ORIGINAL ARTICLE

lajss.org© Latin American Journal of Solids and Structures

www.lajss.org

Physics-Informed Machine Learning for Crashworthiness Prediction of Hexagonal Crash Boxes under Multi-Angle Impacts

Sittha Tongthong[a], **Suphanut Kongwat**[b], **Pattaramon Jongpradist**[a,b*]

[a] Department of Mechanical Engineering, Faculty of Engineering, King Mongkut's University of Technology Thonburi (KMUTT), Bangkok, 10140, Thailand. E-mails: sittha.tongt@kmutt.ac.th; pattaramon.tan@kmutt.ac.th

[b] Future Automotive Structure Research Group (FASt), Mobility and Vehicle Technology Research Center, King Mongkut's University of Technology Thonburi (KMUTT), Bangkok, 10140, Thailand. E-mail: k.suphanut@gmail.com

*Corresponding author

https://orcid.org/10.1590/1679-7825/e9235

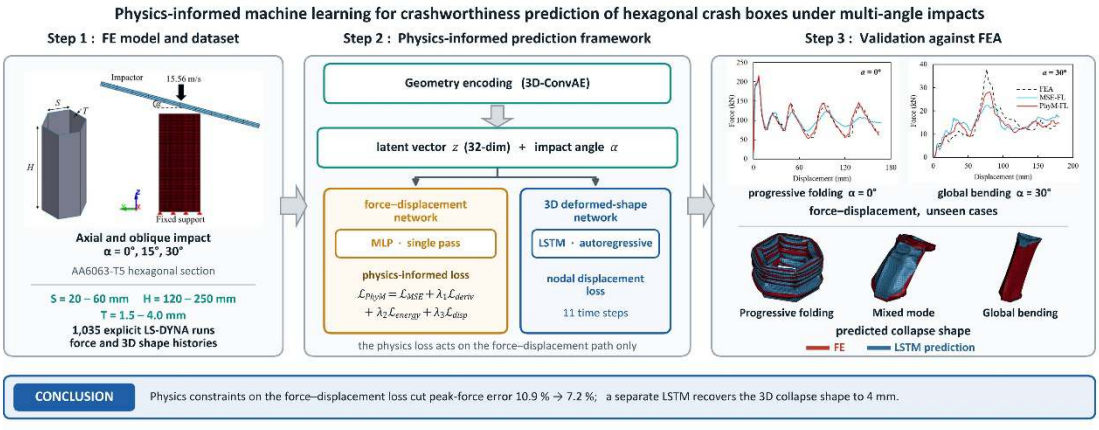
Abstract

A physics-informed machine learning framework is developed to predict temporal force–displacement responses and three-dimensional deformation histories of hexagonal crash boxes under multi-angle loading. To reduce the computational cost of nonlinear explicit finite element simulations, a three-dimensional convolutional autoencoder (3D-ConvAE) encodes crash box geometries into a low-dimensional latent representation, while a long short-term memory (LSTM) network predicts their temporal deformation evolution. A physics-informed loss function incorporating kinematic admissibility and energy conservation is introduced to suppress non-physical responses and improve mechanical consistency. Validation against benchmark numerical data demonstrates high predictive accuracy, with mean errors of 7.24% for peak force and 6.31% for energy absorption. The framework also maintains geometric fidelity across complex failure transitions, including the shift from progressive folding to global bending, achieving a mean Chamfer distance of 4 mm. The proposed approach provides a robust and computationally efficient tool for predicting nonlinear crash responses and supporting crashworthiness optimization.

Keywords

Crashworthiness, Physics-informed machine learning, Hexagonal crash boxes, Spatio-temporal prediction, Oblique impact.

Graphical abstract



Received August 19, 2026. In revised form September 02, 2026. Accepted September 09, 2026. Available online September 11, 2026.
https://doi.org/10.1590/1679-7825/e9235

Latin American Journal of Solids and Structures. ISSN 1679-7825. Copyright © **2026**. This is an Open Access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1 INTRODUCTION

Passenger safety in automotive engineering largely depends on the effectiveness of the vehicle's crumple zone in managing kinetic energy during frontal collisions. Within this system, the crash box serves as the primary energy-absorbing component, designed to convert impact energy into controlled plastic deformation while limiting force transmission to the passenger cabin (Li et al., 2017). Traditionally manufactured as thin-walled tubular structures (Xu et al., 2018), crash boxes have been extensively optimized through geometric modifications, material selection, and structural enhancements to minimize peak impact force and maximize specific energy absorption (Mohammadiha et al., 2015; Ciampaglia et al., 2021). However, real-world accidents frequently involve complex oblique conditions rather than pure axial impact (Abraham et al., 2016; Gulino et al., 2021). Recent studies have demonstrated that the impact angle critically influences crash box performance; specifically, increasing impact angles significantly reduces energy absorption by shifting the deformation mode from stable axial crushing to less efficient global bending (Jongpradist et al., 2024; Zhou et al., 2017; Ma et al., 2021; Liang et al., 2022). Critically, this shift in deformation mode dictates not only the energy absorption efficiency but also the spatial envelope of the collapsed structure. Verifying that the crash box deforms within its allotted packaging space is essential to prevent hazardous intrusion into critical downstream components, such as the radiator or engine block.

Concurrently, machine learning (ML) and data-driven surrogate models have emerged as transformative tools for accelerating complex structural engineering workflows (Khan and Yairi, 2018). In the automotive sector, ML has been widely adopted for applications ranging from autonomous driving (Grigorescu et al., 2020) to predictive maintenance (Zhang et al., 2019) and manufacturing process optimization (Wang et al., 2018). In crashworthiness optimization, ML architectures offer computationally efficient alternatives to high-fidelity finite element analysis (FEA), drastically compressing design iteration cycles while capturing non-linear dynamic responses.

Early data-driven studies primarily employed artificial neural networks (ANNs) to predict scalar crashworthiness indicators, such as initial peak force and total energy absorption (Homsnit et al., 2024; Li et al., 2022). To capture the temporal evolution of the crash, subsequent research utilized recurrent neural networks—particularly Long Short-Term Memory (LSTM) models—which have shown robust capability in modeling time-dependent force-displacement histories (Siami-Namini et al., 2019; Guo et al., 2023; Kohar et al., 2020; Kohar et al., 2021). However, a fundamental limitation of these data-driven approaches lies in their reliance on Mean Squared Error (MSE) as the primary training objective. While MSE effectively minimizes global variance, it inherently penalizes the high-frequency oscillations characteristic of buckling, effectively treating them as noise (Le Guen and Thome, 2019). Consequently, models trained solely on MSE tend to generate over-smoothed response curves that fail to reproduce the sharp local peaks essential for accurate maximum force prediction. This limitation highlights the critical need for loss functions that go beyond simple pointwise matching to explicitly enforce the topological and physical validity of the deformation path.

To address this, Physics-informed Neural Networks (PINNs) have emerged as a powerful deep learning framework. By embedding governing physical laws directly into the loss function, PINNs enforce consistency with partial or ordinary differential equations, enabling data-efficient modeling while preserving fundamental physical principles. This approach has proven effective in domains such as fluid dynamics (Cai et al., 2021), structural analysis (Hu et al., 2024) and the mechanics of origami structures (Liu et al., 2024; Han et al., 2022).

However, the direct application of conventional PINNs to structural crashworthiness presents significant challenges. High-velocity impact dynamics involve severe plastic deformation, localized fracture, and non-linear contact interactions, for which explicit, continuously differentiable governing equations are computationally intractable. Consequently, standard PINN formulations—which explicitly minimize residuals of known PDEs—are ill-suited for dynamic crash simulations. To address these limitations, alternative strategies have been proposed, such as integrating graph neural networks with loss functions derived from domain-specific constitutive relationships rather than PDEs (Wang et al., 2025; Ciampaglia et al., 2024). These studies suggest that developing physics-informed loss functions rooted in engineering knowledge provides a robust pathway for enhancing model reliability in complex mechanical problems where exact equations are missing.

This study presents a comprehensive machine learning framework for predicting the force–displacement response and the full 3D deformation history of automotive crash boxes under axial and oblique impact loading. A Design of Experiments (DoE) strategy is employed to generate a systematic dataset via Finite Element Analysis (FEA), validated against experimental data. The proposed framework introduces three key innovations: (1) a three-dimensional convolutional autoencoder (3D-ConvAE) extracts a compact latent representation of complex crash box geometries directly from CAD models, which is concatenated with the impact angle to serve as the conditioning input; (2) a multilayer perceptron (MLP) predicts the complete force–displacement response in a single forward pass, trained using a novel set of physics-informed loss functions, denoted as $\mathcal{L}_{PhyM}$, specifically formulated to enforce physical constraints such as monotonic displacement and energy conservation; and (3) a long short-term memory (LSTM) model autoregressively predicts the time-dependent sequence of nodal displacements, enabling visual inspection of predicted failure modes and spatial interference.

The model is rigorously evaluated using topology-aware metrics, including Discrete Fréchet distance for force curves alongside Hausdorff and Chamfer distances for 3D geometric fidelity, ensuring accurate prediction of both transient mechanics and dynamic response bounds.

2 METHODOLOGY FOR MACHINE LEARNING FRAMEWORK

The proposed framework is developed to predict both the time-dependent force–displacement response and the three-dimensional deformed shape of hexagonal crash boxes subjected to multi-angle impact loading. As illustrated in Figure 1, the workflow is structured into three interconnected stages: (1) high-fidelity explicit finite element data generation; (2) shared low-dimensional geometric feature extraction via a 3D-ConvAE; and (3) dual-path neural prediction, comprising a Multilayer Perceptron (MLP) for force–displacement mapping and a Long Short-Term Memory (LSTM) network for autoregressive 3D spatial collapse tracking.

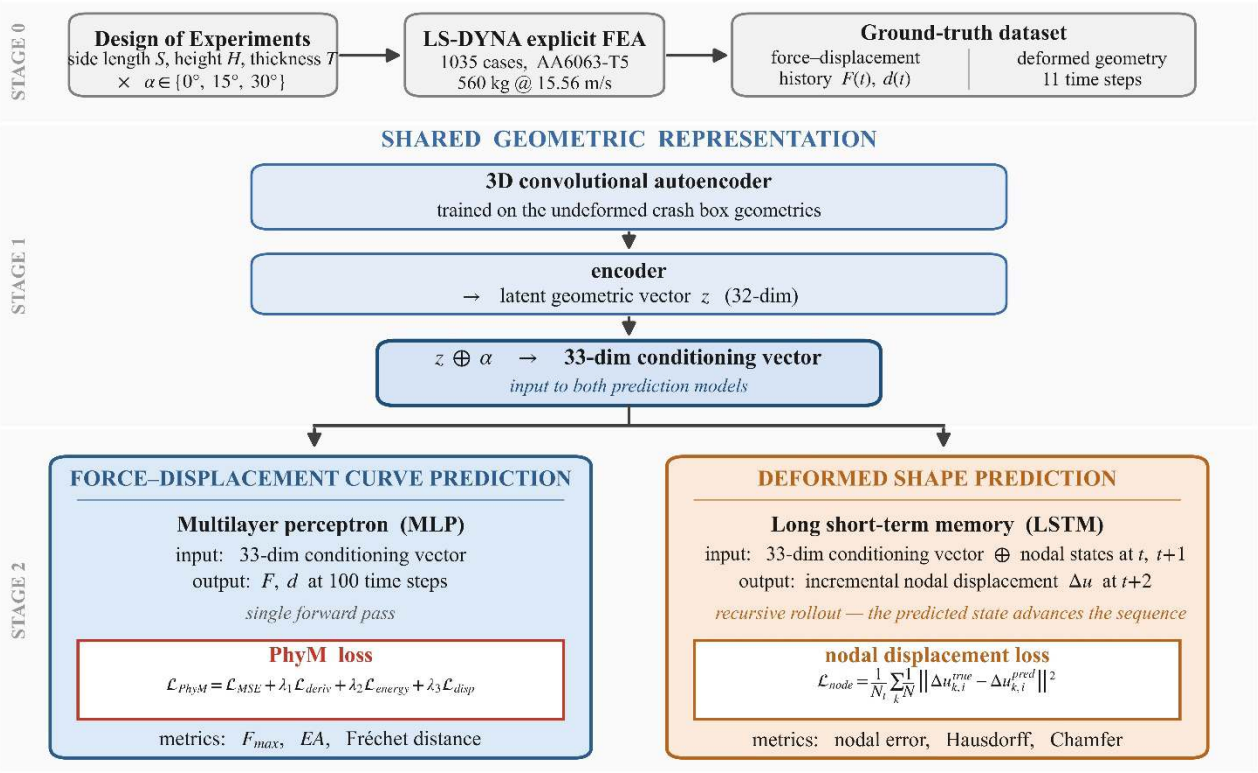


Figure 1 Machine learning framework for predicting force–displacement curves and deformed shapes.

In the initial stage, the geometric parameter space of the hexagonal crash box—defined by side length (*S*), height (*H*), and thickness (*T*)—is systematically sampled using a stratified Design of Experiments (DoE) approach. These geometric parameters are combined with the impact orientation angle (α) to construct the design input vector. High-fidelity non-linear explicit finite element simulations are performed across the sampled parametric space to extract ground-truth crashworthiness targets, specifically the transient force–displacement history $F(d)$ and the temporal deformation sequence of nodal coordinates $\boldsymbol{u}(t)$.

Following dataset generation, the second stage establishes a shared geometric representation. To efficiently process high-dimensional spatial topology without manual feature extraction, a 3D-ConvAE compresses the voxelized CAD models into a low-dimensional latent space. Specifically, the undeformed geometry is encoded into a 32-dimensional latent vector, denoted as $\boldsymbol{z}$. This vector is concatenated with the impact angle α to form a 33-dimensional conditioning vector $[\mathbf{z} \oplus \alpha]$. Supplying this identical conditioning vector to both downstream prediction models ensures physical and geometric consistency across both output domains.

