Template-Free Wavelet-Based Detection of Local Symmetries

Zsuzsanna Püspöki, *Student Member, IEEE*, and Michael Unser, *Fellow, IEEE*

***Abstract*— Our goal is to detect and group different kinds of local symmetries in images in a scale- and rotation-invariant way. We propose an efficient wavelet-based method to determine the order of local symmetry at each location. Our algorithm relies on circular harmonic wavelets which are used to generate steerable wavelet channels corresponding to different symmetry orders. To give a measure of local symmetry, we use the F-test to examine the distribution of the energy across different channels. We provide experimental results on synthetic images, biological micrographs, and electron-microscopy images to demonstrate the performance of the algorithm.**

***Index Terms*— Local symmetries, F-test, circular harmonic wavelets, junction detection.**

I. Introduction

THE ability to detect edges or ridges (that we consider to be twofold symmetries) and local symmetry centers (or symmetric junctions) can be very useful for the quantitative analysis of microscopic images. For example, certain experiments in stem-cell research rely on the accurate detection of cell shape and extracellular structures (like tight junctions) that exhibit polygonal shapes [1]. Also, in polycrystalline materials such as the hexagonal graphene, it is fundamental to detect line defects since they strongly affect the physical and chemical properties of grain boundaries [2], [3]. The application area of methods that permit this kind of analysis is continuously growing [4]–[8] as image analysis is becoming more relevant in data processing.

There exist several approaches dedicated to the detection of junctions. We can separate these approaches into three categories: methods based on detection, grouping, and analysis of edges and gradients [9], [10]; those relying on feature analysis with structure tensors and their derivatives [11], [12]; and approaches based on template matching [13], [14], including some that use steerable filters [15]–[19]. Methods belonging to the first two categories are often used for identifying points of interest in computer-vision applications. They typically distinguish junctions from edges and other varieties of keypoints, sometimes classifying junctions into categories that are of interest to object recognition. For instance, Kovesi [20] is detecting symmetries and anti-symmetries based on local phase. However, he is mostly focusing on bilateral symmetries and ignores rotational ones. Xia *et al.* [5] are using a template-free method for the detection and grouping of junctions. This algorithm relies on the strength of directional derivatives.

Detectors based on templates are mostly designed for specific types of junctions. The rotational alignment of the template with the pattern of interest is usually dealt by considering a series of rotated filters. There is a tradeoff between accuracy and computational cost: to obtain accurate results, a fine angular discretization is required but is computationally very expensive.

A more favorable approach is the application of steerable filters, where one may explore arbitrary rotations with only a small computational overhead. Indeed, it is enough to perform the filtering operations once only, the response for any specific angle being retrieved by the linear combination of few precomputed components. The basics of steerability were formulated by Freeman and Adelson in the early nineties [21] and developed further by Perona [22], Simoncelli and Farid [16], and Unser and Chenouard [23]. To improve the orientation detection of the basic steerable filters, one may impose an optimality criterion on the properties of the detector. Jacob and Unser [17] imposes Canny-like criteria of optimality [24], such as a reasonable approximation of the ideal detector, maximum signal-to-noise ratio, good spatial localization, and reduced oscillations.

In this paper, we propose an algorithm to detect local symmetries, in a template-free fashion: there is need neither to design a specific junction or ridge-shaped template nor to align it with the pattern of interest. Instead, we apply a wavelet transform and determine the desired symmetry orders based on a statistical test on the wavelet coefficients. The algorithm is based on the circular harmonic wavelets [23], [25], which distribute the energy of the signal among a set of angular harmonics. Based on this angular distribution, we propose a measure of symmetry and a hypothesis test for local symmetry at each pixel. Using the noted measure, we also formulate an approximate maximum-likelihood classifier for the orders of local symmetry.

The approach presented here differs from existing methods in several ways. First of all, our framework has a multiresolution detection scheme. Second, since our method is template-free, there is no need to make heavy computations

Manuscript received December 14, 2014; revised April 11, 2015; accepted May 15, 2015. Date of publication May 21, 2015; date of current version June 9, 2015. This work was supported in part by the European Research Council through the European Union Seventh Framework Programme (FP7/2007-2013) under Grant 267439 and in part by the Hasler Stiftung Foundation. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Chun-Shien Lu.

The authors are with the Biomedical Imaging Group, École Polytechnique Fédérale de Lausanne, Lausanne 1015, Switzerland (e-mail: zsuzsanna.puspoki@epfl.ch; michael.unser@epfl.ch).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIP.2015.2436343

1057-7149 © 2015 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

for matching the template to image features of different orientations. Finally, as a result of the statistical formulation, we can provide a measure of symmetry at each location in the image and for arbitrary symmetry orders in a fast and robust way.

The paper is organized as follows: In Section 2, we present the circular harmonic wavelets that define the analytical framework used in this paper. To motivate this formulation, which distinguishes between local radial symmetry and lack thereof, we have a closer look at the statistics of images without symmetric structures in Section 3. Once the formulation is given for a single scale we extend it to multiscale. Then, we examine the distribution of the wavelet coefficients and, based on the results, we formulate a hypothesis test in Section 4 for the detection of local symmetries of any given order. Next, in Section 5, we deal with the problem of classifying local symmetries of different orders. Finally, in Section 6, we present experimental results on synthetic images, biological micrographs, and electron-microscopy images. The mathematical background that is essential to our formulation, the proofs of theorems, and a brief description of the software can be found in Appendices A-C.

II. Concept

A. Notation

We denote the Cartesian and polar representations of the same 2D function f by $f(\boldsymbol{x})$ with $\boldsymbol{x} \in \mathbb{R}^2$ and $f_{\text{pol}}(r, \theta)$ with $r \in \mathbb{R}^+$, $\theta \in [0, 2\pi)$ (similarly in the Fourier domain: $\hat{f}(\boldsymbol{\omega})$ and $\hat{f}_{\text{pol}}(\rho, \phi)$). We denote the adjoint of an operator L by L^* and the complex conjugate of a function f by $\overline{f}$.

B. Isotropic Wavelet Frames

The isotropic wavelet at scale i and location (grid point) $\boldsymbol{x}_0 = 2^i \boldsymbol{k}_0$, $\boldsymbol{k}_0 \in \mathbb{Z}^2$, is written (in 2D) as

$$\psi_{i,\boldsymbol{k}_0}(\boldsymbol{x}) = \psi_i(\boldsymbol{x} - \boldsymbol{x}_0) = \frac{1}{2^i}\psi\left(\frac{\boldsymbol{x} - \boldsymbol{x}_0}{2^i}\right) = \frac{1}{2^i}\psi\left(\frac{\boldsymbol{x}}{2^i} - \boldsymbol{k}_0\right). \tag{1}$$

With isotropic wavelets, we have that $\hat{\psi}(\omega) = h(|\omega|)$ for some function h. Thus, in the Fourier domain,

$$\begin{aligned}\widehat{\psi_i(\cdot - \boldsymbol{x}_0)}(\boldsymbol{\omega}) &= 2^i \hat{\psi}(2^i \boldsymbol{\omega}) \mathrm{e}^{-\mathrm{j}\langle \boldsymbol{x}_0, \boldsymbol{\omega}\rangle} \\ &= 2^i \hat{h}(2^i \rho) \mathrm{e}^{-\mathrm{j}\rho_0 \rho \cos(\phi - \phi_0)},\end{aligned} \tag{2}$$

with $\boldsymbol{x}_0 = \rho_0 \mathrm{e}^{\mathrm{j}\phi_0}$ and $\boldsymbol{\omega} = \rho(\cos\phi, \sin\phi)$ in polar form.

