

Full Length Article



Exploring descriptors for titanium microstructure via digital fingerprints from variational autoencoders

Michael D. White^{a,b,*}, Gowtham Nimmal Haribabu^{a,c}, Jeyapriya Thimukonda Jegadeesan^c, Bikramjit Basu^c, Philip J. Withers^{a,b}, Chris P. Race^{b,d}

^a Department of Materials, University of Manchester, Oxford Road, Manchester, M13 9PL, UK

^b Henry Royce Institute, University of Manchester, Oxford Road, Manchester, M13 9PL, UK

^c Materials Research Centre, Indian Institute of Science, Mathikere, Bangalore, 560012, India

^d Department of Materials Science and Engineering, University of Sheffield, Mappin Street, Sheffield, S1 3JD, UK

ARTICLE INFO

Dataset link: <https://zenodo.org/records/5714384>

Keywords:

Machine learning
Variational autoencoders
Microstructure characterisation
Feature learning
Microstructural fingerprinting

ABSTRACT

Microstructure is key to controlling and understanding the properties of materials, but traditional approaches to describing microstructure capture only a small number of features. We require more complete descriptors of microstructure to enable data-centric approaches to materials discovery, to allow efficient storage of microstructural data and to assist in quality control in metals processing. The concept of *microstructural fingerprinting*, using machine learning (ML) to develop quantitative, low-dimensional descriptors of microstructures, has recently attracted significant attention. However, it is difficult to interpret conclusions drawn by ML algorithms, which are often referred to as “black boxes”. For example, convolutional neural networks (CNNs) can be trained to make predictions about a material from a set of microstructural image data, but the feature space that is learned is often used uncritically and adopted without any validation.

Here we explore the use of variational autoencoders (VAEs), comprising a pair of CNNs, which can be trained to produce microstructural fingerprints in a continuous latent space. The VAE architecture also permits the reconstruction of images from fingerprints, allowing us to explore how key features of microstructure are encoded in the latent space of fingerprints. We develop a VAE architecture based on ResNet18 and train it on two classes of Ti-6Al-4V optical micrographs (bimodal and lamellar) as an example of an industrially important alloy where microstructural control is critical to performance. The latent/feature space of fingerprints learned by the VAE is explored in several ways, including by supplying interpolated and randomly perturbed fingerprints to the trained decoder and via dimensionality reduction to explore the distribution and correlation of microstructural features within the latent space of fingerprints.

We demonstrate that the fingerprints generated via the trained VAE exhibit smooth, interpolable behaviour with stability to local perturbations, supporting their suitability as general purpose descriptors for microstructure. The analysis of computational results uncover that key properties of the microstructures (volume fraction and grain size) are strongly correlated with position in the encoded feature space, supporting the use of VAE fingerprints for quantitative exploration of process–structure–property relationships.

1. Introduction

In the field of metallurgy and materials science, process–structure–property (PSP) linkages are instrumental in guiding material design for targeted applications [1,2]. Despite the importance of PSP linkages, a rigorous mathematical framework is not currently available for systematic analysis in this context [3]. The key problem is that, whilst compositional processing parameters and measured properties are inherently described by numbers, microstructure lacks a comprehensive numerical descriptor. A central impediment is that the characteristic

microstructural features exhibit heterogeneity over a wide range of size, spatial, and temporal scales. From the application standpoint, it is important to identify a subset of salient measures of internal structure that can be tracked through a material’s processing history, that capture the dominant influences on the targeted properties. In conventional microstructural analysis, some quantitative methods are used, but these generally rely on metrics applied to image data, such as grain size, phase fraction and correlation coefficients. This can provide some crucial information, but the metrics which are suitable in each case

* Corresponding author at: Department of Materials, University of Manchester, Oxford Road, Manchester, M13 9PL, UK.
E-mail address: michael.white-3@postgrad.manchester.ac.uk (M.D. White).

will depend on the morphology of the microstructure. For example, interlamellar spacing is useful to describe lamellar microstructures, but is redundant when considering an equiaxed microstructure. In any case, such measures embody only a tiny fraction of the information contained in a microstructural image. A quantitative description of microstructure that is independent of morphology (a microstructural fingerprint) and embodies a full range of features would enable direct comparisons between microstructures with different morphologies, and provide new methods for constructing PSP relationships [3,4], or for quantifying the deviation of a given microstructure from some ideal standard in a quality control procedure.

In the past two decades, reasonable progress has been achieved in the quantification and low-dimensional representation of microstructures, over larger processing and material compositional windows [5–7]. These advances have been possible with the use of concepts and toolsets from data science and informatics [8,9]. The first potential benefit is realised with the use of automated feature engineering of the hierarchical material structures, e.g., establishing low-dimensional representations that provide maximum value in capturing high-fidelity PSP linkages. A systematic and comprehensive quantification of the material structure is now possible, for example by performing statistical analysis of image features [10] or by combining the physics-inspired framework of n -point spatial correlations (n -point statistics) with machine learning approaches, such as principal component analysis (PCA) [11].

A machine learning approach that has not yet been fully explored, in the context of microstructural fingerprinting, is variational autoencoders (VAEs), which comprise a pair of convolutional neural networks (CNNs), referred to individually as the encoder and the decoder. VAEs were first introduced by Kingma and Welling [12]. Initial applications focussed on generating images of individual objects, particularly human faces, utilising datasets such as the CelebA dataset [13]. More recently, attention has shifted towards machine learning applications in materials science, such as quantification of microstructure and the development of new PSP relationships. The key difference is that microstructural image data are such that the entire image field contains potentially meaningful information, rather than an image of a foreground object of interest and some arbitrary background. Attempts have been made to construct VAEs for texture embedding. One example is TextureVAE, which consists of a VGG19 encoder and a decoder comprising 4 convolution blocks [14]. The network was tested on various microstructures and latent dimensions were artificially varied to visualise their effect on the resulting reconstructions. Further exploration of the encoded space has also been shown to carry the potential for material property prediction [15,16]. Conversely, the ability to generate synthetic microstructures from specified properties was demonstrated [17]. Dimensionality reduction of the encoded space has also been shown to provide meaningful visualisation of the space, enabling property prediction from microstructural image data [18] and identification of new processing routes for target orientation distributions [19].

Here, we employ a residual block VAE architecture, based on ResNet18 [20], to encode optical micrograph data from two classes of Ti-6Al-4V; a lamellar microstructure and a bimodal microstructure. The encoded space is explored via paths between learned fingerprints and Gaussian permutations of fingerprints, as a tool for generating artificial microstructures. The latent space is further explored through dimensionality reduction via t -stochastic neighbour embedding (t -SNE), with analysis of morphological metric distributions across the space. Support vector regression (SVR) is then applied to correlate fingerprints contained in the latent space with the same morphological metrics to quantify the trends visualised by the t -SNE.

Our titanium alloy dataset is representative of the microstructures of an important class of industrial materials in which control of microstructure is key to delivering the required in-service performance and in which quality assurance of material (and hence microstructure)

is a critical part of manufacturing processes. Furthermore, the nature of the dataset allows us to evaluate the performance of the VAE with respect to several aspects common to a broad range of metallurgical challenges:

1. the descriptors should transparently handle a range of microstructures (here we have widely varying grain morphology);
2. the space of descriptors (the latent space of the VAE) should be well-behaved, in the sense that the fingerprints should vary smoothly with changes to the microstructure and vice versa;
3. the representation in the latent space should be interpretable in terms of key features of the microstructure (or properties of the material), which is to say that the fingerprints should show statistical correlation with features and properties.

We explore these aspects of the VAE performance for our Ti-6Al-4V microstructure dataset.

2. Materials and methods

2.1. Dataset

The LightForm Ti-6Al-4V alloy bimodal/lamellar dataset (LFTi64BL) is an open access dataset, curated within the LightForm group at the University of Manchester, and can be accessed via Zenodo [21]. The dataset comprises 40 optical 8 bit micrographs of resolution 1292×968 pixels, containing equal numbers of bimodal equiaxed and lamellar microstructures. An example from each classification is shown in Fig. 1.

This dataset contains a relatively small number of images for machine learning purposes. To expand the dataset, patches were extracted with random rotations and reflections applied. The expanded LFTi64BL dataset contains 1000 patches from each image, resulting in 40,000 patches of resolution 256×256 pixels.

2.2. Greyscale normalisation

Each image was white balanced by clipping greyscale intensities outside the 90th percentile and remapping to the range [0, 1]. Fig. 2 shows distributions of mean greyscale intensities across the bimodal and lamellar datasets within LFTi64BL separately, before and after normalisation.

Prior to white-balancing, the variation in mean greyscale intensity is heavily influenced by fluctuations in lighting conditions during image capture. However, after white-balancing, greyscale intensity is more normally distributed across the dataset. This is now indicative of the distribution of phase fractions, with lower mean greyscale intensity corresponding to a higher volume fraction of the β phase.

