LEARNING TOPOLOGY-AWARE REPRESENTATIONS FOR MECHANICS-RELEVANT TPMS METAMATERIALS

Anonymous authors
Paper under double-blind review

ABSTRACT

Learning informative representations of mechanical Triply Periodic Minimal Surface (TPMS) metamaterials is important for mechanics analysis and structural design, yet existing methods have rarely explicitly supervised learned representations with information about global structural connectivity that is relevant to mechanical response. We propose a topology-aware representation learning framework that integrates persistent homology into a DMAE backbone to explicitly encode global structural connectivity. A persistence-guided masking strategy prioritizes patches associated with persistent topological features, while an auxiliary persistent-homology decoder encourages the latent representation to retain information about topological feature counts across connected components, tunnels, and voids. We evaluate the learned representations on two complementary TPMS datasets using family classification, mechanical-property classification, unsupervised clustering, and regression. Across the two datasets, our method demonstrates strong performance in both structural family and mechanical-property classification, improving family classification by up to 3.9% while achieving the best linear-probe performance. More notably, unsupervised property-clustering performance improves by up to 26.7% and 51.5% across the two datasets, while strong regression performance is retained with $R^2 = 0.993$. These results establish persistent topology as an effective inductive bias for representation learning of complex three-dimensional mechanical metamaterials.

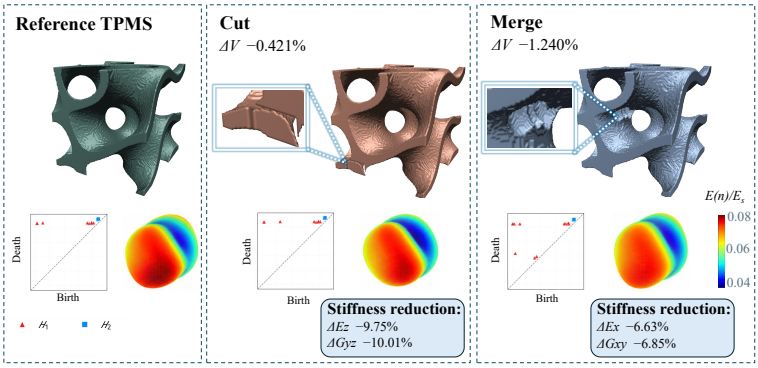

Figure 1: **Localized topological changes can induce disproportionate changes in the mechanical response of TPMS structures.** Starting from the reference TPMS (left), the *Cut* (middle) introduces a physical crack at an existing structural connection, whereas the *Merge* (right) joins two separated tunnels into a single connected passage. Despite only small changes in solid volume fraction ($\Delta V = -0.421\%$ and -1.240%, respectively), both modifications produce clear changes in the persistence diagrams (PDs, summaries of topological features) and directional elastic responses. The *Cut* yields maximum reductions of 9.75% in E_z and 10.01% in G_{yz}, whereas the *Merge* reduces E_x by 6.63% and G_{xy} by 6.85%. Additional topological perturbations, full directional stiffness changes, and detailed analysis are provided in Appendix A.1 and Fig. S1.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

1 INTRODUCTION

Architected meta-materials enable tailored and unconventional physical responses through structural architecture rather than material composition alone, creating opportunities in lightweight structures (Portela et al., 2021), energy absorption (Yang et al., 2022; Fang et al., 2025; Liu et al., 2024b), thermal management (Li et al., 2021), wave control (Shaltout et al., 2019), sensing (Berman et al., 2024), and other multifunctional applications (Hou et al., 2026). Among them, TPMS metamaterials have emerged as a versatile class of architected materials because their smooth, continuous and highly tunable periodic architectures combine interconnected load-bearing and transport pathways with high surface-area-to-volume ratios (Wang et al., 2025), making them attractive for a wide range of applications, particularly in mechanical metamaterial systems (Zhang et al., 2026; Li et al., 2025; Li & Fan, 2018), as illustrated in Appendix A.2.

As increasingly large libraries of mechanical metamaterials are explored computationally (Maurizi et al., 2025; Beatson et al., 2020; Wan et al., 2022), learning compact and informative representations of these complex structures has become increasingly important, and such representations should capture not only geometric similarity but also structural characteristics relevant to mechanical response. Existing studies have learned compact representations for several classes of metamaterials, including periodic microstructures with latent spaces organized by mechanical behavior (Wang et al., 2020b), nature-inspired architected structures optimized with respect to mechanical properties (Shen & Buehler, 2022), and truss-based metamaterials for generation and mechanical-property prediction (Zhan et al., 2025; Zheng et al., 2023). In contrast, TPMS-specific studies have primarily focused on multimodal generation (Makatura et al., 2025), surrogate-based simulation (Chen et al., 2024), and multi-objective mechanical optimization (Yan et al., 2026), while mechanics-relevant TPMS representation learning has not been explicitly studied as a primary objective.

To better learn mechanics-relevant representations of TPMS structures, it is important to capture not only local geometric characteristics, such as wall thickness, curvature, loops, voids, and surface orientation (Qiu et al., 2024), but also global connectivity through how local features are organized across the structure (Chen et al., 2019; Gibson & Ashby, 1997; Deshpande et al., 2001). This is consistent with Maxwell-type rigidity arguments, which link structural connectivity and constraint density to accessible deformation modes and mechanical stability (Pellegrino & Calladine, 1986).

Topological data analysis (TDA) provides a natural way to characterize global structural organization, while persistent homology (PH) captures such topology across multiple scales through persistence diagrams (PDs) (Ghrist, 2008; Edelsbrunner & Harer, 2010; Edelsbrunner et al., 2002; Zomorodian & Carlsson, 2005). This characterization is suitable for TPMS geometries, in which small geometric variations can alter global connectivity and create or eliminate topological features, potentially resulting in different mechanical responses (Wong et al., 2025; Edelsbrunner et al., 2002), as illustrated in Figure. 1. Prior topology-aware methods, as summarized in Table 1, have largely incorporated PH to encode structural connectivity in 2D (Hu et al., 2023; Wang et al., 2020a) or in generic geometric settings (Hofer et al., 2019; Moor et al., 2020; Chen & Gel, 2023; Tola et al., 2025). In materials and mechanics, PH has been used to characterize crystalline compounds (Jiang et al., 2021), structural materials (Tiong et al., 2023), granular media (Mei et al., 2023), and amorphous solids (Minamitani et al., 2025). In architected materials, prior studies have mainly used PH as a geometric descriptor or generative condition (Yan & Lin, 2023; Gao et al., 2025; Du et al., 2026), rather than directly supervising the learned latent representation to encode persistent-topological information. Whether such explicit topology preservation can improve mechanics-relevant representations of complex 3D TPMS structures therefore remains largely unexplored.

Here, we introduce a topology-aware representation learning framework for mechanics-relevant TPMS metamaterials, as illustrated in Figure. 2. Rather than using topology only as an external descriptor or generative constraint, we incorporate persistent-homology information directly into representation learning through two complementary mechanisms. Inspired by masking-based self-supervised representation learning in MAE (He et al., 2022) and DMAE (Wu et al., 2023), we introduce a **persistence-guided masking** strategy that prioritizes topologically informative regions, encouraging the encoder to integrate local geometry with global structural context. In parallel, an auxiliary **persistent-homology prediction head** explicitly supervises the latent representation to retain information about topological feature counts across connected components, tunnels, and voids. In this way, **guided masking determines where the model should focus during representation**

Table 1: Comparison of representative work across metamaterial and TPMS modeling, mechanics-aware methods, representation learning model, and topology-aware structural learning.

Representative work	Scope / focus					Topology-aware		
	Meta.	TPMS	3D struct.	Mech. rel.	Rep. learn.	TDA / PH	Topo head	Topo mask
General metamaterial modeling								
GraphMetaMat (Maurizi et al., 2025)	✓	–	✓	✓	–	–	–	–
Composable Energy Surrogates (Beatson et al., 2020)	✓	–	✓	✓	–	–	–	–
3D Metamaterial Force Control (Wan et al., 2022)	✓	–	✓	✓	–	–	–	–
Metamaterial representation & inverse design								
Neural-Adjoint (Ren et al., 2020)	△	–	–	✓	△	–	–	–
Video Diffusion Metamaterials (Bastek & Kochmann, 2023)	✓	–	–	✓	△	–	–	–
Deep Generative Metamaterial Design (Wang et al., 2020b)	✓	–	–	✓	✓	–	–	–
Nature-Inspired Unsupervised Design (Shen & Buehler, 2022)	✓	–	△	✓	✓	–	–	–
UniMate (Zhan et al., 2025)	✓	–	✓	✓	✓	–	–	–
Generative Truss Metamaterials (Zheng et al., 2023)	✓	–	✓	✓	✓	–	–	–
MetamatBench (Chen et al., 2025)	✓	–	✓	✓	△	–	–	–
Property-Informed Diffusion (Dai et al., 2026)	✓	–	✓	✓	△	–	–	–
TPMS-specific modeling & optimization								
MetaGen (Makatura et al., 2025)	✓	△	✓	✓	△	–	–	–
PNN-TPMS Multiscale FEM (Chen et al., 2024)	✓	✓	✓	✓	–	–	–	–
ML-Optimized TPMS Lattices (Yan et al., 2026)	✓	✓	✓	✓	–	–	–	–
PH-Based Thickness Design (Yan & Lin, 2023)	△	△	✓	–	–	✓	–	–
PH-Driven TPMS Optimization (Gao et al., 2025)	✓	✓	✓	✓	–	✓	–	–
Topology-aware modeling & learning models								
Topo-GenMeta (Du et al., 2026)	✓	–	–	✓	△	✓	–	–
Topological Crystal Representations (Jiang et al., 2021)	–	–	✓	✓	✓	✓	–	–
Void-Topology Failure Prediction (Tiong et al., 2023)	–	–	✓	✓	–	✓	–	–
PH for Granular Transition (Mei et al., 2023)	–	–	✓	✓	–	✓	–	–
PH for Amorphous Solids (Minamitani et al., 2025)	–	–	✓	✓	–	✓	–	–
Conformable Convolution (Yeganeh et al., 2026)	–	–	△	–	–	✓	–	–
Multiparameter PH Representations (Loiseaux et al., 2023)	–	–	△	–	✓	✓	–	–
Point-Level TopoRL (Grande & Schaub, 2025)	–	–	✓	–	–	✓	–	–
Probabilistic Topological Representations (Hu et al., 2023)	–	–	–	–	✓	✓	–	–
Topological Pooling (Chen & Gel, 2023)	–	–	–	–	–	✓	–	–
TopER (Tola et al., 2025)	–	–	–	–	–	✓	–	–
Connectivity-Optimized RL (Hofer et al., 2019)	–	–	–	–	–	✓	✓	–
TopoAE (Moor et al., 2020)	–	–	–	–	–	✓	✓	–
PH-SSL Molecular Learning (Luo et al., 2023)	–	–	✓	–	–	✓	✓	–
Ours	✓	✓	✓	✓	✓	✓	✓	✓

✓ = primary focus; △ = partial or indirect consideration; – = not addressed.
Meta. = metamaterials; TPMS = triply periodic minimal surfaces; TDA/PH = topological data analysis / persistent homology.
The blue sections denote general metamaterial and TPMS-oriented studies, the green section denotes topology-aware studies, and the blue–green rows indicate works belonging to both TPMS-specific and topology-aware categories. The yellow row highlights our method.

learning, while the **PH head constrains how many topological features should be preserved in latent space**. We evaluate representations from geometric and mechanical perspectives and show that topology-aware learning produces more structurally and mechanically informative latent representations for downstream property analysis.

2 Topology-aware representation learning workflow

Our framework extends a 3D DMAE (Wu et al., 2023) based on Vision Transformers (Dosovitskiy et al., 2021) by integrating persistent homology at the masking and latent-representation levels. Sec. 2.1 describes patch-based SDF masking and reconstruction for TPMS. Sec. 2.2 introduces persistence-guided masking, distinct from existing random or adaptive schemes (He et al., 2022; Hinojosa et al., 2024; Liu et al., 2024a). Sec. 2.3 presents auxiliary PH supervision that encourages the latent representation to retain global topological statistics. Architectural details are provided in Appendix A.3.

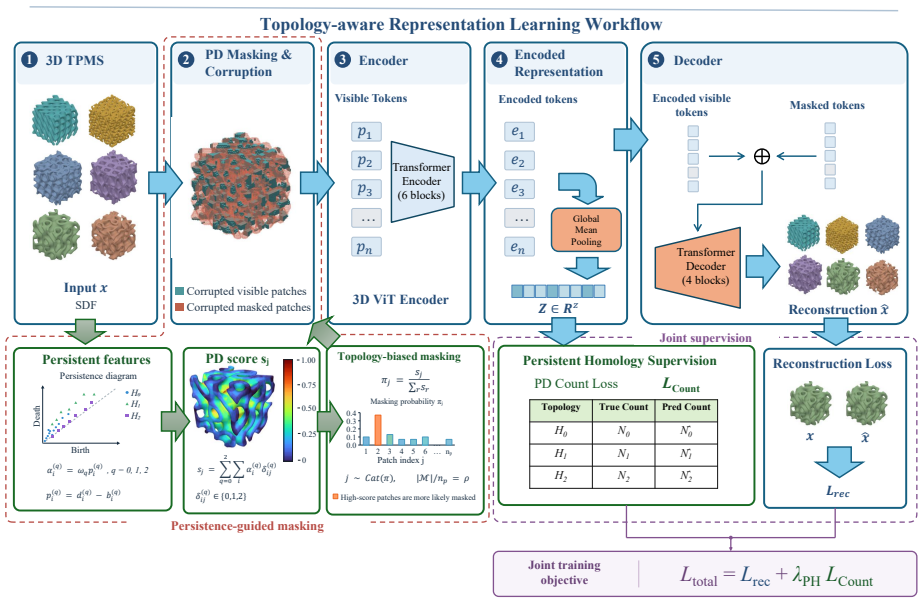

Figure 2: **Overview of the proposed topology-aware representation learning framework.** Starting from voxelized TPMS signed-distance fields, the framework integrates persistence-guided masking and a persistent-homology prediction objective into masked representation learning to preserve topology-relevant information in the latent representation.

