

Article | Received 7 May 2026; Revised 6 July 2026; Accepted 12 July 2026; Published 31 July 2026  
<https://doi.org/10.55092/am20260009>

# Symmetry-imbued Hamiltonian neural operator architecture for stress field prediction of porous metamaterials



Xinyu Lu<sup>1,2</sup>, Kang Wang<sup>1,\*</sup>, Guodong Yi<sup>1,2</sup>, Shuyou Zhang<sup>1,2</sup> and Jianrong Tan<sup>1,2</sup>

<sup>1</sup> State Key Laboratory of Fluid Power and Mechatronic Systems, Zhejiang University, Hangzhou, China

<sup>2</sup> Zhejiang Advanced CNC Machine Tool Technology Innovation Center, Taizhou, China

\* Correspondence author; E-mail: kang.wang@zju.edu.cn.

## Highlights:

- A symmetry-imbued Hamiltonian neural operator architecture is proposed for stress field prediction of porous metamaterials.
- A rotation-steerable lattice projector encodes continuous-angle geometric symmetry to interpret images under arbitrary in-plane rotations.
- A Hamiltonian-Fueled Propagator introduces a symplectic symmetry that encourages global mechanical coherence without relying on pointwise PDE residuals across discontinuous interfaces.

**Abstract:** Predicting stress fields of porous metamaterials under arbitrary rotations is essential for reliable online monitoring in additive manufacturing. However, existing neural operators are limited by two major issues: geometric orientation bias where identical lattices are misinterpreted under rotation, and non-physical stress leakage caused by the failure of Partial Differential Equation (PDE) constraints at discontinuous void-solid interfaces in porous metamaterials. To address these challenges, we propose the Symmetry-Imbued Hamiltonian Neural Operator (SIHNO) that embeds geometric symmetry and an energy-structured inductive bias into a unified neural operator architecture. Specifically, SIHNO introduces a Rotation-Steerable Lattice Projector (RSLP) that lifts lattice images into a continuous-angle geometric symmetry representation. This geometric representation is further coupled with a Hamiltonian-Fueled Propagator (HFP) that replaces local PDE constraints with Hamiltonian-inspired energy propagation. Finally, a slice-aware convolutional decoder reconstructs stress fields based on RSLP and HFP. Comprehensive experiments demonstrate that SIHNO outperforms existing neural operators. The proposed framework provides a robust architecture for stress-field prediction in additive manufacturing.

**Keywords:** additive manufacturing; metamaterial; stress field prediction; Hamiltonian dynamics; neural operator



Copyright©2026 by the authors. Published by ELSP. This work is licensed under a Creative Commons Attribution 4.0 International License, which permits unrestricted use, distribution, and reproduction in any medium provided the original work is properly cited.

## 1. Introduction

Additive manufacturing (AM) enables the fabrication of topology-optimized porous metamaterials, unlocking effective combinations of strength, stiffness, and energy absorption for aerospace, biomedical, and protective applications [1–3]. For such architected metamaterials, the stress field is a critical indicator of load-bearing paths, local stress concentration, and potential failure initiation, and therefore plays a central role in structural reliability assessment and process monitoring. However, stress analysis of porous metamaterials remains computationally expensive because the structure contains complex void-solid interfaces and layer-wise geometric variations. This motivates the development of fast surrogate models that can estimate stress fields directly from as-built or reconstructed topology slices [4,5].

Despite the growing use of operator learning for porous metamaterials, two major issues remain insufficiently addressed [6–8]: (i) Geometric Orientation Bias. In AM process monitoring, the same lattice section of the metamaterial may appear more than once under different in-plane coordinate frames since a metamaterial can be a complex of multiple identical layers that are similar with camera or Computed Tomography (CT) rotation. However, existing feature extraction methods are usually defined on fixed Cartesian grids and tend to recognize the image through rotation-dependent pixel patterns. As a result, conventional methods may interpret the same image under different rotations as distinguished microstructure, leading to degraded stress predictions [9,10]. (ii) Boundary Smoothing and Stress Leakage. Physics-Informed Machine Learning and neural operators often enforce governing laws through pointwise Partial Differential Equation (PDE) residuals [11,12]. Such constraints can be effective in smooth domains with no discontinuities, but become fragile in porous metamaterials where solid and void regions coexist within the same section plane [13,14]. At void-solid interfaces, the elastic modulus changes discontinuously, a pointwise residual formulation in this scenario tends to impose smooth field behavior across the stiffness jump [15]. This mismatch may produce oversmoothing at boundaries and non-physical stress leakage into void regions, blurring the separation between solids and background voids [16].

To jointly resolve geometric orientation bias and boundary stress leakage, we propose Symmetry-Imbued Hamiltonian Neural Operator (SIHNO) by establishing a dual-symmetry mechanistic pathway within the latent space. First, we implement a Rotation-Steerable Lattice Projector (RSLP) that encodes latent space geometric symmetry into rotation-equivariant representations to enable continuous angle prediction. Rather than relying on fixed Cartesian convolutional responses, RSLP extracts lattice images that transform consistently under arbitrary rotations, providing a canonical geometric substrate that preserves geometric symmetry regardless of print orientation. Building upon this distortion-free foundation, we deploy the Hamiltonian-Fueled Propagator (HFP) that formulates global token interactions through a Hamiltonian-inspired latent energy functional. By predicting feature correction residuals from the gradient of the learned energy structure, HFP encourages mechanically coherent long-range coupling without directly imposing pointwise PDE residuals. Our principal contributions are as follows:

- (1) We propose a Rotation-Steerable Lattice Projector, which encodes rotational symmetry of porous metamaterial lattices to project its structures into a rotation-equivariant space, endowing the network with awareness of spatial transformations under arbitrary rotations.
- (2) A Hamiltonian-Fueled Propagator is further introduced to formulate global feature interaction as

Hamiltonian-inspired energy coupling, circumventing stress smoothing and leakage at discontinuous void-solid interfaces in porous metamaterials.

- (3) Through comprehensive experiments, it is demonstrated that SIHNO achieves superior spatial equivariance consistency and accurate stress field prediction across diverse rotated metamaterial topologies, outperforming state-of-the-art neural operators.

## 2. Related works

### 2.1. Spatial representations in neural operators

Neural operators learn resolution invariant mappings between two function spaces, yet when it comes to metamaterial online monitoring, their performance is heavily dictated by how they interpret spatial coordinates and geometric structures [17]. The Fourier Neural Operator (FNO) [18,19] and Neural Spectral Methods (NSM) [20] leverage spectral convolutions to capture spatial correlations in a long range. However, they are fundamentally coordinate-dependent. Their integral kernels are tied to the absolute Cartesian axes of the underlying grid, rendering them incapable of recognizing the same physical object when spatial transformation occurs [21,22]. Similarly, Deep Operator Network (DeepONet) [23] utilizes a branch-trunk architecture to approximate operator mappings, but it treats spatial coordinates as raw, unconstrained inputs. Without an explicit geometric inductive bias, DeepONet also lacks the mechanism to ensure that the learned latent representations are invariant or equivariant to the symmetry group of the material lattice.

More recently, attention-based operators have been introduced to handle complex geometric representations [24]. Transolver [25] proposed physics-aware attention through slicing of tokens, enabling fast and accurate PDE solving on general geometries. Its successor Transolver++ scaled this approach to million-scale meshes while maintaining accuracy [26]. Despite these advances, existing neural operators predominantly learn continuous field mappings without an explicitly encoded geometric symmetry, leaving them vulnerable to geometric orientation in porous metamaterial domains.

