

Reinforcement learning-based toolpath optimisation with 3D U-Net driven rapid thermal prediction

Taekyeong Kim, Dohyeon Kim, Soon-Sung Kwon, Swee Leong Sing, Namhun Kim & Im Doo Jung

To cite this article: Taekyeong Kim, Dohyeon Kim, Soon-Sung Kwon, Swee Leong Sing, Namhun Kim & Im Doo Jung (2026) Reinforcement learning-based toolpath optimisation with 3D U-Net driven rapid thermal prediction, *Virtual and Physical Prototyping*, 21:1, e2627765, DOI: [10.1080/17452759.2026.2627765](https://doi.org/10.1080/17452759.2026.2627765)

To link to this article: <https://doi.org/10.1080/17452759.2026.2627765>



© 2026 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



[View supplementary material](#)



Published online: 15 Feb 2026.



[Submit your article to this journal](#)



Article views: 283



[View related articles](#)



[View Crossmark data](#)

Reinforcement learning-based toolpath optimisation with 3D U-Net driven rapid thermal prediction

Taekyeong Kim ^{a,†}, Dohyeon Kim ^{a,†}, Soon-Sung Kwon ^a, Swee Leong Sing ^b, Namhun Kim ^{a,c} and Im Doo Jung ^{a,c}

^aDepartment of Mechanical Engineering, Ulsan National Institute of Science and Technology, Ulsan, Republic of Korea; ^bDepartment of Mechanical Engineering, National University of Singapore, Singapore, Singapore; ^cArtificial Intelligence Graduate School, Ulsan National Institute of Science and Technology, Ulsan, Republic of Korea

ABSTRACT

Fabricating complex geometries in additive manufacturing induces severe thermal gradients, resulting in significant warpage in the bottom layers. Although optimising toolpath strategies is essential for mitigating these issues, the prohibitive computational costs associated with calculating thermal distributions for complex 3D geometries limit current approaches. To overcome this limitation, we propose a novel framework that accelerates reinforcement learning (RL)-based toolpath optimisation by integrating a 3D U-Net as a surrogate model. The 3D U-Net is trained on data generated via finite difference method (FDM) simulation to enable rapid thermal predictions. This acceleration enables the RL agent to efficiently derive optimal toolpaths to minimise thermal distortion and ensure thermal uniformity. Experimental results indicate that the 3D U-Net accelerates thermal field prediction by reducing computation time by over 99.8% compared to FDM simulations. This rapid inference capability compresses the process from hours to seconds and enables efficient process planning. Based on these fast predictions, the RL-optimised toolpaths achieved a warpage reduction of 95.99% on the pyramidal geometry compared to the conventional toolpath. This study establishes a scalable paradigm for intelligent manufacturing and provides a robust solution to thermal distortion in complex arbitrary geometries.

ARTICLE HISTORY

Received 30 December 2025
Accepted 31 January 2026

KEYWORDS

Reinforcement learning;
additive manufacturing;
thermal distortion; surrogate
model; toolpath
optimisation

1. Introduction

Additive manufacturing (AM) is a revolutionary technology that fabricates three-dimensional (3D) objects through a layer-by-layer deposition process. This approach presents an alternative to conventional manufacturing, enabling the manufacturing of complex geometries while substantially minimising material waste [1]. Furthermore, AM facilitates the production of functional components by embedding sensors for critical applications such as process monitoring [2] and leak detection [3,4].

However, AM processes are inherently characterised by severe thermal gradients, which frequently induce residual stresses and defects. Since AM utilises high-energy sources to rapidly melt and deposit material, the localised heat input generates steep temperature differentials. The subsequent rapid heating and cooling cycles cause thermal expansion and contraction, which

are mechanically constrained by the surrounding solidified layers or substrate [5]. Furthermore, the non-uniform temperature distribution within the part directly governs the evolution of residual stresses [6]. Such severe thermal gradients can compromise the structural integrity of the final component, leading to distortions including warpage and delamination [7–10]. Given these thermal dynamics, the toolpath, or scanning strategy, determines the thermal history of the fabricated part, thereby directly influencing the accumulation of residual stresses and the resulting mechanical properties [11,12]. Consequently, extensive research has been conducted to optimise scanning strategies for distortion mitigation. For instance, a 90° alternating scanning strategy was demonstrated to yield the lowest residual stress among various alternating patterns [13]. Furthermore, substrate deformation was minimised by employing a zigzag pattern aligned with the short edge [14]. In

CONTACT Im Doo Jung  idjung@unist.ac.kr

[†]These authors contributed equally to this work.

 Supplemental data for this article can be accessed online at <https://doi.org/10.1080/17452759.2026.2627765>.

© 2026 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

addition to these conventional approaches, advanced toolpaths such as rotating spiral [15] and S-shaped patterns [16] have been developed to further reduce temperature variations. However, the applicability of existing findings remains restricted to simple geometries and pre-defined patterns, primarily due to the prohibitive computational costs of thermal simulations.

In response to computational constraints, artificial intelligence (AI) has emerged as a pivotal technology, capable of analysing complex, non-linear relationships from large datasets for efficient decision-making. Consequently, AI methods have been extensively employed to optimise process parameters [17,18] and toolpath strategies [19–22] in AM. For instance, convolutional neural networks were utilised to rank generated toolpaths by accelerating simulations in the selective laser sintering process [19], while recurrent neural networks were employed to predict temperature fields induced by each candidate toolpath for subsequent layers [20]. Similarly, artificial neural networks have been applied to suggest toolpath corrections to mitigate void formation at junctions [21], and the k-nearest neighbours method was adopted to achieve uniform temperature distributions by generating optimal toolpaths [22]. Despite these advancements, supervised learning approaches necessitate extensive datasets, leading to significant time constraints on the data acquisition process. Furthermore, the predictive scope is confined to the specific domain of the training data, thereby limiting these models to passive evaluators of predetermined toolpaths rather than active optimisers.

Reinforcement learning (RL) and genetic algorithms (GA) are considered robust alternatives to these limitations, enabling the direct optimisation of process parameters and toolpath patterns. For example, GA was employed for efficient continuous toolpath planning to achieve enhanced thermomechanical properties [23]. RL approaches have been implemented to optimise parameters such as laser power and scan speed [24], and to reduce toolpath planning time for thin-walled structures in wire arc AM [25]. Furthermore, deep RL has been utilised to generate toolpaths to achieve uniform thermal distribution in laser powder bed fusion processes [26]. Nevertheless, these methods are predominantly limited to single layers or simple geometries, lacking generalizability to complex 3D shapes due to computational constraints. In particular, the reliance on finite element method (FEM) simulations for RL evaluation creates a significant computational bottleneck, making each iteration excessively time-consuming. Consequently, replacing FEM simulations with efficient surrogate models, such as random forest [27] and U-Net [28], is essential to accelerate the optimisation process. These models

have been extensively applied to predict microstructure [29], melt pool temperature [30], and residual stress [31]. However, the use of surrogate models to predict the global thermal distribution for toolpath optimisation remains limited, creating a need for a dedicated framework to accelerate this process.

In this study, a novel RL-driven strategy integrated with 3D U-Net was developed to accelerate toolpath optimisation for 3D geometries, aiming to minimise thermal distortion and achieve uniform temperature distribution. For each layer of the given shape, the RL model is trained to select the optimal toolpath that minimises thermal gradients among several candidate patterns. By employing the 3D U-Net to evaluate the temperature distribution instead of computationally expensive simulations, this proposed framework significantly reduces computational complexity and enables the efficient optimisation of complex 3D geometries. The optimised toolpaths were validated through fused filament fabrication (FFF) printing, and their effectiveness was assessed by quantifying the magnitude of warpage in the specimens.

2. Materials and methods

2.1. Overview of pipeline

To mitigate thermal distortion, layer-by-layer toolpath optimisation is essential for ensuring thermal uniformity, as illustrated in Figure 1(a). However, current AM processes employ fixed scanning strategies, such as zigzag patterns, irrespective of the specific 3D geometry, relying on conventional rule-based slicing algorithms. Consequently, applying these conventional algorithms across varying geometries can induce significant thermal gradients, resulting in warpage or deformation (Figure 1(b)).

In this study, we integrate RL and a 3D U-Net model to significantly accelerate the toolpath optimisation for a given 3D geometry (Figure 1(c)). The RL agent is trained to select the optimal toolpath for each layer with the objective of maximising thermal uniformity. To enable efficient state estimation, the 3D U-Net is trained on the finite difference method (FDM)-based simulation datasets to predict the temperature field based on the input toolpath and object geometry. This integrated architecture effectively transforms the traditionally time-consuming simulation feedback loop into a rapid predictive step, making real-time, geometry-aware toolpath optimisation feasible for complex parts. Details regarding the FDM simulations are provided in Section 2.2, followed by the random 3D geometry generation algorithm outlined in Section 2.3.

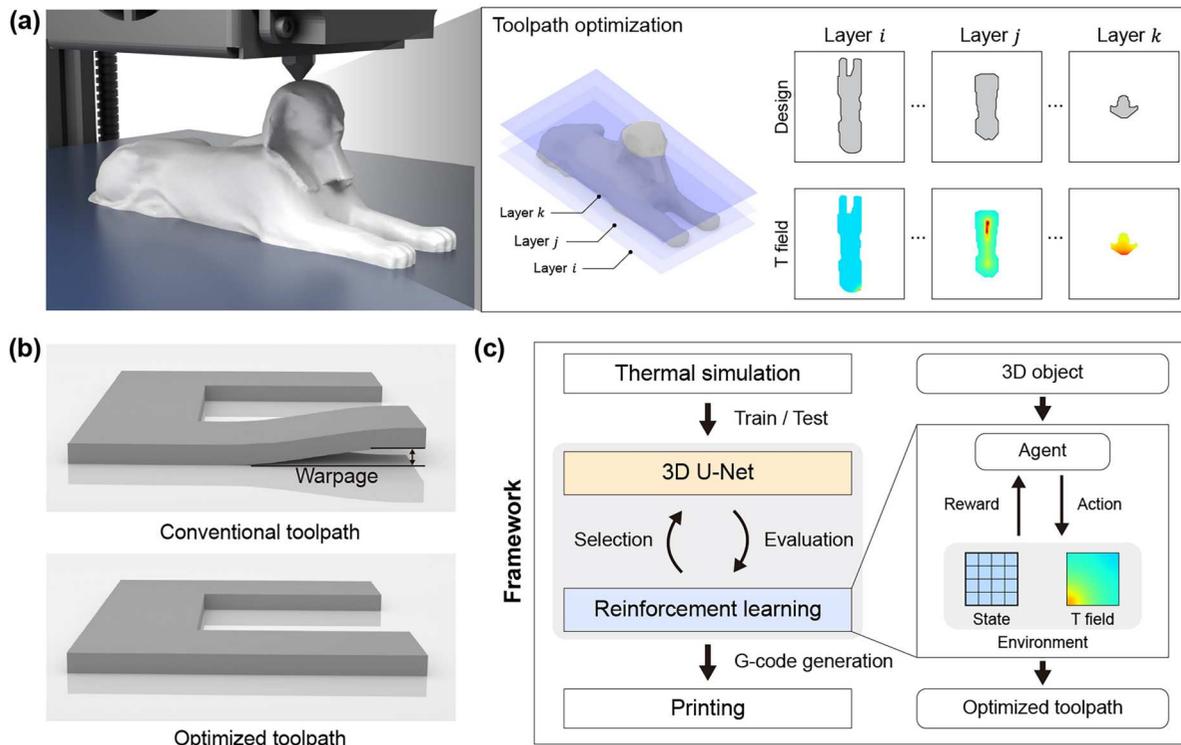


Figure 1. (a) Schematic overview of the layer-by-layer toolpath optimisation strategy. (b) Impact of toolpath strategies: conventional toolpath primarily causes deformation driven by thermal gradients, whereas the RL-optimised toolpaths minimise thermal distortion in the printed object. (c) Optimisation framework integrating RL with a 3D U-Net. The 3D U-Net was trained on FDM-based simulation datasets for accelerated thermal prediction. The RL algorithm leverages these predictions to select optimal toolpaths that ensure thermal homogeneity and mitigate thermal distortion.

