
Local Geometric Transformations in Image Analysis

Zsuzsanna Püspöki

Thèse N° 7014 (juin 2016)

*Thèse présentée à la faculté des sciences et techniques de l'ingénieur
pour l'obtention du grade de docteur ès sciences
et acceptée sur proposition du jury*

Prof. Pierre Vandergheynst, *président*
Prof. Michaël Unser, *directeur de thèse*
Prof. Dimitri Van De Ville, *rapporteur*
Dr. Jean-Christophe Olivo-Marin, *rapporteur*
Dr. Ullrich Köthe, *rapporteur*

École polytechnique fédérale de Lausanne—2016

Cover design by Annette Unser
Printing and binding by Repro-EPFL
Typeset with L^AT_EX
Copyright © 2016 by Zsuzsanna Püspöki
Available at <http://bigwww.epfl.ch/>

Abstract

The characterization of images by geometric features facilitates the precise analysis of the structures found in biological micrographs such as cells, proteins, or tissues. In this thesis, we study image representations that are adapted to local geometric transformations such as rotation, translation, and scaling, with a special emphasis on wavelet representations.

In the first part of the thesis, our main interest is in the analysis of directional patterns and the estimation of their location and orientation. We explore steerable representations that correspond to the notion of rotation. Contrarily to classical pattern matching techniques, they have no need for an a priori discretization of the angle and for matching the filter to the image at each discretized direction. Instead, it is sufficient to apply the filtering only once. Then, the rotated filter for any arbitrary angle can be determined by a systematic and linear transformation of the initial filter.

We derive the Cramér-Rao bounds for steerable filters. They allow us to select the best harmonics for the design of steerable detectors and to identify their optimal radial profile. We propose several ways to construct optimal representations and to build powerful and effective detector schemes; in particular, junctions of coinciding branches with local orientations.

The basic idea of local transformability and the general principles that we utilize to design steerable wavelets can be applied to other geometric transformations. Accordingly, in the second part, we extend our framework to other transformation groups, with a particular interest in scaling. To construct representations in tune with a notion of local scale, we identify the possible solutions for scalable functions and give specific criteria for their applicability to wavelet schemes. Finally, we propose discrete wavelet frames that approximate a continuous wavelet transform. Based on these results, we present a novel wavelet-based image-analysis software that provides a fast and auto-

matic detection of circular patterns, combined with a precise estimation of their size.

Keywords

Circular harmonic wavelets, Cramér-Rao bounds, detection of junctions, directional patterns, estimation of orientation, Fourier multipliers, F-test, image analysis, image representations, local geometric transformations, local symmetries, scalability, steerability, steerable filters, steerable wavelets, tight wavelet frames, transformability

Résumé

La compréhension des caractéristiques géométriques d'une image facilite l'analyse des structures des données microbiologiques, tels que les cellules, les protéines ou les tissus. Dans cette thèse, nous étudions les représentations d'images, en particulier les représentations en ondelettes, qui sont adaptées aux transformations géométriques locales comme les rotations, les translations et les dilatations.

Dans une première partie, nous portons notre attention sur l'analyse des motifs orientés contenus dans une image et sur l'estimation de leur localisation et orientation. Nous explorons des représentations dites orientables qui correspondent à la notion de rotation. Par opposition aux techniques classiques de filtrage par motif, ces représentations ne nécessitent ni la discrétisation de l'angle a priori, ni l'adéquation du filtre à l'image pour chaque direction discrétisée, et il est au contraire suffisant d'appliquer le filtre une seule fois. Ainsi, le filtre auquel on applique une rotation d'un angle arbitraire peut être déterminé par une transformation systématique et linéaire à partir du filtre initial.

Nous établissons les bornes de Cramér-Rao pour les filtres orientables, ce qui nous permet de sélectionner les meilleures harmoniques pour la conception de détecteurs orientables et d'identifier le profil radial optimal. Nous proposons différents moyens de construire une représentation optimale et de construire des méthodes performantes et opérantes pour la détection de motifs, en particulier pour la détection de jonctions dans les images.

L'idée de base des transformations locales et les principes généraux que nous utilisons pour la conception d'ondelettes orientables peuvent être appliqués à d'autres transformations géométriques. Ainsi, nous étendons ces principes à d'autres groupes de transformation dans la seconde partie de cette thèse, avec un intérêt particulier pour les dilatations. Afin de construire des représentations en accord avec la notion de chan-

gement d'échelle local, nous identifions les solutions possibles de fonctions "dilatables" et donnons des critères spécifiques pour leur application à des schémas d'ondelettes. Enfin, nous proposons des ondelettes discrètes redondantes approximant une transformée en ondelette continue. À partir de ces résultats, nous présentons un nouveau software pour l'analyse d'image permettant une détection automatique et rapide de motifs circulaires, ainsi qu'une estimation précise de leur taille.

Mots clefs

Analyse d'image, bornes de Cramér-Rao, détection de jonctions, dilatations, estimation d'orientation, filtres orientables, F-test, motifs orientés, multiplicateurs de Fourier, ondelettes harmoniques circulaires, ondelettes orientables, ondelettes redondantes, représentation d'images, symétries locales, transformations géométriques locales.

Clouds of all different sizes. They come and they go, while the sky remains the same sky always. The clouds are mere guests in the sky that pass away and vanish, leaving behind the sky. The sky both exists and doesn't exist. It has substance and at the same time doesn't. And we merely accept that vast expanse and drink it in.

HARUKI MURAKAMI

*Dedicated to my Grandmother, my Parents,
and to Dr. Erika Taróczy and Dr. Zsolt Nagy*

Acknowledgement

First of all, I would like to express my sincere gratitude to my advisor Prof. Michaël Unser, for the continuous support of my Ph.D studies and related research, and for his patience, motivation, and immense knowledge.

Besides my advisor, I would like to thank the rest of my thesis committee: Prof. Pierre Vandergheynst, Prof. Dimitri Van De Ville, Dr. Jean-Christophe Olivo-Marin, and Dr. Ullrich Köthe for their encouragement, insightful comments, and thought-provoking questions.

I am grateful to Manuelle Mary and Nadia Macor for helping me out with various administrative matters inside and outside EPFL.

I have been very fortunate to be a PhD student in the Biomedical Imaging Group (BIG). I take this opportunity to thank its present and past members: Prof. Arash Amini, Anaïs Badoual, Ayush Bhandari, Dr. Emrah Bostan, Dr. Aurélien Bourquard, Prof. Kunal N. Chaudhury, Dr. Nicolas Chenouard, Dr. Ning Chu, Prof. Adrien Depeursinge, Laurène Donati, Julien Fageot, Dr. Denis Fortun, Emmanuel Froustey, Harshit Gupta, Dr. Ulugbek Kamilov, Dr. Hagai Kirshner, Dr. Stamatis Lefkimmatis, Ramtin Madani, Dr. Michael McCann, Junhong Min, Dr. Ha Nguyen, Dr. Masih Nilchian, Pedram Pad, Dr. Daniel Sage, Daniel Schmitter, Dr. Tomáš Škovránek, Dr. Ferréol Soulez, Dr. Martin Storath Kurt, Dr. Pouya Dehghani Tafti, Dr. Philippe Thévenaz, Virginie Uhlmann, Prof. Dimitri Van De Ville, Dr. Cédric Vonesch, and Dr. John Paul Ward.

Over the last couple of years, I had plenty of collaborations, discussions and brainstorming with several members of the group. Especially, I would like to thank my co-authors, John Paul, Daniel Sage, Arash, Cédric, Virginie, Adrien and Martin for all the work we did together. Also, I would like to express my gratitude to Philippe for his detailed feedback and comments on my writing.

I spent four wonderful years in the lab sharing an office with Julien and Masih. Thank you both, it was great fun to be with you.

I would like to thank Prof. Fred Hamprecht and Dr. Ullrich Köthe for hosting me at the Heidelberg Collaboratory for Image Processing over a month. Also, a big thanks to Dr. Ferran Diego and Dr. Melih Kandemir for sharing their office with me.

I am grateful to all my friends in Lausanne and Zürich, who made Switzerland my second home. In particular, a very special thanks to Daniel for teaching me how to board (and not to break my neck), John Paul for transforming me into a runner, Julien for introducing me to “graph theory”, Martin for making me addicted to “grillade de saumon”, Pouya for accompanying me in our US trip, Ulugbek for the Saturday morning shopping sessions, and Virginie for co-organizing several evenings. Also, I would like to thank all my Hungarian friends from the region, especially Dr. Ágnes Antal, Brigitta Kazi, Gergely Orbán and Dr. Tamás András Tóth for their support, inspiration and friendship. I really appreciate all the good moments that we shared together.

I would like to say special thanks to my friends from Budapest, who are geographically remote, but still part of the most important moments of my life. In particular, I would like to thank Nóra Kamarás, Dr. Nóra Fanna Kovács, Lili Nagy, and Eszter Tagányi for all their love and care.

I am grateful to have several good friends all around the world; and I am very glad that we could maintain our friendship despite the long distance. Especially, many thanks to Dr. Tamás Barbarics, Filip Biljecki, Péter Dormán, Balázs Harmat, Sevara S. Kamilova, István Komáromi, Anna Montvai, and Dr. Rodrigo Ventura.