The wavelets $\psi_{i,\boldsymbol{k}_0}$ form a tight wavelet frame if any finite-energy function f can be decomposed as

$$f = \sum_{i,\boldsymbol{k}_0} \langle f, \psi_{i,\boldsymbol{k}_0}\rangle \psi_{i,\boldsymbol{k}_0}. \tag{3}$$

While (3) is reminiscent of a decomposition in an orthonormal basis, the key difference is that the basis functions are redundant. The key property that ensures that the transform is self-reversible (tight frame) is

$$\sum_i |\hat{h}(2^i \rho)|^2 = 1. \tag{4}$$

In this work we use the Simoncelli's isotropic wavelet [26] defined by its mother wavelet

$$\hat{h}(\rho) = \begin{cases} \cos\left(\frac{\pi}{2}\log_2\left(\frac{2\rho}{\pi}\right)\right), & \frac{\pi}{4} < \rho \le \pi \\ 0, & \text{otherwise.} \end{cases} \tag{5}$$

We apply the multiorder complex Riesz transform on $\psi_{i,\boldsymbol{k}_0}$ to obtain circular harmonic wavelets.

C. Circular Harmonic Wavelet Frames

We now briefly review the circular harmonic wavelets, since they form the backbone of our algorithm, and refer to [23] for more details. We rely on the complex Riesz transform that can be defined in the Fourier domain as

$$\mathscr{R} f(\boldsymbol{x}) \leftrightarrow \mathrm{e}^{\mathrm{j}\phi} \hat{f}_{\text{pol}}(\rho, \phi). \tag{6}$$

The nth-order complex 2D Riesz transform $\mathscr{R}^n$ represents the n-fold iterate of $\mathscr{R}$,

$$\mathscr{R}^n f(\boldsymbol{x}) \leftrightarrow \mathrm{e}^{\mathrm{j}n\phi} \hat{f}_{\text{pol}}(\rho, \phi), \tag{7}$$

whose transfer function is given by

$$\hat{R}^n_{\text{pol}}(\rho, \phi) = \mathrm{e}^{\mathrm{j}n\phi}. \tag{8}$$

The Riesz transform is translation- and scale-invariant, since

$$\forall \boldsymbol{x}_0 \in \mathbb{R}^d, \quad \mathscr{R}\{f(\cdot - \boldsymbol{x}_0)\}(\boldsymbol{x}) = \mathscr{R}\{f(\cdot)\}(\boldsymbol{x} - \boldsymbol{x}_0) \tag{9}$$

$$\forall a \in \mathbb{R}^+ \setminus \{0\}, \quad \mathscr{R}\left\{f\left(\frac{\cdot}{a}\right)\right\}(\boldsymbol{x}) = \mathscr{R}\{f(\cdot)\}\left(\frac{\boldsymbol{x}}{a}\right). \tag{10}$$

The multiorder complex Riesz transform $\mathscr{R}^n$ inherits the properties of the Riesz transform $\mathscr{R}$ (due to its iterative nature), thus it is also translation- and scale-invariant. We refer to [27] for further details on the properties of the higher-order Riesz transform.

One generates circular harmonic wavelet frames by applying the multiorder complex Riesz transform to a primal isotropic function $\psi_{\text{pol}} = \mathscr{F}^{-1}_{\text{2D}}\{\hat{h}\}$ whose shifts and dilations form a wavelet frame.

Using the nth-order Riesz transform, we define the new function $\psi^{(n)} = \mathscr{R}^n \psi$, whose translates and dilates are given by

$$\begin{aligned}\psi^{(n)}_i(\boldsymbol{x} - \boldsymbol{x}_0) &= 2^{-i}\psi^{(n)}\left(\frac{\boldsymbol{x} - \boldsymbol{x}_0}{2^i}\right) \\ &= 2^{-i}\mathscr{R}^n\{\psi\}\left(\frac{\boldsymbol{x} - \boldsymbol{x}_0}{2^i}\right) \\ &= \mathscr{R}^n\left\{2^{-i}\psi\left(\frac{\cdot - \boldsymbol{x}_0}{2^i}\right)\right\}(\boldsymbol{x}) \\ &= \mathscr{R}^n\{\psi_i(\cdot - \boldsymbol{x}_0)\}(\boldsymbol{x}),\end{aligned} \tag{11}$$

where we use the translation- and scale-invariance of $\mathscr{R}^n$. In Fourier, then, we have that

$$\mathscr{F}\left\{\mathscr{R}^n\{\psi_i(\cdot - \boldsymbol{x}_0)\}\right\}(\rho, \phi) = 2^i \hat{h}(2^i\rho)\mathrm{e}^{\mathrm{j}n\phi - \mathrm{j}\rho_0\rho\cos(\phi - \phi_0)}. \tag{12}$$

One can show that the application of the Riesz transform on $\psi_i(\cdot - \boldsymbol{x}_0)$ preserves the tight-frame property [28]. Thus, by choosing N distinct values for the integer n

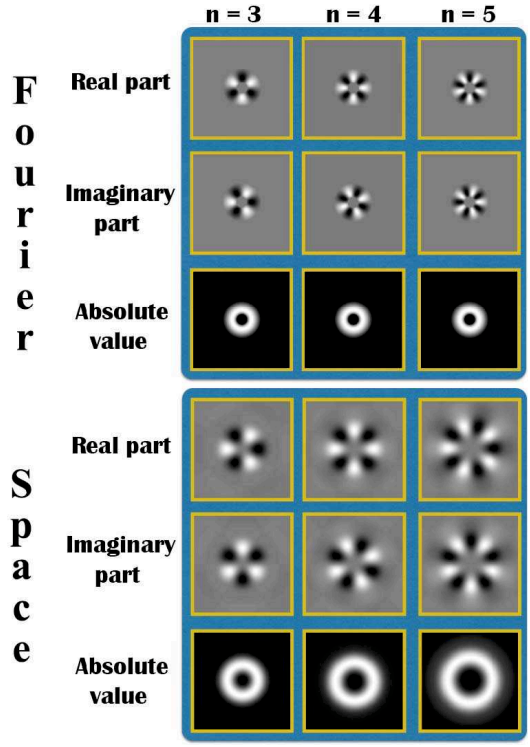

Fig. 1. Circular harmonic wavelets at scale 1. From left to right: nth-order harmonic wavelet, for $n = 3$, $n = 4$, $n = 5$, respectively. From top to bottom: Real part, imaginary part, absolute value, visualized in the Fourier and space domain, respectively.

(distinct set of harmonics), we form a tight frame of steerable wavelets that is similar to the circular harmonic wavelets of Jacovitti and Neri [25], with the difference that the latter ones are non-tight. We note that an nth-order harmonic wavelet has a rotational symmetry of order n around its center, corresponding to the nth-order rotational symmetry of $\mathrm{e}^{\mathrm{j}n\phi}$. We illustrate in Figure 1 the circular harmonic wavelets. As can be observed, in the Fourier domain, the support of the wavelets does not depend on n. By contrast, this is no longer so in the space domain. Specifically, the inverse 2D Fourier transform of $\psi^{(n)}$ is given by

$$\mathscr{F}^{-1}\left\{h\mathrm{e}^{\mathrm{j}n\phi}\right\}(r,\theta) = \frac{\mathrm{j}^n \mathrm{e}^{\mathrm{j}n\theta}}{2\pi}\int_0^\infty \rho h(\rho)\mathrm{J}_n(r\rho)\mathrm{d}\rho \qquad (13)$$

shows that the support is determined by the nth order Bessel function of the first kind (denoted by J_n). Hence, it grows with n.

