Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A physics-guided graph neural framework for real-time reconstruction of mechanical microenvironments in engineered cartilage scaffolds

Yipeng Yue
Tianjin University of Technology
Lilan Gao
Tianjin University of Technology
Bolun Liu
Just Medical Equipment
Yuwei Yang
Tianjin University of Technology
Chunqiu Zhang
zcqarticle@163.com

Tianjin University of Technology

Research Article

Keywords: cartilage tissue engineering, graph neural networks, physics-guided learning, mechanical field reconstruction, real-time digital proxy

Posted Date: November 9th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11275186/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

A physics-guided graph neural framework for real-time reconstruction of mechanical microenvironments in engineered cartilage scaffolds

Yipeng Yue[a], Lilan Gao[a], Bolun Liu[b], Yuwei Yang[a*], Chunqiu Zhang[a*]

[a]Tianjin Key Laboratory for Advanced Mechatronic System Design and Intelligent Control, National Demonstration Center for Experimental Mechanical and Electrical Engineering Education, School of Mechanical Engineering, Tianjin University of Technology, Tianjin 300384, China

[b]Tianjin Key Laboratory of Bone Implant Interface Functionalization and Personality Research, Just Medical Equipment (Tianjin) Co., Ltd., Tianjin, China

First author: Yipeng Yue.

* Corresponding authors: Chunqiu Zhang (e-mail: zcqarticle@163.com) and Yuwei Yang (e-mail: buddhawei@126.com).

Keywords: cartilage tissue engineering; graph neural networks; physics-guided learning; mechanical field reconstruction; real-time digital proxy

ABSTRACT

The spatially heterogeneous mechanical microenvironment within engineered cartilage scaffolds critically regulates cell fate and extracellular matrix maturation, yet its real-time characterization remains limited by the computational burden of high-fidelity finite element analysis (FEA) simulations. This study develops Phys-MRGNet, a topology-conditioned physics-guided graph neural network that rapidly reconstructs the mechanical microenvironment of a 3D-printed cartilage scaffold under varying material states and loading conditions. The network was trained on 2,040 independent finite-element simulation sequences spanning an experimentally constrained viscoelastic parameter domain of the scaffold, with short, medium, and long edges modeled as distinct mechanical relations and a graph-Laplacian-based mechanical field consistency regularizer added without extra learnable parameters. On 20 independently held-out FEA sequences comprising 4.55×10^6 node-step predictions, the model achieved coefficients of determination of 0.9806 and 0.9805 for interior and domain-boundary constitutive states, and a 9% lower hotspot-edge deviation than the variant without the physical term. Inference takes approximately 0.1 s per sequence, roughly 1.2×10^3 times faster

than batch FEA on our hardware. Phys-MRGNet reproduces finite-element strain fields with high accuracy and enables rapid screening of mechanical conditioning strategies for engineered cartilage maturation through a Unity3D digital-proxy prototype, providing a mechanically relevant field descriptor for mechanobiological analysis of scaffold culture.

1. Introduction

Cartilage defects are among the most common musculoskeletal injuries, and scaffold-based tissue engineering remains the mainstream strategy for their repair [1]. Repair outcome is decided by the mechanical microenvironment experienced by seeded cells, which regulates matrix synthesis, phenotype maintenance, and differentiation through mechanotransduction [2]. This microenvironment cannot be observed directly and must be inferred from mechanical analysis. Cartilage and its collagen-based scaffolds are hydrated, time-dependent materials governed by poroelasticity and viscoelasticity alike [3,4], and extrusion printing and bioprinting continue to widen the accessible design space [5,6]. Strain is both the central mechanical cue acting on cells and the computational bottleneck of design optimization: a single finite-element-resolution parameter sweep takes days. In engineered cartilage culture, the key challenge is not merely predicting scalar mechanical properties, but rapidly resolving the spatially heterogeneous mechanical cues experienced by cells across the scaffold microarchitecture.

Among the mechanical descriptors relevant to scaffold-based mechanobiology, spatially resolved strain distributions provide a directly interpretable descriptor of local mechanical stimulation. Therefore, this work targets the maximum principal strain ε_1 at every node and every loading step. For a filamentary lattice compressed between parallel plates, the surface ε_1 is dominated by lateral expansion and filament bending; it is a biologically relevant scalar descriptor of the tensile component experienced by cells adhering to the filament walls, although it does not capture the full cellular mechanical microenvironment. It is also reported node by node by the solver and can be compared with surface strains measured by digital image correlation [7], placing simulation and experiment on closely related physical quantities.

Learning-based surrogates promise to remove this cost, and the limitations of existing approaches define our entry point. Convolutional and fully connected architectures have replaced finite elements in mechanical systems [8,9], and porous additively manufactured constructs have been analyzed computationally [10], but both require structured grid-like inputs; printed scaffolds must be voxelized first, and the surface detail that determines cellular stimulation is the first to be lost. Graph neural networks operate directly on meshes: graph convolution [11], graph attention [12], message passing [13], and inductive or relational variants [14,15] have made irregular domains routine targets for surrogate modeling [16,17], with comprehensive surveys available elsewhere [18,19]. In solid

mechanics, however, these models usually impose a single aggregation rule on all edges, so accuracy drops precisely in the high-gradient regions that matter most. Physics-informed learning embeds governing equations in the training objective [20] and improves consistency while reducing data demand [21,22], but strong partial differential equation (PDE) constraints are difficult to optimize, and existing implementations enforce strong-form residuals on volumetric grids [23,24], whereas the strain of interest here is reported on the surface mesh.

Phys-MRGNet addresses these limitations. It is a physics-guided graph surrogate designed specifically for strain fields of a viscoelastic scaffold on unstructured heterogeneous meshes, tailoring multi-relational attention and a lightweight physical regularizer to the mechanics of the scaffold graph. The architecture alternates geometry-aware message passing with relation-specific attention over three edge-length classes, giving each mechanical scale of the porous lattice its own learned weights, and adds a graph-Laplacian-based mechanical field consistency regularizer (graph-Laplacian strain-continuity penalty) that targets high-gradient errors at negligible cost. Trained on finite element data spanning an experimentally constrained viscoelastic parameter band, the model covers the material states represented across the sampled culture stages from a single training run, and is further integrated into a Unity3D in vitro loading prototype that supports interactive mechanical-state monitoring and screening of mechanical conditioning strategies.

The goal of this work is a topology-conditioned field surrogate that reconstructs the mechanical microenvironment of a predefined scaffold architecture from constitutive state and loading conditions; generalization to unseen scaffold geometries lies outside its scope. The targeted cost reduction makes the design space of loading protocols and material states explorable, turning the established link between scaffold mechanics and cellular mechanotransduction into a concrete engineering exploration problem. The overall workflow of the study is shown in Fig. 1. The work targets rapid reconstruction of high-resolution mechanical fields during cartilage-scaffold culture and loading, serving as a computational engine for bioreactor design, culture-protocol screening, and real-time monitoring of mechanical state.

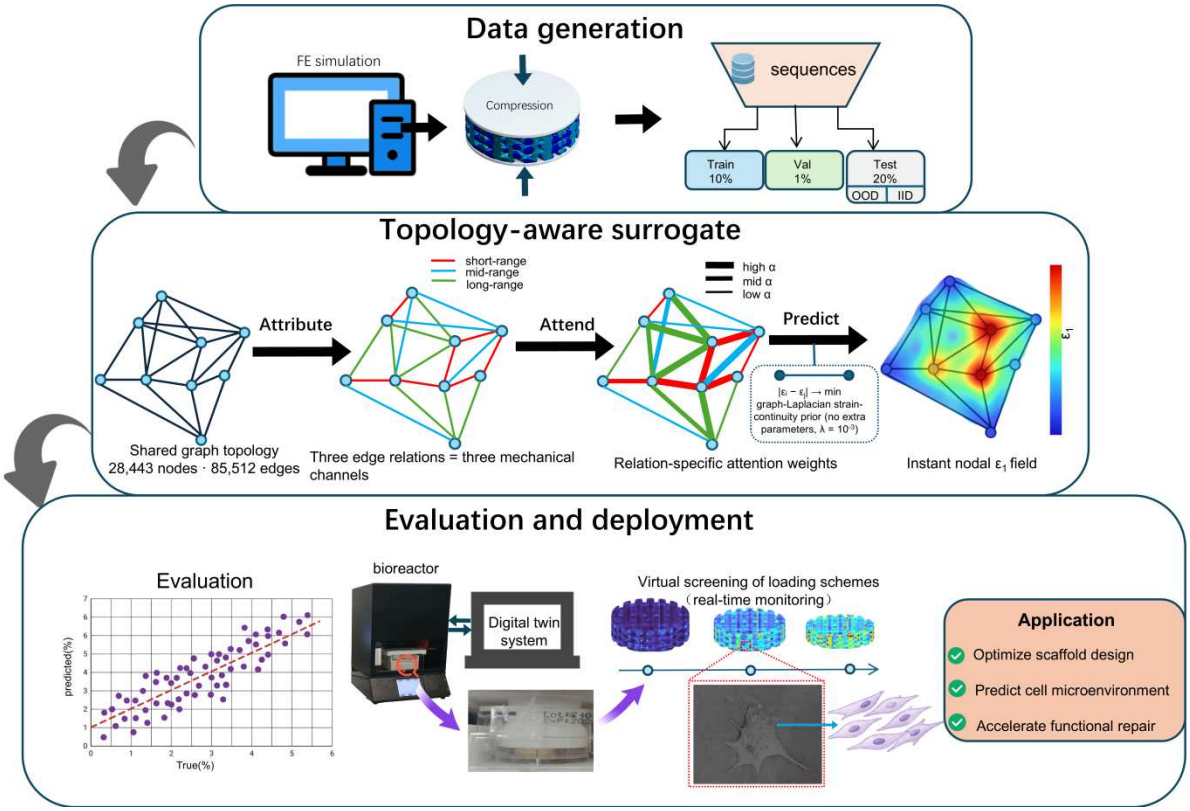


Fig. 1 Overall workflow of the study, from scaffold fabrication, mechanical characterization, and constitutive-domain identification through finite element dataset generation and Phys-MRGNet training to real-time mechanical-state reconstruction and mechanical-conditioning screening in the Unity3D digital proxy.