The final stage comprises two prediction models. For the force–displacement response, an MLP maps the conditioning vector directly to the complete force–displacement curve in a single forward pass. The MLP is trained using a custom Physics-Informed Mathematical loss function ($\mathcal{L}_{PhyM}$), incorporates kinematic bounds and energy conservation

penalties to suppress non-physical numerical oscillations and capture sharp buckling peaks. Concurrently, an LSTM network models the transient 3D deformation sequence autoregressively. At each time step t, the LSTM ingests the conditioning vector alongside nodal states from two preceding time increments $(t-1, t)$ to predict the incremental nodal displacement ($\Delta \boldsymbol{u}$) at the subsequent time step $t+1$, recursively advancing the spatial collapse history. Model performance is rigorously evaluated against explicit FEA benchmark data using scalar metrics (peak force F_{max}, Energy Absorption EA), curve topology metrics (Discrete Fréchet distance), and spatial geometric fidelity metrics (nodal error, Hausdorff distance and Chamfer distance).

3 FINITE ELEMENT MODELLING AND CRASHWORTHINESS ANALYSIS

To generate the high-fidelity ground-truth dataset required for training the physics-guided machine learning framework, explicit non-linear finite element simulations are performed using LS-DYNA. This section details the numerical discretization, constitutive modeling, and experimental validation protocols, followed by an analysis of the transient collapse mechanisms across varying impact angles.

3.1 Finite Element Discretization and Simulation Setup

The hexagonal crash box profile is defined by three primary geometric parameters: side length (S), height (H), and thickness (T), as illustrated in Figure 2. Transient dynamic crash simulations were performed using the explicit non-linear finite element solver LS-DYNA. Simulations were conducted using nominal, idealized geometries to establish a deterministic response baseline across multi-angle impact scenarios. While initial geometric imperfections typically attenuate peak collapse forces, utilizing nominal profiles isolates the pure interaction between impact orientation and cross-sectional parameters by eliminating stochastic geometric noise.

Mesh sensitivity was assessed for three representative configurations with (S, H) = (20, 250), (30, 120), and (60, 185) mm, each with a wall thickness of 3.0 mm. Five element sizes were examined: 2.0, 2.5, 3.0, 4.0, and 5.0 mm. Relative to the 2.0 mm reference mesh, the 3.0 mm mesh yielded maximum absolute relative differences across the three configurations of 2.58% in energy absorption, 0.56% in initial peak force, and 2.58% in mean crushing force. These results support the selected maximum nominal element spacing of 3.0 mm for the evaluated crashworthiness indicators. The quantitative assessment was complemented by qualitative comparisons of folding patterns, focusing on fold wavelength, plastic hinge lines, and lobe formation. To maintain consistent cross-sectional resolution across geometries, each side was discretized into 20 elements, resulting in a nominal side-direction spacing ranging from 1.0 to 3.0 mm over the investigated range of side lengths of 20 to 60 mm.

The dynamic domain is discretized using Belytschko-Tsay shell elements with reduced integration to maximize computational efficiency. The resulting discretized model comprises 10,080 nodes and 9,960 elements. Five integration points are specified through the shell thickness to resolve non-linear bending stress gradients during severe plastic deformation. To suppress zero-energy modes inherent to single-point integration, a stiffness-based hourglass control algorithm is applied. Numerical stability was verified by confirming that the ratio of hourglass energy to total internal energy remained strictly below the 5% threshold throughout the deformation process.

The material response of the AA6063-T5 aluminum alloy is governed by a piecewise linear plasticity constitutive law (*MAT_24). The material is modeled as an isotropic elastoplastic medium with a density of 2,700 kg/m^3, Young's modulus of 68.9 GPa, Poisson's ratio of 0.3, and a yield stress of 180 MPa (Choiron et al., 2019). Post-yield hardening is defined by a tabulated effective stress-plastic strain curve, culminating in an ultimate tensile strength of 245 MPa. Dynamic hardening enhancement for AA6063-T5 has been reported as negligible within relevant strain-rate and temperature ranges (Ran et al., 2024; Wang et al., 2023). Consequently, given the nominal impact velocity regime of 15.56 m/s, strain rate effects were omitted in the constitutive model to reduce complexity without compromising predictive accuracy.

The simulation setup replicates standard dynamic impact testing protocols (Figure 2). The distal end of the crash box is fully clamped. A rigid impactor with a mass of 560 kg is assigned an initial velocity of 15.56 m/s (56 km/h) in the negative axial direction (Hershman, 2001). To investigate crash response sensitivity under oblique loading, three impact angles (α) are evaluated: 0° (pure axial), 15°, and 30°. Contact interfaces are managed using an automatic single-surface algorithm to prevent self-penetration of the folding lobes, while surface-to-surface contact with a friction coefficient of 0.2 is defined between the impactor and the specimen.

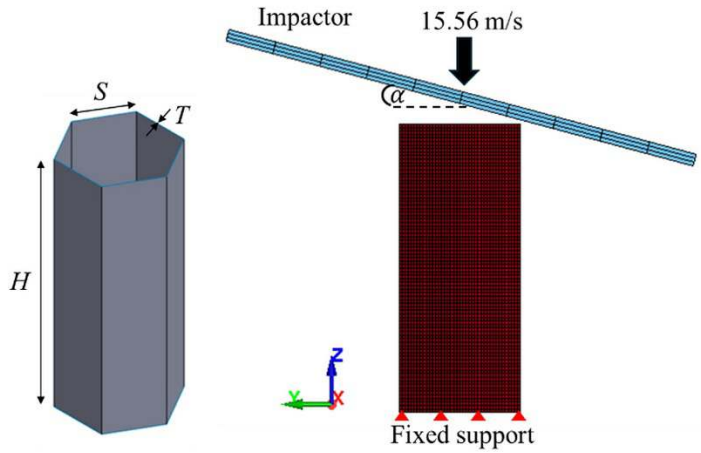


Figure 2 Illustration of the hexagonal crash box geometric parameters and loading conditions.

3.2 Analysis of Deformation Modes and Collapse Mechanisms

Prior to dataset generation, the numerical framework was validated against an axial crash experimental benchmark (Dirgantara et al., 2013) predicting peak force and mean crushing force within a relative error of less than 3%. To validate multi-angle dynamic crushing capabilities, the numerical parameters (element formulation, stiffness hourglass control, and contact friction) were benchmarked against experimental data for oblique crushing of aluminum square tubes at α = 15° (Karantza and Manolakos, 2022). The numerical model predicted the energy absorption, mean crushing force, and specific energy absorption within relative discrepancies of 5.7%, 5.6%, and 5.6%, respectively. The initial peak force was overpredicted by 11.7%, an effect primarily attributed to idealized sharp cross-sectional corners and the omission of initial geometric imperfections that naturally promote premature buckling initiation.

To further assess the consistency of the FE dataset, the mean crushing forces of the 345 axial configurations were compared with analytical predictions based on the Super-Folding Element theory (Wierzbicki and Abramowicz, 1983; Abramowicz and Wierzbicki, 1989). The comparison yielded a MAPE of 9.66% and R^2 = 0.944, with all cases falling within ±20% of the analytical prediction, without calibrating the analytical model to the present FE data. Complementary regression of the FE data yielded wall-thickness and side-width exponents of 1.554 and 0.468, respectively, compared with the theoretical values of 1.600 and 0.400. This agreement demonstrates that the FE dataset captures both the magnitude and the expected geometric scaling of the collapse response. Together with the comparisons against published axial and oblique impact experiments described above, these analytical results support the use of the FE dataset for developing and evaluating the proposed surrogate framework. Targeted experiments on hexagonal specimens would further strengthen the assessment of predictive performance under axial and oblique loading.

Crashworthiness performance is evaluated using three primary metrics: F_{max}, which dictates the initial deceleration severity; *EA*, representing the total work dissipated; and Specific Energy Absorption (*SEA*), ensuring mass-efficient design. Variations in impact angle reveal three distinct topological failure regimes:

1. Axial Progressive Buckling: Predominant at α = 0°, this mode represents the optimal energy dissipation mechanism. As shown in Figure 3a, the deformation is characterized by the formation of symmetric plastic hinges (concertina folding) initiating from the impact end. The response exhibits an initial peak force of approximately 210.1 kN at 7.2 mm displacement, followed by high-frequency oscillations where the wavelength corresponds to the fold length. Stability is maintained throughout the stroke, maximizing plastic work density.
2. Mixed Collapse Mode: At intermediate angles, the structure exhibits a bifurcation behavior (Figure 3b). Initially, local folding lobes attempt to form near the impactor; however, the accumulation of lateral moments eventually overrides the axial stiffness, triggering a delayed transition from local folding to global bending. This results in a peak (81.5 kN) significantly lower than the axial mode. This regime is highly non-linear and sensitive to imperfections, posing the greatest challenge for data-driven surrogates.
3. Global Bending Mode: Dominant at high oblique angles (α = 30°), this mode is characterized by a catastrophic loss of load-carrying capacity (Figure 3c). Rather than progressive folding, a stationary plastic hinge forms near the distal support, causing the structure to collapse laterally akin to an Euler column. In the representative 30° case, the peak force is comparable to the mixed mode (81.5 kN), but the post-peak force decays smoothly without oscillation, resulting in the lowest *SEA*.

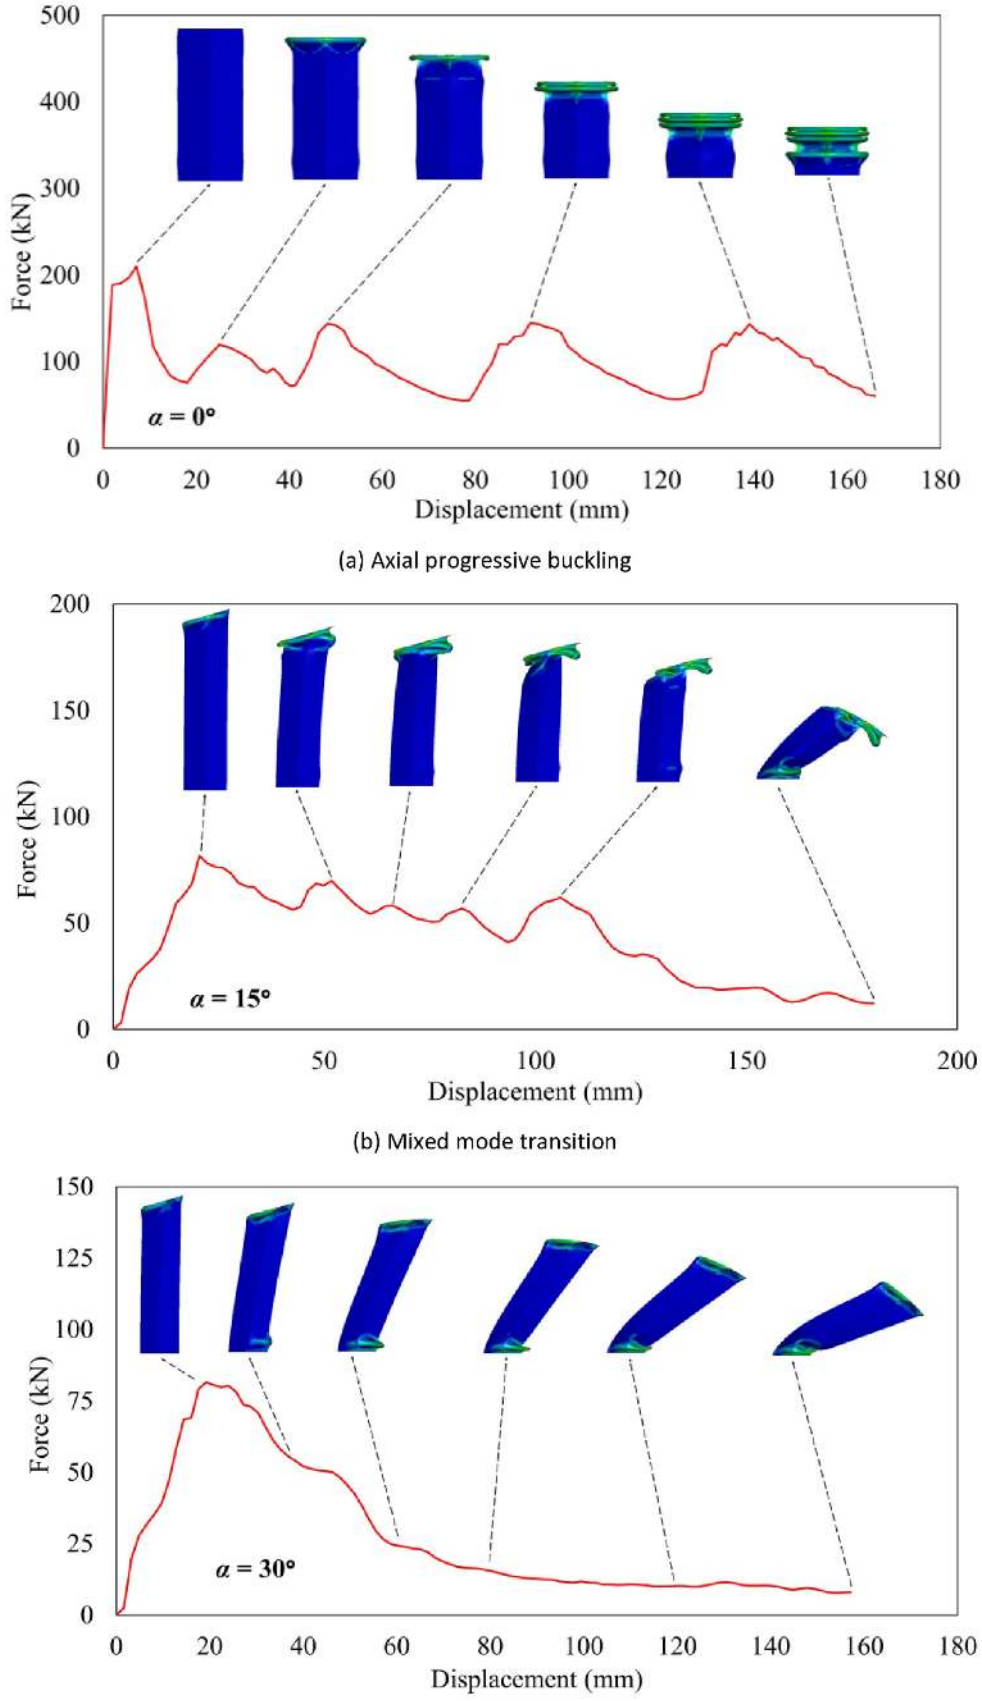


Figure 3 Comparison of deformation modes and force-displacement responses varying impact angles.

