

ARTICLE OPEN



Adaptable physics-based super-resolution for electron backscatter diffraction maps

Devendra K. Jangid^{1,5}, Neal R. Brodnik^{2,5}✉, Michael G. Goebel¹, Amil Khan¹, SaiSidharth Majeti³, McLean P. Echlin⁴, Samantha H. Daly², Tresa M. Pollock⁴ and B. S. Manjunath¹

In computer vision, single-image super-resolution (SISR) has been extensively explored using convolutional neural networks (CNNs) on optical images, but images outside this domain, such as those from scientific experiments, are not well investigated. Experimental data is often gathered using non-optical methods, which alters the metrics for image quality. One such example is electron backscatter diffraction (EBSD), a materials characterization technique that maps crystal arrangement in solid materials, which provides insight into processing, structure, and property relationships. We present a broadly adaptable approach for applying state-of-art SISR networks to generate super-resolved EBSD orientation maps. This approach includes quaternion-based orientation recognition, loss functions that consider rotational effects and crystallographic symmetry, and an inference pipeline to convert network output into established visualization formats for EBSD maps. The ability to generate physically accurate, high-resolution EBSD maps with super-resolution enables high-throughput characterization and broadens the capture capabilities for three-dimensional experimental EBSD datasets.

npj Computational Materials (2022)8:255; <https://doi.org/10.1038/s41524-022-00924-2>

INTRODUCTION

The term image super-resolution is used to describe methods designed to infer high-resolution (HR) image output from low-resolution (LR) input. Since their development, super-resolution methods have been used in applications such as surveillance and security, biometric information identification, remote sensing, astronomy, and medical imaging¹. Generally, these algorithms can be categorized into three groups based on the information available during training: (a) supervised, which have paired LR-HR images during training, (b) semi-supervised, where no LR-HR image pairing is available, and (c) unsupervised, where no ground truth HR is available. Recently, because of their superior performance to traditional methods, various supervised deep CNN architectures using recursive residual blocks^{2,3}, residual connections, and attention-based modules^{4–7} have seen significant use in super-resolution applications.

The approaches for different super-resolution methods vary, but they all share the common goal of producing high-resolution image output that is a clear representation of the low-resolution input in the context of both image content and visual fidelity. Generally, evaluation metrics are centered around the idea that the output image is the product of intensity-based visible-light photography, where the goal is to represent what is seen with the human eye. However for scientific image applications, this idea is often incorrect, since many experimental methods construct images or image maps using electromagnetic information gathered from outside the visible-light range (e.g., X-rays, infrared information), or even not from light at all (e.g., electrons, neutrons), which means the ideas of sharpness and visual clarity have very different meanings. One such example, electron backscatter diffraction (EBSD), used for characterization of crystalline materials, relies on electron diffraction to build maps of material crystallographic information.

EBSD is a scanning electron microscopy technique that maps crystal lattice orientation by analyzing Kikuchi diffraction patterns that are formed when a focused electron beam is scattered by the atomic crystal structure of a material according to Bragg's law. A grid of Kikuchi patterns is collected by scanning the electron beam across the sample surface. These patterns are then indexed to form a grid of orientations, which are commonly represented as images in RGB color space using inverse pole figure (IPF) projections. EBSD datasets are used to determine microstructural properties of materials such as texture, orientation gradients, phase distributions, and point-to-point orientation correlations, all of which are critical for understanding material performance⁸. Furthermore, the fact that EBSD datasets are maps with orientation information, rather than just images with intensity values, changes both the notion of image quality and the motivation for super-resolution methods.

EBSD has two types of resolution: the accuracy with which the EBSD pattern collected at each pixel can be indexed into a crystallographic orientation, and the spacing between pixels in a given mapping. During EBSD mapping, the electron beam must dwell at each point long enough to form a high-quality Kikuchi pattern, which is then indexed into a specific crystallographic orientation^{9–11}. The indexing problem has many salient issues associated with it, among them are the crystallographic differentiation of matrix and precipitate phases, the identification of local strain effects, and the decoupling of overlapping diffraction patterns at grain boundaries. While accurate indexing is critical to EBSD, it is an independent challenge unrelated to super-resolution, as each indexing problem is treated as having no correlation with its neighbors. The lack of assumed spatial correlation separates indexing from other pixel-based problems and makes it ill-suited for SISR. Therefore, the issues associated with indexing accuracy described above are not addressed here.

¹Electrical and Computer Engineering, University California Santa Barbara, Santa Barbara 93106-5050 CA, USA. ²Mechanical Engineering, University California Santa Barbara, Santa Barbara 93106-5050 CA, USA. ³Computer Science, University California Santa Barbara, Santa Barbara 93106-5050 CA, USA. ⁴Materials Department, University California Santa Barbara, Santa Barbara 93106-5050 CA, USA. ⁵These authors contributed equally: Devendra K. Jangid, Neal R. Brodnik. ✉email: nbrodnik@ucsb.edu

