

Predicting peak stresses in microstructured materials using convolutional encoder–decoder learning

Mathematics and Mechanics of Solids
1–22

© The Author(s) 2022

Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/10812865211055504

journals.sagepub.com/home/mms**Ankit Shrivastava***Department of Civil and Environmental Engineering, Carnegie Mellon University, Pittsburgh, PA, USA***Jingxiao Liu***Department of Civil and Environmental Engineering, Stanford University, Stanford, CA, USA***Kaushik Dayal** *Department of Civil and Environmental Engineering; Center for Nonlinear Analysis, Department of Mathematical Sciences; Department of Materials Science and Engineering, Carnegie Mellon University, Pittsburgh, PA, USA***Hae Young Noh***Department of Civil and Environmental Engineering, Stanford University, Stanford, CA, USA*

Received 26 March 2021; accepted 7 October 2021

Abstract

This work presents a machine-learning approach to predict peak-stress clusters in heterogeneous polycrystalline materials. Prior work on using machine learning in the context of mechanics has largely focused on predicting the effective response and overall structure of stress fields. However, their ability to predict peak – which are of critical importance to failure – is unexplored, because the peak-stress clusters occupy a small spatial volume relative to the entire domain, and hence require computationally expensive training. This work develops a deep-learning-based convolutional encoder–decoder method that focuses on predicting peak-stress clusters, specifically on the size and other characteristics of the clusters in the framework of heterogeneous linear elasticity. This method is based on convolutional filters that model local spatial relations between microstructures and stress fields using spatially weighted averaging operations. The model is first trained against linear elastic calculations of stress under applied macroscopic strain in synthetically generated microstructures, which serves as the ground truth. The trained model is then applied to predict the stress field given a (synthetically generated) microstructure and then to detect peak-stress clusters within the predicted stress field. The accuracy of the peak-stress predictions is analyzed using the cosine similarity metric and by comparing the geometric characteristics of the peak-stress clusters against the ground-truth calculations. It is observed that the model is able to learn and predict the geometric details of the peak-stress clusters and, in particular, performed better for higher

Corresponding author:

Hae Young Noh, Department of Civil and Environmental Engineering, Stanford University, Stanford, CA 94305-6104, USA.

Email: noh@stanford.edu

(normalized) values of the peak stress as compared to lower values of the peak stress. These comparisons showed that the proposed method is well-suited to predict the characteristics of peak-stress clusters.

Keywords

Polycrystalline material, microstructures, machine learning, convolutional encoder decoder, von Mises stress, peak-stress clusters

1. Introduction

This paper aims to use a deep-learning-based method to predict peak-stress clusters in heterogeneous linear elastic materials. Specifically, we consider a simplified model of a polycrystalline microstructure as being composed of many grains, and each grain has an anisotropic linear elastic response. The response of each grain is identical, except that the crystal lattice – and hence the elastic stiffness tensor – in each grain has a rotation with respect to some reference grain. In this idealization, the polycrystalline microstructure can be described by a piecewise-constant field $\mathbf{R}(\mathbf{x})$ where $\mathbf{R} \in SO(2)$ is a rotation that maps the lattice in the reference grain to the orientation of the local lattice at the material point located at $\mathbf{x} \in \mathcal{A}$, where $\mathcal{A} \subset \mathbb{R}^2$ is the domain. Then, given the elastic stiffness tensor \mathbf{C}_0 of the reference, the elastic stiffness at a point \mathbf{x} is simply the standard tensor transformation of \mathbf{C}_0 through \mathbf{R} . This provides a piecewise-constant field $\mathbf{C}(\mathbf{x})$ as the heterogeneous stiffness field. This idealization is standard in homogenization, e.g., [1].

The heterogeneity in the material properties generically leads to stress fluctuations and is the subject of numerous works in homogenization, e.g. [2]. The stress fluctuations can lead to stress concentrations and regions of high stress, and it is critical to predict the peak stresses that are the dominant drivers of material failure. Although numerical methods for elasticity, such as the finite element method [3] or Fourier-based schemes [4–7], can solve for the stress field – and the peak stresses as subsequent post-processing – these methods are relatively expensive, particularly if applied to a large number of microstructures drawn from a random distribution. Therefore, machine-learning-based methods have recently been proposed as an alternative that can potentially be less expensive, particularly the process of prediction after training [8–13]. In these methods, the machine-learning models¹ were trained on the data (“ground truth”) obtained from numerical methods to learn the statistical relationship between the microstructure and the corresponding response to applied loads. When used for prediction, these trained models are much faster than the existing numerical schemes.

However, these proposed methods have not been characterized in terms of their ability to predict the peak stresses, as opposed to the effective response that is simply the relation between the applied load and the overall deformation. For instance, an interesting recent contribution is [14] that uses a convolutional long short-term memory (LSTM)-based deep-learning model to predict stress fields in microstructures and then used this to predict the effective stress response (spatially averaged stress) under applied strain. Although the model shows remarkable fidelity and potential, the size of the microstructure input is limited in the number of grains, making it difficult to explore the question of peak stress prediction in a statistical sense. Moreover, training a similar convolutional LSTM model to predict the stress fields on large-sized microstructures will require an extremely large dataset, making it very expensive. Further, an important next step is to characterize the ability of learning-based methods to not only predict averaged response correctly but the peak stresses as well, which are the dominant drivers of material failure.

Although the ability of machine-learning-based methods to predict peak stresses is largely unexplored, it has recently started to receive attention. An important recent contribution is [15], wherein a random forest model was proposed to classify grains in a given microstructure as high stress or not. The input is in the form of specified features such as the number of neighbors, Schmid factors, and so on. Although this approach has shown promising results, it relies on decisions by the user to explicitly provide microstructures’ features as input, that can suffer from user bias. Hence, we propose an approach to this problem that does not require the user to provide explicit features as input.

Specifically, we propose an approach to predict peak-stress clusters in a given microstructure using a combination of a deep-learning-based convolutional encoder–decoder (CED) model and an image processing algorithm. The design of the CED model used to predict stress fields is inspired by convolutional autoencoder (CAE). The autoencoder is an unsupervised deep-learning algorithm that aims to reconstruct a given input. CAE is a specific type of autoencoder that uses convolutional filters to reconstruct their input. The input to




Figure 1. Flow chart of the proposed method to predict peak-stress clusters in a given microstructure. In this work, we use the term “model” to refer to the CED and the word “method” to refer to the complete procedure from predicting stress fields from microstructures using CED to detecting and characterizing peak-stress clusters inside the predicted stress fields.

the CAEs is a single-channel (such as grayscale) or multichannel image (such as RGB images).² As most of the images’ features are locally spatially correlated, the convolutional filters can extract the local features, which a deep fully connected layer cannot capture. Owing to the ability of CAE to learn image–image mapping with arbitrary accuracy, it is heavily used in image compression, de-noising, generation, and segmentation [16, 17].

Typical polycrystalline microstructures under an applied load show peak-stress clusters that are sparse and are governed by the local spatial features of the microstructure as suggested by [18]. For a 2D elasticity problem, a given microstructure and corresponding stress field (von Mises) can be treated as images. Hence, architectures based on convolution operation are appropriate. Therefore, to find a mapping between the microstructure and corresponding stress field, CAE models are used. As the output and input are different for our problem, we refer to our model as CED.

First, the trained CED model is used to predict the stress field for a given microstructure. Then a cluster detection algorithm is used to detect peak-stress clusters within the predicted stress field. The mean squared error (MSE) was used to analyze the accuracy and to critically investigate scenarios where the model makes highly inaccurate predictions. Finally, the accuracy of the peak-stress clusters’ geometric structure was analyzed using the cosine similarity metric and by separately comparing the clusters’ size, shape, orientation, and location in the learning-based predictions against the ground-truth numerical elasticity calculations. Visualization of the filters of the first convolutional layer were used to understand the essential elements of the microstructure as identified by the machine-learning model.

Organization. In Section 2, the basic principles used to design the CED model and algorithm to detect peak-stress clusters are discussed. Section 3 describes the procedure to generate the required dataset and the principle behind designing the architecture of CED. Section 4 discusses the evaluation of our model. In Section 5, the role of filters in the first layer of the CED model is discussed. Concluding remarks are provided in Section 6.

2. CED-based method to predict peak-stress clusters

The peak-stress clusters in a given microstructure are predicted using the approach shown in Figure 1. First, the trained CED model predicts a low-resolution von Mises stress field from a given microstructure. Then model predictions are fed into a cluster detection algorithm to detect peak-stress clusters. The detected peak-stress clusters are then characterized by their location, shape, size, and orientation. The training and designing of the model are explained in later sections.

