Journal of Applied Computational Intelligence

eISSN: 3136-6937

Vol. 4, No. 1, January, 2026

Predicting Lattice Crushing Response with Contact State Gated MeshGraphNets

Vasilica Gheorghiu[1], Rareș Ciobanu[1], Horia Georgescu[2], Livia Kovași[2,*]

[1] *Faculty of Automatic Control and Computers, National University of Science and Technology Politehnica Bucharest, Bucharest, 060042, Romania*

[2] *Faculty of Automatics, Computer Science and Electronics, University of Craiova, Craiova, 200440, Romania*

Received 12 November 2025; Revised 21 March 2026; Accepted 26 March 2026; Published 30 March 2026

Abstract

Lattice crushing involves evolving load paths, irreversible deformation, and intermittent contact between members that are distant in the undeformed mesh. A surrogate restricted to structural connectivity may therefore miss the interactions that govern compaction. This paper proposes a contact state gated MeshGraphNets model for predicting nodal motion and platen reaction during displacement-controlled compression. The formulation combines permanent mesh edges with dynamically updated surface contact edges, using geometric gaps, relative velocities, and edge persistence to distinguish approaching, sustained, and separating interactions. A short rollout objective couples state prediction with reaction consistency and penetration control. The evaluation design considers 480 trajectories spanning three lattice families, with geometry separated training, validation, and test partitions. Across the comparison, median normalized displacement error decreases from 8.4% for a mesh only baseline to 3.2% for the proposed model, while reaction curve error decreases from 12.6% to 4.8%. Maximum penetration decreases from 0.31 to 0.06 mm, and the proposed rollout requires 18 s against a 720 s reference solve. The evaluation tracks both local geometric admissibility and the global force response throughout compaction. The proposed protocol distinguishes contact representation from processor capacity through controlled ablations and evaluates compaction, generalization, and computational cost separately. The formulation establishes a testable route toward contact sensitive surrogates, with separate criteria for prediction accuracy, contact admissibility, and computational efficiency. The present scope is monotonic compression of intact elastoplastic lattices; fracture, unloading, and impact conditions require separate validation.

Keywords: Lattice Crushing, Self-Contact, MeshGraphNets, Contact State Gating, Autoregressive Surrogates

1. Introduction

Architected lattices are used as candidate energy absorbers because their internal geometry can redistribute deformation over a prescribed compression stroke [1]. Relationships between cell topology and progressive collapse are commonly interpreted through bending dominated and stretching dominated deformation mechanisms [2]. In practical absorbers, the response also depends on geometric imperfections, boundary restraint, and the sequence in which individual members lose stability [3]. Studies of graded cellular architectures relate these dependencies to the spatial organization of collapse bands [4]. Engineering assessment requires the reaction history, the extent of compaction, and the location of concentrated deformation alongside stiffness to determine whether a proposed geometry is suitable for a particular loading envelope. Analyses of lattice energy absorption consequently distinguish the initial peak, the crushing

* Corresponding author: kovasi.li@ace.ucv.ro

http://dx.doi.org/10.68406/jaci.2026.vol4iss1n13:174-189

➢ This is an open access article under the CC BY-NC-SA 4.0 license (https://creativecommons.org/licenses/by/4.0/).

© Authors retain all copyrights.

plateau, and the transition toward densification. Repeated numerical evaluation becomes expensive when this assessment is embedded in a search over cell dimensions, relative density, and manufacturing tolerances. A reduced evaluation cost would permit broader design exploration when the surrogate reliably retains the interactions responsible for the force response throughout the compression stroke. Predictions that reproduce early elastic deformation while losing the later contact network provide an incomplete basis for selecting an absorber. This difficulty becomes particularly visible when neighboring layers fold into one another and the active load path changes during compression.

Detailed nonlinear finite element analysis provides a flexible reference framework for such complex contact dominated behavior. Contact formulations for large deformation separate geometric search from the enforcement of normal and tangential interaction constraints [5]. Their computational requirements increase when many surfaces approach, separate, and reconnect during the same loading sequence. Guo et al. [6] additionally emphasized the effects of mesh resolution and constitutive calibration on predicted collapse modes in their numerical investigations of crushing. The discretization and modeling assumptions of the reference calculations remain embedded in a surrogate with reduced inference cost. A compact input output map supports efficient prediction of scalar design targets, whereas reconstruction of local deformation sequences requires a spatially resolved representation of the structure. Reduced representations built around fixed modes face a different limitation when localized folding moves across the specimen [7]. A useful predictive model must therefore accommodate the changing interaction pattern while preserving the reference mesh and its representation of material connectivity throughout deformation. Its applicability should be evaluated separately across topology and friction regimes, with each transfer claim tied to an explicitly defined test range. The central challenge is to represent the transition from separated members to load bearing contact in a form that remains compatible with sequential prediction.

Graph based simulators offer a natural representation of an irregular discretization. Storm et al. [8] developed learned message passing to approximate local interactions while retaining variable particle or mesh connectivity. MeshGraphNets extends this idea to mesh based physical systems and already includes spatial interaction edges for collision and self-contact [9]. The established second edge set provides the starting point for investigating richer contact state information. Lattice crushing introduces a specific representation challenge because small prediction errors can repeatedly create and remove a spatial edge near a contact threshold. Studies of long horizon learned simulation associate such local errors with accumulated drift under autoregressive integration [10]. An abrupt change in neighborhood membership can introduce another source of sensitivity near a contact threshold. Physical admissibility requires separate evaluation of surface overlap and platen reaction consistency alongside the predicted nodal trajectory. Research on physically constrained learned mechanics motivates evaluating these defects alongside predictive error. A targeted model should distinguish an approaching surface from a sustained contact and preserve enough temporal context to avoid unstable switching. Its assessment should isolate that representation choice from changes in network width, training exposure, and baseline contact handling, with geometric contact detection held consistent across the contact capable variants. The foundational MeshGraphNets framework was introduced by Pfaff and colleagues in Learning Mesh-Based Simulation with Graph Networks (ICLR 2021); the recent contact surrogate cited here addresses a related subsequent development.

As lattice members fold, approach adjacent surfaces, and establish new load paths during compression, their evolving contact states provide a direct basis for extending MeshGraphNets toward the prediction of both local deformation and global crushing response. The proposed distinction lies in combining gap dependent messages, bounded persistence information, and hysteretic neighborhood updates within a controlled comparison against both mesh only and ordinary proximity edge models. The comparison examines how contact attributes influence local deformation, global reactions, and accumulated rollout error across geometry separated partitions and controlled architectural variants. The study is organized around the relationship between contact representation, sequential prediction, and engineering observables. Section 2 examines related approaches to crushing analysis, graph-based simulation, and contact sensitive learning. Section 3 defines the reference problem and the graph representation, including the geometry required to construct nonlocal interaction edges. Section 4 describes gated message passing, the training objective, and the rollout procedure. Section 5 presents the evaluation design, analyzes the response data, and examines the relationship between contact representation, generalization, and inference cost. Section 6 states the methodological conclusions, limitations, and future research directions. This organization keeps the background separate from the proposed formulation and reserves quantitative comparisons for the evaluation section. The contribution is the coordinated treatment of contact state and edge persistence within an existing mesh and spatial graph framework.

2. Related Work

2.1. Surrogate modeling of lattice collapse

Surrogates for cellular structures can be organized by the information they predict. Response surface models estimate scalar quantities such as peak load or absorbed energy from a fixed list of design parameters. Work on topology dependent surrogate construction treats these descriptors as a compact approximation of geometric variation [11]. This

approach is efficient when the design family is narrow and the response can be evaluated through a few scalar targets. It becomes less descriptive when different local collapse paths produce similar global energy values. A force history enriches a scalar objective, while spatial localization requires a corresponding field representation. Field prediction is consequently relevant when the engineering decision depends on member interference, collapse band position, or residual clearance. Such tasks require correspondence between the predicted field and the actual discretized geometry. For design screening, the field output can locate member interference and concentrated folding that an integrated energy value alone cannot resolve.

Reduced order approaches retain selected deformation coordinates to reduce computational cost. Research on nonlinear reduced mechanics examines how basis construction and sampling affect accuracy outside the training manifold [12]. For progressive crushing, a basis dominated by smooth deformation may represent the onset of compression adequately while requiring substantial enrichment to describe concentrated folds. Data driven field models offer flexible deformation representations whose ability to resolve local interactions depends on the information supplied to the predictor. Investigations of learning based structural response prediction emphasize the dependence of performance on geometric coverage and the chosen output representation [13]. The remaining challenge concerns preserving mechanically relevant interactions within an efficient numerical approximation. A model must preserve the local information that differentiates competing collapse paths and should be tested on independently withheld geometry combinations.

2.2. Mesh graphs and evolving interactions

A mesh graph preserves a direct connection between simulation nodes and learned states. The encoder processor decoder architecture provides a reusable mechanism for transforming node and edge attributes into predicted state increments [14]. Structural edges transmit information between entities connected by the reference discretization, which is appropriate for local material interactions. Spatial interactions require an additional relation between members that become adjacent during deformation while occupying separate neighborhoods in the reference mesh. Existing mesh based learned simulators use world space edges for this purpose [15]. A specialized contact representation must specify the information conveyed by each link, its activation conditions, and the relationship between contact and structural messages.

