

Denoising of Crystal Orientation Maps

Ralf Hielscher¹, Christian B. Silbermann², and Eric Schmidl²

¹Chemnitz University of Technology, Department of Mathematics,
09107 Chemnitz, Germany

²Chemnitz University of Technology, Department of Mechanical
Engineering, 09107 Chemnitz, Germany

September 24, 2018

This paper compares several well-known methods for denoising orientation data with methods adapted from mathematical image analysis. The latter ones turn out to be much more powerful in terms of preserving low angle grain boundaries and filling holes of non-indexed orientations. We also discuss the effect of denoising to the determination of the kernel average misorientation and the geometrically necessary dislocation density. Synthetic as well as experimental data are considered for this comparison. The examples demonstrate that variational denoising techniques are capable of significantly improving the accuracy of properties derived from EBSD maps.

1 Introduction

Individual crystal orientations are commonly obtained from light microscopy, electron back scatter diffraction (EBSD) by indexing of Kikuchi patterns [25], from automated crystal orientation mapping inside the transmission electron microscope (ACOM-TEM) by indexing spot diffraction patterns [19], or from X-ray diffraction (XRD) [18]. As with any other experimental data these crystal orientations are usually corrupted by errors. In case the error is not systematic but random it is called noise.¹

In this paper we aim at recovering the “true” orientation map from a noisy EBSD or ACOM-TEM orientation map. The key requirement to be able to separate noise from true orientations is the assumption that the noise is spatially independent. This means that the error in one pixel is not correlated to the error in a nearby pixel. In particular, we assume that the orientation map is measured on a sufficiently fine grid such that the true orientations are spatially independent. Directly speaking, our algorithms will not be able to recover one-pixel features but will consider those as noise.

¹In this context, noisy orientation data will be considered “raw data” irrespective of any (automated) data processing related to the aforementioned measurement procedures.

Denosing of EBSD data has already been addressed in several papers [7, 9, 10, 14]. In the current paper we compare well-established methods like the mean filter, the Kuwahara filter [9, 14] and the median filter [10] with variational denoising techniques [1, 4, 6, 21, 23] adapted from mathematical image analysis to the setting of orientation data.

The central idea of variational denoising methods is to determine the true orientations as a solution of a minimization problem that couples the fit of the recovered orientations to the noisy orientations with a measure for the smoothness of the orientation map. This measure of smoothness models our expectations about the orientation map and should favor smooth changes of the crystal orientation disrupted by a few grain and subgrain boundaries.

The paper is divided into two parts: In the first part, we discuss the basic properties of the aforementioned denoising techniques and illustrate them at a synthetic as well as at an experimental EBSD orientation map. In the second part, we discuss the impact of these denoising techniques to various applications, i.e., the determination of the kernel average misorientation, the orientation gradient and the geometrically necessary dislocation density. All these examples demonstrate that variational denoising techniques are capable of significantly improving the accuracy of properties derived from EBSD maps.

All figures and numerical experiments have been done using MTEX 5.1.1 [13] and the script files publicly available at <https://github.com/mtex-toolbox/mtex-paper/tree/master/DenoisingCrystalOrientationMaps>.

1.1 Crystal Orientations and Misorientations

Throughout this paper orientations are denoted by bold capital letters, e.g. \mathbf{O} , and refer to proper rotations that map coordinates with respect to a crystal fixed Euclidean reference frame to coordinates with respect to a specimen fixed Euclidean reference frame. Accordingly, \mathbf{OS} with $\mathbf{S} \in S$ denotes the class of all orientations symmetrically equivalent to \mathbf{O} with respect to the symmetry group S .

Given two orientations $\mathbf{O}_1, \mathbf{O}_2$ its misorientation is defined by $\mathbf{O}_1^{-1}\mathbf{O}_2$. The corresponding class of symmetrically equivalent misorientations becomes $\mathbf{S}_1^{-1}\mathbf{O}_1^{-1}\mathbf{O}_2\mathbf{S}_2$ with $\mathbf{S}_1, \mathbf{S}_2 \in S$. The misorientation $\mathbf{M} = \mathbf{S}_1^{-1}\mathbf{O}_1^{-1}\mathbf{O}_2\mathbf{S}_2$ with the smallest rotational angle is called disorientation. Accordingly, the rotational angle δ and the rotational axis $\boldsymbol{\eta}_C$ of \mathbf{M} are called disorientation angle and disorientation axis. It is important to understand, that according to this definition the disorientation axis $\boldsymbol{\eta}_C$ is with respect to the crystal reference frame and hence only well-defined up to multiplications with crystal symmetries $\mathbf{S} \in S$. However, when translated into specimen coordinates the disorientation axis $\boldsymbol{\eta} = \mathbf{O}_1\mathbf{S}_1\boldsymbol{\eta}_C = \mathbf{O}_2\mathbf{S}_2\boldsymbol{\eta}_C$ is uniquely defined.

Multiplying the disorientation angle δ with the disorientation axis $\boldsymbol{\eta}$ we obtain the so-called disorientation vector

$$\boldsymbol{\theta} = \delta\boldsymbol{\eta}. \quad (1.1)$$

This disorientation vector will play an important role throughout this paper as it allows to represent elements of the curved orientation space locally by ordinary three dimensional vectors. The point is that while many standard mathematical operations, like taking the mean, fail in the curved orientation space, they are well-defined in the three dimensional vector space of disorientation vectors. In our applications we often use the grain mean orientation as the reference orientation \mathbf{O}_2 , which guaranties that orientation \mathbf{O}_1 and the reference orientation




Figure 2.1: Artificial orientation maps colorized by the misorientation angle.

\mathcal{O}_2 are sufficiently close together to allow for a good approximation.²

2 Comparison of Denoising Methods

2.1 An artificial Orientation Map

In this section we construct a very simple synthetic orientation map which we will use to benchmark different denoising methods in the subsequent sections. The synthetic orientation map consists of 100×25 pixels, cf. Fig. 2.1(a). While the left half of the map is perfectly identically oriented (Euler angles $(0^\circ, 0^\circ, 0^\circ)$), the right half has a constant texture gradient of -0.36 degree per pixel. This results in a 18 degree grain boundary separating both halves. In order to simulate a measured orientation map we add different kinds of noise.

When dealing with noisy scalar data $\tilde{f}(\mathbf{x})$ at locations $\mathbf{x} = (x_1, x_2) \in \mathbb{R}^2$ one usually assumes a noise model of the form

$$\tilde{f}(\mathbf{x}) = f(\mathbf{x}) + \varepsilon(\mathbf{x}), \quad (2.1)$$

where $f(\mathbf{x})$ are the noise free data and $\varepsilon(\mathbf{x})$ is the noise, e.g., an independent random sample of a Gaussian distribution. For orientation data $\mathcal{O}(\mathbf{x})$ we have to replace the summation in (2.1) by the multiplication with some random orientations $\mathcal{E}(\mathbf{x})$, i.e.,

$$\tilde{\mathcal{O}}(\mathbf{x}) = \mathcal{E}(\mathbf{x})\mathcal{O}(\mathbf{x}).$$

We assume the orientations $\mathcal{E}(\mathbf{x})$ to be an independent random sample of a probability distribution \mathcal{E} on the space of rotations. Important examples of such distributions follow now.

Bingham- and De la Vallée Poussin-distributed Noise. A generalization of the Gaussian distribution to orientations is the Bingham distribution [2]. Representing the misorientations $\mathcal{E}(\mathbf{x})$ by unit quaternions $\mathbf{q} = (q_0, q_1, q_2, q_e)^T$ the corresponding density function is given by

$$f(\mathbf{q}) = C \exp(\mathbf{q}^T \mathbf{A} \mathbf{q}) \quad (2.2)$$

²Approximating an orientation by its disorientation vector with respect to some reference orientation can be seen as projecting the orientation into the tangential space of the reference orientation. This mapping is often referred to as (matrix) logarithm and the inverse mapping as (matrix) exponential, cf. [16].

with a symmetric 4×4 matrix \mathbf{A} and a normalization constant C such that the integral over f is one. Assuming that the eigenvector corresponding to the largest eigenvalue of A is the orientation $(0, 0, 0)$ and that all other eigenvalues are equal, we obtain anisotropic noise without drift. The variance depends on the fraction between the largest and the smallest eigenvalue.