Rotation-equivariant and steerable operators [27–30] further improve robustness to coordinate transformations by constructing equivariance-aware feature spaces. Such equivariant operators, however, mainly canonicalize geometry and carry no explicit mechanical consistency prior across discontinuous void-solid interfaces, whereas most metamaterial surrogate models predict homogenized or effective quantities rather than resolving boundary accurate sectional stress fields under arbitrary rotation.

### 2.2. Physics embedding in neural architectures

Physics-Informed Neural Networks (PINNs) pioneered the embedding of physical laws into neural networks by penalizing PDE residuals as soft constraints in the loss function [31–33]. This framework has been extended to metamaterial applications [14], yet the fundamental reliance on pointwise PDE residuals renders PINNs ill-suited for porous metamaterials with discontinuous void-solid interfaces, where boundary smoothing and stress leakage inevitably arise since they require computing spatial derivatives across the entire domain. Physics-Informed Neural Operator (PINO) [34] combined the operator learning capability with PDE residual penalties, improving generalization across PDE families.

However, PINO inherits the same pointwise PDE residual limitations at metamaterial discontinuities.

An alternative physics embedding strategy is to enforce structure-preserving laws directly within the network architecture. Hamiltonian Neural Networks (HNN) [35,36] implemented this idea by parameterizing the Hamiltonian function by construction, outperforming unconstrained networks on dynamical systems [37]. However, such frameworks would remain agnostic to spatial transformation since they focus on energy conservation while neglecting the discrete geometric symmetries of the lattice structure, rendering them incapable of maintaining physical consistency under arbitrary rotations.

In summary, neural operators with spatial representations or physical embeddings address complementary but disjoint symmetries. The former interprets geometric orientations, yet lack a mechanical consistent embedding, while the latter is imbued with physics that are not equivariant to image rotation. This gap motivates an architecture that jointly enforces geometric equivariance and an energy-preserving mechanical bias by applying dual symmetries within a single neural operator.

### 3. Methods

#### 3.1. Problem formulation and overall architecture

We formulate the defined metamaterial stress prediction process as learning an operator  $G_\theta : (x, \ell) \mapsto \sigma$  with rotation equivariance and symplectic symmetry, where  $x \in \mathbb{R}^{1 \times H \times W}$  denotes a grayscale or binary topology slice,  $\ell$  is the normalized build-layer index,  $\sigma \in \mathbb{R}^{1 \times H \times W}$  is the corresponding Von Mises stress field in a specified layer.

The central viewpoint of this work is that both pathologies mentioned in the introduction section will be largely mitigated once two complementary symmetries are embedded into the architecture: the geometric symmetry of the square lattice in porous metamaterials, and a symplectic, energy-structured inductive bias inspired by Hamilton's flow. By embedding these symmetries into our neural operator, SIHNO is capable of sustaining high prediction accuracy under arbitrary rotations, allowing it to be a more robust paradigm for stress field prediction in industrial AM processes.


As shown in Figure 1, SIHNO realizes the two-symmetry idea through three sequential blocks: (i) the RSLP module lifts the scalar topology input into a continuous-angle rotation-steerable latent field  $\mathbf{F} \in \mathbb{R}^{d \times H \times W}$ ; (ii) the HFP module provides a Hamiltonian channel by parameterizing a global Hamiltonian  $H_\theta$ , obtaining symplectic corrections from the input map; (iii) the Slice-Aware Convolutional Decoder fuses rotation-steerable features with symplectic corrections and reconstructs the predicted stress field.

In our architecture, the input  $\mathbf{x}$  is first projected into a high-dimensional equivariant feature space via the RSLP module. By parameterizing convolutional kernels with radial profiles and angular harmonic bases, RSLP extracts orientation-aware responses and aggregates them into a rotation-steerable latent representation. This ensures that any orientation of the metamaterial unit cell is represented as a predictable channel permutation rather than a numerical distortion, thereby reducing orientation bias at the source. Simultaneously, a pooled version of the topology is fed into the HFP module to capture the underlying mechanical evolution. Here HFP operates as a latent space dynamical operator that maps it to three relevant branches: the Coordinate ( $Q$ ) branch, the Momentum ( $P$ ) branch, and the Energy ( $H$ ) branch. While both the  $Q$  and  $P$  branches are used to resolve spatial connectivity and force-transfer history, the



Energy branch parameterizes a scalar function  $H_\theta$  representing the system's learned Hamiltonian.

Through an Autograd-based interconnection, the HFP extracts symplectic corrections  $(\frac{\partial H}{\partial \mathbf{p}}, -\frac{\partial H}{\partial \mathbf{q}})$  that govern the latent feature propagation. This stage encourages energy-structured global feature coupling and improves mechanical coherence among distant load-bearing regions. Finally, the equivariant features from the RSLP are fused with the symplectic corrections from the HFP, conditioned on the Layer Index  $\ell$  to account for the depth-dependent stress field in the AM process. By synergizing geometric equivariance with Hamiltonian mechanics, SIHNO is capable of more effectively predict the final elastostatic stress field  $\sigma$  with high physical fidelity, even under complex manufacturing orientations.



**Figure 1.** Overall architecture of SIHNO. A grayscale topology slice is first lifted into a continuous-angle rotation-steerable latent field by RSLP; the latent field is then pooled to a coarse phase-space token grid on which the HFP module parameterizes a learned Hamiltonian  $H_\theta$  and extracts symplectic corrections  $(\Delta \mathbf{q}, \Delta \mathbf{p})$  via autograd. The full-resolution equivariant features and the upsampled symplectic corrections are fused, conditioned on the build-layer index, and decoded by a stack of convolution-attention blocks into the final stress field.

### 3.2. Geometric symmetry: rotation-steerable lattice projector

Porous metamaterial slices often contain lattice motifs that may appear at arbitrary in-plane orientations, while their underlying mechanical role should remain geometrically consistent under rotation. Standard convolutional filters are defined on fixed Cartesian grids and therefore learn orientation-dependent responses, which may cause identical lattice patterns to be mapped to inconsistent latent representations [38–40]. To address this issue, we introduce the RSLP, which embeds continuous in-plane rotational symmetry into the geometric feature extraction stage.

Let  $x \in \mathbb{R}^{1 \times H \times W}$  denote the input topology slice with  $H$  and  $W$  denoting the spatial height and width respectively. As shown in Figure 2a, the continuous in-plane rotation group is denoted as [41]:

$$SO(2) = \{R_\alpha \mid \alpha \in [0, 2\pi)\}, \quad (1)$$


where  $R_\alpha$  is a planar rotation by angle  $\alpha$ . For a feature field  $f$ , the rotation action is defined as:

$$[\mathcal{L}_\alpha f](u) = \rho(\alpha) f(R_\alpha^{-1} u), \quad u \in \mathbb{R}^2, \quad (2)$$

where  $\rho(\alpha)$  is the channel representation. A feature extractor  $\phi$  is rotation-equivariant if

$$\phi(\mathcal{L}_\alpha^{\text{in}} x) = \mathcal{L}_\alpha^{\text{out}} \phi(x), \quad \forall \alpha \in [0, 2\pi). \quad (3)$$

This condition ensures that rotating the input topology induces a predictable transformation of the latent feature field, rather than producing an unrelated feature response.



**Figure 2.** Illustration of the RSLP. **(a)** A porous lattice motif may appear at arbitrary in-plane orientations, motivating continuous-angle rotational modeling based on  $SO(2)$ ; **(b)** RSLP parameterizes convolution kernels using learnable radial bases and angular harmonic bases.

