Fast automated detection of crystal distortion and crystal defects in polycrystal images

Matt Elsey Benedikt Wirth

October 4, 2013

Abstract

Given an image of an atomic crystal, we propose a variational method which at each image location determines the local crystal state and which localizes and characterizes crystal defects. In particular, the local crystal orientation and elastic distortion are detected, as well as dislocations, and grain and twin boundaries. To this end an energy functional is devised whose minimization yields a tensor field G describing the local crystal strain at each point. The desired information about the local crystal state can then be read off from this tensor field; in particular, its curl provides information about grain boundaries and dislocations. As is typical for variational image processing, the energy functional is composed of a fidelity and a regularization term. It has a simple L^2-L^1 type structure so that its minimization can be performed via a split Bregman iteration. GPU parallelization results in short computing times.

1 Introduction

Due to continuing improvements in imaging technology it is nowadays possible to obtain images of materials with atomic resolution, for example, photographs of metallic polycrystals from transmission electron microscopy (TEM). This provides a visualization of the lattice defects in metallic (or other) crystalline structures, which are known to influence or even determine the macroscopic material properties, such as elastic and plastic behavior or yield stresses. It is to be expected that the analysis of such experimental images soon becomes a standard tool to understand and validate models of material properties or of material evolution (such as crystallization, Ostwald ripening, or processes of elastic and plastic deformation).

Images of crystals are often very large, and the local crystal structure and defects are typically tedious and even difficult to spot and characterize with the naked eye. Also, one might have a large number of such images, for example a time series, so that there is a need for a fast automated analysis of crystal images. This article proposes a variational method to automatically localize and characterize crystal defects (grain boundaries, twin boundaries, and dislocations with their Burgers vector) in large crystal images and in reasonable computing time (cf. Fig. 1). The method also provides the local crystal orientation and elastic distortion. It is not only of benefit for the analysis of images from physical experiments, but also for the automatic evaluation of simulated data: Modern physical crystal models such as the phase field crystal [6] combined with modern numerical schemes such as GPU-parallelizable PDE solvers [7] quickly lead to large simulation results.

Our method is based on finding a tensor map $G : \Omega \to \mathbb{R}^{2\times 2}$ which at each point in the given crystal image $u : \Omega \to \mathbb{R}$ encodes the local crystal distortion relative to a fixed reference crystal.

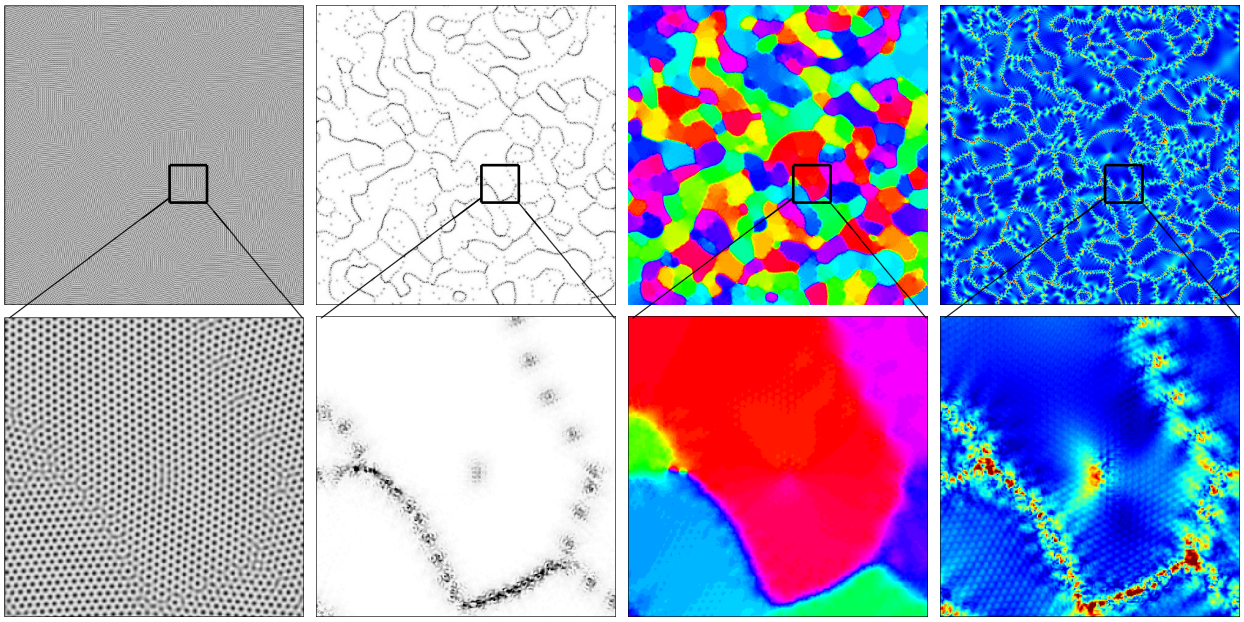

Figure 1: Given a crystal image (left), our method computes a distortion map $G : \Omega \to \mathbb{R}^{2\times 2}$ whose curl magnitude indicates crystal defects (second image; single spots are dislocations, dotted and continuous lines are low and high angle grain boundaries). Polar decomposition of G yields the local crystal orientation (third image, 0 $\pi/3$), and $\operatorname{dist}(G^{-1}, \mathrm{SO}(2))$ can be taken as a measure of crystal strain (right, 0 0.1). Image size is 2048×2048, computing time approximately 10 min.

The curl of this tensor field will serve to identify and characterize crystal defects, while other types of information such as the crystal orientation or deformation can directly be read off from G. The field G will be found by minimizing an L^2-L^1 type energy, for which a GPU-parallelized split Bregman iteration is employed.

In the following we will briefly recall the crystal defects of interest and describe related methods before we introduce and discuss the different terms of the proposed energy functional in Sec. 2. The minimization algorithm is described in Sec. 3, and numerical experiments are presented in Sec. 4.

1.1 Grain boundaries, dislocations, and twin boundaries

In a perfect atomic crystal, atoms are arranged in a Bravais lattice: the atom positions are (in 2D) given by $n_1 a_1 + n_2 a_2$ for two fixed vectors $a_1, a_2 \in \mathbb{R}^2$ and all integers n_1, n_2. A real specimen of a polycrystalline material is typically not homogeneous, but is composed of a number of perfectly crystalline regions (so-called grains) at different crystal orientations, which meet at so-called grain boundaries (Fig. 2 left). Grain boundaries can be characterized by the mismatch angle between the two adjacent grain orientations. Low angle grain boundaries usually consist of a string of point defects in the crystal (cf. Fig. 1).

Another type of crystal defect to be detected by our method are dislocations. In 2D, these are points where a line of atoms terminates in an otherwise perfect crystal (edge dislocation, Fig. 2 middle). If a closed, lattice-aligned curve is drawn around the dislocation, it cannot be mapped onto an undisturbed reference lattice without breaking it up at one point (Fig. 2 middle right). The vector spanning the resulting gap is the Burgers vector. It characterizes the dislocation; for

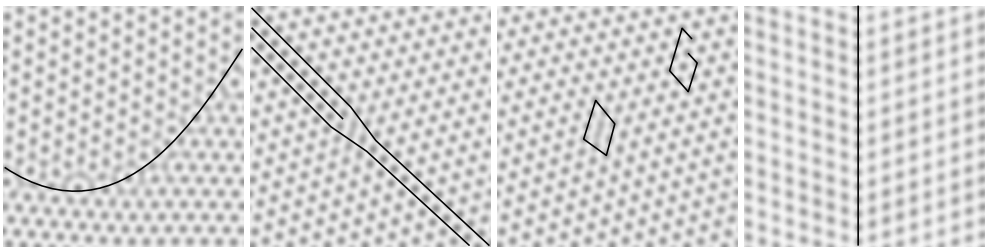

Figure 2: Grain boundary between two grains (left), dislocation with a terminating line of atoms (middle left), same dislocation enclosed by a curve on the lattice and same curve mapped into an undisturbed lattice region (middle right), and a twin boundary between two mirror-symmetric lattice variants.

instance, only dislocations with opposite Burgers vectors can annihilate.

A twin boundary is the final crystal defect of interest (Fig. 2 right). Here, the crystal lattice can occur in two symmetric variants. At a twin boundary, two crystal regions with different crystal variants share a line of atoms along which they meet.

Physicists are interested, for instance, in how defects move and grain boundaries behave or in how the grain size distribution evolves over time. To this end crystal images are often analyzed manually (e. g. [18]). The proposed algorithm aims to alleviate the need for manual analysis of such images.

1.2 Related methods

Several methods for the automatic analysis of crystals do already exist, however, none combines the universality of an image as input (be it a photograph of a pattern, a phase field crystal simulation result, or a TEM image) with the versatility of the crystal distortion as output (from which crystal orientation, defects, and other information can be directly inferred). Indeed, the methods either require the set of all atom positions as input or can only detect grains (and potentially their strain), but ignore defects such as dislocations.

The method in [2] segments a crystal image $u : \Omega \to \mathbb{R}$ into disjoint regions with different crystal orientations, using the Chan–Vese level-set approximation of the piecewise constant Mumford–Shah segmentation. We will later borrow the fitting term from their segmentation functional. The method can be extended to yield also a global deformation $\phi : \Omega \to \mathbb{R}^2$ to account for strained crystals. Due to the many high-dimensional, nonlinearly coupled degrees of freedom (several level set functions, the deformation, the crystal orientation in each region) the functional is difficult and slow to minimize. However, restricting themselves to the mere Mumford–Shah crystal segmentation, the authors in [3] find an efficient, convex reformulation via functional lifting (instead of using the piecewise constant crystal orientation map $\alpha : \Omega \to [0, 2\pi)$ one expresses the functional in terms of the binary function $1_\alpha : \Omega \times [0, 2\pi) \to \{0, 1\}$, $1_\alpha(x, \beta) = 1$ if $\alpha(x) \leq \beta$ and 0 else). An even more efficient version penalizes the segmentation interfaces according to the jump in α, which can be further improved by taking the periodicity of the orientation into account [13] (i. e. measuring the jump in α on S^1 instead of on $[0, 2\pi)$). A corresponding GPU implementation is very fast, but needs lots of memory due to the additional dimension from the functional lifting.

A very fast, but quite ad hoc and thus less flexible approach is to convolve the crystal image with a small wavelet-like patch resembling a unit cell at some chosen orientation [12]. The resulting oscillations in the convolution are smoothed out, yielding a gray value which depends on the local

crystal orientation in the image. The map from gray values to orientations is given by a lookup table that is computed beforehand. Incidentally, dislocations also show up in the convolution via a different color. The method's disadvantage is that it breaks down for strained crystals, it cannot distinguish between orientations which produce the same gray value after convolution, the lookup table has to be recomputed for different images, and the shading may not vary across the image.

A different type of methods uses snapshots from 3D molecular dynamics simulations as input. Here the local crystal state can easily be characterized by identifying the nearest neighbors of an atom. Such methods are not variational and thus relatively fast so that they can sometimes even be used on the fly during molecular dynamics simulations. The algorithm in [14] first detects all defect-free regions and then finds the shortest Burgers circuits around regions with defects, which precisely localizes the dislocation lines and determines their Burgers vectors. A technique of sweeping out tubes around the dislocation lines and triangulating grain boundaries increases the method's robustness [15] (but slows down computations). An approach conceptually similar to ours determines the local crystal strain from the nearest neighbor positions around each atom and then identifies dislocations from the strain curl [1].

In contrast to the above approaches, our method simultaneously finds dislocations, grain boundaries, and crystal distortion, based solely on a crystal image. A previous version of our method was published in [8]. That work is also based on identifying a crystal strain field from which all other information is derived. In the current article, however, we instead consider the inverse of that strain field which turns out to be more natural. Also, as opposed to [8], we will here discuss the mathematical well-posedness of the model and propose a model extension as a rigorous remedy for problems arising from an inherent redundancy in the representation of crystal orientations.

2 Model

This section introduces the variational energy E which we will minimize to obtain the crystal rotation, strain, and defects from a given crystal image $u : \Omega \to \mathbb{R}$. We will first explain how crystal defects can be represented by curl concentrations of a distortion field before suggesting a functional form of E in several steps.