2.1 Patch Masking of TPMS SDFs

Given an TPMS SDF tensor $\mathbf{S} \in \mathbb{R}^{N\times N\times N}$, we partition it into non-overlapping cubic patches of size $p\times p\times p$ using a patchification operator $\mathcal{P}(\cdot)$. Let $g = N/p$ denote the number of patches along each spatial dimension, yielding $n_p = g^3$ patches in total. The resulting clean patches are denoted by $\{\mathbf{P}_j\}_{j=1}^{n_p} = \mathcal{P}(\mathbf{S})$, where $\mathbf{P}_j \in \mathbb{R}^{p\times p\times p}$, and serve as reconstruction targets throughout training.

Following DMAE, Gaussian corruption is first applied to the clean SDF as $\widetilde{\mathbf{S}} = \mathbf{S} + \boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon}$ denotes Gaussian noise, yielding the corresponding corrupted patches $\{\widetilde{\mathbf{P}}_j\}_{j=1}^{n_p} = \mathcal{P}(\widetilde{\mathbf{S}})$. For a masking ratio ρ, we sample a masked set $\mathbb{M} \subset \{1,\dots,n_p\}$ with $|\mathbb{M}| \approx \rho n_p$, while $\mathbb{V} = \{1,\dots,n_p\} \setminus \mathbb{M}$ denotes the visible/unmasked set. In our framework, $\mathbb{M}$ is sampled according to the topology-guided masking distribution described in Sec. 2.2. Only the corrupted visible patches $\{\widetilde{\mathbf{P}}_j : j \in \mathbb{V}\}$ are embedded and processed by the encoder E_θ, producing latent tokens $\{\mathbf{h}_j\}_{j\in\mathbb{V}} = E_\theta(\{\widetilde{\mathbf{P}}_j\}_{j\in\mathbb{V}})$.

For reconstruction, the encoded visible tokens $\mathbf{h}_j$ are combined with a learnable mask token $\mathbf{m}$ at positions $j \in \mathbb{M}$ and restored to their original spatial ordering before being passed to the decoder D_ϕ. The decoder therefore predicts all reconstructed patches as $\{\hat{\mathbf{P}}_j\}_{j=1}^{n_p} = D_\phi(\{\mathbf{h}_j\}_{j\in\mathbb{V}}, \{\mathbf{m}\}_{j\in\mathbb{M}})$. Following DMAE, reconstruction is evaluated over all n_p patches against the clean targets $\mathbf{P}_j$:

$$\mathcal{L}_{\text{rec}} = \frac{1}{n_p}\sum_{j=1}^{n_p} \ell\Big(\hat{\mathbf{P}}_j, \mathbf{P}_j\Big), \tag{1}$$

where ℓ denotes the Smooth L1 loss. For $j \in \mathbb{V}$, the model denoises the corrupted observation $\widetilde{\mathbf{P}}_j$ toward its clean counterpart $\mathbf{P}_j$, whereas for $j \in \mathbb{M}$, it reconstructs $\mathbf{P}_j$ from the surrounding visible context. The combination of Gaussian corruption and topology-guided patch masking therefore encourages the encoder to learn representations that remain informative under both local perturbations and missing structural observations.

2.2 PERSISTENCE-GUIDED TOPOLOGICAL MASKING

Existing MAE-based methods have explored different masking strategies, including uniform random masking in vanilla MAE (He et al., 2022), filtered-noise data-independent masking in ColorMAE (Hinojosa et al., 2024), and distribution-aware adaptive masking in DaM (Liu et al., 2024a). While these strategies modify the spatial distribution of masked patches, they do not explicitly account for persistent topological information. We instead construct a patch-level topological importance score from persistent homology and use it to bias the masking distribution toward topology-relevant regions. Details of the persistent-homology computation and its formulation for TPMS structures are provided in Appendix A.4.

Insights: (1) As shown in Fig. 1 and Appendix A.1, localized topology-changing perturbations can modify structural connectivity and load-transfer pathways, leading to noticeable changes in effective stiffness even when the geometric variation is small. These changes are also reflected in the corresponding persistence diagrams. For example, in the *Cut* case, the perturbation alters the persistence of topological features, coinciding with a clear change in the elastic response, which suggests that persistent homology captures structural variations relevant to mechanical behavior. (2) Within a persistence diagram, features with larger persistence remain present over a wider range of filtration scales and therefore represent more robust topological characteristics (Edelsbrunner et al., 2002; Zomorodian & Carlsson, 2005).

Assumption: We therefore hypothesize that preferentially masking regions associated with highly persistent topological features encourages the model to better learn and reconstruct mechanically relevant structural information. We therefore assign each patch a persistence-weighted topological score and use it to define the masking probability.

Specifically, the topological importance score of patch j is defined as $s_j = \sum_{q=0}^{2} \sum_i \alpha_i^{(q)} \delta_{ij}^{(q)}$, where $\alpha_i^{(q)} = \omega_q p_i^{(q)}$, $p_i^{(q)}$ denotes the persistence of the i-th topological feature in homology dimension q, and ω_q controls the relative contribution of dimension q. The term $\delta_{ij}^{(q)} \in \{0, 1, 2\}$ counts how many of the birth and death critical cells associated with feature i are mapped to patch j. Thus, features with larger persistence contribute more strongly to the score, while ω_q allows different homology dimensions to be weighted separately. Accordingly, a patch receives a larger score when it contains critical cells associated with more persistent topological features.

The patch scores are then normalized into masking probabilities, $\pi_j = s_j / \sum_{r=1}^{n_p} s_r$, for $j \in \{1, \ldots, n_p\}$. Thus, patches with larger π_j are more likely to be selected for masking and are sampled from this distribution until the prescribed masking ratio ρ is reached. Representative PD-score distributions for the baseline structure and four localized topological modifications are visualized in Appendix Fig. S2. Additional PD-score heatmaps for diverse TPMS structures are provided in Appendix A.5 and Fig. S4.

2.3 PERSISTENT-HOMOLOGY SUPERVISION

We introduce an auxiliary PH prediction head operating on the pooled encoder representation. Since the number of retained features can vary substantially across structures, we use log-transformed feature counts as the prediction targets, with $\hat{\mathbf{c}} = [\hat{c}_0, \hat{c}_1, \hat{c}_2]$. The PH supervision loss is defined as,

$$\mathcal{L}_{\mathrm{PH}} = \sum_{q=0}^{2} \ell_{\mathrm{SL1}} \left(\hat{c}_q, \log(1 + n_q) \right), \tag{2}$$

where n_q denotes the number of retained persistence features in homology dimension $q \in \{0, 1, 2\}$, $\log(1 + n_q)$ is the corresponding log-transformed count target, and ℓ_{SL1} denotes the Smooth L1 loss. The logarithmic transformation reduces the dynamic range of the feature counts and stabilizes regression across structures with substantially different topological complexity.

Insights: For Fig. 1 and Table. S1, the *Merged* modification changes the H_1 and H_2 feature counts by -203 and -35, respectively, while also producing pronounced changes in the elastic response. These count variations reflect changes in global topological complexity that are not explicitly supervised by the persistence-guided spatial masking.

Assumption: We therefore hypothesize that explicitly requiring the latent representation to predict global H_0, H_1, and H_2 feature counts can complement persistence-guided masking and encourage mechanics-relevant global topological information to be retained after representation compression.

The complete training objective is

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{rec}} + \lambda_{\text{PD}}\mathcal{L}_{\text{PH}}, \tag{3}$$

where λ_{PD} controls the contribution of PH supervision to the overall training objective.

3 EXPERIMENTS

To evaluate the assumptions introduced in Sec. 2, we assess the learned representations through downstream tasks, ablation studies, and latent-space analysis. We first describe the construction of the two TPMS datasets in Sec. 3.1 and the competing baselines in Sec. 3.2. Sec. 3.3 evaluates structural and mechanics-relevant information through four downstream tasks, while Sec. 3.4 examines the effects of persistence-guided masking and PH supervision. Finally, Sec. 3.5 analyzes how the learned latent space aligns with topology- and mechanics-related descriptors.

3.1 DATASET

We used two complementary TPMS datasets to evaluate representation learning across procedurally generated and continuously interpolated structural design spaces.

MetaDB. The open-source MetaDB dataset represents metamaterial geometries using the MetaDSL procedural representation and expands hand-authored structures through type-aware mutation operators (Makatura et al., 2025). From MetaDB, we select 13 base TPMS types and their mutation-derived variants, resulting in 13,114 structures. In the original MetaDB, each geometry is voxelized and its effective elastic properties are obtained through periodic homogenization using the finite-element framework ProcMeta (Makatura et al., 2025). ProcMeta assumes an isotropic base material with normalized Young's modulus $E = 1$, density $\rho = 1$, and Poisson's ratio $\nu = 0.45$, and provides the resulting homogenized stiffness tensor $\mathbf{C} \in \mathbb{R}^{6\times6}$ in Voigt notation as the mechanical-property label. Further details on the stiffness tensor and the finite-element homogenization procedure are provided in Appendix A.6. We resampled all voxelized geometries to a resolution of 64^3.

Interpolation. We constructed a second dataset based on (Perroni-Scharf et al., 2025) by interpolating eight canonical TPMS primitives. Each primitive was represented by a periodic implicit function $f_i(x, y, z)$, and the interpolated geometry was defined as $F(x, y, z) = \sum_{i=1}^{8} w_i f_i(x, y, z)$.

We obtained the interpolated TPMS geometry from the zero-level set $F(x, y, z) = 0$. For each interpolated geometry, we combined up to six base TPMS families, with each nonzero interpolation weight taking values from 10% to 100% in increments of 10%, while satisfying $\sum_{i=1}^{8} w_i = 1$. Using this interpolation scheme, we generated 18,549 structures and represented each geometry as a 64^3 SDF. We then applied the same periodic homogenization framework, ProcMeta, used for MetaDB to compute the corresponding 6×6 effective stiffness tensor $\mathbf{C}$ as the mechanical-property label.

3.2 BASELINES

We compare our proposed method with geometry-based representation-learning baselines. For general representation learning, we include MVAE (Choy et al., 2019), SVAE (Xiang et al., 2025), VICReg (Bardes et al., 2022), and I-JEPA (Assran et al., 2023). For masking-strategy baselines, we use MAE (He et al., 2022), DMAE (Wu et al., 2023), ColorMAE Hinojosa et al. (2024), and DaM Liu et al. (2024a). Specifically, MAE uses random patch masking, whereas DMAE combines random masking with Gaussian input corruption. ColorMAE generates structured masks by filtering random noise, while DaM adaptively selects masked tokens according to estimated distribution shifts. Additionally, we include ScalableTopo (Wong et al., 2025) as a topology-informed baseline based on scalable persistent-homology regularization.

Table 2: **Task 1 – Family classification.** $\mathrm{F}_{\mathrm{Lin}}$ and $\mathrm{F}_{k\mathrm{NN}}$ denote macro-F1 scores on MetaDB using a linear probe and kNN classifier, respectively. For Interpolation TPMS, $\mathrm{Acc}_{\mathrm{Lin}}^{\mathrm{Str/Rel}}$ and $\mathrm{Acc}_{k\mathrm{NN}}^{\mathrm{Str/Rel}}$ denote strict and relaxed accuracies with respect to the dominant-family labels. Higher is better for all metrics.

MODEL	METADB		INTERPOLATION TPMS			
	$\mathrm{F}_{\mathrm{Lin}}\uparrow$	$\mathrm{F}_{k\mathrm{NN}}\uparrow$	$\mathrm{Acc}_{\mathrm{Lin}}^{\mathrm{Str}}\uparrow$	$\mathrm{Acc}_{\mathrm{Lin}}^{\mathrm{Rel}}\uparrow$	$\mathrm{Acc}_{k\mathrm{NN}}^{\mathrm{Str}}\uparrow$	$\mathrm{Acc}_{k\mathrm{NN}}^{\mathrm{Rel}}\uparrow$
General representation learning						
MVAE	0.009	0.581	0.038	0.051	0.446	0.559
SVAE	0.479	0.955	0.638	0.663	0.787	0.914
VICReg	0.913	0.853	0.263	0.299	0.370	0.456
I-JEPA	0.912	0.902	0.403	0.446	0.497	0.618
Masking strategy baselines						
MAE	0.901	0.907	0.549	0.588	0.643	0.792
DMAE	0.907	0.894	0.724	0.852	0.819	0.932
ColorMAE	0.956	0.952	0.877	0.888	0.821	0.931
DaM	0.952	0.948	0.879	0.893	0.829	0.937
Topology-informed baseline						
ScalableTopo	0.957	0.954	**0.885**	**0.897**	0.813	0.928
OURS	**0.994**	**0.985**	0.867	0.880	**0.831**	**0.939**

3.3 EVALUATION ON DOWNSTREAM TASKS

We evaluate latent representations on four downstream tasks: family classification, mechanical property classification, clustering, and regression. The first probes TPMS structural identity, while the others assess the encoding of effective mechanical properties. All experiments use a fixed 70/15/15 train/validation/test split, with results reported over five random seeds. Models are trained and evaluated on NVIDIA H100 GPUs on Compute Canada clusters.

3.3.1 FAMILY CLASSIFICATION

We evaluate frozen representations using a linear probe and a k-nearest-neighbor (kNN) classifier ($k = 15$). For MetaDB, the 13 base TPMS families serve as single-label targets, with macro-F1 reported as $\mathrm{F}_{\mathrm{Lin}}$ and $\mathrm{F}_{k\mathrm{NN}}$, respectively. For Interpolation TPMS, labels correspond to families with the largest interpolation weight, including ties. Thus, 50% Gyroid–30% FRD–20% PMY is labeled Gyroid, whereas 50% Gyroid–50% Diamond receives both labels. Strict accuracy ($\mathrm{Acc}^{\mathrm{Str}}$) requires exact agreement between predicted and true label sets; relaxed accuracy ($\mathrm{Acc}^{\mathrm{Rel}}$) accepts any nonempty predicted subset of the true labels.

As shown in Table 2, our method achieves the best performance on MetaDB for both linear and kNN probes. On Interpolation TPMS, it achieves the highest kNN accuracy under both strict and relaxed evaluation, while the linear-probe results remain slightly below the best baseline. These results indicate that topology-aware learning preserves family-relevant structural information while promoting a more coherent local organization of structurally similar samples in the latent space.