Finally, I would like to thank my Family for their patience and support, and for the love I have received over these years. Especially, I am really grateful to my parents, Dr. Sándor and Sándorné Püspöki, for providing me with a solid foundation on which I could build. I thank my Grandmother for believing in me, for motivating me, and for her unconditional love that helped me through all difficulties.

Nomenclature

In this chapter we collect the most important notations and terminology of the thesis.

Sets

We use the classical conventions for denoting different sets. The corresponding symbols are summarized in the following table.

Symbol	Description
\mathbb{N}, \mathbb{Z}^+	Non-negative integers, including 0
\mathbb{Z}	Integers
\mathbb{R}	Real numbers
\mathbb{R}^+	Non-negative real numbers
\mathbb{C}	Complex numbers
\mathbb{R}^d	d -dimensional Euclidean space
\mathbb{Z}^d	d -dimensional integers

We use the quantity called Jaccard index for comparing the similarity and diversity of finite sample sets. In the context of the thesis, it concerns the set of the ground truth values A and the set of the detections B .

Definition 0.1 *The **Jaccard index** is defined as*

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|}, \quad (1)$$

in words, the size of the intersection divided by the size of the union of the two sets.

If $J(A, B) = 1$, we have perfect detections. If $J(A, B) = 0$, all of our detections are false positive. Otherwise, the Jaccard index takes values $0 < J(A, B) < 1$, depending on the number of false positive and negative results. In general, we illustrate the Jaccard as percentage, ranging from 0 to 100.

Concerning complex numbers, we denote the imaginary unit with j , such that $j^2 = -1$, and refer to the conjugate of a complex number z by \bar{z} .

Scalars, Vectors and Matrices

We use bold letters to distinguish vectors from scalars, e.g. $\mathbf{x} = (x_1, x_2, \dots, x_d)$, $\mathbf{x} \in \mathbb{R}^d$. We mark matrices by uppercase bold letters, e.g. \mathbf{M} . The elements of the matrices are denoted by subscripted lower case letters, such as $m_{i,j}$.

Spaces

We provide a brief summary on the notations and definitions related to vector spaces.

Symbol	Description
\mathcal{X}, \mathcal{Y}	Generic vector spaces (normed or nuclear)
$L_p, L_p(\mathbb{R}^d)$	Lebesgue space
$L_2(\mathbb{R}^d)$	Hilbert space of finite-energy functions
$\mathcal{S}(\mathbb{R}^d)$	Smooth and rapidly decreasing test functions (Schwartz space)
$\mathcal{D}(\mathbb{R}^d)$	Smooth and compactly supported test functions
$\mathcal{S}'(\mathbb{R}^d)$	Tempered distributions (generalized functions)
$\ell_p, \ell_p(\mathbb{Z}^d)$	Sequence spaces with a finite p norm
$\ell_2(\mathbb{Z}^d)$	Hilbert space of finite energy (square-summable) sequences

Definition 0.2 A **norm** on \mathcal{X} is a map $\mathcal{X} \rightarrow \mathbb{R}$, denoted by $\|\cdot\|$, fulfilling the following properties:

1. $\|f\| \geq 0$ (nonnegativity)
2. $\|af\| = |a| \|f\|$ (positive homogeneity)
3. $\|f + g\| \leq \|f\| + \|g\|$ (triangular inequality)
4. $\|f\| = 0$ implies $f = 0$ (separation of points),

$\forall a \in \mathbb{R}$ (or \mathbb{C}), and $\forall f, g, \in \mathcal{X}$. By relaxing the last requirement we obtain a seminorm. A vector space equipped with a norm is called a **normed space**.

Definition 0.3 A sequence (a_i) on \mathcal{X} is a **Cauchy sequence** iff for any $\epsilon > 0 \exists N \in \mathbb{N}$ such that

$$|a_i - a_j| < \epsilon \quad \forall i, j \geq N. \quad (2)$$

Definition 0.4 A **Banach space** is a complete normed vector space in the sense that every Cauchy sequence in the space converges to a well defined limit that is within the space.

Definition 0.5 The **inner product** is a map $\langle \cdot, \cdot \rangle : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ or \mathbb{C} , fulfilling the following axioms

1. $\langle f, g \rangle = \overline{\langle g, f \rangle}$ (conjugate symmetry)
2. $\langle a(f + g), h \rangle = a\langle f, h \rangle + a\langle g, h \rangle$ (linearity in the first argument)
3. $\langle f, f \rangle \geq 0$, $\langle f, f \rangle = 0$ implies $f = \mathbf{0}$,

$\forall f, g, h \in \mathcal{X}$ and $a \in \mathbb{R}$ or \mathbb{C} .

Definition 0.6 A **Hilbert space** H is a complete normed vector space with the norm defined by the inner product $\|f\| = \sqrt{\langle f, f \rangle}$. A Hilbert space is always a Banach space.

Definition 0.7 The **Lebesgue space** is a space of functions with a finite p -norm $L_p(\mathbb{R}^d)$, where the norms are computed in the sense of

$$\|f\|_p = \left[\int_{\mathbb{R}^d} |f(\mathbf{x})|^p d\mathbf{x} \right]^{1/p} \quad (3)$$

for $(1 \leq p < \infty)$, and

$$\|f\|_\infty = \text{ess sup} \{ |f(x)| : x \in \mathbb{R} \}. \quad (4)$$

Definition 0.8 A subset $Y = \{v_i\}_{i \in I}$ of a finite dimensional vector space \mathcal{X} is called a **basis** for \mathcal{X} if $\mathcal{X} = \text{Span}\{Y\}$ for some index set I and the vectors are linearly independent.

Definition 0.9 A family $\Upsilon = \{v_i\}_{i \in I}$ in a Hilbert space H is called a **frame** if $\forall \mathbf{x} \in H$ $\exists 0 < A \leq B < \infty$, such that

$$A\|\mathbf{x}\|^2 \leq \sum_{i \in I} |\langle v_i, \mathbf{x} \rangle|^2 \leq B\|\mathbf{x}\|^2. \quad (5)$$

A and B are the frame bounds.

Definition 0.10 We call a frame **tight frame** if its frame bounds are the same: $A = B$. A tight frame is a **Parseval frame** if $A = B = 1$.

Definition 0.11 The **Schwartz space** (space of smooth and rapidly decaying test functions) consists of infinitely differentiable functions on \mathbb{R}^d for which all of the seminorms defined by

$$\|f\|_{\mathbf{m}, \mathbf{n}} = \sup_{\mathbf{r} \in \mathbb{R}^d} |\mathbf{r}^{\mathbf{m}} \partial^{\mathbf{n}} f(\mathbf{r})| \quad \forall \mathbf{m}, \mathbf{n} \in \mathbb{N}^d \quad (6)$$

are finite.

Transformations and Operators of Continuously Defined Images

Our theory is developed on continuously defined images $f(\mathbf{x}), \mathbf{x} \in \mathbb{R}^2$. In polar coordinates, we use $\mathbf{x} = r(\cos\theta, \sin\theta)$ with $r \in \mathbb{R}^+$ and $\theta \in [0, 2\pi)$. Notations related to the derivatives of $f(\mathbf{x})$ are summarized in the following.

Symbol	Description
$D_{\mathbf{v}}^n f$	The n th order directional derivative of f along the direction \mathbf{v}
∇f	Gradient vector of f
$\mathbf{J}(\mathbf{x}_0)$	Structure tensor at location \mathbf{x}_0
$\mathbf{H}(\mathbf{x}_0)$	Hessian at location \mathbf{x}_0

Frequently used transforms of $f(\mathbf{x})$ are summarized in the following.

Symbol	Description
$\hat{f} = \mathcal{F}\{f\}$	Fourier transform of f (classical or generalized)
$f = \mathcal{F}^{-1}\{\hat{f}\}$	Inverse Fourier transform of \hat{f}
$\mathcal{H}\{f\}$	Hilbert transform of f

$\mathcal{R}\{f\}$	Complex Riesz transform of f
$\mathcal{R}^n\{f\}$	The n th order complex Riesz transform of f

Definition 0.12 The **Fourier transform** of an $L_1(\mathbb{R}^2)$ function f is denoted by $\mathcal{F}\{f\} = \hat{f}$ and computed according to

$$\hat{f}(\boldsymbol{\omega}) = \int_{\mathbb{R}^2} f(\mathbf{x}) e^{-j\langle \mathbf{x}, \boldsymbol{\omega} \rangle} d\mathbf{x}. \quad (7)$$

In polar coordinates, $\boldsymbol{\omega} = \omega(\cos\varphi, \sin\varphi)$ with $\omega \in \mathbb{R}^+$ and $\varphi \in [0, 2\pi)$. The definition is extended to L_2 (classical) or to \mathcal{S}' (generalised) in the usual ways.

Definition 0.13 The **Hilbert transform** of a 1D function f is defined as

$$\mathcal{H}\{f\}(x) = \mathcal{F}^{-1} \left\{ -\frac{j\omega}{|\omega|} \hat{f} \right\}(x) = \mathcal{F}^{-1} \{-j \operatorname{sgn}(\omega) \hat{f}\}(x). \quad (8)$$

Definition 0.14 Similarly, the **complex Riesz transform** of a function f is defined as

$$\mathcal{R}\{f\}(\mathbf{x}) = \mathcal{F}^{-1} \left\{ \frac{(\omega_x + j\omega_y)}{\|\boldsymbol{\omega}\|} \hat{f} \right\}(\mathbf{x}) = \mathcal{F}^{-1} \{e^{j\varphi} \hat{f}\}(x). \quad (9)$$

The n th order complex Riesz transform of a function f is the n -fold iterate of the complex Riesz transform of f .