D. Local Wavelet Energies

Let S be a set of harmonics, let S_M be the subset of S consisting of harmonics that are integer multiples of M with $S_M = \{n \in S : n = Mk, k \in \mathbb{Z}\} \subset S$, and let $q_{n,i} = \langle \psi_i^n(\cdot - \boldsymbol{x}_0), f\rangle$ be the nth-channel coefficient at the scale and location of interest. We propose to base our analysis on the quantity

$$E_{n,i} = \frac{|q_{n,i}|^2}{\sum_{m\in S}|q_{m,i}|^2} \qquad (14)$$

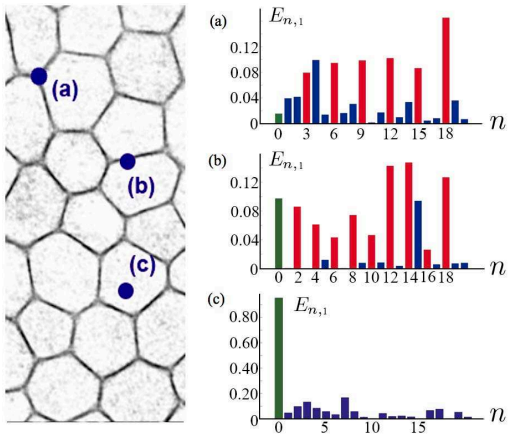

Fig. 2. Illustration of local symmetries in an image at 3 different locations: (a), (b) and (c). The original image is on the left. Subfigure (a) is associated to the normalized wavelet energy vector $E_{n,1}$ (14) at location (a) and scale 1. Location (a) corresponds to a 3-fold symmetric junction, hence every third component has a high amplitude. Subfigure (b) is associated to the vector at location (b) and scale 1. Location (b) corresponds to a 2-fold junction (ridge), hence every second component has a high amplitude. Subfigure (c) is associated to the vector at location (c) and scale 1. Location (c) does not correspond to any particular symmetry order, and we cannot observe any kind of periodic repetition in the energy distribution.

for each $n \in S$ at a given scale and location, which we refer to as normalized wavelet energies. We present in Figure 2 the distribution of normalized-energy wavelet coefficients at a particular scale ($i = 1$) at 3 different locations: three-fold junction (a), ridge (b), non-symmetric point (c). We observe that the "energy" of the wavelet is highly concentrated in M-fold periodic harmonics in case of M-fold symmetry (ridge or symmetric junction) at the location of interest, while no such concentration can be seen in the absence of local symmetries.

III. Wavelet-Domain Statistics

To process the data quantitatively, we propose a hypothesis test to identify energy concentrations in the wavelet domain. We define the null hypothesis as the situation where the distribution of $q_{n,i}$ is an i.i.d. (independent identically distributed) Gaussian, which is a situation where there is no symmetry. In the absence of M-fold symmetries, the average energy for harmonics in S_M must be the same or lower than the average over its complement $S\backslash S_M$. We test this hypothesis against the alternative hypothesis that the image has a local M-fold symmetry (higher concentration of energy in M-fold symmetric harmonics).

We shall now put this approach on a firm statistical basis by deriving the statistical distribution of the wavelet coefficients $q_{n,i}$ under the null hypothesis that the underlying image is the realization of an isotropic random field. Specifically, we shall demonstrate that the circular harmonic wavelet coefficients of such images are uncorrelated on a scale-by-scale basis, which constitutes one of the theoretical contributions of this paper.

We use elements of the theory of generalized stochastic fields to model statistically isotropic images. The discussed concepts are nonstandard because the underlying processes (including fractional Brownian fields) are not stationary, hence

not amenable to a classical treatment. An overview of the definitions and mathematical preliminaries of stochastic fields based on [29] can be found in Appendix A, while the proofs of our results concerning statistically isotropic images can be found in Appendix B.

We shall consider two types of fields: the stationary ones, and the self-similar ones, which are also amenable to a wavelet analysis.

A. Stationary Isotropic Fields

Definition 1 (Wide-Sense Stationary Field of Second Order, w.s.s.): A continuous-space random field s is called wide-sense stationary if

1) *its mean $\mathbb{E}[s(\boldsymbol{x})] = \mu$ is a constant and*
2) *its autocorrelation $R_s(\boldsymbol{x}_1, \boldsymbol{x}_2) = \mathbb{E}[s(\boldsymbol{x}_1)\overline{s}(\boldsymbol{x}_2)]$ depends only on $\boldsymbol{u} = \boldsymbol{x}_1 - \boldsymbol{x}_2$, so that $\mathbb{E}[s(\boldsymbol{x} - \boldsymbol{u})\overline{s}(\boldsymbol{x})] = R_s(\boldsymbol{u})$.*

In the following, we consider w.s.s. fields with zero mean.

Definition 2 (Power Spectrum): The power spectrum of a stationary field is the Fourier transform of its autocorrelation function

$$\Phi_s(\boldsymbol{\omega}) = \mathscr{F}\{R_s\}(\boldsymbol{\omega}). \tag{15}$$

We note that the definition of the power spectrum does only make sense when the stochastic field is stationary with a well-defined autocorrelation. The power spectrum of a statistical field is always real and nonnegative.

Our first result is as follows.

Theorem 1: Let s be a realization of a second-order w.s.s. field with isotropic power spectrum $\Phi_s(\boldsymbol{\omega}) = \Phi_s(\rho)$. Then, the wavelet coefficients $q_{n,i} = \{\langle s, \psi_i^{(n)}(\cdot - \boldsymbol{x}_0)\rangle\}_n$ at a given scale i and position $\boldsymbol{x}_0$ are uncorrelated with zero mean and constant variance

$$\frac{1}{(2\pi)^2}\int_0^\infty \Phi_s\left(\frac{\rho}{2^i}\right)|\hat{h}(\rho)|^2\rho \mathrm{d}\rho. \tag{16}$$

Moreover, when s is Gaussian, the wavelet coefficients q_i are i.i.d. Gaussian within each scale.

B. Self-Similar Isotropic Fields

Definition 3 (Wide-Sense Self-Similarity [29]): A stochastic process s with zero mean is wide-sense self-similar with scaling order H if its autocorrelation function satisfies

$$a^{2H}R_s\left(\frac{\boldsymbol{r}_1}{a}, \frac{\boldsymbol{r}_2}{a}\right) = R_s(\boldsymbol{r}_1, \boldsymbol{r}_2). \tag{17}$$

There are no nontrivial fields that are both wide-sense stationary and wide-sense self-similar. Indeed, if s is wide-sense stationary and self-similar, $R_s(\boldsymbol{r}, \boldsymbol{r}) = R_s(\boldsymbol{0})$, and then $\forall a$, $a^{2H}R_s(\boldsymbol{0}) = R_s(\boldsymbol{0})$. As a result, the variance $\mathrm{Var}(s(\boldsymbol{r})) = R_s(\boldsymbol{r}, \boldsymbol{r})$ either vanishes or is infinite.

However, it is possible to define statistically isotropic self-similar fields with stationary increments. These become stationary after the application of finite differences or analysis functions with vanishing moments. These fields are obtained as solutions of

$$(-\Delta)^{\frac{\gamma}{2}}s = w, \tag{18}$$

where w is a continuous-domain white noise and $(-\Delta)^{\frac{\gamma}{2}}$ is the fractional-Laplacian operator associated with the isotropic Fourier-domain multiplier $\|\boldsymbol{\omega}\|^\gamma$ [29]. In case w is Gaussian, the field we obtain is the 2D generalization of a fractional Brownian motion.

Our formalism allows us to extend our theory to isotropic self-similar fields.

Theorem 2: Let s be a realization of an isotropic self-similar field of order γ and variance σ_0^2. Then, the wavelet coefficients $q_{n,i} = \{\langle s, \psi_i^{(n)}(\cdot - \boldsymbol{x}_0)\rangle\}_n$ at a given scale i and position $\boldsymbol{x}_0$ are uncorrelated with zero mean and constant variance

$$\frac{\sigma_0^2}{(2\pi)^2}4^{i\gamma}\int_0^\infty \frac{1}{\rho^{2\gamma}}|\hat{h}(\rho)|^2\rho \mathrm{d}\rho. \tag{19}$$

Moreover, when s is Gaussian, the wavelet coefficients q_i are i.i.d. Gaussian within each scale.