2.1. Input and ground truth for the model

The CED model’s input is a microstructure of 3D polycrystals in a 2D domain, i.e., we have three Euler angles at every spatial location \mathbf{x} . The input therefore contains information about the orientation and grain boundaries. The ground truth is a low-resolution von Mises stress field. As a large number of samples are required to train the deep-learning model and test its validity, we create a synthetic dataset with known properties. The Dream3D package [19] was used to generate synthetic single-phase microstructures. The package uses a set of predefined statistics for grain geometry and Euler angle distributions to generate these microstructures. A microstructure generated by the package is a three-channel image (see note 2), each channel representing one of the Euler angles, (Φ, Ψ_1, Ψ_2) . The grain boundary structure is a single channel image obtained using the crystal




Figure 2. Images of resolution 128×128 representing Euler angles (Φ , Ψ_1 , Ψ_2) and the grain boundary structure. These images are scaled between $[0, 1]$ before being fed into the model.




Figure 3. (a) von Mises stress field and (b) low-resolution von Mises stress field obtained by piecewise-averaging used as ground truth. This stress field is scaled between $[0, 1]$.

orientation information. Each point of the grain-boundary image is the L_2 -norm of the gradient of Euler angles at that point. The three Euler angle images and the grain-boundary image (Figure 2) are used as the input to the CED model. The primary purpose of including the grain-boundary image was to improve the CED model accuracy. The extra information about the grain boundaries as an input to the model became necessary as the model could not deduce this accurately by itself.

The stress field is computed using linear elasticity for each of the microstructures under an applied macroscopic strain, and the stress field is used to calculate the von Mises stress field. The von Mises stress fields are single-channel images of the same resolution as the microstructures. However, training any deep-learning model with large-sized images will require a vast amount of training data that requires significant computational resources and memory. Hence, we use piecewise-averaged stress fields, which approximates the von Mises stress fields, as shown in Figure 3. These low-resolution stress fields are the ground truth for the model. The piecewise averaging of stress fields preserves the relative character of the peak-stress clusters while making the model training efficient. The CED model architecture is designed considering the size of the microstructures and the ground-truth stress fields.

2.2. Design criteria for CED

The CED model contains two main components, encoder and decoder (Figure 4). Both components consist of a sequence of layers with multiple filters that operate on its input and subsequently pass its output to the next layer. The output from each layer is referred to as the feature map. Details of the operations used in these components are explained in the later sections.

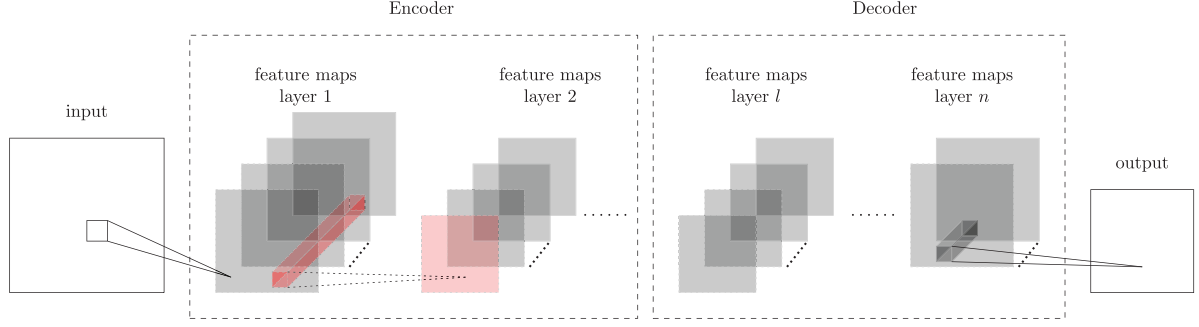


Figure 4. A general framework of the proposed CED model. The input layer consists of a four-channel image representing a microstructure as shown in Figure 2. The output layer consists of a single channel image representing a low-resolution von Mises stress field as shown in Figure 3b. The encoder consists of multiple alternating convolutional and max-pooling layers. The decoder consists of multiple transpose convolution layers.

2.2.1. Encoder. The encoder uses a series of convolutional and max-pooling layers to capture spatial information such as grain boundary, Euler angles, and grain junctions in the neighborhood of a point in the microstructure:

$$\begin{aligned}
 n \times n \times K_{in} &\equiv \text{kernel size of each convolutional filter} \\
 n_{in} \times n_{in} \times K_{in} &\equiv \text{dimension of input to the layer, } \mathbf{C}_{in} \\
 n_{out} &= \frac{n_{in} - n + 2P}{S} + 1 \\
 K_{out} &\equiv \text{number of filters} \\
 n_{out} \times n_{out} \times K_{out} &\equiv \text{dimension of output after convolution, } \mathbf{z} \\
 n_{out} \times n_{out} \times K_{out} &\equiv \text{dimension of output after operating ReLU, } \mathbf{C}_{out}
 \end{aligned} \tag{1}$$

$$\mathbf{z}_{ijk} = \sum_{c=1}^{K_{in}} \sum_{b=1}^n \sum_{a=1}^n \mathbf{w}_{abck} \mathbf{C}_{in(i-a)(j-b)c} + \mathbf{w}_k \quad \forall i, j \in [1, n_{out}] \quad \text{and} \quad k \in [1, K_{out}] \tag{2}$$

\mathbf{w}_{abck} and \mathbf{w}_k are the weights and bias, respectively, of the k th filter in the layer.

It then converts the information into low-resolution with high-description feature maps. A convolutional layer performs convolution operation using multiple filters followed by a nonlinear activation. These filters separately extract the high-level features from the layer input by convolving (weighted dot product) it with a sliding window, as shown in Equation (2). The hyperparameters, such as the number of filters in the layer, size of each filter, padding P on input, and stride S , are predefined. Their choice is explained in Section 2.2.4.

As the convolutional operations are linear, their combination will make the model capture only the linear relationship between the microstructures and the low-resolution stress fields, ignoring any nonlinear relations. Hence, outputs from each convolution operation are fed into a nonlinear activation function. These activation functions make the layer nonlinear, further assisting in extracting the nonlinear features from its input [20]. The output size from the nonlinear activation function is the same as the input size to the function.

For the CED model, ReLU activation

$$\mathbf{C}_{out}_{ijk} = \text{ReLU}(\mathbf{z}_{ijk}) = \max(0, \mathbf{z}_{ijk}) \quad \forall i, j \in [1, n_{out}] \quad \text{and} \quad k \in [1, K_{out}] \tag{3}$$

is used as nonlinear activation. Unlike other activation functions such as Sigmoid and hyperbolic tangent function tanh, ReLU functions do not suffer from the vanishing gradient problem [22]. Here ReLU is also much more computationally efficient [22] making the back-propagation computationally efficient. This combination of convolution and nonlinear activation operations summarizes the information of input (\mathbf{C}_{in}) by detecting its features in the form of feature maps (\mathbf{C}_{out}).

In the encoder, the number of filters in the subsequent layers is increased as more and more layers are added. As the size of the feature maps is decreasing, further reducing the amount of the microstructure information, the number of filters is increased to compensate for the reduction in feature size. However, using a large number

of convolutional filters in deeper layers will increase the number of parameters in the model, which can cause overfitting. Maxpool layer is often used instead of a convolutional layer to avoid a large number of parameters:

$$\begin{aligned}
n_{out} \times n_{out} \times K_{out} &\equiv \text{Dimension of its input, } \mathbf{M_in} \\
m \times m \times 1 &\equiv \text{kernel size of the maxpool filter} \\
m_{out} &= \frac{n_{out} - m + 2P}{S} + 1 \\
\mathbf{M_out}_{ijk} &= \max\{\mathbf{M_in}_{(i-a)(j-b)k} \mid a \in [1, m], b \in [1, m]\} \\
&\quad \forall i, j \in [1, m_{out}] \quad \text{and} \quad k \in [1, C_{out}].
\end{aligned} \tag{4}$$

The maxpool operation, given by Equation (5), separately operates on each of its input $\mathbf{M_in}$ using a sliding window, extracting the most important information while keeping the number of model parameters constant. This operation further reduces the size of its input while keeping important information. During training, maxpool also prevents fluctuations in values of parameters \mathbf{w} in convolutional layers due to a slight change in its input. This further helps in achieving stable convergence while training.

