Proceedings of the IASS Annual Symposium 2025
"The living past as a source of innovation"
27 October– 31 October 2025, México City, México>
<Juan Gerardo Oliva, Juan Ignacio del Cueto, Elisa Drago (eds.)

Machine Learning for Force Geometry: A Homology Model for Stress-Informed Shells

Luis BORUNDA*

*School of Architecture, Virginia Tech
201 Cowgill Hall 1325 Perry St Blacksburg, VA 24061
lborunda@vt.edu

Abstract

As architecture faces rising demands for material efficiency, adaptability, and intelligent systems integration, new computational frameworks are needed to align performance with generative design. This paper presents a machine learning framework for predicting structurally meaningful lattice geometries in freeform architectural shells based on stress input. We introduce a Transformer-based model trained on scalar and directional stress fields to infer reinforcement patterns, producing polyhedral lattices aligned with principal stress trajectories. Operating without templates or rule-based encoding, the model generalizes across varied topologies and boundary conditions. By learning neighborhood relationships and stress flows, it internalizes structural logic beyond local cues, generating fabrication-ready outputs. Inference time is reduced from minutes to milliseconds, enabling a new scale of real-time structural reasoning. This approach bridges simulation and design, positioning AI as a tool for adaptive, performance-driven fabrication in architectural practice.

Keywords: AI-informed Structural Design; Hierarchical Structures; Robotic Fabrication; Biomimetic Structures; Latent structural mapping; Stress-graph isomorphism; Transformer-based structural inference.

1. Introduction

Advances in computational design and robotic manufacturing are reshaping how lightweight, structurally optimized shell structures are realized at an unprecedented scale and complexity. Emerging construction methods now support the material efficiency and geometric complexity necessary for functional grading and bespoke form [1]. This research builds on biomimetic strategies, particularly the adaptive remodeling of cancellous bone, which aligns its internal lattice with principal stress trajectories [2], to develop Artificial Intelligence (AI) -driven generative design methods for architectural shells. By integrating structural simulation, deep learning, and additive manufacturing (AM), we present a model that maps internal force distributions directly to fabrication-ready geometries.

Previous work introduced a stress-geometry isomorphism model that used rule-based segmentation of Finite Element Analysis (FEA) results to generate lattices [1]. However, that approach depended on static quantiles and thresholds, often resulting in discontinuities, high computational cost, and limited scalability. In contrast, we propose a Transformer-based inference pipeline that learns to translate isostatic stress point clouds into hierarchically organized shell structures.

Attention-based inference excels at capturing long-range spatial dependencies [3] and adapting to complex force distributions. Bypassing manual segmentation, the model performs stress inference at unprecedented resolution and scale. It maps local stress magnitudes and principal directions with their global structural logic into geometric configurations, capturing both spatial dependencies and fabrication-aware constraints. Trained on a multi-resolution dataset of simulated shell structures, the system generalizes structural logic across diverse boundary conditions and morphologies, expressing form as an emergent grammar of force.

We hypothesize a learned homological mapping from stress space to the space of lattice topologies, where this space satisfies geometric, structural, and printability constraints. The Transformer approximates this mapping through attention over tokens encoding spatial coordinates and scalar/vector stress features, enabling rapid generation of constructible geometries that reinforce force-intensive regions while conserving material in low-stress zones.

This approach reframes structural design as a predictive task: a continuous, data-driven translation from force fields to form. The resulting system is scalable, fabrication-aware, and structurally intelligent, advancing the integration of AI in architecture from generative assistance to real-time structural inference (**Figure 1**).

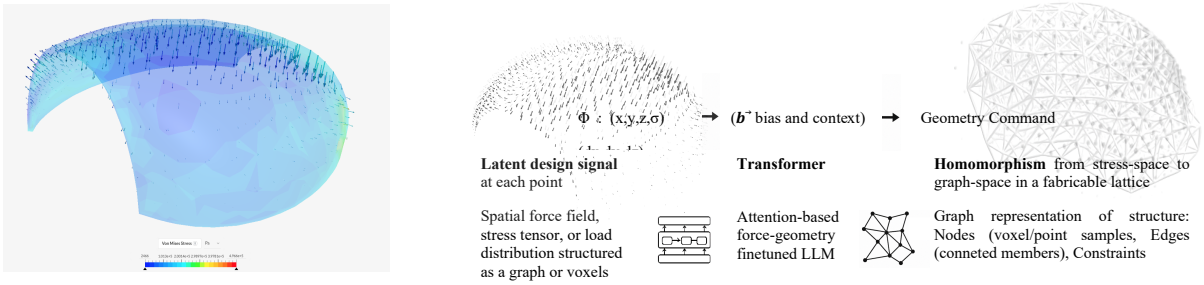

Figure 1: Force Geometry homology strategy. *(Left)* Stress, strain, and deformation trajectories computed via Finite Element Analysis. *(Right)* AI-driven form generation, where the spatial stress field $\boldsymbol{\varphi}(\boldsymbol{x},\boldsymbol{y},\boldsymbol{z},\boldsymbol{\sigma})$ is homologically embedded as attention-derived context and bias vectors ($\boldsymbol{b}^{\rightarrow}$) encoding structural correspondences. This enables a Transformer to infer geometry commands **G**, establishing a learned homomorphism from stress-space to graph-space. The model generates hierarchical, LLM-constructible, 3D-printable lattices that respond structurally to varying stress conditions.

2. Background: From Biomimetic Lattices to Learned Force-Geometry Mapping

2.1. Biomimetic Principles in Hierarchical Structural Design

Biological structures such as trabecular bone exemplify load-adaptive efficiency by distributing material along principal stress trajectories. This principle has informed architected lattices (**Figure 2**), where mechanical performance is governed by geometry rather than bulk material properties. Empirical studies report stiffness-to-weight and fatigue life gains exceeding 100× when internal members align with local stress vectors in biological systems [4], while architectural applications suggest similar benefits through stress-aligned additive manufacturing [5].

