A Robust Algorithm to Calculate Parent β Grain Shapes and Orientations From α Phase Electron Backscatter Diffraction Data in α/β-Titanium Alloys

Alexander Zaitzeff [1], Adam Pilchak[2], Tracy Berman[3], John Allison[3], Selim Esedoglu[1]

Abstract

This paper presents a new model and open-source algorithm for reconstructing prior β phase orientation and grain shapes from measured α-phase electron backscatter diffraction data in α/β titanium alloys. It is based on the image segmentation model of Mumford and Shah, which includes a regularization term that helps create smooth boundaries and overcomes shortcomings of prior reconstruction techniques. Additionally, the new algorithm is resilient to noise. Our algorithm's effectiveness is demonstrated on simulated and real world data.

Keywords: Image Analysis, Phase Transformations, Grain Boundaries

In titanium and titanium alloys, the crystal structure undergoes phase transformation from cubic (bcc) β phase at elevated temperature to hexagonal close packed (hcp) α phase at low temperature.

The distribution of parent β grain sizes and shapes are important with respect to the orientations and sizes of the α colonies that form from those parent grains during cooling from above the β transus which, in turn, affect quasi-static, fatigue, and fracture properties. The problem of inferring the parent β grains, or prior-β grains, from observations of α phase orientations is of great interest because titanium alloys have extensive use within the aerospace industry, and the characteristics of these parent β grains cannot be easily measured directly. Most earlier approaches have focused on the crystallographic relationship between the parent β phase and the possible α orientations formed within it (Burgers orientation relationship, viz., (0001) α || $\{110\}$ β and $\langle 11-20\rangle$ α || $\langle 111\rangle$ β) as well as the misorientations between the 12 possible α variants originating from the same parent β grain. These methods are generally successful,

[1]Department of Mathematics, University of Michigan, Ann Arbor, MI 48109, USA. Corresponding author: Alexander Zaitzeff azaitzef@umich.edu.

[2]Air Force Research Laboratory, Materials and Manufacturing Directorate (AFRL/RXCM), Wright-Patterson Air Force Base, OH 45433 USA

[3]Department of Materials Science and Engineering, University of Michigan, Ann Arbor, MI 48109, USA.

but are subject to issues including the 111 mirror plane that is not included in the rotational symmetry operators as discussed by Cayron [1]. Additionally, in certain instances two colonies on opposite sides of a parent β grain boundary may share one of the five special misorientations between two α colonies with the same parent β grain such that the fortuitously aligned colony can get assigned to its neighboring β grain during reconstruction resulting in the prediction of nonphysical grain shapes and incorrect textures.

In this paper, we propose an automated method for recovering the parent β grain microstructure from α-phase orientations measured with electron backscatter diffraction (EBSD). Our method treats reconstruction as a computer vision image segmentation problem and seeks to recognize regions occupied by distinct objects in a given field of view. The new method is based on the well-known, variational model of Mumford-Shah for image segmentation [2]. In fact, this famous model from computer vision literature was inspired by materials science variational models of grain boundary motion. As a result, it is particularly well-suited to the reconstruction task because it favors physically probable grain shapes.

In this work, we focus on EBSD measurements of the important and widely used titanium alloy, Ti-6Al-4V. To our knowledge, this is the first time a variational image segmentation model has been developed specifically for the β grain reconstruction problem. Previous methods rely on measuring the misorientation angle of potential parent β grains of neighboring α colonies (which our method does not have to find beforehand) [3, 4, 5], or individual voxels (The flood fill algorithm [6]). These methods can produce physically improbable grain shapes. In contrast, BetaTD favors physically reasonable grain shapes. As such, and with its optimization based approach, it also differs from earlier works e.g. [7, 8, 9, 10] that employ a form of region growing. BetaTD is available as an open-source code at `https://github.com/AZaitzeff/betaTD`.

