Computers, Materials & Continua

Tech Science Press

https://doi.org/10.32604/cmc.2026.086951

ARTICLE

Morphology-aware CNN with Statistical Descriptor Regression and Conditional Image Synthesis for Predicting Crack Patterns in Steel Fiber-Reinforced Concrete

Khaja Wahaajuddin Kawkabi, Andre Steve Nziko Talla Nziko, Prince Manyanya Ngangura and Xu Long*

School of Mechanics and Transportation Engineering, Northwestern Polytechnical University, Xi'an, China

*Corresponding Author: Xu Long. Email: xulong@nwpu.edu.cn

Received: 08 June 2026; Accepted: 03 September 2026; Published: 22 October 2026

ABSTRACT: Accurate characterization of crack evolution in steel fiber-reinforced concrete (SFRC) remains a fundamental challenge due to heterogeneous mesostructure and the highly nonlinear nature of fracture processes governed by fiber–matrix interactions. This paper proposes a morphology-aware two-stage framework that reformulates the ill-posed curve-to-image mapping into two well-posed steps. Stage 1 (StatNet) learns to regress an eleven-component vector of crack morphology descriptors directly from uniaxial stress-strain response and mix-design metadata using dedicated encoders. Stage 2 (PatternDecoder) synthesizes high-resolution crack-probability maps conditioned on predicted descriptors and a latent embedding of the stress-strain curve. The framework is trained and validated on 220 meso-scale Abaqus cohesive-zone model simulations with explicit fiber beam elements. The model achieves a global pooled coefficient of $R^2 = 0.940$ across all morphology descriptors on the held-out validation folds, per-sample inference time of 14.57 ms. Comparisons against DirectCNN (direct convolutional neural network) and CGAN (conditional generative adversarial network) baselines confirm the superiority of the proposed approach in terms of structural fidelity, parameter efficiency, and computational cost. Residual diagnostics confirm homoscedasticity and near-normal errors, while gradient sensitivity analysis highlights physically consistent transitions at peak stress. The morphology-aware decomposition transforms an ill-posed mapping into an interpretable, computationally efficient surrogate model that bridges mesoscale fracture mechanisms with rapid crack-pattern assessment.

KEYWORDS: Steel fiber-reinforced concrete; crack pattern prediction; morphology-aware CNN; surrogate modeling; cohesive zone model

1 Introduction

Steel fiber-reinforced concrete (SFRC) has emerged as a high-performance cementitious composite capable of mitigating the intrinsic brittleness of conventional concrete through an efficient crack-bridging mechanism and enhanced post-cracking tensile performance [1,2]. The addition of discrete steel fibers improves key mechanical properties, including energy absorption, fracture toughness, and fatigue resistance, controlling crack initiation and propagation relative to plain concrete [3,4]. These mechanical advantages have positioned SFRC a viable alternative material for demanding engineering applications, such as industrial pavements, tunnel linings, seismic-resistant structures, and blast-resistant elements [5,6]. However, despite these improvements, the mechanical response of SFRC remains fundamentally governed by its heterogeneous meso-scale structure and interfacial characteristics [7,8]. The stochastic nature of fiber

Copyright © 2026 The Authors. Published by Tech Science Press. 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.

distribution and orientation, together with the complex interfacial bond-slip interaction between steel fibers and the cementitious matrix, introduces uncertainty in predicting crack initiation sites, timing, and propagation paths. This uncertainty is intensified under constrained boundary conditions, where local stress distributions from grouting or material expansion further influence cracking behavior [9]. Consequently, conventional deterministic modeling frameworks are insufficient to fully capture the intrinsic unpredictability in the fracture behavior of steel fiber-reinforced concrete (SFRC), thereby posing significant challenges for reliability-based design and performance assessment [10,11].

Computational modeling is essential for understanding SFRC fracture mechanisms, with mesoscale approaches enabling the explicit representation of fibers and their interactions with the surrounding matrix [12,13]. Within this context, cohesive zone models (CZMs) provide a physically grounded framework for simulating fracture processes via traction–separation laws that govern progressive interfacial damage, stiffness degradation, and crack propagation [14,15]. Recent developments have integrated explicit fiber representations, modeled as beam elements within CZM frameworks, allowing simulation of fiber pullout, interfacial debonding, and the resulting toughening mechanisms that distinguish SFRC from plain concrete [16,17]. However, high-fidelity mesoscale simulations with stochastic fiber distributions remain computationally expensive, with Monte Carlo sampling requiring thousands of realizations for uncertainty quantification [18]. This computational cost limits the likelihood of reliability-based assessment and optimization-driven design of SFRC structural systems. Recent studies have demonstrated the efficacious application of data-driven machine learning models to predict the flexural tensile strength of fiber-reinforced concrete [19].

The advent of deep learning offers a complementary paradigm for accelerating computational fracture analysis through surrogate modeling and data-driven prediction frameworks [20–22]. Concurrently, physics-informed neural networks have emerged as a powerful paradigm for solving forward and inverse fracture problems, including brittle fracture and in-plane crack analysis [18,23]. In particular, convolutional neural networks (CNNs) have demonstrated strong performance in crack detection and segmentation tasks, with architectures such as U-Net, DeepCrack, and fully convolutional networks enabling accurate pixel-wise identification of fracture patterns [24,25]. Hybrid CNN-Transformer architectures have recently demonstrated superior performance in pixel-wise crack extraction [26]. In parallel, specialized detection frameworks such as YOLO (You Only Look Once)-based models have been adapted for material defect identification, demonstrating high accuracy in discontinuity recognition [27]. Moreover, generative adversarial networks (GANs) and conditional GANs (CGANs) extend these capabilities by enabling conditional synthesis of crack patterns from material descriptors, reducing computational costs relative to conventional finite element methods [28,29]. More recent data-driven frameworks have successfully employed deep-learned generative models for designing composite microstructures, demonstrating the potential of conditional generation in materials science [30]. However, despite recent advances, most approaches rely on end-to-end mappings from mechanical responses to crack fields, treating the problem as a monolithic regression task. This is fundamentally ill-posed, as identical stress–strain responses may correspond to multiple crack patterns under stochastic mesostructural inconsistency. Consequently, such models suffer from mode collapse and a loss of morphological fidelity, resulting in overly smooth or non-physical crack fields [31–33].

An alternative is to introduce physically meaningful intermediate representations that decompose the ill-posed mapping into well-posed sub-problems [34,35]. In this context, crack morphology descriptors offer a compact yet information-rich representation of fracture patterns, bridging low-dimensional mechanical inputs and high-dimensional spatial crack fields [36]. Recent work by Ajjarapu and Mallikarjunaiah [34] proposed CrackMorph-XAI-Net, demonstrating that morphology-aware learning can extract interpretable

structural features from crack images for engineering assessment. In parallel, hybrid data-driven frameworks combining semi-empirical mechanics with machine learning have demonstrated strong capability to capture complex anisotropic behavior, as illustrated by Long et al. [37] for Cu–Nb nanolaminates through coupled physics–data modeling of size and texture effects on elasto-plasticity. Nevertheless, existing approaches remain largely restricted to either image-level morphology extraction or direct end-to-end crack field synthesis, without establishing a unified mechanics-to-morphology-to-field decomposition. In particular, a structured pipeline that first maps mechanical response to a physically interpretable morphology space and then reconstructs spatial crack patterns conditioned on this representation remains largely unexplored [38,39]. This missing formulation is critical, as it directly addresses the non-uniqueness of stress–strain-to-crack mappings in stochastic SFRC mesoscale systems and enables a controllable reduction of dimensionality while preserving physical interpretability [40].

This paper proposes a morphology-aware two-stage framework for SFRC fracture prediction that transforms the ill-posed stress–strain-to-crack-field mapping into a hierarchical conditional generation problem governed by a low-dimensional manifold crack morphology. This manifold acts as a physically interpretable bottleneck, decoupling the interpretation of global mechanical response from local fracture synthesis and thereby decomposing the inverse problem into two well-posed learning stages. In the first stage, StatNet maps uniaxial stress–strain responses and mix-design descriptors into a compact eleven-dimensional latent space that encodes geometric, topological, and intensity-based characteristics of crack evolution. In the second stage, PatternDecoder reconstructs high-resolution crack-probability fields conditioned on both the inferred morphology descriptors and a shared latent embedding of the mechanical response, thereby enforcing mechanically consistent coupling between macroscopic constitutive behavior and mesoscale fracture topology. This morphology bottleneck reduces solution non-uniqueness, mitigates mode collapse in end-to-end generative models, and improves physical interpretability through a structured dual representation learning scheme that links macroscopic responses to mesoscale damage patterns. The framework is trained on a comprehensive mesoscale dataset generated via cohesive-zone modeling with an explicit fiber-beam representation and realistic bond–slip interactions, enabling physically consistent learning of stochastic fracture evolution in SFRC.

2 Numerical Dataset Generation via Computational Modeling

2.1 Formulation of the Mortar with Cohesive Zone Model

The cohesive zone model is employed to simulate crack initiation and propagation within the mortar matrix [14,15]. Within the finite element framework, the cohesive behavior illustrated in Fig. 1 is incorporated using segment-to-segment cohesive elements, governed by a bilinear traction–separation law [41].

Prior to damage initiation, the nominal tractions are linearly related to displacement jumps through penalty stiffness parameters. The constitutive relations are expressed as

$$\sigma_n = k_{n0} \langle \delta_n \rangle , \sigma_s = k_{s0} \delta_s \tag{1}$$

where k_{n0} and k_{s0} denote the normal and shear penalty stiffnesses, respectively. The Macaulay bracket $\langle\rangle$ ensures that no traction develops under compressive normal separation.

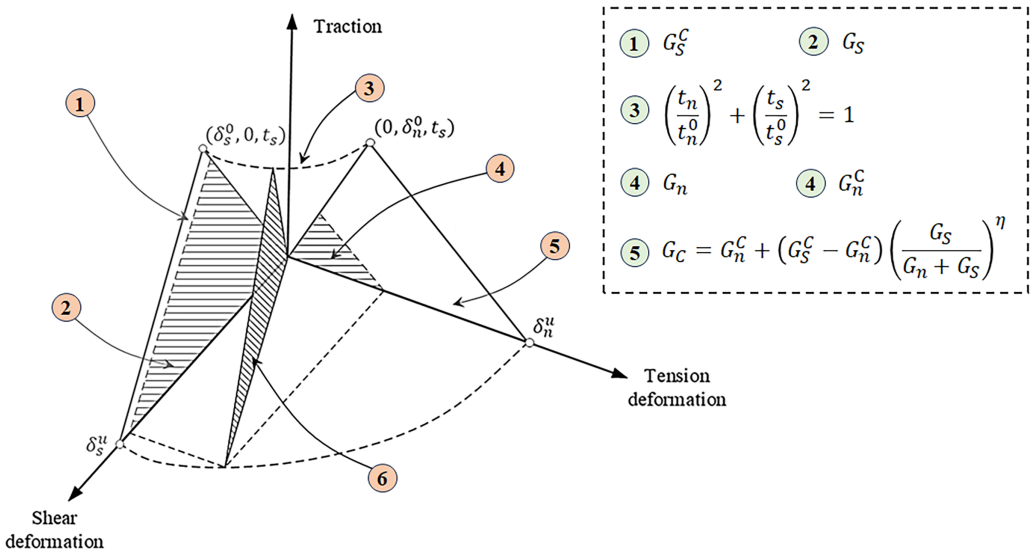


Figure 1: Bilinear mixed-mode traction–separation law for CZM.

Damage initiation is governed by a quadratic nominal stress criterion given by

$$\left(\frac{\langle \sigma_n \rangle}{\sigma_n^0}\right)^2 + \left(\frac{\sigma_s}{\sigma_s^0}\right)^2 = 1 \tag{2}$$

where σ_n^0 and σ_s^0 represent the tensile and shear cohesive strengths, respectively. Following damage initiation, the constitutive response is progressively degraded through a scalar damage variable ($D \in [0,1]$). The degraded tractions are expressed as

$$\sigma_n = (1-D)k_{n0}\langle \delta_n \rangle, \sigma_s = (1-D)k_{s0}\delta_s, D = \frac{\delta_m^f\left(\delta_m^{\max} - \delta_m^o\right)}{\delta_m^{\max}\left(\delta_m^f - \delta_m^o\right)} \tag{3}$$

where $\delta_m = \sqrt{\langle \delta_n \rangle^2 + \delta_s^2}$ is the effective displacement, δ_m^o and δ_m^f are the values at initiation and final failure, respectively. $\delta_m^{\max}$ represents the maximum effective displacement attained during loading. To account for mixed-mode fracture behavior, the critical energy release rate is defined according to the Benzeggagh–Kenane (B–K) criterion

$$G_C = G_C^I + \left(G_C^{II} - G_C^I\right)\left(\frac{G_s}{G_n + G_s}\right)^\eta \tag{4}$$

where G_C^I and G_C^{II} are the mode *I* and mode *II* critical fracture energies, respectively. G_n and G_s denote the work conjugates of normal and shear separations, and $\eta = 1.2$ is adopted for quasi-brittle materials [42].

2.2 Formulation of the Fibers with Embedded Beam Elements and Bond-Slip

To realistically capture the mesoscopic behavior of steel fibers within the SFRC, each fiber is discretized using two-node linear beam elements (B21). A dual-layer fiber representation is introduced to model fiber-matrix interactions while accurately avoiding artificial stiffness contributions. In this framework, two geometrically identical fiber networks are created: a dummy fiber layer and a real fiber layer, both sharing identical nodal coordinates.

The dummy fiber layer is embedded into the host 2D triangular continuum mesh representing the cementitious matrix. It is assigned negligible stiffness and serves solely as a kinematically constrained carrier, ensuring geometric conformity with the surrounding matrix without contributing to the global mechanical response. In contrast, the real fiber layer is a duplicate of the dummy configuration, occupying the same nodal positions but endowed with full mechanical properties of steel fiber (Young's modulus E = 200 GPa, density ρ = 7800 kg/m^3, Poisson's ratio ν = 0.3).

The fiber-matrix interaction is governed by nonlinear bond-slip behavior, enforced at corresponding node pairs between the real and dummy fiber through connector elements. For each node pair, a local coordinate system is defined and aligned with the fiber orientation to ensure that relative motion is accurately resolved along the fiber axis. The connector formulation constrains normal relative displacement to maintain kinematic compatibility while allowing nonlinear tangential slip governed by the bond law.

The bond-slip constitutive law is defined in tabular form, linking the connector shear force F to the relative slip $s = u_{\text{real}} - u_{\text{dummy}}$. The resulting response consists of an initial elastic regime, followed by peak bond strength, post-peak softening, and a residual frictional plateau representing complete debonding. This behavior is expressed as

$$F(s) = \text{Table}\left[(s_1, F_1), (s_2, F_2), \ldots, (s_n, F_n)\right] \tag{5}$$

where the discrete data points are calibrated against an experimental fiber pullout test. The initial slope of the table defines effective bond stiffness, while debonding initiates at the peak force $F_{\max}$ corresponding to slip s_e. The softening regime is captured through the descending branch until full debonding at slip s_f, beyond which a constant frictional resistance F_{fric} is maintained. The total work of bond failure is

$$G_{\text{bond}} = \int_0^{s_f} F(s)ds \tag{6}$$

which is evaluated numerically from the tabulated response.

During implementation, this nonlinear bond-slip law is enforced via connector elements with tabular force-slip definitions, linking each real fiber node to its corresponding dummy layer [43]. The consistent definition of the local coordinate system for each connector ensures correct alignment with fiber orientation, enabling accurate representation of the fiber pullout mechanism. Fig. 2 provides a schematic of dual-layer fiber configuration with connector-based bond modeling.

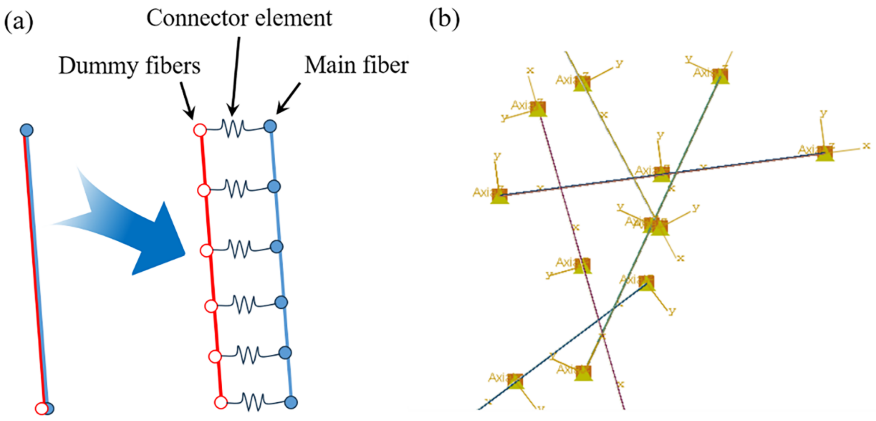


Figure 2: Dual-layer embedded fiber modeling with bond-slip interface. (**a**) Bond slip implementation concept. (**b**) Implementation in the finite element software Abaqus showing node connectivity across layers.