Bone leverages a hierarchical system [6]: a dense cortical shell encloses a compliant, porous lattice that remodels dynamically under load [7-8]. Biomimetic architectural analogues adopt similar strategies such as stress-influenced structures, polyhedral tilings, continuous fiber-reinforced skins, biologically inspired trabecular arrays, and stress-calibrated gridshells [9-12]. These systems embody a form of geometric intelligence, where performance emerges from the spatial logic of material arrangement guided by structural hierarchy, fabrication constraints, and design intent [13–15], yielding lightweight "cellular membranes" shaped by FEA data.

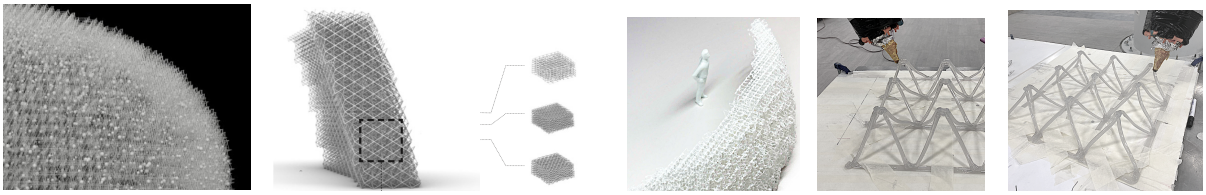

Figure 2: Hierarchical design, material distribution, and robotic additive manufacturing of a shell [16].

2.2. Limitations of Rule-Based Simulation Workflows

Parametric workflows for generating stress-informed structures typically rely on post-processing FEA results using rule-based segmentation, a labor and computationally intensive pipeline. Continuous stress fields are converted into discrete zones labeled high through low stress, based on threshold values. This process is inherently heuristic: thresholds set too high can under-support critical load paths; thresholds set too low negate material efficiency. Moreover, such segmentation does not guarantee smooth transitions in material distribution and may introduce discontinuities in the stress flow. Such forms are typically generated through scripting in computational design environments. Each variation in form or loading condition often requires reparametrization or manual rule adjustment, limiting scalability and introducing redundancy. Additionally, the pipeline incurs significant computational overhead. Stress data must be exported from simulation platforms, translated into geometry software, and re-integrated for fabrication. This multi-step process risks data loss and omission of key features such as principal stress directions, critical data for guiding infill architecture and orientation [1]. Crucially, these workflows lack any form of learning or adaptation. They do not accumulate knowledge across projects, nor can they generalize design strategies. Each application is governed by fixed, manually encoded rules. This impedes the grokking of nuanced or non-intuitive design responses, patterns that might emerge from a broader dataset of structural behavior. These limitations motivate a fundamental shift toward structural inference based on machine learning (ML).

2.3. From Force to Form: A Transformer-Based Alternative

Rather than scripting structural responses, a machine learning model can learn to recognize patterns in force distributions and infer material layouts directly from prior examples. Traditional methods such as topology optimization rely on iterative convergence to approximate efficient designs. In contrast, a trained neural network performs inference in a single forward pass eliminating the need for manual segmentation, static thresholds, or solver-based iterations. Transformer neural networks offer a compelling architecture for this task. Originally developed for sequence modeling in language and vision, Transformers have since been extended to domains such as 3D geometry processing, analytical solving [17], and point cloud classification and segmentation [18]. Their self-attention mechanism effectively captures spatial dependencies across irregular and unordered data, making them well-suited for processing non-uniform FEA outputs. By encoding each stress point as a token with spatial and scalar features, the model learns to infer latent structural logic, effectively "connecting the dots" between local magnitudes and global force trajectories. This enables reasoning over structural paths without predefined rules or case-specific tuning. Unlike rule-based workflows, which must be adapted to each geometry, this approach generalizes across varied forms and load conditions. Shifting from procedural computation to learned inference, the framework reframes structural design as a predictive task: mapping force fields directly to constructible geometry. The following sections detail the model architecture, dataset construction, and evaluation results.

3. Methodology

This investigation develops a machine learning–based structural inference pipeline that replaces manual post-processing of stress fields with data-driven geometric generation. We explore three core capabilities:

- whether Transformers with increasing contextual embeddings can reliably learn the mapping from stress fields to lattice geometries;
- how variations in lattice topology and refinement affect both structural and learning performance;
- how to embed the system within fabrication workflows to produce buildable, adaptive structures.

Together, these define an AI-driven homology model for structure-aware design, where geometry emerges from learned equilibrium patterns, rather than being engineered case-by-case.

We posit the existence of a homological mapping (ϕ: $\mathbb{R}^4 \rightarrow \mathcal{G}$), from the space of spatial-stress descriptors (points in $\mathbb{R}^3$ with associated stress magnitude σ) to a graph space $\mathcal{G}$ of structurally valid lattice topologies. Here, $\mathbb{R}^4$ symbolizes the latent equilibrium-informed state, and $\mathcal{G}$ denotes a set of graphs constrained by geometric and fabrication rules (such as continuity of print paths, non-isolated nodes, and printable member angles) yielding constructible, force-responsive arrays.

The Transformer model acts as a learned surrogate for this mapping. By attending to tokens encoding spatial coordinates, stress magnitude, principal directions, and anisotropy, it captures both local stress features and non-local structural dependencies. It bypasses rule-based segmentation and iterative simulation, offering a data-driven method for generating fabrication-ready, stress-informed geometries. The result is a scalable structural inference engine that generalizes across geometries and load conditions, grounded in learned morphogenetic logic.