Let Ω denote the region in the plane occupied by the titanium alloy, and let $\Sigma_1, \Sigma_2, \ldots, \Sigma_n$ denote the (unknown) regions occupied by the distinct, parent β grains. Let $\beta_1, \beta_2, \ldots, \beta_n$ denote the (also unknown) orientations of those parent β grains. The Σ_i should occupy the entire material Ω, and may border each other but otherwise cannot overlap:

$$\bigcup_{i=1}^{n} \Sigma_i = \Omega, \text{ and } \Sigma_i \cap \Sigma_j = (\partial\Sigma_i) \cap (\partial\Sigma_j) \text{ if } i \neq j. \tag{1}$$

We will consider situations where there may be partial information available from experimental measurements on the α and β phases. To that end, let $\mathcal{A} \subset \Omega$ be the subset of the domain Ω that contains α grain orientation measurements and $\mathcal{B} \subset \Omega$ be the subset that contains β phase orientation measurements.

In order to segment the domain Ω into regions containing distinct β phase grains, we minimize the following energy, denoted E below, which is a variant of the piecewise constant Mumford-Shah energy [2] adapted to the β reconstruction problem, subject to the constraint Eq. (1):

$$E(\Sigma_1,\ldots,\Sigma_n;\beta_1,\ldots,\beta_n) = \sum_{i=1}^{n} \mathrm{Per}(\Sigma_i,\Omega) + \lambda\bigg(\int_{\Sigma_i\cap\mathcal{A}} W(\mathbf{x}) d_\alpha(\beta_i, f_\alpha(\mathbf{x}))d\mathbf{x} + \int_{\Sigma_i\cap\mathcal{B}} W(\mathbf{x}) d_\beta(\beta_i, f_\beta(\mathbf{x}))d\mathbf{x}\bigg). \quad (2)$$

Here, $\mathrm{Per}(\Sigma_i,\Omega)$ denotes the perimeter (length of the boundary) of the i-th phase Σ_i in the entire computational domain Ω. $\mathcal{A}$ and $\mathcal{B}$ are subsets of the domain Ω where we have α or β orientation data available, respectively, given by the functions $f_\alpha(\mathbf{x}) : \mathcal{A} \to SO(3)$ and $f_\beta(\mathbf{x}) : \mathcal{B} \to SO(3)$. $W(\mathbf{x})$ is a weight function on the data. The function $d_\alpha(\cdot,\cdot)$ is the misorientation angle between a β orientation and the closest parent of an α orientation and $d_\beta(\cdot,\cdot)$ is the misorientation angle between two β orientations. These are given by

$$d_\alpha(\beta',\alpha') = \min_{P\in\mathcal{G},T\in\mathcal{T}} \arccos\left(\frac{\mathrm{trace}\big(P\beta'(T\alpha')^{-1}\big)-1}{2}\right)$$

$$d_\beta(\beta_1,\beta_2) = \min_{P\in\mathcal{G}} \arccos\left(\frac{\mathrm{trace}\big(P\beta_1(\beta_2)^{-1}\big)-1}{2}\right)$$

for the cubic symmetry group $\mathcal{G}$ and α-phase to β-phase transformations $\mathcal{T}$. The first term in the energy, often called the *regularization*, controls the expected roughness of the boundary, encouraging shortest, simplest, most regular curves consistent with the measurements, as well as 120° angles at triple junctions. The second term, often called the *fidelity term*, encourages consistency with the measurements (i.e. observations, experimental data). Ours is a weighted one norm penalty on the error, for robustness to outliers. The sole parameter in the model, λ, sets the scale (level of detail) expected in the segmentation, and can be related to e.g. the number of grains expected, and to the level of noise. At the moment this set by the user. An essential benefit of a continuum model such as (2) is resolution independence: The parameter λ need not be adjusted to the sampling frequency used in obtaining the experimental data.