4 DATASET PREPARATION AND FEATURE EXTRACTION

The non-linear dependency of crashworthiness performance on geometric topology and impact obliquity necessitates a high-fidelity dataset that captures the full spectrum of deformation behaviors. This section details the data generation process, moving from a parametric DoE to the extraction of compact latent features via a 3D-ConvAE.

The integrated workflow is illustrated in Figure 4. The framework consists of three sequential phases: (1) parametric generation, where the geometric design space is sampled; (2) nonlinear FEA simulation, generating the ground-truth force (*F*) and deformation (*d*) histories; and (3) geometric encoding, where the voxelized geometries are compressed into latent vectors (**z**) via the 3D-ConvAE to serve as efficient inputs for the surrogate model.

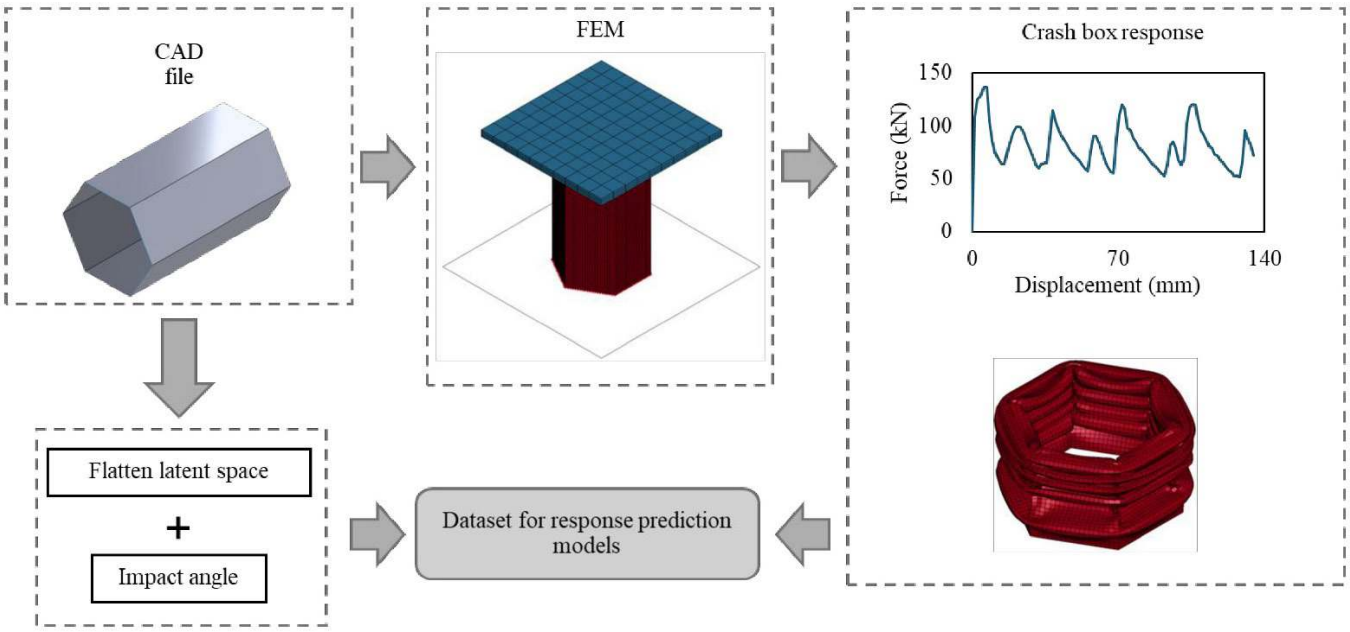


Figure 4 Integrated workflow for dataset generation, FEA simulation, and geometric feature extraction.

4.1 Design of Experiments (DoE) and Sampling Strategy

To ensure robust coverage of the design space while minimizing computational cost, a strategized DoE is employed. The hexagonal crash box is defined by three geometric variables (*S*, *H*, *T*). These parameters are sampled using within standard automotive engineering constraints (Equation 1), covering geometries ranging from compact, thick-walled stiffeners to slender, thin-walled energy absorbers:

$$20\ mm \leq s \leq 60\ mm$$

$$120\ mm \leq H \leq 250\ mm \tag{1}$$

$$1.5\ mm < T < 4.0\ mm$$

This sampling strategy yields a representative subset of 1,000 unique geometric configurations, comprising 765 cases for the training set and 235 cases for the validation set, as illustrated by the blue and red points in Figure 5, respectively. The validation cases are randomly selected values located within the gaps of the geometric parameters used in the training set, thereby generating a well-distributed sample set that thoroughly covers all design variables within their specified bounds. These datasets are used exclusively to train the autoencoder to learn a compact latent representation of the design space. Consequently, the autoencoder is trained solely on geometric variations, independent of loading conditions.

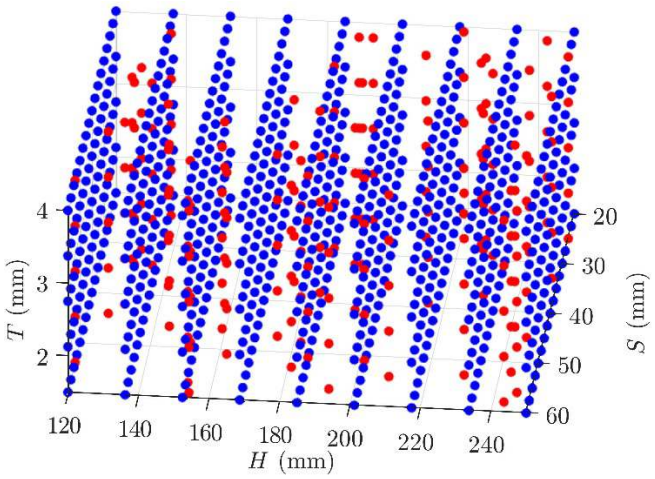


Figure 5 DoE sampling strategy and parameter distribution for each impact angle.

To investigate the influence of design responses on the force–displacement behavior and deformation patterns under axial and oblique impact loadings, a subset of 345 geometric configurations is inclusively selected from the 765 cases in the design space. Each selected configuration is then evaluated under three different impact angles. This approach enables a systematic assessment of the effect of impact angle on the crash response while maintaining geometric consistency across the loading conditions. For each impact angle, the dataset contains 345 distinct geometric configurations. At an impact angle of 0°, all configurations exhibit an axial progressive collapse mode, whereas at 30°, nearly all configurations undergo global bending. Thus, the deformation response at these two angles is predominantly governed by a single failure regime. In contrast, the 15° impact cases encompass all three deformation regimes, comprising 239 axial collapse cases (69.3%), 97 global bending cases (28.1%), and 9 mixed collapse cases (2.6%).

The force–displacement and deformation dataset comprises 1,035 impact cases, of which 828 (80%) are used for training and 207 (20%) for validation. A separate set of 50 configurations selected from the 765-configuration geometric pool is reserved for testing.

4.2 Voxelization Strategy

To enable the application of 3D convolutional operations, the continuous CAD geometries—represented as finite element meshes—must be discretized into structured tensors. While finite element models rely on unstructured Lagrangian meshes (nodes and elements), Convolutional Neural Networks (CNNs) require structured Euclidean grids. Consequently, the bounding box of the crash box domain Ω is discretized into a fixed binary voxel grid with dimensions $N_x \times N_y \times N_z = 32 \times 32 \times 64$.

The voxelization process maps the Cartesian coordinates of the FEA nodes to discrete grid indices. Let the physical dimensions of the bounding box be defined by L_x, L_y, L_z. The voxel resolution (spatial step size) along each axis, denoted as d_x, d_y, d_z, is calculated as:

$$\delta_x = \frac{L_x}{N_x}, \quad \delta_y = \frac{L_y}{N_y}, \quad \delta_z = \frac{L_z}{N_z} \tag{2}$$

For a given finite element node n located at Cartesian coordinates (x_n, y_n, z_n), the corresponding voxel indices (i, j, k) are determined by the floor function mapping:

$$i = \left\lfloor \frac{x_n - x_{min}}{\delta_x} \right\rfloor, j = \left\lfloor \frac{y_n - y_{min}}{\delta_y} \right\rfloor, k = \left\lfloor \frac{z_n - z_{min}}{\delta_z} \right\rfloor \tag{3}$$

where $(x_{min}, y_{min}, z_{min})$ represents the origin of the bounding box, and the indices satisfy $0 \le i \le N_x$, $0 \le j \le N_y$, $0 \le k \le N_z$.

The final voxel occupancy tensor $v(i,j,k)$ is defined as a binary function, where a value of 1 indicates the presence of structural material (node occupancy) and 0 represents void space:

$$v(i,j,k) = \begin{cases} 1 \; if \; \exists n \in Nodes \; such \; that \; n \to (i,j,k) \\ 0 \; otherwise \end{cases} \tag{4}$$

This binary representation captures the global topology and folding complexity of the crash box, while the specific shell thickness parameter (*T*) is preserved as a separate scalar input to the network (as detailed in Section 4.3) to avoid aliasing errors inherent to grid discretization.

4.3 Geometric Feature Extraction via 3D Convolutional Autoencoder

The 3D-ConvAE is employed to map the high-dimensional voxelized geometry into a compact, low-dimensional latent manifold. The network architecture, as illustrated in Figure 6, consists of two symmetric components: an encoder that compresses the spatial hierarchy of the input data and a decoder that reconstructs the original topology (Yang et al., 2016).

The 3D-ConvAE architecture comprises symmetrical encoder and decoder networks engineered to project high-dimensional voxelized finite element geometries onto a low-dimensional latent manifold, from which structural representations are subsequently reconstructed. The encoder comprises five stacked convolutional layers and four max-pooling layers, while the decoder consists of five stacked transposed convolutional layers and four upsampling layers. In addition, both the encoder and decoder include a reshape layer to facilitate the transformation between multidimensional feature maps and the latent vector representation.

The input to the encoder is a four-dimensional tensor of size 32 × 32 × 64 × 1, representing the spatial dimensions (N_x, N_y, N_z) with a single channel corresponding to voxel occupancy derived from the finite element model data. As the data propagates through the network, the encoder progressively performs spatial downsampling while increasing the depth of the feature channels. The encoder employs convolutional filters with Rectified Linear Unit (ReLU) activation functions to extract geometric features at multiple scales, ranging from local edge features in the initial layers to more global topological structures in the deeper layers. Max-pooling layers are used to reduce the spatial resolution by a factor of two at each stage. Following the sequence of convolutional and max-pooling operations, a reshape layer converts the resulting four-dimensional tensor into a one-dimensional vector. This hierarchical compression process distills the geometric information into a 32-dimensional flattened latent vector (*z*), which represents the compact latent space of the geometry. This bottleneck vector serves as a compact "DNA" representation of the crash box, retaining only the most critical geometric variations required for performance prediction.

The decoder is designed to reconstruct the original voxel grid from the 32-dimensional flattened latent vector. Through a sequence of 3D transposed convolutional layers, the spatial dimensions are progressively upsampled using upsampling layers until the original resolution of 32 × 32 × 64 is restored. The final output layer employs a sigmoid activation function, which is suitable for binary data, as it maps the reconstructed voxel values to the range of 0 to 1, representing the probability of voxel occupancy.

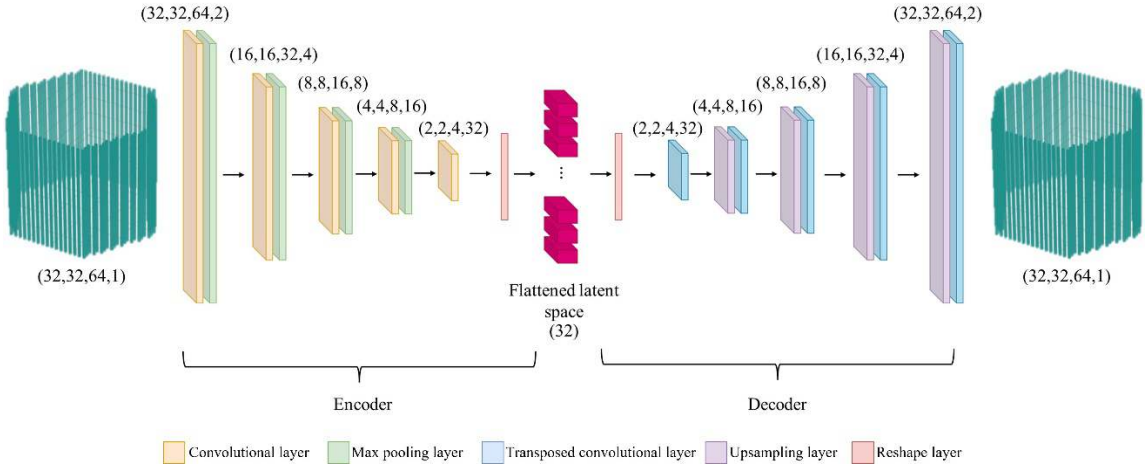


Figure 6 A 3D Convolutional Autoencoder

4.4 Autoencoder Training and Convergence

To ensure generalization, the data was split into training and validation sets as shown in Figure 5. The network was trained for 10,000 epochs using a batch size of 50 utilizing the Adam optimizer with an initial learning rate of 1×10^{-4}. The reconstruction fidelity was optimized by minimizing the Mean Squared Error (MSE) loss between the input occupancy grid v and the reconstructed output $\tilde{v}$:

$$\mathcal{L}_{MSE,v} = \frac{1}{N}\sum_{i=1}^{Nx}\sum_{j=1}^{Ny}\sum_{k=1}^{Nz}(v(i,j,k) - \tilde{v}(i,j,k))^2 \quad (5)$$

where $N = N_x \times N_y \times N_z$ represents the total number of voxels.

