

Enabling geometry and toolpath generalization for machine learning based thermal modeling in laser powder bed fusion

Kahraman G. Demir¹, Tarek Zohdi¹, Grace X. Gu^{1*}

¹*Department of Mechanical Engineering, University of California, Berkeley*

Abstract

A central challenge in machine learning-based thermal modeling for Laser Powder Bed Fusion (LPBF) is achieving reliable generalization across different part geometries and toolpaths. These capabilities are crucial for practical manufacturing, in order to avoid defects caused by poor process planning. Finite element simulations provide relatively accurate results but are too computationally intensive, highlighting the need for fast, generalizable surrogate models. This work presents a U-Net Convolutional Neural Network (CNN) that addresses the generalization challenge by leveraging intelligent feature engineering utilizing signed distance fields to represent geometry, time fields to capture laser scanning dynamics, and the time gradient field to encode heat diffusion patterns. This enables a single trained model to generalize effectively to new part shapes and toolpath orientations. The resulting surrogate model achieves around 1000 times speedup relative to finite element analysis while maintaining high accuracy, enabling real-time thermal prediction for process optimization and intelligent toolpath planning in additive manufacturing.

Keywords: Additive Manufacturing, Laser Powder Bed Fusion, Deep Learning, Temperature Prediction, Thermal Physics Features

Introduction

Laser Powder Bed Fusion (LPBF) has emerged as one of the most versatile and widely adopted metal additive manufacturing technologies, enabling the production of complex geometries with excellent mechanical properties across diverse engineering applications [1]. In LPBF, a high-power laser selectively melts regions of a powder bed according to the cross-section of the part being manufactured. The process involves complex multi-physics phenomena, including rapid heating and cooling cycles, phase transformations,

*Corresponding author: ggu@berkeley.edu

and material consolidation that occur at length scales spanning several orders of magnitude. Central to part quality is the local temperature of the meltpool and the global temperature distribution of the part being made during the LPBF process, both of which are critical in determining microstructure evolution, residual stress formation, and potential defects such as lack of fusion, keyholing, porosity, and surface anomalies [2]. This drives the need to develop methods to control for meltpool temperature and excessive thermal gradients.

This need has driven significant research into process control and monitoring systems. In-situ process monitoring and closed-loop control systems have been developed to dynamically adjust laser parameters based on thermal feedback [3–5]. To name a few works, Kim et al. demonstrated Kalman filtering for thermal state estimation in FDM processes using infrared camera data to improve thermal field monitoring accuracy [6]. Scanning strategy optimization has emerged as particularly critical for thermal management, with recent work demonstrating the significant impact of toolpath orientation on dimensional accuracy and the development of adaptive toolpath generation algorithms for distortion reduction [7–9]. Integration of machine learning frameworks with process optimization has shown promise for autonomous correction systems using computer vision and deep learning in fused filament fabrication (FFF) and broader AM contexts [10, 11]. While understanding and predicting temperature fields is essential for process optimization and quality control, the extreme thermal gradients (up to 10^6 K/m) and rapid cooling rates (up to 10^6 K/s) that characterize the subsurface thermal history make accurate temperature measurement and modeling challenging [12].

Conventional numerical methods enable the measurement and assessment of temperature fields, though they typically require some modeling simplifications. These methods vary from high-fidelity multi-physics solvers to simplified track- and layer-scale models, which capture detailed physics at different computational costs [13]. For instance, Zhang et al. developed a parallel, adaptive 3D heat-transfer framework specifically for LPBF [14]. Similarly, Cheng et al. proposed a Gaussian Process-based extension of the Goldak heat source for FE simulation [15]. In the realm of selective laser sintering (SLS), discrete element approaches provide insight into powder-scale phenomena through coupled formulations [16, 17]. To enhance computational efficiency for specific problems, simplified track- and part-scale analyses have been presented by researchers such as Sharma and Kumar [18], Carraturo et al. [19], and Psihoyos and Lampeas [20]. Furthermore, semi-analytical strategies, which include discrete Green’s functions and spectral graph formulations as demonstrated by Cole et al., offer computationally efficient thermal modeling. These are complemented by grid-free Monte Carlo variants that tackle complex boundary conditions, as shown by Sawhney et al. [21, 22]. Graph-theoretic thermal modeling related to DED was also demonstrated by Riensche et al. [23]. Additionally, Zohdi showcased rapid discrete ray-tracking and knowledge-based heat-kernel methods

that further accelerate laser-material thermal computations, including ultra-fast results using heat-kernel superposition [24–26].

Machine learning techniques have become powerful alternatives that effectively complement conventional methods. Early implementations of neural networks in solving differential equations were conducted by Dissanayake and Lagaris, who demonstrated that feedforward networks could approximate PDE solutions through collocation-based training [27, 28]. Building on these foundations and a myriad of successful efforts since, two influential modern directions have emerged: physics-informed learning, which enforces governing equations during training (e.g., PINNs) [29], and operator learning, which maps problem inputs to full solution fields (e.g., Fourier and physics-informed neural operators) [30, 31]. Additionally, their amalgamations, such as graph-based approaches that learn spatiotemporal dynamics with physics-informed GNNs [32, 33], are gaining traction. In metal additive manufacturing, these methods have been used for applications like temperature prediction, parameter identification [34–37], and toolpath optimization [38]. A prominent theme is geometry representation for cross-domain generalization, encompassing spectral embeddings that capture domain topology [39], implicit field representations with neural operators [40], hybrid graph-mesh approaches [41], direct mesh conditioning [42], learned geometric transformations [43], and explicit geometry encoders [44, 45]. There is a growing need for models that can generalize across different geometries and, in the LPBF context, toolpaths as well.

However, achieving simultaneous generalization across complex geometries and arbitrary toolpaths remains a formidable barrier. The difficulty arises because the thermal field in LPBF is not a static function of geometry alone, but a dynamic response to the specific time-history of the energy input. Conventional Convolutional Neural Networks (CNNs) excel at processing spatial data but often lack the mechanisms to interpret the sequential, vector-based nature of toolpaths relative to the part boundaries. Furthermore, the physics of heat diffusion implies that local thermal behavior is conditioned by the global geometry which varies drastically between different parts. Existing surrogate models often circumvent these issues by fixing the scan strategy or limiting the geometric design space, thereby sacrificing flexibility. To enable true design freedom, a modeling framework must explicitly encode both the spatial constraints of the geometry and the spatiotemporal dynamics of the laser source into a unified format suitable for convolutional networks.

In this work we tackle these needs by introducing a method of encoding geometric and toolpath information for use in machine learning. We utilize signed distance fields (SDF) to encode the part’s geometry which directly pertains to the physics of heat conduction. Smaller features exhibit higher thermal resistance to the surroundings compared to bulkier regions, which have more thermal mass and are able to remain relatively cooler. Smaller features would have smaller SDF values, and vice versa.

Results

System Overview

We introduce a novel way of encoding the toolpath as two separate scalar fields over the lasered top surface. One of them is the time field, $t(\mathbf{x})$, which provides the model with information regarding the directionality of the laser’s motion and the elapsed time of the process. The second field is the magnitude of the spatial gradient of the time field, $\|\nabla_x t\|$, which encodes information on the spatial frequency of the toolpath. Utilizing this feature set to train a U-Net model, we focus on predicting maximum spatiotemporal temperature distribution as it directly correlates with defect formation mechanisms while providing a clear target for assessing the performance of the utilized machine learning model.

To evaluate the effectiveness of the proposed features, we have produced a dataset consisting of more than 500 thermal simulations, encompassing a variety of geometries, based on the methodology described in 1. This section details our approach to dataset generation and provides information on the U-Net architecture, as well as the techniques employed to create the features.

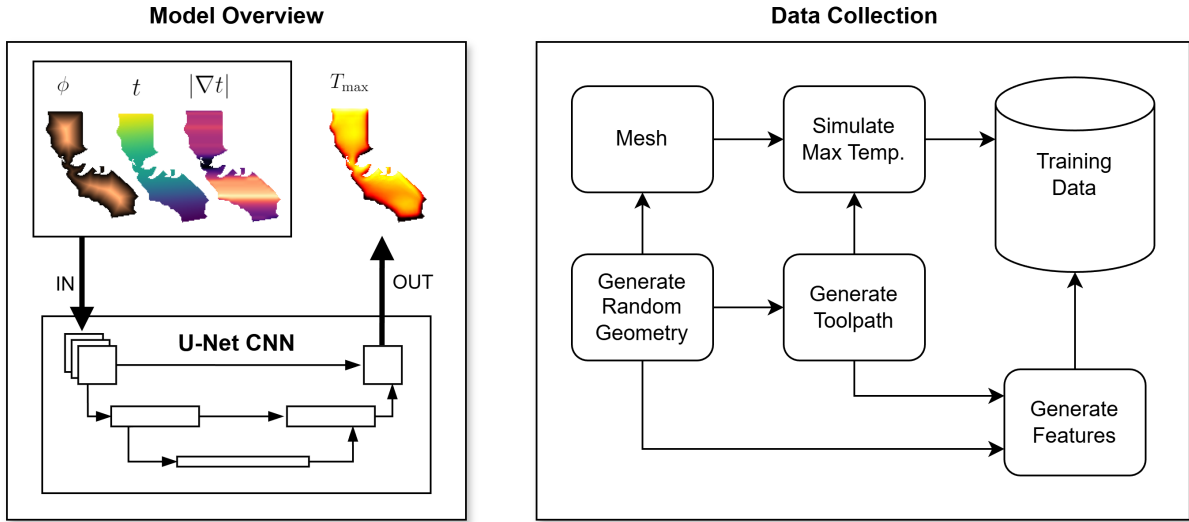


Fig. 1 | **U-Net model and data generation pipeline.** Overview of the complete workflow from geometry generation through feature engineering to neural network training for thermal field prediction.