2.1 Dislocations and grain boundaries as curl concentrations of distortion fields

Initially, consider a perfect, undisturbed and unstrained reference crystal occupying $\mathbb{R}^2$ at a fixed (arbitrarily chosen) orientation. In comparison, a polycrystal is composed of several perfectly crystalline (though potentially strained) regions which are connected by grain boundaries and which may also contain isolated dislocations. Away from these crystal defects, the polycrystal can be interpreted locally as a deformed version of a piece $S \subset \mathbb{R}^2$ of the reference crystal, where the deformation $\phi = (\phi_1, \phi_2)^T : S \to \mathbb{R}^2$ maps each point $\bar{x} \in S$ of the reference crystal onto its deformed position $x = \phi(\bar{x})$ in the polycrystal. The local state of the polycrystal is fully described by the deformation gradient

$$F(x) = D\phi \circ \phi^{-1}(x) = \begin{pmatrix} \partial_{x_1}\phi_1 & \partial_{x_2}\phi_1 \\ \partial_{x_1}\phi_2 & \partial_{x_2}\phi_2 \end{pmatrix} \circ \phi^{-1}(x) \in \mathbb{R}^{2\times 2}\,.$$

At each crystal position $x \in \mathbb{R}^2$ it encodes how the crystal locally is rotated and distorted. Equivalently, we can consider the inverse deformation $\psi = \phi^{-1}$ which transforms the region $\phi(S)$ of the

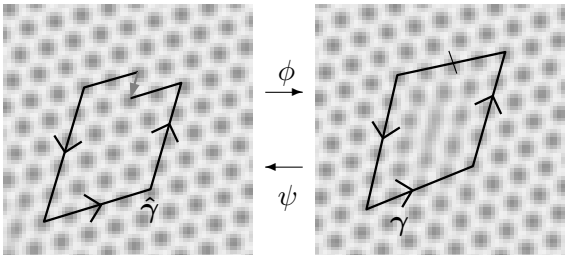

Figure 3: The curve γ around the dislocation (right) can be mapped back onto a curve $\hat{\gamma}$ in the reference crystal via ψ (left). $\hat{\gamma}$ is no longer closed, the gap being the Burgers vector (gray arrow).

polycrystal into the unstrained reference state. Its deformation gradient is given by

$$G(x) = D\psi(x) = D(\phi^{-1})(x) = F^{-1}(x)\,.$$

Not only do F and G encode the local crystal rotation and distortion, but as we will see, they also contain information about crystal defects so that the distortion field G provides a detailed characterization of the polycrystal state. Our analysis of crystal images will be based on extracting such a field G from a given image.

The identification of crystal defects from the field G relies on its interpretation as the gradient of a deformation: Away from defects we can deform any patch of the given crystal back into the chosen fixed reference state. The gradient of the associated deformation ψ is given by G. Consequently, the field G is conservative in defect-free regions, which is equivalent to it having zero curl,

$$0 = \mathrm{curl} G = \begin{pmatrix} \partial_{x_1} G_{12} - \partial_{x_2} G_{11} \\ \partial_{x_1} G_{22} - \partial_{x_2} G_{21} \end{pmatrix}\,.$$

At crystal defects, the interpretation of G as a gradient field breaks down so that grain boundaries and dislocations will be identified with line- and point accumulations of the curl of the distortion field as detailed below.

Within a grain, the deformations ϕ or ψ and the distortion fields F or G typically vary smoothly, while at grain boundaries, where two differently deformed reference crystal patches meet, the deformation ψ obviously is discontinuous. The distortion field G is also discontinuous across the grain boundary, since otherwise the crystal would be in the same state on either side so that it could be interpreted as a coherent grain. Thus grain boundaries are represented by the jump set of the field G, that is, as a network of one-dimensional discontinuity lines (in a real crystal, there is typically a thin transition region instead of a sharp discontinuity).

As for dislocations, let B be a simply connected neighborhood of a dislocation, containing no other dislocation. Let the boundary ∂B be parameterized counterclockwise by $\gamma : [0,1] \to \mathbb{R}^2$. Since ∂B lies in a defect-free region, around each point $x \in \partial B$ we can deform a small neighborhood of the crystal into the reference state. If we do this consistently between adjacent neighborhoods, we obtain an image $\hat{\gamma} : [0,1] \to \mathbb{R}^2$ of the curve γ in the reference crystal (Fig. 3). Note that the curve described by $\hat{\gamma}$ may contain self-intersections and typically is no longer closed. We have

$$\int_B \mathrm{curl} G \,\mathrm{d}x = \int_{\partial B} G n^\perp \,\mathrm{d}x = \int_0^1 G(\gamma(t))\dot{\gamma}(t)\,\mathrm{d}t = \int_0^1 \dot{\hat{\gamma}}(t)\,\mathrm{d}t = \hat{\gamma}(1) - \hat{\gamma}(0)\,,$$

where $n^\perp$ denotes the unit outward normal n to ∂B rotated counterclockwise by $\frac{\pi}{2}$. The vector on the right-hand side is the dislocation's Burgers vector, which is independent of the originally chosen

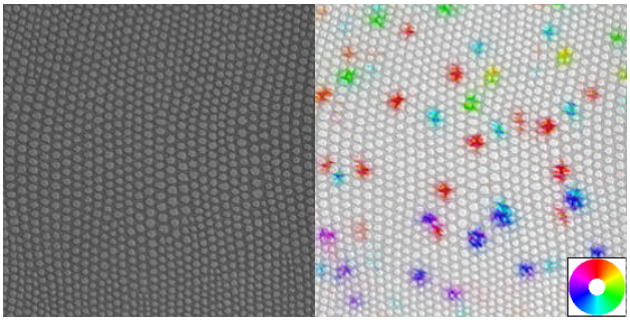

Figure 4: Photograph of a synthetic reptile skin (courtesy `http://www.stoff4you.de`) and regions of curl concentration detected by our method, showing various dislocations and the corresponding Burgers vectors (by color-coding).

parameterization γ. It is also independent of the chosen neighborhood B, as long as B contains only one dislocation. Indeed, had we chosen $\tilde{B}$ instead of B, we would have

$$\int_{\tilde{B}} \mathrm{curl} G \,\mathrm{d}x = \int_B \mathrm{curl} G \,\mathrm{d}x + \int_{\tilde{B}\setminus B} \mathrm{curl} G \,\mathrm{d}x - \int_{B\setminus\tilde{B}} \mathrm{curl} G \,\mathrm{d}x = \int_B \mathrm{curl} G \,\mathrm{d}x$$

due to $\mathrm{curl} G = 0$ on $\tilde{B} \setminus B$ and $B \setminus \tilde{B}$. Thus the field G has a concentration of non-zero curl at a dislocation. If the spatial extent of the dislocation is infinitely small, $\mathrm{curl} G$ will behave like a Dirac delta at this point, while for real crystals, $\mathrm{curl} G$ will be diffused a little over the width of one to a few unit cells. As an example, Figure 4 shows the curl concentrations in the scales pattern of a synthetic reptile skin. These are clearly located in places where an additional row of scales squeezes in between two rows. The direction of $\mathrm{curl} G$ or equivalently of the corresponding Burgers vector is encoded by color. It indicates the direction in which the additional row of scales is inserted. A red dot for instance implies that $\int_{\partial B} G n^\perp \,\mathrm{d}x$ points upwards, thus there is vertical dilation left and compression right of the dislocation so that additional scales are inserted left of it. Only six discrete directions can be identified since a 2D hexagonal lattice can only have simple dislocations with six different Burgers vectors.

Summarizing, the distortion field G describes the crystal rotation and strain at each point, its lines of discontinuity represent the grain boundaries, and point concentrations of $\mathrm{curl} G$ represent dislocations whose Burgers vector is given by the integral of $\mathrm{curl} G$ over a small neighborhood.

2.2 Variational energy

Given an image $u : \Omega \to \mathbb{R}$ of a crystal we aim to extract the crystal distortion field $G : \Omega \to \mathbb{R}^{2\times 2}$. Following the usual paradigm of variational image processing, G will be chosen as the minimizer of an energy that is composed of a fitting or fidelity term $E_{\mathrm{fit}}[G]$, which penalizes a mismatch between the given image and the extracted field G, and a regularization $E_{\mathrm{reg}}[G]$ of G, which incorporates a priori knowledge of the structure of G and correspondingly regularizes G in the case of noisy input data.

The fitting term is taken from [3] and exploits the periodicity of crystals: Consider the perfect unstrained reference crystal from the previous section and let $v_1, \ldots, v_K \in \mathbb{R}^2$ be vectors such that the reference crystal is invariant under translation by v_k, $k = 1, \ldots, K$. For instance, $v_1, \ldots, v_K$ can be chosen as the position vectors of all nearest neighbor atoms relative to the central atom. If

the lattice is locally deformed by an affine transformation $F = G^{-1}$, then the deformed crystal is invariant under translation by $Fv_1, \dots, Fv_K$. Hence, the image gray values $u(x)$ and $u(x+F(x)v_k)$, $k = 1, \dots, K$, should be identical, suggesting the fidelity term

$$E_{\text{fit}}[G] = \int_\Omega \sum_{k=1}^{K} (u(x + G^{-1}(x)v_k) - u(x))^2 \, \mathrm{d}x \,. \tag{1}$$

The fitting term needs to be complemented with a regularizing component in order to reduce the set of local energy minima and to make the extraction of G stable with respect to noise. In our case, we are particularly interested in the identification of crystal defects which are represented as curl concentrations of G. Put differently, curl$G$ is a quantity of interest, and we have the a priori knowledge that it is concentrated on a sparse set. Since L^1-type norms are well-known to promote sparseness we choose the regularization

$$E_{\text{reg}}[G] = \int_\Omega |\text{curl}G| \, \mathrm{d}x = \sup_{\varphi \in C_0^\infty(\Omega;\mathbb{R}^2), |\varphi(x)| \le 1} \int_\Omega \text{tr}\left[G \nabla^\perp (\varphi^T) \right] \, \mathrm{d}x \,,$$

where $\nabla^\perp = (-\partial_{x_2}, \partial_{x_1})^T$ and the integrand $|\text{curl}G|$ has to be interpreted in a similar sense as the integrand $|\nabla G|$ of the total variation seminorm.

Altogether we arrive at the energy

$$E[G] = E_{\text{fit}}[G] + \omega_1 E_{\text{reg}}[G] = \int_\Omega \sum_{k=1}^{K} (u(x + G^{-1}(x)v_k) - u(x))^2 \, \mathrm{d}x + \omega_1 \int_\Omega |\text{curl}G| \, \mathrm{d}x$$

for some weighting parameter $\omega_1 > 0$. Note however, that $G^{-1} \equiv 0$ would globally minimize the fitting term. Thus we have to prevent G from taking infinitely large values. Also, since G is a deformation gradient and hence orientation-preserving, we have to prevent $\det G \le 0$. One can achieve this by adding a coercive energy term

$$E_{\text{coer}}[G] = \int_\Omega f_{\text{coer}}(G) \, \mathrm{d}x \,,$$

where $f_{\text{coer}}(G)$ is zero for moderate values of G and behaves like

$$f_{\text{coer}}(G) \sim |G|^2 \text{ as } |G| \to \infty \qquad \text{and} \qquad f_{\text{coer}}(G) \sim -\log \det G \text{ as } \det G \to 0.$$

Remark 1. In [8] we considered a variant of the above energy which extracts $F = G^{-1}$ instead of G and which is a valid approximation for small crystal strains, that is $G \approx I$. Using $\text{curl}G = \text{curl}F + O((G-I)^2)$, we replaced the above energy by

$$E[F] = \int_\Omega \sum_{k=1}^{K} (u(x + F(x)v_k) - u(x))^2 \, \mathrm{d}x + \omega_1 \int_\Omega |\text{curl}F| \, \mathrm{d}x \,.$$

Here, crystal defects show up as curl concentrations in the field F, however, the slightly less nonlinear form of the energy comes at the expense of misleading results if G or F deviate strongly from the identity: A simple example is given by Figure 5 where $\int_B |\text{curl}F| \, \mathrm{d}x = 0$ even though there is a dislocation at the center of B.