3.3.2 MECHANICAL-PROPERTY PREDICTION

We evaluate mechanics-relevant information in the frozen representations through three downstream tasks. C_{11}, the effective normal stiffness along a principal direction, is used for property classification (Task 2), clustering (Task 3), and regression (Task 4). Figure 3 summarizes representative results across both datasets, with complete numerical results provided in Appendix A.7.

Task 2 – C_{11} property classification. We discretize C_{11} into *Low*, *Medium*, and *High* groups and evaluate linear and kNN ($k = 15$) probes using balanced accuracy and macro-F1. Figure 3 highlights kNN macro-F1 ($F_{k\mathrm{NN}}$): our method achieves the highest score on Interpolation TPMS (0.887) and remains close to the best baseline on MetaDB (0.951 versus 0.954). The linear probe

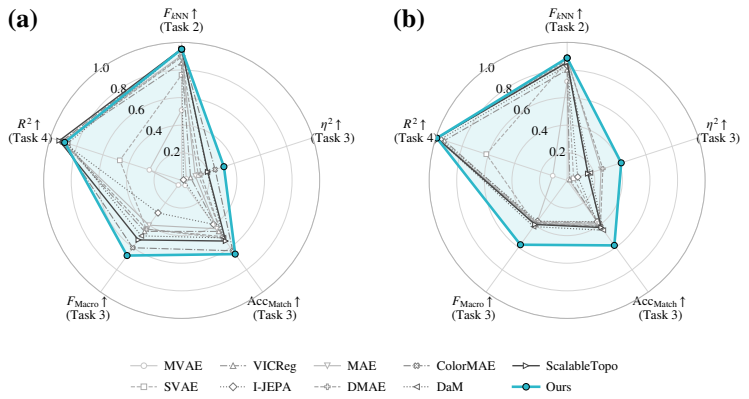

Figure 3: **Summary of mechanical-property downstream performance on (a) MetaDB and (b) Interpolation TPMS.** The radar plots provide a compact comparison of frozen latent representations: $\mathrm{F1}_{k\mathrm{NN}}$ for C_{11} property classification (Task 2), η^2, $\mathrm{Acc}_{\mathrm{Match}}$, and $\mathrm{F}_{\mathrm{Macro}}$ for C_{11} property clustering (Task 3), and R^2 for C_{11} property regression (Task 4). Higher is better for all metrics shown. Detailed numerical results for all metrics and baselines are provided in the Appendix (Tables S2, S3, and S4).

further achieves the highest balanced accuracy and macro-F1 on both datasets, supporting effective discrimination between stiffness groups.

Task 3 – C_{11} property clustering. We assess whether the latent space organizes structures according to C_{11} without supervision using η^2, $\mathrm{Acc}_{\mathrm{Match}}$, and F_{Macro}. As shown in Fig. 3, our method achieves the highest values across all three metrics on both datasets. Relative to the strongest competing baseline, it improves the three metrics by 26.7%, 4.8%, and 11.6% on MetaDB, and by 51.5%, 30.5%, and 39.1% on Interpolation TPMS, respectively. These results demonstrate a substantially stronger mechanics-aware organization of the latent space.

Task 4 – C_{11} property regression. Finally, we assess whether continuous mechanical-property information can be predicted from the frozen representations through C_{11} regression. Our method achieves $R^2 = 0.899$ on MetaDB and $R^2 = 0.993$ on Interpolation TPMS, showing the best performance on the latter while maintaining low prediction errors across both datasets. These results show that the learned representations retain continuous mechanics-relevant information in addition to discrete property organization.

Overall, Tasks 2–4 show that our method consistently preserves mechanics-relevant information, with particularly strong improvements in unsupervised property clustering while maintaining competitive classification and regression performance.

3.4 Ablation Study

Figure 4 examines the individual and combined effects of persistence-guided masking and the auxiliary PD prediction head across both datasets. The full model provides the most consistent overall high performance and achieves the strongest Task 3 clustering results in terms of η^2, ARI, and NMI on both datasets, indicating improved organization of mechanics-relevant information in the latent space. The additional C_{21} (with 21 independent elastic stiffness components in Voigt notation) evaluation shows a similar trend, indicating that the benefit is not restricted to a single elastic-property target. Although individual ablations can perform better on isolated metrics, neither component alone provides the same overall consistency. Definitions of ARI and NMI used for the clustering evaluation are provided in Appendix Sec. A.8.

3.5 Latent-Space Alignment Analysis

To directly assess the learned latent space, we measure the Spearman correlation between pairwise latent distances and four descriptors: the PD-score masking distribution, H_1 and H_2 persistence-

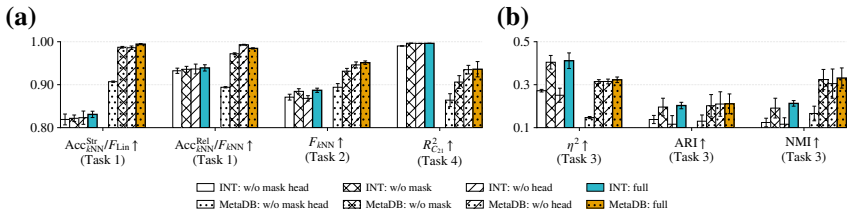

Figure 4: **Ablation study of persistence-guided masking and the auxiliary PD prediction head. (a)** reports representative metrics from Tasks 1, 2, and 4, while **(b)** reports the three Task 3 clustering metrics. We compare the DMAE baseline without either topology-aware mechanism (*w/o mask head*), variants removing either the persistence-guided mask (*w/o mask*) or PD head (*w/o head*), and the full model using both mechanisms (*full*). Complete numerical results across all evaluation metrics and both datasets are provided in Appendix Table S5. Bars show the mean over five seeds, with error bars indicating the corresponding standard deviation. Higher is better for all metrics.

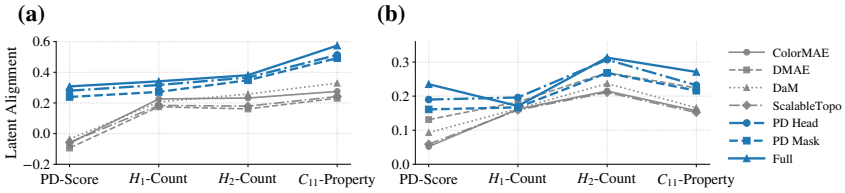

Figure 5: **Direct latent-space alignment on (a) MetaDB and (b) Interpolation TPMS.** PD-Score, H_1-Count, H_2-Count, and C_{11}-Property measure the correspondence between pairwise latent distances and the persistence-guided masking distribution, persistence-feature counts, and mechanical-property differences, respectively. Higher Spearman correlation indicates stronger alignment.

feature counts, and C_{11}. These metrics evaluate whether the latent space reflects the spatial topology distribution used for PD masking, the topological statistics supervised by the PD head, and mechanics-related variation. Detailed definitions of the alignment metrics are provided in Appendix A.9. As shown in Fig. 5, topology-aware learning generally strengthens these latent correspondences across both datasets. On MetaDB, the full model achieves the strongest alignment for all four metrics. On Interpolation TPMS, the full model performs best for PD-score, H_2-count, and C_{11} alignment, while the PD head achieves the strongest H_1-count alignment. These results indicate complementary effects of topology-guided masking and explicit topological supervision.

4 Conclusion

We presented a topology-aware representation learning framework for TPMS metamaterials that explicitly integrates persistent homology into DMAE. The proposed method combines persistence-guided masking, which directs learning toward topologically informative structural regions, with an auxiliary PH prediction head that encourages the latent representation to retain topological information. The learned representations demonstrate strong performance in structural family classification and multiple mechanical-property downstream tasks, including C_{11} classification, clustering, and regression. In particular, the results show improved organization of mechanically meaningful latent structure, especially for continuously interpolated TPMS geometries. Ablation studies reveal complementary roles of the two topology-aware components, while direct latent-space analysis shows stronger correspondence with PD-score distributions, persistence-feature counts, and mechanical properties. Overall, these results demonstrate that persistent topology provides a useful inductive bias for learning compact representations of complex three-dimensional architected materials and provide a foundation for future topology-aware property prediction and generative design.

AI USE STATEMENT

In this work, we used generative AI tools for manuscript polishing as well as code debugging. We have not used generative AI tools for developing the research ideas, methodology, experimental design, or scientific claims, nor to generate datasets or mathematical results. All AI-assisted code modifications were reviewed, executed, and verified by two authors, and all AI-assisted text was carefully reviewed and revised before inclusion in the manuscript. We take full responsibility for the final content of this work, including any text or code produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

To facilitate reproducibility, we provide detailed implementation and evaluation information in the supplementary material. Appendix A.1 presents the controlled topology–mechanics perturbation analysis and persistence-guided masking behavior; Sec. A.3 specifies the model architecture, patchification, encoder–decoder configuration, and training objectives; Sec. A.4 describes the persistent-homology pipeline, including the SDF-based cubical filtration and persistence-diagram construction for TPMS; Sec. A.5 illustrates the topological diversity of the dataset; Sec. A.6 details the finite-element homogenization procedure used to obtain the mechanical-property labels; and Sec. A.7 provides the downstream evaluation protocols, metric definitions, and complete results for Tasks 2–4. Complete ablation results are reported in Table S5, with the additional clustering metrics used in the ablation analysis defined in Sec. A.8. Sec. A.9 provides the definitions of the latent-space alignment measures used in the representation analysis, and Sec. A.10 reports parameter-sensitivity analyses and hyperparameter selection on both Interpolation TPMS and MetaDB.

The MetaDB dataset used in this study is derived from a publicly available open-source dataset, while the Interpolation TPMS dataset constructed for this work will be released on GitHub. The source code and implementation details for the proposed method will also be released through the accompanying GitHub repository.

REFERENCES

Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15619–15629, 2023.

Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regularization for self-supervised learning. In *International Conference on Learning Representations*, 2022.

Adarsh Basavalingappa, Ming Y. Shen, and James R. Lloyd. Modeling the copper microstructure and elastic anisotropy and studying its impact on reliability in nanoscale interconnects. *Mechanics of Advanced Materials and Modern Processes*, 3:6, 2017. doi: 10.1186/s40759-017-0021-5.

Jan-Hendrik Bastek and Dennis M. Kochmann. Inverse design of nonlinear mechanical metamaterials via video denoising diffusion models. *Nature Machine Intelligence*, 5:1466–1475, 2023. doi: 10.1038/s42256-023-00762-x. URL https://www.nature.com/articles/s42256-023-00762-x.

Alex Beatson, Jordan Ash, Geoffrey Roeder, Tianju Xue, and Ryan P. Adams. Learning composable energy surrogates for PDE order reduction. In *Advances in Neural Information Processing Systems*, volume 33, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/0332d694daab22e0e0eaf7a5e88433f9-Abstract.html.

Arielle Berman, Kaiwen Hsiao, Samuel E. Root, Hojung Choi, Daniel Ilyn, Chengyi Xu, Emily Stein, Mark Cutkosky, Joseph M. DeSimone, and Zhenan Bao. Additively manufactured microlattice dielectrics for multiaxial capacitive sensors. *Science Advances*, 10(40):eadq8866, 2024. doi: 10.1126/sciadv.adq8866. URL https://doi.org/10.1126/sciadv.adq8866.

Bea Bleile, Adélie Garin, Teresa Heiss, Kelly Maggs, and Vanessa Robins. The persistent homology of dual digital image constructions. In Ellen Gasparovic, Vanessa Robins, and Katharine Turner (eds.), *Research in Computational Topology 2*, volume 30 of *Association for Women in*

Mathematics Series, pp. 1–21. Springer, Cham, 2022. doi: 10.1007/978-3-030-95519-9_1. URL https://doi.org/10.1007/978-3-030-95519-9_1.

Jianpeng Chen, Wangzhi Zhan, Haohui Wang, Zian Jia, Jingru Gan, Junkai Zhang, Jingyuan Qi, Tingwei Chen, Lifu Huang, Muhao Chen, Ling Li, Wei Wang, and Dawei Zhou. MetamatBench: Integrating heterogeneous data, computational tools, and visual interface for metamaterial discovery. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2*, pp. 5334–5344. Association for Computing Machinery, 2025. doi: 10.1145/3711896.3737416. URL https://doi.org/10.1145/3711896.3737416.

Wen Chen, Seth Watts, Julie A. Jackson, William L. Smith, Daniel A. Tortorelli, and Christopher M. Spadaccini. Stiff isotropic lattices beyond the maxwell criterion. *Science Advances*, 5(9):eaaw1937, 2019. doi: 10.1126/sciadv.aaw1937.

Yan-Zhen Chen, Chu-Hao Wang, Tsung-Yeh Hsieh, Cheng-Che Tung, Po-Yu Chen, and Tsung-Hui Huang. An efficient parameterized neural network enhanced multiscale finite element modeling for triply periodic minimal surface meta-structures and its applications for femur. *Journal of Materials Research and Technology*, 30:6176–6194, 2024. doi: 10.1016/j.jmrt.2024.05.023. URL https://doi.org/10.1016/j.jmrt.2024.05.023.

Yuzhou Chen and Yulia R. Gel. Topological pooling on graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 7096–7103, 2023. doi: 10.1609/aaai.v37i6.25866. URL https://ojs.aaai.org/index.php/AAAI/article/view/25866.

Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3075–3084, 2019.

David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, 2007. doi: 10.1007/s00454-006-1276-5.

Bingxuan Dai, Hongsong Wang, and Jie Gui. Property-informed diffusion-based text-to-microstructure generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 36758–36768, 2026. URL https://openaccess.thecvf.com/content/CVPR2026/html/Dai_Property-Informed_Diffusion-Based_Text-to-Microstructure_Generation_CVPR_2026_paper.html.

Vikram S. Deshpande, Michael F. Ashby, and Norman A. Fleck. Foam topology: Bending versus stretching dominated architectures. *Acta Materialia*, 49(6):1035–1040, 2001. doi: 10.1016/S1359-6454(00)00379-7.

Pawel Dlotko. Cubical complex. In *GUDHI User and Reference Manual*. GUDHI Editorial Board, 3.13.0 edition, 2026. URL https://gudhi.inria.fr/doc/3.13.0/group__cubical__complex.html.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.

Liang Du, Jiangbei Hu, Shengfa Wang, Yu Jiang, Na Lei, Ying He, and Zhongxuan Luo. TopoGenMeta: Generative design of metamaterials based on diffusion model with attention to topology. *Computer-Aided Design*, 190:103977, 2026. doi: 10.1016/j.cad.2025.103977. URL https://doi.org/10.1016/j.cad.2025.103977.