2.3. Variational autoencoders (VAEs)

Variational autoencoders (VAEs) are a tool for encoding image data into a compressed format (or fingerprint) that preserves morphological features. They comprise a pair of convolutional neural networks (CNNs), referred to individually as the encoder and decoder. Fingerprints output by the trained encoder can be fed into downstream tasks such as image classification and property prediction. Fig. 3 provides a visual interpretation of the general VAE architecture.

Suppose we have an image, $x \in \mathbb{R}^{m_1 \times m_2}$. The encoder takes the image x as an input and computes a compressed representation $z \in \mathbb{R}^D$, such that $D \ll m_1 m_2$, where $z = e(x)$ and D is a tunable parameter that denoted the dimensionality of the encoded space. The decoder then takes the encoded vector z as an input and aims to reconstruct the input image, x . We denote the reconstruction as $\hat{x} = d(e(x))$, where $\hat{x} \in \mathbb{R}^{m_1 \times m_2}$. The encoded space learned is continuous and normally distributed.




Fig. 1. Representative optical micrographs of Ti-6Al-4V from LFTi64BL dataset.

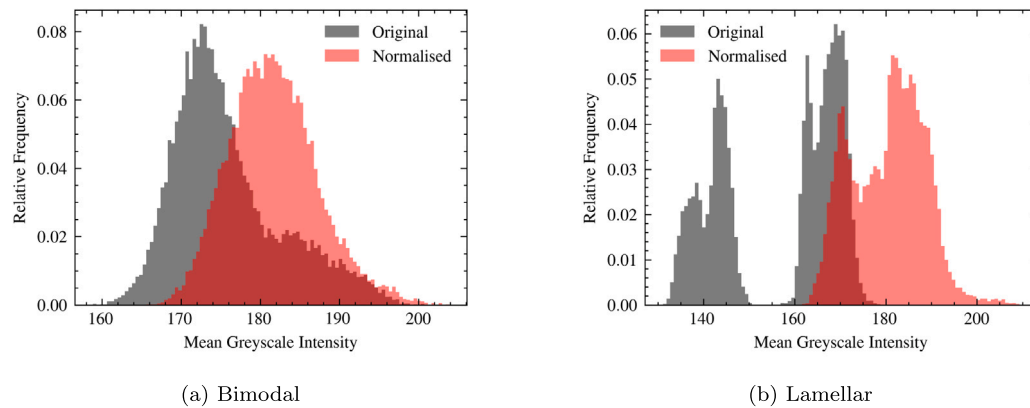


Fig. 2. Distribution of mean greyscale intensities across the LFTi64BL dataset before and after normalisation, for bimodal and lamellar microstructures separately.

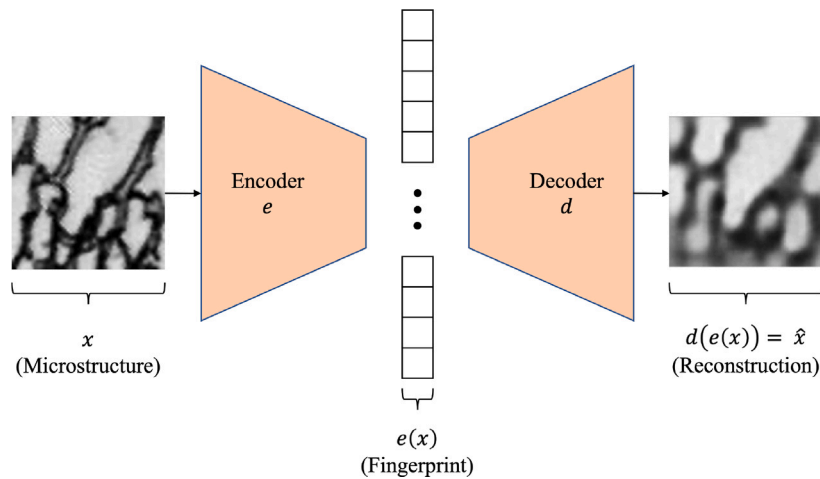


Fig. 3. Schematic of general VAE architecture that takes an image x as input, encodes the image into a fingerprint $e(x)$ and is trained to reconstruct the input image with the mapping $d(e(x))$.

To train a VAE, we must have a set of images, $X = \{x_1, x_2, \dots, x_N\}$, where N is the total number of images in the dataset. The dataset is split into two subsets, $X_{\text{train}} \in \mathbb{R}^{N_{\text{train}} \times m_1 \times m_2}$ and $X_{\text{eval}} \in \mathbb{R}^{N_{\text{eval}} \times m_1 \times m_2}$, where X_{train} is used to train the VAE, X_{eval} is used for evaluating the VAE on unseen data, $N_{\text{train}} + N_{\text{eval}} = N$, $X_{\text{train}} \cup X_{\text{eval}} \equiv X$ and $X_{\text{train}} \cap X_{\text{eval}} \equiv \emptyset$. The training set is split into batches and fed into the VAE, one batch at a time. Ultimately, $d(e(x_i))$ for $x_i \in X_{\text{train}}$ is computed and a loss function is used as a metric for assessing the encoded representation and reconstruction quality. The loss function is then used as an input to an optimiser. The Adam optimiser [22]

is currently the state-of-the-art and is utilised throughout all models discussed herein. The optimiser updates weights and biases in both the encoder and decoder from a single loss function after each batch operation. The model can then be evaluated on X_{eval} to measure the potential for transfer learning, but metrics calculated on these images in the evaluation set are not utilised for updating any weights or biases. Encoded representations, generated from the trained VAE on microstructural image data, can be considered as a signature, or microstructural fingerprint (as described in [3]), and will be referred to as such throughout this paper.

2.4. Loss functions

During VAE training, a loss function is periodically computed on the model, to determine performance and provide the inputs for the optimiser to update the weights and biases in the networks, with the aim of minimising the loss function. This is typically a combination between a measure in the spatial domain of the reconstruction accuracy and imposing a restriction on the encoded space to be normally distributed [12]. To measure how normally distributed the encoded space is, the Kullback–Leibler (KL) divergence, D_{KL} , can be determined between each encoded representation and the unit normal distribution $\mathcal{N}(0, 1)$ [23].

The reconstruction accuracy can be measured in several ways. A popular method is the mean squared error (MSE) [24], which is a distance metric in the spatial domain, given by

$$\text{MSE} = \frac{1}{m} \sum_{j=1}^m (x_j - \hat{x}_j)^2,$$

where $m = m_1 m_2$ is the total number of pixels in the image and the x_j, \hat{x}_j denote pixels in the input and reconstruction, respectively.

Another metric that operates in the spatial domain is the binary cross-entropy (BCE), denoted here as \mathcal{L}_b , which is a measure of negative log likelihood and is given by

$$\mathcal{L}_b = - \sum_{i=1}^m x_i \log \hat{x}_i - \sum_{i=1}^m (1 - x_i) \log(1 - \hat{x}_i).$$

The issue with minimising these loss functions is that they generally result in blurry reconstructions. With the aim of minimising blur, a loss function on the frequency domain was proposed in [25]. This requires calculating the 2D Fourier transform of both the input and reconstruction. The spectral loss, $\mathcal{L}_f(x, \hat{x})$, can then be given by the MSE between the 2D Fourier transforms, i.e.,

$$\mathcal{L}_f(x, \hat{x}) = \frac{1}{m} \sum_{j=1}^m ((\text{Im}\{\mathcal{F}(x)_j\} - \text{Im}\{\mathcal{F}(\hat{x})_j\})^2 + (\text{Re}\{\mathcal{F}(x)_j\} - \text{Re}\{\mathcal{F}(\hat{x})_j\})^2),$$

where \mathcal{F} denotes the 2D FFT, $\text{Im}\{\mathcal{F}\}$ denotes the imaginary part of \mathcal{F} and $\text{Re}\{\mathcal{F}\}$ the real part. The total loss function to be minimised by the optimiser is then given by

$$\mathcal{L} = \alpha \mathcal{L}_b(x, \hat{x}) + (1 - \alpha) \mathcal{L}_f(x, \hat{x}),$$

where $\alpha \in [0, 1]$ is a tunable hyperparameter.

2.5. ResNet18-VAE

ResNet [20] is a deep CNN composed of residual blocks and was initially proposed for classification of the ImageNet dataset [26], which is a dataset containing over 1 million natural images with 1000 classes. As such, the final layer of a standard ResNet is a 1000-dimensional fully connected layer, where the output from each node corresponds to a probability for each ImageNet class label. The depth of the network can be controlled by varying the number of layers within each block and the total number of blocks. Here, we consider ResNet18, which contains 8 residual blocks, each comprised of 2 convolution layers with subsequent ReLU activation and batch normalisation. Each block has the potential to be effectively skipped by summing the input received at each block with the output from that block, after convolution.

