

Adapting U-Net for linear elastic stress estimation in polycrystal Zr microstructures

J. D. Langcaster^a, D. S. Balint^b, M. R. Wenman^a

^a*Department of Materials, Imperial College London, London, UK*

^b*Department of Mechanical Engineering, Imperial College London, London, UK*

1. Abstract

A variant of the U-Net convolutional neural network architecture is proposed to estimate linear elastic compatibility stresses in α -Zr (hcp) polycrystalline grain structures. Training data was generated using VGrain software with a regularity α of 0.73 and uniform random orientation for the grain structures and ABAQUS to evaluate the stress fields using the finite element method. The initial dataset contains 200 samples with 20 held from training for validation. The network gives speedups of around 200x to 6000x using a CPU or GPU, with significant memory savings, compared to finite element analysis with a modest reduction in accuracy of up to 10%. Network performance is not correlated with grain structure regularity or texture, showing generalisation of the network beyond the training set to arbitrary Zr crystal structures. Performance when trained with 200 and 400 samples was measured, finding an improvement in accuracy of approximately 10% when the size of the dataset was doubled.

2. Keywords

Microstructure, Machine Learning, Finite Element Analysis, Zirconium, Nuclear Materials

3. Introduction

Internal microstructural type II stresses [24] in Zr-based nuclear fuel cladding in water-cooled reactors has a significant impact on the onset and progress of hydrogen diffusion and hydride precipitation that may lead to cladding failure. In particular, grain compatibility stresses in Zr-based alloys influence the diffusion of hydrogen within the matrix, leading to the concentration of hydrogen in regions of tensile hydrostatic stresses and allowing the formation of hydrides, causing either embrittlement or delayed hydride cracking (DHC) [15]. As such, there is interest in the modelling of these compatibility stresses to feed into microstructurally sensitive models of hydride precipitation. Currently compatibility stresses can be estimated using finite element (FE) analysis, though such simulations can only feasibly simulate small domains on the order of microns with a few hundred grains. This means that larger scale statistical analysis based upon macroscopic measures of microstructure are impractical even if they are desired for industrial applications. Recent headway by Patel et al. [19] has been made in the microstructural modelling of hydride precipitation in a 2D cross section of the cladding, which incorporates some microstructural features such as grains, crystallographic microhydride precipitation and notch defect stress fields, and can be executed in a few seconds. However, this model lacks awareness of grain compatibility stresses that generally require a full numerical solution, such as by FE methods, resulting in slow execution, which the model was designed to avoid. It is within these restrictions that a machine learning based surrogate can provide an alternative, which is investigated in this paper.

In the last decade, machine learning based methods have been shown to be capable of tackling a wide array of real world problems, such as image domain transfer e.g. segmentation of cell micrographs [21] or diagnosis of COVID-19 using chest x-rays [20]. Very recently, generalised network architectures have been shown to be effective at emulating physical simulations across a variety of domains with very limited data [9], displaying the effectiveness of machine learning techniques for physical problems. Once trained, these algorithms can be orders of magnitude faster than the

simulations they are based on, which can save significant CPU time if the model can generalise beyond the training data.

This paper proposes the use of a convolutional neural network based upon U-Net [21] to accelerate the prediction of the maximum principal stress within an α -Zr hcp grain structure given only a map of grain orientation. Using a machine learning algorithm can bypass the need for an expensive FE or equivalent solver beyond generating training data. Provided the trained algorithm can be generalised beyond the grain structures in the training data then significant speed-ups in calculation can be realised with a small sacrifice in accuracy. A trained network may then be incorporated into a model, such as that by Patel et al. [19] to approximate these stresses to inform the hydride precipitation model without sacrificing the execution speed of the model, although with this work the model is purely elastic and would require further development to incorporate plasticity.

U-Net [21] was originally designed for the segmentation of cancer cell micrographs as an image to image mapping solver. It utilises a convolutional encoder-decoder layout, in which all network layers are parameterised convolutions with downsampling and upsampling steps in the encoder and decoder sections respectively. In addition, the network also feeds outputs from earlier in the network into later layers, which helps the network restore smaller scale details in the output which may be missed during the upsampling step, resulting in sharper output images. U-Net has been extended to several other fields beyond image segmentation, such as solving for the temperature field given a map of sources [28]. U-Net has also been used as a mesh super-resolution method for stress fields [25], where the output of a low resolution FE simulation is refined by a U-Net to match the output of a high resolution FE simulation without the required computation time and memory requirements.

Traditional simulations rely on discretising the partial differential equations that describe the physical behaviour in question. By discretising, the PDEs are transformed into a system of finitely many algebraic equations, amenable to solution by a computer. The finite element method (FEM) is a commonly used discretisation and solution method with significant flexibility in how it divides the input domain into finite elements. Overall, the solver of a PDE system can be expressed as an operator G which maps the input, in this case a map of the orientation or more generally the elasticity tensor of the domain ϕ to the desired output, in this case the stress state σ , such that $G(\phi) = \sigma$. Numerical solvers are an explicitly programmed approximate operator based upon discretisation of the PDE system while the goal of this paper is to propose a neural network that can be trained to approximate the operator G .

Neural networks can vary in size from a handful of parameters for simple classification systems up to hundreds of billions of parameters in the most complex language models [2], but typical convolutional neural networks for parsing spatial data usually have a few million parameters. The basic unit of a neural network, the eponymous neuron, is a simple linear combination of its input vector, \vec{x}^{n-1} , multiplied by a weight matrix, \mathbf{W}^n and added to a bias vector, \vec{b}^n . This linear combination is then passed through a non-linear activation function, σ^n , to produce the output vector, \vec{x}^n , where n is the layer number.

$$\vec{x}^n = \sigma^n \left(\mathbf{W}^n \cdot \vec{x}^{n-1} + \vec{b}^n \right) \quad (1)$$

The activation function must be non-linear in order for the network to learn non-linear relationships in the data. However the degree of non-linearity depends ultimately on the specific data required for the model. While sigmoid functions were often used in earlier neural networks for all layers [22], recently the rectified linear unit (ReLU) [6] has been favoured due to both its performance in fitting the data [8] and computational efficiency. A neural network is then the composition of many of these layers, resulting in a highly non-linear function, which can approximate any arbitrary function, as outlined in the universal approximation theorem [8] [4]. The weight matrices and bias vectors are collectively the parameters of the network that allows the network to fit the specific problem at hand. The overarching network architecture may be encoded collectively as the network's hyperparameters, the optimisation of which is the subject of active research [9] [7], although currently most models are tuned by hand for the specific problem.

The problem of finding suitable parameters for a specific task can be formally encoded as finding the parameters, θ , of a function, f , that maps from the input set, X , to the output set, Y , by means of minimisation of a cost function, C , which is a metric between two points in Y .

$$\theta^* = \min_{\theta} C(f_{\theta}(X), Y) \quad (2)$$

The parameters are typically found using a variant of the gradient descent method via the backpropagation of gradients through to earlier parts of the network [22], taking a small step down the gradient in the parameter space.

$$\theta^{t+1} = \theta^t - \lambda \nabla_{\theta} C(f_{\theta}(X), Y) \quad (3)$$

The size of the step is determined by the learning rate, λ , which may be set at training initiation or changed during runtime to fit the problem.

The cost function has a significant effect on convergence of the network since it determines the cost landscape in which the parameters must be optimised, which can be quite particularly rough for some network architectures [12]. For regression type problems, the mean squared error (MSE) provides a stable measure of difference with the added benefit that the cost is approximately quadratic around local minima, such that the gradient naturally tends towards zero at the minima. As a result, learning rate does not need to be adjusted to allow convergence. Another commonly used cost function is the Kullback-Leibler divergence [10], which measures the relative informational entropy between 2 distributions, making it well suited to dealing with Bayesian Learning algorithms [3], which encode weights as probability distributions rather than single numbers.

As the datasets used for training are often much larger than available GPU memory (VRAM), especially when tracking parameter gradients, the dataset is often broken into smaller batches to perform each step, which introduces some noise into the gradient but allows for faster step times and lower memory requirements.

Connectivity of neurons between layers has significant influence on the characteristics of the network and which tasks the network is suited to tackling. Most networks in early research had each neuron in a layer connected to all neurons in the subsequent layer, known as fully connected (FC) layers. This comes with several disadvantages when applying the network to image or volumetric type data. First, the number of parameters in the network quickly balloons with the number of neurons in each layer, which becomes a significant problem when the input data is, for example, a 1000x1000 pixel image, requiring a million neurons in the input alone. The network can also only work on a predetermined input and output size, so either scaling algorithms must be used or the network must be retrained for each input/output size. As there are unique weights for every point in the domain, translationally invariant rules must be explicitly implemented into the training data, which requires either a very large dataset or careful data augmentation, both of which significantly increase training time.

Inspired by the biology of the human visual system, convolutional layers were introduced by LeCun et al. [11], in which the weights and connections are instead expressed as a convolution between a kernel and the previous layer.

$$\mathbf{x}^n = \sigma^n (\mathbf{W}^n \star \mathbf{x}^{n-1}) \quad (4)$$

Weights are now shared across the entire domain and each point in a layer is only connected to a small neighbourhood of the previous layer, allowing for much easier learning of local, translationally invariant rules, like the differential equations that govern elasticity. Further removing any fully connected layers allows for a fully convolutional neural network [16], which can natively be applied to any input size.

4. Related Work

While machine learning is a field which dates back to the 1950s, it wasn't until the last decade with the advent of Deep Learning and Convolutional Neural Networks that machine learning began to be applied to physical problems due to the now abundant memory and computational speed available.

DeVries et al. [5] applied a traditional fully connected network to approximate full field solutions to viscoelastic problems using only a handful of randomly distributed samples of the solution. The network takes as input the global coordinates of the point of interest and returns the displacement vector at that point, learning the function mapping between these two variables based upon a small number of samples. This allows for significant acceleration over running the full viscoelastic calculation while encoding the solution in a compact, portable form. With the weights and biases and knowledge of the network architecture, the solution can be easily reconstructed using freely available machine learning libraries rather than requiring complex simulation software.

Fully connected networks have also been used to estimate the stress on a patient's aorta given its shape [13]. The position of 5000 points on the surface of the aorta are given as input to an encoder, which generates a shape code, that is then passed through a network that converts it to a stress code that is unpacked by a decoder into the stress at each of the 5000 points on the surface. Both the encoder and decoder were trained in an unsupervised manner as part of an autoencoder, creating an efficient low dimensional representation of the shape and stress data. Not only does this significantly speed up the estimation of stress, the workflow also lends itself to further automation by using machine learning to estimate the shape of the aorta from medical imaging. This alleviates the need for manual construction of a finite element model, a time-consuming and labour intensive task.

Other approaches to stress estimation employ convolutional neural networks to exploit the locality of the governing equations and a natural encoding of geometric features. StressNet and SCSNet proposed by Nie et al. [17] use convolutional layers to estimate the stress in a given cantilever geometry, where StressNet is fully convolutional with some fully connected residual paths while SCSNet uses some fully connected layers in the centre of the network. Prediction error of these networks is around 10% and 2% for SCSNet and StressNet, respectively, while runtimes are significantly faster, taking seconds to calculate hundreds of configurations that would otherwise take hours with a traditional FE analysis approach, allowing for real-time optimisation in structural design.