2.2.2. Decoder. The decoder converts feature maps obtained from the Encoder into low-resolution von Mises stress fields using a series of transpose convolution layers, each performing transpose convolution operation, as shown in Equation (7),

$$\begin{aligned}
n \times n \times C_{in} &\equiv \text{kernel size of each transpose convolutional filter} \\
n_{in} \times n_{in} \times C_{in} &\equiv \text{dimension of layer's input, } \mathbf{A} \\
n_{out} &= (n_{in} - 1) \times S + n - 2P \\
\mathbf{z}_{ijk}^t &= \sum_{c=1}^{C_{in}} \sum_{b=1}^n \sum_{a=1}^n \mathbf{w}_{abck}^t \mathbf{A}_{(i+a)(j+b)c} + \mathbf{w}_k^t \quad \forall i, j \in [1, n_{out}] \quad \text{and} \quad k \in [1, C_{out}]
\end{aligned} \tag{6}$$

\mathbf{w}_{abck}^t and \mathbf{w}_k^t are the weights and bias, respectively, of the k th filter in the layer.

The transpose convolution operation decompresses the information from its input using weighted interpolation. For the same reason as explained in Section 2.2.1, the transpose convolution operation is followed by a nonlinear activation. Here ReLU is used as the activation function for all layers except for the last layer.

In the last transpose convolution layer, the transpose convolution operation is followed by a tanh activation function given by

$$\text{tanh}(\mathbf{z}_{ijk}) = \frac{e^{\mathbf{z}_{ijk}} - e^{-\mathbf{z}_{ijk}}}{e^{\mathbf{z}_{ijk}} + e^{-\mathbf{z}_{ijk}}} \tag{8}$$

The model does not suffer from a vanishing gradient by including tanh. The tanh function is used in the last layer of the model because it is more nonlinear than the ReLU. This helps in closely resembling the ground truth with just a few layers, whereas the ReLU being piecewise linear will require more layers to resemble the same ground truth. The tanh is preferred over the Sigmoid as the former has a stronger gradient than the latter, which reduces the vanishing gradient problem while training. The output from the last layer is the prediction of a low-resolution stress field.

2.2.3. Training model parameters. The MSE function is considered as the objective function for training and validating the model. The MSE represents how close the model predictions are to the ground truth for a given input. First, the complete dataset is divided into training, validation, and test sets. The MSE computed over the training set, the validation set, and the test set are referred to as training error, validation error, and test error, respectively. The goal while training the model is to find optimum parameters, \mathbf{w}^* of the CED that minimizes

the training error (Equation 10),

$$\mathbf{w}^* = \arg \min_{\mathbf{w}} \text{MSE}(\mathbf{w}) \quad (9)$$

$$\text{MSE}(\mathbf{w}) = \mathbb{E}_{i \in [1, N]} [\|\text{CED}(\mathbf{X}_i, \mathbf{w}) - \mathbf{Y}_i\|^2] \quad (10)$$

$\text{CED}(\mathbf{X}_i, \mathbf{w}) \equiv$ the model prediction for a given input \mathbf{X}_i

$\mathbf{Y}_i \equiv$ corresponding ground-truth

$N \equiv$ number of samples.

The training input data \mathbf{X} represents a set of microstructures, and the ground truth \mathbf{Y} represents a set of low-dimensional stress fields for the corresponding input.

The MSE is minimized using the mini-batch gradient descent scheme [21]. Under this scheme, the training dataset is divided into sub-samples, called a batch, with a predefined batch size. The MSE over a batch is used to compute the gradients and update the model parameters. Compared with gradient descent, the mini-batch gradient descent helps avoid local minima due to the noisy update, further helping achieve faster convergence. It also avoids the problem of GPU memory restrictions as suffered by standard gradient descent. The mini-batch gradient is chosen over the stochastic gradient descent to achieve smoother training and avoid expensive latency in CPU–GPU data transfer during every update. As compared with the stochastic gradient descent, the mini-batch gradient descent has a higher likelihood to fall into a local minimum [21]. However, this problem can be controlled by tuning the learning rate and the batch size. These parameters help control the amount of noise added during the gradient update, which helps avoid local minima. For the gradient computations, we have used the ADAM scheme [23]. Compared with other gradient computation schemes, ADAM provides faster convergence by adapting the geometry of the objective function.

During the training, the validation and training errors are evaluated regularly to analyze if the model is underfitting or overfitting. Model hyperparameters are tuned based on underfitting and overfitting as explained in Section 2.2.4. Overfitting is also avoided using an early stopping mechanism whereby the model training is stopped as soon as validation error increases.

2.2.4. Hyperparameter selection for the model. We use the behavior of the validation and training errors during the training to select the model hyperparameters. The different values of the model hyperparameters will lead to a different model architecture.

The hyperparameters such as the number of convolutional, maxpool, transpose convolutional layers, and the number of filters in each layer are governed by the validation error. A large number of layers in the encoder can lead to a significant loss in the information about the input as each layer compresses the information. A large number of layers in the decoder can lead to the inclusion of noise during the weighted interpolation. In both cases, the validation error increases. A small number of layers in both the encoder and the decoder can cause both the training error and the validation error to increase (underfitting).

Similarly, in each layer, a small number of filters will have a small number of parameters, which will increase the chance of the model getting stuck in a local minimum causing high validation error. A large number of filters can either lead to an increase in validation error (overfitting) or no significant reduction in validation error while making training expensive.

The kernel sizes of the filters in all the layers are governed by the validation error, the input dimension, and the ground-truth dimension. The stride and the padding of convolutional filters, maxpool filters, and transpose convolutional filters are governed by Equations (1), (5), and (6), respectively. The stride is always less than or equal to kernel size in convolutional layers to ensure that information from every pixel is considered while making predictions.

Hence, optimum values of the hyperparameters are chosen to achieve a minimum validation error. Finally, the architecture that gives the minimum validation error during the training is used to predict the stress fields.

2.3. Detection of peak-stress clusters

The peak-stress clusters are detected inside the stress fields predicted by the trained CED model. We characterize the peak-stress clusters as regions around a peak in the stress fields with values above a certain threshold. The low-resolution stress fields obtained by averaging von Mises stress fields, see Section 2.1, and the predicted




Figure 5. Stress clusters around three different peaks in a predicted stress field for different thresholds.

stress fields from the model are not smooth, and there are small fluctuations. The stress fields are smoothed using a Gaussian filter to avoid small fluctuations. If the smoothing is not performed, small fluctuations in the stress fields will be detected as peaks, which would lead to the erroneous assignment of these small fluctuations as the actual peak. Further, the algorithm used to detect peak-stress clusters around stress peaks will be computationally expensive.

After the smoothing, the stress fields are fed into a peak detection algorithm which then returns the location of multiple peaks inside stress fields. For every peak, the clusters around it are obtained by assigning the value of 1 if the pixel value is higher than the predefined threshold times the corresponding peak stress. The result is a binary image with multiple disconnected regions representing stress clusters. The stress clusters are then separately labeled using the connected component labeling algorithm as shown in Figure 5.

2.4. Characterizing stress clusters

The peak-stress clusters corresponding to each peak are characterized by their location, size, shape, and orientation:

$$\mathbf{z}_1 X^2 + \mathbf{z}_2 XY + \mathbf{z}_3 Y^2 + \mathbf{z}_4 X + \mathbf{z}_5 Y = 1 \quad (11)$$

$$\arg \min_{\mathbf{z}} \|\mathbf{A}\mathbf{z} - \mathbf{b}\|_2^2 \quad (12)$$

$$\text{where, } \mathbf{z} = [\mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3, \mathbf{z}_4, \mathbf{z}_5]^T$$

$$\mathbf{b} = [1, 1, 1, 1, 1]^T$$

$$\mathbf{A}_i = [X_i^2, X_i Y_i, Y_i^2, X_i, Y_i] \quad \forall i \in [1, n]$$

$$n \equiv \text{number of points inside a cluster}$$

$$\theta = \tan^{-1} \left(\frac{\mathbf{z}_3 - \mathbf{z}_1 - \sqrt{(\mathbf{z}_1 - \mathbf{z}_3)^2 + \mathbf{z}_2^2}}{\mathbf{z}_2} \right) \quad (13)$$

$$\text{AR} = \frac{\mathbf{z}_1 + \mathbf{z}_3 + \sqrt{(\mathbf{z}_1 - \mathbf{z}_3)^2 + \mathbf{z}_2^2}}{\mathbf{z}_1 + \mathbf{z}_3 - \sqrt{(\mathbf{z}_1 - \mathbf{z}_3)^2 + \mathbf{z}_2^2}}. \quad (14)$$

As failure initiates at the peak-stress clusters, detailed information of such clusters is important. This characterization provides geometric details of peak-stress clusters inside a stress field. The location of the corresponding peak characterizes the location of a stress cluster. The ratio of the number of pixels in the cluster to the total number of pixels of the stress field characterizes the cluster size. The shape and orientation of the stress clusters are characterized by the axis ratio and orientation, theta of an equivalent ellipse, as shown in Figure 6. The equivalent ellipse of the peak-stress cluster is obtained by the least-square fitting; see [24] for more details. Equation (12) is minimized to obtain the coefficients \mathbf{z}_i in Equation (11). The ellipse (Figure 6) is centered

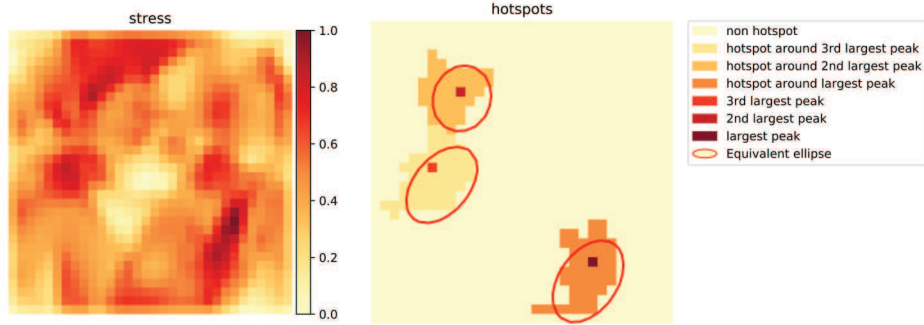


Figure 6. Equivalent ellipse of the stress clusters corresponding to each peak.