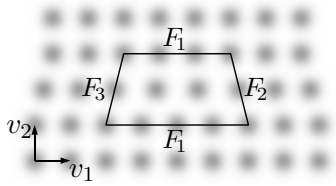

$$F_1 = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$$
$$F_2 = \begin{pmatrix} -1/4 & -1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} 1 & -1/4 \\ 0 & 1 \end{pmatrix}\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$$
$$F_3 = \begin{pmatrix} 1/4 & -1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} 1 & 1/4 \\ 0 & 1 \end{pmatrix}\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$$

Figure 5: Even though the lattice contains a dislocation, the integral of curl$F$ over B (the interior of the trapezoid) gives zero, $\int_B |\mathrm{curl} F|\,\mathrm{d}x = \int_B \mathrm{curl} F\,\mathrm{d}x = \int_{\partial B} F n^\perp\,\mathrm{d}x = 0$, hence the dislocation is ignored. The reason lies in the fact that the approximation $F \approx I$ is not valid since F interprets the crystal to be roughly rotated by $\frac{\pi}{2}$ relative to the reference stencil v_1, v_2.

2.3 Why the curl regularization is insufficient for well-posedness

Unfortunately, the penalization of curl$G$ is not sufficient to guarantee the well-posedness of the lattice segmentation problem. As we will discuss below, this is associated with the high nonlinearity of the fitting term $E_{\mathrm{fit}}[G]$ and in particular with the lack of quasiconvexity of its integrand $f_{\mathrm{fit}}(G,x) = \sum_k (u(x + G^{-1}(x)v_k) - u(x))^2$.

Recall that a function $f : \mathbb{R}^{2\times 2} \to \mathbb{R}$ is called quasiconvex if $f(G) \le \int_{(0,1)^2} f(G + D\varphi)\,\mathrm{d}x$ for all $G \in \mathbb{R}^{2\times 2}, \varphi \in W_0^{1,\infty}((0,1)^2;\mathbb{R}^2)$. A sufficient condition for quasiconvexity is given by rank-1-convexity, meaning that $t \mapsto f(G + tA)$ is convex for all $G \in \mathbb{R}^{2\times 2}$ and for all $A \in \mathbb{R}^{2\times 2}$ with rank 1. The lack of quasiconvexity of f implies that the functional $E^f[\phi] = \int_\Omega f(D\phi)\,\mathrm{d}x$ is not lower semi-continuous in $\phi \in W^{1,p}(\Omega;\mathbb{R}^2)$ under weak convergence. As a consequence, energies involving E^f typically have no minimizers, and almost-minimizers ϕ (for which $E^f[\phi]$ is small) exhibit microstructure. Below, we will demonstrate that the same issues occur in the crystal segmentation problem, since the tensor G can be interpreted as the derivative of a deformation ψ.

To show that f_{fit} is usually non-quasiconvex, assume a noise-free image u to represent a perfect Bravais lattice with lattice vectors $a_1, a_2 \in \mathbb{R}^2$,

$$B = \{n_1 a_1 + n_2 a_2 \ : \ n_1, n_2 \in \mathbb{Z}\}\,,$$

and assume (without loss of generality) compatible stencil vectors $v_1 = a_1, v_2 = a_2, v_3, \dots, v_K \in B$. Let $(a_1|a_2) \in \mathbb{R}^{2\times 2}$ be the matrix with columns a_1, a_2, then

$$f_{\mathrm{fit}}(G,x) = 0 \quad \Leftrightarrow \quad \forall k : G^{-1}v_k \in B \quad \Leftrightarrow \quad G = (a_1|a_2)M^{-1}(a_1|a_2)^{-1} \text{ for some } M \in \mathrm{GL}_2(\mathbb{Z})\,.$$

Hence $t \mapsto f_{\mathrm{fit}}(I + tA, x)$ is zero at $t = 0$ and $t = 1$ for instance for $A = \frac{1}{b}(a_1 \otimes (1-b,c)^T)(a_1|a_2)^{-1}$ with any $b, c \in \mathbb{Z}$. In other words, there is a rank-1 connection between the two minima $G = I$ and $G = I + A$, and thus f_{fit} is not rank-1-convex or quasiconvex.

The implications shall be illustrated in the following concrete example which shows how the formation of microstructure can interfere with the desired segmentation result. Let us assume, the image u exhibits an ℓ-spaced hexagonal lattice (Fig. 6 left) of a particular orientation for $x_1 < 0$ and rotated by $\frac{\pi}{6}$ for $x_1 > 0$ so that the line $x_1 = 0$ represents a grain boundary. As stencil vectors we shall consider the standard set, $v_k = \ell(\cos(\frac{2\pi k}{6}), \sin(\frac{2\pi k}{6}))^T$, $k = 1, \dots, 6$. While $G = I$ minimizes the fitting term for $x_1 < 0$, a rotation by $\frac{\pi}{6}$ yields zero fitting energy for $x_1 > 0$. Note however, that $G = G_1 := \frac{1}{\sqrt{3}} I$ and $G = G_2 := \frac{1}{\sqrt{3}}(e_1|3e_2)$ represent alternative minimizers for $x_1 > 0$, where e_1, e_2 denotes the standard Euclidean basis. Now G_1 and G_2 can be laminated as in Figure 6 right

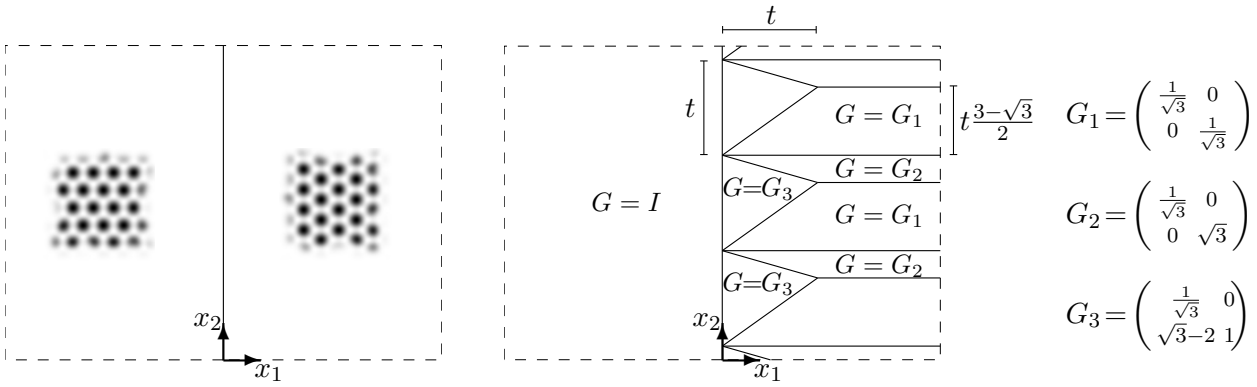

Figure 6: For the two hexagonal crystal grains with grain boundary at $x_1 = 0$ (left) and for the stencil $v_k = \ell(\cos(\frac{2\pi k}{6}), \sin(\frac{2\pi k}{6}))^T$, $k = 1, \dots, 6$, the fitting energy can be made arbitrarily small via lamination (right) with finer and finer lamination width t.

with a period t and a boundary layer at $x_1 = 0$ of width t. It is easily shown that $\mathrm{curl} G = 0$ everywhere, and the fitting term is zero outside the boundary layer. By letting $t \to 0$, we see

$$\inf_G E_{\mathrm{fit}}[G] + \omega_1 E_{\mathrm{reg}}[G] + E_{\mathrm{coer}}[G] = 0\,.$$

Obviously, there is no optimal tensor field G. Furthermore, the near optimal tensor fields G exhibit undesired microstructure for $x_1 > 0$, and the grain boundary does not show up in $\mathrm{curl} G$.

The above lamination construction is only possible due to the non-quasiconvexity of f_{fit}, which enables the tensor field G to maintain a low fitting energy while at the same time G can still be interpreted as the gradient of a continuous deformation, $G = D\psi$ (which is equivalent to $\mathrm{curl} G = 0$). Let us note that even despite a proper initialization near the desired local minimizer G one has to expect minimization problems due to the above issues, especially in the presence of noise. Also, it is not the particular form of our fitting term which is responsible for this behavior, but rather the nonlinearity inherent in the segmentation problem. Indeed, even if the fitting integrand were as simple as $f_{\mathrm{fit}}(G) = \mathrm{dist}^2(G, \mathrm{P})$ (where the point group P denotes the set of rotations and reflections which leave the lattice invariant) it would still be non-quasiconvex. This can be readily seen since $t \mapsto f_{\mathrm{fit}}(G + tA) = \min_{P \in \mathrm{P}} |G + tA - P|^2$ for a rank 1 matrix A is non-convex as the pointwise minimum of quadratic functions.

2.4 Additional H^1-type gradient regularization

The previous discussion implies that the L^1-type curl regularization is not sufficient to prevent the formation of undesired microstructure and to guarantee existence of minimizers. Consequently, we need additional regularization of the full gradient ∇G or parts thereof. There are different alternatives, the simplest of which is $\|\nabla G\|_{L^2}^2$. It is preferable for its strict quadratic convexity, and its smoothing property typically stabilizes the numerical optimization and helps to avoid local minima. Nevertheless, $\|\nabla G\|_{L^2}^2$ seems an inappropriate choice at first glance, since its tendency to smooth out the tensor field G counteracts the concentration of $\mathrm{curl} G$ at grain boundaries and dislocations. Indeed, the TV-seminorm $|G|_{\mathrm{TV}} = \int_\Omega |\nabla G| \, \mathrm{d}x$ instead would still allow sharp grain boundaries across which G jumps, and terms such as $\int_\Omega |\mathrm{div} G| \, \mathrm{d}x$ (to be interpreted in the same weak sense) would even allow point concentrations of $\mathrm{curl} G$ at dislocations. However, it turns out that a regularization with $\|\nabla G\|_{L^2}^2$ experimentally yields the visually most attractive results.

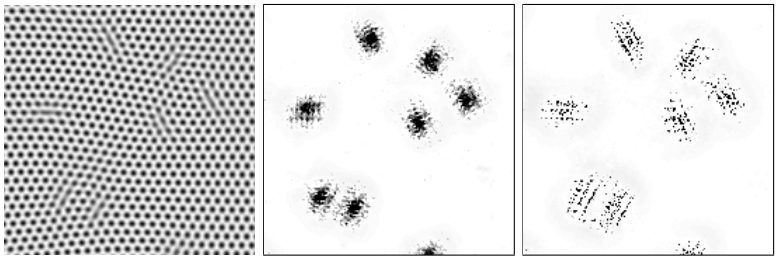

Figure 7: The original crystal image (left) and $|\mathrm{curl} G|$ for the optimal tensor field G, once obtained with additional regularization $\|\nabla G\|_{L^2}^2$ (middle) and once with $\int_\Omega |\mathrm{div} G| \, \mathrm{d}x$ instead (right).

The reason partly lies in the fact that images of crystals have a smallest scale, the interatomic distance. Thus, the grain boundaries and dislocation cores always have a finite, non-zero width. L^1-type regularizations such as $\int_\Omega |\mathrm{div} G| \, \mathrm{d}x$ now indeed lead to very localized concentrations of $\mathrm{curl} G$ in single pixels, however, such pixels are scattered across the whole core region of the crystal defects. Such point clouds of pixels with high $\mathrm{curl} G$ make it difficult to visually separate nearby crystal defects, while an L^2-type regularization smooths out $\mathrm{curl} G$ over the defect cores, which can then easily be distinguished (Fig. 7). Furthermore, inside the grains we expect G to vary gradually, since large grains might be strained in some regions but not in others. $\|\nabla G\|_{L^2}^2$ is the most appropriate regularization for such situations, while L^1-type regularizations would lead to the notorious staircasing effect with piecewise constant G.

The existence of minimizers with this additional regularization is standard by the direct method of the calculus of variations. Indeed, the total energy to be minimized is given by

$$E[G] = E_{\mathrm{fit}}[G] + \omega_1 \|\mathrm{curl} G\|_{L^1} + \omega_2 \|\nabla G\|_{L^2}^2 + E_{\mathrm{coer}}[G] \,,$$

and we have $E[G] \geq 0$ and $E[I] < \infty$. Now along a minimizing sequence G_i, $i \in \mathbb{N}$, the G_i are bounded in $H^1(\Omega; \mathbb{R}^{2\times 2})$ due to the boundedness of the energy, in particular of the terms $\omega_2 \|\nabla G\|_{L^2}^2$ and $E_{\mathrm{coer}}[G]$. Thus, a subsequence converges weakly in $H^1(\Omega; \mathbb{R}^{2\times 2})$ against a tensor field G^*, which is a minimizer of E due to the weak lower semi-continuity of E ($\|\mathrm{curl} G\|_{L^1}$ and $\|\nabla G\|_{L^2}^2$ have convex lower-semicontinuous integrands and are thus lower semi-continuous, and $E_{\mathrm{fit}}[G]$ and $E_{\mathrm{coer}}[G]$ are lower semi-continuous by Fatou's lemma since we may assume pointwise convergence of G_i to G^*, where $\det G^* \neq 0$ almost everywhere due to E_{coer}). Thus we have the following:

Theorem 1. *Given an image $u \in C(\Omega; \mathbb{R})$, the energy $E[G]$ admits minimizers in $H^1(\Omega; \mathbb{R}^{2\times 2})$.*

Let us emphasize that the term $E_{\mathrm{coer}}[G]$ is necessary for theoretical purposes; however we omit it in all of our numerical experiments. The purpose of this term is to prevent G from taking unrealistic values that satisfy the term $E_{\mathrm{fit}}[G]$ but do not have a physical interpretation as a deformation gradient. A similar effect is instead achieved by a local minimization starting from the simple initialization of G described in Section 3.1.

