

000 MAVEN: A MESH-AWARE VOLUMETRIC ENCODING 001 NETWORK FOR SIMULATING 3D FLEXIBLE DEFOR- 002 MATION 003

006 **Anonymous authors**

007 Paper under double-blind review

010 ABSTRACT

013 Deep learning-based approaches, particularly graph neural networks (GNNs),
014 have gained prominence in simulating flexible deformations and contacts of solids,
015 due to their ability to handle unstructured physical fields and nonlinear regression
016 on graph structures. However, existing GNNs commonly represent meshes with
017 graphs built solely from vertices and edges. These approaches tend to overlook
018 higher-dimensional spatial features, e.g. 2D facets and 3D cells, from the original
019 geometry. As a result, it is challenging to accurately capture boundary representa-
020 tions and volumetric characteristics, though this information is critically important
021 for modeling contact interactions and internal physical quantity propagation, par-
022 ticularly under sparse mesh discretization. In this paper, we introduce MAVEN,
023 a **mesh-aware volumetric encoding network** for simulating 3D flexible deforma-
024 tion, which explicitly models geometric mesh elements of higher dimension to
025 achieve a more accurate and natural physical simulation. MAVEN establishes
026 learnable mappings among 3D cells, 2D facets, and vertices, enabling flexible mu-
027 tual transformations. Explicit geometric features are incorporated into the model
028 to alleviate the burden of implicitly learning geometric patterns. Experimental re-
029 sults show that MAVEN consistently achieves state-of-the-art performance across
030 established datasets and a novel metal stretch-bending task featuring large defor-
031 mations and prolonged contacts.

032 1 INTRODUCTION

035 Flexible deformation and contact of solids are prevalent across a wide range of real-world scenarios,
036 ranging from industrial manufacturing (Tao et al., 2025a), aeronautical engineering (Phanden et al.,
037 2021), to nuclear materials (Allen et al., 2012). Accurate modeling of these behaviors and their
038 evolution significantly enhances the understanding of these scenarios. Although many classical
039 numerical solvers, such as Finite Element Methods (Dhatt et al., 2012) and Material Point Methods
040 (Bardenhagen et al., 2004), have been developed for solid systems, achieving the desired accuracy
041 incurs high computational costs, as they rely on fine meshes and small time steps due to low-order
042 approximations and the iterative solution of large linear systems. Recently, deep learning (DL)
043 has rapidly emerged as a powerful tool for efficient physical simulation, showing great potential,
044 particularly in molecular dynamics (Jumper et al., 2021), fluid simulations (Li et al., 2020a), and
045 structural solid deformations (Tao et al., 2025b).

046 Among these DL-based solvers, graph neural networks (GNNs) have demonstrated superior per-
047 formance in the domain of solid deformation, due to their ability to handle dynamic unstructured
048 meshes and perform nonlinear regression on graphs (Sanchez-Gonzalez et al., 2020; Gao et al., 2022;
049 Gladstone et al., 2024). To handle irregular solution domains, existing GNN-based methods input
050 unstructured meshes, where the geometry is discretized into multiple connected simple cells using a
051 regular polyhedron (Figure 1(a)). GNNs abstract mesh vertices into graph nodes to capture internal
052 physical interactions, using edges defined by mesh connectivity, shown in Figure 1(b). Inter-object
053 contact is typically detected via an interaction radius. Mesh edges serve as pathways for propagating
physical quantities, making the mesh structure and its associated graph a central representation of
the physical system.

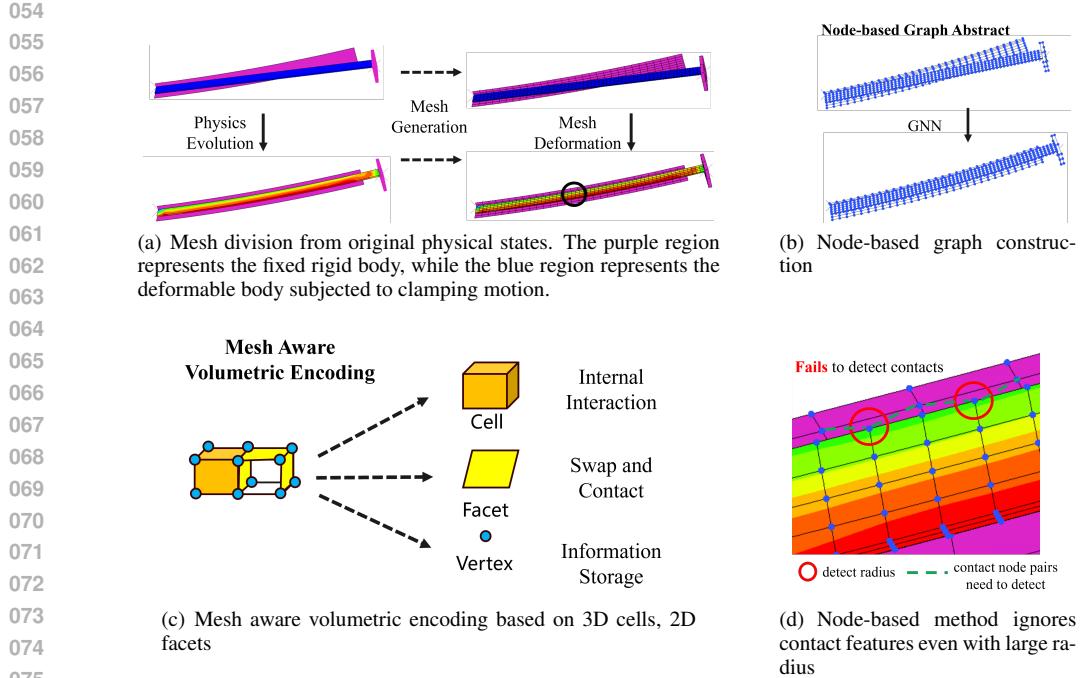


Figure 1: The physical state on the continuous material domain is discretized using structured meshes. Node-based methods construct point-edge graphs from the mesh and apply GNNs for computation. However, such abstraction may overlook contact interactions. A more effective approach should incorporate higher-dimensional geometric structures in the mesh, such as 3D cells and 2D facets, which retain accurate geometric information after discretization.

Although GNNs are effective, their accuracy deteriorates on sparse meshes, which are commonly used for computation efficiency concerns in practice (Allen et al., 2022). As illustrated in Figure 1(d), the distance between the points in GNN differs from the actual distance between surfaces, and this deviation worsens under coarser mesh configurations. As a result, contact information may be missing without an appropriate detection radius. Increasing the radius may help, but, with the cost of computations, there are still no guarantees of complete and accurate contact information. In addition, GNNs model internal propagation in approximating integral kernels (Li et al., 2020b; 2024), which is based on positional vectors and physical variables. However, with coarse meshes, nodes may not be sufficient to adequately sample neighborhood regions, hindering the accurate construction of characteristics of the localized physical field.

These limitations are mainly due to node-based modeling *by only using the vertices*. These methods represent meshes as graphs with edge features encoding distances, but existing approaches relying on topological representations often lose critical spatial features required for physical accuracy. Crucially, in addition to the vertices, the mesh contains a much more comprehensive set of high-dimensional geometric elements, that is, **2D facets** and **3D cells**, as illustrated in Figure 1(c), colored orange and yellow. *Our idea is that such high-dimensional elements like 2D facets and 3D cells elements could be incorporated to enable the model to better characterize geometric structures within 3D continuous space.* With this key incorporation, graph-based contact modeling can explicitly capture boundaries and contact as face-to-face geometries, making approaches more suitable for precise simulations. Additionally, even though in the case of coarse mesh discretization, the model might lead to inaccurate integral approximations based on discretely sampled vertices, volumetric cells could compensate and maintain stable computations by retaining geometric continuity.

To fully exploit high-dimensional geometric elements in mesh-based neural networks, we design a novel framework within an encoder–processor–decoder architecture that explicitly embeds cell and facet elements into the model, thereby enhancing performance under sparse mesh conditions. Technically, we propose MAVEN, a model based on Mesh-Aware Volumetric Encoding, in which we construct explicit nodes for each geometric element in the mesh, including vertices, facets, and cells.

108 During each processing step, the vertex information is encoded into higher-dimensional elements
 109 using learnable position-aware aggregators. The internal interactions and external loads are then
 110 handled at the cell level through facet nodes, allowing information to propagate through the mesh
 111 via its higher-dimensional structures (cells, facets). This cell-facet co-design enables comprehensive
 112 geometric modeling beyond node-based approaches. MAVEN achieves state-of-the-art performance
 113 in extensive evaluations. The main contributions of this paper can be summarized as follows.
 114

- 115 • We propose a paradigm that explicitly incorporates high-dimensional geometric elements
 116 into the simulation of 3D solid systems. This approach enables the model to capture precise
 117 geometric information and maintain stability under sparse discretization conditions.
- 118 • We design MAVEN, a model based on mesh-aware volumetric encoding that captures high-
 119 dimensional geometries by explicitly modeling both cell and facet elements. MAVEN fa-
 120 cilitates data transformation between elements through carefully designed transformation
 121 functions, and propagates information over a cell-facet graph using a modified two-stage
 122 message passing process.
- 123 • We compare MAVEN with state-of-the-art methods on public elastic deformation datasets
 124 and a metal bending problem, with elastic-plastic deformations and a coarse mesh. The
 125 experimental results demonstrate that MAVEN outperforms baselines in solid deformation
 126 with an acceptable computational efficiency.

