

EBSD IMAGE SEGMENTATION USING A PHYSICS-BASED FORWARD MODEL

Se Un Park*, Dennis Wei*, Marc De Graef†, Megna Shah‡, Jeff Simmons‡, Alfred O. Hero*

*University of Michigan

†Carnegie Mellon University

‡Air Force Research Laboratory

ABSTRACT

We propose a segmentation and anomaly detection method for electron backscatter diffraction (EBSD) images. In contrast to conventional methods that require Euler angles to be extracted from diffraction patterns, the proposed method operates on the patterns directly. We use a forward model implemented as a dictionary of diffraction patterns generated by a detailed physics-based simulation of EBSD. The combination of full diffraction patterns and a dictionary allows anomalies to be detected at the same time as grains are segmented, and also increases robustness to noise and instrument blur. The proposed method is demonstrated on a sample of the Ni-base alloy IN100.

Index Terms— Image Segmentation, Dictionary Learning, Electron Backscatter Diffraction (EBSD), Pattern Matching, Materials Science

1. INTRODUCTION

Engineered polycrystalline materials find uses in a variety of aerospace applications. These polycrystalline materials are composed of grains, microscopic domains of approximately uniform crystal orientation (see Fig. 7 for sectional views). To examine this microstructure after the material has been synthesized, optical and electron microscopes (both scanning and transmission) are commonly used. In this paper, we focus on electron backscatter diffraction (EBSD), which offers the ability to measure local crystal orientation and thus identify grains and grain boundaries [1].

In EBSD, a beam of electrons is directed onto a tilted sample of the material. The electrons are scattered by atoms in different crystal planes and are diffracted, generating a characteristic pattern of intersecting lines on a fluorescent screen as seen in Fig. 1. In conventional EBSD analysis, the orientations and widths of selected diffraction lines are measured and matched against a pre-computed index for the material under study. The indexing yields the crystal orientation, commonly described by three Euler angles with respect to a reference frame, for the volume illuminated by the beam. By repeating the process on a grid of points in the sample, an orientation map or image is produced. The image is then segmented into grains by thresholding the Euler angles. The EBSD measurements can also be fused with other modalities in multimodal analyses [2, 3].

In contrast to the indexing approach described above, in this work we use the full diffraction patterns to perform segmentation. This allows us to exploit information that is normally discarded in the indexing process, in particular the characteristics of the background intensity profile. With this additional information, grain boundaries and anomalous points can be detected as explicit classes




Fig. 1. EBSD system. An electron beam is scattered by atoms in different crystal planes in the sample, producing a diffraction pattern on the detector screen. Image taken from [4].

at the same time as grains are segmented. Anomaly detection is an important capability since anomalies may correspond to defects or contaminants in the material that affect its properties. In addition, processing the diffraction patterns directly avoids problems with indexing when the observed pattern is a poor match to the dictionary, as occurs for example at grain boundaries.

Image segmentation techniques have been proposed from various perspectives including clustering methods [5], region-growing methods [6], level set methods [7, 8], and a recently proposed inverse diffusion and expectation-maximization algorithm for materials images [9]. The present work is distinguished by our use of a detailed physics-based model for the material, the EBSD interaction, and the experimental setup. Given this model, diffraction patterns are simulated for a large set of crystal orientations and are collected in a dictionary. Segmentation is done using features that measure similarity between observed diffraction patterns and elements in the dictionary. To the best of our knowledge, this type of dictionary-based segmentation using a physical forward model has not been considered for EBSD images. The advantage of our approach is its greater robustness to instrument blur and noise and its ability to detect anomalies not represented in the dictionary. At the same time, the large size of both the dictionary and the experimental dataset and the high dimension of the diffraction patterns necessitate the use of highly efficient computational methods.

In Section 2, we describe the construction of the diffraction pattern dictionary and the associated observation model. In Section 3, the proposed segmentation and anomaly detection method is discussed. The results of applying the method to a sample of the Ni-base alloy IN100 are presented in Section 4.

