

Original Article

Extracting Grain Orientations from EBSD Patterns of Polycrystalline Materials Using Convolutional Neural Networks

Dipendra Jha¹, Saransh Singh², Reda Al-Bahrani¹, Wei-keng Liao¹, Alok Choudhary¹, Marc De Graef² and Ankit Agrawal¹

¹Department of Electrical Engineering and Computer Science, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208, USA and ²Department of Materials Science and Engineering, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA

Abstract

We present a deep learning approach to the indexing of electron backscatter diffraction (EBSD) patterns. We design and implement a deep convolutional neural network architecture to predict crystal orientation from the EBSD patterns. We design a differentiable approximation to the disorientation function between the predicted crystal orientation and the ground truth; the deep learning model optimizes for the mean disorientation error between the predicted crystal orientation and the ground truth using stochastic gradient descent. The deep learning model is trained using 374,852 EBSD patterns of polycrystalline nickel from simulation and evaluated using 1,000 experimental EBSD patterns of polycrystalline nickel. The deep learning model results in a mean disorientation error of 0.548° compared to 0.652° using dictionary based indexing.

Key words: electron backscatter diffraction, EBSD, deep learning, convolutional neural networks

(Received 29 April 2018; revised 25 July 2018; accepted 17 August 2018)

Introduction

Engineering materials are usually crystalline, most often in polycrystalline form. They consist of multiple “grains” having different crystallographic orientations. While the microscopic properties of these grains are anisotropic due to their crystalline nature, the macroscopic properties of the whole crystal depend on the material’s texture—the relative fractions of each of these grain orientations. Texture also provides information about the thermo-mechanical processing history of materials and can be used to reconstruct the conditions leading to the micro-structure, for example, in geological rocks. Thus, texture is paramount in understanding the processing–structure–property relationships.

Since its development in the early 1990s, automated electron backscatter diffraction (EBSD) has become the primary tool to determine the crystal orientation of crystalline materials across a wide variety of material classes (Adams et al., 1993). The technique provides quantitative information about the grain size, grain boundary character, grain orientation, texture and phase identity of the sample by measuring the angular distribution of backscattered electrons using a combination of a scintillator screen and a charge coupled device camera. The schematic of the EBSD setup is shown in Figure 1. The sample sits at a tilt of σ (typically 70°) with the camera tilted at angle θ_c (typically 0–10°). Electrons travel down from the pole piece and interact with the

specimen at point O . The backscattered yield is measured as a function of direction by the scintillator. A physics-based model can be used to predict the backscattered yield based on the principles of quantum mechanics. Assuming the microscope is parametrized by \mathcal{M} , the geometry of the setup is denoted by \mathcal{G} , and the crystal under investigation is parametrized by \mathcal{C} , the forward model is given by

$$\mathcal{F} \equiv \mathcal{F}(\mathcal{M}, \mathcal{G}, \mathcal{C}). \quad (1)$$

Further details of this model can be found in Callahan & De Graef (2013). An example experimental EBSD pattern of iron with its corresponding physics-based simulation is shown in Figures 2a and 2b, respectively.

There are two major techniques to indexing EBSD patterns, each with its advantages and drawbacks. These include the commercially available Hough-transform based approach (Schwartz et al., 2000) and the newly developed dictionary indexing method (Chen et al., 2015; Wright et al., 2015; Marquardt et al., 2017; Singh & De Graef, 2017). The commercially available solution to the indexing problem uses a feature detection algorithm (Krieger Lassen, 1992). A Hough transform of the diffraction is performed to identify linear features (Kikuchi bands). The angles between the extracted linear features are compared to a precomputed look up table to determine the crystal orientation. This method has been very successful in indexing EBSD patterns and has led to significant advances in materials characterization. However, the performance of this method quickly deteriorates in the presence of noise (Wright et al., 2015; Ram et al., 2017).

Author for correspondence: Ankit Agrawal, E-mail: ankitag@eecs.northwestern.edu
Cite this article: Jha D, Singh S, Al-Bahrani R, Liao WK, Choudhary A, De Graef M and Agrawal A (2018) Extracting Grain Orientations from EBSD Patterns of Polycrystalline Materials Using Convolutional Neural Networks. *Microsc Microanal* 24(5), 497–502. doi: 10.1017/S1431927618015131