127 2 RELATED WORK

128 2.1 LEARNING PHYSICAL SYSTEMS WITH GNNs

132 Recently, the application of Graph Neural Networks (GNNs) for simulating flexible dynamics has
 133 emerged as a promising research direction (Sanchez-Gonzalez et al., 2020; Han et al., 2022a; Shlomi
 134 et al., 2020; Gao et al., 2022). As a baseline and essential method, MGN (Pfaff et al., 2020) rep-
 135 resents meshes as graphs by treating vertices as nodes and using connectivity and proximity-based
 136 edges, within and across objects. It adopts an Encoder-Processor-Decoder architecture, encoding
 137 relative positions as implicit geometric features and learning dynamics via message passing. Later
 138 studies primarily aim to improve message passing through more expressive architectures (Dwivedi
 139 & Bresson, 2020; Shao et al., 2022; Han et al., 2022b), use hierarchical graphs to propagate infor-
 140 mation in various ranges (Cao et al., 2023; Fortunato et al., 2022; Grigorev et al., 2023), and adopt a
 141 hybrid design that integrates both approaches (Yu et al., 2023). Since these methods model physical
 142 features only on vertices, we refer to them as **node-based GNN** approaches.

143 However, graph-based models often neglect essential high-dimensional geometry, particularly in
 144 sparse meshes common in real-world scenarios. For **inter-object contact**, true interactions occur
 145 between surfaces, yet node-based GNNs approximate them via vertex distances, causing errors when
 146 contact regions extend beyond vertex discretization (Allen et al., 2022). For **intra-object dynamics**,
 147 some studies (Li et al., 2024; 2020b) interpret message passing as an approximation of a local kernel
 148 function that performs discrete integration over information from neighboring graph nodes. In the
 149 sparse condition, meshes provide too few nodes to capture local geometry. Consequently, critical
 150 quantities such as volume and surface area are poorly estimated, and these errors propagate through
 151 message passing, leading to significant deviations in predictions.

152 2.2 GEOMETRIC ELEMENTS IN PHYSICAL SIMULATION

153 These limitations stem from node-based modeling that relies solely on mesh vertices, overlooking
 154 the rich set of high-dimensional geometric elements inherently present in the mesh. Mesh repres-
 155 entations also include **3D cell** structures that accurately capture geometric partitioning in the continu-
 156 ous domain, and **2D facet** structures that define boundaries between regions and encode inter-object
 157 contact information. These elements contribute to more stable computations, particularly under
 158 sparse mesh conditions. For example, classical numerical methods (Reddy, 1993; Bardenhagen
 159 et al., 2004) describe physical quantities within volumetric elements by defining a family of shape
 160 functions (Zienkiewicz & Taylor, 2005) that interpolate physical qualities throughout the cell, and
 161 contact penalty terms are imposed on the integrals over the boundary facets. Based on this modeling
 162 approach, numerical methods can maintain controlled errors even under sparse mesh conditions.

162 A limited number of studies have focused on incorporating high-dimensional geometric information
 163 into DL-based physical simulations to improve computational accuracy. PHYMPGN (Zeng et al.,
 164 2025) follows discrete Laplace-Beltrami operators (Reuter et al., 2009), using face areas and cosine
 165 values of neighboring triangles for a single vertex as a broadcast operator between node pairs. This
 166 operator is limited to 2D settings and presents significant challenges when extending to 3D mesh
 167 domains. FIGNet (Allen et al., 2022; Lopez-Guevara et al., 2024) constructs face-to-face edges to
 168 capture contact relationships. Although effective for rigid bodies, these methods still face significant
 169 challenges in modeling internal propagation and dynamic deformations within 3D solids. To address
 170 this, we propose a novel DL-based architecture for 3D dynamic deformation simulation that intrin-
 171 sically integrates high-dimensional geometric structures with hierarchical feature representations.
 172

173 3 METHODOLOGY

175 3.1 PROBLEM DEFINITION

177 The evolution of a Lagrangian system is initiated by an initial material domain Ω_0 , together with a
 178 field function $\mathbf{U}(0, x)$ that defines the initial physical quantities, such as displacement, velocity, and
 179 pressure, at each material point x . At each time $t \in [0, T]$, we focus on the current physical state
 180 $\mathbf{U}(t, x)$ of every point $x \in \Omega_0$. To enable discrete computation, the initial material domain Ω_0 is
 181 partitioned, without overlap or omission, into a set of regular tetrahedra (or hexahedra) cells \mathcal{C} , which
 182 collectively form the **mesh**. The collection of vertices and surface facets from these regular regions
 183 defines the set of vertex \mathcal{V} and the set of facets \mathcal{F} of the mesh, respectively. Physical field information
 184 \mathbf{u}_i^t at time t is stored at each vertex $v_i \in \mathcal{V}$ to approximate the continuous domain, allowing the
 185 state of any point of material to be estimated by interpolation from the values at the vertices of the
 186 corresponding cell. The exact shape of each cell and facet at any given time is determined by the
 187 current state of deformation of its vertices. Excessive distortion or even fracture may occur as a
 188 result. In this work, we primarily consider scenarios in which the mesh does not undergo severe
 189 distortion, which aligns with the assumptions commonly made in industrial simulation settings.

190 The simulation trajectory of a physical system originates from an initial physical field \mathbf{u}^0 gov-
 191 erned by a discretization mesh $\mathcal{M} = \{\mathcal{C}, \mathcal{F}, \mathcal{V}\}$. The input of variable-time external forces
 192 $\{\mathbf{f}^0, \mathbf{f}^{\Delta t}, \dots, \mathbf{f}^{K\Delta t}\}$ acting on the vertices drives the progression of the dynamic state, gen-
 193 erating the sequence of physical evolution $\{\mathbf{u}^0, \mathbf{u}^{\Delta t}, \dots, \mathbf{u}^{K\Delta t}\}$ in discretization evolution time
 194 $0, \Delta t, \dots, K\Delta t = T$. The objective of the simulator is to predict next physical state $\mathbf{u}^{(k+1)\Delta t}$ from
 195 a history of previous states. In this paper, we consider $\{\mathbf{u}^0, \mathbf{u}^{(k-1)\Delta t}, \mathbf{u}^{k\Delta t}, \mathbf{f}^{k\Delta t}\}$ as input states.
 The rollout trajectory can be obtained by applying the simulator to the last prediction iteratively.

197 3.2 MODEL OVERVIEW

199 The overall architecture of MAVEN is illustrated in Figure 2, with the widely adopted "encoder-
 200 processor-decoder" framework (Battaglia et al., 2018; Pfaff et al., 2020). Unlike node-based models,
 201 MAVEN additionally models each element in the cell set $\{\mathcal{C}\}$ and facet set $\{\mathcal{F}\}$ as individual nodes
 202 participating in message passing, thus enhancing the ability to capture high-dimensional geometric
 203 information. First, MAVEN performs feature extraction for all node types. The cell and facet nodes
 204 are initialized with their geometric characteristics, such as 3D volume, surface area, as well as 2D
 205 area and perimeter, while the vertex nodes retain their physical quantities as input features. Sub-
 206 sequently, we employ a stack of processors to model physical interactions within and across solid
 207 objects. In each processor, the cell and facet nodes update their geometric representations from the
 208 nearby vertex nodes via a position-sensitive geometric aggregator. A modified two-stage message
 209 passing is then applied to propagate information during a cell-facet graph constructed by the relative
 210 geometric relationships between elements. Finally, a disaggregation operation distributes the aggre-
 211 gated features back to the vertex nodes, generating smooth intermediate results. The final processor
 212 output is subsequently mapped back to the original domain to produce the predicted results.

212 In MAVEN, the cell and facet focus primarily on diverse features. The facet is a pivotal hub for
 213 information exchange, where external forces, object contacts, and the propagation of internal phys-
 214 ical quantities converge. The diverse information is integrated and subsequently transmitted back to
 215 their respective cells. (Allen et al., 2022) shows that faces can capture contact information more ef-
 fectively, which node-based models might otherwise neglect under sparse conditions. The cells fully

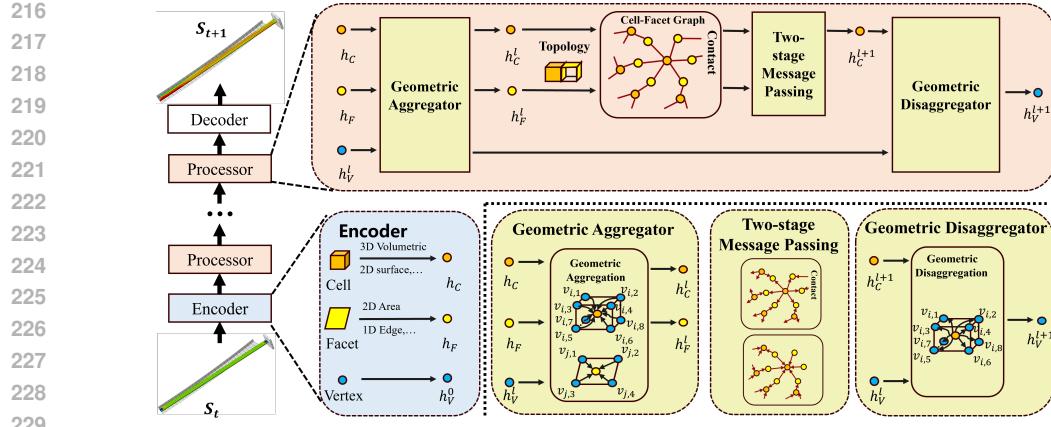


Figure 2: The overall structure of MAVEN. MAVEN follows an encoder–processor–decoder pipeline: it extracts geometric and physical features for vertices, cells, and facets, updates them through position-aware geometric aggregation and refined cell–facet message passing, and finally disaggregates the processed features back to vertices to produce smooth predictions.

characterize the geometric and volumetric properties of the 3D continuum domains. The adoption of volumetric features of adjacent cells as propagation coefficients for vertices significantly enhances the performance of 2D tasks (Zeng et al., 2025), motivating our design of the cell–facet propagation framework, ensuring comprehensive geometric information within the model.

