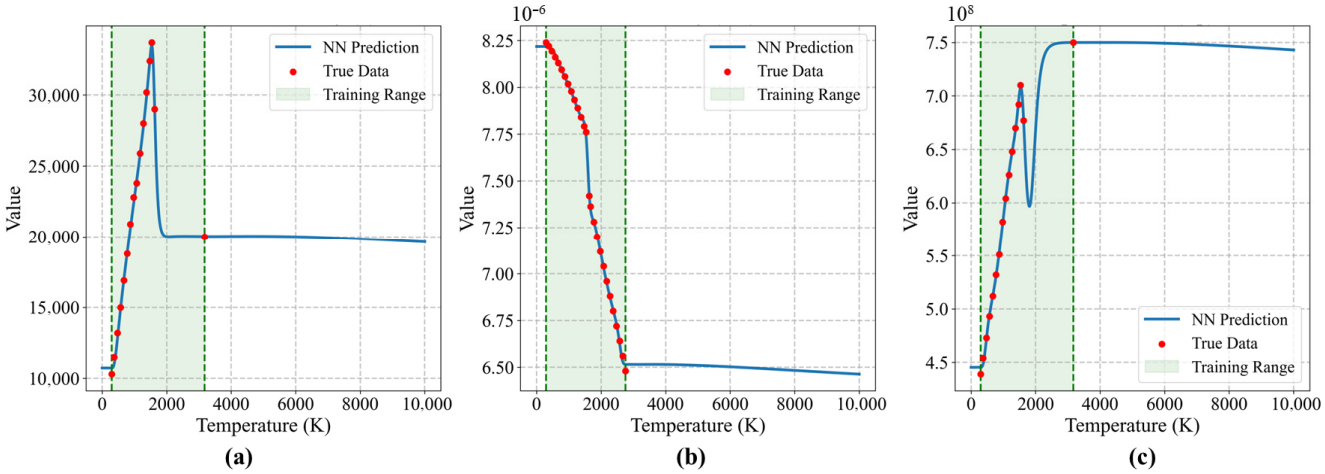


Figure 2. Temperature-dependent thermophysical properties of GH3536 represented by the material surrogate network: (**a**) thermal conductivity; (**b**) density; (**c**) specific heat capacity.

Each MaterialFNN is trained under a composite objective that combines supervised data fidelity with extrapolation regularization:

$$\mathcal{L}_{mat}^{(f)}(\theta_f) = \mathcal{L}_{data}^{(f)}(\theta_f) + \lambda_{ext}\mathcal{L}_{extrap}^{(f)}(\theta_f) \tag{7}$$

The data-fitting term enforces agreement with tabulated measurements $\{T_i, f_i^{ref}\}_{i=1}^{N_{data}}$ across the supervised interval $[T_{min}, T_{max}]$:

$$\mathcal{L}_{data}^{(f)}(\theta_f) = \frac{1}{N_{data}} \sum_{i=1}^{N_{data}} \left(\frac{\hat{f}(T_i;\theta_f) - f_i^{ref}}{\sigma_f} \right)^2 \tag{8}$$

where σ_f normalizes gradient magnitudes across properties with disparate physical units. Because unconstrained network predictions beyond $[T_{min}, T_{max}]$ may introduce spurious

derivative magnitudes into the diffusion operator $\nabla \cdot (k(T)\nabla T)$—a consequential risk given the ambient-to-near-vaporization temperature span traversed during LPBF simulation—the extrapolation regularization term simultaneously penalizes non-zero derivatives and deviations from the nearest boundary reference value at extrapolation samples $\{T_j^{ext}\}_{j=1}^{N_{ext}}$:

$$\mathcal{L}_{extrap}^{(f)}(\theta_f) = \frac{1}{N_{ext}} \sum_{j=1}^{N_{ext}} \left[\left(\frac{d\hat{f}}{dT}\Big|_{T_j^{ext}} \right)^2 + \left(\hat{f}(T_j^{ext}; \theta_f) - f_{bnd}^{ref} \right)^2 \right] \tag{9}$$

where f_{bnd}^{ref} is the tabulated reference value at the nearest training boundary, defined as $f^{ref}(T_{min})$ for $T_j^{ext} < T_{min}$ and $f^{ref}(T_{max})$ for $T_j^{ext} > T_{max}$.

Latent heat at both fusion and vaporization transitions is incorporated via the effective heat capacity method [22,23], augmenting the sensible $C_p(T)$ with Gaussian-distributed enthalpy peaks:

$$C_p^*(T) = C_p(T) + \sum_{i \in \{f,\, v\}} \frac{L_i}{\sqrt{2\pi}\, \sigma_i} \exp\left(-\frac{(T - T_{c,i})^2}{2\sigma_i^2} \right) \tag{10}$$

The solid-liquid transition employs a solidus temperature of 1533.15 K, a liquidus temperature of 1628.15 K, and a latent heat of fusion of 2.7×10^{11} mm^2 s^{-2}; the liquid-vapor transition uses an initial boiling temperature of 2973.15 K, a final boiling temperature of 3173.15 K, and a latent heat of vaporization of 6.9×10^{12} mm^2 s^{-2}. The standard deviation σ_i of the Gaussian distribution, acting as the numerical smoothing parameter, is set proportional to the half-width of each transition interval. This formulation preserves full differentiability of the governing equation across both phase changes, circumventing the need for algorithmic branching [24]. It should be noted that vaporization is modeled here solely as an enthalpy sink; mass loss, recoil pressure, free-surface deformation, and keyhole formation are outside the scope of the present thermal conduction framework. Accordingly, the solid–liquid interface is identified as the 1628.15 K liquidus isotherm extracted from the predicted temperature field, consistent with the melt pool geometry extraction protocol in Sections 3.1 and 3.2.

2.4. VPINN Formulation and Loss Function Architecture

Network Architecture. The solution network $u(\mathbf{x}, t; \theta) : \mathbb{R}^6 \to \mathbb{R}$ maps (t, x, y, z, P, v) to dimensionless temperature $u = T/T_{max}$. To overcome the spectral bias of standard multilayer perceptrons (MLPs) against steep moving-source thermal gradients, the spatiotemporal coordinates $\mathbf{v}_{st} = (t, x, y, z)^T$ are first encoded through fixed Gaussian random Fourier features [25,26]:

$$\gamma(\mathbf{v}) = [\sin(2\pi \mathbf{B}\mathbf{v}_{st}),\ \cos(2\pi \mathbf{B}\mathbf{v}_{st})]^T, \quad \mathbf{B} \in \mathbb{R}^{128 \times 4} \sim \mathcal{N}(0, 10^2) \tag{11}$$

yielding a 256-dimensional spatiotemporal representation. It should be noted that the process parameters (P, v) do not undergo Fourier transformation; instead, they are directly concatenated with the encoded spatiotemporal features. This results in a 258-dimensional input vector that feeds an 8-layer $\times$ 32-neuron MLP with sine activations [27], preserving C^∞ differentiability of the complete computational graph. The architecture is depicted in Figure 3.

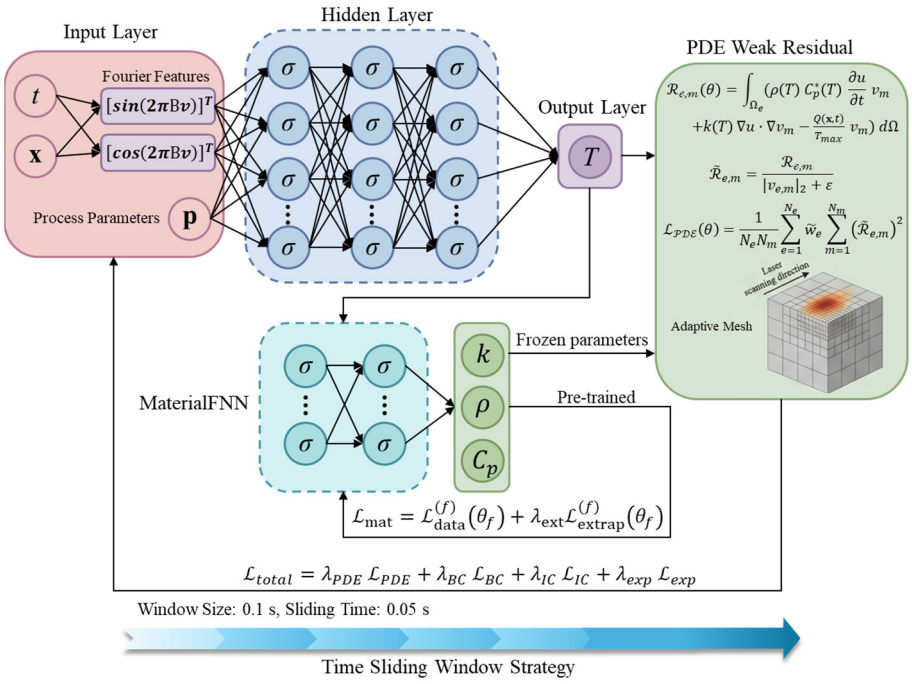


Figure 3. Core architecture of the proposed variational physics-informed learning framework for transient thermal modeling.