As illustrated in Figure 2b, RSLP realizes this property by parameterizing convolutional kernels with steerable circular harmonics. For a convolutional kernel  $\kappa$ , rotation equivariance requires

$$\kappa(R_\alpha u) = \rho_{\text{out}}(\alpha) \kappa(u) \rho_{\text{in}}(\alpha)^{-1}, \quad (4)$$

Instead of learning orientation kernels independently, RSLP learns  $\kappa$  in polar coordinates  $u = (r, \varphi)$  as

$$\kappa(r, \varphi) = \sum_{m \in \mathcal{M}} \sum_{\ell=1}^L b_\ell(r) [A_{m\ell} \cos(m\varphi) + B_{m\ell} \sin(m\varphi)], \quad (5)$$

where  $b_\ell(r)$  denotes a learnable radial basis,  $m$  is the angular frequency,  $M$  is the set of selected harmonic orders,  $L$  is the number of radial basis functions and  $A_{m\ell}, B_{m\ell}$  are learnable channel-mixing coefficients. Under a rotation by  $\alpha$ , the angular component satisfies

$$\cos m(\varphi - \alpha), \quad \sin m(\varphi - \alpha), \quad (6)$$

which can be expressed as a linear transformation of  $\cos(m\varphi)$  and  $\sin(m\varphi)$ . Therefore, the filter response can be continuously steered by transforming its harmonic coefficients, rather than relearning separate filters for different orientations.

Given the input slice  $\ell$ , RSLP first maps it into an orientation-aware feature field through a steerable lifting operation to yield the reconstructed feature map  $\tilde{F}$ :

$$\tilde{F} = \psi_{\text{st}}(x) \in \mathbb{R}^{C \times N_\alpha \times H \times W}, \quad (7)$$

where  $N_\alpha$  denotes the sampled orientation channels used for efficient implementation, and  $C$  is the number of steerable feature channels per orientation. The lifted feature field is then refined by stacked steerable convolutional blocks:

$$\tilde{F}^{(l+1)} = \eta \left( \text{Norm} \left( \mathcal{K}_{\text{st}}^{(l)} * \tilde{F}^{(l)} \right) \right), \quad l = 0, \dots, L_s - 1, \quad (8)$$

where  $\mathcal{K}_{\text{st}}^{(l)}$  is the steerable convolution kernel,  $\eta(\cdot)$  is the nonlinear activation, and  $L_s$  is the number of steerable blocks. Finally, the orientation-aware responses are folded into the latent channel dimension:

$$F_{\text{RSLP}} = \mathcal{A}_\alpha(\tilde{F}) \in \mathbb{R}^{d \times H \times W}, \quad (9)$$

where  $\mathcal{A}_\alpha(\cdot)$  denotes orientation-channel aggregation or reshaping.

The resulting feature field  $F_{\text{RSLP}}$  preserves the rotation structure of the input topology and provides an orientation-consistent representation for subsequent energy-path modeling and stress-field reconstruction. Since the topology slices are discretized on pixel grids, the continuous rotation equivariance is implemented through steerable kernel parameterization and sampled orientation channels, yielding stable arbitrary-angle rotation robustness in practical image-based settings.

### 3.3. Symplectic symmetry: Hamiltonian-fueled propagator

Hamiltonian mechanics provides an energy-structured formulation for describing global mechanical interactions [36,42,43]. In a canonical phase space, the state is represented by conjugate variables  $\mathbf{z} = [\mathbf{q}, \mathbf{p}]^\top$ , and its evolution follows

$$\dot{\mathbf{z}} = \mathbf{J} \nabla_{\mathbf{z}} H, \quad \mathbf{J} = \begin{pmatrix} \mathbf{0}_n & \mathbf{I}_n \\ -\mathbf{I}_n & \mathbf{0}_n \end{pmatrix}, \quad \mathbf{J}^\top = -\mathbf{J}. \quad (10)$$

Equivalently,  $\dot{\mathbf{q}} = \partial H / \partial \mathbf{p}$  and  $\dot{\mathbf{p}} = -\partial H / \partial \mathbf{q}$ . The skew-symmetric structure of  $\mathbf{J}$  yields an energy-structured form since  $\nabla_{\mathbf{z}} H^\top \mathbf{J} \nabla_{\mathbf{z}} H = 0$ .

In this work, the Hamiltonian formulation is introduced as a latent energy-structured constraint rather than an explicit solver of the continuum Hamiltonian system. Compared with pointwise PDE residuals that require spatial differentiation across discontinuous void-solid interfaces, HFP operates on latent degrees of freedom and is therefore more suitable for porous metamaterial domains.

Given the rotation-steerable feature field  $\mathbf{F}_{\text{RSLP}}$ , HFP first maps it into a coarse latent phase space:

$$\mathbf{Q} = \Psi_q(\mathbf{F}_{\text{RSLP}}), \quad \mathbf{P} = \Psi_p(\mathbf{F}_{\text{RSLP}}), \quad \mathbf{Z} = [\mathbf{Q}, \mathbf{P}] \in \mathbb{R}^{N \times 2d}, \quad (11)$$

where  $N$  is the number of latent tokens,  $d$  is the token dimension, and  $\Psi_q$  and  $\Psi_p$  are the learnable projection regarding  $\mathbf{q}$  and  $\mathbf{p}$  respectively. The coordinate tokens  $\mathbf{Q}$  encode lattice configuration, while the momentum tokens  $\mathbf{P}$  describe latent force-transfer tendency among load-bearing regions.

To model non-local mechanical coupling, HFP parameterizes a learnable latent Hamiltonian:

$$H_\theta(\mathbf{Z}) = \sum_{i=1}^N h_\theta(\mathbf{z}_i) + \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N a_{ij} g_\theta(\mathbf{z}_i, \mathbf{z}_j), \quad a_{ij} = \text{Softmax}_j \left( \frac{(\mathbf{W}_q \mathbf{z}_i)^\top (\mathbf{W}_k \mathbf{z}_j)}{\sqrt{d}} \right), \quad (12)$$

where  $\mathbf{z}_i = [\mathbf{q}_i, \mathbf{p}_i]$ ,  $h_\theta$  denotes token-wise latent energy,  $g_\theta$  models pairwise mechanical interaction, and

$a_{ij}$  is the attention-based interaction weight. This formulation allows HFP to construct energy-guided interaction paths among distant struts, junctions, and boundary regions.

The symplectic correction is then obtained by differentiating the learned Hamiltonian:

$$\Delta \mathbf{Z} = \tau \mathbf{J}_d \nabla_{\mathbf{Z}} H_{\theta}(\mathbf{Z}), \quad \mathbf{J}_d = \begin{pmatrix} \mathbf{0}_d & \mathbf{I}_d \\ -\mathbf{I}_d & \mathbf{0}_d \end{pmatrix}, \quad (13)$$

or equivalently,

$$\Delta \mathbf{Q} = \tau \frac{\partial H_{\theta}}{\partial \mathbf{P}}, \quad \Delta \mathbf{P} = -\tau \frac{\partial H_{\theta}}{\partial \mathbf{Q}}, \quad (14)$$

where  $\tau$  is the learnable step size.

The corrected latent state is updated as

$$\hat{\mathbf{Z}} = \mathbf{Z} + \Gamma \odot \Delta \mathbf{Z}, \quad \Gamma = \sigma(\mathbf{W}_g \mathbf{Z}), \quad (15)$$

where  $\Gamma$  adaptively controls the strength of the Hamiltonian correction, and  $\sigma(\cdot)$  is the sigmoid activation.