2. DICTIONARY-BASED FORWARD MODEL

2.1. Construction of Diffraction Pattern Dictionary

We simulate diffraction patterns for a set of 281700 finely-sampled points in orientation (Euler angle) space using a computational

This work was partially supported under the Electronic Imaging component of the ICMD Program at Wright-Patterson AFB, Andrew Rosenberger, PM, under UES subcontract S-875-040-008 and USAF/AFMC grant FA8650-9-D-5037/04.




Fig. 2. Raw diffraction patterns for four representative pixels: (a) grain interior, (b) grain boundary, (c) background-shifted, and (d) high noise.

physics-based forward model. **The quality of this model is demonstrated in [10], where it produces simulated patterns close to experimental patterns.** For each orientation, the process consists of three steps. The first step models the interaction of the electron beam with the sample using the Schrödinger equation with a Bloch wave ansatz [11]. The backscattered electron yield is calculated for a set of directions covering the hemisphere of all possible exit directions. The second step is to interpolate the intensities over the hemisphere onto the pixel locations on the collecting detector. We make use of an equal-area projection of the hemisphere onto a square or hexagonal grid [12] to permit standard bilinear interpolation of the intensities. The third step is to model additional instrument effects such as detector quantum efficiency, Poisson noise, coupling optics, and the point spread function and binning mode of the CCD. The final result is a dictionary of 281700 diffraction patterns, which we represent as a 4800×281700 matrix \mathbf{M} with each column \mathbf{m}_i corresponding to a vectorized 80×60 diffraction pattern.

2.2. Observation model

We now develop a model for experimentally observed diffraction patterns based on the dictionary in Section 2.1. Figs. 2(a) and 2(b) show that typical diffraction patterns consist of intersecting lines superimposed on a background intensity profile. Since the background remains approximately constant as a function of position on the sample, we may express each observed pattern \mathbf{y}_j as

$$\mathbf{y}_j = \alpha_j \mathbf{b}_Y + \tilde{\mathbf{y}}_j, \quad (1)$$

where \mathbf{b}_Y is the common background component, α_j is a scale factor, and $\tilde{\mathbf{y}}_j$ is orthogonal to \mathbf{b}_Y . A similar decomposition $\mathbf{m}_i = \beta_i \mathbf{b}_M + \tilde{\mathbf{m}}_i$ holds for the elements of the dictionary.

Given the orthogonal decomposition in (1) and its counterpart for the dictionary, we postulate the following linear model relating background-subtracted patterns $\tilde{\mathbf{y}}_j$ and $\tilde{\mathbf{m}}_i$:

$$\tilde{\mathbf{y}}_j = \sum_i \tilde{\mathbf{m}}_i a_{ij} + \mathbf{n}_j + \mathbf{z}_j, \quad (2)$$

where \mathbf{n}_j represents measurement noise and \mathbf{z}_j is a gross error representing anomalous patterns that is present only for a small number of indices j . For a pattern $\tilde{\mathbf{y}}_j$ corresponding to a single crystal orientation in the interior of a grain, the number of significant coefficients a_{ij} is expected to be small, corresponding to nearby orientations. For grain boundaries, a cluster of significant coefficients is expected for each adjacent grain. Thus it can be assumed that for non-anomalous patterns $\tilde{\mathbf{y}}_j$, the coefficients $\{a_{ij}\}$ are sparse.

Under the model in (2), an observed pattern \mathbf{y}_j can be characterized by estimating the coefficients $\{a_{ij}\}$, for example using sparse linear regression [13]. However, due to the high dimension of the

dictionary \mathbf{M} (4800×281700), such methods are difficult to implement. In this paper, we limit ourselves to the simpler task of identifying the dictionary elements that provide a close representation of a given observed pattern. This is accomplished by computing normalized inner products with dictionary elements as detailed in Section 3. Section 4 demonstrates the feasibility of this inner product approach for image segmentation. We note that if $\mathbf{n}_j \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\mathbf{z}_j = \mathbf{0}$, and $\tilde{\mathbf{y}}_j$ is known to be well-approximated by a single dictionary element, then inner products are sufficient to identify the most likely dictionary element and estimate the value of the single nonzero a_{ij} . These conditions are approximately satisfied for grain interiors assuming that the dictionary is finely sampled.

