

001
002
003
004
005
006
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046

Multitask Graph Neural Networks for Elastoplastic Response Prediction in Dual-Phase Polycrystals*

Erdem Caliskan¹, Anik Das Anto¹, Massimiliano Lupo Pasini²,
Stephanie TerMaath¹, Reza Abedi^{1*}

¹Department of Mechanical and Aerospace Engineering, University of
Tennessee Knoxville, 1506 Middle Drive, Knoxville, 37916, TN, USA.

²Computational Sciences and Engineering Division, Oak Ridge National
Laboratory, 1 Bethel Valley Rd, Oak Ridge, 37831, TN, USA.

*Corresponding author(s). E-mail(s): rabedi@utk.edu;
Contributing authors: ecalisk1@vols.utk.edu; aanto@vols.utk.edu;
lupopasinim@ornl.gov; stermaat@utk.edu;

Abstract

Microstructure-sensitive prediction of elastoplastic response remains a recurring bottleneck in multiscale damage and fatigue modeling, where large ensembles of statistically distinct polycrystals are required to quantify variability and extreme-value behavior. In this work, we develop a multitask graph neural network (GNN) surrogate that maps dual-phase ferrite–martensite polycrystal microstructures to Statistical Volume Element (SVE)-level elastoplastic quantities of interest. Each SVE is represented as a grain-adjacency graph, with node features encoding phase, geometry, and crystallographic orientation, and edge features encoding relative misorientation. A message-passing graph convolution generates node embeddings, which are pooled into a graph representation and passed to a multitask regression head that jointly predicts 10 scalar QoIs and vector-valued

*This manuscript has been authored in part by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

047 stress–strain responses in orthogonal loading directions across multiple marten-
048 site volume fractions and SVE sizes. Results show high accuracy for scalar QoIs
049 and strong agreement for full stress–strain trajectories, with population envelopes
050 reproducing both median behavior and finite-SVE variability across compositions
051 and partition scales. A unified model trained on pooled volume-fraction data
052 preserves most within-regime accuracy relative to regime-specific models while
053 also capturing the broader cross-regime variation reflected in the pooled test
054 set. Distributional comparisons further demonstrate that the surrogate preserves
055 heterogeneity under SVE partitioning, enabling statistically consistent block-
056 wise random-field construction for mesoscale analyses. Overall, the proposed
057 grain-graph surrogate provides a practical pathway to accelerate ensemble-based
058 studies of constitutive variability relevant to crack nucleation and early-stage
059 damage.

060 **Keywords:** Graph neural networks, Polycrystals, Crystal plasticity, Statistical volume
061 elements, Elastoplastic response, Multitask learning

062
063
064

065 1 Introduction

066
067

068 Polycrystalline metals exhibit a strong microstructure-sensitive elastoplastic response
069 where grain morphology, crystallographic texture, and phase topology jointly shape
070 the onset of yielding, the evolution of hardening, and the directional character of the
071 apparent stiffness. This sensitivity is central to multiscale modeling workflows where
072 macroscopic predictions must reflect not only mean constitutive behavior but also
073 variability, tails of distributions, and spatially localized hot spots that drive early-stage
074 damage. In practice, these requirements lead to ensemble computations over many
075 statistically distinct microstructures, and the cost of repeated high-fidelity simulations
076 becomes the dominant bottleneck.

082 Computational homogenization formalizes this scale-bridging task through the use
083 of *Representative Volume Element* (RVE), which provides an effective representa-
084 tion of the material’s behavior once scale separation is achieved [1–4]. However, in
085 many applications, the relevant domain is not sufficiently large to reach an RVE
086 limit, the RVE size can be prohibitively large (particularly for nonlinear properties),
087 or the RVE may not exist in a strict sense for certain responses such as softening
088
089
090
091
092

and post-peak behavior [5–9]. These limitations motivate *Statistical Volume Elements* (SVEs), which provide a continuum description via *apparent* properties that retain finite-size variability and directional dependence [10–14]. From an engineering standpoint, SVEs offer a controllable compromise: one can trade computational cost against the degree of heterogeneity preserved, while still maintaining a consistent mapping from microstructure to homogenized quantities across the sizes at which variability remains non-negligible [15–17]. Most importantly, by resulting in an inhomogeneous and random coarse-grained representation of the material microstructure, they can systematically relate material aleatory uncertainties to the distribution of macroscopic Quantities of Interest (QoIs). This propagation of statistics across scales cannot be achieved by using RVEs.

A common route to SVE-based mesoscale modeling is the moving-window method, where local apparent properties are obtained by partitioning a larger microstructural domain into a hierarchy of windows [18–22]. This construction naturally supports random-field descriptions of constitutive parameters, which are essential when spatial variability governs macroscopic failure patterns. For example, the ability to resolve intermediate-scale variability (below the RVE limit but above the grain scale) can strongly influence predicted crack paths and failure statistics [9, 15, 23]. For polycrystalline metals, the same logic applies to constitutive variability relevant to crack nucleation, where local yielding and hardening fluctuations can matter even when global averages appear well converged.

Despite their modeling value, SVEs impose a substantial computational burden when evaluated using full-field crystal plasticity. Crystal Plasticity Finite Element Method (CPFEM) simulations remain a standard tool for resolving grain-scale anisotropy and heterogeneous slip, but they are expensive in multi-phase settings and in regimes where large ensembles are required [24, 25]. Moreover, the SVE size needed for plastic descriptors can be substantially larger than that needed for elastic

139 descriptors, implying that the most informative statistics are often the most costly
140 to obtain [17]. This creates a practical gap for multiscale studies that require (i)
141 many SVEs per microstructure for block-wise random-field construction, (ii) many
142 microstructures per material state for uncertainty quantification, and (iii) nonlinear
143 response.

147 Surrogate modeling has therefore become an active strategy to accelerate
148 microstructure-property linkages. Classical reduced-order and data-driven homoge-
149 nization frameworks have demonstrated that nonlinear effective response can be
150 learned from micromechanical simulations when suitable microstructural descrip-
151 tors are available [26, 27]. Machine learning has also been used to boost nonlinear
152 homogenization by coupling data-driven components with established homogeniza-
153 tion formalisms [28]. More recently, deep learning has been used to approximate
154 crystal plasticity response and related field mappings, including approaches that tar-
155 get CPFEM-based predictions or embed learned constitutive components into finite
156 element workflows [29–32]. Neural operator formulations further extend this line of
157 work by learning mappings between function spaces for stress or field prediction,
158 with recent efforts introducing physics-aware constraints, such as preserving equi-
159 librium [33]. While these developments are promising, microstructure representation
160 remains a recurring challenge: image/voxel-based encodings can be data-intensive and
161 resolution-tied, and they struggle to accommodate variable numbers of grains across
162 SVE sizes without careful architectural choices.

174 Graph representations provide a principled alternative because they match the
175 discrete topology of polycrystals. In geometric deep learning, graphs provide a
176 native representation for irregular domains and their relational structure [34–36]. For
177 polycrystalline microstructures, grains can be treated as nodes and grain-boundary
178 neighborhoods as edges, enabling the encoding of crystallographic state at the node

182
183
184

level and misorientation-driven interactions at the edge level. Graph-based GNN surrogates have also demonstrated accurate homogenization in composite microstructures using unstructured topology-derived graphs, including prediction of tensorial stiffness and strength measures across diverse configurations [37]. This representation supports size variability (different grain counts per SVE), preserves physical interpretability, and allows message passing to model short-range mechanisms that dominate heterogeneous response. Recent mechanics-focused studies have further demonstrated the utility of GNNs for material structure-property prediction across diverse settings, including microstructure-informed elasticity and polycrystal-related tasks [38–41]. Graph-based learning has also been explored in fatigue response predictions, where localized values are of primary interest, thereby reinforcing the suitability of grain graphs for capturing rare yet consequential response modes [42]. In addition, open software ecosystems for graph learning and scientific GNN training reduce barriers to scalable surrogate deployment [43–45].

In this work, we develop a deterministic multitask graph neural network surrogate for predicting SVE-level elastoplastic responses in dual-phase polycrystals. Each SVE is represented as a grain-adjacency graph with node features encoding phase, geometry, and crystallographic orientation, and with edge features encoding relative misorientation. The surrogate learns a shared microstructure embedding via Principal Neighbourhood Aggregation (PNA) [46] message passing and predicts both scalar QoIs, including stiffness components, isotropic moduli, yield strengths, and hardening descriptors, as well as vector-valued stress–strain responses in orthogonal loading directions. Beyond pointwise accuracy, we evaluate whether the surrogate preserves finite-SVE variability across sizes and compositions in a manner consistent with SVE-based mesoscale modeling.