Generalized Stochastic Processes

The common notations concerning generalized stochastic processes are summarized in the following.

Symbol	Description
S	Generalized stochastic scalar-field
w	White noise
σ_0^2	Variance of the white noise w
$\Phi_X(\boldsymbol{\omega})$	Power spectrum of a stationary field
$R_S(\mathbf{x}_1, \mathbf{x}_2)$	Autocorrelation of a generalized field S
$(-\Delta)^{\frac{\gamma}{2}}$	Fractional-Laplacian operator associated with the isotropic Fourier-domain multiplier $\ \boldsymbol{\omega}\ ^\gamma$
$\mathcal{B}_S\{\zeta_1, \zeta_2\}$	Correlation form of a generalized field S

A generalized random field compares to a classical random field, in the same way that a generalized function (distribution) compares to a classical function. We characterize stochastic processes (generalized random fields) through inner products with test functions that typically belong to the Schwarz space \mathcal{S} of smooth and rapidly decaying functions.

Definition 0.15 (Generalized random field) A generalized random field S is defined by its “inner products” $\langle S, \zeta \rangle$ with test functions ζ (often $\zeta \in \mathcal{S}$), which form a collection of random variables. These random variables need to fulfil two conditions: For any scalars α, β and test functions ζ_1, ζ_2 , we have

$$\langle S, \alpha\zeta_1 + \beta\zeta_2 \rangle = \alpha\langle S, \zeta_1 \rangle + \beta\langle S, \zeta_2 \rangle \quad (10)$$

Given any N sequences with limits $\lim_{k \rightarrow \infty} \zeta_{1,k} = \zeta_1, \dots, \lim_{k \rightarrow \infty} \zeta_{N,k} = \zeta_N$, and any continuous bounded function f of N variables, we have

$$\lim_{k \rightarrow \infty} \mathbb{E} \{ f(\langle S, \zeta_{1,k} \rangle, \dots, \langle S, \zeta_{N,k} \rangle) \} = \mathbb{E} \{ f(\langle S, \zeta_1 \rangle, \dots, \langle S, \zeta_N \rangle) \}. \quad (11)$$

Definition 0.16 (Correlation form [1].) The correlation form of a generalized field S is defined as

$$\mathcal{B}_S \{ \zeta_1, \zeta_2 \} = \mathbb{E} \left\{ \overline{\langle S, \zeta_1 \rangle} \langle S, \zeta_2 \rangle \right\}, \quad (12)$$

where $\zeta_1, \zeta_2 \in \mathcal{S}$ (Schwartz space), which, in cases of interest to us, can be extended to L_2 .

By the kernel theorem,

$$\mathcal{B}_S \{ \zeta_1, \zeta_2 \} = \int_{\mathbb{R}^2} \int_{\mathbb{R}^2} \overline{\zeta_1(\mathbf{r}_1)} \zeta_2(\mathbf{r}_2) R_S(\mathbf{r}_1, \mathbf{r}_2) d\mathbf{r}_1 d\mathbf{r}_2, \quad (13)$$

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

Example 0.1 (Correlation form of a generalized white noise) The correlation form of a generalized white noise takes the form

$$\mathcal{B}_w \{ \zeta_1, \zeta_2 \} = \sigma_0^2 \langle \zeta_1, \zeta_2 \rangle, \quad (14)$$

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

$$\Phi_w = \sigma_0^2. \quad (15)$$




Figure 1: Connection between the continuous-domain white noise w and the generalized stochastic process S .

We generate our stochastic process S by applying a shaping operator L^{-1} on the continuous-domain white noise or innovation process w . Or, correspondingly, we obtain the white noise by applying a whitening operator L on the stochastic process S .

Definition 0.17 (Innovation process) A stochastic process is called an innovation process if

1. it is stationary, i.e., the random variables $\langle w, \zeta_1 \rangle$ and $\langle w, \zeta_2 \rangle$ are identically distributed, provided ζ_2 is a shifted version of ζ_1 , and
2. it is white in a sense that the random variables $\langle w, \zeta_1 \rangle$ and $\langle w, \zeta_2 \rangle$ are independent, provided ζ_1, ζ_2 are non-overlapping test functions (i.e., $\zeta_1 \zeta_2 \equiv 0$).

The shaping operator determines the correlation structure of the process, and is in an inverse relationship with the whitening operator. Figure 1 illustrates the connection between the continuous-domain white noise w and the stochastic process S . This intuitive description implies that the following equations hold

$$\langle S, \zeta \rangle = \langle L^{-1} w, \zeta \rangle = \langle w, L^{-1*} \zeta \rangle, \quad (16)$$

where L^{-1*} is the adjoint operator of L^{-1} . Depending on the choice of the whitening operator, one can characterize stochastic processes with different properties. In the present thesis, we are interested in wide-sense stationary and wide-sense self-similar fields. The motivation behind this choice is that the power spectrum of many natural

images is isotropic with an $1/|\omega|^\lambda$ type of decay [2, 3], which is consistent with long-range dependencies. Also, this model of background signal fits fluorescence microscopy images well [4], which is relevant to many practical applications of orientation- and size-estimation methods.

Definition 0.18 (Wide-sense stationary field of second order) A continuous-space random field S is called wide-sense stationary (w.s.s.) if

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

In this thesis, we consider w.s.s. fields with zero mean.

Definition 0.19 (Power spectrum). The power spectrum of a stationary field S is the Fourier transform of its autocorrelation function

$$P_X(\omega) = \mathcal{F}\{R_S\}(\omega). \quad (17)$$

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 stochastic field is always real and nonnegative.

Example 0.2 (Correlation form of a zero mean w.s.s. field) The correlation form of a zero mean w.s.s. field is related to its power spectrum Φ_X by

$$\mathcal{B}_S\{\zeta_1, \zeta_2\} = \frac{1}{(2\pi)^2} \int_{\mathbb{R}^2} \overline{\hat{\zeta}_1(\omega)} \hat{\zeta}_2(\omega) \Phi_X(\omega) d\omega. \quad (18)$$

Definition 0.20 (Wide-sense self-similarity [5].) A stochastic process S with zero mean is wide-sense self-similar with scaling order H if its autocorrelation function satisfies

$$a^{2H} R_S\left(\frac{\mathbf{r}_1}{a}, \frac{\mathbf{r}_2}{a}\right) = R_S(\mathbf{r}_1, \mathbf{r}_2). \quad (19)$$

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, $R_S(\mathbf{r}, \mathbf{r}) = R_S(\mathbf{0})$, and then $\forall a, a^{2H} R_S(\mathbf{0}) = R_S(\mathbf{0})$. As a result, the variance $\text{Var}(S(\mathbf{r})) = R_S(\mathbf{r}, \mathbf{r})$ 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 the solution of the fractional stochastic equation

$$(-\Delta)^{\frac{\gamma}{2}} S = w, \quad (20)$$

$(-\Delta)^{\frac{\gamma}{2}}$ is the fractional-Laplacian operator associated with the isotropic Fourier-domain multiplier $\|\boldsymbol{\omega}\|^\gamma$ [5].

The intuitive idea here is to shape or “color” the white noise by an appropriately defined inverse fractional Laplacian, which gives to the solution an inverse-power-law spectrum and the desired (wide-sense) self-similarity and long-range dependencies. In the case where w is Gaussian, the field we obtain is the isotropic 2D generalization of a fractional Brownian motion [6].

Formally, we write that

$$\langle S, \zeta \rangle = \langle w, I_{\gamma,2}^* \zeta \rangle, \quad (21)$$

where $I_{\gamma,2}^* = (-\Delta)^{-\frac{\gamma}{2}}$ is the corrected scale-invariant inverse operator from \mathcal{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 [5], hence the notation and the index 2.)