Variational Weak Form. In contrast to conventional PINNs that minimize pointwise strong-form residuals, the VPINN framework projects the governing PDE onto a finite-dimensional test function space, reducing the differentiability requirement on the network output from second- to first-order spatial derivatives [18]. The spatial domain is decomposed into N_e non-overlapping hexahedral elements via Octree partitioning. The mesh is graded from $h_{max} = 0.2$ mm in the far field to $h_{min} = 0.02$ mm along the scan trajectory. This spatial discretization strategically concentrates quadrature points within the melt pool and heat-affected zone, as depicted in Figure 4 [28]. Within each element, hierarchical Legendre bubble functions:

$$v_m(\xi) = (1 - \xi^2)\, L_p(\xi),\ p = 3,\ N_m = 3^p = 27 \tag{12}$$

serve as test functions, with $(1 - \xi^2)$ enforcing vanishing at element boundaries by construction. The element variational residual, evaluated numerically utilizing a $4 \times 4 \times 4$ Gauss–Legendre quadrature rule yielding 64 integration points per hexahedron to ensure exact integration of polynomial terms, reads [29]:

$$\mathcal{R}_{e,m}(\theta) = \int_{\Omega_e} \left(\rho C_p^* \frac{v}{L_{ref}} \frac{\partial u}{\partial \widetilde{t}}\, v_m + k\, \nabla u \cdot \nabla v_m - \frac{Q(\eta; \mathbf{x}, \widetilde{t})}{T_{max}}\, v_m \right) d\Omega \tag{13}$$

Scale consistency across elements of differing volume is enforced by L^2 normalization: $\widetilde{\mathcal{R}}_{e,m} = \mathcal{R}_{e,m} / (\|v_{e,m}\|_2 + \varepsilon)$, where $\|v_{e,m}\|_2 = \sqrt{\sum_i w_i^q v_{i,m}^2}$. Because the Goldak source concentrates energy deposition within a sub-millimeter region, unweighted optimization disproportionately favors the numerically dominant far-field residuals at the expense of the physically critical melt-pool zone. Element-wise heat-source-aware weighting counteracts this imbalance:

$$w_e = 1 + \beta\, (q_{\mathrm{norm},e})^{\alpha}, \qquad \widetilde{w}_e = w_e / \left(\frac{1}{N_e} \sum\nolimits_e w_e + 10^{-12} \right) \tag{14}$$

The exponent $\alpha = 1$ was chosen as the simplest monotonic mapping that preserves inter-element ranking without adding nonlinearity. The prefactor $\beta = 1$ keeps the maximum element weight within one to two orders of magnitude of the far-field baseline $w_e = 1$, preventing gradient domination by source-adjacent elements. Both values were fixed from a single preliminary run on V01 and reused across all conditions without re-tuning.

Yielding the PDE loss:

$$\mathcal{L}_{PDE}(\theta) = \frac{1}{N_e N_m} \sum_{e=1}^{N_e} \widetilde{w}_e \sum_{m=1}^{N_m} \widetilde{\mathcal{R}}_{e,m}^2 \tag{15}$$

Boundary and initial conditions are enforced via mean-squared-error collocation, contributing losses $\mathcal{L}_{BC}$ and $\mathcal{L}_{IC}$.

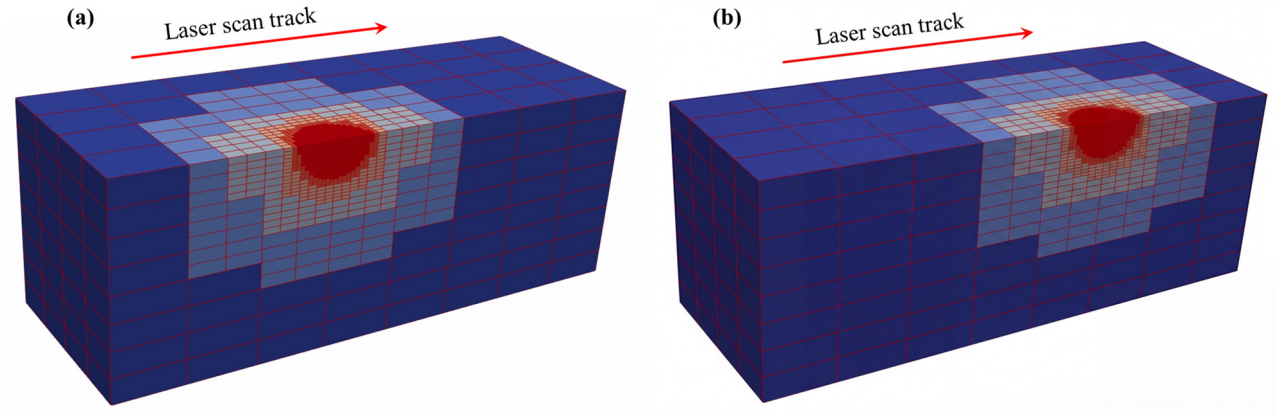


Figure 4. Dynamic quadrature sampling for VPINN weak-form residuals utilizing adaptive mesh refinement along the scan track: (**a**) laser positioned at x = 0 mm; (**b**) laser positioned at x = 1.5 mm.

Physics-Data Co-driven Inverse Calibration and Training Objective. The scalar absorptivity $\phi = \{\eta\}$ is identified jointly with θ by embedding an experimental constraint loss within the training objective, eliminating the iterative forward-solve loops characteristic of FEM calibration. Melt pool boundary coordinates extracted from metallographic cross-sections under the liquidus criterion $T_l = 1628.15$ K (Section 3.1) define N_{exp} target points at quasi-steady-state time t_{ss}:

$$\mathcal{L}_{exp}(\theta, \phi) = \frac{1}{N_{exp}} \sum_{i=1}^{N_{exp}} [u(\mathbf{x}_i, t_{ss}; \theta, \phi) \cdot T_{max} - T_l]^2 \tag{16}$$

Scalar identifiability from $\mathcal{L}_{exp}$ is confirmed by the strict positivity of $\partial W/\partial \eta > 0$ and $\partial D/\partial \eta > 0$ throughout the physical parameter range, a direct consequence of η linearly scaling total absorbed power $P_{\text{eff}} = \eta P$ and thereby governing the spatial extent of all isotherms monotonically. The complete training objective for joint optimization of θ_{temp} and ϕ is:

$$\mathcal{L}_{total}(\theta_{temp}, \phi) = \lambda_{PDE}\, \mathcal{L}_{PDE} + \lambda_{BC}\, \mathcal{L}_{BC} + \lambda_{IC}\, \mathcal{L}_{IC} + \lambda_{exp}\, \mathcal{L}_{exp} \tag{17}$$

To dynamically rebalance competing gradient contributions—particularly the stiff nonlinear PDE residual near phase-change transitions—adaptive weights $\{\lambda_k\}$ are updated at every gradient step via an Inverse Dirichlet strategy utilizing an Exponential Moving Average (EMA) [30,31]:

$$\text{EMA}_k \leftarrow \beta_w\, \text{EMA}_k + (1 - \beta_w)\mathcal{L}_k, \qquad \lambda_k = \text{clip}\left(\frac{1/\text{EMA}_k}{\sum_j 1/\text{EMA}_j},\, w_{min},\, w_{max} \right) \tag{18}$$

where the smoothing factor β_w is uniformly set to 0.9 across all loss components to ensure temporal stability in the adaptive weight evolution. This specific value suppresses erratic gradient fluctuations during the highly non-convex early optimization phase without excessively delaying the algorithmic response to abrupt residual variations near the phase-change interfaces.

Training Protocol. Physical causality is enforced through a sequential time-marching strategy [32] that partitions the normalized temporal domain into overlapping windows of width $\Delta\tau = 0.1$ with stride 0.05: window n is initialized against the converged prediction of window $n-1$ at interface time t_{n-1}, propagating latent heat enthalpy history forward and precluding catastrophic forgetting of early thermal gradients. Each window undergoes two-stage optimization utilizing a full-batch collocation strategy to eliminate stochastic gradient noise. The Adam [33] optimizer, configured with a cosine-annealed learning rate ($\eta_0 = 3 \times 10^{-4}$, 500-step warmup), initially navigates the non-convex landscape for a maximum of 30,000 epochs. Subsequently, the L-BFGS [34] algorithm ($\eta_0 = 0.9$, Strong Wolfe line search) resolves residual stiffness from the nonlinear diffusion term, terminating either when the relative loss change drops below 10^{-8} or after an upper limit of 500,000 iterations. To stabilize the initial optimization dynamics before the EMA scheme activates, the penalty coefficients for the total loss (Equation (17)) were empirically initialized to $\lambda_{PDE} = 1.0$, $\lambda_{BC} = 100$, $\lambda_{IC} = 100$, and $\lambda_{exp} = 1000$, after which the inverse Dirichlet mechanism dynamically constrained the weights within $[10^{-2}, 10^{4}]$. These adaptive weights $\{\lambda_k\}$ frozen during the L-BFGS stage to prevent instability under line-search perturbations.

3. Experimentation and Verification

3.1. Process Parameter Matrix and Melt Pool Measurements

As enumerated in Table 1, five single-track experiments on GH3536 were structured into two orthogonal parametric sweeps: a velocity sweep (1000 to 1500 mm/s) at a fixed laser power of 100 W (cases V01–V03), and a power sweep (100 to 140 W) at a fixed scan speed of 1500 mm/s (cases V03–V05) [35]. Case V03 serves as the shared baseline anchoring cross-sweep consistency [36]. This matrix covers stable conduction-mode operation at the low-power/high-speed end and approaching the conduction-to-transition boundary at the high-power/low-speed end [37]. Tracks were manufactured as isolated scan vectors to preclude inter-track thermal interaction, and cross-sectional analysis was conducted at mid-track locations free of scan-initiation transients, as depicted in Figure 5.

The fusion boundary signatures, specifically width W and depth D, were derived from etched optical micrographs through a multi-stage coordinate extraction and refinement pipeline. To provide high-fidelity supervision for the inverse calibration of absorptivity η, the melt pool contours in the transverse cross-sections were first manually delineated by identifying the characteristic grain structure transition between the fusion zone and the heat-affected zone, as shown in Figure 6 [38]. These discrete, non-uniformly distributed boundary markers were subsequently processed using a cubic B-spline interpolation algorithm to reconstruct a continuous representation of the liquidus interface. To ensure unbiased gradient contributions in the experimental loss function $\mathcal{L}_{exp}$, the resulting splines were re-sampled into a fixed set of N_{exp} equidistant coordinates $\{x_i, z_i\}$ using a chord-length parameterization. This algorithmic treatment effectively filters the stochasticity inherent in manual labeling and ensures that the sparse experimental constraints are distributed with sufficient density to resolve the high-curvature regions of the melt pool base, thereby providing a robust geometric reference for both VPINN training and the subsequent FEM benchmark comparisons.