The methodology consists of three integrated components:

Automated Dataset Generation: Training data is generated by pairing FEA-derived stress fields with algorithmically optimized lattice geometries. Shells are parametrically varied, and stress data is encoded as point clouds with scalar and directional features.

Transformer Training: Transformer models are fine-tuned to infer voxel- or graph-based reinforcement layouts from spatial–stress descriptors. Three progressively enriched versions predict homomorphism bias vectors encoding strut orientation, thickness, and connectivity, conditioned on local stress and global context.

Evaluation and Deployment: Outputs are evaluated for structural accuracy, printability, and performance. Comparative tests are conducted across lattice families (e.g., rhombic dodecahedral vs. hexahedral) and hierarchical refinements. Robotic toolpath integration ensures that predicted geometries are constructible.

3.1 Automated Dataset Generation from FEA and Parametric Design

We developed an automated pipeline for generating paired examples of stress fields and lattice geometries to train the model. Starting from two original shell forms, we applied parametric transformations to create 100 geometries, 50 predominantly linear, 50 non-linear, to ensure broad geometric variation across the training set. Each shell ranged from 5 to 30 meters in span and 1 to 10 meters in height, with curvature profiles and support conditions varied to mimic realistic architectural constraints. Each shell underwent 3–5 loading conditions, including self-weight, uniform pressure, point loads, and asymmetric cases. FEA was automated via SimScale's API-enabled cloud platform, using thermoplastics as base material (Young's modulus ≈ 80 MPa). Simulations were linear or nonlinear depending on expected deformation magnitudes. For each case, the von Mises stress (σ), principal directions ($\vec{v}$), and spatial coordinates (x, y, z) were exported in VTK format. These were normalized, filtered for relevance, and uniformly subsampled into point clouds with 1,000–5,000 points, preserving directional and scalar descriptors to serve as inputs to the model (**Figure 3**).

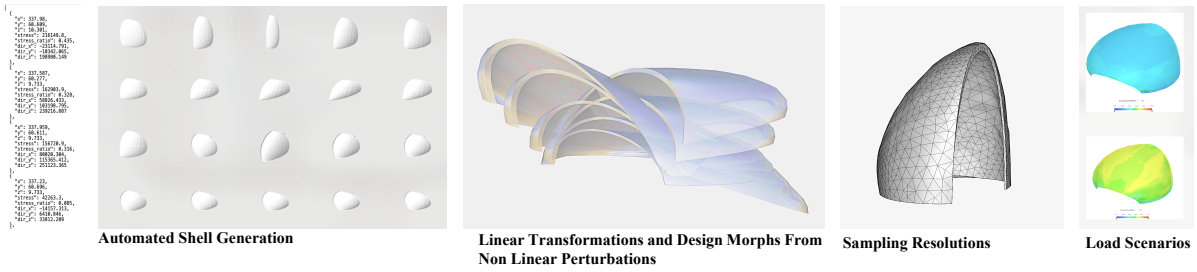

Figure 3: Automated Dataset Generation: Baseline shells are morphed via non-uniform perturbations, with stress fields extracted across varied resolutions and load cases.

3.2 Transformer-Based Model Training for Stress-Informed Design

3.2.1 Dataset Preparation: Stress Field Encoding and Label Generation

Each training sample corresponds to a thin-shell structure under gravitational loading with varying boundary conditions, simulated via finite element analysis (FEA). The resulting stress field is encoded as a point cloud, where each token includes spatial coordinates (x, y, z), von Mises stress (σ), and, when available, additional features such as the maximum principal stress direction vector ($\boldsymbol{v}_1$), anisotropy ratio, stress gradient (∇σ), and geometric context flags (**Table 1**). To regularize the input space, data is discretized into voxel-aligned tokens of the form:

Input token (dict): {x, y, z, σ, $\boldsymbol{v}_1$, ∇σ, anisotropy, flags}

Each voxel is paired with a fabrication-aware label reflecting local structural intent. For evaluation, labels fall into one of four discrete classes tailored to performance on the stress field: **L**: standard lattice strut in low-stress zones; **H**: medium-stress struts with increased deposition via slower print speed; **D**: duplicated members in high-stress regions for added strength; **B**: directional reinforcement aligned with the learned bias vector $\boldsymbol{v}_1$ (Table 1). This embeds both structural reasoning and fabrication constraints.

Model Output (parsed list): [Bias Vector, Bias Magnitude, Label]

Table 1: Dataset and Model Components for Force–Geometry Learning Pipeline. Each stress point is assigned a discrete print instruction derived from normalized stress magnitude and principal direction, enabling fabrication-aware structural reinforcement.

Component	**Description**	**Format / Role**
Input Features	Coordinates with scalar & vector stress features from FEA	(x, y, z, σ, d_x, d_γ, d_z)
Stress Ratio	Normalized von Mises stress with respect to shell maximum	$\sigma / \sigma_{max} \in [0, 1]$
Principal Direction	Unit vector of max principal stress, orienting reinforcement	(d_x, d_γ, d_z)
Label Format	Rule-based voxel classification by stress quantiles & directional bias	Biased Octet Truss, Rhombic Dodecahedon…[dir]
Output Token Format	Lexicographically ordered voxel grid coordinates with print instruction	(i, j, k): Lattice Instruction
Target Geometry	Tessellation with adaptive pattern & fabrication logic	Graph + voxel dict (GH script compatible)

3.2.2 Architecture and Learning Strategy

We adopt a Transformer-based encoder–decoder model with 32 layers and 32 attention heads, adapted from Mistral-7B [19]. The encoder processes the voxelized stress field as a token sequence, capturing both short-range and long-range mechanical dependencies. This architecture allows the model to learn that spatially distant high-stress regions may require coordinated reinforcement. The decoder autoregressively generates output tokens assigning a bias vector and lattice command to each voxel, conditioned on its stress context. We trained three progressively enriched model variants, each incorporating a deeper structural understanding:

- V1 (Principal-Only): Uses the normalized direction of maximum principal stress ($\boldsymbol{v}_1$); encodes directionality only, without magnitude or contextual cues.
- V2 (Magnitude-Weighted): Modulates $\boldsymbol{v}_1$ by local stress magnitude (σ_1), anisotropy, and directional coherence (cosine similarity with neighboring $\boldsymbol{v}_1$), reflecting relative reinforcement intensity.
- V3 (Full Descriptor): Extends V2 by incorporating stress gradients (∇σ) and structural flags (e.g., proximity to edges or curvature), enabling learned behaviors such as adaptive thickening, termination, or bifurcation.