2.3 Numerical Implementation of CZM Insertion and Stochastic Fiber Packing

A structured algorithm is developed to convert a standard 2D triangular mesh into a CZM-ready configuration (Fig. 3, left). All internal edges between adjacent elements are identified and duplicated to generate zero-thickness interfaces. For each active internal edge, the shared nodes are duplicated, and a four-node cohesive element (COH2D4) is inserted to represent inter-element separation. After insertion, node duplication is resolved, element sets are organized for the Abaqus INP file. Steel fibers are modeled as two-node beam elements (B21). Their spatial distribution is generated via a random packing algorithm (Fig. 3, right), enabling controlled placement of macro- and micro-fibers within a prescribed region, including optional heterogeneous density regions. Each fiber is treated as an independent part instance, allowing individual property assignment. A geometric exclusion criterion enforces minimum inter-fiber spacing to avoid overlap and ensure physical realism [44]. The final packing database stores fiber orientations, coordinates, and identifiers for subsequent assembly.

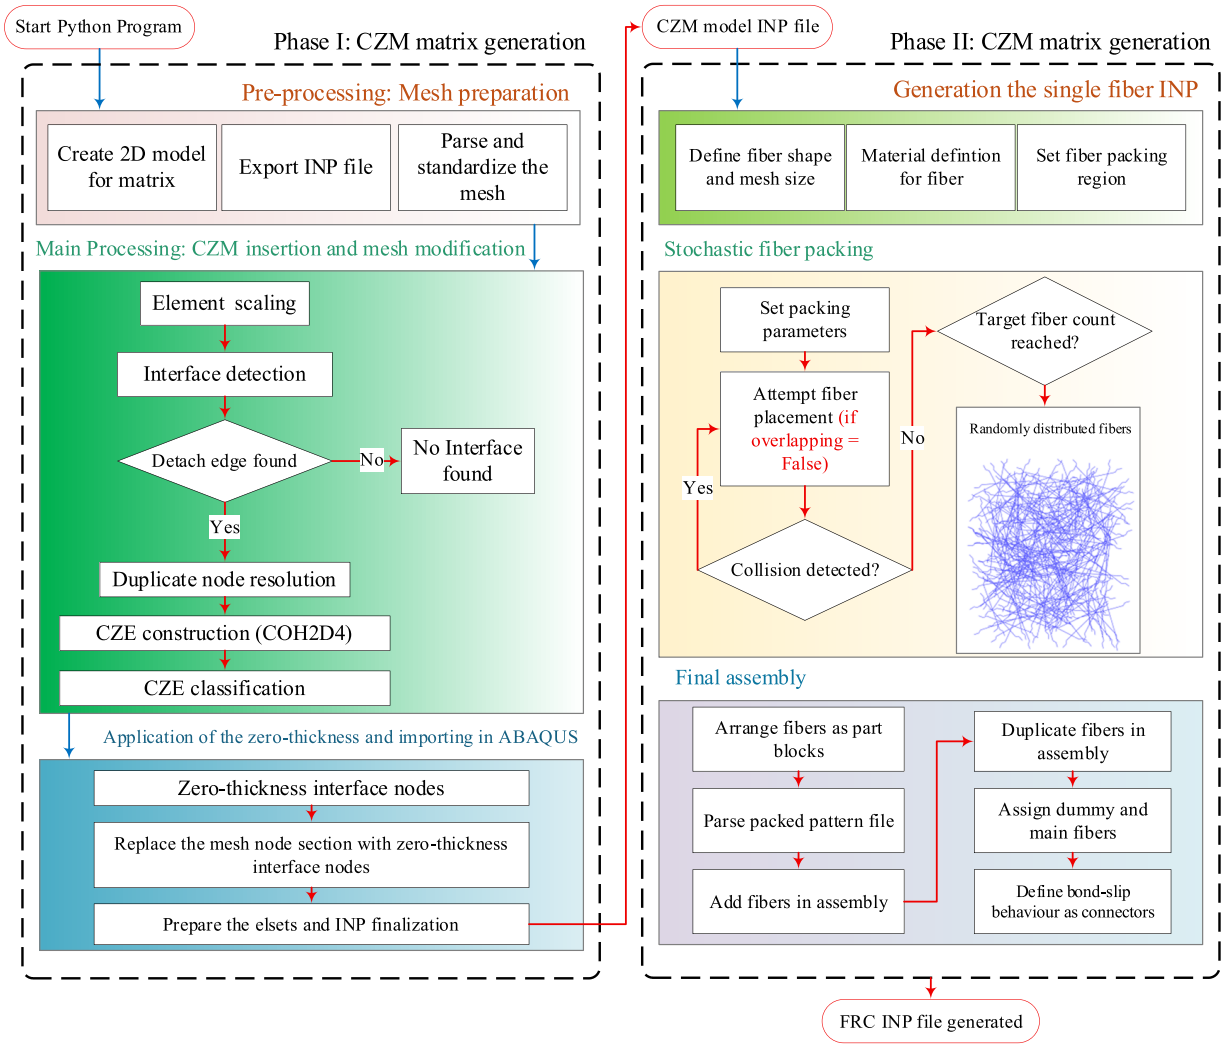


Figure 3: Algorithmic workflow for zero-thickness CZM insertion into the mortar matrix mesh and stochastic packing of discrete macro- and micro-fibers with collision detection.

In the global assembly, each fiber is duplicated into a dummy and a real layer, connected through coincident nodal pairs (Fig. 4). The dummy layer is embedded within the matrix. Fiber-matrix load transfer is governed by nonlinear bond-slip behavior (Section 2.2), enforced via node-to-node connector elements linking each fiber node to its corresponding real fiber node, avoiding unwanted stiffening while preserving accurate kinematic coupling and load transfer.

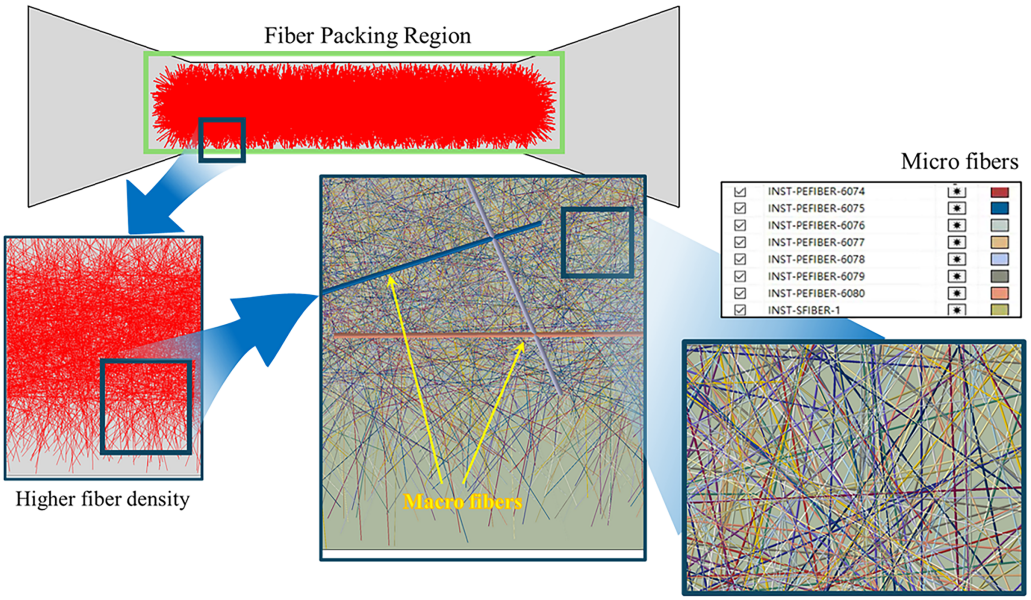


Figure 4: Representative SFRC meso-scale model after complete assembly, showing the mortar matrix with inserted zero-thickness cohesive elements.

Uniaxial tension is applied under quasi-static conditions as shown in Fig. 5. The bottom boundary is fully constrained ($u_x = u_y = 0$), while the top boundary is subjected to a monotonic displacement-controlled loading. The lateral boundaries remain traction-free. The macroscopic stress–strain response is obtained from the reaction force at the loaded boundary.

2.4 Numerical Model Validation and Dataset Specification

The SFRC meso-scale model adopts the dog-bone-shaped specimen geometry reported by Li and Deng [45], as illustrated in Fig. 6a. The specimen total height is 418 mm, with a central gauge section measuring 150 mm in length and 50 mm in width. The top and bottom grip sections are 80 mm high and 100 mm wide to ensure clamping. Rounded fillets of radius 70 mm transition smoothly between the gauge and grip sections to minimize stress concentrations. The bottom grip is fully constrained, while the top grip is subjected to a monotonic displacement-controlled tensile loading at a quasi-static rate of 0.15 mm/min.

The fiber–matrix interaction is governed by the cohesive-type bond–slip formulation introduced in Section 2.2. Since Li and Deng [45] did not report bond-slip parameters, these are calibrated independently against standard pullout tests for steel fibers in ultra-high-performance concrete (UHPC) matrices. The calibrated bond-slip curve is presented in Fig. 7a.

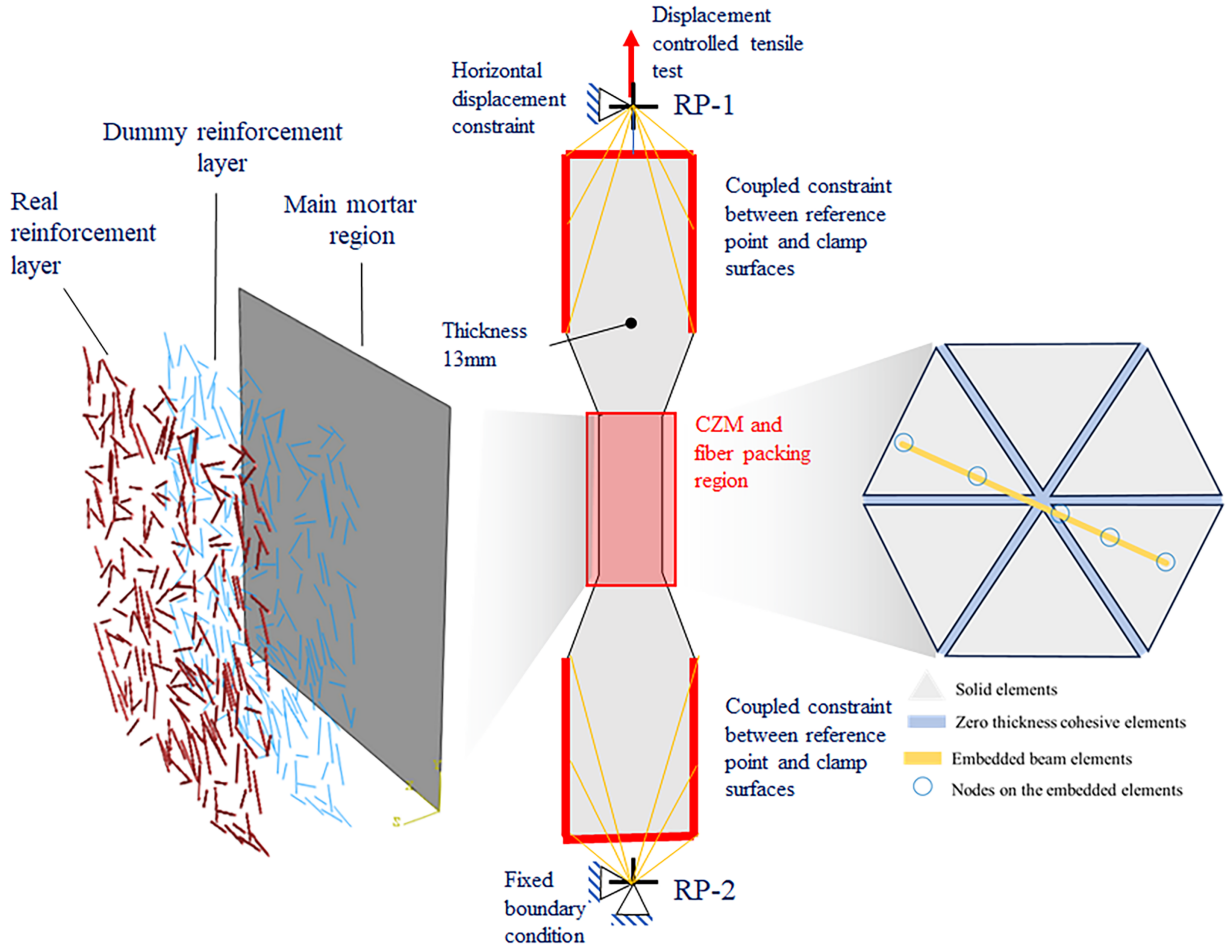


Figure 5: Schematic of uniaxial tension loading and boundary conditions applied to the SFRC specimen.

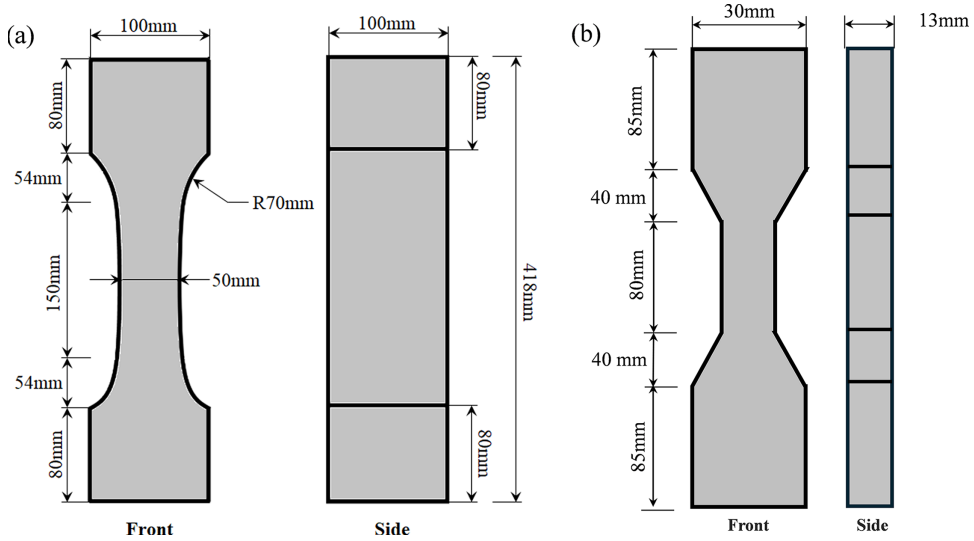


Figure 6: Dog-bone-shaped SFRC specimen geometry. (**a**) Based on the experimental work by Li and Deng [45], and (**b**) based on the experimental work by Chun and Yoo [46].