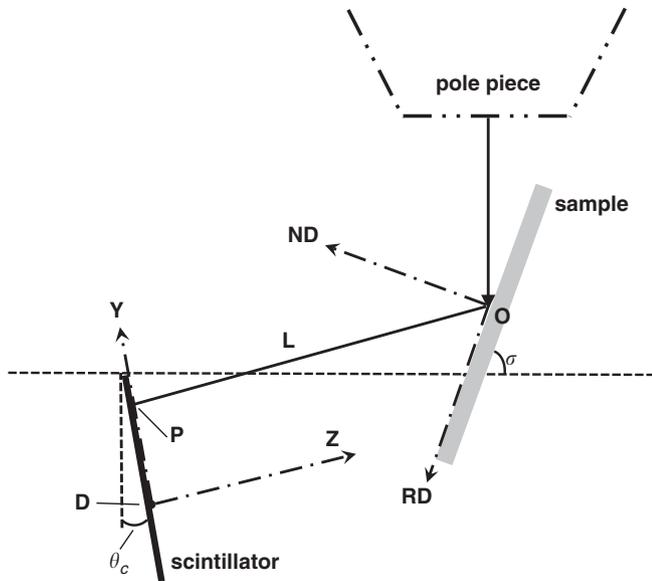


Figure 1. Schematic of the electron backscatter diffraction geometry.

In essence, dictionary based indexing is a nearest neighbor search approach in which the output angles correspond to the orientation angles of the closest EBSD pattern present in dictionary. The distance function used is a dot product as follows:

$$d(\vec{x}_1, \vec{x}_2) = 1 - \frac{\vec{x}_1 \cdot \vec{x}_2}{|\vec{x}_1| |\vec{x}_2|}, \quad (2)$$

where \vec{x} is a vector representing the pixels in the EBSD patterns. The dot product between the pixel intensities in the test sample and each sample in the simulation-based dictionary set is computed and the nearest training sample is used to make the prediction. This method has been shown to be very robust to noise in the diffraction pattern and outperforms the line feature based Hough transform method (Ram et al., 2017) for a wide variety of crystal classes. However, this approach is computationally very expensive, which limits the technique to be an off-line method, and a real time solution to the indexing problem is currently not possible using this approach.

In the present paper, we present a deep learning (LeCun et al., 2015) based model, trained using a simulated diffraction dataset, to predict the crystal orientations for experimental EBSD patterns, such that they have a minimum “disorientation” with respect to their ground truth. Deep learning leverages deep neural networks composed of multiple processing layers to automatically learn the representations of data with multiple levels of abstraction (LeCun et al., 2015). They have achieved great success in the field of computer science with state-of-the-art results in computer vision (Krizhevsky et al., 2012; Szegedy et al., 2017), speech recognition (Mikolov et al., 2011; Deng et al., 2013) and text processing (Sutskever et al., 2014), and are increasingly being used in the relatively nascent field of materials informatics (Agrawal & Choudhary, 2016) for deciphering processing–structure–property relationships.

A convolutional neural network (CNN) is a type of artificial neural network which is composed of convolution layers (LeCun, 2015) in addition to fully connected layers. Since they require minimal preprocessing, they have gained significant attention in fields like computer vision (Krizhevsky et al., 2012; Szegedy et al., 2017), recommender system (Van den Oord et al., 2013) and

natural language processing (Collobert & Weston, 2008). Recently, CNNs have been applied for building models from microstructural data and improving characterization methods (Kondo et al., 2017; Cecen et al., 2018; Ling et al., 2017) and they have been shown to be useful for predicting properties of crystal structures and molecules (Schütt et al., 2018; Wu et al., 2018), detecting cracks in materials/infrastructure images (Gopalakrishnan et al., 2017), and so on. Park et al. (2017) used CNNs for the classification of X-ray diffraction (XRD) patterns in terms of crystal system, extinction group and space group using a large dataset of 150,000 XRD patterns, without any manual feature engineering. Xu & LeBeau (2018) developed CNNs to automatically analyze position averaged convergent beam electron diffraction patterns to extract pattern size, center, rotation, specimen thickness, and specimen tilt, without any need for pre-treating the data. Liu et al. (2016) applied CNNs to learn crystal orientations from simulated EBSD patterns; they built three separate CNN models to individually predict the three Euler angles, but did not take into account the mean disorientation between the predicted and true crystal orientations. Moreover, their models were not tested on experimental data. In this study, our goal is to learn to predict the crystal orientation of experimental EBSD patterns such that they have minimum disorientation with the ground truth.

Building a predictive model for the indexing of EBSD patterns poses two significant challenges. First, we need to minimize the disorientation between the predicted and the ground truth crystal orientations; this requires optimizing for the mean disorientation error which is metric for a highly nonlinear orientation space. Furthermore, this cost function is computationally intensive, making it difficult to manually compute and implement its derivatives with respect to the orientation angles. Therefore, we designed a differentiable approximation to the mean disorientation; it is implemented using TensorFlow (Abadi et al., 2016) and optimized using stochastic gradient descent (Bottou, 1991). The training of the deep learning model was optimized to take advantage of the parallelization available in graphics processing units (GPUs) to process a complete mini-batch.