It should be noted that the primary intention with this workflow is to provide a machine learning framework that could potentially be extended to be trained on higher fidelity models. The reader is referred to the aforementioned high-fidelity modeling approaches to this end. Our coarse finite element model functions as a means of roughly capturing the spatiotemporal distribution of the temperature in various geometries and toolpath orientations.

FEA Simulations

Our thermal model incorporates several key assumptions that simplify the complex multi-physics phenomena in LPBF to be able to generate the necessary training data:

1. Fluid behaviors are neglected within the melt pool, including Marangoni flows, surface tension effects, and other convective heat transfer mechanisms. This assumption focuses the model on conductive heat transfer, which dominates the thermal behavior at the scale of interest.
2. The thermal conduction and radiation through the unsintered powder are neglected. The potential error associated with this assumption can be assessed through Zohdi's [25] approximation where the heat conduction from the consolidated part to the surrounding powder is expressed as:

$$q_p = \mathbb{K}^* \frac{T_s - T_p}{h} A_c \quad (1)$$

where q_p is the heat loss to the surrounding powder, \mathbb{K}^* is the effective heat conduction between the part and the powder, T_s is the temperature of the powder, T_p is the average temperature of the part, h is the length scale or feature size of the part, and A_c is the area of the conductive interface between the part and the powder (*not* to be confused with the cross-sectional area of the part being consolidated). It can be directly inferred that with increasing part complexity (increasing conductive area and decreasing feature sizes) the conduction to the surrounding powder increases. We refer the reader to the Methods section for a deeper analysis for different geometries.

3. Conductive, convective, and radiative heat loss is neglected at all surfaces. The only thermal boundary condition applied is a Dirichlet condition at the bottom of the baseplate, acting as a heat sink at 873 K. This temperature is also the initial temperature of the whole domain.
4. The laser's heat input is modeled as a stationary Gaussian heat source that moves along the predefined toolpath, neglecting beam absorption variations and keyhole formation effects.

The thermal behavior in LPBF processes is governed by the transient heat conduction equation with a moving heat source:

$$\rho c(T) \frac{\partial T}{\partial t} = \nabla \cdot (k \nabla T) + Q_L(\mathbf{x}, t) \quad (2)$$

where ρ is the material density, $c(T) = c_{\text{eff}}(T)$ is the temperature-dependent effective specific heat capacity that incorporates latent heat effects, k is the thermal conductivity, T is the temperature field, and $Q_L(\mathbf{x}, t)$ represents the volumetric laser heat source.

The laser heat source is modeled as a 3D Gaussian distribution that moves along the predefined toolpath:

$$Q_L(\mathbf{x}, t) = \frac{6\sqrt{3}\eta P}{\pi^{3/2}r^3} \exp\left(-\frac{|\mathbf{x} - \mathbf{x}_L(t)|^2}{3r^2}\right) \quad (3)$$

where η is the absorption coefficient, P is the laser power, r is the laser beam radius, d is the layer depth, $\mathbf{x}_L(t)$ is the time-dependent laser position along the toolpath. A unidirectional raster scan pattern was used for all simulations in this work. This pattern scans the laser across the part in parallel lines, all moving in the same direction. The scan lines are spaced apart by the laser diameter with a 25% overlap. Each scan line is made up of points that are spaced apart by the distance the laser travels in a predefined time step. There is no latency introduced as the laser moves between scan lines.

Phase change during melting is handled using the effective specific heat capacity method, which incorporates the latent heat of fusion into the heat capacity term rather than treating it as a separate source term. The effective specific heat capacity is defined as:

$$c_{\text{eff}}(T) = c_{\text{nom}} + c_{\text{app}}(T) \quad (4)$$

where c_{nom} is the sensible heat capacity and $c_{\text{app}}(T)$ accounts for the latent heat of fusion during the melting transition:

$$c_{\text{app}}(T) = \frac{L_f}{\Delta T_m} \exp\left(-4\left(\frac{T - T_m}{\Delta T_m}\right)^2\right) \quad (5)$$

where L_f is the latent heat of fusion, T_m is the melting temperature, and ΔT_m is the melting temperature range over which the phase change occurs. This approach provides a smooth, differentiable representation of the phase change that is also numerical stable. The simulation employs a 2-step backward differentiation formula for time integration:

$$\frac{3T^{n+1} - 4T^n + T^{n-1}}{2\Delta t} = \frac{1}{\rho c_p} (\nabla \cdot (k\nabla T^{n+1}) + Q^{n+1}) \quad (6)$$

where superscripts denote time steps and $\Delta t = 2.5 \times 10^{-5}$ s is the time step size. The weak form of Equation 6 is:

$$\int_{\Omega} \left(\frac{3}{2\Delta t} \rho c_p T^{n+1} v + k \nabla T^{n+1} \cdot \nabla v \right) d\Omega = \int_{\Omega} \left(\frac{2}{\Delta t} \rho c_p T^n v - \frac{1}{2\Delta t} \rho c_p T^{n-1} v + Q^{n+1} v \right) d\Omega \quad (7)$$

for all test functions v in the appropriate function space. FEniCSx was used for the automated solution of every data point [46].

Multi-Scale Meshing Strategy

The domain is discretized into three distinct regions with different mesh resolutions to balance accuracy and computational efficiency, as illustrated in Figure 2 (c) and (d):

The mesh is constructed with three distinct resolutions: the lasered layer uses the finest mesh resolution of 0.1 mm (equivalent to the laser diameter) to accurately capture steep temperature gradients in the melt pool region, the part domain employs a medium mesh resolution of 1.0 mm to resolve thermal diffusion within the solidified material while maintaining computational tractability, and the baseplate is assigned the coarsest mesh resolution of 5.0 mm, as thermal gradients are less severe in this region distant from the laser heat source.

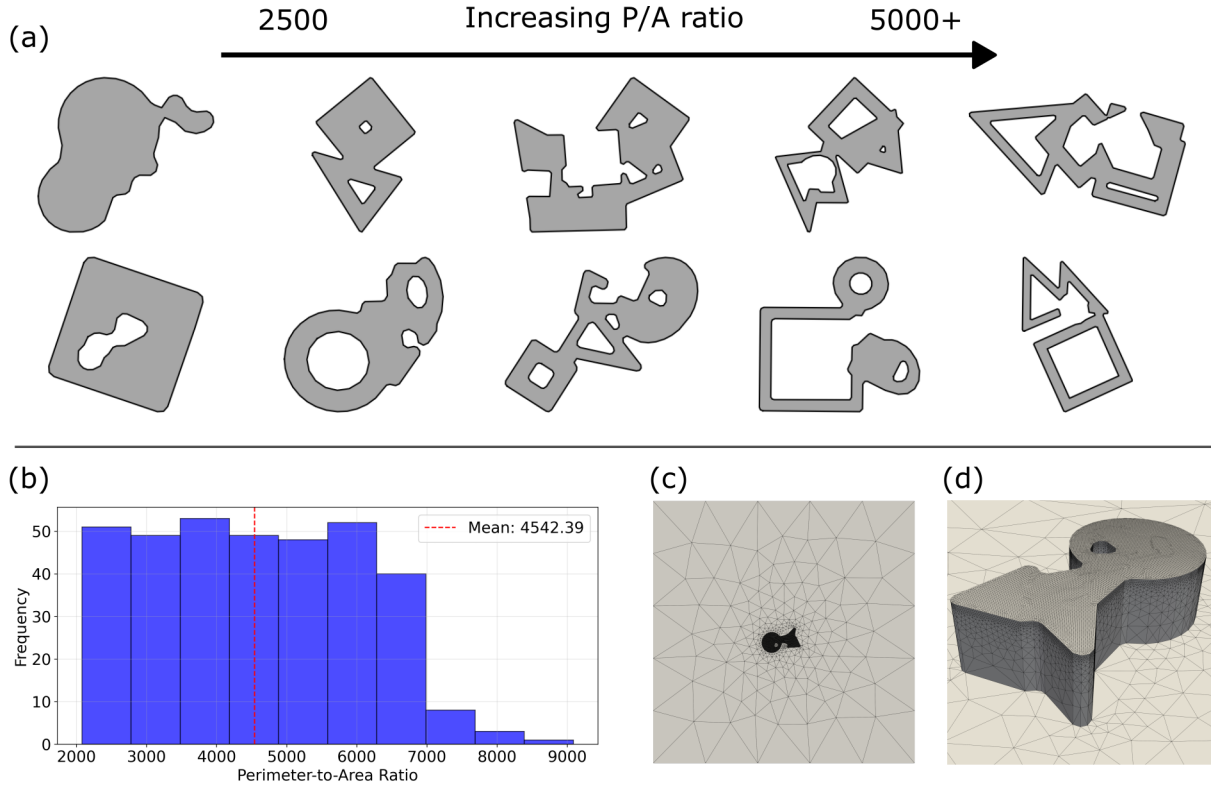


Fig. 2 | **Geometric complexity of generated shapes and meshing strategy.** (a) Examples of various shapes generated with different Perimeter-to-area ratios, (b) histogram distribution of perimeter-to-area ratios in the training dataset, (c) mesh view with baseplate showing multi-scale meshing strategy, and (d) close-up mesh view of the part domain.

