



Denoising of crystal orientation maps

R. Hielscher, C. B. Silbermann, E. Schmidl and Joern Ihlemann

J. Appl. Cryst. (2019). **52**, 984–996



IUCr Journals
CRYSTALLOGRAPHY JOURNALS ONLINE

Copyright © International Union of Crystallography

Author(s) of this article may load this reprint on their own web site or institutional repository provided that this cover page is retained. Republication of this article or its storage in electronic databases other than as specified above is not permitted without prior permission in writing from the IUCr.

For further information see <http://journals.iucr.org/services/authorrights.html>



Denoising of crystal orientation maps

R. Hielscher,^{a*} C. B. Silberman,^b E. Schmidl^a and Joern Ihlemann^b

^aChair of Applied Functional Analysis, Chemnitz University of Technology, 09126 Chemnitz, Germany, and ^bChair of Solid Mechanics, Chemnitz University of Technology, 09126 Chemnitz, Germany. *Correspondence e-mail: ralf.hielscher@mathematik.tu-chemnitz.de

Received 25 September 2018

Accepted 25 June 2019

Edited by Th. Proffen, Oak Ridge National Laboratory, USA

Keywords: electron backscatter diffraction; measurement errors; denoising; inpainting; kernel average misorientation; geometrically necessary dislocations.

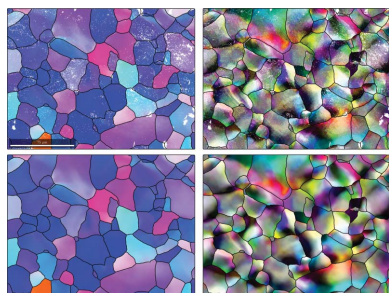
This paper compares several well known sliding-window methods for denoising crystal orientation data with variational methods adapted from mathematical image analysis. The variational methods turn out to be much more powerful in terms of preserving low-angle grain boundaries and filling holes of non-indexed orientations. The effect of denoising on the determination of the kernel average misorientation and the geometrically necessary dislocation density is also discussed. 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 electron backscatter diffraction maps.

1. Introduction

Individual crystal orientations are commonly obtained from light microscopy (Heilbronner & Barrett, 2014; Passchier & Trouw, 2005), from electron backscatter diffraction (EBSD) by indexing of Kikuchi patterns (Zaeerer, 2011), from automated crystal orientation mapping inside a transmission electron microscope (ACOM-TEM) by indexing spot diffraction patterns (Rauch *et al.*, 2010) or from X-ray diffraction (Poulsen & Fu, 2003). As with any other experimental data these crystal orientations are usually corrupted by errors.

In the case of orientation determined from EBSD, several different sources of measurement error superpose (Kamaya, 2011). An imprecise alignment of the specimen or an imprecise pattern center usually results in a systematic error. Errors arising during orientation indexing either by conventional band detection or by pattern-matching methods are often described as random errors (*cf.* Hielscher, Britton & Bartel, 2018). These random errors are caused either by a noisy Kikuchi pattern in combination with a limited resolution in Hough space, for the band-detection methods, or a limited resolution in the orientation space, for pattern-matching approaches. While small errors in orientation maps may appear negligible for the analysis of the texture or the microstructure, they do amplify significantly if the orientation data are used for subsequent calculations of kernel averaged misorientation (KAM), misorientation axes (Prior, 1999; Wilkinson, 2001), orientation curvature tensors or geometrically necessary dislocation (GND) densities (Konijnenberg *et al.*, 2015).

In this paper we aim at reducing the random error in orientation maps. The key requirement to be able to separate the noise from the 'true' orientation map is the assumption that the 'true' orientations are spatially dependent, while the noise is spatially independent. This means that the error in one



pixel is not correlated with the error in some adjacent pixel, while at the same time the orientation map has been measured at a sufficiently fine grid, *i.e.* such that there are no one-pixel features.

Denoising of EBSD data has already been addressed in several papers (Humphreys *et al.*, 2001; Godfrey, 2004; Chen & Kuo, 2010; Gupta & Agnew, 2010; Kamaya, 2010; Wright *et al.*, 2015). In the current paper we compare well established methods like the mean filter, the Kuwahara filter (Humphreys *et al.*, 2001; Godfrey, 2004) and the median filter (Gupta & Agnew, 2010) with variational denoising techniques (Weinmann *et al.*, 2014; Bacak *et al.*, 2015; Bergmann *et al.*, 2016, 2018; Bergmann & Weinmann, 2016) 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 ‘true’ orientation map and should favor smooth changes of the crystal orientation disrupted by a few grain and subgrain boundaries. Note that in this context ‘smoothness’ is directly related to the spatial resolution of the orientation map. If the spatial resolution is too low, microstructural features of interest might have the appearance of noise.

The paper is divided into two parts: In the first part, we discuss the basic properties of the aforementioned denoising techniques and illustrate them with synthetic data as well as with an experimental EBSD orientation map. In the second part, we discuss the impact of these denoising techniques on 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 noisy orientation maps.

All figures and numerical experiments have been done using *MTEX 5.1.1* (Hielscher, Bachmann, Mainprice & Kilian, 2018) and the script files that are 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 $\mathbf{O}_1, \mathbf{O}_2$, their 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 the disorientation. Accord-

ingly, the rotational angle δ and the rotational axis $\boldsymbol{\eta}_C$ of \mathbf{M} are called the 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 disorientation vector

$$\boldsymbol{\theta} = \delta\boldsymbol{\eta}. \quad (1)$$

This disorientation vector will play an important role throughout this paper as it allows us to represent elements of the curved orientation space locally by ordinary three-dimensional vectors. 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 guarantees that orientation \mathbf{O}_1 and the reference orientation \mathbf{O}_2 are sufficiently close together to allow for a good approximation. Approximating 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 the (matrix) logarithm and the inverse mapping as the (matrix) exponential (*cf.* Morawiec, 2004). Note that in contrast to the disorientation axis the computation of the disorientation vector (1) is stable with respect to measurement errors.

2. Comparison of denoising methods

2.1. An artificial orientation map

In this section we construct a 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. 1(a)] and contains several flat regions, a 2° low-angle grain boundary, a small oscillating part with an amplitude of 3° , an 18° high-angle grain boundary and a section with a constant orientation gradient of 0.6° per pixel. For simplicity, all these features have been integrated into the field of the second Euler angle Φ , while the first and third Euler angles are set to $\varphi_1 = \varphi_2 = 0$.

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

$$\tilde{f}(\mathbf{x}) = f(\mathbf{x}) + \varepsilon(\mathbf{x}), \quad (2)$$

where $f(\mathbf{x})$ represents the noise-free data and $\varepsilon(\mathbf{x})$ is the noise, *e.g.* an independent random sample of a Gaussian distribution. For orientation data $\mathbf{O}(\mathbf{x})$ we have to replace the summation in (2) by the multiplication with some random orientations $\boldsymbol{\mathcal{E}}(\mathbf{x})$, *i.e.*

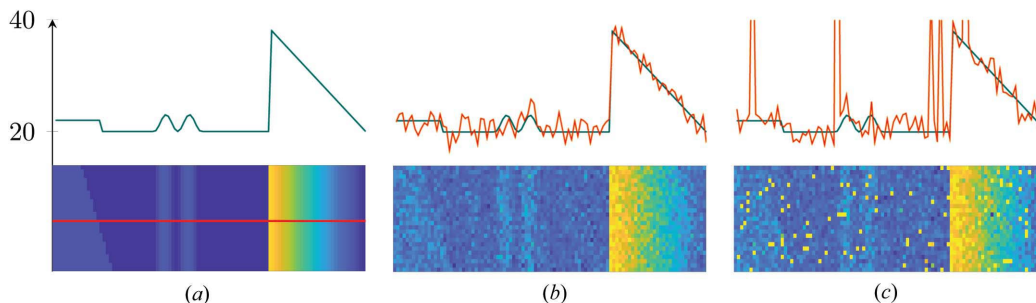


Figure 1

Synthetic orientation maps colorized according to the second Euler angle. (a) Synthetic, noise-free orientation data. (b) After adding De la Vallée Poussin noise with half-width 2° . (c) After adding 5% impulsive noise too. The dimensions of the maps are 100×25 pixels. The upper plots display the second Euler angles of a horizontal section (red line) through the middle of the orientation maps.

$$\tilde{\mathbf{O}}(\mathbf{x}) = \mathcal{E}(\mathbf{x})\mathbf{O}(\mathbf{x}). \quad (3)$$

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.