To convert ResNet18 into an encoder, it was modified by replacing the average pool and 1000-dimensional fully connected layers at the end of the network with a flattening of the final convolution output. This was followed by a 512-dimensional fully connected layer with tanh activation, which is then simultaneously fed into two separate fully connected layers. Each fully connected layer consists of $|z|$ neurons, where $|z|$ is the specified dimension of the encoded space. The dimensionality of the latent space was set to 256 to generate the results presented in

this paper. One of these fully connected layers is utilised as a set of means, μ , whilst the other is treated as a set of standard deviations, σ . These are then combined into the output z as

$$z = \mu + \exp(\sigma \mathcal{N}(0, 1)),$$

where $\mathcal{N}(0, 1)$ denotes the standard normal distribution, with mean 0 and standard deviation 1. This branching into μ and σ , with subsequent combination of the two, is what sets variational autoencoders apart from standard autoencoders. The decoder then essentially mirrors the encoder with transpose convolution layers replacing the standard convolution layers. The architecture for this ResNet18 VAE is provided in Fig. 4.

2.6. Morphological analysis

Morphological measurements were automated for the LFTi64BL dataset as metrics for correlation with fingerprints produced by the VAE. These metrics can then also be plotted as colour maps over a low dimensional space obtained from dimensionality reduction of the fingerprint space, to visualise how such features are distributed. Each metric requires the images to be binarised prior to measurement.

2.6.1. Image binarisation

Binarised images are required to compute certain metrics on the images, such as phase fraction and grain size. A high-pass Gaussian filter was applied to each image, in the Fourier domain, to normalise illumination across the images. The images were then binarised with Otsu's thresholding method [27], before applying area closing to remove noise from all images and α_s laths from the bimodal images. Pixels corresponding to α_p grains are assigned a label, whilst all other pixels are labelled as 0. Fig. 5 provides representative examples of the resulting binary images. The binarised lamellar microstructures are highly accurate, but there is some retention of α_s laths in the bimodal microstructures due to connectivity with the α_p grains.

2.6.2. Phase fraction

Due to the way in which the images were binarised (see Section 2.6.1), the phase fraction is simply determined as the sum of all the pixel values in the binary image divided by the total number of pixels in the image.

2.6.3. Lamellae direction

For the lamellar images in the LFTi64BL dataset, another descriptor we can consider is the dominant direction, or orientation, of the lamellae, relative to the bottom edge of the image. This can be quantified as the mean angle subtended between the elongation direction of each lamella and the bottom edge of the image. An erosion is applied on the binary image to isolate overlapping grains before each grain is labelled with a unique integer. Watershed segmentation is then applied with markers taken from the labelled image and the initial binary image as a mask. Each grain in the watershed image is isolated and eigenvectors are computed for each grain. The primary eigenvector describes the direction of each lamella. The angle between the primary eigenvector and the bottom edge of the image is then calculated for each grain and averaged to provide the metric for lamellae direction of a given image. Fig. 6 shows the erosion and watershed method applied to a representative micrograph and Fig. 7 shows the corresponding eigenvectors for an individual grain.

2.6.4. Bimodal grain size

As a metric for quantifying the bimodal microstructures in LFTi64BL, average grain size measurements were automated following ASTM E1382-97 [28]. Random line scans are applied on the binarised bimodal micrographs and peaks in the derivative of the profile along the line scans are used to detect grain edges. The distance in pixels between peaks in the derivative of the profile line is then converted into a distance in μm .

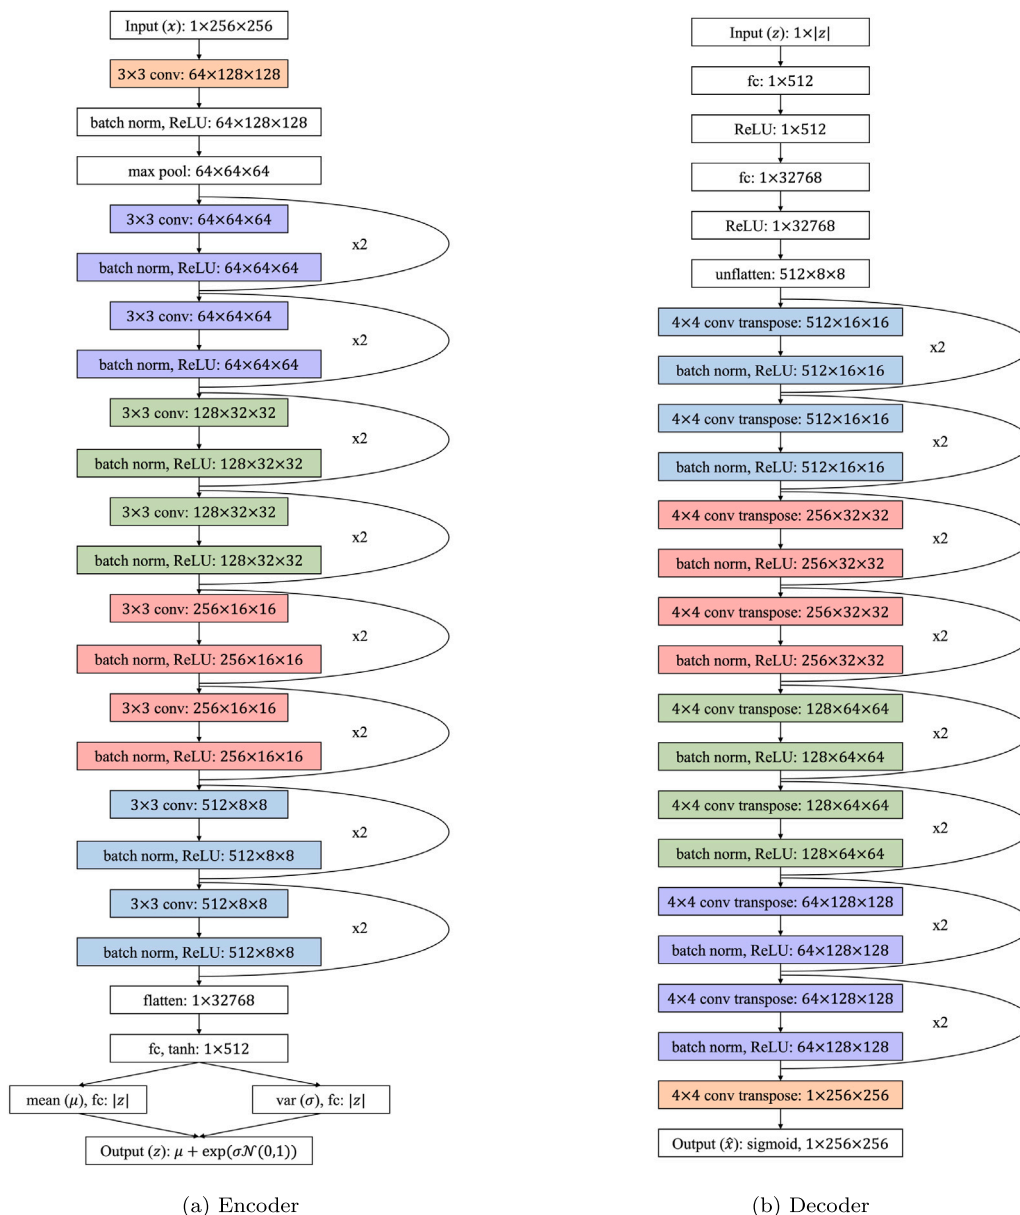


Fig. 4. ResNet18 VAE architecture for the 256 × 256 inputs used in the present study.




Fig. 5. Representative optical micrograph image patches for (a) lamellar and (c) bimodal and corresponding binarised images (b) and (d), respectively.

2.7. Support vector regression (SVR)

Once fingerprints have been extracted from the ResNet18 VAE, regression algorithms can be trained to predict morphological metrics. Here, we use support vector regression (SVR), which is an extension to support vector machines for continuous variables [29]. Gaussian

process regression is a popular alternative to SVR, but relies on relatively low-dimensional inputs compared to the dimensionality of the ResNet18 VAE fingerprints. Fingerprints are randomly split into a training set containing 90% of the fingerprints and a test set containing 10% for input into the SVR. This is repeated 10 times with random splits to perform 10-fold cross-validation.




Fig. 6. Watershed segmentation process applied to a representative micrograph for grain isolation. (a) shows the original micrograph, (b) is the binarised micrograph with area closing applied, (c) illustrates the erosion and labelling and, finally, (d) shows the watershed segmented micrograph, with the colour map corresponding to grain ID.