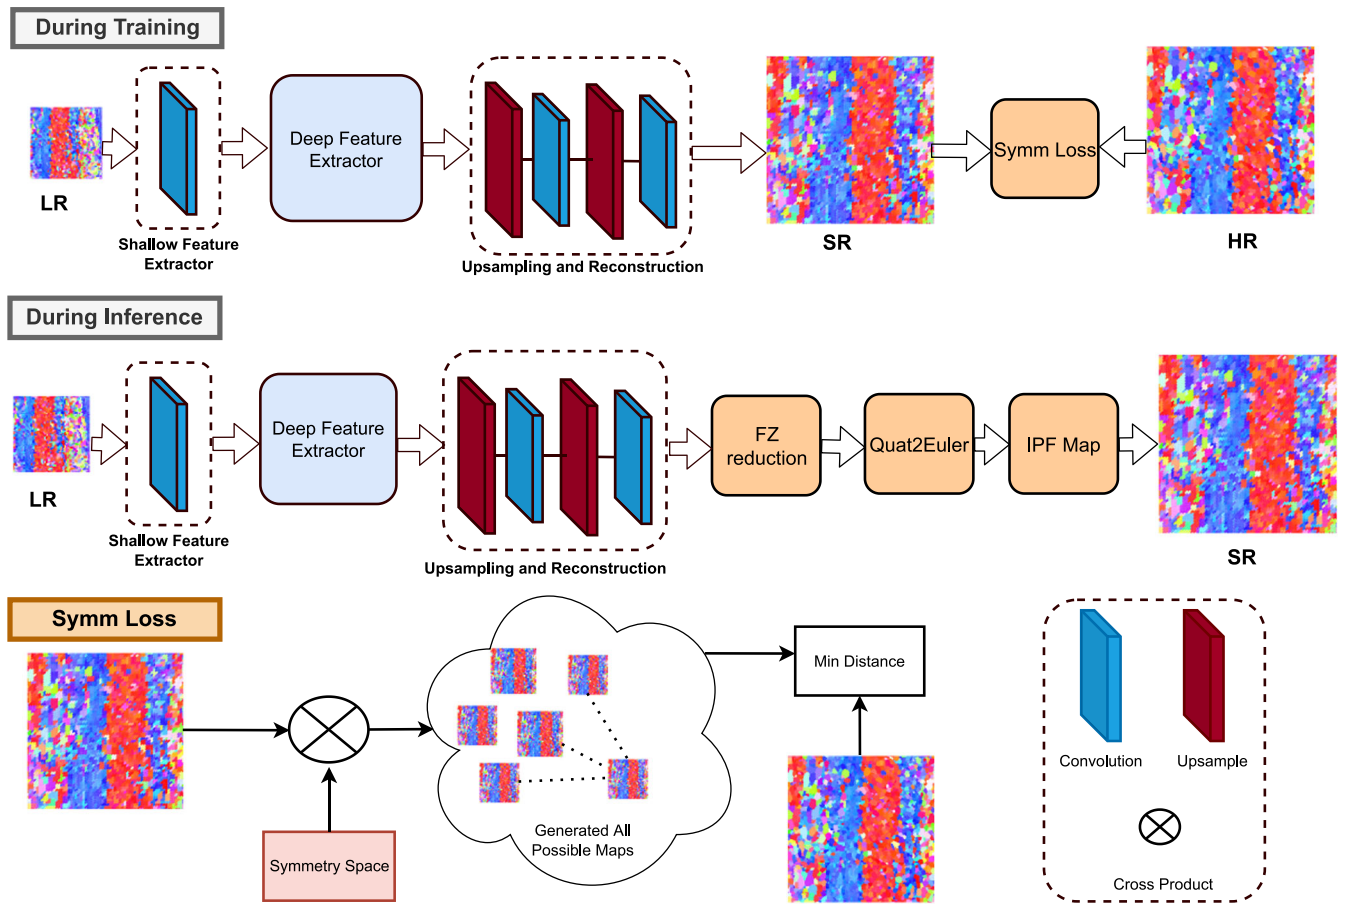


Fig. 1 Network architecture. Training pipeline: a low-resolution EBSD map in quaternion orientation space is given to an image super-resolution network architecture that generates a high-resolution EBSD map in quaternion orientation space. A crystallographic symmetry physics-based loss with L_1 or approximate rotational distance is used during training. Inference pipeline: the image super-resolution network generates a high-resolution EBSD map in quaternion orientation space, which is reduced to a fundamental zone space, and converted to Euler orientation space to visualize in IPF color map. Symmetry loss: takes all possible hexagonal symmetries for the titanium alloy, and computes the minimum distance between all possible generated output and ground truths. The distance can be L_1 or approximate rotational distance.

Instead, we consider the improvement of spatial resolution, which, in experiment, equates to collecting a higher density of data points during mapping. This requirement can lead to long mapping times or force the choice to use a coarser resolution mapping grid for expediency. The necessity to reduce EBSD collection time becomes even more critical when performing many scans during serial-sectioning 3D EBSD measurements¹², where material is removed layer-by-layer with 2D maps collected at each slice, and then stacked into a 3D dataset. In almost all serial-sectioning experiments, the minimum slicing thickness/resolution is much lower than the achievable in-plane imaging resolution, creating anisotropic voxels. Furthermore, poor electrically conductive materials become electrostatically charged by the beam or degrade from beam exposure (e.g., bio materials, polymers), requiring extremely short exposure times and resulting in Kikuchi patterns with weak contrast.

As the demand for greater volume and more detailed resolution material information grows, so too does the demand and expense of EBSD mapping. Spatial resolution in EBSD is of particular interest in the characterization of deformed materials and additive manufacturing¹³, where subgrain misorientation gradients are used to quantify local plastic deformation effects and geometrically necessary dislocation densities^{14–16}. In efforts to improve EBSD resolution and quality, simulations and experimental studies^{17–21} have shown that lowering the electron beam accelerating voltage can significantly improve the spatial resolution of EBSD maps, but map quality and achievable resolution vary

with differing materials and imaging conditions. To improve indexing accuracy, multiple algorithmic approaches have been developed for better Kikuchi pattern mapping^{22,23}, which improves both the precision and accuracy of the orientations shown at each pixel. Machine learning approaches have also been used to accelerate several tasks in the EBSD map construction process, including Kikuchi pattern indexing²⁴, classification²⁵, and crystal identification²⁶. Recently, a residual-based neural network with traditional L_1 loss (ResNet) was used to produce super-resolved EBSD maps from inverse pole figure (IPF) color and Euler angles as an image input²⁷. The desire to accelerate and improve the EBSD mapping process has motivated a wide array of machine learning approaches, but many challenges still exist. One of the most prominent of these is that orientation space is discontinuous and repeating, and the fundamental shape of orientation space changes with the symmetries of the crystal being observed. This makes brute-force network learning with traditional methods highly dependent on the available data for training, and, depending on the orientation, small variations in accuracy can produce dramatically incorrect results.

Given these challenges, we present an adaptable framework for neural-network-based super-resolution of EBSD maps, where all network learning is built around the physics of crystal orientation symmetry. We define a physics-based loss that accounts for crystallographic symmetries, which is used alongside either a traditional L_1 loss metric or a loss based directly on rotational arc lengths, which correspond to conventional misorientation

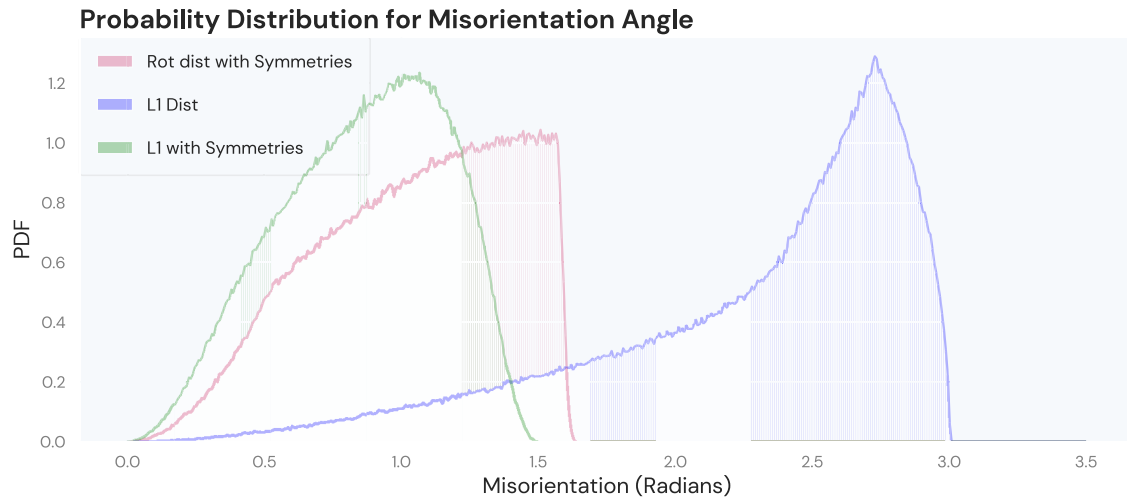


Fig. 2 Orientation loss distributions. Probability distribution of loss distances between pairs of randomly sampled 3D rotation vectors. L_1 distance is shown in blue, L_1 distance accounting for crystal symmetry is shown in green, and approximate rotational distance accounting for crystal symmetry is shown in red.

measurements in crystallography. All super-resolution is done on crystal orientation data expressed as quaternions, meaning each pixel in a given map contains four channels. Using quaternion space allows for complete representation of orientation space, enabling a training approach that is translatable across all 230 crystallographic space groups. As a proof of concept, four state-of-the-art residual and channel-attention networks are used to generate high-resolution EBSD maps from low-resolution input using this physics-based approach. We demonstrate that regardless of network choice, physics-based approaches outperform traditional approaches both qualitatively and quantitatively. This approach has direct application to experimental EBSD measurements of electron beam-sensitive or low-conductivity materials where charge buildup and extended beam exposure are limiting factors, and for 3D EBSD data collection where high out-of-plane imaging resolution is costly. We expect SR-EBSD to accelerate EBSD mapping for defect detection and fast screening of microstructure configurations that limit material properties.

