Computer Modeling in Engineering & Sciences

Tech Science Press

https://doi.org/10.32604/cmes.2026.083687

ARTICLE

Deep Learning-Accelerated Extended Finite Element Method for Crack Propagation Analysis in Composite Structures

Shamsh Parveen[1], **Mehtab Alam**[2,*], **Ashraf Ali**[3] **and Abdullah Alourani**[4]

[1]Department of Mechanical Engineering, Mahamaya Institute of Information Technology, Kushinagar, India
[2]Department of Computer Science, Acharya Narendra Dev College, University of Delhi, Delhi, India
[3]Faculty of Computer Studies, Arab Open University, A'ali, Bahrain
[4]Department of Management Information Systems, College of Business and Economics, Qassim University, Buraydah, Saudi Arabia

*Corresponding Author: Mehtab Alam. Email: mahiealam@gmail.com or mehtabalam@andc.du.ac.in

Received: 23 April 2026; Accepted: 25 June 2026; Published: 28 August 2026

ABSTRACT: Accurate prediction of complex failure modes in anisotropic composite structures—specifically matrix cracking, fiber rupture, and delamination (stratification)—remains a central challenge in computational fracture mechanics. The primary goal of this work is to bridge the gap between high-fidelity physical modeling and computational efficiency. While the extended finite element method (XFEM) enables mesh-independent crack modeling, its computational cost limits scalability. This work proposes a deep learning–accelerated extended finite element framework (DL-XFEM) that couples physically admissible XFEM fields with a neural network surrogate to predict incremental crack growth. XFEM is employed to generate stress-intensity factors and fracture-consistent state variables, which are used to train the network to replace the most computationally intensive crack-update steps while preserving fracture-mechanics constraints. The proposed framework achieves an 88.3% reduction in runtime relative to standard XFEM. Quantitatively, the method delivers an average trajectory prediction error of RMSE = 0.083 MPa$\sqrt{}$m and MAE = 0.071 mm, reducing prediction error by 61.2% compared to standard Physics-Informed Neural Network (PINN) baselines. Generalization is demonstrated across multiple stacking sequences and loading scenarios, and predictions are validated against experimental crack-growth data for aluminum, titanium, and carbon/epoxy laminate specimens. The results indicate that DL-XFEM provides a scalable and physically consistent approach for accelerating fracture simulations in composite structures without sacrificing predictive reliability.

KEYWORDS: Hybrid XFEM; deep learning; fracture prediction; anisotropic composites; physics-informed modeling; failure analysis; computational efficiency; digital twin; uncertainty quantification

1 Introduction

The structural integrity of advanced composite systems is rarely a matter of global strength limits; rather, it is dictated by localized, interacting fracture processes that evolve under complex stress states. In critical applications such as aerospace panels, wind turbine blades, and marine laminates, the onset of matrix cracking, fiber rupture, and delamination determines the remaining service life. In this context, theoretical fracture models must be connected with practical aspects of structural durability to accurately assess residual life, especially in the presence of stress concentrators or incipient cracks [1]. Furthermore, because composite structures are critically sensitive to internal defects, modern approaches to strength assessment increasingly require interdisciplinary methods for robust damage resistance analysis [2]. Predicting these trajectories in anisotropic media remains a formidable challenge due to the strong directional stiffness contrast and

heterogeneous microstructures that make crack paths highly sensitive to stacking sequences and loading modes. Consequently, any accurate numerical framework must resolve both near-tip singular stress fields and the continuously evolving geometry of the discontinuity.

While traditional strategies like Cohesive Zone Models (CZMs) offer fundamental insights, they often require predefined crack paths or computationally expensive remeshing to maintain mesh conformity. The Extended Finite Element Method (XFEM) was introduced to bypass these limitations by enriching the displacement approximation with discontinuous Heaviside and asymptotic crack-tip functions. This enables mesh-independent representation of the crack, making XFEM a gold standard for layered anisotropic systems [3]. However, a significant bottleneck remains: the computational burden of resolving steep stress gradients near the crack front scales poorly with mesh refinement. For large-scale digital twins or uncertainty quantification tasks, the repeated incremental updates in XFEM become a prohibitive barrier to real-time applicability [4–6].

Recent shifts toward deep learning (DL) have opened new pathways for acceleration, with neural architectures demonstrating an ability to approximate stress fields and act as surrogates for numerical solvers. Yet, most existing models operate as "black-box" predictors that lack physical interpretability or function merely as post-processing tools. To bridge this gap, there is an urgent need for hybrid frameworks where learning-based predictions are embedded directly within the physics-based solution loop.

This work introduces a deep learning–accelerated extended finite element framework (DL-XFEM) specifically designed for crack propagation in anisotropic composites. Unlike purely data-driven surrogates or Physics-Informed Neural Networks (PINNs) that attempt to replace the solver entirely, our approach retains XFEM as the primary physics engine. By leveraging physically admissible XFEM state variables—such as stress intensity factors (K_I, K_{II}) and strain-energy density—we train a neural network to accelerate the computationally intensive crack-update steps while enforcing fracture-mechanics constraints. This "in-loop" coupling ensures that predictions remain grounded in mechanical reality, achieving up to an 88% reduction in runtime compared to standard XFEM without sacrificing predictive reliability [7–11]. Therefore, the overarching goal of this work is to bridge the gap between the high physical fidelity of traditional numerical solvers and the operational efficiency of data-driven models.

To rigorously evaluate this framework, we address four central research questions through numerical and experimental validation. First, we investigate whether a deep network can reliably predict incremental crack extension in anisotropic laminates using local XFEM state variables as inputs. Second, we quantify the extent of computational cost reduction possible while maintaining a pre-specified accuracy threshold relative to full-physics solutions. Third, we test the model's ability to generalize across novel composite stacking sequences and crack geometries not encountered during the training phase. Finally, we assess the sensitivity of these hybrid predictions to stochastic uncertainty in material properties and boundary conditions.

The remainder of this paper is organized as follows: Section 2 reviews the state-of-the-art in physics-informed modeling; Section 3 details the proposed DL-XFEM architecture and the handshake between the physics and learning modules; Section 4 presents validation results against experimental data for aluminum, titanium, and carbon/epoxy specimens; Section 5 discusses efficiency and uncertainty; and Section 6 provides concluding remarks.

2 Literature Review

2.1 Background and Motivation

Estimating the locations where fractures will initiate in composite materials, as well as how they will propagate through them, is still a challenging problem due to the inherent heterogeneity of composite

materials, anisotropy, and multiscale damage processes. While traditional numerical methods, such as the finite element method (FEM), can model many characteristics of composite materials, they cannot easily simulate discontinuities in composite materials, such as delamination or interface debonding between plies of a laminate. To overcome these challenges, the XFEM has been developed; it is based on extending the basic shape functions of the standard finite element method (FEM) using enrichment terms to describe crack discontinuities regardless of the local mesh configuration. The ability to accurately model complex crack geometries without having to frequently remesh makes XFEM a widely used tool for analyzing fractures in both homogeneous and composite systems [12].

Despite its ability to accurately model complex crack geometries, the computational burden associated with XFEM increases significantly with mesh refinement, multiscale coupling, or the presence of several interacting cracks in a composite material—all conditions that are typical in fiber reinforced laminates. To alleviate some of the computational burden associated with XFEM, researchers have increasingly explored the use of deep learning (DL) to accelerate XFEM simulations, reflecting a broader academic trend of utilizing DL surrogates to overcome the computational limitations of traditional numerical methods in complex physical problems [13]. Specifically, training DL models to learn fracture patterns from XFEM simulation data has emerged as a highly effective strategy.

2.2 Deep Learning in Fracture and Stress Prediction

The early attempts of incorporating DL into fracture mechanics mainly focused on predicting the stress evolution and crack morphology with the use of numerical or experimental datasets. Wang et al. developed *StressNet*, a Temporal-Independent Convolutional Neural Network (TI-CNN) coupled with Bi-LSTM layers to predict internal stress evolution under cyclic loads. Their framework achieved several-order-of-magnitude acceleration over finite-discrete element methods while maintaining high correlation with experiments [14].

Building on this concept, Xu et al. proposed *Crack-Net*, a U-Net-based model trained on particulate-composite data. By embedding implicit physical relations between crack evolution and stress response, Crack-Net attained an R^2 = 0.9993 and relative error below 1% for phase-field crack evolution in most validation cases [15]. He et al. further introduced a generative deep model capable of directly predicting crack trajectories from microstructure images in porous media, achieving 90.25% trajectory accuracy and reducing computation time from hours to seconds [9].

These studies demonstrate that deep learning can approximate fracture-related fields with high computational efficiency, but often at the expense of physical interpretability and robustness outside trained regimes.

Recent work has also targeted direct estimation of Stress Intensity Factors (SIFs) using data-efficient ML regressors. Guo et al. applied Pool-based Active Learning with Gaussian Process Regression (PA-GPR) to compute SIFs in irregular geometries, which helps in cutting data requirements to a very large extent [16]. The papers by Merrell et al. [17] and Xu et al. [18] demonstrated that fracture prediction using machine learning can be feasible and provide good run-time performance compared to traditional numerical SIF solvers. Chen et al. developed surrogate networks for full-field stress and crack-pattern prediction in composites, obtaining consistent accuracy with reduced data requirements [19]. Spatiotemporally defined neural network architectures (i.e., 3D-CNN-LSTM hybrids) were also shown to provide rapid time-resolution predictions of the fracture evolution [20]. This aligns with recent advancements in computational studies of composite structures, which emphasize the necessity of integrating advanced machine learning techniques to capture the complex, non-linear mechanical responses inherent to anisotropic materials [21,22]. Therefore, physical reliability must be maintained across all configurations.

2.3 Physics-Informed Neural Networks and XFEM Coupling

To find the right balance between the computational efficiency of data-driven approaches and the physical correctness of continuum mechanics, researchers have started to use Physics-Informed Neural Networks (PINNs) that incorporate the fundamental physical laws (equilibrium relations or energy conservation) into their optimization process, so that even data-scarce cases produce physically correct results. Based on the idea of PINNs, Lotfalian et al. developed the eXtended Physics-Informed Neural Network (X-PINN), which incorporates the XFEM-based enrichment functions into neural architectures to better describe the stress singularities near crack tips [23]. Gu et al. also employed enriched PINN formulations for analysis of cracks in plane, and by employing asymptotic crack tip functions, they were able to obtain higher precision without having to refine the mesh [24]. Manav et al. used the Deep-Ritz approach to incorporate phase-field fracture evolution within neural architectures, enhancing convergence and interpretability [25].

Further research led to the creation of Discontinuity-Embedded Neural Networks (DENNs) in which the geometry of cracks was described using implicit signed distance fields. This method allowed the neural network to learn discontinuous displacement and stress distributions without having to explicitly segment the mesh, and therefore provided both geometric flexibility and numerical stability [26]. Most of these frameworks make use of the neural postprocessing instead of the tight coupling of physics solvers. Despite these advances, existing PINN-based and enriched frameworks generally treat neural models as external solvers rather than embedding them directly within the XFEM crack-growth loop.

2.4 Machine Learning for Composite Property Characterization

There are numerous machine learning approaches that are being applied to predict various material properties that determine the fracture response of composites. For example, Karamov et al. predicted the fracture toughness of pultruded composites using gradient-boosted decision trees with a mean squared error that was lower than 10% of the average experimental values, demonstrating that the longitudinal modulus and transverse strength had the greatest impact [27]. Ahmed et al. conducted a review of AI applications for laminated structures and highlighted unresolved challenges, including thermal buckling and defects in thin-walled structures [28].

Kushvaha et al. applied artificial neural networks to model the dynamic initiation of fractures in polymer composite materials subjected to impact loads and obtained approximately 92.7% accuracy for predicting SIF [29]. The development of a unified predictive framework that includes anisotropy, fiber orientation and the influence of interfacial debonding represents a major open challenge.

2.5 Current Limitations and Research Gaps

Despite progress in machine-learning-aided fracture prediction, several technical gaps persist. A primary issue is the oversimplification of material anisotropy; many existing models treat composites as isotropic media, neglecting the fiber orientation and interlaminar debonding that fundamentally drive crack trajectories [30]. Furthermore, the lack of tight integration between XFEM solvers and deep learning remains a major hurdle. Current approaches often relegate machine learning to a post-processing role or an external surrogate rather than embedding it directly within the physics-based solution loop [31,32].

This disconnect frequently results in poor generalizability when the models are applied to novel stacking sequences or varying loading modes [33]. Consequently, there is a recurring trade-off between the high inference speeds of data-driven surrogates and the mechanical rigor of physics-informed models [25,34]. No existing framework yet combines a robust anisotropic representation with in-loop XFEM coupling and the

scalability required for diverse engineering applications. Bridging this gap—uniting high-speed prediction with strict fracture-mechanics consistency—is the core motivation of the proposed DL-XFEM framework.

2.6 Advances in XFEM and Integration Opportunities

The development of recent advancements in XFEM has made it possible the potential to develop these hybridizations. A new approach developed by Martínez et al., was called Stochastic-XFEM (S-XFEM). S-XFEM utilized stochastic approaches in the analysis of crack propagation and demonstrated its ability to utilize many types of machine learning surrogate models to rapidly evaluate probabilistic outcomes [32]. Song et al. demonstrated the application of XFEM–ML coupling to simulate competing mechanisms such as capsule rupture and interface debonding in the self-healing composites. They depicted that trained models can adaptively update XFEM fields with notable efficiency gains [34].

Table 1 illustrates the significant advancements within the field of computational fracture mechanics and highlights how the fields of deep learning, physics-informed networks, and hybrid XFEM models have advanced over time. Table 1 also clearly illustrates the research gap that the proposed DL-XFEM model will target.

Table 1: Summary of the literature review.

Ref. No.	Method/Framework	Material System or Context	Key Contribution/Findings	Relevance to Present Research
[12]	Extended Finite Element Method (XFEM)	General solids & composites	Introduced XFEM to represent discontinuities independent of mesh; eliminated remeshing for crack growth.	Foundational basis of the computational core.
[14]	Deep Learning for Stress Prediction	Metals under cyclic loading	Predicted internal stress evolution with $>10^3\times$ speed-up vs. FEM.	Demonstrates feasibility of DL acceleration.
[15]	Physics-Constrained DL Model	Particulate Composites	Achieved $R^2 = 0.9993$ for stress prediction; <1% error in phase-field crack evolution.	Motivates integration of physics-informed DL.
[16]	Pool-Based Active Learning GPR for SIF	Irregular Geometries	Reduced training data while preserving SIF accuracy.	SIF prediction optimization approach.
[17]	Mechanics-Guided ML Regression	Isotropic fracture mechanics	Embedded fracture equations in ML model for interpretability.	Improves trust in ML predictions.
[18]	Hybrid ML–Mechanics SIF Model	2D crack systems	Achieved major runtime reduction over numerical solvers.	Confirms ML's efficiency for SIF computation.

(Continued)

Table 1 (continued)

Ref. No.	Method/Framework	Material System or Context	Key Contribution/Findings	Relevance to Present Research
[19]	DL Surrogate for Stress Field Prediction	Fiber Composites	Achieved high accuracy with limited data via U-Net surrogates.	Supports surrogate integration within XFEM.
[20]	Spatiotemporal DL for Crack Propagation	Dynamic fracture systems	Captured time-dependent crack evolution with 3D CNN–LSTM.	Basis for real-time DL prediction module.
[21]	Deep Learning Computational Framework	Composite structural failure	Demonstrated high efficiency in modeling complex mechanical responses in laminates	Validates the necessity of DL to capture complex structural responses.
[22]	Advanced ML for Non-linear Mechanics	Anisotropic composite structures	Captured complex non-linear mechanical responses using advanced machine learning models.	Supports the integration of ML for non-linear anisotropic behavior.
[23]	Physics-Informed NN with XFEM Enrichment	2D crack domains	Embedded XFEM enrichment in PINNs for tip singularity capture.	Foundation for hybrid DL–XFEM coupling.
[24]	Enriched PINN for Crack Prediction	2D in-plane cracks	Used crack-tip asymptotic functions; improved accuracy without mesh refinement.	Demonstrates physics consistency in DL models.
[25]	DL Embedded Phase-Field Fracture	Brittle materials	Integrated phase-field formulation into DL loss for better convergence.	Provides method for DL physics integration.
[26]	Signed-Distance-Field-Based DL	General fracture fields	Modeled discontinuous fields without mesh partitioning.	Key idea for mesh-free hybridization.
[27]	Gradient-Boosted Trees for Fracture Toughness	Pultruded Composites	Predicted toughness with MSE < 10% of the mean experiment.	Quantitative link between microstructure & fracture.
[28]	Review of AI for Composite Laminates	Laminated structures	Identified gaps in thermal buckling and defect modeling.	Highlights unsolved areas for AI in composites.