Theorem 1 and Theorem 2 establish the validity of our null hypothesis for a broad class of isotropic processes. Proofs of Theorems 1 and 2 can be found in Appendix B.

The fact that the circular harmonic wavelets are able to decouple such a process is remarkable, because the pixels are highly correlated in space. We also note that this property does not hold for other conventional wavelets.

IV. F-Test for M-Fold Symmetric Junctions

A. Single-Scale Detection

Corresponding to our null hypothesis, according to which the distribution of the wavelet coefficients $q_{n,i}$ are i.i.d. Gaussian, we define the quantity

$$F(\boldsymbol{q}; M) = \frac{\frac{1}{|S_M|}\sum_{n\in S_M}|q_n|^2}{\frac{1}{|S|-|S_M|}\sum_{n\in S\setminus S_M}|q_n|^2} \tag{20}$$

at each point $\boldsymbol{q}$ which we refer to as the test statistic F or F-score. Under the null hypothesis that the wavelet coefficients q_n are i.i.d. Gaussian, the F statistic consists of the ratio of two independent χ^2 distributions (sums of squared i.i.d. Gaussians), with $|S_M|$ and $(|S| - |S_M|)$ degrees of freedom, respectively. It means that F in (20) has an F distribution. This allows us to use the F-test to detect significant local M-fold symmetry centers.

The hypothesis test proceeds as follows:

1) The observed value F_{obs} of the statistic is computed from the wavelet coefficients at the point of interest.
2) The probability that F exceeds F_{obs} is computed under the null hypothesis (one-sided p-value). This is the probability that, under the null hypothesis, we would observe an F score at least as extreme as F_{obs} in the direction of increasing symmetry.
3) The p-value is compared against the significance level α. To reject the null hypothesis at significance level α, we require that

$$pr.(F \geq F_{\mathrm{obs}}) < \alpha, \tag{21}$$

where $pr.$ is the F probability distribution with $|S_M|$ and $(|S| - |S_M|)$ degrees of freedom. Equivalently, we may compute F_α, the value of the statistic such that

$$pr.(F \geq F_\alpha) = \alpha. \tag{22}$$

Then, if $F_{\text{obs}} > F_\alpha$, we conclude that (21) is true, and reject the null hypothesis at level α.

Since the above is a standard right-tailed F test, the value of F_α can be found in tables for different values of α and different degrees of freedom, or computed from the distribution.

B. Multiscale Detection

To obtain multiscale detection, we modify the hypotheses in such a way that we can make detections in one or several scales (i indicates the scale, I is the total number of scales used for detection). In this case, the multiscale statistics are defined as

$$F_i(\boldsymbol{q}; M) = \frac{\frac{1}{|S_M|}\sum_{n\in S_M}|q_{n,i}|^2}{\frac{1}{|S|-|S_M|}\sum_{n\in S\setminus S_M}|q_{n,i}|^2} \tag{23}$$

at each point $\boldsymbol{q}$. Under the null hypothesis that the wavelet coefficients $q_{n,i}$ are i.i.d. Gaussian, all F_i have the same distribution as F in Section IV. A.

Here, the p-value is defined as the probability (under the null hypothesis) that we would observe a symmetry score at least as extreme as the highest score $F_{i,\text{obs}}$ observed at some scale; in other words, p is computed as

$$p = \text{Prob}\bigcup_l \left\{F_l \geq \max_i F_{i,\text{obs}}\right\}, \tag{24}$$

where l runs over all I scales used for detection. An upper bound on p is given by

$$\sum_l \text{Prob}\left\{F_l \geq \max_i F_{i,\text{obs}}\right\} = I \times pr.(F \geq \max_i F_{i,\text{obs}}), \tag{25}$$

where F has the same distribution as all F_i. Thus, to reject the null hypothesis at level α, it is sufficient that

$$pr.(F \geq \max_i F_{i,\text{obs}}) < \frac{\alpha}{I}, \tag{26}$$

or, equivalently, that $\max_i F_{i,\text{obs}} > F_{\alpha/I}$.

V. Classification of Symmetry Orders

The F-score defined earlier can also be used to classify keypoints as one of several symmetry orders. Since the F-score indicates the strength of a local M-fold symmetry, we can compare the scores for different symmetry orders in order to assign a symmetry class to a given point. Our approach is inspired by a maximum a posteriori (MAP) formulation, as discussed below.

Imagine that we had at our disposal the probability $pr.(M = j|\boldsymbol{q})$ of having a j-fold symmetry at the point of interest, given the wavelet coefficients $\boldsymbol{q}$. The MAP estimator of the local order of symmetry at each point would then be

$$\hat{M} = \arg\max_j p(M = j|\boldsymbol{q}), \tag{27}$$

where j spans a preselected set of possible orders of symmetry. In reality, we do not have access to the probability $pr.(M = j|\boldsymbol{q})$ since we lack a realistic probabilistic model for general images. However, a reasonable substitute can be constructed.

As in (20), let us denote by $F(\boldsymbol{q}; M)$ the observed value of the F-score for M-fold symmetries. For this quantity, we shall also use the notation $F_{\text{obs}}^{(M)}$ and denote the correponding random variable by $F^{(M)}$. Recall that, by definition, a higher value of $F(\boldsymbol{q}; M)$ indicates a higher concentration of energy in M-fold symmetric wavelets. It is therefore natural to assume that, for fixed $M = j$, the probability $pr.(M = j|\boldsymbol{q})$ increases monotonically with $F(\boldsymbol{q}; j)$; that is,

$$pr.(M = j|\boldsymbol{q}) = g_j\left(F(\boldsymbol{q}; j)\right), \tag{28}$$

where g_j is some monotonically increasing function.

We can thus construct different classification rules inspired by the MAP approach, by choosing different families of monotonically increasing functions g_j. Different choices can be validated and sorted by their empirical performance, with preference given to simpler formulas. In particular, we consider the following two important choices:

1) The identity function $g_j(u) \equiv u$ for all j. The corresponding classifier is defined by
$$\hat{M} = \arg\max_j F(\boldsymbol{q}; j). \tag{29}$$
The F-scores for all symmetry orders of interest are computed and the local symmetry is assigned as the one with the highest F score.
2) The monotonically increasing function $g_j(u) \equiv -\log p_j(F^{(j)} \geq u)$, where p_j is the probability associated with the j-fold score $F^{(j)}$ under the null hypothesis of the earlier sections. The motivation for this choice is that, in this case, the value
$$g_j\left(F(\boldsymbol{q}; j)\right) = p_j\left(F^{(j)} > F_{\text{obs}}^{(j)}\right) = p_j \tag{30}$$
is the p-value associated with rejecting the null hypothesis of the earlier sections in favor of j-fold symmetry. Consequently, this choice leads to the classification rule
$$\hat{M} = \arg\max_j(-\log p_j) = \arg\min_j p_j, \tag{31}$$
meaning that p-values for all relevant orders of symmetry are computed, and the lowest p-value among them is used to assign a symmetry order to the keypoint.

In practice, both choices lead to robust and highly effective classifiers, as seen in the experimental results of Section VI.

VI. Experimental Results

Our algorithm to detect symmetries has been programmed as a plug-in for ImageJ [30]. Refer to Appendix C for further details related to the plug-in.

In our implementation, we use Simoncelli's isotropic mother wavelets with dyadic scale progressions. We use a predefined set of harmonics for the tests, which is large enough to provide accurate results for the symmetry orders one might test with the software (this case 0:1:24). We note that the number of harmonics can be arbitrarily extended.