This multi-scale meshing strategy ensures adequate spatial resolution where thermal gradients are steepest while avoiding unnecessary computational overhead in regions with smooth temperature variations.

The maximum temperature field $T_{\max}(\mathbf{x})$ is tracked throughout the entire simulation by monitoring the pointwise maximum temperature experienced at each spatial location:

$$T_{\max}(\mathbf{x}) = \max_{t \in [0, t_{\text{final}}]} T(\mathbf{x}, t) \quad (8)$$

This maximum temperature field serves as the target output for the neural network training, as it captures the peak thermal exposure that is a significant factor in determining material microstructure and mechanical properties. To ensure consistent spatial discretization with the input features, the maximum temperatures are sampled at the discrete toolpath points and then interpolated onto the structured 2D grid using Radial Basis Function (RBF) interpolation with linear kernels. This approach provides smooth temperature field representations while maintaining computational efficiency. Each simulation runs for approximately 6-7 minutes per core on an 8-core AMD Ryzen 9 8945HS laptop processor (8 runs in parallel).

Geometry Generation and Meshing

We generate diverse 2D part geometries using a randomized polygon generation algorithm. The algorithm creates complex shapes through boolean operations on primitive geometries to ensure dataset diversity while maintaining manufacturing constraints, such as the minimum feature size, as shown in Algorithm 1.

Algorithm 1 Composite Shape Generation (Simplified)

Input: Shape constraints (area, perimeter, diameter, P/A ratio bounds)

Output: Composite polygon

repeat

Generate 2-4 random primitives (square, ellipse, or triangle), each with random size and transformation

Generate 1-2 smaller random primitives for subtraction

Combine union shapes, then subtract the smaller shapes

Check if resulting shape satisfies all constraints (area, perimeter, diameter, P/A ratio)

until Shape meets all specified constraints or maximum attempts reached

return Final composite shape

To ensure a wide range of geometric complexity in our training dataset, we systematically control the perimeter-to-area (P/A) ratio of generated shapes as shown in Figure 2 (a). The P/A ratio is a quantitative measure of geometric complexity where shapes with higher ratios possess more elaborate or elongated boundaries relative to their enclosed area, while lower ratios correspond to simpler, more compact forms.

For reference, a perfect circle achieves the minimum possible P/A ratio among all shapes, representing the most compact geometry. In contrast, the maximum attainable

P/A ratio is theoretically unbounded unless practical manufacturing constraints are imposed. In our context, the upper bound is determined by the thinnest, longest printable rectangle, which is limited by the laser’s minimum feature size (i.e., the laser diameter). By enforcing these constraints, we ensure that the dataset spans the full spectrum of manufacturable geometric complexity, from simple, compact shapes to highly intricate, elongated forms. We note that the P/A ratio is not scale invariant and therefore the generated shapes are constrained to an area of about 3 mm^2 . A fixed discretization length h is applied uniformly across all geometries. Each shape is mapped onto a structured square grid of 192×192 pixels using a consistent bounding domain to ensure proper scaling relationships.

Figure 2(b) shows the distribution of perimeter-to-area ratios across our training dataset, ensuring representative coverage of geometric complexity.

Features

All features are computed on structured grids of the same N by N dimension. Let $\Omega \subset \mathbb{R}^2$ denote the part cross-section. We construct a square bounding domain $\mathcal{D} = [x_{\min}, x_{\max}] \times [y_{\min}, y_{\max}]$ that encloses Ω by centering a square on the minimum enclosing circle of Ω and adding a small padding. A uniform grid with spacing h (the discretization length) is laid over \mathcal{D} :

$$N = \left\lfloor \frac{x_{\max} - x_{\min}}{h} \right\rfloor + 1 = \left\lfloor \frac{y_{\max} - y_{\min}}{h} \right\rfloor + 1,$$

where $\lfloor \cdot \rfloor$ is the floor function. Each lattice index pair (i, j) maps to the physical coordinate $\mathbf{x}_{ij} = (x_i, y_j)$.

$$x_i = x_{\min} + i h, \quad y_j = y_{\min} + j h, \quad i, j \in \{0, \dots, N - 1\}.$$

We define a binary mask $\mathcal{M}_{ij} = \mathbb{1}_{\{\mathbf{x}_{ij} \in \Omega\}}$ to indicate points inside the part domain. All features are evaluated at the physical coordinates \mathbf{x}_{ij} and assembled into $N \times N$ arrays; the mask \mathcal{M} is used downstream for geometry-aware loss weighting.

From the geometry and toolpath, we compute three features that serve as input channels to our neural network.

The Signed Distance Field (SDF) $\phi(\mathbf{x})$ encodes the geometry by providing the signed distance from any point to the nearest boundary (positive inside the part, negative outside):

$$\phi(\mathbf{x}) = \begin{cases} \min_{\mathbf{b} \in \partial\Omega} \|\mathbf{x} - \mathbf{b}\| & \text{if } \mathbf{x} \in \Omega \\ -\min_{\mathbf{b} \in \partial\Omega} \|\mathbf{x} - \mathbf{b}\| & \text{if } \mathbf{x} \notin \Omega \end{cases} \quad (9)$$

where Ω represents the part domain and $\partial\Omega$ its boundary.

The time field $t(\mathbf{x})$ encodes the temporal information at each point by interpolating the laser scan times from the toolpath onto the grid. We utilize a Radial Basis Function (RBF) interpolator with linear kernels to generate a smooth and spatially continuous time field across the domain.

The interpolated expression takes the form of

$$t(\mathbf{x}) = \sum_{i=1}^{\mathbb{P}} c_i \|\mathbf{x} - \mathbf{x}_i\| + C,$$

where c_i are the linear kernel coefficients solved for at every toolpath point \mathbf{x}_i in \mathbb{P} , defined as the toolpath points, and C is a constant.

The magnitude of the spatial gradient of the time field computed using second-order accurate central finite difference approximations:

$$\|\nabla_x t\| = \sqrt{\left(\frac{\partial t}{\partial x}\right)^2 + \left(\frac{\partial t}{\partial y}\right)^2} \quad (10)$$

where,

$$\frac{\partial t}{\partial x} \approx \frac{t_{i+1,j} - t_{i-1,j}}{2\Delta x}, \quad \frac{\partial t}{\partial y} \approx \frac{t_{i,j+1} - t_{i,j-1}}{2\Delta y}, \quad i, j \in \{0, \dots, N-1\} \quad (11)$$

Figure 3 visualizes these three input channels for representative geometries, demonstrating how each feature encodes different aspects of the thermal physics.

Model Architecture and Training Strategy

We use a U-Net encoder–decoder designed for image-to-image regression with three down-sampling stages, symmetric upsampling, and skip connections. Each stage employs a double-convolution block with 7×7 kernels, batch normalization, and ReLU activations. Downsampling is implemented via $2 \times$ max pooling, and upsampling via $2 \times$ bilinear interpolation, followed by concatenation with the corresponding encoder features. A final 1×1 convolution produces a single-channel temperature map at 192×192 resolution. Operator-level equations are provided in the Methods section.

We employ a masked mean squared error loss that focuses training on regions within the part geometry while ignoring exterior regions:

$$\mathcal{L} = \frac{1}{|\mathcal{M}|} \sum_{i,j \in \mathcal{M}} (T_{pred}^{ij} - T_{target}^{ij})^2 \quad (12)$$

where \mathcal{M} represents the set of pixels within the part interior, and $|\mathcal{M}|$ is the number of interior pixels. This masking ensures that the model only learns temperature predictions




Fig. 3 | **Input feature visualization across different geometries.** Examples of the three input features (signed distance field, time field, and gradient magnitude) computed on geometries with varying complexity levels.

for physically meaningful regions inside the part geometry. No constraints are enforced on exterior regions. The dataset of 500 geometries is split into training (70%) and validation (30%) sets. Training is performed using Adam optimization.

To enable accurate predictions for arbitrary toolpath orientations using a model trained on only one raster angle, we implement a rotation alignment algorithm as shown in Algorithm 2. This algorithm leverages the fact that our model learns the underlying structure of the input-output mapping in a normalized coordinate system, allowing inferences on arbitrary orientations through apriori geometric transformations on the input.

Algorithm 2 Rotation Alignment for Orientation-Generalized Inference

Input: Geometry mask \mathcal{M} , features $(\phi, t, \|\nabla_x t\|)$, query raster angle θ , training raster angle θ_0 , trained model \mathcal{F}

Compute $\Delta\theta \leftarrow \theta - \theta_0$

Rotate inputs by $-\Delta\theta$: $(\phi', t', \|\nabla_x t'\|) \leftarrow \text{Rotate}((\phi, t, \|\nabla_x t\|), -\Delta\theta)$

Predict in the canonical frame: $\hat{T}' \leftarrow \mathcal{F}(\phi', t', \|\nabla_x t'\|)$

Inverse-rotate prediction: $\hat{T} \leftarrow \text{Rotate}(\hat{T}', +\Delta\theta)$

Apply mask: $\hat{T} \leftarrow \hat{T} \odot \mathbf{1}_{\mathcal{M}}$

Output: Orientation-aligned prediction \hat{T}

Simulation Results

The finite element simulations are capable of successfully capturing the spatiotemporal thermal phenomena occurring during LPBF. Figure 4 shows representative simulation outputs including maximum temperature fields, instantaneous temperature distributions, and effective specific heat capacity variations.




Fig. 4 | **Finite element simulations provide spatiotemporal temperature distributions.** (left) maximum temperature field, (center) instantaneous temperature distribution during laser scanning, and (right) effective specific heat capacity incorporating latent heat effects. The simulations successfully capture the steep thermal gradients and phase change behavior characteristic of laser powder bed fusion processes.