A suitable approximation for the anisotropic Bingham distribution is given by the following De la Vallée Poussin (DIVP) distribution [20]:

$$f(\mathbf{q}) = \frac{B\left(\frac{3}{2}, \frac{1}{2}\right)}{B\left(\frac{3}{2}, \kappa + \frac{1}{2}\right)} q_1^{2\kappa}, \quad (2.3)$$

where B denotes the Beta function, q_1 is the real part of the quaternion \mathbf{q} and κ models the variance. Using the DIVP distribution (2.3), the numerically challenging computation of the normalization constant of distribution (2.2) can be omitted. Figure 2.1(b) displays the artificial orientation map corrupted by DIVP-distributed noise with half-width of two degree.

Impulsive Noise. While the generalization of Gaussian noise discussed in the previous paragraph models small deviations from the true orientation due to indexing errors, the noise model in this paragraph models the case that the indexing process results in a completely wrong orientation that is entirely independent of the true orientation. Such noise is called *impulsive noise*. The corresponding distribution involves a parameter $\alpha \in [0, 1]$. Then, with a chance of α the original orientation is replaced by a random orientation and with a chance of $1 - \alpha$ the original orientation is left unchanged. Our artificial orientation map corrupted by Gaussian-like noise and impulsive noise with $\alpha = 0.05$ is depicted in Figure 2.1(c).

2.2 Definition of a Filter

Denosing a noisy orientation map $\tilde{\mathbf{O}}(\mathbf{x}) = \mathcal{E}(\mathbf{x})\mathbf{O}(\mathbf{x})$ means finding an orientation map $\hat{\mathbf{O}}$ that is as close as possible to the “true” orientation map \mathbf{O} . We will call any algorithm that computes out of a noisy orientation map $\tilde{\mathbf{O}}$ a denoised map $\hat{\mathbf{O}}$ a filter \mathcal{F} and write

$$\hat{\mathbf{O}} = \mathcal{F}\tilde{\mathbf{O}}.$$

Accordingly, $\mathcal{F}\tilde{\mathbf{O}}(\mathbf{x})$ denotes the orientation at pixel position \mathbf{x} in the denoised map.

2.3 The Mean Filter

Let \mathbf{x} be a certain pixel position within an EBSD map and $r \in \mathbb{N}$. Then we denote by $\mathcal{N}_r(\mathbf{x})$ the set of all neighboring pixels up to order r and by $\#\mathcal{N}_r(\mathbf{x})$ the number of these neighbors. Now, the so-called mean filter is defined as

$$\mathcal{F}_{\text{mean}}\tilde{\mathbf{O}}(\mathbf{x}) = \frac{1}{\#\mathcal{N}_r(\mathbf{x})} \sum_{\mathbf{x}' \in \mathcal{N}_r(\mathbf{x})} \tilde{\mathbf{O}}(\mathbf{x}') \quad (2.4)$$

which replaces the noisy orientation $\tilde{\mathbf{O}}(\mathbf{x})$ at position $\mathbf{x} = (x_1, x_2)$ with the mean of all orientations $\tilde{\mathbf{O}}(\mathbf{x}')$ in the neighborhood $\mathcal{N}_r(\mathbf{x})$ of \mathbf{x} .

It is important to note that equation (2.4) is formally not correct. The point is that due to the curved non-Euclidean geometry of the orientation space the sum of orientations is




Figure 2.2: Artificial orientation maps corrupted by DIVP noise (all) and impulsive noise (right) after denoising with the mean filter involving first order neighbors (left and right) and third order neighbors (middle).

in general not an orientation anymore. A better option to compute the mean orientation is the largest Eigenvector of the mean orientation tensor as described in [2] or using the disorientation vectors defined in (1.1). However, in order to keep notation simple in the present section we will stick to the notation in (2.4), keeping in mind that its actual computation is a bit more involved.

We can even generalize the mean filter by introducing a weighting function $w: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ which models the impact of an orientation to the mean orientation in dependency of its distance between the points \mathbf{x} and \mathbf{x}' , i.e.,

$$\mathcal{F}_{\text{mean}} \tilde{\mathbf{O}}(\mathbf{x}) = \frac{\sum_{\mathbf{x}' \in \mathcal{N}(\mathbf{x})} w(\|\mathbf{x} - \mathbf{x}'\|) \tilde{\mathbf{O}}(\mathbf{x}')}{\sum_{\mathbf{x}' \in \mathcal{N}(\mathbf{x})} w(\|\mathbf{x} - \mathbf{x}'\|)}. \quad (2.5)$$

Setting the weighting function $\omega(d) = 1$, the filter is independent of the distance $d = \|\mathbf{x} - \mathbf{x}'\|$ to the neighboring pixel at \mathbf{x}' and we end up with formula (2.4). Setting ω to

$$\omega(d) = e^{-d^2/\sigma^2}$$

the corresponding mean filter can be interpreted as the convolution of the noisy image with a Gaussian. The weighting function ω could also be chosen to be dependent on some quality measure of the orientation measurement at position \mathbf{x}' or the misorientation angle between the $\tilde{\mathbf{O}}(\mathbf{x})$ and $\tilde{\mathbf{O}}(\mathbf{x}')$.

The basic motivation for the mean filter is the central limit theorem. As a consequence, the mean filter behaves particularly well if the orientation data are corrupted by Bingham noise. Figure 2.2 shows the application of the mean filter to synthetic orientation data which consist of a uniform part, a grain boundary and a linear gradient, all corrupted by Bingham-distributed noise. We observe that the data are smoothed and the noise is reduced and that this effect increases if the second or third order neighbors are used for averaging too. On the downside, increasing the number of neighbors increases also the blurring to the grain boundary.

While the mean filter behaves well for Gaussian type noise, it cannot handle impulsive noise. This is illustrated in the right subfigure of Fig. 2.2. The reason is that the mean filter averages equally over all orientations and is not able to identify outliers.




Figure 2.3: Artificial orientation maps corrupted by DIVP noise (all) and impulsive noise (right) after denoising with the median filter involving first order neighbors (left) and third order neighbors (middle and right).

2.4 The Median Filter

Let again denote $\mathcal{N}(\mathbf{x})$ the set of neighboring measurement positions to \mathbf{x} and let $\delta(\mathbf{O}, \tilde{\mathbf{O}})$ denote the disorientation angle between two orientations \mathbf{O} and $\tilde{\mathbf{O}}$. Then according to [11] the median filter $\mathcal{F}_{\text{median}}\tilde{\mathbf{O}}(\mathbf{x})$ on some noisy orientation map $\tilde{\mathbf{O}}(\mathbf{x})$ is defined as the orientation $\tilde{\mathbf{O}}(\mathbf{x}')$ with $\mathbf{x}' \in \mathcal{N}^r(\mathbf{x})$ such that it possesses the minimum mean distance to all other orientations $\tilde{\mathbf{O}}(\bar{\mathbf{x}})$, $\bar{\mathbf{x}} \in \mathcal{N}_r(\mathbf{x})$ in the neighborhood, i.e., $\mathbf{x}' \in \mathcal{N}(\mathbf{x})$ is chosen such that for any other choice $\mathbf{x}'' \in \mathcal{N}(\mathbf{x})$ it is satisfied

$$\sum_{\bar{\mathbf{x}} \in \mathcal{N}(\mathbf{x})} \delta(\tilde{\mathbf{O}}(\mathbf{x}'), \tilde{\mathbf{O}}(\bar{\mathbf{x}})) \leq \sum_{\bar{\mathbf{x}} \in \mathcal{N}(\mathbf{x})} \delta(\tilde{\mathbf{O}}(\mathbf{x}''), \tilde{\mathbf{O}}(\bar{\mathbf{x}})). \quad (2.6)$$

Note that the median filter may be extended to be applicable to multiple phases. In this case we set $\delta(\mathbf{O}, \tilde{\mathbf{O}})$ to some large value whenever the orientations \mathbf{O} and $\tilde{\mathbf{O}}$ belong to different phases. Furthermore, the reliability of orientation measurements can be incorporated into the filter using a weight function $\omega(\mathbf{x})$ similarly as for the mean filter. Let us assume that $1/\omega(\mathbf{x})$ models the reliability of the orientation measurement $\tilde{\mathbf{O}}(\mathbf{x})$, i.e. by being the band contrast, the mean angular deviation (MAD), a confidence index or some cross-correlation score. Then, condition (2.6) can be generalized to

$$\sum_{\bar{\mathbf{x}} \in \mathcal{N}(\mathbf{x})} \frac{\omega(\mathbf{x}) \delta(\tilde{\mathbf{O}}(\mathbf{x}'), \tilde{\mathbf{O}}(\bar{\mathbf{x}}))}{\omega(\mathbf{x}') + \omega(\bar{\mathbf{x}})} \leq \sum_{\bar{\mathbf{x}} \in \mathcal{N}(\mathbf{x})} \frac{\omega(\mathbf{x}'') \delta(\tilde{\mathbf{O}}(\mathbf{x}''), \tilde{\mathbf{O}}(\bar{\mathbf{x}}))}{\omega(\mathbf{x}'') + \omega(\bar{\mathbf{x}})}$$

which prefers measurements with smaller expected error to those with higher expected error.