2. Materials and methods

2.1. Constitutive-state-conditioned finite-element dataset

This study develops a material-state-conditioned rapid prediction method for the strain fields of silk fibroin–collagen (SF–Col) composite scaffolds, whose mechanical properties evolve during culture as water is absorbed, the material degrades, and matrix accumulates [25]. Scaffolds with a regular pore architecture were fabricated by 3D bioprinting and mechanically characterized on days 0, 7, 14, 21, and 28, as shown in Fig. 2a; two scaffolds were independently fabricated per time point, one for unconfined compression at 0.5%/s up to 30% strain and one for constant-strain stress relaxation on a CARE M-100 testing system, as shown in Fig. 2b–d. The measured responses were jointly fitted to identify the constitutive parameters that define the material-state space; detailed material characterization was thus used only to establish the constitutive input domain, and the learning task is the reconstruction of the nodal strain field from graph-structured FEA data.

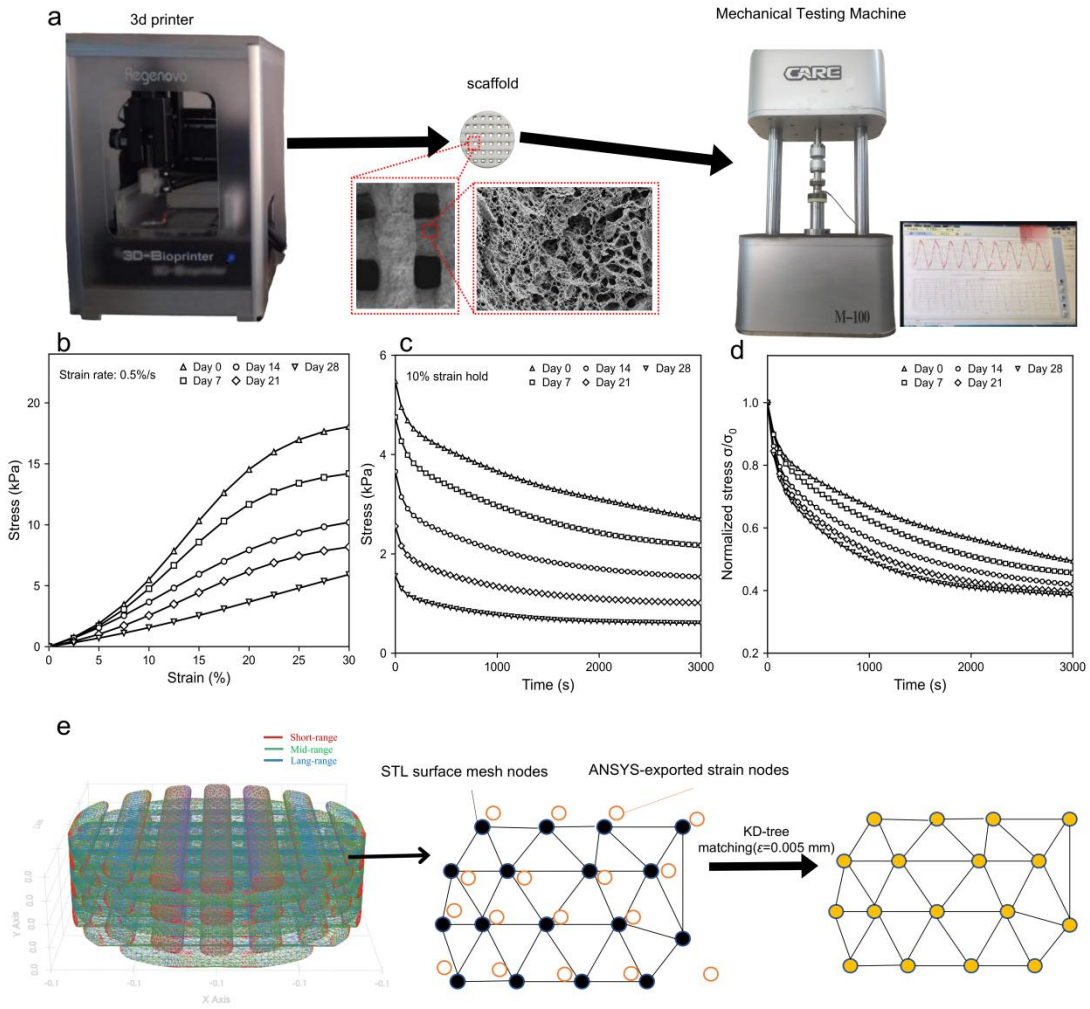


Fig. 2 Fabrication, multi-stage mechanical characterization, constitutive-state identification, and mesh-node alignment of the SF–Col scaffold. (a) Scaffolds were 3D bioprinted and tested in compression and stress relaxation on a CARE M-100 testing system; (b) stress–strain curves at a 0.5%/s strain rate for day 0–28 scaffolds; (c) stress-relaxation curves under a 10% strain hold; (d) normalized relaxation curves; (e) alignment of STL surface-mesh nodes with ANSYS strain nodes by KD-tree matching (0.005 mm tolerance)

FEA was performed for a representative SF–Col scaffold, the 3D-printed construct whose nominal geometry, printing parameters, and porosity are listed in Table 1, under displacement-controlled unconfined compression. A linear ramp from 0 to 1.0 mm was prescribed over an analysis time of 1.0 s, and eight converged states were extracted at $t = 0.125k$ s, $k = 1–8$, corresponding to nominal engineering strains of 4.17–33.3%; each ANSYS analysis step is hereafter called a loading step. Large-deformation kinematics were enabled. The surface is discretized into a triangular mesh of repeating filament topology with 28,443 nodes and 85,512 edges, shared by every simulation.

Table 1 Nominal geometry, printing parameters, and porosity of the SF–Col scaffold. Porosity was computed from the closed volume of the STL surface mesh (42.1 mm^3) relative to the cylindrical envelope

Parameter	Value
Scaffold height H_0	3 mm
Scaffold diameter	10 mm
Filament diameter	1 mm
Design printing gap (nominal pore size)	0.3 mm
Printing angle	90°
Porosity (relative to cylindrical envelope)	approximately 82%
Nominal compressive strain δ/H_0 (δ = 1.0 mm)	33.3%

The prediction target is the maximum principal total strain ε_1 at every node and step, that is, the first principal total strain output by ANSYS node by node, reported in percent strain throughout; ε_1 is the algebraically largest eigenvalue of the strain tensor and can be negative under fully confined compression. The material model contains no plasticity or thermal strain, so total strain equals mechanical strain.

All simulations ran in batch mode in ANSYS Mechanical 2026. The scaffold is treated as a homogeneous, linearly viscoelastic solid whose shear relaxation follows a two-term Prony series,

$$G(t) = G_\infty + G_1 e^{\frac{-t}{\tau_1}} + G_2 e^{\frac{-t}{\tau_2}} \quad (1)$$

decaying from the instantaneous modulus $G_0 = G_\infty + G_1 + G_2$ toward the long-term equilibrium value G_∞, with $E_0 = 2(1 + \nu)G_0$ and $\nu = 0.45$, consistent with the nearly incompressible behavior commonly assumed for hydrated soft biomaterials [26]. The two-term form is standard for the viscoelastic characterization of collagen hydrogels and tissue-engineered scaffolds [26] and provides a compact, identifiable representation of the relaxation spectrum for constitutive-state conditioning.

The sampling domain follows the magnitudes reported for collagen-based hydrogels [26] and is anchored by the experimentally identified envelope across the five sampled culture stages of days 0–28, E_0 20–80 kPa, g_1 0.15–0.30, g_2 0.41–0.45, τ_1 60–80 s, τ_2 850–2000 s, which falls entirely inside the sampling band; the resulting bounds define the constitutive-state space used for FEA sampling.

Strain results for each step are exported by ANSYS on the solver mesh and do not coincide pointwise with the STL surface mesh; KD-tree nearest-neighbor matching at a tolerance of 0.005 mm, as shown in Fig. 2e, defines the common valid node set across the eight steps, and sequences with fewer than 100 valid nodes are discarded. The graph representation, edge relations, and network architecture are described in Section 2.3.

2.2. Data partitioning for constitutive-domain boundary evaluation

Stratified sampling across this experimentally constrained parameter space generated 2,040 independent finite-element simulation sequences, all computed on a representative scaffold architecture, enabling controlled evaluation of constitutive-state generalization while preserving

geometric consistency; 1,632 entered the training pool and 408 the held-out pool. Of these, 200 sequences were used for model training to probe learning efficiency under limited labels, 20 unsampled sequences served for model selection, and the remaining 20 formed the final test set, ten interior and ten boundary sequences comprising 4.55×10^6 node-step observations. The experimentally constrained constitutive parameter space therefore defines the input domain over which the FEA field operator is sampled, while Phys-MRGNet learns a topology-conditioned constitutive-to-field operator that reconstructs the evolving mechanical microenvironment of the scaffold under varying material states and loading conditions.