The question remains what the interplay is between the L^2-norm of the gradient and the L^1-norm of the curl. The strength of the L^2 gradient regularization seemingly renders $\|\mathrm{curl} G\|_{L^1}$ superfluous, the more so because the total curl in a dislocation should be fixed a priori by fitting G to the crystal in the far field. To see this, consider a crystal on $\mathbb{R}^2$ with a dislocation of Burgers vector $\vec{b} = (b, 0)^T$ at the origin. Assuming the crystal to be linearly elastic, the corresponding

deformation tensor field $G_{\vec{b}}$ qualitatively behaves like

$$G_{\vec{b}}(r,\theta) = I + f_{\vec{b}}(r)\begin{pmatrix} -\sin\theta & \cos\theta \\ 0 & 0 \end{pmatrix}, \qquad f_{\vec{b}}(r) = \begin{cases} br/(2\pi\epsilon^2)\,, & r \le \epsilon\,, \\ b/(2\pi r)\,, & r > \epsilon \end{cases}$$

in cylindrical coordinates (r,θ), where ϵ is the radius of the dislocation core (see for instance [5]). Now consider the simplified toy model for automatic dislocation detection,

$$\tilde{E}[G] = \|G - G_{\vec{b}}\|_{L^2}^2 + \omega_1\|\mathrm{curl}G\|_{L^1} + \omega_2\|\nabla G\|_{L^2}^2\,,$$

in which we replaced the complicated fitting term $E_{\mathrm{fit}}[G]$ by the simpler L^2-difference to the true tensor field $G_{\vec{b}}$. Then the L^1 regularization of $\mathrm{curl}G$ is irrelevant:

Lemma 1. *The minimizers of $\tilde{E}$ are independent of the value of $\omega_1 \ge 0$.*

Proof. We will show that any minimizer G satisfies $\|\mathrm{curl}G\|_{L^1} = \|\mathrm{curl}G_{\vec{b}}\|_{L^1}$ so that the minimization of $\tilde{E}$ for different values of ω_1 is equivalent.

First we find $\|\mathrm{curl}G\|_{L^1} \ge \|\mathrm{curl}G_{\vec{b}}\|_{L^1}$: For the minimizing G we have (using the Cauchy–Schwarz inequality)

$$\begin{aligned}\infty > \|G - G_{\vec{b}}\|_{L^2}^2 &= \int_0^\infty \int_{\partial B_r(0)} |G - G_{\vec{b}}|^2 \,\mathrm{d}s\mathrm{d}r \ge \int_0^\infty \tfrac{1}{2\pi r}\big(\int_{\partial B_r(0)} |(G - G_{\vec{b}})n^\perp|\,\mathrm{d}s\big)^2\mathrm{d}r \\ &\ge \int_0^\infty \tfrac{1}{2\pi r}\big|\int_{\partial B_r(0)}(G - G_{\vec{b}})n^\perp\,\mathrm{d}s\big|^2\mathrm{d}r = \int_0^\infty \tfrac{1}{2\pi r}\big|\int_{B_r(0)}\mathrm{curl}G\,\mathrm{d}s - \int_{B_r(0)}\mathrm{curl}G_{\vec{b}}\,\mathrm{d}s\big|^2\mathrm{d}r\end{aligned}$$

and thus $\|\mathrm{curl}G\|_{L^1} \ge |\int_{\mathbb{R}^2}\mathrm{curl}G\,\mathrm{d}s| = |\int_{\mathbb{R}^2}\mathrm{curl}G_{\vec{b}}\,\mathrm{d}s| = \|\mathrm{curl}G_{\vec{b}}\|_{L^1}$ (in the above, $n^\perp$ denotes the unit normal to $\partial B_r(0)$, rotated by $\frac{\pi}{2}$).

Next we consider the minimizer $\hat{G}$ of $\tilde{E}$ for $\omega_1 = 0$ and show $\|\mathrm{curl}\hat{G}\|_{L^1} \le \|\mathrm{curl}G_{\vec{b}}\|_{L^1}$: Indeed, $\hat{G}$ satisfies the Euler–Lagrange equation $-\omega_2\Delta\hat{G} + \hat{G} = G_{\vec{b}}$, and so does the curl, $-\omega_2\Delta(\mathrm{curl}\hat{G}) + \mathrm{curl}\hat{G} = \mathrm{curl}G_{\vec{b}}$. Hence,

$$\begin{aligned}\|\mathrm{curl}G_{\vec{b}}\|_{L^1} &\ge \int_{\mathbb{R}^2}\mathrm{curl}G_{\vec{b}}\cdot\tfrac{\mathrm{curl}\hat{G}}{|\mathrm{curl}\hat{G}|}\,\mathrm{d}x = \int_{\mathbb{R}^2}(-\omega_2\Delta(\mathrm{curl}\hat{G}) + \mathrm{curl}\hat{G})\cdot\tfrac{\mathrm{curl}\hat{G}}{|\mathrm{curl}\hat{G}|}\,\mathrm{d}x \\ &= \int_{\mathbb{R}^2}\mathrm{curl}\hat{G}\cdot\tfrac{\mathrm{curl}\hat{G}}{|\mathrm{curl}\hat{G}|}\,\mathrm{d}x + \omega_2\int_{\mathbb{R}^2}\mathrm{tr}\left(\nabla(\mathrm{curl}\hat{G})^T\nabla\tfrac{\mathrm{curl}\hat{G}}{|\mathrm{curl}\hat{G}|}\right)\mathrm{d}x \ge \|\mathrm{curl}\hat{G}\|_{L^1}\,,\end{aligned}$$

since $\mathrm{tr}\left(\nabla(\mathrm{curl}\hat{G})^T\nabla\frac{\mathrm{curl}\hat{G}}{|\mathrm{curl}\hat{G}|}\right) = \frac{|\nabla(\mathrm{curl}\hat{G})|^2}{|\mathrm{curl}\hat{G}|} - \frac{|(\nabla(\mathrm{curl}\hat{G}))\mathrm{curl}\hat{G}|^2}{|\mathrm{curl}\hat{G}|^3} \ge 0$ (note that $\mathrm{curl}\hat{G}$ is sufficiently smooth and decaying for the above to make sense).

The above implies that $\hat{G}$ globally minimizes $\tilde{E}$ independently of the value of ω_1. □

In addition, the above proof shows that for a given crystal image we may expect our extracted tensor field to be just a smoothed version of the underlying true physical tensor field.

Despite the above, the term $\|\mathrm{curl}G\|_{L^1}$ turns out to be beneficial in numerical experiments, in fact whenever the crystal image contains noise. Just as in the widely used TV $-$ L^2 denoising, the L^1 regularization of $\mathrm{curl}G$ completely eliminates small noise so that in an otherwise curlfree region all curl due to noise is completely suppressed (see Fig. 8).

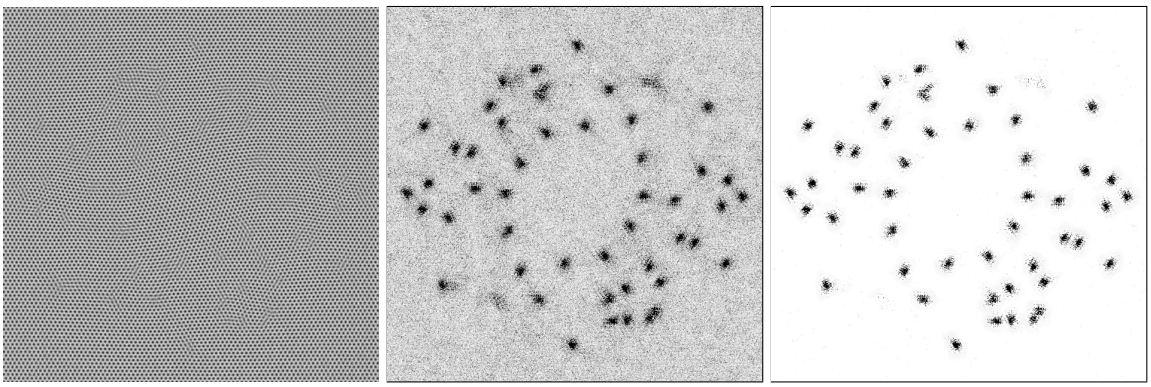

Figure 8: Noisy image of a polycrystal (left; 6% Gaussian noise was added to a phase field crystal simulation result with gray values ranging from 0 to 1, image size is such that interatomic distance is of order 1) and $|\mathrm{curl} G|$ after minimizing $E[G]$ with $\omega_1 = 0$, $\omega_2 = 0.12$ (middle) and $\omega_1 = 0.06$, $\omega_2 = 0.12$ (right).The penalization of $|\mathrm{curl} G|$ obviously suppresses most of the noise.

2.5 Incorporating the point group

The point group P of a Bravais lattice is the set of rotations and reflections which leave the lattice (or equivalently and, for us, more relevant, the corresponding lattice stencil $v_1, \dots, v_K$) invariant. For the example of a hexagonal crystal lattice, P consists of the six rotations R_α by the angles $\alpha = \frac{k\pi}{3}$, $k = 1, \dots, 6$, and six corresponding reflections. The point group introduces an equivalence relation on deformation tensors $G \in \mathbb{R}^{2\times 2}$: Two matrices $G_1, G_2 \in \mathbb{R}^{2\times 2}$ are defined to be equivalent if they only differ by an element of the point group, $G_1 \sim G_2 \Leftrightarrow G_2 = PG_1$ for some $P \in \mathrm{P}$.

If $G_2 \sim G_1$, then the Bravais lattice distorted by G_2^{-1} looks identical to the one distorted by G_1^{-1} so that for the purpose of grain segmentation these matrices should be identified with each other. Our energy does not automatically account for this identification which leads to generic problems when extracting a tensor field G from polycrystal images: Polycrystals frequently contain a closed chain of n adjacent grains connected by low angle grain boundaries. In other words, the corresponding tensor fields G_1 to G_n inside the grains differ only slightly between adjacent grains so that the grain boundaries consist of strings of dislocations in between which orientation and distortion smoothly change from G_j to G_{j+1}, $j = 1, \dots, n-1$ (Fig. 9 left). However, one will often notice a large jump from G_n to G_1 which is incompatible with a low angle grain boundary even though this would be the physically correct description. The reason is that for this grain boundary, G_1 actually has to be interpreted as PG_1 for a particular $P \in \mathrm{P}$ such that PG_1 and G_n can form the correct low angle grain boundary. If this is ignored, a high angle grain boundary with spurious curl concentration shows up in the extracted tensor field (Fig. 9 right).