By construction, the median filter does not introduce new orientations but the orientations of the filtered map are a subset of the noisy orientations. As a consequence, the median filter works very well in removing outliers, e.g., for impulsive noise, cf. Fig. 2.3 (right), and does not blur subgrain boundaries. On the downside, it leads to blocky images which tend to underestimate the true local texture gradient and KAM. These effects are illustrated by Figure 2.3 (left). Another type of median filter that works in a comparable manner was reported in [15].

2.5 The Kuwahara Filter

A second denoising method that is build in tool of many EBSD pattern processing software is the Kuwahara filter. Similarly to the mean and the median filter, the Kuwahara filter is




Figure 2.4: Artificial orientation maps corrupted by DIVP noise (all) and impulsive noise (right) after denoising with the Kuwahara filter involving first neighbors (left) and third neighbors (middle and right).

a sliding window filter that determines the new orientation $\mathcal{F}_{\text{Kuwahara}}\tilde{\mathbf{O}}(\mathbf{x})$ from the noisy orientations $\tilde{\mathbf{O}}(\mathbf{x}')$ within a certain neighborhood $\mathcal{N}(\mathbf{x})$ around the original pixel \mathbf{x} . The idea of the Kuwahara filter is to split this neighborhood $\mathcal{N}(\mathbf{x})$ into smaller regions $\mathcal{N}_i(\mathbf{x})$, $i = 0, \dots, 3$ which are usually taken as the four quadrants if the original region is assumed to be centered at the origin. Next, for each subregion $\mathcal{N}_i(\mathbf{x})$ the mean orientation $\mathbf{m}_i(\mathbf{x})$ and the standard deviation $\sigma_i(\mathbf{x})$ of the misorientation angle to $\mathbf{m}_i(\mathbf{x})$ are computed. Finally, the filtered orientation $\mathcal{F}_{\text{Kuwahara}}\tilde{\mathbf{O}}(\mathbf{x})$ is chosen to be the mean orientation $\mathbf{m}_i(\mathbf{x})$ corresponding to the smallest standard deviation.

Due to the selection of a subregion before taking the mean, the Kuwahara filter shows a good performance at grain boundaries but tends to generate blocky structures in the presence of Gaussian noise, cf. Fig. 2.4. For impulsive noise it behaves quite well if the outliers are not too dense such that there is always a subregion without outlier. However, in order to make the Kuwahara filter effective for Gaussian noise the sliding window has to be chosen about twice as big as for the mean filter, which makes it likely to have an outlier in each subregion. As soon as this happens the outliers generate blocky artifacts in the filtered image, cf. Fig. 2.4 (right).

2.6 Variational Filters

The variational approach to image denoising consists of approximating noisy data $\tilde{\mathbf{O}}(\mathbf{x})$ by orientations $\hat{\mathbf{O}} = \mathcal{F}_{\text{var}}\tilde{\mathbf{O}}(\mathbf{x})$ that minimize a certain energy functional

$$J(\hat{\mathbf{O}}) = \sum_{\mathbf{x}} \delta(\tilde{\mathbf{O}}(\mathbf{x}), \hat{\mathbf{O}}(\mathbf{x}))^2 + \alpha\varphi(\hat{\mathbf{O}}),$$

where the first summand is the data fidelity term, which ensures that the approximating values $\hat{\mathbf{O}}(\mathbf{x})$ are not too far from the measured values $\tilde{\mathbf{O}}(\mathbf{x})$ and the second summand $\varphi(\hat{\mathbf{O}})$, called penalty term, ensures that the approximating image $\hat{\mathbf{O}}$ is more smooth than $\tilde{\mathbf{O}}$. The regularization parameter α allows to adjust the solution between being very close to the original data and being very smooth.

Smoothing Splines. The simplest form of a variational filter uses the square Euclidean norm of the Laplacian as the penalty term, i.e.

$$J(\hat{\mathbf{O}}) = \sum_{\mathbf{x}} \delta(\tilde{\mathbf{O}}(\mathbf{x}), \hat{\mathbf{O}}(\mathbf{x}))^2 + \alpha \left| \Delta \hat{\mathbf{O}}(\mathbf{x}) \right|^2. \quad (2.7)$$




Figure 2.5: Artificial orientation maps corrupted by DIVP noise (all) and impulsive noise (right) after denoising with smoothing splines with $\alpha = 0.2$ (left), $\alpha = 5$ (middle and right).

Solutions of the corresponding minimization problem are known as smoothing splines. However, the generalization of the Laplacian $\Delta \mathbf{O}(\mathbf{x})$ to orientation valued functions $\mathbf{O}(\mathbf{x})$ is not straight forward. Possible definitions have been discussed in [5]. In this work we first approximate the orientation map $\hat{\mathbf{O}}$ locally by a disorientation vector map $\tilde{\boldsymbol{\theta}}(\mathbf{x})$, cf. (1.1), with respect to some reference orientation. Since, the disorientation vector map $\tilde{\boldsymbol{\theta}}(\mathbf{x})$ takes values in \mathbb{R}^3 we can then use the Euclidean Laplacian and apply standard algorithms, e.g. [8] to compute the minimizer $\hat{\boldsymbol{\theta}}$ of the translated problem

$$J(\hat{\boldsymbol{\theta}}) = \sum_{\mathbf{x}} \delta(\tilde{\boldsymbol{\theta}}(\mathbf{x}), \hat{\boldsymbol{\theta}}(\mathbf{x}))^2 + \alpha \left| \Delta \hat{\boldsymbol{\theta}}(\mathbf{x}) \right|^2 \rightarrow \min.$$

Finally, we translate the disorientation vector map $\hat{\boldsymbol{\theta}}(\mathbf{x})$ back into an orientation map $\hat{\mathbf{O}}(\mathbf{x})$. The approximation of orientations by disorientation vectors with respect some reference orientation is fairly good as long as the orientations are close to the reference orientation. Hence, it is appropriate to apply the approximation for each grain separately and choose as the reference orientation the mean orientation of the grain.

In its standard form, smoothing splines behave poorly with respect to impulsive noise. To overcome this restriction the algorithm described in [8] includes an adaptive reweighing scheme that puts a lower weight to function values identified as outliers. Fig. 2.1(c) shows that the adaption to orientation data behaves well with respect to impulsive noise. The smoothing spline algorithm described in [8] also possesses two more features interesting for denoising EBSD data. Firstly, it implements a cross correlation based approach to automatically detect the regularization parameter α . As the automatic detection of this parameter requires several runs of the minimization problem it is recommended to detect the optimal regularization parameter for the largest grain only and use this value for all other grains. Secondly, it allows for weights $\omega(\mathbf{x})$ that can be used to describe the reliability of orientation data at position \mathbf{x} , cf. Subsec. 2.4. Fitting to orientations with low weight is then considered less important in the approximation process than fitting orientations with a high weight.

The results for our synthetic example are shown in Figure 2.5. We observe that the smoothing splines filter has a similar outcome as the mean filter. In particular it shows good results in the uniform and in the linear part but blurs the grain boundary.