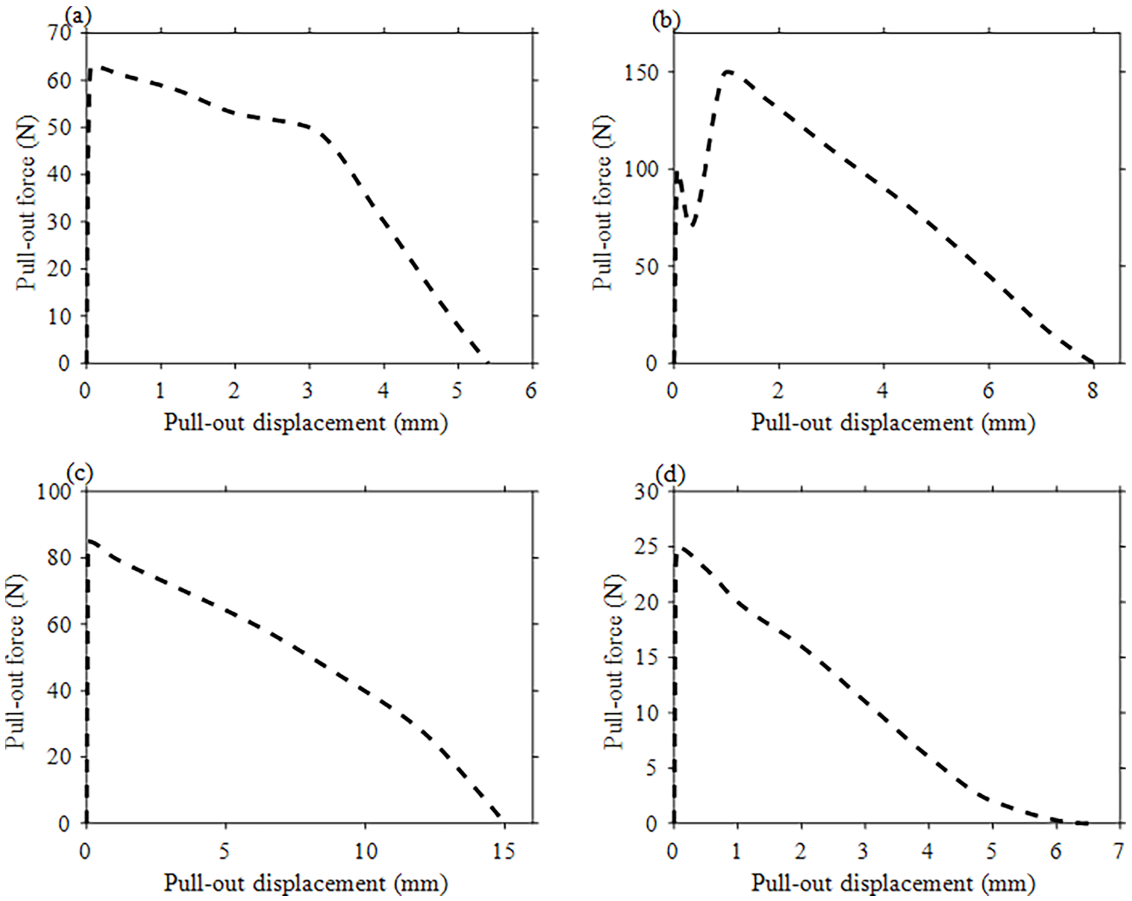


Figure 7: Fiber–matrix bond-slip laws for the dual-layer embedded beam formulation: (**a**) straight fiber for validation based on the study by Li and Deng [45]; (**b**) hooked-end fiber (H), double-peak from debonding plus hook anchorage; (**c**) macro fiber (S), frictional debonding over a long slip length; (**d**) micro fiber (ss), frictional debonding over a short slip length.

To validate the predictive capability of the mesoscale framework beyond stress-strain calibration, we compared simulated crack patterns against experimental results from Chun and Yoo [46] for tension. Four tensile configurations are considered: H20ss00 (2% hooked-end steel fibers), H10ss10 (1.0% hooked-end + 1.0% micro fibers), S20ss00 (2.0% straight macro fibers), and S10ss10 (1.0% straight macro + 1.0% micro fibers). The total fiber volume is constant at 2.0% for all configurations. Hybrid macro fibers (H) are hooked-end, while straight macro fibers (S) are straight steel. Micro fibers (ss) are straight steel of smaller dimensions. FE (Finite element) models employ 2D plane stress elements with a mesh size of 0.2 mm, with explicit fiber representation and a calibrated multi-stage interface model (cohesive zone for debonding, progressive wear-damage friction for post-debonding). Fiber details and material parameters are summarized in Table 1, with boundary geometry of the specimen shown in Fig. 6b. The fiber pullout curves are illustrated in Fig. 7c,d.

The cohesive zone parameters for the mortar matrix are summarized in Table 2 where “B” refers to model from Li and Deng [45] and “A” refers to the model based on the study by Chun and Yoo [46] and adopted from the experimental characterization of hybrid fiber-reinforced UHPC reported by Chun and Yoo [46]. These parameters define the fracture and softening response of the matrix phase within the mesoscale framework.

Table 1: Material properties of steel fibers for single fiber pull-out simulations [46].

Model	Type of Fiber	Fiber Diameter d_f (mm)	Length of Fiber l_f (mm)	Elastic Modulus E_f (GPa)	Tensile Strength f_t (MPa)
Tension	H	0.5	22	200	954
	S	0.3	30	210	2580
	ss	0.2	13	200	2788

Table 2: Mortar matrix properties for CZM validation.

Parameter	Model	Unit	Mortar	Fiber	Cohesive Interface
Density, ρ		**kg/m^3**	**2400**	**7800**	**2400**
Young modulus, E	A	GPa	37.5	200	–
	B		37.5	200	–
Poisson's ratio, ν_{12}	A		0.19	0.3	–
	B		0.19	0.3	–
Elastic stiffness, K_n	A	(MPa/mm)	–	–	10^6
	B		–	–	10^6
Maximum normal stress t_n^0	A	(MPa)			12.5
	B		–	–	4.0
Maximum shear stress t_s^0	A	(MPa)			52.5
	B		–	–	15
Normal mode fracture energy, G_n^C	A	(N/mm)			1.2
	B		–	–	0.6
Shear mode fracture energy, G_s^C	A	(N/mm)			7.25
	B		–	–	1.2
B-K criterion, η	A		–	–	1.2
	B		–	–	1.2

Model validation is performed against the uniaxial tensile experiments [45], as shown in Fig. 8. Three representative mix designs are considered: plain concrete (0% fibers), moderate fiber content (1.3% steel fiber volume fraction), and high fiber content (1.8% steel fiber volume fraction). As reported by Li and Deng [45], the plain matrix exhibits a brittle response characterized by abrupt post-peak softening. At 1.3% fiber content, the response transitions to quasi-ductile behavior with progressive softening and a dominant crack accompanied by secondary branching. At 1.8% fibers content, the material exhibits strain-hardening behavior, with multiple load drops associated with successive fiber pullout events, accompanied by distributed microcracking and multiple major crack paths. The numerical predictions show good agreement with the experimental stress–strain response and crack evolution patterns reported by Li and Deng [45].

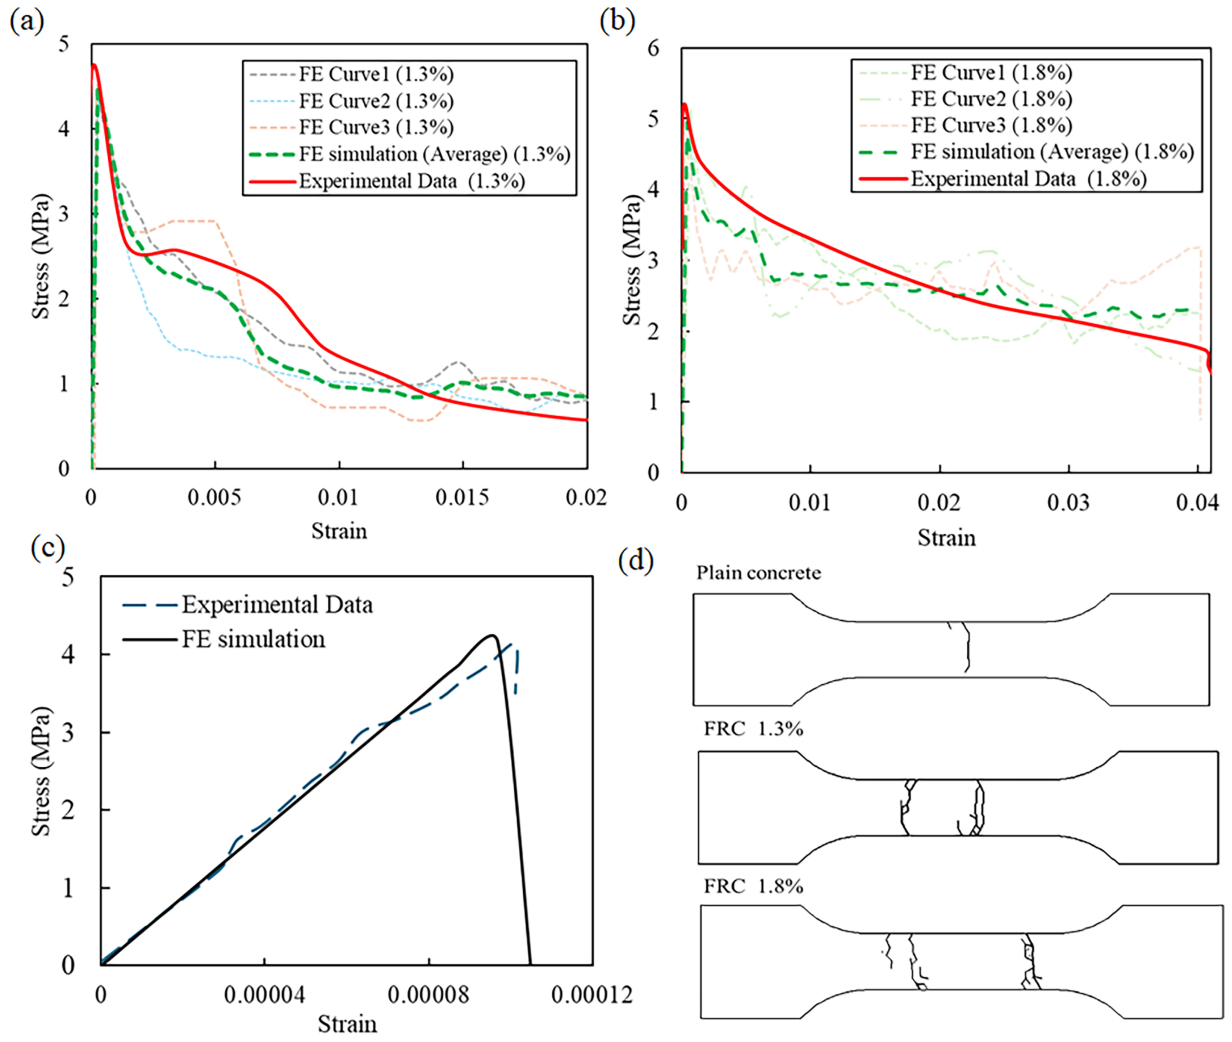


Figure 8: Validation of the meso-scale CZM model under uniaxial tension against experimental results [45]. Left: stress-strain curves for plain concrete (0% fibers), 1.3% steel fiber content, and 1.8% steel fiber content. Right: corresponding crack patterns at failure showing transition from a single dominant crack (plain concrete) to distributed micro-cracking and multiple major cracks (1.8% fibers).

Fig. 9 presents the failure patterns of two representative hybrid fiber configurations: H20ss00 (2% hooked-end steel fibers) and H10ss10 (1% hooked-end steel fibers with 1% micro steel fibers) under uniaxial tension. The H20ss00 configuration exhibits distributed cracking with extensive fiber bridging, while the H10ss10 configuration shows a more localized failure pattern with fewer but wider cracks, where hooked-end fibers provide enhanced mechanical anchorage. These experimental observations are well captured by our mesoscale simulations, confirming the model's ability to reproduce distinct failure mechanisms associated with different fiber geometries. Fig. 10 compares the tensile stress-strain curves from numerical simulations against experimental data for both H-series and S-series configurations. The numerical predictions show excellent agreement with experimental results across all configurations, accurately capturing the elastic modulus, peak tensile strength, and post-peak softening behavior. The close agreement validates the cohesive-zone model parameters, bond-slip constitutive law, and explicit fiber representation implemented in our mesoscale modeling framework.

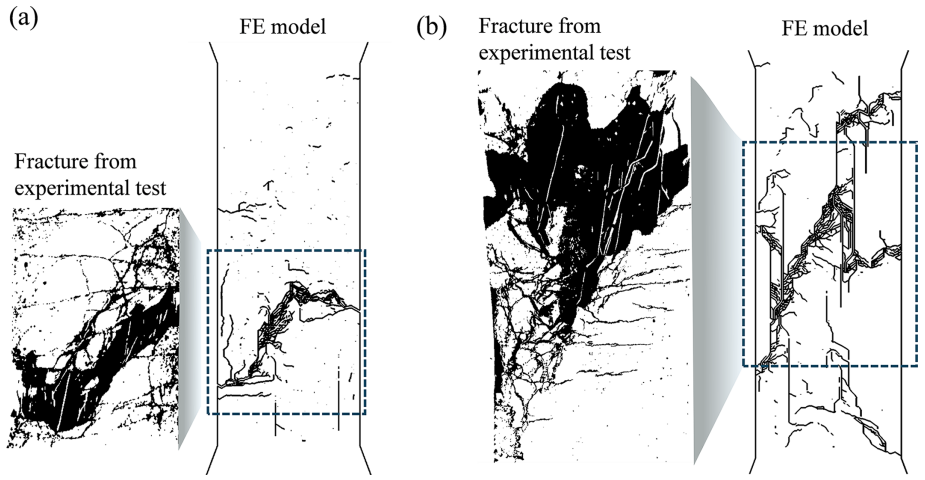


Figure 9: Failure patterns of H20ss00 and H10ss10 hybrid fiber configurations under tension, showing distributed cracking and fiber bridging mechanisms (experimental pattern reported by Chun and Yoo [46]). (**a**) H10ss10 failure mode, and (**b**) H20ss00 failure mode.

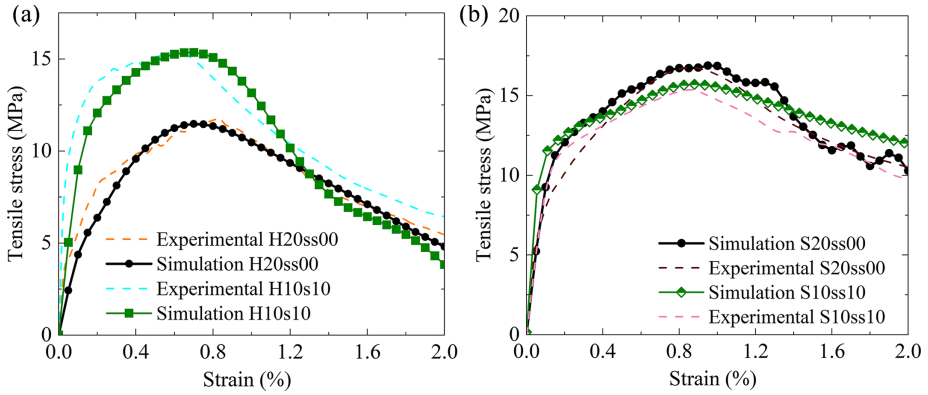


Figure 10: Tensile stress–strain curves for (**a**) H-series (H20ss00, H10ss10) and (**b**) S-series (S20ss00, S10ss10) fiber configurations, comparing simulation results with experimental data reported by Chun and Yoo [46].

The simulation is extended to a parametric study based on the validated specimen geometry. Each configuration in the dataset is systematically encoded using a structured naming convention of the form <F>Ma<M>Mi<m>, where <F> denotes the fiber volume fraction multiplied by 10, <M> is the number of macro-fibers (length 30 mm, diameter 0.3 mm), and <m>, is the number of micro-fibers (length 13 mm, diameter 0.2 mm). For instance, 20Ma150Mi500 represents a specimen with 2.0% total fibers, 220 macro-fibers, and 500 micro-fibers, while 00Ma00Mi00 represents plain concrete. This structured encoding enables automated extraction of mix-design descriptors via a parsing function, that maps each specimen to a normalized feature vector comprising fiber volume fraction (fiber fraction (f/3.0)), macro-fiber count (n_{macro}/300), and micro-fiber count (n_{micro}/1500). This formulation ensures consistent dataset parameterization and facilitates direct integration with the machine learning pipeline.

3 Proposed Morphology-Aware CNN Framework

The proposed framework reformulates the ill-posed curve-to-crack field mapping into two sequential learning tasks. Stage 1 (StatNet) predicts a compact eleven-dimensional vector of crack morphology from the stress–strain response and mix-design parameters. Stage 2 (PatternDecoder) generates a high-resolution

crack probability field conditioned on the predicted morphology descriptors and a latent curve embedding from Stage 1. The overall architecture is illustrated in Fig. 11.