The main contributions of this study are:

- 231 • A grain-graph representation and multitask GNN surrogate for mapping polycrystal
232 SVEs to elastic and plastic homogenized quantities across multiple SVE sizes.
233
- 234 • Joint prediction of scalar properties and full stress–strain trajectories within
235 a shared embedding framework, supporting nonlinear response characterization
236 beyond scalars.
237
- 238 • Validation at the population level (distributions and variability) to assess suitability
239 for SVE-based block-wise random-field construction and ensemble-driven multiscale
240 analyses.
241
242
243
244

245 **2 Dataset Description**

246 The dataset employed in this study is derived from the high-fidelity CPFEM simula-
247 tions reported in our prior work [17]. In that work, the size dependency of homogenized
248 elastic and plastic properties of dual-phase polycrystalline materials was systematically
249 investigated using the SVE framework. The present study uses that validated compu-
250 tational database to construct a structured microstructure-property dataset suitable
251 for deterministic surrogate modeling.
252
253
254
255
256
257
258

259 **2.1 Material System and Microstructure Generation**

260 The material system consists of dual-phase ferrite-martensite polycrystals. Three
261 microstructural configurations were generated corresponding to martensite volume
262 fractions (v_f) of 10% (M1), 45% (M2), and 90% (M3). Each macroscopic microstruc-
263 tural domain has dimensions $400\ \mu\text{m} \times 400\ \mu\text{m} \times 1\ \mu\text{m}$ with an average grain equivalent
264 diameter of $\bar{d} = 5\ \mu\text{m}$. Grains are assumed equiaxed and assigned random crys-
265 tallographic orientations sampled from a uniform orientation distribution function.
266 Phase-dependent elastic constants and crystal plasticity parameters were assigned at
267 the grain level. Microstructures were generated using DREAM.3D [47] and discretized
268 for full-field CPFEM simulations.
269
270
271
272
273
274
275
276

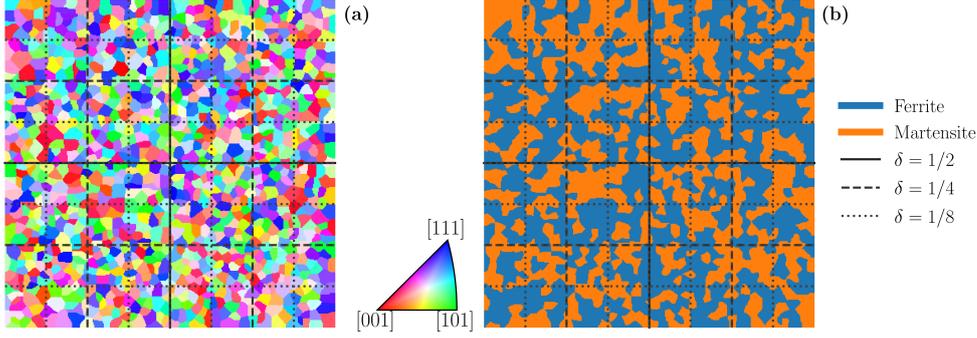


Fig. 1 Example $400\ \mu\text{m} \times 400\ \mu\text{m}$ parent microstructure showing (a) grain orientations colored according to the inverse pole figure (IPF) key and (b) phase distribution for ferrite and martensite. Grid lines indicate the SVE partitioning levels corresponding to $\delta = 1/2, 1/4,$ and $1/8$.

2.2 SVE Extraction

To capture spatial variability and size effects, the SVEs were extracted using a space-filling moving-window partitioning approach. Figure 1 illustrates an example microstructure from $v_f = 0.45$, where the grain orientations are colored using inverse pole figure (IPF) and the partitions used for SVEs are plotted over the domain. The considered SVE sizes are $L_{\text{SVE}} \in \{50, 100, 200, 400\}\ \mu\text{m}$, which tile 30 different $400\ \mu\text{m} \times 400\ \mu\text{m}$ parent domains into $8 \times 8, 4 \times 4, 2 \times 2,$ and 1×1 grids, respectively, without overlap or gaps. These correspond to size $\delta = L_{\text{SVE}}/400\ \mu\text{m} \in \{1/8, 1/4, 1/2, 1\}$. The parameter δ denotes the subdivision ratio relative to the full domain, so decreasing L_{SVE} (larger subdivision) increases the number of SVEs obtained from each parent microstructure and, consequently, the share of samples at that size. In total, there are 30, 120, 480, and 1920 samples per v_f for $\delta = 1, 1/2, 1/4,$ and $1/8$, respectively. All SVEs are non-overlapping and collectively cover the entire domain.

This construction ensures that the dataset explicitly captures composition dependence through the martensite v_f , size dependence through the SVE length scale, and intrinsic statistical variability across different microstructural realizations. As a result,

277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322

323 the database systematically encodes the coupled effects of phase contrast, observa-
 324 tion window size, and microstructural randomness on the homogenized mechanical
 325 response.
 326 response.

327

328 **2.3 CPFEM Simulation and Homogenization**

329

330 Each SVE was analyzed using a full-field crystal plasticity formulation under mixed
 331 boundary conditions (MBC). Three independent loading cases were applied, namely
 332 uniaxial tension in the x direction, uniaxial tension in the y direction, and in-plane
 333 shear loading. The homogenized stress-strain response was obtained by volume averag-
 334 ing the microscopic stress and strain fields. In-plane stresses and strains were expressed
 335 in Voigt notation as
 336

337

338

339

340

341

342

343

344

345

346

347

348

349

350

351

352

353

354

355

356

357

358

359

360

361

362

363

364

365

366

367

368

$$\boldsymbol{\sigma} = \begin{bmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{xy} \end{bmatrix}, \quad \boldsymbol{\gamma} = \begin{bmatrix} \varepsilon_{xx} \\ \varepsilon_{yy} \\ 2\varepsilon_{xy} \end{bmatrix}, \quad (1)$$

and the homogenized elastic constitutive relation was written as

348

349

350

351

352

353

354

355

356

357

358

359

360

361

362

363

364

365

366

367

368

$$\boldsymbol{\sigma} = \mathcal{C} \boldsymbol{\gamma}, \quad (2)$$

where the in-plane stiffness tensor is given by

353

354

355

356

357

358

359

360

361

362

363

364

365

366

367

368

$$\mathcal{C} = \begin{bmatrix} \mathcal{C}_{11} & \mathcal{C}_{12} & \mathcal{C}_{13} \\ \mathcal{C}_{12} & \mathcal{C}_{22} & \mathcal{C}_{23} \\ \mathcal{C}_{13} & \mathcal{C}_{23} & \mathcal{C}_{33} \end{bmatrix}. \quad (3)$$

Here, \mathcal{C}_{33} denotes the in-plane shear stiffness.

361

362

363

364

365

366

367

368

369

370

371

372

373

374

375

376

377

378

379

380

364 **2.4 Output Quantities of Interest**

For each SVE, both elastic and plastic homogenized quantities were reported.

2.4.1 Elastic Properties

The elastic outputs include the in-plane stiffness components $\mathcal{C}_{11}, \mathcal{C}_{12}, \mathcal{C}_{22}, \mathcal{C}_{33}$, and bulk modulus κ together with isotropic shear modulus μ obtained through Reuss angular averaging of \mathcal{C} , *cf.* [5]. Although the full in-plane stiffness tensor in Eq. 3 is computed from the three CPFEM loading cases for every SVE, the present dataset is dominated by the orthotropic components $\{\mathcal{C}_{11}, \mathcal{C}_{12}, \mathcal{C}_{22}, \mathcal{C}_{33}\}$. Because the coupling terms \mathcal{C}_{13} and \mathcal{C}_{23} remain consistently close to zero across the dataset, only these orthotropic-dominant components are retained for learning and reporting.

2.4.2 Plastic Properties

Plastic quantities were extracted from the homogenized stress–strain curves. The yield strength Y was determined using the 0.2% offset criterion, while directional yield strengths Y_x and Y_y were computed for loading along the principal directions. The tail hardening modulus was defined as

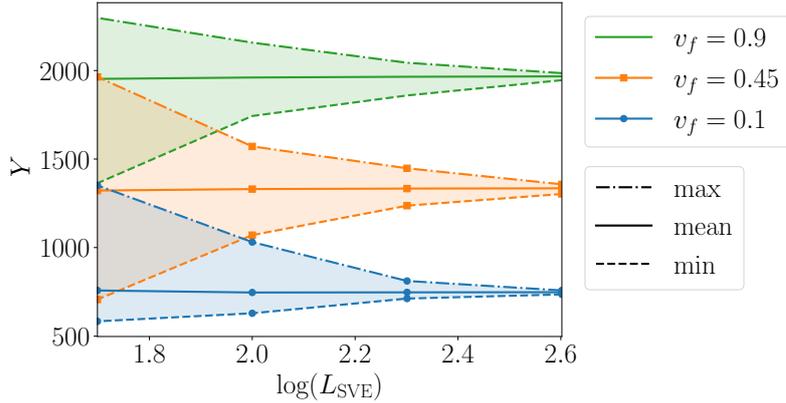
$$H = \frac{1}{2} \left(\frac{d\sigma_{xx}}{d\varepsilon_{xx}} + \frac{d\sigma_{yy}}{d\varepsilon_{yy}} \right), \quad (4)$$

and evaluated as the average slope over the final six increments of the macroscopic stress–strain response. In addition, homogenized stress–strain curves $\sigma_{xx}(\varepsilon)$ and $\sigma_{yy}(\varepsilon)$ were sampled at 24 strain increments to characterize the nonlinear mechanical response.

2.5 Statistical Characteristics

The dataset exhibits several key characteristics. Property variance decreases with increasing SVE size, indicating the convergence toward the representative-volume-element behavior. An example convergence is given in Figure 2. Elastic and plastic

415
416
417
418
419
420
421
422
423
424
425
426
427



428 **Fig. 2** Variation of the average yield strength Y (MPa) with SVE size for three martensite volume
429 fractions plotted against $\log(L_{\text{SVE}})$. Solid lines denote the mean response across SVE realizations,
430 while dashed and dash-dotted lines represent the minimum and maximum values, respectively.

431
432
433
434
435
436
437
438
439
440
441
442
443
444

properties exhibit distinct scaling trends, with plastic quantities, particularly the hard-
ening modulus H , showing larger variability and slower convergence. Intermediate
martensite fractions exhibit increased anisotropy and statistical dispersion. Further-
more, property distributions at small SVE sizes exhibit the highest level of anisotropy
and inhomogeneity. Overall, the dataset spans multiple compositions, length scales,
and mechanical regimes, providing a comprehensive microstructure–property database
suitable for surrogate modeling of SVE-level homogenized response.