Total Variation. The blurring effect of the smoothing splines is mostly due to the use of the square in the penalty term in (2.7). If we replace the Laplacian with the norm of the




Figure 2.6: Artificial orientation maps corrupted by DIVP noise (all) and impulsive noise (right) after denoising by total variation with $\alpha = 0.025$ (left), $\alpha = 0.5$ (middle and right).

gradient and remove the square we end up with the well-known isotropic total variation (TV) functional

$$J(\hat{\mathbf{O}}) = \sum_{\mathbf{x}} \delta(\hat{\mathbf{O}}(\mathbf{x}), \tilde{\mathbf{O}}(\mathbf{x}))^2 + \alpha \|\nabla \hat{\mathbf{O}}(\mathbf{x})\|_2. \quad (2.8)$$

The mathematical challenge in solving the corresponding minimization problem is related to the fact that the penalty term is not differentiable. There is a huge amount of literature describing algorithms that approximately solve (2.8) in the Euclidean setting. Among others, the following algorithms have been generalized to orientation values: an algorithm based on half-quadratic minimization [4] and an algorithm based on proximal mappings [23].

The idea of half-quadratic minimization is to approximate the square root of squares in the TV term by a differentiable functional. Furthermore, we may include a cut-off at a certain threshold angle d such that misorientations larger than the threshold, e.g. at (sub) grain boundaries, are not penalized. An example for such a function is

$$\varphi(\delta) = \begin{cases} \sqrt{\delta^2 + \epsilon} & \omega < d, \\ 0 & \delta \geq d, \end{cases}$$

which for $\epsilon \rightarrow 0$ converges to the TV penalty term.

Applying the total variation filter to our synthetic EBSD map we observe in Fig. 2.6 almost as sharp grain boundaries as the median filter produces and almost as smooth uniform and linear parts as by the mean filter. The only artifact we observe is that the linear portion does not decay at the left hand boundary to the uniform level. On the downside, the total variation filter is not able to distinguish between grain boundaries that exceed the threshold δ and impulsive noise as illustrated in Fig. 2.6 (right).

Infimal Convolution. The drawback of TV regularization is that it priors piecewise constant functions, i.e., the results are prone to staircasing. An idea to overcome this issue is to decompose the image $\mathbf{O}(\mathbf{x}) = \mathbf{U}(\mathbf{x})\mathbf{V}(\mathbf{x})$ into the product of a piecewise constant part $\mathbf{U}(\mathbf{x})$ which is penalized with the TV term and a smooth part $\mathbf{V}(\mathbf{x})$ that is penalized by a higher order differential operator, i.e. the Laplacian. In this setting the functional

$$J(\hat{\mathbf{U}}, \hat{\mathbf{V}}) = \sum_{\mathbf{x}} \delta(\tilde{\mathbf{O}}(\mathbf{x}), \hat{\mathbf{U}}(\mathbf{x})\hat{\mathbf{V}}(\mathbf{x}))^2 + \alpha \|\nabla \hat{\mathbf{U}}(\mathbf{x})\|_2 + \beta \left| \Delta \hat{\mathbf{V}}(\mathbf{x}) \right|. \quad (2.9)$$




Figure 2.7: Artificial orientation maps corrupted by DIVP noise (all) and impulsive noise (right) after denoising by infimal convolution with $\alpha = 0.025$ (left), $\alpha = 0.5$ (middle and right).

to minimize has two arguments \hat{U} , \hat{V} which are punished separately but generate the final orientation map via $\hat{O} = \hat{U}\hat{V}$. Again, half-quadratic minimization and proximal mappings are two suitable methods to find the minimizer of (2.9). The most challenging part in applying infimal convolution is the choice of the two regularization parameters α and β . It is recommended to start with a very high value for β and to choose the parameter α such that the noise is sufficiently reduced. For large β the solution of the infimal convolution functional is almost the same as for the total variation functional and hence we should expect some staircasing. Next, we can reduce the parameter β until the staircasing effect is reduced.

2.7 Morphological Filters and Inpainting

As we have seen in the previous sections most of the denoising methods do not work well with impulsive noise. A possible workaround for this problem is to perform grain reconstruction on the noisy EBSD map and consider all one pixel grains as impulsive noise. Those pixel values are first removed and latter filled by the respective denoising method. This methods is illustrated in Figure 2.8 for the synthetic orientation data given in Fig. 2.1(c).

Infimal convolution is not included into this list as we have not yet implemented inpainting with this method. We observe that all remaining denoising methods show a good performance filling one-pixel holes. Especially, the mean filter, the Kuwahara filter and the total variation method show much better results in comparison to case when the outliers have not been removed Fig. 2.2(right), 2.4(right) and 2.6(right).

This approach can be generalized by removing all grains with a specific shape, e.g. all grains that are only one pixel thick but arbitrary long. While for the sliding window filters the size of the window limits the capability to fill arbitrarily large holes, such a restriction does not persist for variational denoising techniques. This fact makes them extremely powerful for denoising purely-indexed EBSD maps, e.g. from highly-deformed materials.

Up to here, only artificial, synthetic orientation data has been considered. As a particularly challenging example of experimental data we consider now an EBSD map of a real geological Forsterite sample which is publicly available as a standard data in MTEX [13]. As visible in Figure 2.9(a) we find a lot of non-indexed pixels in the map. Non the less, the grain structure is still observable. However, due to the large amount of non-indexed pixels the traditional grain boundary (GB) reconstruction method completely fails to recover the original grain structure, cf. Fig. 2.9(b). In such settings it turns out to be extremely useful to reconstruct the grain structure using the Voronoi decomposition algorithm as described in [3]. This algorithm places the grain boundaries accurately through the non-indexed regions and gives

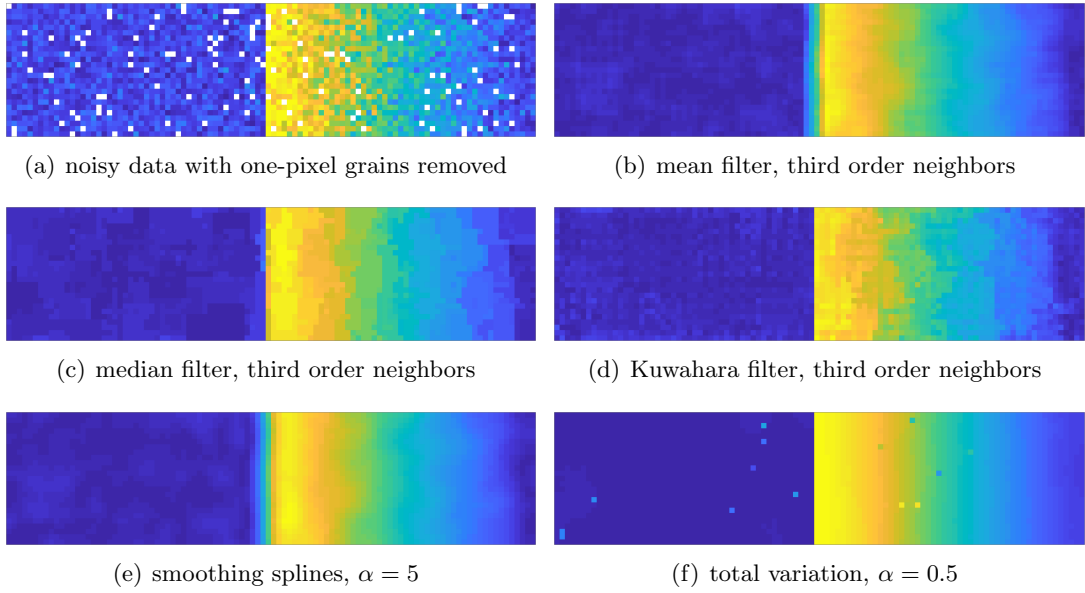


Figure 2.8: Simulated data corrupted by 2 degree DIVP noise and 5 percent impulsive noise where the impulsive noise is removed as one pixel grains and the holes are filled by the inpainting method.