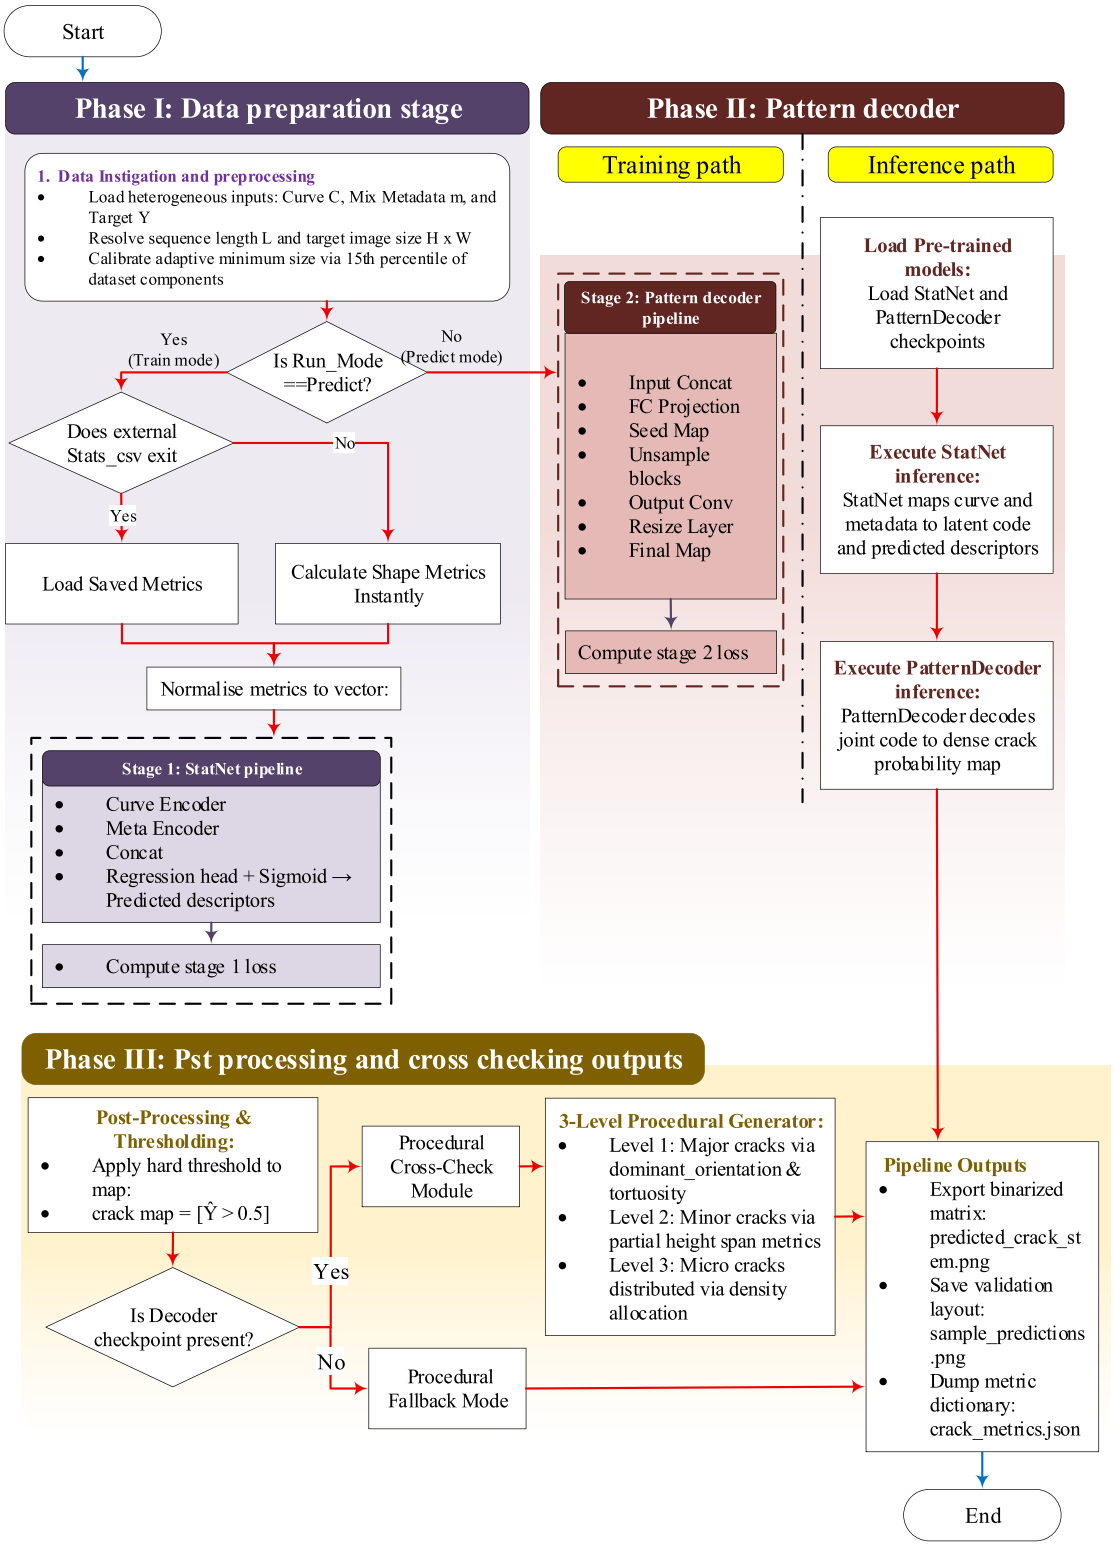


Figure 11: Two-stage CNN framework for SFRC crack prediction. StatNet regresses morphology descriptors; PatternDecoder synthesizes crack maps.

The morphology vector, summarized in Table 3. Comprises major/minor/micro crack counts, mean spacing and standard deviation, tortuosity, branching factor, dominant orientation, orientation entropy, and average crack width. For each specimen, descriptors are either extracted from precomputed statistics or computed from binary crack maps using connected-component labeling, morphological skeletonization, and distance transform. Crack components are classified by normalized size as major (spanning ≥50% of specimen height), minor (10%–50%), or micro (<10%) after applying a minimum size filter. Orientation is mapped from [−90°, 90°] to [0, 1], and orientation entropy is evaluated via Shannon entropy of a 12-bin angular histogram to quantify alignment dispersion. All descriptors are normalized by predefined scaling factors and clipped to [0, 1.5], providing robust regression targets for Stage 1 and conditioning inputs for Stage 2. Each descriptor is normalized by its corresponding scaling constant and clipped to the range [0, 1.5]. The upper bound introduces a controlled tolerance for stochastic excursions beyond the nominal range. The resulting normalized morphology vector is used as the supervised regression target for Stage 1 and as the conditioning input for Stage 2.

Table 3: Eleven-component crack morphology descriptor vector for stage 1 regression.

Index	Descriptor	Symbol	Meaning	Normalization
1	major_crack_count	N_{maj}	Through-going cracks (≥50% height)	8
2	minor_crack_count	N_{min}	Partial cracks (10 to 50 percent height)	45
3	micro_crack_density	M_{mic}	Micro-crack area density	0.018
4	avg_crack_spacing_px	μ_{sp}	Mean horizontal spacing of major cracks	256
5	crack_tortuosity	τ	Path length divided by straight-line length	1.5
6	crack_branching_factor	B	Bifurcations per unit crack length	0.15
7	dominant_orientation	θ	Principal crack angle mapped to (0 to 1)	180
8	avg_crack_width_px	w_b	Mean crack width via distance transform	8
9	std_crack_spacing_px	σ_{sp}	Standard deviation of major crack spacing	128
10	crack_irregularity	I_c	Measure of crack path roughness/deviation from smooth curve	0.5
11	crack_straightness	S_c	Measure of how straight the crack path is (inverse of tortuosity)	1.0

3.1 Stage 1: Morphology Regression Network (StatNet)

StatNet is a dual-encoder regression network designed to process heterogeneous inputs: a stress–strain curve and a low-dimensional mix-design vector. Separate encoders are employed for each modality, and their latent representations are fused prior to regression, ensuring modality-specific feature learning without parameter sharing.

The curve encoder treats the stress-strain response as a two-channel, one-dimensional signal of length L (the number of load steps, typically 100). The encoder consists of four strided convolutional blocks, each comprising a 1D convolution followed by batch normalization and a GELU activation. The first block uses a kernel size of 7 and a stride 2, increasing the channel dimension from 2 to 16. The second block uses a kernel size of 5 and a stride of 2, increasing channels counts from 16 to 32. The third block uses a kernel size of 5 and a stride of 2, increasing channels count from 32 to 64. The fourth block uses a kernel size of 3 and a stride of 2, maintaining 64 channels while further reducing the temporal resolution. After these four blocks, the temporal dimension has been reduced by a factor of 16 relative to the input length. A global average pooling layer collapses the remaining sequence to a single vector of length 64 [47,48]. This architecture

follows established 1D convolutional encoding practices for sequential data. This pooled representation is then passed through a fully connected layer with dropout (rate 0.3) and a ReLU (Rectified Linear Unit) activation, producing the final curve latent $z_c \in \mathbb{R}^{64}$. The curve latent is designed to capture the salient features of the mechanical response (such as peak stress, post-peak softening rate, and the number of load drops) that are most informative for predicting crack morphology.

The mix-design vector $m \in \mathbb{R}^3$ (fiber fraction, macro count, micro count) is encoded using a lightweight multilayer perceptron (MLP) with two hidden layers ($3 \to 32 \to 32 \to 64$), each with ReLU activation and dropout (0.2), yielding the metadata embedding $z_m \in \mathbb{R}^{64}$. The architecture reflects the smooth, low-dimensional nature of the input space. The two embeddings are concatenated into a 128-dimensional feature vector, which passed to a three-layer regression head ($128 \to 64 \to 32 \to 10$). ReLU activations and dropout (0.3) are applied in the hidden layers, and a sigmoid output constrains predictions to $[0, 1]$, consistent with normalized descriptor targets [25,49].

StatNet is trained using a weighted smooth-L1 (Huber) loss over the eleven descriptor components. The smooth-L1 loss is defined as

$$\mathrm{SmoothL1}_\beta(x, y) = \begin{cases} \frac{1}{2}(x - y)^2/\beta, & |x - y| < \beta \\ |x - y| - \beta/2, & \text{otherwise} \end{cases} \tag{7}$$

with the transition parameter set to $\beta = 0.05$. This loss behaves quadratically for small errors, encouraging precise convergence near the optimum, while transitioning to a linear penalty for larger errors, making it robust to occasional outlier descriptor may arise from particularly stochastic crack patterns.

The total Stage 1 loss is

$$\mathcal{L}_{\mathrm{stat}} = \frac{1}{11} \sum_{k-1}^{11} w_k \cdot \mathrm{SmoothL1}_\beta \left(s_k, s_k\right) \tag{8}$$

where w_k are component-specific weights that emphasize major crack count, minor crack count, and dominant orientation due to their dominant influence on downstream crack pattern synthesis. The weight vector is $w = (1.5, 1.2, 1.0, 0.8, 1.0, 0.9, 1.1, 0.7, 0.8, 0.7, 1.0)$, aligned with the descriptor order in Table 3.

3.2 Stage 2: Crack Pattern Decoder (PatternDecoder)

The PatternDecoder synthesizes a dense crack probability map conditioned on two inputs: the curve latent z_c (produced by the frozen StatNet encoder) and the predicted morphology vector $\hat{s}$. The decoder is a conditional generative model that learns the distribution of crack patterns consistent with a given mechanical response and morphology descriptor set.

The conditioning vectors $z_c \in \mathbb{R}^{64}$ and $\hat{s} \in \mathbb{R}^{11}$ are concatenated to form a 74-dimensional code. This code is passed through a fully connected projection stack. The first linear layer maps from 74 to 256 dimensions with ReLU activation. The second linear layer maps from 256 to $32 \times (H/16) \times (W/16)$dimensions, again with ReLU activation. The output of this second layer is reshaped into a low-resolution feature seed of shape $32 \times (H/16) \times (W/16)$, where H and W are the height and width of the target crack image (typically 256×512 pixels). This seed represents a coarse spatial summary of the crack pattern, encoded in 32 feature channels at a resolution sixteen times smaller than the final output.

The seed is then progressively upsampled through four transposed convolution blocks. Each block consists of a transposed 2D convolution with kernel size 4, stride 2, and padding 1, followed by batch normalization and a GELU activation. The first block increases the spatial resolution by a factor of 2 (to $H/8 \times W/8$)

while increasing the channel count from 32 to 64. The second block upsamples to $H/4 \times W/4$ with 32 channels. The third block upsamples to $H/2 \times W/2$ with 16 channels. The fourth block upsamples to $H \times W$ with 8 channels. Following the upsampling blocks, a final 3×3 convolution maps the 8-channel feature map to a single channel. Bilinear interpolation is applied to ensure exact alignment with the target resolution, and a sigmoid activation produces per-pixel crack probabilities in the range $[0, 1]$ [28,50]. This decoding strategy with progressive upsampling and sigmoid output is standard in conditional generative models for image synthesis [28,50]. Recent advancements in diffusion-based generative models have further demonstrated the ability to synthesize complex microstructural features [30].

The crack patterns are highly imbalanced. That is, typically fewer than 10% of pixels belong to cracks, with the vast majority representing intact matrix. A naive binary cross-entropy loss would therefore drive the network to predict a uniformly zero (no crack) output, which achieves low loss but is physically meaningless. To address this, the decoder is trained with a combined objective that balances per-pixel classification accuracy against region overlap [51].

The total Stage 2 loss is

$$\mathcal{L}_{\text{pat}} = \lambda \cdot \text{BCE}(\hat{Y}, Y) + (1-\lambda) \cdot \left(1 - \frac{2\sum_{ij} \hat{Y}_{ij} Y_{ij} + 1}{\sum_{ij} \hat{Y}_{ij} + \sum_{ij} Y_{ij} + 1}\right) \tag{9}$$

where $\lambda = 0.5$ weights the two terms equally. The first term is the standard binary cross-entropy, which evaluates classification accuracy at the pixel level. The second term is the soft Dice loss, where the numerator represents twice the intersection (plus a smoothing term of 1) and the denominator represents the sum of the predicted and ground-truth crack pixel counts (plus the same smoothing term). The soft Dice loss is particularly effective for imbalanced segmentation tasks because it directly optimizes the overlap between predicted and true crack regions, irrespective of the absolute number of crack pixels.

Algorithm 1 summarizes the complete two-stage pipeline. StatNet predicts morphology descriptors from the mechanical response, while PatternDecoder synthesizes the corresponding crack map. The frozen curve latent enables efficient offline precomputation and rapid inference.

Algorithm 1: Morphology-aware crack-pattern prediction

```
Require: training set {(C_i, m_i, s_i, Y_i)}; epochs E1, E2
// - - - - Stage 1: train StatNet f_θ - - - -
for epoch = 1 . . . E1 do
    for each minibatch (C, m, s) do
        z_c ← CurveEncoder(C); z_m ← MetaEncoder(m)
        ŝ ← Sigmoid (Head ([z_c, z_m]))
        L_stat ← (1/10) Σ_k w_k SmoothL1(ŝ_k, s_k)
        θ ← AdamW update on ∇L_stat (clip, AMP)
    end for
end for; freeze θ
// - - - - precompute & cache curve latents - - - -
for each sample i do z_cˆ(i) ← CurveEncoder(C_i) end for
// - - - - Stage 2: train PatternDecoder g_φ - - - -
for epoch = 1 . . . E2 do
    for each minibatch (z_c, s, Y) do
        Ŷ ← Sigmoid (Decoder ([z_c, s]))
```

(Continued)

Algorithm 1 (continued)

L_pat ← λ·BCE($\hat{Y}$, Y) + (1−λ)·Dice($\hat{Y}$,Y)
φ ← AdamW update on ∇L_pat (clip, AMP)
end for
end for
// - - - - Inference for a new specimen (C, m) - - - -
z_c ← CurveEncoder(C); $\hat{s}$ ← f_θ(C, m)
$\hat{Y}$ ← g_φ(z_c, $\hat{s}$)
crack map ← [$\hat{Y}$ > 0.5] (else procedural fallback from $\hat{s}$)
return $\hat{s}$, crack map

3.3 Dataset Split and Validation Protocol

The full dataset comprises 220 meso-scale specimens spanning the fiber configuration space described in Section 2.4 (plain concrete and hooked-end/straight macro- and micro-fiber combinations at varying volume fractions). Model performance is evaluated using stratified 5-fold cross-validation rather than a single fixed hold-out split, to obtain uncertainty estimates on all reported metrics given the moderate dataset size. In each fold, 176 specimens (80%) are used for training and 44 specimens (20%) are held out for validation, with folds rotated so that every specimen serves in the validation set exactly once. Stratification is performed jointly on fiber configuration type (plain, hooked-end, straight) and fiber volume fraction bin, ensuring each fold preserves the overall distribution of these variables and preventing any fold from being dominated by a single configuration family. To prevent data leakage, the train/validation split is performed at the specimen level prior to any feature computation, and all normalization statistics (e.g., feature scaling factors, morphology descriptor normalization constants) are fit exclusively on the training partition of each fold and subsequently applied to the corresponding validation partition. All performance metrics reported in Section 4 are aggregated across the five validation folds rather than computed on a single train/test split.