In general, the appropriate inverse operator of the fractional Laplacian can be obtained from the Taylor series expansion of $\hat{\zeta}$ at the origin by removing a sufficient number of lower order terms:

$$I_{\gamma,2}^* \zeta(\mathbf{x}) = \int_{\mathbb{R}^2} \frac{d\boldsymbol{\omega}}{(2\pi)^2} e^{i\mathbf{x}^T \boldsymbol{\omega}} \frac{\hat{\zeta}(\boldsymbol{\omega}) - \sum_{|\mathbf{k}| \leq \lfloor \gamma - 1 \rfloor} \hat{\zeta}^{(\mathbf{k})}(\mathbf{0}) \frac{\boldsymbol{\omega}^{\mathbf{k}}}{\mathbf{k}!}}{\|\boldsymbol{\omega}\|^\gamma}. \quad (22)$$

Since Fourier-domain derivatives at $\boldsymbol{\omega} = \mathbf{0}$ correspond to moments in the space domain, in the case where ζ is a function whose moments vanish up to degree $\lfloor \gamma - 1 \rfloor$, all the terms in the sum of (22) will be zero, and

$$I_{\gamma,2}^* \zeta(\mathbf{x}) = \int_{\mathbb{R}^2} \frac{d\boldsymbol{\omega}}{(2\pi)^2} e^{i\mathbf{x}^T \boldsymbol{\omega}} \frac{\hat{\zeta}(\boldsymbol{\omega})}{\|\boldsymbol{\omega}\|^\gamma}, \quad (23)$$

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

$$\mathcal{F} \{ I_{\gamma,2}^* \zeta \}(\boldsymbol{\omega}) = \frac{\hat{\zeta}(\boldsymbol{\omega})}{\|\boldsymbol{\omega}\|^\gamma}. \quad (24)$$

Example 0.3 (Correlation form of an isotropic self-similar field) The correlation form of an isotropic self-similar field S , such that $(-\Delta)^{\frac{\gamma}{2}} S = w$ is an innovation field (white noise), is computed as

$$\begin{aligned} \mathcal{B}_S\{\zeta_1, \zeta_2\} &= \mathcal{B}_w \left\{ I_{\gamma,2}^* \zeta_1, I_{\gamma,2}^* \zeta_2 \right\} \\ &= \frac{\sigma_0^2}{(2\pi)^2} \int_0^{2\pi} \int_0^\infty \frac{\hat{\zeta}_1(\omega, \varphi) \overline{\hat{\zeta}_2(\omega, \varphi)}}{\omega^{2\gamma}} \omega d\omega d\varphi \end{aligned} \quad (25)$$

for functions ζ_1 and ζ_2 with sufficiently many vanishing moments. Here, σ_0^2 is the variance of the white noise w .

The correlation form (25) can be compared with the correlation form of an isotropic stationary random field with radial power spectrum (18). It follows that, in our model, the singular distribution $\omega^{-2\gamma}$ can be interpreted as the generalized power spectrum of our background signal.

Contents

Abstract	i
Résumé	iii
Acknowledgement	ix
1 Introduction	1
1.1 Main Contributions	4
1.1.1 Part I.	4
1.1.2 Part II.	5
1.2 Organization of the Thesis	5
I Steerable Wavelet Representations	7
2 Template-Free Directional Image Analysis	9
2.1 Overview	9
2.2 Derivative-Based Approaches	10
2.2.1 Gradient Information and Directional Derivatives	10
2.2.2 The Structure Tensor	11
2.2.3 Higher-Order Directional Structures and the Hessian	12
2.3 Directional Multiscale Approaches	14
2.3.1 Construction of Directional Filters in the Fourier Domain	14
2.3.2 Directional Wavelets with a Fixed Number of Directions	15
2.3.3 Curvelets, Shearlets, Contourlets, and Related Transforms	16

3	Unifying Framework for Steerable Representations	19
3.1	Overview	19
3.2	The Detection Problem at a Glance	20
3.2.1	Model	20
3.2.2	Steerable Representations	20
3.2.3	Particular Examples of Steerable Filters	22
3.2.4	Polar Separable Measurement Functions	24
3.3	Steerable Representations and the Riesz Transform	26
3.3.1	Definition and Properties of the Riesz Transform	26
3.3.2	Connection to the Gradient and Directional Derivatives	27
3.4	Isotropic Steerable Wavelet Frames	29
3.4.1	Radial Wavelet Profiles	29
3.4.2	Circular Harmonic Wavelets	30
3.5	Steerable Wavelets With Angular Selectivity	34
3.5.1	Monogenic Wavelets	34
4	Cramér-Rao Lower Bounds for Steerable Detectors	37
4.1	Overview	37
4.2	The Estimation Problem at a Glance	38
4.2.1	Measurement Framework	38
4.2.2	Background Noise Model	39
4.2.3	Reference Class of Estimators	40
4.2.4	Cramér-Rao Lower Bound	41
4.3	CRLB for an Estimation with Distinct Harmonics	42
4.3.1	Finite Number of Distinct Harmonics	42
4.3.2	Estimation from Best N Measurements	44
4.3.3	Asymptotic Behavior in a Single Scale	45
4.3.4	Results on the Radial Profile	47
4.4	Experiments	48
4.4.1	CRLB for Analytical Junctions	48
4.4.2	CRLB for Non-Polar-Separable Junctions	50
4.4.3	A Proposed Estimator and the CRLB	50
4.5	Extension to Wavelets	54
4.5.1	The Cramér-Rao Lower Bound	55
4.5.2	Asymptotic Behavior for All Scales and Harmonics	57
4.5.3	Experiments with Wavelets	58

4.6	Conclusion	60
4.7	Notes	60
5	Template-free Detection of Symmetries	65
5.1	Overview	65
5.2	Concept	66
5.2.1	Local Wavelet Energies	66
5.2.2	Null Hypothesis	67
5.3	Wavelet-Domain Statistics	69
5.3.1	Stationary Isotropic Fields	69
5.3.2	Self-similar Isotropic Fields	70
5.4	F -Test for M -Fold Symmetric Junctions	70
5.4.1	Single-Scale Detection	70
5.4.2	Multiscale Detection	71
5.5	Classification of Symmetry Orders	72
5.6	Experimental Results	74
5.6.1	Robustness Against Background Signal	74
5.6.2	Rotation Invariance	77
5.6.3	Scale Invariance	77
5.6.4	Micrographs	77
5.7	Conclusion	81
5.8	Notes	82
6	Steerable Wavelet Design	83
6.1	Overview	83
6.2	Concept	84
6.3	Optimal Concentrations of Energy	85
6.4	Design of M -fold-symmetric Wavelets	87
6.4.1	Design of Steerable Wavelets in the Space Domain	88
6.4.2	Design of Steerable Wavelets in the Fourier Domain	89
6.5	Angular Weights and Asymmetric Wavelets	92
6.6	Comparison of Different Designs	96
6.7	Application to Image Analysis	97
6.7.1	Summary of the Algorithm	97
6.7.2	Steering	98
6.7.3	Connecting Junctions	100

6.8	Experimental Results	101
6.8.1	Robustness Against Noise	101
6.8.2	Multiscale Properties	106
6.8.3	Rotation Invariance	110
6.8.4	Practical Applications	111
6.9	Conclusion	112
6.10	Notes	113
II Transformability: Beyond Rotations		115
7	Transformability: Beyond Rotations	117
7.1	Overview	117
7.2	Lie groups and transformation groups	118
7.3	Transformability and Eigenfunctions	120
7.3.1	Characterization of shiftable and scalable families	122
8	Scalable Filters for Tight Wavelet Frames	125
8.1	Overview	125
8.2	Mathematical Formulation	126
8.2.1	Admissible Multipliers and Extended Frames	126
8.2.2	Admissible Multipliers and Transformable Families	128
8.2.3	Families of Dilation Multipliers	129
8.2.4	Adapting Extended Frames	130
8.2.5	Isotropic Wavelets and Extended Frames	132
8.3	Smooth Scaling of Wavelets	132
8.3.1	Example of a Localized Frequency Multiplier	133
8.3.2	Pseudo-Scaling	134
8.3.3	Comments on the Construction	137
8.4	Special Case: One Complex Multiplier	137
8.5	SpotCaliper	140
8.5.1	Steps of the Algorithm	141
8.5.2	Modes and Parameter Settings	141
8.5.3	Post-processing and Visualization	142
8.6	Experimental Results	144
8.6.1	Reference Methods	146

8.6.2	Running Time	146
8.6.3	Robustness Against Background Signal	149
8.6.4	Practical Applications	151
8.7	Conclusion	153
8.8	Notes	154
9	Conclusions	155
9.1	Summary of Results	155
9.2	Outlook	157
	Bibliography	160
	Curriculum Vitæ	179

Chapter 1

Introduction

Advances in microscopy and in medical imaging (computed tomography, X-ray) are shaping the future of biomedical image processing, and in particular, biomedical image analysis. The state of the art in the field is developing rapidly to fulfil the upcoming needs of biomedical and biological research.

Currently, modern microscopes are routinely producing large datasets (2D/3D images over time) that require the storage of several gigabytes per measurement. The need to automatically process this kind of spatio-temporal data is becoming unavoidable; their proper evaluation is no longer manageable manually. This imposes a big challenge in memory handling, data mining, and image processing in general. Many modern approaches in image processing meet this challenge by emphasizing the idea of parsimony or sparsity, through seeking more efficient representations of relevant information in images. In image processing, sparsity implies the assumption that an image can be represented or well approximated in some transform domain using a much smaller number of coefficients than its apparent dimension, which is typically measured in pixels (2D) or voxels (3D). This allows us to compress the image efficiently by expressing it in the suitable transform domain and only keeping the most important coefficients. Such an approach works well if the considered transformation captures the most important features of the image in a few large coefficients. This also means that, those transforms provide a certain measure of locality that allows us to isolate information related to a particular aspect of the image (*e.g.*, frequency, spatial location).

To generate a local representation one can be motivated by two ideas: using a generic

framework and universal methods (*e.g.*, Fourier analysis, wavelets) or dealing with signal-adaptive and specific methods (*e.g.*, local principal/independent component analysis), that are somehow contradictory. Our interest is in representations that lie between these extremes, where the predefined building blocks can be adapted to the local information by applying some transformation. Thus, we maintain generality and universal performance while capturing specific information in a nearly optimal way.