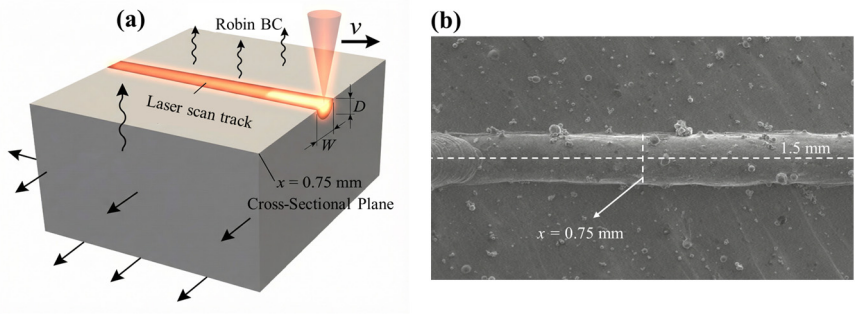


Figure 5. Single-track LPBF configuration: (**a**) schematic of the physical domain and melt-pool dimensions; (**b**) top-surface morphology of the deposited track.

Table 1. Process parameter matrix and experimentally measured melt pool geometric statistics for single-track LPBF experiments on GH3536.

Case	P (W)	v (mm/s)	W (μm)	D (μm)	R (D/W)	Train Role
V01	100	1000	67.123	31.045	0.463	Supervised
V02	100	1250	61.212	25.682	0.420	Validation
V03	100	1500	53.783	21.452	0.399	Validation
V04	120	1500	72.582	38.892	0.536	Supervised
V05	140	1500	75.551	32.067	0.424	Supervised

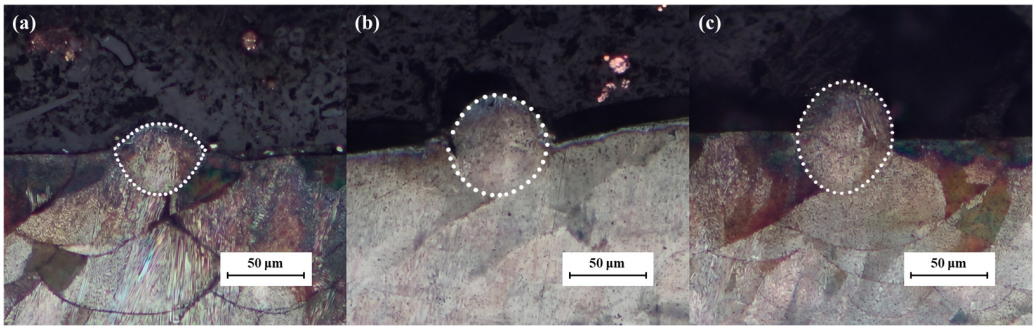


Figure 6. Experimental GH3536 melt-pool geometries serving as target constraints for neural network-based inverse calibration of absorptivity: (**a**) P = 100 W, v = 1000 mm/s; (**b**) P = 120 W, v = 1500 mm/s; (**c**) P = 140 W, v = 1500 mm/s.

3.2. FEM Benchmark

The FEM benchmark was configured using the FEniCS computing platform on a domain $[-1, 1.8] \times [-1, 1] \times [-1, 0]$ mm with governing equations, Robin boundary conditions, Goldak source parameters, and effective specific heat treatment identical to those of Section 2; thermophysical properties were supplied throughout by the pre-trained MaterialFNN, eliminating material representation as a confounding inter-model error source. The domain was discretized with 3D linear hexahedral elements refined to h_{min} = 0.02 mm within a 0.1 mm radial band of the scan path and coarsened to h_{max} = 0.2 mm in the far field, matching the VPINN Octree refinement levels shown in Figure 7. Temporal integration employed an unconditionally stable implicit backward Euler scheme at Δt = 1 μs (source displacement 1.5 μm per increment at maximum scanning velocity), with Newton–Raphson nonlinear iterations converging to a normalized residual tolerance of 10^{-6} per time step; the complete configuration is summarized in Table 2 [39]. Melt pool geometry was extracted from the T_l = 1628.15 K isotherm at t_{ss} in the transverse cross-section, consistent with the extraction protocol of Section 3.1.

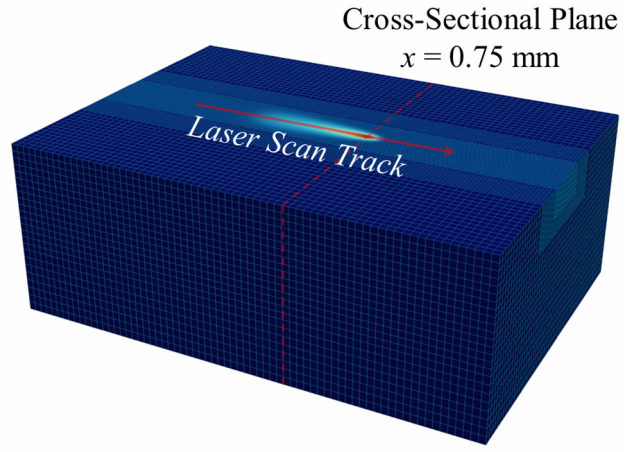


Figure 7. FEM mesh configuration for the 3D transient single-track thermal benchmark, indicating the local refinement zone and the transverse cross-sectional plane at x = 0.75 mm.

Table 2. Finite element analysis configuration for the 3D transient single-track LPBF thermal benchmark.

Parameter	Specification
Computational domain	$[-1, 1.8] \times [-1, 1] \times [-1, 0]$ mm
Temporal integration scheme	Unconditionally stable implicit backward Euler
Near-source element size	0.02 mm
Far-field element size	0.2 mm
Time step	1 μs
Nonlinear solver	Newton–Raphson iteration
Thermophysical property source	Pre-trained MaterialFNN
Heat source model	Goldak double-ellipsoid formulation
Boundary conditions	Robin: surface convection and thermal radiation
Latent heat treatment	Effective heat capacity method with Gaussian regularization
Melt pool geometry extraction criterion	1628.15 K isotherm at quasi-steady time, transverse plane at x = 0.75 mm

3.3. Evaluation Criteria

The VPINN framework is comprehensively evaluated across physical compliance, geometric fidelity, transient thermal accuracy, and computational efficiency. Physics compliance is diagnosed by spatially mapping the strong-form pointwise residual error, $\log_{10}(|\mathrm{Res}|)$, near phase-change interfaces and the moving heat source. Geometric fidelity is assessed by extracting the melt pool width (W), depth (D), and aspect ratio ($R = D/W$) from the quasi-steady liquidus isotherm ($T = T_l$). Deviations are quantified utilizing Mean Absolute Error (MAE) [40] and relative errors [41] against both experimental micrographs (to validate inverse calibration) and FEM benchmarks (to isolate algorithmic discrepancies). Holistic 3D transient thermal fields and enthalpy propagation are quantified against FEM references using the total relative L_2 error via nodal parity analysis and the temporal evolution of the domain-maximum temperature $T_{\max}(t)$. Finally, computational throughput—comprising offline training, inference latency, and total FEM solver times—is benchmarked on a unified hardware platform (NVIDIA RTX 4090 GPU, Intel Xeon Silver 4314 CPU, 270 GB RAM) and averaged over three independent runs to suppress stochastic variance.

4. Results and Discussion

4.1. Convergence and Residual Diagnostics

The comparative logarithmic loss histories in Figure 8 illustrate the convergence dynamics of both the standard strong-form PINN (Figure 8a) and the proposed VPINN (Figure 8b) under the sequential time-window strategy. While both frameworks exhibit a

repeating sawtooth pattern across temporal windows—driven by initial residual reduction under Adam followed by sharp final convergence under L-BFGS—their optimization behaviors diverge substantially. The pointwise residual formulation of the standard PINN introduces severe stiffness into the loss landscape, manifesting as highly unstable, oscillatory training trajectories that require nearly twice the number of global iterations to reach an acceptable threshold. In contrast, the VPINN mitigates these localized gradient spikes through weak-form integration over the finite test spaces, yielding a demonstrably smoother and more stable descent path. Underlying both convergence trajectories is the EMA-based inverse Dirichlet adaptive weighting scheme, which structurally governs the loss component interactions. Because the total loss is computed as a weighted sum where each component is assigned a dynamically adjusted coefficient less than unity, its plotted magnitude consistently appears smaller than $\mathcal{L}_{PDE}$ alone, even as the PDE residual remains the numerically dominant individual term. This mechanism sustains approximate parity among $\mathcal{L}_{PDE}$, $\mathcal{L}_{exp}$, $\mathcal{L}_{BC}$, and $\mathcal{L}_{IC}$ throughout, effectively preventing the numerically stiffer PDE residual from systematically suppressing boundary and data enforcement during the nonlinear diffusion phase. Supported by this balanced optimization environment, the specific loss components exhibit physically meaningful evolution: periodic $\mathcal{L}_{IC}$ resets at window boundaries decay monotonically in amplitude as scanning approaches quasi-steady state, confirming progressive thermal stabilization. Furthermore, the experimental loss $\mathcal{L}_{exp}$ descends along a trajectory closely tracking the total loss, reflecting the network's improving ability to satisfy the sparse melt pool boundary constraints as training progresses.