The simulations demonstrate several key thermal phenomena: steep temperature gradients near the melt pool region, heat accumulation effects from overlapping scan tracks, and thermal diffusion patterns that vary with part geometry. The effective specific heat capacity field clearly shows the phase change region where latent heat absorption occurs.

We observe a consistent inverse correspondence between the gradient magnitude of the time field, $\|\nabla_x t\|$, and the maximum temperature field, T_{\max} . Regions with large $\|\nabla_x t\|$ —indicating sharper changes in laser arrival time or sparse revisitation—tend to exhibit lower T_{\max} , while areas with small $\|\nabla_x t\|$ —corresponding to smoother arrival structure and frequent revisits—show higher T_{\max} . This trend is physically consistent with the heat equation (Eq. 2): $\|\nabla_x t\|$ encodes the spatial structure of energy input, and higher gradients align with reduced effective diffusion to those regions over the scan sequence. The correspondence is generally strong but can be locally modulated by geometric features (e.g., corners, thin ligaments, proximity to boundaries) and the alignment of scan paths relative to edges. These effects can alter local accumulation/dissipation patterns and occasionally attenuate the simple inverse trend, but the global relationship remains evident across geometries.

This perspective also connects to prior scan-structure descriptors such as the Relative Space Time proximity (RSP) map [47], which shares properties with discrete Laplacians on unstructured graphs. Both $\|\nabla_x t\|$ and the RSP quantify spatial variability in laser

revisitation, and thus provide complementary cues about where heat tends to accumulate or dissipate during the process.

Machine Learning Model Performance

The U-Net model demonstrates excellent performance across a wide range of geometric configurations. Overall, model predictions achieve an RMSE of about 2% on average across the entire test set, with maximum per-pixel errors typically limited to 3-4%. A 1% error represents a minimum and maximum absolute error of about 10 K and 20 K respectively given the temperature range of approximately 1000 K to 2000 K. Figure 5(a) presents detailed comparisons between true and predicted temperature fields for various test samples. There is a correlation between the perimeter-to-area ratio and prediction accuracy where more complex geometries are associated with higher prediction errors, as shown in the RMSE analysis plotted against perimeter-to-area ratio in Figure 5(b).

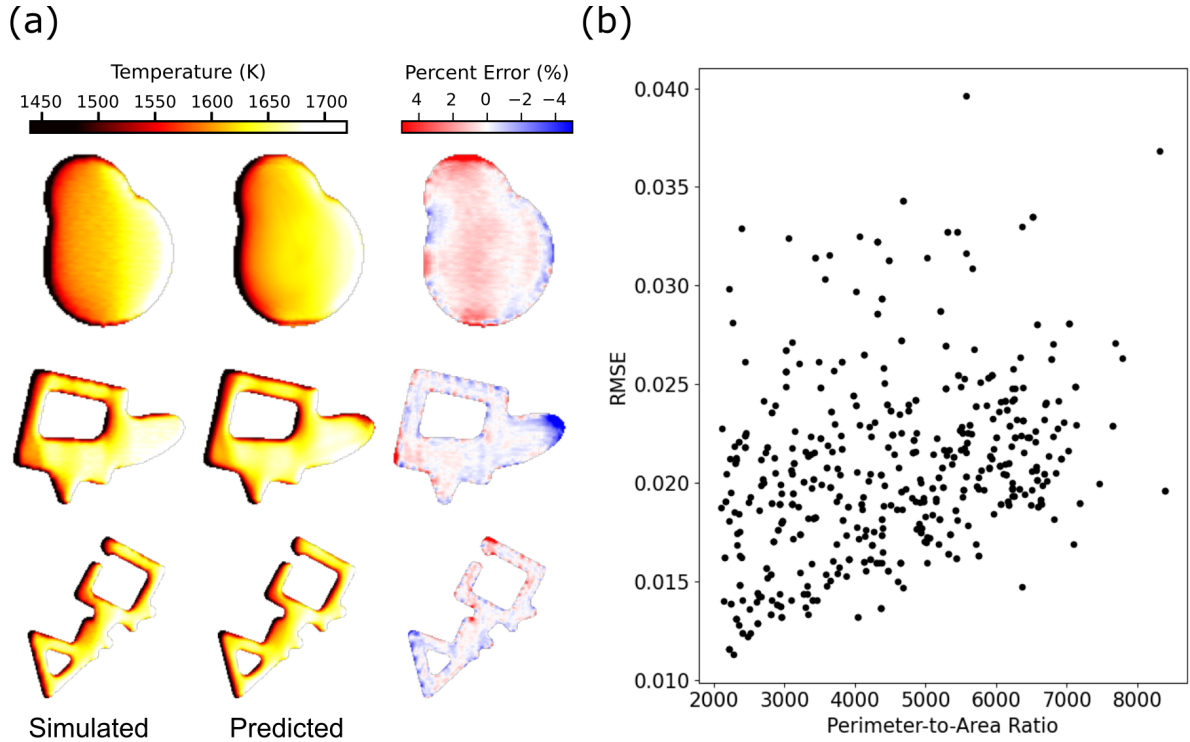


Fig. 5 | **Model prediction accuracy and the effect of geometric complexity.** (a) True vs. predicted temperature fields and error maps for representative test samples, and (b) Root Mean Square Error (RMSE) vs. perimeter-to-area (P/A) ratio.

The contribution of different feature combinations was evaluated in a feature ablation analysis to assess their significances. Results indicate that employing the complete feature set—comprising the Signed Distance Field $\phi(\mathbf{x})$, Time Field $t(\mathbf{x})$, and Gradient Magnitude $\|\nabla_x t\|$ —yields the fastest convergence and lowest training error. Models trained with individual features still achieve reasonable performance, but with noticeably reduced accuracy. Figure 6(a) shows training curves for different combinations of input

features, demonstrating that all three features together contribute to rapid convergence and improved accuracy. The most significant advantage of the complete feature set lies in training efficiency. The full feature set model achieves the same training error in approximately 10 epochs that the single time-based model ($t(\mathbf{x})$ only) requires 35 epochs to reach.

To further illustrate these findings, Figure 6(b) provides a comparison of error maps for different feature set configurations across a range of geometric complexities. For shapes with higher perimeter-to-area ratios, the model trained on the full feature set exhibits noticeably better performance compared to models trained on individual features or partial feature combinations. For shapes with lower perimeter-to-area ratios, the performance differences between feature sets are less pronounced.

The SDF and time gradient field lack information encoding the full temporal nature of the toolpath, and were therefore omitted from the ablation study. The SDF field does not contain any toolpath information and would alone not facilitate temperature predictions for different toolpaths. The time gradient field would not contain information regarding the direction of the toolpath and only encodes localized information resulting from the time field. This deficiency can be observed in the first row in Figure 8 where the left and right of the circular time gradient field exhibit lower magnitudes whereas the opposite is true in the temperature field. This contrast arises from the accumulation of thermal energy over time as the laser rasters from the left to the right.

Computational Efficiency

Table 1: Computational cost scaling with part size

Length (mm)	Input Size	Featurization (ms)	Inference (ms)	Total (ms)
0.5	31×31	13.7	7.4	21.2
1.0	59×59	58.9	9.7	68.6
2.0	116×116	437.0	18.8	455.8
4.0	229×229	8,793.0	50.7	8,843.7
8.0	455×455	98,168.5	210.8	98,379.3

FEA wall-clock (2.0 mm) — 8 cores: ~ 90 s & 1 core: ~ 420 s on an AMD Ryzen 9 8945HS

Notably, neural network inference is rapid, with each forward pass completing in milliseconds and exhibiting only a gradual increase in runtime as input size increases. For instance, for a 2.0 mm part, the neural network produces predictions in 18.8 ms, whereas the corresponding finite element simulation requires approximately 90 seconds on 8 CPU cores. This corresponds to a speedup of approximately 4,800 times (90,000 ms divided by 18.8 ms) for inference alone. When feature computation is included, the total pipeline at 2.0 mm achieves an end-to-end speedup of roughly 200 times (90,000 ms divided by 455.8 ms).