445 3 Graph Neural Network Framework

447 Figure 3 summarizes the grain-graph surrogate used to predict SVE-level elastoplastic
448 QoIs from dual-phase polycrystal microstructures. The framework maps each SVE
449 to a graph, performs message passing to obtain node embeddings that encode local
450 neighborhood interactions, and then aggregates these embeddings into a compact SVE
451 representation that drives a multitask regression head.
452
453
454

455 Each SVE is represented as a graph $G = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of microstruc-
456 tural entities (nodes) and \mathcal{E} is the set of neighborhood relations (edges). We use
457
458
459
460

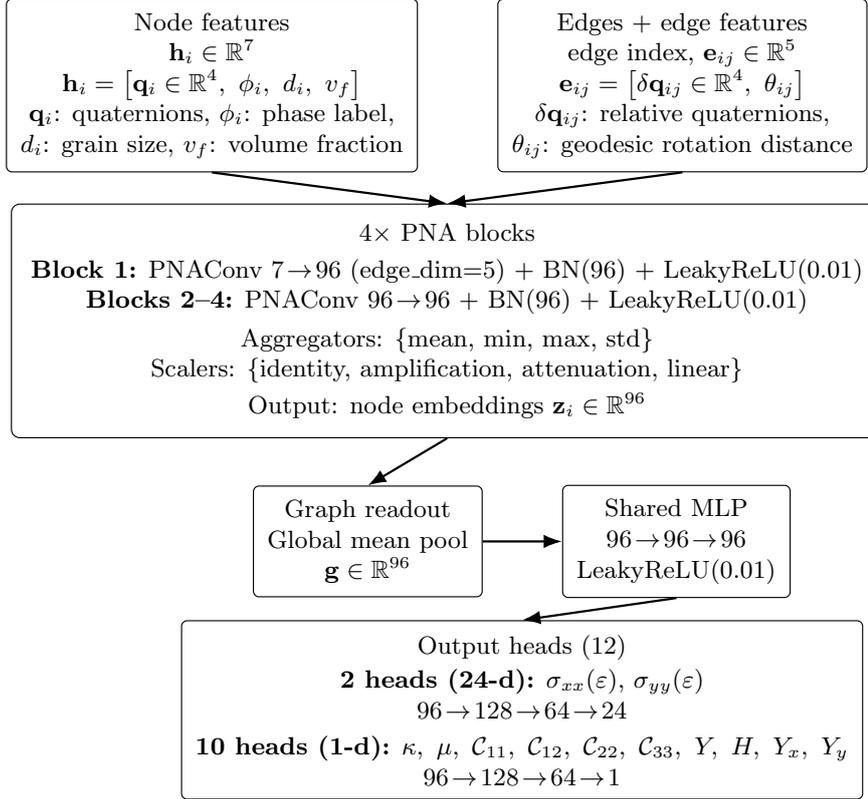


Fig. 3 Simplified PNA-based GNN: node features $\mathbf{h}_i = [\mathbf{q}_i, \phi_i, d_i, v_f]$, edge features $\mathbf{e}_{ij} = [\delta\mathbf{q}_{ij}, \theta_{ij}]$, 4 PNA blocks using the shown aggregators/scalers, global mean pooling, shared Multilayer Perceptron (MLP), and 12 prediction heads.

a grain-based VE representation because it significantly reduces input dimensionality while preserving microstructure. In grain-based graphs, each node corresponds to a grain. This choice provides a physically interpretable representation in which nodes align with crystallographic state variables (orientation, phase) and geometric descriptors (grain size).

Edges encode short-range microstructural interactions. In the grain-graph setting, we connect nodes using adjacency across grain boundaries: two grains i and j are neighbors if they share an interface in the tessellation/segmentation. We treat the graph as undirected in the physical sense, but implement it with two directed edges ($i \rightarrow j$) and ($j \rightarrow i$) for message passing. This construction produces sparse graphs with

507 locality aligned to the mechanisms that drive effective response: Phase connectivity,
 508 crystallographic misorientation, and neighboring grains.

509 Each node $i \in \mathcal{V}$ is assigned a feature vector $\mathbf{h}_i \in \mathbb{R}^7$ (Fig. 3)

$$511 \quad \mathbf{h}_i = [\mathbf{q}_i \in \mathbb{R}^4, \phi_i, d_i, v_f], \quad (5)$$

512 where \mathbf{q}_i is the grain orientation encoded as a unit quaternion¹, ϕ_i is a phase label
 513 (binary encoding for ferrite/martensite), d_i is a grain size descriptor (e.g., number of
 514 elements), and v_f is the SVE-level volume fraction. The scalar v_f is replicated to each
 515 node to provide global compositional context during message passing.

516 Each directed edge $(i, j) \in \mathcal{E}$ carries an edge feature vector $\mathbf{e}_{ij} \in \mathbb{R}^5$

$$517 \quad \mathbf{e}_{ij} = [\delta\mathbf{q}_{ij} \in \mathbb{R}^4, \theta_{ij}], \quad (6)$$

518 where $\delta\mathbf{q}_{ij}$ is the relative quaternion between grains i and j , and θ_{ij} is the geodesic
 519 rotation distance, providing a scalar misorientation measure. This feature design sepa-
 520 rates node-level descriptors (state and geometry) from edge-level descriptors (pairwise
 521 crystallographic relationships), which is particularly relevant for plasticity-sensitive
 522 QoIs.

523 The surrogate is a graph neural network with four message-passing (PNA) lay-
 524 ers. In each layer, every node (grain) updates its internal feature vector by combining
 525 its own information with information coming from its neighboring grains through the
 526 grain-boundary edges. We start from 7 input features per node and 5 features per
 527 edge; the first PNA layer maps the node description from \mathbb{R}^7 to a 96-component hid-
 528 den representation, and the next three layers keep the same 96-dimensional size while

529 ¹Grain orientations are represented with unit quaternions rather than Euler angles because quaternions
 530 provide a smooth, singularity-free description of 3D rotations. In contrast, Euler angles can suffer from
 531 gimbal-lock-type singularities and discontinuities, and their representation is more redundant under crystal
 532 symmetry. For learning on grain graphs, quaternions therefore provide a more stable basis for encoding
 533 orientations and computing relative misorientations between neighboring grains.

progressively refining it as shown in Figure 3. The output of these four layers is a learned 96-dimensional vector $\mathbf{z}_i \in \mathbb{R}^{96}$ for each node, which is called a node embedding, providing a compact numerical summary of the local microstructural context around grain i .

PNA is designed to perform well when different nodes have very different numbers of neighbors. To achieve this, it summarizes each node’s neighborhood in several complementary ways using the aggregators {mean, min, max, std}, and then adjusts these summaries with degree-dependent scalars {identity, amplification, attenuation, linear}. Using multiple summaries, together with degree-aware rescaling, makes the layer more stable and robust across SVEs whose grain-adjacency graphs can vary in connectivity.

Node embeddings are reduced to an SVE representation via global mean pooling

$$\mathbf{g} = \frac{1}{|\mathcal{V}|} \sum_{i \in \mathcal{V}} \mathbf{z}_i, \quad \mathbf{g} \in \mathbb{R}^{96}. \quad (7)$$

Mean pooling provides a stable summary when the number of grains varies across SVEs and partition scales, a significant advantage over CNNs. A shared Multilayer Perceptron (MLP) refines the global mean pooling by adding three more layers and using LeakyReLU (0.01) as the activation function. The shared latent is then mapped to 12 task-specific output heads:

- Two 24-dimensional heads for stress–strain responses, $\sigma_{xx}(\varepsilon)$ and $\sigma_{yy}(\varepsilon)$, each parameterized by $96 \rightarrow 128 \rightarrow 64 \rightarrow 24$.
- Ten scalar heads for $\{\kappa, \mu, C_{11}, C_{12}, C_{22}, C_{33}, Y, H, Y_x, Y_y\}$, each parameterized by $96 \rightarrow 128 \rightarrow 64 \rightarrow 1$.

This structure enforces a shared microstructure embedding while allowing the final mappings to specialize by the QoI type (scalar versus curve).

Training minimizes a multi-output mean squared error (MSE) aggregated over all heads. Let the set of tasks be $\mathcal{T} = \{1, \dots, 12\}$, and let task t have output dimension

553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598

599 d_t (with $d_t = 1$ for scalar QoIs and $d_t = 24$ for stress-strain heads). For a graph G ,
600 predictions are $\hat{\mathbf{y}}^{(t)}(G)$ and targets are $\mathbf{y}^{(t)}(G)$. The combined loss function is

$$602$$

$$603 \mathcal{L} = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \frac{1}{d_t} \left\| \mathbf{y}^{(t)}(G) - \hat{\mathbf{y}}^{(t)}(G) \right\|_2^2. \quad (8)$$

$$604$$

$$605$$

$$606$$

607 The per-task normalization by d_t prevents the 24-dimensional curve heads from domi-
608 nating the optimization purely by dimensionality. In practice, targets are standardized
609 using statistics computed on the training set to reduce scale imbalance across elastic
610 constants and plastic metrics; predictions are then mapped back to physical units for
611 reporting.