Figure 7 presents the loss history which training and validation curves exhibit a rapid initial decay and stabilize at a near-zero error magnitude after approximately 2,000 epochs. The absence of significant divergence between the training and validation loss indicates that the autoencoder has learned a robust representation of the geometric design space without overfitting. This confirms that the 32-dimensional latent vector effectively encapsulates the morphological variations of the crash boxes.

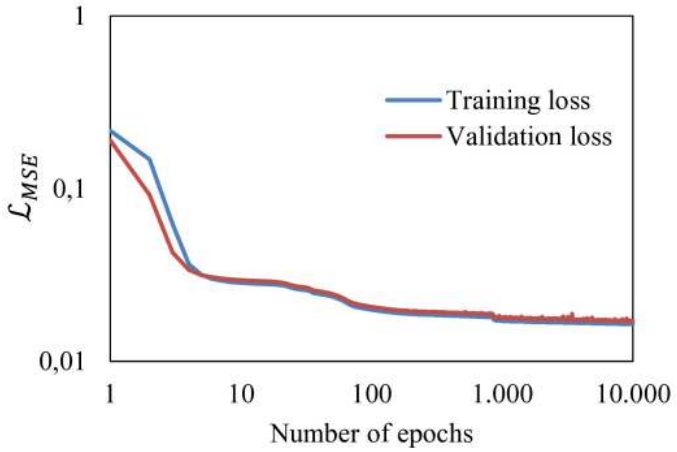


Figure 7 Training and validation loss convergence for the 3D-ConvAE.

5 PHYSICS-INFORMED PREDICTION OF FORCE-DISPLACEMENT RESPONSES

This section details the development and evaluation of the machine learning framework designed to predict the crash pulse. It covers the architecture of the predictive model, the formulation of the novel Physics-informed Mathematical (PhyM) loss function, and the subsequent results comparing the proposed approach against standard baselines.

5.1 Model Architecture and Loss Formulation

The predictive model employs a deep MLP regressor designed to map the static geometric representation to the dynamic crash response under varying loading conditions. The network architecture, summarized in Table 1, accepts a 33-dimensional input vector, which is constructed by concatenating the 32-dimensional latent geometric vector (z) derived from the 3D-ConvAE with the impact angle. This hybrid input strategy allows the model to simultaneously learn the influence of topological design features and external loading directions.

The network comprises two fully connected hidden layers, each containing 1,024 neurons with ReLU activation. This depth and width were specifically chosen to provide sufficient capacity to model the highly non-linear bifurcation behaviors associated with progressive folding and global bending modes. To mitigate overfitting and improve the model's generalization to unseen impact scenarios, L2 regularization is incorporated into the training process by adding a penalty proportional to the squared magnitude of the network weights to the loss function. In this study, an L2 regularization coefficient of 0.001 is adopted into the training loss. The penalty discourages large network weights and limits sensitivity to the training data (Ying, 2019). The output layer consists of 200 neurons, representing the discretized time-history of the crash event: the network predicts the force and displacement over $N_t = 100$ normalized time steps. The network was trained for 5,000 epochs using the Adam optimizer with a learning rate of 1×10^{-4} to ensure stable convergence.

Table 1 Architecture for the force-displacement prediction network

Parameters	Value
Input size	33 (Flattened latent space and impact angle)
Hidden layer	2
Number of neurons	1024
Output size	200 (Force-Displacement steps)
Epochs	5000
Learning rate	0.0001

To ensure physical consistency, the training is guided by a composite loss function. The baseline performance is established using the Mean Squared Error (MSE), which quantifies the pointwise discrepancy between the predicted and ground truth values for both force (*F*) and displacement (*d*). The loss function is expressed as:

$$\mathcal{L}_{MSE} = \frac{1}{2N_t}\sum_{j=1}^{N_t}\left[(F_j^{pred} - F_j^{true})^2 + (d_j^{pred} - d_j^{true})^2\right] \quad (6)$$

where, N_t represents the total number of time steps. The index *j* refers to the discrete time steps $t \in [1, N_t)$. Since the model predicts both force and displacement at each time step, the MSE is calculated by averaging the squared errors over all 2*N* output values.

However, standard MSE often leads to over-smoothed predictions that fail to capture high-frequency phenomena like buckling. To address this, the Physics-Informed Mathematical (PhyM) loss function was introduced, which augments the MSE with three physics-based constraints:

$$\mathcal{L}_{PhyM} = \mathcal{L}_{MSE} + \lambda_1\mathcal{L}_{deriv} + \lambda_2\mathcal{L}_{energy} + \lambda_3\mathcal{L}_{disp} \quad (7)$$

The physical constraints are defined as follows:

1. Energy Conservation ($\mathcal{L}_{energy}$): This constrains the total *EA* calculated as the integral of the force-displacement curve at each time step, to match the ground truth physics.

$$\mathcal{L}_{energy} = \frac{1}{N_t-1}\sum_{j=1}^{N_t-1}\left(EA_j^{pred} - EA_j^{true}\right)^2 \quad (8)$$

2. Kinematic Constraint ($\mathcal{L}_{disp}$): This term penalizes negative increments in the predicted displacement and therefore discourages backward motion in the force–displacement response. It is a soft monotonicity penalty and does not, by itself, guarantee nonnegative absolute displacement or exact satisfaction of all kinematic conditions.

$$\mathcal{L}_{disp} = \frac{1}{N_t-1}\sum_{j=1}^{N_t-1} max\left(-\Delta d_j^{pred}, 0\right) \quad (9)$$

3. Derivative consistency ($\mathcal{L}_{deriv}$): This term penalizes errors in the rate of change of force (ΔF_j) at time step *j*+1 compared to the current time step *j*. By enforcing accuracy in the derivatives, the model is encouraged to capture the sharp stiffness changes associated with fold formation and oscillatory buckling.

$$\mathcal{L}_{deriv} = \frac{1}{N_t-1}\sum_{j=1}^{N_t-1}\left(\Delta F_j^{pred} - \Delta F_j^{true}\right)^2 \quad (10)$$

In the formulations above, $\Delta(\cdot)_t = (\cdot)_{t+1} - (\cdot)_t$ denotes the forward difference between consecutive time steps. All force and displacement variables are normalized to [0, 1] before calculating the loss. The weighting coefficients were selected empirically by examining their effects on the force–displacement predictions. Increasing λ_1 improved reproduction of force fluctuations, with λ_1 = 200 giving the lowest initial peak-force error among the settings examined; further increases did not improve the predictions. For the energy term, λ_2 = 5 provided a compromise between energy-absorption accuracy and force-prediction MSE. Although larger weights could reduce the energy error at some settings, they increased the force-prediction MSE. Increasing λ_3 to 1 substantially reduced the penalty on negative displacement increments, with little additional reduction at higher weights. Accordingly, λ_1 = 200, λ_2 = 5, and λ_3 = 1, were adopted. These values are empirical settings for the present dataset and normalized loss formulation, and their suitability should be reassessed when either changes.

5.2 Evaluation of Input Representation Capabilities

To validate the effectiveness of the 3D-ConvAE for feature extraction, a comparative study was conducted between the proposed latent-based model (MSE-FL) and a baseline parameter-based model (MSE-GP). The primary objective of this comparison is to isolate the influence of geometric input representation on predictive performance; therefore, both models were trained using the identical standard MSE loss function. The MSE-GP baseline utilizes a vector of explicit geometric descriptors specifically side length (*S*), height (*H*), and thickness (*T*) concatenated with the impact angle. In contrast, the proposed MSE-FL model employs the 33-dimensional hybrid vector, combining the 32 latent features derived from the 3D-ConvAE with the impact angle.

The training convergence profiles, illustrated in Figure 8, immediately highlight the advantage of the learned representation. The MSE-FL model achieves a final validation loss of 9.98×10^{-4}, which is approximately 22% lower than that of the MSE-GP model (1.28×10^{-3}). The lower validation loss supports the effectiveness of the learned representation for the present prediction task. By compressing the structure into just three dimensions, the explicit parameter approach loses the topological fidelity required to resolve the highly non-linear bifurcation behaviors inherent in the crash response.

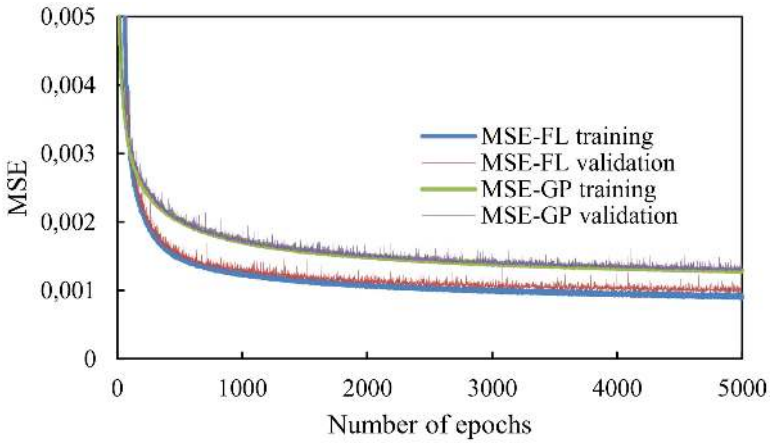


Figure 8 Comparison of MSE loss convergence for MSE-FL model and MSE-GP model

Qualitative analysis of the force–displacement responses for unseen test cases, shown in Figure 9 further corroborates these findings. For the axial impact case (Figure 9a), both models accurately predict the initial peak force; however, a significant divergence occurs in the post-peak regime. The parameter-based prediction tends to flatten out almost immediately after the first peak, failing to track the subsequent peaks and valleys associated with progressive fold formation. In contrast, the latent-based model successfully captures the initial oscillatory fluctuations and maintains the phase of the buckling response for a longer duration, although it also exhibits some signal attenuation in the later stages of the crash event.

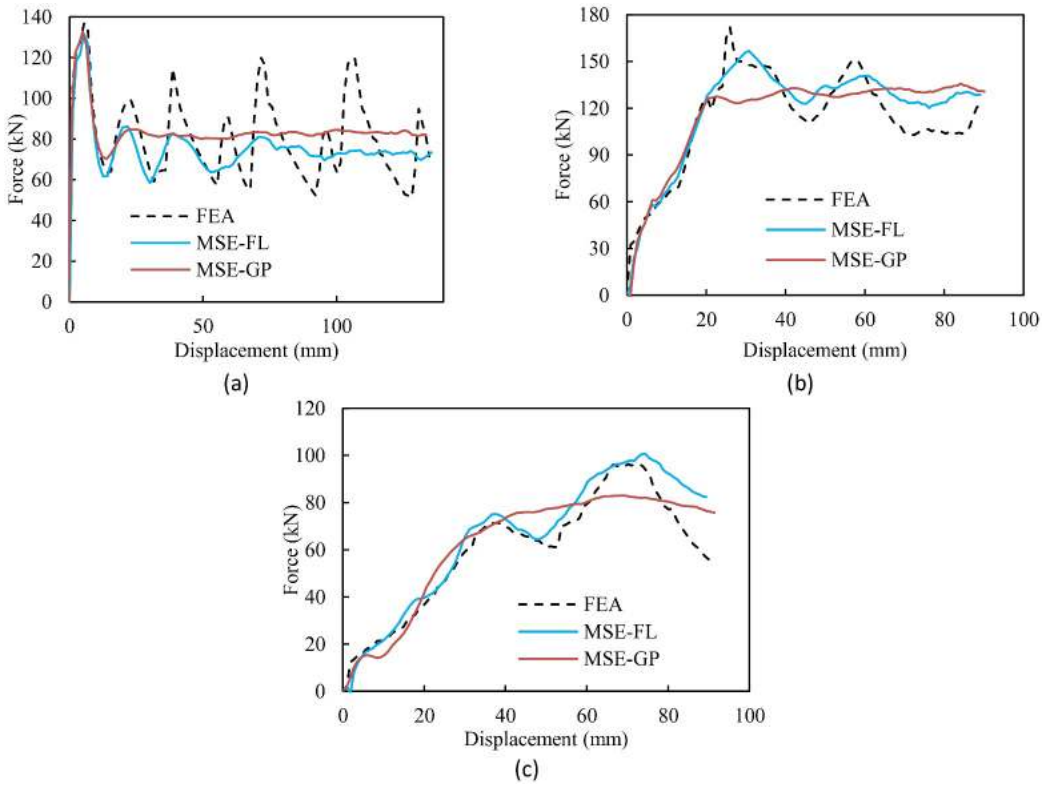


Figure 9 Force–displacement predictions for MSE-FL and MSE-GP model under (a) 0°, (b) 15° and (c) 30° impact angles.

This trend persists under oblique loading conditions Figures 9(b) and 9(c), where the MSE-GP model yields over-smoothed and inaccurate curves. Conversely, the MSE-FL model maintains a tighter correlation with the FEA ground truth, demonstrating that the high-dimensional latent vector provides a more robust, though not infallible, embedding for mapping geometry to mechanical response.

Despite the superiority of the flattened latent space, noticeable local discrepancies remain. As observed in the axial and 15° impact cases, the MSE-FL model captures the general oscillatory trend but occasionally underestimates the amplitude of specific buckling peaks. These residual errors indicate that while the geometric embedding is robust, the reliance on a standard MSE loss function is insufficient to enforce strict physical fidelity. This limitation motivates the development of the physics-informed loss function, which is evaluated in the subsequent section.