Fig. 7. Eigenvectors plotted on an individual lamella for direction measurement, with the primary eigenvector shown in red and the secondary eigenvector plotted in blue.

2.8. *t*-stochastic neighbour embedding (*t*-SNE)

The encoded space learned is high-dimensional (256 dimensions for ResNet18-VAE architecture, outlined in Section 2.5). This makes it difficult to visualise the encoded space in its entirety. To aid in visualising this high-dimensional space, we can perform dimensionality reduction down to 2 or 3 dimensions, which allows us to plot the encoded space and to assess how the microstructures in our training set are distributed. Here, we consider *t*-stochastic neighbour embedding (*t*-SNE) [30].

The fingerprints are represented as a similarity matrix, S , where entries, $S_{i,j}$, denote the probability that z_j is a nearest neighbour of z_i . The aim is then to minimise the distance between $S_{i,j}$ and $S_{j,i}$ [31]. Standard SNE utilises a Gaussian distribution to determine similarity between fingerprints, whereas *t*-SNE relies on the Student's *t* distribution, hence its name. Due to the longer tail of the *t* distribution, relative to a Gaussian distribution, the use of the *t* distribution results in more separated embeddings and helps alleviate the crowding issue often encountered with SNE [32]. Principal component analysis (PCA) is used to initialise the *t*-SNE.

2.9. Traversing the encoded space

One way in which the encoded space can be traversed is through construction of a specific path. Here, we consider a linear path. Two images, x_1 and x_2 , are randomly selected and $e(x_1)$, $e(x_2)$ are computed. A linear path is then constructed from $e(x_1)$ to $e(x_2)$ in the encoded space, according to the following equation.

$$e_n(x_1, x_2) = e(x_1) + \frac{n+1}{N} (e(x_2)), \quad n = 1, \dots, N,$$

where $e_n(x_1, x_2)$ denotes each fingerprint along the path and N is the number of steps along the path. Fingerprints along the path are then supplied to the decoder to output reconstructions.

Another method for generating potentially valid encoded representations of microstructure is to form a random cloud, centred at

a known valid encoded representation, $e(x)$, for some image x . This can be achieved by sampling random Gaussian noise from $\mathcal{N}(0, 1)$ and summing with $e(x)$, i.e.,

$$e_n(x) = e(x) + \gamma \mathcal{N}(\mu, \sigma), \quad n = 1, \dots, N,$$

where $e_n(x)$ is a random neighbour of $e(x)$, γ is a tunable scale factor that controls the noise level, $\mu, \sigma \in \mathbb{R}^D$ denote the element-wise mean and standard deviation for each dimension in the encoded space and N is the number of neighbouring fingerprints to be output.

3. Results

3.1. Reconstruction accuracy

The ResNet18 VAE was trained under two separate regimes. The first included exclusively either lamellar or bimodal microstructures and the second included the full LFTi64BL dataset, containing both lamellar and bimodal microstructures. Fig. 8 shows some example reconstructions from the lamellar dataset after 1000 epochs. When training is restricted to the lamellar dataset, grain boundaries are accurately identified and reconstructed, but there is some smoothing evident in the reconstructions and the interlamellar β appears coarsened. This becomes clearer upon inspection of the morphological metric distributions shown in Fig. 9. The α_p volume fraction is slightly reduced in the reconstructions, with some anomalies between 30% and 50%. The mean lamellae direction, relative to the bottom edge of each image patch, is consistent. This confirms that the α_p grains are oriented correctly in the reconstructions, but there is an increase in the mean lamellae width, which is likely due to the smoothing effect removing smaller grains. There is also an increase in the mean lamellae aspect ratio, owed to the coarsening of the interlamellar β in the reconstructions which results in α_p grains appearing more elongated.

Training the VAE on exclusively bimodal microstructures results in a similar smoothing effect as with the lamellar images, although this appears more pronounced in the bimodal case due to the nature of the fine scale features present in this case. Fig. 10 shows some example reconstructions for the bimodal micrographs. All α_s laths retained in the prior β grains are completely absent from the reconstructions and the outputs are effectively a mask for the α_p grains. This is confirmed to be a feature of the entire set of reconstructions in Fig. 11, which shows a drastic increase in both the α_p volume fraction and mean α_p grain size.

3.2. Traversing the encoded space

The encoded space was explored with the methods discussed in Section 2.9. Fig. 12 shows reconstructions along a linear path between a pair of fingerprints from the training set. A linear path between the fingerprints was constructed and 8 equispaced fingerprints were determined along the path. These fingerprints were supplied to the trained decoder to generate the images in Fig. 12.




Fig. 8. Representative examples of (a) 256×256 patches from the lamellar microstructures and (b) their corresponding reconstructions from the ResNet18 VAE architecture, after 1000 epochs.

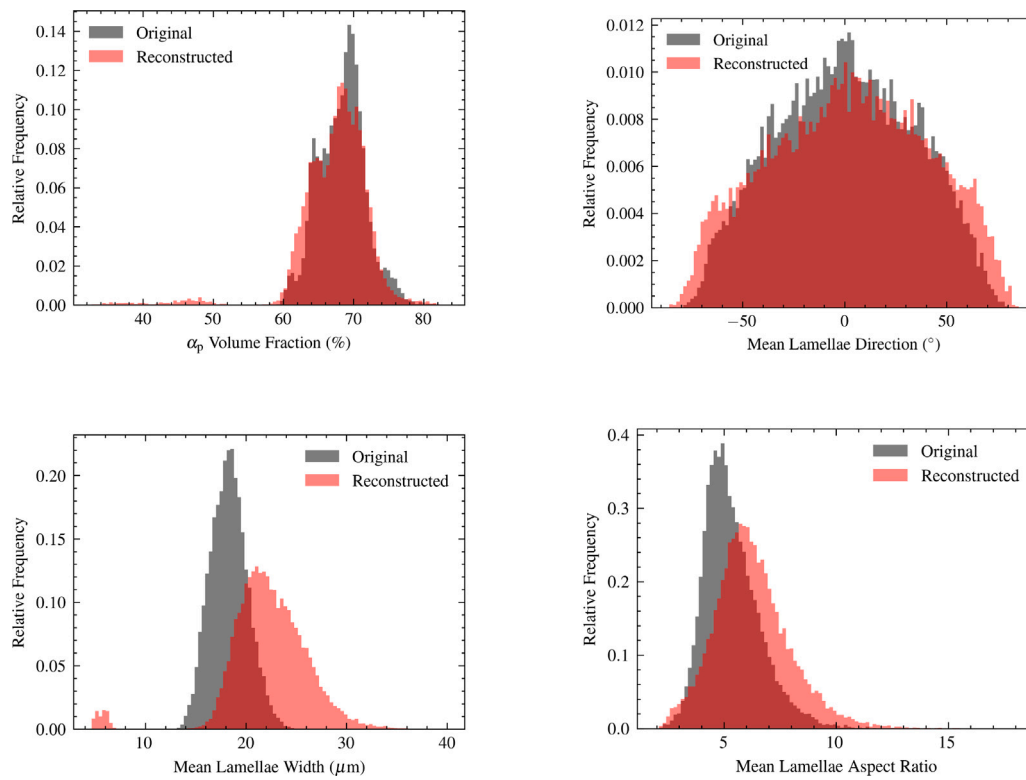


Fig. 9. Morphological metric distributions across the original image patches and their corresponding reconstructions for the lamellar dataset.




Fig. 10. Representative examples of (a) 256×256 patches from the bimodal microstructures and (b) their corresponding reconstructions from the ResNet18 VAE architecture, after 1000 epochs.




Fig. 11. Morphological metric distributions across the original image patches and their corresponding reconstructions for the bimodal dataset.

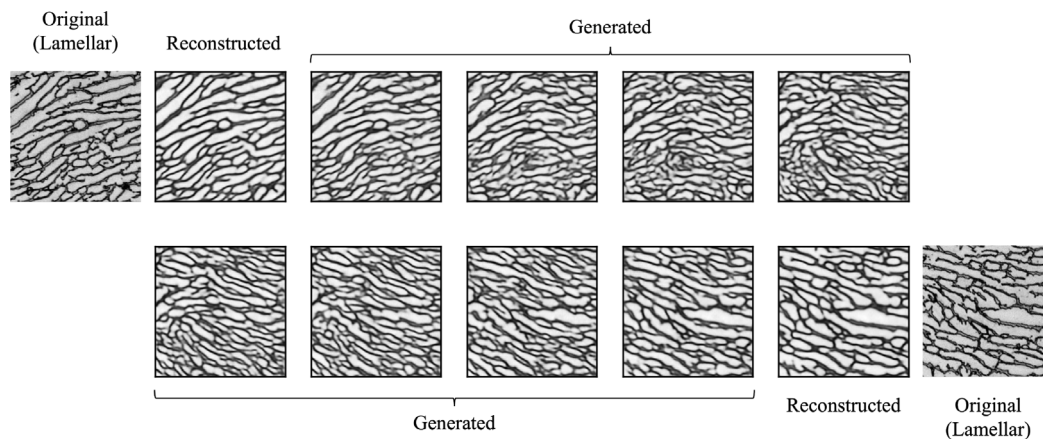


Fig. 12. Reconstructions along linear path through the encoded space learned during training on exclusively lamellar image data from LFTi64BL.