612 To prevent information leakage arising from correlated SVEs extracted from the
613 same parent microstructure, we perform an image-level split. Parent microstructure
614 images are first divided into train/validation/test sets. Then, all SVEs derived from
615 a given parent image are assigned to the same subset, including the full hierarchy of
616 partitions (e.g., the $\delta = 1/1$ SVE and its corresponding $\delta = 1/2$, $\delta = 1/4$, and $\delta = 1/8$
617 sub-SVEs). We target equal representation for each martensite v_f regime within each
618 subset (stratified by v_f), and we preserve representation across SVE sizes through
619 hierarchical assignment. In the representative split used for the main experiments,
620 selection is performed at the image level (e.g., 26 parent images for training and 4
621 parent images for testing), after which all corresponding partition-derived SVEs are
622 included automatically in the same subset.

623 Table 1 summarizes the SVE sampling strategy and the resulting dataset compo-
624 sition. For each martensite $v_f \in \{0.10, 0.45, 0.90\}$, we consider 30 parent microstruc-
625 tures over a $400 \mu\text{m} \times 400 \mu\text{m}$ domain and extract non-overlapping SVEs of L_{SVE}
626 and δ as explained in Section 2. The table also reports the average grain count per
627 SVE, emphasizing that smaller observation windows contain fewer grains and therefore
628 represent more localized microstructural environments. As a representative measure
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644

Table 1 Summary of the SVE dataset (derived from Anto *et al.* [17] for three martensite volume fractions: $v_f = \{0.10, 0.45, 0.90\}$, with 30 parent microstructures per v_f). Coefficient of variance (CoV) of Y is given as an example.

$L_{\text{SVE}} [\mu\text{m}]$	δ	Avg. grain count	Samples per v_f	Share [%]	CoV(Y)
400	1/1	1180	30	1.2	0.006
200	1/2	315	120	4.7	0.023
100	1/4	89	480	18.8	0.051
50	1/8	27	1920	75.3	0.113

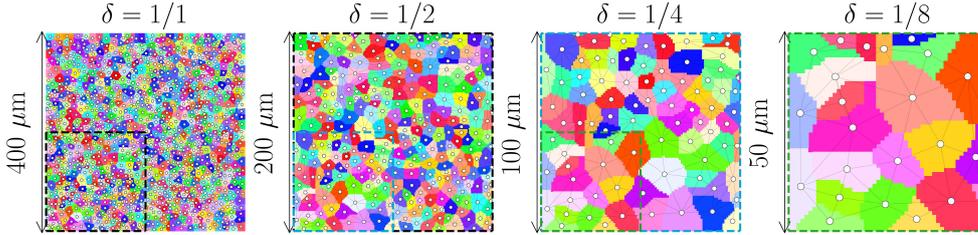


Fig. 4 Example microstructure domains at four SVE sizes together with the corresponding grain-graph representations used as inputs to the GNN model.

of size-dependent scatter, the coefficient of variation (CoV) of the yield strength Y is listed; the monotonic increase in CoV with decreasing L_{SVE} reflects stronger microstructural variability at smaller length scales. Figure 4 illustrates an example domain and corresponding graph from each size referenced in Table 1.

4 Results

This section evaluates the predictive performance and physical consistency of the proposed multitask grain-graph surrogate. We first examine whether a single model trained on multiple martensite volume fractions can maintain accuracy across individual composition regimes (Section 4.1). This analysis establishes that pooling data across volume fractions enables a unified surrogate while preserving most predictive fidelity. We then quantify the accuracy of scalar elastic and plastic quantities of interest using parity plots, highlighting regions where the model achieves strong agreement

691 **Table 2** Effect of pooling volume fractions during training. $R^2_{\text{spec}}(v_f)$ denotes within-regime R^2
692 for a model trained only on regime v_f and evaluated on test SVEs from the same regime.
693 $R^2_{\text{unif}|v_f}$ denotes within-regime R^2 for the unified model \mathcal{M}_{all} evaluated on the test subset
694 restricted to regime v_f . The pooled scores compare (i) concatenated predictions from specialized
695 models (each test sample predicted by its regime-specific model; columns 2 to 7) and (ii) a single
696 unified model evaluated on the concatenated test set (columns 8 and 9).

QoI	R^2_{spec}			$R^2_{\text{unif} v_f}$			$\cup v_f$	
	$v_f=0.1$	$v_f=0.45$	$v_f=0.9$	$v_f=0.1$	$v_f=0.45$	$v_f=0.9$	$R^2_{\text{spec}}^{\text{pool}}$	$R^2_{\text{unif}}^{\text{pool}}$
C_{11}	0.881	0.905	0.761	0.865	0.863	0.751	0.993	0.990
C_{12}	0.734	0.762	0.392	0.704	0.714	0.422	0.969	0.969
C_{22}	0.894	0.910	0.768	0.869	0.880	0.738	0.993	0.991
C_{33}	0.143	0.209	0.075	0.152	0.209	0.103	0.807	0.828
κ	0.966	0.962	0.919	0.931	0.940	0.908	0.998	0.996
μ	0.504	0.564	0.247	0.550	0.565	0.335	0.947	0.950
Y	0.792	0.853	0.586	0.771	0.812	0.611	0.987	0.985
H	0.594	0.745	0.323	0.571	0.648	0.403	0.964	0.966
Y_x	0.779	0.843	0.577	0.738	0.799	0.605	0.986	0.985
Y_y	0.734	0.848	0.561	0.748	0.812	0.591	0.986	0.984

708

709

710

711

712

713

714

715

716

717

718

719

720

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

with CPFEM results (Section 4.2). Next, we evaluate vector-valued stress-strain pre-
dictions in the x and y loading directions using both curve-level metrics and population
envelopes to assess fidelity in nonlinear response and variability (Section 4.3). Finally,
we examine whether the surrogate preserves microstructural heterogeneity across SVE
partitions by comparing predicted and reference distributions and spatial field patterns
(Sections 4.4).

4.1 Volume-fraction generalization

To assess whether a single model can generalize across martensite volume frac-
tions without sacrificing accuracy, we compare two training strategies: (i) v_f -specific
models trained on a single regime, and (ii) a *unified* model trained on the pooled
dataset spanning all v_f . Let \mathcal{D}_{v_f} denote the subset of SVEs with volume fraction
 $v_f \in \{0.10, 0.45, 0.90\}$. We train specialized models \mathcal{M}_{v_f} on \mathcal{D}_{v_f} and a unified model
 \mathcal{M}_{all} on $\bigcup_{v_f} \mathcal{D}_{v_f}$. For evaluation, we report two complementary coefficients of determi-
nation: (a) *within-regime* (conditional) scores, $R^2(\mathcal{M}, \mathcal{D}_{v_f}^{\text{test}})$, computed on test SVEs
restricted to a fixed v_f , and (b) *pooled* scores computed after concatenating all test

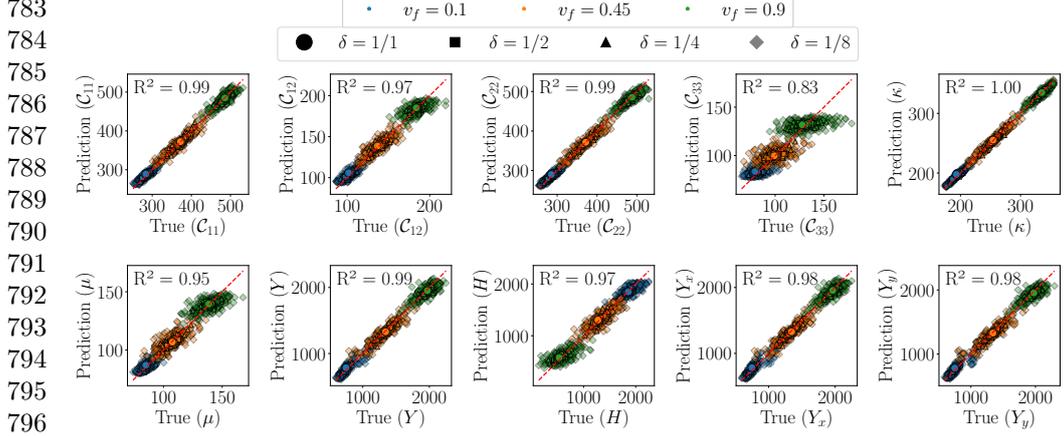
SVEs across regimes. For the pooled score using specialized models, each test sample is predicted by its corresponding \mathcal{M}_{v_f} (a post-hoc regime-aware reference), whereas the unified pooled score uses \mathcal{M}_{all} for all samples.

Before comparing regime-specific and unified training strategies, it is useful to note that direct cross-volume-fraction transfer is generally poor when a model is trained on a single composition regime and evaluated on another. As an illustrative example, Appendix Figure A1 shows parity plots for a model trained only on $v_f = 0.10$ and tested on $v_f = 0.45$, where several scalar QoIs exhibit substantial degradation. This behavior motivates the volume-fraction generalization study below and supports the use of a unified model trained on pooled v_f data.

Table 2 shows that unified training does not degrade performance in a systematic way: the conditional R^2 values of \mathcal{M}_{all} remain close to those of the specialized \mathcal{M}_{v_f} models across most outputs, while eliminating the need to maintain separate models for each composition regime. Moreover, the pooled R^2 from the unified model is comparable to that obtained by combining specialized models, indicating that a single set of learned microstructure features can support accurate predictions across volume fractions. The main exception is the direction-dependent shear component C_{33} , which remains challenging for both strategies; however, even for this case, the unified pooled score is competitive, consistent with the broader observation that shear-dominated response is particularly sensitive to finite-SVE anisotropy effects, which is explained in detail in the next section.