This progression enhances the model's capacity to generalize across shell morphologies and loading asymmetries. All models (**Table 2**) were trained using a JSONL-based *"instruction–input–output"* format and optimized via LoRA (Low-Rank Adaptation) applied to attention and projection layers. Fine-tuning targeted 1–10% of total parameters using 16 low-rank adapters. For real-time fabrication, final models are compiled into GPT-Generated Unified Format (GGUF) for deployment within a Rhinoceros 3D Grasshopper robotic fabrication environment [20]. This workflow aligns with the practical demands of architectural computation, embedding machine-learned structural reasoning into design environments where geometry and fabrication are tightly coupled. The use of JSONL and LoRA ensures the training remains efficient and modular, while Grasshopper provides a direct interpretive layer translating model output into geometric actions, ready for robotic execution. Rather than retrofitting AI to architectural software, the workflow embeds intelligence into existing pipelines for immediate design-to-fabrication translation.

Table 2: Reinforcement Bias Vector Formulations for Structural Learning.

Version Name	**Label Formula**	**Features Used**	**Structural Logic Encoded**	**Learning Value**
(1) Principal direction only	normalize ($\mathbf{v}_1$)	$\mathbf{v}_1$ (max principal stress)	Reinforces dominant axis	Baseline vector regression
(2) Contextual magnitude-weighted	anisotropy $\times \sigma_1$ $\times$ coherence $\times$ normalize($\mathbf{v}_1$)	$\mathbf{v}_1$; σ_1; anisotropy $(\sigma_1 - \sigma_3)/(\sigma_1 + \varepsilon)$; coherence (cosine similarity between local $\mathbf{v}_1$ and neighbors' avg $\mathbf{v}_1$)	Identifying how sharply directed the stress in relative distribution, adds material where direction is strong & locally coherent	Direction and magnitude modulation, improves regressability
(3) Full reinforcement descriptor	Version (2) $+ \nabla\sigma$ + structural flags	All of (2) + $\nabla\sigma$: Stress gradient vector captures sharp stress changes across neighbors + structure labels capture context	Encodes how to split/thicken/terminate based on structure	Best generalization and structural reasoning; learns global structural logic

The approach builds on a lineage of stress-informed design: from Wolff's foundational insights on osseous adaptation [23], through the quantitative formalization of trajectorial logic in trabecular systems [22], to contemporary architectural interpretations in form-found shell structures [21]. Where prior frameworks rely on analytical models or rule-based translation from force to form, the Transformer offers a different framework: learning the mapping from structural behavior to geometry directly from data.

3.3 Ground Truth Inference and Geometric Construction Evaluation

To evaluate the model's generalization capacity, we implemented a dual benchmarking protocol: predictions were generated from both complete and intentionally occluded stress point clouds to simulate sensing noise and real-world partial data scenarios. Comparative analysis on a test set of shells (101–105) allowed us to distinguish between the model's ability to infer structural logic under ideal and incomplete conditions. Each input point cloud was paired with a corresponding "ground truth" lattice geometry, generated by bounding the shell domain and populating it with a base lattice topology (e.g., rhombic dodecahedron or octet truss). Struts were then adapted through a continuous feed rate scaling (0–1) in response to stress magnitude gradient: thickened and duplicated in high-stress zones, reduced or omitted in low-stress regions. A custom path-planning algorithm ensured printable geometry, respecting fabrication constraints such as angle thresholds and continuity.

This dataset generation procedure resulted in approximately 700,000 labeled force–geometry pairs, encoded in graph and voxel formats for downstream training and validation. To ensure learning compatibility, raw FEA stress fields were voxelized and regularized. In regions lacking direct stress data at voxel centroids, we applied a local k-nearest neighbors (kNN) interpolation to estimate von Mises stress and principal direction vectors. This ensured a dense, regularized tensor input compatible with Transformer architectures while preserving spatial mechanical features.

At inference time, the trained Transformer interprets the regularized voxelized stress field as a sequence of tokens encoding spatial coordinates, scalar stress magnitude, and directional descriptors. Each output token assigns a fabrication command to a voxel. Outputs are structured in a JSON-based voxel dictionary and imported into Grasshopper via a custom Python component. Each voxel is instantiated as a printable geometric unit with attributes derived from the predicted bias vector and its associated fabrication directive.

4.1 Results

We evaluated our approach across multiple dimensions: the accuracy of the Transformer's inferred embedded bias vectors, the geometries derived from stress fields, and attributes in terms of material usage and printability. All evaluations were conducted on test cases withheld during training (unseen shell geometries or new load configurations) to assess the generalization capabilities of the model. The Transformer consistently inferred coherent bias vectors that closely aligned with force trajectories derived from FEA when compared to Ground Truth data: when trained with enough data, even under incomplete inputs, the models inferred physically meaningful vectors, internalizing successfully general force logic (**Figure 4**).