2.1.1. Bingham- and De la Vallée Poussin-distributed noise. A generalization of the Gaussian distribution to orientations is the Bingham distribution (Bachmann *et al.*, 2010). To formulate the corresponding density function we will make use of the parameterization of disorientations by quaternions $\mathbf{q} = (q_0, q_1, q_2, q_4)^T$, which are connected to the disorientation angle δ and the disorientation axis $\boldsymbol{\eta}$ by $q_1 = \cos(\delta/2)$ and $(q_2, q_3, q_4)^T = \sin(\delta/2)\boldsymbol{\eta}$. The density function $g(\mathbf{q})$ of the Bingham distribution on the orientation space then reads

$$g(\mathbf{q}) = C \exp(\mathbf{q}^T \mathbf{A} \mathbf{q}) \quad (4)$$

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

A suitable approximation for the anisotropic Bingham distribution is given by the following De la Vallée Poussin (DIVP) distribution (Schaeben, 1999):

$$g(\mathbf{q}) = \frac{B(\frac{3}{2}, \frac{1}{2})}{B(\frac{3}{2}, \kappa + \frac{1}{2})} q_1^{2\kappa}, \quad (5)$$

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 (5), the numerically challenging computation of the normalization constant of distribution (4) can be omitted. Fig. 1(b) displays the artificial orientation map corrupted by DIVP-distributed noise with a half-width of 2° . For simulating noisy orientation data according to (3) we have to draw a random sample $\mathcal{E}(\mathbf{x})$ from a given density function g . This can be done numerically in *MTEX* and is explained in more detail by Hielscher (2013).

2.1.2. Impulsive noise. While the generalization of Gaussian noise discussed in the previous section models small deviations from the true orientation due to indexing errors, the noise model in this section 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 first by Gaussian-like noise and second by impulsive noise with $\alpha = 0.05$ – is depicted in Fig. 1(c).

2.2. Definition of a filter

Denoising 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

$$\hat{\mathbf{O}} = \mathcal{F}\tilde{\mathbf{O}}. \quad (6)$$

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

Implicitly, every filter makes some assumptions about the structure of the noise as well as about the ‘true’ orientation map. In this paper, we do not discuss mathematical details of these assumptions. Instead, we directly demonstrate the effect as we analyze different types of filters with respect to different types of noise and different kinds of features in orientation maps.

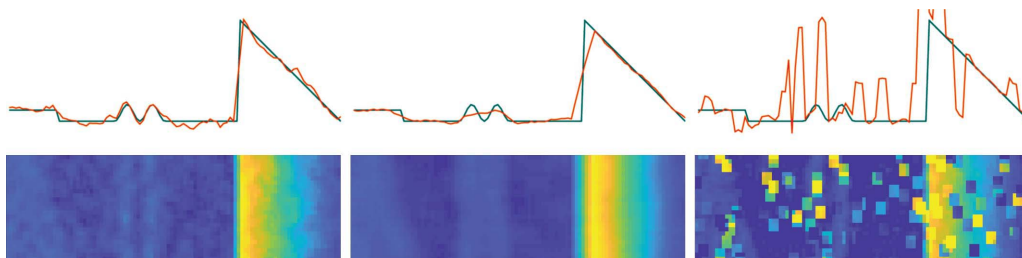
2.3. The mean filter

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 mean filter is defined as

$$\mathcal{F}_{\text{mean}}\tilde{\mathbf{O}}(\mathbf{x}) = \frac{1}{\#\mathcal{N}_r(\mathbf{x})} \sum_{\mathbf{x}' \in \mathcal{N}_r(\mathbf{x})} \tilde{\mathbf{O}}(\mathbf{x}'), \quad (7)$$

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} .

Note that equation (7) is formally not correct. Owing to the curved non-Euclidean geometry of the orientation space, the


Figure 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). The green curves show noise-free data.

sum of orientations is, 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 by Bachmann *et al.* (2010) or the disorientation vectors defined in (1). Both methods have been used for the numerical results presented in this paper. However, in order to keep the notation simple in the present section we will stick to the notation in (7), keeping in mind that its actual computation is slightly more involved.

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

$$\mathcal{F}_{\text{mean}} \tilde{\mathbf{O}}(\mathbf{x}) = \frac{\sum_{\mathbf{x}' \in \mathcal{N}_r(\mathbf{x})} w(\|\mathbf{x} - \mathbf{x}'\|) \tilde{\mathbf{O}}(\mathbf{x}')}{\sum_{\mathbf{x}' \in \mathcal{N}_r(\mathbf{x})} w(\|\mathbf{x} - \mathbf{x}'\|)}. \quad (8)$$

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 obtain formula (7). Setting

$$\omega(d) = \exp(-d^2/\sigma^2), \quad (9)$$

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 disorientation angle between $\tilde{\mathbf{O}}(\mathbf{x})$ 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. Fig. 2 shows the application of the mean filter to our synthetic orientation map Fig. 1(a). 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 also increases the blurring. The blurring effect becomes particularly visible at the grain boundaries and around the small oscillating feature.

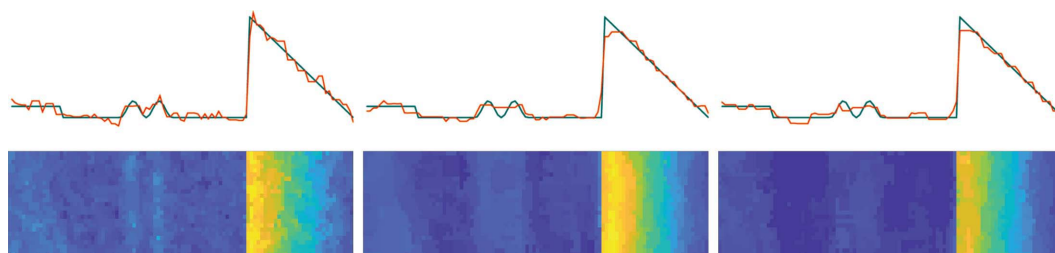
While the mean filter behaves well for Gaussian-type noise, it cannot handle impulsive noise. This is illustrated in the right subfigure of Fig. 2. The reason is that the mean filter averages equally over all orientations and is not able to identify outliers.