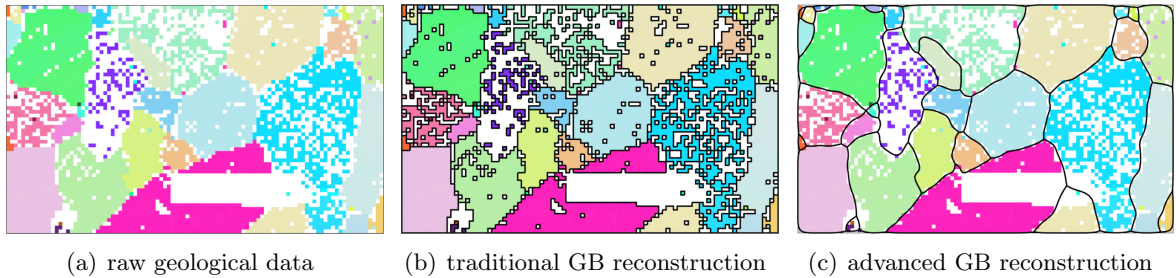


Figure 2.9: Grain boundaries reconstructed from poorly-indexed geological Forsterite data, the standard IPF color map from MTEX is used [13].

a GB reconstruction which is much closer to our expectation, cf. Fig. 2.9(c).

In order to make orientation gradients within the grains visible we colorize the orientations in Figure 2.10 according to their misorientation angle and axis with respect to the grain mean orientation, cf. [22]. We observe some sharp subgrain boundaries in the upper left grain and some horizontal and vertical artifacts in the lower left grains, which most likely originate from the measurement process.

Subsequently, we applied the denoising methods discussed in this paper to this example data set. Since grain boundaries reconstructed by the Voronoi decomposition based algorithm may pass through pixels which then can not be assigned to a certain grain anymore, we decided to leaving all those pixels crossed by a grain boundary unfilled. For the mean, median and Kuwahara filter we have chosen the size of the sliding window sufficiently large such that most of the non-indexed pixels got filled. The resulting denoised maps nicely display the main characteristics of the different denoising methods. The mean filter and the smoothing splines

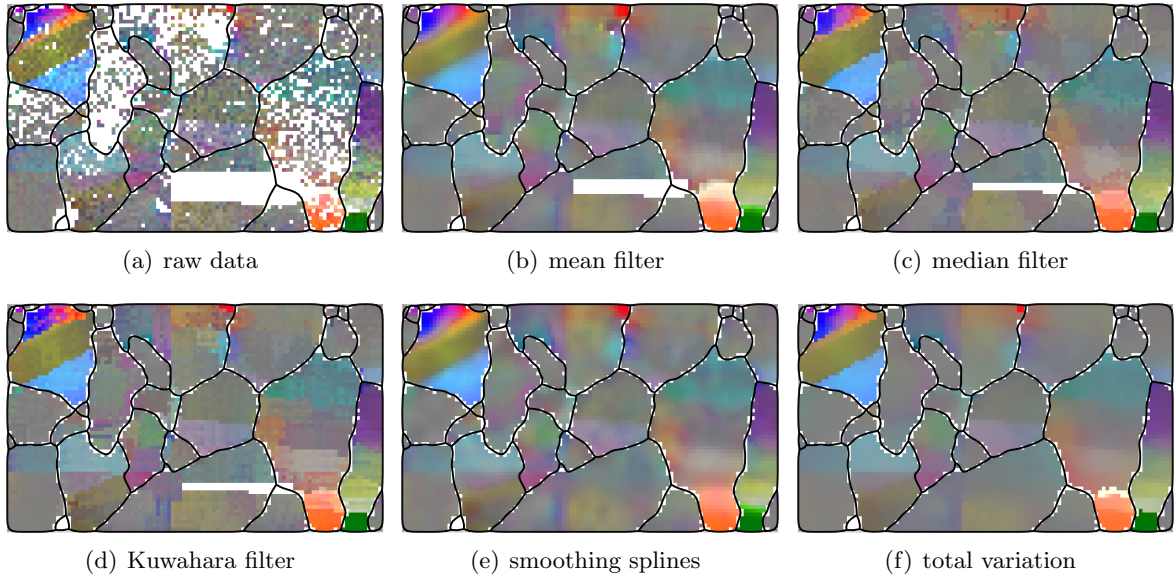


Figure 2.10: Raw and denoised EBSD data of a geological Forsterite sample colored according to misorientation angle and axis with respect to the grain mean orientation: The more the colors are saturated the higher the deviation from the mean orientation.

clearly “smooth away” subgrain boundaries, while median and Kuwahra filter create blocky artifacts and ignore small features as the small yellow region in the upper left grain. The total variation reconstruction performs a very convincing work. Subgrain boundaries remain sharp and regions with many inpaired pixels become very smooth, i.e., the algorithm does not generate an artificial structure. Moreover, the aforementioned inherent artifacts (horizontal and vertical lines) remain clearly visible as such and even become visible in formerly poorly indexed regions, e.g., in the middle grain at the upper boundary.

3 Applications

In this section we discuss the impact of different denoising methods to specific use cases.

3.1 Kernel Average Misorientation Estimation

The kernel average misorientation (KAM) is an isotropic measure for the local texture gradient and is commonly applied to evaluate the local strain or stored energy. Given an orientation map $\mathbf{O}(\mathbf{x})$ we denote by $\mathcal{N}_r^\varepsilon(\mathbf{x})$ the set of r -th order neighboring measurements \mathbf{x}' to \mathbf{x} that have a misorientation angle $\delta(\mathbf{O}(\mathbf{x}), \mathbf{O}(\mathbf{x}'))$ less than a certain threshold ε . Now, the kernel average misorientation is defined as

$$\text{KAM}(\mathbf{x}) = \frac{1}{\#\mathcal{N}_r^\varepsilon(\mathbf{x})} \sum_{\mathbf{x}' \in \mathcal{N}_r^\varepsilon(\mathbf{x})} \delta(\mathbf{O}(\mathbf{x}), \mathbf{O}(\mathbf{x}')).$$

Taking noisy data $\tilde{\mathbf{O}}(\mathbf{x}) = \mathcal{E}(\mathbf{x})\mathbf{O}(\mathbf{x})$ into account the kernel average misorientation KAM can be split into a portion KAM_{true} that originates from the true data and a portion $\text{KAM}_{\text{noise}}$




Figure 3.1: The kernel average misorientation (from third order neighbors) in degree for the exact synthetic data as defined in Section 2 (left) and for the same data corrupted by DIVP noise with half-width $\sigma = 1^\circ$ (right).

that originates from the noise

$$\begin{aligned}
\#\mathcal{N}_r^\varepsilon(\mathbf{x}) \cdot \text{KAM}(\mathbf{x}) &= \sum_{\mathbf{x}' \in \mathcal{N}_r^\varepsilon(\mathbf{x})} \delta(\mathcal{E}(\mathbf{x})\mathbf{O}(\mathbf{x}), \mathcal{E}(\mathbf{x}')\mathbf{O}(\mathbf{x}')) \\
&\leq \sum_{\mathbf{x}' \in \mathcal{N}_r^\varepsilon(\mathbf{x})} \delta(\mathcal{E}(\mathbf{x})\mathbf{O}(\mathbf{x}), \mathcal{E}(\mathbf{x})\mathbf{O}(\mathbf{x}')) + \delta(\mathcal{E}(\mathbf{x})\mathbf{O}(\mathbf{x}'), \mathcal{E}(\mathbf{x}')\mathbf{O}(\mathbf{x}')) \\
&= \sum_{\mathbf{x}' \in \mathcal{N}_r^\varepsilon(\mathbf{x})} \delta(\mathbf{O}(\mathbf{x}), \mathbf{O}(\mathbf{x}')) + \delta(\mathcal{E}(\mathbf{x}), \mathcal{E}(\mathbf{x}')) = \text{KAM}_{\text{true}}(\mathbf{x}) + \text{KAM}_{\text{noise}}(\mathbf{x}).
\end{aligned}$$

This means that the kernel average misorientation of some noisy EBSD map is approximately the kernel average misorientation of the noise-free map plus the mean angular deviation of the noise. It should be noted that in the presence of an orientation gradient the kernel average misorientation increases if the order r of neighbors is increased while the noise component remains constant. This explains why higher order KAM is less sensitive to noise than first order KAM.

Let us demonstrate these effects with the help of our synthetic example from Section 2. In the noise free case the KAM is 0 degree in the uniform portion and 0.5 degree in the linear portion, cf. Fig. 3.1 left. However, if we apply De la Vallée Poussin-distributed noise with half-width $\sigma = 1^\circ$, then the KAM for the noisy data is about 2 degree everywhere on the map, cf. Fig. 3.1 right. In particular, there is no distinction between the region with uniform and linear texture gradient.