4 Results and Discussion

This section presents a quantitative and qualitative evaluation of the proposed morphology-aware framework for predicting crack pattern in SFRC. Convergence behavior and generalization capacity are first examined through training and validation loss trajectories to identify potential overfitting. Subsequently, representative predictive crack fields are compared against numerical simulations to evaluate qualitative accuracy and consistency. The predictive capability for individual morphology descriptors is analyzed through scatter plots, residual distributions, and error histograms. Finally, global performance metrics are reported to validate the central hypothesis that multi-scale damage features can be reliably inferred from macroscopic mechanical responses, followed by a discussion of the framework's limitations and applicability.

4.1 Training Convergence, Loss Analysis, and Computational Efficiency

Fig. 12 presents the training and validation loss curves for both stages of the framework. Stage 1 (StatNet) achieves rapid convergence within the first 30 epochs, with the weighted smooth-L1 loss decreasing from approximately 0.063 to below 0.015 by epoch 200. The validation loss closely tracks the training loss throughout, reaching a minimum of 0.010670, confirming minimal overfitting. Stage 2 converges more slowly, requiring approximately 400 epochs to reach a stable combined BCE-Dice loss. The best validation loss of 0.01525 is achieved at epoch 400, after which mild overfitting is observed. The close alignment between

training and validation trajectories across both stages validates the effectiveness of the sequential training protocol and latent precomputation strategy.

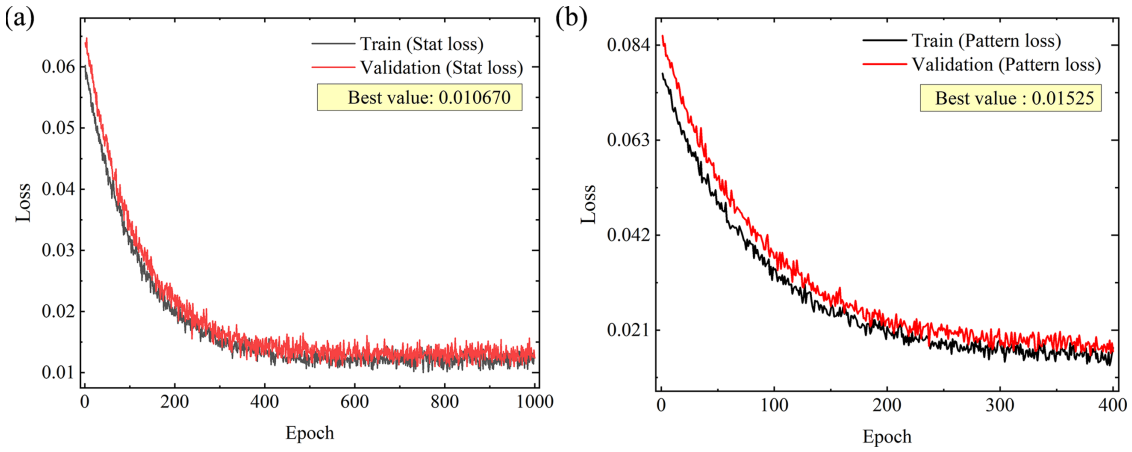


Figure 12: Training and validation loss curves. Stage 1 (StatNet) achieves the best validation loss of 0.01067; Stage 2 (PatternDecoder) reaches 0.01525 at epoch 400.

All training hyperparameters are summarized in Table 4 to ensure full reproducibility of the results. For Stage 1, the AdamW optimizer is used with an initial learning rate of 1×10^{-3} and weight decay of 1×10^{-4}. For Stage 2, a lower initial learning rate of 5×10^{-4} with weight decay of 1×10^{-5} is employed. Both stages use a batch size of 16 and are trained with a cosine annealing learning rate schedule with warmup over the first 10 epochs. Gradient clipping at a maximum norm of 1.0 is applied to stabilize training. Early stopping with a patience of 20 epochs on validation loss is used to prevent overfitting, with the best models obtained at epochs 200 and 400 for StatNet and PatternDecoder, respectively, though StatNet was trained for up to 1000 epochs for final model selection. All experiments were performed with a fixed random seed of 42 for reproducibility.

Table 4: Training hyperparameters for both stages.

Hyperparameter	Stage 1 (StatNet)	Stage 2 (PatternDecoder)
Optimizer	AdamW	AdamW
Initial learning rate	1×10^{-3}	5×10^{-4}
Weight decay	1×10^{-4}	1×10^{-5}
Batch size	16	16
Maximum epochs	1000	400
Best epoch	200	400
Learning rate schedule	Cosine annealing with 10-epoch warmup	Cosine annealing with 10-epoch warmup
Gradient clipping	1	1
Early stopping patience	20 epochs	20 epochs
Random seed	42	42

Before proceeding to the detailed prediction quality assessment, we first establish the computational efficiency of the proposed framework, as this constitutes one of the primary motivations for developing a surrogate modeling approach. The computational cost comparison between the proposed framework and

Abaqus simulations is summarized in Table 5, which reports the wall-clock times for both the trained neural network inference and the full mesoscale Abaqus simulations.

Table 5: Computational cost comparison between proposed framework and Abaqus simulation.

Component	Time	Hardware
StatNet inference (per sample)	6.23 ms	NVIDIA RTX 3090 GPU
PatternDecoder inference (per sample)	8.34 ms	NVIDIA RTX 3090 GPU
Total framework inference (per sample)	14.57 ms	NVIDIA RTX 3090 GPU
Abaqus simulation (per specimen)	2.5 h (9000 s)	Intel Xeon E5-2680 v4 @ 2.4 GHz (16 cores)

The proposed framework achieves a per-sample inference time of 14.57 ms, enabling real-time crack pattern assessment for practical engineering applications. In contrast, a single Abaqus simulation requires approximately 2.5 h, making parametric studies with hundreds of specimens computationally prohibitive. This analysis quantitatively validates the claim that the morphology-aware surrogate model bridges the gap between expensive high-fidelity simulations and practical engineering applications requiring rapid crack pattern assessment.

4.2 Crack Pattern Prediction Quality

Fig. 13 compares predicted crack patterns against ground-truth FEM simulations across six representative specimens spanning a range of fiber/matrix configurations, from sparse single- or double-crack morphologies (e.g., 19Ma99Mi69, 1Ma10Mi9, 1Ma18Mi7) to increasingly complex, densely branched crack networks (26Ma125Mi73, 17Ma4Mi343, 14Ma4Mi88). For the simpler cases, the PatternDecoder correctly reproduces the dominant crack trajectory and overall orientation, with only minor positional offsets visible in the difference maps. As fiber content and crack density increase, the ground-truth patterns exhibit multiple co-existing major cracks with fine branching; the PatternDecoder captures the overall branching topology and spatial distribution of cracks, though the difference maps (Pred–FEM) reveal localized discrepancies of up to ±1.0 in normalized position, concentrated where secondary branches diverge in path or where crack density is highest. These residuals are most pronounced in the two most heavily branched specimens (17Ma4Mi343 and 14Ma4Mi88), reflecting the increased difficulty of resolving fine-scale branch placement as morphological complexity grows, while the dominant crack trunks remain well-aligned across all cases.

4.3 Quantitative Morphology Prediction Performance

Figs. 14 and 15 present scatter plots validating predictive performance across all eleven morphology descriptors. The global prediction performance metrics demonstrate that the machine learning model successfully maps structural loading curves to multi-scale crack network morphology, achieving a high global normalized pooled coefficient of determination (R^2 = 0.940). Evaluating individual parameters in Fig. 14a,b reveals that macro-scale geometric properties and crack width characteristics attain excellent predictive scores, led by the base crack width (w_b), with a maximum R^2 = 0.961 and a Pearson correlation coefficient of r = 0.985. Out of the 11 morphological metrics evaluated, 10 fall within the high-accuracy 'Good' threshold ($R^2 > 0.70$), 1 falls within the 'Fair' tier ($0.40 < R^2 \leq 0.70$), and 0 parameters fall into the 'Poor' tier ($R^2 \leq 0.40$). This stratification indicates that macro-level fracture mechanisms are governed by robust deterministic constraints directly linked to the energy dissipation states of the load curves. The tighter alignment along the 1:1 identity line for w_b in Fig. 14c and the highly concentrated distribution in the

global pooled scatter plot in Fig. 14d confirm that scaling normalization effectively controls scale-dependent calibration errors, preventing the model from exhibiting systematic bias across diverse fracture dimensions.

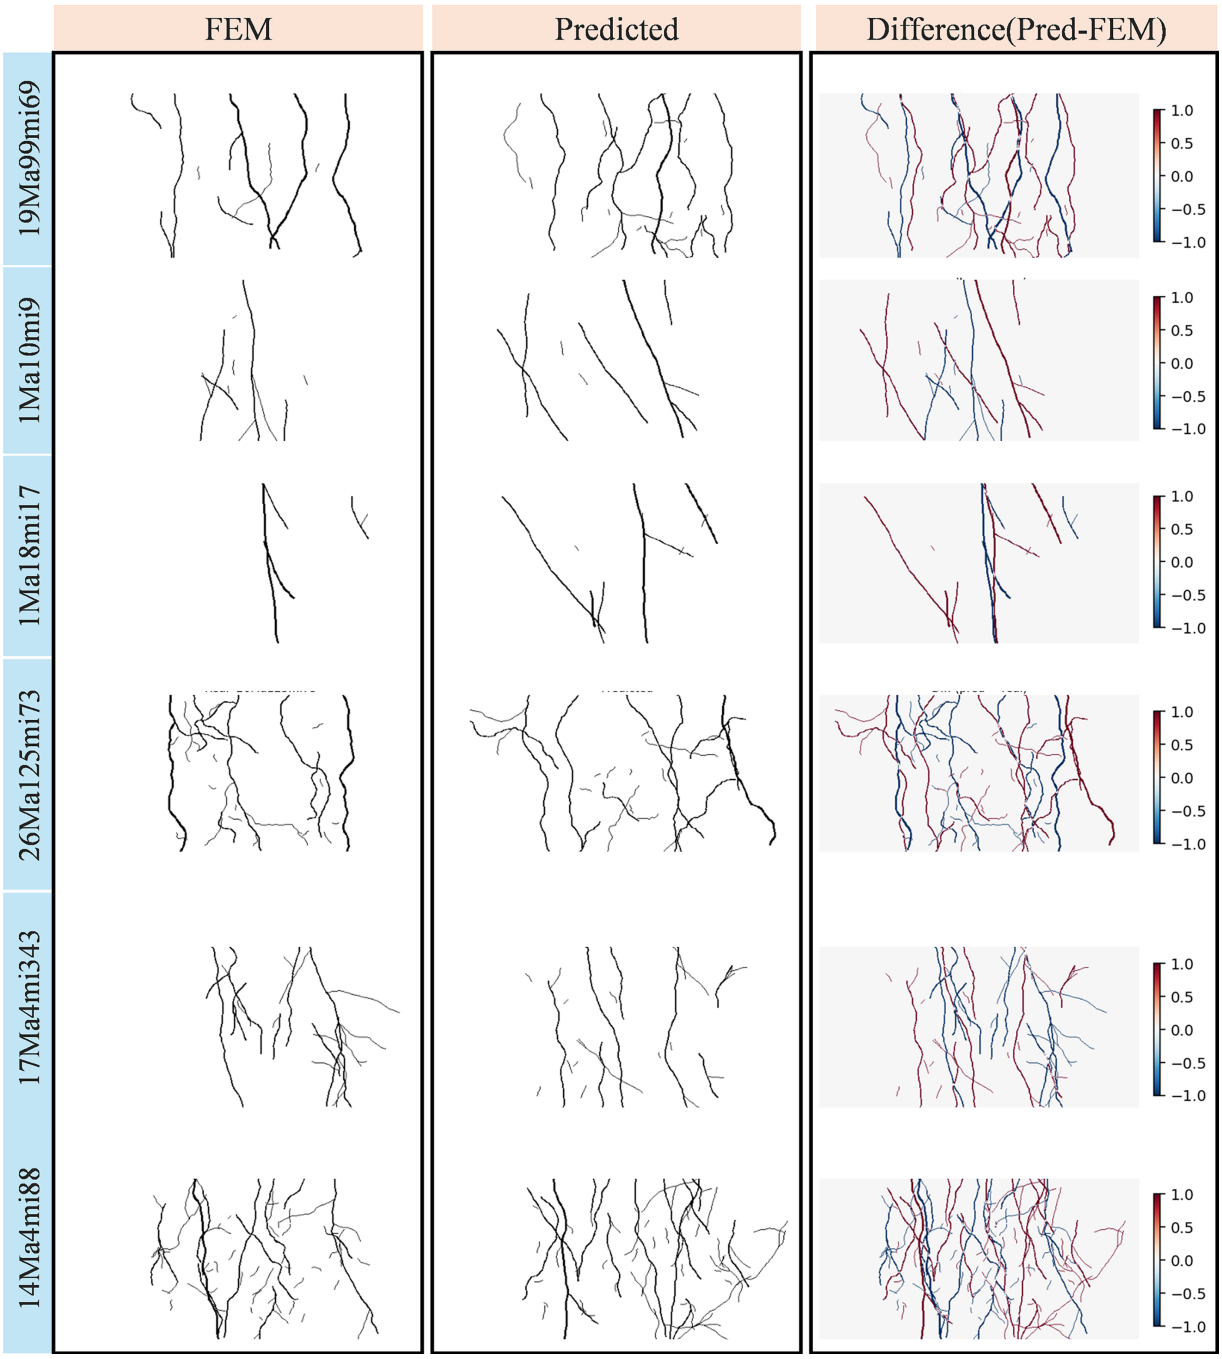


Figure 13: Qualitative crack pattern comparison across six specimens of increasing fiber content and morphological complexity: FEM ground truth (**left**), PatternDecoder prediction (**middle**), and pixel-wise difference map, Pred–FEM (**right**; colorbar range −1.0 to 1.0).

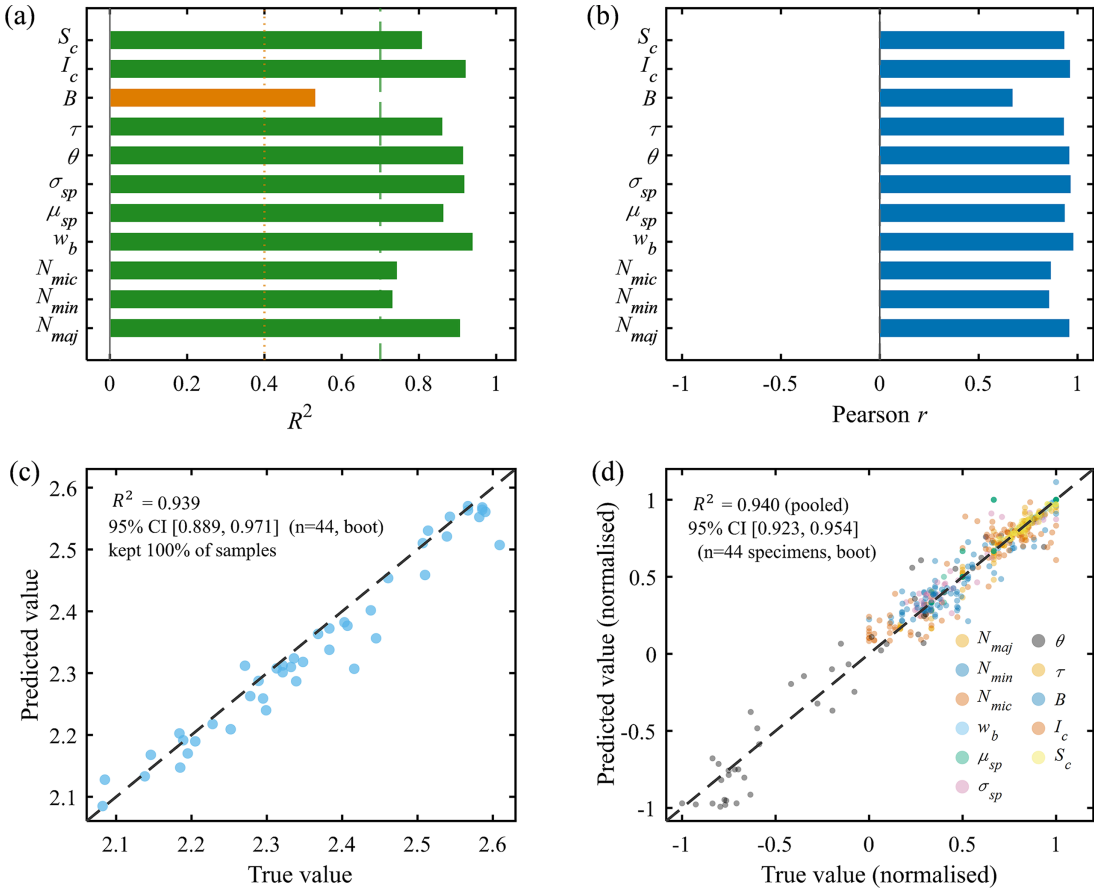


Figure 14: Prediction accuracy of the crack morphology model. (**a**) Coefficient of determination (R^2) categorized by performance tiers. (**b**) Pearson correlation coefficient (r) across parameters. (**c**) True vs. predicted scatter profiles for the top four metrics. (**d**) Global normalized true vs. predicted scatter plot pooling all variables.