Interior and boundary behavior are separated as follows. K-means clustering of the five standardized constitutive parameters divides the training pool into five clusters, and a stratified split assigns 70% of each cluster to the sampling pool from which the 200 training and 20 validation sequences are drawn. Test sequences are taken from the held-out pool in two modes: interior sequences from the bulk of the property distribution, and boundary sequences from the edge of the parameter space, distributed along the first principal component of the material-property vector and outside the central interval covered by the training sequences. Each mode contributes ten sequences, 2,275,440 node-step values, so the boundary set probes material states at the edge of the sampled band while remaining within the experimentally constrained constitutive domain; the comparison therefore concerns robustness to constitutive-domain shift inside the sampled parameter space.

2.3. Phys-MRGNet architecture

The FEA surface mesh is retained as the computational domain of the surrogate; it is represented as an undirected graph $G = (V, E)$, in which each node carries a 17-dimensional feature vector comprising 6 material and loading parameters, 8 geometric descriptors, including coordinates, distance to the centroid, local curvature, and normal components, and 3 topological indices, including degree, mean neighbor distance, and local clustering coefficient, standardized with training-set statistics only. As shown in Fig. 3(a), Phys-MRGNet maps the 17 node features to a per-node scalar output through a four-stage pipeline: a shared encoder projects the inputs to 128 channels, five stacked layers refine the node states in alternation, with the two relational attention layers at positions two and four, and a two-layer output head produces ε_1 per node; all layers share the same mesh connectivity, edge-type labels, and edge geometry.

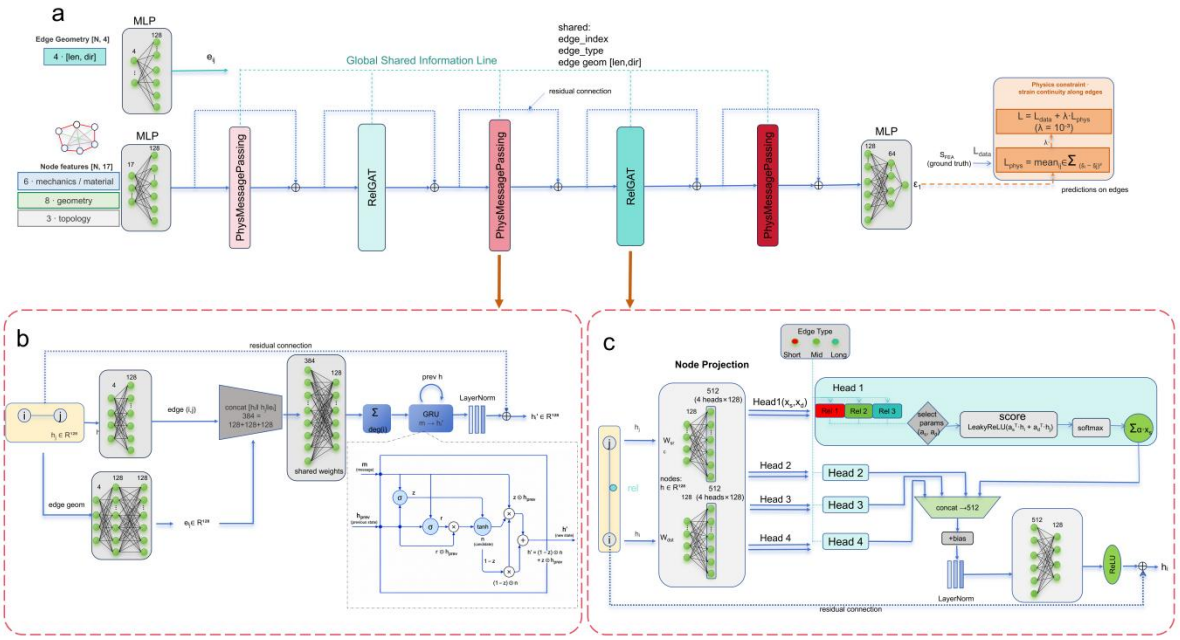


Fig. 3 Overall Phys-MRGNet architecture and the two core layer types. (a) Global data flow: the scaffold graph, 17 node features, edge-type labels, and edge geometry pass through a shared encoder into 128 channels, five alternating layers, and an output head, all sharing the same edge indices, edge types, and edge geometry, with residual connections bypassing every layer and the physics loss terms shown upper right. (b) Geometry-aware message-passing layer: node states h_i and h_j are combined with encoded edge geometry by a shared message MLP, messages are averaged over each node's neighbors, followed by a GRU update, layer normalization, and a residual connection. (c) Multi-relational attention layer: node features are projected into four heads; each head selects one of three relation-specific parameter sets by edge type, aggregates neighbors through a stabilized softmax, and the four heads are concatenated, normalized, and projected back to 128 channels

Edges are split into short, medium, and long classes at the 33.3rd and 66.7th percentiles of the edge-length distribution, corresponding to 0.176 and 0.254 mm and intended to represent local strain transfer, single-filament bending, and cross-filament structural coupling, respectively; this mapping is a mechanics-motivated inductive bias whose usefulness is supported empirically by the ablations. The number of classes follows an accuracy–efficiency trade-off on the validation set: two, three, and four classes give validation R^2 of 0.9760, 0.9806, and 0.9808, so three classes already saturate accuracy. The thresholds are relative percentiles, making the partition scale-independent; encoding the three scales as distinct relations lets the attention layers below learn independent weights per mechanical channel.

The mechanical rationale for this alternation is that not all edges are equivalent: a single aggregation rule does not distinguish the three mechanical scales and loses accuracy precisely in the steepest hotspot regions. Phys-MRGNet therefore alternates geometry-aware message passing with

attention that distinguishes the three edge relations, giving each mechanical scale independent learned weights; alternating the two layer types mitigates oversmoothing on densely connected meshes and reduces cost. The layer placement was selected on the validation set among three candidate configurations, and the scheme with attention at layers two and four performed best and is adopted throughout.

Geometry-aware message passing aggregates information along edges and uses edge geometry explicitly. For each edge (i, j), the edge length and unit direction form a four-dimensional geometric feature encoded into an edge embedding by a small multilayer perceptron (MLP); messages are produced by a shared MLP from the concatenated source state, target state, and edge embedding, and incoming messages at node i are averaged over its degree. Details of the gated update, layer normalization, and residual connection appear in Fig. 3(b).

The layer update is formalized as follows, from layer l to layer l + 1:

$$m_{ij} = \mathrm{MLP}_\theta\left(h_i \parallel h_j \parallel e_{ij}\right), j \in N(i) \tag{2}$$

$$h_i^{(l+1)} = h_i^{(l)} + \mathrm{LN}\left(\mathrm{GRU}\left(h_i^{(l)}, \frac{1}{|N(i)|}\sum_{j\in N(i)} m_{ji}^{(l)}\right)\right) \tag{3}$$

where $N(i)$ is the neighbor set of node i, e_{ij} is the encoded four-dimensional edge geometry feature, MLP_θ is the MLP shared by all edges, and LN denotes layer normalization; the residual connection is implicit in the recurrence for h, preserving an identity path for deep stacking.

Multi-relational attention determines neighbor contributions separately for the three edge relations. Source and target feature maps are expanded into four attention heads by bias-free linear projections; for each edge, each head computes a leaky-ReLU-activated score from the projected states and calls a relation-specific scoring parameter set [15], normalized by a numerically stable softmax over the node's neighborhood; the aggregated heads are concatenated, normalized, and projected back to 128 channels. The computation is shown in Fig. 3(c) and Eqs. (4)–(6); training uses dropout of 0.1, and the optimizer and learning-rate schedule are shared with the baselines, as described in Section 2.5.

For each attention head k, the relation-specific attention is:

$$e_{ij}^{(k)} = \mathrm{LeakyReLU}\left(a_{s,k}{}^{\top}(r_{ij})W_s^{(k)}h_i + a_{d,k}{}^{\top}(r_{ij})W_d^{(k)}h_j\right) \tag{4}$$

$$\alpha_{ij}^{(k)} = \frac{\exp\left(e_{ij}^{(k)}\right)}{\sum_{j'\in N(i)}\exp\left(e_{ij'}^{(k)}\right)} \tag{5}$$

$$h_i^{out} = \left(\sum_{j\in N(i)} \alpha_{ij}^{(k)}\hat{\mathrm{h}}_j^{(k)}\right)_{(k=1)}^{(4)} \tag{6}$$

where $r_{ij} \in$ {short, medium, long} is the relation label of edge (i, j), selecting one attention parameter set and one projection per side; the four head outputs are concatenated in order, followed

by layer normalization, a linear projection from 512 to 128 channels, and ReLU, with a residual connection giving the layer output.