Subsequently, the integrated components of our framework, the 3D U-Net and RL, are fully described in Section 2.4 and Section 2.5, respectively. Lastly, the materials and experimental setup are presented in Section 2.6.

2.2. Finite difference method simulation

FDM simulation is utilised to derive the temperature field from the given toolpath and geometry. The thermal dynamics modelled during deposition include three principal heat transfer mechanisms: the heat source input from the nozzle, convection and radiation to the surrounding environment, and internal conduction between deposited elements and the build platform (Figure 2(a)). The dataset is consequently compiled using a methodology that integrates a random geometry generation algorithm with stochastic toolpath selection, as shown in Figure 2(b). This random geometry algorithm is developed to output diverse 3D objects, including features with minor overhang, which is crucial for achieving robust generalisation during the 3D U-Net training phase (Section 2.3).

Upon importing the 3D model for framework evaluation, the continuous 3D object geometry must be converted into a discrete computational domain suitable for

both the thermal prediction and the RL model, as illustrated in Figure 2(c). This transformation is achieved via voxelization with Trimesh, where the target geometry is partitioned into a uniform grid of volumetric nodes. A voxel size of $0.5 \times 0.5 \times 0.25 \text{ mm}^3$ is utilised for the 128^3 grid to ensure a realistic representation of the physical part geometry. The horizontal resolution of 0.5 mm is selected based on the measured bead width, which exceeds the nozzle diameter due to the die swell effect of the molten material. Additionally, the 0.25 mm voxel height provides a direct one-to-one mapping with the deposition layer height. This process yields a design array that accurately maps the layer-wise topology of the object. Subsequently, this discretized array serves as the primary input format for both the RL and the 3D U-Net surrogate model.

To significantly reduce the substantial computational cost associated with transient thermal simulation, an active element strategy is employed (Figure 2(d)). In this approach, only the nodes corresponding to the deposited material are designated as active elements and are included in the thermal calculation. Once activated through the deposition sequence, these nodes remain in an active state for the remainder of the process. Conversely, nodes representing the

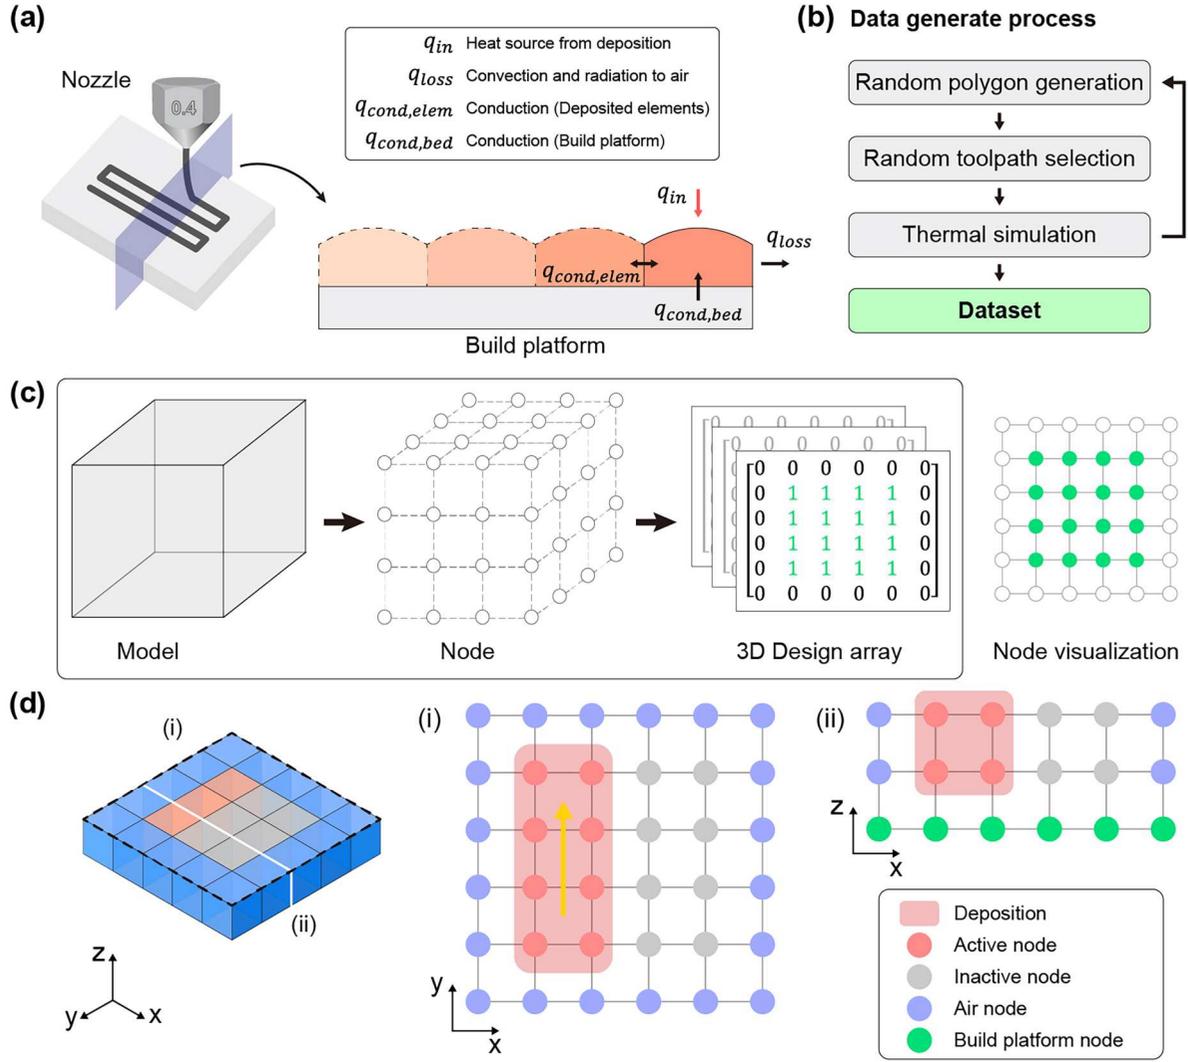


Figure 2. (a) Schematic of thermal dynamics during deposition. The cross-sectional view visualises heat transfer mechanisms: heat input from deposition, internal conduction, and convection/radiation to the environment. (b) Data generation pipeline. Simulations utilising random toolpaths on stochastic geometries are conducted to compile the dataset for 3D U-Net training and validation. (c) Voxelization process where the 3D geometry is discretized into nodes to construct a design array, serving as the input for the RL and 3D U-Net. (d) Active element strategy for computational efficiency. Only deposited material nodes are considered active, while air and platform nodes act as boundary conditions. The XY and XZ cross-sections illustrate the sequential activation of the nodes.

surrounding air and the build platform are designated as inactive and serve solely as thermal boundary conditions. The cross-sectional views (XY and XZ planes) illustrate this sequential activation, where elements are only activated in the simulation domain as the material is deposited layer-by-layer.

The thermal modelling of the FFF process is governed by the transient heat equation based on the principle of energy conservation. The temperature evolution within the domain is expressed as:

$$\rho c_p \frac{\partial T}{\partial t} = k \nabla^2 T + q_{in} - q_{loss} \quad (1)$$

where ρ , c_p , and k are the density, specific heat capacity, and thermal conductivity of the material, respectively.

The term $k \nabla^2 T$ accounts for internal conduction ($q_{cond,elem}$) within the deposited elements. The thermal input, q_{in} , represents the heat source from the nozzle, and the total heat loss, q_{loss} , incorporates both convection (q_{conv}) and radiation (q_{rad}) to the environment. The nozzle heat input (q_{in}) is implemented using a temperature Dirichlet boundary condition maintained at 230 °C. The mapping between deposition time, voxel activation, and energy addition is achieved through a precise spatial and temporal alignment. Specifically, a voxel is activated via an element birth method at the precise moment the nozzle centre reaches its spatial boundaries. Upon activation, the fixed temperature of 230 °C is applied to the voxel for a duration equivalent to the deposition time, which is determined by the ratio of

Table 1. Thermal and physical parameters for simulation.

Material property and parameter	Value	Reference
Thermal conductivity k [W/(m · K)]	0.33	[35]
Specific heat capacity c_p [J/(kg · K)]	1470	
Density ρ [kg/m ³]	945	
Convective heat transfer coefficient h [W/(m ² · K)]	10	
Filament emissivity ϵ	0.92	
Stefan–Boltzmann constant σ [W/(m ² · K ⁴)]	5.67×10^{-8}	–
Printing speed v_p [mm/s]	30	
Travel speed v_t [mm/s]	150	
Ambient temperature T_∞ [°C]	25	
Platform temperature T_p [°C]	80	
Extrusion temperature T_e [°C]	230	

the voxel dimension to the nozzle printing speed. Furthermore, the build platform ($q_{cond,bed}$) is modelled using a fixed temperature Dirichlet boundary condition maintained at 80 °C. This widely established approach reduces computational complexity while accurately resolving the steep thermal gradients at the interface that drive residual stress accumulation, which is validated by high consistency with experimental results [32]. Table 1 summarises the thermal and physical parameters employed for the simulation. To facilitate the numerical simulation and maintain computational efficiency, the following assumptions were made: the printing material is isotropic and homogeneous; thermo-physical properties remain constant and temperature-independent; latent heat effects during phase transitions are neglected; deposited filaments form a continuum with perfect contact, ignoring contact thermal resistance; and heat loss is modelled using constant transfer coefficients with a uniform ambient temperature. The validity of these simplified assumptions is supported by both experimental observations and comparative evaluations with high-fidelity simulations. Existing literature indicates that the influence including temperature-dependent properties and the continuum assumption for deposited filaments on the global thermal trend is negligible within the operational range of the FFF process [32–34]. Furthermore, comparative studies using high-fidelity numerical models confirm that the global thermal distribution and heat dissipation patterns remain consistent despite the simplifications. Further details are provided in Supplementary Section S1 and Figure S1. These results demonstrate that the established modelling framework provides a reliable representation of the thermal behaviour during the material deposition process.

2.3. Random 3D polygon generation

To ensure robust training of the 3D U-Net, a random 3D polygon generation algorithm was developed using

Python libraries including OpenCV, SciPy, and scikit-image (Figure 3(a)). As illustrated in Figure 3(b), the generation process iteratively constructs the 3D object through layer-wise morphological evolution, primarily utilising two stochastic modes: indentation and overhang.

The process begins by initialising the first layer with random active area dimensions. For subsequent layers, the algorithm stochastically selects between two modification modes to introduce geometric variability. In mode 1 (indentation), the algorithm extracts the largest contour of the previous layer and identifies random points along the boundary. Rectangular regions of randomised sizes (w , h) centred at these points are then removed to create indentations. In mode 2 (overhang), the boundary is expanded via morphological dilation using a cross kernel. From the newly generated candidate regions, a random subset is selected and added to the design, constructing an overhanging structure.

To ensure geometric continuity and realistic features, morphological smoothing is applied after each modification: an elliptical kernel is used for indented layers, while a rectangular kernel is applied to expanded layers. A 3×3 cross kernel is utilised during the random geometry generation process to restrict the horizontal expansion to a maximum of one grid unit per layer. This geometric constraint ensures that the resulting overhang angle remains below 60°, which is within the 75° threshold established for unsupported printing in the FFF process [36]. Consequently, structural stability is maintained, and the numerical solver remains highly representative of real-world conditions without unintended thermal biases in the surrogate model. Furthermore, the generated geometries for the dataset feature a height of up to 128 layers and a minimum vertical feature size of 15 layers to ensure numerical stability during the thermal simulation process. Subsequently, connectivity is checked using connected component analysis to retain only the largest component, eliminating isolated regions. This sequence repeats for a pre-defined number of layers or until no geometry remains, resulting in a continuous and diverse 3D polygon geometry.