The fingerprints along this path show a smooth transition along the linear path, which is a direct result of the continuity of the latent space learned by the VAE. These microstructures are synthetic and are not included in the training set, although they do have the same characteristics as the lamellar microstructures in the training set.

Fig. 13 shows the output from the same method applied to the full LFTi64BL dataset, with the linear path defined between a lamellar and a bimodal microstructure. The same smooth transition between the input microstructures is observed, however, intermediate microstructures along the path stray considerably away from the training set, particularly towards the centre of the path. This is to be expected, but confirms that, for VAEs to generate convincing artificial micrographs, it is crucial that the training set be cohesive and not contain drastic variations in microstructure. Otherwise, the latent space constructed is likely to contain microstructures that are not representative of the training set.

To visualise neighbouring microstructures localised around an individual fingerprint, a sample microstructure was randomly selected and supplied to the trained encoder. The fingerprint obtained was perturbed with a small amount of Gaussian noise, as described in Section 2.9 with $\gamma = 0.2$, and provided to the trained decoder to generate a synthetic microstructure. This process was repeated and Fig. 14 shows 10 realisations of the microstructures produced. Each artificial microstructure possesses similar features to the input micrographs in terms of grain morphology and direction of lamellae with respect to the bottom edge of the image. This shows that local fingerprints within the latent space are likely to possess similar microstructural features.

Fig. 15 shows example micrographs constructed with increasing values of γ . As γ increases, images generated become no longer representative of the training dataset. If γ is too small, then there is very little variation in the output microstructures. Fig. 16 shows mean fields across 100 realisations of synthetic microstructures produced

with the same range of γ values as in Fig. 15. These were constructed by averaging the greyscales for each pixel across the set of realisations for each γ value. This confirms that for $\gamma < 0.2$, there is basically no change to the reconstructed microstructure after perturbing the input fingerprint. At $\gamma = 0.2$, there is bias towards retaining some of the larger grains, although the grain boundary does fluctuate slightly and surrounding grains can be subject to change. For larger γ , there is no bias introduced by the input microstructure, but the resulting synthetic microstructures are not representative. Despite synthetic microstructures being discussed here, this method is recommended as a tool for exploring the latent space and visualising the sorts of features that are shared between microstructures within some neighbourhood in the latent space, rather than aiming to produce statistically equivalent but distinct realisations of an input microstructure.

PCA was trained on the fingerprints constructed by the VAE and then used to transform the noise-perturbed fingerprints alongside the original fingerprints (see Fig. 17). This shows the noise-perturbed fingerprints emanating from a single point, which corresponds to the fingerprint to which the noise was applied, and we can see that noise added with a scale factor of $\gamma > 0.5$ are all completely outside the encoded space learned by the VAE.

3.3. Metric distributions across the encoded space

Dimensionality reduction, in the form of *t*-SNE (described in Section 2.8), was also applied to the original 256-dimensional fingerprints to reduce them down to 2-dimensional vectors. This enables fingerprints to be plotted in a 2-dimensional scatter plot to visualise the entire latent space and distribution of metrics across the space. Fig. 18 shows the full LFTi64BL dataset reduced to 2-dimensions. The colour map in this figure that illustrates the microstructure classification, with 0 denoting a bimodal microstructure and 1 denoting lamellar. There is a

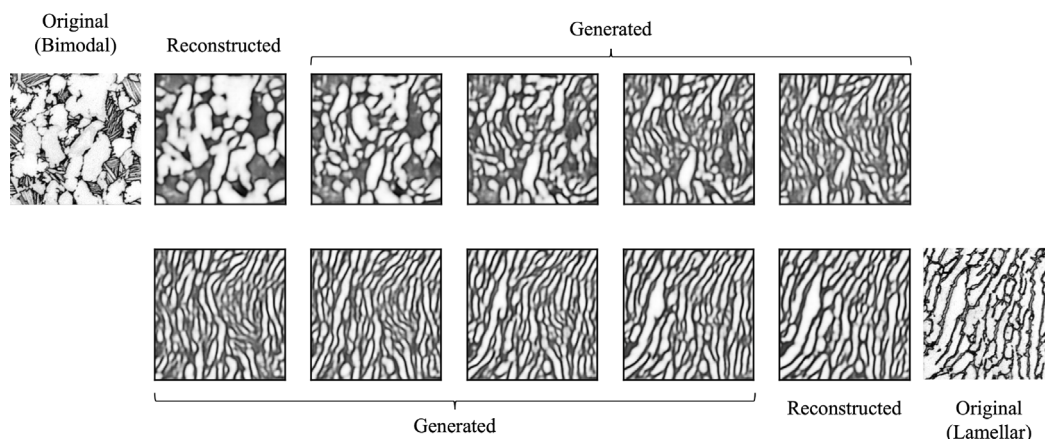


Fig. 13. Reconstructions along linear path through the encoded space learned during training on the full LFTi64BL dataset.

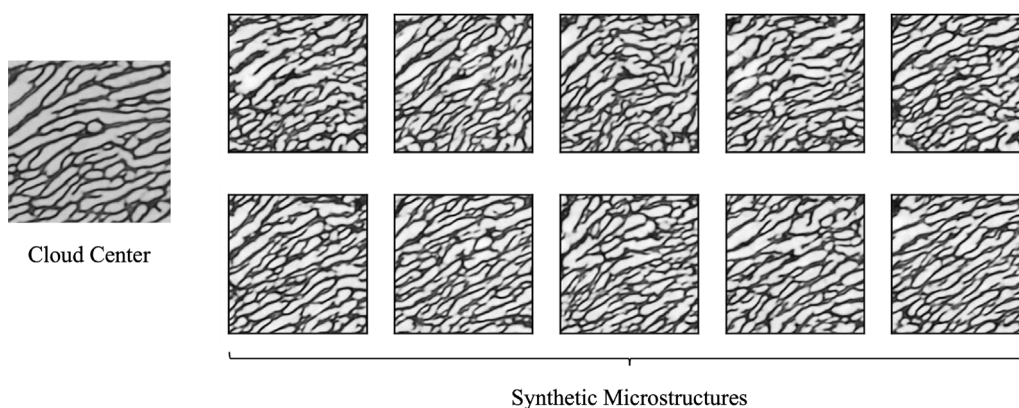


Fig. 14. Synthetic microstructures generated from the addition of unit Gaussian noise, with a scale factor of $\gamma = 0.2$, to a known valid encoded representation, from which morphology is inherited.




Fig. 15. Representative examples of synthetic microstructures generated via the Gaussian cloud method with various γ values, compared with original micrograph.




Fig. 16. Mean fields across 100 realisations of microstructures generated via the Gaussian cloud method with various γ values.

strong clustering of each class, with only a small overlap between them, allowing fingerprints from each class to be easily separated. Training an SVM with 90% of the fingerprints allocated for training and 10% reserved for testing yields a mean classification accuracy of $99.9\% \pm 0.001$ after 10-fold cross-validation.

The same methodology for dimensionality reduction was applied to the lamellar and bimodal fingerprints separately to map distributions of various microstructural features. Figs. 19(a) and 19(b) shows

greyscale intensity from the lamellar micrographs, before and after normalisation.

Prior to normalisation, there is a strong clustering between two groups of images, heavily influenced by illumination during image capture. After normalisation, greyscale intensity is more closely linked to volume fraction and is normally distributed across the dataset. Dimensionality reduction then shows a smooth gradient of greyscale intensity across the latent space. The plots with volume fraction yield




Fig. 17. PCA applied to fingerprints learned by the VAE alongside fingerprints constructed with the Gaussian cloud method. Original fingerprints are shown in grey and the colour map denotes the γ values applied to generate the artificial fingerprints.

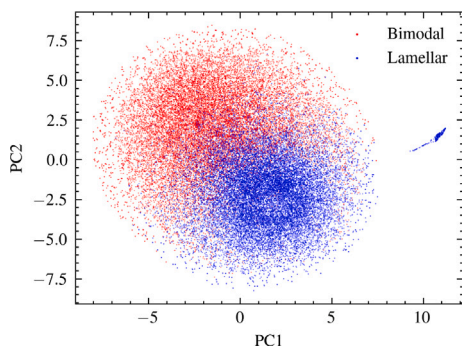


Fig. 18. t -SNE with 2 components applied to encoded representations, where the colour map denotes classifications, with bimodal microstructures in red and lamellar microstructures in blue.

similar results, although less pronounced. Figs. 19(c) and 19(d) shows plots of the encoded spaces for bimodal and lamellar microstructures, trained separately.