3. PROPOSED METHOD FOR SEGMENTATION

In this section, we discuss our method for segmenting experimental diffraction patterns into four classes as shown in Fig. 2. These classes correspond to grain interiors, grain boundaries, and two types of anomalies, one with a shifted background pattern and the other where the background remains roughly centered but the pattern is very noisy. We adopt a decision tree (DT) framework [14] to reorganize this multiple hypothesis test into a hierarchy of binary tests performed at successive nodes of a binary tree. We describe in Sections 3.1 and 3.2 the two dictionary-derived features used in the classifier before presenting the DT in Section 3.3. Both features are based on inner products for computational efficiency.

3.1. Background similarity

The first feature measures the similarity of an observed diffraction pattern to patterns in the dictionary with the background components included. Since the background typically dominates the total energy in the patterns (more than 99.7% in both our dictionary and experimental dataset), this feature essentially measures the similarity between background profiles. We compute the normalized inner product

$$\rho_{ij} = \frac{\mathbf{m}_i^T \mathbf{y}_j}{\|\mathbf{m}_i\| \|\mathbf{y}_j\|} \quad (3)$$

between a given experimental pattern \mathbf{y}_j and all of the dictionary patterns \mathbf{m}_i . We then take the mean inner product

$$\bar{\rho}_j = \frac{1}{M} \sum_{i=1}^M \rho_{ij} \quad (4)$$

to be the measure of background similarity, where $M = 281700$.

3.2. Neighborhood similarity

The second feature is derived from inner products computed after the background components have been removed from both the observations and the dictionary. Toward this end, we use the mean $\bar{\mathbf{y}}$ of the observed patterns as an estimate for the observation background \mathbf{b}_Y , and similarly for the dictionary. The first principal component also yields a very similar background estimate. We then compute the background-subtracted patterns $\tilde{\mathbf{y}}_j$ from (1), where the scale factor α_j is found to be $\mathbf{y}_j^T \mathbf{b}_Y / \mathbf{b}_Y^T \mathbf{b}_Y$. The same is done for the dictionary to determine $\tilde{\mathbf{m}}_i$.

Figures 3(a) and 3(b) show the result of removing the background from the experimental patterns in Figs. 2(a) and 2(b). The diffraction lines that characterize crystal orientation are preserved. Fig. 3(c) shows the background-subtracted dictionary element with the largest normalized inner product with the patterns in Fig. 3(a) and




Fig. 3. (a)(b) Background-compensated patterns corresponding to Figs. 2(a) and 2(b). (c) Best-match dictionary element for both (a) and (b).




Fig. 4. Normalized inner products sorted by decreasing absolute value for four representative pixels - interior, boundary, background-shifted, and high noise. (left: the 100 largest values, right: the 2000 largest values) The initial part of the curves indicates that interior and edge pixels are highly correlated with several dictionary entries, after which the curves steeply decrease. After the 100 largest values, the curves corresponding to the interior and edge pixels appear to converge to a low asymptote, whereas the curves for noisy and background-shifted pixels do not decay much.

3(b), which happens to be the same in both cases. The correspondence between diffraction lines suggests that close matches in orientation space can be found by considering background-compensated normalized inner products.

Based on Fig. 3 and the model in (2), we represent each observed pattern with the dictionary elements having the highest normalized inner products with the observed pattern. Denote by I_j the indices of the dictionary elements chosen for the observation at pixel j ; I_j can be regarded as a proxy for the set of significant coefficients a_{ij} in (2). It is clear from Fig. 3 that more than one representative is required to differentiate between interior patterns (Fig. 3(a)) and boundary patterns (Fig. 3(b)) since the single best matches can be identical. In addition, using multiple representatives improves robustness against noise and discretization of the orientation space.