RESULTS

Framework architecture

Unlike traditional SR architectures that focus on 1 or 3-channel image data, the framework developed in this study, shown in Fig. 1, directly operates on vector-based orientation data, which is expressed as 4-channel quaternions. Quaternions are mathematically robust for crystal orientation representation and allow easy application of symmetry operators between different crystal systems.

The process for this framework follows similar methodologies to experimental EBSD mapping, where a priori crystal symmetry information is used to inform the training and inference processes. For the physics-informed framework, a loss function satisfying space group symmetry requirements is defined, and networks are trained with that loss function on datasets of materials from the corresponding space group. Once trained, the network can infer on EBSD maps for materials that fall under the same space group symmetry. The full framework consists of a shallow feature extractor, a deep-feature extractor, an upscaling and reconstruction module, and the loss function, which dictates symmetry reinforcement.

Shallow feature extractor. This module uses a single convolution layer to reduce spatial size of original EBSD Map and extract shallow features from a given low-resolution 4-channel EBSD map (I_{LR}).

$$F_0 = H_{SF}(I_{LR}) \quad (1)$$

Here, $H_{SF}(\cdot)$ is a single convolution layer of kernel size 3×3 , which has 4 input channels and 128 output channels. The generated shallow features (F_0) are fed to the deep-feature extractor module (H_{DF}).

Deep-feature extractor. To extract essential deep features from EBSD maps, we implemented deep-feature extractors from four well-known single-image super-resolution network architectures. These four architectures employ a variety of recent approaches to the SISR problem: deep residual (EDSR)⁴, channel attention (RCAN)⁵, second-order attention (SAN)⁶, and holistic attention (HAN)⁷ methods. Testing across all four networks enables both a robust analysis of loss functions and a broader understanding of the performance of different architectures in the EBSD-SR problem. It also emphasizes that any deep-feature extractor in existence today or developed in the future can be readily applied to this physics-based learning framework.

$$F_{DF} = H_{DF}(F_0) \quad (2)$$

where, $H_{DF}(\cdot)$ is a deep-feature extractor module, and F_{DF} is a 128 channels feature map, which goes to upscale and reconstruction module.

Upscale and reconstruction module. The extracted deep feature (F_{DF}) uses upscale modules that employ the pixel shuffle operation to initialize new pixels for each resolution increase²⁸. The upscaled feature is then mapped into a super-resolved EBSD map using convolution layers that are proportional to the resolution scaling factor.

$$F_{\uparrow} = H_{\uparrow}(F_{DF}) \quad (3)$$

$$I_{SR} = H_R(F_{\uparrow}) \quad (4)$$

H_{\uparrow} is the upscale module, and H_R is convolution operation of kernel size 3×3 for reconstruction. I_{SR} is the 4 channels generated high-resolution EBSD map in quaternion space.

Loss functions

All networks mentioned in section "Framework architecture" were trained using three different loss functions based on a combination of both established practices for the super-resolution problem and the underlying physics associated with EBSD orientation maps. The losses used are traditional " L_1 loss" and two different physics-guided losses, termed " L_1 with symmetry", and

“approximate rotational distance with symmetry”. A histogram comparison of all loss distances for a sample of random orientation vector pairs is shown in Fig. 2.

L_1 loss. L_1 loss is a standard norm loss that has been widely used in image restoration tasks and has been shown to have advantages over L_2 loss²⁹ in terms of sharpness and visual clarity. There is no underlying physical motivation for using L_1 loss beyond its observed advantages in traditional image restoration tasks, which have established the precedent for its use in the SISR problem. The L_1 loss between generated and ground truth EBSD maps is described using the following equation.

$$L_1 = \frac{1}{N} \sum_{i=1}^N \|H(I_{LR}) - I_{HR}\|_1 \quad (5)$$

where I_{HR} is the ground truth EBSD map, $H(I_{LR})$ is the generated EBSD map, and N is the batch size.

Space group symmetry. The orientations in an EBSD map for a given material can only be understood properly in the context of the space group of that material. These same symmetry operators persist in the EBSD diffraction patterns, and create boundaries in orientation space during pattern indexing, dividing the complete sphere of possible quaternion orientations into repeating subsections. For this reason, the crystallographic relationships associated with pixel rotation values in the EBSD maps were accounted for using what we have termed symmetry loss, as shown in Fig. 1. Symmetries were accounted for according to the space group conventions used to describe crystal symmetry systems. This space group information is provided a priori during training and inference, but this requirement is not considered overly rigorous, as EBSD measurements typically use a priori space group information to simplify the indexing problem. The titanium datasets investigated in this work are part of space group 194, which has a total of 24 symmetries, but only 12 that do not involve a change of handedness. For symmetry-based loss, every pixel value generated by the network is considered as a collection of rotations across all of these symmetries, and the loss distance is calculated as the minimum distance between the ground truth and any value within this collection. For this study, space group symmetry is enforced at the image level. For multi-phase materials, enforcement could also be done at the pixel level through implementation of a phase map.

L_1 loss with symmetry. This loss uses L_1 distance to calculate loss magnitude, but incorporates physics to account for space group symmetry in the EBSD map.

Rotational distance approximation loss with symmetry. Rotational distance loss computes the misorientation angle between the predicted and ground truth EBSD map in the same manner that they would be measured in crystallographic analysis, with approximations to avoid discontinuities. The rotational distance between two quaternions can be computed as the following:

$$\begin{aligned} \theta &= 2\cos^{-1}(\text{Re}(q_1 q_2^*)) \\ &= 2\cos^{-1}(\langle \vec{q}_1, \vec{q}_2 \rangle) \\ &= 2\cos^{-1}\left(1 - \frac{1}{2}\|\vec{q}_1 - \vec{q}_2\|_2^2\right) \\ &= 4\sin^{-1}\left(\frac{1}{2}\|\vec{q}_1 - \vec{q}_2\|_2\right) \\ \theta &= 4\sin^{-1}\left(\frac{d_{\text{euclid}}}{2}\right) \end{aligned} \quad (6)$$

where, $d_{\text{euclid}} = \|q_1 - q_2\|_2$.