The graph representation also introduces resolution and neighborhood sensitivities. A fixed distance threshold can produce different contact degrees when the surface mesh is refined, while a fixed neighbor count can omit physically relevant pairs in compacted regions. Work on multiscale graph simulation addresses the separate problem of transmitting information over long distances in a limited number of processor layers [16]. For lattice crushing, these concerns coexist: local contact may require a changing fine scale neighborhood, whereas global force redistribution can extend across many cells. Processor depth governs the reach of information transfer, while the local edge attributes determine how near surface interactions are distinguished. A meaningful comparison should therefore match latent width and processor depth before attributing an improvement to contact states. It should also include an ordinary proximity edge baseline. Matched width and depth control processor capacity, but different contact edge counts still produce different memory and runtime costs.

2.3. Stability and evaluation of learned mechanics

One step prediction and full trajectory prediction probe different properties. Noise based training and sequential objectives have been used to reduce the distribution shift between clean training states and imperfect predicted states [17]. Their relevance is pronounced near contact because small position errors can alter both the next target and the next graph. A model may have a low average one step error while repeatedly advancing contact onset or delaying separation during rollout. Evaluation must therefore retain the temporal ordering of states and report event sensitive quantities in addition to mean field accuracy. Randomly splitting neighboring frames from the same trajectory would weaken this assessment by exposing the model to nearly identical geometries on both sides of the partition. Geometry separated splits are needed to test whether learned interactions transfer beyond familiar collapse sequences.

Constraint aware learning introduces penalties or architectural restrictions associated with mechanical admissibility. Studies of learned constitutive and structural models distinguish empirical fit from the satisfaction of physical constraints [18]. In the present setting, penetration and boundary reaction consistency provide complementary diagnostics, with equilibrium and contact traction evaluated through dedicated mechanical checks. Excessive penetration weighting may suppress physical deformation, while compensating force errors may preserve integrated energy despite a displaced collapse band. These limitations motivate complementary observables covering field accuracy, global reaction, and local admissibility. The remaining research gap concerns a controlled examination of contact state information under sequential prediction, with matched baselines, geometry separated evaluation, and explicit contact diagnostics. The proposed formulation provides a testable model design for examining this representation gap. A smaller penetration

metric indicates improved geometric admissibility under the chosen tolerance; it does not establish exact contact complementarity.

3. Problem Definition and Contact Graph

3.1. Compression setting and state representation

The compression configuration consists of a finite lattice between a fixed lower platen and a translating upper platen. The reference problem is restricted to monotonic displacement control, an isotropic elastoplastic solid, and moderate loading rates for which inertial contamination can be monitored. The formulation assumes intact members and an isothermal response, preserving material node identity throughout the compression trajectory. The learned state contains nodal position, velocity, a reduced plastic history descriptor, boundary type, and the prescribed loading increment. Material properties and relative density are provided as global conditioning features. The reduced plastic history variable is a learned state descriptor rather than a complete constitutive state, limiting claims about unloading and other path changes.

For a node indexed by i, the initial position is retained separately from the current position. The graph at step t contains permanent structural edges and temporary contact edges. Structural edges follow the actual element connectivity. Surface triangles, surface normals, and member identifiers are retained as auxiliary geometric information for contact search. Eligibility for nonlocal self contact requires distinct opposing surfaces and satisfaction of the reference adjacency exclusion rule. The moving platen is represented by prescribed boundary nodes and analytic surface information. The graph topology and nodal state are expressed by the following equation:

$$G_t = \left(V, E_m \cup E_{c,t}\right), \boldsymbol{s}_{i,t} = \left(\boldsymbol{x}_{i,t}, \boldsymbol{v}_{i,t}, q_{i,t}, b_i\right) \tag{1}$$

In this equation, V denotes the node set, while E_{m} and $E_{\mathrm{c},t}$ encode structural and contact relations, respectively. The state components represent position, velocity, plastic history, and boundary identity. The Lagrangian representation preserves node identity as the material deforms. Position vectors are expressed relative to the lower platen frame to remove irrelevant rigid translation. Length and velocity inputs are normalized using scales fitted exclusively to the training trajectories. These normalization statistics remain fixed during validation and testing.

3.2. Surface aware contact candidates

Contact candidates are generated using a broad phase spatial index followed by a surface based narrow phase. The broad phase identifies nearby surface patches, while the narrow phase computes a signed gap and a local closest point direction. For coarse surface meshes, proximity to a triangle interior requires a surface-based distance calculation to detect interactions between a node and the opposing face. The proposed implementation therefore associates each candidate with a closest point interpolation on the opposing triangle. Messages sent to that point are distributed to its vertices using barycentric weights. Barycentric message interpolation specifies how information enters the graph; it does not by itself impose equal and opposite contact tractions.

A two-threshold rule distinguishes edge creation from edge deletion. A previously absent edge is created when the gap falls below an entry threshold. An existing edge persists until the gap exceeds a larger exit threshold. Both thresholds define positive search distances; physical penetration is assessed separately against the reference tolerance. The interval between them reduces repeated graph switching near the entry distance. A bounded age variable records how long a candidate has persisted, and the normal relative velocity differentiates approach from separation. These quantities condition the learned gate, while the reference contact law defines the mechanical interaction targets. Candidate membership evolves according to the equation Candidate age measures persistence in the search neighborhood and should not be interpreted as the duration of nonzero contact traction.

$$c_{ij,t} = 1\left[g_{ij,t} < r_{in}\right] \vee \left(c_{ij,t-1} \wedge 1\left[g_{ij,t} < r_{out}\right]\right) \tag{2}$$

The Boolean indicator in this equation records candidate membership. The signed gap $g_{ij,t}$ is positive for separation, and the thresholds satisfy $r_{\mathrm{out}} > r_{\mathrm{in}}$. Figure 1 outlines the transformation from a surface mesh to the two edge streams, including adjacency exclusion, closest point evaluation, state assignment, and reciprocal assembly. A cell list is rebuilt after each predicted update, with a conservative search margin for the next prescribed displacement increment. The search margin follows the maximum expected surface motion to retain candidate coverage between saved frames. The main evaluation retains every eligible candidate pair. If the assumed motion bound is exceeded, a candidate can be missed between saved states; the search margin therefore remains a coverage assumption rather than a guarantee.

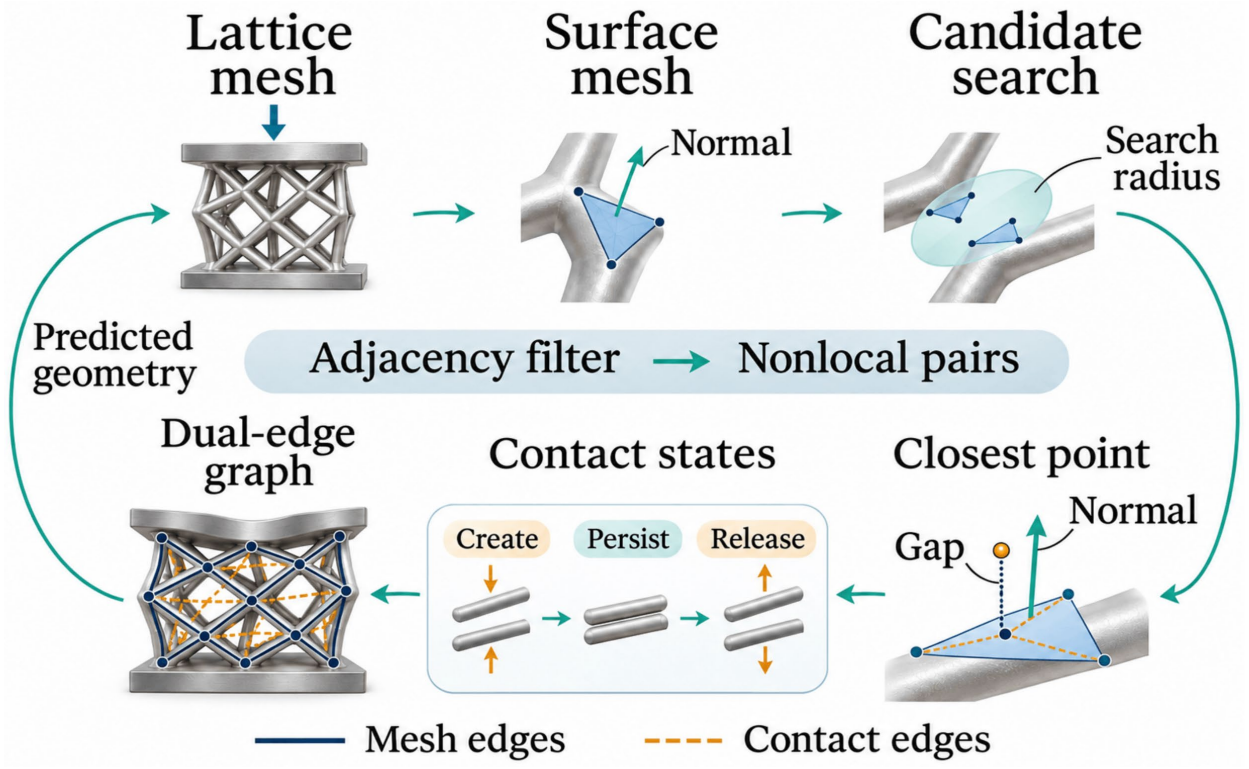


Figure 1. Surface aware graph construction workflow from lattice mesh and platen state to structural edges, geometric contact candidates, hysteretic membership, and contact state attributes

3.3. Targets and engineering observables

The network predicts a displacement increment for each unconstrained node and an auxiliary reaction contribution for platen adjacent nodes. Nodal displacements reconstruct the deformed geometry and update the contact search. Reaction contributions are summed over the loaded boundary to obtain the global compression force. This auxiliary output provides direct reaction supervision, with equilibrium evaluated through separate mechanical diagnostics. The reaction loss and field loss consequently serve distinct purposes. Plastic history is advanced through a bounded decoder increment within the selected scalar history representation and monotonic loading domain. The summed reaction is a supervised boundary observable rather than a force reconstructed from verified internal stress equilibrium.