Table 2 should therefore be interpreted in two complementary ways. The unified model remains competitive with the regime-specific models within individual v_f regimes, but the strongest gains are observed in the pooled evaluation, where the model benefits from learning the broader variation in responses across compositions. Accordingly, the main result is not that pooled training is uniformly superior across all regimes, but that a single unified surrogate preserves most within-regime fidelity

737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782



799
800 while also capturing the dominant cross-regime trends. Here, the pooled R^2 primar-
801 ily measures fidelity across the broader response range obtained by combining all v_f
802 regimes, whereas the conditional R^2 is the stricter measure of within-regime sensitivity
803 at fixed v_f .
804
805
806

807 4.2 Elastic and plastic QoIs

808
809 Figure 5 reports parity plots for the ten scalar QoIs. All stiffness and strength units
810 in the manuscript are MPa. Overall, the surrogate achieves very high accuracy for
811 both elastic and plastic properties, indicating that the learned SVE embedding is
812 sufficiently expressive to resolve microstructure-induced trends beyond simple compo-
813 sitional averaging. The best-performing elastic targets are C_{11} and C_{22} , followed by
814 C_{12} . The derived isotropic bulk modulus κ exhibits the highest agreement among the
815 reported scalars with an R^2 higher than 0.995, while the isotropic shear modulus μ
816 remains strongly correlated ($R^2 = 0.95$). These results suggest that the model cap-
817 tures both the principal stiffness response and derived moduli consistently within a
818 single multitask representation.
819
820
821
822
823
824
825
826
827
828

Plastic scalars show a similarly strong performance. The yield strength Y is predicted with $R^2 = 0.99$ and the hardening descriptor H with $R^2 = 0.97$. Directional yield strengths are predicted with nearly identical accuracy in both loading directions (Y_x : $R^2 = 0.98$, Y_y : $R^2 = 0.98$); this is particularly important for maintaining directional consistency when SVEs exhibit apparent anisotropy due to finite sampling and phase topology. The parity scatter plots for these plastic quantities remain concentrated near the identity line, indicating that the surrogate preserves relative ranking across SVEs and does not systematically bias toward the mean.

Among the elastic scalars, shear modulus C_{33} exhibits the largest dispersion ($R^2 = 0.83$), despite the high accuracy obtained for isotropic shear modulus μ ($R^2 = 0.95$). This contrast suggests that the dominant challenge is not capturing the overall elastic stiffness level, but resolving the direction-dependent shear response that emerges at finite SVE sizes. In particular, C_{33} is more strongly influenced by the microstructure topology, including phase connectivity, interface networks, and local constraint pathways, than the angle-averaged response encoded in μ . At small-to-intermediate SVE sizes, these topology-driven effects can also manifest as apparent anisotropy, even when the parent microstructure is statistically isotropic, thereby increasing sample-to-sample variability and reducing the predictability of a single shear component. The relatively higher dispersion of C_{33} is also influenced by the MBC used in the crystal plasticity simulations. In the adopted formulation, roller-type constraints are applied on certain faces of the SVE. These constraints work well for the normal loading cases that provide C_{11} and C_{22} . However, C_{33} is obtained from the shear loading case, which appears to be more sensitive to this boundary treatment. As a result, the shear response obtained from the SVE simulations may contain additional variability compared with the normal stiffness components². This interpretation is also consistent

²Other types of MBC [48–50] may provide a more accurate shear response. The corresponding C_{33} values could correlate more strongly with the microstructure and therefore show lower dispersion in the surrogate GNN model. However, the original CP dataset used in this work did not include results obtained with such alternative boundary conditions.

829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874

875 with our prior CPFEM study of the same SVE database, where the shear modulus
876 exhibited smaller convergence rates and larger RVE sizes than the bulk modulus, indi-
877 cating stronger sensitivity of shear-dominated quantities to finite-SVE variability and
878 anisotropy [17].
879

880
881 Finally, the evaluation is designed to measure generalization at the parent-image
882 level. The train and test split is performed on microstructure images with stratification
883 across v_f regimes, and all partition-derived SVEs from a given parent image are kept
884 in the same split as explained in Section 3. The high parity agreement across scalars,
885 therefore, reflects transfer to unseen microstructures rather than interpolation among
886 correlated subwindows.
887
888
889

890

891

892

893

4.3 Stress–strain response

894

895

896

897

898

899

900

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

918

919

920

We next evaluate the vector-valued stress–strain targets, $\sigma_{xx}(\varepsilon)$ and $\sigma_{yy}(\varepsilon)$, each repre-
sented as a 24-dimensional output. Figure 6 summarizes population envelopes (median
and 5–95% bands) across the considered volume-fraction regimes and loading direc-
tions. The surrogate reproduces the median curve shape across the full strain range,
including the initial elastic regime, the yield transition, and subsequent hardening.
This indicates that the shared latent embedding captures the coupled features needed
to predict both elastic slope and post-yield evolution.

A key outcome is the agreement in response variability. The predicted 5–95%
envelopes track the observed spread, suggesting that the surrogate preserves hetero-
geneity across SVEs rather than collapsing to a narrow mean response. This matters
for downstream reliability and damage analyses, where variability and extremes (not
only averages) govern initiation and uncertainty. Visually, the agreement is maintained
across all volume-fraction regimes, supporting the model’s ability to represent both
stiffness/strength shifts and changes in scatter with microstructure.

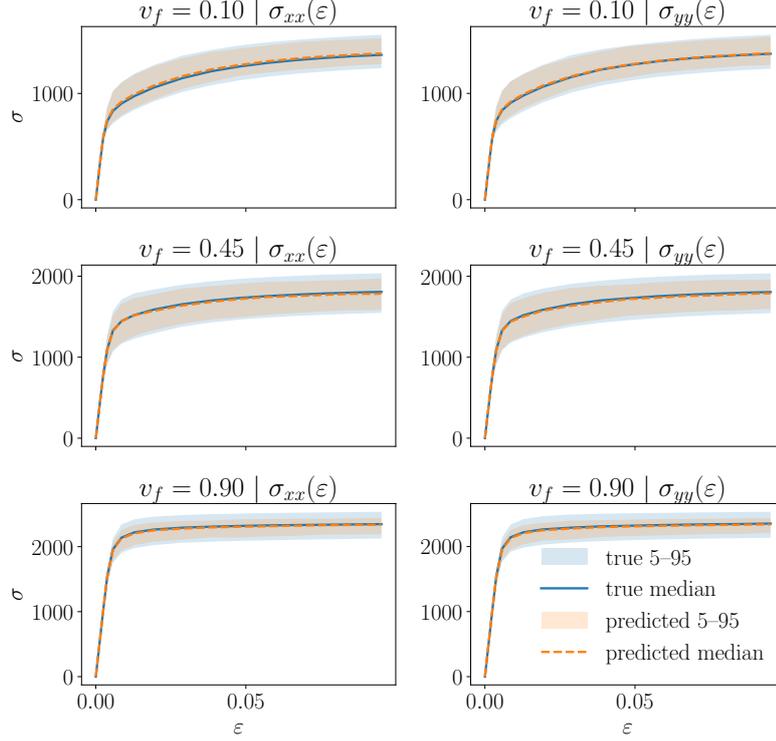


Fig. 6 Population envelopes for predicted and reference stress–strain responses under uniaxial loading in the x and y directions. For each martensite volume-fraction regime ($v_f = 0.10, 0.45,$ and 0.90), the solid curves denote the reference median response and the shaded region indicates the 5–95% interval across SVE realizations.

Quantitative curve metrics are summarized in Table 3. We report R^2 to assess the explained variance of the full discretized response, and MAE/RMSE to quantify the average and quadratic errors in stress magnitude, respectively, across the sampled strain points. To enable comparison across regimes with different stress scales, we also report normalized values relative to the peak reference stress ($\text{NRMSE}_{\text{peak}}$). Finally, the L_∞ error captures the worst-case pointwise deviation across the strain grid, highlighting localized mismatches near the yield surface or in regions of high curvature. Across both loading directions, these metrics indicate comparable accuracy, supporting the conclusion that the surrogate does not over-specialize to a single loading axis and can represent directional stress–strain trajectories within the same multitask model.

921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966

967 **Table 3** Curve-level performance metrics for predicted stress–strain responses in the x
 968 and y loading directions. Metrics are computed over the discretized strain grid for each
 969 response and then aggregated over the evaluation set.

Output	R^2	MAE (MPa)	RMSE (MPa)	NRMSE _{peak} (%)	L_∞ (MPa)
$\sigma_{xx}(\varepsilon)$	0.9561	154.33	194.05	7.17	307.10
$\sigma_{yy}(\varepsilon)$	0.9558	159.94	195.28	7.29	284.60

973
 974

975 4.4 Distributions and spatial consistency

976
 977 Beyond pointwise accuracy, we evaluate whether the surrogate model preserves point-
 978 wise distribution and spatial consistency induced by SVE partitioning. Figure 7
 979 compares predicted and reference probability density functions (PDFs) for represen-
 980 tative scalars (κ , μ , Y , and H) across multiple partition scales. Across regimes, the
 981 predicted PDFs closely follow the reference PDFs in both the value range and the
 982 local maxima, indicating that the model reproduces not only mean shifts but also the
 983 variability structure associated with finite SVEs. This is a nontrivial requirement for
 984 microstructure surrogates, since distributional fidelity is typically more stringent than
 985 the minimization of average error.