While d_{euclid} is Lipschitz, the gradient of θ goes to ∞ as $d_{\text{euclid}} \rightarrow 2$. To address this issue in training a neural network, a linear approximation was computed at $d_{\text{euclid}} = 1.9$, and utilized for points > 1.9 . This can be seen in Fig. 3 as a clamp on the max value the derivative of the function can take on. This loss is the most physically accurate of the three considered, and the distribution of loss values for random rotation vectors for rotational distance loss (shown in red in Fig. 2) matches with the probability distribution of misorientations for hexagonal polycrystals³⁰.

Network output is evaluated using a combination of image quality and domain relevant metrics. Initially, each set of generated images is evaluated using peak signal-to-noise ratio (PSNR) and structure similarity index measure (SSIM). Images are then segmented into individual grain regions using watershed segmentation based on a relative misorientation tolerance. This approach, coupled with domain knowledge, is commonly used to identify and segment grains in EBSD maps.

Qualitative output comparison

To evaluate framework performance, an experimentally gathered dataset of Ti-6Al-4V³¹, discussed in section “Data preprocessing,” was chosen as a candidate dataset for super-resolution testing. This set was of particular interest because of its large number of grains, small grain size, varying local texture, and wide range of

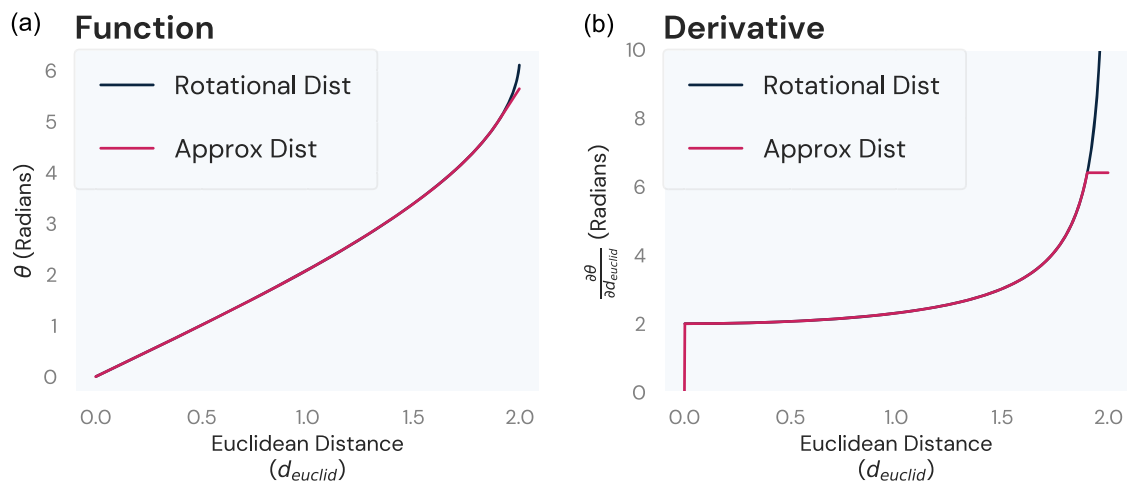


Fig. 3 Approximation of the rotational distance equation. The derivative is not defined at $d_{\text{euclid}} = 2$, so a linear approximation is computed to ensure smooth loss behavior. **a** Plot of rotational distance functions against euclidean distance. **b** Plot showing the derivative of rotational distance against euclidean distance, showing divergence.

represented orientations, all of which make it challenging for the super-resolution problem, as it is desirable to preserve both the gradual orientation gradients within grains as well as the sharp orientation discontinuities at grain boundaries. For training and testing, all LR input was downsampled by a factor of 4 using direct removal of pixel rows and columns to reflect how EBSD resolution would be reduced in actual experiments. When comparing qualitative image results, the reduction from HR to LR by removal of rows and columns from the dataset causes a reduction in visual quality of EBSD maps while offering no possible information inference from subpixel values. This makes shape inference from LR input particularly difficult, as any features removed during downsampling are lost completely. Although it is challenging, our goal is direct application to experimental data, and this form of data loss is exactly what occurs when low-resolution EBSD maps are gathered experimentally. The difficulty associated with this dataset and downsampling is apparent in the traditional bilinear,

bicubic, and nearest neighbor upsampling algorithms, shown in Fig. 4. Bilinear and bicubic approaches produce non-physical results that are the product of interpolations made through quaternion orientation space with no regard for symmetry relationships. Nearest neighbor approximations produce higher resolution visual replicas of the low-resolution input.

A full comparison of image quality on the test set for all networks and losses is shown in Fig. 5. All investigated networks produced significantly better results than traditional algorithms, but there were variations in quality across the different network types and the loss function used. For all networks, grain shapes differed from HR, as many of the fine grain features are lost. This is expected, as the downsampling method removes much of this information, and networks cannot infer on data that does not exist. However, even with different shapes, L_1 and L_1 with symmetry produce grains with smooth contours and a range of facet angles, which is a good reflection of what would be

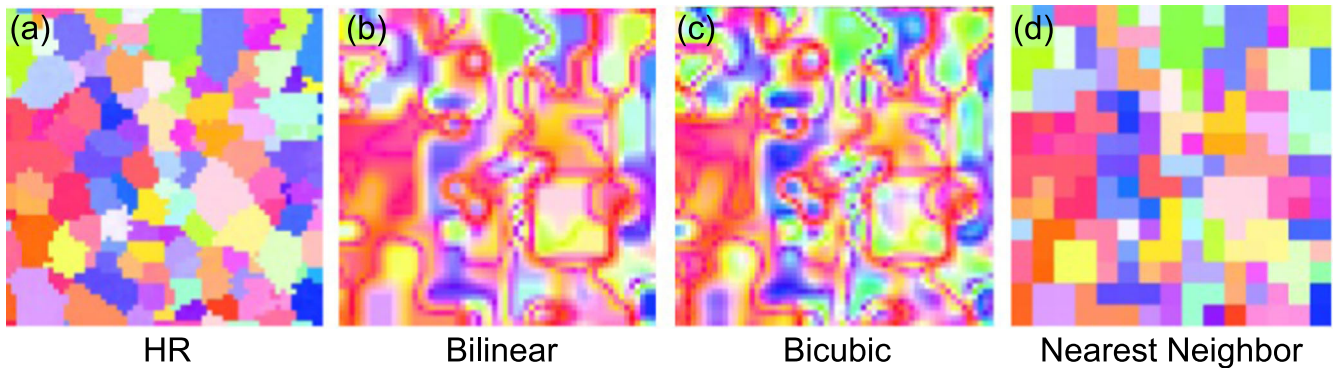


Fig. 4 Comparison of HR ground truth to traditional upsampling. Compared to **a** high-resolution ground truth EBSD map, **b** bilinear, **c** bicubic, and **d** nearest neighbor upsampling produce inferior results. Bilinear and bicubic results are non-physical, and nearest neighbor results are visually identical to LR input.