The second challenge is that the crystal orientation is represented using three Euler angles, which requires learning all three angles simultaneously using a single model, which is different from multi-labeling problems that require predicting different objects present in the input image (Zhou et al., 2008; Briggs et al., 2013; Pham et al., 2015). Most state-of-the-art deep learning architectures are limited to predicting a single output (Krizhevsky et al., 2012; He et al., 2016; Szegedy et al., 2017); existing work on multi-output learning using neural networks has been limited to

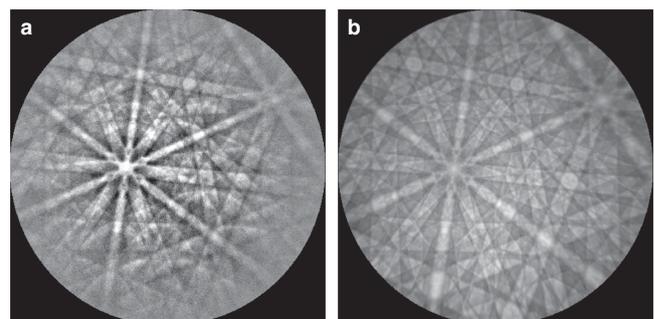


Figure 2. Electron backscatter diffraction pattern from iron (a) experimental and (b) simulation.

shallow feed-forward networks with a single regression layer having multiple outputs (Kang et al., 1996; An et al., 2013; Bezijglov et al., 2016).

We design and implement a novel branched deep CNN optimized for learning multiple outputs; we refer to this model as OMNet. The training and test sets are composed of EBSD patterns of polycrystalline nickel. The simulation dataset used for training the models is composed of 374,852 EBSD patterns. The models are evaluated using a set of 1,000 EBSD patterns from real experiments. The OMNet model outperforms the current dictionary based indexing by 16%, resulting in a mean disorientation of 0.548° compared to 0.652° for the dictionary approach.

Crystal Orientation and Disorientation

The orientation of a crystal is represented by a passive 3D rotation, \mathbf{g} , which maps the specimen’s right-handed Cartesian coordinate frame, $\mathbf{e}^s \equiv (\mathbf{e}_1^s, \mathbf{e}_2^s, \mathbf{e}_3^s)$ onto a right-handed Cartesian coordinate system attached to the crystal, $\mathbf{e}^c \equiv (\mathbf{e}_1^c, \mathbf{e}_2^c, \mathbf{e}_3^c)$, such that $\mathbf{e}_i^c = \mathbf{g}_{ij}\mathbf{e}_j^s$; in this representation, the orientation \mathbf{g} corresponds to a 3×3 special orthogonal matrix, that is, an element of SO(3). There are numerous other representations for orientations, such as the unit quaternion, Rodrigues–Frank vectors, axis-angle pair and cubochoric vector; each with its own distinct properties and advantages. Furthermore, all crystals have certain symmetries associated with them, which lead to degeneracies such that all crystal orientations are not unique. For any crystal, let \mathcal{O}_c represent the set of symmetry operators including the identity operation, with cardinality $\#\mathcal{O}_c = N$. All orientations in the set $\mathcal{O}_c\mathbf{g}$ are equivalent for this crystal symmetry and represent identical orientations.

In the absence of crystal symmetry, the distance metric between two orientations \mathbf{g}_1 and \mathbf{g}_2 is represented by $\mathcal{D}(\mathbf{g}_1, \mathbf{g}_2)$, and is referred to as the misorientation. This metric represents the angle of rotation about some axis to go from one crystal orientation to the other. In this case, the space has simple analytical expressions for the metric tensor as well as smooth and continuous geodesics. The distance metric is given by (assuming \mathbf{g} is in matrix notation)

$$\mathcal{D}(\mathbf{g}_1, \mathbf{g}_2) = \arccos\left(\left(\text{tr}[\mathbf{g}_1^{-1}\mathbf{g}_2] - 1\right) / 2\right). \tag{3}$$

However, in the presence of crystal symmetry, the rotation space becomes degenerate and such an expression is no longer valid. In the presence of crystal symmetry given by the set \mathcal{O}_c , the distance metric referred to as disorientation is given by the following expression (assuming \mathbf{g} and \mathcal{O}_c are both in matrix notation):

$$\mathcal{D}(\mathbf{g}_1, \mathbf{g}_2) = \min_{i,j \in [1,N]} \left\{ \arccos\left(\left(\text{tr}\left[\left(\mathcal{O}_c^i\mathbf{g}_1\right)^{-1}\left(\mathcal{O}_c^j\mathbf{g}_2\right)\right] - 1\right) / 2\right), \right. \\ \left. \arccos\left(\left(\text{tr}\left[\left(\mathcal{O}_c^i\mathbf{g}_2\right)^{-1}\left(\mathcal{O}_c^j\mathbf{g}_1\right)\right] - 1\right) / 2\right) \right\}. \tag{4}$$

This expression gives the minimum angle of rotation, that is, the disorientation, about some axis between any two symmetrically equivalent variants of the two orientations \mathbf{g}_1 and \mathbf{g}_2 .