Field accuracy is evaluated with volume weighted nodal displacement errors normalized by the maximum applied stroke. This fixed normalization avoids unstable percentages near the undeformed state. Force history accuracy is measured relative to the reference force norm over the complete saved sequence. Peak force, plateau force, absorbed energy, and densification onset are retained as separate engineering observables. The plateau interval is fixed before evaluation, and the densification criterion is applied identically to reference and surrogate curves. Maximum penetration is evaluated directly from the reconstructed surface geometry.

Every reference trajectory is treated as one indivisible sample for splitting and uncertainty estimation. Meshes with different realizations of the same design are grouped by design identifier before partitioning. Contact edge inputs comprise geometric and kinematic quantities available during inference; reference normal forces may serve as diagnostic targets. A subsequent traction-based diagnostic should apply a common target definition and input availability policy across all comparators. The proposed outputs therefore rely on geometry, current predicted state, and prescribed loading.

4. Contact State Gated MeshGraphNets

4.1. Dual stream message passing

Separate encoders map structural and contact attributes into a shared latent width. Structural attributes contain reference and current relative positions, their lengths, and a material connectivity tag. Contact attributes contain signed gap, normal direction, tangential relative velocity, normal relative velocity, bounded age, and closest point weights. Node features contain normalized state, boundary identity, and global conditioning variables. The architecture specification uses a latent width of 128, twelve processor blocks, and two hidden layers in each multilayer perceptron. These settings define the common processor configuration used to compare the contact representations.

A scalar gate modulates each contact message using current geometry and persistence information. It is bounded between zero and one and is computed independently from the structural message function. Contact candidates that remain separated can therefore carry a weak anticipatory message, whereas a sustained approaching interaction can produce a stronger one. The gate functions as a bounded message weight, with event probability interpretation requiring a separate calibration experiment. Its dependence on gap, normal relative velocity, and candidate age is represented through the equation A gate value of 0.5 indicates a message weighting factor and carries no probability meaning without an explicitly defined event and calibration test.

$$a_{ij,t} = sigmoid\left[\phi_a\left(g_{ij,t}, v^n_{ij,t}, \tau_{ij,t}\right)\right] \tag{3}$$

The learned function ϕ_a in this equation shares its parameters across contact edges. Tangential state enters the contact message, providing a distinct representation of sliding alongside the normal closure information used by the gate. Deterministic reference adjacency exclusions are applied before gate evaluation. Candidate deletion removes its contribution from aggregation; an active candidate may transmit a weak anticipatory message before physical closure. Structural and contact messages are aggregated separately, with bounded degree scaling applied before the node update to control sensitivity to increasing contact degree during compaction. The residual processor then follows this equation:

$$\boldsymbol{h}_i^{\ell+1} = \boldsymbol{h}_i^{\ell} + \phi_v\left(\boldsymbol{h}_i^{\ell}, \sum_j \boldsymbol{m}_{ij}^{m}, \sum_j a_{ij,t}\boldsymbol{m}_{ij}^{c}\right) \tag{4}$$

Within the processor equation, $\mathbf{h}_i^{\ell}$ denotes the node embedding at processor depth ℓ, and the message superscripts distinguish structural and contact streams. The message aggregates include the degree scaling specified above. A final decoder produces the displacement increment and reaction contribution. Figure 2 presents the inference and training workflow, with two edge encoders, repeated dual stream processing, boundary restoration, geometric checks, and the next contact search connected in temporal order. Reference targets enter the training branch, while inference proceeds from the predicted state and prescribed loading.

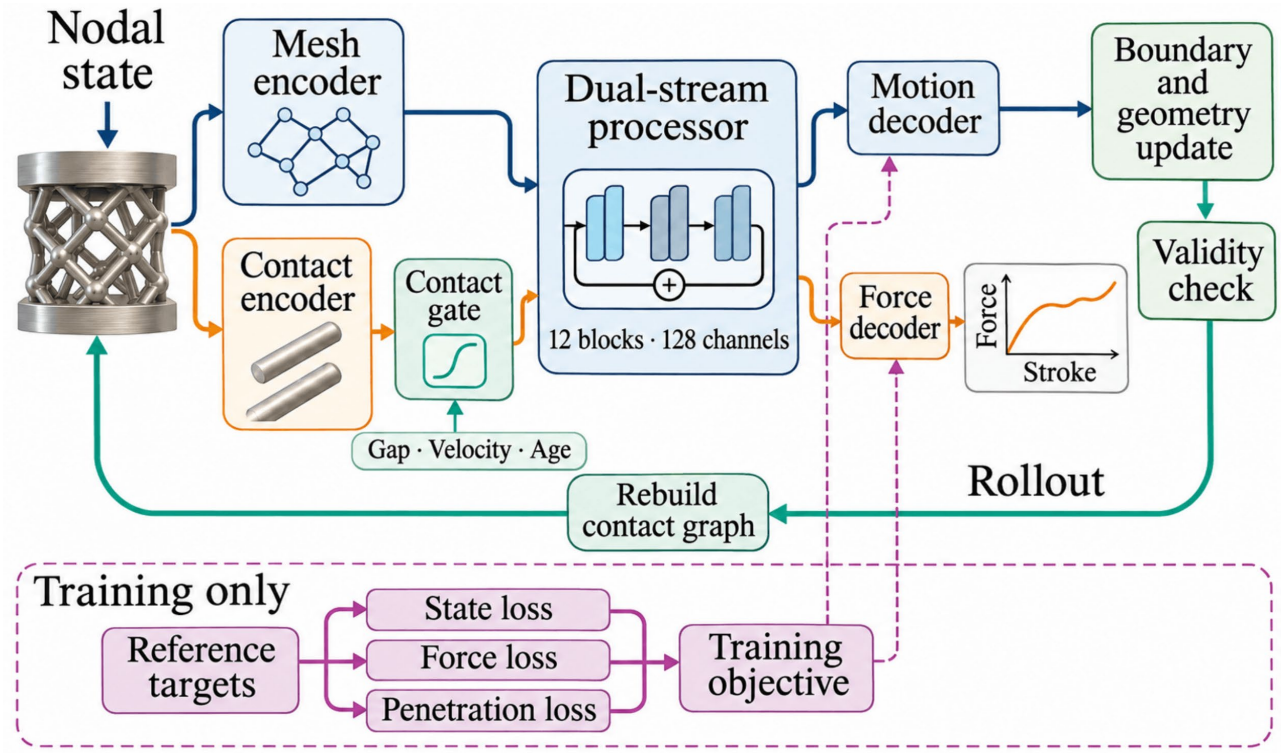


Figure 2. Contact state gated MeshGraphNets workflow with dual encoders, gated processor, displacement and reaction decoders, boundary restoration, rollout feedback, and training only supervision

4.2. Training objectives and temporal exposure

The proposed objective combines state increment accuracy, global reaction consistency, and a geometric penetration penalty over a short-predicted sequence. Volume weighting balances state targets across regions with different mesh resolutions. Force normalization uses a training scale fixed before testing. Penetration is penalized only beyond a small numerical tolerance, because the reference contact algorithm may itself admit controlled overlap. Loss weights are selected on validation trajectories and then frozen throughout test evaluation. These contributions enter the training

objective through the equation The penetration term is a soft regularizer: reducing overlap must be assessed together with displacement and reaction errors to detect an overly restrictive update.

$$\mathcal{L} = \mathcal{L}_{state} + \lambda_f \mathcal{L}_{force} + \lambda_p \mathcal{L}_{pen} \tag{5}$$

The state, force, and penetration terms in this equation represent normalized state error, normalized reaction error, and excess penetration cost, respectively. The coefficients λ_{f} and λ_{p} control their relative contributions. Dimensional consistency follows from normalization before weighting. The proposed schedule starts with one step prediction and increases to eight predicted steps, recomputing contact candidates at each step from predicted positions. Small input perturbations are applied to free node positions and velocities, with boundary states restored immediately afterward. Perturbation scales must remain below geometric feature sizes and be recorded in physical units as well as normalized units. Excessive perturbations can produce geometrically inadmissible training states and distort the learned response near contact. Eight-step exposure reduces the gap between training and inference conditions but does not certify stability over all 101 saved states.