Finally, we look at directionality (discussed in Section 2.6.3) for lamellar microstructures and grain size (Section 2.6.4) for bimodal microstructures as morphological features of interest. Fig. 19(e) shows the t -SNE plot with a colour map corresponding to lamellae direction.

This appears to show random scatter across the encoded space, in contrast to the Gaussian cloud for image generation that seems to reconstruct images with similar direction, when perturbing fingerprints with a small amount of Gaussian noise. Nearest neighbours are shown to have similar directionality, but this is contained within small regions of the latent space and not universal. This random distribution of directionality may be useful in practice, though, as this implies that when encoding the microstructural information in such a manner, directionality can be effectively ignored and sample orientation when imaging would not be of any concern.

Fig. 19(f) shows the t -SNE plot for the bimodal fingerprints with a colour map denoting grain size. Here, we see a gradient of grain size across the latent space, suggesting that nearest neighbours will share similar morphologies.

In each t -SNE reduction, there is a small group of fingerprints that are separated away from the main cluster. There are no discernible differences between the microstructures that result in these fingerprints and those that form the main cluster. Despite these microstructures being included in the training set, their reconstructions after training are effectively just noise. This is a caveat of the VAE training process and not the dimensionality reduction. Removing these outliers from the set of fingerprints and rerunning the t -SNE with the remaining fingerprints results in retention of the main cluster and a complete absence of the secondary cluster. Removing the images that reside in

Table 1
Quantitative analysis of error for SVR predictions.

Class	Metric	Mean percentage error
Lamellar	Greyscale intensity	0.60% \pm 0.01
	Volume fraction α_p (%)	1.83% \pm 0.03
	Directionality ($^\circ$)	148% \pm 29
Bimodal	Greyscale Intensity	0.75% \pm 0.01
	Volume Fraction α_p (%)	2.50% \pm 0.03
	Grain Size (μm)	7.91% \pm 0.12

the secondary cluster from the training set and retraining the VAE from scratch results in a new cluster forming, with a different morphology, that contains a different set of images. Nonetheless, we are able to explore the encoded space of valid reconstructions, using the t -SNE as a guide to identify suitably encoded microstructures.

3.4. Metric predictions from encoded representations

If an approach to microstructural fingerprinting is successfully capturing the “essence” of the microstructure, then we might reasonably expect to be able to recover key features of the microstructure, such as grain size, from the encoded fingerprint. A strong correlation between the fingerprint and a given property also opens the possibility of reversing the inferential process and asking what fingerprints (or which regions of latent space) would correspond to a microstructure exhibiting a given property of interest (with the possibility of a trained VAE then reconstructing an example of the corresponding microstructure).

The predictability of morphological metrics across the encoded space was validated with SVR, described in Section 2.7. The fingerprints output from the ResNet18 VAE are randomly split into 90% training data and 10% test data. This is repeated 10 times to perform 10-fold cross-validation. The SVR is then trained on the training set and outputs for the test set are compared with the true values to measure the prediction accuracy. Fig. 20 shows scatter plots for the SVR predictions across the 10 train–test splits combined, against the true values, with the line $y = x$ plotted in red to illustrate deviation from exact predictions. Table 1 then shows percentage error and standard deviation of the predictions relative to the true values.

4. Discussion

Reconstructions from the proposed ResNet18 VAE were shown to accurately identify grain boundaries, albeit with some smoothing. This information loss is acceptable for correlating with the morphological features discussed, as these features are still captured. However, prediction of properties that rely upon small scale features, such as fine α_s laths, may become hindered. This could be alleviated with higher resolution images and patches with a higher magnification.

The encoded space learned by the VAE is continuous and contains valid microstructural fingerprints that are not included in the training set. Two methods were used to explore the encoded space as a tool for generating artificial microstructures. The linear path method illustrates the continuity of the space and shows a smooth transition between microstructures, when linearly interpolating between two known fingerprints. With a training set containing only lamellar microstructures, reconstructions along this path are completely artificial but still resemble lamellar microstructures. Once the training set is expanded to include both bimodal and lamellar microstructures, the linear interpolation between a bimodal fingerprint and a lamellar fingerprint yields a blend between the two morphologies. These can still be perceived as valid microstructures, but it is important to note that these are not representative of any data in the training set and VAEs trained in such a way should be interpolated carefully.

Gaussian noise was also used to explore the encoded space more locally around an individual fingerprint. Iteratively perturbing a learned

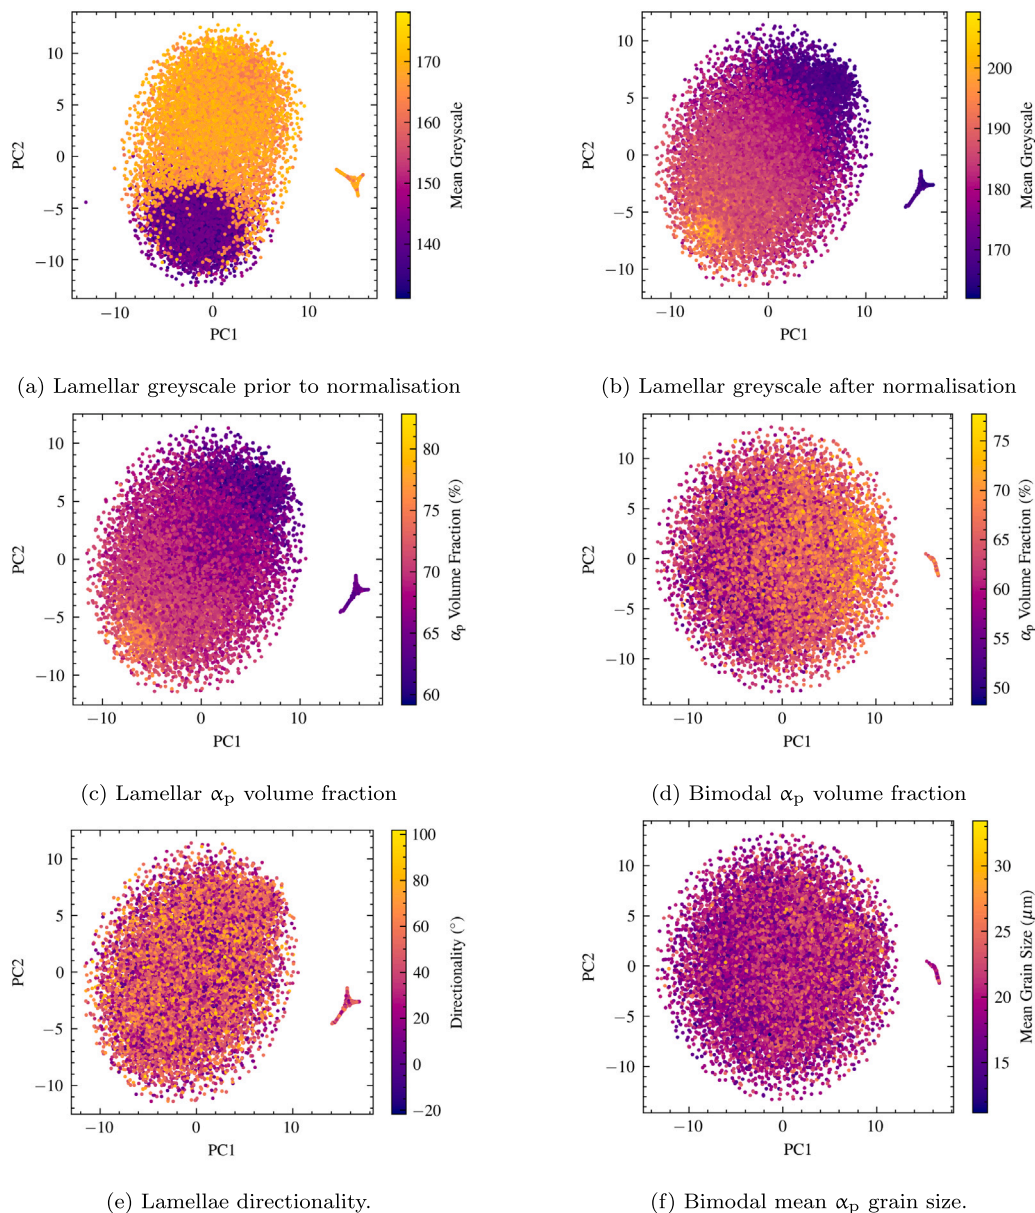


Fig. 19. t -SNE with 2 components applied to encoded representations of the LFTi64BL dataset, with colour maps denoting various morphological metrics.

fingerprint and supplying the output to the trained decoder yields another set of artificially generated microstructures. In this case, the decoder generates microstructures with similar features to the input image, including the lamellae thickness and direction relative to the bottom of the image, provided that a suitably small amount of noise is applied. This behaviour is only observed locally for some features, such as lamellae direction, which is randomly distributed across the encoded space. This may be due to the fact that the VAE was trained on patches with random rotations applied, making the VAE rotationally invariant in this case. This removes any bias towards sample orientation during image capture. Adding large amounts of noise would result in variations of such features.