2.4. Physics-guided training objective

For the smooth and crack-free scaffold considered here, continuity of the displacement field motivates local spatial coherence of the predicted strain field. The term constrains the spatial continuity of the predicted strain field on the scaffold graph and therefore acts as a graph-Laplacian-based mechanical field consistency regularizer; compared with strong-form PINNs that impose equilibrium or constitutive residuals directly [20], it suppresses locally non-smooth prediction artifacts while leaving the FE-supervised mapping untouched. Writing the predicted ε_1 at node i as $\hat{s}_i$, we add to the mean-squared data loss the discrete Dirichlet energy on the mesh graph, the graph counterpart of the smoothness prior in manifold regularization [27], as Eqs. (7) and (8):

$$L = L_{\text{data}} + \lambda L_{\text{phys}}, \tag{7}$$

$$L_{\text{phys}} = \frac{1}{|E|}\sum_{(i,j)\in E}\left(\hat{s}_i - \hat{s}_j\right)^2 = \frac{\hat{\mathbf{s}}^T \mathbf{L}\hat{\mathbf{s}}}{|E|} \tag{8}$$

where $L = D - A$ is the unnormalized graph Laplacian of the scaffold mesh, D the degree matrix, and A the adjacency matrix. The term adds no learnable parameters; it acts on the prediction only.

The form of Eq. (8) fits this geometry: a penalty built on differences of the predicted scalar field between neighboring nodes acts as a prior against nonphysical local fluctuations left by purely data-driven fitting. Because the penalty constrains a geometric property of the predicted field independently of its constitutive content, it remains valid for material states at the edge of the sampled band; Section 3.4 quantifies this.

λ was set by grid search on the validation set over 10^{-4} to 10^{-2}. Validation R^2 and mean hotspot-edge deviation are 0.9778 and 0.426% at $\lambda = 10^{-4}$, 0.9806 and 0.392% at $\lambda = 10^{-3}$, and 0.9785 and 0.404% at $\lambda = 10^{-2}$: too small a weight gives the term too little influence, too large a weight flattens genuine gradients and accuracy drops. $\lambda = 10^{-3}$ is used throughout. The network trains for 50 epochs with a batch size of one sequence, the AdamW optimizer, learning rate 10^{-3}, weight decay 10^{-5}, unit-norm gradient clipping, and a cosine schedule with linear warm-up over the first five epochs; the final model is the checkpoint with the highest validation R^2. Each training sequence carries 8 loading steps on a fixed 28,443-node surface graph, so a single epoch involves 4.55×10^7 node-step observations in the loss, ten times the held-out test set, a key reason the surrogate converges stably under a limited simulation budget.

2.5. Evaluation metrics and baseline configurations

Accuracy is measured by the coefficient of determination R^2, mean absolute error (MAE), and root-mean-square error (RMSE), in percent strain, aggregated over all test nodes and steps. All models are trained with three fixed random seeds, 42, 43, and 44, and reported as mean ± SD; the largest cross-seed SD of R^2 across configurations is 0.0018 and the performance ranking is stable. Because nodes and steps within a sequence are strongly correlated, statistics are kept at two levels: the node-step level only describes error distributions, the sequence level is the primary independent unit of evaluation with $n = 10$ per condition, and the seed level with $n = 3$ only assesses training stability. For the ablation analysis, ΔR^2 and ΔMAE in Fig. 6 denote the mean of paired sequence-level differences across the ten matched test sequences, whereas Table 2 reports the aggregate test-set metrics for each model.

Table 2 Ablation performance (mean ± SD over three random seeds; MAE in % strain). ΔR^2 and ΔMAE are defined in Fig. 6

Model	R^2 (IID)	MAE (%, IID)	R^2 (boundary)	MAE (%, boundary)	**Parameters**
Phys-MRGNet (full)	0.9806 ± 0.0011	0.1853 ± 0.0010	0.9805 ± 0.0012	0.1861 ± 0.0010	910,593
noPhys	0.9776 ± 0.0015	0.1951 ± 0.0020	0.9775 ± 0.0015	0.1954 ± 0.0020	910,593
noGeo	0.9767 ± 0.0012	0.2016 ± 0.0020	0.9766 ± 0.0012	0.2019 ± 0.0020	809,985
noRelGAT	0.9734 ± 0.0017	0.2125 ± 0.0020	0.9731 ± 0.0017	0.2137 ± 0.0020	839,425

The coefficient of determination is defined as:

$$R^2 = 1 - \frac{\sum_n (y_n - \hat{y}_n)^2}{\sum_n (y_n - \bar{y})^2} \tag{9}$$

where y_n and $\hat{y}_n$ are the reference and predicted strains and $\bar{y}$ their reference mean, summed over all pooled node-step samples; MAE and RMSE are the mean absolute value and root mean square over the same residual set and are not repeated.

To characterize extremes, we examine the empirical cumulative distribution of absolute nodal errors and its percentiles. The strain jump of an edge is the absolute difference of its endpoints at a given loading step, $|\varepsilon_{1,i} - \varepsilon_{1,j}|$, and hotspot edges are the 5% of edges with the largest jumps in the FE reference; pooled over all steps the threshold is 2.891% strain. Edge-wise fidelity is the absolute deviation between predicted and reference strain jumps, and the consistency of two hotspot sets is summarized by the Dice coefficient over the top-5% edges; two random 5% edge sets give Dice = 0.05, the chance-level overlap used when interpreting Section 3.2.

Baseline selection follows controlled comparison. GCN aggregates topological neighborhoods only, without any edge information; GAT adds attention weighting on the same aggregation but keeps

a single relation; MPNN injects edge geometry into the messages without distinguishing relation types, the lineage of mesh-based physical simulators [16]. Together the three form a layered contrast from topological aggregation through attention weighting and edge-geometry encoding to relation-specific modeling, locating each inductive bias architecturally, as shown in Fig. 4; the isolated contribution of each component is then judged by the ablations of Section 3.2. For the baseline comparison, GCN, GAT, and MPNN used the same node features, hidden width 128, network depth 5, residual connections, normalization, optimizer, learning-rate schedule, and model-selection protocol as Phys-MRGNet, but were trained with the data loss alone. The strain-continuity regularizer is not part of the baseline training setup; its isolated contribution is evaluated by the noPhys ablation, which retains the full Phys-MRGNet architecture and differs only in the training objective. Each ablation removes one component while keeping depth, width, and training schedule fixed.

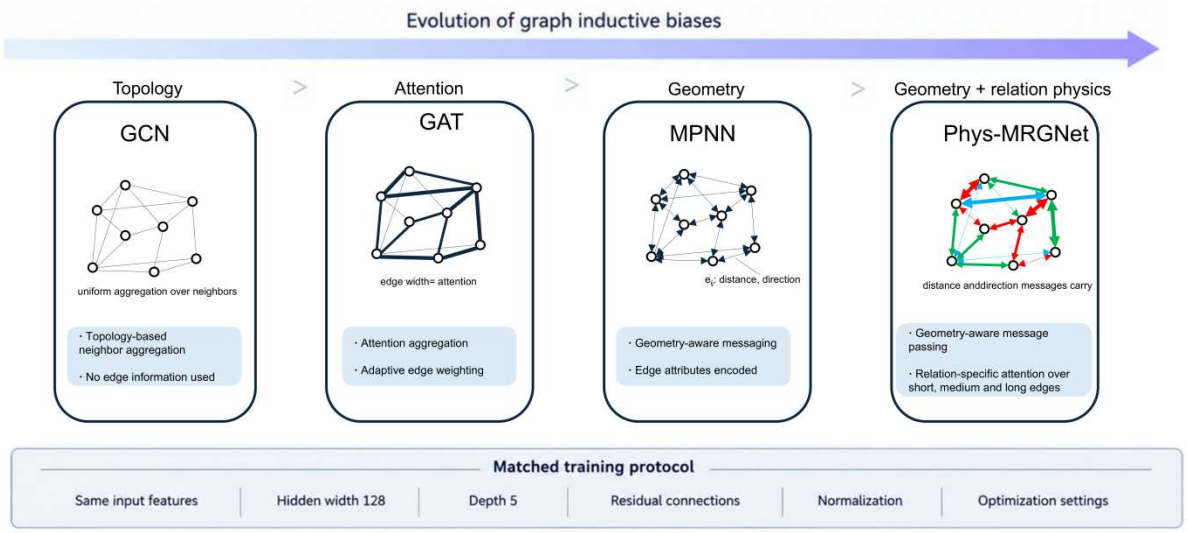


Fig. 4 Graph architectures used in the controlled comparison. Core graph operators of GCN (topology-based neighbor aggregation), GAT (single-relation attention), MPNN (geometry-aware message passing), and Phys-MRGNet (geometry-aware message passing combined with relation-specific attention over short, medium, and long edges); all four models use matched input representations, network depth and width, optimization, and evaluation protocols, each column representing a distinct graph-learning formulation, with the core graph operator defining the intended architectural contrast

3. Results

3.1. Performance relative to baseline graph architectures

Fig. 5 compares prediction quality across the four models as density scatter plots. Across both test conditions, Phys-MRGNet showed the smallest prediction dispersion and the highest sequence-level

R^2 and lowest MAE. The GCN cloud is the most diffuse, GAT narrows markedly, and Phys-MRGNet is the tightest in both Fig. 5(a) and Fig. 5(b). The error distributions remain similar between the two test modes, with no evident shift toward larger prediction errors at the boundary of the sampled constitutive domain. With the sequence as the independent experimental unit, Phys-MRGNet reaches R^2 = 0.9806 ± 0.0009 on the interior test set, with a 95% confidence interval of 0.9801 to 0.9812, MAE = 0.1853% ± 0.0057%, and RMSE = 0.2680% ± 0.0072%; the boundary test set gives R^2 = 0.9805 ± 0.0005, a 95% confidence interval of 0.9802 to 0.9808, MAE = 0.1861% ± 0.0029%, and RMSE = 0.2686% ± 0.0034%.