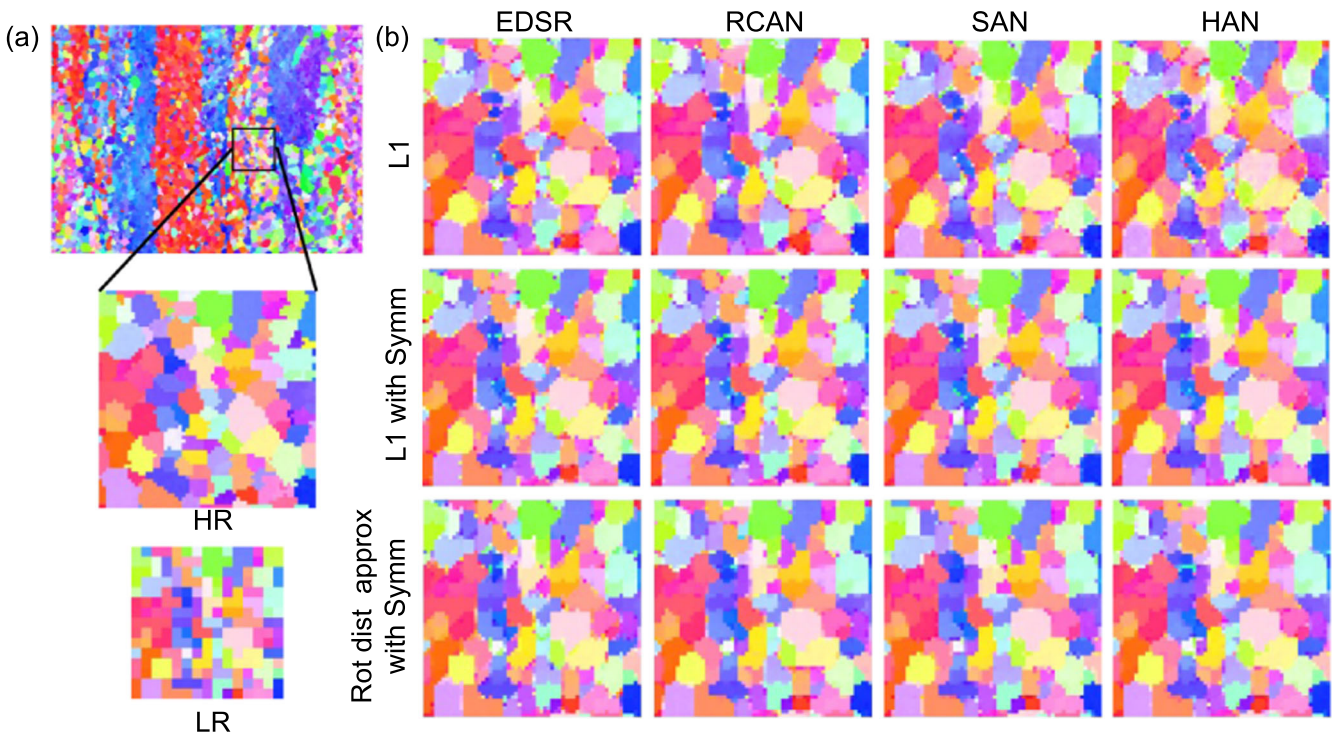


Fig. 5 Visual Comparison of 4x EBSD-SR on a test EBSD map. L_1 has non-physical structure at grain boundaries, and L_1 with symmetry and approximate rotational with symmetry reduce the non-physical structures at grain boundaries. **a** The comparison of LR and HR for a single patch of size 64×64 in an EBSD map. **b** Compares super-resolved EBSD maps for the same patch across different architectures and losses.

Table 1. PSNR/SSIM comparison for 4× super-resolution scaling: physics-based loss consistently outperforms bilinear, bicubic, nearest neighbor, and pure L_1 loss with no physics, regardless of architecture.

	EDSR	RCAN	SAN	HAN
L_1 (no physics)	14.85/0.472	14.89/0.474	14.83/0.465	14.65/ 0.450
L_1 with symmetry	15.05/0.483	15.07/0.484	15.04/0.484	15.13/ 0.497
Rot. dist. approx with symmetry	15.02/0.486	15.24/0.509	15.23/0.510	15.30/ 0.513

Rows represent different loss functions and columns represent different network architectures. Higher number is desired for both PSNR/SSIM. Bold values indicate architecture with best performance for each loss metric.

Table 2. PSNR/SSIM comparison for 4× algorithmic scaling: PSNR/SSIM values for bilinear, bicubic, and nearest neighbor are consistently lower than network-based methods.

	Bicubic	Bilinear	Nearest neighbor
Ti-6Al-4V: (PSNR/SSIM)	11.22/0.211	11.30/0.237	13.25/0.373
Ti-7Al 1% (PSNR/SSIM)	18.47/0.7135	18.66/0.751	22.63/0.823
Ti-7Al 3% (PSNR/SSIM)	19.49/0.7372	19.65/0.773	24.11/0.8328

expected in real materials. However, across all networks, both of these losses produce non-physical image artifacts at the edges of grains. Including physics-based symmetry into L_1 reduces the quantity of these compared to traditional L_1 , but the most dramatic reduction in artifacts occurs for approximate rotational distance with symmetry. Although rotational distance tended to produce slightly more cube-like grain shapes, the overall reduction in visual artifacts makes network output much more physically accurate.

Quantitative output comparison

Quantitative evaluations of peak signal-to-noise ratio (PSNR) and structural similarity index measure (SSIM) across different architectures and losses are shown in Table 1. Comparing between networks, HAN consistently performed best in all cases of physics-based loss. Across all network architectures, the incorporation of physics into the loss function led to improvement in both PSNR and SSIM. Overall, the most physically accurate loss metric, approximate rotational distance, performed best. Bilinear, bicubic, and nearest neighbor scores for Ti-6Al-4V are shown in row 1 of Table 2, but their scores were dramatically lower, and their nonphysicality made further consideration irrelevant.

Table 3 shows the percentage of pixels in the test set identified as single-pixel features in a watershed segmentation with a misorientation tolerance of 10° . With high-resolution ground truth having around 0.2% single-pixel features, this metric can be considered as a close approximation to the percentage of single-pixel artifacts, which appear as salt-and-pepper noise at grain boundaries. Across all architectures, the incorporation of physics-based loss leads to a clear reduction in single-pixel artifacts. Overall, RCAN with rotational distance loss with symmetry had the fewest single-pixel features, but RCAN showed comparatively worse values for L_1 loss with symmetry. These variations in noise are small when compared to the differences between physics-based and non-physics-based loss.

Table 4 shows the percentage difference in number of features relative to LR input, when a watershed segmentation with a misorientation tolerance of 10° is applied. A larger percentage value indicates higher feature counts than the LR input, which correlates to occurrence of noise and spurious features, and arises from regions of noise at grain boundaries or overemphasized misorientation gradients within grains. All network outputs produced greater feature quantities than the input, but L_1 loss

Table 3. Percentage of single-pixel features: rows represent different loss functions and columns represent different network architectures.