Next, we elaborate on each key component of MAVEN. For convenience, in the following description we denote \mathcal{A} as the feature fusion operator that integrates multiple features into a single representation, implemented via multilayer perceptrons (MLP) in practice. Various \mathcal{A} are distinguished using the subscript and superscript notation.

3.3 ENCODER: GEOMETRY-INFORMED FEATURE ENCODING

The MAVEN encoder constructs feature representations for the cell, facet, and vertex nodes while also processing external force conditions and inter-facet contact relationships.

Vertex node encoder For a given vertex node $v_i \in \mathcal{V}$ and its associated physical field quantities $u_{v_i}^t$, we apply standard GNN practices to encode quantities into a high-dimensional latent space to derive the vertex node feature $\mathbf{h}_{v_i}^0$:

$$\mathbf{h}_{v_i}^0 = \mathcal{A}^{\mathcal{V}}(u_{v_i}^t) \quad (1)$$

Cell and facet node encoder Since cells and facets do not possess direct input features, we consider computing their representations from high-dimensional geometric properties. Inspired by the classical finite-volume method (Eymard et al., 2000), we posit that both the internal volume and surface area of a region are critical geometric descriptors. Accordingly, we use volume and total surface area as initial geometric features for each cell, while area and perimeter are used to characterize facet nodes. In addition, we incorporate the initial geometric attributes of each element to enhance the model’s awareness of high-dimensional geometric variations over time. Let $\Omega(c_i)$, $\Sigma(c_i)$ be the volume and surface area of cell c_i , and $\alpha(f_i)$, $\lambda(f_i)$ be the area and perimeter of facet f_i , MAVEN generates cell and facet features \mathbf{h}_{c_i} and \mathbf{h}_{f_i} as follows:

$$\mathbf{h}_{c_i} = \mathcal{A}^{\mathcal{C}}(\Omega(c_i^t), \Sigma(c_i^t), \Omega(c_i^0), \Sigma(c_i^0)), \quad \mathbf{h}_{f_i} = \mathcal{A}^{\mathcal{F}}(\alpha(f_i^t), \lambda(f_i^t), \alpha(f_i^0), \lambda(f_i^0)) \quad (2)$$

Here, all $\mathbf{h}_{v_i}^0$, \mathbf{h}_{c_i} , and \mathbf{h}_{f_i} are projected to the same latent dimension, which is set to 128 in practice.

Facet-to-facet features Instead of constructing edges between vertices, MAVEN establishes contact connections directly between the interacting facets. We improve on (Allen et al., 2022) by applying a simplified Bounding Volume Hierarchy algorithm (Clark, 1976) to detect all pairs of faces within a collision radius r . For two contacting facets f_s and f_r , the translation equivariant vectors of each face edge include: (1) the relative displacement between the center points $\mathbf{d}_{rs}^{\mathcal{F}} = \mathbf{p}_r - \mathbf{p}_s$ on the

270 two faces, (2) the vertices spanning vectors from one face to relative to the center point on another
 271 face $\mathbf{d}_{v_i}^{\mathcal{F}} = \mathbf{x}_{s_i} - \mathbf{p}_s$, $\mathbf{d}_{r_i}^{\mathcal{F}} = \mathbf{x}_{r_i} - \mathbf{p}_r$, and (3) the normal unit vectors of the face of the sender and
 272 receiver faces n_s, n_r . MAVEN generates face-to-face features $\mathbf{h}_{f_s \rightarrow f_r}$ as:

$$\mathbf{h}_{f_s \rightarrow f_r} = \mathcal{A}^{\mathcal{F} \leftrightarrow \mathcal{F}}([\mathbf{d}_{rs}^{\mathcal{F}}, [\mathbf{d}_{s_j}^{\mathcal{F}}]_{s_j \in f_s}, [\mathbf{d}_{r_j}^{\mathcal{F}}]_{r_j \in f_r}, \mathbf{n}_s, \mathbf{n}_r]) \quad (3)$$

273
 274 **External force features** External forces are defined as scripted motions for specific surface vertices,
 275 which means that their next-step positions \mathbf{x}^{t+1} are explicitly provided in the current step. We define
 276 the external force feature for each node $\mathbf{h}_{v_i}^S$ as its relative displacement to the next time step, and
 277 introduce a flag to indicate whether the node is scripted.
 278

$$\mathbf{h}_{v_i}^S = \begin{cases} [1, \mathbf{x}_{v_i}^{t+1} - \mathbf{x}_{v_i}^t], & \text{if } v_i \text{ is scripted} \\ 0, & \text{if } v_i \text{ is not scripted} \end{cases} \quad (4)$$

280 In practice, scripted motions are typically applied over entire surface regions rather than isolated
 281 vertices, making it essential to impose constraints at the surface level. Therefore, we define scripted
 282 features on each facet $\mathbf{h}_{f_i}^S$ by concatenating motion-related features of all its associated vertex nodes.
 283

$$\mathbf{h}_{f_i}^S = \mathcal{A}^S(\text{concat}_{v_j \in f_i}(\mathbf{h}_{v_i}^S)) \quad (5)$$

284 To ensure translational and permutation invariance, we sort the vertices of each facet by their dis-
 285 tances to the facet centroid, thereby enforcing a unique representation.
 286

290 3.4 PROCESSOR

291 All features extracted by the encoders are fed into a processor module composed of L stacked
 292 layers. In the l -th layer, the processor takes the vertex features of the previous layer $\mathbf{h}_{\mathcal{V}}^{l-1}$ and
 293 applies a **geometric aggregator** to incorporate the vertex information into the facet and cell nodes to
 294 generate geometric features $\mathbf{h}_{\mathcal{C}}^l$ and $\mathbf{h}_{\mathcal{F}}^l$. Two-stage message passing is used to propagate physical
 295 information across high-dimensional elements, where messages are first exchanged on facets and
 296 subsequently to cells. Finally, an inverse disaggregation decoder maps the updated cell-level features
 297 back to the vertex nodes for residual updates, producing spatially smooth features $\mathbf{h}_{\mathcal{V}}^{l+1}$ over domain.
 298

299 **Geometric Aggregator** Since vertex features are updated through the processor, it implies that the
 300 features of the cells and facets must also be updated. We update the features of each element by
 301 aggregating the features of all vertices that it contains. A straightforward approach is to concate-
 302 nate the initial features with those of all associated vertices, followed by an aggregation operation.
 303 However, since each element contains a relatively large number of vertices (e.g., eight vertices in
 304 a hexahedron), this approach results in significant computational overhead. Another approach is
 305 to average the features of all vertices, following the conventional GNN. However, this leads to se-
 306 vere homogenization of features between vertices and overlooks the relative geometric relationships
 307 between the nodes in their corresponding cells.

308 Inspired by the shape function (Reddy, 1993) in numerical solvers, which describes physical fields
 309 in a cell using local coordinates, MAVEN constructs aggregation coefficients from each vertex to the
 310 element based on the local coordinate system of the element. These coefficients are shared across
 311 all processor layers. Let $\vec{\mathbf{d}}_{c_i, v_j}$ denote the position vector from the center of the cell c_i to vertices
 312 $v_j \in c_i$, similar to $\vec{\mathbf{d}}_{f_k, v_l}$. Based on the local coordinate system, we employ an MLP to generate a
 313 centered normalized vertex aggregation weight $a_{c_i, v_j} \in \mathbb{R}$ for each cell.
 314

$$a_{c_i, v_0}, \dots, a_{c_i, v_{K-1}} = \text{MLP} \left(\text{concat}_{v \in \{v_0, \dots, v_{K-1}\}} (\vec{\mathbf{d}}_{c_i, v}) \right) \quad \{v_0, \dots, v_{K-1}\} \text{ represents } c_i \quad (6)$$

315
 316 Similarly, coefficients $a_{f_i, v_0}, \dots, a_{f_i, v_{M-1}}$ for each facet f_i can be derived, where K and M denote
 317 the number of associated vertices for each cell and facet, respectively. The coefficients obtained are
 318 then utilized to perform weighted aggregation for element feature update. We also sort the vertices
 319 within each cell to ensure permutation invariance.
 320

$$\mathbf{h}_{c_i}^l = \mathcal{A}_l^{\mathcal{V} \rightarrow \mathcal{C}}(\mathbf{h}_{c_i}, \sum_{v \in \{v_0, \dots, v_{K-1}\}} a_{c_i, v} \mathbf{h}_v^l), \quad \mathbf{h}_{f_i}^l = \mathcal{A}_l^{\mathcal{V} \rightarrow \mathcal{F}}(\mathbf{h}_{f_i}, \sum_{v \in \{v_0, \dots, v_{M-1}\}} a_{f_i, v} \mathbf{h}_v^l) \quad (7)$$

324 **Message passing on cell-facet graph** After extracting features, we construct a bipartite cell-facet
 325 graph $G = (V_G = \{\mathcal{C}, \mathcal{F}\}, E_G)$ to explicitly capture the topological and geometric relationships
 326 between volumetric elements and their boundary interfaces. The E_G contains all pairs (c_i, f_j) for
 327 $f_j \in c_i$. MAVEN performs two distinct message passing steps, each dedicated to the facet and cell
 328 nodes, respectively. In the first stage, information is aggregated through facet nodes, which serve as
 329 'edges', bridging not only adjacent cells but also mediating interactions between external forces or
 330 contacts and the internal dynamics of the object. Inter-object contact interactions are represented on
 331 the facet level, where information from all face-to-face edges is aggregated. To incorporate context
 332 from adjacent cells, we similarly employ a learnable aggregation scheme with a_{f_i, c_j} .

