

Grain-level micromechanical modeling and assessment of fatigue-critical pores using graph neural networks

Luca Loiodice¹, Krzysztof S. Stopka¹, Yixuan Sun^{*2}, Guang Lin², and Michael D. Sangid^{†1}

¹School of Aeronautics and Astronautics, Purdue University, West Lafayette, 47907, IN, USA

²Department of Mathematics and Mechanical Engineering, Purdue University, West Lafayette, 47907, IN, USA

Supplementary Information

Additional information is provided that details and supports the development of the **GFF-MAP** (Grain-level Fatigue Failure – Micromechanical Assessment of Pores) framework.

Virtual microstructure model generation

The workflow used to generate the microstructure models employed in this study is illustrated in Figure S1. Statistically equivalent virtual microstructures (SEVMs) were generated using DREAM.3D [1], based on the log-normal grain size distribution obtained from electron backscatter diffraction characterization of additively manufactured IN718 specimens (mean grain size $50.5 \pm 58.9 \mu\text{m}$). 620 pores were selected from a catalog of approximately 7000 pores extracted from three AM IN718 specimens characterized by X-ray computed tomography, including two specimens containing lack-of-fusion (LOF) defects and one containing keyhole defects. For more information about the material pedigree or the characterization methods, the reader is referred to Ref. [2]. Each pore was inserted at the center of an SEVM as a void phase. For every pore, four distinct SEVMs were simulated in order to evaluate how variations in the surrounding microstructure influence the fatigue behavior and criticality of the pore.




Figure S1: **Workflow for the generation of the microstructure models employed in this study.** Statistically equivalent virtual microstructure (SEVM) models are generated based on AM IN718 grain statistics, pores are sampled from X-ray computed tomography characterization and then each of them is inserted at the center of four unique SEVMs.

^{*}Current affiliation: Argonne National Laboratory, Lemont, 60439, IL, USA

[†]Corresponding author. Email: msangid@purdue.edu

Figure S2 shows the distributions of the equivalent diameter and sphericity for the full pore population characterized by X-ray computed tomography (top row), as well as the corresponding distributions for the subset of 620 pores selected for this study (bottom row). The sampled pores include both lack of fusion (LOF) and keyhole defects and were chosen to (i) span a wide range of equivalent diameters (15-45 μm) and sphericities (0.325-0.650), enabling the analysis of different pore morphologies, and (ii) produce an approximately uniform distribution in both quantities to obtain a balanced dataset. The selected pores were subsequently partitioned into training, validation, and testing sets, as indicated by the colored bars, while aiming to preserve the balanced distributions of equivalent diameter and sphericity across the three subsets.




Figure S2: **Distributions of the equivalent diameters and sphericities for the pores characterized via X-ray computed tomography.** **a, b** display the full pore population, while **c, d** only the subset of pores selected to generate training data for this study. The selected pores were sampled and then subdivided into training, validation, and testing datasets to ensure approximately uniform distributions of size and shape.

Crystal plasticity modeling

Crystal plasticity simulations were performed using an elasto-viscoplastic fast Fourier transform (EVP-FFT) solver for crystal plasticity, originally proposed by Lebensohn et al. [3], built upon the numerical scheme developed by Moulinec and Suquet [4], and using the hybrid parallel implementation introduced by Rovinelli et al. [5]. EVP-FFT computes stress and strain fields that satisfy equilibrium, compatibility, and the following constitutive equation at every material point x .

$$\sigma(x) = C(x) : \varepsilon^{el}(x) = C(x) : (\varepsilon(x) - \varepsilon^{pl}(x)) = C(x) : (\varepsilon(x) - \varepsilon^{pl,t}(x) - \dot{\varepsilon}^{pl}(x)\Delta t) \quad (1)$$

where $C(x)$ is the stiffness tensor, $\varepsilon^{pl,t}$ denotes the plastic strain tensor at time t , and $\dot{\varepsilon}^{pl}$ is the plastic strain rate tensor, which evolves based on the following flow rule:

$$\dot{\varepsilon}^{pl}(x) = \sum_{\lambda=1}^Q M^{\lambda}(x) \dot{\gamma}^{\lambda}(x) = \dot{\gamma}_0 \sum_{\lambda=1}^Q M^{\lambda}(x) \left(\frac{|M^{\lambda}(x) : \sigma(x)|}{\bar{\tau}(\Gamma, x)} \right)^n \text{sgn}(M^{\lambda}(x) : \sigma(x)) \quad (2)$$