The piecewise constant Mumford-Shah model has achieved substantial popularity in segmentation problems, due in part to the geometric regularization imposed on class boundaries that easily allows selecting a desired scale (level of detail) for the segmentation [11, 12]. Related models have been employed for segmenting images of grains e.g. [13, 14] but not in the specific context of parent grain reconstruction. There are a great variety of algorithms for carrying out approximate minimization of Mumford-Shah based variational segmentation models such as (2). Many of them alternate minimization with respect to shapes of the grains Σ_i and their orientations β_i. Typically, an essential difficulty in the optimization – in the form of stiffness of the numerical scheme – comes from the regularization term, which is of higher order than the fidelity term. Our algorithm, BetaTD (Algorithm 1), is based on *threshold dynamics* [15] for quickly simulating gradient descent. This scheme is unconditionally stable, enabling

arbitrarily large gradient descent steps independent of the spatial resolution, and can handle very large numbers of grains efficiently: for instance, it has been used in grain boundary motion simulations with hundreds of thousands of grains in both two and three dimensions. It therefore scales well for high resolution images potentially containing large numbers of parent grains. Nevertheless, we note that optimization literature for cost functions related to the Mumford-Shah functional is vast, and there may be other algorithms that are at least as effective on our new variational model.

Algorithm 1 BetaTD - Algorithm [15] Applied to Eq. (2).

Given a initial partition $\Sigma_1^0, \ldots, \Sigma_N^0$ and a time step Δt.
Let number of time steps, T, be sufficiently large to allow the scheme to reach steady state.
Let $G_{\Delta t}(\mathbf{x}) = \frac{1}{4\pi\Delta t} e^{-\frac{|\mathbf{x}|^2}{4\Delta t}}$ be the heat kernel.
Set $\beta_i = \arg\min_{\beta^*} \sum_{\mathbf{x}\in\Sigma_i^0} d(\beta^*, f(\mathbf{x}))$
for $t \leftarrow 1$ to T **do**
    $S_i(\mathbf{x}) = W(\mathbf{x}) d(\beta_i, f(\mathbf{x}))$.
    $\phi_i^t = -\frac{2}{\sqrt{\Delta t}}(G_{\Delta t} * \mathbb{1}_{\Sigma_i^{t-1}}) + \lambda S_i$.
    $\Sigma_i^t = \{\mathbf{x} : \phi_i^t(\mathbf{x}) = \min_j \phi_j^t(\mathbf{x})\}$
end for

The $\beta_1, \ldots, \beta_n$ need not be updated at every time step. In practice, we update the β_i once the Σ_i have changed sufficiently under BetaTD from the last time the β_i were updated. For given $\Sigma_1, \ldots, \Sigma_n$, we find the optimal β_i, i.e.

$$\beta_i = \arg\min_{\beta^*} \int_{\Sigma_i \cap \mathcal{A}} W(\mathbf{x}) d_\alpha(\beta^*, f_\alpha(\mathbf{x}))\, d\mathbf{x} + \int_{\Sigma_i \cap \mathcal{B}} W(\mathbf{x}) d_\beta(\beta^*, f_\beta(\mathbf{x}))\, d\mathbf{x}$$

using the expectation maximization - von Mises-Fisher algorithm [16].

It should be noted that the algorithm immediately extends to 3d data. Additionally, by simply changing the misorientation angle function, the methods can work for any reconstruction problem or, more generally, any single phase grain segmentation.

The BetaTD algorithm carries out gradient descent for the variational model (2). As such, it can only be expected to find a local minimum. There are more globally – albeit more costly – optimization approaches in the literature, which may be explored in future work. Since BetaTD only finds a local minimum, we run the algorithm multiple times, usually 12 or 16 runs, starting from different initial conditions containing an abundance of possible parent β grains, overestimating their true number. During the energy minimization, unneeded grains typically become empty and are thus effectively removed. Among the multiple runs, we choose the one leading to the lowest final energy, Eq. (2), as the output of the segmentation process. We call this random initialization. We can run the algorithm on the data without cleaning it beforehand, as we do in all the examples in this paper. As such, we use the weight function, denoted as