$$333 \quad \mathbf{h}_{f_i}^{\mathcal{F} \rightarrow \mathcal{F}, l} = \sum_{f_r = f_i} \mathcal{A}^{\mathcal{F} \rightarrow \mathcal{F}}(\mathbf{h}_{f_s \rightarrow f_r}, \mathbf{h}_{f_s}^l), \quad \mathbf{h}_{f_i}^{\rightarrow \mathcal{F}, l} = \mathcal{A}_l^{\rightarrow \mathcal{F}}(\mathbf{h}_{f_i}^S, \mathbf{h}_{f_i}^{\mathcal{F} \rightarrow \mathcal{F}, l}, \mathbf{h}_{f_i}^l, \sum_{(c_j, f_i) \in E_G} a_{f_i, c_j} \mathbf{h}_{c_j}^l) \quad (8)$$

336 In the second stage of message passing, each cell aggregates information from its associated facets.
 337 A similar strategy is adopted, employing symmetric aggregation coefficients $a_{c_i, f_j} = a_{f_i, c_j}$ to com-
 338 bine messages from multiple surfaces.

$$339 \quad \mathbf{h}_{c_i}^{\rightarrow \mathcal{C}, l} = \mathcal{A}_l^{\rightarrow \mathcal{C}}(h_{c_i}^l, \sum_{(c_i, f_j) \in E_G} a_{c_i, f_j} h_{f_j}^{\rightarrow \mathcal{F}, l}) \quad (9)$$

342 **Geometric Disaggregator** Finally, the high-dimensional geometric information encoded in each
 343 cell is retransmitted to its associated vertex nodes using the same symmetric aggregation coefficients
 344 $a_{v_i, c_j} = a_{c_j, v_i}$. A residual connection is applied to update the vertex features. This disaggregation
 345 at the vertex level facilitates a boundary-aware averaging of cell-level information, contributing to
 346 globally smooth predictions across the domain.

$$348 \quad \mathbf{h}_{v_i}^{\rightarrow \mathcal{V}, l} = \mathcal{A}_l^{\rightarrow \mathcal{V}}(h_{v_i}^l, \sum_{v_i \in c_j} a_{v_i, c_j} h_{c_j}^{\rightarrow \mathcal{C}, l}), \quad \mathbf{h}_{v_i}^{l+1} = \mathbf{h}_{v_i}^l + \mathbf{h}_{v_i}^{\rightarrow \mathcal{V}, l} + \text{FFN}(\mathbf{h}_{v_i}^l + \mathbf{h}_{v_i}^{\rightarrow \mathcal{V}, l}) \quad (10)$$

350 where $\text{FFN}(\cdot)$ represents the feed-forward network used in the Transformer (Vaswani et al., 2017).
 351 The next layer uses $\mathbf{h}_{\mathcal{V}}^{l+1}$ as the input vertex features.

353 3.5 DECODER AND UPDATER

355 Our model decodes the features of vertices $\mathbf{h}_{\mathcal{V}}^L$ using an MLP decoder to predict the velocity \hat{x}^{t+1}
 356 and the physical quantities \hat{c}^{t+1} for the next state. The next position \hat{x}^{t+1} is estimated by first-order
 357 integration $\hat{x}^{t+1} = \hat{x}^{t+1} + x^t$

358 **Training Loss** We use the one-step MSE loss as a training objective. Since other physical quantities
 359 may be included, MSE in flexible dynamics is calculated as follows:

$$361 \quad \mathcal{L} = \frac{1}{|\mathcal{V}|} \|x^{t+1} - \hat{x}^{t+1}\|^2 + \frac{1}{|\mathcal{V}|} \|c^{t+1} - \hat{c}^{t+1}\|^2 \quad (11)$$

364 3.6 DISCUSSION

366 Here, we briefly discuss how MAVEN differs from existing approaches.

367 **Compared to classical FEM methods**, MAVEN learns complex, nonlinear physical behaviors di-
 368 rectly from data, avoiding the hand-crafted constitutive models required in FEM. It generalizes
 369 across varying geometries and boundary conditions, enabling faster inference and improved scal-
 370 ability for large-scale simulations.

372 **Compared to basic DL methods such as MGN**, MAVEN propagates physical information through
 373 high-dimensional geometric elements, including cells and facets. This approach introduces a small
 374 amount of additional overhead. However, it enables MAVEN to accurately capture geometric infor-
 375 mation, which enhances the model's awareness of local neighborhoods and significantly improves
 376 its stability under sparse meshing conditions.

377 **Compared to hierarchical approaches**, MAVEN focuses mainly on accurately capturing local
 378 geometric details. Although hierarchical methods are effective at modeling long-range interactions,

they offer limited benefits for precise geometric representation. Moreover, MAVEN can be readily extended to graphs after pooling through an automatic mesh, which we leave for future work.

Compared to PhyMPGN, which depends on the cotangent Laplacian for local geometric, MAVEN avoids the inherent limitations of Laplacian-based operators in 3D, where no symmetric, local, and linearly accurate purely geometric Laplacian exists (Wardetzky et al., 2007). By explicitly modeling high-dimensional geometric elements and constructing operators through message passing, MAVEN provides a more flexible geometric framework, better suited for 3D Lagrangian formulations, and can naturally adapt to other physical settings such as 3D Eulerian formulations.

Compared to FIGNet, which is tailored for rigid-body contact and lacks mechanisms for modeling volumetric physical propagation, MAVEN explicitly represents cells and performs message passing over higher-dimensional geometric elements. This allows MAVEN to capture intra-object dynamics with much higher fidelity, especially under sparse mesh. In our experiments, we treat FIGNet as an ablated variant using only facet-level information, while MAVEN’s joint use of facets and cells yields better accuracy, underscoring the importance of modeling internal geometric structure.

4 EXPERIMENTS

Datasets To evaluate the efficiency of MAVEN, we test our model on datasets with different complexity in 3D solid simulation. The deforming plate (DP) (Pfaff et al., 2020) and cavity grasping (CG) (Linkerh  gner et al., 2023) are typical public datasets for the autoregressive elasticity task. The DP dataset contains relatively **dense tetrahedral meshes**, while the CG data set has **coarser meshes**. To further explore plastic scenarios, we establish a Metal Bending dataset (MBD) inspired by (Clausen et al., 2000) in real-world manufacturing, representing a class of solid deformation involving elastoplastic deformation, large displacements, and **very coarse** hexahedral meshes. In addition to target geometry, we also predict attendant physical quantities in experiments, including inner stress (Stress) and equivalent plastic strain (PEEQ). The rollout steps for all datasets are restricted to between 75 and 125 for a consistent comparison. We briefly present the motion process of the dataset in figure 3. See Appendix B for more details.

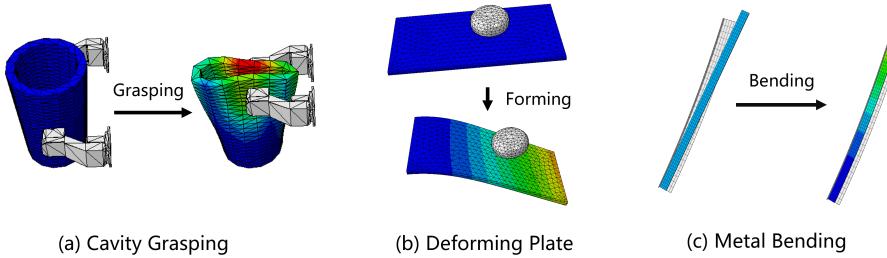


Figure 3: Visual description of the dataset.

Baselines We comprehensively compare MAVEN against baselines with node-based graph simulators. These include classical node-based models MGN (Pfaff et al., 2020) and Graph Transformer (GT) (Yun et al., 2019), as well as hierarchical models HCMT (Yu et al., 2023) and HOOD (Grigorev et al., 2023). To further distinguish the functions between the cell and the facet elements, we adapt FIGNet (Allen et al., 2022) to propagate internal physical quantities through the vertices. The specific settings of the models are provided in Appendix D.

4.1 EXPERIMENTAL RESULTS

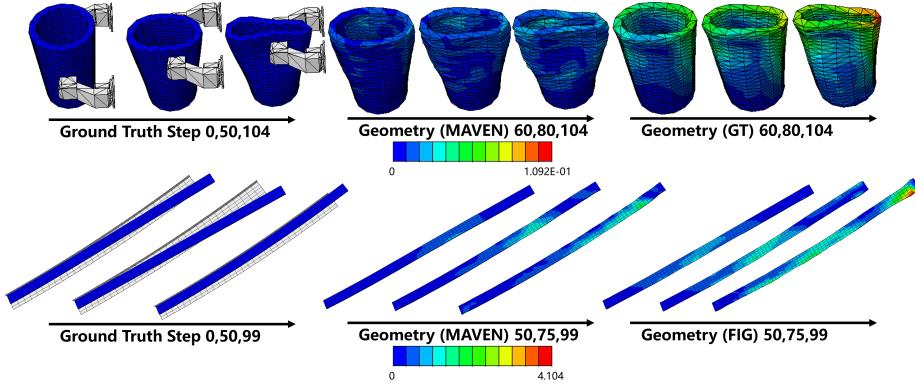
Rollout Results Table 1 shows the rollout results of all models, demonstrating that MAVEN consistently outperforms state-of-the-art methods to predict all physical quantities. From fine-grained to coarse meshes, MAVEN achieves incremental average improvements of 3.41%, 13.07%, and 18.13% across the three datasets, respectively. This indicates that explicitly capturing high-dimensional geometric features is beneficial for physical simulation and becomes even more critical under sparser mesh conditions. In the MBD dataset, geometry-based methods (FIGNet, MAVEN)

432
 433 Table 1: Rollout results($\times 10^3$) for MAVEN and other baselines, with 50-step rollouts and full-
 434 sequence rollouts. Our results are derived by averaging the root mean square error (RMSE) values
 435 computed across all intermediate steps and test datasets.