### 3.4. Slice-aware convolutional decoder

The RSLP and HFP modules provide symmetry-preserving geometric features and global Hamiltonian corrections, respectively. However, these latent representations cannot be directly converted into the final stress field by a plain convolutional decoder. On the one hand, the symplectic corrections extracted by HFP are generated on a coarse token grid and mainly encode global mechanical interactions. Directly upsampling them may lose high-frequency boundary details. On the other hand, conventional convolutional decoders aggregate features only from local neighborhoods, which tends to smooth the discontinuous void-solid interfaces and causes non-physical stress leakage. Therefore, we introduce a Slice-Aware Convolutional Decoder to reconstruct the stress field in a structure-aware manner.

Given the equivariant feature  $F_{\text{RSLP}}$  and the symplectic pair, the decoder input is defined as:

$$Y_0 = \text{Concat}(F_{\text{RSLP}}, U(\phi_s([\Delta \mathbf{Q}, \Delta \mathbf{P}])), \phi_{\ell}(\ell) \otimes \mathbf{1}_{H \times W}), \quad (16)$$

where  $U(\cdot)$  upsamples the symplectic correction to the original resolution,  $F_{\text{RSLP}}$  denotes the rotation-steerable lattice feature, and  $\phi_{\ell}(\ell) \otimes \mathbf{1}_{H \times W}$  injects the layer-wise condition into all spatial locations. A convolutional stem first extracts local boundary-sensitive features:

$$Y_c = \psi_{\text{conv}}(Y_0), \quad y_u = Y_c(:, u), \quad u = 1, \dots, HW, \quad (17)$$

where  $\psi_{\text{conv}}$  is the convolutional stem.

The slice-aware reconstruction module then maps spatial features into  $K$  learnable structural slices:

$$A_{u,k} = \text{Softmax}_k \left( \frac{w_k^{\top} y_u}{\tau} \right), \quad s_k = \frac{\sum_{u=1}^{HW} A_{u,k} y_u}{\sum_{u=1}^{HW} A_{u,k} + \epsilon}, \quad S = [s_1, \dots, s_K]^{\top}. \quad (18)$$

Here,  $A_{u,k}$  denotes the soft membership of spatial position  $u$  to the  $k$ -th structural slice. This formulation allows mechanically correlated regions, such as struts, junctions, and void-solid boundaries, to be grouped even when they are not adjacent in the image plane.

To exchange mechanical information among structural slices, we update each slice token by a similarity-weighted interaction with all other slices. Specifically, for the slice token  $s_i$ , its interaction weight with slice  $s_j$  is defined as

$$\alpha_{ij} = \frac{\exp((W_q s_i)^\top (W_k s_j) / \sqrt{d_a})}{\sum_{m=1}^K \exp((W_q s_i)^\top (W_k s_m) / \sqrt{d_a})}, \quad (19)$$

where  $W_q$  and  $W_k$  are learnable projection matrices. The refined slice token is then obtained by

$$\tilde{s}_i = s_i + W_o \sum_{j=1}^K \alpha_{ij} W_v s_j, \quad i = 1, \dots, K, \quad (20)$$

where  $W_v$  and  $W_o$  are learnable linear mappings.

The refined slice tokens are then fused with the convolutional feature:

$$\tilde{Y}(:, u) = \sum_{k=1}^K A_{u,k} \tilde{s}_k, \quad \hat{\sigma} = \phi_{\text{out}}(Y_c + \tilde{Y}) \in \mathbb{R}^{1 \times H \times W}, \quad (21)$$

where  $\phi_{\text{out}}$  is the output projection head.

In this way, the decoder combines local convolutional refinement with global structure-aware reconstruction. The convolutional path preserves local boundary details, while the slice-aware path models long-range stress transmission along load-bearing structures.

## 4. Validation

### 4.1. Preparations and datasets

To rigorously evaluate the proposed SIHNO framework, a high-fidelity metamaterial dataset was constructed from an architected porous structure that was virtually sliced along the build (Z) direction into layer-wise  $64 \times 64$  images, each a grayscale map distinguishing the solid material matrix from the void regions. Together with the associated stress labels, this yields on the order of  $4 \times 10^5$  pixel-level data points, which were partitioned into a 60%/40% training/testing split at the slice level along the build height using a fixed random seed for reproducibility. The corresponding ground-truth labels are the macroscopic Von Mises stress fields, obtained from finite-element simulations in Abaqus. The stress field ground truth is produced through a standard linear elastostatic analysis in Abaqus. More specifically, the metamaterial body was assigned as Aluminum 6061-T6 with  $E = 68.9$  GPa, and  $\nu = 0.33$ , and we treated the analysis process as a elastostatic problem using four-node bilinear quadrilateral elements (CPS4R). The bottom edge of the model was fully fixed, and a uniform force of 1000 Pa was applied on the top edge of the model. After the static equilibrium is solved, the Von Mises stress field was yielded and extracted for subsequent model training.

All experiments were conducted on a single NVIDIA 4080 Desktop Graphics Processing Unit (GPU) with PyTorch. During the training phase, we used the Adam optimizer with an initial learning rate of  $1 \times 10^{-3}$ , and set the epoch of the training process to 200 with a batch size of 8. L1 norm was selected as the loss function to effectively penalize localized stress concentrations and to enforce sharp spatial gradients along the void-solid interfaces.

To evaluate rotation robustness of all models, we rotate both the input topology slice and the corresponding ground-truth stress field by the same angle. This setting evaluates rotation equivariance of each model. No rotational data augmentation is used during training unless otherwise specified.

#### 4.2. Metrics

To comprehensively evaluate the performance of each model, we employed a set of eight metrics, complemented by an additional stress-leakage diagnosis. Global accuracy and statistical correlation of each model were assessed using Mean Absolute Error (MAE), Relative L2 Norm (Rel.L2), and the Coefficient of Determination ( $R^2$ ). Furthermore, pattern-level consistency were evaluated via Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM), where Higher SSIM and PSNR values indicate that the model preserves the continuous morphological flow and maintains high stress field prediction accuracy.

Crucially, to validate whether the model inherently respects the elastostatic principles, we introduce three additional physical metrics: Gradient RMSE (Grad.RMSE), Laplacian Error (Lap.Err), and Boundary MAE (Bnd.MAE). Grad.RMSE measures the spatial derivative error of first-order, evaluating the model's capacity to deal with sharp stress gradients. Lap.Err quantifies second-order spatial consistency of the predicted stress field, reflecting if the model is capable of suppressing non-physical oscillations and over-smoothed boundaries. Bnd.MAE calculates the absolute error along the void-solid interfaces. Moreover, to directly quantify non-physical stress leakage, we introduce a Void-Region Leakage (VRL) metric defined as the mean absolute deviation of the predicted stress from the stress-free ground-truth reference over the void region, a lower VRL indicates less leakage and hence better performance. These metrics explicitly validate the model's reliability in predicting critical stress concentrations, which are the primary catalysts for structural failure in AM components.

#### 4.3. Stress-field prediction results under fixed and rotated orientations

This subsection reports both quantitative and qualitative stress-field prediction results in two phases. One is on an unrotated test set, and another is under arbitrary in-plane rotational perturbations. We first assess the predictive fidelity and physical consistency of all models on an unrotated testing set, then conduct a second test by applying controlled rotational perturbations to the input images.

In the first phase, the performance of each model is listed in Table 1. It is shown that the proposed SIHNO establishes state-of-the-art performance, securing the best scores across all evaluation metrics. While neural operators like Transolver++ struggles with a high Laplacian Error due to inadequate physical constraints, SIHNO effectively reduces non-physical noise, suppressing prediction errors. Furthermore, in the specific stress magnitude prediction, SIHNO minimizes the Boundary MAE to 0.0532, achieving a 20.3% improvement over PINO.