986
 987 We also assess the spatial consistency of homogenized fields using subdomain maps
 988 shown in Figure 8. The predicted fields reproduce the dominant spatial patterns across
 989 subwindows, and the corresponding error maps remain localized rather than exhibiting
 990 global bias. This behavior suggests that the model learns microstructure-dependent
 991 heterogeneity and spatial consistency that are consistent across partition levels. This
 992 is particularly relevant when partitioned predictions are later assembled into larger
 993 domains within a multiscale method, as in [51].

994
 995 Taken together, the pointwise distribution and field-level comparisons indicate that
 996 the surrogate model accurately represents the underlying material heterogeneity, which
 997 is important for downstream damage and fatigue modeling. In such settings, localized
 998 response patterns and the tails of property distributions can dominate the risk of crack
 999
 1000
 1001
 1002
 1003
 1004
 1005
 1006
 1007
 1008
 1009
 1010
 1011
 1012

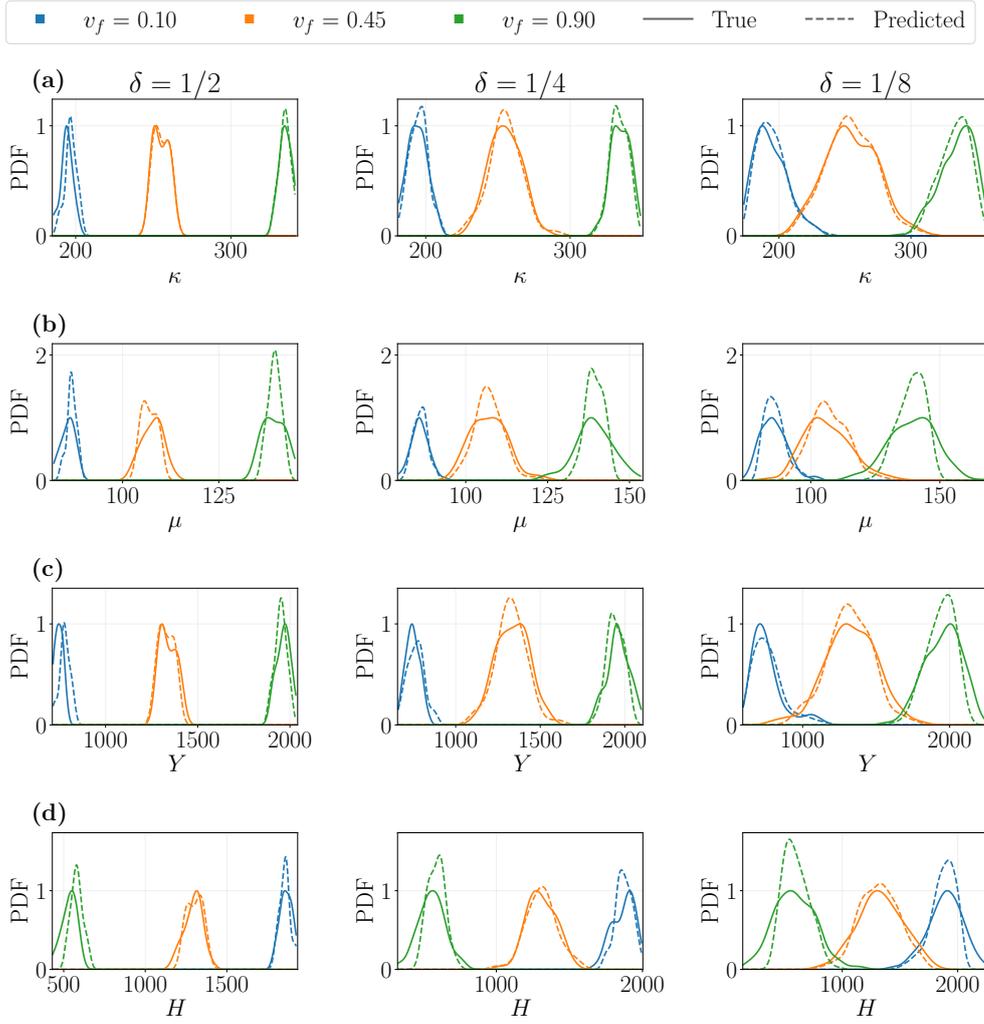
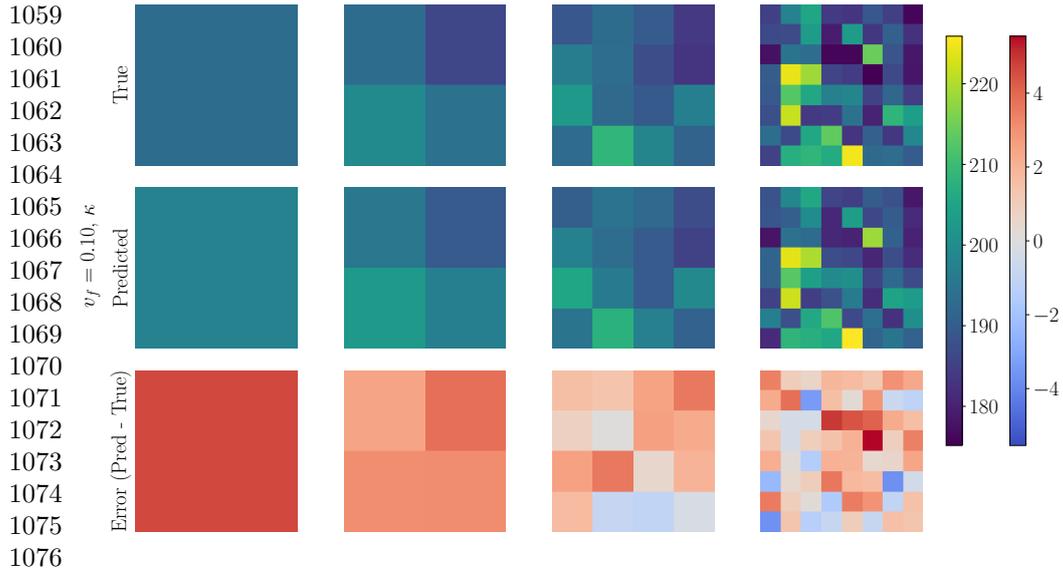


Fig. 7 Comparison of predicted and reference probability density functions (PDFs) for representative scalar QoIs across SVE partition scales and martensite volume fractions. Solid lines denote reference PDFs and dashed lines denote surrogate-predicted PDFs for (a) κ , (b) μ , (c) Y , and (d) H at partition scales $\delta \in \{1/2, 1/4, 1/8\}$.

nucleation, even when global averages appear well predicted. Moreover, retaining distributional fidelity enables the construction of statistically consistent random fields for mesoscale simulations, in which one seeks to populate larger structural domains with spatially varying elastic and fracture constitutive parameters while matching the target

1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058



1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077 **Fig. 8** Spatial consistency of surrogate predictions across SVE partitions. For a representative case,
1078 subdomain maps are shown for the reference field, the surrogate-predicted field, and the corresponding
1079 error field for the selected QoI (κ) at progressively finer partition resolutions.

1080
1081 n-point statistics. Such random-field representations are particularly useful for uncer-
1082 tainty propagation, reliability analysis, and large-area fracture/fatigue assessments,
1083 where explicitly resolving microstructure everywhere is infeasible with direct numeri-
1084 cal simulations (DNS), yet preserving the correct variability and spatial heterogeneity
1085 remains essential and feasible, with the proposed SVE-based upscaling approach.
1086
1087
1088
1089

1090 5 Conclusions

1091
1092
1093 This work developed and evaluated a multitask grain-graph surrogate for predict-
1094 ing SVE-level elastoplastic response from dual-phase polycrystal microstructures. The
1095 main conclusions are as follows:
1096
1097

- 1098 1. **Accurate multitask prediction of coupled elastic and plastic scalars.** Par-
1099 ity analyses demonstrate strong performance across most scalar QoIs, including
1100 stiffness components, isotropic moduli, and plastic descriptors (yield strength and
1101
1102
1103
1104

hardening). The consistent accuracy across both elastic and plastic targets indi-	1105
cates that a shared microstructure embedding can capture coupled trends without	1106
training independent models for each QoI.	1107
	1108
	1109
2. High-fidelity stress–strain trajectory prediction with directional con-	1110
sistency. The surrogate predicts discretized stress–strain responses in both x	1111
and y loading directions with strong agreement in median behavior and variabil-	1112
ity envelopes. Curve-level metrics (Table 3) indicate that the overall error and	1113
explained variance are comparable across directions, supporting stable directional	1114
generalization within a single multitask architecture.	1115
	1116
	1117
	1118
	1119
3. Distributional and spatial agreement under SVE partitioning. Predicted	1120
PDFs closely match the reference PDFs across partition scales and volume-	1121
fraction regimes. Moreover, partition-level field maps reproduce the dominant	1122
spatial patterns and exhibit localized residual errors. These results show that the	1123
model preserves microstructure-driven heterogeneity rather than collapsing to a	1124
mean response, which is essential when downstream analyses depend on pointwise	1125
distribution and spatial consistency.	1126
	1127
	1128
	1129
	1130
	1131
4. Implications for damage and fatigue modeling at larger scales. The	1132
surrogate provides rapid microstructure-conditioned constitutive predictions that	1133
can support high-throughput screening, uncertainty propagation, and statistical	1134
characterization of mesoscale property fields. In particular, the demonstrated dis-	1135
tributional fidelity enables construction of statistically consistent random fields	1136
for larger-domain simulations, where spatially varying constitutive parameters or	1137
strength proxies are required but explicit microstructure resolution everywhere is	1138
impractical. The ability to preserve tail behavior and spatial heterogeneity is partic-	1139
ularly relevant to crack nucleation and early damage processes driven by localized	1140
extremes.	1141
	1142
	1143
	1144
	1145
	1146
	1147
	1148
	1149
	1150