Training batches contain entire short sequences sampled from trajectories, with contact onset and dense compaction deliberately represented. Sampling weights balance precontact frames against contact transitions while maintaining controlled exposure to each trajectory. The protocol proposes five independent initialization seeds and a fixed optimizer schedule. Early stopping uses a validation score combining displacement and force errors, while penetration remains a separate acceptance diagnostic. Independent initialization seeds are required to assess the reproducibility of the proposed representation. Between seed variation and between trajectory variation are reported separately because they describe different uncertainties.

4.3. Rollout control and implementation boundaries

During inference, the prescribed platen increment is applied first, candidate geometry is updated, and the current graph is encoded. The decoder predicts interior displacement increments and auxiliary reactions. Boundary conditions are restored exactly, plastic history is advanced, and the surface state is checked for invalid elements and excessive overlap. The next step uses this predicted state, including its accumulated errors. The main comparison evaluates the direct learned update, enabling geometric diagnostics to capture errors before any corrective projection. For unconstrained nodes, the predicted configuration advances through the displacement update equation Because the reported comparison omits corrective projection, a failed geometric check identifies a limitation of the direct predictor rather than evidence of a corrected feasible state.

$$\boldsymbol{x}_{i,t+1} = \boldsymbol{x}_{i,t} + \widehat{\Delta \boldsymbol{x}}_{i,t} \tag{6}$$

Here, $\widehat{\Delta \mathbf{x}}_{i,t}$ is the decoder prediction for the current loading increment. Applying the update equation produces the geometry used by the subsequent contact search. The computational burden includes broad phase indexing, closest point evaluation, neural processing, and output assembly. Contact search can become the dominant component when many folded surfaces occupy the same region. End to end timing must therefore include geometric processing as well as the neural kernel. The proposed benchmark includes graph updates and device synchronization, records peak memory, and separates model training from amortized inference. Mesh and contact edge counts are recorded at every saved state. Cost is analyzed jointly against mesh and contact edge counts.

Exact angular momentum balance, local constitutive consistency, and contact traction complementarity remain separate verification requirements. Reaction supervision targets the global boundary response, while dedicated constraints or mechanical checks are required to establish these additional properties. Normal orientation errors, thin member underresolution, and rapid changes between saved states are additional failure mechanisms. The proposed implementation records candidate turnover, maximum overlap, and invalid surface counts to make those mechanisms inspectable. The deployment boundary is restricted to the loading, topology, material, and discretization ranges covered by validation.

5. Evaluation and Quantitative Analysis

5.1. Evaluation design and baseline response

The evaluation examines the relationship between contact representation, local deformation, and global crushing response. Table 1 defines an evaluation design containing 480 trajectories, divided into 336 training, 72 validation, and 72 test trajectories by geometry identity. Body centered cubic, octet, and reentrant lattices span relative densities of 0.10–0.30 and five cell sizes. Topology specific interpretations of deformation mechanisms provide the basis for retaining these families as separate evaluation strata [19]. Each trajectory contains 101 saved states, approximately 2500–12000 nodes, and a maximum compression strain of 0.65. Quasi static explicit analysis commonly uses kinetic and internal energy histories to assess inertial contamination [20]. The corresponding reference protocol therefore

requires energy checks and mesh refinement before surrogate errors are evaluated. The four learned variants share processor depth, latent width, normalization, and geometry partitions. Collision benchmark methodology distinguishes geometric detection from mechanical response prediction [21]. Candidate detection is consequently audited independently of displacement and force prediction. Grouped resampling accounts for dependence among observations from the same trajectory [22]. Any subsequent uncertainty assessment should preserve trajectory identity when resampling, with initialization variability examined through independent training seeds. These controls establish common comparison conditions for the response analysis presented below. The quantitative comparisons below should be distinguished from the reference convergence, uncertainty, and transfer checks identified as requirements or future work.

Table 1. Dataset configuration and evaluation settings

Parameter	Specification
Trajectory count	480
Partition	336 train / 72 validation / 72 test
Lattice families	Body centered cubic / octet / reentrant
Relative density	0.10–0.30
Saved states	101 per trajectory
Nodes	2500–12000
Maximum compression strain	0.65
Loading mode	Displacement control

The representative reaction histories in Figure 3(a) separate early stiffness agreement from late compaction sensitivity. The reference reaches 25.8 kN at the final strain station, compared with 25.4 kN for the full model and approximately 20.9 kN for the mesh only variant. Their final differences are 0.4 and 4.9 kN, respectively. The growing discrepancy after a strain of approximately 0.45 is consistent with the role of nonlocal interactions during compaction: folded members establish additional routes for transmitting load as the specimen compacts. Figure 3(b) reports cohort median displacement errors of 8.4%, 5.1%, 3.7%, and 3.2% for the four learned variants. The reduction from the structural graph to the proximity graph is 3.3 percentage points, whereas the additional reduction from proximity edges to the full formulation is 1.9 percentage points. This separation assigns the larger first improvement to access to spatial information and reserves the second comparison for state gating and the objective. Figure 3(c) gives force errors of 12.6%, 7.5%, 5.6%, and 4.8%. The full model has the lowest field and global response errors across the comparison cohort. The representative curves and the cohort metrics describe different aggregation levels, so their numerical differences require interpretation within those respective definitions. The 1.9 percentage point difference combines representation and objective changes and cannot be attributed to gating alone without the corresponding isolated ablation.

Figure 3(d) presents the trajectory level distribution underlying the median displacement errors. The mesh only distribution extends across approximately 6–10%, while the full model values concentrate around 3.2%, with a smaller absolute spread. The distribution resolves the upper error tail and the concentration of cases near an engineering acceptance threshold, complementing the central value. Figure 3(e) compares five reference energies of 75, 95, 115, 135, and 155 J with predicted values of 73, 93, 112, 132, and 151 J. Their signed differences are consistently negative, ranging from minus 2 to minus 4 J. The relative discrepancy remains near a few percent across the energy range, indicating a systematic force deficit across the specimen scales. Figure 3(f) resolves the displacement discrepancy by compression stage. The mesh only error rises from 2.1% early in the stroke to 14.8% late in the stroke, while the full model rises from 1.4% to 5.6%. The widening stage dependent separation is more informative about contact sensitivity than the early error alone. The late full model error exceeds the 5% target despite a lower overall median, identifying late compaction as a critical regime for further improvement. For screening decisions governed by late compaction, the 5.6% stage error should take precedence over the more favorable full-trajectory median.