H. Edelsbrunner, D. Letscher, and A. J. Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 28(4):511–533, 2002.

Herbert Edelsbrunner and John Harer. *Computational Topology: An Introduction*. American Mathematical Society, 2010.

Xin Fang, Dianlong Yu, Jihong Wen, Yifan Dai, Matthew R. Begley, Huajian Gao, and Peter Gumbsch. Large recoverable elastic energy in chiral metamaterials via twist buckling. *Nature*, 639 (8055):639–645, 2025. doi: 10.1038/s41586-025-08658-z. URL https://doi.org/10.1038/s41586-025-08658-z.

fly123. Chinese trunk aircraft C919 3D model. CGTrader, 2014. URL https://www.cgtrader.com/3d-models/aircraft/commercial-aircraft/chinese-trunk-aircraft-c910. Accessed: September 24, 2026.

Depeng Gao, Yuanzhi Zhang, Hongwei Lin, and Qiang Zou. Persistent homology-driven optimization of effective relative density range for triply periodic minimal surfaces. *Computer-Aided Design*, 180:103835, 2025. doi: 10.1016/j.cad.2024.103835. URL https://doi.org/10.1016/j.cad.2024.103835.

Robert Ghrist. Barcodes: The persistent topology of data. *Bulletin of the American Mathematical Society*, 45(1):61–75, 2008.

Lorna J. Gibson and Michael F. Ashby. *Cellular Solids: Structure and Properties*. Cambridge University Press, Cambridge, 2 edition, 1997.

Vincent Peter Grande and Michael T. Schaub. Point-level topological representation learning on point clouds. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 20368–20398. PMLR, 2025. URL https://proceedings.mlr.press/v267/grande25a.html.

Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Doll
'ar, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16000–16009, 2022.

Carlos Hinojosa, Shuming Liu, and Bernard Ghanem. Colormae: Exploring data-independent masking strategies in masked autoencoders, 2024.

Christoph Hofer, Roland Kwitt, Marc Niethammer, and Mandar Dixit. Connectivity-optimized representation learning via persistent homology. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2751–2760. PMLR, 2019. URL https://proceedings.mlr.press/v97/hofer19a.html.

Junming Hou, Long Chen, Xuan Zheng, Jia Wei Wu, Jian Wei You, Zi Xuan Cai, Jiahan Huang, Chenxu Wu, Jian Lin Su, Lianlin Li, Jia Nan Zhang, and Tie Jun Cui. Generative model for information metamaterial design. *Nature Computational Science*, 6(8):870–881, 2026. doi: 10.1038/s43588-026-01025-6. URL https://doi.org/10.1038/s43588-026-01025-6.

Xiaoling Hu, Dimitris Samaras, and Chao Chen. Learning probabilistic topological representations using discrete morse theory. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=cXMHQD-xQas.

Yi Jiang, Dong Chen, Xin Chen, Tangyi Li, Guo-Wei Wei, and Feng Pan. Topological representations of crystalline compounds for the machine-learning prediction of materials properties. *npj Computational Materials*, 7:28, 2021. doi: 10.1038/s41524-021-00493-w. URL https://www.nature.com/articles/s41524-021-00493-w.

Jiuyi Li, Li Wang, Xing He, Jianxiong Liang, and Chaofang Dong. Additively manufactured metallic tpms lattice structures: design strategies, fabrication, multifunctional properties, and applications. *npj Advanced Manufacturing*, 2:45, 2025. doi: 10.1038/s44334-025-00057-6.

Wei Li and Shanhui Fan. Nanophotonic control of thermal radiation for energy applications [invited]. *Optics Express*, 26(12):15995–16021, 2018. doi: 10.1364/OE.26.015995. URL https://doi.org/10.1364/OE.26.015995.

Ying Li, Wei Li, Tiancheng Han, Xu Zheng, Jiaxin Li, Baowen Li, Shanhui Fan, and Cheng-Wei Qiu. Transforming heat transfer with thermal metamaterials and devices. *Nature Reviews Materials*, 6 (6):488–507, 2021.

Jiaming Liu, Ran Xu, Senqiao Yang, Renrui Zhang, Qizhe Zhang, Zehui Chen, Yandong Guo, and Shanghang Zhang. Continual-mae: Adaptive distribution masked autoencoders for continual test-time adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 28653–28663, 2024a.

Wenfeng Liu, Shahram Janbaz, David Dykstra, Bernard Ennis, and Corentin Coulais. Harnessing plasticity in sequential metamaterials for ideal shock absorption. *Nature*, 634(8035):842–847, 2024b. doi: 10.1038/s41586-024-08037-0. URL https://doi.org/10.1038/s41586-024-08037-0.

David Loiseaux, Mathieu Carrière, and Andrew Blumberg. A framework for fast and stable representations of multiparameter persistent homology decompositions. In *Advances in Neural Information Processing Systems*, volume 36, 2023. doi: 10.52202/075280-1553. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/702b67152ec4435795f681865b67999c-Abstract-Conference.html.

Yuankai Luo, Lei Shi, and Veronika Thost. Improving self-supervised molecular representation learning using persistent homology. In *Advances in Neural Information Processing Systems*, volume 36, 2023. doi: 10.52202/075280-1475. URL https://papers.nips.cc/paper/2023/hash/6b555e8552240d6dfe0767146c9ebf36-Abstract-Conference.html.

Liane Makatura, Benjamin Jones, Siyuan Bian, and Wojciech Matusik. MetaGen: A DSL, database, and benchmark for VLM-assisted metamaterial generation, 2025. URL https://arxiv.org/abs/2508.17568.

Marco Maurizi, Derek Xu, Yu-Tong Wang, Desheng Yao, David Hahn, Mourad Oudich, Anish Satpati, Mathieu Bauchy, Wei Wang, Yizhou Sun, Yun Jing, and Xiaoyu Rayne Zheng. Designing metamaterials with programmable nonlinear responses and geometric constraints in graph space. *Nature Machine Intelligence*, 7:1023–1036, 2025. doi: 10.1038/s42256-025-01067-x. URL https://www.nature.com/articles/s42256-025-01067-x.

Jiangzhou Mei, Gang Ma, Jiaying Liu, François Nicot, and Wei Zhou. Modeling shear-induced solid-liquid transition of granular materials using persistent homology. *Journal of the Mechanics and Physics of Solids*, 176:105307, 2023. doi: 10.1016/j.jmps.2023.105307. URL https://doi.org/10.1016/j.jmps.2023.105307.

Emi Minamitani, Takenobu Nakamura, Ippei Obayashi, and Hideyuki Mizuno. Persistent homology elucidates hierarchical structures responsible for mechanical properties in covalent amorphous solids. *Nature Communications*, 16:8226, 2025. doi: 10.1038/s41467-025-63424-z. URL https://www.nature.com/articles/s41467-025-63424-z.

Michael Moor, Max Horn, Bastian Rieck, and Karsten Borgwardt. Topological autoencoders. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 7045–7054, 2020.

Sergio Pellegrino and C. R. Calladine. Matrix analysis of statically and kinematically indeterminate frameworks. *International Journal of Solids and Structures*, 22(4):409–428, 1986.

Maxine Perroni-Scharf, Zachary Ferguson, Thomas Butruille, Carlos M. Portela, and Mina Konaković Luković. Data-efficient discovery of hyperelastic tpms metamaterials with extreme energy dissipation. In *ACM SIGGRAPH 2025 Conference Papers*, 2025. doi: 10.1145/3721238.3730759.

Carlos M. Portela, Bryce W. Edwards, David Veysset, Yuchen Sun, Keith A. Nelson, Dennis M. Kochmann, and Julia R. Greer. Supersonic impact resilience of nanoarchitected carbon. *Nature Materials*, 20(11):1491–1497, 2021. doi: 10.1038/s41563-021-01033-z.

Na Qiu, Yuheng Wan, Yijun Shen, and Jianguang Fang. Experimental and numerical studies on mechanical properties of tpms structures. *International Journal of Mechanical Sciences*, 261:108657, 2024. doi: 10.1016/j.ijmecsci.2023.108657.

Simiao Ren, Willie Padilla, and Jordan Malof. Benchmarking deep inverse models over time, and the neural-adjoint method. In *Advances in Neural Information Processing Systems*, volume 33, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/007ff380ee5ac49ffc34442f5c2a2b86-Abstract.html.

Amr M. Shaltout, Vladimir M. Shalaev, and Mark L. Brongersma. Spatiotemporal light control with active metasurfaces. *Science*, 364:eaat3100, 2019. doi: 10.1126/science.aat3100.

Sabrina Chin-yun Shen and Markus J. Buehler. Nature-inspired architected materials using unsupervised deep learning. *Communications Engineering*, 1:37, 2022. doi: 10.1038/s44172-022-00037-0. URL https://www.nature.com/articles/s44172-022-00037-0.

Leslie Ching Ow Tiong, Gunjick Lee, Gyeong Hoon Yi, Seok Su Sohn, and Donghun Kim. Predicting failure progressions of structural materials via deep learning based on void topology. *Acta Materialia*, 250:118862, 2023. doi: 10.1016/j.actamat.2023.118862. URL https://doi.org/10.1016/j.actamat.2023.118862.

Astrit Tola, Funmilola Mary Taiwo, Cuneyt Akcora, and Baris Coskunuzer. TopER: Topological embeddings in graph representation learning. In *Advances in Neural Information Processing Systems*, volume 38, 2025. doi: 10.52202/085713-2968. URL https://papers.neurips.cc/paper_files/paper/2025/hash/8027ace571384361920665f1d1b69758-Abstract-Conference.html.

Fang Wan, Xiaobo Liu, Ning Guo, Xudong Han, Feng Tian, and Chaoyang Song. Visual learning towards soft robot force control using a 3D metamaterial with differential stiffness. In *Proceedings of the 5th Conference on Robot Learning*, volume 164 of *Proceedings of Machine Learning Research*, pp. 1269–1278. PMLR, 2022. URL https://proceedings.mlr.press/v164/wan22a.html.

Fan Wang, Huidong Liu, Dimitris Samaras, and Chao Chen. Topogan: A topology-aware generative adversarial network. In *Computer Vision – ECCV 2020*, pp. 118–136, 2020a. doi: 10.1007/978-3-030-58580-8_8.

Liqiang Wang, Di Yin, James Utama Surjadi, Junhao Ding, Huangliu Fu, Xin Zhou, Rui Li, Mengxue Chen, Xinxin Li, Xu Song, Johnny C. Ho, and Yang Lu. Harnessing screw dislocations in shell-lattice metamaterials for efficient, stable electrocatalysts. *Nature Communications*, 16:7273, 2025. doi: 10.1038/s41467-025-62489-0. URL https://doi.org/10.1038/s41467-025-62489-0.

Liwei Wang, Yu-Chin Chan, Faez Ahmed, Zhao Liu, Ping Zhu, and Wei Chen. Deep generative modeling for mechanistic-based learning and design of metamaterial systems. *Computer Methods in Applied Mechanics and Engineering*, 372:113377, 2020b. doi: 10.1016/j.cma.2020.113377. URL https://doi.org/10.1016/j.cma.2020.113377.

Yujia Wang, Xuan Zhang, Zihe Li, Huajian Gao, and Xiaoyan Li. Achieving the theoretical limit of strength in shell-based carbon nanolattices. *Proceedings of the National Academy of Sciences*, 119(34):e2119536119, 2022. doi: 10.1073/pnas.2119536119.

Hiu-Tung Wong, Darrick Lee, and Hong Yan. Towards scalable topological regularizers. In *International Conference on Learning Representations*, 2025.

Quanlin Wu, Hang Ye, Yuntian Gu, Huishuai Zhang, Liwei Wang, and Di He. Denoising masked autoencoders help robust classification. In *International Conference on Learning Representations*, 2023.

Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. Structured 3d latents for scalable and versatile 3d generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 21469–21480, 2025. doi: 10.1109/CVPR52734.2025.02000. URL https://doi.org/10.1109/CVPR52734.2025.02000.

Jiacong Yan and Hongwei Lin. Reasonable thickness determination for implicit porous sheet structure using persistent homology. *Computers & Graphics*, 115:236–245, 2023. doi: 10.1016/j.cag.2023.07.022. URL https://doi.org/10.1016/j.cag.2023.07.022.

Ziwei Yan, Jian Zhang, Ruizhi Zhang, Zhiqiang Liu, Guoqiang Luo, Zhonghui Shen, and Qiang Shen. Machine learning-driven optimization for multi-performance trade-offs in triply periodic minimal surface lattice metamaterials. *Thin-Walled Structures*, 218:114040, 2026. doi: 10.1016/j.tws.2025.114040. URL https://doi.org/10.1016/j.tws.2025.114040.

Ting Yang, Hongshun Chen, Zian Jia, Zhifei Deng, Liuni Chen, Emily M. Peterman, James C. Weaver, and Ling Li. A damage-tolerant, dual-scale, single-crystalline microlattice in the knobby starfish, protoreaster nodosus. *Science*, 375(6581):647–652, 2022. doi: 10.1126/science.abj9472. URL https://doi.org/10.1126/science.abj9472.

Yousef Yeganeh, Goktug Guvercin, Nassir Navab, and Azade Farshad. Conformable convolution for topologically constrained learning of complex anatomical structures. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 11982–11990, 2026. doi: 10.1609/aaai.v40i14.38186. URL https://ojs.aaai.org/index.php/AAAI/article/view/38186.

Wangzhi Zhan, Jianpeng Chen, Dongqi Fu, and Dawei Zhou. UniMate: A unified model for mechanical metamaterial generation, property prediction, and condition confirmation. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 74406–74421. PMLR, 2025. URL https://proceedings.mlr.press/v267/zhan25b.html.

Yu Zhang, Shaoyu Zhao, Yingyan Zhang, Chaofeng Lü, and Jie Yang. Multiscale thermo-electro-mechanical performance and stability of tpms piezoelectric metamaterial shells. *International Journal of Mechanical Sciences*, 322:111665, 2026. doi: 10.1016/j.ijmecsci.2026.111665.

Li Zheng, Konstantinos Karapiperis, Siddhant Kumar, and Dennis M. Kochmann. Unifying the design space and optimizing linear and nonlinear truss metamaterials by generative modeling. *Nature Communications*, 14:7563, 2023. doi: 10.1038/s41467-023-42068-x. URL https://www.nature.com/articles/s41467-023-42068-x.

Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005.

A Appendix

A.1 Topological Sensitivity Analysis

To further examine the relationship between localized structural perturbations, persistent topology, and mechanical response, we construct a controlled experiment consisting of one baseline TPMS structure and four locally modified variants. Because the modifications are spatially confined while the overall geometry remains largely unchanged, these examples provide a controlled setting for studying how local connectivity changes are reflected in persistent-homology diagrams, elastic mechanical responses, and the persistence-guided masking heatmaps.

Figure S1 compares four distinct localized modifications derived from the same baseline geometry. *Cut* induces a crack to an existing connectivity pathway; *Single hole* introduces one additional local opening; *Double separate* introduces two nearby but disconnected openings; and *Merged* connects the neighboring openings into a larger merged feature. These perturbations remain spatially localized, enabling comparison of structures with relatively small geometric changes but different alterations to local connectivity.

The resulting mechanical changes are not explained by the amount of removed material alone. The Cut and Single-hole modifications, for example, exhibit similar changes in solid volume fraction, -0.421% and -0.459%, respectively, but substantially different elastic responses. The Cut structure produces reductions of 4.42%, 6.28%, and 9.75% in E_x, E_y, and E_z, respectively, together with reductions of 10.01%, 8.95%, and 2.06% in the corresponding shear moduli. In contrast, the

The Single-hole modification leads to somewhat smaller, but still appreciable, mechanical changes relative to its small reduction in material volume, with E_x and E_y decreasing by 2.81% and 1.13%, respectively, while E_z remains nearly unchanged at $+0.12\%$. Thus, even for comparable and sub-percent changes in solid volume fraction, the mechanical response can vary noticeably depending on where and how the local structural connectivity is modified.

The Double-separate and Merged modifications produce larger changes in solid volume fraction, -1.079% and -1.240%, respectively, and both exhibit larger reductions along selected directions. For Double separate, E_x decreases by 6.46% and the largest shear-modulus reduction reaches 5.60%. The Merged modification shows reductions of 6.63% in E_x and 6.85% in G_{xy}. Although the two cases have comparable overall volume changes and visually similar local modifications, their directional elastic responses are not identical, indicating that the spatial organization and connectivity of the perturbation influence the resulting anisotropic stiffness.

These mechanical changes are accompanied by corresponding differences in the persistence diagrams (explained in Appendix A.4) in Fig. S1. In particular, the localized modifications alter persistent H_1 and H_2 features associated with tunnels and voids, respectively. The comparison therefore shows that relatively small geometric perturbations can simultaneously modify persistent-topological descriptors and directional elastic behavior, motivating the use of persistent homology as an additional structural signal during representation learning.

Figure S2 further examines how the persistence-guided masking signal (key mechanism of the proposed model, described in Sec. 2.2) responds to these perturbations. For the Cut modification, the PD-score distribution changes most noticeably around the interrupted region, whereas the Single-hole and Double-separate modifications exhibit score variations around the newly introduced openings. These results indicate that the masking distribution is sensitive to localized regions in which structural connectivity is altered.

The PD-score maps of the Double-separate and Merged modifications remain relatively similar despite distinguishable differences in their persistence diagrams in Fig. S1. This observation suggests that the spatial masking signal alone may not capture all persistent-topological differences, particularly when structures exhibit similar local importance patterns but differ in their global PH descriptors. The auxiliary PH prediction head therefore provides a complementary learning signal by encouraging PH-derived information to remain decodable from the latent representation. Together, these observations motivate the coordinated use of persistence-guided masking and PH-based latent supervision in the proposed framework.

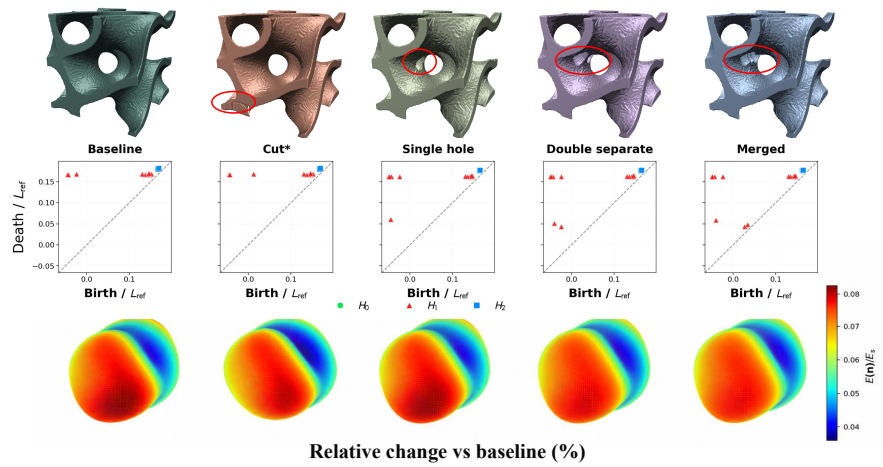

Relative change vs baseline (%)

Case	ΔV	ΔE_x	ΔE_y	ΔE_z	ΔG_1	ΔG_2	ΔG_3
Cut	-0.421%	-4.42%	-6.28%	-9.75%	-10.01%	-8.95%	-2.06%
Single hole	-0.459%	-2.81%	-1.13%	+0.12%	-0.71%	-1.09%	-2.65%
Double separate	-1.079%	-6.46%	-1.89%	-0.13%	-2.23%	-2.73%	-5.60%
Merged	-1.240%	-6.63%	-2.19%	-0.41%	-4.31%	-3.13%	-6.85%

V: Solid volume fraction of the unit cell.
E_x, E_y, E_z: Principal Young's moduli along the x, y, and z directions.
G_1, G_2, G_3: Shear moduli in the three principal shear planes (yz, xz, xy).
$\Delta(\cdot)$: Relative change with respect to the baseline,

$$\Delta(\cdot) = \frac{value - value_{baseline}}{value_{baseline}} \times 100\%.$$

−10% 0% +10%
Relative change

Figure S1: **Topological perturbations and their mechanical responses.** Five geometries are compared: the baseline structure and four locally modified variants (Cut, Single hole, Double separate, and Merged). The top row shows the corresponding geometries and localized modifications, the second row presents their persistence diagrams for H_0, H_1, and H_2, and the third row shows the directional Young's modulus surfaces. The table summarizes the relative changes in solid volume fraction and the six principal elastic moduli with respect to the baseline.

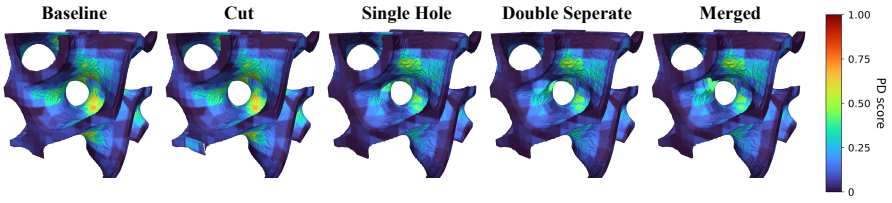

Figure S2: **Persistence masking-guided PD-score distributions under localized topological perturbations.** Normalized PD scores are projected onto the baseline structure and four locally modified variants. Warmer colors indicate regions with larger PD scores and therefore higher probabilities of being selected by the persistence-guided masking strategy.

A.2 Background on TPMS Shell Metamaterials

To provide additional context for the structural systems considered in this work, we first briefly summarize the geometric origin, mechanical characteristics, and representative applications of TPMS shell metamaterials. Figure S3 illustrates that TPMS geometries originate from periodic minimal surfaces that partition space into interpenetrating domains. By introducing a finite thickness around

Table S1: **Persistent-homology feature counts for the controlled topological perturbations.** The table reports the occupied voxel count, voxel occupancy, and the numbers of finite persistence features in H_0, H_1, and H_2. Values in parentheses denote the change relative to the baseline. All structures contain one infinite H_0 feature and no infinite H_1 or H_2 features.

Structure	Occupied voxels	Occupancy (%)	H_0	H_1	H_2
Baseline	162,042	18.315	373	2190	1150
Cut	161,520	18.256	370 (−3)	2194 (+4)	1153 (+3)
Single hole	181,533	20.518	503 (+130)	2012 (−178)	1111 (−39)
Double separate	180,767	20.432	498 (+125)	1982 (−208)	1111 (−39)
Merged	180,465	20.398	498 (+125)	1987 (−203)	1115 (−35)

the minimal surface, these geometries can be converted into shell-based metamaterials with smooth and highly interconnected architectures. Such structural characteristics provide favorable stiffness-to-weight performance while maintaining large surface-area-to-volume ratios and continuous transport pathways. These properties have motivated applications ranging from lightweight load-bearing structures and energy absorption to thermal management, fluid transport, and catalytic systems.

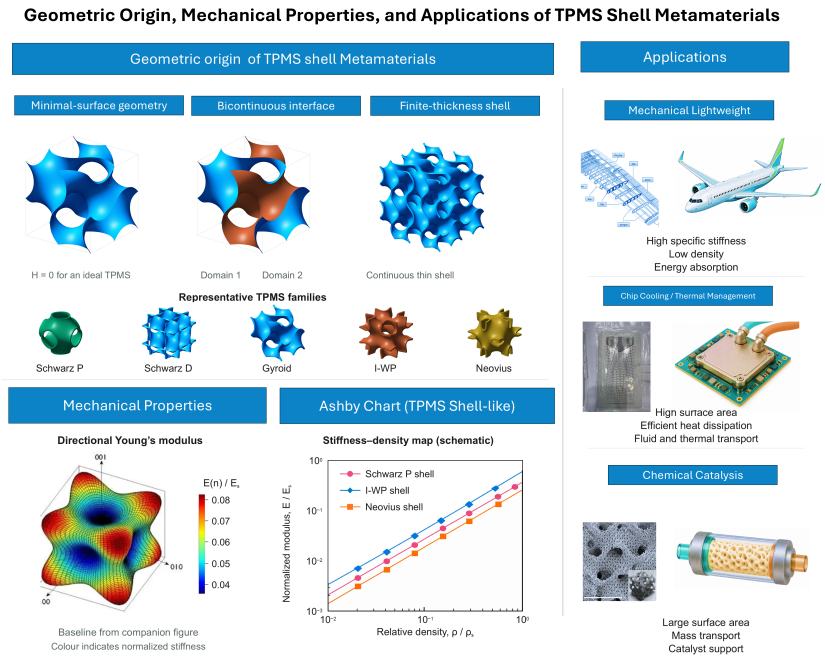

Figure S3: Overview of the geometric origin, mechanical characteristics, and representative applications of TPMS shell metamaterials. TPMS shell architectures originate from smooth periodic minimal surfaces and exhibit interconnected geometries that enable favorable mechanical performance and multifunctional applications, including lightweight structures, thermal management, and catalysis(Wang et al., 2025; 2022; fly123, 2014; Basavalingappa et al., 2017).

A.3 Model Architecture

Our model adopts a 3D adaptation of DMAE backbone, which is built upon a Vision Transformer (ViT) encoder–decoder architecture. Given a single-channel TPMS volume $X \in \mathbb{R}^{1\times64\times64\times64}$, the input is partitioned into non-overlapping $4 \times 4 \times 4$ volumetric patches, resulting in a $16 \times 16 \times 16$ patch grid with $N = 4096$ tokens. Patch embedding is implemented using a 3D convolution with kernel size and stride of 4, and each patch is projected to a $D = 192$-dimensional embedding. A

fixed 3D sinusoidal positional embedding is subsequently added to preserve the spatial location of each volumetric patch.

Following the DMAE/MAE encoding paradigm, masked patches are hiden before the encoder, and only visible patch tokens are processed by the ViT backbone. The encoder consists of six Transformer blocks with an embedding dimension of 192 and six self-attention heads; the feed-forward network in each block expands the feature dimension from 192 to 768 before projecting it back to 192. Each Transformer block uses GELU activation in the feed-forward network, followed by a final LayerNorm after the encoder. For a masking ratio r, the encoder therefore operates on $(1-r)N$ visible tokens and produces:

$$Z_{\text{vis}} \in \mathbb{R}^{(1-r)N \times 192}.$$

The reconstruction branch follows an asymmetric Transformer decoder. The encoded visible tokens are first projected from 192 to a decoder dimension of 128. Learnable mask tokens are inserted at the removed patch locations, and the complete sequence is restored to its original spatial ordering. After adding 3D sinusoidal positional embeddings, the sequence is processed by 4 Transformer decoder blocks with 4 attention heads and an MLP dimension of 512. A final linear prediction layer maps each decoder token to $4^3 = 64$ scalar values, corresponding to one volumetric patch, thereby reconstructing the original $64 \times 64 \times 64$ input volume.

In our topology-aware extension, the DMAE backbone is augmented with an auxiliary persistent-homology prediction head attached to the encoder representation. The shared ViT encoder is jointly optimized by the masked reconstruction and topology-prediction objectives,

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{rec}} + \lambda_{\text{PD}} \mathcal{L}_{\text{topo}}.$$

For downstream representation evaluation, the full volume without masking or noise corruption is processed by the pretrained ViT encoder. The resulting patch-token features are globally averaged to obtain a 192-dimensional structure-level representation, which is subsequently L_2-normalized before downstream evaluation.

A.4 Persistent homology for TPMS shell

We use PH to characterize the multiscale global topology of three-dimensional TPMS structures. Unlike local geometric descriptors, such as surface position, thickness, and curvature, PH captures how topological features—including connected components, tunnels, and voids—emerge and disappear along a nested sequence of spatial thresholds. PDs summarize these features by encoding their birth, death, and persistence across filtration scales.

From signed-distance fields to cubical complexes. Each TPMS structure is represented by an SDF tensor $\mathbf{S} \in \mathbb{R}^{N \times N \times N}$ sampled on a regular three-dimensional domain, where each entry $S_{i,j,k}$ represents the signed distance from the voxel at location (i, j, k) to the TPMS solid–void interface. Negative values correspond to solid regions, whereas positive values correspond to void regions. In our experiments, all SDFs are sampled at a resolution of $N = 64$.