5.3 Evaluations of Physics-Informed Loss Function

To quantify the specific benefit of integrating physical governing laws into the training process, a comparative evaluation was performed between the proposed PhyM-FL model and the MSE-FL baseline. Since both models utilize the identical latent input representation, any divergence in performance can be directly attributed to the influence of the physics-informed loss function ($\mathcal{L}_{PhyM}$) versus the standard Mean Squared Error ($\mathcal{L}_{MSE}$).

Beyond pointwise error, assessing the fidelity of crash pulse predictions requires a metric that accounts for the sequential nature of folding and buckling. We employ the Discrete Fréchet Distance (D_F), a topology-based metric often described as the minimum leash length required to traverse two curves simultaneously (Meng et al., 2025; Aronov et al., 2006). Unlike standard Euclidean error, which compares points at fixed indices, D_F treats the crash response as a continuous path in the 2D Force-Displacement plane. Mathematically, for a ground truth curve C_{true} and a predicted curve C_{pred}, the Fréchet distance is defined as:

$$D_F\left(C_{true}, C_{pred}\right) = \inf_{\phi,\psi} \max_{\tau\in[0,1]} \sqrt{\left(d_{true}\left(\phi(\tau)\right) - d_{pred}\left(\psi(\tau)\right)\right)^2 + \left(F_{true}\left(\phi(\tau)\right) - F_{pred}\left(\psi(\tau)\right)\right)^2} \tag{11}$$

where ϕ and ψ are continuous, non-decreasing reparameterizations that map the unit interval $\tau \in [0, 1]$ to the respective curves, representing the crash progress from 0% to 100%. This metric penalizes deviations in both magnitude (Force) and phase (Displacement), ensuring that a predicted buckling peak is only considered accurate if it occurs at the correct crush depth.

The impact of different loss functions in the machine learning framework with physics-informed constraints is visually evaluated through the force–displacement responses under axial progressive deformation modes, in comparison with the FEA results, as shown in Figure 10. The baseline MSE-FL model, trained solely using the MSE loss, captures the overall trend of the force–displacement response but fails to accurately represent the post-peak behavior following the first peak. The predicted curve becomes excessively smooth and cannot reproduce the fluctuations observed in the FEA results, indicating that a purely data-driven loss function is insufficient to capture the underlying physical crushing behavior. To improve the physical consistency of the predictions, a physics-informed loss function is first introduced by incorporating an energy term ($\mathcal{L}_{energy}$), which constrains the model to preserve the energy absorption represented by the area under the force–displacement curve. However, assigning an excessively large weight to the energy term leads to nonphysical backward displacement. Since displacement in crash scenarios should remain monotonic, a displacement constraint term ($\mathcal{L}_{disp}$) is subsequently introduced to prevent such backward displacement. By appropriately tuning these terms, the predicted force–displacement curve ($\mathcal{L}_{energy} + \mathcal{L}_{disp}$) better captures the energy absorption compared with the baseline model. Although this modification improves the overall response, the post-peak fluctuations remain insufficiently represented. Therefore, a force derivative term ($\mathcal{L}_{deriv}$) is further incorporated into the physics-informed loss function to enable the model to learn the rate of force variation during progressive crushing. With this additional constraint, the proposed PhyM-FL model achieves closer agreement with the FEA results and significantly improves the prediction of the force–displacement response compared with the baseline MSE-FL model.

For the axial collapse regime, the predictions for unseen cases were examined for impact angles of 0° and 15°, as shown in Figure 11. In these scenarios, deformation is primarily governed by progressive folding along the loading direction. For the 0° configuration, both models successfully reproduce the initial peak force and the general oscillatory crushing response. However, the PhyM-FL model exhibits noticeably higher fidelity in capturing the fluctuation pattern of the FEA curve. This improvement is quantified by a significantly lower Fréchet distance of 7.02% for the PhyM-FL model, compared with 16.39% for the MSE-FL baseline. The enhanced curve-matching capability also translates to reduced errors in crashworthiness indicators, particularly the *EA*, for which the error decreases from 3.12% to 0.48%.

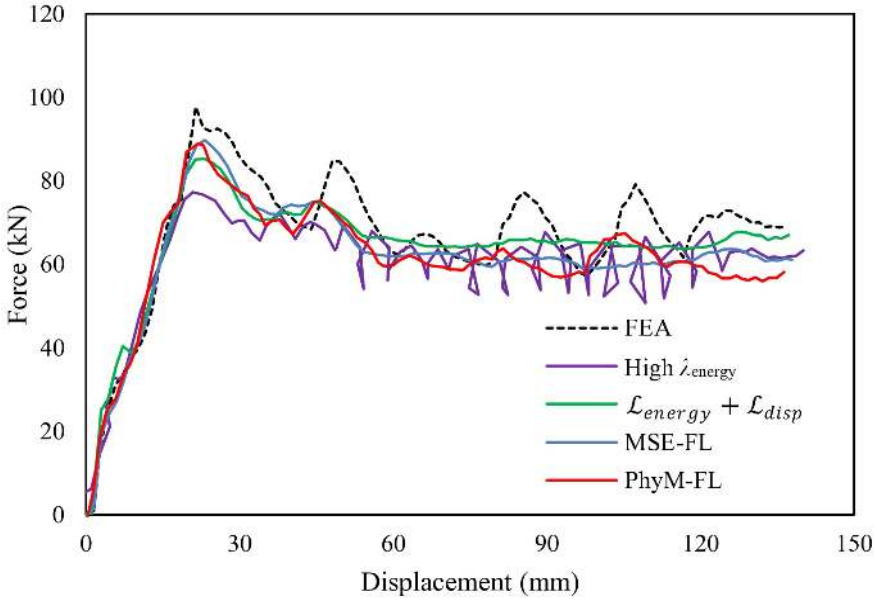


Figure 10 Force-displacement curve with backward displacement

A similar trend is observed for the 15° impact case, where the Fréchet distance decreases from 15.73% (MSE-FL) to 10.20% (PhyM-FL). Additionally, the physics-informed model provides a more accurate prediction of the maximum force (F_{max}), reducing the error from 11.98% to 4.26%. As the impact angle increases to the oblique regime (Figure 12), where global bending modes interact with local buckling, the PhyM-FL model continues to maintain superior structural topology. By enforcing physical and mathematical loss term, it prevents the "smoothing" effect seen in the MSE baseline, effectively capturing the sharp force drops associated with plastic hinge formation. Overall, these results clearly demonstrate that incorporating physical information substantially improves the ability to capture both the global shape of the force–displacement response and key crashworthiness metrics.

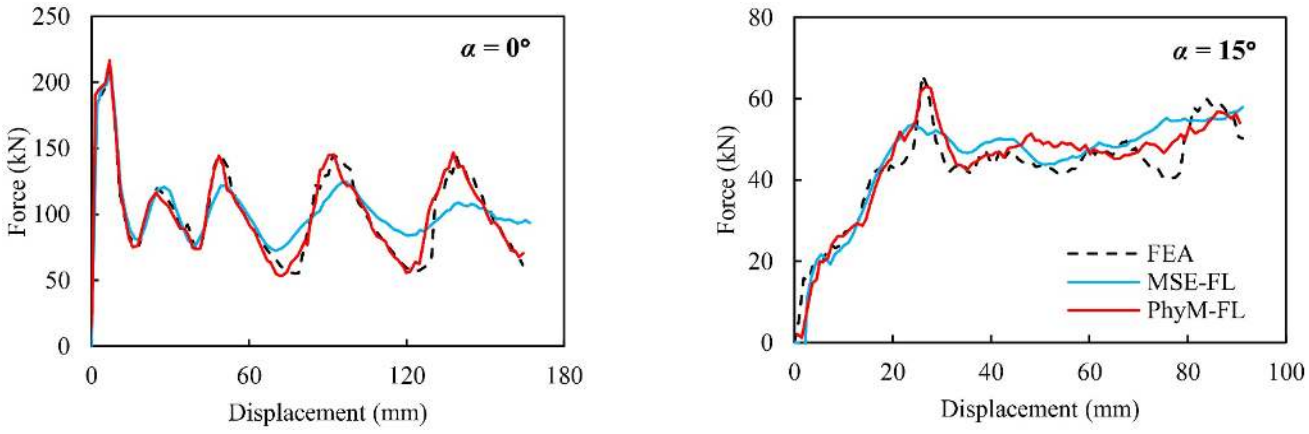


Figure 11 Force-displacement comparison for axial collapse mode.

For the 15° impact angle, the MSE-FL model yields a relatively high Fréchet distance of 21.61%, indicating difficulty in capturing the sharp peak drop and the long decaying tail of the FEA curve. In contrast, the PhyM-FL model reduces the Fréchet distance to 15.52%. With respect to crashworthiness indicators, the PhyM-FL model reduces the prediction errors of F_{max} and EA to 8.51% and 6.94%, respectively, reflecting its improved ability to capture both the global shape and local fluctuations of the bending response.

For the 30° impact case, the PhyM-FL model yields errors of 7.30% and 5.56% for F_{max} and EA, respectively. It is noted that the MSE-FL model achieves a slightly lower Fréchet distance (1.05%) than the PhyM-FL model in this specific instance, suggesting that both loss functions are capable of predicting the overall behavior when the response curve is smooth and lacks high-frequency oscillations.

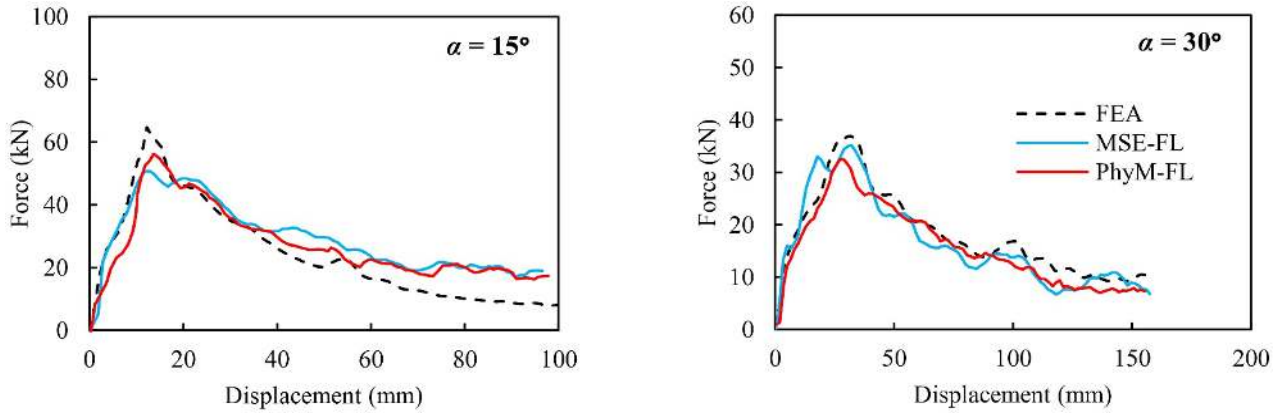


Figure 12 Force-displacement comparison of global bending mode.

Finally, the analysis extended to mixed-mode cases at impact angles of 15° and 30°, illustrated in Figure 13. These scenarios are characterized by the complex interaction between axial folding and global bending mechanisms, resulting in highly irregular force oscillations. For the 15° mixed-mode case, the MSE-FL model tends to over-smooth the predicted curve, failing to capture the combined fluctuation patterns. Consequently, it yields a Fréchet distance of 34.37%, whereas the PhyM-FL model reduces this value to 28.24%, indicating superior agreement with the global shape. While the MSE-FL model achieved a lower F_{max} error in this specific instance (9.31% vs. 13.16%), the PhyM-FL model proved far more reliable for energy prediction, reducing the *EA* error from 18.94% to 13.29%.

A consistent improvement is observed for the 30° mixed-mode configuration, where the PhyM-FL model more accurately reproduces both the peak amplitudes and the oscillatory behavior. Here, the Fréchet distance is reduced to 24.77% (compared with 40.97% for MSE-FL). Significant improvements are also observed in the prediction of F_{max}, where the error drops to 25.36%, notably lower than the 40.97% obtained using the baseline. Overall, these results demonstrate that the PhyM-FL model enhances robustness under complex deformation conditions, enabling a faithful reconstruction of mixed-mode force responses even when competing failure mechanisms interact.

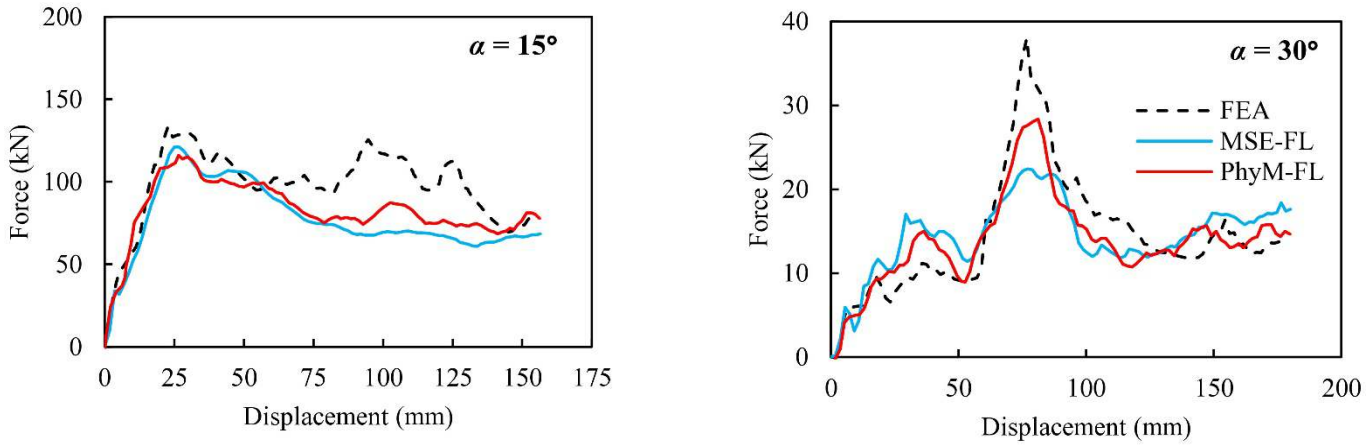


Figure 13 Force-displacement comparison of mixed collapse mode.