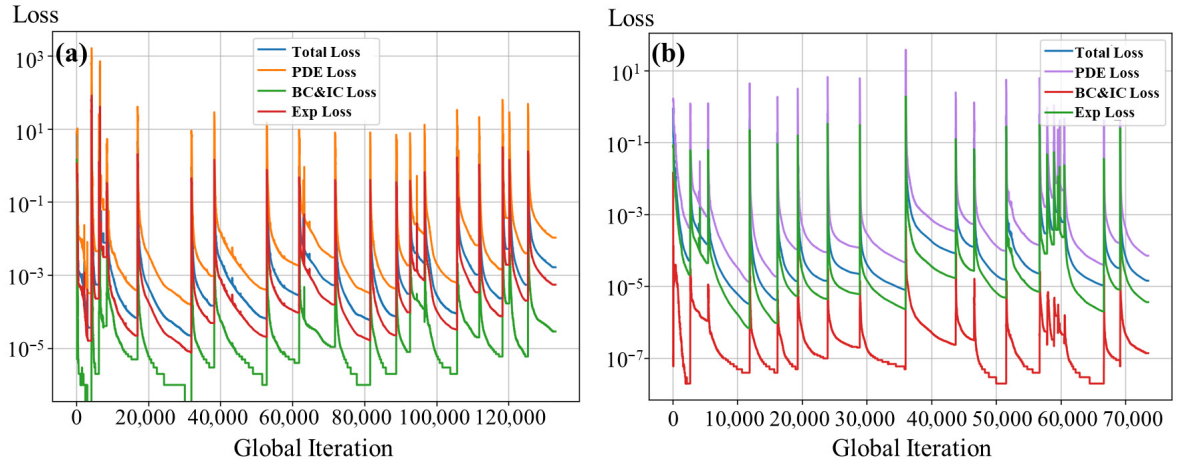


Figure 8. Evolution of the total, PDE, experimental, and boundary/initial-condition losses during time sliding window optimization: (**a**) PINN; (**b**) VPINN.

To assess near-source spatial resolution sensitivity, training loss histories for h_{min} = 0.01 mm and 0.04 mm were compared against the 0.02 mm baseline (Figure 9), holding far-field h_{max} = 0.2 mm. Neither alternative outperforms the baseline. Halving h_{min} to 0.01 mm induces a highly oscillatory training trajectory. The denser near-source quadrature over-samples the steep Goldak heat source and C_p enthalpy peaks, generating conflicting local gradients that destabilize the optimizer. Conversely, coarsening h_{min} to 0.04 mm yields a smoother descent but plateaus at a higher terminal loss. Here, insufficient quadrature points under-resolve the rapid spatial variations of $Q(\mathbf{x},t)$ and $C_p(T)$, embedding inherent approximation errors that further iterations cannot overcome. Thus, h_{min} = 0.02 mm optimally balances quadrature fidelity and optimization stability. Far-field h_{max} variation was excluded, as the accuracy-limiting stiffness is strictly confined to the near-source melt pool.

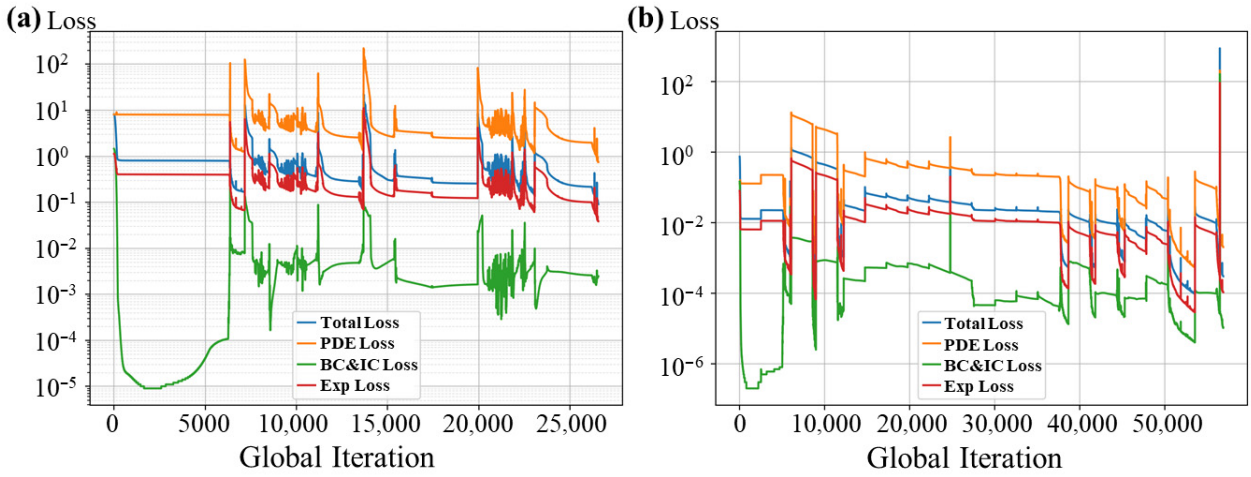


Figure 9. Sensitivity of VPINN training loss histories to near-source element size h_{min}: (**a**) h_{min} = 0.01 mm; (**b**) h_{min} = 0.04 mm.

The strong-form pointwise error $\log_{10}(|\mathrm{Res}|)$ evaluated at $N = 2000$ post-training diagnostic collocation points (strategically sampled with 50% allocated to the immediate vicinity of the heat source and 50% distributed across the far-field substrate) over the entire melt pool domain when the heat source is located at x = 0.75 mm and shown in Figure 10, remains uniformly low in the far-field substrate while concentrating within and immediately adjacent to the melt pool across all five process conditions. This localization reflects the coupled stiffness of the exponentially decaying Goldak source—which compresses the solution's active variation into a sub-millimeter neighborhood of the beam—with the Gaussian-augmented $C_p^*(T)$ nonlinearity at both phase-transition intervals. Although Octree quadrature refinement and heat-source-aware element weighting $\widetilde{w}_e$ concentrate optimization effort toward this zone, the residual floor there measurably exceeds the far-field level. Quantitatively, $\log_{10}(|\mathrm{Res}|) \approx -2$ in the near-source region corresponds to a local temperature relative error on the order of 3%, whereas $\log_{10}(|\mathrm{Res}|) \leq -5$ in the far field implies relative deviations well below 1%, confirming that the equation residual satisfies engineering accuracy requirements across the entire computational domain. This spatial contrast identifies the coupled phase-change and moving-source stiffness as the primary accuracy-limiting mechanism rather than boundary enforcement or inter-window temporal accumulation. This attribution follows directly from the equation structure. The Goldak source (Equation (4)), with semi-axes on the order of 0.05 mm, concentrates virtually all deposited power within a sub-millimeter neighborhood of the beam center. Simultaneously, the effective heat capacity $C_p^*(T)$ (Equation (10)) introduces enthalpy peaks of magnitude 2.7×10^{11} and 6.9×10^{12} mm^2 s^{-2} at the fusion and vaporization transitions. Because the liquidus isotherm coincides with the region of steepest source-induced gradient, both C_p^* and $\partial u/\partial t$ in the residual integrand (Equation (13)) are simultaneously maximized within the same narrow spatial band, producing multiplicative stiffening. In the far field, $Q \to 0$ exponentially and T remains sub-solidus, reducing the problem to smooth linear diffusion with negligible residual. The spatial superposition of these two independently stiff mechanisms accounts for the observed residual localization.

4.2. Calibration Outcome

Joint optimization of θ and $\phi_\eta = \eta$ is supervised exclusively by melt pool contour points from V01, V04, and V05, while V02 and V03 serve as unsupervised validation cases spanning the intermediate range of the velocity sweep. The calibrated η converges to 0.6728, a physically plausible value for GH3536 under conduction-to-transition irradia-

tion. Within the thermal conduction framework of Section 2.1, η should be interpreted as an apparent effective absorptivity that phenomenologically compensates for Marangoni convection, buoyancy-driven flow, and incipient vapor depression absent from the governing equations [42]. This compensatory role becomes more pronounced toward the high-power/low-speed edge of the present parameter matrix; extending ϕ to η, b over a substantially broader process window may be required to decouple energy coupling from geometric penetration.

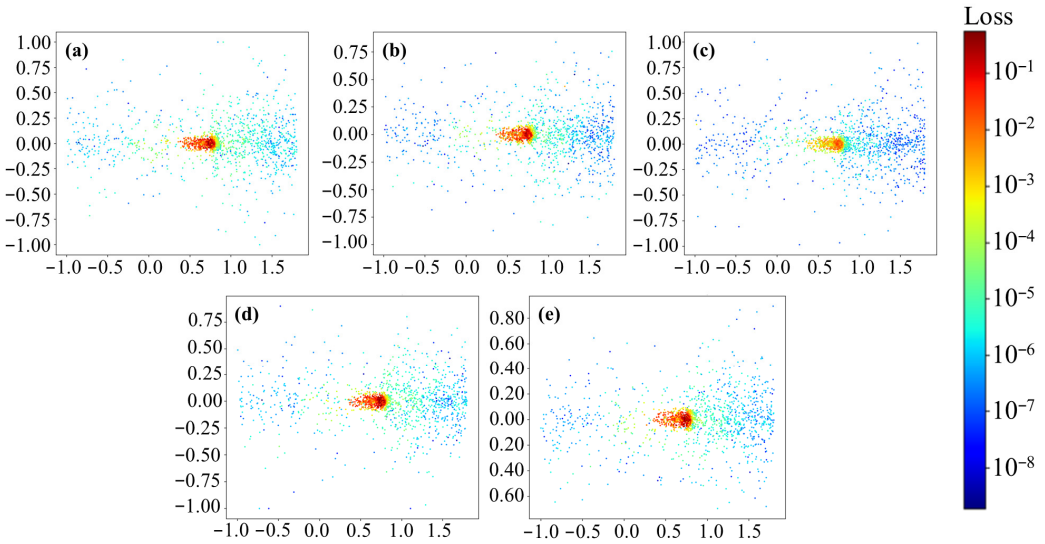


Figure 10. Point-wise spatial distribution of strong-form residuals over the entire melt pool domain when the heat source is located at x = 0.75 mm, for different power and speed combinations: (**a**) 100 W, 1000 mm/s; (**b**) 100 W, 1250 mm/s; (**c**) 100 W, 1500 mm/s; (**d**) 120 W, 1500 mm/s; (**e**) 140 W, 1500 mm/s.