where Q is the total number of slip systems, $\dot{\gamma}^\lambda$ is the shear strain rate on slip system λ , $\dot{\gamma}_0$ is the reference shear strain rate, and n is the rate sensitivity exponent. The Schmid tensor for slip system λ is denoted by $M^\lambda(x)$, while $\bar{\tau}(\Gamma, x)$ represents the critical resolved shear stress (CRSS), which depends on the accumulated shear strain Γ at point x , defined as the sum of shear strains over all slip systems. The resolved shear stress on a slip system λ is related to the stress tensor by:

$$\tau^\lambda(x) = M^\lambda(x) : \sigma(x) = \frac{1}{2} (m^\lambda(x) \otimes n^\lambda(x) + n^\lambda(x) \otimes m^\lambda(x)) : \sigma(x) \quad (3)$$

where m^λ and n^λ denote the slip direction and slip plane normal, respectively. Slip activity leads to strain hardening through dislocation motion. This effect was modeled using a generalized Voce hardening law (Equation 4), selected for its phenomenological nature and limited number of parameters, which helps reduce overall model uncertainty [6].

$$\bar{\tau}(\Gamma, x) = \tau_0 + (\tau_1 + \theta_1 \Gamma(x)) \left[1 - e^{-\frac{\Gamma(x) \theta_0}{\tau_1}} \right] \quad (4)$$

where τ_0 and θ_0 are the initial CRSS and hardening rate, respectively, while τ_1 and θ_1 describe the asymptotic behavior of the material. This crystal plasticity model requires calibration of the reference shear strain rate ($\dot{\gamma}_0$), the rate sensitivity exponent (n), and the hardening parameters (τ_0 , θ_0 , τ_1 , θ_1). These were calibrated against experimental stress-strain curves for AM IN718 using a Monte Carlo approach. For additional details on the calibration process, the reader is referred to Ref. [7].

A non-local averaging (NLA) scheme was employed to regularize micromechanical fields and mitigate spurious values derived from the crystal plasticity simulations, following the methodology presented by Prithivirajan and Sangid [8]. At each material point, a cuboidal averaging volume is defined and rotated to align with each of the modeled slip systems, with its long axis parallel to the slip direction. For each orientation, the micromechanical field is averaged over all voxels within the cuboid that belong to the same grain, and the maximum averaged value across all orientations is taken as the non-local averaged quantity at that point. By constraining the averaging domain to remain within a single grain, the method preserves physically meaningful localization while smoothing noise and maintaining gradients across grain boundaries. The size of the averaging volume was selected as $5 \times 3 \times 3$ voxels (slip direction, in-plane transverse, and out-of-plane directions, respectively), based on the considerations reported by Loiodice et al. [7] for the same material system, ensuring a balance between effective regularization and ability to resolve localized features relevant to fatigue analysis.

As discussed in the main manuscript, the accumulated plastic strain energy density, w^P , was adopted as a fatigue damage indicator parameter. Figure S3 illustrates the evolution of w^P over the 12 applied loading cycles at the fatigue hotspot, defined as the voxel with the highest non-local averaged w^P value in the microstructure at the final simulated loading cycle, for a representative SEVM. The applied uniaxial strain amplitude of 0.4% with $R = 0$, shown in red, leads to a rapid saturation in the growth of w^P . This behavior results from local ratcheting in the material's cyclic response, which causes the growth of w^P to become nearly linear after the first few loading cycles. This characteristic is advantageous for fatigue life prediction studies, in which fatigue life is defined as the number of cycles required for w^P to reach a critical value, w_{cr}^P , since saturation allows for a linear extrapolation of w^P for an arbitrary number of cycles. However, this is beyond the scope of the present work, where w^P is used solely to identify the location of fatigue crack initiation, i.e., the fatigue hotspot within the microstructure.




Figure S3: **Evolution of w^P over 12 loading cycles.** This data refers to the non-local averaged w^P value at the fatigue hotspot of one representative SEVM. The applied uniaxial strain of 0.4% for R=0 is shown in red.