When considering possible adaptations of such schemes, we guide ourselves by the fact that local geometric structures are often repeated throughout natural images; however, each occurrence has typically been deformed by an unknown geometric transformation such as a combination of rotation, translation, and scaling. For example, in microscopic images showing cell cultures, each cell typically has the same shape but their size, orientation and location can differ. Similar observations are made for many types of signals, for instance, those encountered in bioengineering, seismology, or audio processing.

The information that corresponds to the notion of rotation, directionality, and orientation is a key component in the quantitative analysis of images. By those terms, we refer to local directional cues and features that one can identify in natural images. The area of applications based on the detection of orientation is continuously growing as the importance of directionality is becoming more and more evident in image processing. The range of applications spans topics from astronomy [7–9], aerial and satellite imagery [10–12], material sciences [13], to biological and medical applications. Focusing on the last two categories, the palette is quite varied: detection of nodules in the lungs [14] and vessels in retinal fundus images [15, 16] as well as bioimaging [17–19] and neuroimaging [20, 21] are just a few examples. Investigations of collagen in the arterial adventitia also rely on directional analysis [22]. Neuron tracking is of primal importance to understand the development of the brain and requires robust directional image-analysis tools to capture dendrites in 2D and 3D [23]. Further examples include the study of aging of elastin in human cerebral arteries [24], or the estimation of orientation and localization of fluorescent dipoles [25]. The analysis of local directional patterns also includes the detection of ridges and junctions of any order; applications can be found in [26–30]. As a particular example, the accurate detection of cell structures that exhibit polygonal shapes is fundamental in stem cell research [31].

The notion of dilation is fundamental in the detection and size estimation of objects in the image processing of micrographs. Most medical and biological applications relying on cell cultures seek concentration information and accurate image statistics. In molecular biology, cell statistics provide information on the growth rate of microorgan-

isms (that correspond to their division into new cells). In time-lapse microscopy, the analysis can be used to determine the cell lineage. In fluorescence microscopy, the observation of biomolecular interactions in cellular processes requires the quantification of gene expression. The corresponding structures typically have a roundish shape (*e.g.*, nuclei, vesicles, spindle pole bodies). The appearance of circular objects of different size is also common in high-throughput imaging with microarrays. The precise detection of the spots and size measurements are important for understanding *Enzyme-Linked ImmunoSpot* images [32]. Those assays are used to monitor the adaptive immune responses of humans and animals by visualizing antigen-antibody reactions. Accompanying the interaction, the naturally transparent cells are releasing cytokines that color the surrounding substance, typically resulting in blue or black spots. The size of the spots is proportional to the strength of the reaction. Also, in antimicrobial susceptibility testing of significant bacterial isolates, the size of the spots describing the zones of the inhibition areas are measured and used as an indicator of antibiotic-resistance [33].

The common challenge of the applications cited above originates from the necessity of the precise retrieval of geometric information. This motivates us to characterize images through their inherent geometric properties. We treat images as a collection of geometric features that are described by symmetries, angle, orientation, size and location. Instead of describing these transformations through shape dictionaries that capture prior information as hard constraints, we integrate it in a flexible and adaptable way. Hence, we choose a continuous domain formulation, because it overcomes the limitations of dictionary-based approaches that only allow for a finite number of discrete templates. An additional challenge is computational efficiency, since most of the experiments are supposed to run in real time. Therefore, the methods we are looking for must have an implementation that allows for fast processing. Moreover, they should be robust and enable the analysis of a broad class of images. Finally, in order to complete the framework, we intend to characterize the precision of our framework and its applicability to different pattern classes.

In this context, we provide methodologies and frameworks that play a role in the understanding of biological images. More specifically, the focus of this thesis is on local transformable representations and their applications related to wavelet schemes, since they are one of the most efficient and widely used processing tools for feature extraction in image analysis. They provide a unifying framework for decomposing images in a multi-scale hierarchy, and thus, have many successful applications such as contour detection, image filtering, orientation and texture analysis, synthesis and so forth (*e.g.*, [34, 35]).

1.1 Main Contributions

The main contributions of this thesis are categorized into two parts. Part I refers to steerable representations that account for rotation transformations, while in Part II we generalize the characterization of representations related to geometric transformations beyond rotations.

1.1.1 Part I.

- We fully characterize the problem of detecting the location and local orientation of directional patterns in an image. We establish a description of our directional pattern model. Our focus is on steerable representations, which are linked to polar separable detector functions. We explain the property of steerability, and provide a construction framework for steerable filters and wavelets, highlighting the connection with the Riesz transform.
- We derive the fundamental Cramér-Rao lower bounds (CRLB) on the error when estimating local orientations. We develop a formulation based on steerable filters, and use the obtained bounds to provide solutions and guidelines for improving and optimizing the choice of the detector functions. In particular, we address the problem of choosing the best angular frequencies to achieve the lowest CRLB, as well as identifying the optimal radial profile of the detector function. We also discuss the CRLB for different detector schemes, including multiscale approaches, in particular, wavelets.
- We propose an algorithm to detect and group different kinds of local symmetries in images in a multiscale, rotation-invariant, and template-free way. By template-free, we mean that there is need neither to design a specific junction or ridge-shaped template nor to align it with the pattern of interest. Instead, we propose an efficient wavelet-based method to determine the order of local symmetry at each location. Relying on the energy of the wavelet coefficients, we define a measure of local symmetry at each location in the image based on a statistical formulation. Given the measure, we provide a classification rule to distinguish different orders of symmetry.
- We follow a complementary approach that serves to address two objectives not covered by the methods of the previous chapter. First, we target the detection

of junctions and patterns without constraining ourselves to symmetric patterns only. Second, we aim not only at detecting these local patterns, but also at estimating their orientation. These dual objectives are achieved by designing steerable wavelets that are shaped to match a desired angular profile. The approaches we develop are general and can be used for the detection of arbitrary symmetric and asymmetric junctions. We answer the question of optimal design with respect to the angular energy distribution of the wavelets. We propose a framework based on analytical optimization for aligning steerable templates with the orientation of patterns.

1.1.2 Part II.

- The main contribution and the novelty of this part of the thesis is the construction of isotropic wavelet frames that are adaptable under dilation operations. In analogy with steerable wavelets, we propose a general construction of adaptable tight wavelet frames, with an emphasis on scaling operations. In particular, the wavelets that are derived can be “dilated” by a procedure comparable to the operation of steering steerable wavelets. To achieve this, we provide a definition of scalability with respect to discrete wavelet transforms. We identify possible solutions for scaling functions and give specific criteria on their applicability to wavelet schemes. Using scalable filters, our wavelet frames can approximate intermediate scales based on several patterns that can be linearly combined with adjustable coefficients. We use this property to generate disk patterns with continuously varying sizes in between two sequential (discrete) wavelet scales. As a result, we create a novel wavelet-based image-analysis approach with a software implementation in Java, providing fast automatic detection of circular patterns (spots), combined with the precise estimation of their size.

1.2 Organization of the Thesis

Except Chapter 2, 3, and 7, which also contain introductory materials and definitions, the content of the chapters is all based on our original work.

Chapter 2 aims at providing a broad overview of template-free techniques for the directional analysis of images. We review the state of the art in the field and unify our perspective within a common framework using operators. The intent is to provide image-

processing methods that can be deployed in algorithms that analyze biomedical images with improved rotation invariance and high directional sensitivity.

In Chapter 3, we expose the principles behind our directional pattern model and develop a general framework for the detection of their location and orientation. In particular, we define steerable representations that correspond to rotations. We present the Riesz transform and its role in the generation of steerable filters and wavelets. As an example, we give a construction rule on circular harmonic wavelets.

In Chapter 4, we derive the Cramér-Rao lower bound describing the error that is associated to the estimation of the orientation and local features of an image. Based on these bounds, we prove theorems that we use to construct optimal detector functions.

In Chapter 5, we present a method for the identification and classification of local symmetries in biological images. Based on circular harmonic wavelet frames presented in Chapter 3, we propose to detect and group different kinds of local symmetries in images in a multiscale and rotation-invariant way. There, we use the property that the wavelet transform generates steerable wavelet channels corresponding to different symmetry orders and distributes the energy of the signal among a set of angular harmonics.

In Chapter 6, based on the circular harmonic wavelet representations, we create new wavelets schemes whose energy is concentrated along specific angular directions. We develop a method for designing steerable wavelets that can detect local centers of symmetry in images. Based on this design, we then propose an algorithm and a corresponding ImageJ plug-in for estimating the locations and the orientations of junctions in biological micrographs.

In Chapter 7, we generalize the notion of steerability, and identify classes of functions that fulfill this generalized definition for relevant geometric transformations besides rotation (i.e., translation and scaling). We summarize the general definitions and concepts that serve as a basis for the new constructions and results that are developed in Chapter 8. There, we propose the construction of locally scalable filters based on discrete tight wavelet frames that can be transformed to approximate a continuous wavelet transform. As a particular application, we propose a framework for the detection of location and size of circular objects and validate it in realistic scenarios in bioimaging.

In Chapter 9, we conclude on the research and results presented in this thesis and discuss the potential areas of interest for future research related to our work.

Part I

Steerable Wavelet Representations

Chapter 2

State of the Art in Template-Free Directional Image Analysis

2.1 Overview

In the early sixties it was demonstrated that directionality has a key role in visual perception. This result of Hubel and Wiesel [36] initiated a field of research for the following decades [37, 38]. Based on these studies it became clear that the organization of the primary visual cortex makes the perception orientation-sensitive, highlighting the importance of directional cues, such as edges, ridges and corners. Our visual system is able to capture and summarize this information using a small number of visual cells, like orientation-selective feature detectors. It therefore seemed reasonable to benefit from these results and try to mimic some of those ideas. The application area of these methods is continuously growing as image analysis is becoming more relevant in data processing.