A closer inspection of absolute and relative error scales across the target features demonstrates tight validation bounds and stable prediction metrics. As shown in the native-unit comparison in Fig. 15a, the mean absolute error (MAE) remains well-contained, reaching a minimum value of MAE = 0.004 for the branching factor (B), whereas the root mean squared error (RMSE) exhibits its highest value of RMSE = 19.15 for the average crack spacing parameter (S_c). The marginal gap between MAE and RMSE values across most features proves that the underlying regression model is highly robust against large outlier errors. Furthermore, the signed normalized error boxplots in Fig. 15b are symmetrically centered around the zero reference line, confirming that the network does not suffer from systemic over- or under-prediction when quantifying damage propagation. This tracking is reinforced by the cumulative distribution functions (CDFs) in Fig. 15c,d, which show the sharp, steep profiles indicating that over 90% of the validation samples fall safely within narrow acceptable operational tolerances, validating the system's compliance with precision limits.

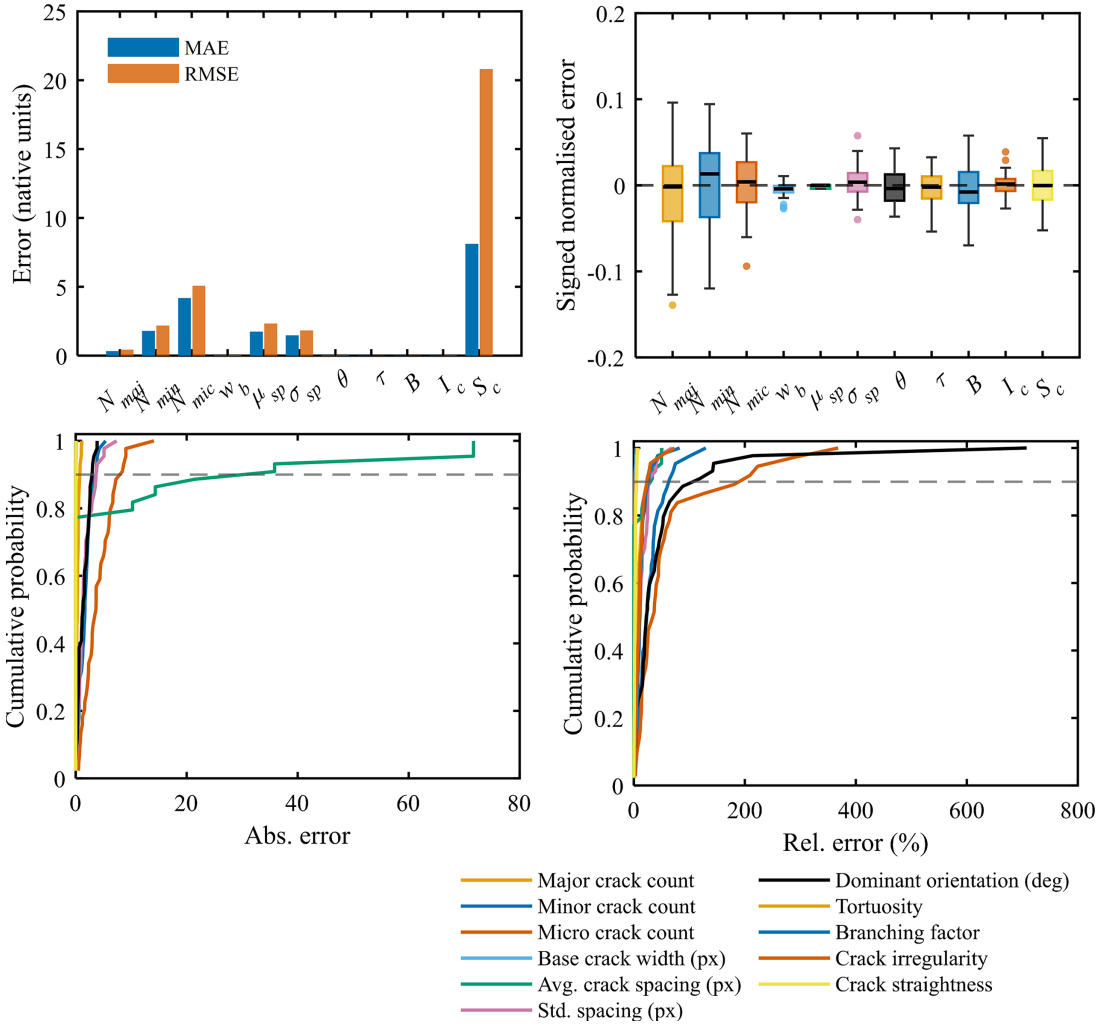


Figure 15: Error quantification and distribution boundaries. (**a**) Comparison of mean absolute error (MAE) and root mean squared error (RMSE) in native units. (**b**) Signed normalized error distributions showing parameter symmetry and outliers. (**c**) Cumulative distribution function (CDF) of absolute errors. (**d**) Cumulative relative error percentage CDF.

4.4 Statistical Homoscedasticity and Trend Stability

To evaluate the model's reliability under unfavorable conditions, residual diagnostic analyses were performed on the four highest-variance metrics isolated by their raw RMSE ranks: average crack spacing (μ), micro crack count (N_{mic}), minor crack count (N_{min}), and spatial standard deviation of spacing (σ_{sp}). Despite these parameters being exposed to localized stochasticity, Fig. 16a–d shows that the prediction residuals remain well contained within the shaded ±1 σ standard deviation bands across the entire spectrum of predicted magnitudes. The red moving-average loess trend lines remain nearly flat and closely aligned with the zero-residual baseline, showing no significant non-linear drifting or structural distortions. This widespread homoscedasticity confirms that the model's error variance is invariant to the absolute magnitude of the cracking features, ensuring uniform prediction accuracy from the initiation of fine micro-cracks up to fully localized macro-fracture bands.

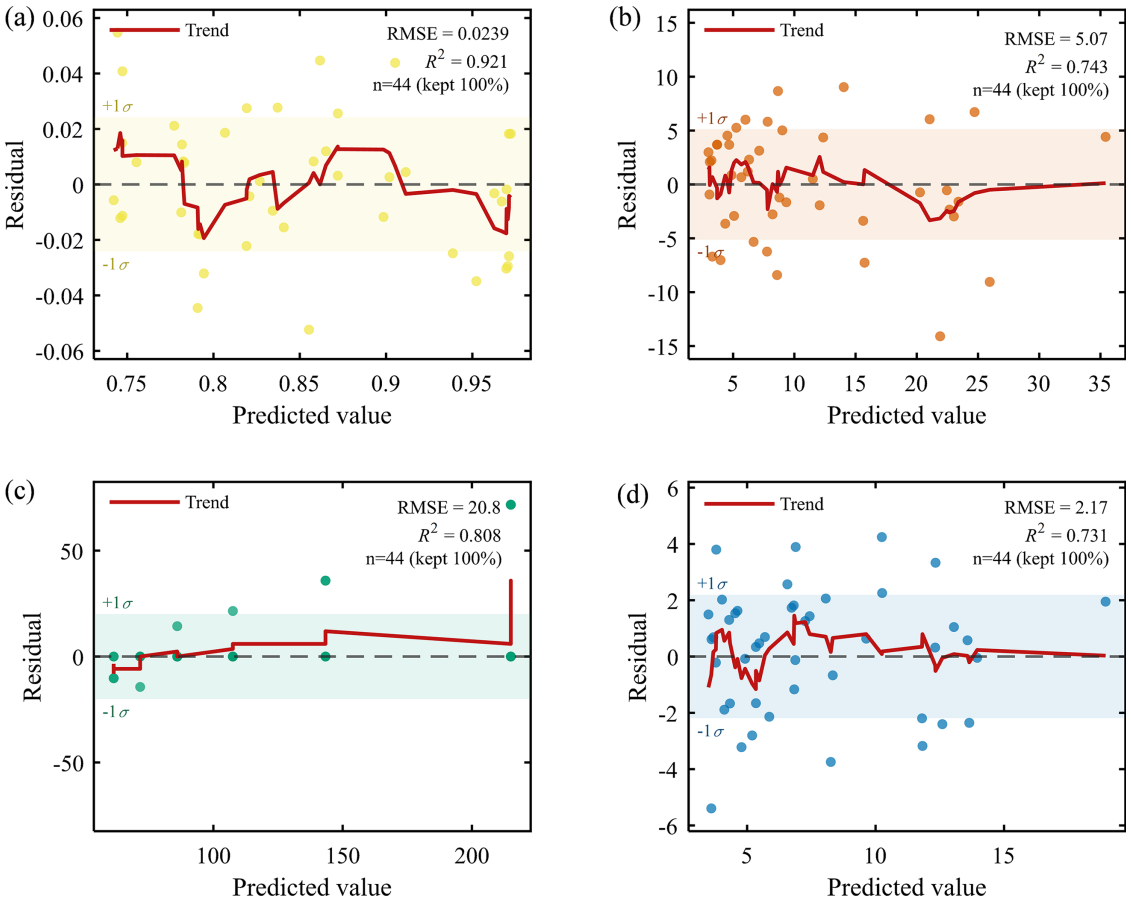


Figure 16: Residual diagnostics for the four parameters with the highest variance. (**a**) Average crack spacing (μ_{sp}). (**b**) Micro crack count (N_{mic}). (**c**) Minor crack count (N_{min}). (**d**) Standard deviation of crack spacing (σ_{sp}). Shaded areas represent $\pm 1\,\sigma$; red curves show the moving average.

The mathematical validity and generalizability of the machine learning architecture are verified through diagnostics of error distributions and feature redundancy checks. The Quantile-Quantile (Q-Q) plots presented in Fig. 17a–c display strong linear alignment between the empirical residual quantiles and the theoretical standard normal distribution line, with minor deviations occurring only at the extreme tails of the distributions. This adherence to normality confirms that the optimization objective successfully captured the deterministic patterns, leaving behind residuals consistent with predominantly random noise within the present dataset. Simultaneously, the cross-correlation matrix in Fig. 17d highlights the model's capacity to learn physical co-dependencies inherent to concrete fracture mechanics; most notably, a dominant positive cross-correlation is detected between average crack spacing μ_{sp} and its spatial standard deviation σ_{sp} (r = 0.97). The model successfully navigates these highly coupled multi-collinear pathways without performance degradation, confirming that it captures physically meaningful relationships within the simulation dataset governing spatial crack network distributions.

4.5 Mechanistic Interpretability and Cross-Compositional Generalization

The sensitivity and adaptability analyses demonstrate that the network captures physically consistent trends in the simulated fracture data rather than memorizing localized training boundaries. Tracing the gradient-based input sensitivity ($\partial input/\partial output$) across the loading path index in Fig. 18a reveals dynamic

shifts: sensitivity metrics remain low during the initial elastic tracking phase, but experience explicit spikes as loading approaches peak stress and transfers into the post-peak softening regime, capturing the physical transition from multi-scale micro-crack coalescence to dominant macro-crack widening. Crucially, as shown in Fig. 18b,c, the model maintains consistent R^2 and MAE distributions when evaluated across three distinct fiber compositions: hybrid (containing both macro- and micro-fibers), macro_only (containing only macro-fibers), and micro_only (containing only micro-fibers). This uniform behavior across varied fiber compositions suggests that the model maintains consistent performance across different fiber compositions within the training distribution and learns the underlying physical relationships governing crack morphology rather than overfitting to specific fiber arrangements. The smooth, anomaly-free profile of the consolidated per-sample normalized error heatmap in Fig. 18d underscores the global stability of the model, establishing it as a reliable tool for real-time crack morphology prediction.

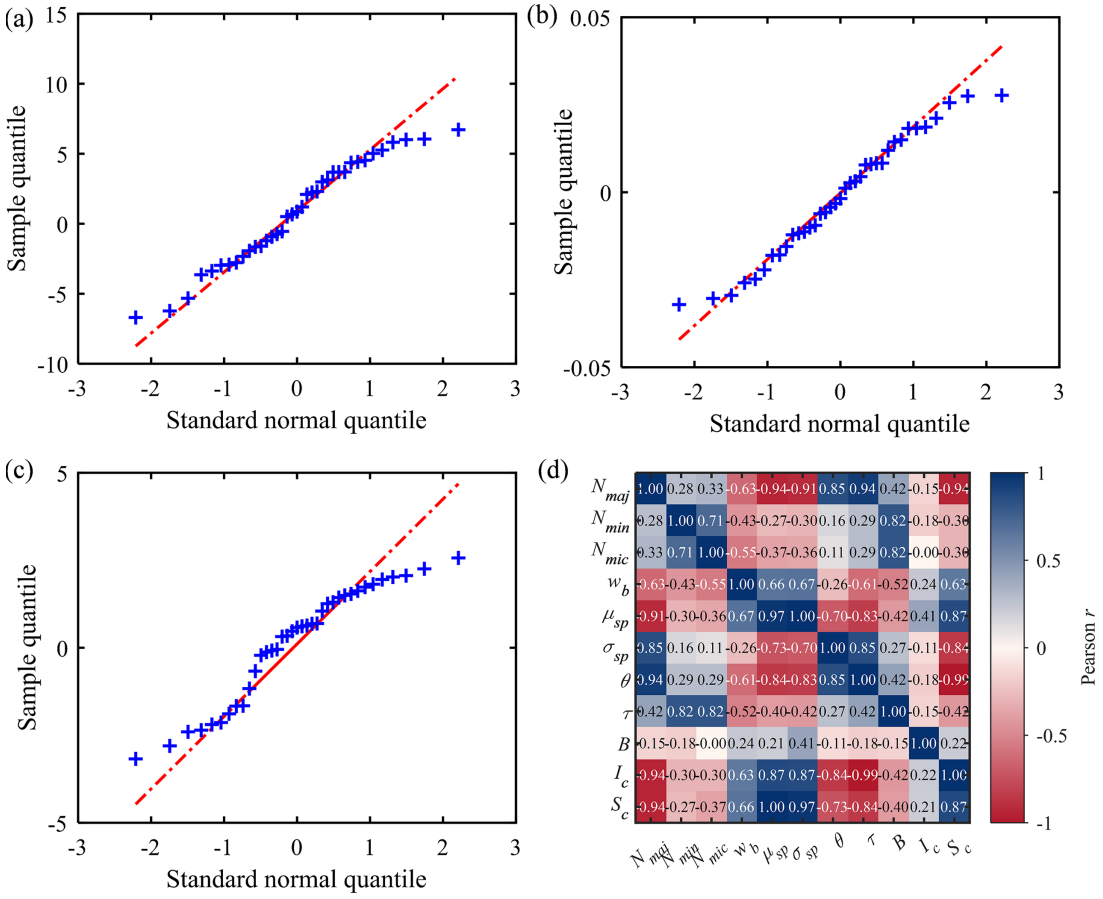


Figure 17: Residual normality compliance and parameter cross-correlations. (**a**–**c**) Quantile-Quantile (Q-Q) plots comparing empirical residual quantiles against a theoretical standard normal distribution for selected metrics. (**d**) Cross-correlation matrix showing Pearson *r* values among the 11 morphological variables.