Table 1. Dream3D StatsGenerator parameters to generate microstructures. The Primary-Equiaxed parameter represents a single-phase material with grains of shape ellipsoid having almost the same length in all directions. The equivalent sphere diameter of grains in microstructures is randomly selected from a truncated lognormal distribution defined using Mu, Sigma, and Min/Max cutoff. The Bin step size further segregates the lognormal distributions into multiple bins, where each bin will have a corresponding shape distribution with default values set by the Preset Statistic parameter. Finally, the Euler angles in all three directions are randomly generated from a uniform distribution.

Preset statistic	Primary Equiaxed	Mu	2.3
Max cutoff	4	Sigma	0.4
Bin step size	1	Min/Max cutoff	4

around the mean of the n pixels with coordinates X and Y inside the cluster. The ellipse’s orientation, θ , and axis ratio (AR) are then obtained using the coefficients \mathbf{z} as shown in Equations (13) and (14). It is important to note that the right-hand side of Equations 11 is set to 1 in order to avoid a trivial solution. Hence, the obtained \mathbf{z} represents the original coefficients of the ellipse scaled by a constant.

3. Numerical experiments

3.1. Generating the dataset


A dataset of 10,000 samples containing the input and the ground truth are generated as explained in Sections 3.1.1 and 3.1.2. The dataset is further divided into 7000 for training, 1000 for validation, and 2000 for testing. Before being used for training or prediction, the dataset is pre-processed.

Initially, the model was trained using a dataset with 40,000 samples. However, the test error using 40,000 samples was 0.0188618, whereas, for 10,000 samples, it was 0.0197408. Hence, increasing the data did not improve the accuracy significantly but increased the training time by a factor of four, which acted as a bottleneck to the tuning of model architectures. It also required a large amount of memory to store the data and results for post-prediction analysis. Hence, in the later sections, the model training, the hyperparameter tuning, and the evaluation are performed using the dataset with 10,000 samples.

3.1.1. Generating input for the model. For the current problem, microstructures “represented by the Euler angles” of size $128 \times 128 \times 3$ were generated using Dream3D with parameters in Table 1. Many microstructures generated with the dimensions of 64 and 32 were similar in terms of the grain distribution. The dimension of 128 ensured that the generated microstructures are sufficiently different from each other. Next, the information about the grain boundaries inside a microstructure was obtained using the L_2 -norm of the gradient of the Euler angles for each pixel inside the microstructure. The grain boundary information, along with the microstructure, was used as an input. The final dimension of the input was $128 \times 128 \times 4$ (Figure 2).

Table 2. Normalized values of the elastic stiffness tensor components for cubic symmetry

c_{1111}	c_{1122}	c_{2323}
1.0000	0.7209	0.4477

**Figure 7.** A typical microstructure sample generated using DREAM3D in the XY plane. Here \mathcal{A} and $\partial\mathcal{A}$ represent domain and boundary of the microstructure.

3.1.2. *Generating ground truth for the model.* Linear elasticity (Equation 15),

$$\begin{aligned}
 \nabla \cdot \boldsymbol{\sigma}(\mathbf{x}) &= 0 & \forall \mathbf{x} \in \mathcal{A} \\
 \boldsymbol{\sigma}(\mathbf{x}) &= \mathbf{C}(\mathbf{x}) : \boldsymbol{\epsilon}(\mathbf{x}) & \forall \mathbf{x} \in \mathcal{A} \\
 \boldsymbol{\epsilon}(\mathbf{x}) &= \mathbf{e}(\mathbf{u}^*(\mathbf{x})) + \mathbf{E} & \forall \mathbf{x} \in \mathcal{A} \\
 \mathbf{e}(\mathbf{u}^*(\mathbf{x})) &= \frac{\nabla \mathbf{u}^*(\mathbf{x}) + \nabla \mathbf{u}^{*\top}(\mathbf{x})}{2} & \forall \mathbf{x} \in \mathcal{A} \\
 \mathbf{u}^*(\mathbf{x}) & & \text{Periodic in } \mathcal{A}
 \end{aligned} \tag{15}$$

$$\mathbf{E} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 10^{-4} \end{bmatrix}$$

where \mathbf{E} is the applied macroscopic strain, $\mathbf{u}^*(\mathbf{x})$ is the periodic fluctuation displacement field, and $\boldsymbol{\epsilon}(\mathbf{x})$ is the strain field, was applied to generate the von Mises stress fields using a Fourier method [6]. The non-dimensionalized fourth-order elasticity tensor considered represents copper with cubic symmetry (Table 2).

The size in pixels of the von Mises stress field is $128 \times 128 \times 1$ (Figure 3a). For the reasons explained in Section 2.1, a moving average filter of size 4×4 with stride 4 and no padding was operated on the von Mises stress field to obtain a stress field of size $32 \times 32 \times 1$ (Figure 3b). This piecewise averaged stress field is the ground truth for the CED model. Models that reconstructed the ground truth of size $64 \times 64 \times 1$ and $128 \times 128 \times 1$ had a much higher reconstruction error as compared with ground-truth of size $32 \times 32 \times 1$.

3.1.3. *Preprocessing input and ground truth.* Before using the microstructures as the input and the low-resolution von Mises stress fields as the ground truth for training and prediction, respectively, their values are scaled to lie between $[0, 1]$. The scaling is performed for each image and each channel separately. As the model is nonlinear, the effect of the linear scaling of the input and ground truth during the gradient update of the parameters will be nonlinear. Hence, the scaled input and ground truth ensure that the pixels with very large values do not completely dominate during the computation of weights. For example, in the unscaled low-resolution stress fields, the pixel values range $[1 \times 10^{-4}, 13.68]$. Hence, the pixels with values close to 13.68 will dominate the training compared to the pixel with values 1×10^{-4} . However, when scaled between $[0, 1]$, the pixels with high values will not dominate as much as they would have during the unscaled data Figure 7.

3.2. CED architecture


Various architectures of the CED model were designed based on the governing criteria discussed in Section 2.2. The architectures were trained based on criteria discussed in Section 2.2.3 and validation error was evaluated to

Table 3. Hyperparameters for optimization of MSE. One epoch means that each sample in the training set has been used once to update the model parameters.

Batch size	100	Base learning rate	0.0001
Number of epochs	4000	Maximum learning rate	0.1

Table 4. The architecture of the CED model. The input to all layers (including the model input) was not padded with any values as it reduced the model accuracy due to misleading information. The number of filters was tuned by varying it as a multiple of r . However, for evaluation and visualization purposes, $r = 1$ is considered.