Machine learning approaches have also been used to aid modelling of a variety of microstructural phenomena. A deep learning approach has been applied to estimating stress and strain fields in a two phase composite with plastic deformation by Yang et al. [26]. In this work, a generative-adversarial network (GAN) was trained to produce the von Mises stress and plastic strain in a given composite sample following a compressive loading-unloading test. A GAN, two components are required. A generator network, in this case a variant of U-Net, is tasked with producing a predicted stress/strain field, while a discriminator network, in this case PatchGAN, is trained to detect whether a given stress/strain field is real data or generated by the generator network, with both networks trained in tandem. The composite samples consist of an 8x8 image consisting of two phases, soft and hard, to produce a 256x256 map of the stress and strain in the domain. FE analysis was used as the ground truth.

Incorporating a CNN into the physics-informed framework allows for the direct transfer from the input configuration to the output fields, so can be used as a general solver for a specific PDE. Zhao et al. [28] uses a network based on U-Net [21], which takes a heat source field as input and outputs the associated temperature field, solving the heat transfer equation generally for any arbitrary source field. The hard boundary constraint method was used to include the boundary conditions as an intermediate step in the loss calculation. As the boundary conditions are not included as a network input, the network can solve arbitrary source fields but is limited to a single boundary configuration. While that limits this framework as a fully general solver, many problems such as tensile tests assume the same boundary conditions with only variation in the internal structure.

In this work a 3D grain structure, with a variety of textures, under uniaxial elastic tensile loading is used to train a U-net convolutional network. The output of the network is a microscale stress map incorporating the grain to grain compatibility stresses. The intention is to feed this map into the model of Patel et al. [19] for prediction of hydride precipitation to improve its predictive ability without comprising its great advantage of speed, once plasticity is taken into account.

5. Methodology

5.1. Dataset Generation

Grain structures were simulated using controlled Poisson Voronoi tessellation using the VGrain software package [27] [29]. A domain size of 10 μm x 10 μm x 10 μm was used with average volumetric grain size of 10 μm^3 and 20 μm^3 , corresponding to a linear grain size of 2.7 μm and 3.4 μm and resulting in 100 or 50 grains per simulation respectively. There is a tradeoff between capturing the geometric complexity which requires more grains and the runtime of the simulations, which grows with the number of grains. The grain structures were generated with regularity, α , of 0.73, where regularity is defined as in Zhang et al. [27] and parameters from Zhu et al. [29] due to the 3D structure.

$$\alpha = A \left(\frac{c}{c_m} - \frac{c_0}{c_m} \right)^{k+nc/c_m}, \quad c_0 \leq c \leq c_m \quad (5)$$

$$P_r(c) = \int_{D_L/D_{\text{mean}}}^{D_R/D_{\text{mean}}} \frac{c^c}{\Gamma(c)} x^{c-1} \exp(-cx) dx \quad (6)$$

where P_r , D_L , D_R and D_{mean} are input parameters and c is implicitly calculated from equation (6). $\Gamma(x)$ is the gamma function which extends the definition of the factorial to the real numbers. c_0 , c_m , A , n and k are empirical constants provided by Zhang et al. [27]. $\alpha = 0$ corresponds to a pure Voronoi tessellation while $\alpha = 1$ results in the most regular structure, for example a hexagonal lattice in the 2D case. While typical Zircaloy has a grain size of approximately 10 μm , in the linear elastic case the problem can be non-dimensionalised, so dependence on the length scale can be omitted from the network, although there will be some dependence on the resolution of the voxelised representation due to the U-Net architecture.

VGrain generates Python scripts that can be imported into the ABAQUS FE analysis software for simulation of the stress response of polycrystal structures. The built-in functionality for the material properties was used with

Table 1: Empirical constants for Controlled Poisson Voronoi Tessellation (CPVT).

Constant	Value
c_0	3.555
c_m	47.524
A	0.738895
k	0.323911
n	-0.414367

the fully anisotropic elasticity tensor. The Python scripts were then modified to include the elastic stiffness tensor of Zr material properties at 298 K from Tromans' review of HCP elasticity [23]. Orientation was incorporated using ABAQUS' built-in coordinate orientation functionality, with orientations generated from a uniform distribution over $SO(3)$, the 3D rotation space. This was chosen to give the neural network the widest range of potential grain misorientations for learning the stress response without bloating the training set size, allowing for a smaller training set size and faster training.

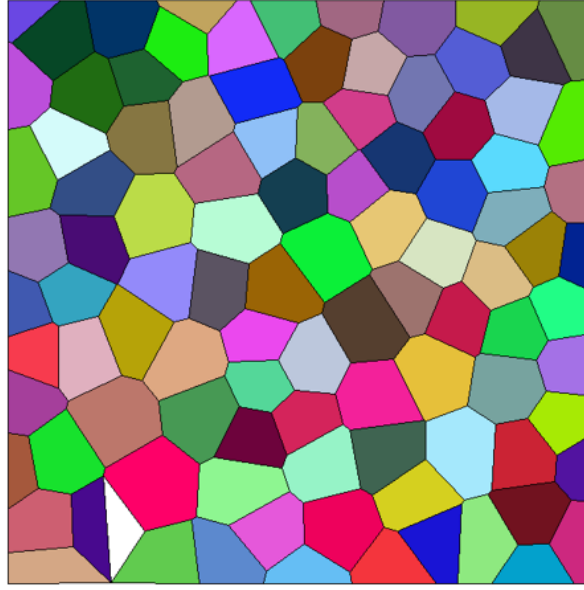


Figure 1: Example of VGRAIN structure. The colours of each grain correspond to the Euler angles in the ZXZ basis, with red, green and blue channels mapped to the respective Euler angle.

Meshes in ABAQUS were generated using unstructured quadratic tetrahedral elements (C3D10) as this allowed for automatic mesh generation for all simulations without human input, which would be unfeasible for larger dataset sizes. Elements were generated using a seed spacing of $0.3 \mu\text{m}$ along grain edges, resulting in around 2 million elements per simulation. The simulations were loaded using a uniform stress of 1 MPa tensile along the positive X/Transverse face of the domain with encastre constraints along the negative x/transverse, using the linear elastic implicit solver. The chosen load is arbitrary thanks to the linearity of the elastic stress equations, so the load and resulting stress field can be made dimensionless.

$$C = L\hat{C} \quad (7)$$

$$\sigma = C : \varepsilon = (L\hat{C}) : \varepsilon = L(\hat{C} : \varepsilon) = L\hat{\sigma} \quad (8)$$

$$\nabla \cdot \sigma = \nabla \cdot (L\hat{\sigma}) = L(\nabla \cdot \hat{\sigma}) = 0 \rightarrow \nabla \cdot \hat{\sigma} = 0 \quad (9)$$

where C is the elastic stiffness tensor, σ is the stress, L is the applied load and variables with a hat are the dimensionless counterparts.

The ABAQUS simulations were run on a workstation with a 2.3 GHz quad core Xeon and 64 GB of RAM, taking around 43 hours to complete with each sample taking around 13 minutes. GPGPU acceleration was not used as it caused erratic changes in the time taken to complete, so it was unclear if doing so would have benefited simulation times overall.

Once the stresses had been calculated in ABAQUS, the output database files were exported to the VTK file format

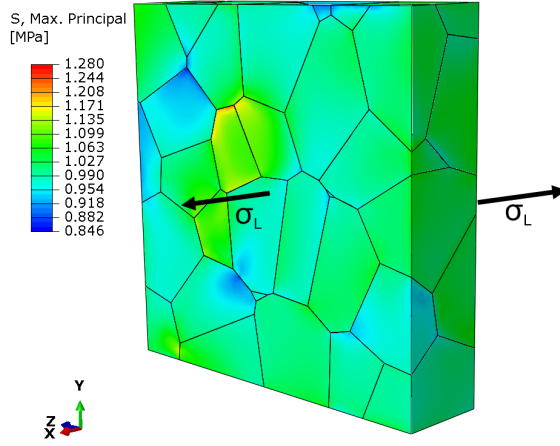


Figure 2: Example of FEA results from ABAQUS, showing the maximum principal stress of a cut through from the model. Locations of stress concentrations and depressions can be seen at the grain boundaries of up to 28% of the initial loading stress of 1 MPa.

using code based on the work of Liu et al.[14] to be voxelised in the Paraview [1] data visualisation software using the "Resample to Image" function. A resolution of 128^3 voxels was used as a compromise between resolving stress gradients and the memory footprint of the samples. The samples were then saved to the VTI format for import into Python. 200 samples were generated, with 180 used for gradient descent training and 20 used for validation during training. While this is a relatively limited dataset size for most machine learning applications, initial experimentation found much better than expected performance and therefore this size was used going forward.

An additional 165 samples were generated for validating the trained network against a variety of microstructural regularities and textures. Five distinct textures were tested, with the Kearns factors used in Table 2, with 33 samples per texture. The samples were split across 11 different regularities, $\alpha = 0.0$ to 0.9 in 0.1 increments, with an additional three samples at $\alpha = 0.73$ to match the training data. None of the samples in this validation set were used in the training of the network. Non-uniform textures were generated by rejection sampling of the crystal c -axis using a double-poled Bingham-Kent distribution on the unit hemisphere, with the rotation matrix completed by taking perpendicular vectors, rotated by a uniformly random angle around the crystal c -axis. The resulting pole figures for these are displayed in 6.

Table 2: Kearns factors for the textures tested where R , T and N correspond to the y /rolling, x /transverse and z /normal directions respectively.

Texture	f_R	f_T	f_N
Strongly Normal	0.10	0.10	0.80
Weakly Normal	0.25	0.25	0.50
Uniform	0.33	0.33	0.33
Weakly Transverse	0.25	0.50	0.25
Strongly Transverse	0.10	0.80	0.10

A further 200 samples were generated for evaluating the effect of increasing the training dataset size on the performance of the network, maintaining the 9:1 ratio of training to validation samples. These samples were all generated with a uniform random texture, a regularity, α , of 0.7 and average volumetric grain size of $10 \mu\text{m}^3$, corresponding to a linear size of $2.7 \mu\text{m}$.

Crystal orientations are encoded using a mapping in \mathbb{R}^4 of the c -axis vector on the projective hemisphere, \mathbb{RP}^2 . Due to the symmetry of HCP crystal systems, rotation around the c -axis does not change the elastic properties therefore does not change the stress response so can be neglected. The following map was used to map the vector into \mathbb{R}^4 so equivalence relations are reflected in the numerical representation i.e. a rotation by π and 3π should yield the same representation and so should antipodal points. The following mapping was used:

$$r_1 = xy \quad (10)$$

$$r_2 = yz \quad (11)$$

$$r_3 = zx \quad (12)$$

$$r_4 = z^2 - (x^2 + y^2)/4 \quad (13)$$

where r_i are the resulting vector components and x , y and z are the components of the c -axis on the unit hemisphere.

This map produces a unique embedding for all points on the unit hemisphere so will not result in any orientations resolving to the same vector. This representation was found to improve the networks convergence so shortened training time compared to the axis representation or the use of Euler angles. The mapping is inspired by the projection of a Veronese surface into \mathbb{R}^4 , though the specific form of this map was chosen to make the distance between the mapping of two points similar to the angle between two points on the unit hemisphere.

Only the maximum principal stress is used in this work as this was considered the most useful parameter for cracking models. The full stress tensor may be modelled in future work using six channels in the output for each unique stress component. Prior to input into the network, stress is converted to a relative basis to ensure that values are $\sim O(1)$

$$\hat{\sigma} = \frac{\sigma_{\max} - \sigma_L}{\sigma_L} \quad (14)$$

where $\hat{\sigma}$ is the relative stress, σ_{\max} is the maximum principal stress of that voxel and σ_L is the applied stress.