To determine the number of dictionary elements required, we examine in Fig. 4 the decay of the sorted normalized inner products for the four different types of patterns. Fig. 4 suggests that a few tens of dictionary elements are sufficient to capture the nearest neighbors for interior and boundary patterns. For concreteness, we choose the number of representatives $|I_j|$ to be 40.

Given the sets I_j , we define a dictionary-based measure of similarity between a pixel j and its neighbors in a 3×3 neighborhood \mathcal{N}_j centered at but not including j . Specifically, the similarity between patterns at pixels j and $k \in \mathcal{N}_j$ is given by $s_{jk} = |I_j \cap I_k|$, and the average similarity of pixel j with its neighbors is

$$s_j = \frac{\sum_{k \in \mathcal{N}_j} s_{jk}}{|\mathcal{N}_j|}. \quad (5)$$

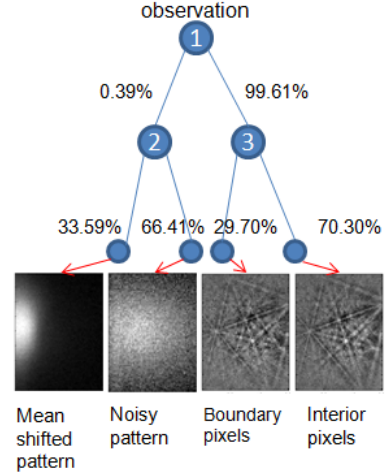


Fig. 5. Decision tree. The background similarity criterion (4) is used at nodes 1 and 2 while the neighborhood similarity criterion (5) is used at node 3. The division of the population at each parent node is shown above the branches. Representative diffraction patterns are also shown.

3.3. Decision Tree

Our proposed DT classifier is illustrated in Fig. 5. We first test at node 1 the binary hypothesis that a pixel is an anomaly or not by thresholding the mean uncompensated inner product $\bar{\rho}_j$ in (4), because anomalous pixels have low $\bar{\rho}_j$. If the test is passed then we test at node 2 whether the pixel is a background-shifted anomaly or high-noise anomaly using a second threshold on $\bar{\rho}_j$. The shift in background implies that the value of $\bar{\rho}_j$ is significantly lower than for the other anomaly. Finally if the first hypothesis test fails, we test between interior pixels and boundary pixels at node 3 using the neighborhood intersection similarity measure (5), which is computed from background-compensated diffraction patterns. Interior pixels have high similarity values because they share many common neighbors in the dictionary with adjacent pixels. On the other hand, boundary pixels can be viewed as mixtures of adjacent interior pixels and therefore have lower similarities because the set of nearest neighbors is only partially shared with those on either side of the boundary.

The only tuning parameters in our DT classifier are the decision thresholds on the similarity measures used at each node of the tree. When there is ground truth data available these thresholds can be selected by cross-validated minimization of nodal deviance measures such as mis-classification rate, Gini index, or entropy [15]. **Parameter selection is generally more difficult however when the learning is unsupervised as in our case.** Our application of DT classifiers differs from previous imaging applications, e.g., land cover classification in remote sensing [16, 17], in several important ways. First, our classifier is a hybrid DT in that it uses different features (background-compensated and uncompensated patterns) and similarity measures (inner products and neighborhood intersections). Second, unlike standard non-parametric DT classifiers, our DT is informed by a physics model through the generated dictionary of diffraction patterns. As a consequence, the tree structure of our DT classifier is fixed and need not be learned from the data.




Fig. 6. (a) Empirical histogram of the background similarity measure $\bar{\rho}_j$ between 0.99 and 1. The threshold t_2 at node 2 in Fig. 5 is set to 0.996. (b) Empirical histogram of the neighborhood similarity measure s_j . The threshold at node 3 in Fig. 5 is set to the 0.3 quantile at 31.25.

4. RESULTS

We apply the method described in Section 3 to EBSD data obtained from a section of the Ni-base alloy IN100. The dimensions of the section are 512 x 384 pixels. The diffraction pattern corresponding to each pixel in the sample is of size 80 x 60 pixels.