Gas-phase padding sensitivity study

Given the requirement of EVP-FFT to enforce periodic boundary conditions in the input microstructures, the SEVMs were padded with a gas-phase along the X and Y directions and extruded along the Z direction, in order to simulate the presence of free surfaces and mitigate boundary effects within the region of interest of the microstructures. Rovinelli et al. [9] showed that the effect of periodicity completely vanishes in less than 10 voxels from the bounding box limit of the microstructures. Nonetheless, a sensitivity study was conducted to determine the sufficient padding thickness required to obtain accurate crystal plasticity predictions.

One of the 2480 SEVMs was simulated using three different padding thicknesses (4, 8, and 12 voxels), and the results were compared to assess potential differences. The grain-level $w^{P,max}$ micromechanical field was used as a representative variable for this analysis, given its role in identifying the fatigue hotspot location in this study. Figure S4 presents parity plots comparing the predicted $w^{P,max}$ values for the 8-voxel and 4-voxel padding cases (Figure S4a), and for the 8-voxel and 12-voxel padding cases (Figure S4b).




Figure S4: **Comparison of grain-level $w^{P,max}$ computed by crystal plasticity simulations for SEVMs with varying number of voxels used for padding.** **a** provides a comparison between the 8-voxel and 4-voxel padding cases, while **b** provides a comparison between the 8-voxel and 12-voxel padding cases.

Figure S4a shows that the 4-voxel padding case tends to under-predict $w^{P,max}$ in several grains, likely due to periodicity effects influencing grains near the boundaries, resulting in $R^2 = 0.647$. In contrast, Figure S4b demonstrates strong agreement between the 8-voxel and 12-voxel padding cases, with predictions aligning closely and yielding $R^2 = 0.986$. Specifically, in regard to identifying the fatigue hotspot grain, the predictions from the 8-voxel and 12-voxel cases are in accordance, while the 4-voxel case provides a different prediction.

These results indicate that padding the SEVMs with 8 voxels is sufficient to mitigate inaccuracies arising from the imposed periodic boundary conditions, while increasing the padding to 12 voxels provides only marginal improvement. Therefore, an 8-voxel padding thickness was adopted for all SEVMs analyzed in this study.

Graph neural network modeling

For each GNN model developed in this work, a systematic optimization study was performed to identify the most effective architecture, hyperparameters, and combinations of node and edge features. This process was performed using the Weights & Biases (wandb) experiment-tracking platform [10], which enabled automated execution and comparison of thousands of training runs spanning all specified parameter combinations.

The study first focused on identifying the optimal set of node and edge features, exploring all combinations of the following descriptors:

Node features:

- Crystallographic orientation (Euler angles, quaternions, directional stiffness)
- Size (equivalent diameter, volume, number of voxels)
- Morphology (sphericity, surface-to-volume)
- Local morphological changes induced by pore insertion (change in volume, compactness, and concavity of grains in the pore’s neighborhood).

Edge features:

- Relative position (Eulerian distance or full distance vector)
- Crystallographic misorientation
- Shared surface area.

Pores were also represented as nodes with the same feature structure but they were identified by a binary flag, and the properties not applicable to them (e.g., crystallographic orientation) were set to zero.

Once the optimal combination of node and edge features was identified for each GNN model, hyperparameter optimization was conducted. This involved selecting the most suitable loss function, optimizer, and learning-rate scheduler, followed by fine-tuning the learning rate, batch size, number of hidden layers, number of hidden features, and dropout rate. A detailed summary of the final configurations adopted for each GNN model is provided in Table S1. The diversity in optimal architectures, loss functions, and finalized feature sets across models highlights the necessity of such a systematic optimization procedure to tailor each GNN to its specific prediction task and to maximize overall predictive performance.

Table S1: Optimal GNN architecture, hyperparameters, and finalized features for the five GNN models.