For all three supervised conditions—Figure 11a,d,e—VPINN predictions of W, D, and $R = D/W$ extracted from the $T_l = 1628.15$ K isotherm at t_{ss} reproduce both experimental measurements and FEM reference solutions with close three-way agreement. Using the same globally calibrated η, the model captures the expected trend that increasing scan speed contracts the melt pool during the velocity sweep (V01–V03 at fixed P), while increasing laser power expands the melt pool during the power sweep (V03–V05 at fixed v). For the unsupervised conditions V02 and V03—Figure 11b,c, fully bracketed in (P, v) space by V01 and V04—predicted W and D track the FEM-consistent contraction trend along the velocity sweep within acceptable deviations; R preserves the correct inter-condition ordering across all five cases, with the residual inter-method discrepancy attributable to the absence of geometric supervision at intermediate velocities and the asymmetric sensitivity of Gaussian-regularized $C_p^*(T)$ to depth-direction isotherm displacement. The three-condition calibration thus sustains predictive fidelity across the complete five-condition parametric matrix without condition-specific re-tuning.

4.3. Transient Thermal Accumulation and Causal Training Effects

The sequential time-window strategy reproduces two hallmark signatures of single-track thermal accumulation that global space-time minimization systematically suppresses. For P = 100 W, v = 1000 mm/s (Figure 12), the trailing thermal wake broadens progressively across successive windows: rising far-field temperature ahead of the advancing beam elevates the effective preheat state, producing a quasi-steady melt pool measurably wider and deeper than the early-scan transient—a feature recoverable only when enthalpy history is propagated forward as a hard window-boundary constraint rather than re-optimized

globally. The domain-maximum temperature $T_{max}(t)$ provides a complementary scalar diagnostic: the predicted steep ascent from ambient decelerates into an asymptotic plateau whose onset coincides precisely with the decay of periodic window-reset amplitude in the loss history (Figure 8). The VPINN reproduces both the transient ascent rate and the asymptotic level within FEM accuracy bounds, confirming that the causal sliding-window mechanism faithfully resolves the full multi-timescale thermal response without accumulating temporal drift across window interfaces.

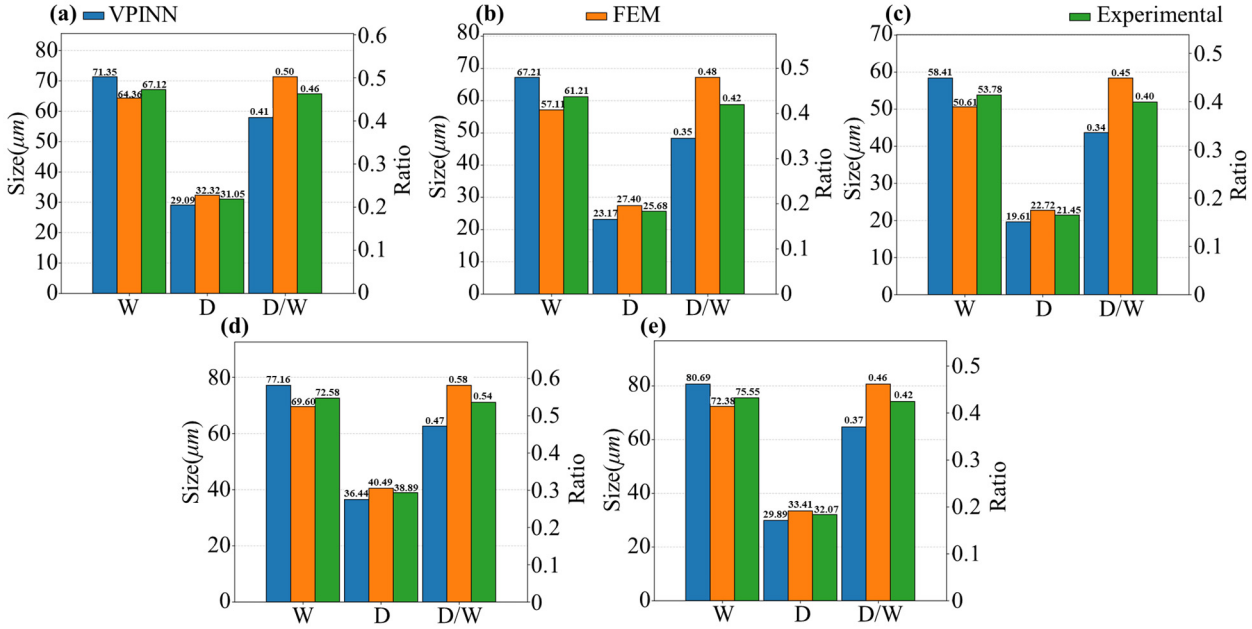


Figure 11. Comparison of predicted melt pool width (W), depth (D), and aspect ratio (D/W) among VPINN, FEM benchmark, and experimental measurements for different process conditions: (**a**) 100 W, 1000 mm/s; (**b**) 100 W, 1250 mm/s; (**c**) 100 W, 1500 mm/s; (**d**) 120 W, 1500 mm/s; (**e**) 140 W, 1500 mm/s.

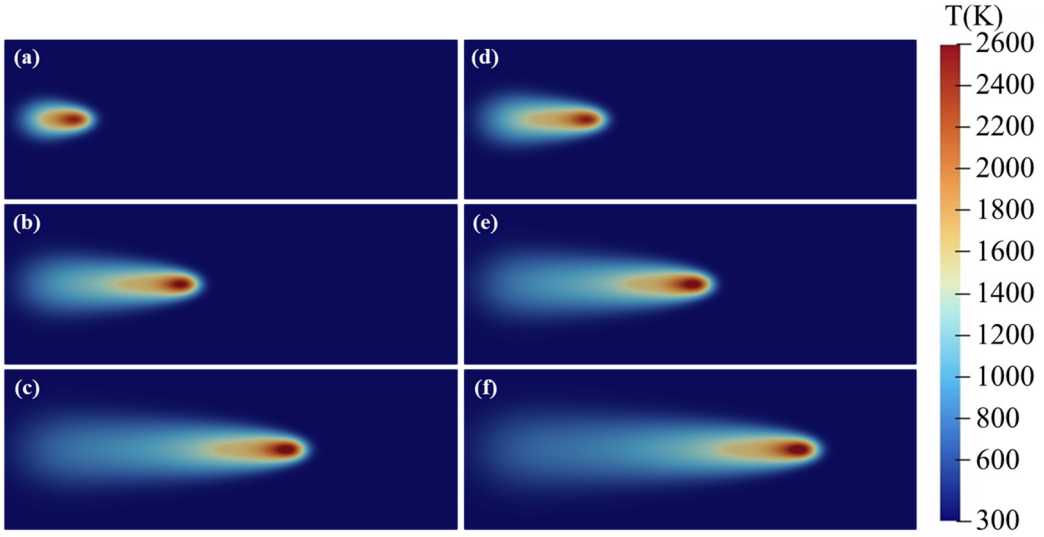


Figure 12. Transient temperature field evolution during single-track scanning (P = 100 W, v = 1000 mm/s) at successive time steps: (**a**) t = 0.0002 s; (**b**) t = 0.0004 s; (**c**) t = 0.0006 s; (**d**) t = 0.0008 s; (**e**) t = 0.0010 s; (**f**) t = 0.0012 s.

4.4. FEM Comparison Across Five Process Conditions

Because the two frameworks share identical governing equations, thermophysical representations, calibrated Goldak parameters, and boundary conditions, the VPINN–FEM

comparison certifies algorithmic fidelity—that the variational surrogate solves the shared mathematical problem with accuracy competitive with implicit FEM—rather than constituting independent physical validation, which is separately established by the experimental isotherm comparisons of Section 4.2.

Across the velocity sweep (V01–V03, fixed $P = 100$ W) and the power sweep (V03–V05, fixed $v = 1500$ mm/s), the VPINN correctly captures the progressive contraction of the near-source thermal envelope with increasing scan speed (Figure 13) and the monotonic isosurface expansion with increasing laser power (Figure 14), without condition-specific re-tuning of ϕ. In both sweeps, elevated relative errors are confined to the immediate melt-pool vicinity and trailing wake, while far-field regions remain below a few percent—spatial error topology fully consistent with the residual localization reported in Section 4.1 and attributable to the irreducible coupled stiffness of the moving source and phase-change nonlinearity rather than to boundary enforcement or inter-window accumulation.