This chapter aims at providing a broad overview of state-of-the-art techniques in the field for the directional analysis of images. We start with classical methods such as directional-gradient and the structure tensor. Then, we discuss how these methods can be improved with respect to robustness, invariance to geometric transformations (with a particular interest in scaling), and computational cost. We move forward to higher degrees of directional selectivity and discuss Hessian-based detection schemes. We give

a brief overview of Fourier filters, directional wavelets, curvelets, and shearlets.

We focus on the continuous domain setup for explaining the relevant concepts because it allows for convenient, compact, and intuitive formulation. It primarily involves differential and convolution operators (smoothing filters and wavelets) that are acting on continuously defined images, $f(\mathbf{x}), \mathbf{x} \in \mathbb{R}^2$. The final transcription of a continuous domain formula into an algorithm requires the discretization of the underlying filters which can be achieved using standard techniques. For instance, partial derivatives can be closely approximated using finite differences, while there are well-established techniques for computing wavelets using digital filters¹.

2.2 Derivative-Based Approaches

2.2.1 Gradient Information and Directional Derivatives

Some of the earliest and simplest techniques in image analysis to account for orientation rely on gradient information. Intuitively, the direction of the gradient corresponds to the direction of steepest ascent. The local direction of an image f at \mathbf{x}_0 can be estimated in terms of the direction orthogonal to its gradient. A direction is specified in \mathbb{R}^2 by a unit vector $\mathbf{v} = (v_1, v_2) \in \mathbb{R}^2$ with $\|\mathbf{v}\| = 1$. The first-order directional derivative $D_{\mathbf{v}}f$ along the direction \mathbf{v} can be expressed in terms of the gradient as

$$D_{\mathbf{v}}f(\mathbf{x}_0) = \lim_{h \rightarrow 0} \frac{f(\mathbf{x}_0) - f(\mathbf{x}_0 - h\mathbf{v})}{h} = \langle \mathbf{v}, \nabla f(\mathbf{x}_0) \rangle, \quad (2.1)$$

where the right hand side is the inner product between \mathbf{v} and the gradient vector $\nabla f(\mathbf{x}_0)$ evaluated at \mathbf{x}_0 . We note that (2.1) is maximum when \mathbf{v} is collinear with $\nabla f(\mathbf{x}_0)$ (by the Cauchy - Schwartz inequality). Conversely, $D_{\mathbf{v}_0}f(\mathbf{x}_0)$ vanishes when $\mathbf{v}_0 \perp \nabla f$, so that \mathbf{v}_0 provides us with a local estimate of the directionality of the image.

Gradient-based orientation estimators are frequently used as they can be discretized and implemented easily. However, the gradient-based estimation of the orientations is sensitive to noise. The robustness can be improved by smoothing the image by a Gaussian kernel before taking the derivative. A still very popular method based on gradients is Canny's classical edge detector [40].

¹The content in this chapter is based on the article [39]: Z. Püspöki, M. Storath, D. Sage, M. Unser, "Transforms and Operators for Directional Bioimage Analysis: A Survey", *Advances in Anatomy, Embryology and Cell Biology*, in press, 2016.

2.2.2 The Structure Tensor

The estimation of the local orientation using derivatives can be made more robust by using the structure tensor [41]. The structure tensor is a matrix derived from the gradient of the image and can be interpreted as a localized covariance matrix of the gradient. Since the pioneering work of Förstner [42], Bigün and Granlund [43], and Harris and Stephens [44], the structure tensor has become a tool for the analysis of low-level features, in particular for corner and edge detection as well as texture analysis. In 2D, the structure tensor at location \mathbf{x}_0 is defined by

$$\mathbf{J}(\mathbf{x}_0) = \int_{\mathbb{R}^2} \tilde{w}(\mathbf{x} - \mathbf{x}_0) (\nabla f(\mathbf{x})) \nabla^T f(\mathbf{x}) dx_1 dx_2, \quad (2.2)$$

where \tilde{w} is a non-negative isotropic observation window (e.g., a Gaussian) centered at \mathbf{x}_0 . More explicitly, the (2×2) matrix $\mathbf{J}(\mathbf{x}_0)$ reads

$$\mathbf{J}(\mathbf{x}_0) = \int_{\mathbb{R}^2} \tilde{w}(\mathbf{x} - \mathbf{x}_0) \begin{pmatrix} f_{x_1}^2(\mathbf{x}) & f_{x_1}(\mathbf{x})f_{x_2}(\mathbf{x}) \\ f_{x_2}(\mathbf{x})f_{x_1}(\mathbf{x}) & f_{x_2}^2(\mathbf{x}) \end{pmatrix} dx_1 dx_2 \quad (2.3)$$

$$= \begin{pmatrix} (\tilde{w} * f_{x_1}^2)(\mathbf{x}_0) & (\tilde{w} * f_{x_1}f_{x_2})(\mathbf{x}_0) \\ (\tilde{w} * f_{x_2}f_{x_1})(\mathbf{x}_0) & (\tilde{w} * f_{x_2}^2)(\mathbf{x}_0) \end{pmatrix}, \quad (2.4)$$

where $\tilde{w} * f$ denotes the convolution of \tilde{w} and f , and the partial derivative of f with respect to variable x_i is denoted by f_{x_i} . This reveals that \mathbf{J} is a smoothed version of

$$\begin{pmatrix} f_{x_1}^2(\mathbf{x}) & f_{x_1}(\mathbf{x})f_{x_2}(\mathbf{x}) \\ f_{x_2}(\mathbf{x})f_{x_1}(\mathbf{x}) & f_{x_2}^2(\mathbf{x}) \end{pmatrix}. \quad (2.5)$$

The eigenvalues of the structure tensor are noted λ_{\max} and λ_{\min} , with $\lambda_{\min}, \lambda_{\max} \in \mathbb{R}$. They carry information about the distribution of the gradient within the window \tilde{w} . Depending on the eigenvalues, one can discriminate between homogeneous regions, rotational symmetric regions without a predominant direction, and regions with dominant directions. For this purpose, two measures are defined, the so called energy E and the coherence C . The energy is defined based on the eigenvalues of the structure tensor as $E = |\lambda_1| + |\lambda_2|$. If $E \approx 0$, which corresponds to $\lambda_{\max} = \lambda_{\min} \approx 0$, then the region is homogeneous. If $E \gg 0$, then the characteristics of the structure are determined by the coherency information. The coherency information C is defined as

$$0 \leq C = \frac{\lambda_{\max} - \lambda_{\min}}{\lambda_{\max} + \lambda_{\min}} = \frac{\sqrt{(J_{22} - J_{11})^2 + 4J_{12}^2}}{J_{22} + J_{11}} \leq 1, \quad (2.6)$$

where J_{ij} denotes elements of the structure tensor. If $C \approx 0$, which corresponds to $\lambda_{\max} \approx \lambda_{\min}$, then the region is rotationally symmetric without predominant direction, the structure has no orientation. On the contrary, a coherency close to 1 indicates that the structure in the image is locally 1D.

The energy of the derivative in the direction \mathbf{v} can be expressed as

$$\|D_{\mathbf{v}}f\|_{\tilde{w}}^2 = \langle \mathbf{v}^T \nabla f, \mathbf{v}^T \nabla f \rangle_{\tilde{w}} = \mathbf{v}^T \mathbf{J} \mathbf{v}. \quad (2.7)$$

This means that, in the window centered around \mathbf{x}_0 , the dominant orientation of the neighborhood can be computed by

$$\mathbf{v}_1 = \arg \min_{\|\mathbf{v}\|=1} \|D_{\mathbf{v}}f\|_{\tilde{w}}^2. \quad (2.8)$$

We interpret $\|D_{\mathbf{v}}f\|_{\tilde{w}}^2$ as the average energy in the window defined by \tilde{w} and centered at \mathbf{x}_0 . Moreover, $D_{\mathbf{v}}f = \langle \nabla f, \mathbf{v} \rangle$ is the derivative in the direction of \mathbf{v} . The minimizing argument corresponds to the eigenvector with the smallest eigenvalue of the structure tensor at \mathbf{x}_0 . The dominant orientation of the pattern in the local window \tilde{w} is computed as $\mathbf{v}_1 = (\cos\theta, \sin\theta)$, with $\theta = \frac{1}{2} \arctan\left(\frac{2J_{12}}{J_{22}-J_{11}}\right)$.

In the 3D shape estimation of DNA molecules from stereo cryo-electron micrographs [45], the authors took advantage of the structure tensor. Other applications can be found in the works of Köthe [46] and Bigün *et al.* [47].

While simple and computationally efficient, the structure-tensor method has drawbacks: it only takes into account one specific scale, the localization accuracy for corners is low, and the integration of edge and corner detection is ad hoc (e.g., Harris' corner detector [44]).