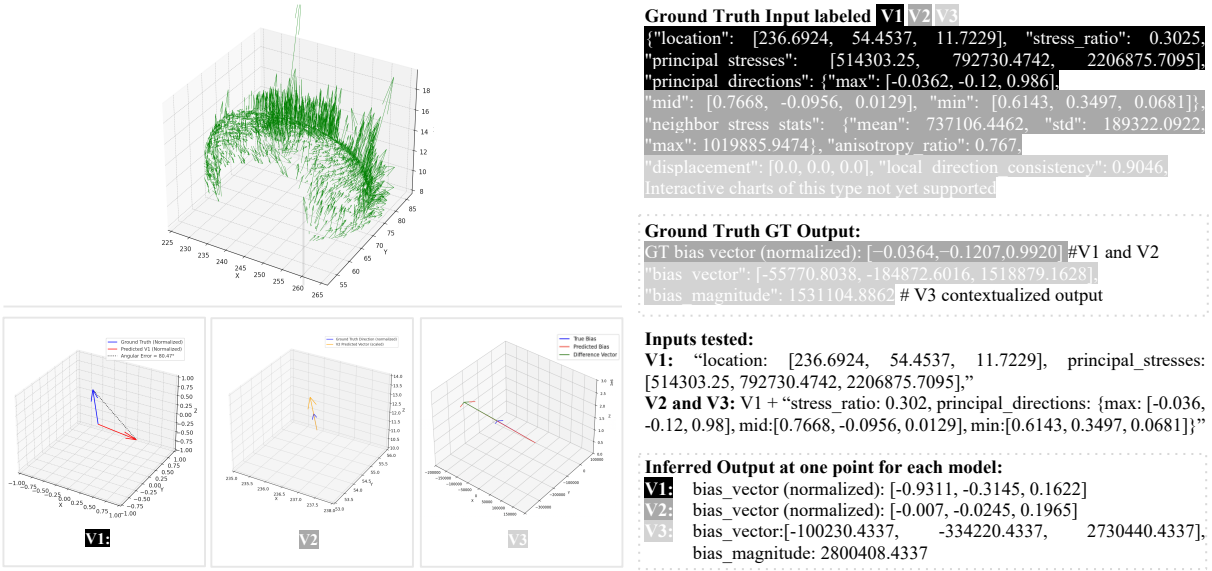

Figure 4: 3D Bias Vector Comparison at One Point per Machine Learning Model.

As illustrated in **Figure 5**, the model inferred a three-dimensional lattice structure, with stress-aligned reinforcement emerging along curved trajectories that bridge the two supported ends. These diagonals are densified where stress concentrations are greatest, producing a sweeping band of reinforced struts that fluidly spans the shell's surface. In contrast, low-stress zones are resolved with lighter, orthogonally oriented members, contributing to geometric porosity and material economy.

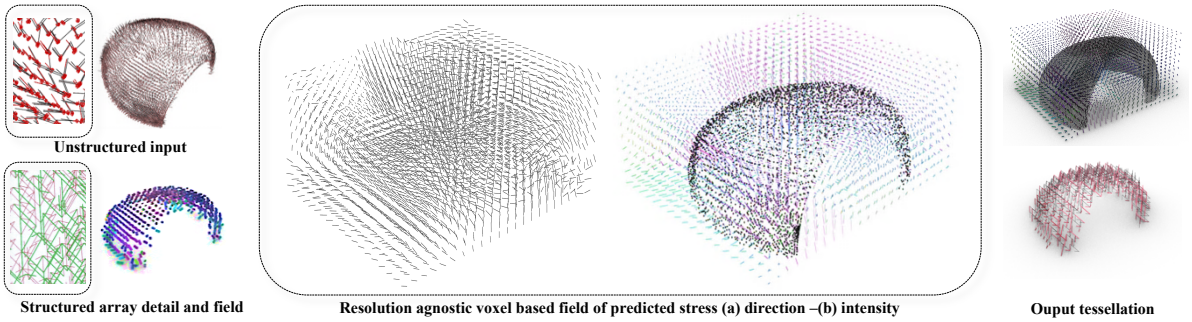

Figure 5: Generalization and robustness of structural inference on a freeform shell. This figure demonstrates the system's ability to interpret sparse, occluded data from an unseen geometry and generate a coherent, multi-resolution structural lattice. (Left) Conceptual visualizations of the input and the learned logic. The top panel shows the sparse, unstructured FEA-derived stress input simulating occluded data. The bottom panel shows the lattice connectivity logic that emerges from the inferred vector field. (Middle) The model-predicted directional field. This structured field of bias vectors and normalized magnitudes represents a fabrication-aware spatial stress interpretation learned from the training data. (Right) Fabrication-ordered polyhedral lattices at 150 cm (top) and 15 cm (bottom) resolution, demonstrating the system's resolution-aware structural adaptation.

The configuration emerged purely from learned structural patterns, with no explicit templates or rule-based encoding, underscoring the model's capacity to generalize and internalize structural logic. To assess inference quality, we compared the predicted lattice geometries against algorithmically generated reference tessellations at 15 cm and 150 cm voxel resolutions. Evaluation was based on voxel-level overlap with a stress-informed reference volume, as well as continuity and directional alignment. The resulting lattices displayed a high degree of coherence with the underlying stress directions and maintained topological continuity in over 90% of cases. Minor discrepancies typically involved redundant members or local misalignments, but the system consistently produced structurally plausible configurations aligned with the overall stress behavior. Token prediction accuracy reached 98%, and with minimal post-processing corrections, all output sequences were structurally valid.