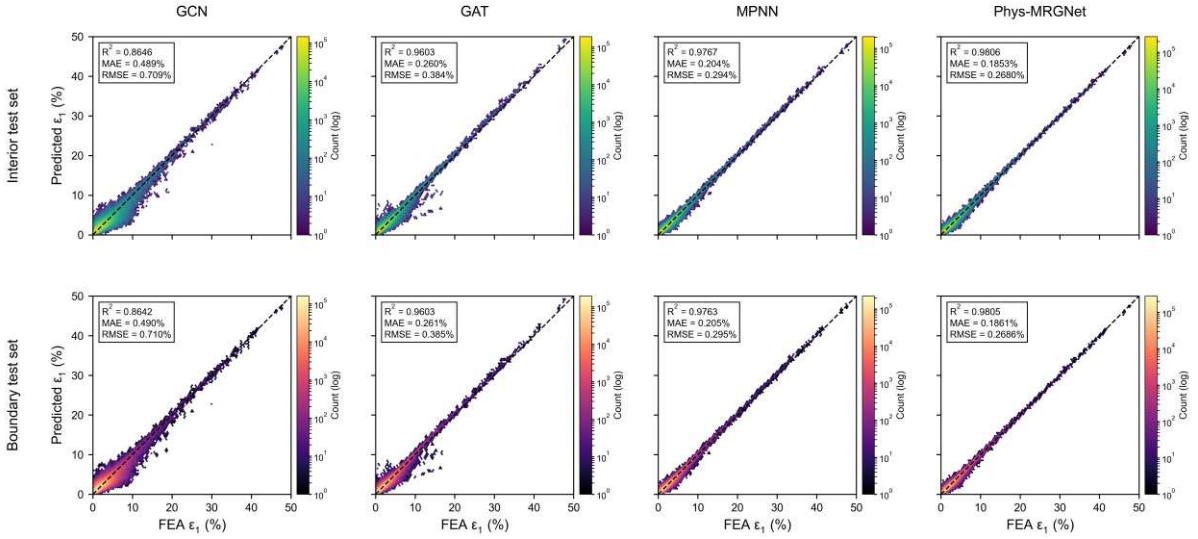


Fig. 5 Density scatter plots of predicted surface ε_1 (%) versus FE reference for GCN, GAT, MPNN, and Phys-MRGNet (all nodes and steps pooled; point density on a log scale). (a) interior test set; (b) boundary test set

3.2. Component ablation analysis

Fig. 6 and Table 2 present the ablation from two complementary angles. Fig. 6(a) compares the relative performance changes associated with each modification, expressed as ΔR^2 relative to the full model: removing relation-specific attention, noRelGAT, costs the most at −0.0072, followed by removing edge geometry, noGeo, at −0.0039, and the physical term, noPhys, at −0.0030. Fig. 6(b) shows the matching MAE penalties, which grow from 0.0098 to 0.0272 in the same order. Because all variants share training seeds and test sequences, the error bars in Fig. 6(a) and (b) show the standard deviation of the paired differences computed sequence by sequence. Fig. 6(c) tracks hotspot-edge overlap through loading as $\Delta Dice$ relative to the full model: noPhys rises slightly above the full model only at the first step and falls systematically below from the second onward, while noRelGAT shows the largest gap and does not close by the end of loading. The two test conditions, solid for interior and

dotted for boundary, nearly coincide, with boundary ΔMAE of 0.0093 to 0.0276, indicating that the ordering of component effects does not depend on the distance between test and training distributions.

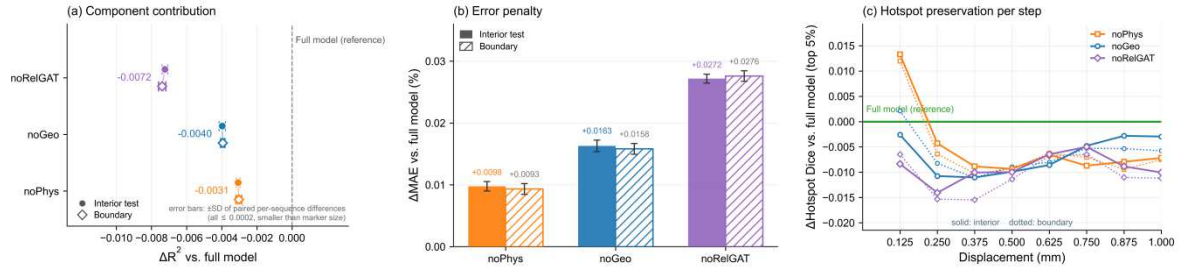


Fig. 6 Mechanistic decomposition of the ablation study. (a) ΔR^2 of each ablated variant relative to the full model (filled circles: interior; open diamonds: boundary; error bars: standard deviation of the paired differences across the ten matched test sequences). (b) ΔMAE penalty relative to the full model (solid bars: interior; hatched bars: boundary; error bars: standard deviation of the paired differences across the ten matched test sequences). (c) Hotspot-edge Dice difference (ΔDice, top 5%) relative to the full model versus applied displacement (solid: interior; dotted: boundary); the green zero line is the full-model baseline; noPhys exceeds the baseline only at the first step, while noRelGAT shows the largest gap

The cross-group comparisons provide supporting evidence and do not isolate individual components, because noGeo and MPNN differ in graph operator, training objective, and parameter count. On the interior set the noGeo variant reaches $R^2 = 0.9767 \pm 0.0012$, closely similar to the MPNN baseline (0.9767 ± 0.0013) at the seed level, despite 138,240 more parameters; the two separate only on the boundary set, noGeo higher by 0.0003, 0.9766 versus 0.9763. noRelGAT keeps geometric message passing but removes the relational split, and scores below the MPNN baseline overall (0.9734 versus 0.9767) despite 167,680 more parameters. GAT, with 1,007,105 parameters, is 10.6% larger than Phys-MRGNet, with 910,593, yet 0.0203 lower in R^2; noGeo is 20.6% larger than MPNN with no accuracy change. Within the tested configurations, the full model shows its largest gain when geometric message encoding and relation-specific attention are combined: neither component reproduces the full benefit on its own, and this interaction is the central observation of the section.

3.3. Selective action of the physical regularizer

The regularizer adds no learnable parameters; noPhys matches the full model at 910,593, so its accuracy gain costs no model capacity. Its effect is shown in Fig. 7(a): the cumulative distribution of absolute nodal error for the full model lies below that of the variant without the physical term everywhere, with the gap growing monotonically from approximately 2% at the median to 8.9% at the 99th percentile. A regularizer that improved the field uniformly would shift the whole curve down by a constant factor; instead, the distribution body barely moves while the tail contracts. This behavior

follows from the structure of the penalty: its gradient is proportional to the local inter-nodal difference, so corrective pressure concentrates where the predicted field is least smooth; a uniform weight therefore does not imply uniform action.

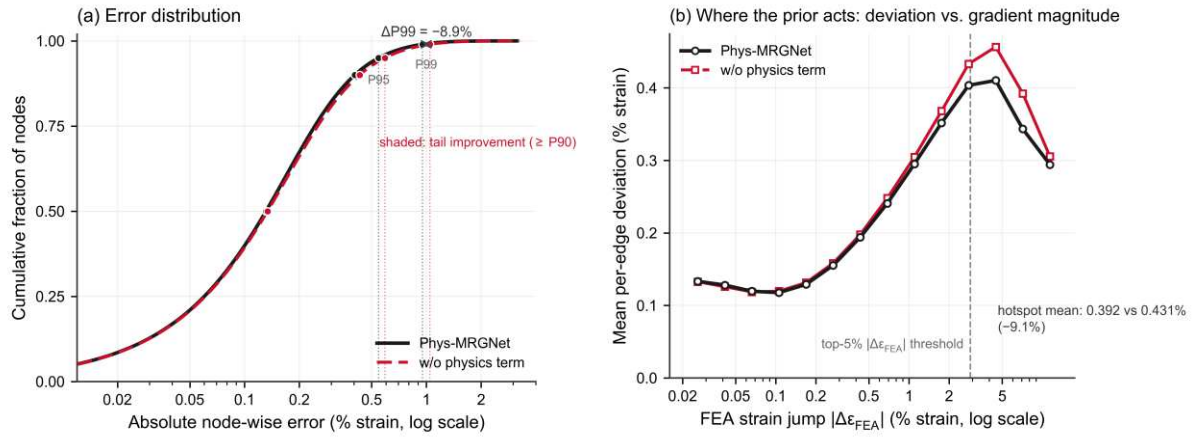


Fig. 7 Where the strain-continuity prior acts. (a) Cumulative distributions of absolute nodal error for the full model and the no-physical-term variant (log horizontal axis); shaded region: improvement above P90. (b) Per-edge deviation versus the magnitude of the FE strain jump (binned mean ± SEM); the dashed vertical line marks the top-5% hotspot-edge threshold, where the improvement of the prior concentrates

Fig. 7(b) bins the per-edge deviation by the magnitude of the strain jump: both models rise with gradient magnitude, but the variant without the physical term is systematically higher for $|\Delta\varepsilon|$ above approximately 1%, with the largest differences in the 2–3% band. Averaged over all edges, adding the physical term reduces the deviation from the FE strain jumps from 0.229% to 0.222%; on the high-gradient hotspot edges the same comparison gives 0.431% versus 0.392%, a 9% reduction, three times larger. The term thus selectively suppresses local residuals where the field is steepest. That selectivity matters because mechanotransduction links strain stimuli and stress-relaxation behavior to matrix synthesis and spreading in chondrocytes [28,29], and errors in gradient magnitude could affect downstream analyses of spatially heterogeneous mechanical cues.