To implement the F-test and compute the α-values, we take advantage of the Beta Class implementation in the jsc.distributions.Beta library.[1] To evaluate the performance of

[1] http://www.jsc.nildram.co.uk/api/jsc/distributions/Beta.html

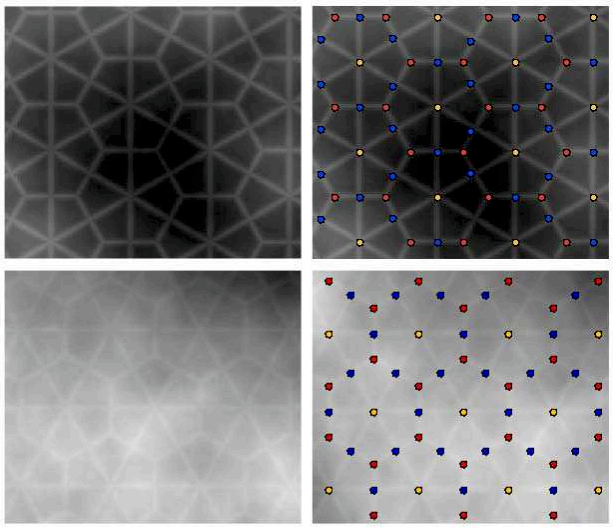

Fig. 3. From left to right and from top to bottom: test image corrupted with isotropic Brownian motion (mean 0, standard deviation 10); corresponding detections; image corrupted with isotropic Brownian motion (mean 0, standard deviation 50); corresponding detections.

Fig. 4. Jaccard index under isotropic Brownian motion, as a function of its standard deviation σ.

Fig. 5. Detection error in the sense of root-mean-square error (RMSE) under isotropic Brownian motion, as a function of its standard deviation σ.

the algorithm we use a variety of test images, synthetic and actual microscopic images (real micrographs). We note that, in our formulation, ridges can be considered as 2-fold symmetries or junctions.

A. Robustness Against Background Signal

It has been observed that the power spectrum of many natural images is isotropic with an $\frac{1}{\|\omega\|^s}$ type of decay that is consistent with a fractional Brownian motion model [31], [32]. Experimentally, this model of background signal fits fluorescence microscopy images well [33]. To make the detection task more realistic, we have used this model to generate a background signal that is added to the patterns of interest to simulates the autofluorescence that is present in real micrographs.

First, we generate a series of (1024 $\times$ 1024) test images populated by 3-, 4-, and 6-fold junctions at fixed locations. The ground-truth data contains 610 3-fold, 921 4-fold, and 304 6-fold symmetric junctions, respectively. We corrupt the images with isotropic Brownian motion, with a mean of zero and standard deviation of 0 to 100. An illustration of the test images can be seen in Figure 3. To separate distinct effects, in this experiment we use only the first wavelet scale to make detections.

In the evaluation phase we use the Hungarian algorithm to match the detections with the nodes of the original grid. The detections are accepted if they are no further than 5 pixel away from the original nodes. Otherwise, they are counted as false-positive results.

To make a quantitative evaluation we compute the Jaccard index, and the position root mean square error (RMSE). The RMSE error is computed for the matched detections. An illustration of the results (crop of one of the test images) can be seen in Figure 3, the Jaccard index in Figure 4, and the RMS errors in Figure 5. Based on the graphs, we claim that our algorithm performs well under a wide range of background noise.

Fig. 6. Jaccard index under isotropic Brownian motion of zero mean and standard deviation 10 for rotations γ from 0 to 360 degrees.

B. Rotation Invariance

We generate another series of (1024 $\times$ 1024) test images. Similarly to the previous case, we build a synthetic grid and rotate it around its center with 1 degree steps from 0 to 360 degrees. We apply isotropic Brownian motion (of zero mean and standard deviation 10, as illustrated on Figure 3) on these images. The ground-truth data contains approximately 600 3-fold, 900 4-fold, and 300 6-fold symmetric junctions. We note that the number of reference nodes varies as some of the nodes will be rotated out of the bounds.

We compute the Jaccard index (Figure 6), and the RMS error of the detections (Figure 7). Based on the graphs, we confirm that our method is essentially rotation invariant and performs well independently of the orientation of the junctions. The small fluctuations in the values are due to numerical errors.

Fig. 7. Detection error in the sense of root-mean-square error (RMSE) under isotropic Brownian motion of zero mean and standard deviation 10 for rotations γ from 0 to 360 degrees.

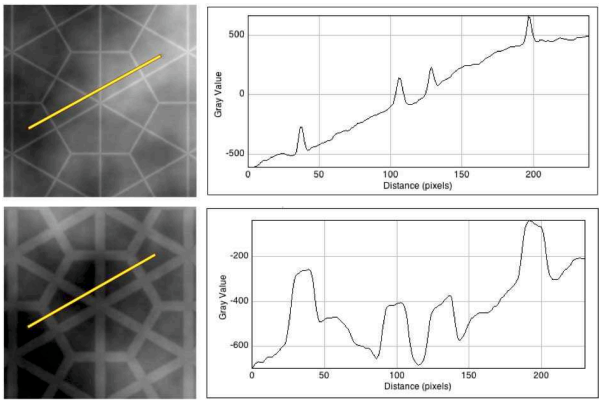

Fig. 8. Test images of different line thicknesses, corrupted with isotropic Brownian motion (mean 0, standard deviation 10). From left to right: test image; profile plot along the corresponding bright line. From top to bottom: line thickness $l = 1.0$, line thickness $l = 10.0$, respectively.

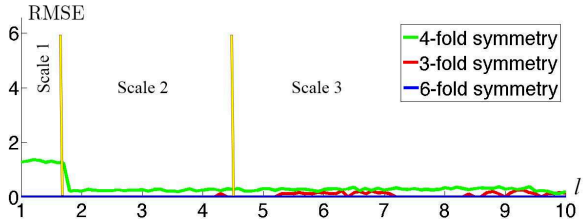

Fig. 9. Position error in the sense of root-mean-square error (RMSE) under isotropic Brownian motion of zero mean and standard deviation 10 for line thicknesses 1 to 15 with steps 0.1. Scale 1, 2, and 3 indicates the scale at which the detections are made.

C. Scale Invariance

Similarly to the previous experiments, we generate a series of (1024×1024) test images. We again build a synthetic grid and change the profile of the ridges from thickness level 1 to 15 with steps 0.1. An illustration of the test images for different line thicknesses corrupted with Brownian motion is shown in Figure 8. Here, the difficulty is twofold. First, the background intensity is continuously changing in a wide range. Second, the edges to detect do not always appear as clear peaks, but they can form a plateau of varying width. One can find in Figure 9 the average position error of the detections under isotropic Brownian motion, as a function of the thickness level. Thin lines are detected at the first wavelet scale; as we increase the width of the lines we obtain the detections at higher scales. The detection is accurate, independently of the scale. The algorithm performs well for different symmetry orders.

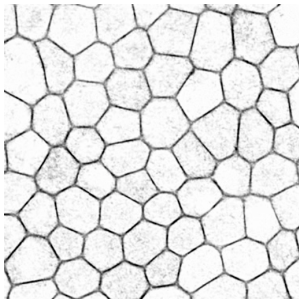

Fig. 10. Hexagonal embryonic stem cells in light microscopy [1]. Detection of two-fold symmetries or ridges and three-fold junctions. Subfigures: original image, detected ridges, detected three-fold junctions, combined image.

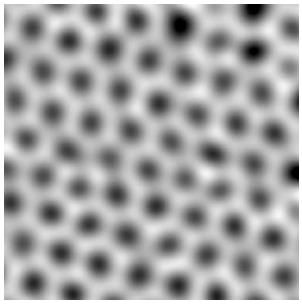

Fig. 11. Atomic-resolution ADF-STEM image of hexagonal graphene crystals [2]. Detection of two-fold symmetries or ridges (dark grey) and three-fold junctions (light grey). Subfigures: original image, combined RGB image.

D. Micrographs

Figure 10 features hexagonal embryonic stem cells imaged by fluorescence microscopy with 3-fold symmetric junctions. As a result of the imaging technique, the test image contains Poisson noise. We were interested in detecting the outlines of the cells, containing ridges and junction points. Results shown for their detection are indeed satisfying and consistent with the regularity of the structure.