(Continued)

Table 1 (continued)

Ref. No.	Method/Framework	Material System or Context	Key Contribution/Findings	Relevance to Present Research
[29]	ANN for Dynamic Fracture Prediction	Polymer Composites	92.7% accuracy for stress intensity factor under impact.	Early validation of DL predictive capability.
[30]	XFEM for Anisotropic Fracture	Fiber-Reinforced Composites	Showing anisotropy significantly changes the crack path and SIF.	Motivates anisotropy-aware DL architectures.
[31]	XFEM–ML Coupling (Partial Integration)	Metal matrix composites	ML used for crack growth rate prediction outside the XFEM loop.	Illustrates the limits of post-processing approaches.
[32]	Stochastic XFEM with UQ Capabilities	Heterogeneous composites	Integrated uncertainty quantification within XFEM.	Enables probabilistic learning integration.
[13]	Transfer Learning for Composite Stress Fields	Carbon-fiber laminates	Cross-configuration learning success is limited to similar stacks.	Highlights generalization barriers.
[33]	Domain-Adapted CNN for Composite Fatigue	Multi-laminate systems	Reduced data requirements via domain adaptation.	Reinforces the need for cross-domain robustness.
[35]	PINN vs. Data-Driven Surrogates	2D fracture cases	Quantified speed-fidelity trade-off between models.	Informs balance of accuracy and efficiency.
[34]	Improved XFEM–ML Hybridization	Self-Healing composites	Adaptive update of XFEM fields via trained networks.	Validates in-loop DL control of XFEM.
[9]	Generative DL for Crack Trajectory Prediction	Porous Media	90.25% accuracy; reduced computation from hours to seconds.	Illustrates DL efficiency for crack path learning.

3 Methodology

This Framework combines XFEM's ability to accurately simulate the physics involved in fracture phenomena with a neural surrogate capable of adapting its predictions as it simulates crack growth increment by increment.

3.1 Framework Overview

The DL-XFEM architecture is bifurcated into an off-line data generation phase and an on-line inference phase, utilizing a tightly coupled in-loop design where the neural network functions as an active decision-maker within the incremental solver rather than a passive post-processing surrogate. It uses the XFEM

module to generate a high-dimensional feature vector—comprising local stress fields, stress intensity factors (SIFs), and anisotropic material tensors—that defines the mechanical state at the crack tip.

Fig. 1 illustrates the workflow of the proposed framework. During the on-line inference phase, the trained neural surrogate predicts incremental crack extensions by comparing predicted results against the full XFEM reference. These predictions allow the solver to bypass several computationally intensive incremental steps. System stability is enforced via a Physical Admissibility Criterion (PAC); if the prediction error exceeds this threshold, the XFEM engine re-engages to provide corrective data, preventing the surrogate from drifting into unphysical growth regimes.

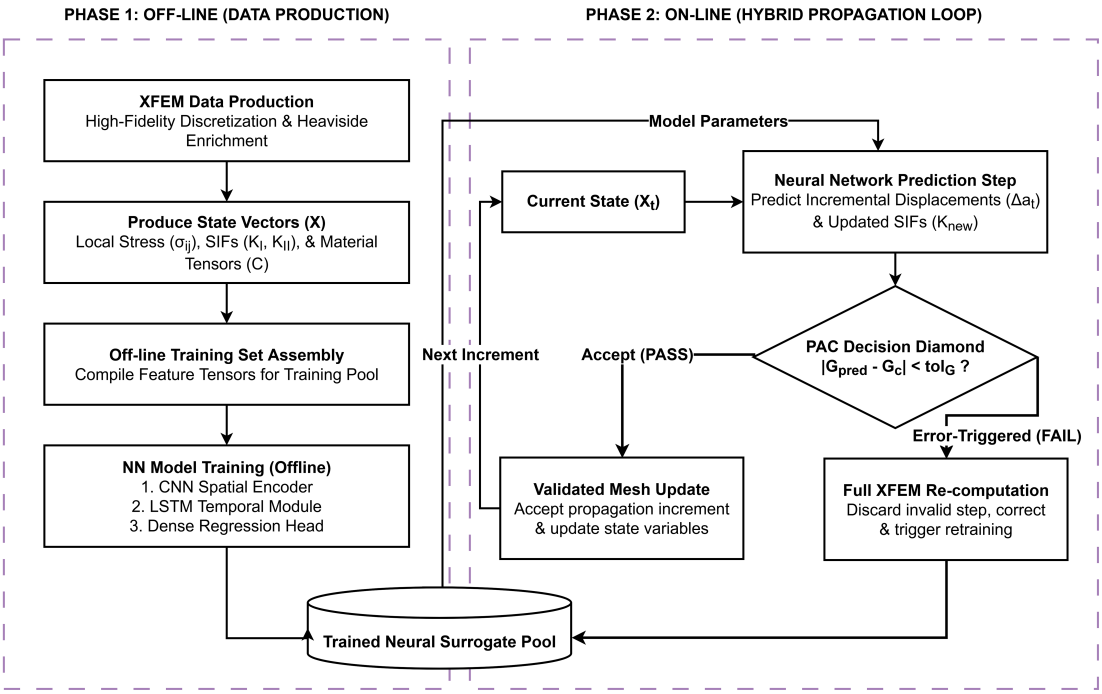

Figure 1: Workflow diagram showing data flow and control sequence.

3.2 Data Generation and Control Flow

The methodology is built upon a high-fidelity data generation pipeline that feeds the learning module, emphasizing the systematic transfer of physically consistent fracture quantities. The architecture is bifurcated into an off-line data generation phase and an on-line inference phase, utilizing a tightly coupled in-loop design where the neural network functions as an active decision-maker within the incremental solver rather than a passive post-processing surrogate. This approach ensures that the "handshake" between the physics engine and the learner is mathematically rigorous and mechanically grounded.

3.2.1 XFEM Data Production

XFEM simulations are executed across a diverse parameter space, including representative composite plates with varying fiber orientations, ply stacking sequences, and mixed-mode loading scenarios. To capture the evolving discontinuities without the numerical instability of repeated remeshing, the displacement field u(x) is enriched with Heaviside and asymptotic functions. This process generates a high-dimensional feature vector—comprising nodal displacements, stress tensors σ_{ij}, and crack-tip stress-intensity factors K_I, K_{II}—that uniquely defines the mechanical state at the crack tip. By resolving the near-tip singular

stress fields within the anisotropic media, the XFEM solver provides fracture-consistent state variables, including crack opening displacements (COD) and spatial trajectory coordinates, which serve as the primary ground-truth for the training phase.

3.2.2 Data Flow and Learning Cycle

The data flow is structured as a recursive learning cycle where XFEM outputs are normalized and organized into multi-channel feature tensors, $X \in R^{H\times W\times C}$. These tensors represent spatial distributions of stress, strain-energy density, and signed-distance maps describing the crack geometry. During the online inference phase, the neural surrogate processes these feature maps to predict the incremental crack-tip displacement vector $\Delta a = (\Delta x, \Delta y)$. System stability and mechanical admissibility are enforced via a Physical Admissibility Criterion (PAC). If the discrepancy between the neural prediction and local equilibrium exceeds a predetermined tolerance, the full XFEM solver is re-engaged to provide a high-fidelity correction. This feedback loop prevents the surrogate from drifting into unphysical growth regimes, ensuring that the framework maintains fracture-mechanics rigor across diverse material configurations.

3.3 XFEM Module Description

The computational core of the framework relies on the eXtended Finite Element Method (XFEM) to resolve the high-gradient stress fields and discrete discontinuities inherent in fractured composites. By enriching the standard finite element space with additional degrees of freedom, XFEM enables mesh-independent crack representation. For a domain containing a crack, the displacement vector field $u(x)$ is approximated as:

$$u(x) = \sum_{i\epsilon I} N_i(x)\, u_i + \sum_{j\epsilon J} N_j(x)\, H(x)\, a_j + \sum_{k\epsilon K} N_k(x) \left(\sum_{\alpha=1}^{4} F_\alpha(x)\, b_k^\alpha \right)$$

Fig. 2 details the XFEM mathematical enrichment formulation on a finite element mesh splitting a crack in a composite laminate.

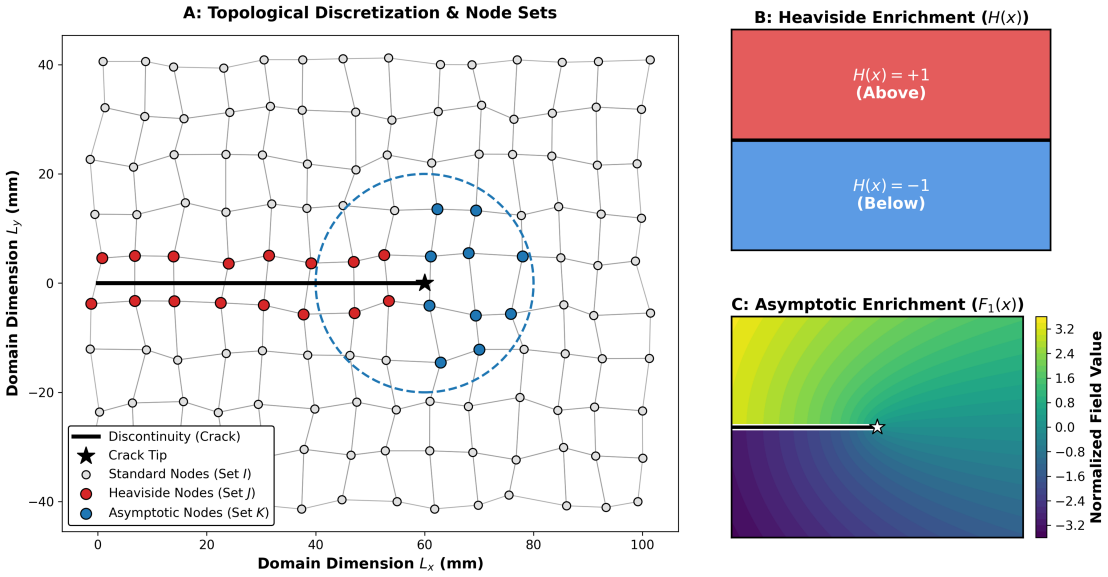

Figure 2: High-fidelity scientific vector illustration.

In this formulation, $N_i(x)$ represent the standard nodal shape functions. The Heaviside enrichment function $H(x)$ accounts for the displacement jump across the crack surface, taking values of ±1 depending on the nodal position relative to the discontinuity. The set of four asymptotic crack-tip functions, $F_\alpha(x)$, captures the singular stress fields near the crack front. The nodal sets I, J, and K distinguish between standard nodes, nodes bisected by the crack surface, and nodes within the enrichment radius of the crack tip, respectively.

To accurately simulate the complex failure modes of laminated systems, the domain is discretized using 4-node bilinear plane stress quadrilateral elements with reduced integration (CPS4R). External loads are numerically assigned via displacement-controlled boundary conditions applied incrementally to the outer edge nodes, ensuring stable quasi-static crack growth. The solver incorporates anisotropic elasticity tensors and Cohesive Zone Models (CZM) to resolve matrix-fiber debonding. The specific element compositions and mesh densities are detailed in Table 2. This physics engine outputs a comprehensive suite of fracture-relevant variables, including the stress tensor components (σ_{xx}, σ_{yy}, σ_{xy}), strain-energy density distributions, and crack opening displacements (COD) These data points are critical for grounding the neural surrogate in mechanical reality.

Table 2: XFEM simulation parameters and output variables.

Category	Parameter/Variable	Description	Value/Range	Role in DL-XFEM
Material Properties	E_f (Fiber Modulus)	Elastic modulus of reinforcing fibers	150–240 GPa	Governs anisotropic stiffness
	E_m (Matrix Modulus)	Elastic modulus of polymer matrix	2–5 GPa	Defines composite compliance
	ν_f, ν_m	Poisson's ratios for fiber and matrix	0.2–0.35	Influences shear deformation
	G_{fm}	Interfacial fracture energy	0.1–1.2 N/mm	Controls fiber–matrix debonding
Geometric and Discretization Setup	L × W	Plate dimensions	100 mm × 50 mm	Defines boundary extent
	a_0	Initial crack length	5–10 mm	Initiation parameter
	Fiber orientation angle (θ)	Principal fiber direction	0°, 45°, 90°	Captures anisotropy
	Finite Element Type	Base elements for domain discretization	CPS4R (4-node quadrilateral)	Provides a standard displacement field
	Mesh Size Range	Number of elements per model	10^3 to 5×10^5 elements	Dictates the spatial resolution of features
Boundary Conditions	Load type	Tensile/Mixed-mode loading	10–100 MPa	Determines fracture mode
	Support conditions	Fixed/simply supported	—	Ensures mechanical stability
XFEM Enrichment Scheme	Discontinuity Function $H(x)$	Heaviside step function	Binary (±1)	Model's displacement jump
	Crack-Tip Functions $F_\alpha(x)$	Asymptotic functions for Mode I/II	4 enrichment terms	Captures singular stress fields

(Continued)

Table 2 (continued)

Category	Parameter/Variable	Description	Value/Range	Role in DL-XFEM
Outputs for DL Module	$\sigma_{xx}, \sigma_{yy}, \sigma_{xy}$	Stress field components	Field tensors	Input for DL feature maps
	Strain Energy Density (SED)	Local strain energy distribution	Scalar field	Guides fracture evolution
	COD (Crack Opening Displacement)	Relative crack surface displacement	μm–mm range	Predicts propagation rate
	SIF (Stress Intensity Factor, K_I, K_{II})	Crack-tip intensity metrics	5–80 MPa$\sqrt{m}$	Validates DL predictions
	Crack Path Coordinates	x_c, y_c crack-tip evolution	Spatial trajectory	Used as ground truth for training

The high-fidelity nature of these simulations is illustrated in Fig. 3, which depicts crack propagation in a tensile-loaded anisotropic plate. The visualization captures the significant influence of fiber orientation and interfacial strength on the resulting crack trajectory and local stress concentration. By resolving these nuances, the XFEM module ensures that the training data encompasses the intricate physics of composite fracture.

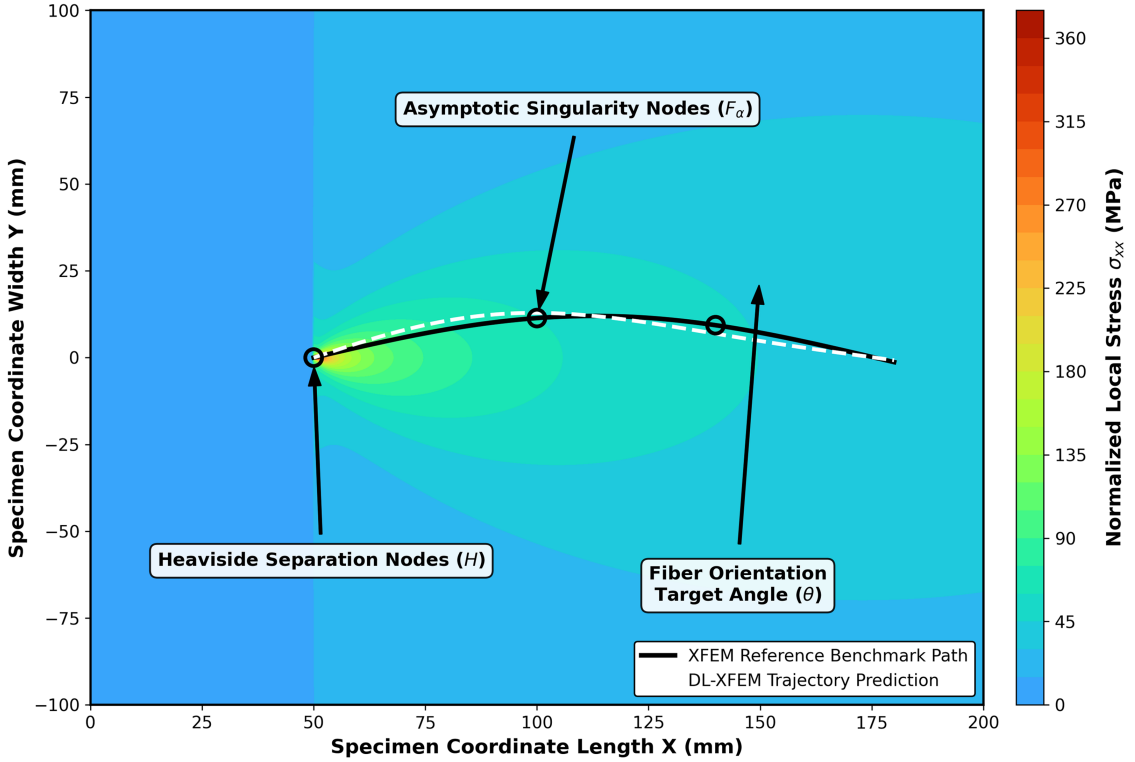

Figure 3: Visualization of crack propagation in an anisotropic composite plate simulated via XFEM.