Components	Operations in layers	Kernel size	Stride	Number of filters for $r = 1$
Encoder	Convolution + ReLU	$8 \times 8 \times 4$	8	64
	Maxpool	$2 \times 2 \times 1$	2	64
	Convolution + ReLU	$3 \times 3 \times 64$	1	16
	Maxpool2D	$2 \times 2 \times 1$	1	16
Decoder	Transpose convolution + ReLU	$2 \times 2 \times 16$	2	64
	Transpose convolution + ReLU	$3 \times 3 \times 64$	1	8
	Transpose convolution + tanh	$10 \times 10 \times 8$	2	1

**Figure 8.** Validation loss for different values of the multiplier r . The number of filters is varied for each layer by multiplying r with the number of filters in Table 4. The validation loss is for $r = 3$ is 0.0178987 and $r = 9$ is 0.017834. This suggests that the validation loss remains almost constant for $r > 3$, hence $r = 3$ is used as the optimal value.

tune the model hyperparameters. The hyperparameters used for training the architectures are listed in Table 3. Parameters such as the number of layers and kernel size of the filters were tuned based on the criteria discussed in Section 2.2.4. The input of all layers was not padded as it implies fictitious information from outside the domain, causing the CED model to be less accurate. Without maxpool layers, the convergence of training loss was unstable; hence, it was necessary to include it in the CED architecture. Once the number of layers, the function of each layer, and the kernel size of the filters were finalized, the number of filters for each layer, except for the output layer, was tuned by varying them as a multiple of variable r , see Table 4. The validation error against the variable r is shown in Figure 8. With increase in r , the validation loss decreases until $r = 3$ and then remains almost constant, i.e., around 0.0178; hence, $r = 3$ is chosen as the optimal value. As there is no significant difference in the performance between the model architecture for $r = 1$ and $r = 3$, the architecture with $r = 1$ was chosen as the final CED architecture as it had a smaller number of filters in the first layer, making the analysis of the output of the first layer easier, see Section 5. PyTorch [25] was used for designing and training the model.




Figure 9. Histogram of MSE for 2000 microstructure test samples. The horizontal axis represents the MSE of each bin, and the vertical axis gives the percentage of samples for corresponding bins. The parameters μ and σ represent the mean and the standard deviation of the distribution, respectively. Around 97.1% of test samples have MSE below 0.0337. Hence, the average per-pixel error is around 3.29×10^{-05} . The grain size inside the microstructures plays a major role in the accuracy of the model prediction.

4. Evaluation

The trained CED model with architecture as described in Table 4 was used to analyze 2000 separate test samples previously unseen during training. The performance of the peak-stress cluster prediction approach was analyzed using the MSE, cosine similarity metric, and by comparing clusters characteristics. First, we determined the microstructures from the test samples for which the model prediction is highly inaccurate. To do this, we characterized the microstructures in terms of their average grain diameter, average axis ratio, and average Euler angles. We then compared them against the MSE. Next, we excluded the microstructures with very high MSE while analyzing the accuracy inside the predicted stress fields. The cluster detection in such microstructures was not useful due to high inaccuracy in stress fields prediction. The accuracy of the clusters in stress fields was analyzed in two ways: first, by evaluating cosine similarity between rescaled predicted and ground-truth stress fields; and, second, by comparing the characteristics of the peak-stress clusters detected in both prediction and ground truth.

4.1. Error analysis of the predicted stress fields using MSE

The histogram of MSE for different microstructures is shown in Figure 9. The microstructures with average, largest, and smallest MSE and their corresponding model stress field prediction and ground truth are shown in Figure 10. The MSE was studied against the microstructure's average grain diameter and average grain aspect ratio, as shown in Figure 11. For Figure 11a, first the histogram of the average grain diameter is estimated. The bins of the histogram are represented as a set $S = \{\text{Bin}_1, \text{Bin}_2, \text{Bin}_3, \dots, \text{Bin}_n\}$. Each bin in the histogram represents microstructures with similar average grain diameters. The average of MSE for all the microstructures that corresponds to Bin_i , for all $i \in [1, n]$, is computed and represented as MSE_i , for all $i \in [1, n]$. These averaged MSE for each bin are plotted against the corresponding bin average grain diameter as shown in Figure 11a. The main reason to study the plots for each bin was to understand the effect of grains size on the model performance. If the grain size significantly affects the model performance, it should be visible in the plots. Figure 11a shows that the MSE increases with an increase in average grain diameter of the microstructure. It suggests the prediction is better for microstructure with smaller grains. Similar behavior was observed when the MSE in each individual grain of the microstructures was analyzed against the corresponding grain diameter. Figure 12a shows that the error in the prediction of the stress distribution in larger grains was higher than in smaller grains. However, if the grain size is too small, then the error increases, as shown in Figure 12b.

The model performs better around medium-sized grains (size ranging from [0.2, 0.4]) as grain junctions and grain boundaries provide sufficient feature information for prediction. In areas around larger grains, due to large homogeneous regions, the number of features available locally to the model is not enough to make predictions; hence, the model predicts stress with higher error. In the case of smaller grains, there are multiple sharp jumps




Figure 10. The (a) best (with MSE of 0.0066), (b) average (with MSE of 0.0204), (c) and worst (with MSE of 0.1811) cases of low-resolution stress field prediction for the corresponding microstructure by the CED model. The prediction suffers from high inaccuracy for microstructures with large grains (relative to the computational domain).

in stress values around grains due to grain boundaries, which the model cannot capture effectively, leading to a higher error compared with medium-sized grains. This complete analysis suggests that if a microstructure consists mainly of medium-sized grains, then the model prediction error will be small. This behavior can also be observed by visual inspection, as shown in Figure 10.

Similarly, the effect of the Euler angles and the aspect ratio for each microstructure was compared against the MSE. The effect of the aspect ratio on MSE is shown in Figure 11b. The prediction is better for microstructures with grains that are closer to circular rather than elongated. Figure 13 shows that the variation of the MSE remains almost the same for Euler angles in all three directions except for small outliers. This suggests that Euler angles do not affect the model performance.

4.2. Error analysis of the predicted stress fields using cosine similarity

As the MSE can range from $[0, \infty)$, it is difficult to use this to interpret the accuracy of the predicted stress fields when compared with the ground truth. Hence, the cosine similarity (Equation (16)) is also used to analyze the




Figure 11. Plots of (a) average grain size and (b) average aspect ratio against MSE. The histogram of average grain size of the microstructures are computed. For each bin of the histogram, the MSE and the average grain size of the microstructures belonging to the bin is collected and averaged. The horizontal and vertical axes in (a) represents the MSE and the average grain size (averaged over the histogram bins), respectively. A similar calculation is performed to obtain (b). The grain diameter is normalized by the microstructure size, i.e., by 128. The MSE increases with mean grain size and mean grain aspect ratio, implying that the prediction was less accurate for microstructure with larger grains or elongated grains.




Figure 12. Plots of grain diameter against MSE. The plot is obtained in the same way as explained in Figure 11. In (a), we consider all grains in the test set with diameter in the range $[0, 1.12]$, and (b) considers only the subset of grains with diameter in the range $[0, 0.4]$. Plot (b) provides an understanding of the variation of MSE for smaller grains. The MSE is high for very small grains and very large grains.




Figure 13. Plots of average Euler angle against average MSE. Each point represents a bin of Euler angle histogram. The average of Euler angle and MSE is calculated over a bin. Except for a few outliers, the variation of MSE with respect to Euler angle in all three directions is essentially constant. The plot shows that model performance is independent of the Euler angle.




Figure 14. Histogram of cosine similarity measures for test images. The horizontal and vertical axes represent the cosine similarity and percentage of samples, respectively. The parameters μ and σ represent the mean and the standard deviation, respectively, of the corresponding distributions. The cosine similarity ranges between $[-0.0908, 0.9217]$ indicating that for some microstructures, high-stress regions in predicted stress fields are overlapping with low-stress regions in the ground truth or vice versa.

model accuracy. If \mathbf{A} and \mathbf{B} represent the predicted and the ground-truth stress field for a given microstructure, the cosine similarity is given by

$$\text{Cosine similarity}(\mathbf{A}, \mathbf{B}) = \frac{\mathbf{A} \cdot \mathbf{B}}{\|\mathbf{A}\| \|\mathbf{B}\|}. \quad (16)$$

Before computing the cosine similarity, we scale the predicted and ground-truth stress fields between $[-1, 1]$. The cosine similarity will range between $[-1, 1]$, with the value of 1 representing perfectly similar images. If the cosine similarity is computed when the pixel values of both the predicted stress field and the ground truth range between $[0, 1]$, as explained in the Section 3.1.3, it would not provide sufficient information about the accuracy. If the low-stress regions in predicted stress fields coincide with peak-stress regions in the ground truth, the cosine similarity will be low but is still positive. By scaling between $[-1, 1]$, we were able to distinguish the high-stress regions from the low-stress regions by making values of the low-stress regions negative. If the high-stress regions in the ground truth overlap with the low-stress regions in the predicted stress fields (or vice versa), the product will be negative, acting as a penalty to the cosine similarity. Further, the cosine similarity metric includes the norm of images in its denominator, which normalizes the stress values in each image. This helps in analyzing how accurately the hotspots are overlapping, irrespective of the hotspot magnitudes. In Figure 14, it can be seen that some cosine similarity values are negative, which suggests that the predictions in such images are completely wrong. One such case is shown in Figure 10c. On further investigation, it was observed that the images with the negative cosine similarity values also had very high MSE. Hence, this metric is a good indicator to compare the distribution of the stress fields irrespective of their magnitude. The predicted stress fields may not be perfectly similar; however, high positive values indicate that peak-stress clusters are approximately overlapping, and for almost all images, peak-stress clusters do not overlap with the low-stress regions.