In Figure 3.2(a) we have plotted horizontal profiles of the KAM of our synthetic data corrupted by DIVP noise with the half-width varying from $\sigma = 0.1^\circ$ to $\sigma = 1^\circ$. We observe that in the presence of noise only a local gradient is observable which is four times as large as the half-width of the noise distribution.

Next, we apply the total variation filter to the synthetic data corrupted with DIVP noise with half-width $\sigma = 1^\circ$. We may expect that the KAM reduces if the regularization parameter α is increased. This effect is illustrated in Fig. 3.2(b) and 3.3. The decrease of the KAM with increasing α appears quite moderate. For all choices of α between 0.1 and 1 we obtain a much better approximation of the true KAM than from the noisy data.

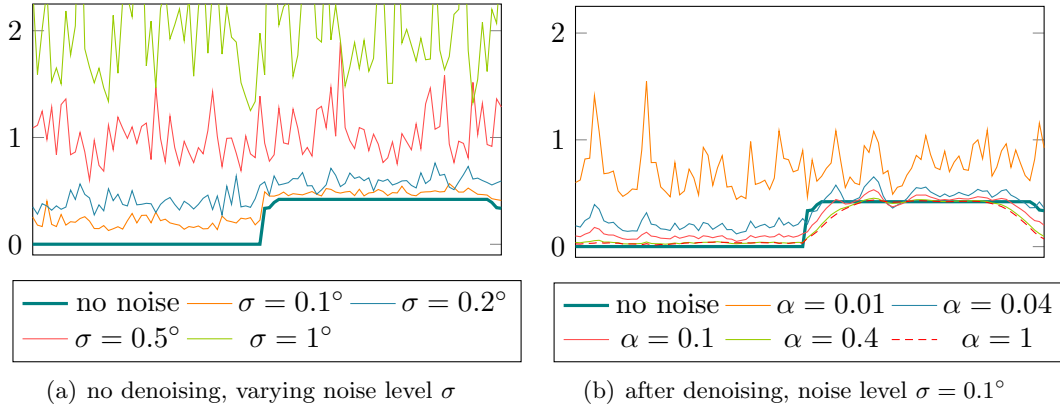


Figure 3.2: Horizontal KAM profiles of the synthetic data corrupted by DIVP noise. In (a) the noise level σ varies and no denoising is applied. In (b) noise level is $\sigma = 1^\circ$ and KAM is computed from data denoised by total variation using different regularization parameters α .




Figure 3.3: KAM maps computed from data denoised by total variation using different regularization parameters α .

3.2 Gradient and Curvature Tensor Estimation

The orientation gradient or curvature tensor describes the local change of orientation within a specimen. Let us consider the local orientation $\mathbf{O}(\mathbf{x})$ as a smooth function of the spatial variable $\mathbf{x} = (x_1, x_2, x_3)$. Then, the partial derivative with respect to the a -th coordinate is defined as

$$\frac{\partial}{\partial x_a} \mathbf{O}(\mathbf{x}) = \lim_{h \rightarrow 0} \frac{\mathbf{O}(\mathbf{x} + h \mathbf{e}_a) - \mathbf{O}(\mathbf{x})}{h}. \quad (3.1)$$

It is important to understand that the matrix $\frac{\partial}{\partial x_a} \mathbf{O}(\mathbf{x})$ just like the orientation $\mathbf{O}(\mathbf{x})$ mixes crystal and specimen coordinates. Representing $\frac{\partial}{\partial x_a} \mathbf{O}(\mathbf{x})$ entirely with respect to the specimen reference frame leads to the skew-symmetric matrix

$$\mathbf{W}^a(\mathbf{x}) = \left(\frac{\partial}{\partial x_a} \mathbf{O}(\mathbf{x}) \right) \mathbf{O}(\mathbf{x})^\top, \quad (3.2)$$

which is related to the lattice misorientation vector $\boldsymbol{\theta}^a(\mathbf{x})$ per unit distance by the equation

$$\theta_i^a(\mathbf{x}) = \frac{1}{2} \epsilon_{ijk} W_{jk}^a(\mathbf{x}), \quad (3.3)$$

where ϵ_{ijk} is Levi-Civita's permutation symbol and we have assumed summation over all indices appearing twice. The vector $\boldsymbol{\theta}^a(\mathbf{x})$ can be interpreted as the misorientation axis between the orientations $\mathbf{O}(\mathbf{x} + h \mathbf{e}_a)$ and $\mathbf{O}(\mathbf{x})$ in specimen coordinates scaled by the misorientation angle per distance h for the limiting process $h \rightarrow 0$, cf. [16].

Following the latter characterization, the vectors $\boldsymbol{\theta}^{x_1}(\mathbf{x}^{m,n})$ and $\boldsymbol{\theta}^{x_2}(\mathbf{x}^{m,n})$ can be determined approximately from discrete orientation data $\mathbf{O}(\mathbf{x}^{m,n})$ at some two-dimensional regular grid $\mathbf{x}^{m,n}$ ($m = 1, 2, \dots; n = 1, 2, \dots$) by the products

$$\begin{aligned}\boldsymbol{\theta}^{x_1}(\mathbf{x}^{m,n}) &\approx \frac{\delta(\mathbf{O}(\mathbf{x}^{m,n}), \mathbf{O}(\mathbf{x}^{m+1,n}))}{\Delta x_1} \cdot \boldsymbol{\eta}(\mathbf{O}(\mathbf{x}^{m,n}), \mathbf{O}(\mathbf{x}^{m+1,n})), \\ \boldsymbol{\theta}^{x_2}(\mathbf{x}^{m,n}) &\approx \frac{\delta(\mathbf{O}(\mathbf{x}^{m,n}), \mathbf{O}(\mathbf{x}^{m,n+1}))}{\Delta x_2} \cdot \boldsymbol{\eta}(\mathbf{O}(\mathbf{x}^{m,n}), \mathbf{O}(\mathbf{x}^{m,n+1})),\end{aligned}$$

where $\Delta x_1, \Delta x_2$ are the step sizes of the measurement grid, $\delta(\mathbf{O}_1, \mathbf{O}_2)$ is the misorientation angle and $\boldsymbol{\eta}(\mathbf{O}_1, \mathbf{O}_2)$ is the misorientation axis between two orientations in specimen coordinates (cf. Sec. 1.1).³ It should be noted that other approximations of $\boldsymbol{\theta}^{x_1}(\mathbf{x}^{m,n})$ are possible, e.g. symmetric misorientations between $\mathbf{O}(\mathbf{x}^{m+1,n})$ and $\mathbf{O}(\mathbf{x}^{m-1,n})$.

The lattice misorientation vectors per unit distance given by $\boldsymbol{\theta}^{x_1}, \boldsymbol{\theta}^{x_2}, \boldsymbol{\theta}^{x_3}$ with respect to the specimen directions x_1, x_2 and x_3 define the lattice curvature

$$\boldsymbol{\kappa}(\mathbf{x}) = \begin{pmatrix} \theta_1^{x_1}(\mathbf{x}) & \theta_1^{x_2}(\mathbf{x}) & \theta_1^{x_3}(\mathbf{x}) \\ \theta_2^{x_1}(\mathbf{x}) & \theta_2^{x_2}(\mathbf{x}) & \theta_2^{x_3}(\mathbf{x}) \\ \theta_3^{x_1}(\mathbf{x}) & \theta_3^{x_2}(\mathbf{x}) & \theta_3^{x_3}(\mathbf{x}) \end{pmatrix}. \quad (3.4)$$

In the case of orientation data measured on a two-dimensional grid $\boldsymbol{\theta}^{x_3}$ remains unknown.

As an example we consider an EBSD data set of interstitial free (IF) steel as presented in [12]. The orientations colored by an inverse pole figure (IPF) map are displayed in Figure 3.4(a). The white pixels correspond to non-indexed orientations. Figure 3.4(b) displays the same data but colored according to the misorientation axis and the misorientation angle with respect to the grain mean orientation as explained in [22]. The latter one clearly shows that the orientation measurements are effected by some noise. Applying a total variation based filter for noise removal we end up with the orientation map depicted in the Figures 3.4(c) and 3.4(d).

Although the raw data appears to be acceptable, cf. Fig. 3.4(a), it is still too noisy to compute meaningful values for the curvature as shown in Figure 3.5(a). The reason is that taking the orientation gradient significantly amplifies the noise in the data. However, when the curvature is computed from the denoised data it may allow for interpretation, cf. Fig. 3.5(b).