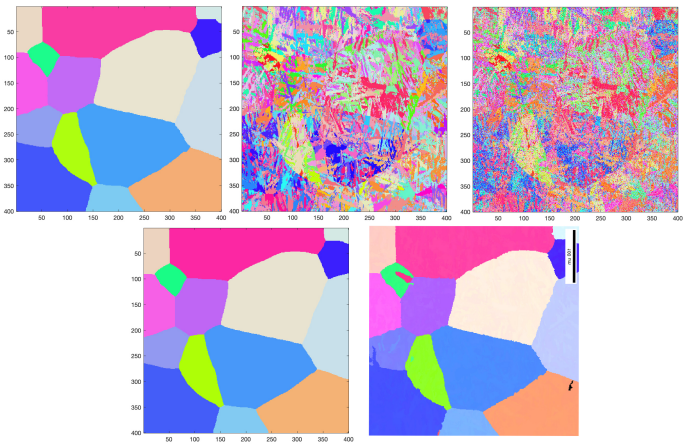

Figure 1: Simulated data set. The simulated parent β grains and simulated α colonies are shown top left and top center respectively. The top right shows the noisy α phase orientations. The bottom left shows the predicted β grains shapes and orientations from BetaTD and the bottom right shows the output of algorithm used in [4] for comparison.

$W(x)$ in Eq. (2), to indicate how confident that we are that a measurement is correct. This is naturally given by a function of the confidence index. Based on experience with working with TI-alloyed data, we used

$$W(\mathbf{x}) = \max\left\{ \min\left\{0\ ,\ \frac{\mathrm{CI}(\mathbf{x}) - .03}{.07}\right\}\ ,\ 1\right\}.$$

where $\mathrm{CI}(\mathbf{x})$ is the confidence index at data point $\mathbf{x}$. If clean up is done on the data beforehand, a weight function is not needed.

We begin by running BetaTD on a synthetic (simulated) data set. This analysis starts from a simulated parent β microstructure as shown in the top left of Fig. 1. From this microstructure, we generate many α lathes for each β grain as shown in the top center of Fig. 1. The orientation of each α lathe is produced by randomly picking one of the twelve variants from the α/β relationship. Then we randomly rotate 10% of orientations to simulate misidentifying the orientation. Following this, 5% of the orientation data is corrupted and given a confidence index of zero. We generate noise by adding random values from a mean zero, variance $\frac{1}{100}$ normal to the Bunge Euler angles of the data. This results in a noisy image of α colonies generated from a "ground truth" image of 15 β grains, shown in the top right of Fig. 1. The parent β microstructure predicted by BetaTD is shown in the bottom left image in Fig. 1. For comparison, we ran the algorithm from [4] on the data (shown in the bottom right). The predicted microstructure from BetaTD had 0.26% of its orientations differ by 3 degrees or more from the ground truth compared to the algorithm from [4] which had 0.99%.

To see how our algorithm scales with the size and number of grains we tiled our simulated data set by a factor of 2, 3 and 4 so the resulting images were

Voxels	Number of Grains	Average Time (in min)
400^2	15	11.2
800^2	60	48.4
1200^2	105	116.9
1600^2	240	209.8

Table 1: Timings of tiled simulated data sets.

Data Set	Size	Chosen λ	Notes
1	436 by 831	50	Highly Textured
2	620 by 1095	75	Textured
3	256 by 256	100	Few α colonies per β grain

Table 2: Data Sets

4, 9 and 16 times bigger respectively. Table 1 reports the average time to run BetaTD from a single random initial condition over 24 different attempts. The algorithm scales almost linearly.

Now, we report the results of the algorithm used to predict the parent β microstructure on three real EBSD data sets of various levels of difficulty. The data sets are summarized in Table 2.

To choose the parameter λ in Eq. (2), we ran the model on a subset of the data with a variety of different values of λ. Based on the output, we chose which λ to run on the full data set.

Fig. 2 shows the result on data set 1, a particularly clean image of α colonies that has resulted from a fairly low number (around 25) of parent β grains. We used random initialization. The segmentation by BetaTD creates realistic looking grains and segments the data well. Note that the orientation seems to vary across the β grain in the top left of the image. Because our model assumes that each grain has constant orientation, it finds an extra boundary with a misorientation angle of 1.36 degrees. Our energy, Eq. (2), can be adapted to allow the orientation to vary smoothly across each parent β grain, which is left for future work. We also note the detection of additional grains compared with the results of the algorithm used in [4].

Fig. 3 shows the result on data set 2, which contains substantially more parent β grains. Compared to the algorithm used in [4], shown on the bottom, our algorithm gives cleaner boundaries and finds far fewer spurious grains. We found our initial orientations and boundaries for BetaTD with a version of the algorithm in [4].