To directly verify the suppression of non-physical stress leakage, we further compare the VRL in Table 2. SIHNO attains the lowest VRL among all models, confirming that the symplectic symmetry brought by HFP can reduce stress from leaking into void regions across discontinuous interfaces.



**Table 1.** Standard quantitative comparison on the metamaterial stress prediction task without rotation. **Bold** marks the best value and underline indicates the second best.

| Model             | MAE ↓         | Rel. $L^2$ ↓  | PSNR (dB) ↑  | $R^2$ ↑       | SSIM ↑        | Lap. Err ↓    | Grad. RMSE ↓  | Boundary MAE ↓ |
|-------------------|---------------|---------------|--------------|---------------|---------------|---------------|---------------|----------------|
| Transolver++ [26] | 0.0706        | 18.70%        | 18.30        | 0.6642        | 0.8549        | 0.2416        | 0.0867        | 0.0800         |
| FNO [18]          | 0.0538        | 12.72%        | <u>24.26</u> | 0.6866        | 0.8776        | <u>0.1157</u> | <u>0.0471</u> | 0.0674         |
| PINO [34]         | <u>0.0529</u> | <u>12.56%</u> | 24.16        | <u>0.6922</u> | <u>0.8801</u> | 0.1264        | 0.0492        | <u>0.0668</u>  |
| NSM [20]          | 0.1072        | 21.53%        | 17.49        | 0.5428        | 0.7615        | 0.2750        | 0.1151        | 0.1224         |
| PINN [31]         | 0.1771        | 31.10%        | 13.95        | 0.1619        | 0.2895        | 0.3205        | 0.1495        | 0.1912         |
| HNN [35]          | 0.1291        | 26.04%        | 15.77        | 0.3886        | 0.6982        | 0.3181        | 0.1278        | 0.1524         |
| SIHNO(Ours)       | <b>0.0438</b> | <b>11.28%</b> | <b>25.16</b> | <b>0.7323</b> | <b>0.8842</b> | <b>0.0972</b> | <b>0.0389</b> | <b>0.0532</b>  |


**Table 2.** VRL across all models. **Bold** marks the best value.

| Model | SIHNO (Ours)  | Transolver++ [26] | FNO [18] | PINO [34] | NSM [20] | PINN [31] | HNN [35] |
|-------|---------------|-------------------|----------|-----------|----------|-----------|----------|
| VRL ↓ | <b>0.0079</b> | 0.0110            | 0.0121   | 0.0168    | 0.1089   | 0.2996    | 0.1170   |


Figure 3 and Figure 4 displays the visual inspection of the predicted stress contours that corroborates the quantitative metrics. It is clear that explicit PDE embedded models like PINN fails to capture the localized stress concentrations, validating that point-wise PDE soft-constraints collapse across porous metamaterials. Similarly, NSM and HNN exhibit severe over-smoothing phenomena. In sharp contrast, SIHNO reconstructs the peak stress zones precisely at the discontinuities. By capturing these critical energy paths, SIHNO demonstrates its superior capability for stress prediction of porous structures.

We further evaluate the predictive robustness under arbitrary rotational perturbations.


Figure 5 illustrates the visual comparative results across different spatial orientations and Figure 6 the gradient map, the flaw of conventional operators is well exposed here. Notably, FNO and PINO suffer a major degradation in performance. As the input topology rotates, their predicted stress contours severely distort and fail to align with the altered geometric boundaries, generating greater errors.




**Figure 3.** Qualitative comparison of predicted Von Mises stress fields on the standard testing set with no rotation.



**Figure 4.** Qualitative comparison of predicted spatial stress gradients on the standard testing set. This figure highlights the models' capacity to resolve high-frequency boundary information and localized load-bearing paths in the material.



**Figure 5.** Qualitative comparison of stress field reconstruction across various surrogate models at representative rotation angles.



**Figure 6.** Qualitative comparison of predicted stress gradients under rotational perturbations.



This visual breakdown confirms that these standard operators memorize absolute pixel coordinates and lack intrinsic geometric symmetry that is encoded in SIHNO. Furthermore, models like NSM and PINN exhibit predictive collapse, outputting chaotic and non-physical noise distributions that lose the topological identity of the metamaterial. By lifting the raw images into a rotation-equivariant latent manifold, our model is equipped with the capability to recognize intrinsic mechanical properties even under arbitrary orientations.


Table 3 further supports the visual breakdowns, providing a rigorous numerical benchmark. At non-orthogonal rotation angles like  $10^\circ$  and  $45^\circ$ , the performance of standard models degrades greatly. FNO and PINO have degraded greatly, proving that their latent representations are tied to the absolute spatial grid, and may not suffice for rotation equivariance.

**Table 3.** Quantitative comparison under arbitrary rotations. **Bold** marks the best value of each metric and underline indicates the second best.


| Angle       | Model              | MAE ↓         | Rel. $L^2$ ↓  | PSNR (dB) ↑  | $R^2$ ↑       | SSIM ↑        | Lap. Err ↓    | Grad. RMSE ↓  | Boundary MAE ↓ |
|-------------|--------------------|---------------|---------------|--------------|---------------|---------------|---------------|---------------|----------------|
| $10^\circ$  | Transolver++ [26]  | <u>0.1000</u> | <u>22.14%</u> | <u>16.99</u> | <u>0.6393</u> | <u>0.8376</u> | <u>0.2211</u> | <u>0.0852</u> | <u>0.0813</u>  |
|             | FNO [18]           | 0.2158        | 43.75%        | 10.96        | 0.2290        | 0.5404        | 0.3184        | 0.1458        | 0.1821         |
|             | PINO [34]          | 0.2227        | 44.63%        | 10.79        | 0.2237        | 0.5322        | 0.2935        | 0.1373        | 0.1888         |
|             | NSM [20]           | 0.1977        | 34.62%        | 13.10        | 0.0382        | 0.1258        | 0.2400        | 0.1325        | 0.1680         |
|             | PINN [31]          | 0.2119        | 37.89%        | 12.22        | 0.0275        | 0.0920        | 0.2927        | 0.1535        | 0.1971         |
|             | HNN [35]           | 0.2025        | 36.48%        | 12.61        | 0.0232        | 0.1310        | 0.2565        | 0.1354        | 0.1822         |
|             | <b>SIHNO(Ours)</b> | <b>0.0634</b> | <b>15.60%</b> | <b>20.85</b> | <b>0.7630</b> | <b>0.8633</b> | <b>0.1487</b> | <b>0.0544</b> | <b>0.0354</b>  |
| $45^\circ$  | Transolver++ [26]  | <u>0.1674</u> | <u>31.74%</u> | <u>14.02</u> | <u>0.5324</u> | <u>0.7540</u> | 0.2958        | <u>0.1021</u> | <u>0.1621</u>  |
|             | FNO [18]           | 0.2262        | 45.83%        | 10.86        | 0.2092        | 0.5555        | 0.3190        | 0.1395        | 0.1985         |
|             | PINO [34]          | 0.2208        | 44.98%        | 10.93        | 0.2393        | 0.5609        | 0.3175        | 0.1399        | 0.1925         |
|             | NSM [20]           | 0.2032        | 36.48%        | 12.77        | 0.0158        | 0.1941        | <u>0.2430</u> | 0.1347        | 0.1685         |
|             | PINN [31]          | 0.2245        | 41.39%        | 11.52        | 0.0059        | 0.1627        | 0.3185        | 0.1638        | 0.2044         |
|             | HNN [35]           | 0.2037        | 36.81%        | 12.61        | 0.0156        | 0.1914        | 0.2488        | 0.1352        | 0.1851         |
|             | <b>SIHNO(Ours)</b> | <b>0.0953</b> | <b>21.22%</b> | <b>17.46</b> | <b>0.5926</b> | <b>0.7912</b> | <b>0.1785</b> | <b>0.0827</b> | <b>0.0763</b>  |
| $270^\circ$ | Transolver++ [26]  | 0.0699        | 18.34%        | 18.59        | 0.7178        | 0.8558        | 0.2346        | 0.0851        | 0.0417         |
|             | FNO [18]           | 0.0546        | 12.89%        | 24.00        | 0.8672        | 0.8769        | <u>0.1174</u> | <u>0.0479</u> | 0.0349         |
|             | PINO [34]          | <u>0.0532</u> | <u>12.64%</u> | <u>24.07</u> | <u>0.8702</u> | <u>0.8792</u> | 0.1263        | 0.0493        | <u>0.0330</u>  |
|             | NSM [20]           | 0.1070        | 21.56%        | 17.42        | 0.6776        | 0.7664        | 0.2777        | 0.1153        | 0.1077         |
|             | PINN [31]          | 0.1768        | 30.76%        | 13.93        | 0.2405        | 0.3403        | 0.3084        | 0.1440        | 0.1849         |
|             | HNN [35]           | 0.1279        | 25.79%        | 15.88        | 0.5938        | 0.7034        | 0.3078        | 0.1249        | 0.1312         |
|             | <b>SIHNO(Ours)</b> | <b>0.0438</b> | <b>11.28%</b> | <b>25.16</b> | <b>0.8892</b> | <b>0.8842</b> | <b>0.0973</b> | <b>0.0389</b> | <b>0.0233</b>  |