To evaluate the impact of reinforcement bias formulation, we compared over three model variants (>2200 datapoints each) encoding progressively more structural information. **Figure 6** presents the cumulative distribution function (CDF) of angular error, showing that V3 surpasses both V1 and V2 in directional consistency, 94.3% of V3's predictions fall below 0.6°, compared to 72.6% for V1 and 51.9% for V2. The plateau in V2's curve reflects its susceptibility to directional ambiguity, often alternating between 0° and 180° under uncertain inputs. In contrast, V3's steep and smooth CDF curve confirms its ability to generalize structural logic and maintain precision across varied load scenarios.

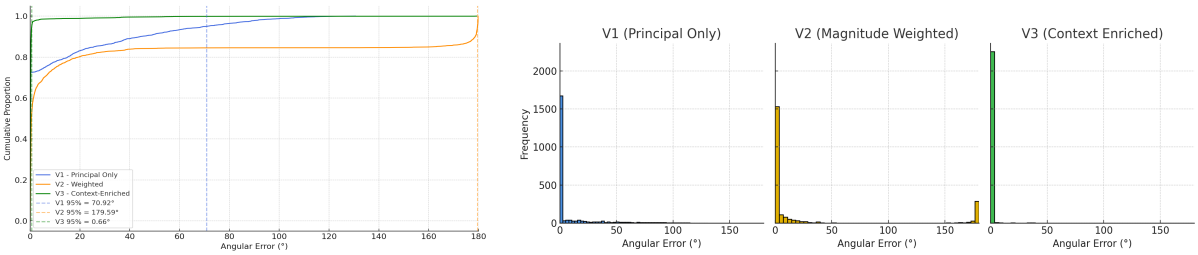

Figure 6. (Left) Cumulative Distribution of Angular Error Across Model Variants. The plot shows the proportion of predicted reinforcement vectors whose angular deviation from ground truth falls below each threshold. A steeper curve indicates greater directional accuracy and consistency. (Right) Angular Error Histogram.

Despite the large vector magnitudes, both V2 and V3 demonstrated strong directional accuracy. V3 achieved the best overall performance, with a median angular error of 0.14°, a mean of 0.83°, and a 95th percentile of just 0.66°, indicating high consistency across diverse conditions. V2 followed with a median of 0.51°, but a mean error of 30.4° and a 95th percentile approaching 180°, reflecting instability in ambiguous stress zones where directional flips occurred. V1, based solely on principal direction vectors, returned a median error of 0.0° due to frequent perfect alignment under ideal conditions, yet exhibited a mean error of 10.7° and a 95th percentile of 70.9°, revealing a skewed distribution and reduced reliability under complex or irregular stress fields.

To test robustness, we introduced occlusions into the input stress fields. V1 frequently defaulted to global directions under incomplete data. V2 maintained angular accuracy but under-predicted reinforcement in ambiguous zones. V3 preserved both orientation and magnitude, maintaining structural alignment across noise and topological variation. Its strong performance under occlusion and novel shell configurations confirms the generalizability of its learned structural logic. **Figure 7** illustrates the 3D spatial distribution of angular similarity (IoU > 0.5) across the shell for each model variant. **V1** shows scattered regions of lower precision, particularly at boundaries and curvature changes. **V2** exhibits stronger directional consistency across the surface. **V3** displays the most uniform accuracy, maintaining high IoU even in geometrically complex regions.

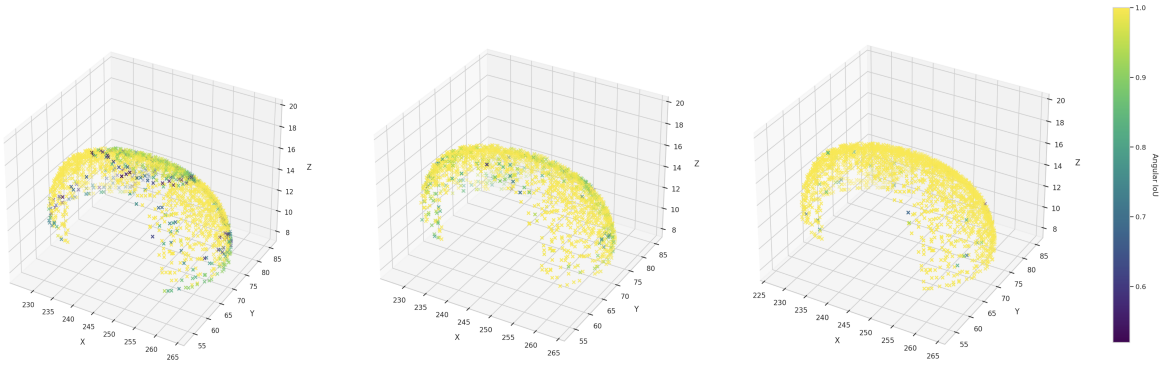

Figure 7. Directional Accuracy and Robustness Across Model Variants. 3D IoU Angular Similarity plots for V1 (Principal-Only), V2 (Magnitude-Weighted), and V3 (Context-Enriched) variants, filtered for IoU > 0.5. The colormap indicates directional alignment accuracy between predicted and ground-truth reinforcement vectors.

Table 3: Summary of Bias Vector Variant Comparison

Variant	Features Used	Median IoU	Angular Error (°)	Correlation R	Robustness	Notes
V1	$\boldsymbol{v}_1$ only	0.73	10.7° (mean)	–0.015	Degrades sharply with incomplete data	Simple alignment with stress axes
V2	$\boldsymbol{v}_1$ × σ_1 × Anisotropy × Coherence	0.81	30.4° (mean)	0.17	Stable, underestimates in noisy regions	Learns spatial–magnitude coupling
V3	V2 + $\nabla\sigma$ + Context Flags	0.84	0.83° (mean)	0.80	Strongest recovery and adaptation	Context-aware, predicts hierarchy, avoids errors

5. Discussion

Table 3 highlights how progressively enriched structural features enhance model performance. V1, based solely on the direction of maximum principal stress ($\boldsymbol{v}_1$), exhibited no correlation with stress magnitude (R = –0.015). While it performed adequately in ideal conditions, its uniform reinforcement strategy failed under ambiguity. V2 introduced anisotropy, directional coherence, and stress magnitude as weighting factors, improving directional accuracy (0.51° median error) and modestly raising correlation (R = 0.17). However, it continued to falter in regions with complex or conflicting stress cues. V3 integrated additional features (stress gradients, curvature, and edge proximity) yielding substantial gains: 0.14° median angular error, sub-0.66° at the 95th percentile, and a strong correlation with magnitude (R = 0.80), indicating reliable generalization across diverse structural scenarios.