Deep Learning Approach

The ideal data driven approach for building a predictive model would be to train a machine learning model on the EBSD patterns from experiments. However, experiments are generally expensive

and yield a relatively small number of diffraction patterns; in our case, we have selected 1,000 “experimental” diffraction patterns. Instead, we leverage the ability to simulate realistic EBSD patterns for training, such that the model can predict the crystal orientations for the experimental EBSD patterns with minimum disorientation with respect to the true orientations.

The training and test datasets are composed of EBSD patterns of polycrystalline nickel. The training dataset contains two simulated EBSD pattern dictionaries, one generated with a cubochoric sampling of $N=100$ samples along the cubic semi axis, the other with $N=50$ (see Singh & De Graef (2016) for details). The first dataset has 333,227 patterns, and the second has 41,625 patterns. Combining them, a total of 374,852 patterns were used as the training set without any data augmentation. As the pattern pixel values range from [0, 255], they were rescaled to the range of [0, 1]. The performance of the models was evaluated by indexing 1,000 simulated patterns with known orientations. The disorientation between the predicted and known Euler angles provides the efficacy of the approach. It is important to note that the microscope and diffraction geometry for the training and test set were identical. There exist two main challenges associated with developing such a predictive model. We discuss these challenges along with how we tackle them below.

Optimizing the Mean Disorientation Error

This problem requires optimizing the mean disorientation error between the predicted and the true crystal orientations. This is a challenging task for two reasons: first, the disorientation is the distance metric of a non-Euclidean manifold. In the absence of symmetries in orientation, this metric is easily computed using analytical expressions. However, the presence of crystal symmetries introduces degeneracies in the space resulting in discontinuities in the gradient of the disorientation metric with respect to the input orientations. This renders the disorientation function inappropriate for optimization using stochastic gradient descent for any deep learning model.

Second, the original disorientation algorithm is computationally intensive; for 432 cubic rotational symmetry, it takes one pair of predicted and true Euler angles and computes their 24×24 symmetrically equivalent orientations to find the disorientation (equation (4)). It would be extremely cumbersome to manually compute and implement its derivatives with respect to the input Euler angles. Hence, it is infeasible to train any predictive model using the disorientation function in feasible time.

The disorientation is computed using equation (4). It consists of 1,152 evaluations, each step computing the symmetrically equivalent orientation pairs, followed by two computations to determine the required angle of rotation between them. The disorientation computation contains the $\arccos(x)$ function which is undefined for values outside its domain of $[-1, 1]$. We approximated it by putting an upper bound of 1 to the magnitude of all the values passed to the $\arccos(x)$ function.

We implemented a differentiable approximation of the mean disorientation by building a computational tensor graph using TensorFlow (Abadi et al., 2016). We leveraged its auto-differentiation support for computing the gradients of the mean disorientation error with respect to the Euler angles. The mean disorientation error was optimized by training a deep learning model using the stochastic gradient descent algorithm (Bottou, 1991). When the minibatch size was 64, the sequential algorithm involved $64 \times 1,152 = 73,728$