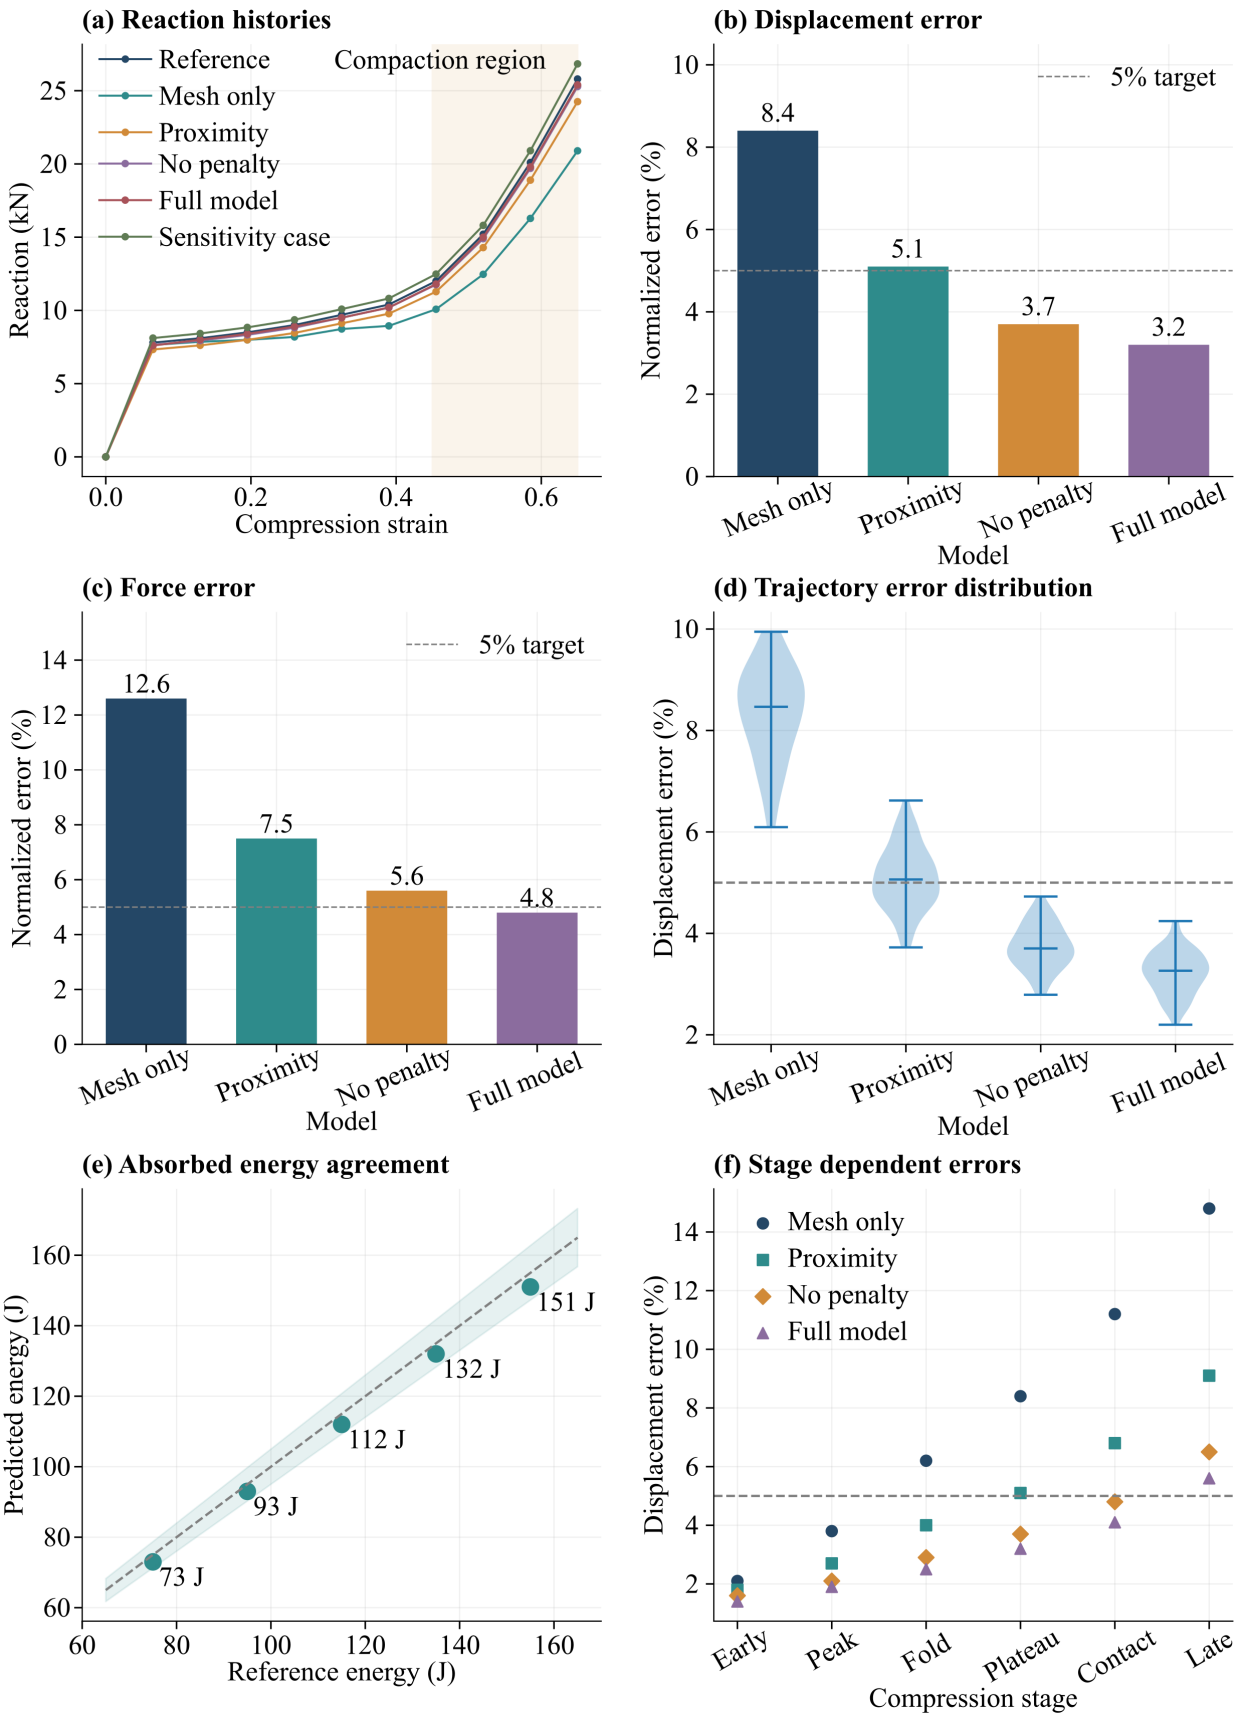


Figure 3. Baseline response comparison: (a) reaction histories; (b) displacement error by model; (c) force error by model; (d) displacement error distributions; (e) absorbed energy agreement; (f) error by compression stage

Table 2 connects the field and force comparisons to geometric admissibility and inference effort. Maximum penetration decreases from 0.31 mm for the mesh only model to 0.18 mm with proximity edges, 0.12 mm with gating without the penalty, and 0.06 mm for the full model. The last change halves the maximum overlap while displacement

error changes from 3.7% to 3.2%. Joint evaluation of penetration and displacement distinguishes improved contact admissibility from excessive suppression of the physical deformation. The reported pair of changes illustrates the desired direction of a balanced objective, with the acceptance decision still dependent on local thickness and reference contact tolerance. Runtime increases from 12 to 18 s across the same sequence. The additional 6 s represents a 50% increase over the mesh only surrogate, reflecting the cost of candidate construction and contact processing. For an eventual solver-based evaluation, overlap persistence and affected surface fraction should accompany the maximum. A single coarse patch may control the extreme value, whereas repeated moderate overlap across many members may have a larger influence on the reaction plateau. The test records should retain those locations and connect them to the corresponding force residuals at the same compression stations. The 0.06 mm maximum cannot establish acceptability without the corresponding member thickness, reference tolerance, and spatial extent of overlap.

Table 2. Median prediction errors, maximum penetration, and full trajectory inference times

Model	Displ. error %	Force error %	Max overlap mm	Time s
Mesh only	8.4	12.6	0.31	12
Proximity edges	5.1	7.5	0.18	16
Gated no penalty	3.7	5.6	0.12	17
Full proposed	3.2	4.8	0.06	18

5.2. Contact mechanisms and search sensitivity

The contact study separates geometric neighborhood construction, message content, and temporal persistence. Established assessments of cellular crushing distinguish plateau resistance from densification using explicitly defined response criteria [23]. The proposed comparison fixes those criteria across all variants to preserve a common assessment interval for contact quality. Contact enforcement formulations also introduce algorithm dependent tolerances that influence admissible overlap [24]. Penetration measurements should therefore be interpreted against the reference tolerance and minimum member thickness. The five controlled contact variants use distance only attributes, gap and velocity attributes, gap and history attributes, gated messages without hysteresis, and the complete representation. These variants share the surface search and prescribed loading sequence. Hysteretic switching uses distinct activation and release thresholds to reduce repeated state changes around a transition [25]. Here, that mechanism controls candidate membership, while the gate controls the magnitude of an existing message. The mechanical contribution of a retained edge depends on its gate: small weights limit its influence, whereas persistently high weights can delay separation. Calibration methodology compares predicted event scores with observed event frequencies [26]. Its application to an auxiliary contact classifier is examined separately from the gate, whose primary role is message modulation.

Figure 4(a) places six variants against five diagnostic quantities. Displacement errors span 3.2–8.4%, force errors span 4.8–12.6%, and the annotated turnover values decrease from 24% for the structural reference diagnostic to 8% for the full contact formulation. The structural baseline entry records turnover from an accompanying geometric candidate diagnostic, while its predictor operates solely on structural edges. Colors are normalized within each diagnostic, allowing the relative ordering to remain visible despite the different units; the raw annotations preserve the magnitude needed for engineering interpretation. The full model combines the lowest displacement, force, and penetration values with the highest runtime in the matrix at 18 s. This ordering highlights the computational cost of retaining and processing contact state. Figure 4(b) examines penetration distributions for the five controlled variants, with target central values of 0.16, 0.12, 0.10, 0.08, and 0.06 mm. The distance only example lies above the 0.10 mm tolerance through much of its distribution. Adding motion or persistence shifts the distribution toward that threshold, while the complete representation places its central region below it. The comparison concerns typical penetration across the evaluation sample and is distinct from the trajectory maximum reported in Table 2. A typical penetration statistic can fall below the tolerance while an individual trajectory maximum still exceeds it; the two summaries are not interchangeable.

Figure 4(c) evaluates an auxiliary event score at five confidence bins. Predicted confidence values of 0.10, 0.30, 0.50, 0.70, and 0.90 correspond to observed fractions of 0.13, 0.29, 0.50, 0.69, and 0.87. The largest absolute calibration difference is 0.03 at the lowest and highest bins. The calibration residual changes sign across the confidence range, with event frequency slightly underestimated in the lowest bin and overestimated in the highest bin. The displayed error bars show the variation within each confidence bin in the calibration comparison. These auxiliary diagnostics address whether a score can support event interpretation; the mechanical effect of a gated message additionally depends on its vector content and subsequent aggregation. Figure 4(d) relates six entry radii from 0.6 to 1.6 mm to displacement errors of 5.8%, 4.3%, 3.5%, 3.2%, 3.4%, and 3.9%. Candidate counts rise from 12 thousand to 76 thousand over the sweep. Increasing the radius from 0.6 to 1.2 mm improves the displacement metric by 2.6 percentage points while increasing candidates by 26 thousand. Extending it from 1.2 to 1.6 mm adds another 38 thousand candidates and raises error by 0.7

percentage points, illustrating how greater neighborhood coverage can also increase irrelevant interactions. Without bin counts and a specified interval construction, the plot supports descriptive calibration assessment rather than a statistical calibration guarantee.