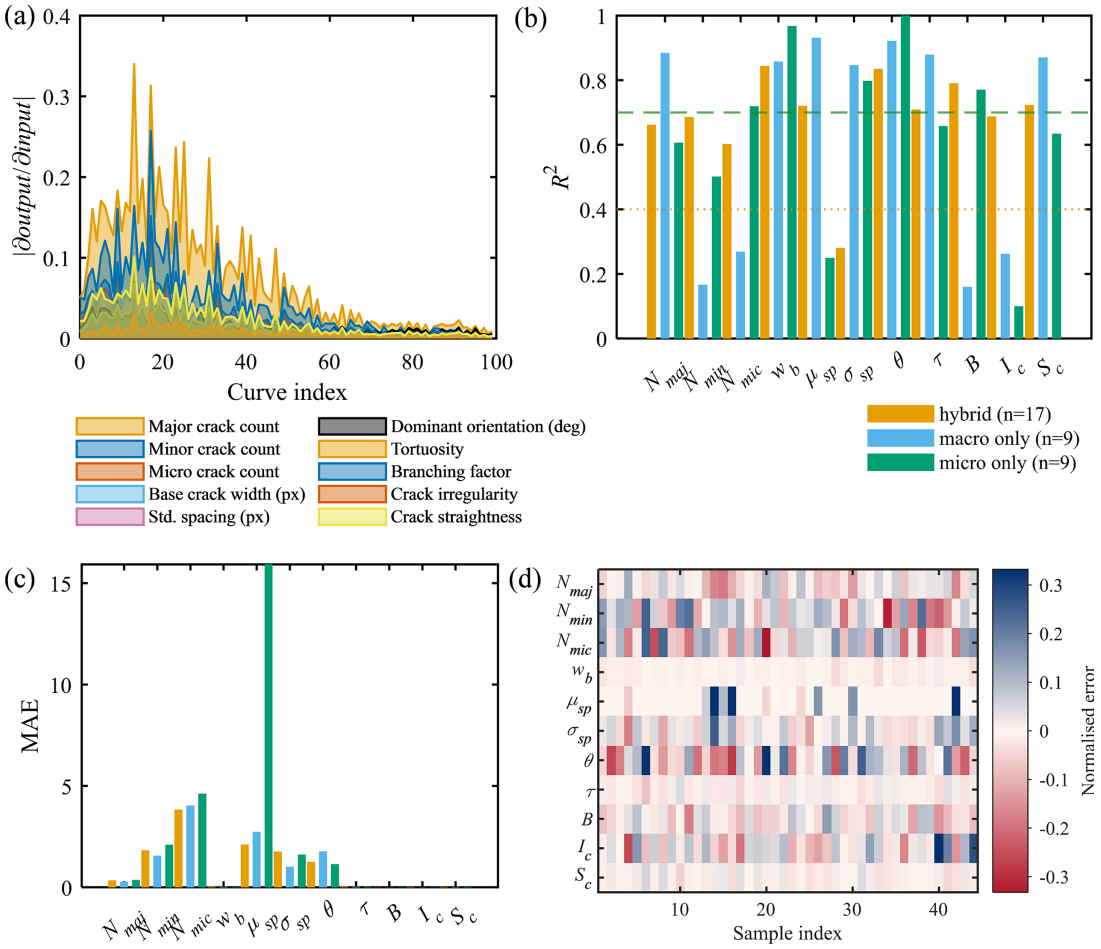


Figure 18: Model interpretability and performance generalization across specimen configurations. (**a**) Gradient-based input sensitivity ($\partial output/\partial input$) traced across the loading path. (**b**) R^2 Performance grouped by specimen geometry. (**c**) MAE variations across specimen types. (**d**) Consolidated per-sample normalized error heatmap.

4.6 Comparison with End-to-End Baselines

To validate the necessity of the proposed morphology-aware two-stage decomposition, we compare our framework against two representative end-to-end baselines that map stress-strain curves directly to crack probability maps without an intermediate morphological descriptor bottleneck: (i) a DirectCNN, an encoder-decoder architecture trained to regress crack images directly from input curves, and (ii) a CGAN, a conditional generative adversarial network trained with the same input-output pairing. Both baselines were trained and evaluated on the same dataset and specimen splits as the proposed method to ensure a fair comparison.

Fig. 19 presents representative crack morphology predictions for two specimens alongside the ground-truth patterns. The proposed method closely reproduces the branching topology, curvature, and continuity of the original crack networks. In contrast, DirectCNN outputs are dominated by dense, disconnected speckle-like artifacts that fail to capture coherent crack paths, consistent with the oversmoothing and detail-collapse behavior expected from direct regression under an ill-posed mapping. CGAN predictions are visually closer to plausible crack-like structures than DirectCNN, but the generated networks remain thinner, more

fragmented, and less faithful to the specific branching topology of the ground truth, reflecting the difficulty of conditioning adversarial generation on low-dimensional curve inputs alone.

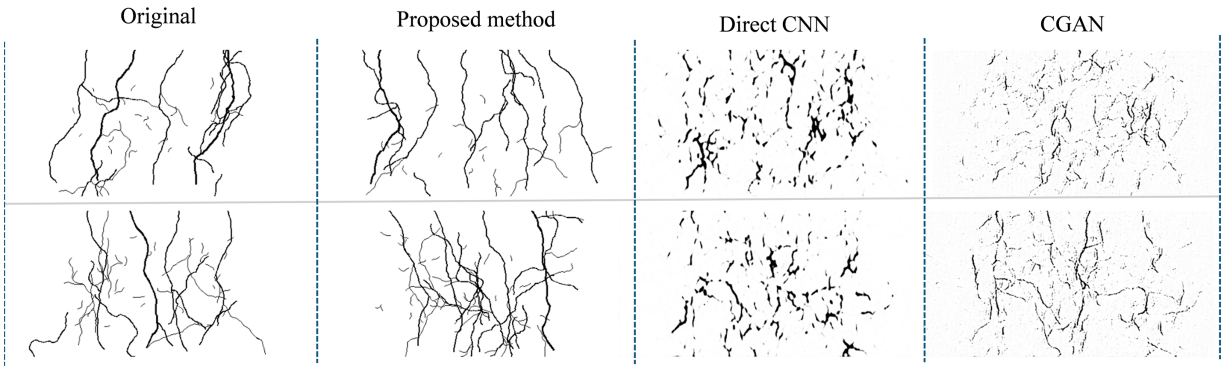


Figure 19: Qualitative comparison of predicted crack morphology for two representative specimens (rows). From left to right: ground-truth original crack pattern, proposed method, DirectCNN baseline, and CGAN baseline.

Fig. 20 and Table 6 summarize the computational cost of the three approaches. The proposed framework requires only 0.178×10^6 parameters, a 22.6-fold reduction relative to both DirectCNN and CGAN, which each require approximately 4.03×10^6 parameters (Fig. 20a). This compactness translates directly into training efficiency: the proposed method converges in 5.67 min, vs. 15.28 min for DirectCNN and 57.36 min for CGAN (Fig. 20b) the latter reflecting the added overhead of adversarial optimization. At inference, the proposed method also achieves the lowest per-sample latency (14.57 ms), compared to 24.27 ms for DirectCNN and 26.67 ms for CGAN (Fig. 20c), a meaningful advantage for high-throughput or design-loop applications.

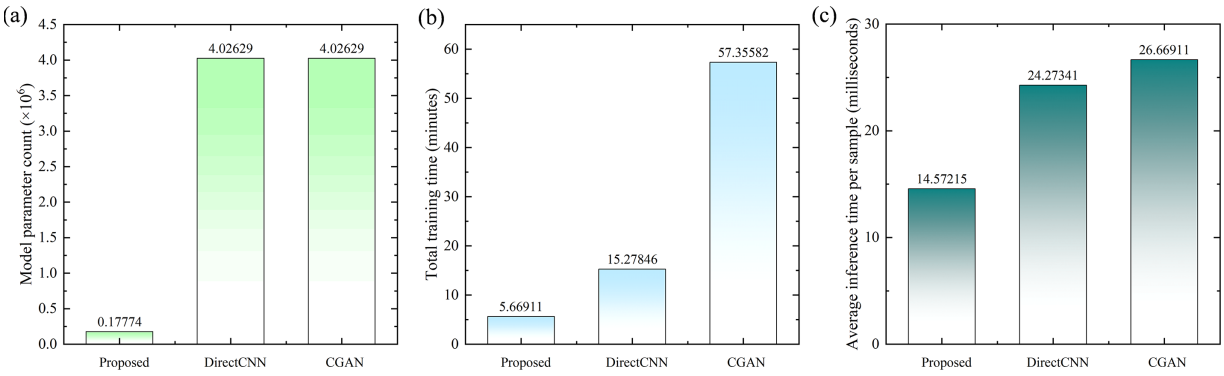


Figure 20: Computational cost comparison between the proposed framework and baseline models. (**a**) Number of trainable model parameters ($\times 10^6$). (**b**) Total training time (minutes). (**c**) Average inference time per sample (milliseconds).

Fig. 21 reports three complementary image quality metrics. The proposed method attains the highest structural similarity index (SSIM = 0.722), outperforming DirectCNN (0.701) and, more substantially, CGAN (0.603), indicating better preservation of global structural agreement with the ground truth (Fig. 21b). In terms of detail sharpness, quantified via the variance of the Laplacian, the proposed method and CGAN perform comparably (0.159 and 0.149, respectively), while DirectCNN collapses to 0.057 (Fig. 21a) a roughly 2.6-fold reduction that quantitatively corroborates the oversmoothed, low-detail outputs observed

qualitatively. Pixel-wise mean squared error is similar across all three methods (0.057, 0.062, and 0.053 for the proposed method, DirectCNN, and CGAN, respectively; Fig. 21c), indicating that raw pixel-level accuracy alone is insufficient to distinguish model quality and underscoring the importance of structure- and detail-sensitive metrics such as SSIM and Laplacian variance in this task.

Table 6: Summarizing the quantitative comparison between the proposed method and the DirectCNN and CGAN baselines across all evaluation metrics.

Metric	Proposed	DirectCNN	CGAN
Parameters ($\times 10^6$)	0.178	4.03	4.03
Training time (min)	5.67	15.28	57.36
Inference time (ms)	14.57	24.27	26.67
SSIM	0.722	0.701	0.603
Laplacian variance	0.159	0.057	0.149
MSE	0.057	0.062	0.053

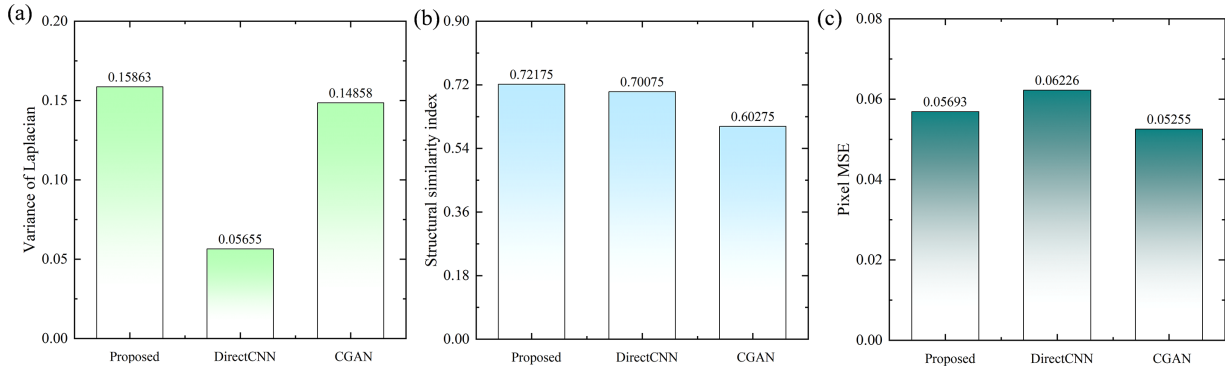


Figure 21: Quantitative image quality comparison between the proposed framework and baseline models. (**a**) Variance of Laplacian, reflecting edge/detail sharpness. (**b**) Structural similarity index (SSIM) relative to ground truth. (**c**) Pixel-wise mean squared error (MSE).

Taken together, these results indicate that direct end-to-end mappings (whether discriminative (DirectCNN) or generative (CGAN)) struggle to recover fine-grained, topologically coherent crack morphology from stress-strain curves alone, either collapsing toward oversmoothed outputs or producing structurally plausible but locally inaccurate patterns. The proposed morphology-aware decomposition instead achieves superior structural fidelity and detail preservation while requiring an order-of-magnitude fewer parameters and lower computational cost at both training and inference time, supporting the central motivation for the two-stage design adopted in this work.

It is important to clarify, however, that while the morphology bottleneck does not fully resolve the inherent non-uniqueness of the stress-strain-to-crack mapping (which would require probabilistic modeling to produce distributions over plausible crack patterns) it does transform this ill-posed problem into a more tractable form by imposing strong regularization through a physically interpretable low-dimensional manifold. This represents a substantial improvement over direct end-to-end regression and generative baselines in terms of stability, interpretability, and physical consistency within the present simulation framework. Addressing the full stochastic nature of the problem through probabilistic formulations, such

as variational inference or conditional diffusion models, remains an important direction for future work, as recent studies have highlighted the challenges and opportunities in AI-based crack detection [52].

4.7 Component Contribution Analysis

The ablation study systematically dismantles our single model to identify which components are truly essential for performance. Table 7 and Fig. 22a show that the complete model achieves strong descriptor prediction accuracy, and every ablation causes measurable degradation, confirming that all components contribute positively. Removing the metadata inputs causes the largest drop among structural ablations (0.74), suggesting that auxiliary information about the crack generation process provides critical context that the curve latent alone cannot capture. Freezing the curve encoder (0.80) and removing the curve latent entirely (0.79) have milder effects, indicating that the descriptor prediction pathway is relatively robust to perturbations in the encoding stage. The loss function ablations reveal that our default Huber configuration is well-tuned: uniform weighting drops performance to 0.76, while aggressive beta weighting severely hurts accuracy (0.66), demonstrating that the loss landscape requires careful balancing to maintain prediction stability across all descriptors simultaneously.

Table 7: Results of ablation study by the complete model to achieve strong descriptor prediction accuracy with all components contributing positively.

Variant	Mean R^2	Profile r	Δ R^2	Δr
Full model	0.810	0.600	Ref	Ref
w/o meta	0.740	0.520	−0.070	−0.080
w/o curve	0.790	0.570	−0.020	−0.030
Freeze curve	0.800	0.590	−0.010	−0.010
Descr. (4)	0.800	0.580	−0.010	−0.020
Descr. (7)	0.820	0.610	+0.010	+0.010
Huber w = 1	0.760	0.550	−0.050	−0.050
Huber $\beta = 1$	0.660	0.470	−0.150	−0.130
Pixel BCE	n/a	0.300	–	–
Pixel Dice	n/a	0.270	–	–
Pixel BCE + Dice	n/a	0.310	–	–
GT descr.	1.000	0.640	+0.190	+0.040
Shuffled descr.	n/a	0.240	–	–

The most telling result appears in panel Fig. 22b, where the same model variants are evaluated on the downstream task of rendered image fidelity. Here, the pixel-decoder variants perform at chance levels, while all morphology-bottleneck variants substantially outperform them. This stark contrast reveals that the descriptor bottleneck is not merely a convenient intermediate representation but a necessary inductive bias for the model to generate physically realistic crack patterns. Even when descriptor predictions are imperfect (as in the loss-ablated variants), the model still produces better renderings than when it attempts to learn pixel-level mappings directly. The per-descriptor heatmap in panel Fig. 22c further localizes these effects, showing that certain morphological features are robustly predicted across all ablations, while others (orientation and branching) are more sensitive to the removal of metadata. Finally, panel Fig. 22d quantifies the cumulative cost of each ablation, revealing that the pixel decoder is the single most damaging modification.