Once such a problem is identified, our model can be appropriately adapted by introducing a set of cuts S which are smooth one-dimensional lines that can also form a network (we will later describe a heuristic method of how to find such a set S). Across S we allow the tensor field G to jump between different representatives of the same equivalence class of matrices, that is, we redefine our variational energy as

$$E[G] = E_{\mathrm{fit}}[G] + \int_{\Omega\setminus S} \omega_1 |\mathrm{curl} G| + \omega_2 |\nabla G|^2 \,\mathrm{d}x + E_{\mathrm{coer}}[G] \tag{2}$$

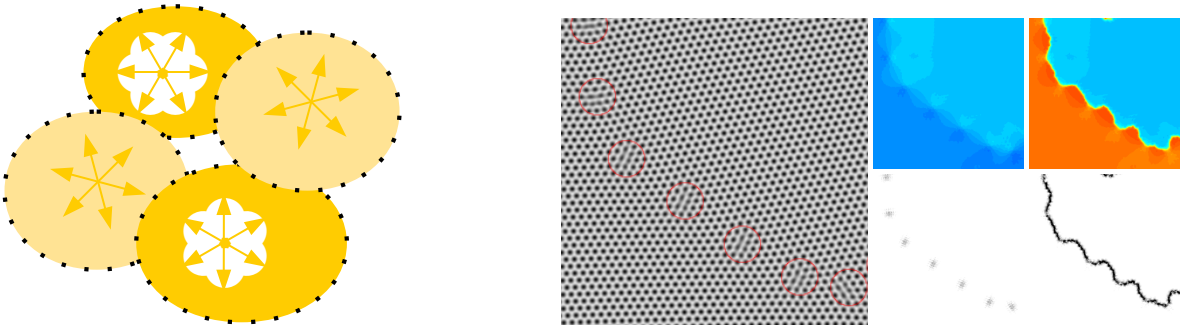

Figure 9: Left: Schematic of a chain of adjacent grains in a polycrystal which form pairwise low angle grain boundaries. The arrows indicate the local lattice stencil, and the local distortion tensor G in each grain approximately is a counterclockwise rotation R_α by the angle α, with $\alpha = 0, \frac{\pi}{12}, \frac{\pi}{6}, \frac{\pi}{4}$ along the chain, starting from the top grain in counterclockwise direction. However, to describe the boundary between the top and the top right grain as a low angle boundary, the angle $\alpha = 0$ would have to be interpreted as the equivalent $\alpha = \frac{\pi}{3}$. Right: Crystal image with low angle grain boundary (dislocations manually circled in red) as well as local crystal orientation (top) and $|\mathrm{curl} G|$ (bottom) for the extracted tensor field G, once with the bottom grain interpreted as oriented at $\alpha = 0$ (left), once with $\alpha = \frac{\pi}{3}$ (right). The interpretation $\alpha = \frac{\pi}{3}$ leads to spurious curl concentration and the detection of a non-physical high angle grain boundary, shown by the sharp change in color in the upper plot and solid curve indicating curl concentration in the lower plot. In the interpretation $\alpha = 0$ a low angle grain boundary with isolated dislocations is correctly identified.

which is to be minimized under the constraint

$$G^- \sim G^+ \text{ almost everywhere on } S\,, \tag{3}$$

where G^- and G^+ respectively represent the value of the tensor field G on either side of S. Hence, if S cuts a grain in two pieces, half of the grain may be ascribed a distortion $P_1 G$ while the other half may be assigned a distortion $P_2 G$, where $P_1, P_2 \in \mathrm{P}$ will be selected by the energy minimization in such a way that all grain boundaries are correctly captured.

3 Algorithm

Here we describe the algorithm used to detect distortion and crystal defects in images based on minimizing (2) subject to (3). In the exposition we will drop the E_{coer} term from the variational energy. As previously mentioned, initializing the minimization scheme with a properly chosen G_0 gives good results in the absence of the E_{coer} term. Subsequently, we will discuss the initialization of G_0, the solution of (2) in the case that $S = \emptyset$, a heuristic scheme to choose a good candidate for S, and, lastly, how to obtain a local minimum of (2) for nonempty S, subject to the constraint (3).

3.1 Initialization of G as rotations

The images that are appropriate inputs to this algorithm consist primarily of perfectly crystalline regions. In these regions, given the appropriate stencil $\{v_k\}$, G can be correctly described by a simple rotation in the absence of strain. For this reason, we initialize the algorithm with the initial

guess $G_0(x)$ chosen to be a simple rotation matrix at each x:

$$G_0(x) = \begin{pmatrix} \cos\alpha(x) & -\sin\alpha(x) \\ \sin\alpha(x) & \cos\alpha(x) \end{pmatrix} := R_{\alpha(x)},$$

where $\alpha(x) \in [\theta, \theta + (2\pi)/K)$ for some arbitrarily fixed $\theta \in \mathbb{R}$ is chosen to minimize the integrand of (1). In practice, we use an exhaustive search over a set of p linearly-spaced angles $\{\alpha_j\}$ to find a minimizing angle α_ℓ in this set; then choose α^* to minimize the quadratic interpolation of $\sum_k (u(x+R_\alpha^{-1}v_k)-u(x))^2$ at $\alpha = \alpha_\ell, \alpha_{\ell\pm1}$, (where the subscript $(\ell\pm1)$ is to be interpreted $\mod p$). In practice, we find it sufficient to choose $p \sim 24/K$. This is a very simple initialization which does not enforce any continuity of $G_0(x)$, however numerical results indicate that it is sufficient even in the absence of the E_{coer} term.

In the strained crystal case, the above suffices if the strain is small. If the strain M is approximately known, that is, the deformed crystal position x is expected to be obtained from the reference crystal position $\bar{x}$ by $x = MR^{-1}\bar{x}$, or $\bar{x} = RM^{-1}x$, then it can be accounted for by testing strained rotations $R_{\alpha j}M^{-1}$ rather than R_{α_j} alone in the previous procedure. In most applications, one would not expect a large, unknown strain; however, if this case does arise, one could include the term E_{coer} in the energy for the reasons previously discussed. We will not discuss this possibility further in this work.

3.2 Unconstrained minimization via split Bregman iteration

First, we discuss the local minimization of (2) in the case that $S = \emptyset$. In this case, the constraint (3) is trivially satisfied. Due to the L^2–L^1 character of the energy (2), we apply the split Bregman algorithm to obtain a local minimizer.

The split Bregman iteration is obtained from the Bregman iteration, which we describe first. For simplicity, we consider only linear constraints here, though the Bregman iteration may be applied to more general convex constraints. Given a convex functional $E_c[v]$ subject to a linear constraint

$$\min_v E_c[v], \text{ subject to } Av = b\,, \tag{4}$$

the Bregman iteration [4, 11] replaces this problem with an unconstrained formulation:

$$\min_v E_c[v] + \lambda\|Av - b\|_2^2, \tag{5}$$

with a parameter $\lambda > 0$. In general, the solutions to (4) and (5) do not agree except in the limit $\lambda \to \infty$. The advantage of the Bregman iteration (described subsequently) is that if solutions to $Av = b$ exist, the Bregman iteration applied to (5) converges to some v for which $Av = b$, as proven in [11]. Furthermore, if one of the iterative solutions v^* found by the Bregman iteration has the property that $Av^* = b$, then v^* solves (4) [9]. Often only a few steps of the Bregman iteration are required to obtain the minimizing v, though this is highly dependent on the specific problem to be solved and the choice of initial guess v_0. The Bregman iteration is quite simple:

$$\begin{aligned} v^{k+1} &= \operatorname*{argmin}_v E_c[v] + \lambda\|Av - b^k\|_2^2\,, \\ b^{k+1} &= b^k + b - Av^{k+1}\,. \end{aligned}$$

This iteration is simple to perform as long as one can solve the minimization sub-problem updating v^k to v^{k+1} efficiently.

Given an energy

$$E_c[v] = \|\zeta[v]\|_1 + H[v] \tag{6}$$

with both $\|\zeta[v]\|_1$ and $H[v]$ convex, the split Bregman formulation repurposes the Bregman iteration to minimize this L^2–L^1 type energy [9]. The idea is to introduce a new variable ψ which corresponds to the L^1-type term and then to add the appropriate constraint term to enforce this correspondence

$$\min_{v,\psi} \|\psi\|_1 + H[v] \quad \text{subject to } \zeta[v] = \psi.$$

This minimization problem is solved by the Bregman iteration, which now appears as

$$\begin{aligned} \{v^{k+1}, \psi^{k+1}\} &= \operatorname*{argmin}_{v,\psi} \|\psi\|_1 + H[v] + \lambda\|\psi - \zeta[v] - b^k\|_2^2\,, \\ b^{k+1} &= b^k - \psi^{k+1} + \zeta[v^{k+1}]\,. \end{aligned} \tag{7}$$

The "split" in "split Bregman" refers to the suggestion to solve the minimization subproblem by a single alternating minimization step on v^{k+1} and ψ^{k+1}, which allows for v^{k+1} to depend *only* on L^2-type terms, and gives an explicit form for ψ^{k+1}:

$$v^{k+1} = \operatorname*{argmin}_{v} H[v] + \lambda\|\psi^k - \zeta[v] - b^k\|_2^2\,, \tag{8}$$

$$\psi^{k+1} = \text{shrink}(\zeta[v^{k+1}] + b^k, \tfrac{2}{\lambda})\,, \tag{9}$$

$$b^{k+1} = b^k - \psi^{k+1} + \zeta[v^{k+1}]\,, \tag{10}$$

with

$$\text{shrink}(z, \omega) := \tfrac{z}{|z|} \max(|z| - \omega, 0)\,.$$

The advantage of the split Bregman formulation is that the inner minimization problem now has a strictly L^2-type convex energy, while the original problem contains both L^1- and L^2-type terms. As split Bregman differs from the original Bregman iteration only in the minimization procedure for (7) (but not in the formulation of the problem), if a minimizer ψ^* to (7) with $\psi^* = \zeta[v]$ is found, this ψ^* is a minimizer of (6).

Though the energy $E_{\text{fit}}[G]$ is nonconvex, we apply the split Bregman scheme:

$$\begin{aligned} H[G] &= E_{\text{fit}}[G] + \omega_2\|\nabla G\|_2^2\,, \\ \zeta[G] &= \omega_1 \text{curl}(G) \end{aligned}$$

and apply the update (8)–(10). While the convergence theorem is not valid for nonconvex $E_{\text{fit}}[G]$, the scheme still appears to be quite successful at finding a local minimum for (2.2). More details on the computational cost are presented in Section 3.5.

We utilize the Fletcher–Reeves nonlinear conjugate gradient method with Armijo stepsize control to perform the update (8), as methods incorporating the Hessian will not provide much benefit given a nonconvex fitting energy E_{fit} but have significantly increased computational cost and memory requirements. We utilize biquadratic interpolation on the nine nearest grid cells to x to interpolate the image value $u(x)$. This interpolation is not even continuous for arbitrary data u, much less twice continuously differentiable; however, for reasonably well-resolved images, the interpolation is at least approximately smooth.

3.3 Choosing the cut set S automatically

The cut set S must be chosen before the minimization problem defined by (2) and the constraint (3) may be solved. Here we present a heuristic algorithm for choosing a good set S that eliminates the problem shown in Figure 9. The only user input to this algorithm is the set of adjacent atom sites $\{v_k\}$, which is already required by the energy calculation.

To motivate this algorithm for choosing S, imagine an unstrained polycrystal. Away from dislocations and grain boundaries, G can be described by a single, constant rotation within each individual grain. Choose $\theta = 0$ as the lower bound of the range that $\alpha(x)$ is chosen from. The initialization will choose the angles $\alpha(x) \in [0, 2\pi/K)$. A low angle grain boundary, with one grain oriented as $\alpha_1 = \varepsilon_1$ and the other as $\alpha_2 = -\varepsilon_2$, with $0 < \varepsilon_1, \varepsilon_2 \ll 1$, will initially be assigned the matrices $G_0^- = R_{\varepsilon_1}$ on one side and $G_0^+ = R_{2\pi/K-\varepsilon_2}$ on the other due to the angle bounds on the initial condition. The minimization algorithm will not be able to resolve this jump in G. The boundary between these grains will be detected to be a region of high curl$G$, rather than being correctly interpreted as a low-angle grain boundary. However, a different choice of θ, for example $\theta = -\pi/K$, alleviates this problem: the matrices $G_0^- = R_{\varepsilon_1}$ and $G_0^+ = R_{-\varepsilon_2}$ will be detected by the initialization scheme.

In general, for any choice of θ, a low angle grain boundary with orientation angles slightly greater than and less than θ will be misinterpreted. To avoid this problem we obtain local minimizers from two different initializations, one with the choice $\theta = 0$ and one with $\theta = -\pi/K$, and we define the cut set S to be the interface between regions where each of the two solutions is preferred. See Algorithm 1 for details. Note that, while Algorithm 1 produces a set S that is the union of closed curves, S may, in general, also contain open curves. We use the choice of S and the initial guess G_0 obtained from Algorithm 1 as input to the variational problem defined by (2) and (3) to obtain the final tensor field G. The expense of Algorithm 1 is not prohibitive, as the work to compute G_1 and G_2 ensures that the guess G_0 is quite good and only limited additional work is needed to evolve G_0 to the final G.