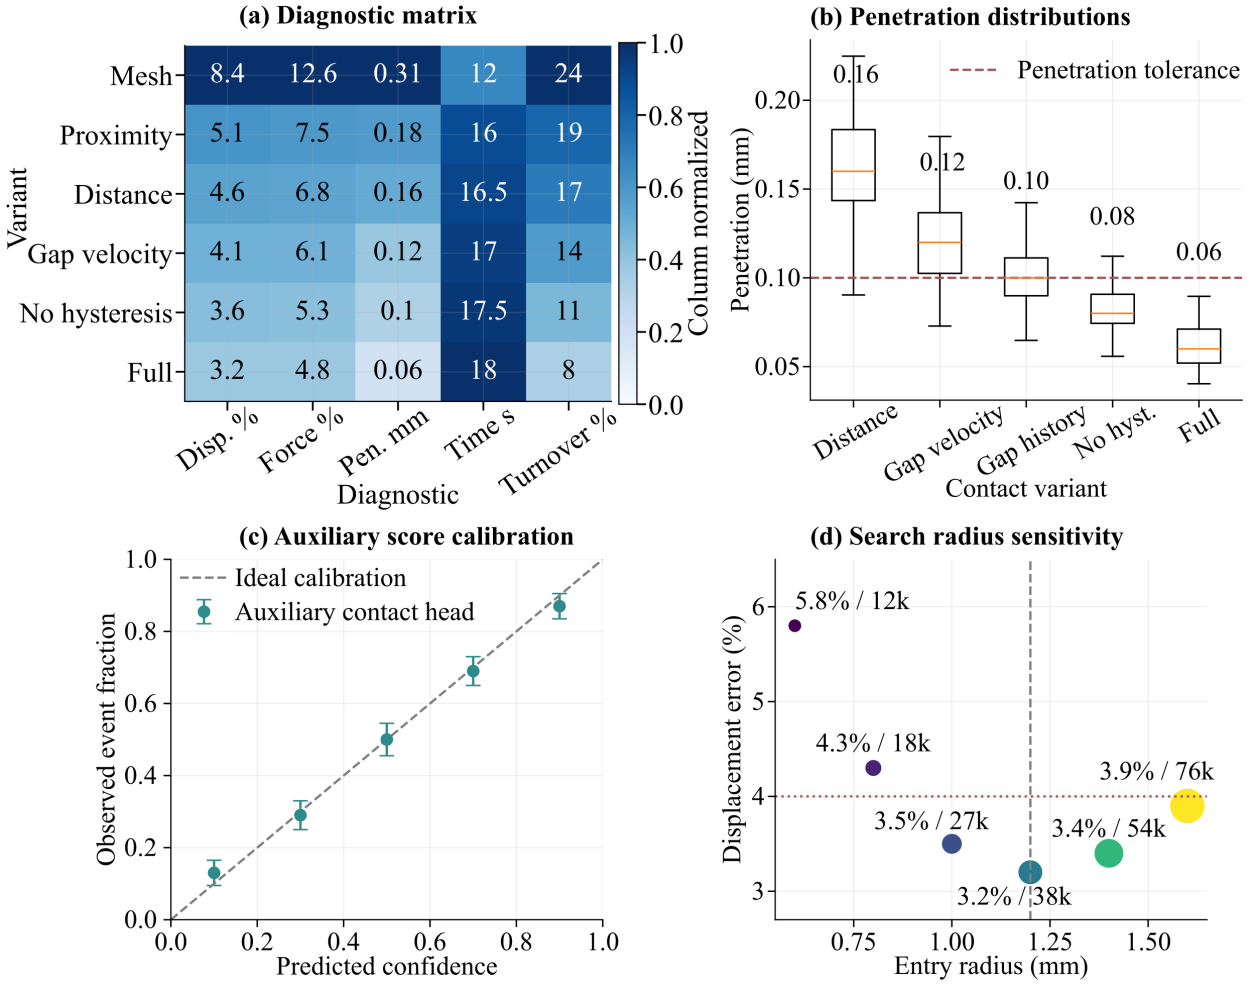


Figure 4. Contact mechanism sensitivity: (a) annotated diagnostic heatmap; (b) penetration distributions; (c) auxiliary contact score calibration; (d) entry radius tradeoff with candidate count represented by marker size

The radius example identifies a candidate tradeoff that should be examined at several mesh resolutions. A fixed physical distance encloses more surface nodes after refinement, changing both contact degree and the size of the aggregated message. Degree scaling controls latent magnitude growth, while surface discretization governs closest point geometry and normal accuracy. The planned follow up should therefore compare physical gap coverage, candidate count, and force error at a common specimen geometry. Persistence and hysteresis also require assessment under temporal refinement. Saving states too far apart can merge approach, closure, and separation into one interval, leaving candidate age poorly resolved. Halving the displacement increment would provide a direct test of whether the inferred contact states remain stable under finer sampling. A factorial ablation could cross gating, age input, hysteresis, and penetration weighting to identify redundant or interacting mechanisms. Within each comparison, initialization seeds and training exposure should remain matched. Gate distributions should then be stratified by approaching, sustained, and separating surface pairs. Saturation across all three states would indicate limited sensitivity of the contact channel to the intended state distinctions. The corresponding force residuals should be inspected around individual entry and release events, where a global average can conceal switching induced errors. The factorial study and temporal refinement described here are proposed follow-up analyses, not additional completed validation results.

5.3. Transfer behavior and computational scaling

Transfer is examined separately for relative density, lattice family, mesh resolution, and interaction conditions. Frictional contact laws couple tangential resistance to the local normal interaction and can alter folding sequences in cellular compression [27]. Friction is therefore fixed within each geometry comparison and varied through an independent sensitivity campaign. Distribution shift analysis distinguishes parameter interpolation from changes in the governing input distribution [28]. The density example follows that distinction by marking 0.30 as the upper training boundary and treating 0.35 as an extension. Figure 5(a) gives median displacement errors of 2.8%, 3.0%, 3.2%, 3.6%,

4.4%, and 6.9% across density levels of 0.10–0.35. Error increases by 1.6 percentage points within the training range and by another 2.5 percentage points over the final extension alone. The sharper final change motivates separate reporting of extrapolation. The displayed variation also expands toward the highest density, where the error bar reaches approximately 7.8%. A future reference campaign should inspect whether this increase coincides with earlier densification, more active surface pairs, or a different collapse mode. Those mechanisms require recorded fields and event histories, since density and contact count can change together. The 0.35 density point is an extension case outside the nominal 0.10–0.30 design range in Table 1 and is interpreted separately from within-range comparisons.

Figure 5(b) compares four force histories on a held-out geometry, with final values ranging from approximately 20.8 to 25.8 kN. The full model ends at 25.4 kN, while the proximity graph reaches approximately 23.5 kN. The difference becomes most apparent in the rising compaction branch, where the predicted contact network must sustain increasing resistance. The shaded band around the full model visualizes trajectory spread, enabling comparison of local response variability with the separation between the predicted and reference force histories. Figure 5(c) compares six processor settings with inference times of 8, 12, 16, 18, 25, and 32 s and displacement errors of 7.1%, 5.4%, 4.1%, 3.2%, 3.0%, and 2.9%. Pareto analysis retains separate accuracy and cost objectives when comparing numerical methods [29]. In this accuracy and cost tradeoff, increasing runtime from 8 to 18 s reduces error by 3.9 percentage points, whereas increasing it from 18 to 32 s yields only another 0.3 percentage points. The 12-block setting falls within the 20 s time budget and below the 4% error line. The larger settings supply a smaller marginal improvement and should be examined for memory demand and robustness before their additional cost is justified. The 20 s budget and 4% error line are comparison criteria for the displayed tradeoff, not universal limits for acceptable absorber design.

Figure 5(d) gives signed energy discrepancy distributions across five density groups, with central values from minus 1.8% to plus 2.1%. The change in sign indicates a density dependent bias: the lower density predictions tend to underestimate absorbed energy, while the higher density predictions tend to overestimate it. The shaded acceptance interval extends from minus 3% to plus 3%. The central values lie within the acceptance interval, while the distribution tails identify individual cases that exceed it and require trajectory level assessment.