	Regression				Classification
	$\sigma^{ZZ,avg}$	$\sigma^{ZZ,max}$	$w^{P,avg}$	$w^{P,max}$	Critical pores
Model Architecture					
Loss Function	MSE	MSE	Smooth L1	Smooth L1	Focal Loss
Hidden Channels	256	256	256	128	128
Number of Layers	4	4	5	5	3
Hyperparameters					
Learning Rate	0.00040	0.00776	0.00160	0.00786	0.00030
Batch Size	16	8	8	8	32
Dropout Rate	0.115	0.187	0.225	0.113	0.162
Alpha (α) - Focal loss	—	—	—	—	0.82
Gamma (γ) - Focal loss	—	—	—	—	5.0
Features					
Node Features	Set A	Set B	Set B	Set A	Set C
Edge Features	Set X	Set Y	Set Y	Set Y	Set Y

Set A: Quaternions, Equivalent diameters, Sphericity

Set B: Quaternions, Sphericity, Surface/volume ratio

Set C: Sphericity, Surface/volume ratio, Volume change *, Concavity change *, Compactness change *

Set X: Distance vector, Shared surface area, Misorientation

Set Y: Distance vector, Shared surface area

Prediction of grain-level micromechanical fields

A detailed comparison between the predictions of the regression GNN models and the corresponding ground truth values obtained from the crystal plasticity simulations is presented in Figure S5 for $\sigma^{ZZ,avg}$, Figure S6 for $\sigma^{ZZ,max}$, Figure S7 for $w^{P,avg}$, and Figure S8 for $w^{P,max}$.




Figure S5: **Comparison of $\sigma^{ZZ,avg}$ GNN predictions against crystal plasticity predictions over the testing dataset.** **a, b** include all grains from all microstructures in the testing dataset, while **c, d** include only the subset of grains that immediately neighbor the pore.




Figure S6: **Comparison of $\sigma^{ZZ,max}$ GNN predictions against crystal plasticity predictions over the testing dataset.** **a, b** include all grains from all microstructures in the testing dataset, while **c, d** include only the subset of grains that immediately neighbor the pore.




Figure S7: **Comparison of $w^{P,avg}$ GNN predictions against crystal plasticity predictions over the testing dataset.** **a, b** include all grains from all microstructures in the testing dataset, while **c, d** include only the subset of grains that immediately neighbor the pore.




Figure S8: **Comparison of $w_{GNN}^{P,max}$ GNN predictions against crystal plasticity predictions over the testing dataset.** **a, b** include all grains from all microstructures in the testing dataset, while **c, d** include only the subset of grains that immediately neighbor the pore.

These figures present the same results discussed in the main manuscript, but with a higher level of detail. In addition to the heatmap parity plots, each figure also includes a scatter representation, allowing individual data points to be visualized explicitly. This enables a clearer visualization of outliers and systematic deviations from the ideal 45° line. A systematic decrease in model performance is observed when transitioning from $\sigma^{ZZ,avg}$ to $\sigma^{ZZ,max}$, and subsequently to $w^{P,avg}$ and $w^{P,max}$. As the predicted quantities become increasingly localized and sensitive to micromechanical heterogeneity, the predictions become more challenging, resulting in a broader scatter and a larger deviation of the data points from the ideal parity line. This behavior is especially pronounced for $w^{P,avg}$ and $w^{P,max}$, where the higher number of outliers is more clearly revealed in the scatter plots. Additionally, the GNN models tend to exhibit reduced predictive accuracy for grains experiencing negligible plasticity, where w^P values are close to zero, and particularly below $10^{-2} MJ/m^3$, as evidenced by a shift to the left in the cloud of data points.

In addition to the parity plot analysis, and to complement the visualizations shown in the main manuscript for $\sigma^{ZZ,avg}$ and $w^{P,max}$, the performance of the GNN models for predicting $\sigma^{ZZ,max}$ and $w^{P,avg}$ is examined with spatial visualizations on a representative microstructure from the testing dataset. Figure S9 presents a) the $\sigma^{ZZ,max}$ computed via crystal plasticity simulations, b) the corresponding GNN predictions, and c) the grain-wise difference between the two. For each quantity, a three-dimensional view of the free surfaces is shown in the top row, while a two-dimensional slice through the mid-plane normal to the Z-direction, highlighting the region surrounding the pore, is shown in the bottom row. An analogous visualization is provided in Figure S10 for $w^{P,avg}$. These spatial comparisons provide additional insight into where the GNN predictions are accurate and where discrepancies arise, complementing the information conveyed by the parity plots. These figures illustrate how prediction difficulty increases for more localized micromechanical fields and puts into perspective the trends observed in the regression analysis. Consistent with the findings for $\sigma^{ZZ,avg}$ and $w^{P,max}$ discussed in the main manuscript, the largest prediction errors are typically concentrated in grains immediately adjacent to the pore and near free surfaces, where steep stress and deformation gradients pose a greater challenge for data-driven prediction models.