5.2. Network Architecture

U-Net [21] was used as the basis for the neural network as it has been successfully adapted to a variety of problems beyond medical image segmentation while also being simple to implement. Several modifications have been made to allow 3D orientation maps to be fed in. The network is formed of blocks of two consecutive convolutions followed by either a max pooling or a transposed convolution for the encoding and decoding sides respectively, with channel concatenations between the encoder and decoder sides, allowing spatial information to filter forward across the network. 3D convolutions with a kernel size of $3 \times 3 \times 3$ throughout the network with zero padding of depth 1 to ensure that a size of 128^3 is maintained across the network. Downsampling on the encoder side is implemented using max pooling in $2 \times 2 \times 2$ regions, resulting in a halving of domain size along all dimensions. Upsampling on the decoder side is implemented using transposed convolutions with a kernel size of $2 \times 2 \times 2$ and a stride of 2, resulting in a doubling of domain size along all dimensions. The network is organised into blocks of two convolutions followed by an up/downsampling operation, with concatenation of downsampling inputs and upsampling outputs to feed spatial information into later parts of the network. The network is implemented using the PyTorch deep learning library with the CUDA 10.2 backend [18].

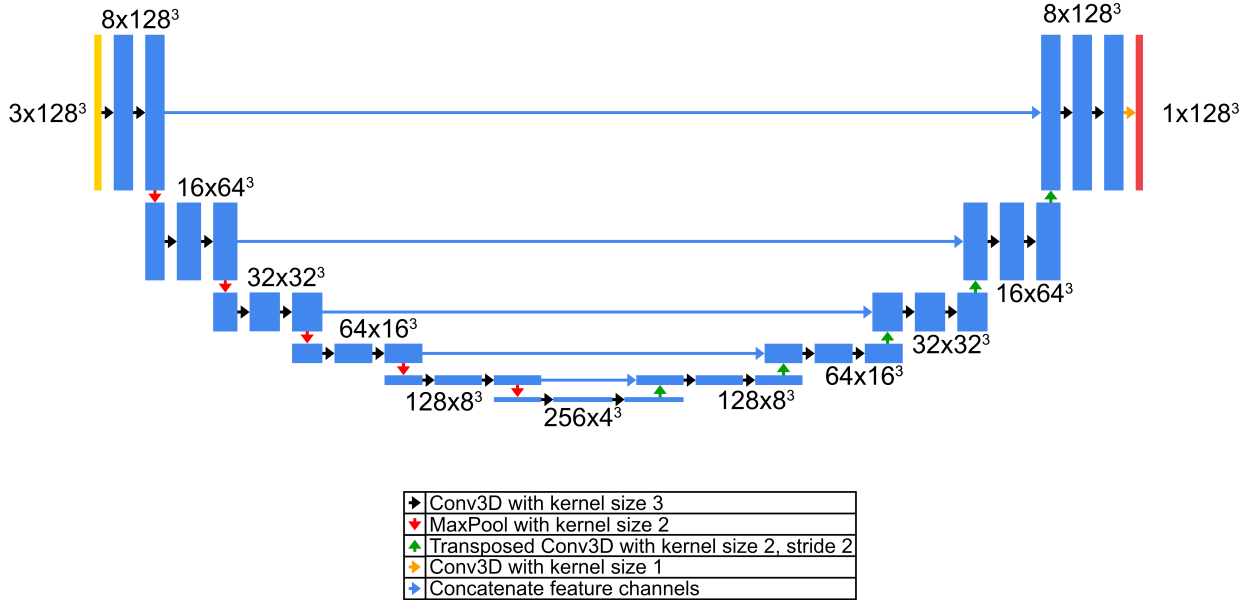


Figure 3: U-Net architecture used in this work. $N \times M^3$ indicates a layer has N feature channels and a voxel resolution of M .

An additional network architecture was also used to test the viability of reduced resolution training data. The network has the same overall architecture but the number of channels in the hidden layers has been doubled thanks to the memory freed from the smaller input and output tensors. Increasing the number of channels increases the expressibility of the network at the expense of increased memory usage and computation time and can improve convergence during training. This network was trained on 64^3 voxelisations of the same simulations, so any differences in network performance are due to the change in architecture and the limitations of a lower resolution.

5.3. Training

The network was trained for 250 epochs with a batch size of 10 to keep the VRAM usage within 16 GB. Network weights were optimised using PyTorch’s Adam optimiser, which combines standard gradient descent with a momentum term to prevent suboptimal solutions due to shallow local minima. The cost function used was a weighted sum of MSE in the absolute value of the stress and MSE in the spatial gradient of the stress, estimated using the 3D Sobel operator. The weighting between the absolute stress and the spatial gradient, β , increases from 0 for the first epoch to 0.5 for the last epoch in order to ease training of the network, since spatial gradients are very noisy at the beginning of training but become stable as training progresses. A factor κ is tuned manually to ensure that the terms both have similar magnitudes, with $\kappa = 60$ used for this work.

$$C(\hat{y}|x, y) = (1 - \beta)\text{MSE}(\hat{y}, y) + \beta\text{MSE}(\nabla_{\text{Sobel}}\hat{y}, \nabla_{\text{Sobel}}y)/\kappa \quad (15)$$

$$h = [1, 2, 1] \quad (16)$$

$$h' = [1, 0, -1] \quad (17)$$

$$H_{x,ijk} = h'_i h_j h_k \quad (18)$$

$$H_{y,ijk} = h_i h'_j h_k \quad (19)$$

$$H_{z,ijk} = h_i h_j h'_k \quad (20)$$

$$\nabla_{\text{Sobel},l}y = H_l \star y \quad (21)$$

Training was performed using an nVidia Quadro RTX 5000 16 GB and took around 5 hours to complete. As the spatial gradient of the stress is important in diffusion problems, the spatial gradient was monitored during validation as well as the stress itself.

Figure 4 shows that while error in the training stress continues to decrease until around 200 epochs, the error in the validation stress stops improving at around 75 epochs, suggesting error in the result is driven by a lack of training data as opposed to network saturation. This is unsurprising given the limited size of the training dataset so it is expected some configurations had insufficient representation in the dataset.

5.4. Error Estimation

Prediction errors in the output for each dataset are estimated by running the respective dataset through the trained neural network to produce the corresponding predictions. The absolute value of the difference for each sample between the prediction and the ground truth is then taken for further analysis, which will be referred to as the absolute residual. Mean errors for each dataset are calculated by taking the mean of the absolute residual for each sample while 99th percentile and maximum errors are calculated by sorting each sample into a list and taking the corresponding value. Additionally, basis errors are calculated using the assumption of the entire domain being composed of a single homogeneous crystal and finding the absolute residual.

For hypothesis testing, the mean of the means is taken then tested with the Shapiro-Wilkes test to ensure normality before a one-tailed t-test is used to quantify confidence that including the network is on average an improvement over assuming a homogeneous crystal for downstream processing.

6. Results and Discussion

6.1. 200 samples with 128^3 resolution

Figure 5 shows the output of the network and demonstrates a modest accuracy sacrifice in return, as shown in Figure 5.d for a significant speed-up compared to a full FEA treatment. In most cases, the network correctly identifies areas of higher/lower stress than the load with similar extremes of stress. In particular the network is very effective at identifying the stress gradients in the structure, displaying significantly lower error rates compared to the neglecting grain compatibility stresses. The bulk of the error in the output occurs due to a slight underestimate of around 2% of the most extreme stresses, likely because these occur at the grain boundaries where the data is discontinuous, which are difficult for the network to model.

The largest errors in the network are driven by errors across grain boundaries, mainly due to the discontinuities between stresses across the boundaries. This causes the network to assign a stress according to the neighbouring grain. There are also large errors, on the order of 10%, driven by the failure of the network to identify some

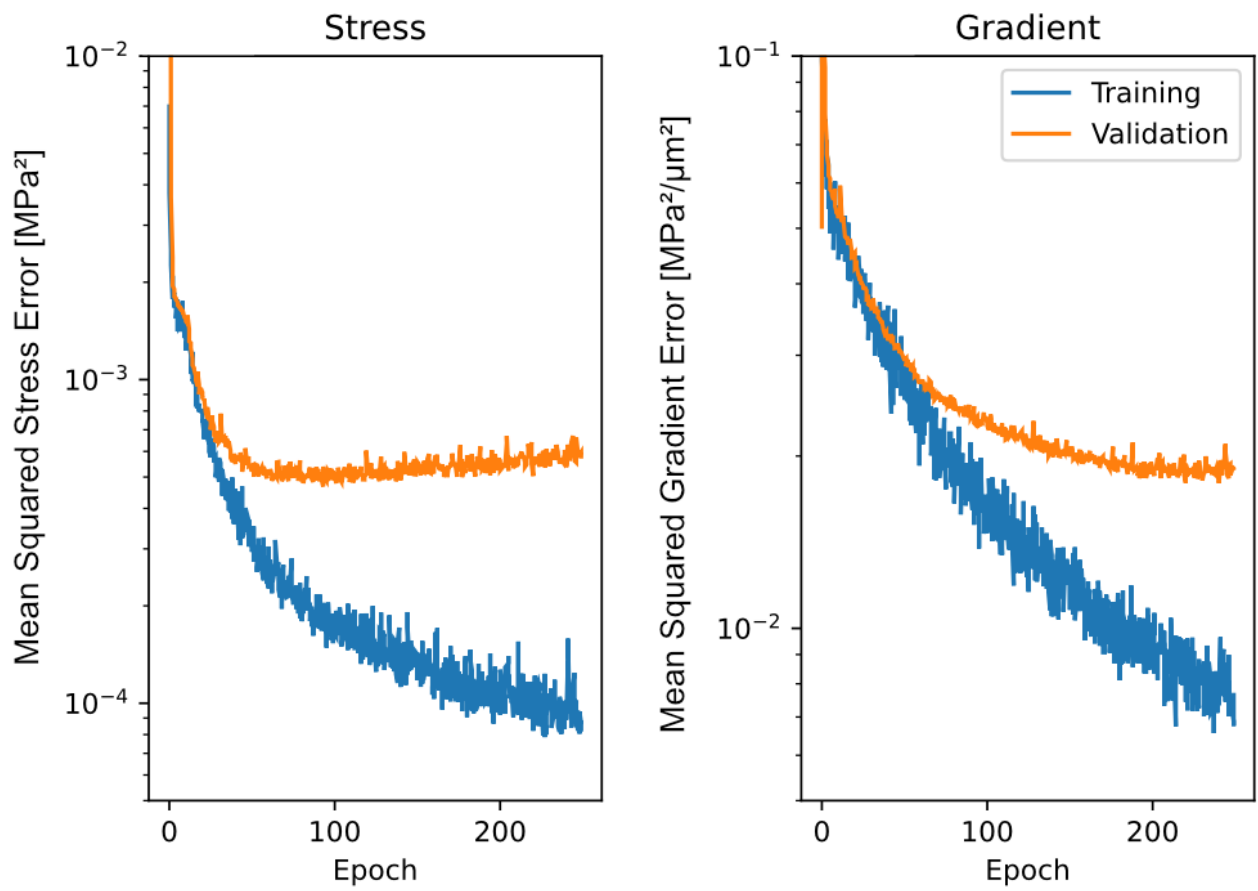


Figure 4: Mean squared error of the network output and spatial gradient of network output.

grain configurations as generators of large stress deviance, likely caused by a lack of representation of the particular environment in the training set.