2.4. 3d U-Net surrogate model for temperature field prediction

The results obtained from the FDM simulation are utilised as the input data for the 3D U-Net, which consists of four distinct information channels: toolpath traversal order, previous temperature field, shape mask, and layer index, as illustrated in Figure 4(a). The toolpath traversal order indicates the sequence in which the print

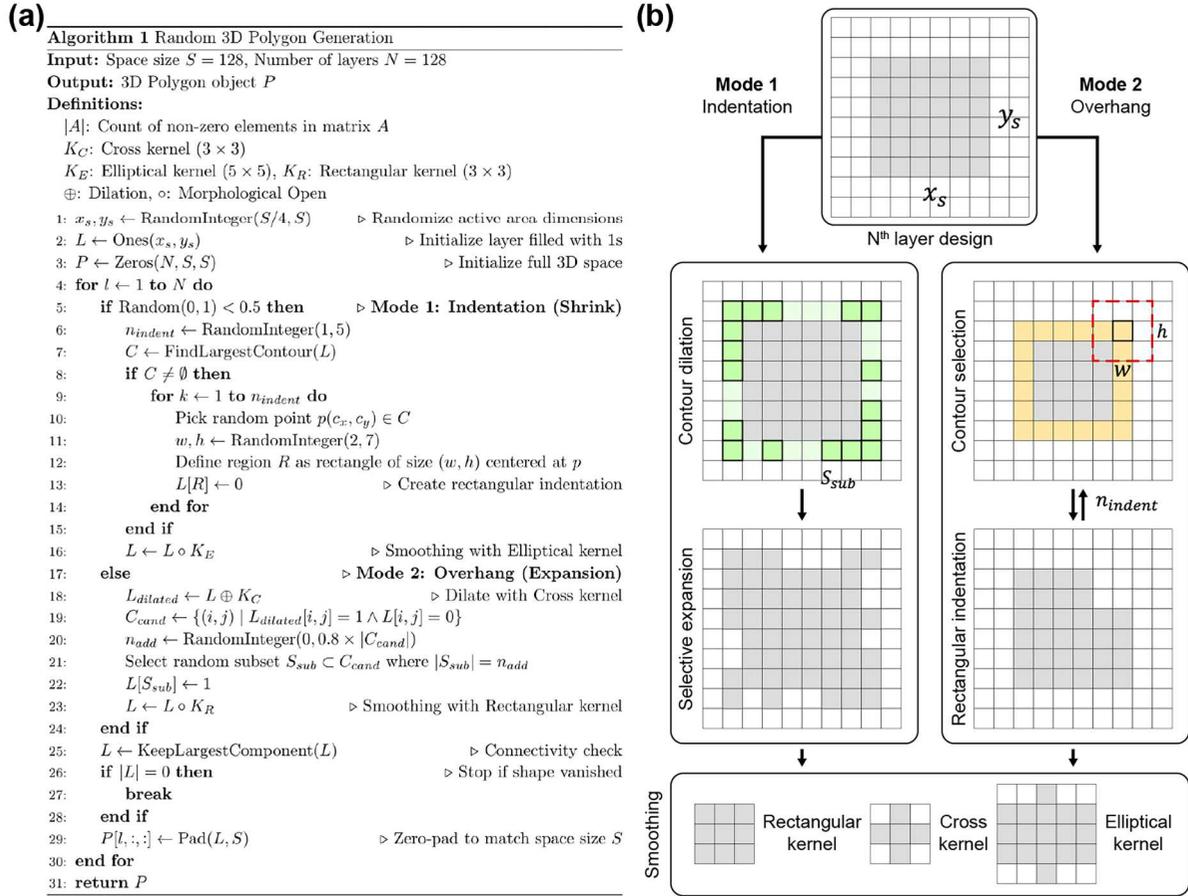


Figure 3. Overview of the random 3D polygon generation method. (a) Pseudocode detailing the procedural generation logic, which iterates through layers to construct the 3D geometry. (b) Schematic of the shape evolution process within a single layer. The algorithm stochastically selects between two modification modes: Indentation (Mode 1), which creates inward features by removing rectangular regions, and Overhang (Mode 2), which expands the boundary via morphological dilation and selective addition. Morphological kernels are utilised for the smoothing process to maintain geometric continuity.

nozzle deposits material across the layer, which is crucial for modelling the transient thermal history. The previous temperature field provides the thermal state of the part inherited from the fabrication of all preceding layers, serving as a critical boundary condition for the current thermal prediction. The initial temperature for any newly activated layer or element that has not been previously deposited is set to the ambient temperature. The temperature values are scaled using min–max normalisation to the range of 0–1 to ensure numerical stability during the training process. The shape mask represents the region of the material deposition for each layer. Lastly, the layer index serves as a depth indicator, providing the model with contextual information regarding the current height within the entire 3D geometry. All numerical input values are mapped and visualised using a grayscale colormap for consistent data representation.

To reconstruct the thermal field, 3D U-Net can be employed as the surrogate model designed to rapidly

predict the transient temperature field of the printed part. This 3D U-Net adopts an encoder-decoder architecture, leveraging 3D convolutional layers to effectively process the volumetric input data and capture the complex spatial correlations within the printed part [28]. Furthermore, the incorporation of residual modules, where the input is added directly to the output of the convolutional blocks, and skip connections significantly enhances the network’s ability to propagate crucial gradient information between the encoder and decoder levels [37]. Nevertheless, to address the limitations of convolutional operation restricted to local receptive fields, squeeze and excitation (SE) blocks were integrated to comprehensively consider both spatial feature information and channel dependencies [38]. These blocks explicitly capture the interdependencies between channels to adaptively recalibrate feature responses. This mechanism enables the network to selectively emphasise informative features while suppressing less relevant ones, thereby enhancing the

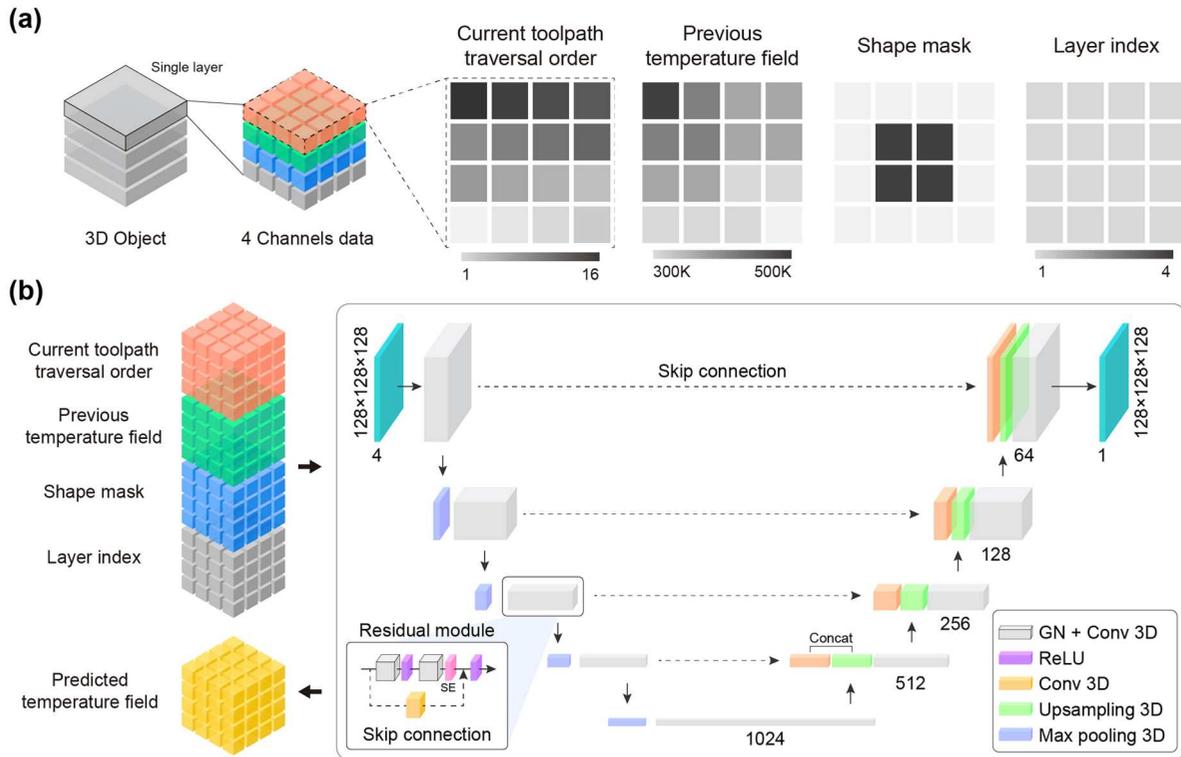


Figure 4. (a) Input representation for the 3D U-Net with four channels: toolpath traversal order, previous temperature field, shape mask, and layer index. Values are visualised using a grayscale colormap. (b) Architecture of the proposed 3D U-Net. The model utilises the four-channel input to predict the subsequent temperature field, integrating residual modules and skip connections.

representational capacity required for accurate thermal field reconstruction. As this structural architecture enhances training stability and model performance, the residual SE U-Net is employed to ensure precise reconstruction of the temperature field (Figure 4(b)).

In this study, the 3D U-Net utilised a four-channel input and produced a single-channel output, representing the predicted temperature field. A total of 80 geometries created via a random algorithm were employed to generate 4,236 samples for the training dataset, while 20 distinct geometries were utilised to produce 927 samples for the validation dataset. The network's depth was defined by its feature map configuration, with the number of feature channels in the convolutional layers progressing from 64, 128, 256, 512, and 1024 across the encoding path. Training was conducted over a maximum of 100,000 iterations. The optimisation process employed the AdamW optimiser with an initial learning rate of 2×10^{-4} . To ensure stable convergence, the ReduceLROnPlateau scheduler was employed to dynamically regulate the learning rate based on validation performance assessed every 200 iterations. In this configuration, the learning rate was multiplied by a factor of 0.1 if the validation loss failed to improve for 20 consecutive intervals, with a minimum threshold set at 10^{-7} . The model's performance was evaluated

using the mean squared error (MSE) loss function. The 3D U-Net was trained on an Ubuntu 22.04 server equipped with an AMD Ryzen Threadripper 3990X CPU and three NVIDIA RTX 3090 GPUs.

2.5. Reinforcement learning for toolpath optimisation

As shown in Figure 5(a), RL is characterised by the interaction of five essential elements: the agent, environment, state, action, and reward [39]. Within this framework, the agent serves as the decision-maker interacting with the environment. At each time step, the agent observes the current state, executes an action based on its policy, and receives a reward which serves as a feedback signal. Through a process of trial and error within this Markov decision process framework [40], the agent iteratively refines its actions based on the received feedback to maximise the expected cumulative reward. In this study, the proximal policy optimization (PPO) algorithm was employed, a robust policy gradient method that utilises a clipping mechanism to prevent drastic policy updates. This algorithm was selected for its demonstrated stability and superior computational efficiency in handling complex control tasks compared to alternative RL approaches [41]. Here, the

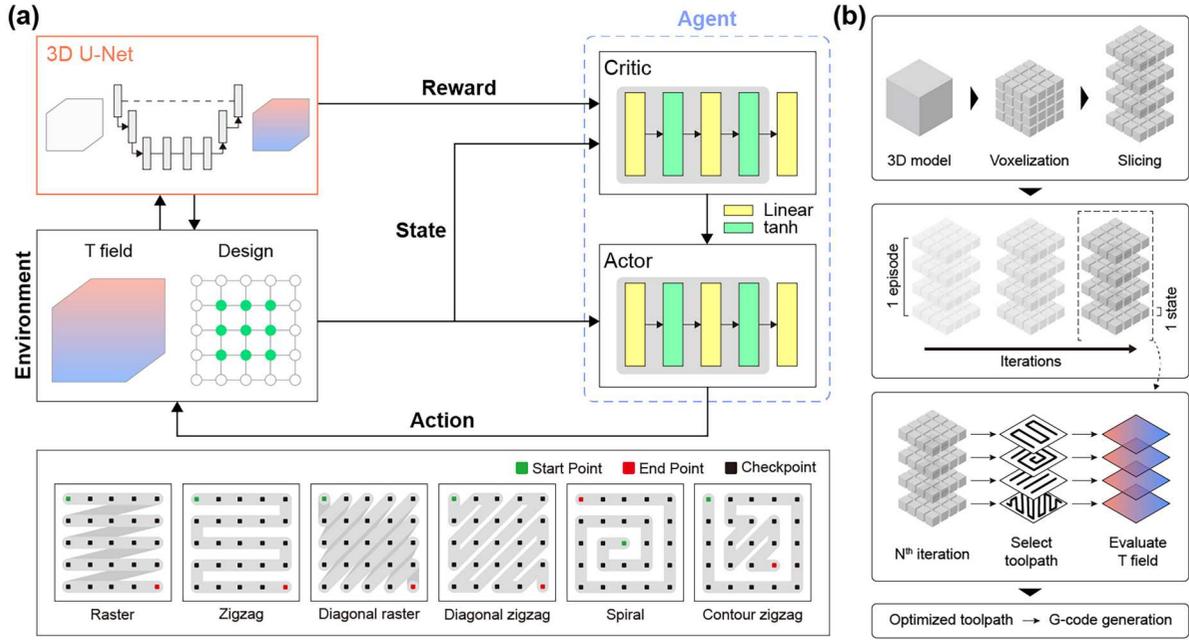


Figure 5. (a) Schematic of the RL framework. The 3D U-Net serves as an environment surrogate, evaluating the agent's selected action and returning the updated temperature field and state. The agent learns to identify optimal toolpaths to minimise thermal distortion across all layers. The action space is defined by 38 distinct toolpath patterns. (b) RL training workflow. Following voxelization and slicing, the agent optimises toolpaths layer-by-layer. The training process iterates for a predetermined number of episodes, with the number of steps per episode corresponding to the total layer count. Upon completion of training, the optimised toolpaths are generated as G-code files for 3D printing.