To enable persistent-homology analysis, we construct a cubical complex $\mathbb{K}$ from the regular SDF grid using GUDHI (Dlotko, 2026). The complex is composed of cubical cells with dimensions $d \in \{0, 1, 2, 3\}$. The three-dimensional cells $c^{(3)}_{i,j,k}$ correspond to cubes, while the lower-dimensional cells $c^{(2)}$, $c^{(1)}$, and $c^{(0)}$ denote faces, edges, and vertices, respectively. Each $c^{(3)}_{i,j,k} \in \mathbb{K}$ is associated with the SDF entry $S_{i,j,k}$.

We define the filtration field by taking the absolute value of the SDF,

$$F_{i,j,k} = |S_{i,j,k}|. \tag{4}$$

Here, $F_{i,j,k}$ denotes the unsigned distance from the voxel to the solid–void interface. We define the filtration function $f : \mathbb{K} \to \mathbb{R}$ to assign a filtration value to each cell in $\mathbb{K}$, with $f(c^{(3)}_{i,j,k}) = F_{i,j,k}$ for each cube $c^{(3)}_{i,j,k}$. For a lower-dimensional cell $c^{(d)}$, where $d \in \{0, 1, 2\}$, its filtration value is

set to the minimum value among all incident three-dimensional cubes containing it, i.e., $f(c^{(d)}) = \min_{c^{(3)}_{i,j,k} \supseteq c^{(d)}} f(c^{(3)}_{i,j,k})$. Such an assignment ensures closure under inclusion, so whenever a cube enters the filtration, all of its boundary cells are already present.

For a threshold $t \geq 0$, the sublevel complex $\mathbb{K}_t$ contains all cells $c^{(d)} \in \mathbb{K}$ satisfying $f(c^{(d)}) \leq t$, where $d \in \{0, 1, 2, 3\}$. As t increases, cells with larger filtration values are progressively included, corresponding to regions farther from the solid–void interface. The resulting sublevel complexes form a nested filtration,

$$\mathbb{K}_{t_1} \subseteq \mathbb{K}_{t_2} \subseteq \cdots \subseteq \mathbb{K}_{t_m}, \qquad 0 \leq t_1 < t_2 < \cdots < t_m. \tag{5}$$

At each filtration level t, let $C_q(\mathbb{K}_t; \mathbb{Z}_2)$ denote the q-dimensional chain group over the finite field $\mathbb{Z}_2 = \{0, 1\}$, where arithmetic is performed modulo 2. The group $C_q(\mathbb{K}_t; \mathbb{Z}_2)$ consists of linear combinations of the q-dimensional cells in $\mathbb{K}_t$, with $q \in \{0, 1, 2, 3\}$. The boundary operator ∂_q maps each q-dimensional chain to its $(q-1)$-dimensional boundary,

$$\partial_q : C_q(\mathbb{K}_t; \mathbb{Z}_2) \to C_{q-1}(\mathbb{K}_t; \mathbb{Z}_2). \tag{6}$$

Here, ∂_q encodes the incidence relations between cells of adjacent dimensions; for example, the boundary of a face consists of its edges, whereas the boundary of a cube consists of its faces. A q-dimensional cycle is a chain with zero boundary, whereas a q-dimensional boundary is a chain that itself forms the boundary of a $(q+1)$-dimensional chain. The corresponding cycle and boundary groups are defined as

$$Z_q(\mathbb{K}_t) = \ker(\partial_q), \qquad B_q(\mathbb{K}_t) = \mathrm{im}(\partial_{q+1}), \tag{7}$$

where $\ker(\partial_q)$ contains the q-dimensional chains mapped to zero by ∂_q, and $\mathrm{im}(\partial_{q+1})$ contains the q-dimensional chains generated as boundaries of $(q+1)$-dimensional chains.

The q-th homology group is then defined as

$$H_q(\mathbb{K}_t; \mathbb{Z}_2) = Z_q(\mathbb{K}_t)/B_q(\mathbb{K}_t). \tag{8}$$

The quotient removes cycles that are boundaries of higher-dimensional chains, leaving independent homological features represented by $H_q(\mathbb{K}_t; \mathbb{Z}_2)$ (Moor et al., 2020; Bleile et al., 2022; Edelsbrunner et al., 2002).

Topological features. At each filtration level t, the number of independent homological features in dimension q is quantified by the Betti number

$$\beta_q(t) = \dim_{\mathbb{Z}_2} H_q(\mathbb{K}_t; \mathbb{Z}_2). \tag{9}$$

For the three-dimensional cubical complexes considered here, $\beta_0(t)$, $\beta_1(t)$, and $\beta_2(t)$ count connected components, independent tunnels, and voids present at filtration level t, respectively (Edelsbrunner et al., 2002).

Persistence diagrams. Persistent homology tracks the appearance and disappearance of these topological features as the filtration threshold increases. Although a three-dimensional cubical complex can in principle support homology dimensions up to $q = 3$, we consider only $q \in \{0, 1, 2\}$. Here, H_0, H_1, and H_2 characterize connected components, tunnels, and voids, respectively, whereas H_3 corresponds to closed three-dimensional cycles without boundary, which do not arise in the finite TPMS domains considered here.

A feature in homology dimension $q \in \{0, 1, 2\}$ is represented by a birth–death pair $(b_i^{(q)}, d_i^{(q)})$, where $b_i^{(q)}$ and $d_i^{(q)}$ denote the filtration thresholds at which the i-th feature appears and disappears, respectively. Its persistence is defined as

$$p_i^{(q)} = d_i^{(q)} - b_i^{(q)}. \tag{10}$$

Only finite persistence intervals with persistence exceeding the prescribed threshold $\epsilon = 0.005$ are retained. For each homology dimension, the retained birth–death pairs form the corresponding persistence diagram,

$$\mathrm{PD}_q = \{(b_i^{(q)}, d_i^{(q)})\}_{i=1}^{n_q}, \qquad q \in \{0, 1, 2\}, \tag{11}$$

where n_q denotes the number of retained topological features in homology dimension q. Accordingly, PD_0, PD_1, and PD_2 summarize the multiscale persistence of connected components, tunnels, and voids, respectively. Features with larger persistence remain present over a broader range of filtration thresholds and therefore correspond to more persistent topological structures (Zomorodian & Carlsson, 2005; Cohen-Steiner et al., 2007).

A.5 TOPOLOGICAL DIVERSITY OF THE DATASET

To illustrate the topological diversity of the dataset, we selected 20 representative TPMS samples spanning a broad range of persistence-diagram characteristics. The samples were chosen to exhibit distinct combinations of topological feature counts in H_0, H_1, and H_2, together with substantial variation in feature persistence and birth–death distributions. This selection therefore covers diverse connected-component, tunnel, and cavity structures, providing a representative set of geometries with clearly differentiated persistent-homology signatures.

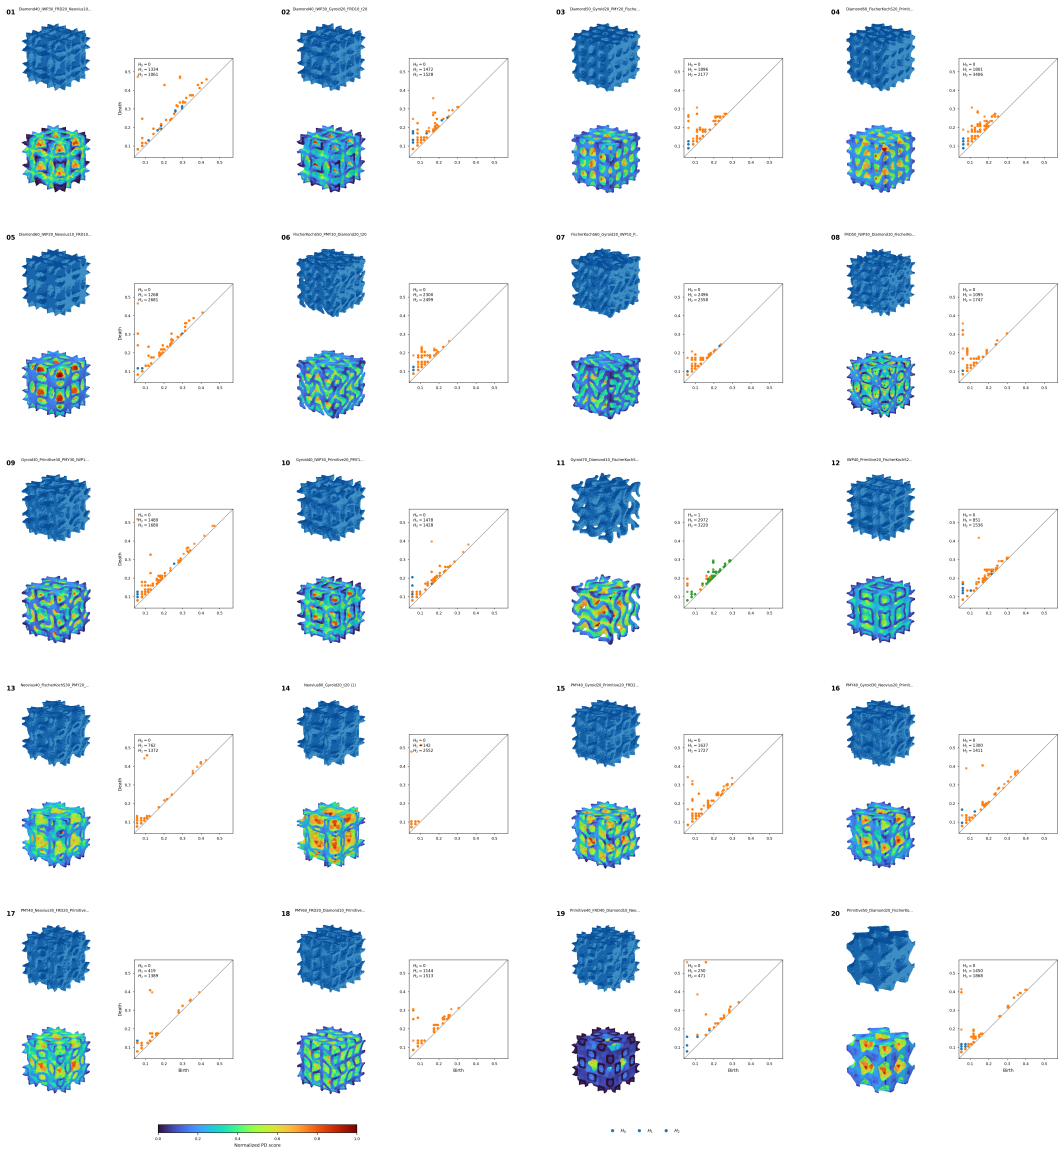

Figure S4: **Representative TPMS structures with diverse persistent-homology signatures.** Twenty TPMS samples were selected to span distinct persistence-diagram characteristics. Each voxelized geometry is shown alongside its corresponding persistence diagram, with the selected set exhibiting different combinations of H_0, H_1, and H_2 feature counts as well as substantial variation in their birth–death distributions. This selection provides a representative visualization of the broad topological diversity present in the TPMS dataset.

A.6 FINITE-ELEMENT HOMOGENIZATION

The effective elastic properties of each TPMS structure were evaluated using a voxel-based finite-element homogenization method under small-strain linear elasticity. Each geometry was discretized on a regular voxel grid and treated as a periodic unit cell. Periodic boundary conditions were imposed through periodic node mapping, and six independent macroscopic unit-strain states, corresponding to the Voigt components 11, 22, 33, 23, 31, and 12, were applied to determine the homogenized elastic response. The resulting stress responses were assembled into the effective 6×6 stiffness matrix $\mathbf{C}$, which contains 21 independent components due to elastic symmetry. The same voxel discretization, constituent-material properties, and homogenization settings were used for the ground-truth, reconstructed, and generated structures, ensuring that differences in the computed mechanical response arose from changes in geometry rather than from the numerical setup. Representative TPMS geometries and their corresponding homogenized stiffness matrices $\mathbf{C}$ are shown below to illustrate the mechanical responses obtained from the finite-element homogenization procedure.

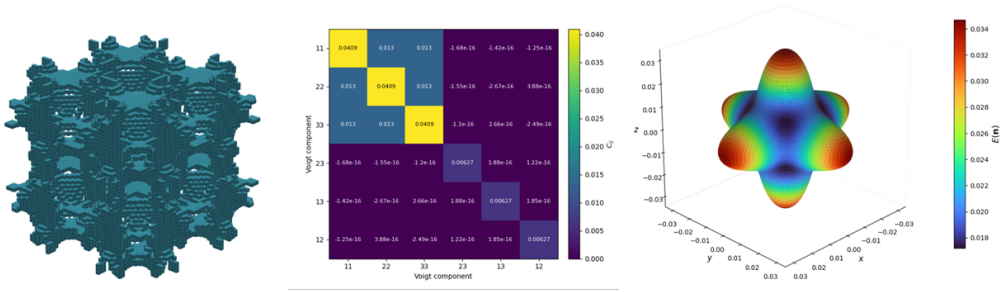

Figure S5: **Representative finite-element homogenization example.** A representative TPMS structure is shown together with its homogenized elastic response obtained from the voxel-based finite-element analysis. The corresponding 6×6 effective stiffness matrix $\mathbf{C}$ summarizes the anisotropic elastic properties of the structure, while the three-dimensional directional Young's modulus surface $E(\mathbf{n})$, derived from the compliance matrix $\mathbf{S} = \mathbf{C}^{-1}$, visualizes the variation in effective stiffness with loading direction.

A.7 DOWNSTREAM TASKS

A.7.1 TASK 2: C_{11} ELASTIC-PROPERTY CLASSIFICATION

We evaluate whether the frozen representations separate structures according to their mechanical-property levels by formulating C_{11} prediction as a three-class classification problem. The continuous C_{11} values are divided into *Low*, *Medium*, and *High* groups based only on the training-set statistics. Two lightweight probes are applied to the frozen latent representations: a k-nearest-neighbor (kNN) classifier with $k = 15$ and a linear logistic-regression classifier. No encoder parameters are updated during downstream evaluation.

For both probes, we report balanced accuracy and macro-F1. For compactness, balanced accuracy is denoted by Acc in Table S2. For C classes, balanced accuracy is defined as

$$\mathrm{Acc} = \frac{1}{C} \sum_{c=1}^{C} \frac{\mathrm{TP}_c}{\mathrm{TP}_c + \mathrm{FN}_c}, \tag{12}$$

where TP_c and FN_c denote the numbers of true-positive and false-negative samples for class c, respectively. Thus, Acc corresponds to the mean recall across classes and prevents a dominant property group from disproportionately affecting the reported performance.