Figure 10 visualizes the effect of Algorithm 1. The solutions G_1 and G_2 are computed for $\theta = 0$ and $\theta = \pi/6$, respectively, with $S = \emptyset$. The set Σ and the interface $S = \partial\Sigma$ are computed as described in Algorithm 1, and G_0 is used to initialize the minimization of (2). The resulting tensor field G is shown at right. The top row shows the local crystal orientations obtained from polar decomposition of the tensor fields. The bottom row shows the norm of the curl of the tensor fields. Observe that jumps in G are permitted across S without penalty as long as $G^+ \sim G^-$ across S.

3.4 Numerical implementation on GPU

Here we describe the numerical implementation used to solve the variational problem (2)–(3). We indicate the numerical discretization of the domain Ω, how the discrete version of the energy is calculated, how the cut set S is handled numerically, and describe the minimization procedure and its implementation on the GPU in MATLAB.

For simplicity of exposition, we assume a square domain $\Omega = [0, L)^2 \subset \mathbb{R}^2$ and, for computational simplicity, we use periodic boundary conditions. Numerical tests indicate that for non-periodic input images u the evolution of G is sufficiently local that discontinuities in u across the periodic boundary only affect G within a few atomic spacings of $\partial\Omega$. Ω is discretized on a $(n \times n)$ rectangular grid with uniform grid spacing $\Delta x = n/L$. We denote discrete grid locations and function values by $\tilde{f}_{ij} = \tilde{f}(x_{ij}) = f(j\Delta x, i\Delta x)$.

Algorithm 1 Heuristic algorithm for choosing a good cut set S and initalization G_0.

Choose constants C, σ, $\tau > 0$ with σ, $\tau \propto \ell$, the atomic spacing.

1. Locally minimize (2) with $S = \emptyset$, using $\theta = 0$ and $\theta = -\pi/K$ as the parameter choice for initialization as described in Section 3.1, which yields local minima G_1 and G_2.

2. If $\int_\Omega |\mathrm{curl}(G_2)| < \int_\Omega |\mathrm{curl}(G_1)|$, relabel G_2 as G_1 and G_1 as G_2.

3. Compute $|\mathrm{curl}(G_1)| * \Psi_\sigma$ and $|\mathrm{curl}(G_2)| * \Psi_\sigma$, where Ψ_σ is a Gaussian kernel of width σ.

4. Define the region $\Sigma_2' = \{x : |(\mathrm{curl}(G_1)| * \Psi_\sigma)(x) - (|\mathrm{curl}(G_2)| * \Psi_\sigma)(x) > C\}$ and define $\Sigma_2 = \mathrm{Dil}(\Sigma_2', \tau) = \{x : \mathrm{dist}(x, \Sigma_2') < \tau\}$.
   Define $\Sigma_1 = \{x : |(\mathrm{curl}(G_2)| * \Psi_\sigma)(x) - (|\mathrm{curl}(G_1)| * \Psi_\sigma)(x) > C\}$.

5. Set $\Sigma = \Sigma_2 \setminus \Sigma_1$ and $S = \partial\Sigma$. Choose

$$G_0(x) = \begin{cases} G_1(x), & x \in \Sigma^c \\ G_2(x), & x \in \Sigma. \end{cases}$$

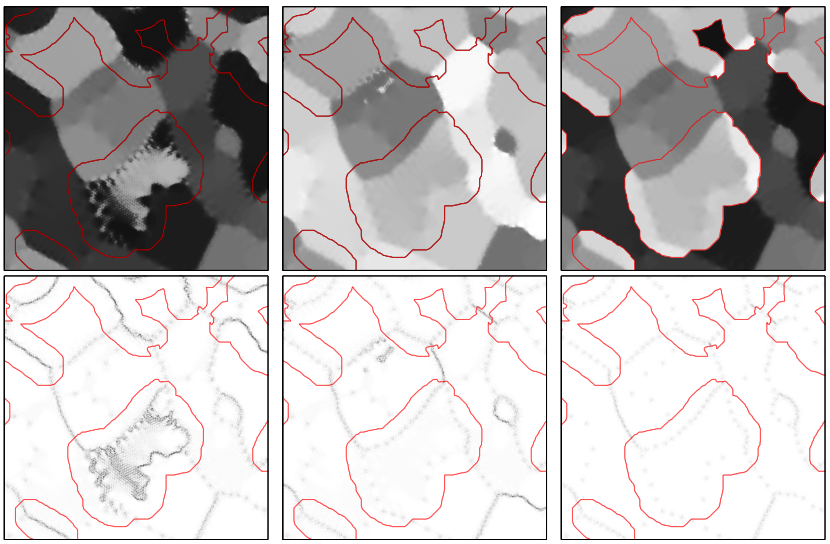

Figure 10: Top row: Rotation matrices associated with G_1, G_2, and result G. Bottom row: $|\mathrm{curl}(G_1)|$, $|\mathrm{curl}(G_2)|$, $|\mathrm{curl}(G)|$. The cut set S, obtained from Algorithm 1, is shown in red. G_1 and G_2 both exhibit regions of spurious curl due to a spatially inconsistent identification of crystal orientations. G, however, which is allowed to jump across S as long as $G^- \sim G^+$ across S, does not suffer from that deficiency.

The fitting energy E_{fit} is discretized as

$$\tilde{E}_{\text{fit}}[\tilde{G}] = \Delta x^2 \sum_{i,j=0}^{n-1} \sum_{k=1}^{K} \left(u(x_{ij} + \tilde{G}_{ij}^{-1} v_k) - \tilde{u}_{ij} \right)^2 .$$

In order to compute $u(x_{ij} + \tilde{G}_{ij}^{-1} v_k)$ in terms of the discretized image $\tilde{u}$, we apply biquadratic interpolation. We note that biquadratic interpolation is not even continuous in general, but gives good numerical results for reasonably well-resolved image data u. The minimization process also requires the gradient $\nabla_x u(x_{ij} + \tilde{G}_{ij}^{-1} v_k)$, which is computed exactly for the biquadratic polynomial fitted to each grid cell.

In order to handle the regularizing terms of (2) numerically, we simultaneously enforce the constraint that $G^- \sim G^+$ across S and handle the integration over $\Omega \setminus S$ by replacing the usual discretized partial derivatives

$$\frac{\partial \tilde{G}_{ij}}{\partial x_h} = \frac{\tilde{G}(y_{ij;h}) - \tilde{G}_{ij}}{\Delta x} ,$$

where $y_{ij;h} = x_{ij} + e_h \Delta x$ and e_h is the h^{th} standard Euclidean basis vector, by

$$\frac{\hat{\partial} \tilde{G}_{ij}}{\hat{\partial} x_h} = \begin{cases} \frac{\tilde{G}(y_{ij;h}) - \tilde{G}_{ij}}{\Delta x} , & (x_{ij}, y_{ij;h}) \notin \tilde{S} , \\ \frac{R(x_{ij}, y_{ij;h}) \tilde{G}(y) - \tilde{G}_{ij}}{\Delta x} , & (x_{ij}, y_{ij;h}) \in \tilde{S} , \end{cases}$$

where $(x, y) \in \tilde{S}$ if and only if $\gamma x + (1 - \gamma) y \in S$ for some $\gamma \in [0, 1)$ and $R(x, y)$ is defined by

$$R(x, y) = \underset{P \in \mathrm{P}}{\operatorname{argmin}} \| P \tilde{G}(y) - \tilde{G}(x) \|_2 .$$

Then the regularizing terms in (2) are computed as

$$\int_{\Omega \setminus S} \omega_1 |\mathrm{curl} G| + \omega_2 |\nabla G|^2 dx \approx \Delta x^2 \sum_{i,j=0}^{n-1} \omega_1 |\hat{\mathrm{curl}} \tilde{G}_{ij}| + \omega_2 |\hat{\nabla} \tilde{G}_{ij}|^2 ,$$

with

$$|\hat{\mathrm{curl}} \tilde{G}_{ij}| = \left[\left(\frac{\hat{\partial} \tilde{G}_{ij}^{(12)}}{\hat{\partial} x_1} - \frac{\hat{\partial} \tilde{G}_{ij}^{(11)}}{\hat{\partial} x_2} \right)^2 + \left(\frac{\hat{\partial} \tilde{G}_{ij}^{(22)}}{\hat{\partial} x_1} - \frac{\hat{\partial} \tilde{G}_{ij}^{(21)}}{\hat{\partial} x_2} \right)^2 \right]^{1/2}$$

and

$$|\hat{\nabla} \tilde{G}_{ij}|^2 = \sum_{\alpha, \beta = 1}^{2} \left(\frac{\hat{\partial} \tilde{G}_{ij}^{(\alpha\beta)}}{\hat{\partial} x_1} \right)^2 + \left(\frac{\hat{\partial} \tilde{G}_{ij}^{(\alpha\beta)}}{\hat{\partial} x_2} \right)^2 ,$$

where $\tilde{G}_{ij}^{(\alpha\beta)}$ denotes the (α, β) entry of the matrix $\tilde{G}_{ij}$.

3.5 Computational effort

For images as in Figure 1 we generally found a total of ten split Bregman steps to be sufficient given an appropriate choice of G_0. If Algorithm 1 is used to compute G_0, we suggest taking twenty more split Bregman steps, ten each to produce G_1 and G_2. For each Bregman step, we perform five

Table 1: GPU Computation times in MATLAB for sections of image in Figure 11, left.

image size	total time	split Bregman	NCG descent	inter-polation
2048^2	87.4 s	83.3 s	73.1 s	20.5 s
1024^2	40.9 s	38.1 s	33.1 s	8.73 s
512^2	14.0 s	11.4 s	9.88 s	3.34 s
256^2	11.7 s	8.97 s	7.80 s	3.40 s
128^2	11.1 s	8.93 s	7.81 s	3.60 s

nonlinear conjugate gradient steps for (8). The code is implemented in MATLAB. The included basic GPU support is utilized without any further optimizations. Computation times on a Tesla S10 GPU are displayed in Table 1 for different image sizes. The computational time is approximately three times longer if the cost of computing G_0 via Algorithm 1 is included.

For an image with 2048^2 pixels, the minimization takes approximately one and a half minutes, with the vast majority of the time (over 95%) spent inside the split Bregman iteration. Most of the iteration time is spent in the nonlinear conjugate gradient steps which in turn spend a third of their time interpolating the image u to evaluate the fitting energy and its gradient. We observe a strongly sublinear scaling with image size due to the high GPU parallelization. With more careful memory allocation or larger GPU memory (the GPU employed had 4 GB of memory available), the runtimes are short enough that it should be possible to apply this algorithm to even larger images reasonably quickly.

The interpolation used to evaluate $u(x + G^{-1}(x)v_k)$ at each pixel x and each stencil vector v_k appears to be the compuational bottleneck. This is to be expected; such a computation requires highly unstructured memory access which is not naturally achieved by the highly parallel GPU architecture. An optimized version of this algorithm implemented outside of MATLAB should allow for significant speedup; alternatively, strong performance may be attainable in a shared memory distributed computing paradigm such as OpenMP.

4 Results

Here we present numerical results of our method for different types of input data, both from computer simulations and real images. We examine (1) a time series of output from a phase field crystal simulation (demonstrating temporal consistency of the output from our algorithm), (2) a synthetically strained crystal image (showing the ability of our algorithm to detect strain and also that the presence of isolated dislocations can decrease total strain in a crystal), (3) a variety of other crystal types and a real HRTEM image (indicating the robustness of our method), and (4) a series of evolutions closely related to the phase field crystal evolution, measuring the coarsening rate of the evolutions directly from concentrations of curl measured using our algorithm.

4.1 Coarsening simulations

Defects in atomic crystals are associated with a physical energy (in particular with elastic energy due to the crystal strain induced by the defects). Their time evolution is typically governed by dissipative, energy-driven dynamics: The crystal defects move and partially annihilate, thereby

decreasing the total energy. This becomes visible as a coarsening phenomenon; over time, the average crystal grain size increases. The process can also be simulated, for instance using phase field crystals (PFC, [6]). Figure 11 shows the result of our method applied to four consecutive snapshots in time from such a coarsening simulation. From top to bottom, local crystal orientation, volume strain, elastic strain, and defect locations are depicted. The crystal orientation is given by the rotation angle of the rotation matrix R closest to $F = G^{-1}$ (which can be computed via polar decomposition $F = R\sqrt{F^T F}$), the local volume change relative to the perfect crystal is $\det F - 1$, elastic strain can be measured as $\mathrm{dist}(G^{-1}, \mathrm{SO}(2))$, and the defect locations are places with high curl magnitude $|\mathrm{curl} G|$, as previously discussed. Grain boundaries and dislocations are clearly visible (some low angle grain boundaries show up as strings of dislocations). These are also the places, where elastic and volume strain are concentrated. Single dislocations produce a dipole in the visualization of volume strain with volume dilation on one side and compression on the opposite side, so that the direction of the corresponding Burgers vector can be inferred in an intuitive manner. It is obvious that the grains grow over time, reducing the amount of defects, though some isolated dislocations seem to survive inside large grains.