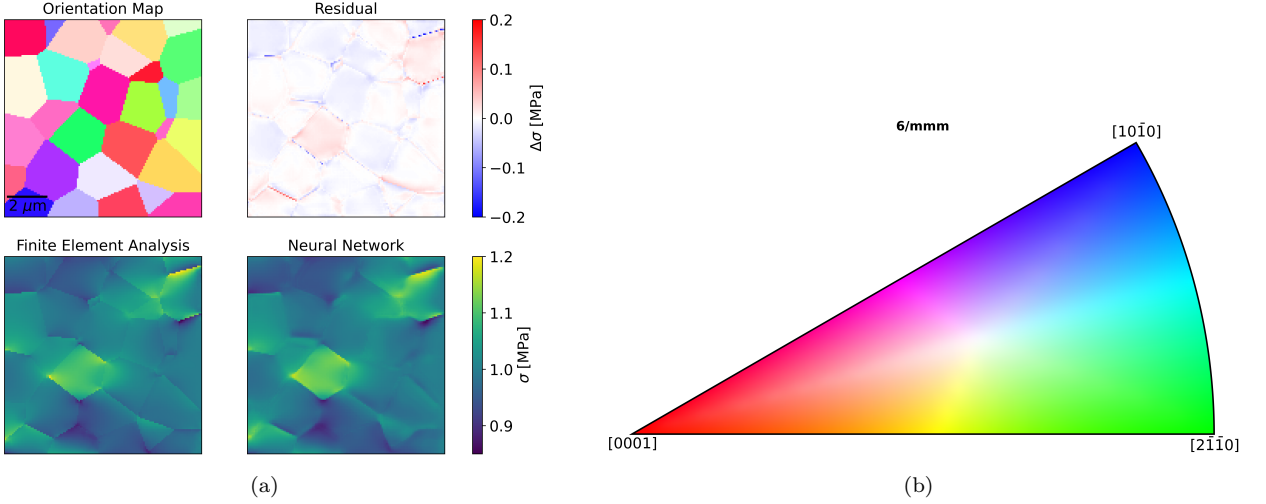


Figure 5: (a) Example of a network prediction where max. principal stresses shown are in MPa and cross-sections of the full 3D output are shown for legibility. (b) Inverse pole figure colour key

Table 3: Time taken for different processors to run the network on the dataset with a comparison to FEA.

Processor	Total Time (s)	Average Time (s)	Speed Up
Machine 1: Quadro RTX 5000	25	0.125	6226x
Machine 2: GeForce GTX 1050	77	0.385	2022x
Machine 2: i5 8300H	997	4.99	156x
Machine 1: Xeon Silver 4112	1693	8.47	92x
Machine 1: FEA on Xeon (Baseline)	155660	778	1x

To estimate the speed-up of the network, the time taken to execute the network across the entire dataset on two machines, Machine 1 on which the network was trained and Machine 2, which was a consumer grade laptop. Table 3 shows that using neural network inference is significantly faster than running a full FEA treatment, requiring around 1% of the time on the same CPU; both the FEA and the neural network were configured to use all 4 cores of the CPU. Introducing hardware acceleration for network evaluation improves time significantly further, with an order of magnitude improvement on an entry level GPU and further improvement with more powerful hardware.

6.2. Generalisation

Performance of the network was evaluated on datasets with different grain size regularities and orientation distributions to test generalisation of the network outside the training dataset. Since the network was trained using only a single regularity and orientation distribution, it was not guaranteed that the network would be applicable to the broader space of polycrystal combinations.

Figure 7 shows that the network error is insensitive to the regularity of the input grain structure for all textures, greatly expanding the range of inputs with which the network is applicable. Mean errors are essentially equal across all regularities, while 99th percentile deviation is not correlated with regularity at all, so is likely an artefact of small sample sizes. Insensitivity to the larger-scale structure of the input suggests that the network focuses on the local neighbourhood around a point to determine the stress in the output.

There are some significant differences in absolute error, however this is reflected in significant differences in the overall stress distribution due to texture. Figure 8 shows how the error populations differ with respect to texture, showing significant differences between the distributions. Some of this difference can be attributed to differences in stress distribution due to texture, with mean errors less than half of the mean stress and 99 percentile errors fairing much better in all cases except for strongly transverse case. Gradient errors display a significant reduction in error, particularly in the strongly transverse ($f_R=0.10$, $f_T=0.80$, $f_N=0.10$) texture, making this network particularly useful for further calculation involving stress gradients, for example diffusion problems. Figure 9 shows the errors relative

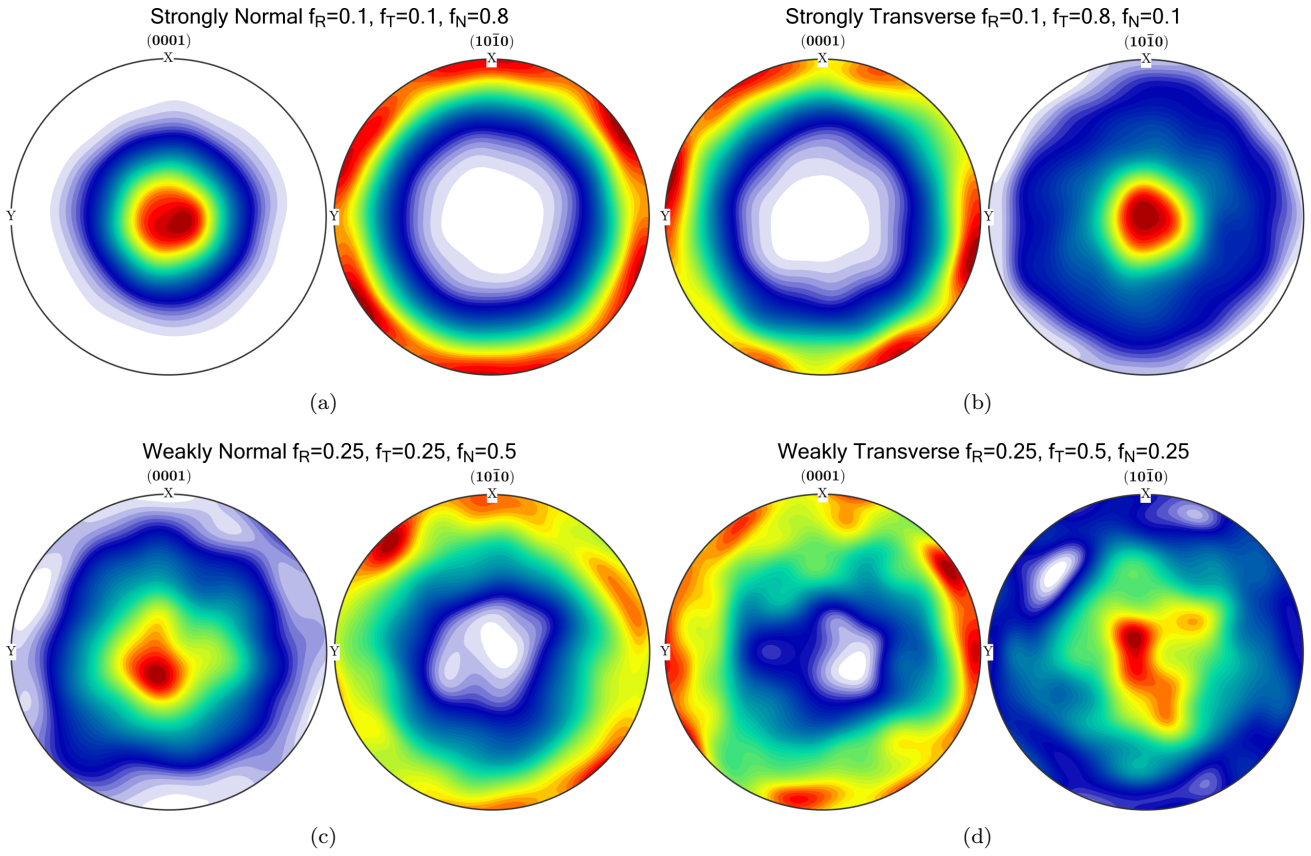


Figure 6: Pole figures for the non-uniformly textured datasets.

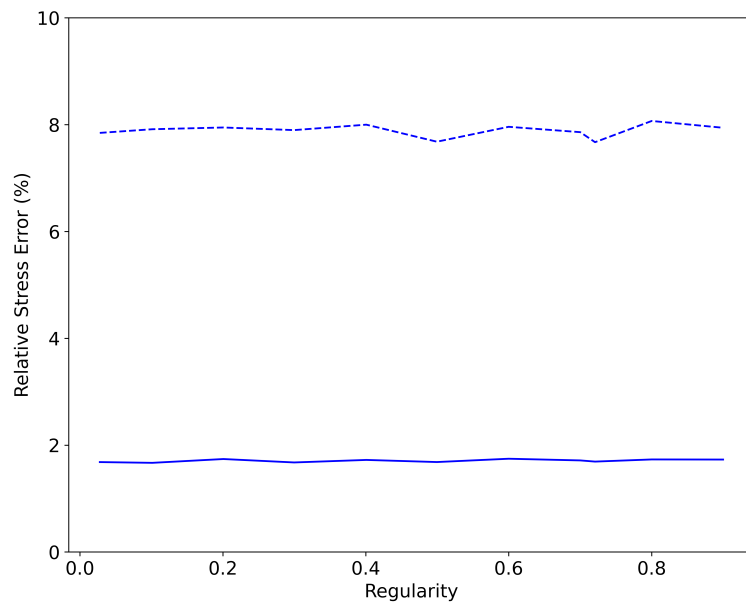


Figure 7: Graph of mean (solid) and 99th percentile (dashed) relative error for all textures against regularity.

to the ground truth stress, with some differences in the relative error between textures, with an increase in relative error from the Strongly Normal to Strongly Transverse texture, with the Uniform texture error being the median. These differences show that the network accuracy does depend on the texture of the input sample, though even in the worst case the errors are still below 10% of the applied stress for 99% of the input domain. Overall, given that the network was only trained using samples with random textures ($f_R=0.33$, $f_T=0.33$, $f_N=0.33$), comparable error distributions between the different textures suggests that the network can generalise to textures beyond that which it has been trained with. This significantly reduces the data generation needed to fully cover the input parameter space.