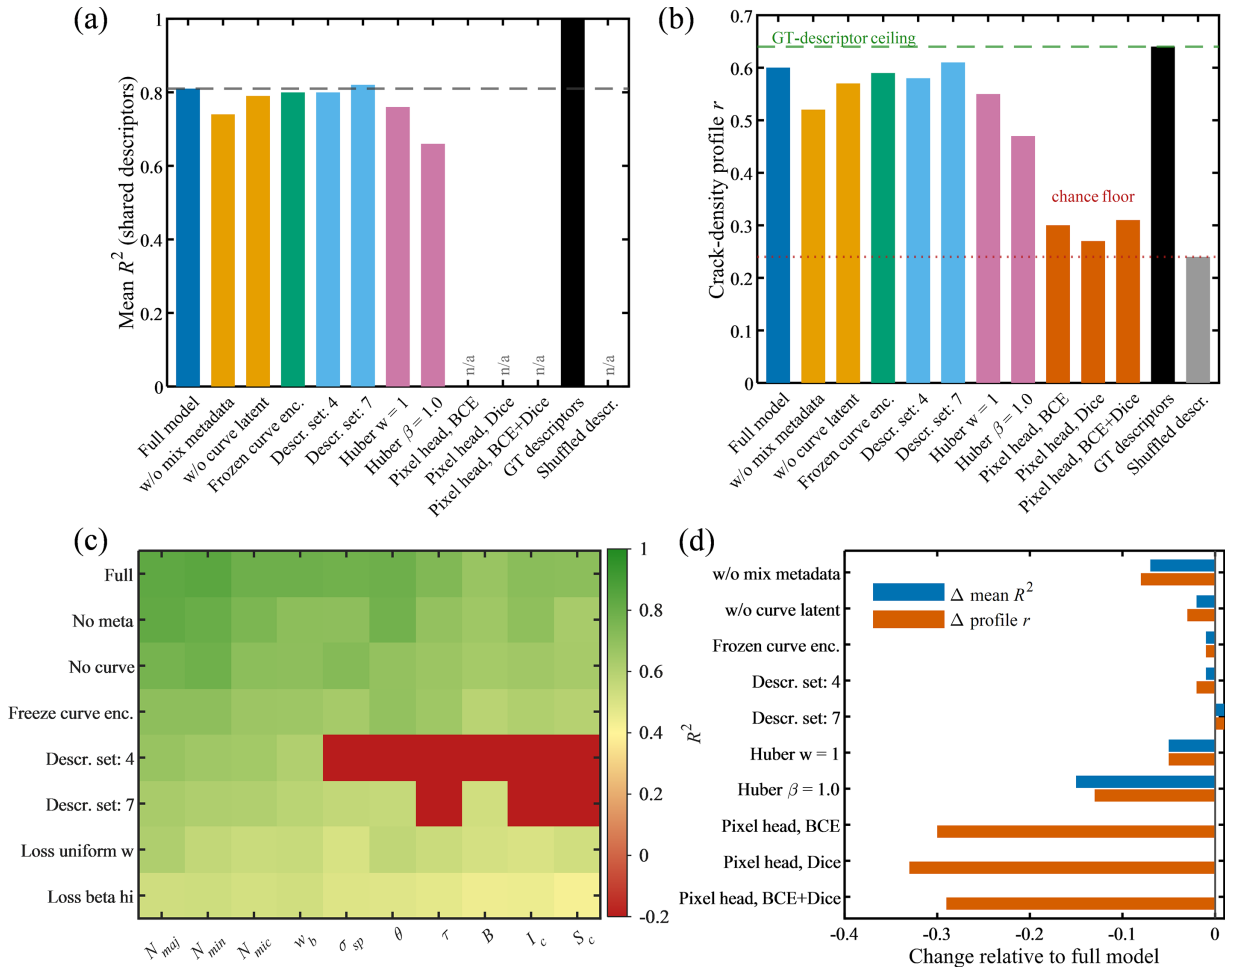


Figure 22: Ablation study of model components. (**a**) Mean R^2 for shared descriptors. (**b**) Profile correlation r. (**c**) Per-descriptor R^2 heatmap. (**d**) Degradation from full model. Morphology-bottleneck variants outperform pixel decoders, confirming the bottleneck's critical role.

5 Conclusions

This work introduced a morphology-aware two-stage framework for predicting SFRC crack patterns from uniaxial stress-strain responses and mix-design metadata. By decomposing the ill-posed curve-to-image regression into sequential morphology regression (StatNet) and conditional pattern synthesis (PatternDecoder), the framework achieved a global pooled R^2 of 0.940 across eleven morphology descriptors on the validation set. Key findings are summarized as follows:

(1) The framework achieves high predictive accuracy for morphological descriptors, with R^2 = 0.961 for crack width and R^2 > 0.70 for ten of the eleven descriptors evaluated. Residual diagnostics confirm homoscedasticity and near-normal error distributions, indicating robust performance within the simulation dataset.

(2) Gradient sensitivity analysis reveals activation spikes at peak stress, capturing the microcrack to macrocrack transition. The network captures physically meaningful relationships rather than superficial correlations, though this observation is limited to the present simulation dataset.

(3) Consistent performance across hybrid, macro-only, and micro-only compositions indicates generalization within the training distribution. Future work requires experimental validation, multiaxial loading,

3D extension, larger datasets, and probabilistic formulations to address the inherent non-uniqueness of crack pattern evolution.

The proposed framework establishes a viable pathway toward real-time, interpretable surrogate modeling for fiber-reinforced composites, bridging the gap between expensive high-fidelity simulations and practical engineering applications. By transforming an ill-posed inverse problem into a tractable, morphology-aware decomposition, this approach enables rapid assessment of crack patterns from readily available mechanical test data, supporting reliability-based design and parametric optimization of SFRC structures without prohibitive computational costs. While the morphology bottleneck does not fully resolve the inherent non-uniqueness of the stress-strain-to-crack mapping, it substantially improves tractability by imposing strong regularization through a physically interpretable low-dimensional manifold. However, it is important to acknowledge that all conclusions are drawn from the 220-simulation dataset generated using the validated mesoscale CZM framework; validation against experimental data and extension to different loading conditions or fiber types would be necessary to establish broader generalizability. Furthermore, the current mesoscale model assumes 2D plane-stress conditions and restricts fiber orientations to the in-plane direction, representing a simplification of the true 3D fiber distribution in SFRC. Future work will address these limitations by extending to 3D modeling, multiaxial loading conditions, larger experimental datasets, and probabilistic formulations to capture the full stochastic nature of crack pattern evolution and quantify prediction uncertainty.

Acknowledgement: The authors would like to express their sincere gratitude to all individuals who contributed to this research.

Funding Statement: This work was supported by the National Natural Science Foundation of China (No. 52475166), and National Key Research and Development Program of China (No. 2024YFE0204900).

Author Contributions: Khaja Wahaajuddin Kawkabi: Writing original draft, visualization, validation, methodology, investigation, conceptualization. Andre Steve Nziko Talla Nziko: conceptualization, methodology. Prince Manyanya Ngangura: Writing review & editing. Xu Long: Writing review & editing, supervision, resources, visualization, funding acquisition. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: Data will be made available on request.

Ethics Approval: Not applicable.

Conflicts of Interest: Given his role as Guest Editor of this special issue, Xu Long had no involvement in the peer review of this article and had no access to information regarding its peer review. Full responsibility for the editorial process for this article was delegated to another journal editor. The authors declare no conflicts of interest.

References

1. Bentur A, Mindess S. Fibre reinforced cementitious composites. 2nd ed. Boca Raton, FL, USA: CRC Press; 2007.
2. Banthia N, Sappakittipakorn M. Toughness enhancement in steel fiber reinforced concrete through fiber hybridization. Cem Concr Res. 2007;37(9):1366–72. doi:10.1016/j.cemconres.2007.05.005.
3. Li VC, Leung CKY. Steady-state and multiple cracking of short random fiber composites. J Eng Mech. 1992;118(11):2246–64. doi:10.1061/(asce)0733-9399(1992)118:.
4. Naaman AE, Reinhardt HW. Proposed classification of HPFRC composites based on their tensile response. Mater Struct. 2006;39(5):547–55. doi:10.1617/s11527-006-9103-2.
5. Shah AA, Ribakov Y. Recent trends in steel fibered high-strength concrete. Mater Des. 2011;32(8–9):4122–51. doi:10.1016/j.matdes.2011.03.030.

6. Smarzewski P, Barnat-Hunek D. Effect of fiber hybridization on durability related properties of ultra-high performance concrete. Int J Concr Struct Mater. 2017;11(2):315–25. doi:10.1007/s40069-017-0195-6.
7. Wille K, Naaman AE, Parra-Montesinos GJ. Ultra-high performance concrete with compressive strength exceeding 150 MPa (22 ksi): a simpler way. ACI Mater J. 2011;108(1):46. doi:10.14359/51664215.
8. Li J, Zhang L, Li Y, Sun C. Mesoscale modeling of steel fiber reinforced concrete using geometric entity expansion and point-line topology. Materials. 2026;19(8):1508. doi:10.3390/ma19081508.
9. Chen Q, Cheng W, Zou B, Kong B, Deng Y, Long X. Experimental study on additional stress induced by grouting with polyurethane-modified cementitious materials under confined conditions. J Polym Mater. 2025;42(2):463–75. doi:10.32604/jpm.2025.064257.
10. Schoen SM. Reliability analysis and design of RC and SFRC structures considering polymorphic uncertainties. [cited 2026 Jan 1]. Available from: https://hss-opus.ub.ruhr-uni-bochum.de/opus4/frontdoor/index/index/docId/13335.
11. Barros JAO, Cunha VMCF, Ribeiro AF, Antunes JAB. Post-cracking behaviour of steel fibre reinforced concrete. Mater Struct. 2005;38(1):47–56. doi:10.1007/BF02480574.
12. Zhang S, Hamed E, Amin A, Gilbert RI. Meso-mechanical modelling and analysis of the creep of steel-fibre-reinforced concrete. Mag Concr Res. 2025;77(19–20):1186–203. doi:10.1680/jmacr.25.00096.
13. Tiberti G, Minelli F, Plizzari G. Cracking behavior in reinforced concrete members with steel fibers: a comprehensive experimental study. Cem Concr Res. 2015;68(3):24–34. doi:10.1016/j.cemconres.2014.10.011.
14. Barenblatt GI. The mathematical theory of equilibrium cracks in brittle fracture. In: Advances in applied mechanics. Vol. 7. Amsterdam, The Netherlands: Elsevier; 1962. p. 55–129. doi:10.1016/s0065-2156(08)70121-2.
15. Hillerborg A, Modéer M, Petersson PE. Analysis of crack formation and crack growth in concrete by means of fracture mechanics and finite elements. Cem Concr Res. 1976;6(6):773–81. doi:10.1016/0008-8846(76)90007-7.
16. Ren W, Yang Z, Sharma R, Zhang C, Withers PJ. Two-dimensional X-ray CT image based meso-scale fracture modelling of concrete. Eng Fract Mech. 2015;133:24–39. doi:10.1016/j.engfracmech.2014.10.016.
17. Wu JY, Nguyen VP, Nguyen CT, Sutula D, Sinaie S, Bordas SP. Phase-field modeling of fracture. In: Advances in applied mechanics. Vol. 53. Cambridge, MA, USA: Academic Press; 2020.
18. Goswami S, Yin M, Yu Y, Karniadakis GE. A physics-informed variational DeepONet for predicting crack path in quasi-brittle materials. Comput Meth Appl Mech Eng. 2022;391(18):114587. doi:10.1016/j.cma.2022.114587.
19. Abbas YM, Khan MI. Efficacious application of data-driven machine learning models for predicting and optimizing the flexural tensile strength of fiber-reinforced concrete. Structures. 2024;64:106574. doi:10.1016/j.istruc.2024.106574.
20. 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. doi:10.1016/j.jcp.2018.10.045.
21. Lu L, Meng X, Mao Z, Karniadakis GE. DeepXDE: a deep learning library for solving differential equations. arXiv:1907.04502. 2019.
22. Chellam JAI, Mansingh BB, Alex DS, Binoj JS. Prediction and validation of mechanical properties of *Areca catechu/Tamarindus indica* fruit fiber with nano coconut shell powder reinforced hybrid composites. J Polym Mater. 2025;42(3):773–94. doi:10.32604/jpm.2025.069295.
23. Ghaffari Motlagh Y, Jimack PK, de Borst R. Deep learning phase-field model for brittle fractures. Int J Numer Methods Eng. 2023;124(3):620–38. doi:10.1002/nme.7135.
24. Long J, Shelhamer E, Darrell T. Fully convolutional networks for semantic segmentation. In: Proceedings of the 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR); 2015 Jun 7–12; Boston, MA, USA. p. 3431–40. doi:10.1109/cvpr.2015.7298965.
25. Liu Y, Yao J, Lu X, Xie R, Li L. DeepCrack: a deep hierarchical feature learning architecture for crack segmentation. Neurocomputing. 2019;338(12):139–53. doi:10.1016/j.neucom.2019.01.036.
26. Quan J, Ge B, Wang M. CrackViT: a unified CNN-transformer model for pixel-level crack extraction. Neural Comput Appl. 2023;35(15):10957–73. doi:10.1007/s00521-023-08277-7.

27. Zhou Z, Yuan M. SteelGuard-yolo: steel surface defect detection network based on improved YOLOv5s. Rev Int Métodos Numér Cálc Diseño Ing. 2024;40(2):29. doi:10.23967/j.rimni.2024.05.011.
28. Isola P, Zhu JY, Zhou T, Efros AA. Image-to-image translation with conditional adversarial networks. In: Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR); 2017 Jul 21–26; Honolulu, HI, USA. p. 5967–76. doi:10.1109/cvpr.2017.632.
29. Lin X, Wu J, Liang S. A CGAN framework for predicting crack patterns and stress-strain behavior in concrete random media. Comput Model Eng Sci. 2025;145(1):215–39. doi:10.32604/cmes.2025.070846.
30. Wang Y, Zhang P, Xiang Y, Hou J, Chen X, Wang X, et al. CNN-GAN based microstructure design for biphasic composites with balanced modulus and frequency. Comput Mater Sci. 2026;265(35):114513. doi:10.1016/j.commatsci.2026.114513.
31. Arjovsky M, Chintala S, Bottou L. Wasserstein GAN. arXiv:1701.07875. 2017.
32. Salimans T, Goodfellow I, Zaremba W, Cheung V, Radford A, Chen X. Improved techniques for training GANs. arXiv:1606.03498. 2016.
33. Mirza M, Osindero S. Conditional generative adversarial nets. arXiv:1411.1784. 2014.
34. Ajjarapu SSP, Mallikarjunaiah SM. CrackMorph-XAI-net: a topology-preserving and explainable framework for automated crack morphology. arXiv:2605.08265. 2026.
35. Gdoutos EE. Fracture mechanics: an introduction. 3rd ed. Cham, Switzerland: Springer International Publishing; 2020. doi:10.1007/978-3-030-35098-7.
36. Fujita Y, Hamamoto Y. A robust automatic crack detection method from noisy concrete surfaces. Mach Vis Appl. 2011;22(2):245–54. doi:10.1007/s00138-009-0244-5.
37. Long X, Mohamed B, Gu T, Abdelkibir B, Guo Q. Effects of size and texture on anisotropic elasto-plasticity in Cu–Nb nanolaminates by integrating semi-empirical and data-driven approaches. Mech Mater. 2026;217(13):105670. doi:10.1016/j.mechmat.2026.105670.
38. Barredo Arrieta A, Díaz-Rodríguez N, Del Ser J, Bennetot A, Tabik S, Barbado A, et al. Explainable artificial intelligence (XAI): concepts, taxonomies, opportunities and challenges toward responsible AI. Inf Fusion. 2020;58(3):82–115. doi:10.1016/j.inffus.2019.12.012.
39. Guidotti R, Monreale A, Ruggieri S, Turini F, Pedreschi D, Giannotti F. A survey of methods for explaining black box models. arXiv:1802.01933. 2018.
40. Lu L, Jin P, Karniadakis GE. DeepONet: learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. arXiv:1910.03193. 2019.
41. Park K, Paulino GH, Roesler JR. A unified potential-based cohesive model of mixed-mode fracture. J Mech Phys Solids. 2009;57(6):891–908. doi:10.1016/j.jmps.2008.10.003.
42. Benzeggagh ML, Kenane M. Measurement of mixed-mode delamination fracture toughness of unidirectional glass/epoxy composites with mixed-mode bending apparatus. Compos Sci Technol. 1996;56(4):439–49. doi:10.1016/0266-3538(96)00005-X.
43. Dassault Systèmes Simulia Corp. Abaqus analysis user's guide. Paris, France: Dassault Systèmes Simulia Corp; 2016.
44. Feder J. Random sequential adsorption. J Theor Biol. 1980;87(2):237–54. doi:10.1016/0022-5193(80)90358-6.
45. Li J, Deng Z. Tensile behavior of hybrid fiber-reinforced ultra-high-performance concrete. Front Mater. 2021;8:769579. doi:10.3389/fmats.2021.769579.
46. Chun B, Yoo DY. Hybrid effect of macro and micro steel fibers on the pullout and tensile behaviors of ultra-high-performance concrete. Compos Part B Eng. 2019;162(10):344–60. doi:10.1016/j.compositesb.2018.11.026.
47. Ige AO, Sibiya M. State-of-the-art in 1D convolutional neural networks: a survey. IEEE Access. 2024;12(1):144082–105. doi:10.1109/ACCESS.2024.3433513.
48. Kiranyaz S, Avci O, Abdeljaber O, Ince T, Gabbouj M, Inman DJ. 1D convolutional neural networks and applications: a survey. Mech Syst Signal Process. 2021;151:107398. doi:10.1016/j.ymssp.2020.107398.
49. Zou Q, Zhang Z, Li Q, Qi X, Wang Q, Wang S. DeepCrack: learning hierarchical convolutional features for crack detection. IEEE Trans Image Process. 2019;28(3):1498–512. doi:10.1109/TIP.2018.2878966.
50. Goodfellow IJ, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, et al. Generative adversarial networks. arXiv:1406.2661. 2014.

51. Milletari F, Navab N, Ahmadi SA. V-net: fully convolutional neural networks for volumetric medical image segmentation. arXiv:1606.04797. 2016.
52. Zhong Y, Li J, Liu G, Hao H. Probabilistic quantification of cracks in fiber-reinforced concrete beams by using Bayesian neural networks. Struct Control Heal Monit. 2026;2026(1):1898553. doi:10.1155/stc/1898553.