2.4. The median filter

Let us again denote by $\mathcal{N}(\mathbf{x})$ 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 Gupta & Agnew (2010), 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 $\mathbf{x}' \in \mathcal{N}(\mathbf{x})$ such that it possesses the minimum mean distance to all other orientations $\tilde{\mathbf{O}}(\tilde{\mathbf{x}})$, $\tilde{\mathbf{x}} \in \mathcal{N}(\mathbf{x})$, in the neighborhood, *i.e.* $\mathbf{x}' \in \mathcal{N}(\mathbf{x})$ is chosen such that for any other choice $\mathbf{x}'' \in \mathcal{N}(\mathbf{x})$ the following inequality is satisfied:

$$\sum_{\tilde{\mathbf{x}} \in \mathcal{N}(\mathbf{x})} \delta[\tilde{\mathbf{O}}(\mathbf{x}'), \tilde{\mathbf{O}}(\tilde{\mathbf{x}})] \leq \sum_{\tilde{\mathbf{x}} \in \mathcal{N}(\mathbf{x})} \delta[\tilde{\mathbf{O}}(\mathbf{x}''), \tilde{\mathbf{O}}(\tilde{\mathbf{x}})]. \quad (10)$$

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


Figure 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).

angular deviation, a confidence index or some cross-correlation score. Then, condition (10) can be generalized to

$$\sum_{\tilde{\mathbf{x}} \in \mathcal{N}(\mathbf{x})} \frac{\omega(\mathbf{x}) \delta[\tilde{\mathbf{O}}(\mathbf{x}'), \tilde{\mathbf{O}}(\tilde{\mathbf{x}})]}{\omega(\mathbf{x}') + \omega(\tilde{\mathbf{x}})} \leq \sum_{\tilde{\mathbf{x}} \in \mathcal{N}(\mathbf{x})} \frac{\omega(\mathbf{x}'') \delta[\tilde{\mathbf{O}}(\mathbf{x}''), \tilde{\mathbf{O}}(\tilde{\mathbf{x}})]}{\omega(\mathbf{x}'') + \omega(\tilde{\mathbf{x}})}, \quad (11)$$

which favors measurements with smaller expected error over 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. 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 Fig. 3 (left). Another type of median filter that works in a comparable manner was reported by Kobler *et al.* (2013).

2.5. The Kuwahara filter

A second denoising method that is a built-in tool of much EBSD-pattern-processing software is the Kuwahara filter. Similarly to the mean and the median filters, the Kuwahara filter is 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 disorientation 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.

Because a subregion is selected before taking the mean, the Kuwahara filter shows good performance at grain boundaries but tends to generate blocky structures in the presence of Gaussian noise (*cf.* Fig. 4). The resulting maps show more details than those generated by the mean and median filters but also appear more noisy. For impulsive noise, the Kuwahara filter behaves well if the outliers are not too dense such that there is always a subregion without outliers. However, in order to make the Kuwahara filter effective for Gaussian noise, the sliding window has to be chosen to be about twice as large as

for the mean filter, which makes it likely to have at least one outlier in each subregion. As soon as this happens, the outliers generate blocky artifacts in the filtered image (*cf.* Fig. 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

$$J(\hat{\mathbf{O}}) = \sum_{\mathbf{x}} \delta[\tilde{\mathbf{O}}(\mathbf{x}), \hat{\mathbf{O}}(\mathbf{x})]^2 + \alpha \varphi(\hat{\mathbf{O}}). \quad (12)$$

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 the penalty term, ensures that the approximating image $\hat{\mathbf{O}}$ is more smooth than $\tilde{\mathbf{O}}$. The regularization parameter α allows one to adjust the solution between being very close to the original data and being very smooth.

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

$$J(\hat{\mathbf{O}}) = \sum_{\mathbf{x}} \delta[\tilde{\mathbf{O}}(\mathbf{x}), \hat{\mathbf{O}}(\mathbf{x})]^2 + \alpha \|\Delta \hat{\mathbf{O}}(\mathbf{x})\|^2. \quad (13)$$

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 straightforward. Possible definitions have been discussed by Bergmann *et al.* (2014). In this work we first approximate the orientation map $\tilde{\mathbf{O}}$ locally by a disorientation vector map $\tilde{\boldsymbol{\theta}}(\mathbf{x})$ [*cf.* equation (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.* Garcia, 2010) to compute the minimizer $\hat{\boldsymbol{\theta}}$ of the translated problem:

$$J(\hat{\boldsymbol{\theta}}) = \sum_{\mathbf{x}} \delta[\tilde{\boldsymbol{\theta}}(\mathbf{x}), \hat{\boldsymbol{\theta}}(\mathbf{x})]^2 + \alpha \|\Delta \hat{\boldsymbol{\theta}}(\mathbf{x})\|^2 \rightarrow \min. \quad (14)$$

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.

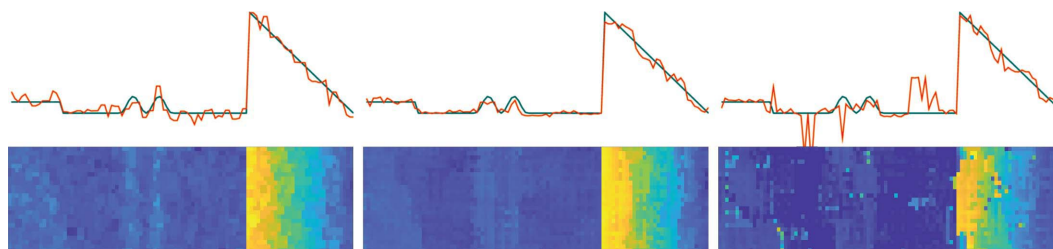


Figure 4

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

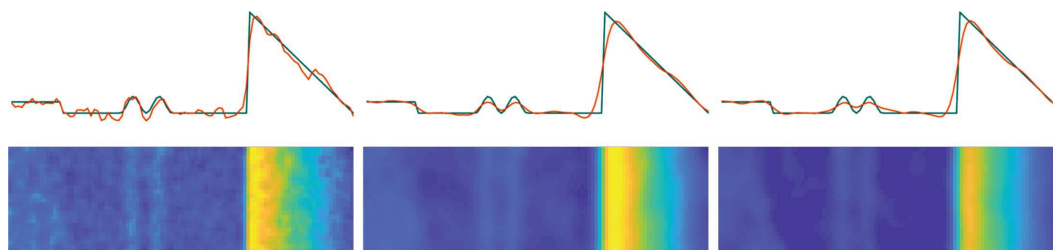


Figure 5

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

In their standard form, smoothing splines behave poorly with respect to impulsive noise. To overcome this restriction, the algorithm described by Garcia (2010) includes an adaptive reweighting scheme that assigns a lower weight to function values identified as outliers. Fig. 1(c) shows that the adaption to orientation data behaves well with respect to impulsive noise. The smoothing spline algorithm described by Garcia (2010) 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. Section 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 Fig. 5. We observe that the smoothing splines filter behaves similarly to the mean filter, *i.e.* it is very good at recovering the orientation gradient and the oscillating feature but significantly blurs the grain boundary. The overall performance appears to be better than that of the mean filter.

2.6.2. Total variation. The blurring effect of the smoothing splines is mostly due to the use of the square in the penalty term in (13). If we replace the Laplacian with the norm of the gradient and remove the square we obtain the well known isotropic total variation (TV) functional

$$J(\hat{\mathbf{O}}) = \sum_{\mathbf{x}} \delta[\hat{\mathbf{O}}(\mathbf{x}), \tilde{\mathbf{O}}(\mathbf{x})]^2 + \alpha \|\nabla \hat{\mathbf{O}}(\mathbf{x})\|_2. \quad (15)$$

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 (15) in the Euclidean setting. Among others, the following algorithms have been generalized to orientation values: an algorithm based on half-quadratic minimization (Bergmann *et al.*, 2016) and an algorithm based on proximal mappings (Weinmann *et al.*, 2014).