3.3 Dislocation Density Estimation

The dislocation density tensor characterizes the dislocation state of the crystal lattice. To be precise, $\boldsymbol{\alpha}$ how big the local closure failure of the crystal is and how much dislocations are geometrically necessary in order to preserve the compatibility. Under the assumptions of small plastic deformations and negligible lattice strain gradients the *lattice curvature tensor* $\boldsymbol{\kappa}(\mathbf{x})$ at a position \mathbf{x} can be related to the dislocation density tensor $\boldsymbol{\alpha}(\mathbf{x})$ by

$$\boldsymbol{\alpha} = \boldsymbol{\kappa}^T - \text{tr}(\boldsymbol{\kappa})\mathbf{I} \quad \Leftrightarrow \quad \boldsymbol{\kappa} = \boldsymbol{\alpha}^T - \frac{1}{2} \text{tr}(\boldsymbol{\alpha})\mathbf{I} \quad (3.5)$$

³We prefer this definition of a discrete gradient as equ. (3.2) together with a finite difference approximation of equ. (3.1) does not lead to a skew symmetric matrix.

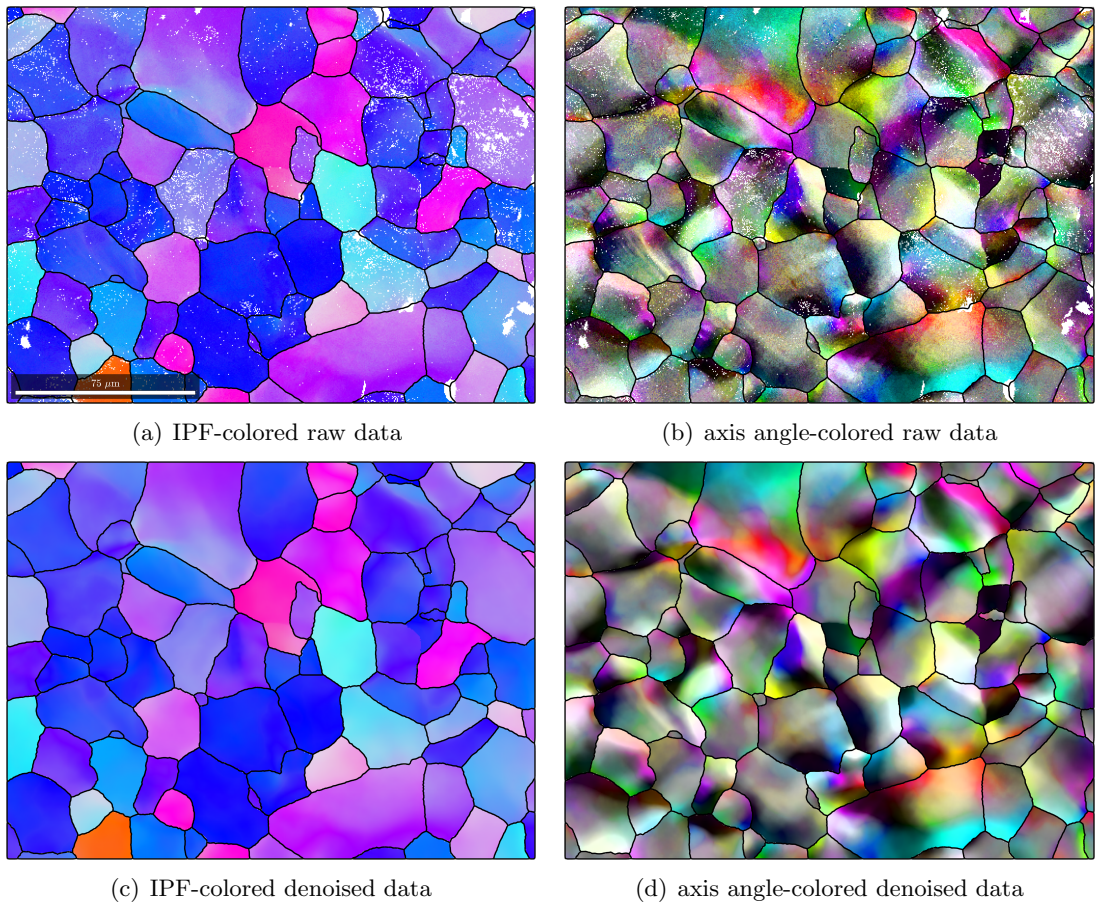


Figure 3.4: Orientation map of interstitial free steel [12] before and after denoising.

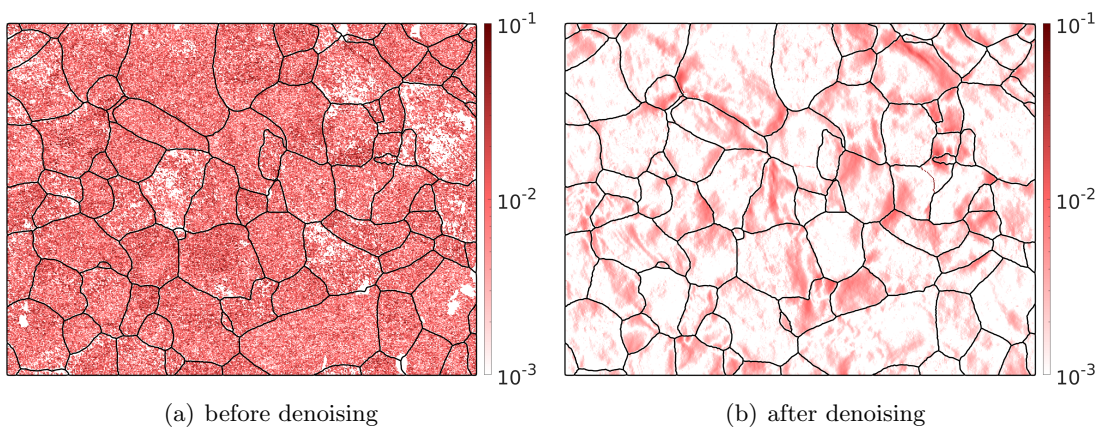


Figure 3.5: First coefficient κ_{11} (unit $1/\mu\text{m}$) of the lattice curvature tensor of the EBSD data depicted in Fig. 3.4(a) (left) and in Fig. 3.4(c) (right).




Figure 3.6: GND density computed (a) from the raw (Hough based) EBSD map, (b) after denoising and (c) using high resolution EBSD data [12].

cf. [17]. Given that only the first two columns of the lattice curvature tensor $\boldsymbol{\kappa}$ are known the following coefficients of the dislocation density tensor can be computed: $\alpha_{12}, \alpha_{13}, \alpha_{21}, \alpha_{23}, \alpha_{33}$ as well as the difference $\alpha_{11} - \alpha_{22}$.

One approach to determine the missing coefficients in the lattice curvature tensor is to represent the dislocation density tensor as the weighted sum

$$\boldsymbol{\alpha}(\mathbf{x}) = \sum_{k=1}^K \rho_k(\mathbf{x}) \mathbf{b}_k \otimes \mathbf{l}_k \quad (3.6)$$

of screw and edge dislocations at preferred slip systems described by its Burgers vectors \mathbf{b}_k and line tangent vectors \mathbf{l}_k .⁴ Using the relationship (3.5) we obtain for the lattice curvature tensor the representation

$$\boldsymbol{\kappa}(\mathbf{x}) = \sum_{k=1}^K \rho_k(\mathbf{x}) \left(\mathbf{l}_k \otimes \mathbf{b}_k - \frac{1}{2} \text{tr}(\mathbf{b}_k \otimes \mathbf{l}_k) \mathbf{I} \right). \quad (3.7)$$

However, even in the case of complete knowledge of the lattice curvature tensor $\boldsymbol{\kappa}$ the decomposition (3.7) can not be uniquely recovered, given that sufficiently many slip systems are considered. For this reason one imposes as an additional condition that some measure of the total energy of dislocations, e.g.

$$U(\mathbf{x}) = \sum_{k=1}^K \rho_k(\mathbf{x}) u_k.$$

should be minimized at each position \mathbf{x} , cf. [17, 24], where u_k denote the energy of the k -th dislocation type (measured per dislocation density).