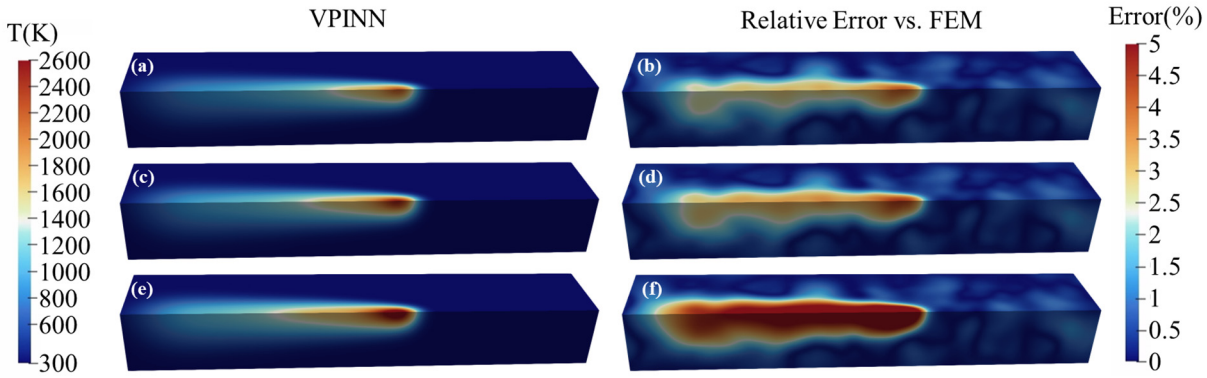


Figure 13. 3D VPINN temperature field predictions and pointwise relative errors against FEM at a constant scanning speed of v = 1500 mm/s: (**a**,**b**) P = 100 W; (**c**,**d**) P = 120 W; (**e**,**f**) P = 140 W.

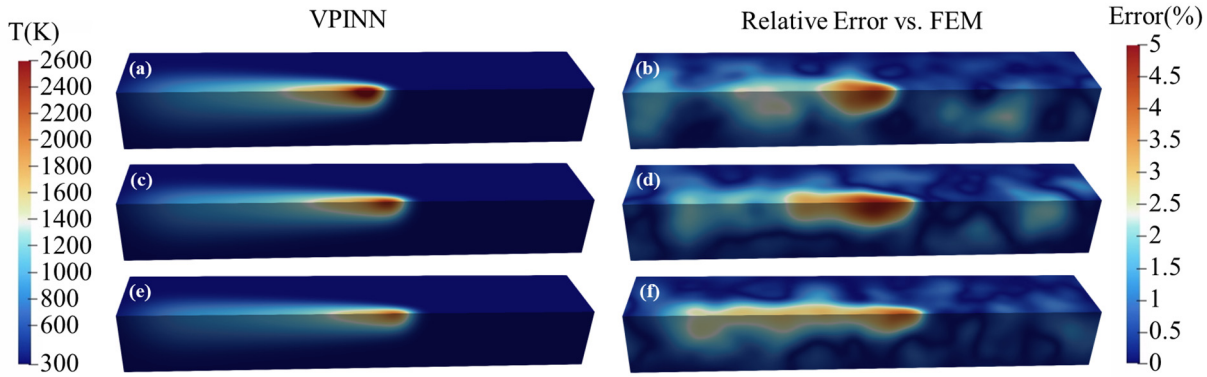


Figure 14. 3D VPINN temperature field predictions and pointwise relative errors against FEM at a constant laser power of P = 100 W: (**a**,**b**) v = 1000 mm/s; (**c**,**d**) v = 1250 mm/s; (**e**,**f**) v = 1500 mm/s.

Nodal parity analysis across all five conditions, shown in Figure 15, yields total relative L_2 errors on the order of a few percent, with no systematic inter-condition trend in error magnitude or distribution across either the velocity or power sweep. Deviations from the parity diagonal increase exclusively at the highest temperatures, consistent with the elevated sensitivity of the solution near the heat-source core and the $C_p^*(T)$ phase-transition peaks, and no sub-liquidus bias is detectable across the parametric matrix, confirming accurate representation of the conduction-dominated far field throughout.

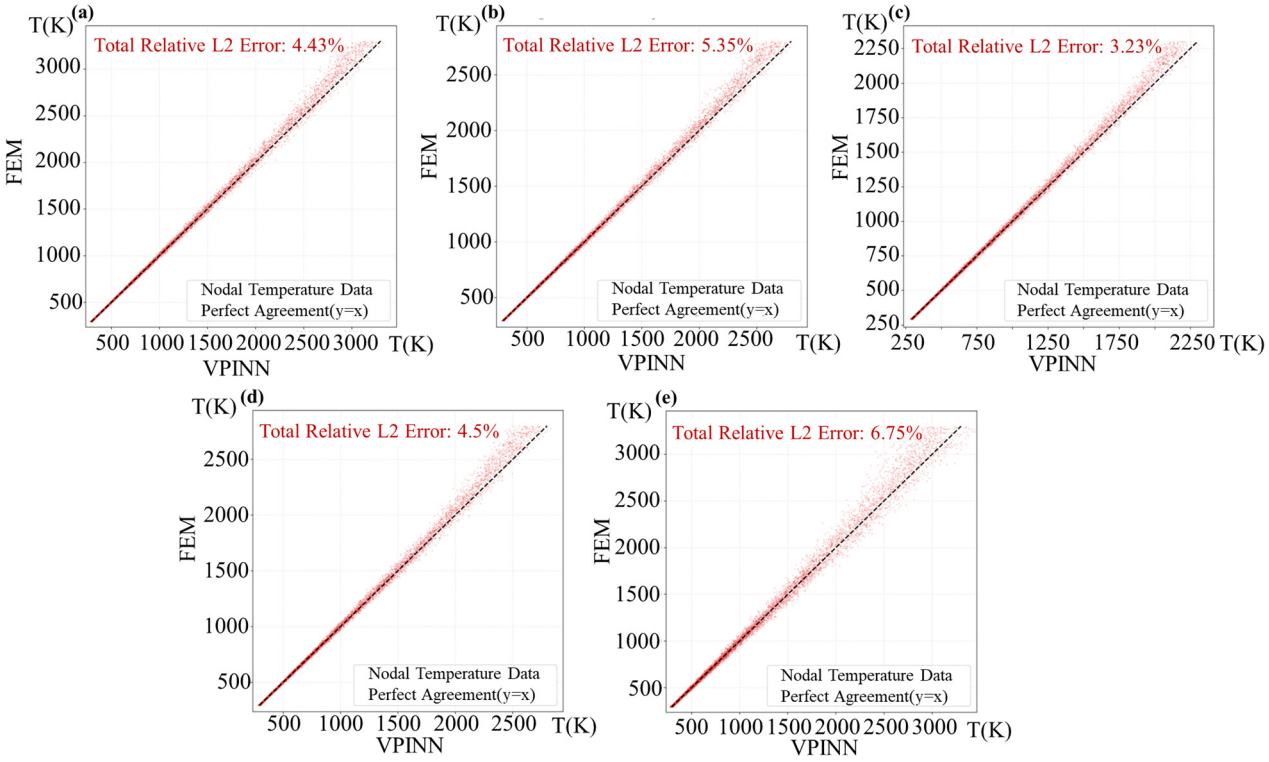


Figure 15. Nodal temperature parity analysis between VPINN predictions and FEM references across different power and speed combinations: (**a**) 100 W, 1000 mm/s; (**b**) 100 W, 1250 mm/s; (**c**) 100 W, 1500 mm/s; (**d**) 120 W, 1500 mm/s; (**e**) 140 W, 1500 mm/s.

4.5. Hyperparameter Sensitivity and Computational Efficiency

Figure 16a evaluates the sequential time-marching hyperparameters: window width $\Delta\tau$ and sliding stride $\delta\tau$. Configurations lacking temporal overlap $\delta\tau \geq \Delta\tau$ fail to propagate latent heat enthalpy history, causing divergence or unacceptably high terminal losses. Conversely, $\Delta\tau = 0.30$ s exhausts GPU memory due to excessive spatiotemporal quadrature points. Reliable convergence uniquely requires a 50% overlap ($\delta\tau = 0.05$ s). The configuration ($\Delta\tau = 0.10$ s, $\delta\tau = 0.05$ s) converges in ~70,000 epochs (Loss $< 1.37 \times 10^{-3}$), whereas $\Delta\tau = 0.20$ s requires 2.3 times more iterations without meaningful accuracy gains. Thus, ($\Delta\tau = 0.10$ s, $\delta\tau = 0.05$ s) was adopted to optimally balance melt-pool transient resolution, stable enthalpy transfer, and computational cost.

Figure 16b examines the Gaussian smoothing parameter σ for the effective heat capacity $C_p^*(T)$ (Equation (10)). Decreasing σ narrows the phase-change temperature band and amplifies the enthalpy peak, intensifying the localized stiffness of the residual integrand (Equation (13)). Below $\sigma \approx 15$, this ill-conditioned loss landscape causes both VPINN and PINN to diverge. Above this threshold, VPINN convergence iterations decrease progressively as σ increases (from ~200,000 at $\sigma = 15$ to <100,000 at $\sigma = 30$), yielding relative L_2 errors between 4.85% and 6.30%. PINNs demand a higher σ threshold and consistently require more iterations than VPINNs at matched σ, highlighting the gradient-smoothing advantage of variational weak-form integration over pointwise strong-form evaluation. However, excessively large σ introduces systematic modeling errors by unphysically smearing the latent heat contribution into the sub-solidus conduction regime.