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 disorientations larger than the threshold, *e.g.* at (sub) grain boundaries, are not penalized. An example of such a function is

$$\varphi(\delta) = \begin{cases} (\delta^2 + \epsilon)^{1/2} & \delta < d, \\ 0 & \delta \geq d, \end{cases} \quad (16)$$

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. 6 that it can be seen as an improved version of the median filter. Similarly to the median filter it reproduces sharp grain boundaries but tends to blocky images. Both low- and high-angle grain boundaries are recovered very well. On the downside, one can see some staircasing in the linear gradient part. Also, with increasing regularization parameter the oscillating feature is replaced by a flat feature. We may conclude that the total variation filter does a very good job for orientation maps with sharp grain and subgrain boundaries but has problems recovering small oscillating features. In its current implementation the total variation filter is also not able to distinguish between grain boundaries that

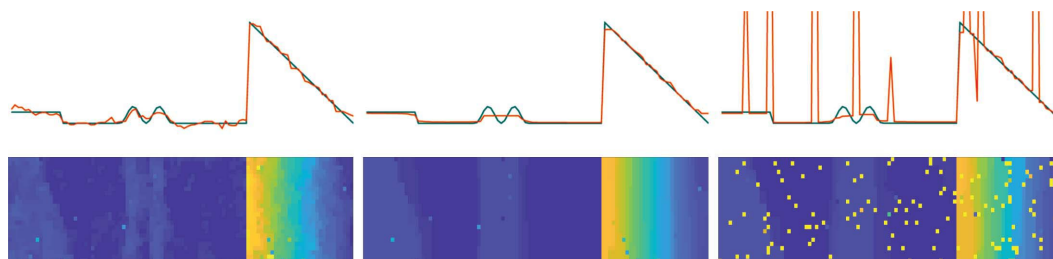


Figure 6

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

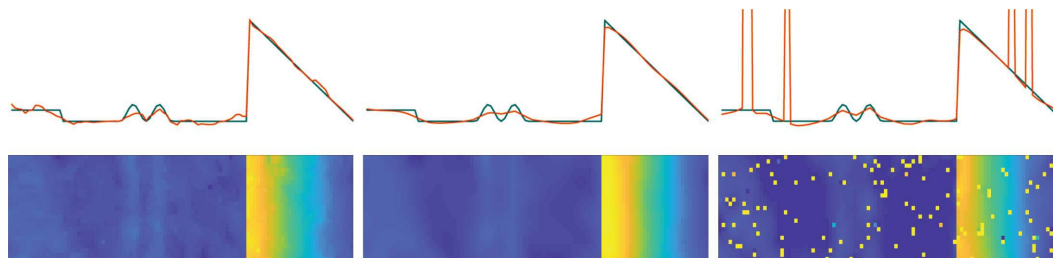


Figure 7

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

exceed the threshold δ and impulsive noise, as illustrated in Fig. 6 (right).

2.6.3. Infimal convolution. The drawback of TV regularization is that it prioritizes 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, as suggested by Bergmann *et al.* (2018). In this setting the functional

$$J(\hat{\mathbf{U}}, \hat{\mathbf{V}}) = \sum_{\mathbf{x}} \delta[\tilde{\mathbf{O}}(\mathbf{x}), \hat{\mathbf{U}}(\mathbf{x})\hat{\mathbf{V}}(\mathbf{x})]^2 + \alpha \|\nabla \hat{\mathbf{U}}(\mathbf{x})\|_2 + \beta \|\Delta \hat{\mathbf{V}}(\mathbf{x})\|. \quad (17)$$

to minimize has two arguments $\hat{\mathbf{U}}, \hat{\mathbf{V}}$, which are penalized separately but generate the final orientation map via $\hat{\mathbf{O}} = \hat{\mathbf{U}}\hat{\mathbf{V}}$. Again, half-quadratic minimization and proximal mappings are two suitable methods to find the minimizer of (17). 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. Then, we can reduce the parameter β until the staircasing effect is reduced.

Comparing the infimal convolution filter in Fig. 7 with the total variation filter in Fig. 6, we observe that the high-angle grain boundary is preserved with similar sharpness, while the small oscillating feature appears less blocky, but the low-angle grain boundary is severely smoothed.

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 subsequently filled by the respective denoising method. This method is illustrated in Fig. 8 for the synthetic orientation data given in Fig. 1(c).

Infimal convolution is not included in this list as we have not yet implemented inpainting with this method. We observe that

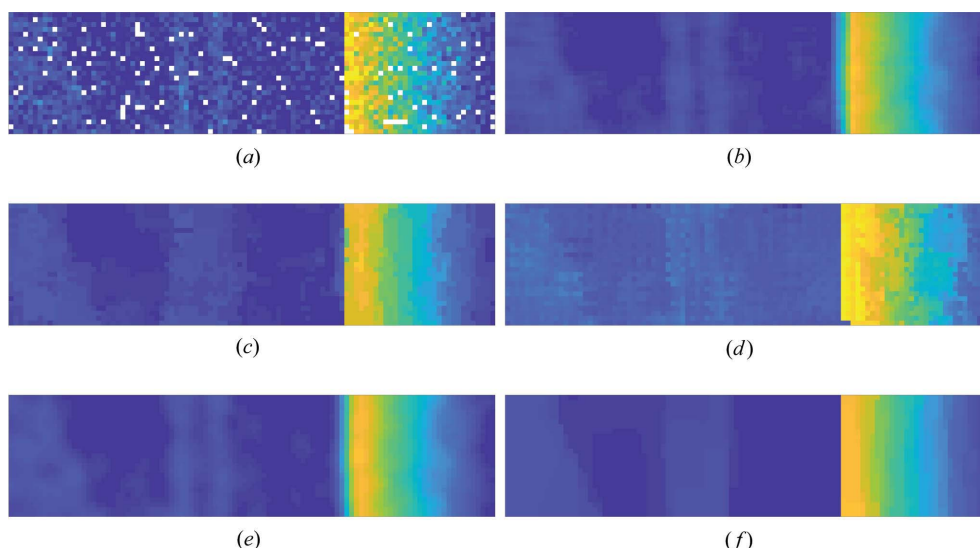


Figure 8

Simulated data corrupted by 2° DIVP noise and 5% impulsive noise, where the impulsive noise is removed as one-pixel grains and the holes are filled by the inpainting method. (a) Noisy data with one-pixel grains removed, (b) mean filter, third-order neighbors, (c) median filter, third-order neighbors, (d) Kuwahara filter, third-order neighbors, (e) smoothing splines, $\alpha = 5$ and (f) total variation, $\alpha = 0.5$.

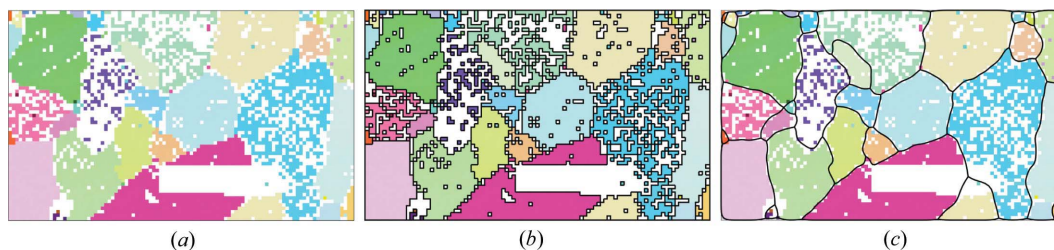


Figure 9

Grain boundaries reconstructed from poorly indexed geological forsterite data: (a) raw geological data, (b) traditional GB reconstruction and (c) advanced GB reconstruction. The standard IPF color map from *MTEX* is used (Hielscher, Bachmann, Mainprice & Kilian, 2018).

all remaining denoising methods show good performance for filling one-pixel holes. In particular, the mean filter, the Kuwahara filter and the total variation method show much better results in comparison to the case when the outliers have not been removed [cf. Figs. 2 (right), 4 (right) and 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 arbitrarily 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 apply for variational denoising techniques. This fact makes them extremely powerful for denoising purely indexed EBSD maps, *e.g.* from highly deformed materials.

To this point, only artificial, synthetic orientation data have 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 set in *MTEX* (Hielscher, Bachmann, Mainprice & Kilian, 2018). As visible in Fig. 9(a), we find a lot of non-indexed pixels in the map. Nonetheless, the grain structure is still observable. However, owing to the large number of non-indexed pixels, the traditional grain boundary (GB) reconstruction method completely fails to recover the original grain

structure [cf. Fig. 9(a)]. In such settings it turns out to be extremely useful to reconstruct the grain structure using the Voronoi decomposition algorithm as described by Bachmann *et al.* (2011). This algorithm places the grain boundaries accurately through the non-indexed regions and gives a GB reconstruction which is much closer to our expectation [cf. Fig. 9(c)].