Applying this approach to the curvatures computed in the previous section we obtain the GND densities depicted in Figure 3.6. The GND densities computed before denoising are completely corrupt while the GND densities computed after denoising seem physically reasonable and, in the first place, interpretable.

In [12] the Kikuchi pattern of the sample discussed above have also been analyzed by a cross-correlation method. The resulting high resolution EBSD data has then been used

⁴From the physical point of view, we seek for an equivalent composition of elementary dislocation segments such that the same dislocation state is obtained.

to compute the GND densities depicted in Figure 3.6(c). Although the latter ones offers much more detail they still compare well with the GND densities computed from the Hough transform-based EBSD data after denoising. We also observe that blank regions in Figure 3.6(c) which correspond to non-indexed orientations correspond to regions of high GND in 3.6(b), which is reasonable as those regions are hard to index.

3.4 Denoising High Resolution EBSD Data

High angular resolution EBSD data originate from improved methods for interpreting Kikuchi pattern and most commonly from correlation or dictionary based methods. In those methods the orientation $\mathbf{O}(\mathbf{x})$ of a Kikuchi pattern $P(\mathbf{x})$ at position \mathbf{x} is determined as the solution of the maximization problem

$$C(P(\mathbf{x}), M(\mathbf{O}(\mathbf{x}))) \rightarrow \max$$

where $M(\mathbf{O}(\mathbf{x}))$ denotes a simulated master pattern rotated in orientation $\mathbf{O}(\mathbf{x})$ and $C(P, M)$ denotes the cross-correlation between a measured pattern P with a master pattern M . Such methods have been proven to be much more accurate than traditional Hough transform based indexing methods and, hence, denoising methods have not been considered yet. However, it happens quite often that the maximization problem leads to multiple peaks and the indexing algorithm has to decide for one particular maximum. In those cases a combination with the denoising methods presented in this paper could serve as an additional criteria for choosing the right orientation and increase the accuracy of the methods.

More precisely, we suggest to combine the cross-correlation term with a regularization term in the following form

$$\sum_{\mathbf{x}} C(P(\mathbf{x}), M(\mathbf{O}(\mathbf{x}))) - \lambda \|\nabla \mathbf{O}(\mathbf{x})\|_2 \rightarrow \max ,$$

which is essentially the minimization problem (2.8) but with the correlation term replacing the misorientation to a orientation determined by another indexing method. According to this approach orientations are chosen such that the simulated pattern have a high correlation with the measured pattern and simultaneously minimize the total variation functional, i.e. vary smoothly in spatial domain.

The disadvantage of this approach is that it is computational very expensive. It is subject to future research to find algorithms that manage to solve this optimization problem in reasonable time.

Acknowledgments

The authors want to thank B. Britton, Imperial College London, for kindly providing the EBSD data set of interstitial free steel discussed in the sections 3.2 and 3.3. Furthermore, the authors gratefully acknowledge the help of J. Ihlemann, T. Lampke and D. Dietrich.

This research was supported by German Science Foundation (DFG) within the Collaborative Research Center SFB 692 HALS.

References

- [1] M. Bačák, R. Bergmann, G. Steidl, and A. Weinmann. A second order non-smooth variational model for restoring manifold-valued images. *Preprint Univ. Kaiserslautern*, 2015.
- [2] F. Bachmann, R. Hielscher, P. E. Jupp, W. Pantleon, H. Schaeben, and E. Wegert. Inferential statistics of electron backscatter diffraction data from within individual crystalline grains. *J. Appl. Cryst.*, 43:1338 – 1355, 2010.
- [3] F. Bachmann, R. Hielscher, and H. Schaeben. Grain detection from 2d and 3d EBSD data – specification of the MTEX algorithm. *Ultramicroscopy*, 111:1720 – 1733, 2011.
- [4] R. Bergmann, R. H. Chan, R. Hielscher, J. Persch, and G. Steidl. Restoration of manifold-valued images by half-quadratic minimization. *Inverse Prob. Imaging*, 10:281–304, 2016.
- [5] R. Bergmann, F. Laus, G. Steidl, and A. Weinmann. Second order differences of cyclic data and applications in variational denoising. *SIAM Journal on Imaging Sciences*, 7(4):2916–2953, 2014.
- [6] R. Bergmann and A. Weinmann. A second-order tv-type approach for inpainting and denoising higher dimensional combined cyclic and vector space data. *Journal of Mathematical Imaging and Vision*, 55(3):401–427, Jul 2016.
- [7] D. Chen and J.-C. Kuo. Bilateral filter based orientation smoothing of ebsd data. *Ultramicroscopy*, 110, 06 2010.
- [8] D. Garcia. Robust smoothing of gridded data in one and higher dimensions with missing values. *Computational Statistics & Data Analysis*, 54:1167 – 1178, 2010.
- [9] A. Godfrey. Edge preservation near triple junctions during orientation averaging of ebsp data. *Scripta Materialia - SCRIPTA MATER*, 50:1097–1101, 04 2004.
- [10] V. K. Gupta and S. R. Agnew. A simple algorithm to eliminate ambiguities in ebsd orientation map visualization and analyses: Application to fatigue crack-tips/wakes in aluminum alloys. *Microscopy and Microanalysis*, 16:831, 2010.
- [11] V. K. Gupta and S. R. Agnew. A Simple Algorithm to Eliminate Ambiguities in EBSD Orientation Map Visualization and Analyses: Application to Fatigue Crack-Tips/Wakes in Aluminum Alloys. *Microscopy and Microanalysis*, 16:831–841, 2010.
- [12] J. L. R. Hickey, S. Rouland, and T. B. Britton. Heterogeneous local plastic deformation revealed using in-situ tensile testing and high angular resolution electron backscatter diffraction. *arXiv*, 2018.
- [13] R. Hielscher, F. Bachmann, D. Mainprice, and R. Kilian. MTEX 5.1 - A texture calculation toolbox. <http://mtex-toolbox.github.io>, 2018.
- [14] F. J. Humphreys, P. S. Bate, and P. J. Hurley. Orientation averaging of electron backscattered diffraction data. *Journal of microscopy*, 201:50–8, 02 2001.

- [15] A. Kobler, A. Kashiwar, H. Hahn, and C. Kübel. Combination of in situ straining and ACOM TEM: A novel method for analysis of plastic deformation of nanocrystalline metals. *Ultramicroscopy*, 128:68–81, Jan. 2013.
- [16] A. Morawiec. *Orientations and Rotations*. Springer, 2004.
- [17] W. Pantleon. Resolving the geometrically necessary dislocation content by conventional electron backscattering diffraction. *Scripta Materialia*, 58:p994–997, 2008.
- [18] H. F. Poulsen and X. Fu. Generation of grain maps by an algebraic reconstruction technique. *Journal of Applied Crystallography*, 36(4):1062–1068, July 2003.
- [19] E. Rauch, J. Portillo, S. Nicolopoulos, D. Bultreys, S. Rouvimov, and P. Moeck. Automated nanocrystal orientation and phase mapping in the transmission electron microscope on the basis of precession electron diffraction. *Zeitschrift für Kristallographie*, 225(2-3):103–109, 2010.
- [20] H. Schaeben. The de la Vallée Poussin standard orientation density function. *Textures and Microstructures*, 33:365 – 373, 1999.
- [21] S. Setzer and G. Steidl. Variational methods with higher order derivatives in image processing. In *Approximation XII: San Antonio 2007*, pages 360–385, 2008.
- [22] K. Thomsen, K. Mehnert, P. W. Trimby, and A. Gholinia. Quaternion-based disorientation coloring of orientation maps. *Ultramicroscopy*, 182:62 – 67, 2017.
- [23] A. Weinmann, L. Demaret, and M. Storath. Total variation regularization for manifold-valued data. *SIAM Journal on Imaging Sciences*, 7(4):2226–2257, 2014.
- [24] A. J. Wilkinson and D. Randman. Determination of elastic strain fields and geometrically necessary dislocation distributions near nanoindents using electron back scatter diffraction. *Philosophical Magazine*, 90(9):1159–1177, 2010.
- [25] S. Zaefferer. A critical review of orientation microscopy in SEM and TEM. *Crystal Research and Technology*, 46(6):607–628, June 2011.