PPO framework employs an actor-critic architecture with feature extractors to process high-dimensional 3D thermal fields and 2D layer designs (Figure S2(a)). These modules utilise downsampling units to extract essential spatial features while progressively reducing data dimensions, as shown in Figure S2(b). Residual connections are integrated within these blocks to mitigate vanishing gradients and preserve original spatial information throughout the network. Additionally, group normalisation provides a consistent normalisation mechanism independent of batch configuration to ensure robust training stability and convergence in high dimensional observation spaces [42]. The resulting feature vectors are concatenated for the actor and critic networks to facilitate thermal-guided toolpath optimisation.

A 3D U-Net provides a reward to the critic network for evaluating thermal uniformity. The reward is calculated at each step as the negative standard deviation of the temperature values across all activated nodes:

$$R_{step} = -\sqrt{\frac{1}{N} \sum_{i=1}^N (T_i - \bar{T})^2}$$

where T_i denotes the temperature of the i -th activated node, \bar{T} represents the mean temperature of all nodes activated at the current step, and N is the total number of these nodes. The total reward for a single

episode is defined as the cumulative sum of these individual step rewards:

$$R_{episode} = \sum R_{step}$$

Each episode within this framework consists of a sequence of steps equivalent to the number of layers in the target geometry. The reward function incorporates global temperature distribution data to promote thermal uniformity throughout the fabrication process. Metrics such as the min-max range or spatial gradient norms are excluded as their representation is limited to localised thermal fluctuations rather than global uniformity. Standard deviation is selected as the primary metric because it remains relatively scale-invariant across varying numbers of activated nodes, unlike the variance or L_2 norm. The L_2 norm, being scale-dependent, tends to place excessive emphasis on rewards during the subsequent stages of the process as the node quantity increases. Given that warpage is often critical in the initial layers, such a bias may hinder the optimisation objective with the physical requirements of the build. This selection also prevents numerical instabilities, such as vanishing or exploding gradients, during the PPO training process. Furthermore, a negative reward structure is implemented to prevent premature convergence to sub-optimal behaviours. Both the

generalisation capability and behavioural diversity of the agent are enhanced by defining the objective as penalty minimisation, which encourages continuous exploration of the policy space [43]. Based on this assessment, the actor network determines the optimal toolpath from a discrete action space of 38 distinct patterns (Table S1; Green and red markers indicating the start and end points, respectively).

The toolpath optimisation process begins with voxelization and slicing, which are essential for converting the 3D geometry into layer-wise inputs for the RL model (Figure 5(b)). The agent is trained in 500,000 time steps, where each episode comprises a sequence of steps equal to the total number of sliced layers. Performance evaluation is conducted every 50 episodes, and model parameters are exported whenever the moving average reward of the last 100 episodes shows improvement. In the early stages where fewer than 100 episodes have been completed, the average reward is calculated using all available episodes. The observation space is defined with dimensions of $128 \times 128 \times 128$ (Height \times Width \times Depth). The temperature field and design data are transformed and convolved to process the input features. For the PPO hyperparameters, the learning rate was set to 0.0003 to control the gradient descent step size. To ensure training stability, the clipping range was fixed at 0.2, thereby constraining the magnitude of policy updates.

2.6. Materials and experimental process

The experiments were conducted using an Ultimaker S3 printer, which operates based on the material extrusion technique, commonly known as FFF. Acrylonitrile butadiene styrene (ABS) was employed as the printing material. To ensure experimental consistency, the 3D printing procedures and process variables were maintained at constant values, with the toolpath pattern being the only exception. The key process parameters were set as follows: a nozzle diameter of 0.4 mm, a nozzle temperature of 230°C, a platform temperature of 80°C, a filament diameter of 2.85 mm, and a layer thickness of 0.25 mm. Key thermal and mechanical properties of the ABS material, essential for the FDM simulations, are listed in Table 1. Even when the platform temperature was maintained above 80°C geometric deviations were observed in the form of either warpage or compressive deformation (Figure S3). Consequently, optimising the toolpath is essential to mitigate these issues beyond the limitations of thermal control.

A total of 38 distinct toolpath patterns were included for selection, comprising fundamental strategies such as raster, zigzag, and spiral, along with established hybrid

patterns like contour-zigzag combinations (Table S1). These basic scanning strategies have been extensively analysed for their impact on thermal and mechanical properties in AM [44,45]. In this study, the contour with zigzag pattern was designated as the representative conventional toolpath. Consequently, the experimental evaluation focuses on comparing the performance of this conventional toolpath against the RL-optimised toolpath generated by our proposed framework. The G-code defining the selected toolpath and geometry was programmatically generated using Python and exported into the requisite G-code format. Since the toolpath fundamentally consists of the sequential coordinates of the print nozzle, the G-code generation process involved mapping these coordinate sequences directly into the commands for the 3D printing environment. To adapt to arbitrary designs, the pre-defined toolpath masks are instantiated through a spatial intersection process where the patterns are clipped to the target geometry (Figure S4). The extracted cells are re-indexed to maintain the original spatial logic, ensuring a consistent path sequence regardless of the shape's complexity. Furthermore, the boundary contours are extracted using OpenCV's contour detection functions and integrated into the toolpath order to ensure high geometric fidelity. This mechanism ensures robust adaptability by addressing complex geometries without gaps.

3. Results and discussion

Rectangular, U-shaped, and pyramidal geometries were selected as target shapes for analysis, as they induce significantly different temperature gradients depending on the toolpath patterns. The dimensions of these objects are illustrated in Figure 6(a). The selected objects are utilised not only for warpage analysis due to their considerable aspect ratio [9], but also feature a toolpath that traverses an extensive distance within a single layer, concentrating the heat from the nozzle in a specific area. This enables a comparison of the degree of warpage resulting from the toolpath.

Figure 6(b) visualises representative input features, presented as a specific cross-sectional slice for clarity, while the model inherently processes volumetric data to fully capture the inter-layer thermal history. The input tensor integrates four distinct channels, the current toolpath traversal order, previous temperature field, shape mask, and layer index, enabling the network to correlate geometric and temporal deposition parameters with the resulting temperature distribution. Specifically, the toolpath traversal order of the layer currently being deposited is replicated across every layer of

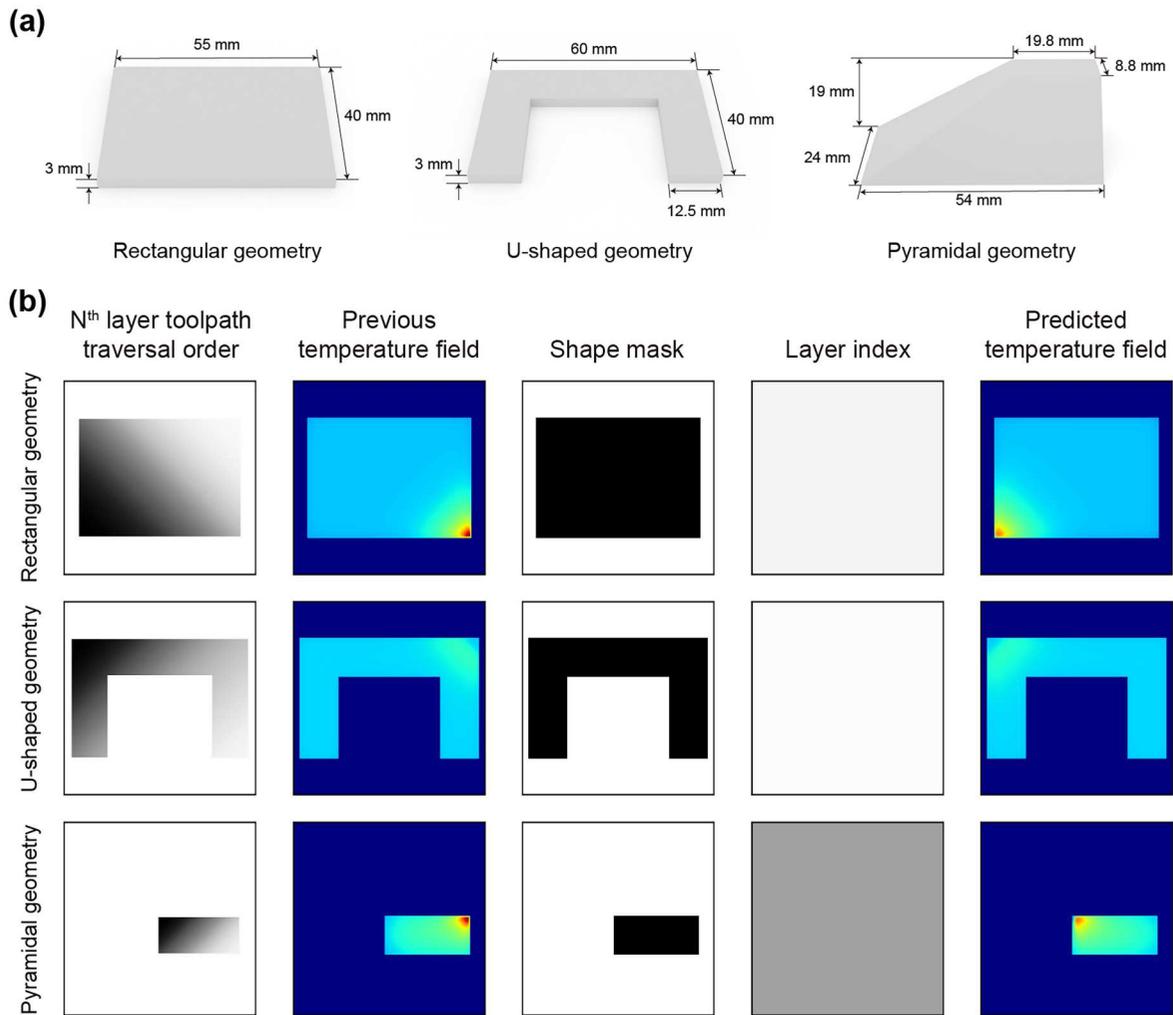


Figure 6. Target geometries for analysis and corresponding training data samples. (a) Geometric specifications and dimensions of the rectangular, U-shaped, and pyramidal geometries selected for comparative evaluation. (b) Representative data instances used for the validation of the 3D U-Net model. For clarification, specific cross-sectional slice for each geometry is depicted, although the actual input consists of volumetric data spanning multiple layers. Each data sample comprises four distinct feature channels: current toolpath traversal order, previous temperature field, shape mask, and layer index.

the 3D input volume, ensuring that the entire input stack contains uniform information representing the current deposition step. Consequently, the 2D toolpath sequence of the topmost layer is replicated along the depth dimension of the input tensor, allowing the 3D U-Net to consistently receive the current toolpath data across all spatial layers for each prediction step. This configuration allows the model to capture the global thermal context, reflecting the physical phenomenon that heat input from the filament significantly influences the thermal states of underlying layers via conduction.