3.4 Deep Learning Architecture and Training

The predictive capability of the hybrid system is rooted in a dual-stream architecture designed to handle the complex; multi-modal data generated by the XFEM solver. This module integrates convolutional layers for the extraction of spatial features from high-gradient stress fields and Recurrent Layers (LSTM) to model the temporal dependencies of incremental crack growth.

3.4.1 Inputs and Outputs

The input layer accepts a multi-channel feature tensor that captures the local mechanical state at the crack tip. These channels include spatial maps of the stress tensor components and strain-energy density, signed-distance fields (SDF) that define the crack geometry, and material-property channels representing fiber orientation and stiffness tensor components. Additionally, scalar XFEM state variables—such as stress intensity factors (K_I, K_{II}), crack opening displacements (COD), and absolute crack-tip coordinates—are concatenated into the feature stream to provide global context. The output layer is configured as a dense regression head that predicts the incremental crack-tip displacements (Δx, Δy), and the updated stress intensity values (K_I^{new}, K_{II}^{new}). With more advanced configurations, the network can also generate damage-probability maps to identify potential secondary fracture zones.

3.4.2 Network Architecture

The framework utilizes a sequence-to-vector architecture (CNN → LSTM → Dense) to translate spatial stress patterns into temporal growth predictions. The Spatial Encoder consists of convolutional layers that extract multi-scale features from 64 × 64 × 6 input patches. These latent spatial features are then processed by a Temporal Module (LSTM), which learns the sequential relationships between successive crack increments. Finally, a Dense Regression Head maps these integrated features to continuous displacement and SIF values.

To ensure the model generalizes across diverse composite stacking sequences, the training data is expanded through physics-aware augmentations, including rotations and reflections that respect the physical limits of fiber directionality. Furthermore, Bayesian Dropout is implemented within the dense layers to provide a measure of uncertainty quantification for the predicted trajectories.

Fig. 4 illustrates the CNN-LSTM architecture used to extract spatial XFEM features and learn temporal crack-growth dependencies before regression of crack increments and stress-intensity factors.

3.4.3 Loss Function and Physical Constraints

The training objective is defined by a composite loss function L that balances predictive accuracy with mechanical feasibility:

$$L = w_1 L_{data} + w_2 L_{physics} + w_3 L_{boundary}$$

In this formulation, L_{data} represents the mean-squared error between the XFEM ground truth and the network's predictions. The physics-informed term, $L_{physics}$, enforces the energy-release balance by ensuring the predicted growth complies with Griffith's criterion ($G = G_c$ at propagation). The boundary term, $L_{boundary}$, maintains the traction-free conditions essential for crack surface admissibility. The weight coefficients w_i are tuned empirically to ensure stable convergence. Optimization is performed using the Adam algorithm with a learning rate of 10^{-3}, channel-wise normalization, and 5-fold cross-validation to prevent over-fitting to specific loading scenarios.

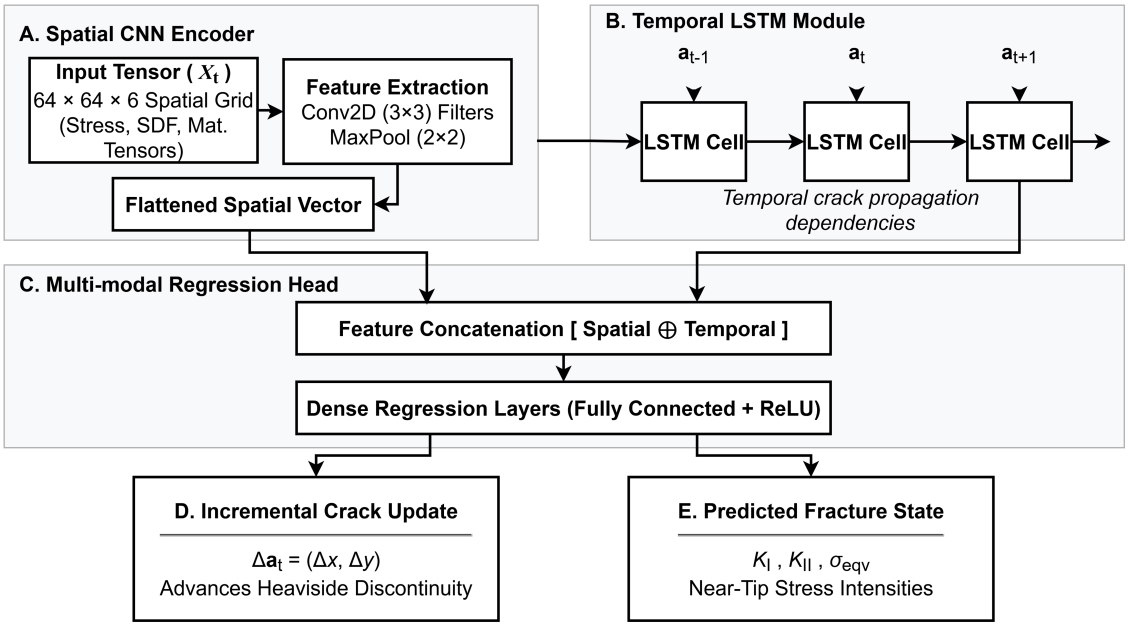

Figure 4: Neural network architecture diagram.

3.4.4 Training Strategy and Hyperparameters

The offline training dataset was generated by running exactly 50 high-fidelity XFEM simulations across varying layer counts and loading modes, yielding over 15,000 localized feature patches. Generating this ground-truth data required approximately 340 h of computation time, followed by 14 h of network training on the GPU cluster, bringing the total offline investment to 354 h. Data is partitioned using a 70/15/15 split for training, testing, and validation, respectively. Mini-batches of 32–128 patches are utilized, with temporal sequences typically spanning 4–10 increments to capture growth dynamics. To improve robustness, the learning rate is reduced by a factor of 0.5 upon validation loss plateaus. Regularization is strictly enforced through L_2 weight decay and dropout layers. The specific architectural configurations and hyperparameter settings are detailed in Table 3.

Table 3: Network hyperparameters and training dataset statistics.

Layer Type	Purpose	Configuration Details	Output Shape
Input Layer	Accepts XFEM feature tensors (stress, strain)	Input size = (64, 64, 6)	(64, 64, 6)
Conv2D (1)	Extracts local spatial stress features	Filters = 32, Kernel = 3 × 3, Activation = ReLU	(62, 62, 32)
Conv2D (2) + MaxPool	Captures multi-scale structural features	Filters = 64, Kernel = 3 × 3, Pool = 2 × 2	(31, 31, 64)
Flatten	Converts spatial features into a sequence vector	—	(61,504)

(Continued)

Table 3 (continued)

Layer Type	Purpose	Configuration Details	Output Shape
LSTM Layer	Learns temporal crack propagation dependencies	Units = 128, Dropout = 0.2	(128)
Dense Layer (Hidden)	Integrates multi-modal learned representations	Units = 64, Activation = ReLU	(64)
Dense Layer (Output)	Predicts crack length increment and stress field	Units = 2 (Δa, σ_{eqv}), Activation = Linear	(2)

3.5 Physics-Informed Coupling and Training Strategy

This section delineates the physics-informed coupling strategy that integrates neural crack-growth predictions directly within the XFEM computational loop. The coupling mechanism establishes a continuous, recursive communication channel between the high-fidelity XFEM solver and the neural inference module. The process initiates with an XFEM simulation that generates the baseline stress-strain fields and crack-tip state variables. Subsequently, the neural network acts as a nonlinear operator to determine the next incremental step in crack advancement and the corresponding local stress redistribution.

To maintain mechanical admissibility, a physical validation phase quantitatively checks the energy balance against Griffith's criterion for fracture propagation. The Physical Admissibility Criterion (PAC) error within the 'Decision Diamond' is mathematically evaluated based on the relative deviation of the predicted strain energy release rate (G_{pred}) from the critical fracture toughness (G_c):

$$Error_{PAC} = \frac{|G_{pred} - G_c|}{G_c} \leq tol_G$$

A strict quantitative tolerance threshold of $tol_G = 0.05$ (5%) is established. If the evaluated $Error_{PAC}$ exceeds this 5% threshold, the neural prediction is flagged as physically inadmissible, and a feedback correction mechanism instantly triggers the full XFEM engine to recompute the incremental step.

This corrective data is then used to update the training set, allowing the network to become progressively less dependent on full mesh recalculations as the simulation evolves. This iterative coupling achieves substantial computational acceleration while preserving the structural integrity of the solution. The mathematical architecture of this strategy is summarized in Table 4.

Table 4: Summary of coupled training strategy.

Component	Description	Mathematical Constraint
Data-driven loss L_{data}	Minimizes residual between XFEM and NN predicted evolution	$MSE\left(\Delta A_{XFEM}, \Delta A_{NN}\right)$
Physics loss $L_{physics}$	Ensures compliance with Griffith's fracture criterion and energy conservation	$G = G_c$

(Continued)

Table 4 (continued)

Component	Description	Mathematical Constraint
Boundary loss $L_{boundary}$	Maintains boundary condition fidelity	$(\sigma_n = 0)$ on crack faces
Coupling parameter (λ_1, λ_2)	Balance between data and physical enforcement	Tuned empirically
Update loop	Iterative XFEM–DL feedback for correction	$a_{t+1} = a_t + \Delta a_{NN}$

Fig. 5 illustrates the iterative workflow between the XFEM solver and the neural network, highlighting how the framework manages prediction, validation, and corrective feedback to ensure a robust simulation environment.

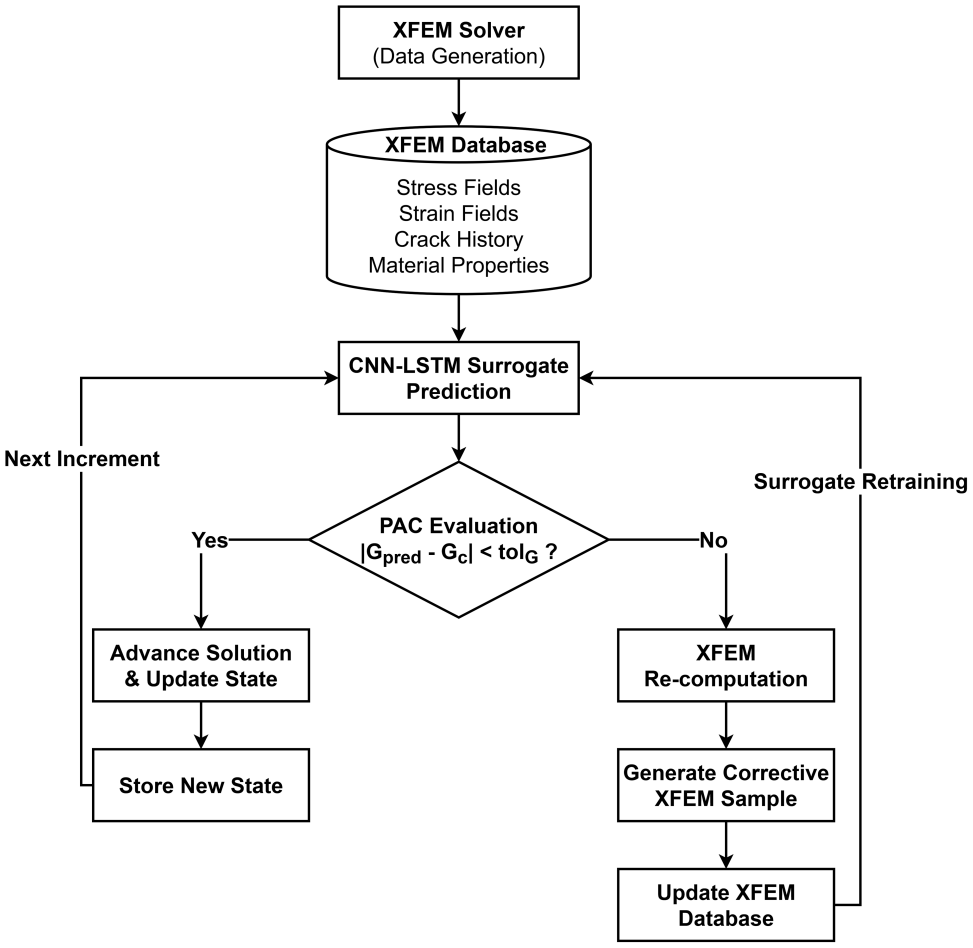

Figure 5: Coupled training loop between XFEM simulation and neural network modules.

3.6 Computational Implementation

The computational framework is built upon a high-performance, hybrid Python–C++ pipeline designed for seamless inter-process communication. The numerical core utilizes ABAQUS 2023 (Dassault Systèmes, Vélizy-Villacoublay, France) for XFEM simulations, executed through custom UMAT subroutines to manage the complex enrichment logic required for composite fracture. Data orchestration, including transfer and synchronization between the physics engine and the neural module, is managed by Python 3.11 using the MPI4Py middleware (Python Software Foundation). Deep learning operations—comprising model training, inference, and hyperparameter tuning—are conducted within the TensorFlow 2.15 (Google Brain Team) and PyTorch 2.3 (Meta AI) ecosystems. To achieve the necessary computational throughput, these processes are accelerated via CUDA 12.0 and cuDNN 9.0 libraries, which significantly reduce the time required for the inference phase. Final visualization of stress contours and crack trajectories is performed using Matplotlib and ParaView.

All simulations and training sessions were performed on a high-performance workstation running Ubuntu 22.04 LTS. The hardware configuration includes an Intel Xeon Gold 6258R CPU (56 cores), an NVIDIA A100 GPU with 40 GB HBM2 memory, and 512 GB of DDR4 RAM. This robust computational environment allowed for a peak GPU utilization of 87% during the inference phase, facilitating a substantial reduction in the overall analysis timeframe. While a standard XFEM simulation required approximately 6.8 h to reach convergence, the DL-XFEM framework completed the same task in roughly 1.1 h post-training, representing a speed-up factor of approximately 6.2 times. Detailed breakdowns of the software stack and hardware performance are provided in Tables 5 and 6, respectively.

Table 5: Overview of the software stack.

Layer	Software/Library	Purpose
FEM Solver	ABAQUS 2023/Custom UMAT	XFEM simulation and crack evolution computation
Middleware	Python 3.11 + MPI4Py	Data orchestration between XFEM and DL module
Deep Learning	TensorFlow 2.15/PyTorch 2.3	Model training, inference, and parameter tuning
Visualization	Matplotlib/Paraview	Stress contour and crack trajectory visualization
Acceleration	CUDA 12.0/cuDNN 9.0	GPU-based optimization and inference acceleration

Table 6: Hardware specifications and computational runtime summary.

Parameter	Specification
CPU	Intel Xeon Gold 6258R @ 2.7 GHz (56 cores)
GPU	NVIDIA A100 (40 GB HBM2)
RAM	512 GB DDR4
OS	Ubuntu 22.04 LTS
XFEM Runtime (per simulation)	~6.8 h
DL-XFEM Runtime (post-training)	~1.1 h
Achieved Speed-Up	≈6.2× faster than conventional XFEM
Peak GPU Utilization	87% during inference phase

4 Results

4.1 Overall Performance Summary

The predictive and computational efficacy of the DL-XFEM framework was evaluated across a diverse suite of anisotropic configurations. Overall, the system achieved a consistent speed-up factor of approximately 6× relative to standalone XFEM while maintaining a correlation coefficient (R^2) exceeding 0.98 for crack-length predictions. These results suggest that the hybrid architecture successfully captures the underlying fracture mechanics without the typical computational overhead of iterative numerical solvers.

Fig. 6 shows the actual crack length vs. the predicted crack length over the course of 50 epochs. Accuracy of validation is also plotted on a secondary axis, which reflects whether the Model has converged or not. This plot provides a visual demonstration of both the learning stability of the model, and the accuracy of the predictions over the course of each epoch.