Dimensionality reduction via t -SNE was performed to reduce the fingerprints to 2 dimensions. This enables them to be plotted in a 2-dimensional scatter plot to visualise the distribution of morphological features across the encoded space. The fingerprint position in the

encoded space appears to depend significantly on the volume fraction of α_p , relative to other morphological metrics. This was confirmed by SVR, which was trained to predict volume fraction with an average percentage error of $1.83\% \pm 0.03$ for the lamellar dataset and $2.50\% \pm 0.03$ for the bimodal dataset. Grain size was also highly correlated with the encoded representations and SVR was able to predict grain size with reasonable accuracy, giving an average percentage error of $7.91\% \pm 0.12$. Directionality of lamellae appeared to be distributed randomly across the encoded space from the t -SNE plots and this was also confirmed with SVR, which was unable to learn any trends in the data and repeatedly predicted the mean direction for each fingerprint, resulting in an average percentage error of $148\% \pm 29$. In practice, the random nature of the distribution of lamellae direction could be a useful feature, as this implies that sample orientation under the microscope can be effectively ignored when capturing a dataset for VAE without any bias being introduced.

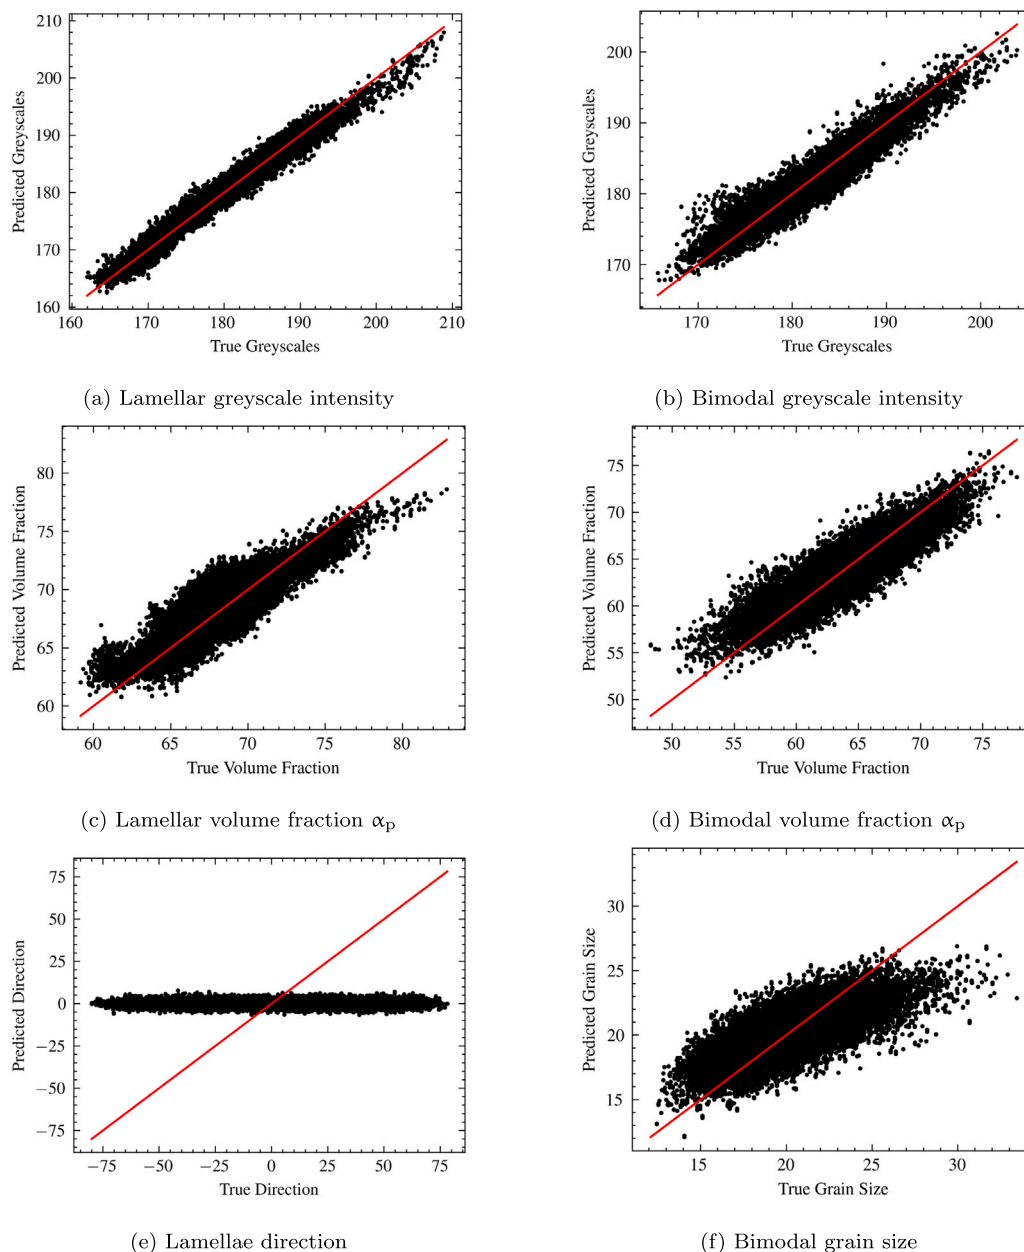


Fig. 20. Scatter plots to illustrate combined 10-fold cross validation SVR predictions of various morphological metrics with the line $y = x$ added to highlight where true predictions should lie.

5. Conclusions

The latent space of microstructural fingerprints generated from variational autoencoders (VAEs) has been explored to further our understanding of how VAEs encode feature information from microstructural image data.

- We show that a residual block VAE architecture based on ResNet18 is able to produce accurate reconstructions of microstructures (up to some smoothing).
- Interpolation of fingerprints along a linear path in latent space and random perturbations about fingerprints of input microstructures resulted in the generation of plausible synthetic microstructures, demonstrating the suitability of the VAE for smooth representation of microstructure.

- Fingerprints constructed by the trained VAE were shown to correlate with morphological features, including α_p volume fraction and grain size, using support vector regression (SVR) with 10-fold cross-validation.
- Principal component analysis (PCA) was shown to provide useful insight into the amount of noise that can be added before the synthetic fingerprints become disconnected from the learned latent space.

Our study based on a set of micrographs of titanium alloy microstructure demonstrates that a VAE can encode material microstructures to produce fingerprints that exhibit several of the key features required in a general purpose fingerprint. Important next steps will be to test this approach on a broader class of microstructures and to apply it to datasets exhibiting variation in both microstructure and measured properties, to allow exploration of the suitability of VAEs

for predicting microstructure–process–property relationships in an explainable framework. We also note that generative adversarial networks (GANs) are found to generate high-fidelity synthetic microstructures, with strong statistical similarity to the training data, even in a limited data regime [33], with possible extensions to reconstructing 3D volumes [34]. It would be interesting to consider synthetic images output from a GAN as inputs to the VAE to determine how the generated images are distributed through the encoded space and assess their potential for bolstering morphology prediction.

CRedit authorship contribution statement

Michael D. White: Writing – review & editing, Writing – original draft, Validation, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Gowtham Nimmal Haribabu:** Writing – review & editing, Validation, Formal analysis. **Jeyapriya Thimukonda Jegadeesan:** Writing – review & editing, Validation, Formal analysis. **Bikramjit Basu:** Writing – review & editing, Supervision. **Philip J. Withers:** Writing – review & editing, Supervision, Conceptualization. **Chris P. Race:** Writing – review & editing, Supervision, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Michael D White reports financial support was provided by Engineering and Physical Sciences Research Council. Michael D White reports financial support was provided by Scheme for Promotion of Academic and Research Collaboration. Michael D White reports financial support was provided by Science Foundation Ireland. Chris P Race reports financial support was provided by The Royal Society. Philip Withers reports financial support was provided by Engineering and Physical Sciences Research Council. Michael D White reports a relationship with Rolls-Royce plc that includes: funding grants. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

All data used is available to download through Zenodo via the following link: <https://zenodo.org/records/5714384>.