1151 Overall, the proposed grain-graph surrogate demonstrates that physics-informed
1152 microstructure descriptors combined with message passing and multitask learning
1153 can accurately reproduce both scalar homogenized properties and full nonlinear
1154 stress–strain trajectories, while maintaining the variability structure induced by finite
1155 SVEs. These capabilities align with the needs of modern fracture and fatigue mod-
1156 eling workflows that increasingly rely on AI-assisted coarse-graining and large-scale
1157 computational studies.
1158
1159
1160

1161
1162
1163

1164 **Author contribution**

1165
1166 **Erdem Caliskan:** Writing – original draft, Visualization, Investigation, Validation,
1167 Methodology, Software, Data curation, Formal analysis. **Anik Das Anto:** Writing –
1168 review and editing, Data curation, Investigation, Validation, Methodology, Software.
1169
1170 **Massimiliano Lupo Pasini:** Writing – review and editing, Conceptualization, Fund-
1171 ing acquisition, Methodology. **Stephanie TerMaath:** Writing – review and editing,
1172 Conceptualization, Project administration, Funding acquisition, Methodology. **Reza**
1173 **Abedi:** Writing – review and editing, Supervision, Conceptualization, Formal analysis,
1174 Project administration, Funding acquisition, Methodology.
1175
1176
1177
1178
1179

1180
1181

1182 **Code availability**

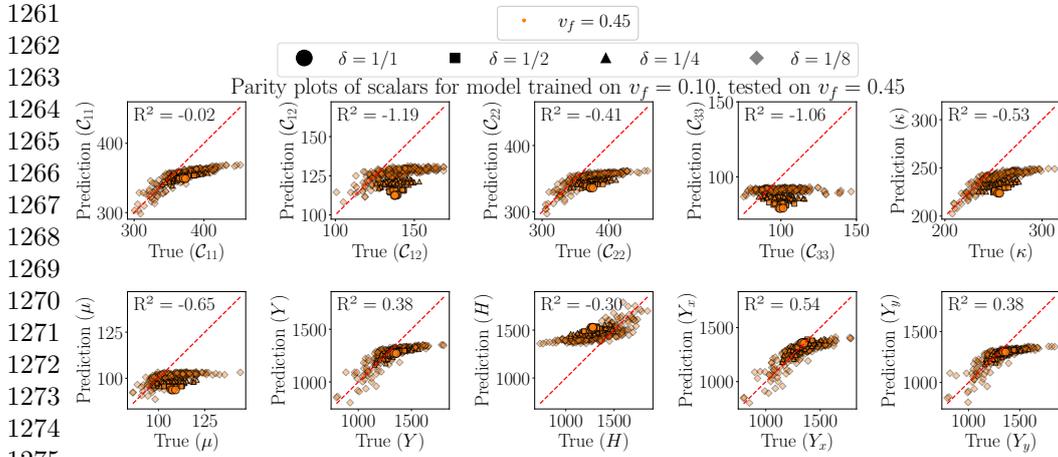
1183
1184 The open-source HydraGNN library [44], used to generate the numerical results
1185 described in this article, can be accessed at the following GitHub repository: https://github.com/ORNL/HydraGNN/tree/cp2d_paper. The scripts used in this work are
1186 located under examples/cp2d.
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196

Data availability	1197
	1198
The datasets generated and analyzed during this study have been deposited in Zenodo under the reserved DOI 10.5281/zenodo.19041349. The repository will be made publicly available upon acceptance/publication of this manuscript.	1199
	1200
	1201
	1202
	1203
	1204
	1205
Ethics declaration	1206
	1207
Not applicable.	1208
	1209
	1210
Declaration of Competing Interest	1211
	1212
	1213
Erdem Caliskan reports that financial support was provided by Oak Ridge Institute for Science and Education (ORISE). The other authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.	1214
	1215
	1216
	1217
	1218
	1219
	1220
	1221
Acknowledgements	1222
	1223
	1224
Massimiliano Lupo Pasini and Erdem Caliskan thank Dr. Vladimir Protopopescu for his valuable feedback in the preparation of the manuscript.	1225
	1226
	1227
This research is sponsored by the Artificial Intelligence Initiative as part of the Laboratory Directed Research and Development (LDRD) Program of Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the US Department of Energy under contract DE-AC05-00OR22725. This work used resources of the Oak Ridge Leadership Computing Facility, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725, under ASCR Leadership Computing Challenge (ALCC) award LRN070.	1228
	1229
	1230
	1231
	1232
	1233
	1234
	1235
	1236
	1237
	1238
	1239
	1240
	1241
	1242

1243 This research was supported in part by the Air Force Office of Scientific Research
 1244 (AFOSR) under grant number FA9550-22-1-0359. Erdem Caliskan acknowledges par-
 1245 tial support by an appointment to the Oak Ridge National Laboratory GRO Program,
 1246 sponsored by the U.S. Department of Energy and administered by the Oak Ridge
 1247 Institute for Science and Education (ORISE).

1251 The authors would also like to thank the University of Tennessee Infrastruc-
 1252 ture for Scientific Applications and Advanced Computing (ISAAC). A portion of the
 1253 computation for this work was performed with ISAAC computational resources.

1258 Appendix A Cross-volume-fraction transfer example



1276 **Fig. A1** Parity plots for scalar QoIs for a model trained on $v_f = 0.10$ and evaluated on $v_f = 0.45$.
 1277 Marker shape denotes the SVE partition scale δ . The noticeable loss of accuracy across several outputs
 1278 indicates limited direct transferability between composition regimes, motivating the pooled volume-
 1279 fraction training strategy discussed in Section 4.1.

1283 References

- 1284
 1285
 1286 [1] K. Sab, On the homogenization and the simulation of random materials. *Eur. J.*
 1287 *Mech., A/Solids* **11**(5), 585–607 (1992)
 1288

[2] M. Ostoja-Starzewski, X. Du, Z. Khisaeva, W. Li, Comparisons of the size of the representative volume element in elastic, plastic, thermoelastic, and permeable random microstructures. <i>International Journal for Multiscale Computational Engineering</i> 5 (2) (2007)	1289 1290 1291 1292 1293 1294 1295
[3] S. Nemat-Nasser, M. Hori, <i>Micromechanics: overall properties of heterogeneous materials</i> , vol. 37 (Elsevier, Amsterdam, The Netherlands, 2013)	1296 1297 1298 1299
[4] S. Stapleton, L. Appel, J.W. Simon, S. Reese, Representative volume element for parallel fiber bundles: Model and size convergence. <i>Composites Part A: Applied Science and Manufacturing</i> 87 , 170–185 (2016)	1300 1301 1302 1303 1304 1305
[5] M. Ostoja-Starzewski, Random field models of heterogeneous materials. <i>International Journal of Solids and Structures</i> 35 (19), 2429–2455 (1998)	1306 1307 1308 1309
[6] T. Kanit, S. Forest, I. Galliet, V. Mounoury, D. Jeulin, Determination of the size of the representative volume element for random composites: statistical and numerical approach. <i>International Journal of solids and structures</i> 40 (13-14), 3647–3679 (2003)	1310 1311 1312 1313 1314 1315 1316 1317
[7] I. Gitman, H. Askes, L. Sluys, Representative volume: Existence and size determination. <i>Engineering fracture mechanics</i> 74 (16), 2518–2534 (2007)	1318 1319 1320 1321
[8] C. Pélissou, J. Baccou, Y. Monerie, F. Perales, Determination of the size of the representative volume element for random quasi-brittle composites. <i>International Journal of Solids and Structures</i> 46 (14-15), 2842–2855 (2009)	1322 1323 1324 1325 1326 1327
[9] R. Abedi, J. Garrard, K. Acton, Elastic and strength properties of statistical volume elements: Determination of isotropic and homogeneous size limits. <i>Computers & structures</i> 277 , 106959 (2023)	1328 1329 1330 1331 1332 1333 1334

- 1335 [10] S.C. Baxter, L.L. Graham, Characterization of random composites using moving-
1336 window technique. *Journal of Engineering Mechanics* **126**(4), 389–397 (2000)
1337
1338
- 1339 [11] L. Huyse, M.A. Maes, Random field modeling of elastic properties using homog-
1340 enization. *Journal of Engineering Mechanics* **127**(1), 27–36 (2001)
1341
1342
- 1343 [12] M. Ostoja-Starzewski, Material spatial randomness: From statistical to repre-
1344 sentative volume element. *Probabilistic engineering mechanics* **21**(2), 112–132
1345 (2006)
1346
1347
1348
- 1349 [13] N. Tregger, D. Corr, L. Graham-Brady, S. Shah, Modeling the effect of mesoscale
1350 randomness on concrete fracture. *Probabilistic Engineering Mechanics* **21**(3),
1351 217–225 (2006)
1352
1353
1354
- 1355 [14] M. Salmi, F. Auslender, M. Bornert, M. Fogli, Apparent and effective mechanical
1356 properties of linear matrix-inclusion random composites: Improved bounds for
1357 the effective behavior. *International Journal of Solids and Structures* **49**(10),
1358 1195–1211 (2012)
1359
1360
1361
- 1362 [15] J.M. Garrard, R. Abedi, Statistical volume element averaging scheme for fracture
1363 of quasi-brittle materials. *Computers and Geotechnics* **117**, 103229 (2020)
1364
1365
1366
- 1367 [16] J.M. Garrard, R. Abedi, Statistical volume elements for the characterization
1368 of angle-dependent fracture strengths in anisotropic microcracked materials.
1369 *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part B:*
1370 *Mechanical Engineering* **6**(2), 021008 (2020)
1371
1372
1373
- 1374 [17] A.D. Anto, R. Fleishel, S. TerMaath, R. Abedi, Size dependency of elastic
1375 and plastic properties of metallic polycrystals using statistical volume elements.
1376 *Applied Sciences* **14**(18), 8207 (2024)
1377
1378
1379
1380