Regarding the thermal inputs, the temperature field from the previous step is utilised as an implicit representation of the cumulative thermal history. The existing temperature distribution enables the model to account for the transition from initial heating dynamics to quasi-steady behaviour without an explicit time variable.

The shape mask represents the geometric profile of individual layers to define the cumulative morphology of the part. Furthermore, the layer index encodes spatial height information to distinguish vertical progression. The combination of the layer index and the toolpath traversal order establishes the necessary spatiotemporal context for the 3D U-Net to predict accurate temperature fields based on the instantaneous thermal-geometric state. As demonstrated in Figures 7 and 8, this approach enables high-fidelity prediction of thermal distributions with minimal error throughout the entire fabrication process.

An ablation study was conducted to verify the effectiveness of residual connections and SE blocks using unseen data distinct from the validation set. This evaluation compared the predictive performance of the basic U-Net, residual U-Net, and residual SE U-Net against the

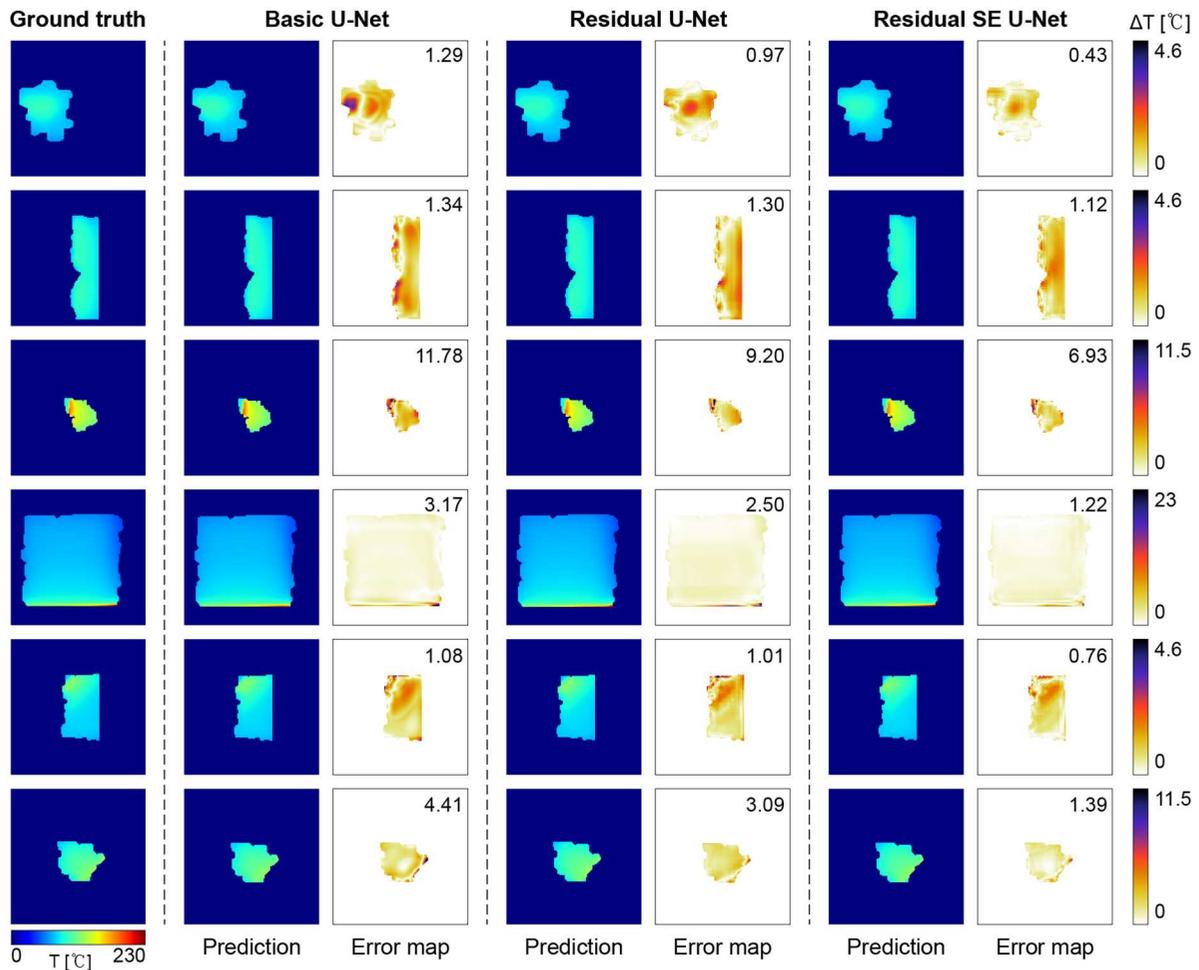


Figure 7. Quantitative comparison of predicted temperature fields: ground truth, basic U-Net, residual U-Net, and residual SE U-Net. The prediction fidelity is evaluated using the MSE, indicated at the top of each error map. The comparison reveals that incorporating residual connections and SE blocks significantly improves the prediction fidelity, with the residual SE U-Net showing the lowest deviation from the simulation data.

ground truth thermal fields (Figure 7). The basic U-Net exhibits relatively high deviations, particularly in regions with high thermal gradients, as evidenced by the corresponding error map. In contrast, the integration of residual connections significantly mitigates these deviations. The residual U-Net demonstrates improved boundary definition and reduced error magnitudes compared to the basic architecture. For quantitative comparison, the MSE was utilised to evaluate the prediction fidelity. The basic U-Net exhibited an MSE ranging from 1.08 to 11.78, whereas the incorporation of residual connections reduced this range to 0.97 to 9.2, corresponding to a 21.9% improvement in the maximum MSE. Furthermore, the residual SE U-Net achieved the lowest MSE range of 0.43 to 6.93, representing a substantial reduction of 41.2% in the maximum MSE compared to the basic U-Net.

The average and maximum prediction errors for each model are summarised in Table 2. Regarding the average

prediction error, the basic U-Net exhibited a range of 0.89 to 2.94°C whereas the residual SE U-Net achieved a reduced range of 0.50 to 2.16°C, demonstrating improved overall accuracy. In terms of maximum error, the basic U-Net showed large fluctuations, with the maximum error varying from 2.70°C to a peak of 60.95°C. In contrast, the residual SE U-Net effectively minimised these peak deviations, narrowing the maximum error range to 1.93 to 14.67°C. This reduction in the upper bound confirms that the proposed model maintains consistent predictive performance even in complex regions where the basic U-Net exhibits large errors. These results demonstrate that the integration of residual connections and SE blocks significantly reduces prediction errors against the simulation data, leading to accurate thermal field reconstruction. Consequently, the residual SE U-Net architecture was adopted as the 3D U-Net model for this study to ensure high-fidelity thermal predictions.

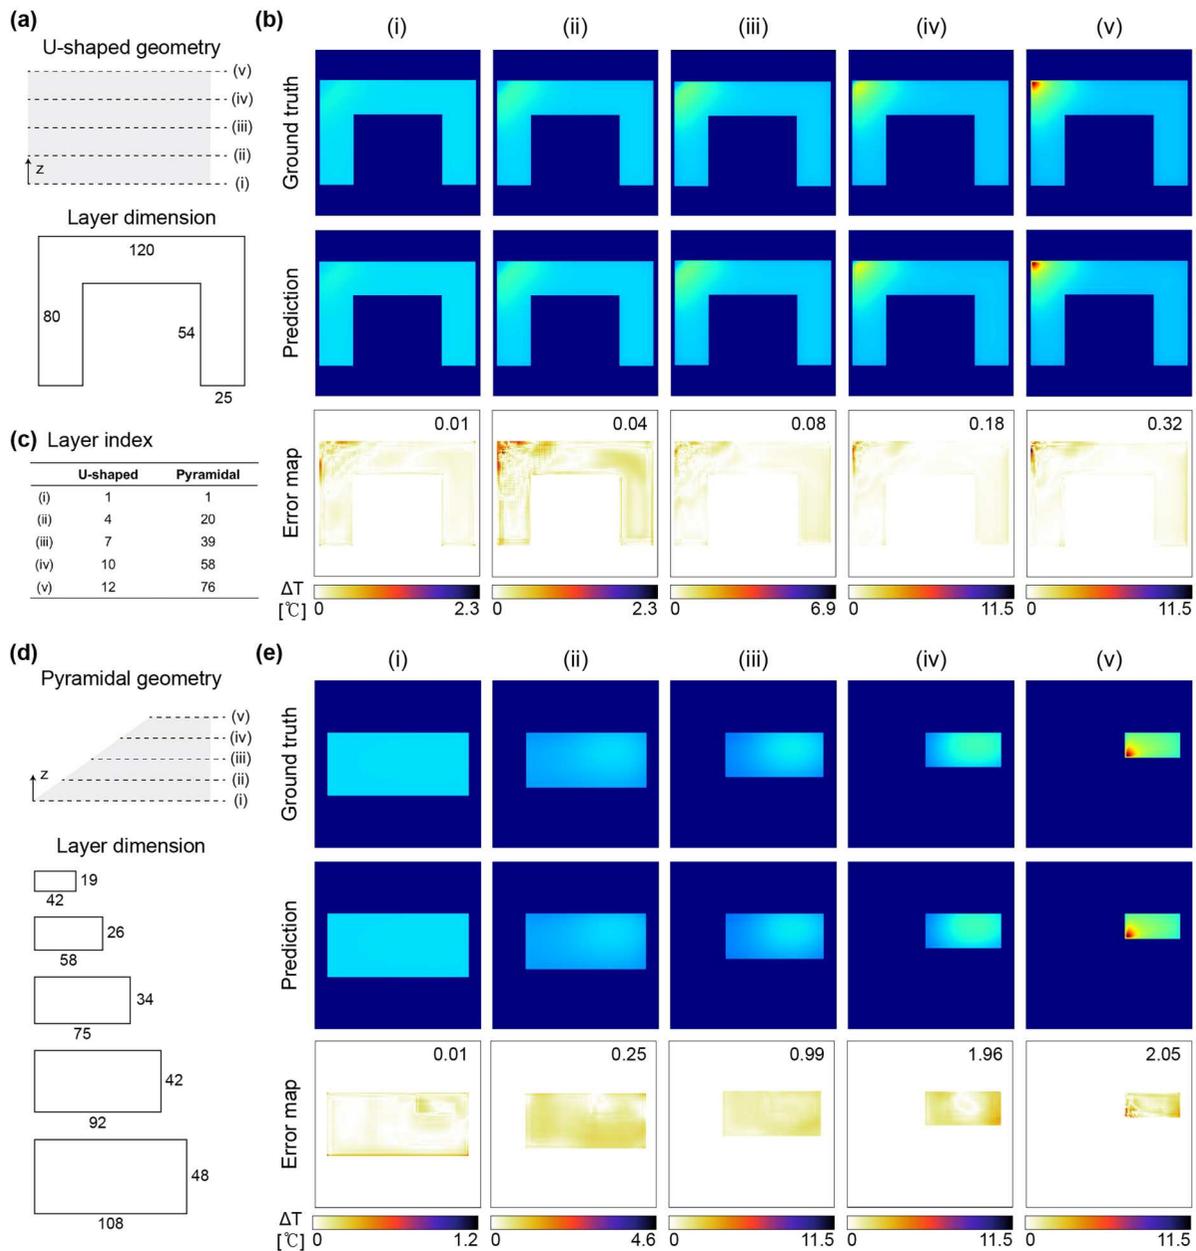


Figure 8. Layer-wise comparison of temperature distributions between the FDM simulation and the 3D U-Net prediction for different geometries. (a, d) Schematics of the U-shaped and pyramidal geometries, respectively, indicating five representative build heights (i–v) and the corresponding dimensions for each layer cross-section. (c) Table listing the layer indices corresponding to the representative heights for both geometries. (b, e) Detailed comparison of thermal fields at the selected heights for each geometry: ground truth, model prediction, and error map. The prediction fidelity is quantified by the MSE displayed at the top of each error map, demonstrating that the proposed model maintains high accuracy across varying layer geometries.