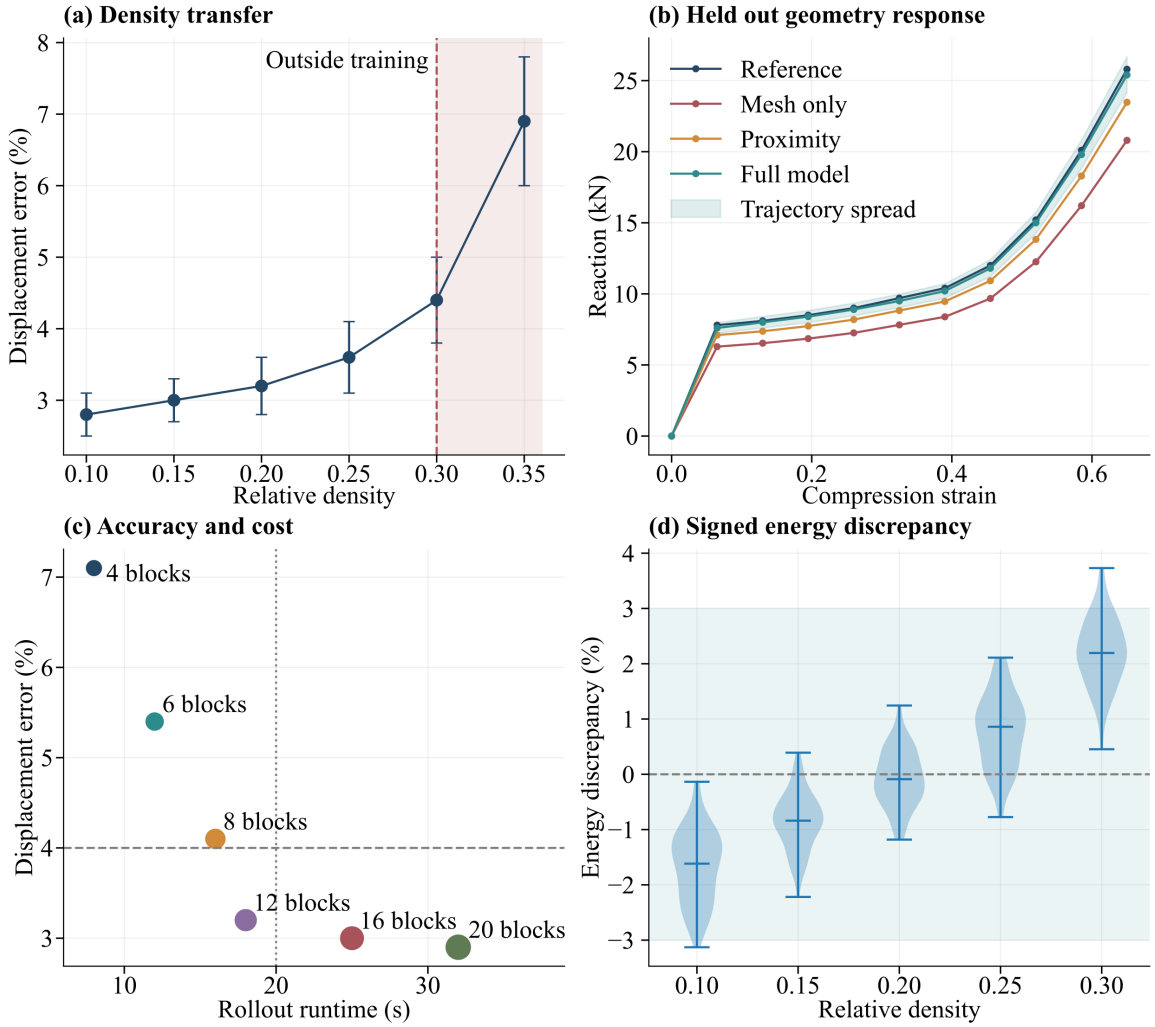


Figure 5. Transfer and efficiency analysis: (a) density dependent displacement error; (b) held out reaction histories; (c) accuracy versus runtime tradeoff; (d) signed absorbed energy discrepancy distributions

Figure 6 separates runtime into contact search, neural processing, and output assembly at five graph sizes from 2500 to 12000 nodes. Total times increase from 8 to 37 s, while search time rises from 2 to 13 s. The search share therefore grows from 25% to approximately 35%, reflecting the increasing burden of surface neighborhood construction. Benchmark methodology includes initialization, synchronization, and data movement when these operations affect end to end execution [30]. The 6000-node case requires 18 s against a 720 s reference solve, producing a ratio of forty under the specified end to end timing scope. The comparison includes graph rebuilding, neural evaluation, and output assembly within the trajectory timing. The ratio is descriptive of the reported timing pair; reproducible performance claims additionally require hardware, precision, software versions, and repeated timing statistics.

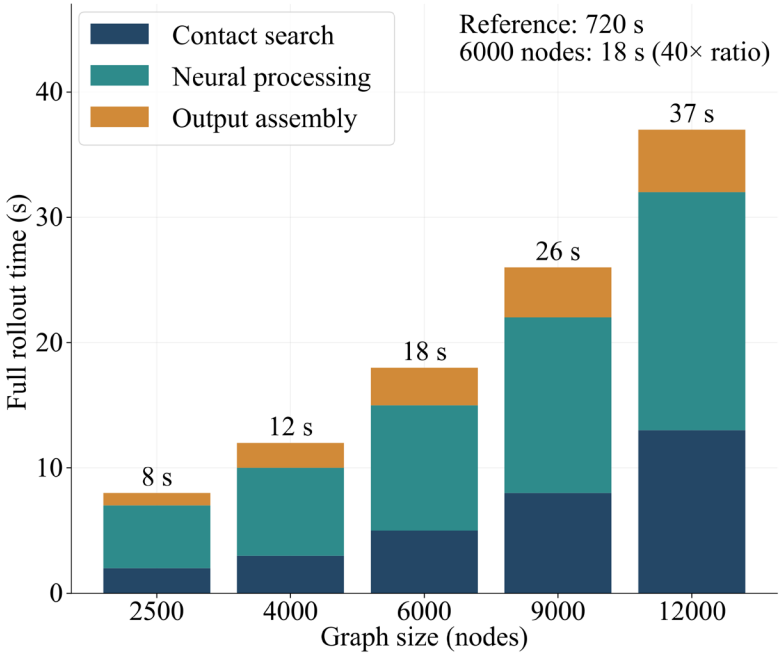


Figure 6. Runtime composition across five graph sizes with search, neural processing, and output assembly components, total time annotations, and a reference timing note

The scaling profile warrants inspection of the compacted configurations that produce the largest candidate neighborhoods. Between the 6000 and 12000 node examples, total runtime rises from 18 to 37 s, while contact search rises from 5 to 13 s. A targeted scaling assessment should hold the material node count approximately fixed while increasing compaction to span the dense candidate regime. Peak memory, closest point evaluations, and time per processor block should be recorded at each saved state. A second test should refine thin surface members while preserving the physical search distances, allowing geometric coverage and computational growth to be examined together. Candidate truncation, if required by memory limits, should be accompanied by the omitted pair fraction and the resulting overlap locations. Failure accounting should retain trajectories with excessive penetration, invalid surface elements, or numerical overflow. Their position within the density and topology ranges is needed to establish a usable acceptance boundary. For uncertain cases, a reference solver fallback would add cost to the surrogate runtime, so an engineering deployment study should report both the acceptance fraction and the average end to end evaluation time under that policy.

6. Conclusion

This study develops a contact state gated MeshGraphNets framework for lattice crushing prediction, with the evaluation indicating that dynamic surface interactions improve the representation of deformation, platen reaction, and geometric admissibility. These findings describe the reported comparison settings and do not establish uniform accuracy across untested lattice or contact regimes.

1. Prediction accuracy and contact admissibility. The full model achieves a median normalized displacement error of 3.2% and a reaction curve error of 4.8%, compared with 8.4% and 12.6%, respectively, for the mesh only baseline. Both error measures decrease by approximately 61.9%. Maximum penetration decreases from 0.31 to 0.06 mm, corresponding to an 80.6% reduction. These changes associate improved field prediction with closer agreement in global response and reduced surface overlap.

2. Contact representation and computational performance. Ordinary proximity edges reduce displacement error from 8.4% to 5.1%, while contact state gating and geometric regularization further reduce it to 3.2%. The additional improvement of 1.9 percentage points supports distinguishing contact availability from the information carried by

contact messages. The difference becomes more pronounced during compaction: late-stage displacement error decreases from 14.8% for the mesh only model to 5.6% for the full formulation. A complete rollout requires 18 s against a 720 s reference solve, yielding a fortyfold ratio under the specified timing scope.

3. Methodological contribution and applicability. The framework combines surface-aware candidate construction, gap- and velocity-dependent gating, candidate persistence, and hysteretic edge updates. These mechanisms address the evolving interaction topology during progressive crushing. Applicability is bounded by monotonic compression, intact members, the selected constitutive representation, and the validated geometry range. Exact equilibrium and contact complementarity require dedicated mechanical verification, while mesh resolution and temporal sampling influence candidate coverage and contact-state accuracy.

Future development should prioritize richer history representations, refinement assessment, and uncertainty-based acceptance criteria to support reliable screening of unfamiliar lattice designs. Late compaction and the density extension are immediate priorities because the reported errors increase most visibly in those regimes.

7. Declarations

7.1 Author Contributions

Vasilica Gheorghiu contributes to conceptualization, methodology, software, validation, analysis, data collection, draft preparation. Rareșa Ciobanu contributes to validation, analysis, investigation, data collection, draft preparation. Horia Georgescu contributes to conceptualization, validation, analysis. Livia Kovași to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. All authors have read and agreed with the manuscript before its submission and publication.

7.2 Funding and Acknowledgments

This research received no specific financial support from any funding agency.