436 437 Model	438 Rollout	CG		DP		MBD		
		439 Pos	440 Pos	441 Stress	442 Pos	443 Stress	444 PEEQ	
439 MGN	50	6.38 \pm 0.39	14.01 \pm 0.087	24,495,364 \pm 793,284	686.19 \pm 70.51	12508.07 \pm 392.56	0.24 \pm 0.045	
	ALL	16.89 \pm 0.49	23.65 \pm 0.19	30,623,890 \pm 457,279	2012.16 \pm 299.60	9737.58 \pm 287.61	1.45 \pm 0.060	
441 GT	50	6.15 \pm 0.37	14.72 \pm 0.30	24,384,076 \pm 915,030	678.87 \pm 39.99	10368.29 \pm 1737.58	0.41 \pm 0.060	
	ALL	16.69 \pm 0.62	26.77 \pm 0.52	32,171,330 \pm 224,721	1406.61 \pm 72.84	14255.72 \pm 1203.51	2.07 \pm 0.083	
443 HCMT	50	6.14 \pm 0.29	14.46 \pm 0.47	22,335,358 \pm 663,289	851.86 \pm 52.23	17940.56 \pm 2662.00	0.45 \pm 0.010	
	ALL	16.87 \pm 0.24	24.94 \pm 0.76	30,317,188 \pm 457,279	2003.30 \pm 77.71	11539.27 \pm 834.80	1.30 \pm 0.16	
445 HOOD	50	6.96 \pm 0.083	14.27 \pm 0.32	23,474,653 \pm 259,738	623.57 \pm 23.47	11739.02 \pm 982.37	0.37 \pm 0.078	
	ALL	18.84 \pm 0.85	24.01 \pm 0.30	30,941,529 \pm 683,294	1762.41 \pm 35.92	8352.52 \pm 482.60	1.56 \pm 0.093	
447 FIGNet	50	6.26 \pm 0.086	14.74 \pm 0.18	23,926,010 \pm 93,458	515.17 \pm 67.13	5583.71 \pm 1060.30	0.22 \pm 0.12	
	ALL	17.59 \pm 0.51	26.51 \pm 0.23	31,491,198 \pm 237,542	1030.57 \pm 15.90	5402.31 \pm 805.40	1.09 \pm 0.52	
449 MAVEN	50	5.21 \pm 0.0056	13.78 \pm 0.17	21,657,348 \pm 170,228	276.73 \pm 44.46	4901.56 \pm 46.06	0.20 \pm 0.068	
	ALL	15.41 \pm 0.11	23.41 \pm 0.32	27,907,490 \pm 158,020	810.42 \pm 24.08	4776.72 \pm 71.20	1.01 \pm 0.024	
451 Improv.		13.07%	1.33%	5.49%	33.82%	11.90%	8.67%	

452
 453 significantly outperform node-based approaches. In addition, the cell-element-aware architecture
 454 enables MAVEN to better capture variations in three-dimensional volumes compared to FIGNet.

455
 456 **Visualization** Figure 4 presents the visualization results. Compared with other methods, MAVEN
 457 adopts a cell-based propagation approach, allowing physical contact information to be transmitted
 458 more effectively throughout the entire deformable body. As a result, MAVEN achieves lower errors
 459 even in regions that are far from the deformation part. At the same time, the use of facet-based
 460 contact detection allows MAVEN and FIGNet to maintain stable contact even under particularly
 461 coarse meshes. More visualization results and analyses can be found in Appendix F.



476 Figure 4: Visualization of error maps. The first and second rows respectively show sample visual-
 477 izations from cavity grasping and metal bending datasets.

4.2 ABLATION STUDY

481 We validate two key components in our model, specifically, the aggregators that explicitly compute
 482 geometric features, and the feature aggregation method based on local geometric coordinates. We
 483 conducted systematic ablation studies on the MBD and CG to evaluate the contributions of different
 484 component models, testing three models: 1) Our full model; 2) **Model A**, which replaces geometric-
 485 based aggregation coefficients with a degree-averaging approach, **is used to validate the effectiveness**
 486 **of our geometric aggregation strategy based on local coordinate systems**; 3) **Model B** that replaces

486 geometric input features of 3D cells and facets with zero padding, is used to assess the importance of
 487 explicitly computing geometric features; and 4) **Model C**, which removes explicit modeling of cell
 488 and facet nodes by averaging their precomputed geometric features onto neighboring vertex nodes
 489 and propagating them through standard message passing, is used to evaluate the role of explicitly
 490 representing higher-dimensional geometric elements.

Model	Rollout	CG		MBD	
		Pos	Pos	Stress	Stress
Ours	50	5.21	276.73	4901.56	4901.56
	ALL	15.41	810.42	4776.72	4776.72
Model A	50	6.26	402.79	5302.67	5302.67
	ALL	17.45	926.71	6683.94	6683.94
Model B	50	5.41	391.40	6839.26	6839.26
	ALL	15.933	1652.31	6680.39	6680.39
Model C	50	6.05	632.71	9802.55	9802.55
	ALL	17.08	1680.20	10375.86	10375.86

Table 2: Ablation results on CG and MBD.

506 Table 4.2 shows the averaged error results in the test datasets. We observe that replacing geometry-
 507 aware aggregation with simple averaging (Model A) performs significantly worse on the generally
 508 sparse CG dataset. This indicates that such sparsity requires capturing detailed intra-element geom-
 509 etry to surpass standard node-based methods. Similarly, omitting explicit geometric features (Model
 510 B) leads to severe degradation on the highly sparse MBD dataset. This suggests that in extremely
 511 sparse settings, traditional GNNs struggle to infer local geometric structure implicitly and therefore
 512 fail to accurately capture the underlying geometric topology. For Model C, which does not explicitly
 513 model high-dimensional geometric elements, its performance on both datasets is close to traditional
 514 node-based methods. This suggests that adding geometric features alone, without modeling higher-
 515 dimensional topology, is insufficient for nodes to fully capture surrounding geometric structure.

516 These ablation results collectively demonstrate the effectiveness and necessity of MAVEN’s de-
 517 sign choices, including the explicit modeling of higher-dimensional geometric elements, the explicit
 518 computation of geometric features, and the use of geometry-aware aggregation for message updates.

5 CONCLUSION AND FUTURE WORK

523 We propose MAVEN, an architecture that models mesh geometry with high-dimensional features
 524 and learnable aggregation to simulate 3D solid contact and deformation on coarse meshes, outper-
 525 forming baselines in physical propagation and contact representation. Remaining challenges in-
 526 clude **stabilizing meshes against excessive distortion** and **extending the approach to arbitrary-
 527 dimensional meshes**, which define directions for future work.

529 REFERENCES

531 G Abaqus. Abaqus 6.11. *Dassault Systemes Simulia Corporation, Providence, RI, USA*, 3:73, 2011.

532 Kelsey R Allen, Yulia Rubanova, Tatiana Lopez-Guevara, William Whitney, Alvaro Sanchez-
 533 Gonzalez, Peter Battaglia, and Tobias Pfaff. Learning rigid dynamics with face interaction graph
 534 networks. *arXiv preprint arXiv:2212.03574*, 2022.

536 Todd R Allen, Roger E Stoller, and Prof Shinsuke Yamanaka. Comprehensive nuclear materials.
 537 Technical report, Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States), 2012.

539 Scott G Bardenhagen, Edward M Kober, et al. The generalized interpolation material point method.
 540 *Computer Modeling in Engineering and Sciences*, 5(6):477–496, 2004.

540 Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi,
 541 Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al.
 542 Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*,
 543 2018.

544 Yadi Cao, Menglei Chai, Minchen Li, and Chenfanfu Jiang. Efficient learning of mesh-based phys-
 545 ical simulation with bi-stride multi-scale graph neural network. In *International conference on*
 546 *machine learning*, pp. 3541–3558. PMLR, 2023.

547 James H Clark. Hierarchical geometric models for visible surface algorithms. *Communications of*
 548 *the ACM*, 19(10):547–554, 1976.

549 Arild H Clausen, Odd S Hopperstad, and Magnus Langseth. Stretch bending of aluminium extru-
 550 sions for car bumpers. *Journal of Materials Processing Technology*, 102(1-3):241–248, 2000.

551 Gouri Dhatt, Emmanuel Lefrançois, and Gilbert Touzot. *Finite element method*. John Wiley & Sons,
 552 2012.

553 Vijay Prakash Dwivedi and Xavier Bresson. A generalization of transformer networks to graphs.
 554 *arXiv preprint arXiv:2012.09699*, 2020.

555 Robert Eymard, Thierry Gallouët, and Raphaèle Herbin. Finite volume methods. *Handbook of*
 556 *numerical analysis*, 7:713–1018, 2000.

557 Meire Fortunato, Tobias Pfaff, Peter Wirnsberger, Alexander Pritzel, and Peter Battaglia. Multiscale
 558 meshgraphnets. *arXiv preprint arXiv:2210.00612*, 2022.

559 Han Gao, Matthew J Zahr, and Jian-Xun Wang. Physics-informed graph neural galerkin networks:
 560 A unified framework for solving pde-governed forward and inverse problems. *Computer Methods*
 561 *in Applied Mechanics and Engineering*, 390:114502, 2022.

562 Rini Jasmine Gladstone, Helia Rahmani, Vishvas Suryakumar, Hadi Meidani, Marta D’Elia, and
 563 Ahmad Zareei. Mesh-based gnn surrogates for time-independent pdes. *Scientific reports*, 14(1):
 564 3394, 2024.

565 Artur Grigorev, Michael J Black, and Otmar Hilliges. Hood: Hierarchical graphs for generalized
 566 modelling of clothing dynamics. In *Proceedings of the IEEE/CVF Conference on Computer*
 567 *Vision and Pattern Recognition*, pp. 16965–16974, 2023.

568 Jiaqi Han, Wenbing Huang, Hengbo Ma, Jiachen Li, Josh Tenenbaum, and Chuang Gan. Learning
 569 physical dynamics with subequivariant graph neural networks. *Advances in Neural Information*
 570 *Processing Systems*, 35:26256–26268, 2022a.