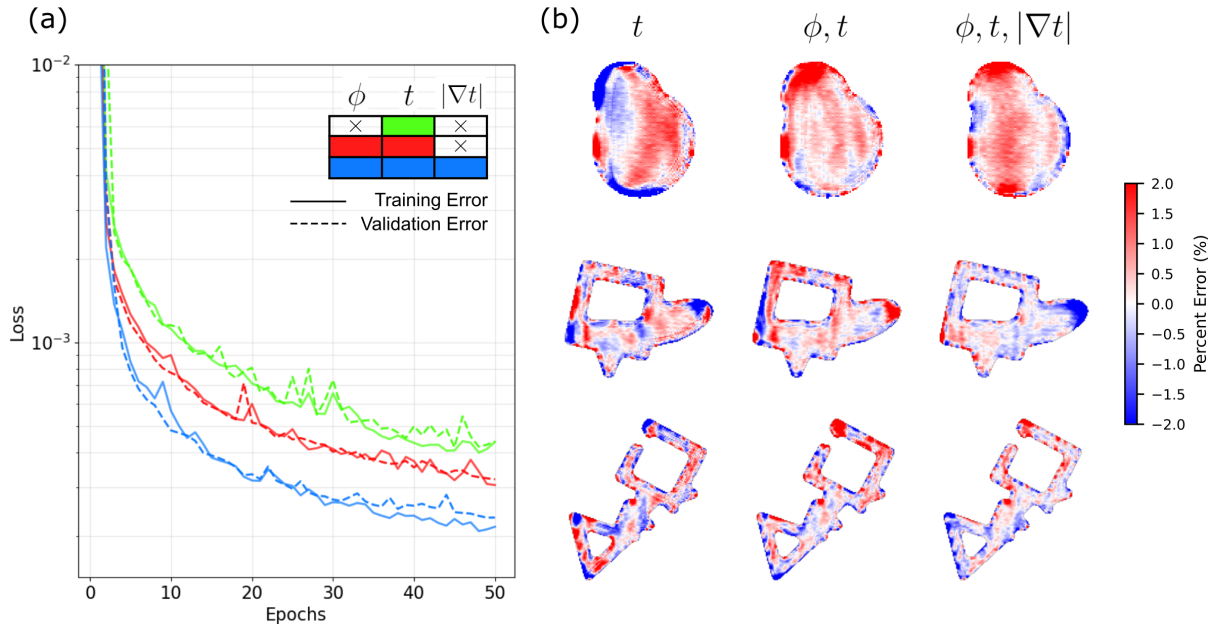


Fig. 6 | **Feature ablation study demonstrates improved training efficiency in using all features.** (a) Training error for models trained on different feature sets. Each row in the top-right table refers to a training set, with crossed boxes indicating the absence of their column’s feature. (b) Percentage error plots for different shapes where each column represents the errors associated with the predictions of models trained on different feature sets.

However, Table 1 also shows that featurization dominates total runtime and scales steeply with resolution (e.g., 98.2 s at 455×455). This cost arises from constructing the features which currently involve dense computations over the full grid and toolpath. Thus, while the learned model is not the bottleneck, the pre-processing stage can erase much of the inference-time benefit as domains grow, making the present pipeline insufficiently scalable for large parts without additional systems-level optimization.

There are several general strategies to address the featurization bottleneck without modifying the trained model. These include employing adaptive or multi-scale feature computation to focus resources where they are most needed, partitioning the computational domain to enable more efficient processing with parallelization and leveraging hardware acceleration for key operations. Future work may explore these approaches to maintain rapid inference while ensuring that feature generation remains efficient and scalable as problem sizes increase.

Geometric Connectivity: Gaps and Thin Bridges

Results show that the model reliably captures the effect of geometric connectivity on thermal behavior. In regions where small gaps interrupt the material domain, there is no conduction path, and the temperature disparity between adjacent points across the

gap is larger than it would be if the region were connected. Conversely, thin bridges establish a conductive link that facilitates heat transfer along narrow ligaments, which often manifests as relatively elevated local temperatures and more spatially coherent thermal structure compared to disconnected gaps.

Figure 7 illustrates these phenomena. Subfigure (a) shows three representative shapes with narrow gaps, predictions faithfully reflect stronger temperature contrasts across the non-connected regions. Subfigure (b) shows three shapes with thin bridges, temperatures remain comparatively higher along the connecting ligaments, consistent with the presence of a conduction pathway. These behaviors indicate that the network has internalized domain connectivity through the features, enabling robust generalization across topological variations. The discrepancy between connected and unconnected regions in this context will heavily depend on the significance of the heat conduction through the unsintered powder in these regions. The FE model here overestimates this discrepancy, leading the CNN to learn and reproduce these sharp temperature contrasts.

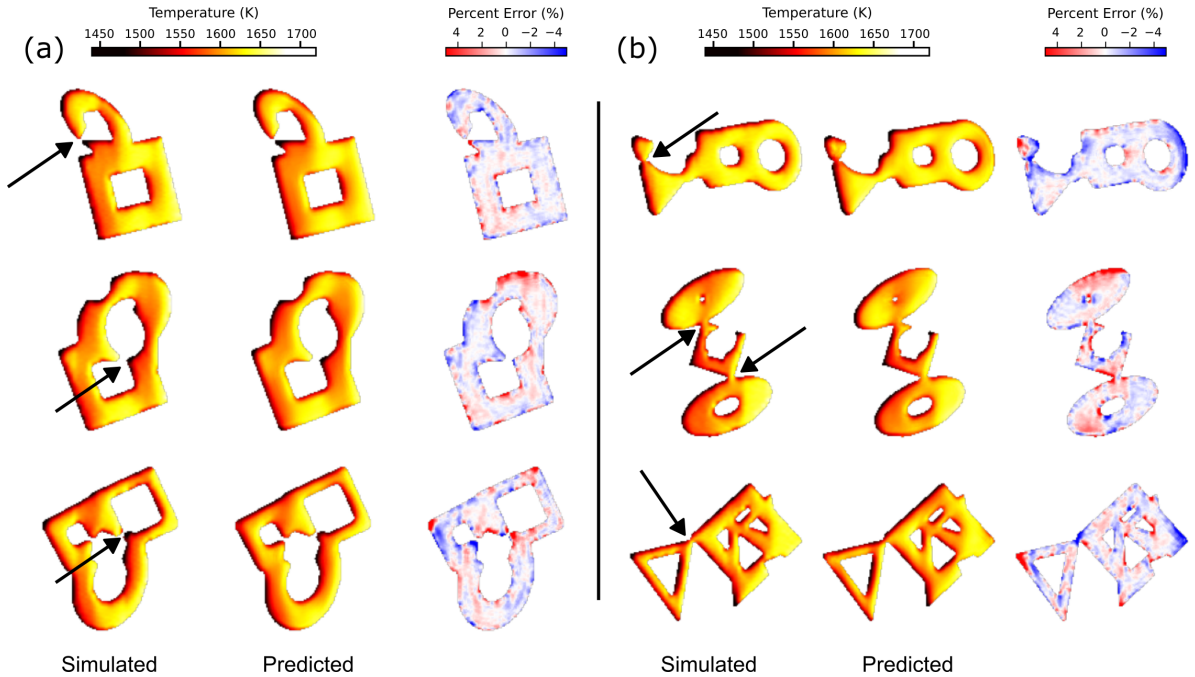


Fig. 7 | **The model captures intricate geometric features accurately.** (a) Three shapes with small gaps: predictions exhibit larger temperature disparities across non-conductive separations. (b) Three shapes with thin bridges: predictions show relatively elevated temperatures and coherent thermal structure along the narrow conductive ligaments.

Generalization to Spatially Varying Scan Speed

Although the network was trained using fixed laser scan speed, we evaluated its ability to generalize to spatially varying scan speeds by introducing a sinusoidal speed modulation

over the toolpath. We selected a simple circular geometry (radius 1 mm) and generated toolpaths with spatially varying speed defined as

$$v(y) = v_0 + A \sin\left(\frac{2\pi y}{\lambda}\right), \quad A < v_0,$$

with base speed $v_0 = 1.0$ m/s and amplitude A chosen to maintain strictly positive speed across the domain. We tested a baseline case that is effectively constant speed (very large λ) and three spatial wavelengths relative to a characteristic length $L = 2$ mm: $\lambda \in \{L/1.5, L/4, L/8\}$. Toolpaths were discretized at $\Delta s = 20$ μ m.

For each wavelength, we used the model to predict the maximum temperature. Figure 8 presents the results for each case in three columns: (left) the imposed speed field, (middle) the normalized time-field gradient magnitude $\|\nabla_x t\|$, and (right) the predicted temperature field, all masked to the part interior.

Two clear patterns emerge from this analysis. Regions with higher time-field gradients $\|\nabla_x t\|$ consistently exhibit lower predicted peak temperatures, with this monotonic inverse trend becoming more pronounced as the spatial modulation frequency increases (smaller λ). Despite never seeing spatially varying speed during training, the model correctly responds to the modulation via the time and gradient inputs, indicating that it has learned physically meaningful relationships.

Rapid Toolpath Orientation Optimization

One of the key advantages of our approach is its ability to enable rapid exploration of toolpath parameters for process optimization. We demonstrate this capability through a comprehensive study examining toolpath orientation effects on thermal behavior. We select three representative shapes with low, medium, and high perimeter-to-area ratios and performed a comprehensive analysis involving neural network predictions at 100 different raster angles (0° to 360°), FEA simulations at 10 different raster angles for validation, and comparison of mean and standard deviation of temperature distributions.

The optimization results reveal distinct geometric dependencies that provide both validation of the model and insights into thermal physics. Figure 9 (a) shows remarkably flat mean temperature and standard deviation curves across all raster angles for the low perimeter-to-area ratio shape. This behavior provides strong confidence in the model predictions, as the nearly circular geometry should theoretically exhibit identical thermal responses regardless of toolpath orientation due to rotational symmetry. The flat response confirms that the neural network correctly captures this symmetry.