2.2.3 Higher-Order Directional Structures and the Hessian

To capture higher-order directional structures, the gradient information is replaced by higher-order derivatives. In general, an n th-order directional derivative is associated with n directions. Taking all of these to be the same, the directional derivative of order n in \mathbb{R}^2 is defined as

$$D_{\mathbf{v}}^n f(\mathbf{x}) = \sum_{k=0}^n \binom{n}{k} v_1^k v_2^{n-k} \partial_{x_1}^k \partial_{x_2}^{n-k} f(\mathbf{x}), \quad (2.9)$$

which is a linear combination of partial derivatives of order n . More specifically, if we fix $n = 2$ and the unit vector $\mathbf{v}_\theta = (\cos\theta, \sin\theta)$, we obtain

$$D_{\mathbf{v}_\theta}^2 f(\mathbf{x}) = \cos^2(\theta) \partial_{x_1}^2 f(\mathbf{x}) + 2 \cos(\theta) \sin(\theta) \partial_{x_1} \partial_{x_2} f(\mathbf{x}) + \sin^2(\theta) \partial_{x_2}^2 f(\mathbf{x}). \quad (2.10)$$

Xia *et al.* [28] use a template-free method for the detection and grouping of junctions based on the strength of directional derivatives.

The Hessian filter is a square matrix of second-order partial derivatives of a function. For example, in 2D, the smoothed Hessian matrix, useful for ridge detection at location \mathbf{x}_0 , can be written as

$$\mathbf{H}(\mathbf{x}_0) = \begin{pmatrix} (\tilde{w}_{11} * f)(\mathbf{x}_0) & (\tilde{w}_{12} * f)(\mathbf{x}_0) \\ (\tilde{w}_{21} * f)(\mathbf{x}_0) & (\tilde{w}_{22} * f)(\mathbf{x}_0) \end{pmatrix}, \quad (2.11)$$

where \tilde{w} is a smoothing kernel and $\tilde{w}_{ij} = \partial_{x_i} \partial_{x_j} \tilde{w}$ denotes its derivatives with respect to the coordinates x_i and x_j . In the window centered around \mathbf{x}_0 , the dominant orientation of the ridge is

$$\mathbf{v}_2 = \arg \min_{\|\mathbf{v}\|=1} (\mathbf{v}^T \mathbf{H} \mathbf{v}). \quad (2.12)$$

The minimizing argument corresponds to the eigenvector with the smallest eigenvalue of the Hessian at \mathbf{x}_0 . The eigenvectors of the Hessian are orthogonal to each other, so the eigenvector with the largest eigenvalue corresponds to the direction orthogonal to the ridge.

A sample application of the Hessian filter is vessel enhancement. In [48], Frangi *et al.* define a measure called vesselness which corresponds to the likelihood of an image region to contain vessels or other image ridges. The vesselness measure is derived based on the eigenvalues of the Hessian filter. In 2D, a vessel is detected when one of the eigenvalues is close to zero ($\lambda_1 \approx 0$) and the other one is much larger $|\lambda_2| \gg |\lambda_1|$. The direction of the ridge is given by the eigenvector of the Hessian filter output corresponding to λ_1 . The authors define the measure of vesselness as

$$V = \begin{cases} 0, & \text{if } \lambda_1 > 0 \\ \exp\left(-\frac{(\lambda_1/\lambda_2)^2}{2\beta_1}\right) \left(1 - \exp\left(-\frac{\lambda_1^2 + \lambda_2^2}{2\beta_2}\right)\right), & \text{otherwise,} \end{cases} \quad (2.13)$$

where β_1 and β_2 control the sensitivity of the filter². Alternative vesselness measures based on the Hessian have been proposed by Lorenz *et al.* [49] and Sato *et al.* [50].

²Plugin available at <http://fiji.sc/Frangi/>

2.3 Directional Multiscale Approaches

In natural images, oriented patterns are typically living on different scales, for example, thin and thick blood vessels. To analyze them, methods that extract oriented structures separately at different scales are required. The classical tools for multiscale analysis are wavelets. In a nutshell, a wavelet is a bandpass filter that responds almost exclusively to features of a certain scale.

In the present thesis, we focus on wavelet bases with dyadic scale progressions. We note that, however, there exist bases with other integer dilation factors [51, 52], and also with arbitrary real [53] and rational dilation factors [54, 55]. For implementation, we use regular rectangular sampling grids, for other solutions, like quincunx wavelets, we refer to [56] and [57].

The separable wavelet transform that is commonly used is computationally very efficient but provides only limited directional information and suffers from inherent directional bias. Its operation consists of filtering with 1D wavelets with respect to the horizontal and vertical directions. As a result, two pure orientations (vertical and horizontal), and a mixed channel of diagonal directions are extracted. Using the dual-tree complex wavelet transform [58]³, one can increase the number of directions to six while retaining the computational efficiency of the separable wavelet transform. (We refer to [59] for a detailed treatment of this transform.) Next, we describe how to construct wavelets with an even higher orientational selectivity at the price of higher computational cost.

2.3.1 Construction of Directional Filters in the Fourier Domain

For junction and line detection, it is often desirable to design linear filters that respond exclusively to structures within a narrow angular range. Methods based on the Fourier transform provide an effective way to construct such orientation-selective filters. The key property for directional analysis is that rotations in the spatial domain propagate as rotations to the Fourier domain. Formally, we write that

$$f(\mathbf{R}_\theta \mathbf{x}) \xleftrightarrow{\mathcal{F}} \hat{f}(\mathbf{R}_\theta \boldsymbol{\omega}), \quad (2.14)$$

where \mathbf{R}_θ denotes a rotation by the angle θ .

³Available at <http://eeweb.poly.edu/iselesni/WaveletSoftware/>

The construction is based on a filter ϕ whose Fourier transform $\hat{\phi}$ is supported on a wedge around the ω_1 axis. In order to avoid favoring special orientations, one typically requires that $\hat{\phi}$ be non-negative and that it form (at least approximately) a partition of unity of the Fourier plane under rotation, like

$$\sum_{\theta_m} |\hat{\phi}(\mathbf{R}_{\theta_m} \boldsymbol{\omega})|^2 = 1, \text{ for all } \boldsymbol{\omega} \in \mathbb{R}^2 \setminus \{\mathbf{0}\}. \quad (2.15)$$

Here, $\theta_1, \dots, \theta_n$ are arbitrary orientations which are typically selected to be equidistant, with $\theta_m = (m-1)\pi/n$. To get filters that are well localized in the spatial domain, one chooses $\hat{\phi}$ to be a smooth function; for example the Meyer window function [51, 60]. A directionally filtered image f_{θ_m} can be easily computed by rotating the window $\hat{\phi}$ by θ_m and multiplying it with the Fourier transform \hat{f} of the image, and by transforming back to the spatial domain. This is written

$$f_{\theta_m}(\mathbf{x}) = \mathcal{F}^{-1} \{ \hat{\phi}(\mathbf{R}_{\theta_m} \cdot) \hat{f} \}(\mathbf{x}). \quad (2.16)$$

(We refer to [61] for filterings based on convolutions in the spatial domain.) The resulting image f_{θ_m} contains structures that are oriented along the direction θ_m . The local orientation θ is given by the orientation of the minimum filter response

$$\theta(\mathbf{x}) = \arg \min_{\theta_m} |f_{\theta_m}(\mathbf{x})|. \quad (2.17)$$

Such directional filters have been used in fingerprint enhancement [62] and in crossing-preserving smoothing of images [63, 64].

2.3.2 Directional Wavelets with a Fixed Number of Directions

Now we augment the directional filters by scale-selectivity. Our starting point is the windowing function of (2.15). The simplest way to construct a directional wavelet transform is to partition the Fourier domain into dyadic frequency bands (“octaves”). To ensure a complete covering of the frequency plane, we postulate again nonnegativity and a partition-of-unity property of the form

$$\sum_{i \in \mathbb{Z}} \sum_{\theta_m} |\hat{\phi}(2^{-i} \mathbf{R}_{\theta_m} \boldsymbol{\omega})|^2 = 1, \text{ for } \boldsymbol{\omega} \in \mathbb{R}^2 \setminus \{\mathbf{0}\}. \quad (2.18)$$

Classical examples of this type are the Gabor wavelets that cover the frequency plane using Gaussian windows which approximate (rescaled) partition-of-unity [65, 66]. These

serve as a model for filters in the mammalian visual system [67, 68]. Alternative constructions are Cauchy wavelets [69] or constructions based on the Meyer window functions [51, 60]. We refer to Vandergheynst and Gobbers [70] and Jacques *et al.* [71] for further information on the design of directional wavelets. In particular, sharply direction-selective Cauchy wavelets have been used for symmetry detection [69].

A further possible direction is the extension of directional wavelet transforms to nonuniform lattices such as polar grids or general graphs [72–74].

Kovesi accounts for detecting symmetries and anti-symmetries based on local phase information provided by wavelets constructed with complex valued Log Gabor functions [75]. However, his work is mostly focusing on bilateral symmetries and ignores rotational ones.

2.3.3 Curvelets, Shearlets, Contourlets, and Related Transforms

Curvelets, contourlets, and shearlets are 2D generalizations of directional wavelets that are aimed at the analysis of contours and smooth curves. However, they are equipped with an anisotropic scaling which follows a parabolic scaling law. This models the scaling behavior of sufficiently smooth curves in 2D. As a result, curvelets, contourlets, and shearlets of finer scales are more and more elongated, corresponding to an increase in their directional selectivity.