571 Xu Han, Han Gao, Tobias Pfaff, Jian-Xun Wang, and Li-Ping Liu. Predicting physics in mesh-
 572 reduced space with temporal attention. *arXiv preprint arXiv:2201.09113*, 2022b.

573 John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger,
 574 Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate
 575 protein structure prediction with alphafold. *nature*, 596(7873):583–589, 2021.

576 Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*,
 577 2014.

578 Zhihao Li, Haoze Song, Di Xiao, Zhilu Lai, and Wei Wang. Harnessing scale and physics:
 579 A multi-graph neural operator framework for pdes on arbitrary geometries. *arXiv preprint*
 580 *arXiv:2411.15178*, 2024.

581 Zongyi Li, Nikola Kovachki, Kamayr Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, An-
 582 drew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential
 583 equations. *arXiv preprint arXiv:2010.08895*, 2020a.

584 Zongyi Li, Nikola Kovachki, Kamayr Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, An-
 585 drew Stuart, and Anima Anandkumar. Neural operator: Graph kernel network for partial differ-
 586 ential equations. *arXiv preprint arXiv:2003.03485*, 2020b.

594 Jonas Linkerh  gner, Niklas Freymuth, Paul Maria Scheikl, Franziska Mathis-Ullrich, and Ger-
 595 hard Neumann. Grounding graph network simulators using physical sensor observations. *arXiv*
 596 *preprint arXiv:2302.11864*, 2023.

597

598 Tatiana Lopez-Guevara, Yulia Rubanova, William F Whitney, Tobias Pfaff, Kimberly Stachenfeld,
 599 and Kelsey R Allen. Scaling face interaction graph networks to real world scenes. *arXiv preprint*
 600 *arXiv:2401.11985*, 2024.

601 Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter Battaglia. Learning mesh-based
 602 simulation with graph networks. In *International conference on learning representations*, 2020.

603

604 Rakesh Kumar Phanden, Priavrat Sharma, and Anubhav Dubey. A review on simulation in digital
 605 twin for aerospace, manufacturing and robotics. *Materials today: proceedings*, 38:174–178,
 606 2021.

607 Junuthula Narasimha Reddy. An introduction to the finite element method. *New York*, 27(14), 1993.

608

609 Martin Reuter, Silvia Biasotti, Daniela Giorgi, Giuseppe Patan  , and Michela Spagnuolo. Discrete
 610 laplace–beltrami operators for shape analysis and segmentation. *Computers & Graphics*, 33(3):
 611 381–390, 2009.

612 Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter
 613 Battaglia. Learning to simulate complex physics with graph networks. In *International conference*
 614 *on machine learning*, pp. 8459–8468. PMLR, 2020.

615

616 Yidi Shao, Chen Change Loy, and Bo Dai. Transformer with implicit edges for particle-based
 617 physics simulation. In *European conference on computer vision*, pp. 549–564. Springer, 2022.

618 Jonathan Shlomi, Peter Battaglia, and Jean-Roch Vlimant. Graph neural networks in particle
 619 physics. *Machine Learning: Science and Technology*, 2(2):021001, 2020.

620

621 Shilong Tao, Zhe Feng, Haonan Sun, Zhanxing Zhu, and Yunhuai Liu. Ladeep: A deep
 622 learning-based surrogate model for large deformation of elastic-plastic solids. *arXiv preprint*
 623 *arXiv:2506.06001*, 2025a.

624

625 Shilong Tao, Zhe Feng, Haonan Sun, Zhanxing Zhu, and Yunhuai Liu. Unisoma: A unified
 626 transformer-based solver for multi-solid systems. *arXiv preprint arXiv:2506.06021*, 2025b.

627

628 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
 629   ukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural informa-*
630 tion processing systems, 30, 2017.

631

632 Max Wardetzky, Saurabh Mathur, Felix K  lberer, and Eitan Grinspun. Discrete laplace operators: no
 633 free lunch. In *Symposium on Geometry processing*, volume 33, pp. 37. Aire-la-Ville, Switzerland,
 634 2007.

635

636 Youn-Yeol Yu, Jeongwhan Choi, Woojin Cho, Kookjin Lee, Nayong Kim, Kiseok Chang, Chang-
 637 Seung Woo, Ilho Kim, Seok-Woo Lee, Joon-Young Yang, et al. Learning flexible body collision
 638 dynamics with hierarchical contact mesh transformer. *arXiv preprint arXiv:2312.12467*, 2023.

639

640 Seongjun Yun, Minbyul Jeong, Raehyun Kim, Jaewoo Kang, and Hyunwoo J Kim. Graph trans-
 641 former networks. *Advances in neural information processing systems*, 32, 2019.

642

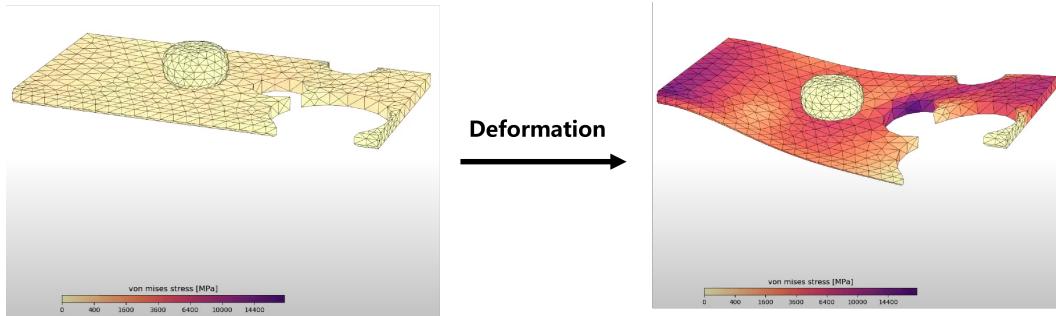
643 Bocheng Zeng, Qi Wang, Mengtao Yan, Yang Liu, Ruizhi Chengze, Yi Zhang, Hongsheng Liu,
 644 Zidong Wang, and Hao Sun. PhyMPGN: Physics-encoded message passing graph network for
 645 spatiotemporal PDE systems. In *The Thirteenth International Conference on Learning Represen-*
646 tations, 2025. URL <https://openreview.net/forum?id=fU8H41zk1m>.

647

Olgierd Cecil Zienkiewicz and Robert Leroy Taylor. *The finite element method for solid and struc-*
648 tural mechanics. Elsevier, 2005.

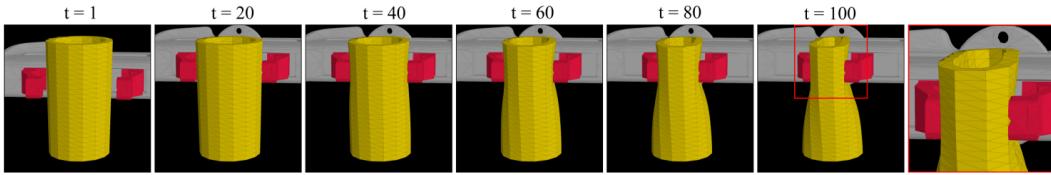
648 A ALGORITHMIC DETAILS OF MAVEN
649650 For completeness of the model, this chapter elaborates on the algorithmic details of MAVEN.
651652 **Generating face-to-face edges** To accelerate the construction of facet-facet edge sets, we employ
653 a coarse-grained bounding box approach where each facet \mathcal{F}_i is enclosed by an axis-aligned minimum
654 bounding rectangle (AAMBR) $[x_i^{\min}, x_i^{\max}] \times [y_i^{\min}, y_i^{\max}] \times [z_i^{\min}, z_i^{\max}]$. We perform
655 per-dimension intersection tests on the x, y, and z axes, then combine the results through logical
656 AND to determine the final intersection relationship, with the single-dimension intersection algo-
657 rithm detailed in Algorithm 1.
658659 **Algorithm 1:** Detecting 1-D intersection of segments

660 **Input:** segments set $\mathcal{S} = \{[x_i^{\min}, x_i^{\max}]\}$, a detect radius ε
 661 **Output:** detected intersection pairs $I = \{(i, j) | (i, j) \text{ is within radius } \varepsilon\}$
 662 initialization: $I = \emptyset, S' = \{[x_i^{\min} - \varepsilon, x_i^{\max} + \varepsilon]\}, A = \emptyset, B = \emptyset$
 663 generate events: events = $\{(x_i^{\min}, 0, i) \cup (x_i^{\max}, 2, i) \cup \{(x_i^{\min} - \varepsilon, 1, i) \cup (x_i^{\max} + \varepsilon, 3, i)\}$
 664 sort all event tuples in ascending order;
 665 **for** (x, type, id) in events **do**
 666 **if** $\text{type} == 0$ **then**
 667 $A = A \cup \{id\}$
 668 $I = I \cup \{(id, x) | x \in B\}$
 669 **end**
 670 **if** $\text{type} == 1$ **then**
 671 $B = B \cup \{id\}$
 672 $I = I \cup \{(x, id) | x \in A\}$
 673 **end**
 674 **if** $\text{type} == 2$ **then**
 675 $A = A \setminus \{id\}$
 676 **end**
 677 **if** $\text{type} == 3$ **then**
 678 $B = B \setminus \{id\}$
 679 **end**
 680 **end**
 681 **return** I ;

683 B DATASET
684697 Figure 5: Description of Deforming Plate from (Pfaff et al., 2020).
698699 **Deforming Plate Dataset** (Pfaff et al., 2020) This dataset contains 1,200 3D dynamic simulations of
700 a hyperelastic deformable plate pressed by a rigid solid (Fig.5). Each sample records the geometry
701 of the plate and the internal stress, with an average of 1,271 points per simulation. In our setup, we
unroll this dataset with a step size of 100 iterations to maintain consistent sequence lengths across all