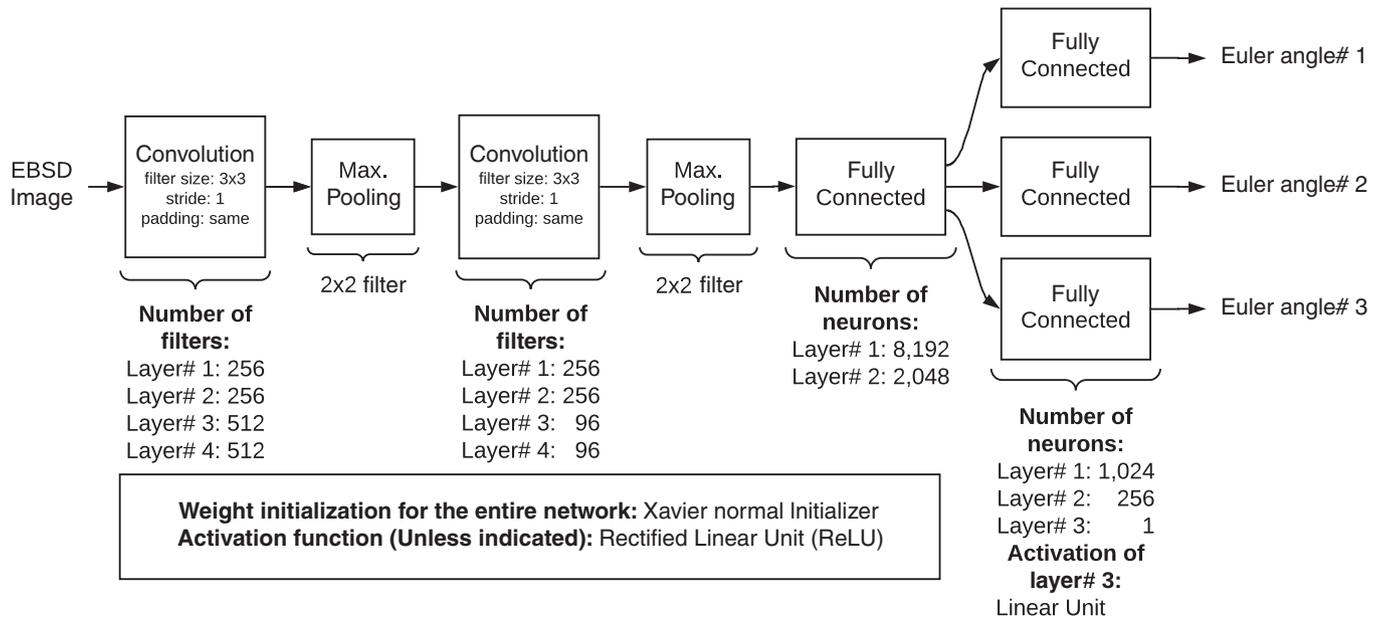


Figure 3. OMNet: CNN architectures for learning multiple outputs.

computations for the misorientation between the symmetrically equivalent predicted and ground truth crystal orientations. It was very costly both in terms of processing time and memory transfer; it took around 24 h to train our model for one epoch using a Titan X GPU with 12 GB memory. This made it impractical to train a deep learning model using the sequential implementation in feasible time. We optimized it to process one mini-batch so that it could leverage the parallelization available in GPUs.

The models are evaluated using the mean disorientation error and mean symmetrically equivalent orientation absolute error (MSEAE). The MSEAE is computed by considering the periodicity of orientation angles as follows:

$$\text{MSEAE}_1(\vec{y}_s, \hat{y}_s) = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^3 | \vec{y}_{sj} - \hat{y}_{sj} |, \quad (5)$$

$$\text{MSEAE}_2(\vec{y}_s, \hat{y}_s) = \text{MSEAE}_1(\vec{y}_s, \hat{y}_s) \bmod (2\pi), \quad (6)$$

$$\text{MSEAE}(\vec{y}_s, \hat{y}_s) = \begin{cases} \text{MSEAE}_2(\vec{y}_s, \hat{y}_s), & \text{if } \text{MSEAE}_2(\vec{y}_s, \hat{y}_s) \leq \pi \\ 2\pi - \text{MSEAE}_2(\vec{y}_s, \hat{y}_s), & \text{else} \end{cases} \quad (7)$$

where \vec{y}_s and \hat{y}_s are Euler angle triplets of the symmetrically equivalent true and predicted orientations with minimum disorientation.

CNN Architectures for Learning Multiple Outputs

Optimizing for the mean disorientation requires learning the crystal orientation angles using a single model such that they can be used to optimize the mean disorientation error. The conventional approach to learning multiple outputs would be to train individual models for learning each output. Since the three Euler angles are correlated with the crystal orientation, we have to learn all the three orientation angles simultaneously using a single model such that they can be leveraged for optimizing for mean

disorientation. Existing multi-output learning using neural networks has been limited to shallow feed-forward networks having a regression layer with multiple outputs (Kang et al., 1996; An et al., 2013; Bezioglov et al., 2016).

We explored several design approaches for learning multiple outputs. Figure 3 demonstrates a novel CNN model architecture—a branched model with individual and independent model components for each output. The first four model components are composed of multiple convolution layers and max pooling. Convolution layers capture the locally correlated features present in the input EBSD patterns; they learn the high level abstract features from the inputs. As the three outputs require similar learning capability, the branched model is composed of three classifier branches containing equal numbers of layers and parameters.

Each branch leverages around two million parameters independent of each other which can be optimized to learn the individual outputs. As the outputs are correlated with each other, they all share the same convolution outputs and the first fully connected component. The convolutional layers are the computationally expensive components; they extract high level features from the inputs that are required for learning all outputs. Sharing the convolutional layers keeps the computational cost comparable to the model having a single regression layer with multiple outputs. The branching technique is currently used in the inception model architectures, but for a different reason (Szegedy et al., 2016). The point where the model starts branching can have a significant impact on the model performance. For the training, we explored branching at different layers, but it was limited by the available GPU memory. Since the model architecture is designed to optimize for learning multiple outputs, we refer to this architecture as OMNet.