Let us emphasize that to the authors' knowledge there exists no comparable method to identify defects and strain in a crystal image. The methods in [3, 13] can only yield a piecewise constant approximation to just the top row in Figure 11 (which obviously is very unsatisfactory near defects), where from the computational viewpoint the required memory and complexity scales like the number of allowed discrete orientations. The method in [2] is also restricted to piecewise constant orientations, but it additionally provides a global crystal deformation from which a strain (e. g. second and third row) could be derived. However, a global smooth deformation is inappropriate to capture the very localized distortions near defects, in particular, strain near dislocations cannot be found since it cannot be represented by a continuous deformation (cf. Sec. 2.1). The method in [12] can only provide the top row and must fail as soon as the crystal contains stronger distortions (as e. g. in the next section). Finally, the methods analyzing molecular dynamics [1, 14, 15] cannot be applied to images.

4.2 Strained crystals

Our proposed method yields the crystal strain field G as its primary output, which is of particular advantage when deformed or non-equilibrated crystals are analyzed. A simple experiment serves for illustration: At each numerical time step in a PFC simulation we artificially shift a horizontal strip inside the material by one pixel, thus creating a drag force on the adjacent crystal regions. Figure 12 shows the effect, depicting the amount of shear (given by G_{12}) and the elastic strain $\mathrm{dist}(G^{-1}, \mathrm{SO}(2))$. We can now compare the result of this experiment for two different crystals, one with defects and one without. A histogram of the elastic strain immediately reveals that—though the peak strain is higher in the crystal with defects—the strain on average is higher in the configuration without defects, so the defects have a relaxing effect on the rest of the crystal.

4.3 Different crystal types

The method also works robustly on experimental images. Figure 13 shows a photograph of a bubble raft, a number of bubbles floating on a liquid surface. Such bubbles typically form crystalline patterns with many characteristics of atomic lattices for which reason they are sometimes examined as a very illustrative model of defect movement in atomic crystals. Our method identifies a grain

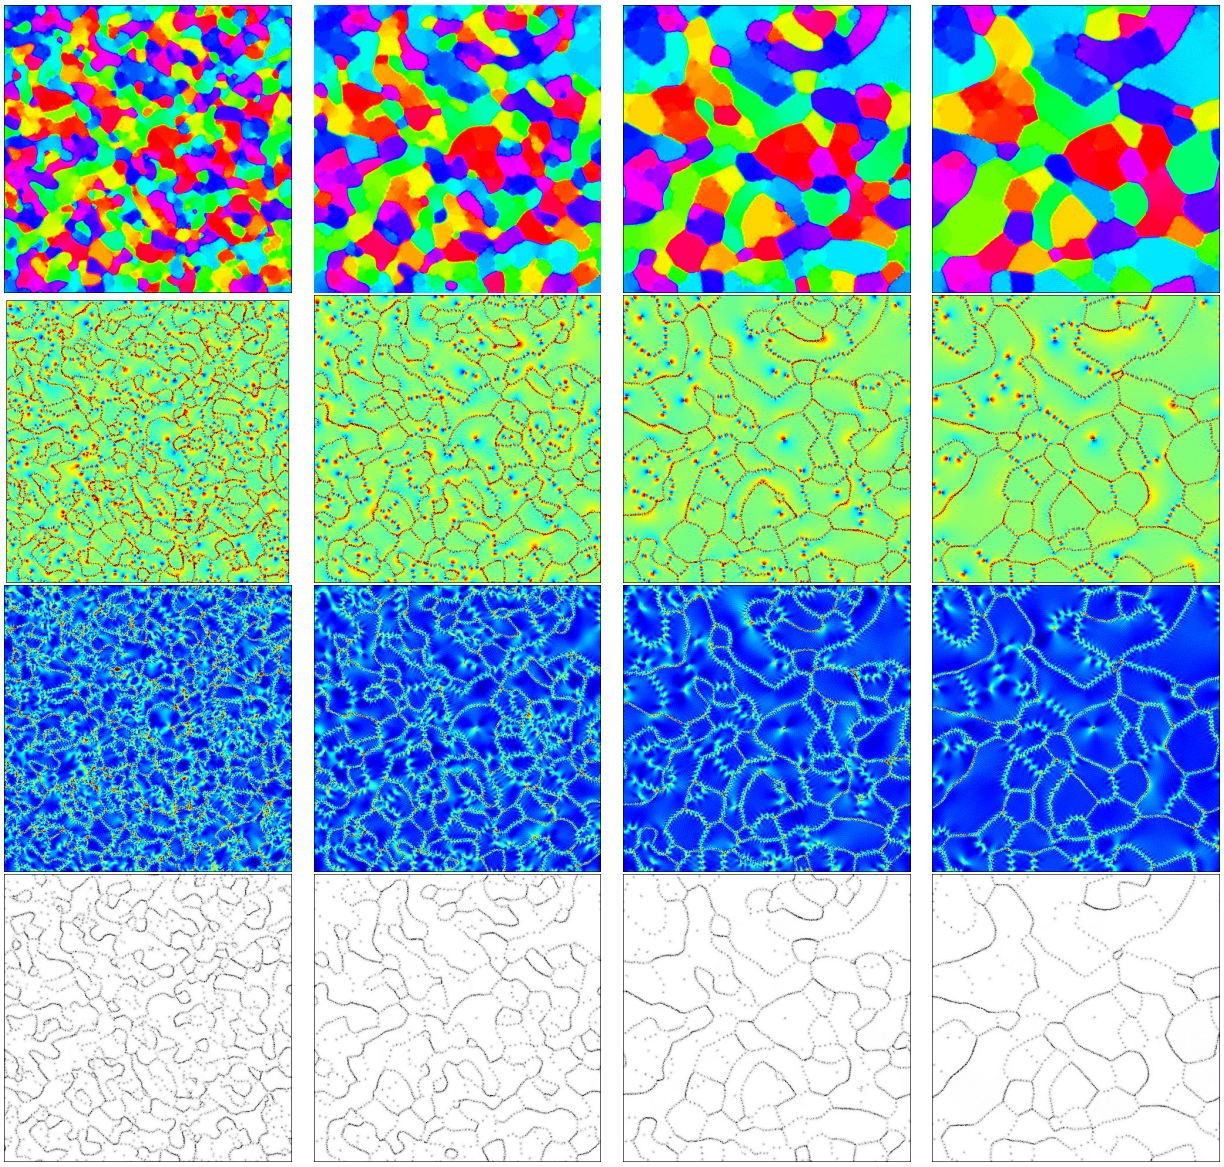

Figure 11: Results of our algorithm applied to different snapshots of a phase field crystal (PFC) coarsening simulation. From top to bottom we show the time evolution of local crystal orientation (0 $\pi/3$), local volume distortion of the crystal measured as $\det G^{-1} - 1$ (-0.05 0.05), elastic strain measured as $\mathrm{dist}(G^{-1}, \mathrm{SO}(2))$ (0 0.1), and $|\mathrm{curl} G|$.

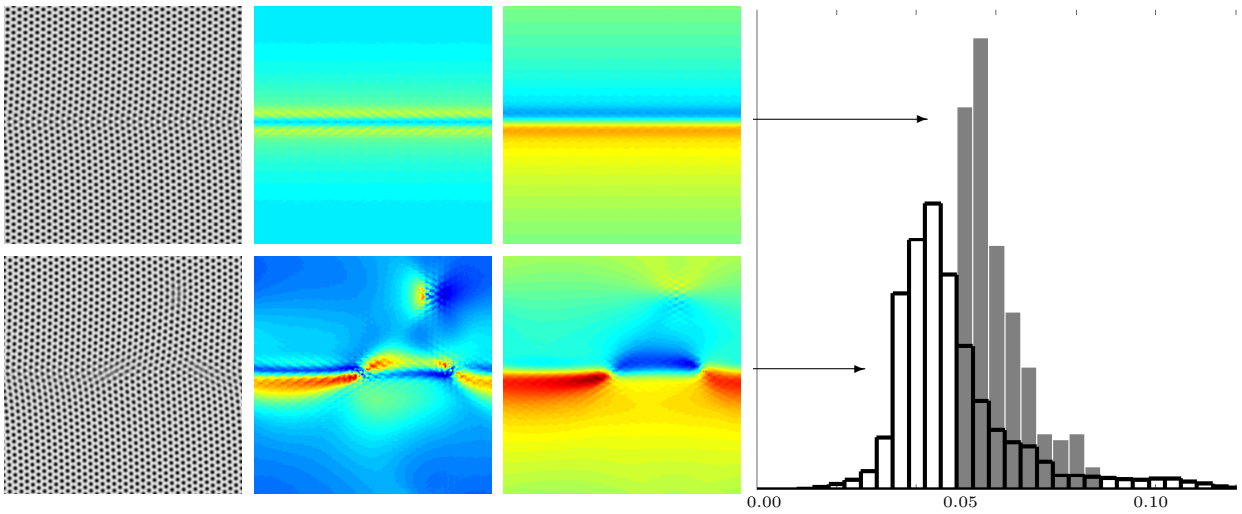

Figure 12: PFC simulation of a sheared crystal without (top) and with (bottom) defects, showing the crystal image (left), the elastic strain $\mathrm{dist}(G^{-1}, \mathrm{SO}(2))$ (middle, 0 0.15), and G_{12} (right, -0.2 0.2). The rightmost graph shows a histogram of $\mathrm{dist}(G^{-1}, \mathrm{SO}(2))$. The crystal with defects has a higher peak strain, but lower overall strain.

boundary as well as isolated dislocations, even though the reflections and shadows vary strongly across the image.

Figure 14 shows a high resolution TEM image of a nanocrystalline palladium thin film, courtesy of Nick Schryvers, from the work [16]. Despite the noise and the different shading across the image, the method finds a concentration of $\mathrm{curl}G$ at the image center, corresponding to a dislocation. It also identifies two distinct regions of different crystal orientation, even though the angle difference in between is only 12°. These two orientations belong to symmetric crystal variants so that the interface represents a twin boundary. Note that twin boundaries do not produce concentrations in $\mathrm{curl}G$ as grain boundaries do. In fact, $\mathrm{curl}G$ is zero across a twin boundary since the strain or deformation gradient fields on either side are rank-one-connected.

The synthetic reptile skin from Figure 4 represents an equally challenging input, since the hexagonal scales pattern contains lots of dislocations and the size of the grid cells varies considerably across the image. This makes an initialization of G with rotations seem rather inappropriate, nevertheless it is sufficient for the minimization to identify all dislocations correctly. Of course, for the initialization one could in principle also vary the stencil size besides its orientation to start the optimization closer to the desired minimum.

Lastly, in Figure 15 we show the analysis of a crystal image from a PFC simulation which exhibits a square lattice instead of the six-fold rotational symmetry of hexagonal crystals. The underlying PFC energy is a modification of the standard PFC energy as presented in [17, Sec. VI]. It was originally proposed in [10] in the context of the Swift–Hohenberg equation. An appropriate stencil for this image is given by $v_1 = (\ell, 0)^T$, $v_2 = (0, \ell)^T$, $v_3 = (-\ell, 0)^T$, $v_4 = (0, -\ell)^T$ with ℓ being the interatomic distance. Again, grain boundaries and dislocations are properly identified. In this particular simulation, the grains are more or less aligned to each other, varying mostly by only 15° in orientation.