Figure 7 summarizes the eight metrics with error bars over the testing microstructures. The error bars are calculated as the standard deviation (SD) of the performance. In each panel, the star (★) highlights the best performance among the compared models. Across all eight metrics, SIHNO performs the best, and the error bars remain comparable to those of the baselines, indicating stable behavior across datasets.

Figure 8 extends our evaluation to a continuous  $360^\circ$  spectrum. This trajectory exposes the geometric vulnerability of conventional physics-embedded models in that they exhibit drastic V-shaped error spikes, reaching acceptable levels only at orthogonal increments due to grid realignment. In contrast, SIHNO maintains the lowest error envelope across all arbitrary angles, yielding reliable stress predictions.




**Figure 7.** Quantitative comparison across all metrics. Each bar indicates the mean performance over test microstructures, and error bars show the corresponding standard deviation. For each metric, the star (★) marks the model with the best mean value.



**Figure 8.** Evaluation of rotational robustness over a continuous 360° orientation spectrum.

To further evaluate the model's robustness against varying process parameters in additive manufacturing, Figure 9 plots the per-layer MAE over the full  $Z$  range. Here each  $Z$  level sets a different local boundary state for the slice. As depicted, SIHNO consistently maintains the lowest error envelope throughout the majority of the build process, demonstrating improved stability. Notably, almost all models encounter a sharp error surge near the geometric extreme at  $Z \approx 8.5$ , which corresponds to the complex stress concentrations at the top boundary layers. Even in this high-gradient regime, SIHNO provides more constrained and reliable inferences compared to pure data-driven baselines. This performance proves that by embedding both spatial symmetry and symplectic symmetry into a neural operator, the model can decode the implicit mechanical evolution along the each rotation angle, fulfilling the requirements for high-fidelity AM monitoring.




**Figure 9.** Evaluation of cross-layer generalization performance across the build height. SIHNO maintains the lowest error envelope throughout the majority of the build process, outperforming other baselines.

#### 4.4. Equivariance and noise robustness

To isolate the geometric symmetry induced by the RSLP module, we additionally quantify the intrinsic rotation equivariance of each surrogate. For a given in-plane rotation  $R_\theta$ , we compute the discrepancy between rotating the prediction and re-predicting from the rotated input averaged over the test set with a validity mask that excludes the zero-padded corners introduced by off-grid rotation.

As shown in Figure 10, the coordinate-tied operators incur large equivariance errors that collapse to near-zero only at the orthogonal angles ( $90^\circ, 180^\circ, 270^\circ$ ), producing the characteristic V-shaped profile of grid-locked models. Meanwhile, SIHNO maintains the lowest equivariance error across the entire spectrum, and it further contracts to near-zero at the D4 angles, directly evidencing that RSLP endows the operator with near-equivariant behaviour rather than fitting rotated samples.

Practical AM monitoring also introduces acquisition and preprocessing noise into the input slices. To assess robustness, we perturb the test inputs with additive Gaussian noise at three signal-to-noise ratios (40 dB, 30 dB, and 20 dB) and re-evaluate the accuracy. Table 4 report the resulting MAE, where SIHNO retains the lowest error at every noise level, rising only from 0.0448 (clean) to 0.0711 at 20 dB. This indicates that the symplectic symmetry represented by HFP acts as an implicit regularizer that suppresses high-frequency perturbations and discontinuities, reducing the effect of noise injections.



**Figure 10.** Rotation-equivariance error as a function of the input rotation angle  $\theta$ . SIHNO sustains the lowest error envelope at all angles and contracts to near-zero at the orthogonal (D4) angles.

**Table 4.** MAE under additive Gaussian input noise at three signal-to-noise ratios, evaluated at inference time without retraining. **Bold** marks the best value in each column, and underline indicates the second best.

| Model               | Clean         | 40 dB         | 30 dB         | 20 dB         |
|---------------------|---------------|---------------|---------------|---------------|
| FNO [18]            | 0.0538        | 0.0543        | 0.0576        | 0.0784        |
| PINO [34]           | <u>0.0529</u> | <u>0.0533</u> | <u>0.0561</u> | <u>0.0746</u> |
| NSM [20]            | 0.1072        | 0.1074        | 0.1079        | 0.1124        |
| PINN [31]           | 0.1999        | 0.2000        | 0.1998        | 0.2004        |
| HNN [35]            | 0.1291        | 0.1292        | 0.1295        | 0.1311        |
| <b>SIHNO (Ours)</b> | <b>0.0448</b> | <b>0.0453</b> | <b>0.0486</b> | <b>0.0711</b> |

#### 4.5. Ablation study

To dissect the individual contributions of geometric symmetry and symplectic symmetry within our framework, we conducted a comprehensive ablation study. We evaluate the full SIHNO model against two ablated variants: (i) the w/o RSLP variant, where the rotation-steerable geometric projector is replaced by a standard convolutional stem, and (ii) the w/o HFP variant, where the Hamiltonian latent coupling module is removed. The predictive performance of these models was quantified across three representative rotation angles at  $10^\circ$ ,  $45^\circ$ , and  $270^\circ$  respectively. We reuse the same eight evaluation metrics in the previous experiments, covering global accuracy, structural fidelity, and physics-aware high-frequency errors. The result is displayed in Table 5.