4.3. Error analysis of clusters' geometric characteristics inside predicted stress fields

The accuracy of peak-stress cluster detection in the predicted stress fields is also studied by comparing their characteristics such as location, size, shape, and orientation with clusters in the ground truth. For the current work, clusters around the three largest peaks in the microstructure are compared. As per our definition of clusters, the characteristics will change with the threshold, as shown in Figure 5. Threshold values of $\{0.5, 0.7, 0.8, 0.9\}$ were considered for our analysis. This analysis helps us understand what kind of clusters, depending on peak and threshold, can be detected inside the stress fields predicted by our CED model more accurately.

4.3.1. Comparing location of peak-stress clusters. The peak-stress cluster location is defined as the location of the peak-stress surrounded by the cluster, see Section 2.4; therefore, the cluster location is not affected by the threshold. The cluster location accuracy is analyzed using the Euclidean distance between stress peaks in the




Figure 15. Each box plot represents the distribution of the distance (error in location) between the largest, second-largest, and third-largest peaks in the ground-truth and predicted stress fields. The model predicts the location of higher-intensity peaks more accurately.

predicted stress fields and stress peaks of the same rank in the ground-truth. Figure 15 shows the distribution of distances in the test samples. It shows that the location of clusters around the largest peaks in predicted stress fields is more accurate. As the intensity of stress values in these clusters is large, the squared error for these peaks will be larger if the model predicts values inaccurately. Hence, the model tries to learn by mainly focusing on the values at these peaks. Therefore, the prediction of the location of the highest peaks is much better.

4.3.2. Comparing size of peak-stress clusters. The size of clusters in the predicted stress fields was compared with clusters in the ground truth using the absolute relative difference in the area, δ_a , as defined in the following:

$$\delta_a = \left| \frac{\text{Area}_{\text{groundtruth}} - \text{Area}_{\text{predicted}}}{\text{Area}_{\text{groundtruth}}} \right|. \quad (17)$$

Here $\text{Area}_{\text{groundtruth}}$ refers to area of a peak-stress cluster in a ground-truth stress field. As shown in Figure 16, with an increase in the threshold, the relative difference increases. In addition, as the threshold increases, the clusters become smaller in size. As δ_a is inversely proportional to the size of the cluster in ground truth, δ_a can be high for small clusters even with a small error in the detection of the number of pixels in the clusters in the predicted stress field as compared with the ground truth. As the higher threshold provides details about more concentrated clusters around the peak, the comparison shows that the model does not capture concentrated details around the peaks; instead, it captures the overall spread of the peak-stress clusters.

On comparing δ_a for clusters around the different peaks and same threshold, the model performs well where there are stress clusters around smaller peaks in the ground truth. In the stress fields of most microstructures, these clusters are of moderate intensity but more spread out. Hence, the details of such clusters are predicted more accurately by the model. For similar reasons, the δ_a is lowest for the clusters around the third largest peaks and for all threshold values.

The model's ability to predict the size of larger clusters more accurately can be attributed to the objective function MSE. While training, the model emphasizes predicting a large number of pixels of stress clusters with average magnitude more accurately instead of a much smaller number of pixels but with relatively higher magnitude.

4.3.3. Comparing the orientation of peak-stress clusters. The accuracy of cluster orientation, θ , was compared using the absolute difference, $\delta_\theta := |\theta_{\text{groundtruth}} - \theta_{\text{predicted}}|$, between the orientation of cluster in predicted stress fields and the ground truth.

Among the clusters defined using threshold values of {0.5, 0.7, 0.8}, the average δ_θ is relatively small for the largest peak, as shown in Figure 17a. This suggests that the model tries to capture the direction of the spread of clusters around higher peaks more accurately and gives lower priority to around lower peaks. This behavior can be attributed to the objective function MSE, where the model prioritizes learning high-intensity stress values. For a threshold value of {0.9}, the behavior is the opposite; as the clusters are small, adding or removing even a single pixel in the clusters can change the orientation of the equivalent ellipse by a significant amount causing the high error.




Figure 16. The mean absolute relative error of the areas of the detected clusters in the predicted stress fields for different thresholds. The model predicts the size of larger clusters more accurately compared with smaller clusters. The errors are lower for clusters with smaller thresholds and around smaller peaks since their area is relatively larger in ground-truth stress fields.

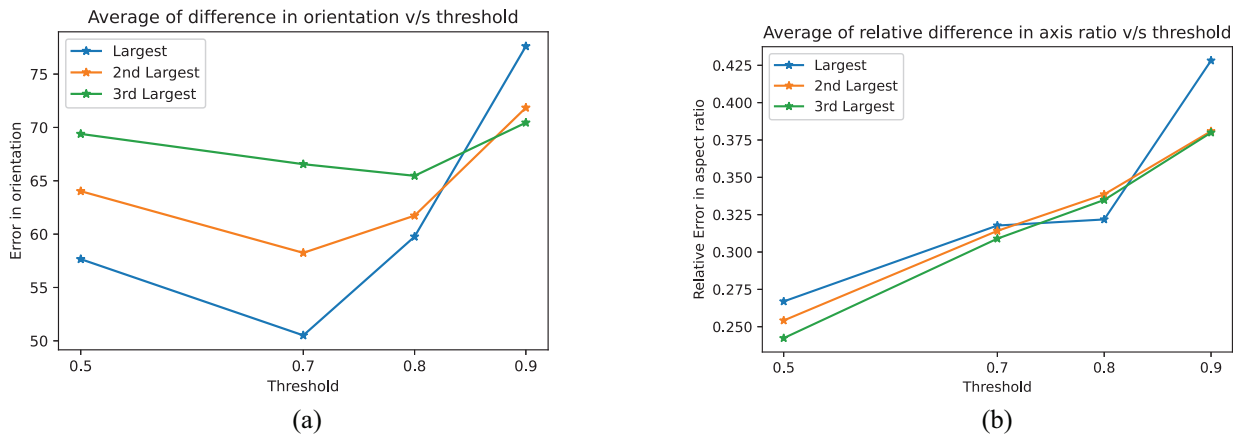


Figure 17. Mean absolute difference of (a) the orientation and (b) the mean absolute relative difference of the axis ratio of stress clusters in predicted and ground-truth stress fields for different thresholds. The orientation and shape of clusters defined using the threshold of 0.7 or less are fairly accurate, and clusters details around the largest peak are fairly accurate.

Although this definition is very sensitive, this metric is still useful because it tells us that the orientation of clusters around larger peaks is more accurate and that the model is prioritizing clusters around the large peaks. Due to the sensitivity of this metric, an unacceptable range would be 90° or more. Thus, for example, if a peak-stress cluster predicted by the model is oriented horizontally, whereas, in the ground truth, it is oriented vertically, this prediction would be completely unacceptable. Other, less-sensitive metrics could provide a better picture of the orientation accuracy.

This analysis suggests that the orientation of clusters defined using a threshold of 0.7 or less is fairly accurately predicted using our approach.

4.3.4. Comparing the shapes of peak-stress clusters. The shape was compared using the relative absolute difference in axis ratio, $\delta_{ar} := \left| \frac{AR_{\text{groundtruth}} - AR_{\text{predicted}}}{AR_{\text{groundtruth}}} \right|$, where $AR_{\text{groundtruth}}$ is the axis ratio of a peak-stress cluster in a ground-truth stress field.

As shown in Figure 17b, δ_{ar} is almost the same for all three clusters irrespective of the threshold used to define them. However, the error increases with an increase in the threshold. As the clusters with lower thresholds are larger in size, for the same reason as in Section 4.3.2, the axis ratio is more accurate for these clusters. With the increase in threshold, the size decreases, and even a slight shift of a pixel in clusters can change the axis ratio by a considerable value, hence the high error. This analysis suggests that the axis ratio can be more accurately estimated for all three clusters defined using the threshold of 0.8 or less using our approach.