In order to make orientation gradients within the grains visible, we colorize the orientations in Fig. 10 according to their disorientation angle and axis with respect to the grain mean orientation (cf. Thomsen *et al.*, 2017). We observe some sharp subgrain boundaries in the upper-left grain and some horizontal and vertical artifacts in the lower-left grains, which probably 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 cannot be assigned to a certain grain anymore, we decided to leave all those pixels crossed by a grain boundary unfilled. For the mean, median and Kuwahara filters we have chosen the size of the sliding window to be sufficiently large that most of the non-indexed pixels get filled. The resulting denoised maps nicely display the

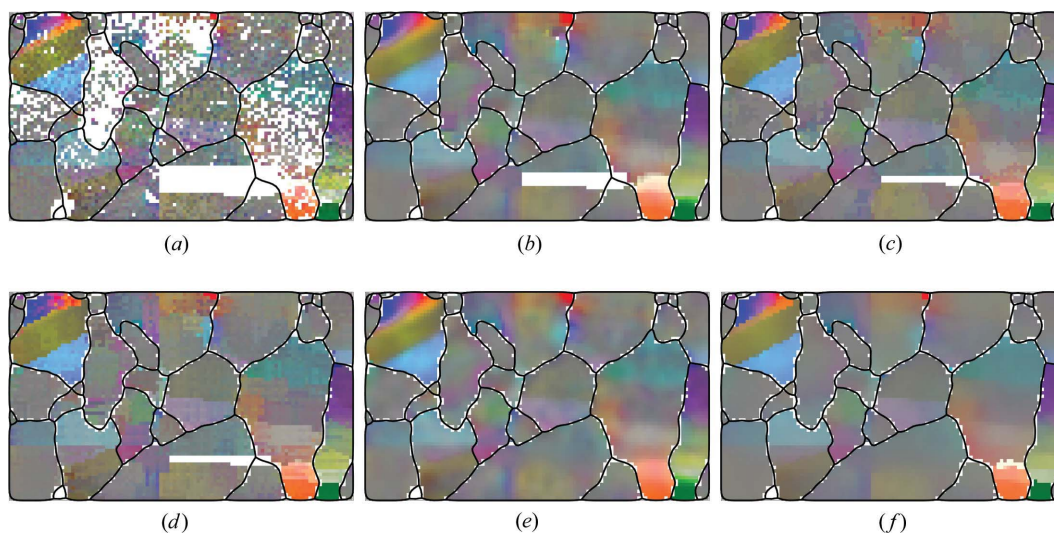


Figure 10

Raw and denoised EBSD data of a geological forsterite sample colored according to disorientation angle and axis with respect to the grain mean orientation (Thomsen *et al.*, 2017): (a) raw data, (b) mean filter, (c) median filter, (d) Kuwahara filter, (e) smoothing splines and (f) total variation. The more the colors are saturated, the higher the deviation from the mean orientation. Fully saturated colors correspond to a disorientation angle of at least 2.5° .

main characteristics of the different denoising methods. The mean filter and the smoothing splines clearly ‘smooth away’ subgrain boundaries, while the median and Kuwahra filters create blocky artifacts and ignore small features such as the small yellow region in the upper-left grain. The total variation reconstruction performs a very convincing job. Subgrain boundaries remain sharp and regions with many inpainted 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 application of different denoising methods to specific use cases.

3.1. Kernel average misorientation estimation

The 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 disorientation angle $\delta[\mathbf{O}(\mathbf{x}), \mathbf{O}(\mathbf{x}')] < \varepsilon$. Now, the kernel average misorientation is defined as

$$\text{KAM}(\mathbf{x}) = \frac{1}{\#\mathcal{N}_r^\varepsilon(\mathbf{x})} \sum_{\mathbf{x}' \in \mathcal{N}_r^\varepsilon(\mathbf{x})} \delta[\mathbf{O}(\mathbf{x}), \mathbf{O}(\mathbf{x}')]. \quad (18)$$

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

$$\begin{aligned} \#\mathcal{N}_r^\varepsilon(\mathbf{x}) \text{KAM}(\mathbf{x}) &= \sum_{\mathbf{x}' \in \mathcal{N}_r^\varepsilon(\mathbf{x})} \delta[\mathcal{E}(\mathbf{x})\mathbf{O}(\mathbf{x}), \mathcal{E}(\mathbf{x}')\mathbf{O}(\mathbf{x}')] \\ &\leq \sum_{\mathbf{x}' \in \mathcal{N}_r^\varepsilon(\mathbf{x})} \delta[\mathcal{E}(\mathbf{x})\mathbf{O}(\mathbf{x}), \mathcal{E}(\mathbf{x})\mathbf{O}(\mathbf{x}')] \\ &\quad + \delta[\mathcal{E}(\mathbf{x})\mathbf{O}(\mathbf{x}'), \mathcal{E}(\mathbf{x}')\mathbf{O}(\mathbf{x}')] \\ &= \sum_{\mathbf{x}' \in \mathcal{N}_r^\varepsilon(\mathbf{x})} \delta[\mathbf{O}(\mathbf{x}), \mathbf{O}(\mathbf{x}')] + \delta[\mathcal{E}(\mathbf{x}), \mathcal{E}(\mathbf{x}')] \\ &= \text{KAM}_{\text{true}}(\mathbf{x}) + \text{KAM}_{\text{noise}}(\mathbf{x}). \end{aligned} \quad (19)$$

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. Note 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° in the regions of constant orientation, 0.5° in the right-hand region with linear orientation gradient and rises up to 1° at the small oscillating feature [cf. Fig. 11(a)]. Next, we applied De la Vallée Poussin-distributed noise with half-width $\sigma = 1^\circ$ and computed the KAM from the noisy data [cf. Fig. 11(b)]. The KAM of the noisy data is about 2° everywhere. In particular, no features of the ‘true’ KAM map are observable anymore.

In Fig. 12(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, orientation gradients are only observable if they are about four times as large as the half-width of the noise distribution.

Next, we applied the infimal convolution variation filter to the synthetic data corrupted by DIVP noise with half-width

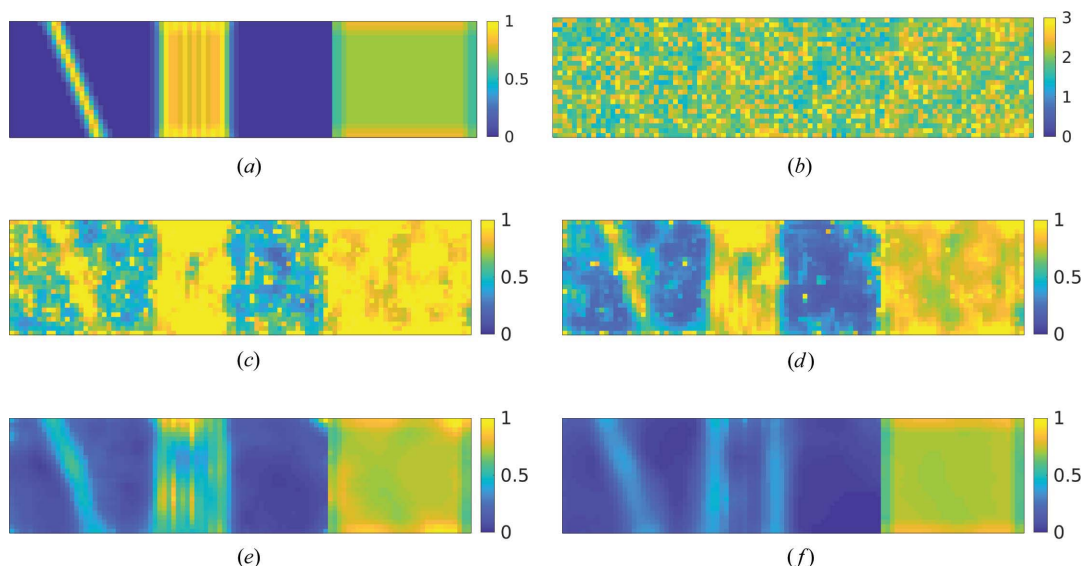


Figure 11

The kernel average misorientation (from third-order neighbors) in degrees (a) for the exact synthetic data as defined in Section 2 (no noise), (b) for the same data corrupted by DIVP noise with half-width $\sigma = 1^\circ$, and (c)–(f) after denoising with the infimal convolution filter with different regularization parameters α, β [(c) $\alpha = \beta = 0.0075$, (d) $\alpha = \beta = 0.01$, (e) $\alpha = \beta = 0.2$ and (f) $\alpha = \beta = 0.4$].