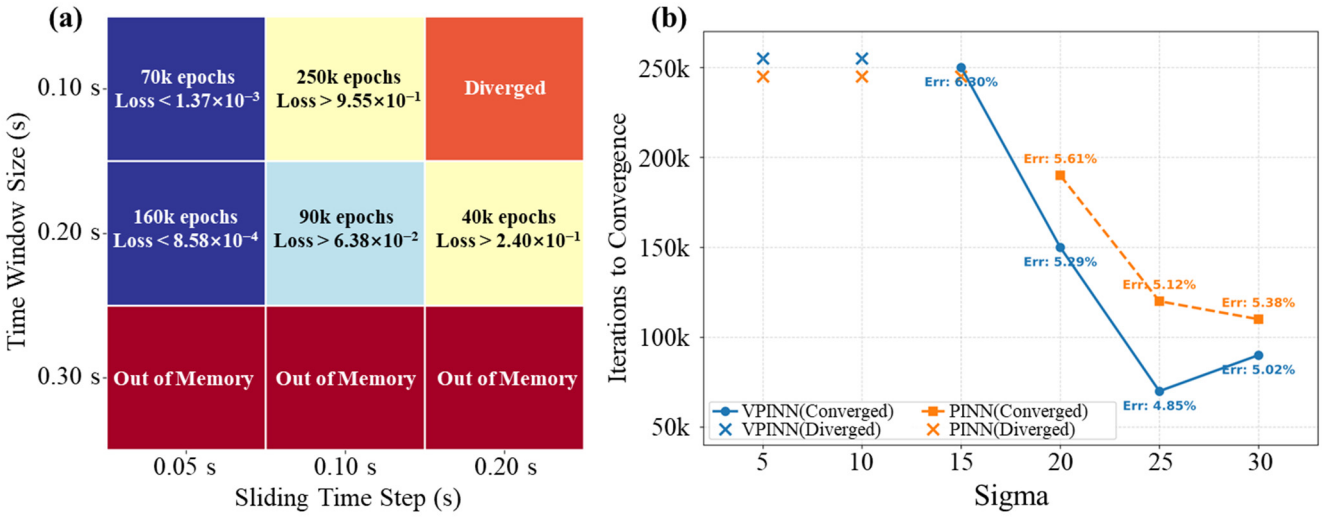


Figure 16. Hyperparameter sensitivity analysis: (**a**) Training outcome matrix evaluating combinations of window width $\Delta\tau$ and sliding stride $\delta\tau$. (**b**) Effect of the Gaussian smoothing parameter σ on convergence iterations for VPINN and PINN, with relative L_2 errors annotated for converged cases.

All wall-clock times were recorded on a single NVIDIA RTX 4090 GPU (Intel Xeon Silver 4314, 2.40 GHz, 270 GB RAM) and averaged over three runs. Because (P, v) enter the solution network as explicit input coordinates, a single training run—totaling 15,323 s across all temporal windows and all five process conditions—resolves the complete parametric space, amortizing the offline investment over every subsequent query. Per-case FEM wall-clock time on the same platform ranges from 3000 to 4000 s; surrogate inference for a full scan-trajectory field reconstruction requires 4–5 s per condition, yielding a per-query speedup of two to three orders of magnitude. The amortized cost per evaluated condition therefore decreases monotonically with the number of queries, crossing below the FEM baseline at 5 conditions, as quantified in Figure 17. For the five-condition parametric matrix examined here, the trained surrogate delivers the complete set of thermal field predictions at a fraction of the cumulative FEM cost, establishing the framework as computationally competitive for high-throughput parametric exploration and real-time digital twin integration [43] where the trained model is repeatedly queried at new (P, v) operating points without retraining.

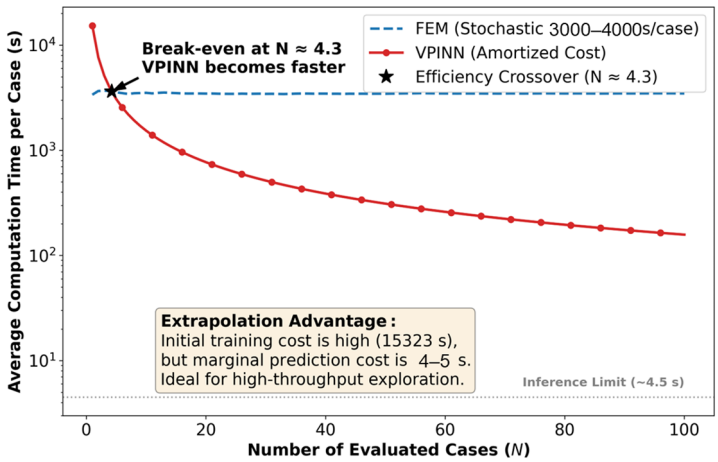


Figure 17. Computational efficiency analysis illustrating the amortized per-case evaluation cost of VPINN compared to the FEM baseline.

5. Conclusions

In this study, a VPINN framework was developed for three-dimensional transient melt pool thermal modeling in a GH3536 single-track LPBF process. The framework successfully integrates a continuously differentiable Goldak double-ellipsoid moving heat source, temperature-dependent thermophysical property surrogates, and an effective heat-capacity treatment of latent heat within a weak-form residual-minimization scheme. The principal findings and the necessary modeling simplifications of this research are summarized as follows:

High-Fidelity Geometric Prediction via Inverse Calibration: Using sparse experimental melt pool profiles for supervision, the framework successfully inferred an effective absorptivity parameter. Across a five-condition parametric matrix spanning laser power, the predicted melt pool width, depth, and aspect ratio agreed closely with cross-sectional optical micrograph measurements, maintaining accuracy across two held-out conditions (V02 and V03) that received no geometric supervision during training, though both lie within or near the boundary of the calibrated parameter range. By operating within a pure continuum thermal conduction framework, the model achieves profound computational efficiency. In this context, the calibrated effective absorptivity acts as a phenomenological lump parameter, elegantly compensating for unmodeled fluid dynamics while maintaining geometric predictive accuracy.

Nodal Temperature Accuracy vs. FEM: The VPINN surrogate demonstrated strong algorithmic fidelity in resolving steep moving-source thermal gradients. Nodal temperature parity analysis revealed total relative L_2 errors ranging from 3.23% to 6.75% across the varied power and speed conditions when compared directly against the implicit FEM benchmarks.

Significant Computational Acceleration: Surrogate inference reduced the per-condition query time to merely 4–5 s, delivering a speedup of two to three orders of magnitude relative to the 3000–4000 s required per FEM evaluation. Although the model requires an initial offline training investment of 15,323 s, the amortized computational cost breaks even with the FEM baseline after evaluating approximately five conditions. Given this current spatial and temporal scope, the model is highly optimized for isolated single-track analysis, which unavoidably excludes the inter-track and interlayer thermal accumulation effects required for larger-scale simulations.

Future work should build upon this efficient mathematical framework by prioritizing the integration of thermally coupled fluid dynamics to smoothly extend the model's mechanistic validity into keyhole-mode operation. Furthermore, expanding the methodology to encompass multi-track and multi-layer thermal accumulation regimes, alongside the simultaneous identification of multiple heat source geometry parameters using richer experimental supervision, will further enhance the model's capacity to independently constrain complex melt pool morphologies across a broader process window.

Author Contributions: Conceptualization, Z.X. and X.W. (Xin Wang); methodology, Z.X., X.W. (Xin Wang) and Y.M.; software, Z.X. and M.W.; validation, Z.X., M.W. and X.W. (Xianglong Wang); formal analysis, Z.X. and Y.M.; investigation, Z.X. and X.W. (Xianglong Wang); resources, X.W. (Xin Wang) and Y.M.; data curation, Z.X. and X.W. (Xianglong Wang); writing—original draft preparation, Z.X.; writing—review and editing, X.W. (Xin Wang), Y.M., M.W. and X.W. (Xianglong Wang); visualization, Z.X. and M.W.; supervision, X.W. (Xin Wang) and Y.M.; project administration, X.W. (Xin Wang) and Y.M.; funding acquisition, X.W. (Xin Wang) and Y.M. All authors have read and agreed to the published version of the manuscript.

Funding: This work was financially supported by the National Major Science and Technology Projects of China. (Grant No. 2024ZD0608100), the National Natural Science Foundation of China (Grant No. 52405437), the China Postdoctoral Science Foundation (Grant No. 2024M761107), the Beijing Natural Science Foundation (Grant No. L245024), the Technology Development Project of China

Academy of Machinery Science & Technology Group Co., Ltd., and the Zhengzhou Outstanding Young Scientific and Technological Talent Program.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The original contributions presented in the study are included in the article; further inquiries can be directed to the corresponding author.

Conflicts of Interest: Authors Zhenghao Xu, Xin Wang, Yuan Meng, Mingwei Wang, and Xianglong Wang were employed by China Machinery Institute of Advanced Materials (Zhengzhou) Co., Ltd., Beijing National Innovation Institute of Lightweight Co.,Ltd., and China Academy of Machinery Science and Technology Group. The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