	EDSR	RCAN	SAN	HAN
L_1 (no physics)	5.7	5.5	6.0	7.4
L_1 with symmetry	4.6	4.9	4.4	2.7
Rot. dist. approx with symmetry	1.8	0.9	1.3	1.4

Lower values indicate better performance, with ground truth containing approximately 0.2% single-pixel features. Physics-based loss reduces noise, leading to lower single-pixel feature counts.

Bold values indicate architecture with best performance for each loss metric.

Table 4. Percent feature difference over input: rows represent different loss functions and columns represent different network architectures.

	EDSR	RCAN	SAN	HAN
L_1 (no physics)	522.2	475.7	530.0	628.2
L_1 with symmetry	433.3	444.4	415.5	284.1
Rot. dist. approx with symmetry	244.3	152.7	166.4	159.3

Lower values indicate better performance. Physics-based loss reduces noise and spurious features, which keeps the overall feature count closer to the amount expected based on input.

Bold values indicate architecture with best performance for each loss metric.

without symmetry caused the greatest number of spurious features, having over 450% more features than the initial input. Once symmetry is introduced, performance improves across all networks, with HAN producing the fewest excess features for L_1 loss with symmetry, and RCAN producing the fewest excess features for approximate rotational distance with symmetry.

Additional material systems

To verify the robustness of the SR-EBSD approach across different materials datasets, the holistic attention network (HAN), which showed consistently strong performance, was trained on two additional experimental datasets of Ti-7Al³², which were plastically deformed to 1% and 3% strain respectively, and are discussed in section “Data preprocessing”. These two datasets contain much larger grains than the Ti-6Al-4V set, and also exhibit texture due to material processing. The HAN was trained on all three datasets together rather than on each set individually. Results are shown in Table 5, and comparison to traditional upscaling algorithms for Ti-7Al is shown in rows 2 and 3 of Table 2.

When comparing PSNR/SSIM values, the values for both Ti-7Al datasets are much higher than those for the Ti-6Al-4V set. This is likely due to a combination of grain size and texture differences between sets, with the Ti-6Al-4V set having smaller grains and microtextured regions. Much larger grains makes shape inference

Table 5. PSNR/SSIM comparison for 4 × super-resolution scaling across different materials: Ti-6Al-4V, Ti-7Al 1%, and Ti-7Al 3% for HAN network: physics-based loss consistently outperforms bilinear, bicubic, nearest neighbor, and pure L_1 loss with no physics, and columns represent different titanium datasets.

	Ti-6Al-4V	Ti-7Al 3%	Ti-7Al 1%
L_1 (no physics)	14.94/0.478	26.66/0.865	25.37/0.852
L_1 with symmetry	15.19/0.497	26.82/0.869	25.38/0.856
Rot. dist. approx with symmetry	15.35/0.5264	27.25/0.881	25.72/0.871

Higher number is desired for both PSNR/SSIM.

at a 4× scale reduction less difficult and strong material texture makes the range of orientations in the dataset narrower, both of which reduce the burden on the network during the learning/inference process. Regardless of dataset difficulty, rotational distance loss with symmetry consistently produces the highest quality results.

DISCUSSION

The results demonstrate that regardless of architecture or approach, the incorporation of domain-related physics into the training process leads to better results for the SISR problem for scientific data. In both qualitative and quantitative evaluation across every metric considered, physics-based loss consistently outperformed traditional L_1 loss regardless of the network used. When comparing physics-based losses specifically, rotational distance consistently outperformed all other losses on all quantitative metrics, though it should be noted that L_1 with symmetry exhibited well-behaved grain shapes in qualitative evaluation while also having quantitative results superior to traditional methods. For the architectures studied here, comparison across all evaluation metrics shows that attention-based models outperform purely residual architectures in the EBSD-SR problem, with holistic attention (HAN) exhibiting the best performance. This result is somewhat expected, as the layer attention module and channel-spatial attention modules present in the HAN architecture provide additional ability to learn spatial correlations across different channels and layers⁷. For this reason, the HAN architecture also exhibited the strongest performance on benchmark datasets for the traditional single-image super-resolution problem⁷.

The results of this investigation present a strong case for the benefits of an adaptable learning approach that can be easily applied to future architectures. Although HAN was the best performing network overall, it still was not able to consistently outperform other networks across all metrics presented here. Furthermore, when considering the range of possible microstructures that can exist across 230 different space groups and their multi-phase combinations, it is likely there will never be a single architecture that consistently performs best. The physics-based approach presented here improves performance and presents a path for EBSD super-resolution to keep pace with developments across the SISR field as a whole. This approach can be readily extended to other materials and microstructures using phase masks labeled by space group, accounting for each of their respective symmetries using the methods described here. Going forward, this approach to physics-informed EBSD super-resolution can be used in high-throughput EBSD experiments for the generation of larger, more robust datasets, and also to make higher resolution 3D datasets when combined with asymmetric serial-sectioning approaches (higher resolution in x , and y , lower resolution in z). With spatial super-resolution, the number of time-

consuming pattern gathering steps can be reduced. EBSD patterns in 2D, which would normally take minutes or hours to gather can be super-resolved in seconds, and 3D EBSD patterns, which would normally take days or weeks to gather can be super-resolved in minutes. These super-resolution tools can accelerate the materials development process while ensuring that all network learning occurs in the domain boundaries established by physics and crystallography.

In conclusion, we designed an adaptable, physics-guided approach to the super-resolution problem for EBSD orientation maps that employs several deep-feature extraction methods from existing single-image super-resolution architectures, as well as losses accounting for crystal symmetry and rotation physics. Unlike existing SR methods, which operate on scalar image data, the training pipeline is implemented in quaternion orientation space. The inference pipeline produces quaternion output that is converted into Euler angle representation and colored based on IPF projection conventions. Qualitative and quantitative image analysis demonstrate that networks with physics-based learning consistently outperform both traditional upscaling algorithms and analogous network approaches that do not employ physics. Accounting for crystal symmetry in learning leads to increases in PSNR and SSIM, and also reduces single-pixel artifacts and spurious visual features. L_1 loss with symmetry produces well-behaved grain shapes, and approximate rotational distance with symmetry greatly reduces the occurrence of noise and visual artifacts. The presented framework can be readily applied to future super-resolution network architectures.

METHODS

Orientation representation

The focus of this investigation centers on orientation vector maps that describe crystalline domains, which are fundamentally anisotropic and periodic in nature. Orientations for each pixel within the network learning environment are expressed in terms of quaternions of the form $q = q_0 + \hat{i}q_1 + \hat{j}q_2 + \hat{k}q_3$. For orientation representation, quaternions are beneficial due to their lack of ambiguity with respect to orientation representation and crystal symmetry. To avoid redundancy in quaternion space (between q and $-q$), all orientations are expressed with their scalar q_0 as positive. For visualization according to established conventions, quaternions are reduced to the Rodrigues space fundamental zone based on space group symmetry, converted into Euler angles, and projected using inverse pole figure (IPF) projection using the open-source Dream3D software³³.