Figure 11 contains an atomic-resolution ADF-STEM image of hexagonal graphene crystals with 3-fold symmetric junctions. The detection task is more challenging than in the embryonic stem-cell image due to blurring effects. Here the image quality is quite low. We intend to detect the node areas (light grey), and the connecting regions (dark grey). In this case, the method yielded again good results.

VII. Conclusion

We presented a wavelet-based framework for the detection and classification of the centers of local symmetry in images. The attractive features of the algorithm are (i) the multiscale approach and the rotation invariance, which make it possible to obtain a precise estimate of the junctions across scales, at arbitrary orientations; (ii) robustness; and (iii) speed. Also, since any combination of arbitrary orders of symmetry may be tested without the need to design specific templates, detecting arbitrary symmetry orders is possible without any computational overhead.

Appendix

The Appendix contains the proofs of Theorem 1 and 2, which are fundamental to our work. They establish the validity of our null hypothesis (independence of wavelet coefficients) for a broad class of isotropic processes.

A. Mathematical Preliminaries of Stochastic Processes

We now present the nonstandard preliminary concepts from the theory of generalized stochastic processes that are required in Section B. To prove the results, we shall work with correlation forms, which are a generalization of the concept of autocorrelation.

Definition 4 (Correlation Form [34]): The correlation form of a generalized field s is

$$\mathscr{B}_s\{\varphi_1, \varphi_2\} = \mathbb{E}\{\langle s, \varphi_1\rangle\overline{\langle s, \varphi_2\rangle}\} \tag{32}$$

for $\varphi_1, \varphi_2 \in \mathscr{S}$ (Schwartz space), which, in cases of interest to us, can be extended to L_2.

By the kernel theorem,

$$\mathscr{B}_s\{\varphi_1, \varphi_2\} = \int_{\mathbb{R}^2}\int_{\mathbb{R}^2} \varphi_1(\boldsymbol{r}_1)\overline{\varphi_2(\boldsymbol{r}_2)}R_s(\boldsymbol{r}_1, \boldsymbol{r}_2)\mathrm{d}\boldsymbol{r}_1\mathrm{d}\boldsymbol{r}_2, \tag{33}$$

where $R_s(\boldsymbol{r}_1, \boldsymbol{r}_2) \in \mathscr{S}'(\mathbb{R}^2 \times \mathbb{R}^2)$ (space of tempered distributions) is the generalized correlation function of s.

We are interested in the following cases: (i) generalized white noise with finite variance; (ii) w.s.s. fields with zero mean; (iii) fields whitened by $\mathrm{L} = (-\Delta)^{\frac{\gamma}{2}}$, as an example of a self-similar field with stationary increments.

(i) The correlation form of a generalized white noise takes the form

$$\mathscr{B}_w\{\varphi_1, \varphi_2\} = \sigma_0^2\langle\varphi_1, \varphi_2\rangle, \tag{34}$$

where σ_0^2 is called the variance of the noise w. Thus, the autocorrelation of the noise is $\sigma_0^2\delta$ (where δ is the Dirac impulse). Its power spectrum is flat, with

$$\Phi_w \equiv \sigma_0^2. \tag{35}$$

(ii) The correlation form of a zero mean w.s.s. field is related to its power spectrum Φ_s by

$$\mathscr{B}_s\{\varphi_1, \varphi_2\} = \frac{1}{(2\pi)^2}\int_{\mathbb{R}^2} \hat{\varphi}_1(\boldsymbol{\omega})\overline{\hat{\varphi}_2(\boldsymbol{\omega})}\Phi_s(\boldsymbol{\omega})\mathrm{d}\boldsymbol{\omega}. \tag{36}$$

(iii) We are considering a self-similar field s such that $(-\Delta)^{\frac{\gamma}{2}}s = w$ is an innovation field (white noise), or, formally,

$$\langle s, \varphi\rangle = \langle w, I_{\gamma,2}^*\varphi\rangle, \tag{37}$$

where $I_{\gamma,2}^* = (-\Delta)^{\frac{-\gamma}{2}}$ is the corrected scale-invariant inverse operator from $\mathscr{S}$ to L_2 of $(-\Delta)^{\frac{\gamma}{2}}$ with an appropriate correction for the singularity at $\boldsymbol{\omega} = \mathbf{0}$. (The operator here is a special case of a family introduced in [29], hence the notation and the index 2.)
In general, the inverse operator of the fractional Laplacian can be obtained from the Taylor series expansion of φ at the origin by removing a sufficient number of lower order terms:

$$I_{\gamma,2}^*\varphi(\boldsymbol{x}) = \int_{\mathbb{R}^2} \frac{\mathrm{d}\boldsymbol{\omega}}{(2\pi)^2}\mathrm{e}^{\mathrm{j}\boldsymbol{x}^T\boldsymbol{\omega}}\frac{\hat{\varphi}(\boldsymbol{\omega}) - \sum_{|\boldsymbol{k}|\le\lfloor\gamma-1\rfloor}\hat{\varphi}^{(\boldsymbol{k})}(\mathbf{0})\frac{\boldsymbol{\omega}^{\boldsymbol{k}}}{\boldsymbol{k}!}}{\|\boldsymbol{\omega}\|^\gamma}. \tag{38}$$

In the case where $\varphi = \psi$ is a wavelet function whose moments vanish up to degree $\lfloor\gamma - 1\rfloor$, all the terms in the sum of (38) will be zero, and

$$I_{\gamma,2}^*\psi(\boldsymbol{x}) = \int_{\mathbb{R}^2} \frac{\mathrm{d}\boldsymbol{\omega}}{(2\pi)^2}\mathrm{e}^{\mathrm{j}\boldsymbol{x}^T\boldsymbol{\omega}}\frac{\hat{\psi}(\boldsymbol{\omega})}{\|\boldsymbol{\omega}\|^\gamma}, \tag{39}$$

which corresponds to a shift-invariant operator whose Fourier-domain description is

$$\widehat{I_{\gamma,2}^*\psi}(\boldsymbol{\omega}) = \frac{\hat{\psi}(\boldsymbol{\omega})}{\|\boldsymbol{\omega}\|^\gamma}. \tag{40}$$

As a consequence of (37), and using (40) together with Parseval's identity to rewrite (34) in the Fourier-domain, we obtain the correlation form of the isotropic self-similar field as

$$\begin{aligned}\mathscr{B}_s\{\psi_1, \psi_2\} &= \mathscr{B}_w\{I_{\gamma,2}^*\psi_1, I_{\gamma,2}^*\psi_2\}\\ &= \frac{\sigma_0^2}{(2\pi)^2}\int_{\mathbb{R}^2}\frac{\hat{\psi}_1(\boldsymbol{\omega})\hat{\psi}_2(\boldsymbol{\omega})}{\|\boldsymbol{\omega}\|^{2\gamma}}\mathrm{d}\boldsymbol{\omega}\end{aligned} \tag{41}$$

for functions ψ_1 and ψ_2 with sufficiently many vanishing moments.