1. Trejos-Taborda, J.; Reyes-Osorio, L.; Garza, C.; del Carmen Zambrano-Robledo, P.; Lopez-Botello, O. Finite Element Modeling of Melt Pool Dynamics in Laser Powder Bed Fusion of 316L Stainless Steel. *Int. J. Adv. Manuf. Technol.* **2022**, *120*, 3947–3961. [CrossRef]
2. Psihoyos, H.O.; Lampeas, G.N. Density-Based Optimization of the Laser Powder Bed Fusion Process Based on a Modelling Framework. *Alloys* **2023**, *2*, 55–76. [CrossRef]
3. Zhang, T.; Lin, X.; Qin, Y.; Zhu, D.; Wang, J.; Zhang, C.; Bai, Y. Multiphysics Modeling of Heat Transfer and Melt Pool Thermo-Fluid Dynamics in Laser-Based Powder Bed Fusion of Metals. *Materials* **2025**, *18*, 3183. [CrossRef] [PubMed]
4. Iveković, A.; Montero-Sistiaga, M.L.; Vleugels, J.; Kruth, J.-P.; Vanmeensel, K. Crack Mitigation in Laser Powder Bed Fusion Processed Hastelloy X Using a Combined Numerical-Experimental Approach. *J. Alloys Compd.* **2021**, *864*, 158803. [CrossRef]
5. Kusano, M.; Watanabe, M. Heat Source Model Development for Thermal Analysis of Laser Powder Bed Fusion Using Bayesian Optimization and Machine Learning. *Integr. Mater. Manuf. Innov.* **2024**, *13*, 288–304. [CrossRef]
6. Hosseini, E.; Scheel, P.; Müller, O.; Molinaro, R.; Mishra, S. Single-Track Thermal Analysis of Laser Powder Bed Fusion Process: Parametric Solution through Physics-Informed Neural Networks. *Comput. Methods Appl. Mech. Eng.* **2023**, *410*, 116019. [CrossRef]
7. Elahi, S.M.; Leonor, J.P.; Wu, R.Y.; Wagner, G.J. Efficient Part-Scale Thermal Modeling of Laser Powder Bed Fusion via a Multilevel Finite Element Framework. *Addit. Manuf.* **2025**, *109*, 104897. [CrossRef]
8. Hemmasian, A.; Ogoke, F.; Akbari, P.; Malen, J.; Beuth, J.; Barati Farimani, A. Surrogate Modeling of Melt Pool Temperature Field Using Deep Learning. *Addit. Manuf. Lett.* **2023**, *5*, 100123. [CrossRef]
9. Raissi, M.; Perdikaris, P.; Karniadakis, G.E. Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations. *J. Comput. Phys.* **2019**, *378*, 686–707. [CrossRef]
10. Karniadakis, G.E.; Kevrekidis, I.G.; Lu, L.; Perdikaris, P.; Wang, S.; Yang, L. Physics-Informed Machine Learning. *Nat. Rev. Phys.* **2021**, *3*, 422–440. [CrossRef]
11. Liao, S.; Xue, T.; Jeong, J.; Webster, S.; Ehmann, K.; Cao, J. Hybrid Thermal Modeling of Additive Manufacturing Processes Using Physics-Informed Neural Networks for Temperature Prediction and Parameter Identification. *Comput. Mech.* **2023**, *72*, 499–512. [CrossRef]
12. Jiang, F.; Xia, M.; Hu, Y. Physics-Informed Machine Learning for Accurate Prediction of Temperature and Melt Pool Dimension in Metal Additive Manufacturing. *3D Print. Addit. Manuf.* **2024**, *11*, e1679–e1689. [CrossRef] [PubMed]
13. Sharma, R.; Guo, Y.B.; Raissi, M.; Guo, W.G. Physics-Informed Machine Learning of Argon Gas-Driven Melt Pool Dynamics. *J. Manuf. Sci. Eng.* **2024**, *146*, 081008. [CrossRef]
14. Zhu, Q.; Lu, Z.; Hu, Y. Transfer Learning-Enhanced Physics Informed Neural Network for Accurate Melt Pool Prediction in Laser Melting. *Adv. Manuf.* **2025**, *2*, 0001. [CrossRef]
15. Zhu, Q.; Lu, Z.; Hu, Y. A Reality-Augmented Adaptive Physics Informed Machine Learning Method for Efficient Heat Transfer Prediction in Laser Melting. *J. Manuf. Process.* **2024**, *124*, 444–457. [CrossRef]
16. Safari, H.; Wessels, H. Physics-Informed Surrogates for Temperature Prediction of Multi-Tracks in Laser Powder Bed Fusion. *arXiv* **2025**. [CrossRef]
17. Kharazmi, E.; Zhang, Z.; Karniadakis, G.E. Variational Physics-Informed Neural Networks for Solving Partial Differential Equations. *arXiv* **2019**. [CrossRef]
18. Kharazmi, E.; Zhang, Z.; Karniadakis, G.E.M. Hp-VPINNs: Variational Physics-Informed Neural Networks with Domain Decomposition. *Comput. Methods Appl. Mech. Eng.* **2021**, *374*, 113547. [CrossRef]
19. Goldak, J.; Chakravarti, A.; Bibby, M. A New Finite Element Model for Welding Heat Sources. *Metall. Trans. B* **1984**, *15*, 299–305. [CrossRef]

20. Lu, J.; Zheng, H.; Ji, X.; Guan, Y.; Wang, Z.; Cheng, J.; Zhang, W. Crack Characteristics Analysis and Mechanisms in GH3536 Alloy Manufactured by Laser Powder Bed Fusion. *Eng. Fail. Anal.* **2024**, *162*, 108382. [CrossRef]
21. Wang, S.; Teng, Y.; Perdikaris, P. Understanding and Mitigating Gradient Flow Pathologies in Physics-Informed Neural Networks. *SIAM J. Sci. Comput.* **2021**, *43*, A3055–A3081. [CrossRef]
22. Voller, V.R.; Swaminathan, C.R.; Thomas, B.G. Fixed Grid Techniques for Phase Change Problems: A Review. *Int. J. Numer. Methods Eng.* **1990**, *30*, 875–898. [CrossRef]
23. Patel, R.G.; Manickam, I.; Trask, N.A.; Wood, M.A.; Lee, M.; Tomas, I.; Cyr, E.C. Thermodynamically Consistent Physics-Informed Neural Networks for Hyperbolic Systems. *J. Comput. Phys.* **2022**, *449*, 110754. [CrossRef]
24. Tavakoli, R. Smooth Modeling of Solidification Based on the Latent Heat Evolution Approach. *Int. J. Adv. Manuf. Technol.* **2017**, *88*, 3041–3052. [CrossRef]
25. Tancik, M.; Srinivasan, P.; Mildenhall, B.; Fridovich-Keil, S.; Raghavan, N.; Singhal, U.; Ramamoorthi, R.; Barron, J.; Ng, R. Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains. In *Proceedings of the Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Red Hook, NY, USA, 2020; Volume 33, pp. 7537–7547.
26. Wang, S.; Wang, H.; Perdikaris, P. On the Eigenvector Bias of Fourier Feature Networks: From Regression to Solving Multi-Scale PDEs with Physics-Informed Neural Networks. *Comput. Methods Appl. Mech. Eng.* **2021**, *384*, 113938. [CrossRef]
27. Sitzmann, V.; Martel, J.; Bergman, A.; Lindell, D.; Wetzstein, G. Implicit Neural Representations with Periodic Activation Functions. In *Proceedings of the Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Red Hook, NY, USA, 2020; Volume 33, pp. 7462–7473.
28. Burstedde, C.; Wilcox, L.C.; Ghattas, O. P4est: Scalable Algorithms for Parallel Adaptive Mesh Refinement on Forests of Octrees. *SIAM J. Sci. Comput.* **2011**, *33*, 1103–1133. [CrossRef]
29. Golub, G.H.; Welsch, J.H. Calculation of Gauss Quadrature Rules. *Math. Comput.* **1969**, *23*, 221–230. [CrossRef]
30. Liu, Y.; Cai, L.; Chen, Y.; Wang, B. Physics-Informed Neural Networks Based on Adaptive Weighted Loss Functions for Hamilton-Jacobi Equations. *Math. Biosci. Eng.* **2022**, *19*, 12866–12896. [CrossRef] [PubMed]
31. McClenny, L.D.; Braga-Neto, U.M. Self-Adaptive Physics-Informed Neural Networks. *J. Comput. Phys.* **2023**, *474*, 111722. [CrossRef]
32. Wang, S.; Sankaran, S.; Perdikaris, P. Respecting Causality for Training Physics-Informed Neural Networks. *Comput. Methods Appl. Mech. Eng.* **2024**, *421*, 116813. [CrossRef]
33. Zhang, Z. Improved Adam Optimizer for Deep Neural Networks. In Proceedings of the 2018 IEEE/ACM 26th International Symposium on Quality of Service (IWQoS), Banff, AB, Canada, 4–6 June 2018; pp. 1–2.
34. Liu, D.C.; Nocedal, J. On the Limited Memory BFGS Method for Large Scale Optimization. *Math. Program.* **1989**, *45*, 503–528. [CrossRef]
35. Chen, N.; Zheng, D.; Niu, P.; Li, R.; Yuan, T. Laser Powder Bed Fusion of GH3536 Nickel-Based Superalloys: Processing Parameters, Microstructure and Mechanical Properties. *Mater. Charact.* **2023**, *202*, 113018. [CrossRef]
36. Weissbach, R.; Penny, R.W.; Hart, A.J. A Rapid Experimental Workflow for Studying Melt Track Scaling in Laser Powder Bed Fusion Using High-Precision Metal Template Substrates. *Prog. Addit. Manuf.* **2025**, *10*, 8695–8706. [CrossRef]
37. Patel, S.; Vlasea, M. Melting Modes in Laser Powder Bed Fusion. *Materialia* **2020**, *9*, 100591. [CrossRef]
38. Tang, J.; Wróbel, R.; Scheel, P.; Gaechter, W.; Leinenbach, C.; Hosseini, E. The Role of Process Parameters and Printing Position on Meltpool Variations in LPBF Hastelloy X: Insights into Laser-Plume Interaction. *Addit. Manuf. Lett.* **2024**, *9*, 100203. [CrossRef]
39. Zienkiewicz, O.C.; Taylor, R.L.; Zhu, J.Z. *The Finite Element Method: Its Basis and Fundamentals*; Butterworth-Heinemann: Oxford, UK, 2013; ISBN 978-1-85617-633-0.
40. Willmott, C.J.; Matsuura, K. Advantages of the Mean Absolute Error (MAE) over the Root Mean Square Error (RMSE) in Assessing Average Model Performance. *Clim. Res.* **2005**, *30*, 79–82. [CrossRef]
41. de Myttenaere, A.; Golden, B.; Le Grand, B.; Rossi, F. Mean Absolute Percentage Error for Regression Models. *Neurocomputing* **2016**, *192*, 38–48. [CrossRef]
42. Ren, Z.; Gao, L.; Clark, S.J.; Fezzaa, K.; Shevchenko, P.; Choi, A.; Everhart, W.; Rollett, A.D.; Chen, L.; Sun, T. Machine Learning–Aided Real-Time Detection of Keyhole Pore Generation in Laser Powder Bed Fusion. *Science* **2023**, *379*, 89–94. [CrossRef] [PubMed]
43. Łach, Ł.; Svyetlichnyy, D. Advanced Numerical Modeling of Powder Bed Fusion: From Physics-Based Simulations to AI-Augmented Digital Twins. *Materials* **2026**, *19*, 426. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.