**Table 5.** Component ablation under input rotation. **Bold** marks the best value in each column, and underline indicates the second best.

| Angle | Variant  | MAE ↓         | Rel. $L^2$ ↓  | PSNR (dB) ↑  | $R^2$ ↑       | SSIM ↑        | Lap. Err ↓    | Grad. RMSE ↓  | Boundary MAE ↓ |
|-------|----------|---------------|---------------|--------------|---------------|---------------|---------------|---------------|----------------|
| 10°   | w/o HFP  | 0.0727        | 18.91%        | 18.95        | 0.6920        | 0.8430        | 0.1625        | 0.0590        | 0.0394         |
|       | w/o RSLP | <u>0.0693</u> | <u>16.37%</u> | <u>19.03</u> | <u>0.7578</u> | <u>0.8552</u> | <u>0.1555</u> | <u>0.0561</u> | <u>0.0383</u>  |
|       | SIHNO    | <b>0.0634</b> | <b>15.60%</b> | <b>20.85</b> | <b>0.7630</b> | <b>0.8633</b> | <b>0.1487</b> | <b>0.0544</b> | <b>0.0354</b>  |
| 45°   | w/o HFP  | 0.1142        | 25.88%        | 15.82        | 0.5236        | 0.7598        | 0.1974        | 0.0959        | 0.0875         |
|       | w/o RSLP | <u>0.1074</u> | <u>22.52%</u> | <u>16.83</u> | <u>0.5342</u> | <u>0.7780</u> | <u>0.1898</u> | <u>0.0880</u> | <u>0.0818</u>  |
|       | SIHNO    | <b>0.0953</b> | <b>21.22%</b> | <b>17.46</b> | <b>0.5926</b> | <b>0.7912</b> | <b>0.1785</b> | <b>0.0827</b> | <b>0.0763</b>  |
| 270°  | w/o HFP  | <u>0.0453</u> | <u>11.48%</u> | <u>24.95</u> | 0.8790        | <u>0.8804</u> | <u>0.0982</u> | <u>0.0397</u> | <u>0.0240</u>  |
|       | w/o RSLP | 0.0474        | 11.84%        | 24.71        | <u>0.8827</u> | 0.8739        | 0.1006        | 0.0409        | 0.0277         |
|       | SIHNO    | <b>0.0438</b> | <b>11.28%</b> | <b>25.16</b> | <b>0.8892</b> | <b>0.8842</b> | <b>0.0973</b> | <b>0.0389</b> | <b>0.0233</b>  |

The necessity of the dual-symmetry architecture is displayed clearly under variational rotational perturbations. At angles of severe grid-misalignment like 45°, removing either component incurs great penalty. Specifically, the w/o RSLP variant exhibits degraded global metrics, confirming that without geometric canonicalization, error brought by rotation corrupts the latent mapping inside the model. Conversely, eliminating the HFP module pushes the Boundary MAE greatly, indicating that without symplectic symmetry, the model fails to maintain local force equilibrium at complex porous interfaces. This comparison provides detailed illustrations of interface fidelity and stress leakage suppression. RSLP preserves rotation equivariance under arbitrary orientations, while HFP maintains stress prediction.

In summary, the ablation study comprehensively demonstrates that neither rotation-steerable geometric encoding nor Hamiltonian latent coupling alone is sufficient to deal with the complex physics of architected metamaterials under arbitrary image rotations. SIHNO acts as a symmetry regularizer that yields comprehensive and robust stress prediction results for porous metamaterials.

#### 4.6. Computational cost

Beyond predictive accuracy, practical deployment for online AM monitoring also requires favorable computational cost. In terms of computational cost, SIHNO contains approximately  $1.25 \times 10^6$  trainable parameters (1 253 137 ), and performs inference in about 1.9 ms per topology slice, which is orders of magnitude faster than the corresponding Abaqus simulation and thus supports online AM monitoring. Table 6 reports a comparison of model size, inference latency, peak GPU memory, and accuracy across all baselines on identical hardware. SIHNO maintains millisecond speed while attaining the lowest error.

**Table 6.** Computational cost on identical hardware. Inference latency is measured per  $64 \times 64$  slice after GPU warm-up. **Bold** marks the best MAE.

| Model               | Params    | Latency (ms) | GPU mem (MB) | MAE ↓         |
|---------------------|-----------|--------------|--------------|---------------|
| FNO [18]            | 1 188 385 | 1.64         | 57.8         | 0.0538        |
| PINO [34]           | 1 188 300 | 1.65         | 57.1         | 0.0529        |
| NSM [20]            | 1 581 553 | 0.64         | 56.4         | 0.1072        |
| PINN [31]           | 643 393   | 0.72         | 64.4         | 0.1771        |
| HNN [35]            | 1 201 409 | 0.95         | 53.3         | 0.1291        |
| <b>SIHNO (Ours)</b> | 1 253 137 | 1.91         | 69.5         | <b>0.0438</b> |

## 5. Conclusion

This work presented SIHNO, a symmetry-imbued Hamiltonian neural operator for stress-field prediction of porous metamaterials. The proposed Rotation-Steerable Lattice Projector encodes continuous-angle rotational symmetry through steerable convolutional kernels, reducing orientation bias under arbitrary in-plane rotations. The Hamiltonian-Fueled Propagator introduces a Hamiltonian-constructed latent bias that promotes global mechanical coherence without imposing pointwise PDE residuals, while the slice-aware convolutional decoder reconstructs sharp stress fields near discontinuous void-solid interfaces. Comprehensive experiments show that SIHNO achieves improved prediction accuracy, off-grid rotation robustness, and boundary stress fidelity compared with other neural operator baselines.

## Data availability statement

The data that support the findings of this study are not publicly available due to confidentiality, but are available from the corresponding author upon reasonable request.

## Declaration of generative AI and AI-assisted technologies

During the preparation of this manuscript, the authors used OpenAI ChatGPT 5.5 only to improve language and readability. The authors take full responsibility for the content of the manuscript.

## Acknowledgments

This research received no external funding.

## Authors' contribution

Conceptualization, Xinyu Lu and Kang Wang; methodology, Xinyu Lu; validation, Xinyu Lu; formal analysis, Xinyu Lu and Kang Wang; investigation, Xinyu Lu; resources, Kang Wang, Guodong Yi, Shuyou Zhang and Jianrong Tan; data curation, Xinyu Lu and Kang Wang; writing—original draft preparation, Xinyu Lu; writing—review and editing, Xinyu Lu and Kang Wang; visualization, Xinyu Lu and Kang Wang; supervision, Kang Wang, Guodong Yi, Shuyou Zhang and Jianrong Tan. All authors have read and agreed to the published version of the manuscript.

## Conflicts of interest

The authors declare that they have no conflict of interest.

## References

- [1] Zhou L, Miller J, Vezza J, Mayster M, Raffay M, *et al.* Additive manufacturing: a comprehensive review. *Sensors* 2024, 24(9):2668.