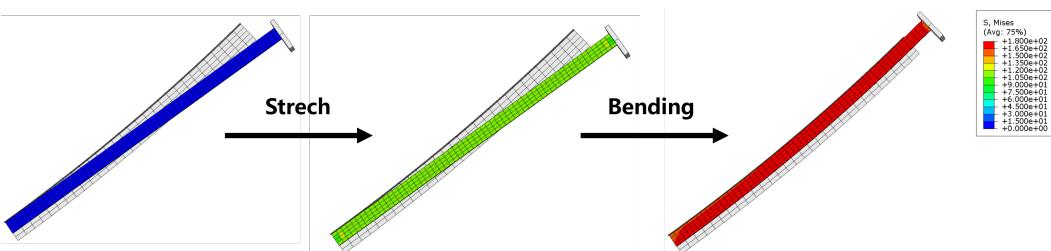
702 datasets. Following the protocol of the original article, we split the data into 1,000 training samples,
 703 100 validation samples, and 100 test samples.
 704

705 **Cavity Grasping** (Linkerh  gner et al., 2023) The dataset comprises a three-dimensional dynamic
 706 simulation of deformable cavities subjected to gripping by a rigid gripper (Fig. 6), containing a total
 707 of 840 samples. The gripper is modeled as two rigid bodies, corresponding to its two jaws, which
 708 undergo motion in opposing directions. The deformable objects are cone-shaped cavities generated
 709 with randomly assigned radii in the range [50, 87.5]. Their material properties are specified as
 710 elastic, with Poisson’s ratios drawn cyclically from the set -0.9, 0.0, 0.49. This dataset, primarily
 711 designed for autoregressive modeling tasks, provides temporal trajectories over 105 simulation steps.
 712 Each sample contains 1,386 points. Following the strategy of the original article, 600 samples are
 713 used for training, 120 for validation and 120 for testing.
 714

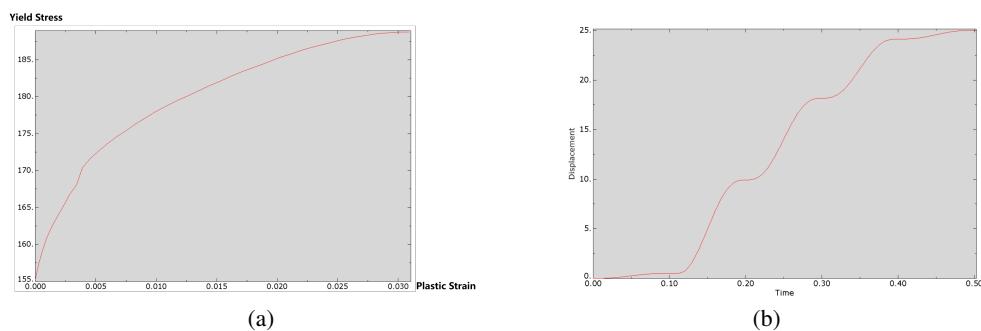


720 Figure 6: Description of Cavity Grasping from (Linkerh  gner et al., 2023)
 721

722 **Metal Bending Dataset** To rigorously validate the capability of MAVEN, we designed an industri-
 723 ally inspired test scenario featuring large deformations, coarse mesh discretization, and elastoplastic
 724 material behavior. This configuration mimics challenging real-world engineering applications such
 725 as metal-forming processes (Clausen et al., 2000), where computational methods must simultane-
 726 ously handle geometric nonlinearity, material nonlinearity, and under-resolved meshes while main-
 727 taining physical accuracy and numerical stability. The results are calculated by ABAQUS software
 728 (Abaqus, 2011).
 729



738 Figure 7: Description of Metal Bending.
 739



752 Figure 8: (a) True Stress-Strain Curve of Aluminum Profile. (b) Exemplar Motion Trajectory for
 753 Clamp Mechanism.
 754

755 As illustrated in Fig. 7, this scenario involves clamping a straight aluminum profile using a special-
 756 ized device, where the profile is first stretched beyond its elastic limit to induce plastic yielding, and

756 then progressively pressed against a curved steel die through controlled displacement; the resulting
 757 compressive contact forces generate permanent plastic deformation to achieve the prescribed target
 758 geometry. The straight metal component is modeled as a slender component of size $2 \times 12 \times 200 \text{ mm}^3$
 759 to accurately replicate field conditions. For discretization, we employ a $1 \times 3 \times 5 \text{ mm}^3$ mesh grid
 760 throughout the component. The material properties are configured as measured aluminum profile
 761 characteristics, with a Poisson's ratio of 0.37, Young's modulus of 69,000, and the true stress-strain
 762 curve depicted in Fig. 8(a). The 3D rigid die geometry is determined by its cross-sectional profile
 763 and a characteristic guiding curve. To ensure continuous contact between the aluminum workpiece
 764 and the die during bending operations, this guiding curve must maintain convexity and smoothness.
 765 We construct the curve by combining two circular arcs lying in the XY and XZ planes, respectively,
 766 each defining principal curvatures, and then synthesizing them into a composite spatial curve. In
 767 the ellipse $\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1 (a > b > 0)$ with a focal point at $(c, 0)$, we employ three distinct uniform
 768 distributions to control the morphology of the die, which are $\frac{c}{a} \sim \mathcal{U}[0.1, 0.3]$, $\frac{b}{a} \sim \mathcal{U}[0.1, 0.3]$
 769 and $a \sim \mathcal{U}[170, 190] (\text{unit:mm})$. The clamping tool's motion trajectory is generated by applying a
 770 classical evolute algorithm along the characteristic curve (see Fig. 8(b) for an illustrative example).
 771 The dataset contains an average of 1163 nodes per sample, with rollout lengths varying between 75
 772 and 125 timestep. We generated a total of 1000 trajectories, divided into 800 for training, 100 for
 773 validation, and 100 for testing.

774 C METRICS

775 In our autoregressive framework, consistent with the baseline methods under comparison, we em-
 776 ploy the Root Mean Square Error (RMSE) as the evaluation metric. Given the predicted physical
 777 quantity \hat{y}_i and ground truth value y_i of current N vertices, the RMSE is calculated as:

$$778 \quad RMSE = \sqrt{\frac{1}{n} \|\hat{y}_i - y_i\|^2} \quad (12)$$

779 For positions, we take the distance between two points as the sole comparative physical quantity. To
 780 evaluate the performance across multiple trajectories and frames, we report the final error metric as
 781 the average value computed over all frames for each individual trajectory:

$$782 \quad ERROR = \frac{\sum_{i=1}^M \sum_{j=1}^{T_i} RMSE_{i,j}}{\sum_{i=1}^M T_i} \quad (13)$$

783 The number of trajectories is denoted by M , where T_i represents the length (number of timesteps)
 784 of each individual trajectory.

785 D IMPLEMENTATION

786 **Model Hyperparameters** To ensure a fair comparison, all models were evaluated under similar
 787 parameter budgets and computational costs, with detailed configurations provided in Table 3. Due
 788 to the larger parameter count per block in MAVEN and FIGNet, these models employed fewer
 789 propagation layers. For HCMT and HOOD, while the optimal hierarchy layer count originally
 790 reported was 5 layers on the deforming plate dataset, we reduced it to 3 layers in our implementation
 791 because the hierarchical partitioning algorithm failed to correctly subdivide new meshes beyond the
 792 fifth layer. For all datasets, we adopted a uniform batch size of 8 in all experiments.

793 **Training Implementation** All models were uniformly trained in 1M steps. We used Adam opti-
 794 mizer (Kingma, 2014), with a learning rate decreasing from 10^{-4} to 10^{-5} . All experiments were
 795 conducted using a single RTX 3090 24GB GPU and repeated three times for the calculation of the
 796 standard deviation. All models were trained using Mean Squared Error (MSE) as the loss function,
 797 where each physical quantity was first normalized via Gaussian standardization and then summed
 798 directly to compute the final loss. This standardized approach ensured a fair and controlled compar-
 799 ison to objectively assess the relative performance of each method.

800 **Dataset Input and Detect Parameter** The inputs and outputs of each dataset, as well as the cor-
 801 responding node-based contact detection radius r_W and face-based contact detection radius r_F , are
 802 shown in Table 4. To ensure fairness, the node-based model and our model are able to detect nearly
 803 the same number of contact edges.

810
811 Table 3: Key hyperparameters and parameter numbers of models.
812

		Deforming Plate	Metal Bending	Cavity Grasping	
813 814 815 816 817 818 819 820 821 822 823 824 825 826 827 828 829 830 831 832 833 834 835 836	MGN	HIDDENS LAYERS	[128, 128] 15	[128, 128] 15	[128, 128] 15
		PARAMETERS(M)	3.85M	3.85M	3.85M
813 814 815 816 817 818 819 820 821 822 823 824 825 826 827 828 829 830 831 832 833 834 835 836	HCMT	HIDDENS LAYERS	[144, 144] 12+3	[144, 144] 6+9	[144, 144] 10+5
		PARAMETERS(M)	3.24M	3.24M	3.11M
813 814 815 816 817 818 819 820 821 822 823 824 825 826 827 828 829 830 831 832 833 834 835 836	FIGNet	HIDDENS LAYERS	[96, 96] 15	[96, 96] 15	[96, 96] 15
		PARAMETERS(M)	3.08M	3.48M	3.05M
813 814 815 816 817 818 819 820 821 822 823 824 825 826 827 828 829 830 831 832 833 834 835 836	HOOD	HIDDENS LAYERS	[128, 128] 12+3	[128, 128] 6+9	[128, 128] 10+5
		PARAMETERS(M)	3.56M	3.16M	3.47M
813 814 815 816 817 818 819 820 821 822 823 824 825 826 827 828 829 830 831 832 833 834 835 836	GT	HIDDENS LAYERS	[128, 128] 15	[128, 128] 15	[128, 128] 15
		PARAMETERS(M)	3.64M	3.61M	3.57M
813 814 815 816 817 818 819 820 821 822 823 824 825 826 827 828 829 830 831 832 833 834 835 836	MAVEN	HIDDENS LAYERS	[96, 96] 15	[96, 96] 15	[96, 96] 15
		PARAMETERS(M)	3.11M	3.15M	3.08M

834
835 Table 4: Model input, output and contact detection parameters for dataset. S denotes stress, and P
836 denotes PEEQ.