In contrast, Figure 9 (b) demonstrates clear raster angle dependencies for the medium complexity shape, with distinct peaks and valleys indicating optimal and suboptimal orientations. This intermediate complexity geometry exhibits strong preferences for specific

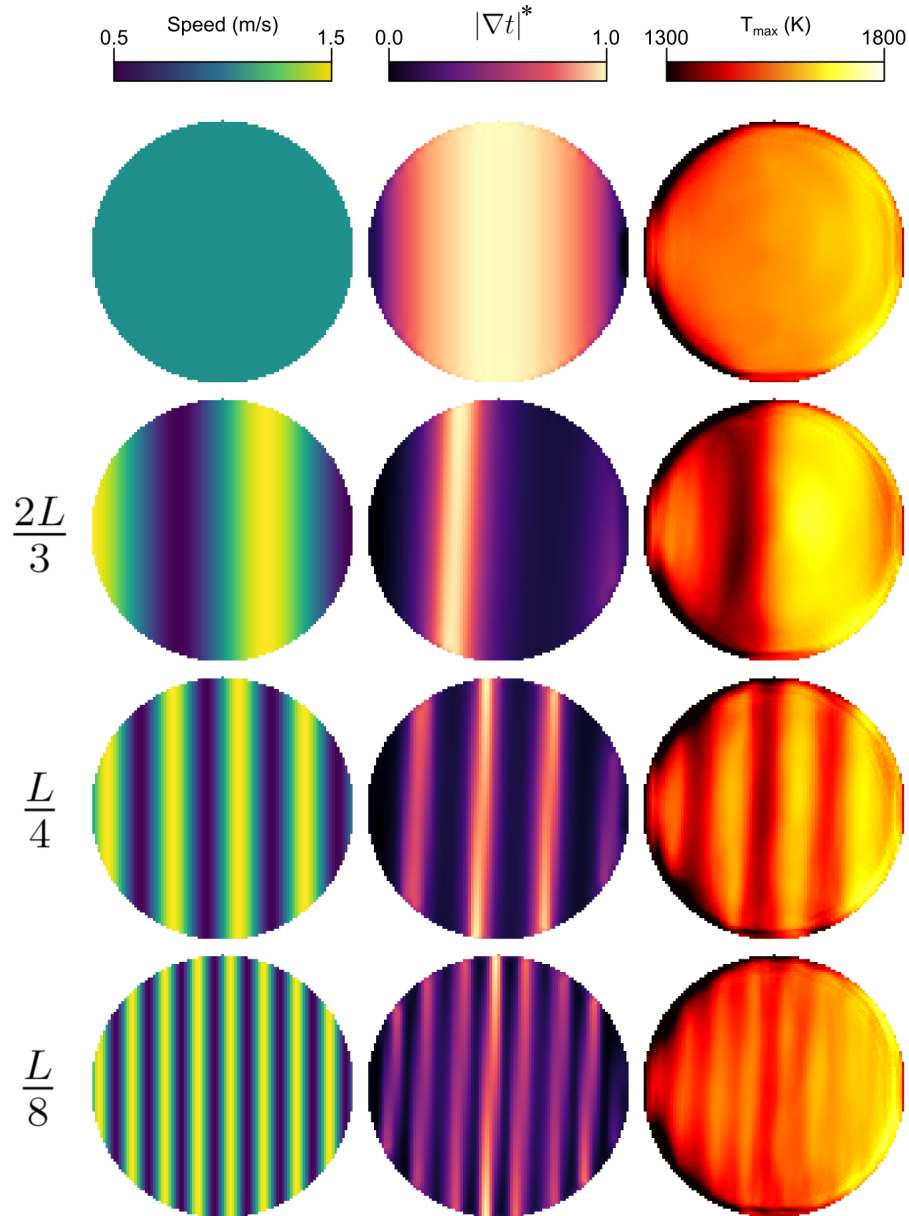


Fig. 8 | **Generalization to spatially varying scan speed.** Rows sweep from an effectively constant-speed baseline to progressively higher spatial frequency (shorter wavelength) modulation. Columns show the imposed speed field (left), normalized time-field gradient magnitude $\|\nabla_x t\|^*$ (middle), and the predicted maximum temperature (T_{\max}) (right). Higher $\|\nabla_x t\|^*$ correlates with lower predicted peak temperatures, demonstrating generalization to variable-speed toolpaths.

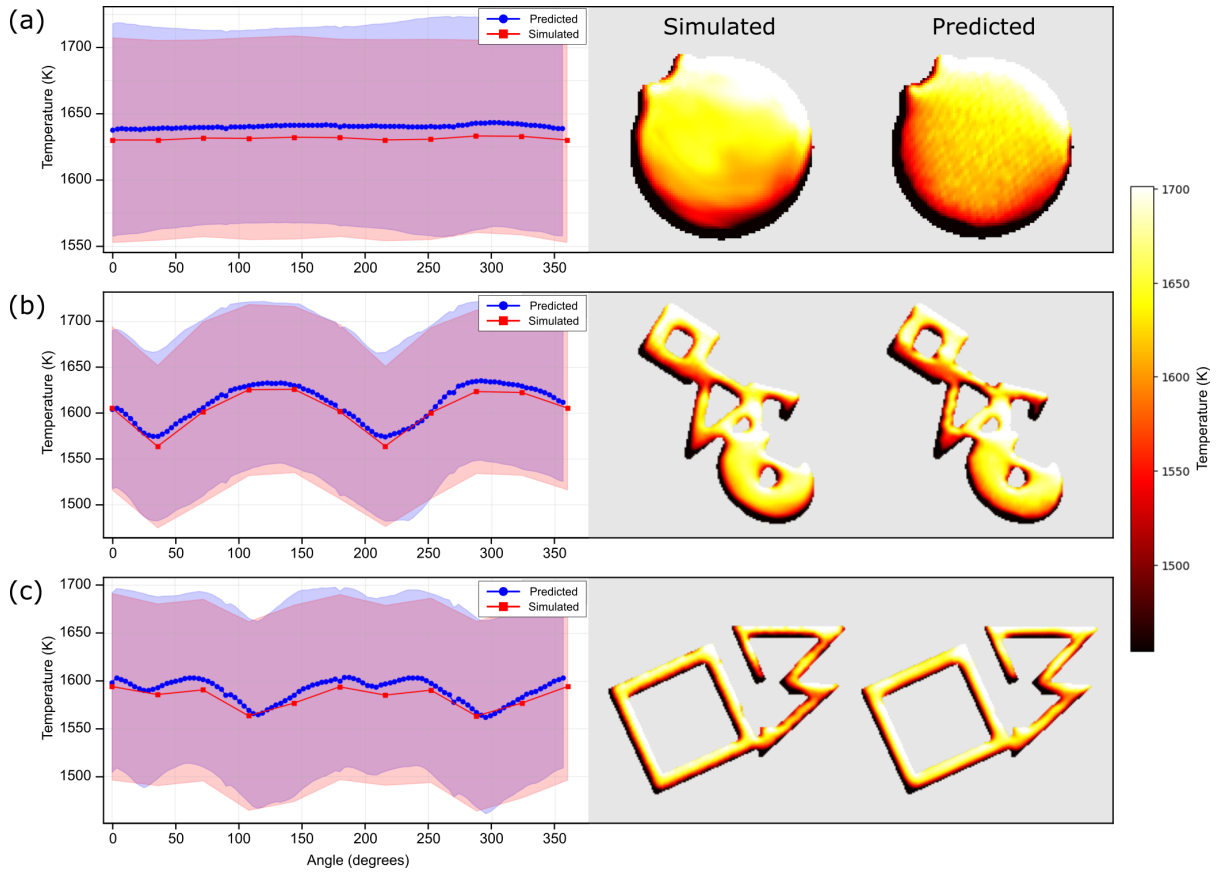


Fig. 9 | **Toolpath orientation optimization across geometric complexity levels.** (a) low perimeter-to-area ratio shape (simple geometry), (b) medium perimeter-to-area ratio shape (moderate complexity), and (c) high perimeter-to-area ratio shape (complex geometry). Points on plots are mean values for an orientation and the filled-bounds represent the spread of the temperature distribution.

raster angles which can be explained by the diagonal orientation of the spread of the shape. Toolpaths oriented along this spread lead to smaller temperatures and larger temperatures when perpendicular to it.

Figure 9 (c) presents a different behavior for the highly complex geometry with the highest perimeter-to-area ratio. While temperature variations occur across different raster angles, the curves do not demonstrate strong preferences for any particular orientation. This result reflects the geometric complexity: when part boundaries become sufficiently intricate, local thermal effects from detailed geometric features begin to dominate over global toolpath orientation effects.

The results demonstrate good agreement between CNN predictions and FEA simulations across all geometric complexities, with both trends and magnitudes captured accurately. This validation confirms the reliability of the neural network model for toolpath optimization applications. The rapid inference capability enables practical process optimization, making real-time exploration of toolpath parameters feasible for manufacturing workflows while providing physically meaningful insights.

Discussion

The key contribution of this research is a feature engineering methodology that effectively encodes geometry and laser toolpath which allows machine learning models to generalize to unseen geometries providing rapid temperature prediction capability for toolpath optimization. The provided framework can be extended to other problems involving moving point sources, which includes most additive manufacturing processes but could also potentially be applied to other processes such as metal casting where the relative spatio-temporal relationships between points in the flow are critical for predicting flow patterns for preventing defects such as cold shuts and misruns.

Several key opportunities for future research directions are evident following this work. Firstly, the FEA model utilized in this study contains simplifying assumptions which were necessary for the computational tractability of generating the training dataset. Considering the results demonstrated in this paper, future studies can utilize this framework to train larger models on higher-fidelity simulation data or even experimental data obtained from thermal imaging. Regardless of the data source, experimental validation remains essential to verify model predictions against real thermal measurements in actual LPBF processes for the practical application of this framework.

Secondly, overhangs are not addressed in the present work, though they are critical geometric features when it comes to diffusion of heat throughout a part. Shelf-like features, for example, are prone to higher temperatures. Future research should investigate incorporating overhang characteristics through additional input channels that capture geometric support relationships, thermal bridging effects, and powder bed interactions.

This could involve developing specialized feature engineering approaches for unsupported regions and their thermal implications, a simple example being an array where every pixel quantifies the projected distance from the top layer to the bottom most feature in sight of the projected point.