To enable a comprehensive evaluation of the model's engineering utility, correlation plots for the key crashworthiness indicators, including F_{max}, *EA*, and maximum displacement (d_{max}), are presented in Figure 14 for the training, validation, and unseen datasets. The PhyM-FL model demonstrates strong correlation with the ground truth, achieving coefficients of determination (R^2) of 0.9803, 0.9925, and 0.9964 for F_{max}, *EA*, and d_{max}, respectively. To complement these aggregate measures, prediction accuracy is also assessed using a ±10% relative-error band. Despite the high R^2 values, 30.2% of F_{max} predictions and 18.2% of *EA* predictions fall outside this band, illustrating that close agreement in the overall response variation does not necessarily imply small relative errors for every case. Compared with MSE-FL, the clearest improvement is obtained for F_{max}: R^2 increases from 0.9712 to 0.9803, while the proportion outside the band decreases from 39.6% to 30.2%, a reduction of 9.4 percentage points. For *EA*, a modest improvement is observed, with R^2 increasing from 0.9922 to 0.9925 and the proportion outside the band decreasing from 18.8% to 18.2%.

Both models retain all d_{max} predictions within the ±10% band. These results indicate that the principal benefit of the PhyM loss is improved peak-force prediction, accompanied by a modest improvement in energy absorption and maintained displacement accuracy.

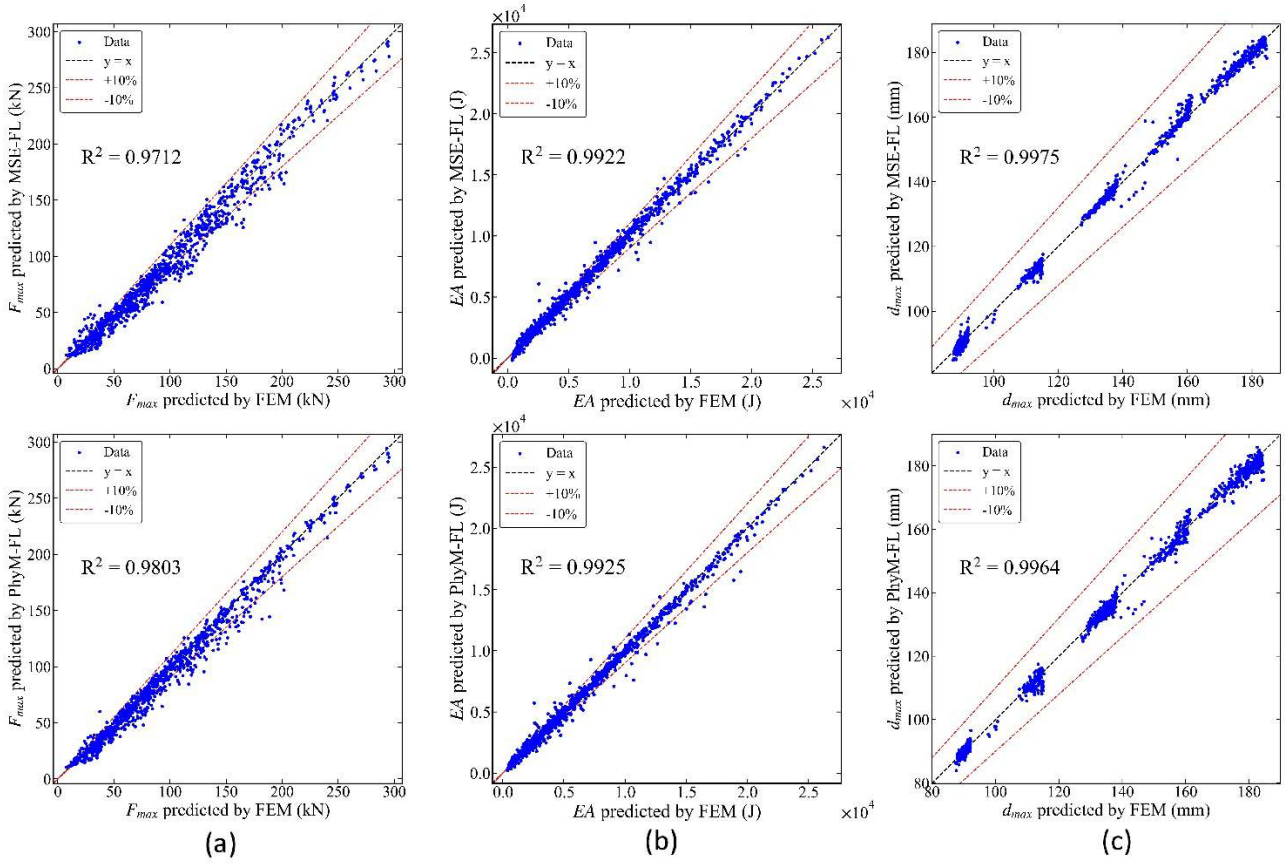


Figure 14 Correlation plots for crashworthiness indicators using MSE-FL and PhyM-FL models: (a) F_{max}, (b) *EA,* (c) d_{max}.

To quantitatively assess the global fidelity of the predicted crash pulses, Figure 15 presents histograms of the Fréchet distance errors for the unseen test set. These distributions highlight distinct differences in prediction accuracy between the physics-informed and baseline approaches.

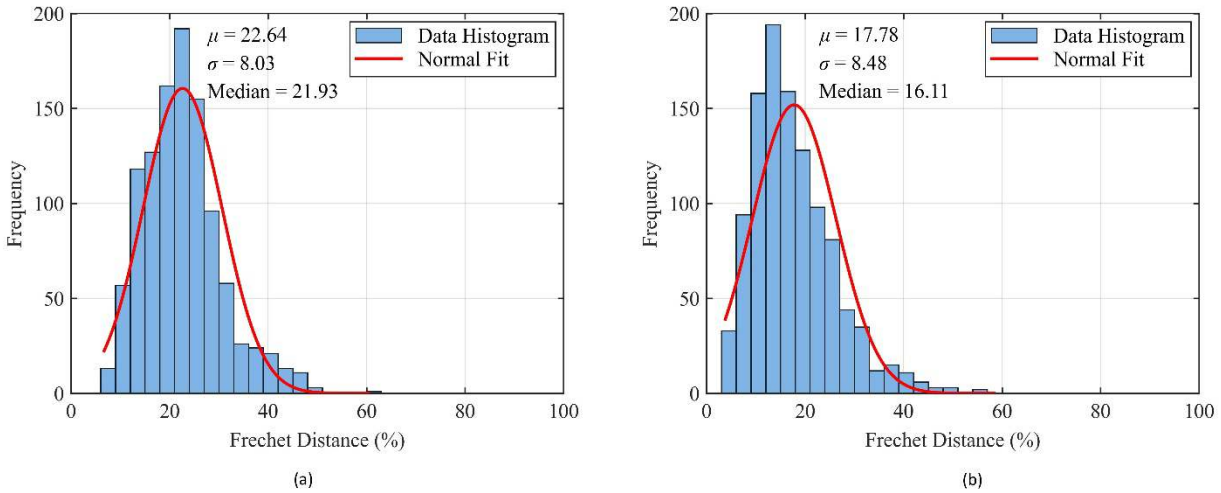


Figure 15 Distribution of Fréchet distance errors for (a) MSE-FL and (b) PhyM-FL models.

As shown in Figure 15a, the MSE-FL model exhibits higher mean and median Fréchet distances of 22.64% and 21.93%, respectively. Although it maintains a relatively tighter error distribution (σ = 8.03%), this apparent consistency comes at the expense of accuracy, as the MSE loss tends to smooth out critical oscillatory features in order to minimize overall variance. In contrast, the PhyM-FL model (Figure 15b) achieves lower mean and median errors of 17.78% and 16.11%, respectively, indicating improved overall prediction accuracy of the force–displacement response. However, its distribution is slightly more dispersed (σ = 8.48%), suggesting a modest increase in variability despite better fidelity in capturing the underlying physical behavior

Collectively, these results demonstrate that incorporating physical governing laws through the PhyM loss function yields a significant improvement in capturing both the maximum force and the global structural response (Fréchet distance), while maintaining comparable reliability for energy absorption and displacement predictions.

6 PREDICTIONS OF TEMPORAL 3D DEFORMATION HISTORY

6.1 LSTM-based Sequential Learning Framework

Predicting the full 3D deformation history involves mapping the temporal evolution of highly non-linear buckling modes. Unlike the scalar force-displacement prediction, this task requires determining the instantaneous spatial coordinates of all finite element nodes at each time step. To address the challenge of temporal dependency in dynamic crushing, an LSTM network is employed due to its superior capability in capturing long-range dependencies and mitigating the vanishing gradient problem inherent in standard Recurrent Neural Networks (RNNs) (Greff et al., 2016).

The architecture, illustrated in Figure 16, is designed to predict incremental nodal displacements rather than absolute coordinates. This differential formulation stabilizes the learning process by focusing on the frame-to-frame motion (Δu) instead of the global position, which can vary significantly across different impact angles. The input to the LSTM network consists of a sequence of historical nodal positions arranged within a sliding temporal window of length (w). At each time step (t), the input tensor contains the spatial coordinates of all tracked nodes, typically structured with dimensions of (w, $N \times 3$), where w is set to 2, N denotes the number of nodes, and the three channels correspond to the x, y, and z directions. This sequential representation enables the network to capture short-term temporal dependencies in the deformation history.

Within the LSTM architecture, the first LSTM layer processes this sequence using gated memory cells—comprising input, forget, and output gates—to filter irrelevant historical information while retaining meaningful deformation features. This layer primarily captures short-range temporal patterns, such as the onset of folding and axial compression behavior. The second and third stacked LSTM layers progressively refine this representation, extracting higher-level temporal features associated with nonlinear buckling, mode interaction, and accumulated plastic deformation. Stacking multiple LSTM layers therefore enables hierarchical learning of increasingly complex folding mechanics. Following the recurrent layers, a Flatten layer reshapes the final temporal feature map into a one-dimensional vector, facilitating the transition from sequential modeling to regression. A dense layer with a hyperbolic tangent (tanh) activation function is then employed to capture nonlinear plastic deformation behavior. Finally, the output layer regresses the incremental nodal displacements $\Delta u_{t+1,i}$ along the x, y, and z directions, enabling continuous prediction of the deformation progression.

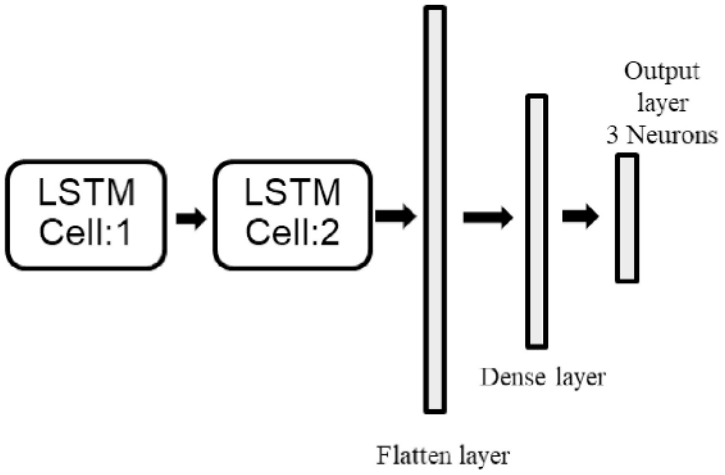


Figure 16 The LSTM architecture.

6.2 Loss Function and Fidelity Metrics

The training objective is designed to minimize the discrepancy between the predicted incremental displacements and the ground truth values derived from FEA. To ensure stability during the recursive prediction of long sequences, the network is trained to learn the incremental nodal displacements rather than absolute coordinates. This approach reduces the dynamic range of the output variables and enforces temporal continuity. The incremental displacement vector for the *i*-th node at time *t*+1, denoted as $\Delta \boldsymbol{u}_{t+1,i}$, is defined as:

$$\Delta \boldsymbol{u}_{t+1,i} = \boldsymbol{x}_{t+1,i} - \boldsymbol{x}_{t,i} \tag{12}$$

where $\boldsymbol{x}_{t,i}$ represents the coordinate vector of node *i* at time step *t*. Consequently, the Mean Squared Error (MSE) of these increments serves as the primary loss function ($\mathcal{L}_{disp}$) for backpropagation:

$$\mathcal{L}_{node} = \frac{1}{N_t}\sum_{k=1}^{N_t}\frac{1}{N}\left\|\Delta u_{k,i}^{true} - \Delta u_{k,i}^{pred}\right\|^2 \tag{13}$$

where N_t the total time steps in the crash event, *N* denotes the total number of mesh nodes, and ‖·‖ denotes the Euclidean norm. To quantitatively evaluate the prediction accuracy, the nodal error, maximum nodal error, and average nodal error (Kohar et al., 2021) are employed as evaluation metrics.

While MSE ensures pointwise accuracy, it may not fully capture the topological quality of the deformed shape, especially under large deformations where mesh sliding occurs. To strictly evaluate the geometric validity, two set-to-set distance metrics are introduced, i.e., Hausdorff Distance (*HD*) and Chamfer Distance (*CD*) (Yang et al., 2018; Avbelj et al., 2014; Javaheri et al., 2020).