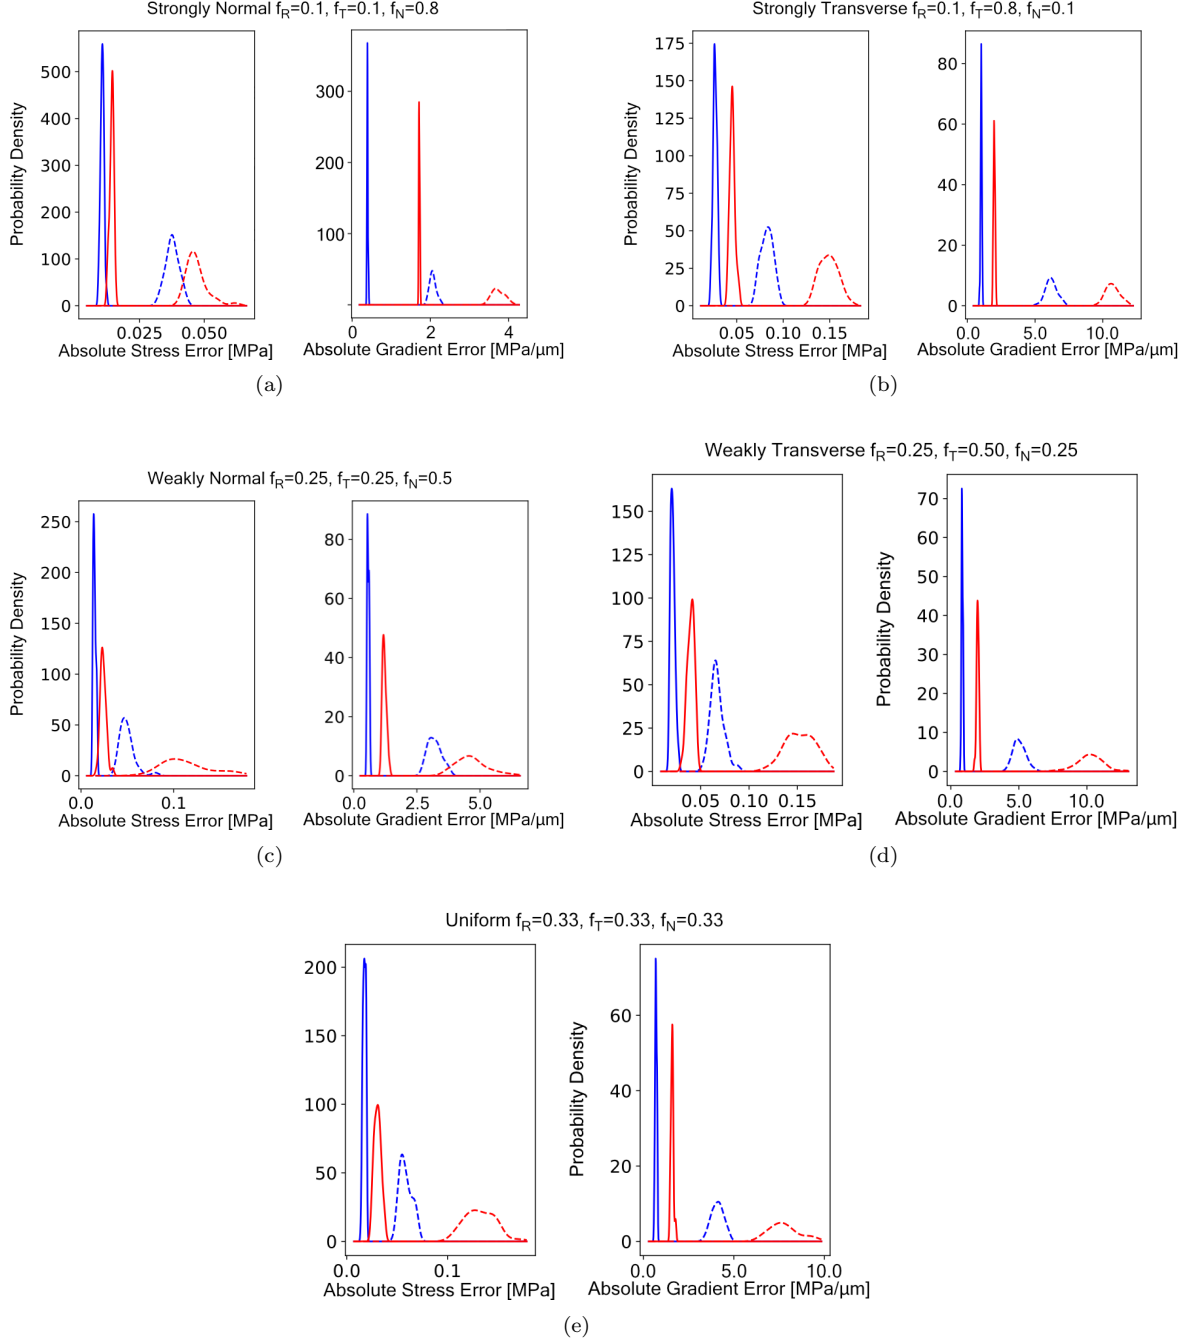


Figure 8: Probability density distributions of mean (solid) and 99th percentile (dashed) absolute errors of the stress and spatial gradients of the stress. The blue lines are the distributions of inference errors while red lines are the distributions of the absolute stress deviation from the loading stress.

A one-tailed t-test analysis was performed on mean errors of each texture to ensure statistical confidence that the mean error was less than the basis case. Since the mean of the mean errors was used, the Central Limit Theorem ensures that the criterion will be close to normally distributed so a t-test is valid and this was confirmed using the Shapiro-Wilkes test at the 1% significance level. For both the stress and the stress gradient and for all textures, the

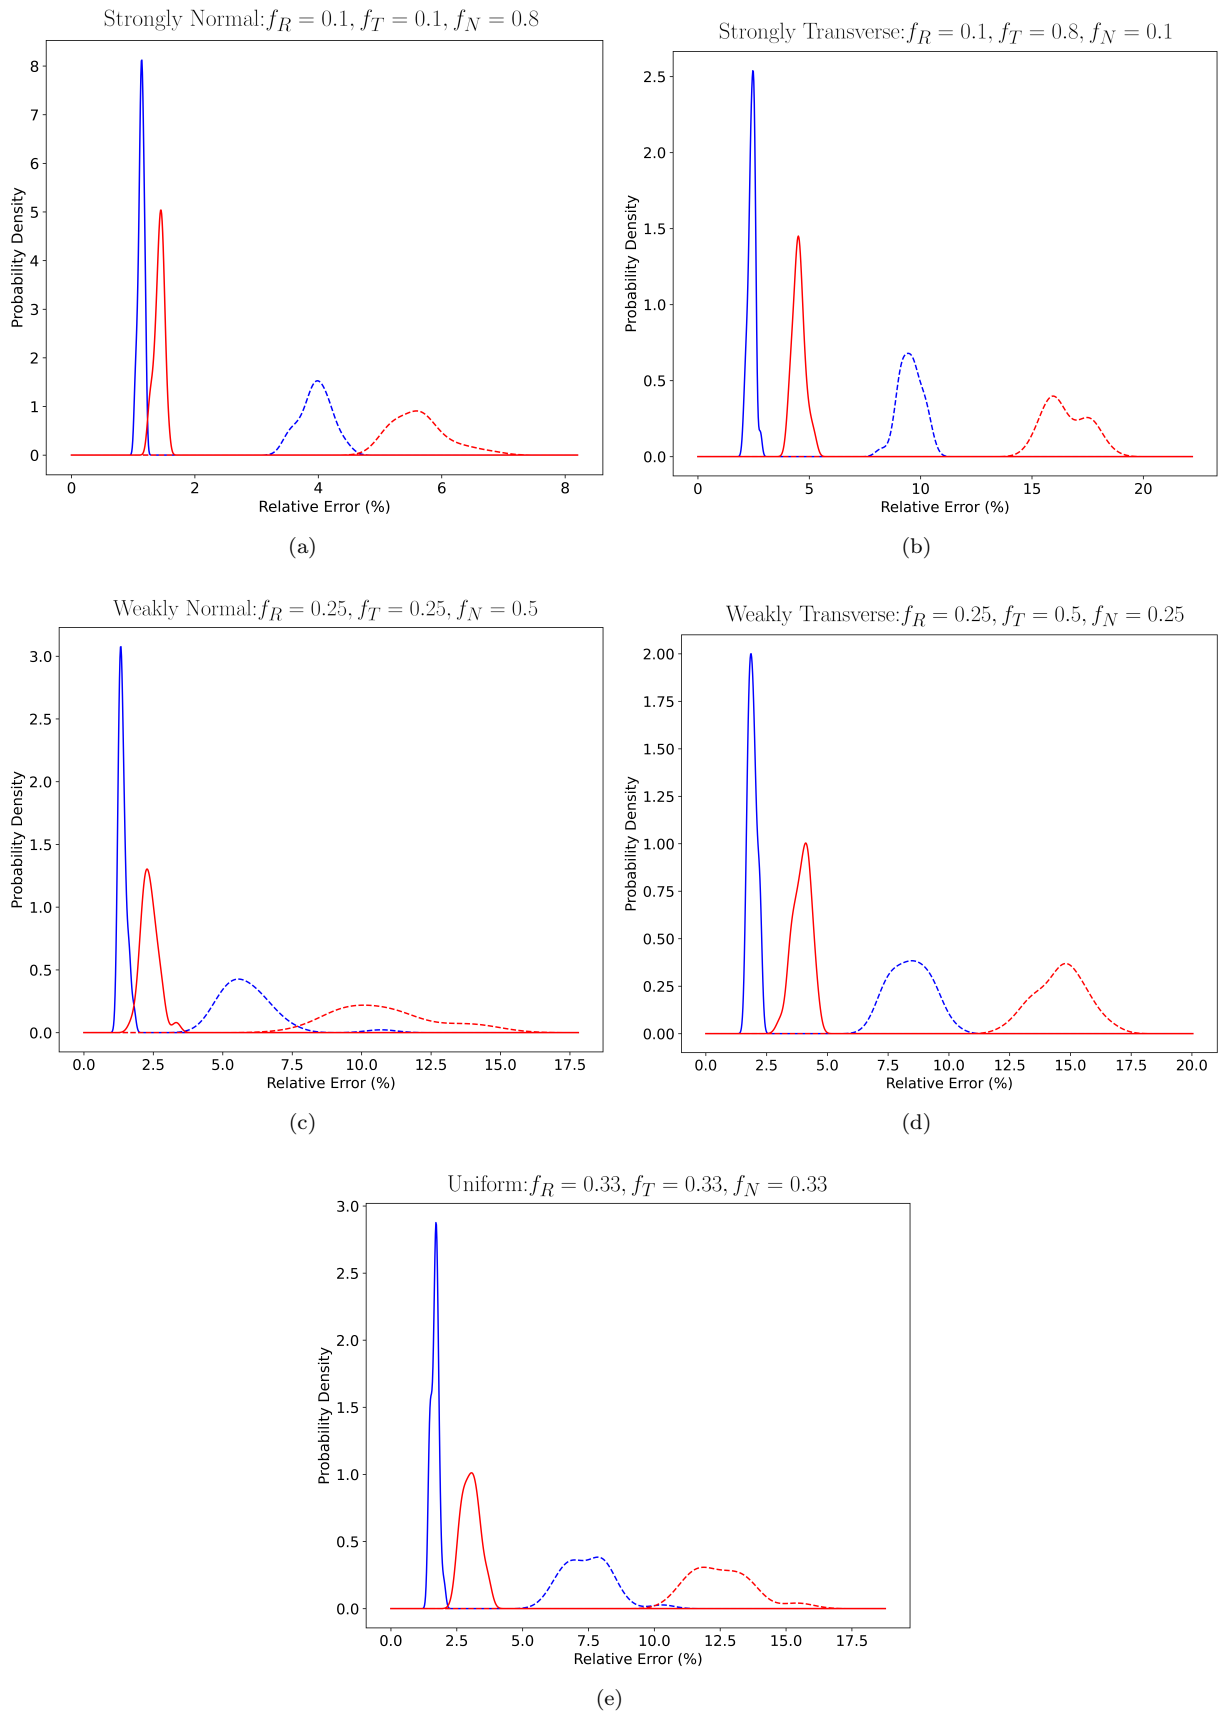


Figure 9: Probability density distributions of mean (solid) and 99th percentile (dashed) relative errors of the stress. The blue lines are the distributions of inference errors while red lines are the distributions of the absolute stress deviation from the loading stress.

mean error is statistically significantly better than the base case, with p-values at least as small as 1×10^{-19} . For the 99th percentile errors the Shapiro-Wilkes test was used to determine normality, however the distributions were not normal so t-test analysis is not valid. However, as can be seen from Fig. 8, the 99th percentile errors are significantly better than the basis case for all but the the strongly transverse texture stress. For the strongly transverse case, the 99th Percentile errors are still better in the majority of samples.