To evaluate the model’s adaptability to complex geometric features, a layer-wise validation was conducted using U-shaped geometry, as illustrated in Figure 8(a). The 3D U-Net predictions demonstrated exceptional stability across all layers as shown in Figure 8(b). The model accurately reproduced the local heat concentration at the corners, maintaining low MSE values ranging from 0.01 to 0.32. Specifically at the uppermost layer (v), where the highest MSE of 0.32 was recorded, the model achieved a low average error of 0.29°C, with

a peak error of 11.52°C. The corresponding layer indices for the five representative build heights (i–v), selected across the build direction, are listed in Figure 8(c).

In addition, a pyramidal geometry was analysed to evaluate performance under varying cross-sectional areas, as shown in Figure 8(d). The 3D U-Net predictions exhibit a high degree of consistency with the ground truth thermal fields across all layers, effectively capturing the thermal gradients despite the reduction in cross-

Table 2. Quantitative comparison of thermal prediction errors.

Model architecture	Average prediction error [°C]		Maximum prediction error [°C]	
	Min	Max	Min	Max
Basic U-Net	0.89	2.94	2.70	60.95
Residual U-Net	0.84	2.25	3.10	24.04
Residual SE U-Net	0.50	2.16	1.93	14.67

sectional area (Figure 8(e)). For all layers, the model maintains high fidelity with MSE values remaining below 2.05. Specifically, the peak error across all layers is limited to 6.72°C, with an average error of 1.17°C. This demonstrates that the model maintains robust predictive performance regardless of geometric complexity.

The generalisation capability of the surrogate model was further validated using an out-of-distribution test set consisting of multi-branched geometries, holes, and disconnected islands. These topological features were explicitly excluded from the training data. As illustrated in Figure S5, the model maintains high predictive fidelity across these unseen features. The recorded MSE values confirm its robust performance and applicability to diverse geometric configurations.

To determine the optimal network architecture for high-fidelity thermal field reconstruction, a comparative study was conducted by varying the network depth and the number of feature channels, as detailed in Figure 9(a). Three distinct configurations were evaluated to identify the balance between model capacity and learning efficiency: Case (i) utilises a deep structure ($N = 5$) with a high channel dimensions ranging from 64 to 1024; Case (ii) retains the depth ($N = 5$) but reduces the channel width by half (32 to 512) to evaluate the impact of parameter reduction; and Case (iii) employs a shallower architecture ($N = 4$) with channel dimensions (64 to 512). For a comparative evaluation of predictive fidelity, the MSE was calculated across the test dataset, comprising 821 data derived from 10 random geometries. The model configured as Case (i) achieved the lowest MSE of 0.71, demonstrating the superior capability of the deep architecture with expanded channel dimension in reconstructing the thermal field. The model in Case (ii) followed with an MSE of 0.91, showing comparable performance but slightly higher error due to the reduced channel dimension. In contrast, the shallower architecture of Case (iii) recorded the highest error of 1.07. These comparative results indicate that while network depth is the primary factor governing prediction accuracy, sufficient channel dimension is also essential for minimising errors. Consequently, Case (i) was selected as the backbone for the proposed 3D U-Net to maximise predictive fidelity and capture thermal variations.

Figure 9(b) presents the inference time comparison between the FDM simulation and the 3D U-Net model, utilising a logarithmic scale. Model inference is performed on a system equipped with an AMD Ryzen Threadripper 3990X CPU and an NVIDIA RTX 3090 GPU. Once training is complete, the inference latency is approximately 0.3 s per layer. This latency remains consistent as the RL and 3D U-Net models utilise fixed-size input tensors for each layer to ensure a uniform computational load regardless of the geometric complexity of the part. Our proposed approach significantly reduced the inference time for the rectangular, U-shaped, and pyramidal geometries from hours (4,066 s, 1,861 s, 12,773 s) to seconds (3.61, 3.63, 23.1 s). This corresponds to a consistent reduction of over 99.8% for all geometries. The advantage of the proposed 3D U-Net model is attributed to its scalability regarding the geometry size. In conventional FDM simulations, the computational cost increases exponentially as the build volume and height increase, because the number of mesh elements and time steps for thermal calculation grows proportionally. The simulation for the large-scale pyramidal geometry takes approximately 3.5 h. In contrast, the inference time of the U-Net remains on the order of seconds, showing only a marginal increase even for larger geometries. Since the computational complexity of the trained network is primarily determined by the network architecture rather than the sequential temporal integration of physics solvers, the proposed method is particularly advantageous for predicting the thermal history of large-scale industrial parts.

The training stability and convergence of the RL model, exemplified by the U-shaped geometry, are illustrated in Figure 9(c). The cumulative reward exhibits a distinct increasing trend as the number of episodes increases, indicating that the agent progressively improves its policy. Subsequently, the reward values stabilise after approximately 3,000 episodes, confirming that the model has successfully converged to an optimal policy for effective thermal management. To ensure the highest degree of optimisation fidelity, this study adopts a one-policy-per-geometry strategy necessitated by the complex nonlinear coupling between thermal evolution and localised geometric features. While such part-specific optimisation was previously regarded as computationally prohibitive, the proposed surrogate-based framework reduces the training duration to within several hours. This advancement facilitates a focus on individualised precision over broad generalisation to promote uniform thermal distributions.

In addition, a simulation-based evaluation was conducted to compare the efficacy of the RL-optimised toolpaths against a conventional toolpath across three

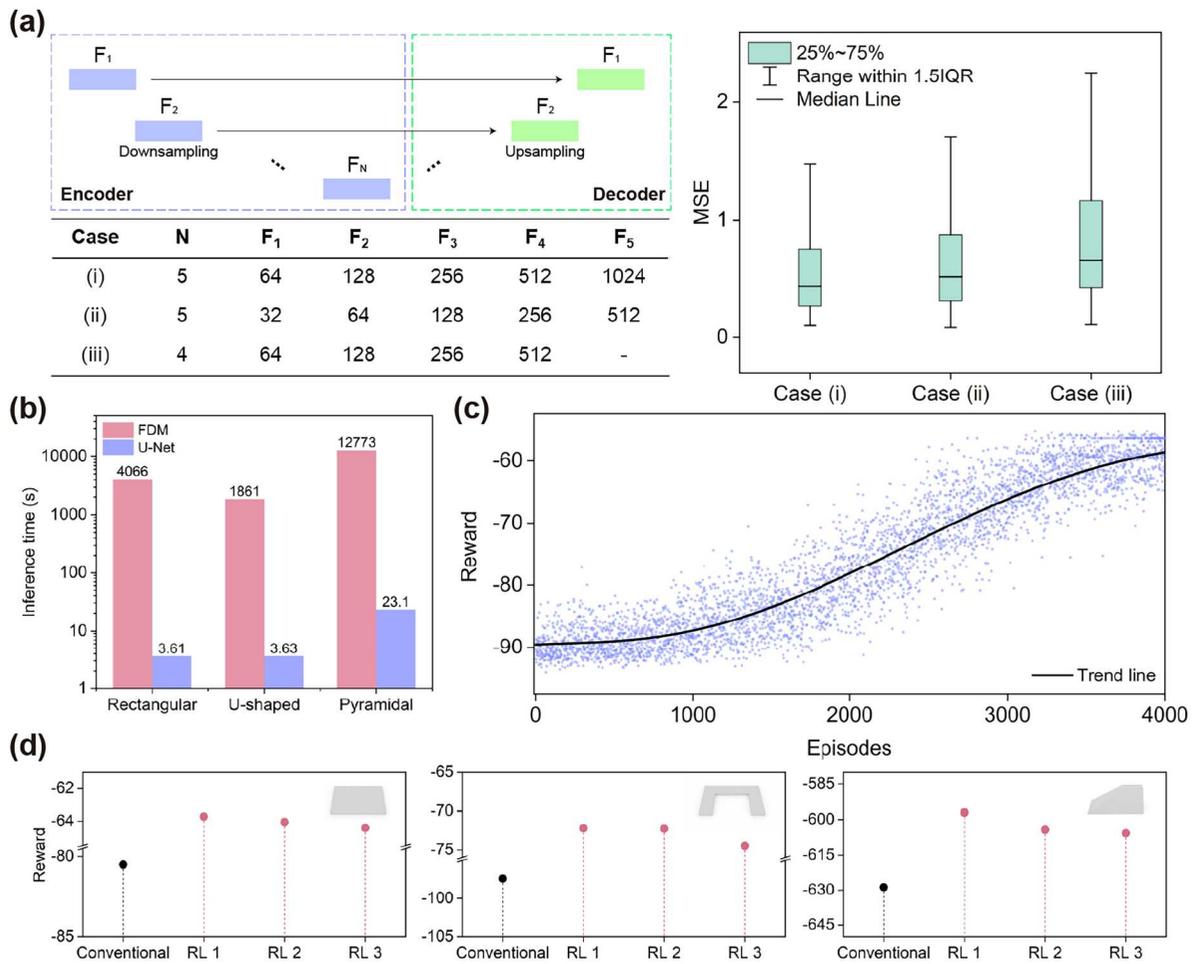


Figure 9. Validation of model training and optimisation performance. (a) Evaluation of 3D U-Net training performance across different architectural parameters. (b) Computational efficiency analysis demonstrating the acceleration from FDM-based hours to 3D U-Net-based seconds for three distinct geometries. (c) Learning curves for the RL model. (d) Simulation-based toolpath evaluation comparing the thermal uniformity of the conventional toolpath and the RL-optimised toolpaths.

distinct geometries (Figure 9(d)). The reward metric serves as a quantitative indicator of thermal uniformity, where a higher value represents a more uniform thermal distribution. For the rectangular geometry, the conventional toolpath recorded a score of -80.49 , while the RL-optimised toolpath achieved higher rewards of -63.72 (RL 1), -64.05 (RL 2), and -64.40 (RL 3). A similar trend was observed for the U-shaped geometry. The conventional toolpath recorded a score of -97.39 , while the RL-optimised toolpath achieved higher rewards of -72.18 (RL 1), -72.25 (RL 2), and -74.49 (RL 3). Similarly, for the pyramidal geometry, the conventional toolpath recorded a score of -628.76 , while the RL-optimised toolpath achieved higher rewards of -596.99 (RL 1), -604.17 (RL 2), and -605.65 (RL 3). In conclusion, these results confirm that the RL-driven strategy consistently outperforms the conventional approach, ensuring superior thermal uniformity regardless of geometric complexity.

The thermal deviation was evaluated across specific cross-sectional planes for each geometry to visualise the heat accumulation and dissipation. Thermal deviation was defined as the difference between the spatially averaged temperature and the local temperature at each point. As shown in Figure 10(a), the rectangular geometry fabricated with the conventional toolpath exhibits a significant thermal concentration, particularly in the corner region. A steep temperature gradient is observed, with a temperature deviation of 7.27 . In contrast, the RL-optimised toolpath demonstrates a much more uniform thermal distribution, achieving a temperature deviation of 3.97 . This corresponds to a 45.39% reduction in temperature deviation compared to the conventional toolpath.