7.3 Institutional Review Board Statement

Not applicable.

7.4 Informed Consent Statement

Not applicable.

7.5 Date Availability

Not applicable.

7.6 Declaration of Competing Interest

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

7.7 Declaration of generative AI and AI-assisted technologies

During the preparation of this manuscript, the authors utilized ChatGPT for linguistic polishing, logical organization and draft revision. The generative AI tool was not involved in research design, data collection, data analysis, or the equation of core conclusions. All outputs generated by AI were reviewed, revised and verified manually by the authors. The authors take full responsibility for all content of this manuscript.

8. References

[1] Bieler, S., & Weinberg, K. (2025). Energy absorption of sustainable lattice structures under static compression. Meccanica, 60(8), 2533–2547. https://doi.org/10.1007/s11012-025-02003-4

[2] Yu, Z., Savinov, R., Matura, M., Zhang, P., & Shi, J. (2025). Current research status on advanced lattice structures for impact and energy absorption applications: A systematic review. Thin-Walled Structures, 215, Article 113490. https://doi.org/10.1016/j.tws.2025.113490

[3] Ding, J., Ma, Q., Li, X., Zhang, L., Yang, H., Qu, S., Wang, M. Y., Zhai, W., Gao, H., & Song, X. (2024). Imperfection-enabled strengthening of ultra-lightweight lattice materials. Advanced Science, 11(41), Article 2402727. https://doi.org/10.1002/advs.202402727

[4] Daynes, S., & Feih, S. (2025). Functionally graded lattice structures with tailored stiffness and energy absorption. International Journal of Mechanical Sciences, 285, Article 109861. https://doi.org/10.1016/j.ijmecsci.2024.109861

[5] Maddio, P. D., Sinatra, R., & Cammarata, A. (2026). A novel 2D penalty formulation for frictionless contact. Multibody System Dynamics, 67(1), 221–245. https://doi.org/10.1007/s11044-025-10119-4

[6] Guo, Z., Yang, F., Li, L., & Wu, J. (2024). Bio-inspired curved-elliptical lattice structures for enhanced mechanical performance and deformation stability. Materials, 17(17), Article 4191. https://doi.org/10.3390/ma17174191

[7] Scheunemann, L., & Faust, E. (2026). A manifold learning approach to nonlinear model order reduction of quasi-static problems in solid mechanics: Proof of concept and parameter study. Computational Mechanics, 77(4), 983–1023. https://doi.org/10.1007/s00466-025-02696-0

[8] Storm, J., Rocha, I. B. C. M., & van der Meer, F. P. (2024). A microstructure-based graph neural network for accelerating multiscale simulations. Computer Methods in Applied Mechanics and Engineering, 427, Article 117001. https://doi.org/10.1016/j.cma.2024.117001

[9] Dubey, V. K., Haese, C. E., Gültekin, O., Dalton, D., Rausch, M. K., & Fuhg, J. (2026). Graph neural network surrogates for contacting deformable bodies with necessary and sufficient contact detection. Computer Methods in Applied Mechanics and Engineering, 448, Article 118413. https://doi.org/10.1016/j.cma.2025.118413

[10] Prakash, S. K. A., & Tucker, C. (2024). Probabilistic graph networks for learning physics simulations. Journal of Computational Physics, 513, Article 113137. https://doi.org/10.1016/j.jcp.2024.113137

[11] Gongora, A. E., Friedman, C., Newton, D. K., Yee, T. D., Doorenbos, Z., Giera, B., Duoss, E. B., Han, T. Y.-J., Sullivan, K., & Rodriguez, J. N. (2024). Accelerating the design of lattice structures using machine learning. Scientific Reports, 14(1), Article 13703. https://doi.org/10.1038/s41598-024-63204-7

[12] Vlachas, K., Garland, A., Quinn, D. D., & Chatzi, E. (2024). Parametric reduced-order modeling for component-oriented treatment and localized nonlinear feature inclusion. Nonlinear Dynamics, 112(5), 3399–3420. https://doi.org/10.1007/s11071-023-09213-z

[13] Kneifl, J., Fehr, J., Brunton, S. L., & Kutz, J. N. (2025). Multi-hierarchical surrogate learning for explicit structural dynamical systems using graph convolutional neural networks. Computational Mechanics, 75(3), 1115–1135. https://doi.org/10.1007/s00466-024-02553-6

[14] Würth, T., Freymuth, N., Zimmerling, C., Neumann, G., & Kärger, L. (2024). Physics-informed MeshGraphNets (PI-MGNs): Neural finite element solvers for non-stationary and nonlinear simulations on arbitrary meshes. Computer Methods in Applied Mechanics and Engineering, 429, Article 117102. https://doi.org/10.1016/j.cma.2024.117102

[15] Gladstone, R. J., Rahmani, H., Suryakumar, V., Meidani, H., D'Elia, M., & Zareei, A. (2024). Mesh-based GNN surrogates for time-independent PDEs. Scientific Reports, 14(1), Article 3394. https://doi.org/10.1038/s41598-024-53185-y

[16] Perera, R., & Agrawal, V. (2024). Multiscale graph neural networks with adaptive mesh refinement for accelerating mesh-based simulations. Computer Methods in Applied Mechanics and Engineering, 429, Article 117152. https://doi.org/10.1016/j.cma.2024.117152

[17] Pegolotti, L., Pfaller, M. R., Rubio, N. L., Ding, K., Brugarolas Brufau, R., Darve, E., & Marsden, A. L. (2024). Learning reduced-order models for cardiovascular simulations with graph neural networks. Computers in Biology and Medicine, 168, Article 107676. https://doi.org/10.1016/j.compbiomed.2023.107676

[18] Dornheim, J., Morand, L., Nallani, H. J., & Helm, D. (2024). Neural networks for constitutive modeling: From universal function approximators to advanced models and the integration of physics. Archives of Computational Methods in Engineering, 31(2), 1097–1127. https://doi.org/10.1007/s11831-023-10009-y

[19] Bernard, A. R., & ElSayed, M. S. A. (2024). Design, manufacturing, and analysis of periodic three-dimensional cellular materials for energy absorption applications: A critical review. Materials, 17(10), Article 2181. https://doi.org/10.3390/ma17102181

[20] Li, L., Yang, F., Guo, Z., Wu, J., Lu, G., & Fan, H. (2024). Truss-plate hybrid lattice metamaterials with broadband vibration attenuation and enhanced energy absorption. Virtual and Physical Prototyping, 19(1), Article e2345386. https://doi.org/10.1080/17452759.2024.2345386

[21] Dai, X., Kövecses, J., & Teichmann, M. (2024). Dynamic simulation and collision detection for flexible mechanical systems with contact using the floating frame of reference formulation. Journal of Computational and Nonlinear Dynamics, 19(11), Article 111005. https://doi.org/10.1115/1.4066329

[22] Fatoretto, M. B., Demétrio, C. G. B., Brophy, C., & de Andrade Moral, R. (2026). Bootstrap confidence intervals for generalized linear mixed models in dose-response analysis. Journal of Statistical Computation and Simulation, 96(14), 3630–3653. https://doi.org/10.1080/00949655.2026.2677023

[23] Bozkurt, I. (2024). Determination of deformation behaviors and energy absorption of lightweight composite lattice cylinders with different structures. International Journal of Polymer Science, 2024(1), Article 2280726. https://doi.org/10.1155/2024/2280726

[24] Du, Y., Li, Y., Coros, S., & Thomaszewski, B. (2024). Robust and artefact-free deformable contact with smooth surface representations. Computer Graphics Forum, 43(8), Article e15187. https://doi.org/10.1111/cgf.15187

[25] Raymundo, C., Pérez-Zuñiga, G., & Yupanqui, I. (2026). Phase-dependent nonlinear MPC for stair climbing: Edge-pivot control with a double-pendulum model. Results in Control and Optimization, 22, Article 100675. https://doi.org/10.1016/j.rico.2026.100675

[26] Kängsepp, M., Valk, K., & Kull, M. (2025). On the usefulness of the fit-on-test view on evaluating calibration of classifiers. Machine Learning, 114(4), Article 105. https://doi.org/10.1007/s10994-024-06652-6

[27] Gao, J., & McShane, G. J. (2026). Using heterogeneous fold patterns to manipulate the compressive collapse mechanisms of ductile origami-inspired metamaterials. International Journal of Mechanical Sciences, 314, Article 111377. https://doi.org/10.1016/j.ijmecsci.2026.111377

[28] Ma, J. (2024). A survey of out-of-distribution generalization for graph machine learning from a causal view. AI Magazine, 45(4), 537–548. https://doi.org/10.1002/aaai.12202

[29] Phan, Q. M., & Luong, N. H. (2024). Parameter-less Pareto local search for multi-objective neural architecture search with the Interleaved Multi-start Scheme. Swarm and Evolutionary Computation, 87, Article 101573. https://doi.org/10.1016/j.swevo.2024.101573

[30] Geiß, A., Hovi, T., Calotoiu, A., & Wolf, F. (2026). Validating the performance of GPU ports using differential performance models. Future Generation Computer Systems, 174, Article 108018. https://doi.org/10.1016/j.future.2025.108018