These results affirm a central premise: structural performance emerges from context, not isolated local values. Attention-based sequence prediction provides a robust framework for encoding this context. Unlike heuristic or uniformly weighted models, the Transformer architecture dynamically prioritizes salient features, enabling it to capture spatial dependencies and adapt reinforcement across topological transitions such as branching, tapering, or thickening. This aligns with recent work demonstrating that supernode representations can retain sufficient structural information from large-scale FEA datasets to support Transformer-based inference in biomechanical systems [24]. These representations are resolution-agnostic, capable of adapting to both coarse and fine discretizations through upsampling or downsampling while preserving essential structural features. Building on this insight, our approach proposes a parallel opportunity: leveraging compact, graph-based representations within architectural applications to enable scalable, context-rich learned inference.

5.1. Conceptual Contribution: From Replication to Generalizable Structural Inference

Though trained on heuristic-generated labels, the model does not replicate static threshold logic. Instead, it learns to infer reinforcement layouts by reading contextual patterns in stress magnitude, directionality, and spatial flow. This shifts the task from segmentation to structural inference: geometry is not imposed but predicted as a learned response to force conditions. By internalizing equilibrium-like logic, the model moves beyond function matching toward interpretive design reasoning, representing a conceptual shift in how structure can be computationally encoded.

Notably, V3 exhibits robustness under partial or occluded input, maintaining alignment even when full boundary conditions are not available. In contrast, V1 and V2 often fail to generalize under incomplete data, defaulting to erroneous patterns. This suggests that V3 does not memorize local cues, instead it learns spatial relationships essential for scalable, context-aware design systems. Its performance across metrics (angular error, scalar correlation, IoU), and input robustness, positions it as a viable alternative to segmentation-based workflows. The model's ability to generalize to unseen geometries and maintain continuity from partial inputs validates that it has learned to infer underlying structural principles, moving beyond the static, rule-based logic of its heuristic training data. By offering a direct, fabrication-ready translation from structural fields to reinforcement geometry, it bypasses traditional post-processing pipelines. A mean angular deviation under 1° indicates high-fidelity alignment suitable for additive manufacturing, where tolerances up to 5° are typically acceptable without structural compromise.

Finally, these model variants serve as epistemic probes into what kinds of structural knowledge can be encoded by neural inference. V1 delineates the limitations of direction-only models. V2 begins to encode scalar expressiveness but lacks global understanding. V3, by embedding spatial context and stress transitions, demonstrates that higher-order structural logic can be learned. Rather than reacting to isolated features, it recognizes systemic patterns bridging the gap between physical behavior and generative geometry.

6. Conclusion

We present a machine learning–based force-geometry homology model that learns beyond local feature interpolation by internalizing the logic of structural equilibrium. Unlike traditional topology optimization, which relies on iterative FEM solvers, our method enables near-instantaneous inference once trained. The workflow of (1) Analysis, (2) Transformer Inference, and (3) FEA-based Validation preserves engineering rigor while drastically accelerating design feedback, ensuring the structural feasibility of every generated output.

By delegating geometric generation to a model trained on structural examples, this system eliminates the need for manually scripted reinforcement rules, streamlining the bridge between simulation and fabrication. The model exhibits a form of design intelligence akin to metaheuristic reasoning (i.e., morphogenetic algorithms), instead of through iterative search, by learning structural principles from data. This allows it to interpret stress trajectories as continuous spatial guides for material distribution, generating biomimetic, hierarchical lattices responsive to both directional flow and scalar intensity. In tests with previously unseen shell typologies and supports, the model adapted its output accordingly, thickening material near anchors, tapering across spans, demonstrating an internalized understanding of structural behavior.

Nonetheless, limitations persist. Accuracy depends on the diversity of training data, and unrepresented load conditions may yield unrealistic geometries. Accuracy depends on the diversity of the training data; while our parametric approach generated a broad set of geometries, future work should. The model currently lacks explicit enforcement of constraints like buckling resistance or static equilibrium, though our FEA post-validation mitigates this risk. Future work could embed physics-informed modules or hybrid inference-solvers to increase robustness. Critically, this approach enables new design workflows. Fast, differentiable inference allows for coupling with multi-objective optimization, inverse design, or real-time form-finding. Native compatibility with environments like Rhino/Grasshopper supports direct deployment in design practice.

Beyond architectural shells, the method shows potential for construction-scale additive manufacturing. A Transformer trained on stress–geometry pairs can inform robotic printing of functionally graded lattice structures that minimize material while maximizing strength. As open-source deployment tools evolve, such models will operate offline or onsite, enabling adaptive robotic fabrication. In conclusion, this research advances a learning-based model of structural inference. By training a Transformer to translate force data into geometry, we replace rule-based segmentation with continuous, context-aware inference. The system enables materially efficient, structurally expressive, and fabrication-ready design proposing a shift toward architectural systems conceived not just by force, but through it.

Acknowledgements

The authors acknowledge support from Virginia Tech and the College of Architecture, Arts, and Design. This research also benefited from academic API access provided by SimScale [25].