In the case of the U-shaped structure (Figure 10(b)), the geometry inherently induces localised heat accumulation, resulting in a thermal imbalance where heat is concentrated asymmetrically. Insufficient heat dissipation

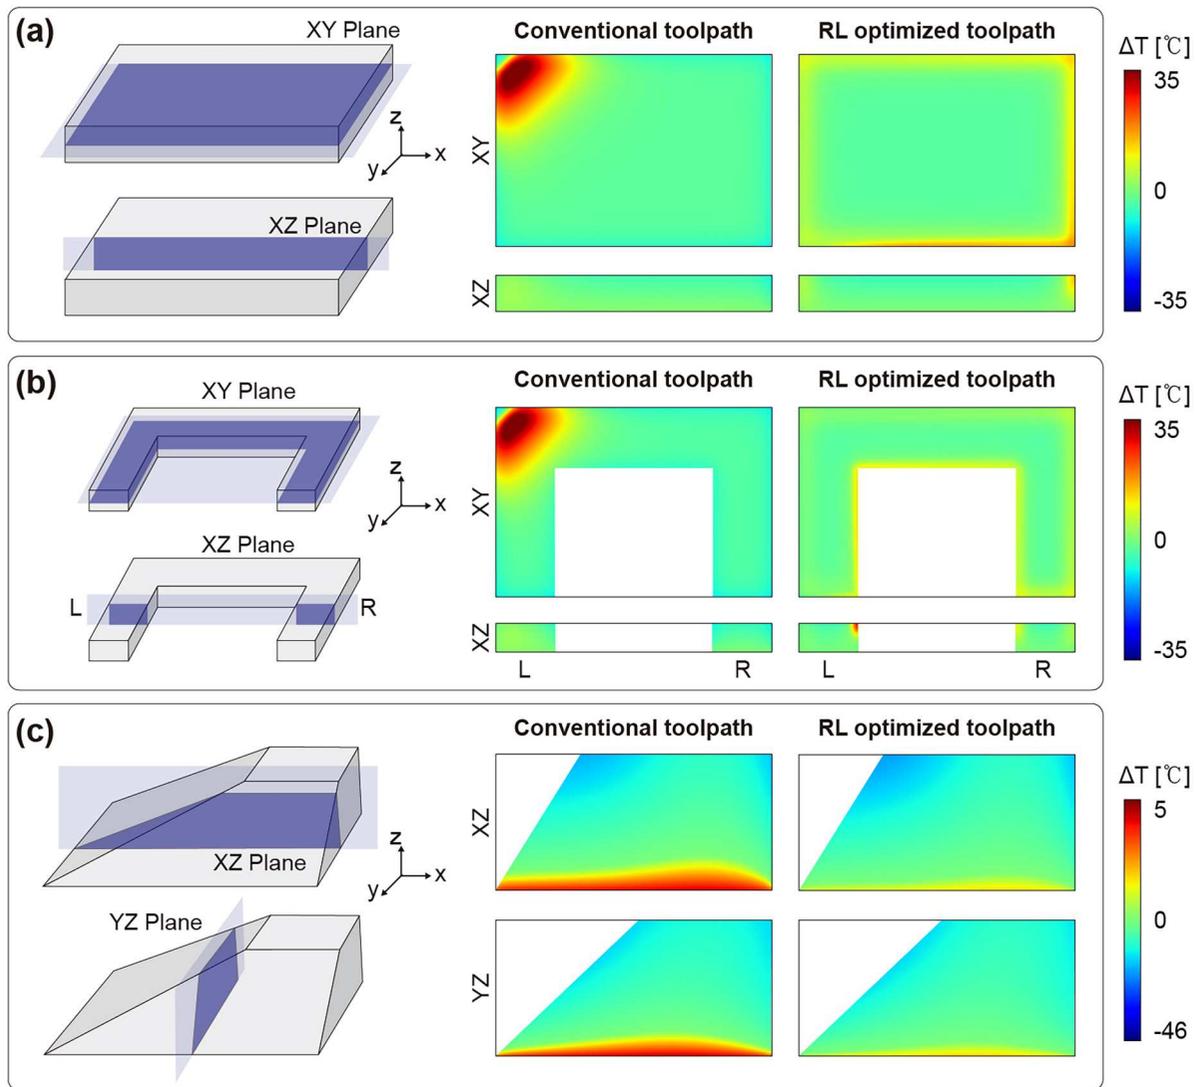


Figure 10. Comparative analysis of internal thermal deviation between conventional and RL-optimised toolpaths. The thermal deviation is evaluated on specific cross-sectional planes defined for each geometry: (a) the XY and XZ planes for the rectangular shape, (b) the XY and XZ planes for the U-shaped structure, and (c) the XZ and YZ planes for the pyramidal geometry.

caused by the conventional toolpath resulted in a distinct temperature deviation of 8.79. In contrast, the optimised toolpath effectively regulated the thermal history, maintaining a balanced thermal distribution between the left and right sections and reducing the temperature deviation to 3.03. Compared with the conventional toolpath, the optimised toolpath reduced the temperature deviation by 65.53%. This demonstrates the RL model's capability to handle complex geometric features where thermal imbalance often occurs.

The pyramidal geometry (Figure 10(c)) is characterised by a continuously decreasing cross-sectional area along the build height. The conventional toolpath caused excessive heat accumulation in a localised area, reaching a temperature deviation of 4.87. In contrast, the RL-optimised toolpath reduced the temperature

deviation to 4.06, which corresponds to a 16.63% reduction compared to the conventional toolpath.

To verify the effectiveness of the proposed framework, experiments were conducted on three distinct geometries: rectangular, U-shaped, and pyramidal structures. The performance of the RL-optimised toolpaths was evaluated against a conventional toolpath. Warpage was quantified by measuring the vertical displacement at two critical points, denoted as (i) and (ii), for each specimen. The magnitude of warpage was assessed by measuring the maximum lifting height at the corners of the geometries. These quantitative measurements were performed using ImageJ to ensure precise evaluation of the vertical deflection. While Figure S6 presents a comprehensive comparison of the commercial toolpath against all three RL toolpaths,

Figure 11 highlights the direct comparison between the conventional toolpath and the specific RL-optimised toolpath (RL 1) that achieved the minimum warpage.

For the rectangular geometry presented in Figure 11(a), the specimen fabricated via the conventional toolpath exhibited substantial deformation totalling 2.04 mm due to accumulated thermal stress. In contrast, the RL-optimised toolpaths effectively mitigated this issue, with RL 1 and RL 2 achieving improvement rates of 85.78% and 84.31%, respectively. This result validates

the model's thermal management capability in simple geometries. Regarding the U-shaped geometry shown in Figure 11(b), the conventional toolpath resulted in a total displacement of 4.14 mm characterised by asymmetric lifting, a distortion minimised by the RL-optimised toolpaths as evidenced by the 83.82% improvement in RL 1 which confirms the agent's adaptability to discontinuous cross-sections. In the case of the pyramidal geometry (Figure 11(c)) where varying cross-sectional areas complicate heat dissipation, the

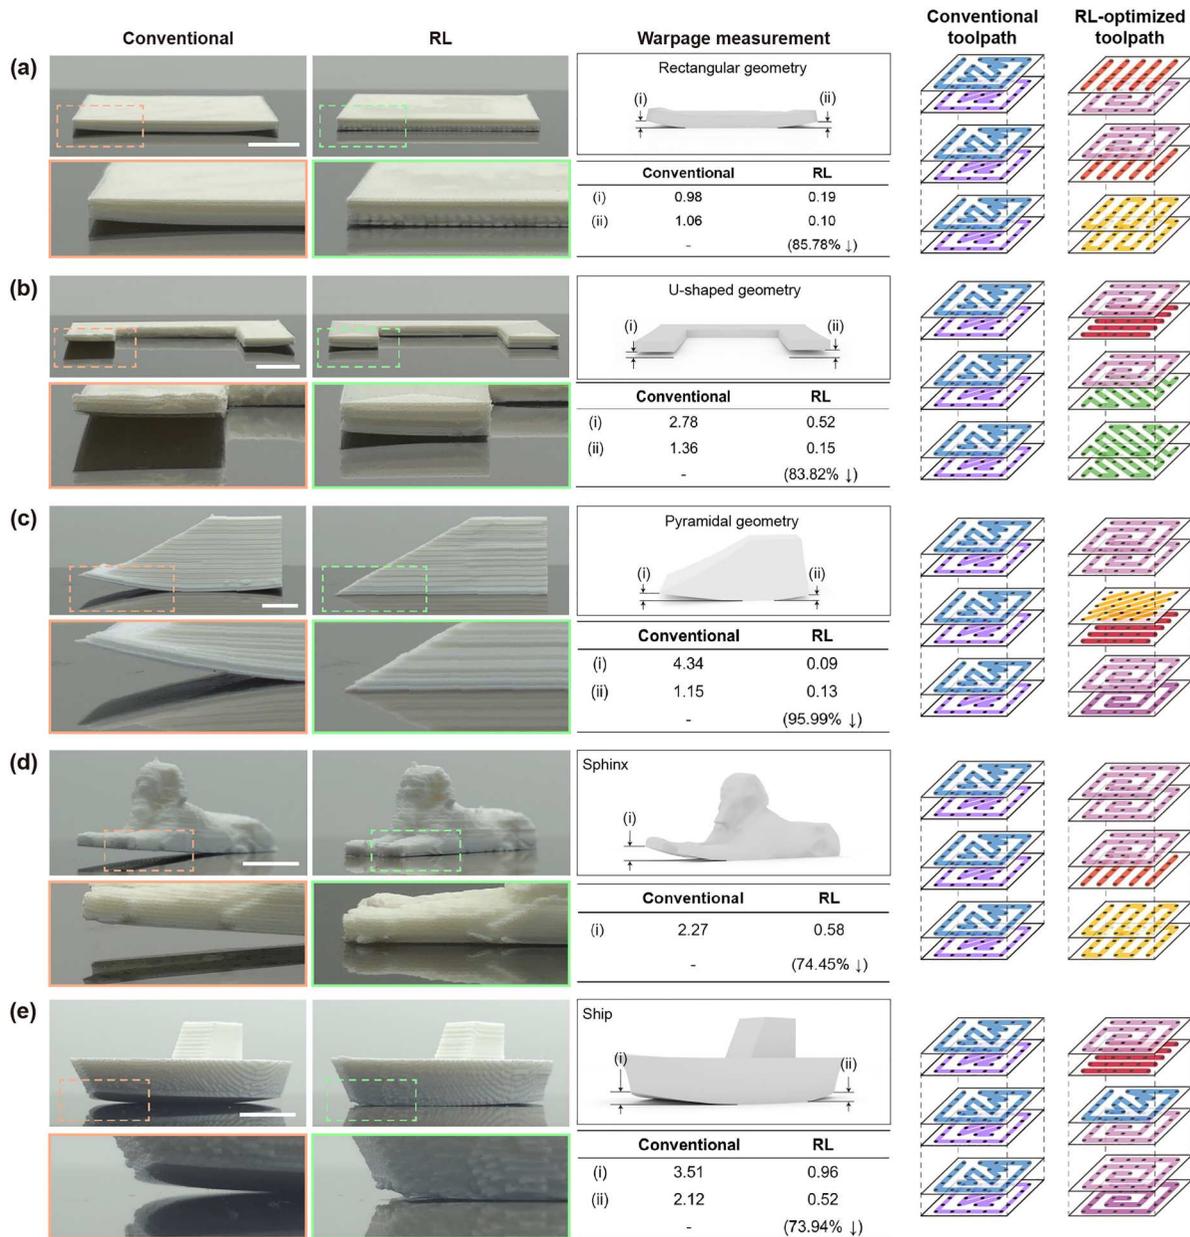


Figure 11. Experimental validation of warpage reduction using the RL-optimized toolpath. Comparison of printed samples for (a) rectangular, (b) U-shaped, and (c) pyramidal geometries. The performance of the conventional toolpath is compared against the RL-optimized toolpath. Schematic of distortion measurement locations (i, ii) and tables listing displacement values in millimetres and reduction rates of RL methods compared to the conventional toolpath. Evaluation of model generalizability on complex arbitrary geometries: (d) Sphinx and (e) ship models. (Scale bar: 1 mm).

conventional toolpath yielded a total warpage of 5.49 mm, while the RL-optimised toolpaths preserved geometric fidelity by reducing displacements to 0.22 mm (RL 1), 0.24 mm (RL 2), and 0.94 mm (RL 3), corresponding to improvements of 95.99%, 95.63%, and 82.88%, respectively. Finally, to verify the generalisation capability of the proposed method, the validation was extended to complex arbitrary geometries: the Sphinx and ship models (Figures 11(d and e)). For the Sphinx geometry, the conventional toolpath resulted in a maximum displacement of 2.27 mm. In contrast, the RL-optimised toolpath significantly mitigated this deformation to 0.58 mm, achieving an improvement rate of 74.45%. A similar improvement was observed for the ship geometry, where the warpage decreased from 5.68 mm to 1.48 mm, representing a 73.94% reduction. These results demonstrate that the effectiveness of the RL-driven strategy is not limited to simple primitives but extends robustly to diverse freeform shapes.