For each class c, precision and recall are defined as

Table S2: **Task 2 – Low/medium/high C_{11} elastic-property classification.** Acc denotes balanced accuracy and F denotes macro-F1. The subscripts kNN and Lin correspond to the kNN and linear probes, respectively. Higher is better for all metrics.

MODEL	METADB				INTERPOLATION TPMS			
	$\mathrm{Acc}_{k\mathrm{NN}}\uparrow$	$\mathrm{F}_{k\mathrm{NN}}\uparrow$	$\mathrm{Acc}_{\mathrm{Lin}}\uparrow$	$\mathrm{F}_{\mathrm{Lin}}\uparrow$	$\mathrm{Acc}_{k\mathrm{NN}}\uparrow$	$\mathrm{F}_{k\mathrm{NN}}\uparrow$	$\mathrm{Acc}_{\mathrm{Lin}}\uparrow$	$\mathrm{F}_{\mathrm{Lin}}\uparrow$
General representation learning								
MVAE	0.509	0.509	0.631	0.631	0.719	0.717	0.512	0.511
SVAE	0.766	0.765	0.793	0.793	0.863	0.863	0.698	0.697
VICReg	0.853	0.853	0.913	0.913	0.799	0.801	0.835	0.837
I-JEPA	0.902	0.902	0.911	0.912	0.831	0.833	0.863	0.864
Masking strategy baselines								
MAE	0.907	0.907	0.901	0.901	0.883	0.883	0.896	0.896
DMAE	0.893	0.894	0.907	0.907	0.870	0.871	0.923	0.923
ColorMAE	0.952	0.952	0.957	0.956	0.852	0.853	0.918	0.918
DaM	0.948	0.948	0.952	0.952	0.857	0.858	0.912	0.912
Topology-informed baseline								
ScalableTopo	**0.954**	**0.954**	0.960	0.957	0.856	0.856	0.916	0.916
OURS	0.951	0.951	**0.964**	**0.964**	**0.887**	**0.887**	**0.924**	**0.925**

$$P_c = \frac{\mathrm{TP}_c}{\mathrm{TP}_c + \mathrm{FP}_c}, \qquad R_c = \frac{\mathrm{TP}_c}{\mathrm{TP}_c + \mathrm{FN}_c}, \tag{13}$$

where FP_c denotes the number of false-positive samples. The class-wise F1 score is

$$F_{1,c} = \frac{2P_cR_c}{P_c + R_c}, \tag{14}$$

and the macro-F1 score is

$$\mathrm{F} = \frac{1}{C}\sum_{c=1}^{C} F_{1,c}. \tag{15}$$

Macro-F1 assigns equal importance to all property groups and therefore jointly reflects class-wise precision and recall. The subscripts kNN and Lin distinguish the kNN and linear probes, respectively.

As reported in Table S2, our method achieves the best performance across all four classification metrics on the Interpolation TPMS dataset, reaching $\mathrm{F}k\mathrm{NN} = 0.887$ and $\mathrm{FLin} = 0.925$. On MetaDB, it achieves the highest linear-probe performance, with $\mathrm{F}_{\mathrm{Lin}} = 0.964$, while remaining within 0.003 of the best kNN results. These results indicate that the learned representations retain mechanically relevant information that can be effectively predicted by simple downstream probes.

A.7.2 TASK 3: C_{11} PROPERTY CLUSTERING

We next examine whether mechanically similar structures are naturally organized together in the frozen latent space without using C_{11} labels to train the representation. The resulting latent clusters are evaluated from two complementary perspectives: their ability to explain continuous variations in C_{11} and their agreement with the discrete *Low*/*Medium*/*High* property groups.

We first report the correlation ratio η^2, which measures the fraction of the total C_{11} variation explained by the latent clustering. Let y_i denote the continuous C_{11} value of sample i, $\bar{y}$ denote the global mean, and let cluster k contain n_k samples with mean property value $\bar{y}_k$. The correlation ratio is defined as

$$\eta^2 = \frac{\sum_{k=1}^{K} n_k \left(\bar{y}_k - \bar{y}\right)^2}{\sum_{i=1}^{N} \left(y_i - \bar{y}\right)^2}, \tag{16}$$

where N is the total number of samples and K is the number of clusters. The numerator measures the between-cluster variation in C_{11}, whereas the denominator measures the total variation across the dataset. Therefore, a larger η^2 indicates that samples assigned to different latent clusters exhibit more distinct mechanical-property distributions.

Because cluster indices are arbitrary, direct comparison between predicted cluster IDs and reference property classes is not meaningful. We therefore determine an optimal one-to-one mapping π between cluster indices and reference classes and compute the matched clustering accuracy as

$$\mathrm{Acc}_{\mathrm{Match}} = \max_{\pi} \frac{1}{N} \sum_{i=1}^{N} \mathbb{I}\left[\pi(z_i) = y_i^{\mathrm{cls}}\right], \tag{17}$$

where z_i is the predicted cluster assignment, y_i^{cls} is the reference *Low/Medium/High* property label, π denotes the cluster-to-class assignment, and $\mathbb{I}[\cdot]$ is the indicator function. The optimal mapping removes the arbitrary permutation of cluster identities, so $\mathrm{Acc}_{\mathrm{Match}}$ reflects the fraction of samples whose latent clusters agree with the corresponding property groups.

After applying the same optimal cluster-to-class mapping, we additionally compute the macro-F1 score,

$$\mathrm{F}_{\mathrm{Macro}} = \frac{1}{C} \sum_{c=1}^{C} \frac{2P_c R_c}{P_c + R_c}, \tag{18}$$

where P_c and R_c denote precision and recall for property class c. Unlike matched accuracy, $\mathrm{F}_{\mathrm{Macro}}$ weights all property classes equally and is therefore more sensitive to whether the clustering quality is consistent across different C_{11} ranges.

As shown in Table S3, our method achieves the best clustering performance across all three metrics on both datasets. On MetaDB, it reaches $\eta^2 = 0.323$, $\mathrm{Acc}_{\mathrm{Match}} = 0.660$, and $\mathrm{F}_{\mathrm{Macro}} = 0.673$, outperforming the strongest competing baseline for each metric. On the Interpolation TPMS dataset, it similarly achieves the highest $\eta^2 = 0.412$, $\mathrm{Acc}_{\mathrm{Match}} = 0.582$, and $\mathrm{F}_{\mathrm{Macro}} = 0.576$, indicating that the learned latent space is more strongly organized according to continuous C_{11} variation and the corresponding property groups.

A.7.3 Task 4: C_{11} Property Regression

Finally, we evaluate whether the continuous mechanical-property value can be quantitatively recovered from the frozen representation. A regression probe is trained on the latent features to predict C_{11}, while the representation encoder remains fixed. We evaluate the predictions using mean absolute error (MAE), root mean squared error (RMSE), and the coefficient of determination (R^2).

Let y_i and $\hat{y}_i$ denote the ground-truth and predicted C_{11} values of sample i, respectively. The mean absolute error is defined as

$$\mathrm{MAE} = \frac{1}{N} \sum_{i=1}^{N} \left|y_i - \hat{y}_i\right|, \tag{19}$$

where N is the number of evaluated samples. MAE measures the average absolute prediction error in the same units as C_{11} and weights all errors linearly.

The root mean squared error is defined as

Table S3: **Task 3 – C_{11} property clustering.** η^2 measures the fraction of continuous C_{11} variation explained by the latent clusters. $\mathrm{Acc}_{\mathrm{Match}}$ and $\mathrm{F}_{\mathrm{Macro}}$ evaluate agreement with the discrete property groups after optimal cluster-to-class matching. Higher is better for all metrics.

MODEL	METADB			INTERPOLATION TPMS		
	$\eta^2 \uparrow$	$\mathrm{Acc}_{\mathrm{Match}} \uparrow$	$\mathrm{F}_{\mathrm{Macro}} \uparrow$	$\eta^2 \uparrow$	$\mathrm{Acc}_{\mathrm{Match}} \uparrow$	$\mathrm{F}_{\mathrm{Macro}} \uparrow$
General representation learning						
MVAE	0.036	0.044	0.041	0.014	0.372	0.367
SVAE	0.127	0.519	0.436	0.049	0.399	0.394
VICReg	0.068	0.457	0.457	0.020	0.382	0.377
I-JEPA	0.014	0.393	0.291	0.081	0.370	0.366
Masking strategy baselines						
MAE	0.104	0.433	0.401	0.258	0.391	0.380
DMAE	0.147	0.504	0.436	0.272	0.410	0.383
ColorMAE	0.255	0.630	0.603	0.162	0.413	0.393
DaM	0.196	0.514	0.497	0.177	0.446	0.414
Topology-informed baseline						
ScalableTopo	0.199	0.543	0.533	0.158	0.422	0.395
OURS	**0.323**	**0.660**	**0.673**	**0.412**	**0.582**	**0.576**

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}. \tag{20}$$

Because the prediction errors are squared before averaging, RMSE places a larger penalty on samples with large regression errors and is therefore more sensitive to outliers than MAE.

We further report the coefficient of determination,

$$R^2 = 1 - \frac{\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{N}(y_i - \bar{y})^2}, \tag{21}$$

where $\bar{y}$ denotes the mean ground-truth C_{11} value. The numerator represents the residual sum of squares of the regression predictions, while the denominator represents the total variation in the ground-truth properties. An R^2 value closer to 1 indicates that a larger fraction of the variation in C_{11} can be recovered from the frozen latent representation.

As shown in Table S4, our method achieves competitive regression performance on MetaDB, with $\mathrm{MAE} = 0.009$, $\mathrm{RMSE} = 0.047$, and $R^2 = 0.899$. On the Interpolation TPMS dataset, it achieves the lowest MAE (0.003) and RMSE (0.005), while matching the highest R^2 of 0.993. These results show that the learned representation preserves fine-grained continuous mechanical-property information in addition to the discrete and clustering-level organization evaluated in Tasks 2 and 3.

A.8 CLUSTERING METRICS IN ABLATION STUDY

In addition to η^2, we evaluate the alignment between the unsupervised K-means clusters and the discretized C_{11} property groups using the Adjusted Rand Index (ARI) and Normalized Mutual Information (NMI).

Let n_{ij} denote the number of samples shared by predicted cluster i and property class j, with row sums $a_i = \sum_j n_{ij}$, column sums $b_j = \sum_i n_{ij}$, and total sample size n. The ARI is defined as

Table S4: **Task 4 – C_{11} property regression.** MAE measures the average absolute prediction error, RMSE places greater emphasis on large errors, and R^2 measures the fraction of C_{11} variation explained by the regression prediction. Lower is better for MAE and RMSE, while higher is better for R^2.

MODEL	METADB			INTERPOLATION TPMS		
	MAE ↓	RMSE ↓	R^2 ↑	MAE ↓	RMSE ↓	R^2 ↑
General representation learning						
MVAE	0.049	0.078	0.246	0.039	0.050	0.110
SVAE	0.028	0.051	0.471	0.024	0.031	0.616
VICReg	0.011	0.049	0.903	0.010	0.014	0.935
I-JEPA	0.010	0.048	0.907	0.007	0.010	0.965
Masking strategy baselines						
MAE	0.010	0.047	0.909	0.005	0.007	0.984
DMAE	0.009	0.050	0.899	**0.003**	**0.005**	**0.993**
ColorMAE	**0.007**	0.046	0.915	**0.003**	**0.005**	0.992
DaM	0.010	0.057	0.869	0.004	**0.005**	0.991
Topology-informed baseline						
ScalableTopo	0.008	**0.041**	**0.932**	**0.003**	**0.005**	0.992
OURS	0.009	0.047	0.899	**0.003**	**0.005**	**0.993**

Table S5: **Ablation study on Interpolation TPMS and MetaDB.** Results are reported as mean ± standard deviation over five random seeds. For Interpolation TPMS, Task 1 reports strict and relaxed kNN family-classification accuracy; for MetaDB, Task 1 reports macro-F1 from the linear and kNN probes. Task 2 reports $F_{k\mathrm{NN}}$ for C_{11} property classification; Task 3 reports η^2, ARI, and NMI for C_{11} property clustering; and Task 4 reports R^2 for C_{11} and the full stiffness tensor C_{21}. Higher is better for all metrics.

Interpolation TPMS

	Task 1		Task 2	Task 3			Task 4	
Ablation	$\mathrm{Acc}^{\mathrm{Str}}_{k\mathrm{NN}}$ ↑	$\mathrm{Acc}^{\mathrm{Rel}}_{k\mathrm{NN}}$ ↑	$F_{k\mathrm{NN}}$ ↑	η^2 ↑	ARI ↑	NMI ↑	$R^2_{C_{11}}$ ↑	$R^2_{C_{21}}$ ↑
w/o mask & head	0.8191 ± 0.0126	0.9324 ± 0.0061	0.8710 ± 0.0068	0.2721 ± 0.0060	0.138 ± 0.019	0.124 ± 0.020	$\mathbf{0.9925 \pm 0.0003}$	0.9901 ± 0.0002
w/o mask	0.8218 ± 0.0077	0.9358 ± 0.0069	0.8845 ± 0.0062	0.4045 ± 0.0317	0.196 ± 0.041	0.191 ± 0.046	0.9925 ± 0.0005	$\mathbf{0.9964 \pm 0.0003}$
w/o head	0.8232 ± 0.0155	0.9365 ± 0.0115	0.8682 ± 0.0062	0.2508 ± 0.0328	0.116 ± 0.040	0.116 ± 0.030	0.9917 ± 0.0002	0.9961 ± 0.0003
Full	$\mathbf{0.8309 \pm 0.0074}$	$\mathbf{0.9391 \pm 0.0073}$	$\mathbf{0.8870 \pm 0.0049}$	$\mathbf{0.4116 \pm 0.0363}$	$\mathbf{0.203 \pm 0.015}$	$\mathbf{0.213 \pm 0.013}$	0.9923 ± 0.0006	0.9962 ± 0.0003

MetaDB

	Task 1		Task 2	Task 3			Task 4	
Ablation	F_{Lin} ↑	$F_{k\mathrm{NN}}$ ↑	$F_{k\mathrm{NN}}$ ↑	η^2 ↑	ARI ↑	NMI ↑	$R^2_{C_{11}}$ ↑	$R^2_{C_{21}}$ ↑
w/o mask & head	0.9070 ± 0.0016	0.8940 ± 0.0015	0.8940 ± 0.0086	0.1470 ± 0.0050	0.130 ± 0.029	0.166 ± 0.034	0.8990 ± 0.0130	0.8640 ± 0.0150
w/o mask	0.9872 ± 0.0019	0.9720 ± 0.0021	0.9314 ± 0.0065	0.3146 ± 0.0088	0.202 ± 0.052	0.324 ± 0.047	0.9066 ± 0.0181	0.9061 ± 0.0149
w/o head	0.9860 ± 0.0040	$\mathbf{0.9930 \pm 0.0010}$	0.9460 ± 0.0070	0.3150 ± 0.0110	0.210 ± 0.057	0.305 ± 0.068	$\mathbf{0.9370 \pm 0.0200}$	0.9350 ± 0.0100
Full	$\mathbf{0.9944 \pm 0.0014}$	0.9845 ± 0.0014	$\mathbf{0.9513 \pm 0.0040}$	$\mathbf{0.3228 \pm 0.0130}$	$\mathbf{0.211 \pm 0.046}$	$\mathbf{0.331 \pm 0.047}$	0.8991 ± 0.0191	$\mathbf{0.9360 \pm 0.0178}$

$$\mathrm{ARI} = \frac{\displaystyle\sum_{ij}\binom{n_{ij}}{2} - \frac{\displaystyle\sum_i \binom{a_i}{2}\sum_j \binom{b_j}{2}}{\binom{n}{2}}}{\displaystyle\frac{1}{2}\left[\sum_i \binom{a_i}{2} + \sum_j \binom{b_j}{2}\right] - \frac{\displaystyle\sum_i \binom{a_i}{2}\sum_j \binom{b_j}{2}}{\binom{n}{2}}}.$$

ARI corrects the agreement between the two partitions for chance, with larger values indicating stronger correspondence.