Finally, Fig. 4 shows the result on data set 3. This data set did not contain sufficient α orientation information to determine the β grain orientation uniquely. Despite this, the algorithm picks out the same boundaries when run with different random analytical initial boundaries. The results indicate that the algorithm can pick out the boundaries of parent β grains even with insufficient α orientation data to determine the parent β grain orientation.

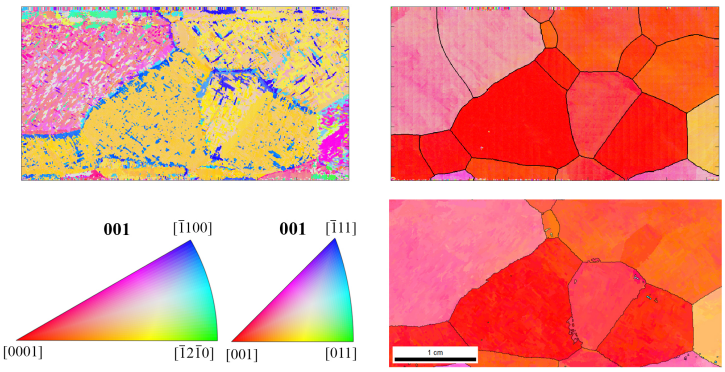

Figure 2: Data set 1. Top Left: The input α platelet orientations. Top Right: BetaTD predicted boundaries and in order to visually evaluate the results better, we choose to color each pixel by picking among the six possible β-parents of each α orientation the one closest to the β_i found by BetaTD. Bottom left: The IPF keys of the α phase (left) and β phase (right) orientations for the images in this paper. Bottom Right: Predicted results of the algorithm presented in [4].

In this article, we introduced a new variational model, Eq. (2), for reconstructing parent β grains from α colonies based on the segmentation method of Mumford-Shah. The model is particularly well suited to the β reconstruction problem because it is biased towards producing segmentations in which each region resembles a plausible grain, with a regular, clean, sharp boundary across which the crystallographic orientation changes discontinuously. This translated into much more localized and well defined boundaries than what has been obtained with the previous pointwise reconstruction approach [6] or α grain reconstruction approaches [3, 4, 5]. We demonstrated these properties on simulated and real images with satisfactory results.

While this is an encouraging first step, there remains more to be done. Future work could implement a piecewise smooth Mumford-Shah (as opposed to the piecewise constant version) which would give better segmentation when the orientation across each parent β grains is not nearly constant. Additionally, the algorithm could be extended to non-Burgers orientation relation microstructures. The code for BetaTD is available at `https://github.com/AZaitzeff/betaTD`.

References

[1] C. Cayron, Groupoid of orientational variants, Acta Crystallogr. Sect. A: Found. of Crystallogr. 62 (1) (2006) 21–40.

[2] D. Mumford, J. Shah, Optimal approximations by piecewise smooth functions and associated variational problems, Commun. on Pure and Appl. Math. 42 (5) (1989) 577–685.

[3] L. Germain, N. Gey, M. Humbert, Reliability of reconstructed β-orientation maps in titanium alloys, Ultramicrosc. 107 (12) (2007) 1129–1135.

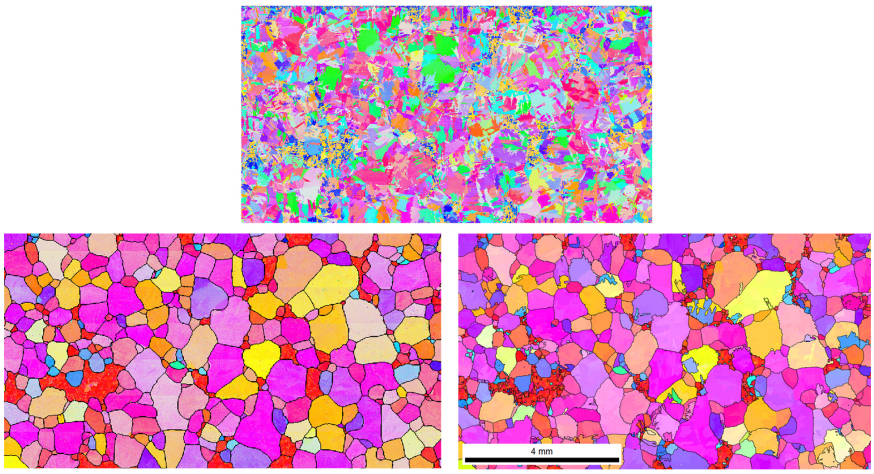

Figure 3: Data set 2. Top: The input α orientations. Bottom Left: Proposed β orientations and boundaries (left). Bottom Right: Results of the algorithm used in [4].