Dataset	Input	Output	r_W	r_F	noise
Deforming Plate	$type_i, x_i^t - x_i^{t-1}, S_i^t$	$x_i^{t+1} - x_i^t, S_i^{t+1} - S_i^t$	0.03	0.01	0.003
Cavity Grasping	$type_i, x_i^t - x_i^{t-1},$	$x_i^{t+1} - x_i^t$	0.1	0.05	0.01
Metal Bending	$type_i, x_i^t - x_i^{t-1}, S_i^t, P_i^t$	$x_i^{t+1} - x_i^t, S_i^{t+1} - S_i^t, P_i^{t+1} - P_i^t$	1	0.3	0.1

843
844 E MODEL EFFICIENCY845
846 We briefly discuss the computational efficiency of MAVEN. Table 5 shows the runtime performance
847 of various baselines on two datasets. The computational overhead of MAVEN primarily stems
848 from geometric feature computation and inter-node mapping operations, particularly on hexahedral
849 meshes. However, MAVEN still achieves an efficiency improvement of 2922.66% over the Abaqus
850 simulators (712.44ms per step) in the metal bending dataset.851
852 Table 5: The inference time per step(ms) for each model on three dataset

Dataset	MAVEN	MGN	FIGNet	HCMT	HOOD	GT
CG	25.48	18.37	24.15	55.05	51.37	22.02
DP	48.27	39.73	43.08	64.94	57.28	54.83
MBD	23.57	17.42	17.08	43.68	36.45	22.34

858
859 F VISUALIZATION860
861 Fig. 9 and 10 present the complete visualization results for each dataset. It is worth noting that in
862 the Metal Bending dataset, message-passing-based methods MGN and HOOD exhibit severe mesh
863 distortions, especially at contact regions. In contrast, graph-attention-based methods HCMT and GT
better preserve mesh shapes, though their overall structures become distorted. However, all of these

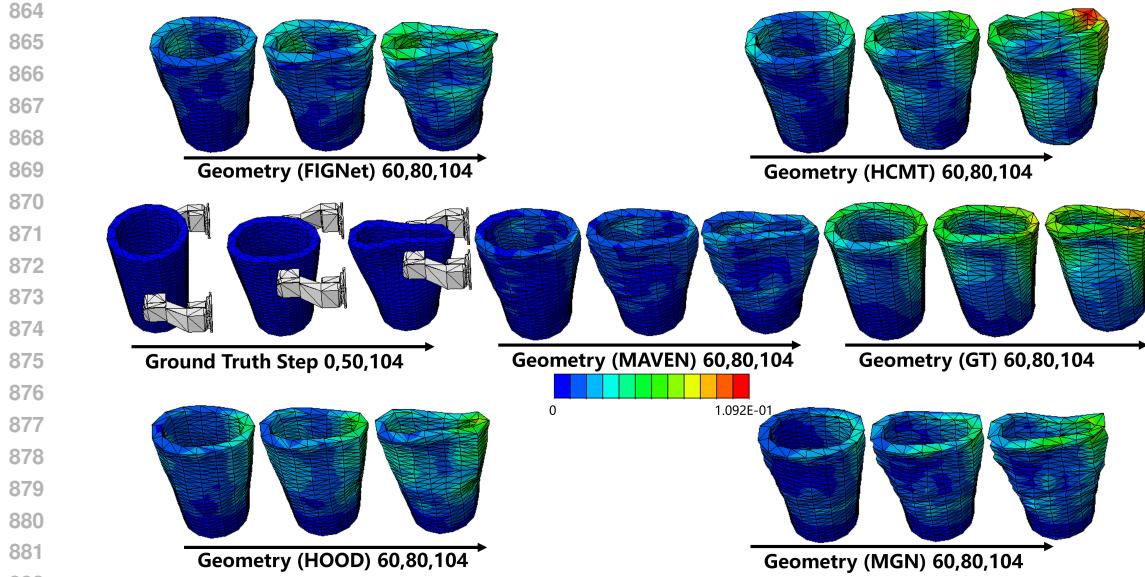


Figure 9: Visualization of distance error map on Cavity Grasping.

887 methods suffer from severe interpenetration issues, as further illustrated in the rollout animations
888 provided in the Supplementary Material. We believe that this phenomenon arises because individual
889 nodes cannot accurately identify their intended contact regions. Enlarging the detection radius intro-
890 duces many irrelevant points into the contact set, which, to some extent, is also exacerbated by the
891 large discrepancies in mesh lengths across the x, y, and z axes. More specifically, near the fixed end,
892 the mesh exhibits severe interpenetration. Closer to the moving end, the deformable body is heavily
893 constrained by the lower rigid body, preventing it from generating the correct motion and causing it
894 to stagnate. This indicates that facet-based contact detection is of critical importance when dealing
895 with sparse meshes.

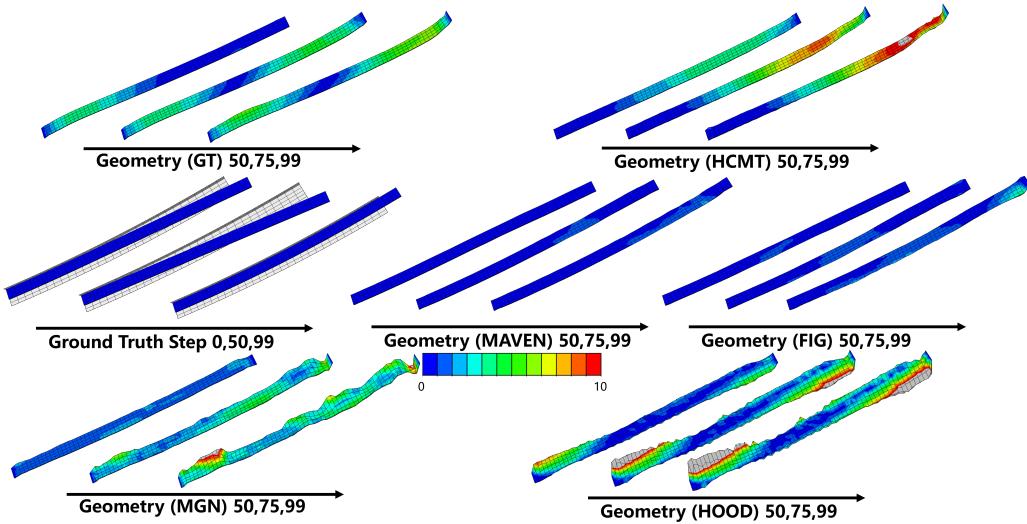


Figure 10: Visualization of distance error map on Metal Bending dataset. Gray indicates that the error at this location exceeds the given error upper bound.

917

918 **G LIMITATION AND FUTURE WORK**
919920 In this section, we briefly discuss the limitations of MAVEN and outline several promising directions
921 for future work.
922923 **The dependence on mesh quality.** Since MAVEN explicitly models cell-level information, the ge-
924 ometric quality of the initial mesh can substantially affect its performance. This behavior resembles
925 that of traditional numerical methods (e.g., FEM, FVM) rather than node-based GNN models. In
926 our experiments on the Cavity Grasping dataset, we compared meshes generated by a basic triangu-
927 lation procedure with those produced by a higher-quality meshing algorithm. The results show that
928 MAVEN’s performance degrades markedly when operating on low-quality meshes. However, most
929 existing deep learning datasets do not provide their original mesh discretizations, and meshes recon-
930 structed from point clouds typically exhibit very low geometric quality. Consequently, generating
931 more datasets with high-quality meshes, as well as reducing sensitivity to mesh quality, constitutes
932 an important direction for our future work.
933934 **Support long-range interaction.** While MAVEN is fundamentally designed as a local operator
935 that explicitly leverages high-dimensional geometric structures (facets and cells), it does not yet na-
936 tively capture long-range interactions. Existing graph pooling strategies (e.g., kNN clustering, BFS
937 grouping, global slicing) can, in principle, be directly applied to MAVEN cell–facet graph to expand
938 the receptive field. However, these methods often introduce significant computational overhead
939 and offer limited benefit in modeling a global high-dimensional geometric structure. Developing
940 a geometry-aware hierarchical extension that supports efficient long-range information propagation
941 therefore represents an important direction for future work.
942943 **Extend to boarder range of systems.** MAVEN is designed primarily for elastoplastic solid sim-
944 ulation on arbitrary 3D volumetric meshes, where explicit modeling of facets and cells provides
945 high geometric fidelity. Although MAVEN can be adapted to thin-shell or surface-based geometries
946 through appropriate redefinition of geometric features, and can extend to Eulerian formulations us-
947 ing fixed meshes and standard boundary-condition encodings, these adaptations require additional
948 geometry-aware considerations and suitable datasets. Extending MAVEN into a unified framework
949 capable of supporting surface-based systems, thin-shell structures, and Eulerian physical simulations
950 on sparse and irregular 3D meshes remains an important direction for future work.
951952 **H STATEMENT ON THE USE OF LARGE LANGUAGE MODELS**
953954 We declare that LLMs are mainly used in this paper to improve the clarity and fluency of the text, and
955 to a limited extent for code generation in technically mature modules, in order to reduce unnecessary
956 repetitive work.
957958
959
960
961
962
963
964
965
966
967
968
969
970
971