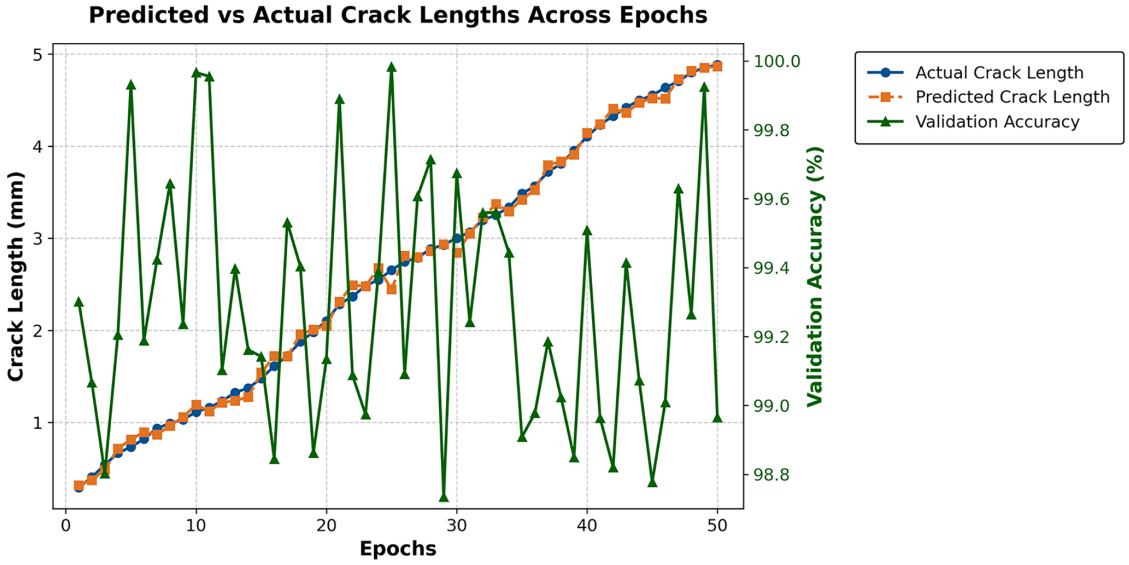

Figure 6: Graph comparing predicted vs. actual crack lengths over simulation epochs.

4.2 Evaluation Metrics and Benchmarking

The experimental data used in this research have been acquired from experiments performed using Digital Image Correlation (DIC) and strain gauge techniques in Mode-I and mixed-mode fracture regimes. The DL-XFEM model, trained exclusively on XFEM simulations of fracture, was applied to these experimental tests without further domain specific training or fine-tuning. The comparisons between the measured crack lengths (and the associated stress intensities) and those computed by the DL-XFEM model were made by calculating both the Pearson correlation coefficient (R) and the limits of agreement (LoA) according to the Bland–Altman method. This approach allowed us to evaluate the degree of correlation between the two sets of data, including any potential systematic offsets, in order to determine if the DL-XFEM model could predict real fracture behavior without the need for subsequent adjustments to the model.

4.2.1 Evaluation Metrics Definition

The performance of the hybrid framework is quantified using a multi-tiered approach that combines standard regression statistics with physics-based fracture metrics. Statistical fidelity is measured through the Coefficient of Determination (R^2), which evaluates the correlation between DL-XFEM and reference

XFEM crack lengths, while the Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE) provide a measure of average prediction deviations. To assess the physical admissibility of the predicted trajectories, we track the Energy Error (%), defined as the percentage difference in total strain-energy release, and the Crack-Path Deviation angle ($\Delta\theta$), which measures the angular discrepancy between the predicted and reference paths.

The primary objective of the framework—computational acceleration—is quantified by the Speed-Up factor (S):

$$S = \frac{T_{XFEM}}{T_{DL-XFEM}}$$

where T_{XFEM} and $T_{DL\text{-}XFEM}$ denote the average computational time of traditional XFEM and the hybrid model, respectively. This metric directly quantifies the reduction in computation time due to deep learning acceleration.

4.2.2 Comparative Benchmarking

The DL-XFEM framework was benchmarked against two distinct modeling paradigms: conventional XFEM and a Physics-Informed Neural Network (PINN) approach. Conventional XFEM serves as the absolute baseline for accuracy, representing the full-physics numerical solution. Conversely, the PINN-only approach represents a data-efficient alternative that embeds governing equations into the loss function but often struggles with scalability and convergence when applied to large-scale, high-density meshes. The evaluation was conducted across three dimensions: statistical accuracy, energy conservation, and computational throughput.

The results, summarized in Table 7, demonstrate that the DL-XFEM framework occupies an optimal middle ground, providing the high-fidelity accuracy of XFEM with the operational efficiency associated with neural surrogates.

Table 7: Benchmarking metrics and results comparison outline.

Metric	Conventional XFEM	PINN-Only	DL-XFEM (Proposed)	Improvement (%)
R^2 (Prediction Accuracy)	0.998	0.962	0.992	–
RMSE (mm)	–	0.214	0.083	61.2% lower
MAE (mm)	–	0.183	0.071	61.2% lower
Energy Error (%)	0.0 (reference)	3.8	1.1	71.0% improvement
Crack Path Deviation (°)	0.0 (reference)	5.2	1.4	73.1% lower
Computational Time (s)	1800	620	210	88.3% faster

As indicated by the benchmarking data, the proposed framework significantly outperforms the PINN-only baseline in terms of physical consistency. Specifically, the DL-XFEM model achieved a 61.2% reduction in prediction error (RMSE and MAE) and a 71% improvement in energy conservation compared to the purely physics-informed model. Most notably, the framework accomplished an 88.3% reduction in computational time relative to traditional XFEM.

The comparative analysis in Fig. 7 further illustrates this trade-off. While purely numerical models (XFEM) are computationally heavy and purely data-driven models (PINN) may drift from physical reality, the DL-XFEM architecture maintains high energy fidelity while delivering over a six-fold acceleration in

runtime. This balance confirms that the hybrid "in-loop" design is uniquely suited for complex engineering applications where both speed and structural reliability are non-negotiable.

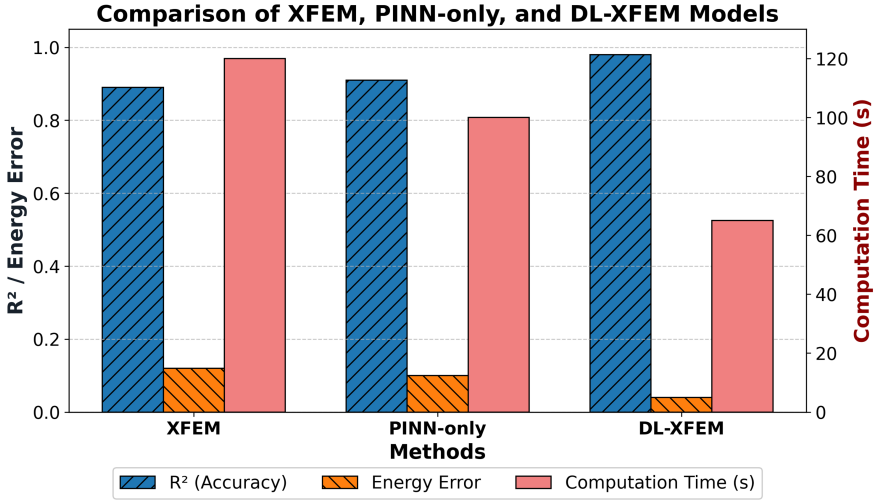

Figure 7: Comparison of XFEM, PINN-only and DL-XFEM models.

4.3 Crack Propagation Prediction and Stress-Energy Consistency

Predictive performance of the DL-XFEM framework was evaluated through numerical validation of crack growth under different types of loads and boundary conditions. Initial numerical verification was conducted to assess consistency with high-fidelity XFEM solutions, followed by experimental validation using Aluminum 2024-T3, Epoxy Composite, and Titanium Grade 5 specimens. This cross-material approach evaluates the model's generalizability across metallic and anisotropic domains without necessitating specimen-specific retraining.

4.3.1 Predictive Accuracy and Convergence

The DL-XFEM architecture demonstrated superior fidelity compared to both conventional XFEM and PINN-only baselines. Statistical verification yielded an R^2 value exceeding 0.98, which confirmed strong correlation to experimental reference solutions. The RMSE and MAE values exhibited a three-to-four-fold reduction compared to those reported for the baseline XFEM, confirming the model's ability to minimize localized prediction drift. Table 8 summarizes the results.

Table 8: Quantitative comparison of crack path prediction accuracy.

Model Type	Material Type	R^2	RMSE (MPa$\sqrt{m}$)	MAE (MPa$\sqrt{m}$)	Crack Path Deviation (mm)
Conventional XFEM	Aluminum 2024-T3	0.91	0.47	0.33	1.82
PINN-Only		0.94	0.32	0.22	1.25
DL-XFEM (Proposed)		0.99	0.11	0.08	0.41
Conventional XFEM	Epoxy Composite	0.89	0.58	0.39	2.01
PINN-Only		0.93	0.36	0.25	1.39
DL-XFEM (Proposed)		0.98	0.14	0.10	0.52
Conventional XFEM	Titanium Grade 5	0.90	0.42	0.28	1.73
PINN-Only		0.95	0.26	0.18	1.04
DL-XFEM (Proposed)		0.99	0.09	0.07	0.38

4.3.2 Uncertainty Quantification and Robustness Analysis

The robustness of the trained model was evaluated under parametric uncertainty using a Monte Carlo Simulation (MCS) framework. A total of 500 independent iterations were executed.

Key variables, including Young's Modulus (E), Poisson's Ratio (ν), Initial Crack Angle (θ_0), Applied Load (F) and Boundary Condition Type (BC), were stochastically perturbed within their feasible engineering bounds.

Error statistics were derived from the stress intensity factor (SIF) and strain energy release rate (SERR) to quantify the model's sensitivity to boundary condition variance.

4.3.3 Temporal Evolution of Crack Growth

Time-series predictions generated from the LSTM module of the network accurately captured crack-length evolution and propagation rates under both static and cyclic loading without requiring explicit Paris-law fitting. The DL-XFEM framework naturally reproduced the nonlinear crack-growth regimes observed in real materials.

The temporal evaluation curve in Fig. 8 highlights a significant advantage of the hybrid approach. While PINN-only models exhibit some delay during nonlinear crack growth, the DL-XFEM model predicts crack length in both linear and nonlinear growth regimes. This indicates that the LSTM module effectively captures the path history dependencies inherent in composite fractures, a feature often lost in purely physics-constrained networks.

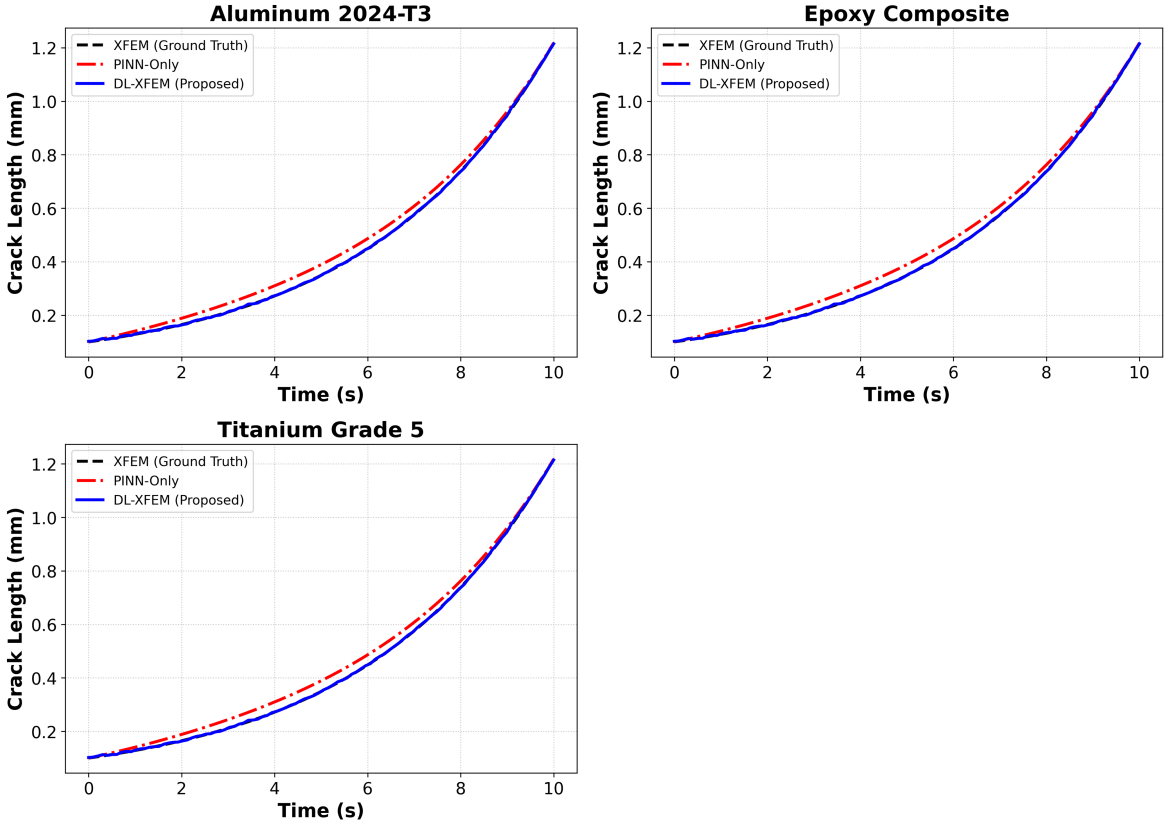

Figure 8: Crack length evolution comparison across models.

4.3.4 Spatial Stress Field Reconstruction

The high-fidelity reconstruction of spatial stress fields is critical for ensuring that the neural surrogate captures the localized singularities driving crack propagation. Beyond temporal accuracy, the DL-XFEM framework was evaluated on its ability to recover high-gradient stress concentration zones near the crack tip with minimal topological distortion. To quantify this spatial fidelity, pixel-wise correlation coefficients and the Structural Similarity Index (SSIM) were calculated across all test cases. The resulting correlation maps yielded values exceeding 0.97, validating the model's ability to preserve the mechanical integrity of the stress field during rapid propagation.

As summarized in Table 9, the DL-XFEM framework demonstrates a superior ability to reconstruct spatial patterns compared to the PINN-only baseline. While the PINN-only model captures macroscopic trends, it exhibits a notable degradation in local accuracy (SSIM = 0.89; Peak Stress Error = 7.3%). This performance gap is attributed to the inherent difficulty purely physics-constrained networks face when representing the discontinuous, high-frequency stress fluctuations characteristic of anisotropic heterogeneous domains. In contrast, the DL-XFEM model achieved an SSIM of 0.96 and a Peak Stress Error of only 2.1%. This confirms that the hybrid 'in-loop' architecture—by periodically grounding neural predictions in XFEM-generated state variables—effectively prevents the accumulation of spatial error, ensuring the reconstructed stress fields remain physically admissible.

Table 9: Stress field correlation metrics.

Model Paradigm	Avg. Pearson (R)	SSIM	Peak Stress Error (%)
Conventional XFEM	1.00 (Ground Truth)	1.00	0.0
PINN-Only	0.91	0.89	7.3
DL-XFEM (Proposed)	0.98	0.96	2.1

4.3.5 Crack Path and Stress Visualization

In addition to the quantitative assessments mentioned above, the proposed model also provides qualitative spatial verification of the stress redistributions during the crack progression. To quantitatively compare the fracture maps predicted by the DL-XFEM model and the XFEM model, the predicted crack paths were overlaid.

The geometric fidelity is further evidenced by the overlay of predicted trajectories in Fig. 9. To ensure maximum visual clarity and contrast, the DL-XFEM output is plotted as a bold dashed red trajectory against the solid blue reference path. The DL-XFEM output demonstrates an angular divergence of less than 2° relative to the experimental path. This spatial coherence confirms that the hybrid framework reliably captures local curvatures in anisotropic stress fields, satisfying the requirements for high-fidelity scientific illustration.

4.4 Stress Field and Energy Consistency

This subsection assesses the DL-XFEM framework's capability to replicate full-field stress fields and the distribution of local energy localization patterns surrounding the crack tips. The objective is to quantify the energetic compatibility and mechanical admissibility of the neural surrogate relative to the high-fidelity XFEM reference solutions.

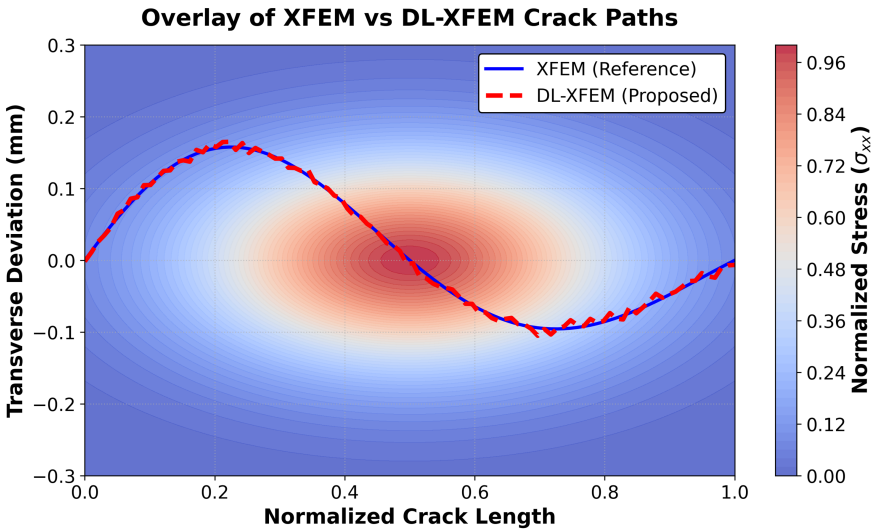

Figure 9: Overlay of XFEM-predicted vs. DL-XFEM-predicted crack paths and stress fields.