Acknowledgements

This research was supported through funding from the Scheme for Promotion of Academic and Research Collaboration (SPARC) grant MHRD-18-0015. MDW was supported by the University of Manchester and Rolls-Royce plc under the Engineering and Physical Sciences Research Council (EPSRC), United Kingdom grant EP/S022635/1 and Science Foundation Ireland (SFI) grant 18/EPSC-CDT/3584. CPR was supported by a University Research Fellowship of the Royal Society. PJW was supported by the Henry Royce Institute for Advanced Materials, funded through EPSRC grants EP/R00661X/1, EP/S019367/1, EP/P025021/1 and EP/P025498/1.

References

- [1] B. Basu, N. Gowtham, Y. Xiao, S.R. Kalidindi, K.W. Leong, Biomaterialomics: Data science-driven pathways to develop fourth-generation biomaterials, *Acta Biomater.* 143 (2022) 1–25, <http://dx.doi.org/10.1016/j.actbio.2022.02.027>.
- [2] E.A. Holm, R. Cohn, N. Gao, A.R. Kitahara, T.P. Matson, B. Lei, S.R. Yarasi, Overview: Computer vision and machine learning for microstructural characterization and analysis, *Metall. Mater. Trans. A* 51 (2020) 5985–5999, <http://dx.doi.org/10.1007/s11661-020-06008-4>.
- [3] T.L. Burnett, P.J. Withers, Completing the picture through correlative characterization, *Nature Mater.* 18 (2019) 1041–1049, <http://dx.doi.org/10.1038/s41563-019-0402-8>.
- [4] B.L. DeCost, E.A. Holm, A computer vision approach for automated analysis and classification of microstructural image data, *Comput. Mater. Sci.* 110 (2015) 126–133, <http://dx.doi.org/10.1016/j.commatsci.2015.08.011>.
- [5] M.I. Latypov, S.R. Kalidindi, Data-driven reduced order models for effective yield strength and partitioning of strain in multiphase materials, *J. Comput. Phys.* 346 (2017) 242–261, <http://dx.doi.org/10.1016/j.jcp.2017.06.013>.
- [6] Y.C. Yabansu, P. Steinmetz, J. Hötzer, S.R. Kalidindi, B. Nestler, Extraction of reduced-order process-structure linkages from phase-field simulations, *Acta Mater.* 124 (2017) 182–194, <http://dx.doi.org/10.1016/j.actamat.2016.10.071>.
- [7] E. Popova, T.M. Rodgers, X. Gong, A. Cecen, J.D. Madison, S.R. Kalidindi, Process-structure linkages using a data science approach: Application to simulated additive manufacturing data, *Integr. Mater. Manuf. Innov.* 6 (2017) 54–68, <http://dx.doi.org/10.1007/s40192-017-0088-1>.
- [8] K. Rajan, Materials informatics: The materials gene and big data, *Annu. Rev. Mater. Res.* 45 (2015) 153–169, <http://dx.doi.org/10.1146/annurev-matsci-070214-021132>.
- [9] D.L. McDowell, R.A. LeSar, The need for microstructure informatics in process-structure–property relations, *MRS Bull.* 41 (2016) 587–593, <http://dx.doi.org/10.1557/mrs.2016.163>.
- [10] M. White, A. Tarakanov, P. Withers, C. Race, K. Law, Digital fingerprinting of microstructures, *Comput. Mater. Sci.* 218 (2023) 111985, <http://dx.doi.org/10.1016/j.commatsci.2022.111985>.
- [11] D.T. Fullwood, S.R. Niezgoda, B.L. Adams, S.R. Kalidindi, Microstructure sensitive design for performance optimization, *Prog. Mater. Sci.* 55 (2010) 477–562, <http://dx.doi.org/10.1016/j.pmatsci.2009.08.002>.
- [12] D.P. Kingma, M. Welling, Auto-encoding variational Bayes, 2013, <http://dx.doi.org/10.48550/arXiv.1312.6114>.
- [13] Z. Liu, P. Luo, X. Wang, X. Tang, Deep learning face attributes in the wild, 2014, <http://dx.doi.org/10.48550/arXiv.1411.7766>.
- [14] A. Sardeshmukh, S. Reddy, B.P. Gautham, P. Bhattacharyya, TextureVAE: Learning interpretable representations of material microstructures using variational autoencoders, 2021.
- [15] R. Cang, H. Li, H. Yao, Y. Jiao, Y. Ren, Improving direct physical properties prediction of heterogeneous materials from imaging data via convolutional neural network and a morphology-aware generative model, *Comput. Mater. Sci.* 150 (2018) 212–221, <http://dx.doi.org/10.1016/j.commatsci.2018.03.074>.
- [16] Y. Kim, H.K. Park, J. Jung, P. Asghari-Rad, S. Lee, J.Y. Kim, H.G. Jung, H.S. Kim, Exploration of optimal microstructure and mechanical properties in continuous microstructure space using a variational autoencoder, *Mater. Des.* 202 (2021) 109544, <http://dx.doi.org/10.1016/j.matdes.2021.109544>.
- [17] H.S. Stein, D. Guevarra, P.F. Newhouse, E. Soedarmadji, J.M. Gregoire, Machine learning of optical properties of materials – predicting spectra from images and images from spectra, *Chem. Sci.* 10 (2019) 47–55, <http://dx.doi.org/10.1039/C8SC03077D>.
- [18] Y. Pathak, K.S. Juneja, G. Varma, M. Ehara, U.D. Priyakumar, Deep learning enabled inorganic material generator, *Phys. Chem. Chem. Phys.* 22 (2020) 26935–26943, <http://dx.doi.org/10.1039/D0CP03508D>.
- [19] S. Sundar, V. Sundararaghavan, Database development and exploration of microstructure versus process relationships using variational autoencoders, 2020, <http://dx.doi.org/10.48550/arXiv.2001.09171>.
- [20] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, 2015, <http://dx.doi.org/10.48550/arXiv.1512.03385>.
- [21] M. White, C. Daniel, J.Q. da Fonseca, X. Zeng, N. Byers, B. Karnasiewicz, Lamellar and Bi-Modal Ti-64 microstructure images, 2021, <http://dx.doi.org/10.5281/zenodo.5714384>, URL <https://zenodo.org/records/5714384>.
- [22] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, 2014, <http://dx.doi.org/10.48550/arXiv.1412.6980>.
- [23] D.P. Kingma, M. Welling, An introduction to variational autoencoders, *Found. Trends Mach. Learn.* 12 (2019) 307–392, <http://dx.doi.org/10.1561/2200000056>.
- [24] R. Yu, A tutorial on VAEs: From Bayes’ rule to lossless compression, 2020, <http://dx.doi.org/10.48550/arXiv.2006.10273>.
- [25] S. Björk, J.N. Myhre, T.H. Johansen, Simpler is better: spectral regularization and up-sampling techniques for variational autoencoders, 2022, <http://dx.doi.org/10.48550/arXiv.2201.07544>.
- [26] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, ImageNet: A large-scale hierarchical image database, *IEEE* (2009) 248–255, <http://dx.doi.org/10.1109/CVPR.2009.5206848>.
- [27] S.L. Bangare, A. Dubal, P.S. Bangare, S. Patil, Reviewing Otsu’s method for image thresholding, *Int. J. Appl. Eng. Res.* 10 (2015) 21777–21783, <http://dx.doi.org/10.37622/IJAER/10.9.2015.21777-21783>.

- [28] E. Committee, Test methods for determining average grain size using semiautomatic and automatic image analysis, 2015, <http://dx.doi.org/10.1520/E1382-97R15>.
- [29] M. Awad, R. Khanna, Support vector regression, Apress (2015) 67–80, http://dx.doi.org/10.1007/978-1-4302-5990-9_4.
- [30] R. Silva, P. Melo-Pinto, T-SNE: A study on reducing the dimensionality of hyperspectral data for the regression problem of estimating oenological parameters, *Artif. Intell. Agric.* 7 (2023) 58–68, <http://dx.doi.org/10.1016/j.aiia.2023.02.003>.
- [31] G.E. Hinton, S. Roweis, *Stochastic Neighbor Embedding*, MIT Press, 2002.
- [32] L. van der Maaten, G. Hinton, Visualizing data using t-SNE, *J. Mach. Learn. Res.* 9 (2008) 2579–2605.
- [33] N.H. Gowtham, T.J. Jeyapriya, C. Bhattacharya, B. Basu, A deep adversarial approach for the generation of synthetic titanium alloy microstructures with limited training data, *Comput. Mater. Sci.* 230 (2023) 112512, <http://dx.doi.org/10.1016/j.commatsci.2023.112512>.
- [34] R. Bostanabad, Reconstruction of 3D microstructures from 2D images via transfer learning, *Comput. Aided Des.* 128 (2020) 102906, <http://dx.doi.org/10.1016/j.cad.2020.102906>.