3.4. Field-prediction fidelity at the constitutive-domain boundary

Fig. 8(a) shows predicted and reference fields for one boundary sequence at three displacements, 0.125, 0.5, and 1.0 mm. The spatial structure is faithfully reproduced: high-strain bands along the filament layout, shoulder regions where strain peaks, and the inward progression of steep gradients from the loaded surface are all captured. The residual deviations visible in the field maps are consistent with the aggregate R^2 of 0.9805 on the boundary set, indicating that the spatial structure protected by the physical term survives in the predictions.

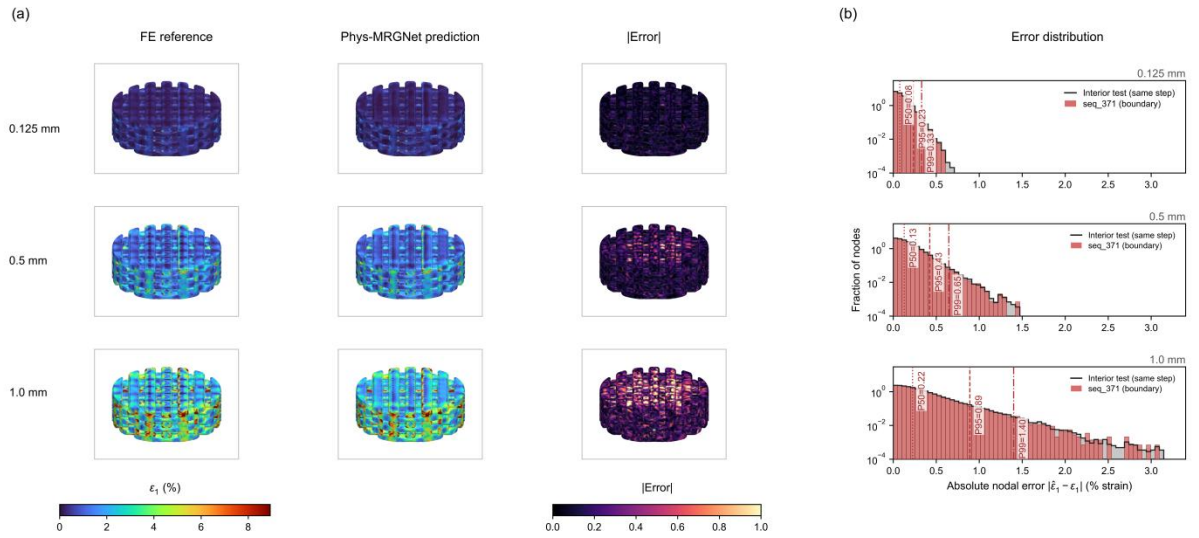


Fig. 8 Field-prediction fidelity at the constitutive-domain boundary. (a) Surface ε_1 fields of the boundary sequence (seq_371) at three steps (0.125, 0.5, and 1.0 mm); each row shows, from left to right, the FE reference, the Phys-MRGNet prediction, and the absolute error, with color bars for ε_1 (%) and absolute error (%). (b) Per-node absolute-error distributions for the same steps (red fills) against the overall interior-test distribution at the same load steps (grey step lines), with P50, P95, and P99 marked (log vertical axis)

Beyond aggregate metrics, the nodal error distribution gives more direct evidence on boundary behavior. As shown in Fig. 8(b), the per-node absolute-error distribution of the boundary set closely matches that of the interior set at every percentile: medians of 0.13% in both, 95th and 99th percentiles of 0.55% and 0.95%, and a maximum error of 3.14% strain at a filament junction in the most strained step. Moving the material state to the edge of the parameter band introduces no additional error tail. This distribution-level evidence corroborates the aggregate R^2 of Section 3.1: the mapping learned at the edge of the band has the same error composition as in the data-rich interior.

4. Discussion

The practical value of the surrogate is the combined trade between accuracy and computational cost. A complete eight-step finite element solution takes approximately 2 min; under our hardware and software configuration Phys-MRGNet infers the same quantity on a single RTX 4090 in approximately 0.1 s, a speed-up of roughly 1.2×10^3 that is specific to this configuration and not a cross-platform universal. Training is a one-off offline cost, with the full seven-configuration suite taking approximately 6 h on the same resources. Timing compares the complete nonlinear equilibrium iterations of one eight-step FE sequence with single-sequence forward inference in PyTorch 2.0.1 using FP32 precision, including graph construction and feature normalization, averaged over warmed-

up runs. This shift moves parametric optimization from days to interactive time and makes virtual screening of mechanical conditioning strategies possible in place of physical trial and error.

The central architectural contribution is to treat edge heterogeneity as an explicit relational inductive bias rather than as an auxiliary geometric descriptor: geometry-aware message passing encodes how information is transmitted along an edge, whereas relation-specific attention controls the contribution of distinct mechanical scales. This separation provides a structured representation of local and nonlocal interactions within the scaffold mesh and accounts for the complementary gains observed in the controlled comparisons of Section 3.2, which cannot be explained by parameter count alone. The practical benefit of the work is improved fidelity in local high-gradient regions without added learnable parameters, together with fast field queries within the material-state domain at near-real-time inference cost.

The main effect of the physical regularizer lies not in global R^2 but in fidelity at the high-gradient tail: it has a comparatively modest effect on the bulk of the error distribution while cutting hotspot-edge deviation by 9%. High strain-gradient regions are where mechanotransduction acts on matrix synthesis and cell spreading [28,30]; improved fidelity there yields more faithful high-gradient descriptions for mechanobiological analysis. From a graph-spectral perspective, the Laplacian term acts as a field-level smoothness prior that preferentially suppresses high-frequency spatial artifacts while preserving the large-scale strain structure required by the FE-supervised mapping. The light regularizer also preserves the training stability of the purely data-driven model and avoids the optimization difficulties of strong PDE constraints.

That the model shows no measurable degradation beyond the seed spread at the edge of the training distribution follows from three design choices: the constitutive parameters enter as node features, so the network learns a topology-conditioned constitutive-state-to-field operator; the physical regularization is material-agnostic and keeps constraining field smoothness in data-sparse regions; and relaxation enters the field through the smooth mechanism of Poisson's-ratio drift. One model therefore covers the material-state range represented across the sampled culture stages without retraining, which suits online monitoring. Phys-MRGNet is thus better viewed as a protocol- and constitutive-state-conditioned spatial field operator defined for a specified scaffold topology. The boundary experiment quantifies robustness to constitutive-domain shift while holding the scaffold architecture fixed; this distribution-level stability is consistent with reports that graph-network surrogates generalize across loading conditions in continuum mechanics [31]. Generalization to unseen scaffold geometries is a distinct question and was not examined here, since all sequences share the single representative scaffold architecture defined in Section 2.1; extending the operator to new geometries, longer relaxation histories, or volumetric fields will require additional validation.

The learned map is conditioned by the loading protocol. Under the prescribed displacement ramp, each of the eight states is determined by the applied displacement together with the constitutive parameters, so the network predicts them without explicit temporal recurrence, capturing the state

variables that fix the field under this protocol. Within the present ramp the viscoelastic response stays close to its instantaneous limit, and richer loading histories, including varying rates and hold periods, would call for additional state inputs. The same pattern, a scaffold graph plus a constitutive state yielding a nodal field, offers a template for adaptation to other structure–material systems.

4.1. Real-time deployment of the field surrogate

The deployment chain is summarized in Fig. 9 and should be read as a real-time digital proxy whose present validation covers inference speed. Fig. 9(a) shows the custom bioreactor built for displacement-controlled loading of cartilage scaffolds. Fig. 9(b) shows the coupled software side, where the trained network is exported from PyTorch to ONNX and executed by the Unity neural-network inference backend, so the field is computed live at run time. Digital twins are established in manufacturing [32,33] and are extending into healthcare [34]; the present prototype supplies the real-time inference core such systems require. The sensor streams are currently monitoring channels and are not yet assimilated into the surrogate, so the prototype closes the inference-to-visualization loop, whereas the sensor-to-state estimation loop remains a future extension. Recursive estimators coupling finite elements with real-time measurements provide assimilation templates [35], the remaining components of closed-loop control exist in bioprocess practice [36,37,38], and as culture data accumulate these streams can join the node feature space and calibrate constitutive parameters online [39].