All models were implemented using Python and TensorFlow (Abadi et al., 2016). They are trained using Titan X GPUs with 12 GB memory. An extensive search was carried out to tune the hyperparameters, such as learning rate, optimization algorithm, momentum and learning rate decay. We used a batch size of 64 and trained using Adam (Kingma & Ba, 2014) for 100 epochs with a patience of 10. We searched through several CNN model architectures and loss functions—conventional loss functions, followed by mean disorientation error, and finally their hybrids.

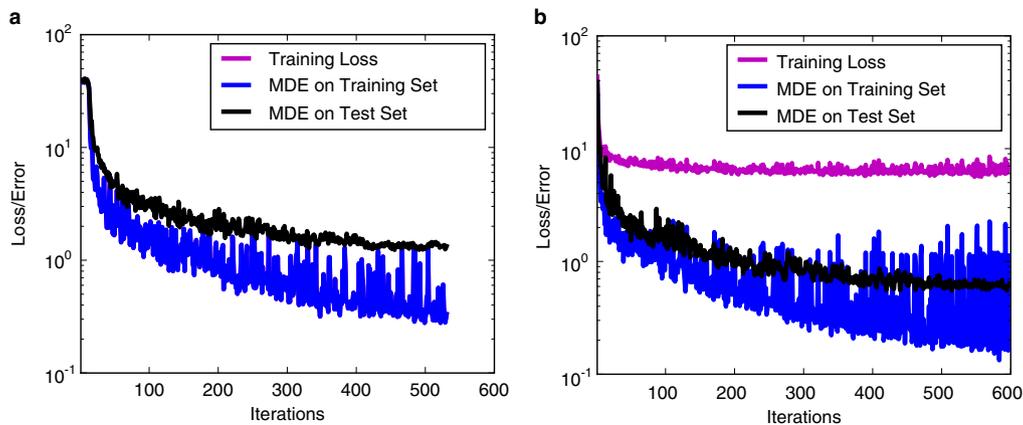


Figure 4. Loss and mean disorientation error using different loss functions. **a:** The training loss and mean disorientation error (MDE) on training set and test set for MDE as the loss function. **b:** The loss and MDE for the hybrid loss function of the sum of mean absolute error and MDE.

Results

We optimized using the mean disorientation error as the loss function as shown in Figure 4a. Since our goal was to optimize for mean disorientation error, we expected the algorithm to result in an improved mean disorientation error on both sets. However, the mean disorientation error alone as the loss function did not perform well; it achieved a mean disorientation error of 1.224° on the test set (experimental EBSD patterns). We observed a lot of oscillations in the training loss curve compared to while using conventional loss functions. This may be due to the mean disorientation error being computed using the symmetrically equivalent orientations rather than the actual outputs (equation (4)). Optimizing for the mean disorientation error increased the model training time by around 30–40%.

Since optimizing for mean disorientation did not perform well, we designed and experimented with several hybrid loss functions, combining the mean disorientation with the conventional loss functions such as mean absolute error (MAE) and mean squared error. We assigned different weights to the constituent losses, even some conditional loss functions that optimized for the Euler

angles first, followed by optimizing for the mean disorientation or a hybrid loss. The best loss function was the sum of the MAE and the mean disorientation error, as shown in Figure 4b and Table 1. The mean disorientation error decreased steadily with time, although the total loss became almost constant. The mean disorientation error on the training set is computed using only the current mini-batch; hence, the mean disorientation error curve for training set has oscillations. The model achieved a mean disorientation error of 0.548° on the experimental EBSD patterns, about 16% better than using dictionary based indexing.

Summary

It is one of the goals of EBSD analysis to obtain grain orientations that are as accurate as possible. Having the ability to determine grain orientations to within a fraction of a degree makes it possible to perform quantitative comparisons between experimental data sets and predictive micro-structure models, in particular for cases in which the material contains large amounts of plastic deformation. Current commercially available EBSD indexing solutions do not perform well when the material is heavily deformed. The recent dictionary indexing approach performs significantly better, but suffers from a high computational cost which makes the approach unfeasible as a real-time indexing solution. The deep learning approach described in this paper has the potential to have a great impact on the field of materials science by providing an indexing approach that is both rapid and accurate, once the training process has been completed.

If deep learning based indexing in real time becomes possible, then this would have a significant impact on the field of materials characterization by providing a faster and more accurate indexing approach than is currently available commercially. While this paper establishes the efficacy of neural networks in learning orientations from EBSD patterns for pristine simulated patterns, the characteristics of such a network in the presence of noise still need to be established. The current model was trained with orientations as the only variable and assuming that the microscope geometry is precisely known; this is seldom the case for real experiments. New strategies leveraging techniques in incremental learning might prove to be useful for training the model for different detector geometry parameters. Finally, the current model can also be extended to other electron diffraction modalities such as the scanning electron microscope based electron channeling patterns and transmission Kikuchi diffraction modality as well as the transmission electron microscope based precession electron diffraction patterns.