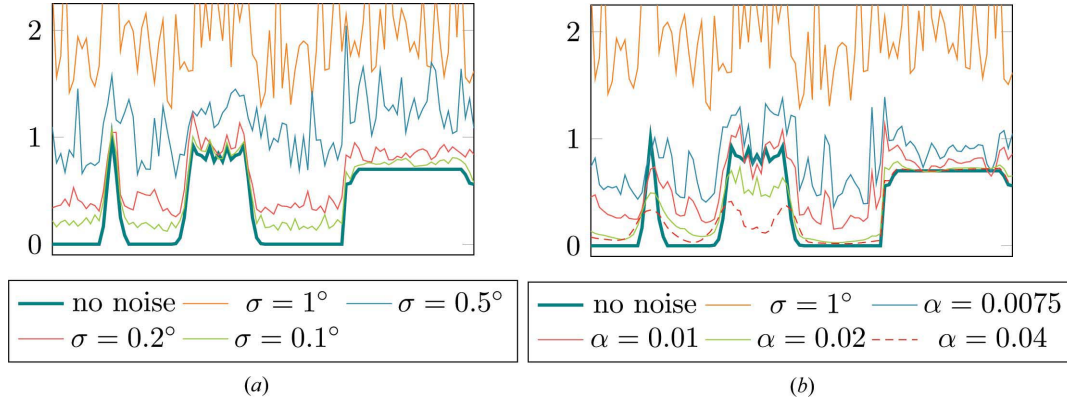


Figure 12

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) the noise level is $\sigma = 1^\circ$ and the KAM is computed from data denoised by total variation using different regularization parameters α .

$\sigma = 1^\circ$. Figs. 11(c)–11(f) and 12(b) illustrate the dependency of the KAM computed from the denoised orientation data with respect to different regularization parameters α , β . For $\alpha = 0.0075$ and $\alpha = 0.01$ the KAM remains noisy but makes the features of the ‘true’ KAM map visible. For $\alpha = 0.02$ and $\alpha = 0.04$ the KAM becomes more and more smooth. While in the large region with linear texture gradient the ‘true’ KAM is approximated better and better, we observe underestimation of the KAM in the small oscillating region. However, for all the choices of the regularization parameter the resulting KAM maps are more accurate than the KAM map computed from the noisy data directly.

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 \mathbf{e}_a is defined as

$$\frac{\partial}{\partial x_a} \mathbf{O}(\mathbf{x}) = \lim_{h \rightarrow 0} \frac{\mathbf{O}(\mathbf{x} + h\mathbf{e}_a) - \mathbf{O}(\mathbf{x})}{h}. \quad (20)$$

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 frames. Representing $\frac{\partial}{\partial x_a} \mathbf{O}(\mathbf{x})$ entirely with respect to the specimen reference frame leads to the skew-symmetric matrix

$$\mathbf{W}^a(\mathbf{x}) = \left[\frac{\partial}{\partial x_a} \mathbf{O}(\mathbf{x}) \right] \mathbf{O}(\mathbf{x})^T. \quad (21)$$

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

$$\boldsymbol{\theta}_i^a(\mathbf{x}) = \frac{1}{2} \epsilon_{ijk} W_{jk}^a(\mathbf{x}), \quad (22)$$

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 disorientation axis between the orientations $\mathbf{O}(\mathbf{x} + h\mathbf{e}_a)$ and $\mathbf{O}(\mathbf{x})$ in specimen coordinates scaled by the disorientation angle per distance h for the limiting process $h \rightarrow 0$ (cf. Morawiec, 2004).

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 $\mathbf{O}(\mathbf{x}^{m,n})$ on some two-dimensional regular grid $\mathbf{x}^{m,n}$ ($m = 1, 2, \dots; n = 1, 2, \dots$) by the products

$$\begin{aligned} \boldsymbol{\theta}^{x_1}(\mathbf{x}^{m,n}) &\simeq \frac{\delta[\mathbf{O}(\mathbf{x}^{m,n}), \mathbf{O}(\mathbf{x}^{m+1,n})]}{\Delta x_1} \boldsymbol{\eta}[\mathbf{O}(\mathbf{x}^{m,n}), \mathbf{O}(\mathbf{x}^{m+1,n})], \\ \boldsymbol{\theta}^{x_2}(\mathbf{x}^{m,n}) &\simeq \frac{\delta[\mathbf{O}(\mathbf{x}^{m,n}), \mathbf{O}(\mathbf{x}^{m,n+1})]}{\Delta x_2} \boldsymbol{\eta}[\mathbf{O}(\mathbf{x}^{m,n}), \mathbf{O}(\mathbf{x}^{m,n+1})], \end{aligned} \quad (23)$$

where Δx_1 , Δx_2 are the step sizes of the measurement grid, $\delta(\mathbf{O}_1, \mathbf{O}_2)$ is the disorientation angle and $\boldsymbol{\eta}(\mathbf{O}_1, \mathbf{O}_2)$ is the disorientation axis between two orientations in specimen coordinates (cf. Section 1.1).¹ Note that other approximations of $\boldsymbol{\theta}^{x_1}(\mathbf{x}^{m,n})$ are possible, e.g. symmetric disorientations between $\mathbf{O}(\mathbf{x}^{m+1,n})$ and $\mathbf{O}(\mathbf{x}^{m-1,n})$.

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

$$\boldsymbol{\kappa}(\mathbf{x}) = \begin{bmatrix} \boldsymbol{\theta}_1^{x_1}(\mathbf{x}) & \boldsymbol{\theta}_1^{x_2}(\mathbf{x}) & \boldsymbol{\theta}_1^{x_3}(\mathbf{x}) \\ \boldsymbol{\theta}_2^{x_1}(\mathbf{x}) & \boldsymbol{\theta}_2^{x_2}(\mathbf{x}) & \boldsymbol{\theta}_2^{x_3}(\mathbf{x}) \\ \boldsymbol{\theta}_3^{x_1}(\mathbf{x}) & \boldsymbol{\theta}_3^{x_2}(\mathbf{x}) & \boldsymbol{\theta}_3^{x_3}(\mathbf{x}) \end{bmatrix}. \quad (24)$$

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 steel as presented by Hickey *et al.* (2018). The orientations colorized by an inverse pole figure (IPF) map are displayed in Fig. 13(a). The white pixels correspond to non-indexed orientations. Fig. 13(b) displays the same data but colorized according to the disorientation axis and the disorientation angle with respect to the grain mean orientation as explained by Thomsen *et al.* (2017). Fig. 13(b) clearly shows that the orientation measurements are affected by some noise. Applying a total-variation-based filter for noise removal, we obtain the orientation maps depicted in Figs. 13(c) and 13(d).

¹ We prefer this definition of a discrete gradient as equation (21) together with a finite difference approximation of equation (20) does not lead to a skew-symmetric matrix.