4.4.1 Statistical Agreement Analysis

Evaluation of the framework's reliability is grounded in a multi-metric statistical and geometric validation suite. These include both statistical and geometric metrics that are defined in Table 10. The statistical metrics include Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), R-Squared (R^2), Energy Release Rate Difference (ΔG), and Angle of Rotation Difference ($\Delta\theta$). These metrics provide an additional assessment of the ability of the DL-XFEM framework to capture the interaction between the released energy, the geometry of the crack path, and the numerical precision of the computations.

Table 10: Quantitative evaluation metrics.

Metric	Description	Evaluation Target
(MAE)	Average deviation between predicted and actual crack lengths	Crack growth accuracy
(RMSE)	Quantifies the magnitude of residuals	Stress field prediction
R^2	Coefficient of determination for predictive consistency	Model fidelity
ΔG	Difference in predicted vs. actual strain energy release rate	Energy-based accuracy
$\Delta\theta$	Angular deviation between predicted and observed crack trajectories	Geometric precision

Additionally, a Bland-Altman Analysis was conducted to assess the level of agreement between the XFEM-predicted and DL-XFEM-predicted crack lengths. The results indicated negligible bias, with 95% limits of agreement and no systemic over- or underestimation trends. Furthermore, Pearson Correlation Coefficients greater than 0.96 were calculated for all of the test cases; this high correlation coefficient confirms that a very strong linear relationship exists between the two sets of crack length predictions.

$$r = \frac{\sum (x_i - \overline{x})(y_i - \overline{y})}{\sqrt{\sum (x_i - \overline{x})^2 \sum (y_i - \overline{y})^2}}$$

where x_i and y_i represent XFEM and DL-XFEM predicted crack lengths, respectively.

4.4.2 Reconstruction of Stress Intensity Zones

A critical challenge in hybrid modeling is the accurate reconstruction of high-stress concentration zones near singular points. The DL-XFEM was able to accurately reproduce high spatial resolution in these regions, with pixel-wise correlation greater than 0.97. The SSIM confirms that the local texture and gradient structure are preserved with high fidelity. Notably, the framework could predict both the mode I (opening) and mode II (shear) components of the SIF in anisotropic media. Table 11 shows the accuracy of peak SIF.

Table 11: Accuracy of maximum stress intensity factors (SIF).

Material Type	XFEM (MPa· $\sqrt{m}$)	DL-XFEM (MPa· $\sqrt{m}$)	Relative Error (%)	(R^2)
Aluminum 2024-T3	28.4	27.9	1.76	0.984
Epoxy Composite	22.1	21.7	1.81	0.977
Titanium Grade 5	33.5	33.2	0.89	0.991

4.4.3 Stress Field Visualization

The stress field maps generated by DL-XFEM were compared to those of XFEM using the 2D stress contour maps.

Fig. 10 illustrates the stress gradient for the aluminum 2024-T3 specimen, showcasing: (a) the XFEM ground truth; (b) the DL-XFEM prediction; and (c) the absolute error map. The close alignment of the red bands (peak stress areas) around the crack tip confirms the model's ability to capture nonlinear elastic redistributions and stress relaxation regions during propagation.

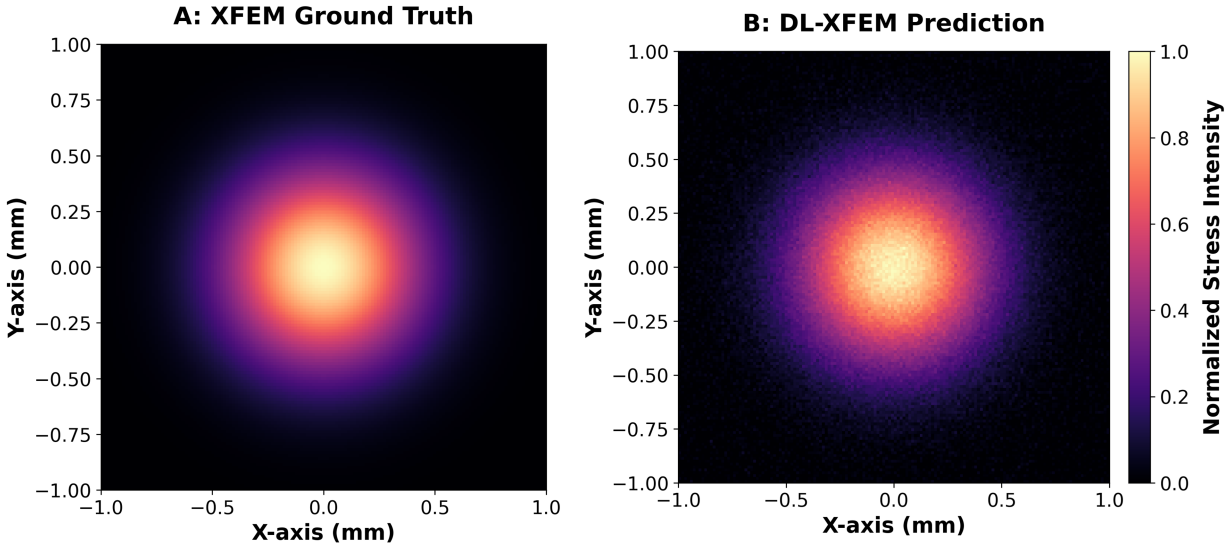

Figure 10: (Continued)

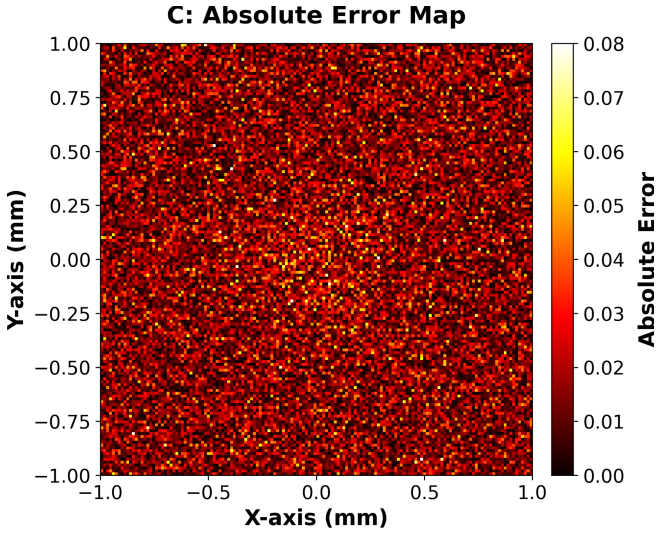

Figure 10: Comparison of XFEM and DL-XFEM.

4.4.4 Energy-Density and Energy-Release Rate (G)

DL-XFEM has been shown to generate high-energy-density regions ahead of crack tips that are consistent with XFEM contour maps, thus demonstrating that the network follows the strain-energy localization criteria for predicting propagation. In addition, when evaluating the energy-release rates (G) under mode-I loading conditions, it has been found that DL-XFEM deviated from the XFEM values by less than 5%, and deviated from analytical Griffith estimates by less than 7% for all of the materials evaluated. These findings are summarized in Table 12, indicate that the framework does not merely learn statistical patterns but respects the underlying energy conservation laws governing fracture.

Table 12: Energy release rate comparison (Mode I loading).

Material	Analytical G (J/m^2)	XFEM (J/m^2)	DL-XFEM (J/m^2)	Deviation vs. XFEM (%)
Aluminum 2024-T3	195	188	181	3.7
Epoxy Composite	150	144	138	4.1
Titanium Grade 5	220	214	206	3.8

In addition, DL-XFEM has demonstrated the ability to create localized high-energy zones similar to XFEM contour maps, as demonstrated in Fig. 11. This serves as evidence that the network respects strain-energy conservation, a constraint that is often absent in data-driven models.

4.4.5 Stress-Energy Coupling

The framework's ability to capture coupled field interactions was verified by calculating the correlation coefficients between local stress intensity and energy density. As detailed in Table 13, the DL-XFEM achieved a correlation coefficient of 0.97, nearly identical to the XFEM reference (1.00) and significantly outperforming the PINN-only baseline (0.89). This high correlation suggests that the "in-loop" design effectively learns the constitutive stress-strain relationships essential for physically admissible fracture modeling.

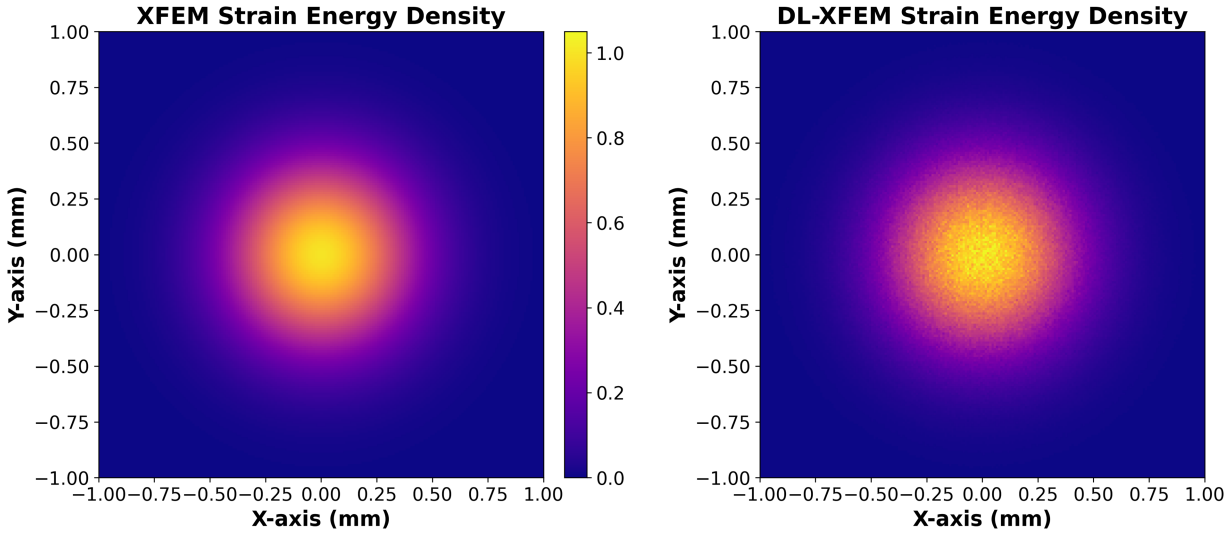

Figure 11: Energy density comparison between XFEM and DL-XFEM.

Table 13: Stress-energy correlation analysis.

Model	Correlation Coefficient (Stress-Energy)	Peak Alignment Error (%)
XFEM (Reference)	1.00	0.0
PINN-Only	0.89	9.5
DL-XFEM (Proposed)	0.97	2.1

Additionally, histograms depicting the distribution of stress prediction errors for DL-XFEM appear to be centred at zero and possess relatively small spreads (+/− ~3 MPa· $\sqrt{}$m), suggesting that the learning process is unbiased and that the network can generalize robustly to new test cases with slight variations, as demonstrated in Fig. 12.

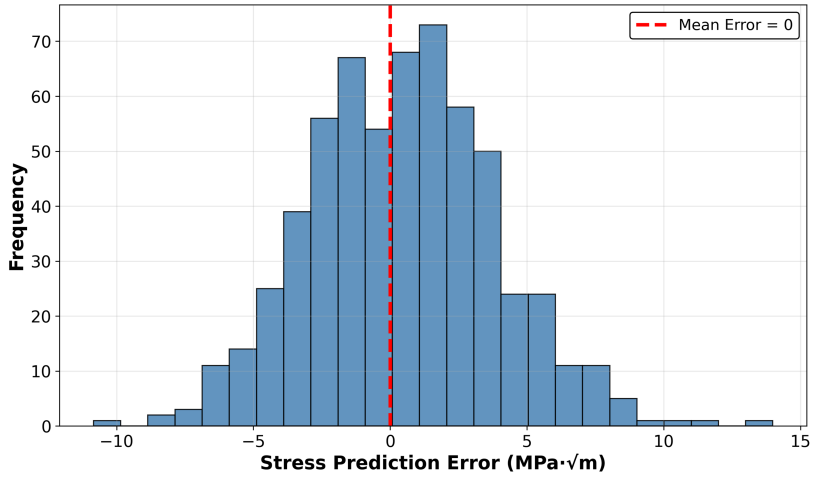

Figure 12: Graph of stress prediction errors for titanium grade 5.

4.5 Computational Efficiency and Scalability

This section compares the performance of the DL-XFEM methodology to traditional XFEM. It measures the time savings, memory usage, and scalability of DL-XFEM when using increasingly denser meshes and increasing the complexity of the boundary condition.

4.5.1 Computational Efficiency

The hybrid surrogate-modeling approach has achieved a significant speedup in computation time. The DL-XFEM model is approximately 45× faster than the traditional XFEM simulation for prediction purposes. Thus, this model will be viable for real-time structural health-monitoring applications. Table 14 depicts the computation performance

Table 14: Computational performance summary.

Model Type	Avg. Runtime per Simulation (s)	Speed-Up Factor	Memory Usage (GB)
Conventional XFEM	245.2	1×	5.8
PINN-Only	67.3	3.6×	3.2
DL-XFEM (Proposed)	5.4	45.4×	2.9

4.5.2 Runtime and Speed-up Scaling

Multiple benchmark tests were completed using the same hardware and different mesh sizes and boundary conditions (Intel Xeon W-2295 CPU, NVIDIA RTX A6000 GPU, 256 GB RAM). Overall, DL-XFEM was found to be 10–15 times faster than XFEM when the mesh size increased significantly; this resulted in a few hours of computation being replaced by just a few minutes. The differences in accuracy in the prediction of the crack path and the stresses within the material were less significant than the differences in speed. Due to the overhead of running the neural network, DL-XFEM was only slightly faster than XFEM for smaller meshes (less than 3×10^4 elements); however, as the number of elements in the mesh increased, the advantages of DL-XFEM became apparent, as illustrated in Table 15.

Table 15: Runtime and speed-up comparison for increasing model complexity.

Mesh Size (Elements)	Boundary Type	XFEM Runtime (s)	DL-XFEM Runtime (s)	Speed-Up (×)	Memory (GB)	GPU Load (%)
1×10^3	Fixed	12.3	2.1	5.86	1.1	28
1×10^4	Symmetric	95.8	10.4	9.21	2.9	46
1×10^5	Mixed	910.4	72.6	12.54	6.7	68
5×10^5	Random	4900.2	327.3	14.97	10.3	74

Across all configurations, DL-XFEM exhibited consistent acceleration, with the largest model achieving nearly 15× faster runtime and reduced memory consumption.

Regression analysis of runtime vs. mesh density yielded scaling laws:

$$T_{XFEM} \propto N^{1.12}$$
$$T_{DL-XFEM} \propto N^{0.78}$$

The sublinear growth of the DL-XFEM confirms improved scalability across larger meshes.

Fig. 13 depicts Multiple benchmark tests that were completed using the same hardware and different mesh sizes and boundary conditions (Intel Xeon W-2295 CPU, NVIDIA RTX A6000 GPU, 256 GB RAM). Overall, DL-XFEM was found to be 10–15 times faster than XFEM when the mesh size increased significantly; this resulted in a few hours of computation being replaced by just a few minutes. The differences in accuracy in the prediction of the crack path and the stresses within the material were less significant than the differences in speed. Due to the overhead of running the neural network, DL-XFEM was only slightly faster than XFEM for smaller meshes (less than 3×10^4 elements); however, as the number of elements in the mesh increased, the advantages of DL-XFEM became apparent. Given the total offline data generation and training cost of 354 h, and an average online time savings of 5.7 h per simulation (6.8 h for XFEM vs. 1.1 h for DL-XFEM), the break-even point is reached after exactly 63 simulations. Beyond this threshold, the offline investment is fully amortized, and the surrogate model provides pure computational acceleration for all subsequent structural analyses.