The Hausdorff Distance (*HD*) quantifies the maximum mismatch between the two surfaces, serving as an indicator of the worst-case outlier:

$$HD\left(S_{true}, S_{pred}\right) = max\left\{\sup_{x\in S_{true}}\inf_{y\in S_{pred}}\|\boldsymbol{x}-\boldsymbol{y}\|, \sup_{y\in S_{pred}}\inf_{x\in S_{true}}\|\boldsymbol{x}-\boldsymbol{y}\|\right\} \tag{14}$$

Complementarily, the Chamfer Distance (*CD*) provides a robust assessment of the average global shape similarity. Unlike HD which focuses on the extreme outlier, *CD* computes the sum of squared distances between each point and its nearest neighbor in the corresponding set, as shown in Equation (15). The terms are normalized by the total number of nodes in the true set ($|S_{true}|$) and predicted set ($|S_{pred}|$), respectively, providing a mean error value that represents the overall topological accuracy:

$$CD\left(S_{true}, S_{pred}\right) = \frac{1}{|S_{true}|}\sum_{x\in S_{true}}\min_{y\in S_{pred}}|x-y|^2 + \frac{1}{|S_{pred}|}\sum_{y\in S_{pred}}\min_{x\in S_{true}}|x-y|^2 \tag{15}$$

By combining these metrics, the framework ensures that the predicted crash modes are not only visually plausible but also geometrically accurate for engineering validation.

6.3 Analysis of Spatial Error Distribution

The LSTM network was trained for 100 epochs using the data partition described in Section 4.1. The predicted deformation results consist of 11 deformed configurations corresponding to 11 distinct time steps, as exemplary illustrated in Figure 17 at *t* = 3, 6, and 10. In the case of axial collapse mode (0°), the deformation initiates at the impactor interface and propagates downward, forming distinct plastic hinges. The LSTM model accurately captures this progressive failure mode with errors remaining negligible in the undeformed bottom region. However, localized error concentrations are observed specifically at the active plastic hinges in the middle-to-top section with the average nodal error of 5.53 mm for *t* = 10. This localization occurs because the high curvature gradients at the fold lines present a challenge for reconstruction from the latent space, yet the overall topology remains consistent with the physical folding logic.

In contrast, the 15° impact case exhibits a mixed-mode deformation characterized by asymmetric axial folding followed by a transition to global instability. While the model captures the overall collapse behavior, the error distribution becomes more dispersed across the crushing zone compared with the axial case, with an average nodal error

of 6.01 mm at t = 10. This indicates that the model encounters greater difficulty in accurately tracking the bifurcation point at which the structure transitions from progressive folding to global bending.

Under the 30° impact condition, the structure undergoes global bending. In this case, the model successfully predicts the formation of a stationary plastic hinge near the support, with an average nodal error of 1.43 mm for the entire crash box at the final time step. This relatively low error compared with the other two modes is attributed to the simpler deformation behavior associated with the global bending mode.

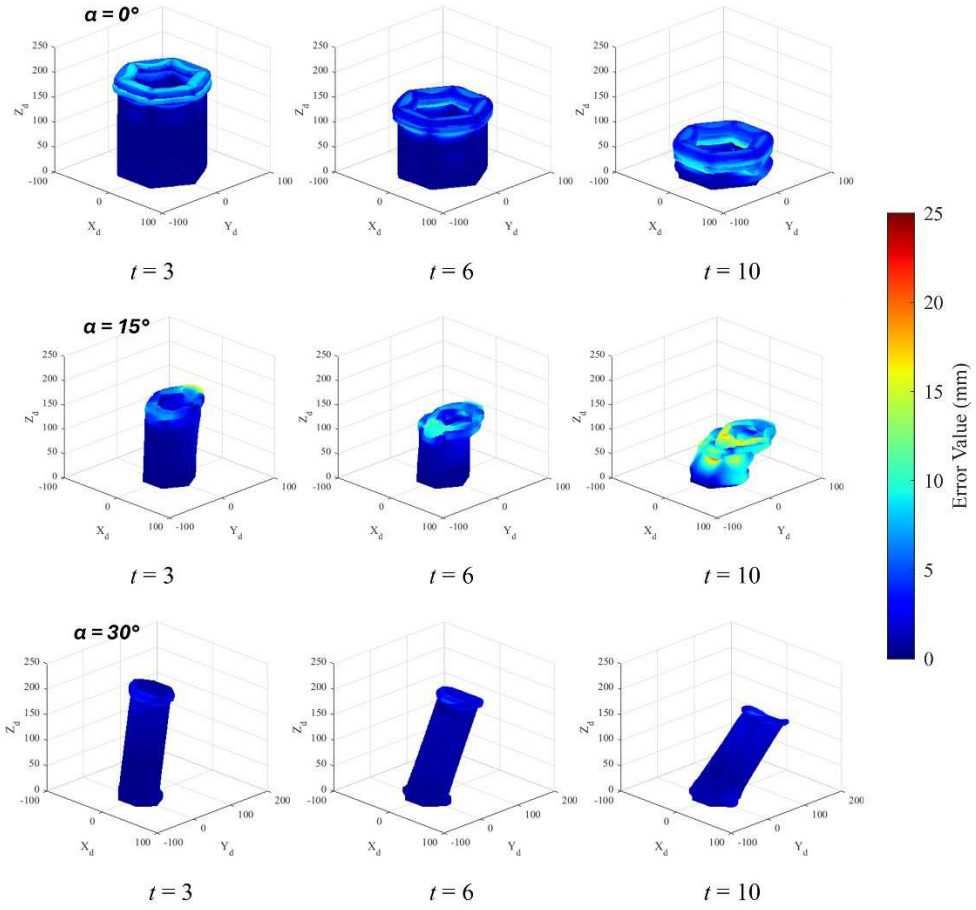


Figure 17 Nodal error distribution at different impact angles.

6.4 Generalization Capability and Error Analysis

To rigorously evaluate the generalization capability of the proposed LSTM framework, the models are tested on 50 unseen crash box configurations that are distinct from those used in the training set. In this evaluation, the *HD* and *CD* are employed instead of nodal error. These shape-based error metrics quantify the discrepancy between two-point sets using nearest-neighbor distances and therefore provide a more comprehensive assessment of geometric similarity.

The predictive performance is further categorized according to the three dominant deformation modes—axial progressive collapse, mixed collapse, and global bending—illustrated in Figure 18. As shown in Figure 18(a), the average shape errors exhibit a monotonic increase in both *HD* and *CD*, reflecting the accumulation of discrepancies as deformation severity increases. At the early stages of deformation, the errors remain relatively low due to the simpler deformation characteristics. As the crash progresses toward the final time step, the average *HD* and *CD* increase to approximately 10 mm and 4 mm, respectively. This increase is primarily driven by the emergence of localized complexities in fold formation as well as the development of plastic hinges in the crash box structures.

Complementarily, Figure 18(b) presents the frequency distribution of the final *CD* values for the 50 unseen datasets at the final time step. The histogram shows that most samples corresponding to axial progressive

deformation are concentrated within the low *CD* range of approximately 1–4 mm, indicating that the predicted shapes closely match the true deformation shapes and that the model performs consistently for this deformation mode. For global bending, the distribution is slightly broader; although many samples still fall within the range of 2–5 mm, the variability is greater than that observed for axial collapse, suggesting moderate prediction accuracy. In contrast, mixed collapse exhibits the widest spread, with several *CD* values extending beyond 3 mm and reaching up to 14 mm, and many results distributed in the range of approximately 6–14 mm. This trend indicates that the model encounters greater difficulty in accurately capturing the deformation characteristics associated with the mixed deformation mode.

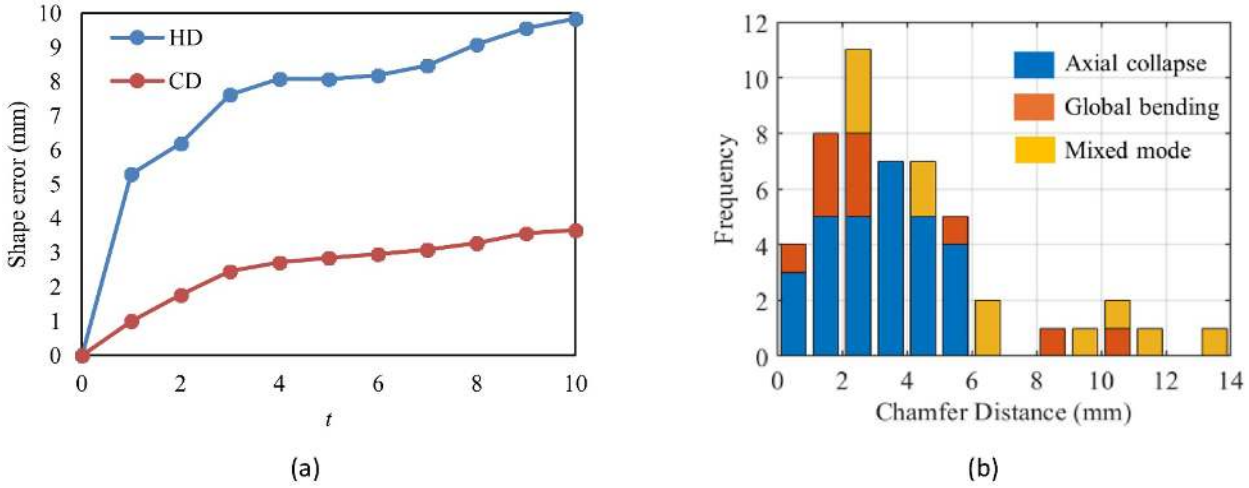


Figure 18 Shape error evaluation of deformation prediction: (a) *HD* and *CD* across deformation stages (T), (b) frequency distribution of *CD* at final time step.

Overall, the results suggest that prediction performance varies across deformation modes. The most reliable shape reconstruction is achieved for axial progressive collapse, moderate performance is observed for global bending, and larger geometric deviations occur more frequently for the mixed collapse mode.

To gain a deeper understanding, representative unseen cases from each collapse mode are compared in Figure 19, with FE results shown in red and LSTM predictions in blue. Agreement is quantified using HD and CD. For axial collapse mode, the *HD* and *CD* are 9.6 mm and 2.7 mm, respectively, for the 0° impact, and improve to 4.4 mm and 1.7 mm for the 15° impact. Given the total crash box height of 239 mm, a maximum deviation of 9.6 mm is relatively minor, confirming that the LSTM effectively captures the periodic nature of concertina folding without significant topological distortion.

For the global bending mode, typically observed under 30° loading, the deformation path is highly nonlinear, with buckling predominantly localized in the lower region. Despite this complexity, the model captures the macroscopic bending behavior with reasonable accuracy. The error metrics for unseen global bending cases (*HD*: 6.1 mm, *CD*: 1.8 mm) are comparable to those of the axial cases. While noticeable deviations occur at the top portion of the crash box, directly subjected to impact by the impactor and the distal ends due to rigid-body rotation, the low *CD* values indicate that the overall shape volume is well-preserved, suggesting that the autoencoder effectively learns the low-frequency components of the global bending deformation field.

However, the mixed-mode deformation presents the most significant challenge, involving a stochastic transition from initial axial folding to subsequent global buckling. At a 15° impact angle, notable discrepancies arise near the upper folding region, resulting in a peak *HD* of 32.7 mm and a *CD* of 11.2 mm for specific geometries for S = 34 mm, H = 248 mm and T = 2.6 mm. Notably, a significant disparity is observed between the high HD (indicating severe local outliers) and the relatively low *CD* (indicating good global shape agreement). This discrepancy reveals that while the model successfully captures the macroscopic collapse mode, it struggles to precisely localize the stochastic bifurcation point where the structure transitions from symmetric folding to global bending. These elevated errors are attributed to mode interaction, where directional biases in the early folding stage of 38.3 mm near the top end propagate and amplify during the global buckling phase. Conversely, at a 30° impact angle, the reduced folding complexity leads to improved performance (*HD*: 9.2 mm, *CD*: 4.3 mm). Although the mixed-mode predictions exhibit higher relative errors, the model still correctly identifies the failure tendency, which is critical for safety assessment.

The proposed LSTM framework demonstrates strong generalization capability in predicting crash box deformation, particularly for axial collapse and global bending modes, where low *HD* and *CD* values indicate good agreement with the FE results. However, its performance deteriorates under mixed-mode deformation due to the increased geometric complexity arising from the simultaneous occurrence of axial folding and global buckling. This interaction leads to larger deviations, especially in regions characterized by severe localized folding. The reduced accuracy in mixed-mode cases can also be attributed to the relatively limited number of such samples in the training dataset compared to axial collapse and global bending cases. This imbalance restricts the model's ability to learn representative mixed-mode features, resulting in higher prediction errors. Overall, the model achieves accurate predictions across most deformation scenarios; nevertheless, further improvements in capturing complex folding mechanisms would enhance its robustness and reliability for mixed-mode crash conditions.



Figure 19 Comparison of deformed shapes between predicted deformation and FE results.

7 CONCLUSIONS

A physics-informed deep learning framework was established for the crashworthiness prediction of hexagonal crash boxes under multi-angle impact conditions. By integrating a 3D-ConvAE with a Long Short-Term Memory (LSTM) network, the proposed architecture successfully mapped voxelized geometric inputs to high-dimensional spatio-temporal responses. A primary contribution of this research is the compression of complex 3D CAD geometries into a compact 32-dimensional latent space. It was demonstrated that utilizing this latent representation as an input feature significantly enhances predictive fidelity compared to traditional raw parametric data, particularly in capturing localized buckling onset.

Regarding the force–displacement response, the implementation of a physics-informed loss function (PhyM) provided a rigorous constraint mechanism over traditional Mean Squared Error (MSE) approaches. By enforcing force derivatives and energy conservation principles, the PhyM-enhanced model exhibited superior capability in capturing high-frequency oscillatory buckling patterns. Quantitative validation indicated a marked improvement in performance, with errors in peak force F_{max}, *EA*, and Fréchet distance being reduced from 10.86%, 8.23%, and 22.64% to 7.24%, 6.31%, and 17.78%, respectively. These improvements demonstrate the benefit of incorporating physics-informed loss terms in the prediction of complex crushing responses.