Finally, an important direction is the extension of the approach to full three-dimensional temperature field prediction. Predicting 3D temperature distributions would provide more comprehensive thermal information throughout the part volume, but would require the development of 3D feature representations, potentially more sophisticated network architectures such as 3D U-Nets or neural operators, efficient volumetric data handling strategies, and addressing inter-layer temperature gradients.

As it stands, the rapid inference capability our approach enables provides several practical applications in LPBF. The model provides a tool for quickly assessing thermal hotspots in different cross-sections during build preparation, allowing designers to identify potential quality issues before physical printing and potentially mitigating them. Most significantly, the millisecond inference times make the model suitable for rapid toolpath optimization, a capability that has been elusive with traditional simulation approaches due to their computational intensity. The demonstrated toolpath orientation optimization capability could be extended to other process parameters like laser speed and beam properties, enabling larger parameter space explorations.

Methods

FEA Simulation Parameters

Finite element time step: $\Delta t = 2.5 \times 10^{-5}$ s (calculated as $0.25 \times d_L/v$).

Material properties are based on Stainless Steel 316L: $\rho = 7900$ kg/m³, $k = 20.0$ W/(m·K), $c_p = 500$ J/(kg·K), $T_m = 1673.15$ K, $L = 270,000$ J/kg.

Laser properties: $P = 25$ W, $d_L = 0.1$ mm, $v = 1.0$ m/s, $\alpha = 1.0$, $\eta_{overlap} = 25\%$.

Error associated with bulk material approximation:

Framing the approximation in light of the heat input from the laser, Q_L , we can express the relative influence on the system by the laser and the powder conduction by,

$$\frac{Q_L}{q_p} = \frac{hQ_L}{\mathbb{K}^*(T_s - T_p)A_c} \quad (13)$$

Assuming that $Q_L \approx O(10^3)$, $\mathbb{K}^* \approx O(10^1)$ and $T_s - T_p \approx O(10^3)$, and $A_c \propto P$ where P is the perimeter of the parts cross sectional area, the ratio can be expressed as:

$$\frac{Q_L}{q_p} \propto \frac{10h}{P} \quad (14)$$

This relation can be evaluated for the lowest and highest perimeter-to-area ratios practically possible. At the minimum, a circle's perimeter and characteristic conductive lengths are,

$$P_{circle} = 2\pi r, \quad h_{circle} \approx 2r \quad (15)$$

For a constant area, A and evaluating the variables with respect to the areas,

$$A_{circle} = A = \pi r^2 \implies r = \sqrt{\frac{A}{\pi}} \implies P_{circle} = 2\pi \sqrt{\frac{A}{\pi}} \quad (16)$$

$$h_{circle} = 2r = 2\sqrt{\frac{A}{\pi}} \quad (17)$$

For the circle, the ratio becomes:

$$\left. \frac{Q_L}{q_p} \right|_{circle} \propto \frac{10 \cdot 2\sqrt{\frac{A}{\pi}}}{2\pi \sqrt{\frac{A}{\pi}}} = \frac{20\sqrt{\frac{A}{\pi}}}{2\pi \sqrt{\frac{A}{\pi}}} = \frac{10}{\pi} \approx 3.18 \quad (18)$$

At the maximum, a thin rectangle with length L and width $2r_{laser}$ (diameter of laser for smallest possible feature size) has perimeter and characteristic conductive length of,

$$P_{rect} = 2L + 4r_{laser}, \quad h_{rect} \approx 2r_{laser} \quad (19)$$

For the same constant area A ,

$$A_{rect} = A = L \cdot 2r_{laser} \implies L = \frac{A}{2r_{laser}} \implies P_{rect} = 2 \cdot \frac{A}{2r_{laser}} + 4r_{laser} = \frac{A}{r_{laser}} + 4r_{laser} \quad (20)$$

$$h_{rect} = 2r_{laser} \quad (21)$$

For the thin rectangle, the ratio becomes:

$$\left. \frac{Q_L}{q_p} \right|_{rect} \propto \frac{20r_{laser}}{\frac{A}{r_{laser}} + 4r_{laser}} \quad (22)$$

For the thin rectangle case, we can simplify the expression by considering that for practical LPBF geometries, the area A is typically much larger than $4r_{laser}^2$. This allows us to

approximate:

$$\left. \frac{Q_L}{q_p} \right|_{rect} \propto \frac{20r_{laser}}{\frac{A}{r_{laser}} + 4r_{laser}} \approx \frac{20r_{laser}}{\frac{A}{r_{laser}}} = \frac{20r_{laser}^2}{A} \quad (23)$$

Thus, the ratio of laser heat input to powder conduction varies between these geometric extremes:

$$\frac{20r_{laser}^2}{A} \lesssim \frac{Q_L}{q_p} \lesssim \frac{10}{\pi} \quad (24)$$

For typical LPBF part dimensions with $A \sim 10^{-5} \text{ m}^2$ and $r_{laser} = 50 \times 10^{-6} \text{ m}$, the bounds in 24 evaluate to:

$$0.005 \lesssim \frac{Q_L}{q_p} \lesssim 3.18 \quad (25)$$

This analysis indicates that the bulk material approximation potentially introduces large amounts of error for thin features experiencing relatively smaller laser-to-conduction ratios compared to compact geometries. This methodology is adopted from [25].

Network Architecture Details

This appendix provides the operator-level description and diagram of the U-Net used in this work. Let the input feature tensor be $\mathbf{X} \in \mathbb{R}^{C_{in} \times H \times W}$ with $C_{in} = 3$ (SDF, time field, $\|\nabla t\|$) and $(H, W) = (192, 192)$.

We first define the double-convolution block, which is a convolutional block $\text{CBR}(\cdot)$ defined as:

$$\text{CBR}(\mathbf{X}) = \sigma(\text{BN}(\text{Conv}_{k=7}(\mathbf{X}))). \quad (26)$$

where $\sigma(\cdot)$ is ReLU and $\text{BN}(\cdot)$ is batch normalization. Padding preserves spatial dimensions.

The double-convolution block applies this operation twice:

$$\text{DoubleConv}(\mathbf{Z}) = \text{CBR}(\text{CBR}(\mathbf{Z})). \quad (27)$$

The downsampling operation is defined as $\text{Down}(\mathbf{Z}) = \text{DoubleConv}(\text{MaxPool}_{2 \times 2}(\mathbf{Z}))$.

Given the encoder feature \mathbf{E} and decoder feature \mathbf{D} , the upsampling operation is defined as:

$$\text{Up}(\mathbf{D}, \mathbf{E}) = \text{DoubleConv}([\text{Upsample}_{\times 2}(\mathbf{D}), \mathbf{E}]), \quad (28)$$

where $\text{Upsample}_{\times 2}$ is bilinear interpolation and Concat is channel-wise concatenation.

The encoder-decoder is defined as $\mathbf{H}_0 = \text{DoubleConv}(\mathbf{X})$ with channels $C = [16, 32, 64, 128]$ across stages. The encoder stacks $\mathbf{H}_1 = \text{Down}(\mathbf{H}_0)$, $\mathbf{H}_2 = \text{Down}(\mathbf{H}_1)$, $\mathbf{H}_3 = \text{Down}(\mathbf{H}_2)$. The decoder computes $\mathbf{D}_1 = \text{Up}(\mathbf{H}_3, \mathbf{H}_2)$, $\mathbf{D}_2 = \text{Up}(\mathbf{D}_1, \mathbf{H}_1)$, $\mathbf{D}_3 = \text{Up}(\mathbf{D}_2, \mathbf{H}_0)$.

The output is defined as $\mathbf{Y} = \text{Conv}_{1 \times 1}(\mathbf{D}_3)$, where $\mathbf{Y} \in \mathbb{R}^{1 \times H \times W}$.

Training Configuration

The dataset of 500 geometries was split into training and validation sets using a 70/30 ratio, with all features and outputs represented at 192×192 resolution. Training was performed using the Adam optimizer with a learning rate of $\eta = 2.5 \times 10^{-4}$ and a batch size of 16. Early stopping was applied after 50 epochs to prevent overfitting. All computations were performed on an NVIDIA RTX 4060 Laptop GPU. Input features and output temperatures were normalized to facilitate training: the signed distance field was scaled to $[0, 0.001]$ m, the time field to $[0, 0.08]$ s, the gradient magnitude to $[0, 0.002]$ s^{-1} , and the temperature field to $[1000, 2000]$ K.

Acknowledgements

We acknowledge support from Alfred P. Sloan Foundation and the Savio computational cluster resource provided by the Berkeley Research Computing program.

Data Availability

The data was generated using the code available at github.com/kahraman-gd/geopath_lpbfd.

Code Availability

The code used to generate the data and train the machine learning model is available at github.com/kahraman-gd/geopath_lpbfd.

Competing Interests

GXG is an Editorial Board Member of npj Artificial Intelligence. GXG was not involved in the journal's review of, or decisions related to, this manuscript. The authors have no other competing interests.

References

1. DebRoy, T. *et al.* Additive manufacturing of metallic components – Process, structure and properties. *Progress in Materials Science* **92**, 112–224. <https://www.sciencedirect.com/science/article/pii/S0079642517301172> (2018).