For predicted clusters U and property classes V, NMI is computed from their mutual information,

$$I(U;V) = \sum_{i,j} p_{ij} \log \frac{p_{ij}}{p_i p_j},$$

and normalized by the entropies of the two partitions:

$$\mathrm{NMI} = \frac{I(U;V)}{\sqrt{H(U)H(V)}},$$

where

$$H(U) = -\sum_i p_i \log p_i, \qquad H(V) = -\sum_j p_j \log p_j.$$

NMI ranges from 0 to 1, with higher values indicating stronger agreement between the clustering structure and the C_{11} property groups.

A.9 Latent-Space Alignment Metrics

We evaluate whether pairwise distances in the learned latent space reflect topology- and mechanics-related differences between TPMS structures. For each model, the latent vectors are standardized using statistics computed from the corresponding training split. For a held-out pair of structures (i, j), the latent distance is defined as

$$d_z(i,j) = \|\widetilde{\mathbf{z}}_i - \widetilde{\mathbf{z}}_j\|_2 ,$$

where $\widetilde{\mathbf{z}}$ denotes the standardized latent representation.

For each descriptor, we compute a corresponding pairwise distance $d_t(i, j)$ and quantify its alignment with the latent distance using the Spearman rank correlation

$$\rho_s = 1 - \frac{6\sum_{m=1}^{M}(R_m^z - R_m^t)^2}{M(M^2-1)},$$

where M is the number of evaluated structure pairs, and R_m^z and R_m^t denote the ranks of the latent distance and descriptor distance for pair m, respectively. Higher ρ_s indicates stronger correspondence between the learned latent geometry and the corresponding structural descriptor.

PD-Score Alignment. Let $s_{i,k}$ denote the PD score of patch k in structure i. Following the masking procedure, the scores are normalized within each structure as

$$\pi_{i,k} = \frac{s_{i,k}}{\sum_{r=1}^{n_p} s_{i,r}},$$

where n_p is the number of patches. The pairwise PD-score distance is

$$d_{\mathrm{PD}}(i,j) = \|\boldsymbol{\pi}_i - \boldsymbol{\pi}_j\|_2 .$$

PD-Score Alignment is then defined as

$$\rho_{\mathrm{PD}} = \rho_s \left(d_z, d_{\mathrm{PD}}\right).$$

H_1- and H_2-Count Alignment. Let $n_i^{(q)}$ denote the number of persistence features in homology dimension $q \in \{1, 2\}$ for structure i. The pairwise count distance is

$$d_{H_q}(i,j) = \left| n_i^{(q)} - n_j^{(q)} \right|.$$

The corresponding alignment is

$$\rho_{H_q} = \rho_s \left(d_z, d_{H_q}\right), \qquad q \in \{1,2\}.$$

C_{11}-Property Alignment. The scalar stiffness component C_{11} is standardized using statistics from the training split, yielding $\widetilde{C}_{11}$. The pairwise property distance is

$$d_{C_{11}}(i,j) = \left| \widetilde{C}_{11,i} - \widetilde{C}_{11,j} \right|.$$

The property alignment is defined as

$$\rho_{C_{11}} = \rho_s \left(d_z, d_{C_{11}}\right).$$

Table S6: **Parameter sensitivity analysis on the Interpolation TPMS dataset.** We evaluate the sensitivity to masking ratio ρ, PD-loss weight λ_{PD}, and patch size p. For the masking-ratio sweep, $\lambda_{\mathrm{PD}} = 0.02$ and $p = 4$ are fixed; for the PD-loss-weight sweep, $\rho = 0.75$ and $p = 4$ are fixed; and for the patch-size sweep, $\rho = 0.75$ and $\lambda_{\mathrm{PD}} = 0.05$ are fixed. Bold values indicate the best performance within each sensitivity block. † denotes the value adopted in the main configuration.

		Task 1: Family classification				Task 2: C_{11} classification		Task 3: C_{11} clustering			Task 4: Regression	
Parameter	Value	$\mathrm{Acc}^{\mathrm{Str}}_{\mathrm{Lin}}\uparrow$	$\mathrm{Acc}^{\mathrm{Rel}}_{\mathrm{Lin}}\uparrow$	$\mathrm{Acc}^{\mathrm{Str}}_{k\mathrm{NN}}\uparrow$	$\mathrm{Acc}^{\mathrm{Rel}}_{k\mathrm{NN}}\uparrow$	$F_{k\mathrm{NN}}\uparrow$	$F_{\mathrm{Lin}}\uparrow$	$\eta^2\uparrow$	$\mathrm{Acc}_{\mathrm{Match}}\uparrow$	$F_{\mathrm{Macro}}\uparrow$	$R^2_{C_{11}}\uparrow$	$R^2_{C_{21}}\uparrow$
Mask ratio ρ	0.55	0.8376	0.8484	0.8275	0.9317	**0.8869**	**0.9218**	**0.3901**	0.5459	0.5349	**0.99225**	0.99628
	0.65	0.8606	0.8750	0.8272	0.9360	0.8846	0.9157	0.3134	0.5493	**0.5473**	0.99185	**0.99636**
	0.75†	0.8764	0.8879	0.8347	0.9396	0.8793	0.9075	0.3166	**0.5508**	0.5465	0.99190	0.99633
	0.85	**0.8908**	**0.9023**	**0.8491**	**0.9450**	0.8728	0.9108	0.1500	0.4409	0.4412	0.99204	0.99613
λ_{PD}	0.01	**0.8872**	**0.8994**	0.8279	0.9310	0.8778	**0.9210**	0.0972	0.4518	0.4529	0.99184	0.99610
	0.02	0.8764	0.8879	**0.8347**	**0.9396**	0.8793	0.9075	0.3166	0.5508	0.5465	0.99190	**0.99633**
	0.05†	0.8660	0.8807	0.8315	0.9364	0.8847	0.9165	**0.4612**	**0.6127**	**0.6097**	0.99194	0.99620
	0.10	0.8591	0.8768	0.8225	0.9299	**0.8891**	0.9162	0.4560	0.6075	0.6055	**0.99221**	0.99604
	0.20	0.8365	0.8545	0.8178	0.9271	0.8887	0.9116	0.3129	0.5445	0.5381	0.99206	0.99583
Patch size p	4†	**0.8660**	**0.8807**	**0.8315**	**0.9364**	**0.8847**	**0.9165**	0.4612	0.6127	0.6097	**0.99194**	**0.99620**
	8	0.8383	0.8548	0.8246	0.9332	0.8789	0.9137	**0.4818**	**0.6277**	**0.6252**	0.99110	0.99587
	16	0.8146	0.8311	0.8124	0.9238	0.8810	0.9131	0.4550	0.6137	0.6087	0.99010	0.99567

A.10 Parameter Sensitivity and Hyperparameter Selection

We further examine the sensitivity of the proposed framework to three key hyperparameters: the patch size p, the masking ratio ρ, and the effective-loss weight λ_{eff}, which are most relavant to our proposed mechanisms. The results on the Interpolation TPMS dataset are summarized in Table S6. Each parameter is varied independently while the remaining relevant settings are fixed.

For the patch size, $p = 4$ achieves the best results in Task 1,2,and 4. The only exception is Task 3, where $p = 8$ yields slightly better clustering performance, improving η^2, $\mathrm{Acc}_{\mathrm{Match}}$, and F_{Macro} by only 0.0206, 0.0150, and 0.0155, respectively, compared with $p = 4$. Given these relatively small gains in Task 3 and the consistently better performance of $p = 4$ in Tasks 1, 2, and 4, we adopt $p = 4$ for the main experiments.

The masking ratio shows a clearer trade-off between family-level structural discrimination and property-related organization of the latent space. Increasing the masking ratio from $\rho = 0.75$ to 0.85 improves the four Task 1 metrics by up to 0.0144, indicating stronger family discrimination under more aggressive masking. However, this comes at a substantially larger cost in Task 3: η^2 decreases from 0.3166 to 0.1500, while $\mathrm{Acc}_{\mathrm{Match}}$ and F_{Macro} decrease from 0.5508 and 0.5465 to 0.4409 and 0.4412, respectively. The kNN property-classification F1 in Task 2 also decreases from 0.8793 to 0.8728. Conversely, lower masking ratios improve some property-related metrics but lead to noticeably weaker Task 1 performance; for example, reducing ρ from 0.75 to 0.55 lowers the two linear-probe family metrics by approximately 0.039. We therefore choose $\rho = 0.75$ as a balanced operating point, retaining strong structural discrimination without the substantial degradation in clustering observed at $\rho = 0.85$.

A similar trade-off is observed for the PD-loss weight λ_{PD}. While $\lambda_{\mathrm{PD}} = 0.01$ gives the best linear family-classification performance, $\lambda_{\mathrm{PD}} = 0.05$ improves η^2, $\mathrm{Acc}_{\mathrm{Match}}$, and F_{Macro} in Task 3 by 0.3640, 0.1609, and 0.1568, respectively, at the cost of only 0.0212 and 0.0187 in the two linear Task 1 metrics. Further increasing λ_{PD} to 0.10 yields only marginal gains in Task 2 and Task 4 while slightly reducing Task 1 and Task 3 performance. We therefore select $\lambda_{\mathrm{PD}} = 0.05$ as a balanced setting across tasks.

Based on above sensitivity results and analysis, we use $p = 4$, $\rho = 0.75$, and $\lambda_{\mathrm{PD}} = 0.05$ in the main experiments.

Table S7: **Parameter sensitivity analysis on the MetaDB dataset.** We evaluate the sensitivity to the PD-loss weight λ_{PD}, while fixing the masking ratio to $\rho = 0.75$ and patch size to $p = 4$. Bold values indicate the best performance within the sensitivity sweep. † denotes the value adopted in the main configuration.

		Task 1: Family classification		Task 2: C_{11} classification				Task 3: C_{11} clustering			Task 4: Regression	
Parameter	Value	$F_{\mathrm{Lin}} \uparrow$	$F_{k\mathrm{NN}} \uparrow$	$\mathrm{BA}_{\mathrm{Lin}} \uparrow$	$F_{\mathrm{Lin}} \uparrow$	$\mathrm{BA}_{k\mathrm{NN}} \uparrow$	$F_{k\mathrm{NN}} \uparrow$	$\eta^2 \uparrow$	$\mathrm{Acc}_{\mathrm{Match}} \uparrow$	$F_{\mathrm{Macro}} \uparrow$	$R^2_{C_{11}} \uparrow$	$R^2_{C_{21}} \uparrow$
λ_{PD}	0.005	**0.9945**	0.9844	0.9569	0.9569	0.9513	0.9512	**0.3235**	0.6593	0.6723	0.9026	0.9169
	0.010†	0.9944	0.9845	**0.9641**	**0.9640**	0.9513	0.9512	0.3228	0.6598	0.6728	0.8991	**0.9360**
	0.015	0.9930	**0.9877**	0.9600	0.9599	0.9539	0.9537	0.3217	0.6598	0.6728	0.8565	0.9033
	0.020	0.9931	0.9840	0.9635	0.9635	0.9539	0.9537	0.3202	0.6687	0.6809	0.9007	0.9192
	0.030	0.9937	0.9827	0.9560	0.9560	0.9503	0.9502	0.3073	0.5811	0.5603	0.8753	0.9122
	0.100	0.9924	0.9859	0.9559	0.9559	0.9513	0.9511	0.3073	0.5801	0.5595	0.8826	0.8981
	0.150	0.9911	0.9842	0.9600	0.9600	0.9523	0.9522	0.3012	0.5274	0.4989	**0.9031**	0.9004
	0.200	0.9927	0.9848	0.9575	0.9574	0.9503	0.9501	0.3204	**0.6703**	**0.6823**	0.8898	0.9125
	0.225	0.9924	0.9869	0.9569	0.9569	**0.9544**	**0.9542**	0.3057	0.5157	0.4999	0.8849	0.9106
	0.250	0.9921	0.9868	0.9590	0.9589	0.9518	0.9517	0.2510	0.4522	0.4389	0.8858	0.8995
	0.275	0.9890	0.9854	0.9560	0.9559	0.9533	0.9532	0.2002	0.5125	0.4799	0.8947	0.9034
	0.300	0.9900	0.9863	0.9566	0.9564	0.9528	0.9527	0.2032	0.5217	0.5044	0.9009	0.9078
	0.350	0.9906	0.9848	0.9605	0.9605	0.9508	0.9506	0.1078	0.4711	0.4550	0.8520	0.9039
	0.400	0.9923	0.9838	0.9584	0.9585	0.9498	0.9497	0.1092	0.4636	0.4492	0.9024	0.8950