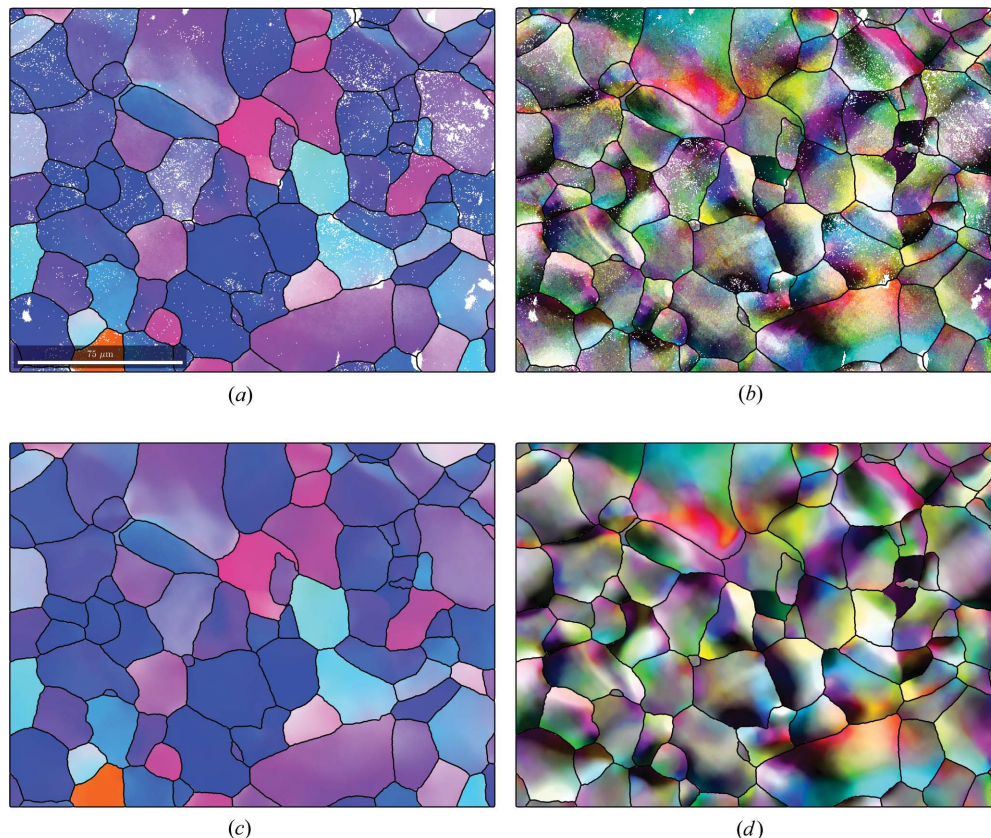


Figure 13

Orientation maps of interstitial-free steel (Hickey *et al.*, 2018) before and after denoising by the total variation method. (a) IPF-colored raw data, (b) axis-angle-colored raw data, (c) IPF-colored denoised data and (d) axis-angle-colored denoised data.

Although the raw data appear to be acceptable [*cf.* Fig. 13(a)], they are still too noisy to compute meaningful values for the curvature, as shown in Fig. 14(a). The reason is that computing the orientation gradient significantly amplifies the noise in the data. Applying a denoising filter before computing the curvature improves the accuracy considerably [*cf.* Fig. 14(b)].

3.3. Dislocation density estimation

The dislocation density tensor characterizes the dislocation state of the crystal lattice. To be precise, \mathbf{a} quantifies the local closure failure of the crystal and describes to what extent dislocations are geometrically necessary in order to preserve the compatibility. Under the assumptions of small lattice rotations (*i.e.* the rotation gradients are additive) and negligible elastic lattice strain gradients, the lattice curvature tensor $\boldsymbol{\kappa}(\mathbf{x})$ at a position \mathbf{x} can be related to the dislocation density tensor $\mathbf{a}(\mathbf{x})$ by

$$\mathbf{a} = \boldsymbol{\kappa}^T - \text{tr}(\boldsymbol{\kappa})\mathbf{I} \Leftrightarrow \boldsymbol{\kappa} = \mathbf{a}^T - \frac{1}{2}\text{tr}(\mathbf{a})\mathbf{I} \quad (25)$$

(*cf.* Pantleon, 2008). 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: α_{12} , α_{13} , α_{21} , α_{23} , α_{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

$$\mathbf{a}(\mathbf{x}) = \sum_{k=1}^K \rho_k(\mathbf{x}) \mathbf{b}_k \otimes \mathbf{l}_k \quad (26)$$

of screw and edge dislocations at preferred slip systems described by its Burgers vectors \mathbf{b}_k and line tangent vectors \mathbf{l}_k , where the corresponding dislocation densities ρ_k are strictly positive.² Using the relationship (25) we obtain for the lattice curvature tensor the representation

$$\boldsymbol{\kappa}(\mathbf{x}) = \sum_{k=1}^K \rho_k(\mathbf{x}) [\mathbf{l}_k \otimes \mathbf{b}_k - \frac{1}{2}\text{tr}(\mathbf{b}_k \otimes \mathbf{l}_k)\mathbf{I}]. \quad (27)$$

However, even in the case of complete knowledge of the lattice curvature tensor $\boldsymbol{\kappa}$, the decomposition (27) cannot be uniquely recovered, provided 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.*

$$U(\mathbf{x}) = \sum_{k=1}^K \rho_k(\mathbf{x}) u_k \quad \text{with} \quad \rho_k(\mathbf{x}) \geq 0, \quad (28)$$

should be minimized at each position \mathbf{x} (*cf.* Pantleon, 2008; Wilkinson & Randman, 2010), where u_k denotes the energy of the k th dislocation type (measured per unit dislocation density and per unit volume).

² From the physical point of view, we seek an equivalent composition of elementary dislocation segments such that the same dislocation state is obtained.

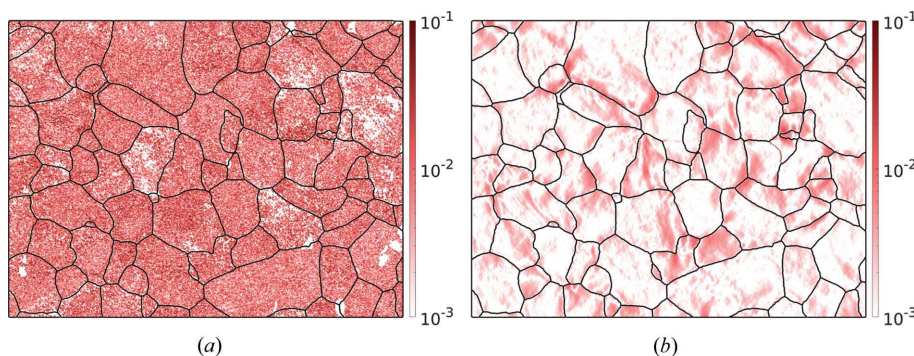


Figure 14

The coefficient κ_{11} (unit μm^{-1}) of the lattice curvature tensor of the EBSD data depicted in Fig. 13(a) (before denoising) (a) and in Fig. 13(b) (after denoising) (b).

Applying this approach to the curvatures computed in the previous section we obtain the total GND density as depicted in Fig. 15. The GND densities computed before denoising are completely corrupted, while the GND densities (and thus the total density/energy) computed after denoising seem physically reasonable and, in the first place, interpretable.

Hickey *et al.* (2018) also analyzed the Kikuchi pattern of the sample discussed above by a cross-correlation method. The resulting high-angular-resolution EBSD data were used to compute the total GND density depicted in Fig. 15(c). Although Fig. 15(c) offers much more detail, these GND densities still compare well with the total GND density computed from the Hough-transform-based EBSD data after denoising. We also observe that blank regions in Fig. 15(c), due to non-indexed orientations, correspond to regions of high GND density in Fig. 15(b). This is reasonable as those regions are hard to index.

In future research, the presented application to the estimation of the dislocation density (tensor) should be extended to the geometrically nonlinear case when lattice rotations may become large and relation (25) no longer holds.