2. Mostafaei, A. *et al.* Defects and anomalies in powder bed fusion metal additive manufacturing. *Current Opinion in Solid State and Materials Science* **26**, 100974. <https://www.sciencedirect.com/science/article/pii/S1359028621000772> (2022).
3. Wang, R., Standfield, B., Dou, C., Law, A. C. & Kong, Z. J. Real-time process monitoring and closed-loop control on laser power via a customized laser powder bed fusion platform. *Additive Manufacturing* **66**, 103449. <https://www.sciencedirect.com/science/article/pii/S2214860423000623> (2023).
4. Taherkhani, K., Cantzler, G., Eischer, C. & Toyserkani, E. Development of control systems for laser powder bed fusion. *The International Journal of Advanced Manufacturing Technology* **129**, 5493–5514 (2023).
5. Riensche, A. *et al.* Rapid, Autonomous, and Shape-Agnostic Physics-Guided Thermal History Control to Improve Part Quality in Laser Powder Bed Fusion Additive Manufacturing (2024).
6. Kim, Y., Alcantara, D. & Zohdi, T. I. Thermal state estimation of fused deposition modeling in additive manufacturing processes using Kalman filters. *International Journal for Numerical Methods in Engineering* **123**, 3021–3033. <https://onlinelibrary.wiley.com/doi/10.1002/nme.6490> (2022).
7. Dar, J., Ponsot, A. G., Jolma, C. J. & Lin, D. A review on scan strategies in laser-based metal additive manufacturing. *Journal of Materials Research and Technology* **36**, 5425–5467 (2025).
8. Krasniqi, M., Löffler, F. & Tutsch, R. Influence of scanning strategies on dimensional accuracy in laser powder bed fusion. *Measurement: Sensors* **38**, 101840 (2025).
9. Qin, M. *et al.* Adaptive toolpath generation for distortion reduction in laser powder bed fusion process. *Additive Manufacturing* **64**, 103432 (2023).
10. Jin, Z., Zhang, Z. & Gu, G. X. Autonomous in-situ correction of fused deposition modeling printers using computer vision and deep learning. *Manufacturing Letters* **22**, 11–15 (2019).
11. Wang, X. Q., Jin, Z., Zheng, B. & Gu, G. X. Transformer-Based Approach for Printing Quality Recognition in Fused Filament Fabrication. *npj Advanced Manufacturing* **2**, 15. ISSN: 3004-8621 (Apr. 19, 2025).
12. Thampy, V. *et al.* Subsurface Cooling Rates and Microstructural Response during Laser Based Metal Additive Manufacturing. *Scientific Reports* **10**, 1981. <https://www.nature.com/articles/s41598-020-58598-z> (2020).

13. Meier, C. *et al.* Physics-based modeling and predictive simulation of powder bed fusion additive manufacturing across length scales. *GAMM-Mitteilungen* **44**, e202100014 (2021).
14. Zhang, Z.-D. *et al.* *3-Dimensional Heat Transfer Modeling for Laser Powder Bed Fusion Additive Manufacturing Using Parallel Computing and Adaptive Mesh* 2022.
15. Cheng, J. *et al.* A Gaussian Process-Based extended Goldak heat source model for finite element simulation of laser powder bed fusion additive manufacturing process. *Computational Materials Science* **244**, 113185 (2024).
16. Ganeriwala, R. & Zohdi, T. I. A coupled discrete element-finite difference model of selective laser sintering. *Granular Matter* **18**, 21 (2016).
17. Ganeriwala, R. & Zohdi, T. I. Multiphysics Modeling and Simulation of Selective Laser Sintering Manufacturing Processes. *Procedia CIRP* **14**, 299–304 (2014).
18. Sharma, R. & Kumar, A. Track-Scale Simulations of Selective Laser Melting to Investigate Development and Mitigation of Thermal Stresses. *Lasers in Manufacturing and Materials Processing* **6**, 464–492 (2019).
19. Carraturo, M., Viguerie, A., Reali, A. & Auricchio, F. *Two-level Method Part-scale Thermal Analysis of Laser Powder Bed Fusion Additive Manufacturing* 2021.
20. Psihoyos, H. O. & Lampeas, G. N. Efficient thermomechanical modelling of Laser Powder Bed Fusion additive manufacturing process with emphasis on parts residual stress fields. *AIMS Materials Science* **9**, 455–480 (2022).
21. Cole, K. D., Riensche, A. & Rao, P. K. Discrete Green’s functions and spectral graph theory for computationally efficient thermal modeling. *International Journal of Heat and Mass Transfer* **183**, 122112 (2022).
22. Sawhney, R., Miller, B., Gkioulekas, I. & Crane, K. Walk on Stars: A Grid-Free Monte Carlo Method for PDEs with Neumann Boundary Conditions. *ACM Transactions on Graphics* **42**, 1–20 (2023).
23. Riensche, A. *et al.* Thermal modeling of directed energy deposition additive manufacturing using graph theory. *Rapid Prototyping Journal* **29**, 324–343 (2023).
24. Zohdi, T. I. Rapid Simulation of Laser Processing of Discrete Particulate Materials. *Archives of Computational Methods in Engineering* **20**, 309–325 (2013).
25. Zohdi, T. I. Ultra-fast laser-patterning computation for advanced manufacturing of powdered materials exploiting knowledge-based heat-kernels. *Computer Methods in Applied Mechanics and Engineering* **343**, 234–248 (2019).
26. Zohdi, T. I. Construction of a rapid simulation design tool for thermal responses to laser-induced feature patterns. *Computational Mechanics* **62**, 393–409 (2018).

27. Dissanayake, M. W. M. G. & Phan-Thien, N. Neural-network-based approximations for solving partial differential equations. *Communications in Numerical Methods in Engineering* **10**, 195–201. <https://onlinelibrary.wiley.com/doi/abs/10.1002/cnm.1640100303> (1994).
28. Lagaris, I. E., Likas, A. & Fotiadis, D. I. Artificial neural networks for solving ordinary and partial differential equations. *IEEE Transactions on Neural Networks* **9**, 987–1000. <https://ieeexplore.ieee.org/document/712178/> (1998).
29. Raissi, M., Perdikaris, P. & Karniadakis, G. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics* **378**, 686–707 (2019).
30. Li, Z. *et al.* *Fourier Neural Operator for Parametric Partial Differential Equations* 2021.
31. Li, Z. *et al.* Physics-Informed Neural Operator for Learning Partial Differential Equations. *ACM / IMS Journal of Data Science* **1**, 1–27 (2024).
32. Sanchez-Gonzalez, A. *et al.* *Learning to Simulate Complex Physics with Graph Networks* 2020.
33. Zhang, H. *et al.* *Combining physics-informed graph neural network and finite difference for solving forward and inverse spatiotemporal PDEs* 2024.
34. Chen, J. *et al.* Accelerating Thermal Simulations in Additive Manufacturing by Training Physics-Informed Neural Networks With Randomly Synthesized Data. *Journal of Computing and Information Science in Engineering* **24** (2023).
35. Liao, S. *et al.* Hybrid thermal modeling of additive manufacturing processes using physics-informed neural networks for temperature prediction and parameter identification. *Computational Mechanics* **72**, 499–512 (2023).
36. Zhu, Q., Liu, Z. & Yan, J. Machine learning for metal additive manufacturing: predicting temperature and melt pool fluid dynamics using physics-informed neural networks. *Computational Mechanics* **67**, 619–635 (2021).
37. Sajadi, P., Dehaghani, M. R., Tang, Y. & Wang, G. G. *Real-Time 2D Temperature Field Prediction in Metal Additive Manufacturing Using Physics-Informed Neural Networks* 2024.
38. Kim, D. H. & Zohdi, T. I. Tool path optimization of selective laser sintering processes using deep learning. *Computational Mechanics* **69**, 383–401 (2022).
39. Sahli Costabal, F., Pezzuto, S. & Perdikaris, P. Delta-PINNs: physics-informed neural networks on complex geometries. OpenReview. <https://openreview.net/forum?id=5P96KWeULzE> (2022).

40. Li, Z. *et al.* *Geometry-Informed Neural Operator for Large-Scale 3D PDEs* 2023. arXiv: 2309.00583 [cs.LG]. <http://arxiv.org/abs/2309.00583>.
41. Chenaud, M., Magoulès, F. & Alves, J. Physics-Informed Graph-Mesh Networks for PDEs: A hybrid approach for complex problems. *Advances in Engineering Software* **197**, 103758. <http://arxiv.org/abs/2410.02819> (2024).
42. Cameron, S., Pretorius, A. & Roberts, S. *Nonparametric Boundary Geometry in Physics Informed Deep Learning* preprint. 2024.
43. Yin, M. *et al.* A scalable framework for learning the geometry-dependent solution operators of partial differential equations. *Nature Computational Science* **4**, 928–940. <https://www.nature.com/articles/s43588-024-00732-2> (2024).
44. Zhong, W. & Meidani, H. *Physics-Informed Geometry-Aware Neural Operator* 2024. arXiv: 2408.01600 [cs.LG]. <http://arxiv.org/abs/2408.01600>.
45. Li, H., Miao, Y., Khodaei, Z. S. & Aliabadi, M. H. Finite-PINN: A physics-informed neural network with finite geometric encoding for solid mechanics. *Journal of the Mechanics and Physics of Solids* **203**, 106222. <https://www.sciencedirect.com/science/article/pii/S002250962500198X> (2025).
46. Baratta, I. A. *et al.* DOLFINx: The next Generation FEniCS Problem Solving Environment. <https://zenodo.org/records/10447666> (2025) (Dec. 31, 2023). Pre-published.
47. Demir, K., Zhang, Z., Ben-Artzy, A., Hosemann, P. & Gu, G. X. Laser scan strategy descriptor for defect prognosis in metal additive manufacturing using neural networks. *Journal of Manufacturing Processes* **67**, 628–634 (2021).