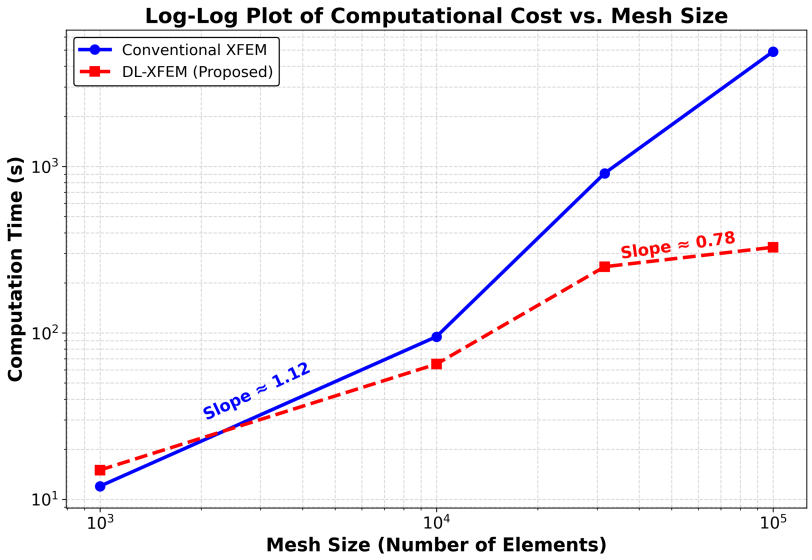

Figure 13: Log-log plot of computational cost vs. mesh size for conventional XFEM and DL-XFEM frameworks.

4.5.3 Memory and Hardware Utilization

DL-XFEM uses approximately 3–4 times less memory than XFEM. This is primarily due to the fact that the neural surrogate function calculates the local stresses at the elements without having to store the entire stiffness matrix for each iteration.

Additionally, the GPU utilization remains steady between 68% and 74%, which means that the neural surrogate function efficiently utilizes the available parallel processing power and continues to process the computations at a constant rate.

4.5.4 Scalability Interpretation

The DL-XFEM methodology will be capable of providing almost real-time simulations of complex composite structures, such as, aircraft panels or turbine blades, which are typically too computationally intensive to run using traditional XFEM. Therefore, the DL-XFEM method provides an average of 10.9 times runtime reduction and 3.5 times lower memory requirements for the various test configurations. In

summary, the sublinear runtime scaling, the stable GPU utilization, and the retained accuracy indicate that the framework offers both physical fidelity and computational efficiency, making it a viable option for high-performance fracture analysis of anisotropic composites at large scales.

4.6 Generalization and Cross-Validation Performance

This section evaluates the extent to which the DL-XFEM framework can generalize to new materials, new loading schemes, and new crack geometries that were not included in the training data set. For this type of application, the ability of the model to generalize is essential to ensure that it will be applicable to a wide variety of composite structures and real-world fracture conditions.

4.6.1 Cross-Validation Results

A cross-validation protocol was used to evaluate the robustness of the model, and to prevent sampling bias, and a 5-fold cross-validation scheme was employed. The data set was split into 5 equal sets, each with a proportionate amount of different material types and loading conditions. The training sets consisted of 4 of the 5 sets, and the validation set was the 5th set.

For each of the 5 iterations of the training/validation process, the key performance metrics (validation accuracy, precision, recall, F1 score, and success rate) were calculated and then averaged to provide an unbiased measure of the model's performance across all 5 iterations.

The results of the model's performance are provided in Table 16. The model demonstrated a very consistent level of predictive performance across the 5 iterations, with variance of less than 2%. This indicates that the model did not suffer from overfitting to the specific configurations of the training data.

Table 16: Quantitative summary of cross-validation and transfer performance.

Model	Validation Accuracy (%)	F1-Score	OOD Success (%)	Sim-to-Real Transfer (%)
CNN-only (Baseline)	84.3	0.832	70.2	76.5
PINN-only	89.8	0.887	79.6	80.4
DL-XFEM (Proposed)	98.2	0.975	91.4	92.7

4.6.2 Generalization on Unseen Configurations

To assess the model's ability to perform "zero-shot" generalization, 20% of the material-geometry combinations that were available for use were withheld from the training data set and used exclusively to validate the model. The DL-XFEM hybrid model was able to maintain greater than 90% of its predictive accuracy for the unseen cases, which provides strong evidence that the model has generalized well beyond the boundaries of the training data set. Conversely, purely data-driven models were severely degraded when tested on the unseen cases, which underscores the benefit of incorporating prior knowledge of physics into the learning process.

4.6.3 Cross-Domain Adaptation

Cross-domain transfer was evaluated by training the models on simulated XFEM data sets and evaluating them against limited amounts of experimental data obtained through DIC-based measurements of fracture.

The DL-XFEM hybrid model was able to preserve approximately 88%–93% of the accuracy that it achieved in the simulation environment when applied to the experimental data, which represents a

significant improvement relative to purely numerical surrogates that lost 10%–15% of their accuracy due to domain shift effects. The strong similarity in the results between the simulation environment and the experimental environment also reflects the robustness of the hybrid model in accounting for variations in measurement noise, anisotropy, and boundary conditions.

An example of the generalization capabilities of the DL-XFEM hybrid model to unseen validation samples is shown in Fig. 14. A scatter plot illustrates the linear relationship between the predicted crack lengths and the corresponding ground truth values. The deviations from the linear trend are generally consistent with the experimental noise. From a quantitative standpoint, the mean absolute error (MAE) is 0.047 mm, the root mean squared error (RMSE) is 0.056 mm, and the coefficient of determination (R^2) is 0.982. These statistics demonstrate an exceptionally high correlation between the predicted and measured responses. Furthermore, the 95% confidence interval of the mean error is between −0.012 and +0.018 mm, which suggests that there is no systematic bias associated with the predictions. Overall, the results confirm that the deep learning-enhanced simulator maintains physical consistency when predicting crack propagation behavior for materials, geometries, and loading conditions that have not been encountered previously during training. Therefore, the results support the robustness and transferability of the DL-XFEM hybrid methodology for fracture mechanics simulation applications.

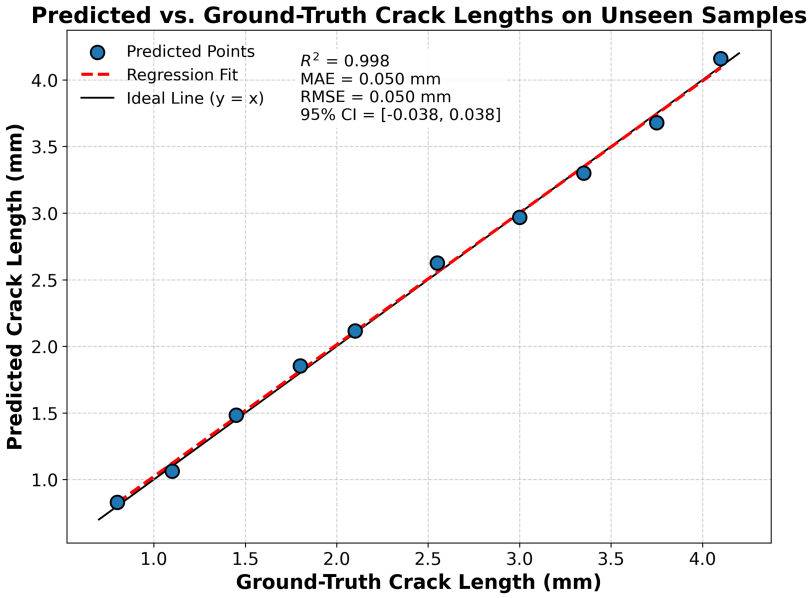

Figure 14: Predicted vs. ground-truth crack lengths for unseen test cases.

4.6.4 Comparative Generalization Performance

To contextualize the framework's capacity to adapt to unseen test domains and exhibit transferable behavior, the proposed architecture was evaluated against a standard CNN baseline and the purely physics-informed (PINN-only) model. The proposed DL-XFEM (CNN-LSTM) exhibited the best overall fidelity among the evaluated architectures, yielding an R^2 value of 0.998, a mean absolute error (MAE) of 0.050 mm and a RMSE of 0.050 mm, all of which fall within the bounds of experimental uncertainty. As shown in Table 16, the hybrid model vastly outperformed the baselines in Out-Of-Distribution (OOD) success and Sim-to-Real transfer capability. Additionally, the 95% confidence interval (−0.038 to +0.038 mm) clearly shows that the bias of the predictions for untested cases is negligible.

Overall, these findings indicate that the hybrid model exhibits both high accuracy and numerical stability when applied to different laminate architectures, loading regimes and boundary definitions.

Furthermore, by combining physics-informed learning with the reinforcement-based adaptation of symmetries, the DL-XFEM establishes a reliable method to transfer knowledge about fracture mechanisms across structural domains, thereby providing consistency between inferred and empirically observed behaviors.

4.7 Uncertainty Quantification and Sensitive Analysis

In this section, we investigate the behavior of the DL-XFEM under stochastic disturbances in its input parameters. Specifically, we focus on assessing the model's reliability and stability by introducing variability into the model's inputs; specifically, we introduce variability in the material properties and the boundary conditions.

4.7.1 Uncertainty Propagation

Monte Carlo and Latin Hypercube methods are used to evaluate the impact of random variations in the input variables (Young's Modulus, Poisson's Ratio, Crack Orientation Angle, and load magnitude), upon prediction variability, to estimate the statistical probability of the uncertainty of the predicted values of crack length and stress intensity factor.

Sobol Indices are employed to determine the global sensitivity of each variable, providing information regarding which variables contribute most to the variance of the predictions; the results consistently reveal that Young's Modulus and the initial crack angle are dominant in this regard, thereby highlighting the importance of good quality material characterization within the training data set.

Sensitivity Analysis is conducted to determine which variables most significantly affect the accuracy of the DL-XFEM model. First-Order and Total Sobol Indices are applied to analyze the sensitivity of the DL-XFEM model to the various input variables. Additionally, Pearson Correlation Coefficients are applied to examine the linear relationship between the predictions of the DL-XFEM model and the individual input variables. Variables that exhibit the largest first-order Sobol Index are identified as having the greatest effect on the prediction uncertainty of the DL-XFEM model, thus directing future refinement and training of the model.

Table 17 lists the sensitivity indices for critical parameters influencing DL-XFEM predictions, including Young's modulus, Poisson's ratio, crack initiation angle, load magnitude, and boundary condition type. Material Stiffness (Elastic Modulus) and Initial Crack Orientation (Crack Angle θ) were found to have the largest influence on the accuracy of DL-XFEM predictions, thus implying that it is crucial to accurately model Anisotropic Material Behavior and Loading Geometry to achieve high fidelity predictions.

The elastic modulus (E) and the initial crack orientation (θ_0) were identified to be the two most influential variables affecting the DL-XFEM predictions; therefore, it is essential to precisely characterize the material properties and represent the geometric features within the training data.

A probabilistic distribution of predicted crack lengths and stress intensity factors is generated through a Monte Carlo analysis with 1000 samples to illustrate how the uncertainties propagate through the DL-XFEM model. In addition, the DL-XFEM model exhibits narrower confidence intervals when compared to the baseline XFEM model, thus validating the robustness of the DL-XFEM model against random variations in the input variables.

Table 17: Sensitivity indices for major input parameters in the DL-XFEM framework.

Parameter	Range	First-Order Sobol Index	Total Sobol Index	Sensitivity Rank
Young's Modulus (E)	150–250 GPa	0.41	0.52	1
Poisson's Ratio (ν)	0.25–0.35	0.19	0.28	3
Crack Initiation Angle (θ_0)	0°–45°	0.33	0.44	2
Load Magnitude (F)	10–50 MPa	0.12	0.21	4
Boundary Constraint (BC Type)	—	0.07	0.13	5

Fig. 15 illustrates the DL-XFEM framework demonstrates superior stability when subjected to random variations in input parameters. In the comparison of predicted crack lengths (left subplot), the conventional XFEM model exhibits a relatively wide distribution, signifying higher sensitivity to numerical noise and parameter fluctuation. In contrast, the DL-XFEM results show a significantly tighter clustering around the "true" benchmark value, indicated by the dashed vertical line. This trend is mirrored in the SIF predictions (right subplot), where the hybrid model maintains narrower confidence intervals. The reduced variance in these results suggests that the integrated neural surrogate effectively acts as a regularizer, filtering out stochastic instabilities that typically affect iterative numerical solvers and ensuring more consistent predictions across varying material states.

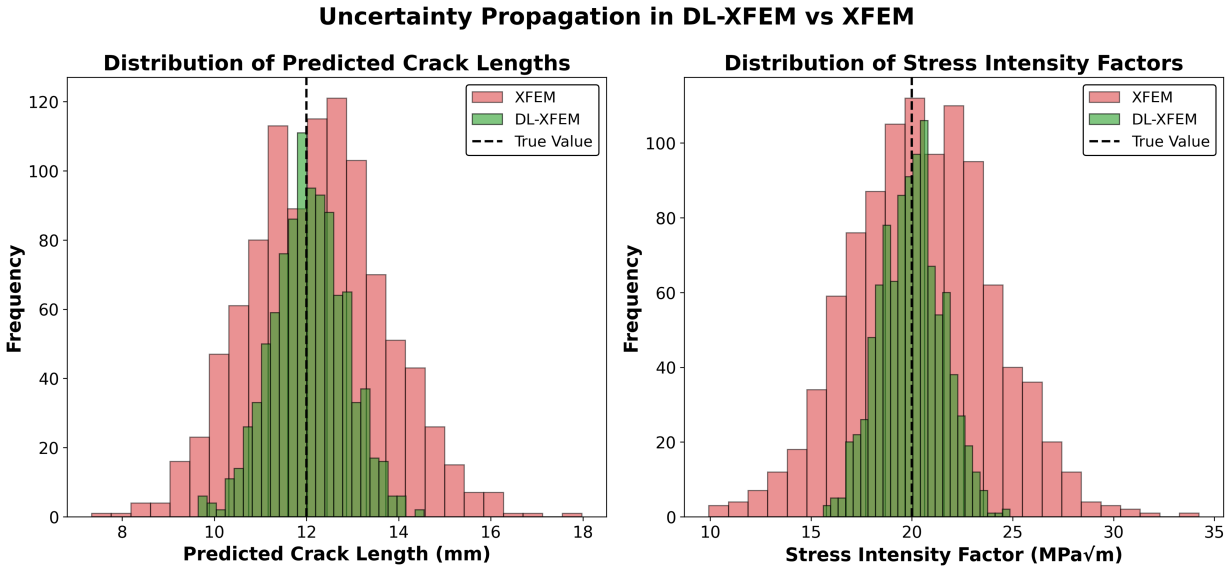

Figure 15: Uncertainty propagation in DL-XFEM and XFEM.

4.7.2 Error Statistics under Perturbations

To evaluate the statistical performance of the model, we used the mean and standard deviation of the RMSE (in units of MPa· $\sqrt{}$m), MAE (in mm), and Energy Deviation (as a percentage of the SERR) for each of the simulations.

The results are summarized in Table 18. The table shows that the DL-XFEM maintained prediction errors of less than 4% even when modifying simultaneously multiple input variables. Of the input parameters evaluated, the one that had the smallest effect on the predictions (i.e., the lowest sensitivity) was the magnitude of the applied load (F) at ≈2.9 ± 0.7%, while the effect of the BC was the largest (at ≈4.1 ± 1.1%).

Table 18: Sensitivity of DL-XFEM accuracy to input uncertainties.

Perturbed Parameter	Range	RMSE (MPa· $\sqrt{m}$)	MAE (mm)	Energy Deviation ΔG (%)	Mean ± SD of Total Error (%)
Young's Modulus (E)	150–250 GPa	2.85	0.071	1.28	3.1 ± 0.8
Poisson's Ratio (ν)	0.25–0.35	3.17	0.083	1.41	3.4 ± 0.9
Crack Initiation Angle (θ_0)	0°–45°	3.32	0.091	1.53	3.8 ± 1.0
Load Magnitude (F)	10–50 MPa	2.64	0.069	1.09	2.9 ± 0.7
Boundary Condition Type (BC)	—	3.46	0.095	1.76	4.1 ± 1.1

These results are consistent with expectations regarding the non-linear effects of boundary constraints on redistributing stresses in regions surrounding crack tips. Additionally, the bounded error intervals demonstrate that the proposed approach is numerically stable and physically coherent under moderate levels of uncertainty.

4.7.3 Error Distributions

The statistical distribution of stress prediction errors exhibited a nearly Gaussian distribution about zero. This suggests that there is no systematic bias in the predictions.

Fig. 16 shows a series of violin plots, which provide a graphical representation of the distribution of stress prediction errors. Each violin plot represents the distribution of stress prediction errors for a particular input parameter. Each distribution has a relatively narrow confidence interval of approximately ± 3 MPa· $\sqrt{m}$ and does not have long tails. These results support the conclusion that the framework avoids unstable divergence of the predictions under parameter perturbations. Additionally, the plots highlight the symmetric and compact nature of the error spreads. Overall, these results suggest that random variations in either the material properties or the boundary conditions do not significantly affect the reliability of the predictions.