Table 4: T-test statistic and p-value for the mean stress and stress gradient being less then the mean of the basis case.

Value	Texture	Kearns	Statistic	p-Value
Stress	Strongly Transverse	$f_R=0.10, f_T=0.80, f_N=0.10$	-30.9	1.22×10^{-38}
	Weakly Transverse	$f_R=0.25, f_T=0.50, f_N=0.25$	-25.8	4.36×10^{-28}
	Uniform	$f_R=0.33, f_T=0.33, f_N=0.33$	-20.7	1.37×10^{-24}
	Weakly Normal	$f_R=0.25, f_T=0.25, f_N=0.50$	-15.4	8.10×10^{-20}
	Strongly Normal	$f_R=0.10, f_T=0.10, f_N=0.80$	-20.8	6.24×10^{-30}
Stress Gradient	Strongly Transverse	$f_R=0.10, f_T=0.80, f_N=0.10$	-65.3	9.07×10^{-58}
	Weakly Transverse	$f_R=0.25, f_T=0.50, f_N=0.25$	-66.2	1.59×10^{-53}
	Uniform	$f_R=0.33, f_T=0.33, f_N=0.33$	-95.0	1.46×10^{-67}
	Weakly Normal	$f_R=0.25, f_T=0.25, f_N=0.50$	-26.3	2.77×10^{-29}
	Strongly Normal	$f_R=0.10, f_T=0.10, f_N=0.80$	-199	3.68×10^{-83}

6.3. Evaluation of network trained on 64^3 samples applied to 128^3 samples

With volumetric data, reductions in resolution can yield significant memory and computation time reductions with a much smaller corresponding reduction of linear resolution. A reduction from 128^3 to 64^3 voxel resolution results in 1/8th the memory footprint while still retaining an adequate resolution of features in the samples, facilitating larger batch sizes during training. This stabilises the gradient during gradient descent and reduces the number of steps per epoch, resulting in much faster training times. Furthermore, the additional memory freed also allows for the training of larger networks than was feasible previously.

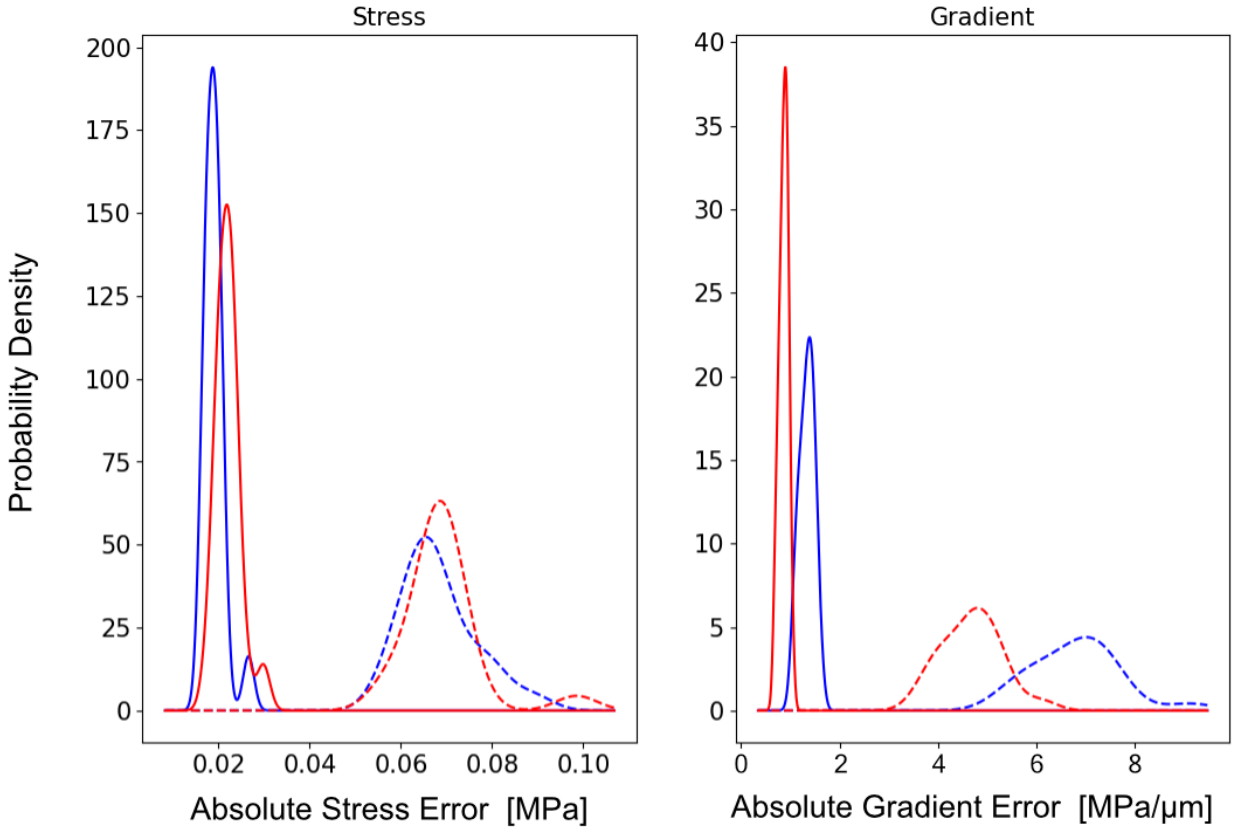


Figure 10: Probability density distribution of errors in stress and stress gradient for networks trained on 64^3 samples applied to 64^3 samples (blue) and 128^3 samples (red). Mean error is solid line while 99th percentile error is dashed.

The network maintains the ability to infer the stress distribution when evaluated on higher resolution samples than it was trained with. There is some reduction in accuracy, $O(1\%)$, suggesting that there is some scaling bias, which may be addressed by including voxel coordinates in the orientation map input. On the other hand, prediction of the stress gradient improves when evaluating the network on higher resolution samples. This is possibly due to the higher resolution allowing for easier reproduction of sharp or discontinuous features at grain boundaries, reducing error in these areas.

With some modifications, scale and resolution invariance would allow the network to be applied to much larger domain sizes than would be otherwise feasible thanks to the significant reduction in memory usage and computation time. Running a full FEA on a $10 \times 10 \times 10 \mu\text{m}^3$ domain with 100 grains typically consumes around 50 GB of RAM when using a mesh density that guaranteed convergence after a mesh sensitivity analysis. Meshes were seeded using a size of $0.3 \mu\text{m}$ resulting in about 300,000 elements. We assumed no mesh adaptivity as this increases complexity for the user as well as increased time to generate the mesh, although it depends on the software being used. Using FEA for larger domains would require either significantly more memory, a reduction in mesh resolution, and therefore accuracy, or paging the simulation out to storage which would considerably increase computation time. By contrast, estimation using the network requires around 200 MB of VRAM for the model itself and an additional 150 MB per 128^3 sample. With a 16 GB GPU, this means a domain of $40 \times 40 \times 40 \mu\text{m}^3$ with the same resolution, or an $80 \times 80 \times 80 \mu\text{m}^3$ domain with half resolution, like in the 64^3 case. This could be run entirely on the GPU, using around 9.6 GB of VRAM. Running the network on the CPU would increase inference times but would allow the network to access more memory, allowing for even larger domain sizes, possibly up to $140 \times 140 \times 140 \mu\text{m}^3$ within the same 50 GB memory footprint of the FEA models, allowing for estimation for a domain 2744x larger than would be possible using FEA on the same hardware.

6.4. Comparison between 200 and 400 samples at 64^3 resolution

As 200 samples is a small dataset, in the context of machine learning, an additional training run including a further 200 samples was run. Doubling the size of the dataset presents a greater variety of scenarios for the network to learn on, easing training and preventing failure cases. 64^3 was used as the sample resolution so larger batch sizes could be used to speed up training.

Table 5 shows that training with the additional samples results in an increase of approximately 15% in model accuracy compared to the 200 sample model for both the stress and stress gradient. Figure 11 shows the mean and 99th percentile error distributions for the predicted stress and stress gradient for 200 (blue) and 400 (red) training examples, showing a consistent decrease in the errors. Additionally, the network output is much clearer with the larger dataset, representing features around grain boundaries more accurately. Maximum errors are relatively unchanged, possibly due to the maximum errors being driven by singularities in the FEA simulations. While error statistics are not significantly affected, qualitatively, the resulting estimates are much clearer with the additional training data. It’s possible that this may lead to improvements in downstream simulations so further work should investigate this effect on simulation accuracy.

Table 5: Mean errors in stress and stress gradient for networks trained on 200 samples (ZircNet-200) and 400 samples (ZircNet-400).

Error Statistic		ZircNet-200	ZircNet-400	% Difference
Average Absolute Stress (MPa)	Mean	0.0213	0.0185	-13.1%
	99th Percentile	0.0815	0.0706	-13.4%
	Max	0.2518	0.2552	+1.3%
Average Absolute Stress Gradient (MPa/ μm)	Mean	0.1568	0.1326	-15.4%
	99th Percentile	0.7781	0.6735	-13.4%
	Max	3.5420	3.4034	-3.9%

6.5. Failure Cases

Occasionally the network will fail to identify a grain that generates significant stress, leading to a large underestimate in the compatibility stress deviation. It is likely that these failure cases reflect a lack of representation of similar environments in the training set. Conversely the network also misidentifies a grain as a significant stress generator that does not significantly deviate. Again, this is likely due to a lack of representation in the training set. These errors persist through a doubling in the number of channels in the network though they are reduced in severity, suggesting that this is an issue in the size of the training dataset.