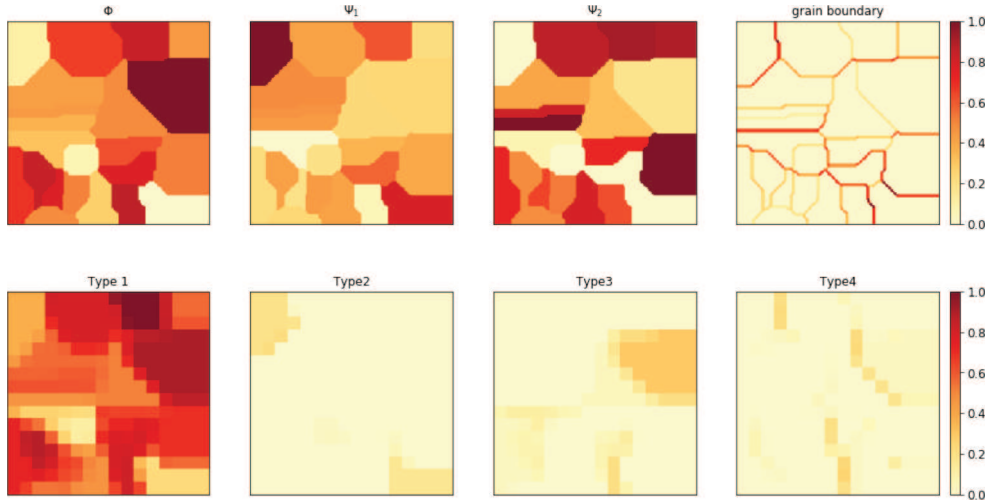


Figure 18. The first row shows the scaled input to the CED model. The output of filters in the first layer of the model after activation was analyzed and grouped into five types based on their behavior. The second row shows the output of selected filters belonging to the first four types. The output from the filters belonging to the fifth type remains zero, and hence is not shown here. Type 1 filters remain active in all regions. Type 2 and Type 3 remain active only in regions with relatively high and low Euler angles, respectively. Type 4 remains active only in regions where grain boundaries in specific orientation are present irrespective of Euler angle values.

The complete analysis for cluster characteristics provides the understanding that the peak-stress clusters around the largest peak defined using a 0.7 or less threshold can accurately be estimated using our approach. The stress regions around lower peaks and with a higher threshold may not be that accurate.

4.4. Computational cost of the CED model

The computational cost of the CED model was analyzed by estimating the time taken to train the model on the training set and the time taken by the trained model to predict the low-resolution stress fields for the complete microstructure dataset. A CPU machine with the configuration Intel i7-3770 CPU @ 3.40 GHz was used to analyze the computational cost.

The machine took 8 hours, 32 minutes, and 14 seconds to train the CED model using 8000 training samples and training parameters as listed in Table 3. For predicting the low-resolution stress fields for 10,000 microstructures using the trained CED model, the machine took 1 minute and 25 seconds. Using the EVPFFT package [6], the machine took 10 hours, 34 minutes, and 36 seconds to simulate stress for the same 10,000 microstructures. The package was parallelized, and two processors were used to perform the simulations.

The CED model is around 448 times more computationally efficient as compared with the EVPFFT package. However, these comparisons are limited only to CPU machines, and the efficiency can be different on GPU machines. It is also important to note that training the machine-learning model requires a relatively large memory compared with the EVPFFT package. This is because the model needs to load the complete dataset into RAM before training. A data-efficient training scheme can help reduce the requirement for large RAM.

5. Feature visualization

The output of the filters from the first layer of the CED model was analyzed to investigate the microstructure features extracted by the model while predicting the stress fields. In this work, the analysis was limited to just the first convolutional layer, as the filter behavior in the successive layers does not directly depend on the microstructure but is influenced by information extracted by the predecessor layers. As listed in Table 4, there are 64 filters in first layers. Each filter outputs a feature map for a given microstructure.

Figure 18 shows the features extracted by selected filters of the model from a given microstructure and its grain boundary information. As the microstructure had many complexities in such as grain boundaries, Euler angles, and grain junctions, we also study the behavior of these filters on simple manually designed microstructures. In the first case, we considered a single crystal microstructure with the fixed Euler angle values.




Figure 19. Manually designed bicrystal microstructures with the same values (normalized) of (Φ, Ψ_1, Ψ_2) but different grain boundary orientation. The output of the filters in the first layer of the model was studied to analyze the role of grain boundary orientation.




Figure 20. Activation plots with varying Euler angle values. Filter 16 represents Type 2, Filter 22 represents Type 3, and Filter 32 represents Type 1. Type 1 filters capture all the grains irrespective of their Euler angles, Type 2 filters detect grains with relatively larger Euler angles, and Type 3 filters detect grains with relatively smaller Euler angles.

The values were varied from 0 to 1 since the scaled microstructure is fed into the model. In the second case, we considered a bicrystal with Euler angle values (normalized) but rotated the grain boundaries, as shown in Figure 19. The response of the filters was then studied. We classified filters based on their behavior as follows.

- **Type 1:** These filters are active at every point in the microstructure irrespective of relative Euler angle values and grain boundary orientation. The output after the ReLU activation from these filters will always be non-zero, as shown in Figure 20. It means that these filters could extract information from all combinations of Euler angles and grain boundaries. However, the magnitude of feature extraction varies with Euler angles and grain boundary orientation.
- **Type 2:** These filters are active only in regions with large Euler angles, see Figure 20. Their response decreases with a decrease in Euler angle values. The response also varies with grain boundary orientation around grains with larger Euler angles. They are unable to detect grain boundaries around grain with smaller Euler angles.
- **Type 3:** These filters are active only in the presence of small Euler angles, see Figure 20. Their response will decrease and then becomes 0, with an increase in Euler angle values. Their response varied with the orientation of grain boundaries but cannot detect grain boundaries around grains with large Euler angles. It shows that the filters were able to extract features only with relatively small Euler angle values.
- **Type 4:** These filters remain inactive for most cases but are active only in the presence of grain boundary in a specific orientation or due to significant relative difference between Φ , Ψ_1 , and Ψ_2 , as shown in Figure 21. It shows that the filters mainly focus on grain boundary detections, as shown in Figure 18.
- **Type 5:** These filters remain inactive in every case. The parameters of these filters are negative; hence, the output after the ReLU activation from these filters will always be zero.

This observation is consistent with the observations in Figure 18. As the microstructure values before being fed to the model are scaled from 0 to 1, the behavior of filters explains that the model is concerned with only relative variations in microstructure features, and it does not care about the values. It is important to note that the size of the single-crystal and bicrystal microstructures used to analyze the filter behavior is 8×8 . It represents a small section of a large microstructure. As discussed previously, the model performance on single-crystals and bicrystals of size 128×128 are highly inaccurate; hence visualization on such microstructures can lead to wrong interpretations. Moreover, the output of the filters in the first layer on the microstructure of the size 8×8 was a scalar value, which was easy to analyze.




Figure 21. Activation plot of Filter 52 representing Type 4. The left figure shows activation of the filter when bicrystal microstructures with the same Euler angles but different orientations of grain boundary, as shown in Figure 19 were fed into the model. On the right, the filter activation when monocrystalline microstructures with varying Euler angle values were provided to the model. These two figures show that the filter mainly focuses on detecting the grain boundaries and remain inactive in the regions with no grain boundary.




Figure 22. Percentage change in test error when a filter with a given index (plotted on the horizontal axis) was removed from the model. The plot shows that filter 0 and 32 play the most crucial role in stress field prediction as the increase in test error after removing these filters is higher than the other filters.

Finally, we study the importance of each filter in predicting the stress distribution. For this, we select a filter, set its parameters to 0, and predict the stress distribution on the test set. Setting the parameter to 0 is equivalent to removing those filters, as the output from these filters will not influence the output of successive layers. The test loss obtained after removing a particular filter from the model is compared against the original test loss. Figure 22 shows the percentage change in loss. The error in the model prediction on test samples increased by almost 80% on removing filter 32. This filter also remains active at every point in the microstructure, detecting every feature. Hence, this filter is the most important. Similarly, filter 0 is the next most important filter.

6. Conclusion

In this work, we have presented a method to predict peak-stress clusters in polycrystalline microstructures under a macroscopic strain. First, stress fields were predicted using a CED model for given microstructures. A combination of a peak detection algorithm and a connected component labeling algorithm was used to detect peak-stress clusters in the predicted stress fields. The method was demonstrated for microstructures of pixel size 128×128 , and peak-stress clusters were detected in the stress field of resolution 32×32 predicted by the CED model. The model was trained using 7000 microstructure samples. The performance of the model was analyzed for 2000 test samples by comparing predicted fields with the ground-truth using mean squared metrics for the overall stress field. The accuracy of detected peak-stress clusters was analyzed using the cosine similarity metric and comparing their geometric characteristics.

On investigating the MSE against the average grain sizes, average grain shapes, and average Euler angles of microstructures, it was observed that the CED model performed well for samples with average (normalized) grain size in the range of [0.2, 0.4]. The MSE for the microstructures with the average grains of size outside this range was high. The cosine similarity metric was used to analyze the overall accuracy of peak-stress clusters inside the predicted stress field. Further, the location, size, shape, and orientation of peak stress clusters of threshold values of {0.5, 0.6, 0.7, 0.9} was individually compared. This analysis showed that the geometric characteristics such as size and axis ratio of the peak-stress clusters around the highest peaks are more accurate. The orientation of lower threshold clusters around larger peaks is more accurate, and the location of the clusters around the largest peak is most accurate.