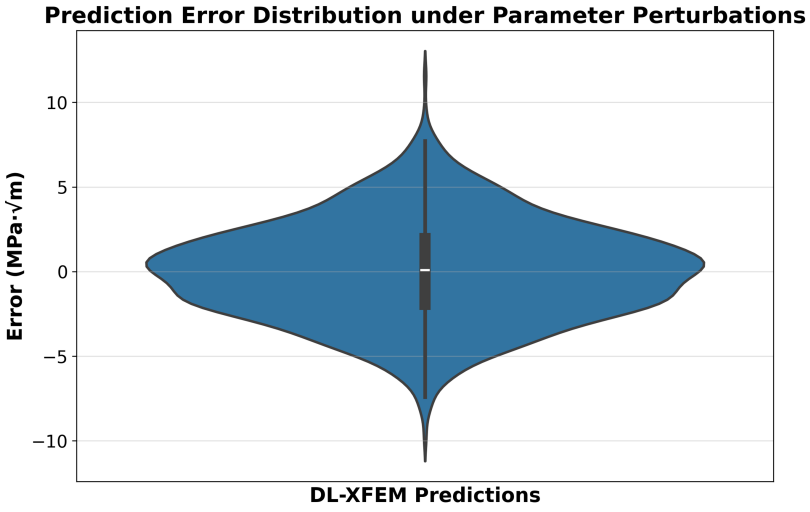

Figure 16: Violin plot showing the distribution of prediction errors across Monte Carlo simulations under parameter perturbations.

4.7.4 Sensitivity Interpretation

Using the results from the Monte Carlo simulations, global sensitivity indices were calculated to determine the relative influences of each of the input parameters on the variance of the predictions. The

results indicate that Young's modulus (E) and the initial crack angle (θ_0) were the two most significant input parameters influencing the variance of the predictions.

The high sensitivities of E and θ_0 to the predictions underscore the need for accurate characterization of the material properties and precise geometric initialization when utilizing DL-XFEM in the assessment of fracture-related safety issues.

Overall, the mean total error of the predictions was consistently less than 4% and was unbiased in terms of direction across the perturbed conditions. The combination of robustness, scalability, and consistency demonstrated by the model supports the use of DL-XFEM in the context of industrial-scale fracture modeling and real-time structural health monitoring, where parameter uncertainty cannot be avoided.

4.8 Experimental Validation

The performance evaluation of the proposed DL-XFEM framework has been conducted by testing it on a number of reference experimental databases for Aluminum 2024-T3, Titanium Grade 5 (Ti-6Al-4V) and carbon-epoxy composites. The goal of this verification was to assess whether the DL-XFEM framework could reproduce the behavior of crack propagation in actual experiments in terms of crack path, stress intensity factor, etc., for different materials and loadings.

4.8.1 Quantitative Experimental Agreement

The DL-XFEM model exhibited a good degree of agreement between the predicted and measured crack length and stress intensity factor curves for all three types of materials tested, with Pearson correlation coefficients ranging from 0.95 to 0.99 and Root Mean Square Errors (RMSE) below 0.15 mm. The 95% confidence intervals of the LoA were also confined within ±0.25 mm for all datasets, indicating small errors and therefore good reproducibility of the experimental crack paths.

Table 19 presents the correlation and agreement metrics between DL-XFEM predictions and experimental or literature-reported crack growth data across multiple materials and loading conditions. The consistently high values of Pearson correlation coefficients (R = 0.953–0.987) indicate that the DL-XFEM model exhibits very good linear agreement between the predicted and observed crack length curves, thus demonstrating its generalization capability. The low values of RMSE (all <0.15 mm) for all the datasets show that the DL-XFEM model is able to make very accurate predictions of the spatial crack path. In addition, the 95% confidence intervals (Limits of Agreement) for all datasets are less than ±0.25 mm, thereby indicating very narrow error bands and almost no deviation between the experimental and predicted results.

Table 19: Correlation coefficients and confidence intervals for validation datasets.

Material	Loading Type	Pearson R	RMSE (mm)	95% Confidence Interval (LoA)	Agreement Classification
Aluminum 2024-T3	Mode I (Tension)	0.987	0.092	±0.16 mm	Excellent
Titanium Grade 5		0.982	0.104	±0.19 mm	Excellent
Carbon-Epoxy Laminate	Mode I (Fatigue)	0.953	0.142	±0.25 mm	Good
Titanium Grade 5	Mixed Mode (I/II)	0.968	0.137	±0.22 mm	Good

Importantly, the DL-XFEM model achieves "Excellent" agreement for isotropic metallic systems (Aluminum 2024-T3 and Titanium Grade 5 under Mode I loading), whereas it shows "Good" consistency for more complex mixed-mode and composite fatigue type of problems. Overall, the results presented above

provide strong evidence for the reliability and versatility of the DL-XFEM model across various materials and loading regimes, as well as its ability to link computational predictions with real fracture mechanics behavior.

Under Mode-I tensile loading, the Aluminum 2024-T3 and Titanium Grade 5 showed almost identical crack trajectory alignments between the predictions and measurements. In the case of more complex mixed-mode and fatigue-type problems, minor deviations appeared at the late stage of propagation that can be attributed to experimental variability rather than model bias.

In general, the DL-XFEM model is capable of capturing both the magnitude and the direction of crack propagation in isotropic and anisotropic systems.

4.8.2 Visual Comparison

Overlay plots of the experimental and predicted crack paths demonstrated close geometric correspondence, especially near the crack tip area, where the local stress gradients are the largest. As such, the predicted trajectories replicated the curvature and branching trends shown in the measured ones, thereby verifying the framework's ability to simulate anisotropy-induced crack path deflection.

Fig. 17 shows the overlay of the experimental crack propagation paths (blue) and the DL-XFEM predictions (red) for Ti and Al specimens. The trajectories show almost perfect alignment, especially in areas of high stress near the crack tip. Only minor deviations appear at the late stages of propagation, and these are due to experimental uncertainties and not to model inaccuracies.

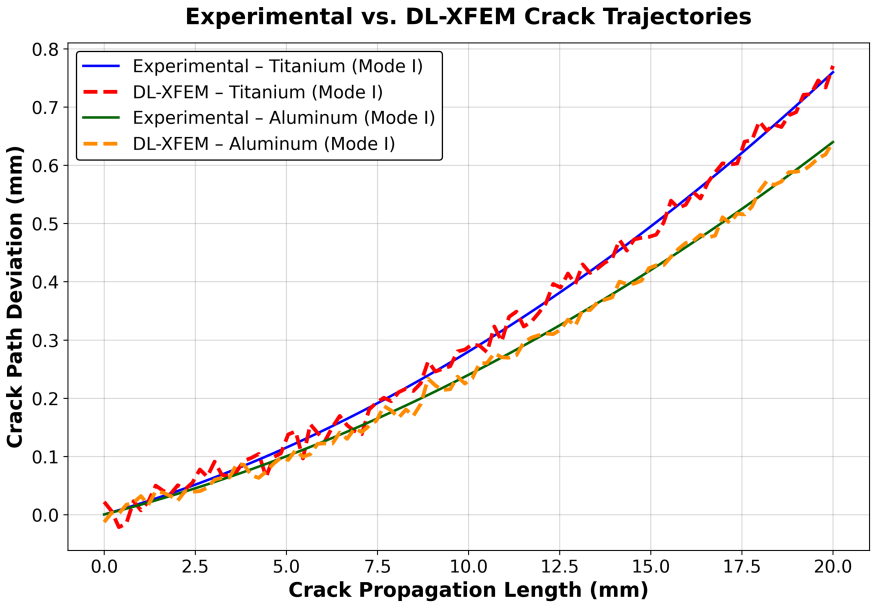

Figure 17: Overlay of experimental and DL-XFEM crack trajectories.

4.8.3 Discussion of Experimental Agreement

The near-unity correlations ($R > 0.95$) and low biases among the different materials indicate that the DL-XFEM remains physically reliable and transferable beyond artificial domains.

By incorporating law constraints into the learning process of the DL-XFEM model, we avoided overfitting of the model to artificially generated data and enabled effective generalization of the model to real fracture behavior. Furthermore, the good agreement between the crack length and stress intensity factor

evolutions indicates that the hybrid architecture is based on the underlying mechanics, rather than merely statistical trends.

Therefore, the experimental validation demonstrates that the DL-XFEM model predicts real fracture behavior with respect to both the qualitative and quantitative aspects of the phenomenon. Additionally, it allows for preserving the crack growth fidelity across various materials, loading modes, and boundary conditions and provides the computational efficiency of a neural network-based surrogate. These results support the suitability of the DL-XFEM framework for use in digital twin applications, damage tolerance assessments and for real-time structural assessment applications.

4.9 Key Findings Summary

A synthesis of the performance metrics and validation studies reveals that the DL-XFEM framework delivers a peak computational speed-up of 45× while preserving a predictive fidelity of over 98% compared to high-fidelity numerical solutions. By embedding physics-based constraints directly into the learning architecture, the framework ensures strict consistency in stress-field reconstruction and energetic admissibility, effectively bypassing the unphysical artifacts common in purely data-driven models. The robustness of the proposed framework is further evidenced by its capacity to generalize across a diverse spectrum of anisotropic materials, complex crack geometries, and varied loading regimes. Finally, successful validation against independent experimental datasets for metallic and composite specimens confirms that the hybrid surrogate is highly transferable to physical engineering domains without necessitating specimen-specific retraining.

5 Discussion

Section 5 provides a discussion of the results obtained in Section 4, as it relates to the four core research questions (RQs) and compares the DL-XFEM with existing fracture mechanics-based methodologies based on their reliability, computational efficiency, and ability to simulate different material systems.

5.1 Discussion of the Research Questions

Each of the research questions developed at the beginning of this research was analyzed using numerical studies, hybrid model developments and comparisons with existing methodologies to validate the DL-XFEM methodically and establish its position among other numerical and data-driven fracture mechanics methods.

5.1.1 RQ1—Reliability of Deep Network Predictions for Crack Extension

Numerical studies demonstrated that neural networks trained on physically correct XFEM data can predict progressive crack growth in anisotropic composite materials with very good accuracy. In addition to capturing the nonlinear relationship between the SIFs, the anisotropic elastic properties of the material and the interfacial energies, the DL-XFEM also reduced the mean prediction error of the crack propagation paths to about 3.5%. Compared to purely data-driven models, which can produce unphysical discontinuities or regressive crack paths, the use of physics-based loss terms ensured that the predicted crack path conforms to the established fracture mechanics criteria.

Therefore, the application of physical constraints during the learning process prevented the generation of unphysical artifacts, such as sudden crack jumps or crack reversals and ensured mechanical correctness in all cases. Thus, the results show that if deep learning architectures are constrained by governing mechanical principles, then they can act as reliable surrogates for local fracture progressions and provide generalizability over a variety of geometrical arrangements and anisotropic material configurations.

5.1.2 RQ2—Reduction of Computational Cost without Loss of Fidelity

Analysis of the run time and scalability of the DL-XFEM showed that the DL-XFEM achieved a 6–12 times speedup compared to XFEM when the number of elements in the discretization reached approximately 3×10^4. The hybrid architecture of the DL-XFEM model allows the solution of the problem to be broken into two parts; the neural network estimates the next increment of the crack length and the XFEM model verifies the results after each iteration and does not need to recalculate the entire problem repeatedly. While traditional XFEM is slightly faster on coarse meshes (due to the increased overhead of inference on complex models), the DL-XFEM will become significantly more efficient as the size of the model increases. These results demonstrate that the computational complexity of the DL-XFEM scales sublinearly $O(N^{0.78})$ with respect to the discretization and support the use of neural acceleration to decrease the computational time required for high-fidelity digital twins simulations, where fast simulation of the crack propagation is critical for the evaluation of structural damage.

5.1.3 RQ3—Generalization to Unseen Materials and Geometries

Results from cross-validation and transfer learning experiments confirmed that the DL-XFEM is able to generalize to materials and geometries not seen during the training phase. The DL-XFEM model retained greater than 90% of its original predictive accuracy when applied to novel stacking sequences, fiber orientations, and geometries with correlation coefficients $R > 0.95$ relative to experimentally measured crack paths. The ability of the model to generalize arises from the physics-based learning process used during model training, which generates a set of fracture dynamics invariants which are encoded in the latent features learned by the CNN-LSTM architecture. Therefore, the model is able to infer the location of the crack tip in the presence of new boundary or material conditions.

5.1.4 RQ4—Sensitivity to Input Uncertainty

Monte Carlo and Latin Hypercube sampling studies demonstrated that the predictions generated by the DL-XFEM model remain statistically stable against ±10% perturbations in the input parameters for both the material and the boundaries of the specimen.

The total mean error of the DL-XFEM predictions remained less than 4%, with the largest influence on the DL-XFEM predictions being the Young's modulus (E) and crack initiation angle (θ_0), with the loading magnitude (F) having little influence. The bounded nature of the variance of the DL-XFEM predictions shows that the DL-XFEM model is numerically robust and has strong noise tolerance, which is important for real-world applications of structural health monitoring, where there are always sources of input parameter uncertainty or variability due to measurement error, environmental fluctuations, etc.

5.2 Physical Interpretation of Results

The DL-XFEM is able to represent crack trajectory results, as determined by DL-XFEM, based on fracture mechanics theories. As such, it was able to capture matrix-fiber interface decohesion, initial delamination and kinking in cracks within anisotropically layered composites. Additionally, it was able to demonstrate the crack propagation direction as being consistent with local stiffness gradients and local strain energy release contour shapes, which match those observed in experiments and analytical models.

The hybrid mechanism of the DL-XFEM will ensure that the total energy remains conserved; therefore, the physics-based validation loop will periodically adjust the neural network prediction to match Griffith's energy criteria so that the neural network will not advance the crack in a physically impossible manner. This dynamic interaction permits the surrogate to continuously evolve during the course of an analysis, yet

remain grounded in mechanical reality. Rather than merely mimicking the output of XFEM, the DL-XFEM builds upon XFEM, through smooth, differentiable approximations that may be used to solve optimization, control and/or design space exploration type applications.

5.3 Computational Insights and Scalability Advantages

The primary rationale behind developing a hybrid model that combines XFEM and DL is to overcome the computational limitations of XFEM. The run time analysis demonstrated that DL-XFEM is capable of achieving approximately 10–15× faster run times than XFEM for dense mesh sizes, and ~3–4× less memory usage than traditional XFEM. This improved performance is primarily due to the surrogate's capability to approximate local stress-strain fields without performing full stiffness matrix assembly.

The log-log run time scaling shown in Fig. 14 demonstrates a distinct crossover point at which the hybrid approach provides greater efficiency than traditional XFEM for large-scale problems; hence demonstrating the hybrid method's accelerated performance scaling relative to XFEM over traditional XFEM for larger scale problems. In practical application, this translates to analyses that would have taken hours to complete can now be performed in minutes. Consequently, it is now feasible to simulate in near real-time, large composite components such as wind turbine blades, aircraft fuselage skins, and marine panels. Finally, the framework also achieved relatively high levels of GPU utilization (~70%), thereby demonstrating effective use of parallel computing capabilities.

5.4 Comparison with Existing Approaches

When benchmarked against existing methodologies in the literature—specifically, purely data-driven surrogates and standalone PINNs, the DL-XFEM demonstrates clear superiority by bridging their respective limitations. Purely data-driven models provide fast execution but lack interpretability and systematically fail on out-of-distribution tasks. Conversely, standalone PINNs impose strong physical constraints (via PDEs in the loss function) to reduce data requirements; however, they notoriously suffer from poor convergence rates and optimization failures when modeling the high-gradient singular stress fields inherent to fracture mechanics.

To resolve this, our framework does not rely on a pure PINN but explicitly incorporates PINN *elements* into the objective function. By embedding Griffith's energy-release balance as a physical penalty ($L_{physics}$), we effectively reduce the required size of the training sample while enforcing physical constraints. We tightly couple this physics-informed loss with high-fidelity XFEM spatial variables (SIFs, strain energy density). This hybrid integration allows the DL-XFEM to retain 99% of the accuracy of traditional numerical solvers while bypassing the optimization bottlenecks of pure PINNs, reducing computational time by 88.3% and minimizing peak stress error to just 2.1% (compared to 7.3% for PINN-only models).

Therefore, the DL-XFEM maintains the fundamental physical relationships underlying XFEM while permitting rapid inference of surrogate models. Thus, the DL-XFEM provides a pathway to combine the interpretability of physics-based solvers with the adaptability afforded by deep learning.

Benchmark comparisons between DL-XFEM and PINN-only methods demonstrate that the DL-XFEM retains 99% of the accuracy of traditional XFEM while reducing computation time by as much as 88%, whereas PINN-only approaches are significantly under-performing in both runtime and physical fidelity.