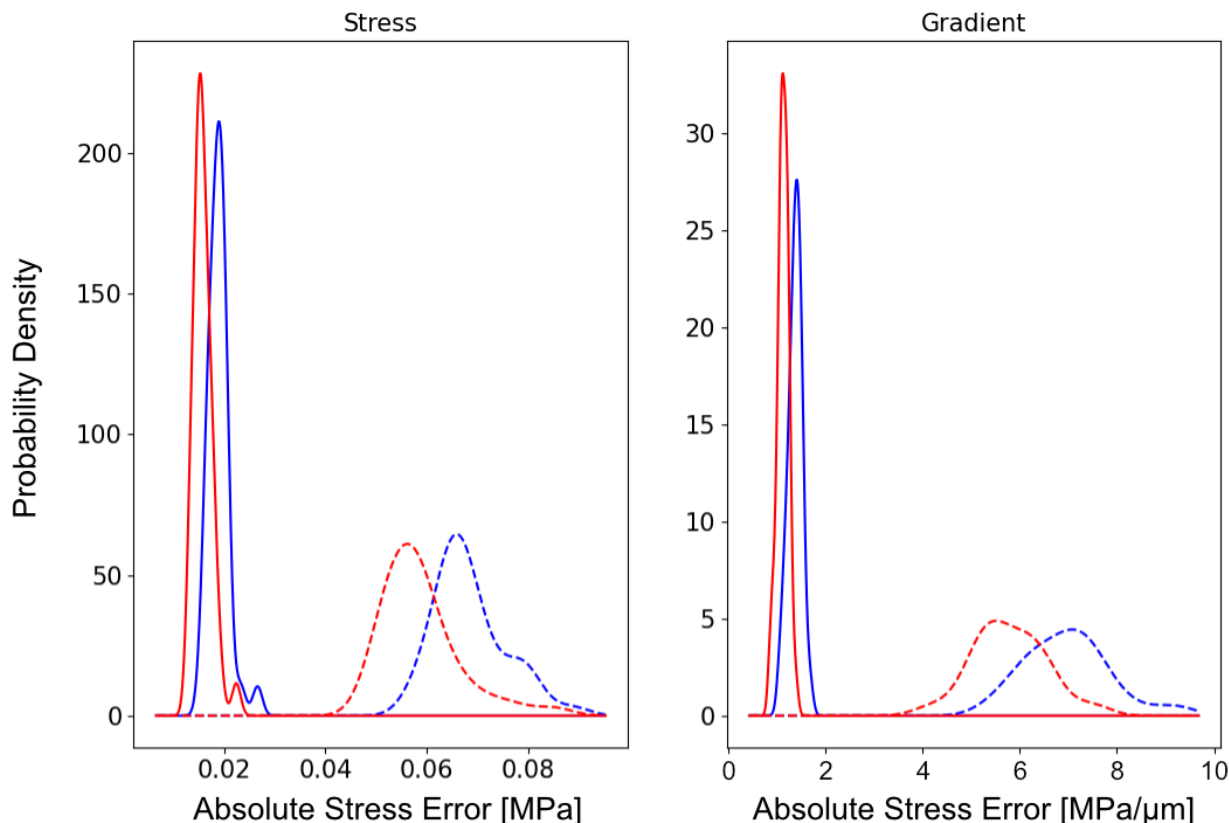


Figure 11: Probability density distribution of errors in stress and stress gradient for networks trained on 200 samples (blue) and 400 samples (red). Mean error is solid while 99th percentile error is dashed.

In the false positive example shown in figure 12, the grains surrounding the false positive stress grain as well as the grain itself are oriented approximately with the loading stress parallel to the basal plane, so the stiffness with respect to the load will be approximately equal, however the orientation embedding is significantly different.

Analogously, for the false negative failure case, shown in figure 13, the top two grains are both oriented with the basal plane parallel to the loading stress while the bottom grain is oriented with the prismatic axis parallel to the loading stress, so the stiffness with respect to the load will be different. In this case the orientation embeddings are all different, so it is unclear why the network fails to detect a deviation.

Taking these failure cases into account, future improvements on this work could investigate embeddings which better reflect the stiffness dependence on orientation with respect to the loading axis. Direct input of the stiffness tensor may also be a useful direction to follow, with the caveat that it would require significantly more input channels to the network.

Another intriguing failure to match the training data are in single points in which FEA predicts a large localised stress. These points tend to occur where the mesh elements are of low quality, with small angle tetrahedra prone to numerical instability, so it is possible that these are singularities introduced by the solver, which the network does not replicate, suggesting that a neural network approach is more resilient to numerical singularities. Additionally, in real Zr microstructures, plasticity effects mean that significant stress concentrations are unlikely to exist. This means the smoothness is positive feature for adaptation of the model to emulating plasticity simulations.

6.6. Limitations of current model

Currently a “naive” learning approach is used which does not exploit additional constraints which can be imposed on the stress field from the physics of the system. By exploiting the underlying PDEs as an additional cost term, the dependency on data could be vastly reduced compared to “naive” training methods, greatly expanding the generality achievable by these methods.

Additionally, the current network architecture fails to converge when attempting to train with the full stress

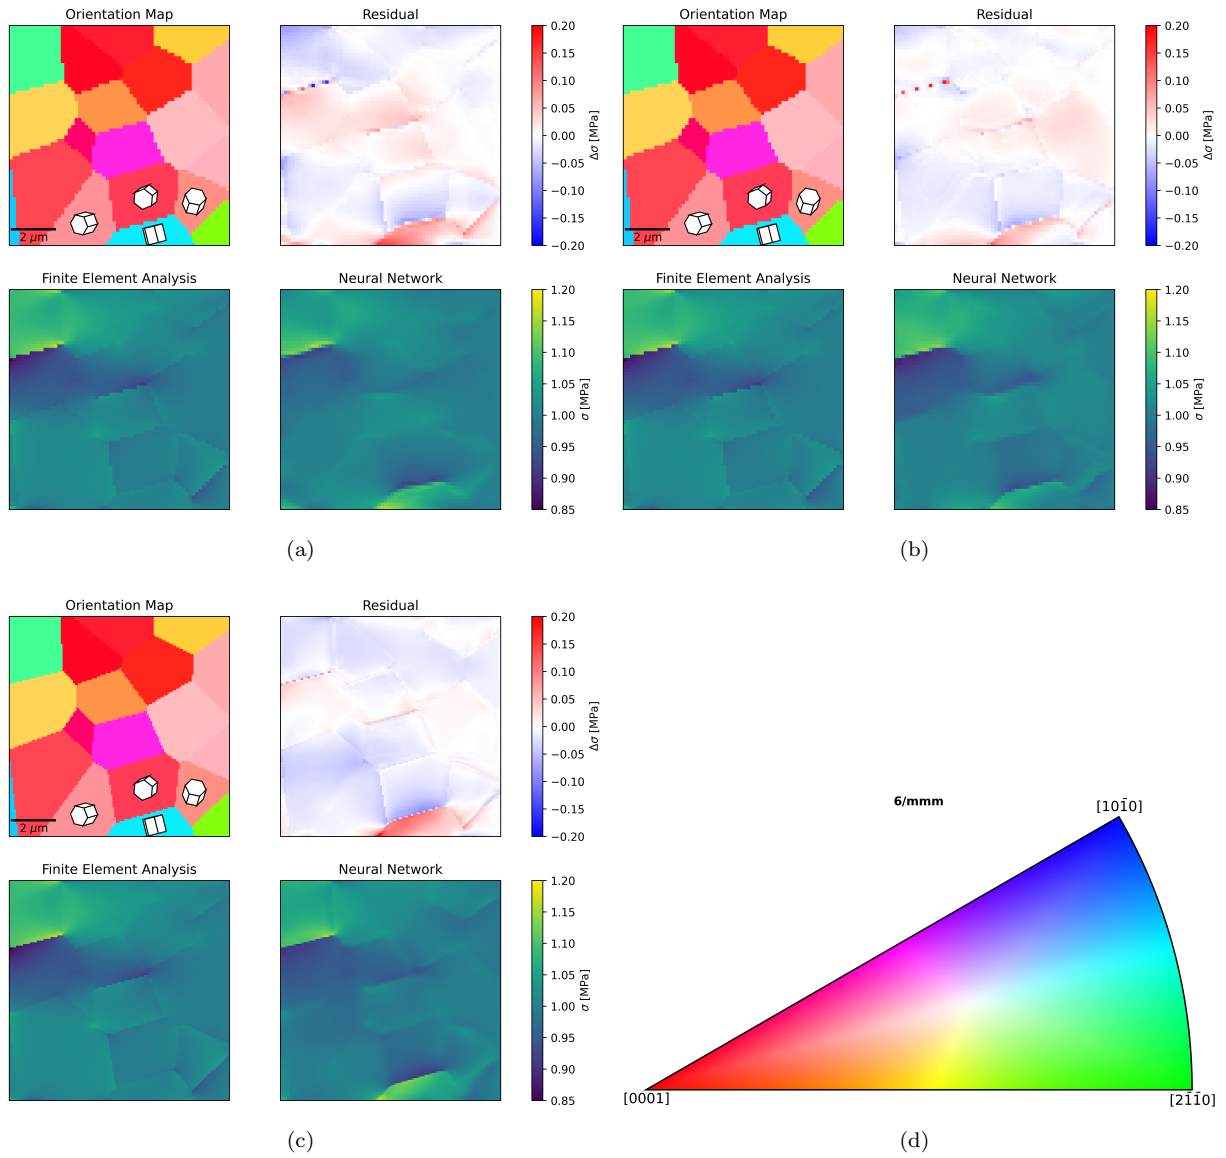


Figure 12: A false positive failure case for the network trained on (a) 200 samples at 64^3 resolution, (b) 400 samples at 64^3 resolution, and (c) 200 samples at 128^3 resolution. The colour key for the orientation map is displayed in (d). Stresses are the max principal stresses in the network prediction and ground truth case.

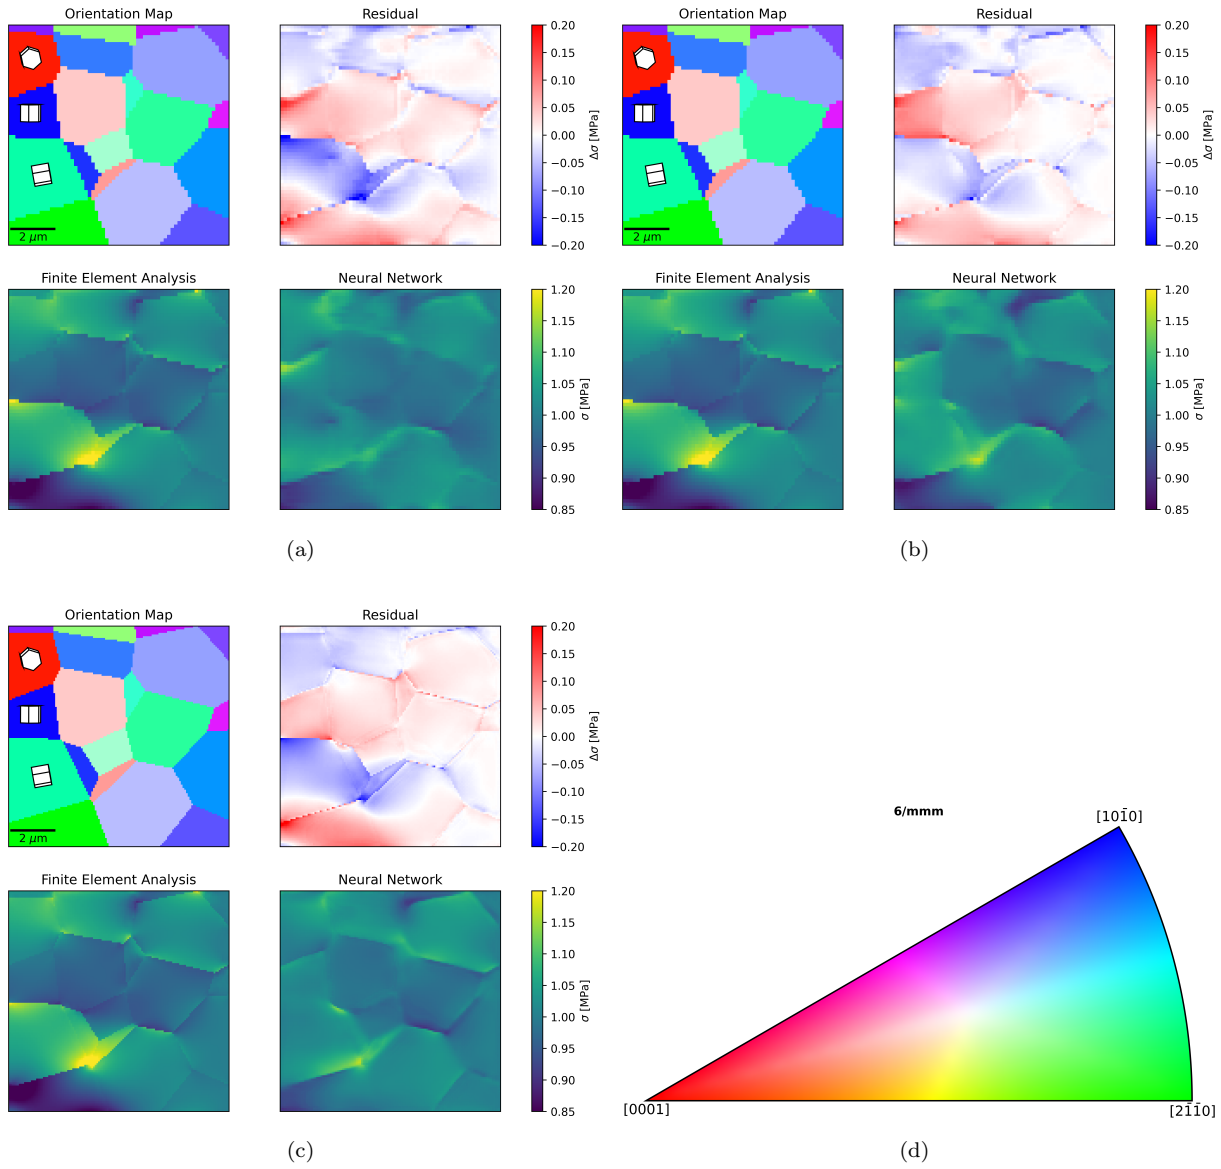


Figure 13: A false negative failure case for the network trained on (a) 200 samples at 64^3 resolution, (b) 400 samples at 64^3 resolution, and (c) 200 samples at 128^3 resolution. The colour key for the orientation map is displayed in (d). Stresses are the max principal stresses in the network prediction and ground truth case.