Pipeline

In order to maximize applicability of EBSD super-resolution to materials research, network output must be interpretable based on established crystallography conventions. To facilitate this, we designed an inference pipeline where network output can be converted to the visualization space used and accepted within the field, as shown in Fig. 1. During training, we use a physics-based symmetry loss within existing image super-resolution network architectures, but modified to have the appropriate number of input and output channels. The input to the network is a LR EBSD map, and the generated output is a HR EBSD map, both in the quaternion domain. The physics-based loss, which gives the network information about crystal symmetry, is computed in quaternion orientation space. During inference, the network outputs a HR EBSD map in quaternion space, which is then reduced to the fundamental zone in Rodrigues vector space before being converted into Euler space and projected using inverse pole figure (IPF) color to visualize orientations. IPF color maps are generated with the commercially available open-source Dream3D software³³.

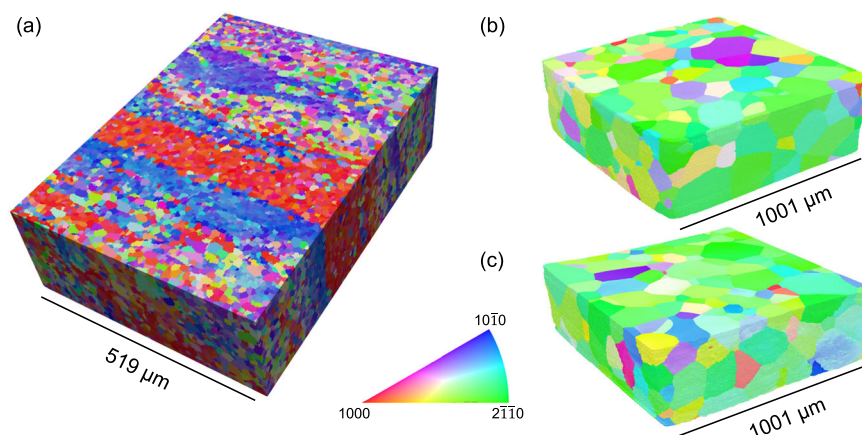


Fig. 6 Rendering of 3D EBSD datasets Investigated. Shown in IPF coloring of the titanium alloys, **a** Ti-6Al-4V and **b** Ti-7Al mechanically loaded in tension to 1% strain and **c** 3%, used for network training, testing and validation. Dataset details are available elsewhere^{31,32}.

Data preprocessing

High-resolution (HR) ground truth. The ground truth data is experimental 3D EBSD data gathered from two titanium alloys, Ti-6Al-4V and Ti-7Al, using a commercially available rapid-serial-sectioning electron microscope known as the Tribeam^{12,34}. The Ti-6Al-4V dataset, shown in Fig. 6a, is of pixel size $346 \times 142 \times 471 \times 4$, where the last dimension is the quaternion component. Analogously, the Ti-7Al shown in Fig. 6b, c are of size $770 \times 674 \times 132 \times 4$ and $770 \times 770 \times 224 \times 4$ pixels, respectively, with all edges cropped to produce a perfect parallelepiped volume. Each voxel in the Ti-6Al-4V set has resolution of $1.5 \times 1.5 \times 1.5 \mu\text{m}$, and in both Ti-7Al sets, each voxel has a resolution of $1.3 \times 1.3 \times 1.3 \mu\text{m}$.

These titanium alloys are composed primarily of the hexagonal close packed grains. In total, the Ti-6Al-4V dataset contains about 57,000 grains, visible in the IPF maps as regions of different color. The Ti-7Al material has larger grain size, with 500–1000 grains in each dataset. The datasets were proportionally divided into training, validation, and test subsets in ratios of 65%, 15%, and 20%, respectively. The sample volume was sectioned such that each subset contains an equivalent fraction of orthogonal images from each face of the sample to avoid any bias due to material anisotropy. During training, each volume was broken into two-dimensional patches of size 64×64 pixels.

Low-resolution (LR) input. Low-resolution EBSD maps used as network input are downscaled versions of the high-resolution ground truth. However, because of how EBSD information is gathered, these images are not downscaled using pixel averaging. Instead the low-resolution EBSD maps are produced by removal of rows and columns from the high-resolution ground truth with a downscale factor of 4 ($LR = \frac{1}{4}HR$). This is done to imitate the beam raster steps that would occur in an EBSD experiment with lower resolution, where a lower resolution would not influence the electron beam-material interaction volume at each pixel, but rather lead to greater raster distance between consecutive pixels of the same size.

Network implementation and output evaluation

We use a learning rate of 0.0002, Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.99$, ReLU activation, batch size of 4 and downscaling factor of 4. The patch size of HR images is 64. The framework is implemented in PyTorch and trained on NVIDIA Tesla V100 GPU. For a batch size of 4 and patch size of 64 across the 3 datasets in this work, the training time for each network for 2000 epochs was ~60 h. Once training is complete, inference time for a given input

is on the order of less than one second for an imaging area that would normally take about 10 min to gather manually.

For performance evaluation using watershed segmentation, we use a misorientation tolerance of 10 degrees. This is considered a conservative tolerance for feature identification, making it well suited for identification of image artifacts. Watershed segmentation was performed using Dream3D³³.

DATA AVAILABILITY

Pre-trained versions of the network modules produced in this paper are publicly available through the BisQue cyberinfrastructure at https://bisque2.ece.ucsb.edu/client_service/. Material datasets will be available by request at discretion of the authors.

CODE AVAILABILITY

Architecture code is publicly accessible at <https://github.com/UCSB-VRL/EBSD-Supersolution>.

Received: 3 February 2022; Accepted: 1 November 2022;
Published online: 15 December 2022

REFERENCES

- Wang, Z., Chen, J. & Hoi, S. C. Deep learning for image super-resolution: a survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **43**, 3365–3387 (2021).
- Kim, J., Lee, J. K. & Lee, K. M. Deeply-recursive convolutional network for image super-resolution. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1637–1645 (IEEE, 2016).
- Tai, Y., Yang, J. & Liu, X. Image super-resolution via deep recursive residual network. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3147–3155 (IEEE, 2017).
- Lim, B., Son, S., Kim, H., Nah, S. & Mu Lee, K. Enhanced deep residual networks for single image super-resolution. In: *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 136–144 (IEEE, 2017).
- Zhang, Y. et al. Image super-resolution using very deep residual channel attention networks. In: *Proceedings of the European conference on computer vision (ECCV)*, 286–301 (IEEE, 2018).
- Dai, T., Cai, J., Zhang, Y., Xia, S.-T. & Zhang, L. Second-order attention network for single image super-resolution. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11065–11074 (IEEE, 2019).
- Niu, B. et al. Single image super-resolution via a holistic attention network. In: *European Conference on Computer Vision*, 191–207 (Springer, 2020).
- Schwartz, A. J., Kumar, M., Adams, B. L. & Field, D. P. (eds.) *Electron Backscatter Diffraction in Materials Science* (Springer US, 2009).
- Lassen, K. Automatic high-precision measurements of the location and width of kikuchi bands in electron backscatter diffraction patterns. *J. Microsc.* **190**, 375–391 (1998).