To evaluate the operational efficacy of the proposed method, the total fabrication times for RL-optimised and conventional toolpaths were compared. As illustrated in Figure S7, the RL-optimised toolpaths consistently demonstrated enhanced efficiency across all tested geometries. The total printing duration was reduced by a maximum of 25%. These results indicate that the RL agent promotes thermal uniformity without compromising productivity. This is achieved through the generation of efficient toolpaths that fulfil both thermal and temporal requirements. The reliability of the optimisation is further supported by the observation that localised peak errors at geometric boundaries remain distinct from the primary regions where warpage predominantly occurs. This distribution ensures that the surrogate model provides accurate thermal data for the optimisation of global uniformity (Figure S8). While the current framework is optimised for pre-fabrication planning to achieve global thermal consistency, the consistent per-layer inference speed of 0.3 s demonstrates significant potential for real-time applications. The integration of environmental sensors and the further refinement of the optimisation loop may allow this approach to evolve into a near-real-time adaptive system. Such a system would be capable of responding to unpredictable thermal variations during the FFF process.

4. Conclusion

This study presented a novel framework that integrates a 3D U-Net surrogate model with RL to address severe thermal gradients and the resulting bottom-layer warpage in complex geometries. The 3D U-Net

architecture replaced computationally intensive FDM simulations and reduced the inference time by over 99.8%. This acceleration compressed the thermal field prediction time from hours to seconds and facilitated rapid iterative optimisation. Utilising these efficient predictions, the RL agent generated optimal toolpaths designed to minimise thermal distortion. Experimental validation demonstrated the robustness of the framework across rectangular, U-shaped, and pyramidal geometries. In particular, the proposed approach achieved a substantial warpage reduction of 95.99% on the pyramidal geometry compared to the conventional toolpath. Consequently, this work establishes a scalable approach for intelligent manufacturing and overcomes the computational barriers of thermal management in complex geometries.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was supported by the Technology development Program (No. RS-2022-TI022487(S3248116)) funded by the Ministry of SMEs and Startups (MSS, Korea), the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2025-25440533, RS-2025-2453204, RS-2024-00416891, RS-2023-00211636), Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. RS-2020-II201336, Artificial Intelligence Graduate School Program (UNIST)), and the Human Resources Development of the Korea Institute of Energy Technology Evaluation and Planning (KETEP) grant funded by the Korea government Ministry of Knowledge Economy (No. RS-2023-00240918).

ORCID

Taekyeong Kim  <http://orcid.org/0000-0002-4016-4742>

Dohyeon Kim  <http://orcid.org/0009-0006-8527-2444>

Soon-Sung Kwon  <http://orcid.org/0009-0003-9152-7309>

Im Doo Jung  <http://orcid.org/0000-0003-0883-1848>

Data availability statement

The data that support the findings of this study are available from the corresponding author upon reasonable request.

References

- [1] Zhang D, Sun S, Qiu D, et al. Metal alloys for fusion-based additive manufacturing. *Adv Eng Mater.* 2018;20:1700952. doi:10.1002/adem.201700952
- [2] Park S, Kim T, Kim KM, et al. Quick dimensional inspection for continuous welding and assembly using machine

- learning-powered smart jig. *J Manuf Syst* 2025;82:478–496. doi:10.1016/j.jmsy.2025.07.001
- [3] Kim H, Lee J, Kim T, et al. Advanced thermal fluid leakage detection system with machine learning algorithm for pipe-in-pipe structure. *Case Stud Thermal Eng.* 2023;42:102747. doi:10.1016/j.csite.2023.102747
- [4] Seo J, Kim T, You K, et al. High quality large-scale nickel-rich layered oxides precursor co-precipitation via domain adaptation-based machine learning. *InfoMat.* 2025; e70031. doi:10.1002/inf2.70031
- [5] Li C, Liu Z, Fang X, et al. Residual stress in metal additive manufacturing. *Procedia Cirp.* 2018;71:348–353. doi:10.1016/j.procir.2018.05.039
- [6] Childs T, Berzins M, Ryder G, et al. Selective laser sintering of an amorphous polymer—simulations and experiments. *Proc Inst Mech Eng Part B J Eng Manuf.* 1999;213:333–349. doi:10.1243/0954405991516822
- [7] Mani M, Feng S, Brandon L, et al. Measurement science needs for real-time control of additive manufacturing powder-bed fusion processes. In: *Additive manufacturing handbook*. CRC Press; 2017. p. 629–652.
- [8] Daly M, Tarfaoui M, Chihi M, et al. FDM technology and the effect of printing parameters on the tensile strength of ABS parts. *Int J Adv Manuf Technol.* 2023;126:5307–5323.
- [9] Armillotta A, Bellotti M, Cavallaro M. Warpage of FDM parts: experimental tests and analytic model. *Robot Comput Integr Manuf.* 2018;50:140–152. doi:10.1016/j.rcim.2017.09.007
- [10] Kuo C-C, Wu Y-R, Li M-H, et al. Minimizing warpage of ABS prototypes built with low-cost fused deposition modeling machine using developed closed-chamber and optimal process parameters. *Int J Adv Manuf Technol* 2019;101:593–602. doi:10.1007/s00170-018-2969-7
- [11] Yan W, Yue Z, Feng J. Study on the role of deposition path in electron beam freeform fabrication process. *Rapid Prototyp J.* 2017;23:1057–1068. doi:10.1108/RPJ-03-2016-0043
- [12] Zhang Y, Jing H, Xu L, et al. Effects of different scanning patterns on nickel alloy-directed energy deposition based on thermal analysis. *Virtual Phys Prototyp.* 2021;16:S98–S115. doi:10.1080/17452759.2021.1896173
- [13] Ali H, Ghadbeigi H, Mumtaz K. Effect of scanning strategies on residual stress and mechanical properties of selective laser melted Ti6Al4V. *Mater Sci Eng A.* 2018;712:175–187. doi:10.1016/j.msea.2017.11.103
- [14] Ding D, Huang R, Liu T, et al. Effects of path patterns on residual stresses and deformations of directed energy deposition-arc built blocks. *Virtual Phys Prototyp.* 2024;19:e2296128. doi:10.1080/17452759.2023.2296128
- [15] Kim SI, Hart AJ. A spiral laser scanning routine for powder bed fusion inspired by natural predator-prey behaviour. *Virtual Phys Prototyp.* 2022;17:239–255. doi:10.1080/17452759.2022.2031232
- [16] Sun L, Ren X, He J, et al. Numerical investigation of a novel pattern for reducing residual stress in metal additive manufacturing. *J Mater Sci Technol.* 2021;67:11–22. doi:10.1016/j.jmst.2020.05.080
- [17] Kim T, Kim JG, Park S, et al. Virtual surface morphology generation of Ti-6Al-4V directed energy deposition via conditional generative adversarial network. *Virtual Phys Prototyp.* 2023;18:e2124921. doi:10.1080/17452759.2022.2124921
- [18] Baek AMC, Park E, Seong M, et al. Multi-objective robust parameter optimization using the extended and weighted k-means (EWK-means) clustering in laser powder bed fusion (LPBF). *Expert Syst Appl.* 2024;236:121349. doi:10.1016/j.eswa.2023.121349
- [19] Kim D, Zohdi T. Tool path optimization of selective laser sintering processes using deep learning. *Comput Mech* 2022;69:383–401. doi:10.1007/s00466-021-02079-1
- [20] Ren K, Chew Y, Liu N, et al. Integrated numerical modeling and deep learning for multi-layer cube deposition planning in laser aided additive manufacturing. *Virtual Phys Prototyp.* 2021;16:318–332. doi:10.1080/17452759.2021.1922714
- [21] Nguyen L, Buhl J, Bambach M. Continuous Eulerian tool path strategies for wire-arc additive manufacturing of rib-web structures with machine-learning-based adaptive void filling. *Addit Manuf.* 2020;35:101265. doi:10.1016/j.addma.2020.101265
- [22] Qin M, Qu S, Ding J, et al. Adaptive toolpath generation for distortion reduction in laser powder bed fusion process. *Addit Manuf.* 2023;64:103432.
- [23] Zhou Z, Shen H, Lin J, et al. Continuous tool-path planning for optimizing thermo-mechanical properties in wire-arc additive manufacturing: an evolutionary method. *J Manuf Process* 2022;83:354–373. doi:10.1016/j.jmapro.2022.09.009
- [24] Dharmadhikari S, Menon N, Basak A. A reinforcement learning approach for process parameter optimization in additive manufacturing. *Addit Manuf.* 2023;71:103556.
- [25] Petrik J, Bambach M. Reinforcement learning and optimization based path planning for thin-walled structures in wire arc additive manufacturing. *J Manuf Process* 2023;93:75–89. doi:10.1016/j.jmapro.2023.03.013
- [26] Qin M, Ding J, Qu S, et al. Deep reinforcement learning based toolpath generation for thermal uniformity in laser powder bed fusion process. *Addit Manuf.* 2024;79:103937.
- [27] Breiman L. Random forests. *Mach Learn.* 2001;45:5–32. doi:10.1023/A:1010933404324
- [28] Çiçek Ö, Abdulkadir A, Lienkamp SS, et al. 3D U-Net: learning dense volumetric segmentation from sparse annotation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer; 2016. p. 424–432.
- [29] Senthilnathan A, Nath P, Mahadevan S, et al. Surrogate modeling of microstructure prediction in additive manufacturing. *Comput Mater Sci.* 2025;247:113536.
- [30] Hemmasian A, Ogoke F, Akbari P, et al. Surrogate modeling of melt pool temperature field using deep learning. *Addit Manuf Lett.* 2023;5:100123. doi:10.1016/j.addlet.2023.100123
- [31] Dong G, Wong JC, Lestandi L, et al. A part-scale, feature-based surrogate model for residual stresses in the laser powder bed fusion process. *J Mater Process Technol.* 2022;304:117541. doi:10.1016/j.jmatprotec.2022.117541
- [32] Mosleh N, Esfandeh M, Dariushi S. Simulation of temperature profile in fused filament fabrication 3D printing method. *Rapid Prototyp J.* 2024;30:134–144. doi:10.1108/RPJ-02-2023-0067
- [33] Ramos N, Mittermeier C, Kiendl J. Experimental and numerical investigations on heat transfer in fused

- filament fabrication 3D-printed specimens. *Int J Adv Manuf Technol* 2022;118:1367–1381. doi:10.1007/s00170-021-07760-6
- [34] D’Amico A, Peterson AM. An adaptable FEA simulation of material extrusion additive manufacturing heat transfer in 3D. *Addit Manuf.* 2018;21:422–430.
- [35] Serdeczny MP, Comminal R, Mollah MT, et al. Numerical modeling of the polymer flow through the hot-end in filament-based material extrusion additive manufacturing. *Addit Manuf.* 2020;36:101454.
- [36] Mittal Y, Agarwal V, Yadav D, et al. A mechanistic model for overhang limits in additive manufacturing. *Prog Addit Manuf.* 2025: 1–20. doi:10.1007/s40964-025-01154-w
- [37] Lee K, Zung J, Li P, et al. Superhuman accuracy on the SNEMI3D connectomics challenge. *arXiv preprint arXiv:1706.00120*; 2017.
- [38] Roy AG, Navab N, Wachinger C. Concurrent spatial and channel ‘squeeze & excitation’ in fully convolutional networks. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer; 2018. p. 421–429.
- [39] Sutton RS. *Reinforcement learning: an introduction, A Bradford book*. Cambridge, MA: The MIT Press; 2018.
- [40] van Otterlo M, Wiering M. Reinforcement learning and Markov decision processes. In: Wiering M, van Otterlo M, editors. *Reinforcement Learning: State-of-the-Art*. Berlin, Heidelberg: Springer Berlin Heidelberg; 2012. p. 3–42.
- [41] Schulman J, Wolski F, Dhariwal P, et al. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*; 2017.
- [42] Wu Y, He K. Group normalization. In: *Proceedings of the European conference on computer vision (ECCV)*; 2018. p. 3–19.
- [43] Sun H, Han L, Yang R, et al. Exploit reward shifting in value-based deep-rl: optimistic curiosity-based exploration and conservative exploitation via linear reward shaping. *Adv Neural Inf Process Syst.* 2022;35:37719–37734.
- [44] Sun X, Mazur M, Cheng C-T. A review of void reduction strategies in material extrusion-based additive manufacturing. *Addit Manuf.* 2023;67:103463.
- [45] Zhao D, Guo W. Shape and performance controlled advanced design for additive manufacturing: a review of slicing and path planning. *J Manuf Sci Eng.* 2020;010801. doi:10.1115/1.4045055