tensor as the output. While it is possible to train an individual network for each component, this would increase computational time linearly with the number of components and may limit correlation discovery between components. The network also fails to converge when attempting to train with the displacement map as the output, limiting the effectiveness of physics informed learning as displacement boundary conditions cannot be enforced.

7. Conclusions

The network architecture and dataset generation combination presented is a proof of concept for a more complete system for industrial implementation. Currently, the network output is limited to the maximum principal stress, however, a full stress and strain tensor combination would be much more useful for practical problems and is the immediate focus on expansion of this work.

Generalisation of the network beyond microstructure characteristics represented in the training set has been demonstrated, greatly expanding its applicability beyond the training set. Both grain size distributions and textures beyond the training set resulted in similar inference errors, around 1% on average, and within 10% for 99% of the simulation volume, when compared to those inside the training set.

The network has demonstrated that it can capture the overall form of the microscale stress based upon only a crystal orientation map as input while requiring up to 6000x less time and 6500x less memory compared to traditional FEA methods, allowing for the possibility of simulating much larger domain sizes than would be possible with FE.

This work uses a “naive” learning approach which does not utilise the constraining PDEs of linear elasticity. Despite the lack of explicit physics, the model is still able to capture the principal stress fields present in the microstructure with reasonable accuracy.

8. Acknowledgements

This work was completed with funding from EPSRC Centre for Doctoral Training in Nuclear Energy Futures (EP/S023844/1) and Rolls-Royce Plc. (UK).

9. Data Availability

The processed data and models are available on request.

References

- [1] U. Ayachit. *The ParaView Guide: A Parallel Visualization Application*. Kitware, 2015. ISBN 978-1930934306.
- [2] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. Language Models are Few-Shot Learners. *Advances in Neural Information Processing Systems*, 2020-December, 5 2020. URL <https://arxiv.org/abs/2005.14165v4>.
- [3] W. L. Buntine and A. S. Weigend. Bayesian Back-Propagation. *Complex Systems*, 5(6):603–643, 1991. URL https://www.complex-systems.com/abstracts/v05_i06_a04/.
- [4] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems*, 2(4):303–314, 12 1989. ISSN 09324194. doi: 10.1007/BF02551274. URL <https://link.springer.com/article/10.1007/BF02551274>.
- [5] P. M. R. DeVries, T. B. Thompson, and B. J. Meade. Enabling large-scale viscoelastic calculations via neural network acceleration. *Geophysical Research Letters*, 44(6):2662–2669, 3 2017. ISSN 00948276. doi: 10.1002/2017GL072716. URL <http://doi.wiley.com/10.1002/2017GL072716>.

- [6] X. Glorot, A. Bordes, and Y. Bengio. Deep Sparse Rectifier Neural Networks. In G. Gordon, D. Dunson, and M. Dudík, editors, *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 315–323, Fort Lauderdale, FL, USA, 4 2011. JMLR Workshop and Conference Proceedings. URL <https://proceedings.mlr.press/v15/glorot11a.html>.
- [7] X. Gong, S. Chang, Y. Jiang, and Z. Wang. AutoGAN: Neural Architecture Search for Generative Adversarial Networks. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3223–3233. IEEE, 10 2019. ISBN 978-1-7281-4803-8. doi: 10.1109/ICCV.2019.00332. URL <https://ieeexplore.ieee.org/document/9010885/>.
- [8] B. Hanin and M. Sellke. Approximating Continuous Functions by ReLU Nets of Minimal Width. 10 2017. URL <http://arxiv.org/abs/1710.11278>.
- [9] M. F. Kasim, D. Watson-Parris, L. Deaconu, S. Oliver, P. Hatfield, D. H. Froula, G. Gregori, M. Jarvis, S. Khatiwala, J. Korenaga, J. Topp-Mugglestone, E. Viezzer, and S. M. Vinko. Building high accuracy emulators for scientific simulations with deep neural architecture search. *Machine Learning: Science and Technology*, 3 (1):015013, 3 2022. ISSN 2632-2153. doi: 10.1088/2632-2153/ac3ffa. URL <https://iopscience.iop.org/article/10.1088/2632-2153/ac3ffa>.
- [10] S. Kullback and R. A. Leibler. On Information and Sufficiency. *The Annals of Mathematical Statistics*, 22(1): 79–86, 3 1951. ISSN 0003-4851. doi: 10.1214/aoms/1177729694. URL <https://projecteuclid.org/euclid.aoms/1177729694>.
- [11] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2323, 1998. ISSN 00189219. doi: 10.1109/5.726791.
- [12] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein. Visualizing the Loss Landscape of Neural Nets. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31 (NeurIPS 2018)*, pages 6389–6399, Montréal, 2018. Curran Associates, Inc. ISBN 9781510884472. URL <https://proceedings.neurips.cc/paper/2018/file/a41b3bb3e6b050b6c9067c67f663b915-Paper.pdf>.
- [13] L. Liang, M. Liu, C. Martin, and W. Sun. A deep learning approach to estimate stress distribution: a fast and accurate surrogate of finite-element analysis. *Journal of the Royal Society Interface*, 15(138), 1 2018. ISSN 17425662. doi: 10.1098/rsif.2017.0844.
- [14] Q. Liu, J. Li, and J. Liu. ParaView visualization of Abaqus output on the mechanical deformation of complex microstructures. *Computers & Geosciences*, 99:135–144, 2 2017. ISSN 0098-3004. doi: 10.1016/J.CAGEO.2016.11.008.
- [15] Y. Liu, W. Wan, and F. P. Dunne. Characterisation and modelling of micro- and macroscale creep and strain rate sensitivity in Zircaloy-4. *Materials Science and Engineering: A*, 840:142981, 4 2022. ISSN 0921-5093. doi: 10.1016/J.MSEA.2022.142981.
- [16] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 07-12-June-2015, pages 3431–3440. IEEE Computer Society, 10 2015. ISBN 9781467369640. doi: 10.1109/CVPR.2015.7298965.
- [17] Z. Nie, H. Jiang, and L. B. Kara. Stress Field Prediction in Cantilevered Structures Using Convolutional Neural Networks. *Journal of Computing and Information Science in Engineering*, 20(1), 3 2020. ISSN 1530-9827. doi: 10.1115/1.4044097.
- [18] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [19] M. Patel, L. Reali, A. P. Sutton, D. S. Balint, and M. R. Wenman. A fast efficient multi-scale approach to modelling the development of hydride microstructures in zirconium alloys. *Computational Materials Science*, 190:110279, 4 2021. ISSN 0927-0256. doi: 10.1016/J.COMMATSCI.2021.110279.

- [20] M. M. Rahaman, C. Li, Y. Yao, F. Kulwa, M. A. Rahman, Q. Wang, S. Qi, F. Kong, X. Zhu, and X. Zhao. Identification of COVID-19 samples from chest X-Ray images using deep learning: A comparison of transfer learning approaches. *Journal of X-Ray Science and Technology*, 28(5):821–839, 1 2020. ISSN 0895-3996. doi: 10.3233/XST-200715.
- [21] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 9351, pages 234–241. Springer Verlag, 2015. ISBN 9783319245737. doi: 10.1007/978-3-319-24574-4{_}28.
- [22] D. E. Rumelhart, G. E. Hint, and R. J. Williams. Learning Internal Representations by Error Propagation. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 1: Foundations*, page 318–362. MIT Press, Cambridge, MA, USA, 1986. ISBN 026268053X.
- [23] D. Tromans. Elastic Anisotropy of HCP Metal Crystals and Polycrystals. *Int. J. Res. Rev. Appl. Sci.*, 6(4):14, 2011. URL https://www.arpapress.com/Volumes/Vol6Issue4/IJRRAS_6_4_14.pdf.
- [24] P. J. Withers and H. K. Bhadeshia. Residual stress. Part 1 – Measurement techniques. <http://dx.doi.org/10.1179/026708301101509980>, 17(4):355–365, 2001. ISSN 02670836. doi: 10.1179/026708301101509980. URL <https://www.tandfonline.com/doi/abs/10.1179/026708301101509980>.
- [25] Q. Xu, Z. Nie, H. Xu, H. Zhou, H. R. Attar, N. Li, F. Xie, and X.-J. Liu. SuperMeshing: A New Deep Learning Architecture for Increasing the Mesh Density of Physical Fields in Metal Forming Numerical Simulation. *Journal of Applied Mechanics*, 89(1):11002–11003, 1 2022. ISSN 0021-8936. doi: 10.1115/1.4052195. URL <https://asmedigitalcollection.asme.org/appliedmechanics/article/89/1/011002/1115883/SuperMeshing-A-New-Deep-Learning-Architecture-for>.
- [26] Z. Yang, C.-H. Yu, and M. J. Buehler. Deep learning model to predict complex stress and strain fields in hierarchical composites. *Science Advances*, 7(15), 4 2021. doi: 10.1126/SCIADV.ABD7416. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8034856/>.
- [27] P. Zhang, D. Balint, and J. Lin. An integrated scheme for crystal plasticity analysis: Virtual grain structure generation. *Computational Materials Science*, 50(10):2854–2864, 8 2011. ISSN 09270256. doi: 10.1016/j.commatsci.2011.04.041.
- [28] X. Zhao, Z. Gong, Y. Zhang, W. Yao, and X. Chen. Physics-informed convolutional neural networks for temperature field prediction of heat source layout without labeled data. *Engineering Applications of Artificial Intelligence*, 117:105516, 1 2023. ISSN 0952-1976. doi: 10.1016/J.ENGAPPAL.2022.105516. URL <http://arxiv.org/abs/2109.12482>.
- [29] H. X. Zhu, P. Zhang, D. Balint, S. M. Thorpe, J. A. Elliott, A. H. Windle, and J. Lin. The effects of regularity on the geometrical properties of Voronoi tessellations. *Physica A: Statistical Mechanics and its Applications*, 406:42–58, 7 2014. ISSN 03784371. doi: 10.1016/j.physa.2014.03.012.