10. Chen, Y. H. et al. A dictionary approach to electron backscatter diffraction indexing. *Microsc. Microanalysis* **21**, 739–752 (2015).
11. Jackson, M. A., Pascal, E. & De Graef, M. Dictionary indexing of electron backscatter diffraction patterns: a hands-on tutorial. *Integrating Mater. Manuf. Innov.* **8**, 226–246 (2019).
12. Echlin, M. P., Burnett, T. L., Polonsky, A. T., Pollock, T. M. & Withers, P. J. Serial sectioning in the SEM for three dimensional materials science. *Curr. Opin. Solid State Mater. Sci.* **24**, 100817 (2020).
13. Sames, W. J., List, F. A., Pannala, S., Dehoff, R. R. & Babu, S. S. The metallurgy and processing science of metal additive manufacturing. *Int. Mater. Rev.* **61**, 315–360 (2016).
14. Brewer, L. N., Field, D. P. & Merriman, C. C. Mapping and assessing plastic deformation using EBSD. In: *Electron Backscatter Diffraction in Materials Science*, 251–262 (Springer US, 2009).
15. Witzten, W. A., Polonsky, A. T., Pollock, T. M. & Beyerlein, I. J. Three-dimensional maps of geometrically necessary dislocation densities in additively manufactured Ni-based superalloy IN718. *Int. J. Plasticity* **131**, 102709 (2020).
16. Witzten, W. A. et al. Boundary characterization using 3d mapping of geometrically necessary dislocations in AM Ta microstructure. *J. Mater. Sci.* <https://doi.org/10.1007/s10853-022-07074-2> (2022).
17. Ren, S., Kenik, E., Alexander, K. & Goyal, A. Exploring spatial resolution in electron back-scattered diffraction experiments via monte carlo simulation. *Microsc. Microanalysis* **4**, 15–22 (1998).
18. Chen, D., Kuo, J.-C. & Wu, W.-T. Effect of microscopic parameters on EBSD spatial resolution. *Ultramicroscopy* **111**, 1488–1494 (2011).
19. Keller, R. & Geiss, R. Transmission EBSD from 10 nm domains in a scanning electron microscope. *J. Microsc.* **245**, 245–251 (2011).
20. Steinmetz, D. R. & Zaefferer, S. Towards ultrahigh resolution EBSD by low accelerating voltage. *Mater. Sci. Technol.* **26**, 640–645 (2010).
21. Tripathi, A. & Zaefferer, S. On the resolution of EBSD across atomic density and accelerating voltage with a particular focus on the light metal magnesium. *Ultramicroscopy* **207**, 112828 (2019).
22. Jackson, M. A., Pascal, E. & De Graef, M. Dictionary indexing of electron backscatter diffraction patterns: a hands-on tutorial. *Integrating Mater. Manuf. Innov.* **8**, 226–246 (2019).
23. Lenthe, W., Singh, S. & De Graef, M. A spherical harmonic transform approach to the indexing of electron back-scattered diffraction patterns. *Ultramicroscopy* **207**, 112841 (2019).
24. Ding, Z., Zhu, C. & De Graef, M. Determining crystallographic orientation via hybrid convolutional neural network. *Mater. Charact.* **178**, 111213 (2021).
25. Kaufmann, K., Lane, H., Liu, X. & Vecchio, K. S. Efficient few-shot machine learning for classification of EBSD patterns. *Sci. Rep.* **11** <https://doi.org/10.1038/s41598-021-87557-5> (2021).
26. Kaufmann, K. et al. Crystal symmetry determination in electron diffraction using machine learning. *Science* **367**, 564–568 (2020).
27. Jung, J. et al. Super-resolving material microstructure image via deep learning for microstructure characterization and mechanical behavior analysis. *npj Comput. Mater.* **7** <https://doi.org/10.1038/s41524-021-00568-8> (2021).
28. Shi, W. et al. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1874–1883 (IEEE, 2016).
29. Zhao, H., Gallo, O., Frosio, I. & Kautz, J. Loss functions for image restoration with neural networks. *IEEE Trans. Comput. Imaging* **3**, 47–57 (2017).
30. Skrytnyy, V. I. & Yaltsev, V. N. Relative misorientations of crystals. *IOP Conf. Ser.: Mater. Sci. Eng.* **130**, 012059 (2016).
31. Hémerly, S. et al. A 3D analysis of the onset of slip activity in relation to the degree of micro-texture in Ti-6Al-4V. *Acta Materialia* **181**, 36–48 (2019).
32. Charpagne, M., Stinville, J. C., Polonsky, A. T., Echlin, M. P. & Pollock, T. M. A multi-modal data merging framework for correlative investigation of strain localization in three dimensions. *JOM* **73**, 3263–3271 (2021).
33. Groeber, M. A. & Jackson, M. A. DREAM.3D: a digital representation environment for the analysis of microstructure in 3D. *Integr. Mater. Manuf. Innov.* **3**, 56–72 (2014).
34. Echlin, M. P., Straw, M., Randolph, S., Filevich, J. & Pollock, T. M. The TriBeam system: Femtosecond laser ablation in situ SEM. *Mater. Charact.* **100**, 1–12 (2015).

ACKNOWLEDGEMENTS

This research is supported in part by NSF awards number 1934641 and 1664172. The authors gratefully acknowledge Patrick Callahan, Toby Francis, Andrew Polonsky, and Joseph Wendorf for collection of the 3D Ti-6Al-4V and Ti-7Al datasets. The MRL Shared Experimental Facilities are supported by the MRSEC Program of the NSF under Award No. DMR 1720256; a member of the NSF-funded Materials Research Facilities Network (www.mrfn.org). Use was also made of computational facilities purchased with funds from the National Science Foundation (CNS-1725797) and administered by the Center for Scientific Computing (CSC). The CSC is supported by the California NanoSystems Institute and the Materials Research Science and Engineering Center (MRSEC; NSF DMR 1720256) at UC Santa Barbara. Use was made of the computational facilities purchased with funds from the National Science Foundation CC* Compute grant (OAC-1925717) and administered by the Center for Scientific Computing (CSC). The ONR Grant N00014-19-2129 is also acknowledged for the titanium datasets.

AUTHOR CONTRIBUTIONS

D.K.J.: development of framework, design and training of networks, qualitative and quantitative evaluation of results N.R.B.: development of framework, design of loss functions, preparation of datasets for network training, qualitative and quantitative evaluation of results M.G.G.: development of framework, design of loss functions, quantitative evaluation of results A.K.: development of framework, design and training of networks S.M.: development of framework, design and training of networks M.P.E.: preparation of datasets, qualitative and quantitative evaluation of results S.H.D. (Co-PI): development of framework, dataset preparation approaches, and material evaluation approaches T.M.P. (Co-PI): development of framework, dataset preparation approaches, and material evaluation approaches B.S.M. (Co-PI): development of framework, architecture design, and training approaches. All authors contributed to the writing of this manuscript with writing efforts being led by N.R.B., D.K.J., and M.P.E.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Correspondence and requests for materials should be addressed to Neal R. Brodnik.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022