The LSTM framework effectively reconstructed the 3D deformation history across highly non-linear buckling regimes. Geometric fidelity, evaluated via Hausdorff and Chamfer distances, confirmed the model's ability to generalize to unseen impact scenarios. While axial progressive collapse and global bending modes were predicted with high stability, the mixed collapse mode exhibited the largest error distribution. This highlights the increased complexity of capturing coupled deformation mechanisms where stochastic transitions occur. Overall, this framework demonstrates that physics-informed surrogate models can replace computationally intensive explicit FEA for the rapid design and optimization of advanced thin-walled energy absorbers, offering a path toward real-time crashworthiness evaluation in automotive engineering.

The present models were trained for a fixed impact velocity and a single material, and their predictive accuracy under other loading conditions and material systems remains to be established. The PhyM-FL framework could be extended by incorporating impact velocity and relevant material properties as additional inputs alongside the geometric latent vector and impact angle. Such an extension would require an expanded FE dataset and targeted experimental validation to evaluate predictive accuracy across different deformation and collapse mechanisms. These developments could broaden the framework's applicability and support efficient crash-box design and optimization across a wider range of loading conditions and material systems.

Acknowledgments

The authors gratefully acknowledge the financial support provided by the Petchra Pra Jom Klao Master's Degree Scholarship from King Mongkut's University of Technology Thonburi and King Mongkut's University of Technology Thonburi (KMUTT), Thailand Science Research and Innovation (TSRI), and National Science, Research and Innovation Fund (NSRF) Fiscal year 2027.

Authors' Contributions: Conceptualization, S Tongthong, S Kongwat, and P Jongpradist; Methodology, S Tongthong, S Kongwat, and P Jongpradist; Investigation, S Tongthong, S Kongwat, and P Jongpradist; Writing - original draft, S Tongthong, S Kongwat, and P Jongpradist; Writing - review & editing, S Kongwat and P Jongpradist; Funding acquisition, P Jongpradist; Resources, S Tongthong; Supervision, S Kongwat and P Jongpradist.

Data availability statement: The data presented in this study are available on request from the corresponding author.

Editor: Marcilio Alves.

References

Abraham, N., Ghosh, B., Simms, C., Thomson, R., & Amato, G. (2016). Assessment of the impact speed and angle conditions for the EN1317 barrier tests. International journal of crashworthiness, 21(3), 211-221. https://doi.org/10.1080/13588265.2016.1164444

Abramowicz, W., & Wierzbicki, T. (1989). Axial crushing of multicorner sheet metal columns. Journal of Applied Mechanics, 56(1), 113–120. https://doi.org/10.1115/1.3176030

Aronov, B., Har-Peled, S., Knauer, C., Wang, Y., & Wenk, C. (2006). Fréchet distance for curves, revisited. In European symposium on algorithms (pp. 52-63). Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/11841036_8

Avbelj, J., Müller, R., & Bamler, R. (2014). A metric for polygon comparison and building extraction evaluation. IEEE Geoscience and Remote Sensing Letters, 12(1), 170-174. https://doi.org/10.1109/LGRS.2014.2330695

Cai, S., Wang, Z., Wang, S., Perdikaris, P., & Karniadakis, G. E. (2021). Physics-informed neural networks for heat transfer problems. Journal of Heat Transfer, 143(6), 060801. https://doi.org/10.1115/1.4050542

Choiron, M. A., Yaqin, M. A., Mahzan, S., & Zin, M. (2019). Energy absorption and deformation pattern on two segment crash box with rubber connection variation. Journal of Advanced Research in Fluid Mechanics and Thermal Sciences, 55(1), 20-28.

Ciampaglia, A., Fiumarella, D., Niutta, C. B., Ciardiello, R., & Belingardi, G. (2021). Impact response of an origami-shaped composite crash box: Experimental analysis and numerical optimization. Composite Structures, 256. https://doi.org/10.1016/j.compstruct.2020.113093

Ciampaglia, A., Fiumarella, D., Boursier Niutta, C., Ciardiello, R., & Belingardi, G. (2024). Physics based data driven method for the crashworthiness design of origami composite tubes. International Journal of Mechanics and Materials in Design, 20(3), 445-462. https://doi.org/10.1007/s10999-023-09685-2

Dirgantara, T., Gunawan, L., Putra, I. S., Sitompul, S. A., & Jusuf, A. (2013). Numerical and experimental impact analysis of square crash box structure with holes. Applied Mechanics and Materials, 393, 447-452. https://doi.org/10.4028/www.scientific.net/AMM.393.447

Greff, K., Srivastava, R. K., Koutník, J., Steunebrink, B. R., & Schmidhuber, J. (2016). LSTM: A search space odyssey. IEEE transactions on neural networks and learning systems, 28(10), 2222-2232. https://doi.org/10.1109/TNNLS.2016.2582924

Grigorescu, S., Trasnea, B., Cocias, T., & Macesanu, G. (2020). A survey of deep learning techniques for autonomous driving. Journal of Field Robotics, 37(3), 362–386. https://doi.org/10.1002/rob.21918

Gulino, M. S., Di Gangi, L., Sortino, A., & Vangi, D. (2021). Injury risk assessment based on pre-crash variables: The role of closing velocity and impact eccentricity. Accident Analysis & Prevention, 150. https://doi.org/10.1016/j.aap.2020.105864

Guo, W., Xu, P., Yang, C., Guo, J., Yang, L., & Yao, S. (2023). Machine learning-based crashworthiness optimization for the square cone energy-absorbing structure of the subway vehicle. Structural and multidisciplinary Optimization, 66(8), 182. https://doi.org/10.1007/s00158-023-03629-2

Han, H., Sorokin, V., Tang, L., & Cao, D. (2022). Lightweight origami isolators with deployable mechanism and quasi-zero-stiffness property. Aerospace Science and Technology, 121. https://doi.org/10.1016/j.ast.2021.107319

Hershman, L. L. (2001). The US new car assessment program (NCAP): Past, present and future. National Highway Traffic Safety Administration.

Homsnit, T., Jongpradist, P., Kongwat, S., Jongpradist, P., & Thongchom, C. (2024). Crashworthiness design of an automotive S-rail using ANN-based optimization to enhance performance and safety. Structural and Multidisciplinary Optimization, 67(6), 93. https://doi.org/10.1007/s00158-024-03803-0

Hu, H., Qi, L., & Chao, X. (2024). Physics-informed Neural Networks (PINN) for computational solid mechanics: Numerical frameworks and applications. Thin-Walled Structures, 205. https://doi.org/10.1016/j.tws.2024.112495

Javaheri, A., Brites, C., Pereira, F., & Ascenso, J. (2020). A generalized Hausdorff distance based quality metric for point cloud geometry. In 2020 Twelfth international conference on quality of multimedia experience (QoMEX) (pp. 1-6). IEEE. https://doi.org/10.1109/QoMEX48832.2020.9123087

Jongpradist, P., Tongthong, S., Kongwat, S., Ruangjirakit, K., Thongchom, C., & Hasegawa, H. (2024). Optimizing functionally graded hexagonal crash boxes with honeycomb filler for enhanced crashworthiness. In Structures (Vol. 59, p. 105775). Elsevier. https://doi.org/10.1016/j.istruc.2023.105775

Karantza, K. D., & Manolakos, D. E. (2022). Crashworthiness analysis of square aluminum tubes subjected to oblique impact: Experimental and numerical study on the initial contact effect. Metals, 12(11), 1862. https://doi.org/10.3390/met12111862

Khan, S., & Yairi, T. (2018). A review on the application of deep learning in system health management. Mechanical Systems and Signal Processing, 107, 241–265. https://doi.org/10.1016/j.ymssp.2017.11.024

Kohar, C. P., Connolly, D. S., Liusko, T., & Inal, K. (2020). Using artificial intelligence to aid vehicle lightweighting in crashworthiness with aluminum. In MATEC Web of conferences (Vol. 326, p. 01006). EDP Sciences. https://doi.org/10.1051/matecconf/202032601006

Kohar, C. P., Greve, L., Eller, T. K., Connolly, D. S., & Inal, K. (2021). A machine learning framework for accelerating the design process using CAE simulations: An application to finite element analysis in structural crashworthiness. Computer Methods in Applied Mechanics and Engineering, 385. https://doi.org/10.1016/j.cma.2021.114008

Le Guen, V., & Thome, N. (2019). Shape and time distortion loss for training deep time series forecasting models. Advances in neural information processing systems, 32. https://doi.org/10.48550/arXiv.1909.09020

Li, Z., Yu, Q., Zhao, X., Yu, M., Shi, P., & Yan, C. (2017). Crashworthiness and lightweight optimization to applied multiple materials and foam-filled front end structure of auto-body. Advances in Mechanical Engineering, 9(8). https://doi.org/10.1177/1687814017702806

Li, Z., Ma, W., Zhu, H., Deng, G., Hou, L., Xu, P., & Yao, S. (2022). Energy absorption prediction and optimization of corrugation-reinforced multicell square tubes based on machine learning. Mechanics of Advanced Materials and Structures, 29(26), 5511-5529. https://doi.org/10.1080/15376494.2021.1958032

Liang, H., Hao, W., Sun, H., Pu, Y., Zhao, Y., & Ma, F. (2022). On design of novel bionic bamboo tubes for multiple compression load cases. International Journal of Mechanical Sciences, 218. https://doi.org/10.1016/j.ijmecsci.2022.107067

Liu, C. X., Wang, X., Liu, W., Yang, Y. F., Yu, G. L., & Liu, Z. (2024). A physics-informed neural network for Kresling origami structures. International Journal of Mechanical Sciences, 269. https://doi.org/10.1016/j.ijmecsci.2024.109080

Ma, F., Liang, H., Pu, Y., Wang, Q., & Zhao, Y. (2021). Multi-objective optimization of crash box filled with three-dimensional cellular structure under multi-angle impact loading. Proceedings of the Institution of Mechanical Engineers, Part D: Journal of automobile engineering, 235(9), 2397-2412. https://doi.org/10.1177/0954407021998174

Meng, X., Qin, Y., Xie, L., He, P., & Zhu, L. (2025). Hierarchical Prediction of Subway-Induced Ground Settlement Based on Waveform Characteristics and Machine Learning with Applications to Building Safety. Buildings, 15(18). https://doi.org/10.3390/buildings15183390

Mohammadiha, O., Beheshti, H., & Aboutalebi, F. H. (2015). Multi-objective optimisation of functionally graded honeycomb filled crash boxes under oblique impact loading. International journal of crashworthiness, 20(1), 44-59. https://doi.org/10.1080/13588265.2014.970398

Ran, H., Huang, H., & Deng, X. (2024). Crashworthiness of double-gradient hierarchical multi-cell hexagonal tubes under multi-load impacts. Archives of Civil and Mechanical Engineering, 25(1), 30. https://doi.org/10.1007/s43452-024-01087-1

Siami-Namini, S., Tavakoli, N., & Namin, A. S. (2019). The performance of LSTM and BiLSTM in forecasting time series. In 2019 IEEE International conference on big data (Big Data) (pp. 3285-3292). IEEE. https://doi.org/10.1109/BigData47090.2019.9005997

Wang, H., Tan, P.F.M., Foo, Y.Z., Owyong, S., Lu, Z., Chen, H.T., Nguyen, T.H., & Lin, A. (2025). Integrating graph neural networks with physics-informed loss function for mechanical response prediction of hollow concrete structures with morphed honeycomb configurations. Materials & Design, 251. https://doi.org/10.1016/j.matdes.2025.113659

Wang, J., Ma, Y., Zhang, L., Gao, R. X., & Wu, D. (2018). Deep learning for smart manufacturing: Methods and applications. Journal of Manufacturing Systems, 48, 144–156. https://doi.org/10.1016/j.jmsy.2018.01.003

Wang, X., Jiang, T., Ji, C., Zhao, C., Wang, Y., & Gao, Z. (2023). Mechanical properties and constitutive modeling of 6063-T5 aluminum alloy over a wide range of strain rates and temperatures. International Journal of Thermophysics, 44(9), 133. https://doi.org/10.1007/s10765-023-03239-6

Wierzbicki, T., & Abramowicz, W. (1983). On the crushing mechanics of thin-walled structures. Journal of Applied Mechanics, 50(4a), 727–734. https://doi.org/10.1115/1.3167137

Xu, X., Zhang, Y., Wang, J., Jiang, F., & Wang, C. H. (2018). Crashworthiness design of novel hierarchical hexagonal columns. Composite Structures, 194, 36-48. https://doi.org/10.1016/j.compstruct.2018.03.099

Yang, Y., Wu, Q. J., & Wang, Y. (2016). Autoencoder with invertible functions for dimension reduction and image reconstruction. IEEE Transactions on Systems, Man, and Cybernetics: Systems, 48(7), 1065-1079. https://doi.org/10.1109/TSMC.2016.2637279

Yang, Y., Feng, C., Shen, Y., & Tian, D. (2018). Foldingnet: Point cloud auto-encoder via deep grid deformation. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 206-215). https://doi.org/10.1109/CVPR.2018.00029

Ying, X. (2019). An overview of overfitting and its solutions. In Journal of physics: Conference series (Vol. 1168, No. 2, p. 022022). IOP Publishing. https://doi.org/10.1088/1742-6596/1168/2/022022

Zhang, W., Yang, D., & Wang, H. (2019). Data-driven methods for predictive maintenance of industrial equipment: A survey. IEEE Systems Journal, 13(3), 2213–2227. https://doi.org/10.1109/JSYST.2019.2905565

Zhou, C., Jiang, L., Tian, K., Bi, X., & Wang, B. (2017). Origami crash boxes subjected to dynamic oblique loading. Journal of Applied Mechanics, 84(9), 091006. https://doi.org/10.1115/1.4037160