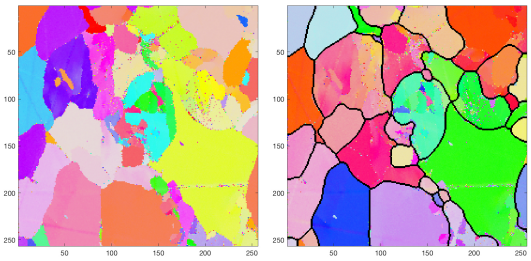

Figure 4: Data set 3. Left: The input α orientations. Right: BetaTD predicted β orientations and boundaries

[4] A. Pilchak, J. Williams, Microstructure and texture evolution during friction stir processing of fully lamellar ti-6al-4v, Metall. and Mater. Trans. A 42 (3) (2011) 773–794.

[5] J. Tiley, A. Shiveley, A. Pilchak, P. Shade, M. Groeber, 3d reconstruction of prior β grains in friction stir–processed ti–6al–4v, J. of Microsc. 255 (2) (2014) 71–77.

[6] E. Wielewski, D. Menasche, P. Callahan, R. Suter, Three-dimensional α colony characterization and prior-β grain reconstruction of a lamellar ti–6al–4v specimen using near-field high-energy x-ray diffraction microscopy, J. of Appl. Cryst. 48 (4) (2015) 1165–1171.

[7] G. Miyamoto, N. Iwata, N. Takayama, T. Furuhara, Mapping the parent austenite orientation reconstructed from the orientation of martensite

by ebsd and its application to ausformed martensite., Acta Materialia 58 (2010) 6393–6403.

[8] L. Germain, S. R. Dey, M. Humbert, N. Gey, Determination of parent orientation maps in advanced titanium-based alloys., J. of Microsc. 227 (2007) 284–291.

[9] L. Germain, P. Blaineau, N. Gey, M. Humbert, Current approaches for reconstructing the parent microtexture from that inherited by phase transformation., Mater. Sci. Forum 702–703 (2012) 846–849.

[10] L. Germain, N. Gey, R. Mercier, P. Blaineau, M. Humbert, An advanced approach to reconstructing parent orientation maps in the case of approximate orientation relations: Application to steels., Acta Materialia 60 (2012) 4551–4562.

[11] T. F. Chan, L. A. Vese, Active contours without edges, IEEE Trans. on Image Process. 10 (2) (2001) 266–277.

[12] L. A. Vese, T. F. Chan, A multiphase level set framework for image segmentation using the Mumford and Shah model, Int. J. of Comput. Vis. 50 (3) (2002) 271–293.

[13] J. Waggoner, Y. Zhou, J. Simmons, A. Salem, M. De Graef, S. Wang, Interactive grain segmentation using graph cut algorithms, in: Proceedings of SPIE (Computational Imaging XI), Vol. 8657, 2013.

[14] J. Waggoner, Y. Zhou, J. Simmons, M. De Graef, S. Wang, Graph-cut based interactive segmentation of 3d materials-science images., Mach. Vis. and Appl. 25 (2014) 1615–1629.

[15] S. Esedoglu, F. Otto, Threshold dynamics for networks with arbitrary surface tensions, Commun. on Pure and Appl. Math. 68 (5) (2015) 808–864.

[16] Y.-H. Chen, D. Wei, G. Newstadt, M. DeGraef, J. Simmons, A. Hero, Statistical estimation and clustering of group-invariant orientation parameters, in: Information Fusion (Fusion), 2015 18th International Conference on, IEEE, 2015, pp. 719–726.