- [18] L. Graham, S. Baxter, Simulation of local material properties based on moving-window gmc. *Probabilistic Engineering Mechanics* **16**(4), 295–305 (2001) 1381
1382
1383
1384
- [19] L. Graham-Brady, E. Siragy, S. Baxter, Analysis of heterogeneous composites based on moving-window techniques. *Journal of engineering mechanics* **129**(9), 1054–1064 (2003) 1385
1386
1387
1388
1389
1390
- [20] K. Acton, L. Graham-Brady, *Fitting an Anisotropic Yield Surface Using the Generalized Method of Cells* (Springer Netherlands, Dordrecht, 2009), pp. 27–41 1391
1392
1393
1394
- [21] K. Acton, L. Graham-Brady, Meso-scale modeling of plasticity in composites. *Computer methods in applied mechanics and engineering* **198**(9-12), 920–932 (2009) 1395
1396
1397
1398
1399
1400
- [22] K. Acton, L. Graham-Brady, Elastoplastic mesoscale homogenization of composite materials. *Journal of engineering mechanics* **136**(5), 613–624 (2010) 1401
1402
1403
1404
- [23] K.A. Acton, S.C. Baxter, B. Bahmani, P.L. Clarke, R. Abedi, Voronoi tessellation based statistical volume element characterization for use in fracture modeling. *Computer Methods in Applied Mechanics and Engineering* **336**, 135–155 (2018) 1405
1406
1407
1408
1409
1410
- [24] A. Asgharzadeh, T. Park, S. Nazari Tiji, F. Pourboghrat, A computationally efficient multiscale, multi-phase modeling approach based on cpem to assess the effect of second phase particles on mechanical properties. *Crystals* **13**(8), 1199 (2023) 1411
1412
1413
1414
1415
1416
1417
1418
- [25] J.N. Fuhg, L. van Wees, M. Obstalecki, P. Shade, N. Bouklas, M. Kasemer, Machine-learning convex and texture-dependent macroscopic yield from crystal plasticity simulations. *Materialia* **23**, 101446 (2022) 1419
1420
1421
1422
1423
1424
1425
1426

- 1427 [26] M.I. Latypov, S.R. Kalidindi, Data-driven reduced order models for effective
1428 yield strength and partitioning of strain in multiphase materials. *Journal of*
1429 *Computational Physics* **346**, 242–261 (2017)
1430
1431
1432
1433 [27] M.I. Latypov, L.S. Toth, S.R. Kalidindi, Materials knowledge system for nonlinear
1434 composites. *Computer Methods in Applied Mechanics and Engineering* **346**,
1435 180–196 (2019)
1436
1437
1438
1439 [28] M. Tannous, C. Ghnatios, O. Castelnau, P.P. Castañeda, F. Chinesta, Machine
1440 learning-boosted nonlinear homogenization. *Mechanics of Materials* **201**, 105229
1441 (2025)
1442
1443
1444 [29] J.N. Heidenreich, M.B. Gorji, D. Mohr, Modeling structure-property relation-
1445 ships with convolutional neural networks: Yield surface prediction based on
1446 microstructure images. *International Journal of Plasticity* **163**, 103506 (2023)
1447
1448
1449
1450 [30] Y. He, Y. Heider, B. Markert, Embedding an ann-based crystal plasticity model
1451 into the finite element framework using an abaqus user-material subroutine. *arXiv*
1452 preprint arXiv:2410.08214 (2024)
1453
1454
1455
1456 [31] G.D. Pasparakis, L. Graham-Brady, M.D. Shields, Bayesian neural networks for
1457 predicting uncertainty in full-field material response. *Computer Methods in*
1458 *Applied Mechanics and Engineering* **433**, 117486 (2025)
1459
1460
1461
1462 [32] Y. Mao, S. Keshavarz, M.N.T. Kilic, K. Wang, Y. Li, A.C. Reid, W.k. Liao,
1463 A. Choudhary, A. Agrawal, A deep learning-based crystal plasticity finite element
1464 model. *Scripta Materialia* **254**, 116315 (2025)
1465
1466
1467
1468 [33] M.S. Khorrami, P.K. Goyal, J.R. Mianroodi, B. Svendsen, P. Benner, D. Raabe,
1469 Divergence-free neural operators for stress field modeling in polycrystalline
1470 materials. *arXiv preprint arXiv:2408.15408* (2024)
1471
1472

- [34] M.M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, P. Vandergheynst, Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine* **34**(4), 18–42 (2017) 1473
1474
1475
1476
1477
1478
- [35] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, P.S. Yu, A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* **32**(1), 4–24 (2020) 1479
1480
1481
1482
1483
1484
- [36] B. Sanchez-Lengeling, E. Reif, A. Pearce, A.B. Wiltschko, A gentle introduction to graph neural networks. *Distill* **6**(9), e33 (2021) 1485
1486
1487
1488
- [37] E. Caliskan, R. Abedi, M.L. Pasini, Graph neural networks for mechanical property prediction of 2d fiber composites. *Materials & Design* p. 114500 (2025) 1489
1490
1491
1492
- [38] Y. Zhao, H. Li, H. Zhou, H.R. Attar, T. Pfaff, N. Li, A review of graph neural network applications in mechanics-related domains. *Artificial Intelligence Review* **57**(11), 315 (2024) 1493
1494
1495
1496
1497
1498
- [39] M. Dai, M.F. Demirel, Y. Liang, J.M. Hu, Graph neural networks for an accurate and interpretable prediction of the properties of polycrystalline materials. *npj Computational Materials* **7**(1), 103 (2021) 1499
1500
1501
1502
1503
1504
- [40] M. Dai, M.F. Demirel, X. Liu, Y. Liang, J.M. Hu, Graph neural network for predicting the effective properties of polycrystalline materials: A comprehensive analysis. *Computational Materials Science* **230**, 112461 (2023) 1505
1506
1507
1508
1509
- [41] Y. Qin, S. DeWitt, B. Radhakrishnan, G. Biros, Graingnn: A dynamic graph neural network for predicting 3d grain microstructure. *Journal of Computational Physics* **510**, 113061 (2024) 1510
1511
1512
1513
1514
1515
1516
1517
1518

- 1519 [42] G.J. Sim, M.G. Lee, M.I. Latypov, Fip-gnn: Graph neural networks for scalable
1520 prediction of grain-level fatigue indicator parameters. *Scripta Materialia* **255**,
1521 116407 (2025)
1522
1523
1524
- 1525 [43] M. Fey, J.E. Lenssen, Fast graph representation learning with pytorch geometric.
1526 arXiv preprint arXiv:1903.02428 (2019)
1527
1528
- 1529 [44] M. Lupo Pasini, J.Y. Choi, K. Mehta, P. Zhang, R. Weaver, A. Chowdhury, C. Li,
1530 Z. Ye, J. Baker, Hydragnn v4. 0. Tech. rep., Oak Ridge National Laboratory
1531 (ORNL), Oak Ridge, TN (United States) (2025)
1532
1533
1534
- 1535 [45] M.L. Pasini, G.S. Jung, S. Irle, Graph neural networks predict energetic and
1536 mechanical properties for models of solid solution metal alloy phases. *Computa-
1537 tional Materials Science* **224**, 112141 (2023)
1538
1539
- 1540 [46] G. Corso, L. Cavalleri, D. Beaini, P. Liò, P. Veličković, Principal neighbourhood
1541 aggregation for graph nets. *Advances in neural information processing systems*
1542 **33**, 13260–13271 (2020)
1543
1544
1545
- 1546 [47] M.A. Groeber, M.A. Jackson, Dream. 3d: a digital representation environment
1547 for the analysis of microstructure in 3d. *Integrating materials and manufacturing
1548 innovation* **3**(1), 56–72 (2014)
1549
1550
1551
- 1552 [48] C. Huet, Application of variational concepts to size effects in elastic heterogeneous
1553 bodies. *Journal of the Mechanics and Physics of Solids* **38**(6), 813–841 (1990)
1554
1555
- 1556 [49] S. Hazanov, C. Huet, Order relationships for boundary conditions effect in hetero-
1557 geneous bodies smaller than the representative volume. *Journal of the Mechanics
1558 and Physics of Solids* **42**(12), 1995–2011 (1994)
1559
1560
1561
1562
1563
1564

[50] S. Hazanov, M. Amieur, On overall properties of elastic heterogeneous bodies smaller than the representative volume. <i>International Journal of Engineering Science</i> 33 (9), 1289–1301 (1995)	1565 1566 1567 1568 1569 1570
[51] B. Bahmani, M. Yang, A. Nagarajan, P.L. Clarke, S. Soghrati, R. Abedi, Automated homogenization-based fracture analysis: Effects of sive size and boundary condition. <i>Computer Methods in Applied Mechanics and Engineering</i> 345 , 701–727 (2019)	1571 1572 1573 1574 1575 1576 1577 1578 1579 1580 1581 1582 1583 1584 1585 1586 1587 1588 1589 1590 1591 1592 1593 1594 1595 1596 1597 1598 1599 1600 1601 1602 1603 1604 1605 1606 1607 1608 1609 1610