References

[1] L. Borunda and J. Anaya, “Hierarchical structures, computational design, and digital 3D printing,” *J. Int. Assoc. Shell Spatial Struct.*, vol. 64, no. 1, pp. 5–18, 2022. DOI: 10.20898/j.iass.2022.015.

[2] T. M. Keaveny, E. F. Morgan, G. L. Niebur, and O. C. Yeh, “Biomechanics of trabecular bone,” *Annu. Rev. Biomed. Eng.*, vol. 3, no. 1, pp. 307–333, 2001. DOI: 10.1146/annurev.bioeng.3.1.307.

[3] Q. Liu, V. Zhong, H. Meidani, D. Abueidda, S. Koric, and P. Geubelle, “Geometry-Informed Neural Operator Transformer,” *arXiv Prepr. arXiv:2305.15512*, May 2023.

[4] A. M. Torres *et al.*, “Bone-inspired microarchitectures achieve enhanced fatigue life,” *Proc. Natl. Acad. Sci. USA*, vol. 116, no. 49, pp. 24457–24462, 2019. DOI: 10.1073/pnas.1905814116.

[5] C. T. Tam and C. T. Mueller, "Additive manufacturing along principal stress lines," *3D Print. Addit. Manuf.*, vol. 4, no. 2, pp. 63–72, 2017. DOI: 10.1089/3dp.2017.0001.

[6] P. Fratzl and R. Weinkamer, "Nature's hierarchical materials," *Prog. Mater. Sci.*, vol. 52, no. 8, pp. 1263–1334, 2007. DOI: 10.1016/j.pmatsci.2007.06.001.

[7] E. R. Dumont, "Bone density and the lightweight skeletons of birds," *Proc. R. Soc. B Biol. Sci.*, vol. 277, no. 1691, pp. 2193–2198, 2010. DOI: 10.1098/rspb.2010.0117.

[8] E. Novitskaya, C. J. Ruestes, M. M. Porter, and M. A. Lubarda, "Reinforcements in avian wing bones: Experiments, analysis, and modeling," *Acta Biomater.*, vol. 60, pp. 267–280, 2017. DOI: 10.1016/j.actbio.2017.07.038.

[9] J. F. V. Vincent, O. A. Bogatyreva, N. R. Bogatyrev, A. Bowyer, and A.-K. Pahl, "Biomimetics: its practice and theory," *J. R. Soc. Interface*, vol. 3, no. 9, pp. 471–482, Sept. 2006.

[10] S. Bodea et al., "Buga fibre pavilion: Towards robotically-fabricated composite building structures," in *Fabricate 2020: Making Resilient Architecture*, J. Burry, J. Sabin, and B. Sheil, M. Skavara, Eds. London, U.K.: UCL Press, 2020, pp. 70–77.R.

[11] Naboni, A. Kunic, and L. Breseghello, "Computational design, engineering and manufacturing of a material-efficient 3D printed lattice structure," *Int. J. Archit. Comput.*, vol. 18, no. 3, pp. 251–264, 2020. DOI: 10.1177/1478077120945981.

[12] G. Pasquarelli *et al.*, "Additive manufacturing revolutionizes lightweight gridshells," *Struct. Eng. Int.*, vol. 30, no. 1, pp. 28–35, 2020. DOI: 10.1080/10168664.2019.1691661.

[13] M. Pelanconi and A. Ortona, "Review on the design approaches of cellular architectures produced by additive manufacturing," in *Ind. Addit. Manuf.: Proc. AMPA2020*, M. Meboldt and C. Klahn, Eds. Cham: Springer, 2020, pp. 52–64.

[14] M. Ashby, "The properties of foams and lattices," *Philos. Trans. R. Soc. A*, vol. 364, no. 1838, pp. 15–30, 2006. DOI: 10.1098/rsta.2005.1678.

[15] A. Edmondson, "Fuller's tensegrity and octet truss concepts," *Synergetics J.*, 1987.

[16] L. Borunda, "Impresión digital 3D. Diseño y fabricación digital de superficies continuas," Ph.D. dissertation, Univ. Politécnica de Madrid, 2021.

[17] Z. Li et al., "Geometry-informed neural operator for large-scale 3D PDEs," in *Adv. Neural Inf. Process. Syst.*, vol. 36, 2023.

[18] X. Yu *et al.*, "PoinTr: Diverse point cloud completion with geometry-aware transformers," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, 2021, pp. 12478–12487.

[19] Ollama, "Mistral," Ollama, 2024. [Online]. Available: https://ollama.com/library/mistral.

[20] M. Ladrón de Guevara, L. R. Borunda, and R. Krishnamurti, "Multi-resolution in architecture as a design driver for additive manufacturing applications," *Int. J. Archit. Comput.*, vol. 18, no. 3, pp. 244–261, 2020. DOI: 10.1177/1478077120924802.

[21] J. Wolff, *Das Gesetz der Transformation der Knochen*, Berlin: A. Hirschwald, 1892.

[22] J. G. Skedros and S. L. Baucom, "Mathematical analysis of trabecular trajectories in apparent trajectorial structures," *J. Theor. Biol.*, vol. 244, no. 1, pp. 15–45, 2007.

[23] P. Block and S. Adriaenssens, "Shell structures for architecture: Form finding and optimization," *Struct. Eng. Int.*, vol. 23, no. 2, pp. 196–208, 2013.

[24] F. Greifeneder, W. Fenz, B. Alkin, J. Brandstetter, and P. Moser, "Real-time prediction of brain deformation in surgical simulators using transformer-based surrogates," *Proc. 2025 Annual Modeling and Simulation Conference (ANNSIM'25)*, Madrid, Spain, 2025.

[25] SimScale GmbH, *SimScale: Cloud-Based Engineering Simulation*. [Online]. Available: https://www.simscale.com.