Table 1. Mean Disorientation Error (MDE) and Mean Symmetrically Equivalent Orientation Absolute Error (MSEAE) Using Different Models and Loss Functions.

Model	Loss function	Simulation data		Experimental data
		MDE	MDE	MSEAE
Dictionary-based indexing	–	–	0.652°	[0.6592°, 0.3534°, 0.6484°]
Deep learning (OMNet)	MAE	0.064°	0.596°	–
Deep learning (OMNet)	MSE	0.292°	1.285°	–
Deep learning (OMNet)	MDE	0.272°	1.224°	–
Deep learning (OMNet)	MAE + MDE	0.132°	0.548°	[0.7155°, 0.2194°, 0.7066°]
Deep learning (OMNet)	MSE + MDE	0.171°	0.658°	–

MAE, mean absolute error; MSE, mean squared error.

Acknowledgments. This work is supported primarily by the AFOSR MURI award FA9550-12-1-0458. Partial support is also acknowledged from the following grants: NIST award 70NANB14H012; NSF award CCF-1409601; DOE awards DE-SC0007456 and DE-SC0014330.

References

- Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, Corrado GS, Davis A, Dean J, Devin M, Ghemawat S, Goodfellow I, Harp A, Irving G, Isard M, Jia Y, Jozefowicz R, Kaiser L, Kudlur M, Levenberg J, Mane D, Monga R, Moore S, Murray D, Olah C, Schuster M, Shlens J, Steiner B, Sutskever I, Talwar K, Tucker P, Vanhoucke V, Vasudevan V, Viegas F, Vinyals O, Warden P, Wattenberg M, Wicke M, Yu Y and Zheng X (2016) Tensorflow: Large-scale machine learning on heterogeneous distributed systems. arXiv preprint:160304467.
- Adams BL, Wright SI and Kunze K (1993) Orientation imaging: The emergence of a new microscopy. *Metall Trans A Phys Metall Mater Sci* **24**, 819–831.
- Agrawal A and Choudhary A (2016) Perspective: Materials informatics and big data: Realization of the fourth paradigm of science in materials science. *APL Mater* **4**, 053208.
- An N, Zhao W, Wang J, Shang D and Zhao E (2013) Using multi-output feedforward neural network with empirical mode decomposition based signal filtering for electricity demand forecasting. *Energy* **49**, 279–288.
- Beziglov A, Blanton B and Santiago R (2016) Multi-output artificial neural network for storm surge prediction in north carolina. arXiv preprint: 160907378.
- Bottou L (1991) Stochastic gradient learning in neural networks. In *Proceedings of Neuro-Nimes 91, 4th International Conference on Neural Networks and their Applications*. Nanterre, France: EC2.
- Briggs F, Fern XZ and Raich R (2013) Context-aware miml instance annotation. In *2013 IEEE 13th International Conference on Data Mining (ICDM)*, pp. 41–50. Piscataway, NJ: IEEE.
- Callahan PG and De Graef M (2013) Dynamical electron backscatter diffraction patterns. part I: Pattern simulations. *Microsc Microanal* **19**, 1255–1265.
- Cecen A, Dai H, Yabansu YC, Kalidindi SR and Song L (2018) Material structure–property linkages using three-dimensional convolutional neural networks. *Acta Mater* **146**, 76–84.
- Chen YH, Park SU, Wei D, Newstadt G, Jackson MA, Simmons JP, De Graef M and Hero AO (2015) A dictionary approach to electron backscatter diffraction indexing. *Microsc Microanal* **21**, 739–752.
- Collobert R and Weston J (2008) A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th International Conference on Machine Learning*, Lawrence N and Reid M (eds.), pp. 160–167. Helsinki, Finland: PMLR.
- Deng L, Li J, Huang JT, Yao K, Yu D, Seide F, Seltzer M, Zweig G, He X, Williams J, Gong Y and Acero A (2013) Recent advances in deep learning for speech research at Microsoft. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8604–8608. Piscataway, NJ: IEEE.
- Gopalakrishnan K, Khaitan SK, Choudhary A and Agrawal A (2017) Deep convolutional neural networks with transfer learning for computer vision-based data-driven pavement distress detection. *Constr Build Mater* **157**, 322–330.
- He K, Zhang X, Ren S and Sun J (2016) Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778. Piscataway, NJ: IEEE.
- Kang K, Oh JH, Kwon C and Park Y (1996) Generalization in a two-layer neural network with multiple outputs. *Phys Rev E* **54**, 1811–1815.
- Kingma D and Ba J (2014) Adam: A method for stochastic optimization. arXiv preprint:1412.6980.
- Kondo R, Yamakawa S, Masuoka Y, Tajima S and Asahi R (2017) Microstructure recognition using convolutional neural networks for prediction of ionic conductivity in ceramics. *Acta Mater* **141**, 29–38.