B. Proofs for the Theorems on Statistically Isotropic Images

Proof of Theorem 1: For a given scale i and position $\boldsymbol{x}_0$, we consider the two harmonics n_1, n_2. We use (36) to write the correlation form of a zero-mean w.s.s. field which, by definition, is the same as the covariance of the corresponding

wavelet coefficients. We write

$$\begin{aligned}
&\mathscr{B}_s\{\psi_i^{(n_1)}(\cdot-\boldsymbol{x}_0),\psi_i^{(n_2)}(\cdot-\boldsymbol{x}_0)\}\\
&\quad=\frac{2^{2i}}{(2\pi)^2}\int\hat{\psi}^{(n_1)}(2^i\boldsymbol{\omega})\overline{\hat{\psi}^{(n_2)}(2^i\boldsymbol{\omega})}\Phi_s(\boldsymbol{\omega})\mathrm{d}\boldsymbol{\omega}\\
&\quad=\frac{2^{2i}}{(2\pi)^2}\int_0^\infty\int_{-\pi}^{\pi}\Phi_s(\rho)|\hat{h}(2^i\rho)|^2\mathrm{e}^{\mathrm{j}n_1\phi}\mathrm{e}^{-\mathrm{j}n_2\phi}\rho\mathrm{d}\rho\mathrm{d}\phi\\
&\quad=\frac{1}{(2\pi)^2}\int_0^\infty\int_{-\pi}^{\pi}\Phi_s\left(\frac{\rho}{2^i}\right)|\hat{h}(\rho)|^2\mathrm{e}^{\mathrm{j}n_1\phi}\mathrm{e}^{-\mathrm{j}n_2\phi}\rho\mathrm{d}\rho\mathrm{d}\phi\\
&\quad=\frac{1}{(2\pi)^2}\int_0^\infty\Phi_s\left(\frac{\rho}{2^i}\right)|\hat{h}(\rho)|^2\rho\mathrm{d}\rho\int_{-\pi}^{\pi}\mathrm{e}^{\mathrm{j}(n_1-n_2)\phi}\mathrm{d}\phi\\
&\quad=C_i\int_{-\pi}^{\pi}\mathrm{e}^{\mathrm{j}(n_1-n_2)\phi}\mathrm{d}\phi, \qquad (42)
\end{aligned}$$

where we took advantage of the properties of the Fourier transform and a change of variables. We used the fact that the radial part is fixed and identical for all components. The last integral is 1 if $n_1 = n_2$ and vanishes otherwise. This means that the wavelet coefficients $\{\langle s, \psi_i^{(n)}(\cdot - \boldsymbol{x}_0)\rangle\}_n$ are uncorrelated with common variance

$$\frac{1}{(2\pi)^2}\int_0^\infty\Phi_s\left(\frac{\rho}{2^i}\right)|\hat{h}(\rho)|^2\rho\mathrm{d}\rho. \qquad (43)$$

In the Gaussian case, uncorrelation is equivalent to independence. ■

Proof of Theorem 2: As in the proof of Theorem 1, the covariance of the wavelet coefficients $\langle s, \psi_i^{(n_1)}(\cdot-\boldsymbol{x}_0)\rangle$ and $\langle s, \psi_i^{(n_2)}(\cdot - \boldsymbol{x}_0)\rangle$ is given by the correlation form $\mathscr{B}_s\{\psi_i^{(n_1)}(\cdot - \boldsymbol{x}_0), \psi_i^{(n_2)}(\cdot - \boldsymbol{x}_0)\}$. Since all of the moments of our wavelets (which are based on the Simoncelli design) vanish, we can make use of (41) and (12) to find the covariance of the wavelet coefficients. Rewritten in polar coordinates, it reads

$$\begin{aligned}
&\mathscr{B}_s\{\psi_i^{(n_1)}(\cdot-\boldsymbol{x}_0),\psi_i^{(n_2)}(\cdot-\boldsymbol{x}_0)\}\\
&\quad=\frac{2^{2i}\sigma_0^2}{(2\pi)^2}\int_0^\infty\int_{-\pi}^{\pi}\frac{|\hat{h}(2^i\rho)|^2}{\rho^{2\gamma}}\mathrm{e}^{\mathrm{j}(n_1-n_2)\phi}\rho\mathrm{d}\rho\mathrm{d}\phi\\
&\quad=\frac{\sigma_0^2}{(2\pi)^2}\frac{2^{2i}}{2^{2i}}\int_0^\infty\int_{-\pi}^{\pi}\frac{4^{i\gamma}}{\rho^{2\gamma}}|\hat{h}(\rho)|^2\mathrm{e}^{\mathrm{j}(n_1-n_2)\phi}\rho\mathrm{d}\rho\mathrm{d}\phi\\
&\quad=\frac{\sigma_0^2}{(2\pi)^2}4^{i\gamma}\int_0^\infty\frac{1}{\rho^{2\gamma}}|\hat{h}(\rho)|^2\rho\mathrm{d}\rho\int_{-\pi}^{\pi}\mathrm{e}^{\mathrm{j}(n_1-n_2)\phi}\mathrm{d}\phi\\
&\quad=C_i\int_{-\pi}^{\pi}\mathrm{e}^{\mathrm{j}(n_1-n_2)\phi}\mathrm{d}\phi.
\end{aligned}$$

The last integral is 1 if $n_1 = n_2$ and vanishes otherwise. This means that the wavelet coefficients $\{\langle s, \psi_i^{(n)}(\cdot - \boldsymbol{x}_0)\rangle\}_n$ are uncorrelated with common variance

$$\frac{\sigma_0^2}{(2\pi)^2}4^{i\gamma}\int_0^\infty\frac{1}{\rho^{2\gamma}}|\hat{h}(\rho)|^2\rho\mathrm{d}\rho. \qquad (44)$$

In the Gaussian case, this implies that the wavelet coefficients are i.i.d. Gaussian. ■

C. Description of the Software

Our algorithm has been programmed as a plug-in for ImageJ, which is a public-domain image-processing platform. It is implemented in Java.

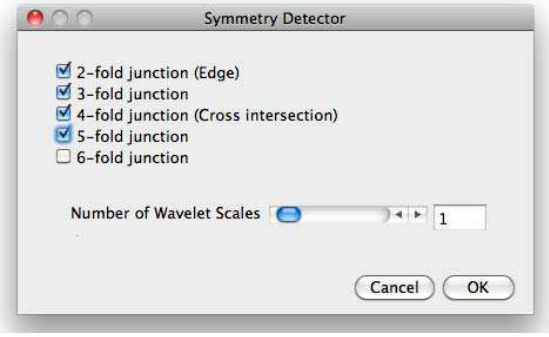

Fig. 12. Graphical user interface of the detector of symmetry.

The options of the program can be controlled by the settings window (Figure 12). This allows the user to determine 1) the different orders of symmetry to test, namely: ridge, three-fold junction, cross intersection, 5-fold junction, and 6-fold junction; and 2) the number of wavelet scales.

The user is informed about the different steps of the computation process by a log window. When the detection and the classification terminate, a colored image pops up. The different colors mark the different types of junctions. The typical running time of the algorithm for two different orders of symmetry for an image of size (512 × 512) pixels is around a few seconds on a 2.4 GHz MAC.

REFERENCES

[1] F. Osakada *et al.*, "Toward the generation of rod and cone photoreceptors from mouse, monkey and human embryonic stem cells," *Nature Biotechnol.*, vol. 26, no. 2, pp. 215–224, Feb. 2008.

[2] P. Y. Huang *et al.*, "Grains and grain boundaries in single-layer graphene atomic patchwork quilts," *Nature*, vol. 469, no. 7330, pp. 389–392, Jan. 2011.

[3] L. Zhao *et al.* (Aug. 2010). "The atomic-scale growth of large-area monolayer graphene on single-crystal copper substrates." [Online]. Available: http://arxiv.org/abs/1008.3542

[4] F. Zhao, P. R. S. Mendonca, R. Bhotika, and J. V. Miller, "Model-based junction detection algorithm with applications to lung nodule detection," in *Proc. 4th IEEE Int. Symp. Biomed. Imag., Nano Macro (ISBI)*, Apr. 2007, pp. 504–507.

[5] G.-S. Xia, J. Delon, and Y. Gousseau, "Accurate junction detection and characterization in natural images," *Int. J. Comput. Vis.*, vol. 106, no. 1, pp. 31–56, Jan. 2014.

[6] M. Maire, P. Arbelaez, C. Fowlkes, and J. Malik, "Using contours to detect and localize junctions in natural images," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2008, pp. 1–8.