In this example, ground truth class labels are not available. Thus to set the thresholds in the DT classifier in Fig. 5, we consider the empirical distributions (obtained from the experimental sample) of the background similarity $\bar{\rho}_j$ (4) and the neighborhood similarity s_j (5). The distribution of $\bar{\rho}_j$ can be divided into three parts by inspection: a uniform distribution with a small number of samples over $[\rho_{\min}, t_1]$, an increasing tail over $[t_1, t_2]$, and a large peak over $[t_2, 1]$. We choose the first threshold $t_1 = 0.815$ to separate background-shifted patterns at node 1. From Fig. 6(a), which shows the histogram of $\bar{\rho}_j$ between 0.99 and 1, we select $t_2 = 0.996$ as the threshold to classify high-noise patterns at node 2. Lastly, at node 3 we use $|I_j| = 40$ dictionary representatives and consider the empirical histogram of the neighborhood similarity s_j shown in Fig. 6(b). We set the threshold at the 0.3 quantile at 31.25, which corresponds to the knee in the histogram. **Choosing a threshold by fitting a 2-component Gaussian mixture model to Fig. 6(b) yields a very similar value.** The percentages in Fig. 5 indicate the division of the parent population at each node resulting from the above thresholds. Note that the percentages below node 3 are not 30% and 70% exactly because the neighborhood similarities are quantized.

In Fig. 7(a), we show a neighborhood similarity map computed according to (5) for $|I_j| = 40$. As expected, interior pixels have larger similarity values while boundary pixels and anomalies have smaller similarities since they are less likely to share common neighbors in the dictionary with spatially adjacent pixels. Similarity maps evaluated with different numbers of nearest neighbors, $|I_j| = 20, 40, 200,$ and 2000 , yield equally good performance visually but are not presented due to limited space. For comparison, we show a similarity map in Fig. 7(b) computed using only the observed patterns, specifically by averaging the uncompensated inner products between a pattern at a given pixel and patterns at adjacent pixels. Comparing Figs. 7(b) and 7(a), it can be seen that the map in Fig. 7(b) has blurrier edges near the upper boundary of the sample. This suggests that the dictionary-based approach has a denoising effect that results in sharper segmentations.

Fig. 7(c) shows the segmented image produced by our DT classifier using the thresholds specified above. In terms of segmenting grains, our result is consistent with the image in Fig. 7(d) from a




Fig. 7. (a) Neighborhood similarity map with $|I_j| = 40$. The grayscale runs between the 0.05 and 0.95 quantiles of the similarity values to suppress the visual effect of outliers. (b) Similarity map derived from observed patterns only. The upper parts of the sample are blurry and the boundary structure is not as clear as in the dictionary-based map (a). (c) Segmentation result from our decision tree in Fig. 5. Grain interiors in white, boundaries in black, noisy pixels in red, and background-shifted pixels in blue. (d) Segmentation result from DREAM.3D. Black clusters represent pixels that cannot be classified.

conventional segmentation algorithm that uses Euler angles, specifically as implemented in the software DREAM.3D [18]. However, the proposed method also identifies anomalies (colored red and blue in Fig. 7(c)) that are either not detected (black clusters) or misclassified as one or more grains in Fig. 7(d). This is due to the fact that the conventional segmentation algorithm is not aware of anomalies and has difficulty extracting meaningful Euler angles from anomalous diffraction patterns, as indicated by the black clusters.

5. CONCLUSION AND FUTURE WORK

We have proposed a segmentation and anomaly detection method for EBSD images that is novel in two respects: the use of full diffraction patterns instead of Euler angles or equivalent representations of orientation, and the construction of a dictionary of reference diffraction patterns through a physics-based forward model. Our decision tree classifier uses features that measure similarity between observed patterns and dictionary elements and is efficient enough to handle large sets of high-dimensional patterns. The method is able to not only segment grain interiors from boundaries but also discover and identify two types of anomalies.

Future work will focus on estimating the coefficients a_{ij} in the observation model (2) in order to characterize grain orientations and super-resolve boundaries.