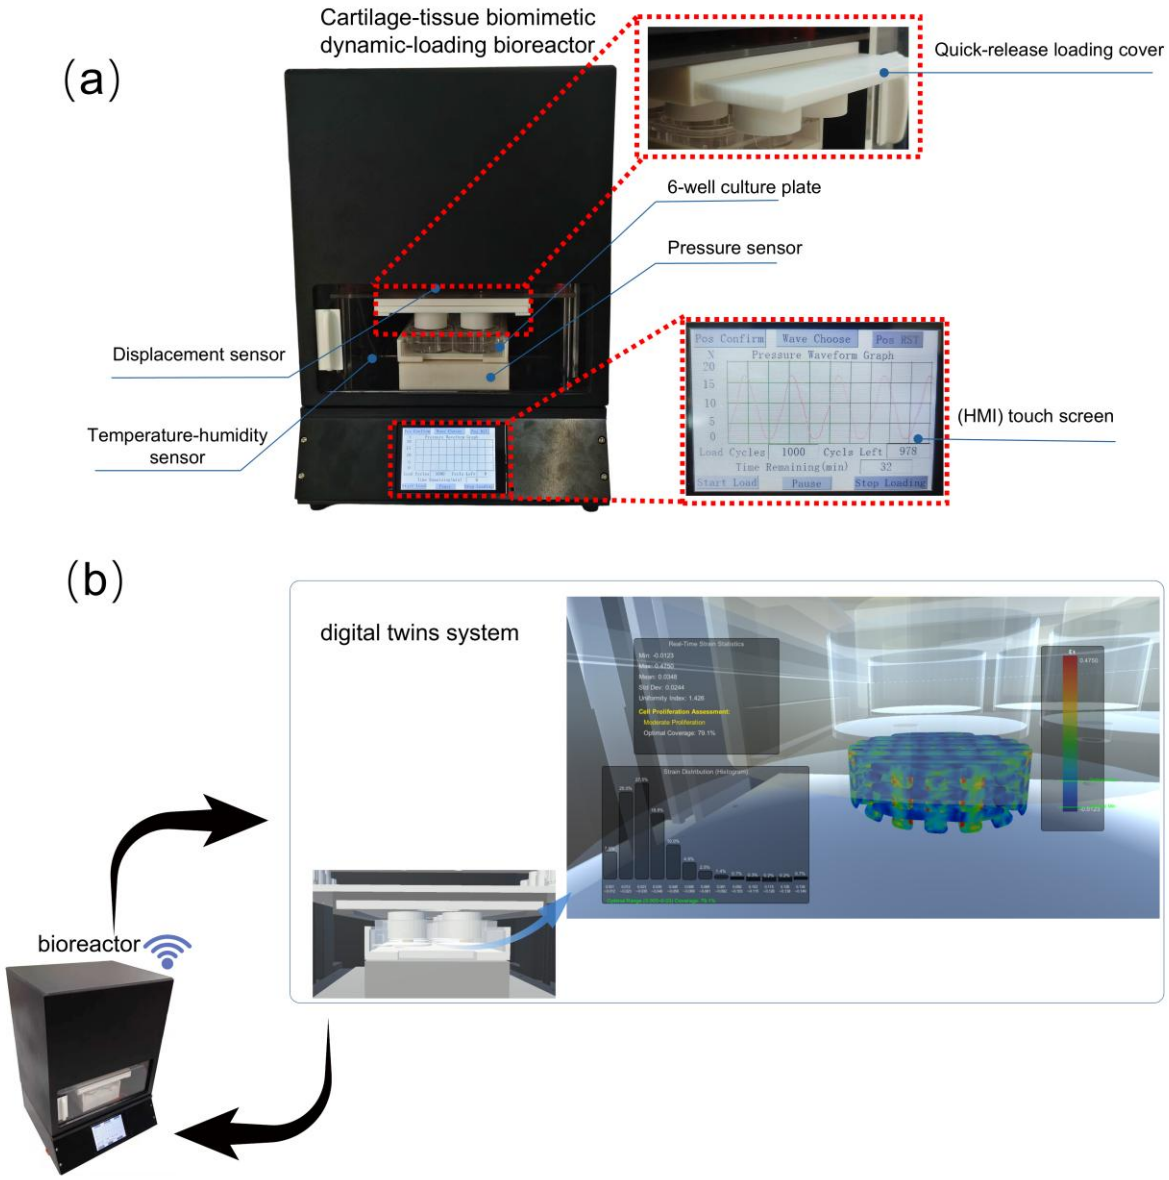


Fig. 9 Custom bioreactor for displacement-controlled cartilage-scaffold loading coupled to a real-time digital proxy. (a) The bioreactor integrates actuation, computation, and interaction modules: displacement and pressure sensors report the loading state in real time, temperature and humidity sensors monitor the culture environment, an HMI touchscreen supports local protocol setting and control, and a quick-release lid with a 6-well plate simplifies handling. (b) The bioreactor communicates wirelessly with the real-time digital proxy, which mirrors the scaffold ε_1 field and strain statistics in virtual space and provides an interactive mechanical-state simulation interface, enabling remote monitoring and virtual screening of mechanical conditioning

One forward pass over the scaffold graph for a single loading step takes approximately 12 ms on the same RTX 4090, comfortably inside a real-time frame budget; the mesh, connectivity, and edge labels are loaded once at startup. The interface exposes the two inputs

the surrogate was trained on, applied displacement and the five constitutive parameters, as controls, and the ε_1 field together with strain statistics is recomputed and recolored in the virtual space as the controls move, so a user can sweep a loading protocol or a range of material states and watch the strain distribution respond.

4.2. Scope and outlook

The scope of this work is bounded in three respects, each with a clear extension path. Modeling: the constitutive law is a two-term Prony linear viscoelastic model on a single filamentary lattice topology; volumetric relaxation, creep, load control, and time-varying loading can be brought in by extending the input space and the constitutive form, while the graph backbone supports incremental fine-tuning [40] and transfer along established lattice-property surrogate routes [41,42,43,44]. Validation: accuracy is referenced to high-fidelity finite element results, and comparison against measured scaffold responses can be established through DIC/DVC experiments [45,46], for which mechanical loading itself alters matrix deposition and distribution [47]. Regularization: the physical penalty currently uses a fixed, uniformly applied weight, and adaptive or edge-class-scaled weighting is the direct next step.

Future work should therefore proceed along three directions: validation against measured scaffold responses and transfer to additional printed geometries; temporal modeling across steps and extension to volumetric meshes, which would move predictions from surface ε_1 toward the interior quantities a complete mechanobiological model needs [48,49]; and sensor-assisted constitutive-state estimation, which would broaden the model toward closed-loop bioreactor applications.

5. Conclusions

This work presented Phys-MRGNet, a physics-guided graph neural network that reconstructs high-resolution mechanical landscapes associated with cellular mechanostimulation during scaffold conditioning, resolved as nodal ε_1 fields directly on the unstructured surface mesh of the investigated scaffold system. Across the held-out sequences, the model retained finite-element-level accuracy for both interior and domain-boundary material states. The ablations showed that relation-aware attention contributes the largest single gain, but only together with geometric message encoding, and that the strain-continuity term concentrates its improvement in high-gradient hotspot edges without learnable parameters. Deployed through ONNX and Unity3D, the framework enables rapid exploration of mechanical conditioning strategies for engineered cartilage maturation.

References

1. Armiento AR, Stoddart MJ, Alini M, Eglin D (2018) Biomaterials for articular cartilage tissue engineering: Learning from biology. Acta Biomater 65:1–20. https://doi.org/10.1016/j.actbio.2017.11.021

2. O'Conor CJ, Case N, Guilak F (2013) Mechanical regulation of chondrogenesis. Stem Cell Res Ther 4(4):61. https://doi.org/10.1186/scrt211

3. Ateshian GA, Warden WH, Kim JJ, Grelsamer RP, Mow VC (1997) Finite deformation biphasic material properties of bovine articular cartilage from confined compression experiments. J Biomech 30(11–12):1157–1164. https://doi.org/10.1016/S0021-9290(97)85606-0

4. Mow VC, Kuei SC, Lai WM, Armstrong CG (1980) Biphasic creep and stress relaxation of articular cartilage in compression: theory and experiments. J Biomech Eng 102(1):73–84. https://doi.org/10.1115/1.3138202

5. Boretti G, Amirfallah A, Edmunds KJ, Hamzehpour H, Sigurjónsson ÓE (2024) Advancing cartilage tissue engineering: a review of 3D bioprinting approaches and bioink properties. Tissue Eng Part B Rev. https://doi.org/10.1089/ten.teb.2024.0168

6. Hutmacher DW (2000) Scaffolds in tissue engineering bone and cartilage. Biomaterials 21(24):2529–2543. https://doi.org/10.1016/S0142-9612(00)00121-6

7. Pan B, Qian K, Xie H, Asundi A (2009) Two-dimensional digital image correlation for in-plane displacement and strain measurement: a review. Meas Sci Technol 20(6):062001. https://doi.org/10.1088/0957-0233/20/6/062001

8. Hashemi A, Jang J, Beheshti J (2023) A machine learning-based surrogate finite element model for estimating dynamic response of mechanical systems. IEEE Access 11:54509–54525. https://doi.org/10.1109/ACCESS.2023.3282453

9. LeCun Y, Bengio Y, Hinton G (2015) Deep learning. Nature 521(7553):436–444. https://doi.org/10.1038/nature14539

10. Zhang Z, Li H, Fogel GR, Liao Z, Li Y, Liu W (2018) Biomechanical analysis of porous additive manufactured cages for lateral lumbar interbody fusion: a finite element analysis. World Neurosurg 111:e581–e591. https://doi.org/10.1016/j.wneu.2017.12.127

11. Kipf TN, Welling M (2017) Semi-supervised classification with graph convolutional networks. ICLR 2017. arXiv:1609.02907. https://arxiv.org/abs/1609.02907

12. Veličković P, Cucurull G, Casanova A, Romero A, Liò P, Bengio Y (2018) Graph attention networks. ICLR 2018. arXiv:1710.10903. https://arxiv.org/abs/1710.10903

13. Gilmer J, Schoenholz SS, Riley PF, Vinyals O, Dahl GE (2017) Neural message passing for quantum chemistry. ICML 2017. arXiv:1704.01212. https://arxiv.org/abs/1704.01212

14. Hamilton WL, Ying R, Leskovec J (2017) Inductive representation learning on large graphs. NeurIPS 2017. arXiv:1706.02216. https://arxiv.org/abs/1706.02216