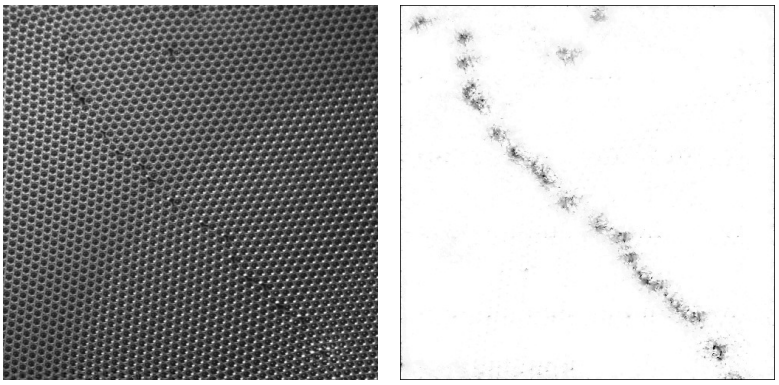

Figure 13: Photograph of a bubble raft (courtesy Don Stone, copyright 2009 Board of Regents of the University of Wisconsin System, `http://homepages.cae.wisc.edu/ ~stone/bubble%20raft%20movies.htm`) and detected regions of curl concentration, showing a grain boundary and single dislocations.

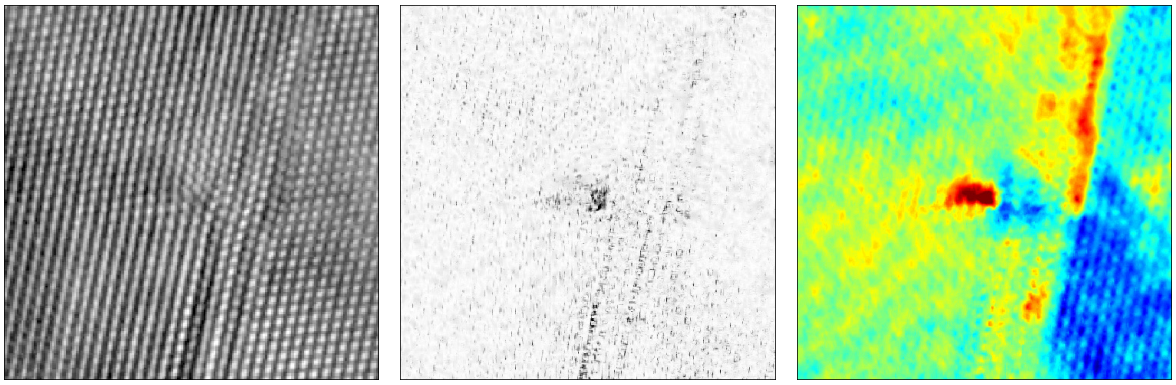

Figure 14: HRTEM image of a nanocrystalline palladium thin film from [16], detected regions of curl concentration showing an isolated dislocation, and the crystal orientation obtained from polar decomposition ($14°$ $26°$) showing a strong twin boundary.

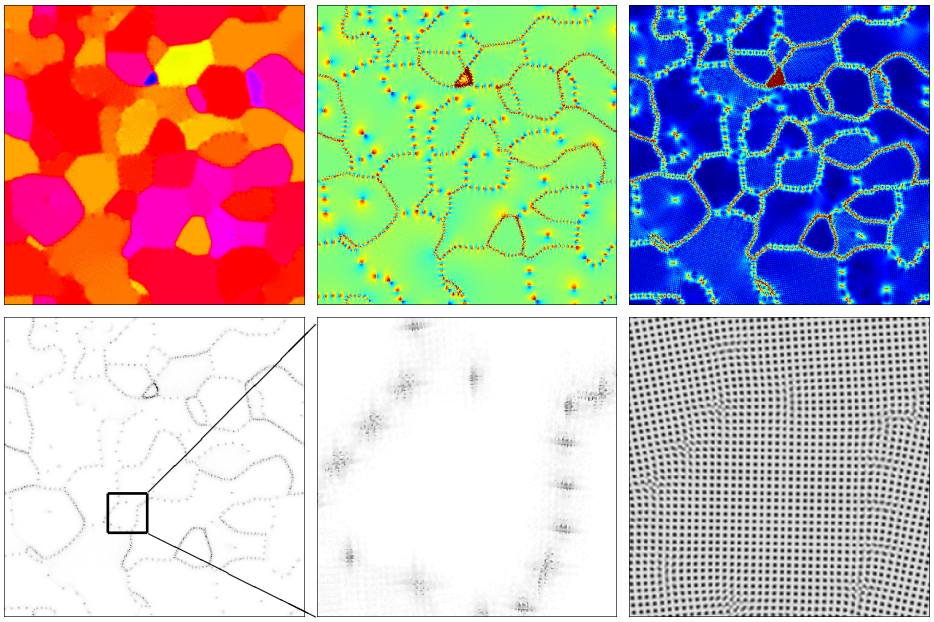

Figure 15: Results of our algorithm applied to microstructure with cubic symmetry. We show the local crystal orientation (0 $\pi/2$, top left), local volume distortion of the crystal measured as $\det G^{-1} - 1$ (-0.05 0.05, top middle), elastic strain measured as $\mathrm{dist}(G^{-1}, \mathrm{SO}(2))$ (0 0.05, top right), $|\mathrm{curl} G|$ (bottom left), a zoom-in on $|\mathrm{curl} G|$ (bottom middle), and a zoom in on the microstructure (bottom right).

4.4 Simple application: Scaling law for characteristic lengths

The tool devised in this article makes it easy to extract macroscopic characteristics from polycrystals, such as the average length scale of a grain. Such macroscopic characteristics could be used, for example, to validate physical models by comparing corresponding simulations with physical experiments. As a simple example, consider the model of crystal evolution given by the $H^{-\beta}$ gradient descent of the following variant of the PFC energy,

$$J[u] = \int_\Omega \frac{1}{2}[(-\Delta)^{-\alpha}(\Delta u + u)]^2 + \frac{1}{4}(u^2 - \delta)^2 \,\mathrm{d}x\,,$$

where $\alpha, \beta \in \mathbb{R}$. This yields a time series of images as in Figure 1 left. To obtain the distance of each pixel to the nearest defect (of which the average over all pixels is the average length scale) we postprocess our data as follows. $|\mathrm{curl} G|$ is first smoothed with a Gaussian kernel, and then a distance map is computed to all pixels where the smoothed curl has a larger value than a specified cutoff.

Figure 16 shows that different values of α and β yield the same coarsening rate in time, even though the final polycrystalline configurations look slightly different. The average length scale seems to behave roughly like $t^{\frac{1}{3}}$ or even $t^{\frac{1}{4}}$. Figure 16 additionally shows the distribution of atoms (or pixels) to the nearest crystal defects, and this distribution seems to be rather stable.

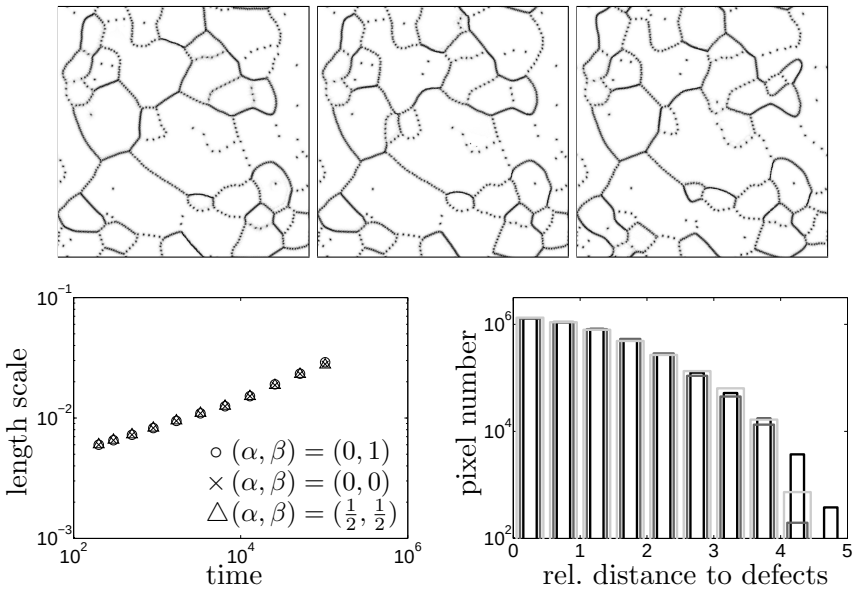

Figure 16: Top: Final polycrystal for the evolution of the PFC variant with $(\alpha, \beta) = (0, 1)$, $(0, 0)$, and $(\frac{1}{2}, \frac{1}{2})$. Bottom left: Average length scale as a function of time for the same parameters. Bottom right: Distribution of the distance to the nearest defect, in units of the average length scale, for $(\alpha, \beta) = (0, 1)$ at the initial time (black), the middle time (dark gray), and the final time of the simulation (light gray).

Acknowledgements

The authors thank Bob Kohn for numerous helpful discussions and suggestions, and Nick Schryvers for supplying the HRTEM image. M.E. also thanks the Max Planck Institute for Mathematics in the Natural Sciences for their hospitality while this work was ongoing.

References

[1] C. Begau, J. Hua, and A. Hartmaier. A novel approach to study dislocation density tensors and lattice rotation patterns in atomistic simulations. *Journal of the Mechanics and Physics of Solids*, 60(4):711–722, 2012.

[2] Benjamin Berkels, Andreas Rätz, Martin Rumpf, and Axel Voigt. Extracting grain boundaries and macroscopic deformations from images on atomic scale. *J. Sci. Comput.*, 35(1):1–23, 2008.

[3] Markus Boerdgen, Benjamin Berkels, Martin Rumpf, and Daniel Cremers. Convex relaxation for grain segmentation at atomic scale. In Dieter Fellner, editor, *VMV 2010 - Vision, Modeling & Visualization*, pages 179–186. Eurographics Association, 2010.

[4] L. M. Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Comput. Math. Math. Phys.*, 7(2):200–217, 1967.

[5] L. De Luca, A. Garroni, and M. Ponsiglione. γ-convergence analysis of systems of edge dislocations: the self energy regime. *Arch. Ration. Mech. Anal.*, 2012.

[6] K. R. Elder and Martin Grant. Modeling elastic and plastic deformations in nonequilibrium processing using phase field crystals. *Phys. Rev. E*, 70:051605, 2004.

[7] M. Elsey and B. Wirth. A simple and ecient scheme for phase eld crystal simulation. *ESAIM: Math. Mod. Num. Anal.*, 2012.

[8] M. Elsey and B. Wirth. Segmentation of crystal defects via local analysis of crystal distortion. In *8'th International Conference on Computer Vision Theory and Applications (VISAPP'13)*, 2013.

[9] Tom Goldstein and Stanley Osher. The split Bregman method for L1–regularized problems. *SIAM J. Imag. Sci.*, 2(2):323–343, 2009.

[10] Ron Lifshitz and Dean M. Petrich. Theoretical model for Faraday waves with multiple-frequency forcing. *Phys. Rev. Lett.*, 79(7):1261–1264, 1997.

[11] Stanley Osher, Martin Burger, Donald Goldfarb, Jinjun Xu, and Wotao Yin. An iterative regularization method for total variation-based image restoration. *Multiscale Model. Simul.*, 4(2):460–489, 2005.

[12] H. M. Singer and I. Singer. Analysis and visualization of multiply oriented lattice structures by a two-dimensional continuous wavelet transform. *Phys. Rev. E*, 74:031103, Sep 2006.

[13] E. Strekalovskiy and D. Cremers. Total variation for cyclic structures: Convex relaxation and efficient minimization. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1905–1911, June 2011.

[14] Alexander Stukowski and Karsten Albe. Dislocation detection algorithm for atomistic simulations. *Modelling and Simulation in Materials Science and Engineering*, 18(2):025016, 2010.

[15] Alexander Stukowski and Karsten Albe. Extracting dislocations and non-dislocation crystal defects from atomistic simulation data. *Modelling and Simulation in Materials Science and Engineering*, 18(8):085001, 2010.

[16] B. Wang, H. Idrissi, M. Galceran, M.S. Colla, S. Turner, S. Hui, J.P. Raskin, T. Pardoen, S. Godet, and D. Schryvers. Advanced TEM investigation of the plasticity mechanisms in nancrystalline freestanding palladium films with nanoscale twins. *Int. J. Plast.*, 37:140–156, 2012.

[17] Kuo-An Wu, Ari Adland, and Alain Karma. Phase-field-crystal model for fcc ordering. *Phys. Rev. E*, 81:061601, 2010.

[18] Kuo-An Wu and Peter W. Voorhees. Phase field crystal simulations of nanocrystalline grain growth in two dimensions. *Acta Materialia*, 60(1):407–419, 2012.