Figure S9: **Visualization of the $\sigma^{ZZ,max}$ field predicted by crystal plasticity and by the GNN, and the corresponding grain-wise difference, for a representative microstructure from the testing dataset.** **a** Crystal plasticity simulations, **b** GNN model predictions, and **c** Difference between the predictions and simulations. Three-dimensional views of the SEVM are presented in the top row, and two-dimensional slices of the mid-plane normal to the Z -direction, highlighting the region surrounding the pore, are shown in the bottom row.




Figure S10: **Visualization of the $w^{P,avg}$ field predicted by crystal plasticity and by the GNN, and the corresponding grain-wise difference, for a representative microstructure from the testing dataset.** **a** Crystal plasticity simulations, **b** GNN model predictions, and **c** Difference between the predictions and simulations. Three-dimensional views of the SEVM are presented in the top row, and two-dimensional slices of the mid-plane normal to the Z -direction, highlighting the region surrounding the pore, are shown in the bottom row.

References

- [1] Groeber, M. A. & Jackson, M. A. DREAM.3D: A Digital Representation Environment for the Analysis of Microstructure in 3D. *Integrating Materials and Manufacturing Innovation* **3**, 56–72 (2014). URL <https://doi.org/10.1186/2193-9772-3-5>.
- [2] Stopka, K. S. *et al.* Intentionally seeding pores in additively manufactured alloy 718: Process parameters, microstructure, defects, and fatigue. *Additive Manufacturing* **66**, 103450 (2023). URL <https://www.sciencedirect.com/science/article/pii/S2214860423000635>.
- [3] Lebensohn, R. A., Kanjarla, A. K. & Eisenlohr, P. An elasto-viscoplastic formulation based on fast Fourier transforms for the prediction of micromechanical fields in polycrystalline materials. *International Journal of Plasticity* **32-33**, 59–69 (2012). URL <https://www.sciencedirect.com/science/article/pii/S0749641911001951>.
- [4] Moulinec, H. & Suquet, P. A numerical method for computing the overall response of nonlinear composites with complex microstructure. *Computer Methods in Applied Mechanics and Engineering* **157**, 69–94 (1998). URL <https://linkinghub.elsevier.com/retrieve/pii/S0045782597002181>.
- [5] Rovinelli, A., Proudhon, H., Lebensohn, R. A. & Sangid, M. D. Assessing the reliability of fast Fourier transform-based crystal plasticity simulations of a polycrystalline material near a crack tip. *International Journal of Solids and Structures* **184**, 153–166 (2020). URL <https://linkinghub.elsevier.com/retrieve/pii/S0020768319301052>.
- [6] Tome, C., Canova, G. R., Kocks, U. F., Christodoulou, N. & Jonas, J. J. The relation between macroscopic and microscopic strain hardening in F.C.C. polycrystals. *Acta Metallurgica* **32**, 1637–1653 (1984). URL <https://www.sciencedirect.com/science/article/pii/0001616084902220>.
- [7] Loiodice, L., Stopka, K. S. & Sangid, M. D. A Workflow to Accelerate Microstructure-Sensitive Fatigue Life Predictions. *Advanced Engineering Materials* **n/a**, e202502138 (2025). URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/adem.202502138>. eprint: <https://advanced.onlinelibrary.wiley.com/doi/pdf/10.1002/adem.202502138>.
- [8] Prithivirajan, V. & Sangid, M. D. The role of defects and critical pore size analysis in the fatigue response of additively manufactured IN718 via crystal plasticity. *Materials & Design* **150**, 139–153 (2018). URL <https://linkinghub.elsevier.com/retrieve/pii/S0264127518302946>.
- [9] Rovinelli, A. *et al.* Assessing reliability of fatigue indicator parameters for small crack growth via a probabilistic framework. *Modelling and Simulation in Materials Science and Engineering* **25**, 045010 (2017). URL <https://dx.doi.org/10.1088/1361-651X/aa6c45>.
- [10] Biewald, L. Experiment tracking with weights and biases (2020). URL <https://www.wandb.com/>. Software available from wandb.com.