The authors thank Lawrence Drummy, Michael Uchic and Michael Groeber of AFRL and Mike Jackson of UES for their helpful comments as this research evolved.

6. REFERENCES

- [1] A.J. Schwartz, M. Kumar, B.L. Adams, and D.P. Field, *Electron Backscatter Diffraction in Materials Science*, Springer, 2009.
- [2] H. Nordmark, M. Di Sabatino, M. Acciarri, J. Libal, S. Binetti, E. J. Ovrelid, J. C. Walmsley, and R. Holmestad, “EBIC, EBSD and TEM study of grain boundaries in multicrystalline silicon cast from metallurgical feedstock,” in *Photovoltaic Specialists Conference, 2008. PVSC '08. 33rd IEEE*, may 2008, pp. 1–6.
- [3] Yun Wang, H. Kimura, Y. Akiniwa, and K. Tanaka, “Ebsd-afm hybrid analysis of fatigue slip system and crack initiation in polycrystalline metal under cyclic torsional loading,” in *Micro-NanoMechatronics and Human Science, 2005 IEEE International Symposium on*, nov. 2005, pp. 223–228.
- [4] Center for Nanoscience and Nanotechnology, Hebrew University of Jerusalem, “High Resolution Scanning Electron Microscope Sirion: Basics and Tutorials,” http://nanoscience.huji.ac.il/unc/sem_basics.htm, 2013, accessed February 5, 2013.
- [5] Z. Wu and R. Leahy, “An optimal graph theoretic approach to data clustering: theory and its application to image segmentation,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 15, no. 11, pp. 1101–1113, nov 1993.
- [6] L.G. Shapiro and G.C. Stockman, *Computer vision*, Prentice Hall, 2001.
- [7] Stanley Osher and Nikos Paragios, *Geometric Level Set Methods in Imaging, Vision, and Graphics*, Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2003.
- [8] T.F. Chan and L.A. Vese, “Active contours without edges,” *Image Processing, IEEE Transactions on*, vol. 10, no. 2, pp. 266–277, feb 2001.
- [9] Hsiao-Chiang Chuang, L.M. Huffman, M.L. Comer, J.P. Simmons, and I. Pollak, “An automated segmentation for nickel-based superalloy,” in *Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on*, oct. 2008, pp. 2280–2283.
- [10] P. G. Callahan and M. De Graef, “Dynamical EBSD patterns: I. pattern simulations,” *Microsc. Microanal.*, 2013, in press.
- [11] A. Winkelmann, C. Trager-Cowan, F. Sweeney, A. P. Day, and P. Parbrook, “Many-beam dynamical simulation of electron backscatter diffraction patterns,” *Ultramicroscopy*, vol. 107, pp. 414–421, 2007.
- [12] D. Roşca, “New uniform grids on the sphere,” *Astronomy and Astrophysics*, vol. 520, no. A63, pp. 1–4, Oct. 2010.
- [13] R. Tibshirani, “Regression shrinkage and selection via the Lasso,” *J. R. Statist. Soc. B*, vol. 58, no. 1, pp. 267–288, 1996.
- [14] L. Breiman, J. Friedman, C. J. Stone, and R. A. Olshen, *Classification and regression trees*, Chapman & Hall/CRC, 1984.
- [15] J. Friedman, T. Hastie, and R. Tibshirani, *The Elements of Statistical Learning*, Springer Series in Statistics, 2009.
- [16] M. A. Friedl and C. E. Brodley, “Decision tree classification of land cover from remotely sensed data,” *Remote Sensing of Environment*, vol. 61, no. 3, pp. 399–409, 1997.
- [17] M. Pal and P. M. Mather, “An assessment of the effectiveness of decision tree methods for land cover classification,” *Remote Sensing of Environment*, vol. 86, no. 4, pp. 554–565, 2003.
- [18] “DREAM.3D: Digital Representation Environment for Analyzing Microstructure in 3D,” <http://dream3d.bluequartz.net>, 2013, accessed February 5, 2013.