The reason for this good performance is because the model prioritized learning the characteristics of very high-stress values. Further, peak-stress clusters defined using a threshold value of at most 0.7 are more accurate than stress clusters with a threshold value of more than 0.7. The reason for this is that the model was able to better capture the large-scale structure of peak stress clusters and missed the finer-scale details. The overall analysis provides insight into the microstructure features that are most important in determining the peak stresses.

The proposed CED model was trained and further evaluated to predict low-resolution von Mises stress fields. With a slight change in hyperparameters, the model can also be trained to predict other quantities such as the principal stresses. Although our method shows promising results, a CED model with a different architecture can potentially provide more accurate predictions. For instance, a deep-learning architecture that can predict stress fields with the same resolution as the microstructure, instead of a coarsened stress field, will likely be more effective in providing accurate predictions. However, the dimensionality will be significantly larger, and hence training and prediction will be vastly more expensive. Alternatively, a CED model with a fully connected layer can be used to analyze the compressed information in the lower dimension [26], to go beyond the current work wherein we did not use fully connected layers as it suffered from early overfitting.

An important issue that we did not address is the evolution of peak stresses under evolving load. A CED model combined with LSTM architecture can potentially be designed to predict peak-stress time evolution in large systems without significant recomputation at each time.

Finally, a very important challenge in machine-learning-based models is not only to make predictions but also to understand the physics underpinning the predictions. Various visualization algorithms such as smoothgrad [27] and guided back-propagation [28] can be used to investigate the specific features of the microstructure that play a significant role in the formation of peak-stress clusters.

Research data

The DREAM3D software to generate microstructure is available publicly at dream3d.bluequartz.net/. A version of the code developed for this work is available at github.com/ankitvaibhava/Peak_Stress_in_Microstructures.

Acknowledgements

We thank Pittsburgh Supercomputing Center for XSEDE computing resources, Ankita Mangal and Vahid Tari for help with learning to use Dream3D and EVPFFT, and Professor Pei Zhang, Professor Susu Xu, Dr Mostafa Mirshekari, Dr Jonathon Fagert, and Dr Rajat Arora for useful discussions.


Funding

The authors disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported by the National Science Foundation [grant numbers DMREF 1921857 and CMMI MOMS 1635407], Army Research Office [grant number MURI W911NF-19-1-0245], and Office of Naval Research [grant number N00014-18-1-2528].

Notes

1. In machine learning, the statistical expression that expresses a relationship between input variable and response variable is referred to as a “model.” It is often distinguished from the word “method,” which describes the entire process of preprocessing the input, making a prediction using the model given the input, and post-processing the model prediction for further application.
2. In this work, we refer to a 3D array of size $H \times W \times C$ as a multichannel image, where C is the number of channels. The input to CED model is a four channel image with size $H = 128$, $W = 128$, and $C = 4$. Similarly ground truth for the model is a single-channel image with size $H = 32$, $W = 32$, and $C = 1$. While discussing dimensions $H \times W$, we refer to them as image resolution.

ORCID iD

Kaushik Dayal  <https://orcid.org/0000-0002-0516-3066>

References

- [1] Liu, Y and Castañeda, PP. Homogenization estimates for the average behavior and field fluctuations in cubic and hexagonal viscoplastic polycrystals. *J Mech Phys Solids* 2004; 52(5): 1175–1211.
- [2] Lipton, R, Sinz, P, and Stuebner, M. Uncertain loading and quantifying maximum energy concentration within composite structures. *J Computat Phys* 2016; 325: 38–52.
- [3] Bathe, KJ. *Finite Element Procedures*. Klaus-Jurgen Bathe, 2006.
- [4] De Geus, T, Vondřejc, J, Zeman, J, et al. Finite strain fft-based non-linear solvers made simple. *Comput Meth Appl Mech Eng* 2017; 318: 412–430.
- [5] Zeman, J, de Geus, TW, Vondřejc, J, et al. A finite element perspective on nonlinear FFT-based micromechanical simulations. *Int J Numer Meth Eng* 2017; 111(10): 903–926.
- [6] Lebensohn, RA, Kanjarla, AK, and Eisenlohr, P. An elasto-viscoplastic formulation based on fast Fourier transforms for the prediction of micromechanical fields in polycrystalline materials. *Int J Plasticity* 2012; 32-33: 59–69.
- [7] Peng, X, Nepal, D, and Dayal, K. Effective response of heterogeneous materials using the recursive projection method. *Comput Meth Appl Mech Eng* 2020; 364: 112946.
- [8] Mianroodi, JR, Siboni, NH, and Raabe, D. Teaching solid mechanics to artificial intelligence - a fast solver for heterogeneous materials. *NPJ Computat Mater* 2021; 7(1): 1–10.
- [9] Mao, Y, He, Q, and Zhao, X. Designing complex architected materials with generative adversarial networks. *Sci Adv* 2020; 6(17): eaaz4169.
- [10] Yang, Z, Li, X, Catherine Brinson, L, et al. Microstructural materials design via deep adversarial learning methodology. *J Mech Des* 2018; 140: 11.
- [11] Lubbers, N, Lookman, T, and Barros, K. Inferring low-dimensional microstructure representations using convolutional neural networks. *Phys Rev E* 2017; 96(5): 052111.
- [12] Li, X, Zhang, Y, Zhao, H, et al. A transfer learning approach for microstructure reconstruction and structure-property predictions. *Sci Rep* 2018; 8(1): 1–13.
- [13] Mozaffar, M, Bostanabad, R, Chen, W, et al. Deep learning predicts path-dependent plasticity. *Proc Natl Acad Sci USA* 2019; 116(52): 26414–26420.
- [14] Frankel, A, Tachida, K, and Jones, R. Prediction of the evolution of the stress field of polycrystals undergoing elastic-plastic deformation with a hybrid neural network model. *Machine Learning: Science and Technology* 2020; in press.
- [15] Mangal, A, and Holm, EA. Applied machine learning to predict stress hotspots I: Face centered cubic materials. *Int J Plasticity* 2018; 111: 122–134.
- [16] Ronneberger, O, Fischer, P, and Brox, T. U-NET: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Berlin: Springer, pp. 234–241.
- [17] Badrinarayanan, V, Kendall, A, and Cipolla, R. SegNET: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 2017; 39(12): 2481–2495.
- [18] Rollett, AD, Lebensohn, RA, Groeber, M, et al. Stress hot spots in viscoplastic deformation of polycrystals. *Modell Sim Mater Sci Eng* 2010; 18(7): 074005.
- [19] Groeber, MA, and Jackson, MA. DREAM.3D: a digital representation environment for the analysis of microstructure in 3D. *Integrat Mater Manuf Innovation* 2014; 3(1): 56–72.
- [20] Nair, V, and Hinton, GE. Rectified linear units improve restricted boltzmann machines. In: *Proceedings of the 27th International Conference on International Conference on Machine Learning (ICML'10)*. Madison, WI: Omnipress, pp. 807–814.
- [21] Goodfellow, I, Bengio, Y, and Courville, A. *Deep Learning*. Cambridge, MA: MIT Press, 2016.
- [22] Qiu, S, Xu, X, and Cai, B. FreLu: Flexible rectified linear units for improving convolutional neural networks. In: *2018 24th International Conference on Pattern Recognition (ICPR)*. IEEE, pp. 1223–1228.
- [23] Kingma, DP, and Ba, J. Adam: A method for stochastic optimization. In: *3rd International Conference on Learning Representations (ICLR 2015)*, San Diego, CA, May 7–9, 2015.
- [24] Rosin, PL. A note on the least squares fitting of ellipses. *Pattern Recognition Letters* 1993; 14(10): 799–808.
- [25] Paszke, A, Gross, S, Massa, F, et al. PyTorch: An imperative style, high-performance deep learning library. In: *Advances in Neural Information Processing Systems*, pp. 8026–8037.
- [26] Guo, X, Liu, X, Zhu, E, et al. Deep clustering with convolutional autoencoders. In: *International Conference on Neural Information Processing*. Berlin: Springer, pp. 373–382.
- [27] Smilkov, D, Thorat, N, Kim, B, et al. Smoothgrad: removing noise by adding noise. *arXiv e-prints* 2017; 1706.
- [28] Springenberg, JT, Dosovitskiy, A, Brox, T, et al. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:14126806*, 2014.