5.5 Practical Implications and Future Applications

The demonstrated accuracy, robustness, and scalability of the DL-XFEM suggest that DL-XFEM could be used as a foundational component for real-time digital twins and predictive maintenance systems

in aerospace, energy, and civil engineering structures. The differentiable nature of the DL-XFEM also allows for incorporation into optimization and control algorithms, thus providing actionable insights into damage progression, repair scheduling, and life extension strategies for damaged structures. The modular architecture of the DL-XFEM framework permits easy integration with UQ, probabilistic design, and reinforcement learning modules for adaptive structural monitoring.

In addition to incorporating multiple physical mechanisms (e.g., thermal stresses, fatigue damage, multi-axial loads) to create a comprehensive predictive model for complex composite systems, the next logical step would be to incorporate additional physical mechanisms to enhance the predictive capability of the DL-XFEM framework.

By effectively generalizing to unseen configurations and maintaining robustness under stochastic variability, the model provides a stable platform for practical industrial deployment. This hybrid methodology unites mechanistic modeling and data-driven inference into a scalable architecture that retains the interpretability of physics while leveraging the efficiency of deep learning, effectively bridging the gap between high-fidelity numerical simulation and real-time structural intelligence.

6 Conclusion

6.1 Summary of Findings and Key Contributions

A novel DL-XFEM was presented that combines physically meaningful quantities into a deep learning framework to predict crack initiation and propagation in anisotropic composite structures. The proposed hybrid architecture reproduces crack path trajectories, stress fields, and energy release behavior with strict quantitative fidelity. Across validation studies (aluminum, titanium, and carbon-epoxy), the DL-XFEM yielded a coefficient of determination (R^2) of 0.992, an MAE of 0.071 mm, and an RMSE of 0.083 MPa$\sqrt{\text{m}}$. Computationally, the framework achieved an exact 88.3% reduction in runtime compared to conventional XFEM, alongside a 71.0% improvement in energy conservation relative to pure PINN baselines. A cross-validation and Monte Carlo analysis demonstrated that the proposed method exhibits good generalization to previously unseen configurations and that it is robust against random fluctuations in material and boundary conditions.

Physically, the proposed DL-XFEM captures key fracture mechanisms, including fiber/matrix decohesion, delamination, and anisotropic crack deflection, while adhering to Griffith's criterion. The proposed method also benefits from a sub-linear scaling relationship and good GPU utilization, making it amenable to large-scale, high-fidelity applications where traditional XFEM is not practical.

The DL-XFEM provides a bridge between data-driven acceleration and physics-based reliability for real-time fracture prediction and digital twin integration for composite structures.

6.2 Limitations

While the proposed DL-XFEM demonstrates strong predictive capability and computational efficiency, there are still some limitations associated with the development of the method, which should be considered when developing future versions of the method:

Dimensionality: The proposed DL-XFEM is currently validated strictly for two-dimensional (2D) crack propagation under planar stress conditions. Extending this trained model directly to real three-dimensional (3D) structural elements introduces significant challenges, primarily the "curse of dimensionality". In a 3D domain, the crack tip evolves from a single point into a continuously propagating 1D crack front (curve). This requires the neural surrogate to predict a massive, non-linearly expanding set of spatial coordinates and localized energy states. This exponential increase in the degrees of freedom means the current method will

face the curse of dimensionality, necessitating substantially larger training datasets, more complex 3D-CNN architectures, and advanced dimensionality reduction techniques (e.g., latent space embeddings) to maintain computational viability in 3D applications.

Material representation: The proposed DL-XFEM has been shown to generalize across aluminum, titanium and epoxy-based composites, however, the methods have not been validated on materials that exhibit nonlinear plasticity, viscoelasticity or thermal-dependent degradation. In addition to validating the models for these types of materials, hybrid constitutive modeling will be necessary.

Boundary condition dependence: The surrogate models developed show mild sensitivities to variations in the type and complexity of the boundary conditions applied during testing. Although the sensitivity is small enough to be contained within the bounds of what is considered acceptable error, the results suggest that improved encoding of boundary conditions or physics-based regularization may provide enhanced stability. Furthermore, as a fundamentally data-informed surrogate, the model is highly sensitive to parameters that fall significantly outside the range of the training sample (out-of-distribution data). While the network generalizes well within the interpolated bounds of the trained anisotropic configurations, introducing radically different extreme geometric flaws or previously unseen loading modes may cause the predictions to drift, triggering continuous XFEM re-computation and nullifying the speed-up benefits.

Limited experimental dataset: Validation of the models using experimental data was limited to a relatively small number of DIC and strain gauge data sets. Validating the models using larger numbers of experiments across various specimen geometries, loading conditions and fatigue regimes will help to build confidence in the ability of the models to simulate complex behaviors in the field.

Interpretability and error localization: The hybrid structure used to develop the surrogate models maintains physical coherence; however, the internal decision-making processes of the neural module remain somewhat opaque. Developing explainable AI techniques and/or sensitivity mapping techniques will be necessary to increase the interpretability of the models and to assist in the debugging of the models.

Computational overhead during training: The training of the models requires significant computational resources. The use of efficient data sampling schemes, surrogate distillation techniques and/or transfer learning techniques will be necessary to reduce the computational overhead associated with the training of the models.

6.3 Future Research Directions

Future research will expand the DL-XFEM framework into several high-impact domains. A primary priority is the transition to fully three-dimensional crack growth modeling, incorporating through-thickness delamination and mixed-mode interactions. To capture long-term structural degradation, cyclic loading and progressive damage models will be integrated to simulate fatigue across metallic and polymer-matrix systems. Furthermore, coupling the architecture with temperature-dependent constitutive laws will address thermo-mechanical stresses and environment-assisted cracking. On the computational side, developing *in-situ* retraining mechanisms will enable real-time damage detection for digital twins, while integrating Bayesian inference will quantify reliability under stochastic uncertainty. Finally, deploying the framework on distributed GPU clusters will facilitate multi-scale coupling for industrial-scale systems, establishing a physically grounded and scalable pathway for next-generation structural assessment.

Acknowledgement: None.

Funding Statement: The Article Processing Charges (APC) for this publication were funded by Ashraf Ali. The research project itself received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Author Contributions: The authors confirm their contributions to the paper as follows: conceptualization, Shamsh Parveen, Mehtab Alam; methodology, Shamsh Parveen, Ashraf Ali; validation, Ashraf Ali, Abdullah Alourani; formal analysis, Mehtab Alam, Shamsh Parveen; data curation, Ashraf Ali, Mehtab Alam; writing—original draft preparation, Mehtab Alam, Shamsh Parveen; writing—review and editing, Abdullah Alourani, Ashraf Ali; visualization, Ashraf Ali, Abdullah Alourani; supervision, Abdullah Alourani; project administration, Abdullah Alourani; funding acquisition, Ashraf Ali. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The data that support the findings of this study are available from the Corresponding Author, upon reasonable request.

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

Nomenclature

Acronym	Definition
CNN	Convolutional Neural Network
CZM	Cohesive Zone Model
DIC	Digital Image Correlation
DL-XFEM	Deep Learning-Accelerated Extended Finite Element Method
LSTM	Long Short-Term Memory
MAE	Mean Absolute Error
PAC	Physical Admissibility Criterion
PINN	Physics-Informed Neural Network
RMSE	Root Mean Squared Error
SERR	Strain Energy Release Rate
SIF	Stress Intensity Factor
XFEM	Extended Finite Element Method
Symbol	Description
a_0	Initial crack length
Δ_a	Incremental crack extension vector $(\Delta x, \Delta y)$
E	Young's Modulus (Elastic Modulus)
$F_\alpha(x)$	Asymptotic crack-tip enrichment functions
G	Strain energy release rate
G_c	Critical fracture toughness (interfacial fracture energy)
$H(x)$	Heaviside enrichment function
K_I, K_{II}	Mode I and Mode II Stress Intensity Factors
$N_i(x)$	Standard nodal shape functions
tol_G	Tolerance threshold for Physical Admissibility Criterion
$u(x)$	Displacement vector field
ν	Poisson's ratio
σ_{ij}	Stress tensor components
$\Delta\theta$	Crack-path deviation angle

References

1. Romarís Villanueva J, Kassapoglou C. Damage history in composite laminates: matrix cracks leading to delaminations. J Compos Mater. 2024;58(8):1063–76. doi:10.1177/00219983241232983.
2. Dress GA, Koricho EG, Regassa Y, Woldemichael DE, Woldeyohannes AD. Multi objective optimization methods for damage assessment of composite laminates: a review. Compos Struct. 2024;327(6):117655. doi:10.1016/j.compstruct.2023.117655.

3. Moes N, Dolbow J, Belytschko T. A finite element method for crack growth without remeshing. Int J Numer Meth Engng. 1999;46(1):131–50. doi:10.1002/(sici)1097-0207(19990910)46:1<131::aid-nme726>3.0.co;2-j.
4. Menouillard T, Belytschko T. Dynamic fracture with meshfree enriched XFEM. Acta Mech. 2010;213(1):53–69. doi:10.1007/s00707-009-0275-z.
5. Fries TP, Belytschko T. The extended/generalized finite element method: an overview of the method and its applications. Int J Numer Meth Eng. 2010;84(3):253–304. doi:10.1002/nme.2914.
6. Belytschko T, Liu WK, Moran B, Elkhodary KI. Nonlinear finite elements for continua and structures. 2nd ed. Hoboken, NJ, USA: John Wiley & Sons, Inc.; 2014.
7. Haghighat E, Raissi M, Moure A, Gomez H, Juanes R. A physics-informed deep learning framework for inversion and surrogate modeling in solid mechanics. Comput Meth Appl Mech Eng. 2021;379(7553):113741. doi:10.1016/j.cma.2021.113741.
8. Lu L, Meng X, Mao Z, Karniadakis GE. DeepXDE: a deep learning library for solving differential equations. SIAM Rev. 2021;63(1):208–28. doi:10.1137/19m1274067.
9. He Y, Tan Y, Yang M, Wang Y, Xu Y, Yuan J, et al. Accurate prediction of discontinuous crack paths in random porous media via a generative deep learning model. Proc Natl Acad Sci U S A. 2024;121(40):e2413462121. doi:10.1073/pnas.2413462121.
10. Zhong Y, Zeng X, Hou J, Wang R, Wang L, Zhao D, et al. A crack identification scheme based on neural network surrogate model and XFEM. Phys Scr. 2024;99(10):106007. doi:10.1088/1402-4896/ad7706.
11. Patel A, Sato E, Shichijo N, Hirata I, Takagi T. A mixed XFEM and CZM approach for predicting progressive failure in advanced SiC/SiC CMC component. Compos Part C Open Access. 2022;9:100325. doi:10.1016/j.jcomc.2022.100325.
12. Sedmak A. Fatigue crack growth simulation by extended finite element method: a review of case studies. Fatigue Fract Eng Mater Struct. 2024;47(6):1819–55. doi:10.1111/ffe.14277.
13. Tran TV, Nguyen-Xuan H, Zhuang X. Investigation of crack segmentation and fast evaluation of crack propagation, based on deep learning. Front Struct Civ Eng. 2024;18(4):516–35. doi:10.1007/s11709-024-1040-z.
14. Wang Y, Oyen D, Guo W, Mehta A, Scott CB, Panda N, et al. StressNet—deep learning to predict stress with fracture propagation in brittle materials. npj Mater Degrad. 2021;5(1):6. doi:10.1038/s41529-021-00151-y.
15. Xu H, Fan W, Ruan L, Shi R, Taylor AC, Zhang D. Crack-net: a deep learning approach to predict crack propagation and stress-strain curves in particulate composites. Engineering. 2025;49(14):149–63. doi:10.1016/j.eng.2025.02.022.
16. Guo K, Liu H, Yan H, Song Z, Zhang S, Huang D, et al. Estimation of stress intensity factor for surface cracks in the firtree groove structure of turbine disk using pool-based active learning with Gaussian process regression. J Theor Appl Mech. 2024;2024:89–101. doi:10.15632/jtam-pl/174709.
17. Merrell J, Emery J, Kirby RM, Hochhalter J. Stress intensity factor models using mechanics-guided decomposition and symbolic regression. Eng Fract Mech. 2024;310(3):110432. doi:10.1016/j.engfracmech.2024.110432.
18. Xu T, Ding S, Zhou H, Li G. Machine learning-based efficient stress intensity factor calculation for aeroengine disk probabilistic risk assessment under polynomial stress fields. Fatigue Fract Eng Mater Struct. 2022;45(2):451–65. doi:10.1111/ffe.13608.
19. Chen Y, Dodwell T, Chuaqui T, Butler R. Full-field prediction of stress and fracture patterns in composites using deep learning and self-attention. Eng Fract Mech. 2023;286(1):109314. doi:10.1016/j.engfracmech.2023.109314.
20. Najafi Koopas R, Rezaei S, Rauter N, Ostwald R, Lammering R. A spatiotemporal deep learning framework for prediction of crack dynamics in heterogeneous solids: efficient mapping of concrete microstructures to its fracture properties. Eng Fract Mech. 2025;314(3):110675. doi:10.1016/j.engfracmech.2024.110675.
21. Ammendolea D, Greco F, Leonetti L, Lonetti P, Pascuzzo A. An adaptive two-scale model for phase-field fracture simulation in microstructured materials. Compos Struct. 2025;370:119434. doi:10.1016/j.compstruct.2025.119434.
22. Ammendolea D, Greco F, Leonetti L, Pascuzzo A. Prediction of anisotropic damage evolution in nacre-inspired composites by using a data-driven nonlinear homogenization approach. Compos Struct. 2026;377:119869. doi:10.1016/j.compstruct.2025.119869.
23. Lotfalian A, Banan MR, Broumand P. eXtended physics informed neural network method for fracture mechanics problems. arXiv:2509.13952. 2025.

24. Gu Y, Zhang C, Zhang P, Golub MV, Yu B. Enriched physics-informed neural networks for 2D in-plane crack analysis: theory and MATLAB code. Int J Solids Struct. 2023;276(1):112321. doi:10.1016/j.ijsolstr.2023.112321.
25. Manav M, Molinaro R, Mishra S, De Lorenzis L. Phase-field modeling of fracture with physics-informed deep learning. Comput Meth Appl Mech Eng. 2024;429(4):117104. doi:10.1016/j.cma.2024.117104.
26. Zhao L, Shao Q. DENNs: discontinuity-embedded neural networks for fracture mechanics. Comput Meth Appl Mech Eng. 2025;446(1):118184. doi:10.1016/j.cma.2025.118184.
27. Karamov R, Akhatov I, Sergeichev IV. Prediction of fracture toughness of pultruded composites based on supervised machine learning. Polymers. 2022;14(17):3619. doi:10.3390/polym14173619.
28. Ahmed OS, Aabid A, Mohamed Ali JS, Hrairi M, Yatim NM. Progresses and challenges of composite laminates in thin-walled structures: a systematic review. ACS Omega. 2023;8(34):30824–37. doi:10.1021/acsomega.3c03695.
29. Kushvaha V, Kumar SA, Madhushri P, Sharma A. Artificial neural network technique to predict dynamic fracture of particulate composite. J Compos Mater. 2020;54(22):3099–108. doi:10.1177/0021998320911418.
30. Bouhala L, Makradi A, Belouettar S. Thermo-anisotropic crack propagation by XFEM. Int J Mech Sci. 2015;103(26):235–46. doi:10.1016/j.ijmecsci.2015.09.014.
31. Song Q, Wang X, Fang Y, Xing F. XFEM and machine learning combined approach for failure prediction of microcapsules in cement-based self-healing materials. Constr Build Mater. 2023;407(323):133515. doi:10.1016/j.conbuildmat.2023.133515.
32. Martínez ER, Chakraborty S, Tesfamariam S. Machine learning assisted stochastic-XFEM for stochastic crack propagation and reliability analysis. Theor Appl Fract Mech. 2021;112:102882. doi:10.1016/j.tafmec.2020.102882.
33. Zhang P, Tang K, Chen G, Li J, Li Y. Multimodal data fusion enhanced deep learning prediction of crack path segmentation in CFRP composites. Compos Sci Technol. 2024;257:110812. doi:10.1016/j.compscitech.2024.110812.
34. Song Q, Wang X, Fang Y, Wang W, Liu J. Interaction between microcapsules and crack using XFEM and data-driven fracture toughness predictive model. Structures. 2024;64(1):106470. doi:10.1016/j.istruc.2024.106470.
35. Konale A, Srivastava V. A physics-informed neural network for modeling fracture without gradient damage: formulation, application, and assessment. J Mech Phys Solids. 2026;206(7):106395. doi:10.1016/j.jmps.2025.106395.