3.4. Denoising high-resolution EBSD data

High-angular-resolution EBSD data originate from improved methods for interpreting Kikuchi patterns 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

$$C\{P(\mathbf{x}), M[\mathbf{O}(\mathbf{x})]\} \rightarrow \max, \quad (29)$$

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 and a master pattern M . Such methods have 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 often that the maximization problem leads to multiple peaks and the indexing algorithm has to choose one particular maximum. In those cases a combination with the denoising methods presented in this paper could serve as an additional criterion 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:

$$\sum_{\mathbf{x}} -C\{P(\mathbf{x}), M[\mathbf{O}(\mathbf{x})]\} + \lambda \|\nabla \mathbf{O}(\mathbf{x})\|_2 \rightarrow \min, \quad (30)$$

which is essentially the minimization problem (15) but with the correlation term replacing the disorientation to an orientation determined by another indexing method. According to this approach, orientations are chosen such that the simulated pattern has a high correlation with the measured pattern and, simultaneously, minimizes the total variation functional, *i.e.* varies smoothly in the spatial domain.

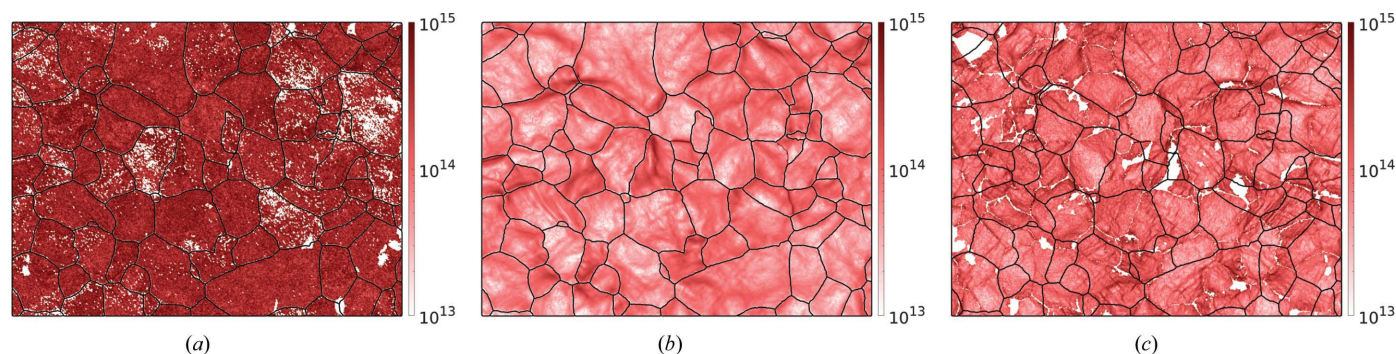


Figure 15

Total GND density $\sum \rho_k$ (unit $1/\text{m}^2$) computed (a) from the raw (Hough-based) EBSD map, (b) after denoising and (c) using high-angular-resolution EBSD data (Hickey *et al.*, 2018).

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

4. Conclusions and discussion

In the present paper we have shown that denoising filters provide a very powerful tool to increase the accuracy of noisy orientation maps. This is particularly true if the orientation data are used for subsequent calculations of kernel averaged misorientations, misorientation axes, orientation curvature tensors or geometrically necessary dislocation densities.

The main risk associated with denoising filters is to ‘over-smooth’ the data. In order to avoid this, the denoising method and the regularization parameters have to be selected carefully.

In future research, even more advanced techniques from mathematical image analysis, like wavelet decomposition, could be adapted to the setting of orientation maps to improve the distinction between small microstructural features and real noise.

Acknowledgements

The authors want to thank B. Britton, Imperial College London, for kindly proving the EBSD data set of interstitial-free steel discussed in Sections 3.2 and 3.3. His criticisms and suggestions significantly helped us to improve the quality of the paper. Furthermore, the authors gratefully acknowledge the help of Dr Dagmar Dietrich.

Funding information

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

References

- Bacak, M., Bergmann, R., Steidl, G. & Weinmann, A. (2015). Preprint. University of Kaiserslautern, Germany.
 Bachmann, F., Hielscher, R., Jupp, P. E., Pantleon, W., Schaeben, H. & Wegert, E. (2010). *J. Appl. Cryst.* **43**, 1338–1355.

- Bachmann, F., Hielscher, R. & Schaeben, H. (2011). *Ultramicroscopy*, **111**, 1720–1733.
 Bergmann, R., Chan, R. H., Hielscher, R., Persch, J. & Steidl, G. (2016). *Inverse Prob. Imaging*, **10**, 281–304.
 Bergmann, R., Fitschen, J. H., Persch, J. & Steidl, G. (2018). *J. Math. Imaging Vis.* **60**, 1459–1481.
 Bergmann, R., Laus, F., Steidl, G. & Weinmann, A. (2014). *SIAM J. Imaging Sci.* **7**, 2916–2953.
 Bergmann, R. & Weinmann, A. (2016). *J. Math. Imaging Vis.* **55**, 401–427.
 Chen, D. & Kuo, J.-C. (2010). *Ultramicroscopy*, **110**, 1297–1305.
 Garcia, D. (2010). *Comput. Stat. Data Anal.* **54**, 1167–1178.
 Godfrey, A. (2004). *Scr. Mater.* **50**, 1097–1101.
 Gupta, V. K. & Agnew, S. R. (2010). *Microsc. Microanal.* **16**, 831–841.
 Heilbronner, R. & Barrett, S. (2014). *Image Analysis in Earth Sciences: Microstructures and Textures of Earth Materials*. Berlin, Heidelberg: Springer.
 Hickey, J. L. R., Rouland, S. & Britton, T. B. (2018). *arXiv*: 1807.02017.
 Hielscher, R. (2013). *J. Multivariate Anal.* **119**, 119–143.
 Hielscher, R., Bachmann, F., Mainprice, D. & Kilian, R. (2018). *MTEX 5.1*, <http://mtex-toolbox.github.io>.
 Hielscher, R., Britton, B. & Bartel, F. (2018). *arXiv*: 1810.03211.
 Humphreys, F. J., Bate, P. S. & Hurley, P. J. (2001). *J. Microsc.* **201**, 50–58.
 Kamaya, M. (2010). *Mater. Trans.* **51**, 1516–1520.
 Kamaya, M. (2011). *Ultramicroscopy*, **111**, 1189–1199.
 Kobler, A., Kashiwar, A., Hahn, H. & Kübel, C. (2013). *Ultramicroscopy*, **128**, 68–81.
 Konijnenberg, P., Zaefferer, S. & Raabe, D. (2015). *Acta Mater.* **99**, 402–414.
 Morawiec, A. (2004). *Orientations Rotations*. Berlin, Heidelberg: Springer.
 Pantleon, W. (2008). *Scr. Mater.* **58**, 994–997.
 Passchier, C. W. & Trouw, R. A. (2005). *Microtectonics*. Berlin, Heidelberg: Springer.
 Poulsen, H. F. & Fu, X. (2003). *J. Appl. Cryst.* **36**, 1062–1068.
 Prior, D. J. (1999). *J. Microsc.* **195**, 217–225.
 Rauch, E., Portillo, J., Nicolopoulos, S., Bultreys, D., Rouvimov, S. & Moeck, P. (2010). *Z. Kristallogr.* **225**, 103–109.
 Schaeben, H. (1999). *Textures Microstruct.* **33**, 365–373.
 Thomsen, K., Mehnert, K., Trimby, P. W. & Gholinia, A. (2017). *Ultramicroscopy*, **182**, 62–67.
 Weinmann, A., Demaret, L. & Storath, M. (2014). *SIAM J. Imaging Sci.* **7**, 2226–2257.
 Wilkinson, A. J. (2001). *Scr. Mater.* **44**, 2379–2385.
 Wilkinson, A. J. & Randman, D. (2010). *Philos. Mag.* **90**, 1159–1177.
 Wright, S. I., Nowell, M. M., Lindeman, S. P., Camus, P. P., De Graef, M. & Jackson, M. A. (2015). *Ultramicroscopy*, **159**, 81–94.
 Zaefferer, S. (2011). *Cryst. Res. Technol.* **46**, 607–628.