[7] T.-A. Pham, M. Delalandre, S. Barrat, and J.-Y. Ramel, "Accurate junction detection and reconstruction in line-drawing images," in *Proc. 21st Int. Conf. Pattern Recognit. (ICPR)*, Nov. 2012, pp. 693–696.

[8] W. Liang and H. Yunan, "Vision-aided navigation for aircrafts based on road junction detection," in *Proc. IEEE Int. Conf. Intell. Comput. Intell. Syst. (ICIS)*, vol. 4. Nov. 2009, pp. 164–169.

[9] D. J. Beymer, "Finding junctions using the image gradient," Artif. Intell. Lab., Massachusetts Inst. Technol., Cambridge, MA, USA, Tech. Rep. AIM-1266, Dec. 1991, pp. 1–46.

[10] T.-A. Pham, M. Delalandre, S. Barrat, and J. Ramel, "A robust approach for local interest point detection in line-drawing images," in *Proc. 10th IAPR Int. Workshop Document Anal. Syst. (DAS)*, Mar. 2012, pp. 79–84.

[11] U. Köthe, "Edge and junction detection with an improved structure tensor," in *Pattern Recognition*, vol. 2781. Berlin, Germany: Springer-Verlag, 2003, pp. 25–32.

[12] J. Bigun, T. Bigun, and K. Nilsson, "Recognition by symmetry derivatives and the generalized structure tensor," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 12, pp. 1590–1605, Dec. 2004.

[13] L. Parida, D. Geiger, and R. Hummel, "Junctions: Detection, classification, and reconstruction," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 20, no. 7, pp. 687–698, Jul. 1998.
[14] F. Escolano, P. Suau, and B. Bonev, *Information Theory in Computer Vision and Pattern Recognition*. London, U.K.: Springer-Verlag, 2009.
[15] W. T. Freeman, "Steerable filters and local analysis of image structure," Ph.D. dissertation, MIT Media Lab., Massachusetts Inst. Technol., Cambridge, MA, USA, 1992.
[16] E. P. Simoncelli and H. Farid, "Steerable wedge filters for local orientation analysis," *IEEE Trans. Image Process.*, vol. 5, no. 9, pp. 1377–1382, Sep. 1996.
[17] M. Jacob and M. Unser, "Design of steerable filters for feature detection using Canny-like criteria," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 8, pp. 1007–1019, Aug. 2004.
[18] G. González, F. Aguet, F. Fleuret, M. Unser, and P. Fua, "Steerable features for statistical 3D dendrite detection," in *Proc. 12th Int. Conf. Med. Image Comput. Comput.-Assisted Intervent. (MICCAI)*, vol. 5762. Sep. 2009, pp. 625–632.
[19] M. Muehlich, D. Friedrich, and T. Aach, "Design and implementation of multisteerable matched filters," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 2, pp. 279–291, Feb. 2012.
[20] P. Kovesi, "Symmetry and asymmetry from local phase," in *Proc. 10th Austral. Joint Conf. Artif. Intell.*, 1997, pp. 2–4.
[21] W. T. Freeman and E. H. Adelson, "The design and use of steerable filters," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 9, pp. 891–906, Sep. 1991.
[22] P. Perona, "Steerable-scalable kernels for edge detection and junction analysis," in *Proc. 2nd Eur. Conf. Comput. Vis. (ECCV)*, Dec. 1992, pp. 663–672.
[23] M. Unser and N. Chenouard, "A unifying parametric framework for 2D steerable wavelet transforms," *SIAM J. Imag. Sci.*, vol. 6, no. 1, pp. 102–135, 2013.
[24] J. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-8, no. 6, pp. 679–698, Nov. 1986.
[25] G. Jacovitti and A. Neri, "Multiresolution circular harmonic decomposition," *IEEE Trans. Signal Process.*, vol. 48, no. 11, pp. 3242–3247, Nov. 2000.
[26] J. Portilla and E. P. Simoncelli, "A parametric texture model based on joint statistics of complex wavelet coefficients," *Int. J. Comput. Vis.*, vol. 40, no. 1, pp. 49–70, Oct. 2000.
[27] M. Unser and D. Van De Ville, "Wavelet steerability and the higher-order Riesz transform," *IEEE Trans. Image Process.*, vol. 19, no. 3, pp. 636–652, Mar. 2010.
[28] M. Unser, N. Chenouard, and D. Van De Ville, "Steerable pyramids and tight wavelet frames in $L_2(\mathbb{R}^d)$," *IEEE Trans. Image Process.*, vol. 20, no. 10, pp. 2705–2721, Oct. 2011.
[29] M. Unser and P. D. Tafti, *An Introduction to Sparse Stochastic Processes*. Cambridge, U.K.: Cambridge Univ. Press, 2014.
[30] M. D. Abràmoff, P. J. Magalhães, and S. J. Ram, "Image processing with ImageJ," *Biophoton. Int.*, vol. 11, no. 7, pp. 36–43, Jul. 2004.
[31] B. Pesquet-Popescu and J. L. Vehel, "Stochastic fractal models for image processing," *IEEE Signal Process. Mag.*, vol. 19, no. 5, pp. 48–62, Sep. 2002.
[32] A. P. Pentland, "Fractal-based description of natural scenes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-6, no. 6, pp. 661–674, Nov. 1984.
[33] D. Sage, F. R. Neumann, F. Hediger, S. M. Gasser, and M. Unser, "Automatic tracking of individual fluorescence particles: Application to the study of chromosome dynamics," *IEEE Trans. Image Process.*, vol. 14, no. 9, pp. 1372–1383, Sep. 2005.
[34] I. M. Gel'fand and N. Y. Vilenkin, *Generalized Functions: Applications of Harmonic Analysis*, vol. 4. New York, NY, USA: Academic, 1964.

Zsuzsanna Püspöki was born in Budapest, Hungary, in 1986. She received the M.Sc./Diploma (Hons.) degree in electrical engineering from the Budapest University of Technology and Economics, in 2011. She is currently pursuing the Ph.D. degree with the Biomedical Imaging Group, École Polytechnique Fédérale de Lausanne, under the direction of M. Unser. She is working on problems related to image analysis, wavelet theory, and computer vision. Her research interests include the analysis of local geometric transformations and symmetries, and directional analysis in images.

Michael Unser (M'89–SM'94–F'99) received the M.S. (*summa cum laude*) and Ph.D. degrees in electrical engineering from the École Polytechnique Fédérale de Lausanne (EPFL), Switzerland, in 1981 and 1984, respectively. From 1985 to 1997, he was a Scientist with the National Institutes of Health, Bethesda, USA. He has a strong interest in sampling theories, multiresolution algorithms, wavelets, and the use of splines for image processing. He has authored 200 journal papers on those topics, and is one of ISI's highly cited authors in engineering. He is currently a Full Professor and the Director of the Biomedical Imaging Group with EPFL. His main research area is biomedical image processing. He was the Associate Editor-in-Chief of the IEEE TRANSACTIONS ON MEDICAL IMAGING (2003–2005), and served as an Associate Editor of the IEEE TRANSACTIONS ON MEDICAL IMAGING (1999–2002 and 2006–2007), the IEEE TRANSACTIONS ON IMAGE PROCESSING (1992–1995), and the IEEE SIGNAL PROCESSING LETTERS (1994–1998). He is currently a member of the Editorial Boards of *Foundations and Trends in Signal Processing*, and *Sampling Theory in Signal and Image Processing*. He co-organized the first IEEE International Symposium on Biomedical Imaging (2002), and was the Founding Chair of the Technical Committee of the IEEE-SP Society on Bio Imaging and Signal Processing. He was a recipient of the 1995 and 2003 best paper awards, the 2000 Magazine Award, and two IEEE Technical Achievement Awards (2008 SPS and 2010 EMBS). He is a fellow of the European Association for Signal Processing and a member of the Swiss Academy of Engineering Sciences.