- [2] Inayathullah S, Buddala R. Review of machine learning applications in additive manufacturing. *Results Eng.* 2025, 25:103676.
- [3] Gao S, Sun Y, Xi L, Zhao T, Huang Y, *et al.* Recent advances in machine learning-assisted design of additive manufacturing metastructures: a review. *Compos. Struct.* 2025, 372:119525.
- [4] Soori M, Jough FKG, Dastres R, Arezoo B. Additive manufacturing modification by artificial intelligence, machine learning, and deep learning: a review. *Addit. Manuf. Front.* 2025, 4(2):200198.
- [5] Kim N, Lee D, Yu J, Cho SW, Lee D, *et al.* Toward a robust and generalizable metamaterial foundation model. *npj Comput. Mater.* 2025, 12(1):54.
- [6] Du Y, Mukherjee T, Li R, Hou Z, Dutta S, *et al.* A review of deep learning in metal additive manufacturing: impact on process, structure, and properties. *Prog. Mater. Sci.* 2026, 157:101587.
- [7] Qin J, Hu F, Liu Y, Witherell P, Wang CC, *et al.* Research and application of machine learning for additive manufacturing. *Addit. Manuf.* 2022, 52:102691.
- [8] Wang C, Tan X, Tor SB, Lim C. Machine learning in additive manufacturing: state-of-the-art and perspectives. *Addit. Manuf.* 2020, 36:101538.
- [9] Yang T, Jin Y, Lu Z, Liu T, Wang T, *et al.* Metamaterial acoustic lens-enabled noncontact real-time monitoring of mechanical properties of metal structures manufactured by additive friction stir deposition. *J. Mater. Process. Technol.* 2025, 341:118913.
- [10] Colosimo BM, Garghetti F, Grasso M, Pagani L. On-line inspection of lattice structures and metamaterials via *in-situ* imaging in additive manufacturing. *Addit. Manuf.* 2024, 95:104538.
- [11] Luo K, Zhao J, Wang Y, Li J, Wen J, *et al.* Physics-informed neural networks for PDE problems: a comprehensive review. *Artif. Intell. Rev.* 2025, 58(10):323.
- [12] Zhang W, Suo W, Song J, Cao W. Physics-informed neural networks (PINNs) as intelligent computing technique for solving partial differential equations: limitation and future prospects. *Sci. China-Phys. Mech. Astron.* 2026, 69(1):214602.
- [13] Chen Y, Lu L, Karniadakis GE, Dal Negro L. Physics-informed neural networks for inverse problems in nano-optics and metamaterials. *Opt. Express* 2020, 28(8):11618–11633.
- [14] Chen K, Dong X, Gao P, Chen Q, Peng Z, *et al.* Physics-informed neural networks for topological metamaterial design and mechanical applications. *Int. J. Mech. Sci.* 2025, 301:110489.
- [15] Hezaveh SB, Jahanbakhsh A, Firuznia R, Borzooei M, Chang L, *et al.* PINN-guided porous hierarchical structure for extreme thin film evaporation. *Int. J. Heat Mass Transf.* 2026, 256:128027.
- [16] Liu Q, Koric S, Abueidda D, Meidani H, Geubelle P. Toward signed distance function based metamaterial design: neural operator transformer for forward prediction and diffusion model for inverse design. *Comput. Methods Appl. Mech. Eng.* 2025, 446:118316.
- [17] Azizzadenesheli K, Kovachki N, Li Z, Liu-Schiaffini M, Kossaiji J, *et al.* Neural operators for accelerating scientific simulations and design. *Nat. Rev. Phys.* 2024, 6(5):320–328.
- [18] Li Z, Kovachki N, Azizzadenesheli K, Liu B, Bhattacharya K, *et al.* Fourier neural operator for parametric partial differential equations. *arXiv* 2020, arXiv:2010.08895.
- [19] Li Z, Huang DZ, Liu B, Anandkumar A. Fourier neural operator with learned deformations for pdes on general geometries. *J. Mach. Learn. Res.* 2023, 24(388):1–26.

- [20] Du Y, Chalapathi N, Krishnapriyan AS. Neural spectral methods: self-supervised learning in the spectral domain. In *Proceedings of The Twelfth International Conference on Learning Representations (ICLR 2024)*, Vienna, Austria, May 7–11, 2024, pp. 29197–29214.
- [21] Cranmer M, Greydanus S, Hoyer S, Battaglia P, Spergel D, *et al.* Lagrangian neural networks. *arXiv* 2020, arXiv:2003.04630.
- [22] Hildebrand S, Klinge S. Comparison of neural FEM and neural operator methods for applications in solid mechanics. *Neural Comput. Appl.* 2024, 36(27):16657–16682.
- [23] Lu L, Jin P, Karniadakis GE. DeepONet: learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv* 2019, arXiv:1910.03193.
- [24] Qiao W, Xu Y, Li H. Lie group convolution neural networks with scale-rotation equivariance. *Neural Networks* 2025, 183:106980.
- [25] Wu H, Luo H, Wang H, Wang J, Long M. Transolver: a fast transformer solver for PDEs on general geometries. In *Proceedings of the 41st International Conference on Machine Learning (ICML 2024)*, Vienna, Austria, July 21–27, 2024, pp. 53681–53705.
- [26] Luo H, Wu H, Zhou H, Xing L, Di Y, *et al.* Transolver++: an accurate neural solver for PDEs on million-scale geometries. In *Proceedings of the 42nd International Conference on Machine Learning (ICML 2025)*, Vancouver, Canada, July 13–19, 2025, pp. 41432–41449.
- [27] Helwig J, Zhang X, Fu C, Kurtin J, Wojtowysch S, *et al.* Group equivariant fourier neural operators for partial differential equations. *arXiv* 2023, arXiv:2306.05697.
- [28] Fu J, Xie Q, Meng D, Xu Z. Rotation equivariant proximal operator for deep unfolding methods in image restoration. *IEEE Trans. Pattern Anal. Mach. Intell.* 2024, 46(10):6577–6593.
- [29] Shen P, Herbst M, Viswanathan V. Rotation equivariant operators for machine learning on scalar and vector fields. *arXiv* 2021, arXiv:2108.09541.
- [30] Kim S, Song J, Shin S, Cho G, Kim S, *et al.* EqGINO: equivariant geometry-informed Fourier neural operators for 3D PDEs. In *Proceedings of the 43rd International Conference on Machine Learning (ICML 2026)*, Seoul, Republic of Korea, July 6–11, 2026.
- [31] Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* 2019, 378:686–707.
- [32] Farea A, Yli-Harja O, Emmert-Streib F. Understanding physics-informed neural networks: techniques, applications, trends, and challenges. *AI* 2024, 5(3):1534–1557.
- [33] Toscano JD, Oommen V, Varghese AJ, Zou Z, Ahmadi Daryakenari N, *et al.* From PINNs to PIKANs: recent advances in physics-informed machine learning. *Mach. Learn. Comput. Sci. Eng.* 2025, 1(1):15.
- [34] Li Z, Zheng H, Kovachki N, Jin D, Chen H, *et al.* Physics-informed neural operator for learning partial differential equations. *ACM/IMS J. Data Sci.* 2024, 1(3):1–27.
- [35] Greydanus S, Dzamba M, Yosinski J. Hamiltonian neural networks. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS 2019)*, Vancouver, Canada, December 8–14, 2019, pp. 15353–15363.

- [36] David M, Méhats F. Symplectic learning for hamiltonian neural networks. *J. Comput. Phys.* 2023, 494:112495.
- [37] Tanaka Y, Yaguchi T, Iwata T, Ueda N. Energy-consistent neural operators for Hamiltonian and dissipative partial differential equations. In *Proceedings of the 28th International Conference on Artificial Intelligence and Statistics (AISTATS 2025)*, Mai Khao, Thailand, May 3–5, 2025, pp. 1882–1890.
- [38] Du Y, Zhang J, Zeevi T, Dvornek NC, Onofrey JA. SRE-Conv: symmetric rotation equivariant convolution for biomedical image classification. In *Proceedings of 2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI)*, Houston, USA, April 14–17, 2025, pp. 1–5.
- [39] Yang Y, Dasmahapatra S, Mahmoodi S. Rotation-scale equivariant steerable filters. In *Proceedings of Medical Imaging with Deep Learning (MIDL 2023)*, Nashville, USA, July 10–12, 2023, pp. 695–715.
- [40] Li Y, Qiu Y, Chen Y, He L, Lin Z. Affine equivariant networks based on differential invariants. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Seattle, USA, June 17–21, 2024, pp. 5546–5556.
- [41] Weiler M, Cesa G. General E(2)-equivariant steerable CNNs. In *Proceedings of Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, Vancouver, Canada, December 8–14, 2019, pp. 14334–14345.
- [42] Cheng X, Wang L, Cao Y, Chen C. Learning non-separable Hamiltonian systems with pseudo-symplectic neural network. *J. Comput. Phys.* 2025, 550:114630.
- [43] Horn P, Ulibarrena VS, Koren B, Zwart SP. A generalized framework of neural networks for Hamiltonian systems. *J. Comput. Phys.* 2025, 521:113536.