15. Schlichtkrull M, Kipf TN, Bloem P, van den Berg R, Titov I, Welling M (2018) Modeling relational data with graph convolutional networks. ESWC 2018:593–607. https://doi.org/10.1007/978-3-319-93417-4_38
16. Pfaff T, Fortunato M, Sanchez-Gonzalez A, Battaglia PW (2021) Learning mesh-based simulation with graph networks. ICLR 2021. arXiv:2010.03409. https://arxiv.org/abs/2010.03409
17. Sanchez-Gonzalez A, Godwin J, Pfaff T, Ying R, Leskovec J, Battaglia P (2020) Learning to simulate complex physics with graph networks. ICML 2020. arXiv:2002.09405. https://arxiv.org/abs/2002.09405
18. Wu Z, Pan S, Chen F, Long G, Zhang C, Yu PS (2021) A comprehensive survey on graph neural networks. IEEE Trans Neural Netw Learn Syst 32(1):4–24. https://doi.org/10.1109/TNNLS.2020.2978386
19. Zhou J, Cui G, Hu S, Zhang Z, Yang C, Liu Z, Wang L, Li C, Sun M (2020) Graph neural networks: a review of methods and applications. AI Open 1:57–81. https://doi.org/10.1016/j.aiopen.2021.01.001
20. Raissi M, Perdikaris P, Karniadakis GE (2019) Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. J Comput Phys 378:686–707. https://doi.org/10.1016/j.jcp.2018.10.045
21. Cuomo S, Di Cola VS, Giampaolo F, Rozza G, Raissi M, Piccialli F (2022) Scientific machine learning through physics-informed neural networks: where we are and what's next. J Sci Comput 92(3):88. https://doi.org/10.1007/s10915-022-01939-z
22. Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L (2021) Physics-informed machine learning. Nat Rev Phys 3(6):422–440. https://doi.org/10.1038/s42254-021-00314-5
23. Li H, Miao Y, Sharif Khodaei Z, Aliabadi MH (2025) Finite-PINN: a physics-informed neural network with finite geometric encoding for solid mechanics. J Mech Phys Solids 203:106222. https://doi.org/10.1016/j.jmps.2025.106222
24. Sunil P, Sills RB (2026) FE-PINNs: finite-element-based physics-informed neural networks for surrogate modeling. APL Mach Learn 4(1):016106. https://doi.org/10.1063/5.0299671
25. Gemmiti CV, Guldberg RE (2006) Fluid flow increases type II collagen deposition and tensile mechanical properties in bioreactor-grown tissue-engineered cartilage. Tissue Eng 12(3):469–479. https://doi.org/10.1089/ten.2006.12.469
26. Meghezi S, Drouin B, Mantovani D (2017) Collagen hydrogel-based scaffolds for vascular tissue regeneration: mechanical and viscoelastic characterization. In: Characterization of Polymeric Biomaterials. pp 397–439. https://doi.org/10.1016/B978-0-08-100737-2.00016-9
27. Belkin M, Niyogi P, Sindhwani V (2006) Manifold regularization: a geometric framework for learning from labeled and unlabeled examples. J Mach Learn Res 7:2399–2434. https://jmlr.org/papers/v7/belkin06a.html

587 28. Wann AKT, Zuo N, Haycraft CJ, Jensen CG, Poole CA, McGlashan SR, Knight MM (2012)
588 Primary cilia mediate mechanotransduction through control of ATP-induced Ca^{2+} signaling in
589 compressed chondrocytes. FASEB J 26(4):1663–1671. https://doi.org/10.1096/fj.11-193649
590 29. Yu W, Sun W, Chen H, Wang J, Xue B, Cao Y (2022) Gradual stress relaxation of hydrogel
591 regulates cell spreading. Int J Mol Sci 23(9):5170. https://doi.org/10.3390/ijms23095170
592 30. Uzieliene I, Bironaite D, Bernotas P, Sobolev A, Bernotiene E (2021) Mechanotransducive
593 biomimetic systems for chondrogenic differentiation in vitro. Int J Mol Sci 22(18):9690.
594 https://doi.org/10.3390/ijms22189690
595 31. Brandstetter J, Worrall D, Welling M (2022) Message passing neural PDE solvers. ICLR 2022.
596 arXiv:2202.03376. https://arxiv.org/abs/2202.03376
597 32. Liu M, Fang S, Dong H, Xu C (2021) Review of digital twin about concepts, technologies, and
598 industrial applications. J Manuf Syst 58:346–361. https://doi.org/10.1016/j.jmsy.2020.06.017
599 33. Tao F, Zhang H, Liu A, Nee AYC (2019) Digital twin in industry: state-of-the-art. IEEE Trans Ind
600 Inform 15(4):2405–2415. https://doi.org/10.1109/TII.2018.2873186
601 34. Kamel Boulos MN, Zhang P (2021) Digital twins: from personalised medicine to precision public
602 health. J Pers Med 11(8):745. https://doi.org/10.3390/jpm11080745
603 35. Xie H, Song J, Zhong Y, Gu C (2020) Kalman filter finite element method for real-time soft tissue
604 modeling. IEEE Access 8:53471–53483. https://doi.org/10.1109/ACCESS.2020.2981400
605 36. Biziato D, Rologi E, Blaha B, Watts S, McCoy RJ, Gaddum N (2018) Application of the
606 Stratophase Ranger system and automated closed-loop feeding control for T-cell therapy
607 bioprocessing. Cytotherapy 20(5):S108–S109. https://doi.org/10.1016/j.jcyt.2018.02.320
608 37. Charelli L, Balbino TA (2025) Mechanically active organotypic-on-chip devices for dynamic cell
609 culture. In: Smart Organ-on-Chip Devices. pp 73–84. https://doi.org/10.1016/B978-0-443-13403-
610 6.00004-0
611 38. Song JH, Lee SM, Yoo KH (2018) Label-free and real-time monitoring of human mesenchymal
612 stem cell differentiation in 2D and 3D cell culture systems using impedance cell sensors. RSC Adv
613 8(54):31246–31254. https://doi.org/10.1039/C8RA05273E
614 39. Butler KT, Davies DW, Cartwright H, Isayev O, Walsh A (2018) Machine learning for molecular
615 and materials science. Nature 559(7715):547–555. https://doi.org/10.1038/s41586-018-0337-2
616 40. Perera R, Agrawal V (2023) A generalized machine learning framework for brittle crack problems
617 using transfer learning and graph neural networks. Mech Mater 181:104639.
618 https://doi.org/10.1016/j.mechmat.2023.104639
619 41. Bastek JH, Kumar S, Telgen B, Glaesener RN, Kochmann DM (2022) Inverting the structure–
620 property map of truss metamaterials by deep learning. Proc Natl Acad Sci USA 119(1):e2111505119.
621 https://doi.org/10.1073/pnas.2111505119
622 42. Lee S, Zhang Z, Gu GX (2022) Generative machine learning algorithm for lattice structures with
623 superior mechanical properties. Mater Horiz 9(3):952–960. https://doi.org/10.1039/D1MH01792F

624 43. Ma S, Tang Q, Liu Y, Feng Q (2022) Prediction of mechanical properties of three-dimensional
625 printed lattice structures through machine learning. J Comput Inf Sci Eng 22(3):031008.
626 https://doi.org/10.1115/1.4053077
627 44. Yüksel N, Eren O, Börklü HR, Sezer HK (2024) Mechanical properties of additively
628 manufactured lattice structures designed by deep learning. Thin-Walled Struct 196:111475.
629 https://doi.org/10.1016/j.tws.2023.111475
630 45. Gao Z, Desai JP (2010) Estimating zero-strain states of very soft tissue under gravity loading
631 using digital image correlation. Med Image Anal 14(2):126–137.
632 https://doi.org/10.1016/j.media.2009.11.002
633 46. Madi K, Tozzi G, Zhang QH, Tong J, Cossey A, Au A, Hollis D, Hild F (2013) Computation of
634 full-field displacements in a scaffold implant using digital volume correlation and finite element
635 analysis. Med Eng Phys 35(9):1298–1312. https://doi.org/10.1016/j.medengphy.2013.02.001
636 47. Kock LM, Ito K, van Donkelaar CC (2013) Sliding indentation enhances collagen content and
637 depth-dependent matrix distribution in tissue-engineered cartilage constructs. Tissue Eng Part A
638 19(17–18):1949–1959. https://doi.org/10.1089/ten.tea.2012.0688
639 48. Kelly DJ, Prendergast PJ (2005) Mechano-regulation of stem cell differentiation and tissue
640 regeneration in osteochondral defects. J Biomech 38(7):1413–1422.
641 https://doi.org/10.1016/j.jbiomech.2004.06.026
642 49. Prendergast PJ, Huiskes R, Søballe K (1997) Biophysical stimuli on cells during tissue
643 differentiation at implant interfaces. J Biomech 30(6):539–548. https://doi.org/10.1016/S0021-
644 9290(96)00140-6

Statements and Declarations

Funding

This work was supported by the National Natural Science Foundation of China [grant numbers 32271371, U23A6009, 12572351].

Competing interests

The authors have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The FEA simulation data and processed datasets supporting the findings of this study are available from the corresponding author upon reasonable request.

Author contributions

Yipeng Yue: Conceptualization, Methodology, Investigation, Visualization, Writing – original draft. Lilan Gao: Supervision, Writing – review and editing. Bolun Liu: Resources. Yuwei Yang: Conceptualization, Supervision, Funding acquisition, Writing – review and editing. Chunqiu Zhang: Conceptualization, Project administration, Supervision, Funding acquisition, Writing – review and editing.

Ethics approval

Not applicable. This article does not contain any studies with human participants or animals performed by any of the authors.