- Krieger Lassen N (1992) Automatic crystal orientation determination from ebps. *Micron Microsc Acta* **6**, 191–192.
- Krizhevsky A, Sutskever I and Hinton GE (2012) Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems* **25**, pp. 1097–1105. San Francisco, CA: Morgan Kaufmann Publishers Inc.
- LeCun Y (2015) Lenet-5, convolutional neural networks. <http://yannle.com/exdb/lenet> (retrieved March 15, 2018).
- LeCun Y, Bengio Y and Hinton G (2015) Deep learning. *Nature* **521**, 436–444.
- Ling J, Hutchinson M, Antono E and Decost B (2017) Building data-driven models with microstructural images: Generalization and interpretability. *Mater Dis* **10**, 19–28.
- Liu R, Agrawal A, Liao WK, Choudhary A and De Graef M (2016) Materials discovery: Understanding polycrystals from large-scale electron patterns. In *2016 IEEE International Conference on Big Data (Big Data)*, December 5–8, Washington DC, pp. 2261–2269. Piscataway, NJ: IEEE.
- Marquardt K, De Graef M, Singh S, Marquardt H, Rosenthal A and Koizumi S (2017) Quantitative electron backscatter diffraction (EBSD) data analyses using the dictionary indexing (DI) approach: Overcoming indexing difficulties on geological materials. *Am Mineral* **102**, 1843–1855.
- Mikolov T, Deoras A, Povey D, Burget L and Černocký J (2011) Strategies for training large scale neural network language models. In *2011 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, December 11–15, Waikoloa, HI, pp. 196–201. Piscataway, NJ: IEEE.
- Park WB, Chung J, Jung J, Sohn K, Singh SP, Pyo M, Shin N and Sohn KS (2017) Classification of crystal structure using a convolutional neural network. *IUCrJ* **4**, 486–494.
- Pham A, Raich R, Fern X and Arriaga JP (2015) Multi-instance multi-label learning in the presence of novel class instances. In *International Conference on Machine Learning*, Vol. 37, Lawrence N and Reid M (eds.), pp. 2427–2435. Lille, France: PMLR.
- Ram F, Wright S, Singh S and Graef MD (2017) Error analysis of the crystal orientations obtained by the dictionary approach to EBSD indexing. *Ultramicroscopy* **181**, 17–26.
- Schütt KT, Sauceda HE, Kindermans PJ, Tkatchenko A and Müller KR (2018) Schnet – A deep learning architecture for molecules and materials. *J Chem Phys* **148**, 241722.
- Schwartz A, Kumar M, Adams B and Field D (eds.) (2000) *Electron Backscatter Diffraction in Materials Science*, 2nd ed. New York, NY: Springer.
- Singh S and De Graef M (2016) Orientation sampling for dictionary-based diffraction pattern indexing methods. *Model Simul Mater Sci Eng* **24**, 085024.
- Singh S and De Graef M (2017) Dictionary indexing of electron channeling patterns. *Microsc Microanal* **23**, 1–10.
- Sutskever I, Vinyals O and Le QV (2014) Sequence to sequence learning with neural networks. In *Advances in neural information processing systems* **27**, pp. 3104–3112. San Francisco, CA: Morgan Kaufmann Publishers Inc.
- Szegedy C, Ioffe S, Vanhoucke V and Alemi AA (2017) Inception-v4, inception-resnet and the impact of residual connections on learning. *AAAI* **4**, 12.
- Szegedy C, Vanhoucke V, Ioffe S, Shlens J and Wojna Z (2016) Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2818–2826. Piscataway, NJ: IEEE.
- Van den Oord A, Dieleman S and Schrauwen B (2013) Deep content-based music recommendation. In *Advances in neural information processing systems* **26**, pp. 2643–2651. San Francisco, CA: Morgan Kaufmann Publishers Inc.
- Wright SI, Nowell MM, Lindeman SP, Camus PP, Graef MD and Jackson MA (2015) Introduction and comparison of new EBSD post-processing methodologies. *Ultramicroscopy* **159**, 81–94.
- Wu Z, Ramsundar B, Feinberg EN, Gomes J, Geniesse C, Pappu AS, Leswing K and Pande V (2018) MoleculeNet: A benchmark for molecular machine learning. *Chem Sci* **9**, 513–530.
- Xu W and LeBeau JM (2018) A deep convolutional neural network to analyze position averaged convergent beam electron diffraction patterns. *Ultramicroscopy* **188**, 59–69.
- Zhou ZH, Zhang ML, Huang SJ and Li YF (2008) MIML: A framework for learning with ambiguous objects. *CORR abs/0808.3231*, 112.