Among these transforms, curvelets were the first to be defined in the continuous domain [76]. In their case, directionality is provided by rotations; all curvelet atoms within a single scale are shifts or rotations of a single pattern. Similarly to directional wavelets, their construction corresponds to a division of the frequency plane into rings (representing scales), where each ring is further divided into wedges (corresponding to orientations). However, contrary to directional wavelets, where each scale has the same number of wedges, in the case of curvelets, the number of wedges (or orientations) doubles every other scale.

Contourlets are inspired by curvelets, but attempt to overcome the difficulties associated with the discretization through a discrete-domain definition in terms of filter-banks [77]. As a consequence, for contourlets, the frequency domain divisions are based on wedges on square-shaped rings instead of circular ones.

Finally, comparable to curvelets, shearlets also start from a continuous-domain definition which is then discretized. The main difference between those methods is that, for shearlets, directionality is defined in terms of shear transformations instead of rotations [78–80]. The reason behind this choice is that, unlike rotations, shears can be

perfectly discretized on a Cartesian grid for integer shear parameters.

Applications of the previously mentioned methods include texture classification of tissues in computed tomography [81], texture analysis [82], image denoising [83], contrast enhancement [84], and reconstruction in limited-angle tomography [85]. Furthermore, they are closely related to a mathematically rigorous notion of the orientation of image features, the so-called wavefront set [86, 87]. Loosely speaking, the wavefront set is the collection of all edges along with their normal directions. This concept is used for the geometric separation of points from curvilinear structures, for instance, to separate spines and dendrites [88] and for edge detection with resolution of overlaying edges [79, 89, 90].

A thorough description of wavelets, contourlets, shearlets and other related transforms can be found in [91]. Relevant software packages implementing these transforms are the Matlab toolboxes CurveLab⁴, ShearLab⁵, FFST⁶, and the 2D Shearlet Toolbox⁷.

⁴Available at <http://www.curvelet.org/>

⁵Available at <http://www.shearlab.org/>

⁶Available at <http://www.mathematik.uni-kl.de/imagepro/members/haeuser/ffst/>

⁷Available at <http://www.math.uh.edu/~dlabate/software.html>

Chapter 3

Unifying Framework for Steerable Representations

3.1 Overview

For the purpose of detecting or enhancing a given type of directional pattern (edge, junction, ridge, corner), a natural inclination is to try to match it against similar patterns. The simplest way to do that is to construct a template and try to align it with the pattern of interest. Usually, such algorithms rely on the discretization of the orientation. To obtain accurate results, a fine discretization is required. In general, Fourier filters and wavelet transforms are computationally expensive in this role because a full 2D filter operation has to be computed for each discretized direction.

However, an important exception is provided by steerable filters, where one may perform arbitrary (continuous) rotations and optimizations with a substantially reduced computational overhead. The basics of steerability were formulated by Freeman and Adelson in the early nineties [92–94] and developed further by Perona [95], Simoncelli and Farid [96], and Unser and Van de Ville [97], Unser and Chenouard [98], and Ward *et al.* [99, 100].

In this chapter, we formulate the problem of detecting the location and local orientation of directional patterns in an image. These local patterns in particular include junctions of any order, and in the interest of applications, we shall use the term junc-

tion to refer to any such pattern. But our derivation equally applies to any other kind of directional pattern. In the following, we give a description of our junction model. Our focus is on steerable representations, which are linked to polar separable detector functions. We explain the property of steerability, and provide a construction framework for steerable filters and wavelets, highlighting the connection with the Riesz transform.

3.2 The Detection Problem at a Glance

3.2.1 Model

Our interest in the first part of this thesis lies in estimating the orientation of local directional patterns in an image. Without loss of generality, the pattern is assumed to be centered at $\mathbf{x} = \mathbf{0}$. Let $J(\mathbf{x})$ denote the general shape of the pattern of interest, and $J_{\theta^*}(\mathbf{x}) = J(\mathbf{R}_{\theta^*}\mathbf{x})$ its rotated version with unknown orientation θ^* that is to be estimated. Our local image model is then

$$I_{\text{loc}}(\mathbf{x}) = J(\mathbf{R}_{\theta^*}\mathbf{x}) + S(\mathbf{x}) \quad (3.1)$$

or, in polar coordinates,

$$I_{\text{loc}}(r, \theta) = J(r, \theta + \theta^*) + S(r, \theta), \quad (3.2)$$

where S denotes the background signal which, for our purposes is considered to be noise.

The observations (measurements, or coefficients) used in estimating θ^* are of the form

$$q_\alpha = \langle I_{\text{loc}}, \xi_\alpha \rangle = \langle J_{\theta^*}, \xi_\alpha \rangle + \langle S, \xi_\alpha \rangle, \quad (3.3)$$

where $\{\xi_\alpha\}$ is a set of steerable measurement functions indexed by α (introduced in the next subsections). The problem studied here is to estimate the local orientation angle θ^* from the steerable measurements q_α .

In the following, we give a description of the steerable measurement functions of the model. The noise term will play an important role in the next chapters; it will be fully described in Chapter 4.

3.2.2 Steerable Representations

The justification for the use of steerable coefficients as measurements in (3.3) comes from the theory of local steerable representations due to Freeman and Adelson [94].

One intuitive way to interpret the estimation problem is that we are trying to find the angle of the rotated junction $J(r, \theta + \theta^*)$, by matching it against all possible rotations of a filter $f(r, \theta)$. This is computationally feasible if f admits a steerable representation in the sense of Definition 3.1.

Definition 3.1 *A function f on the plane is steerable in the finite basis $\{\xi_\alpha\}$ if, for any rotation matrix \mathbf{R}_{θ_0} , we can find coefficients $\{c_\alpha(\theta_0)\}$ such that*

$$f_{\theta_0}(\mathbf{x}) = f(\mathbf{R}_{\theta_0}\mathbf{x}) = \sum_{\alpha} c_{\alpha}(\theta_0)\xi_{\alpha}(\mathbf{x}). \quad (3.4)$$

A family of functions $\{\xi_{\alpha}\}$ is steerable if its linear span is invariant under arbitrary rotations.

Using vector notation, where $\xi = (\xi_{\alpha})_{\alpha}$ denotes the column vector of basis functions, the last part of Definition 3.1 is equivalent to saying that for each θ_0 , a matrix $\mathbf{L}(\theta_0)$ exists such that

$$\xi(r, \theta + \theta_0) = \mathbf{L}(\theta_0)\xi(r, \theta). \quad (3.5)$$

An illustrative example of such a family is $\{\cos(\theta), \sin(\theta)\}$, whose rotations can be written as

$$\begin{pmatrix} \cos(\theta + \theta_0) \\ \sin(\theta + \theta_0) \end{pmatrix} = \begin{pmatrix} \cos(\theta_0) & -\sin(\theta_0) \\ \sin(\theta_0) & \cos(\theta_0) \end{pmatrix} \begin{pmatrix} \cos(\theta) \\ \sin(\theta) \end{pmatrix}, \quad (3.6)$$

which is a weighted sum of the unrotated functions.

For any arbitrary angle θ_0 , we find the response for the rotated filter $f_{\theta_0} = f(\mathbf{R}_{\theta_0}\cdot)$ using

$$\langle I_{\text{loc}}, f_{\theta_0} \rangle = \langle I_{\text{loc}}, \sum_{\alpha} c_{\alpha}(\theta_0)\xi_{\alpha} \rangle = \sum_{\alpha} c_{\alpha}(\theta_0)q_{\alpha} \quad (3.7)$$

with q_{α} as defined in (3.3). In practice, it means that, one needs to compute q_{α} of (3.3) only once. One can then deduce all desired responses for arbitrarily rotated filters through simple linear combinations.

We note that the Fourier transform of a steerable representation is also steerable. This follows from the observation that a rotation by θ_0 in space is equivalent to a rotation by θ_0 in the Fourier domain. By taking the Fourier transform of (3.4), we find that

$$\hat{f}(\omega, \varphi + \theta_0) = \sum_{\alpha} c_{\alpha}(\theta_0)\hat{\xi}_{\alpha}(\omega, \varphi). \quad (3.8)$$

Any estimator based on this framework is then a mapping E that takes the measurements q_α and returns an estimate $\tilde{\theta}$ of θ^* :

$$E : (q_\alpha) \mapsto \tilde{\theta}. \quad (3.9)$$

The goal is to maximize $\langle I_{\text{loc}}, f(\cdot + \theta_0) \rangle$ over θ_0 :

$$\tilde{\theta} = \underset{\theta_0 \in [0, 2\pi)}{\operatorname{argmax}} |\langle I_{\text{loc}}, f(\cdot + \theta_0) \rangle|. \quad (3.10)$$

3.2.3 Particular Examples of Steerable Filters

We illustrate the previously described ideas with particular examples.

First, we consider the Gaussian $g(r, \theta) = e^{-r^2/2}$, whose Fourier transform is computed as $\hat{g}(\omega, \varphi) = 2\pi e^{-\omega^2/2}$. In the Fourier domain, its second-order derivatives can be written as

$$\mathcal{F}\{\partial_1^2 g\} = -\omega^2 \hat{g}(\omega, \varphi) \cos^2 \varphi, \quad (3.11)$$

$$\mathcal{F}\{\partial_1 \partial_2 g\} = -\omega^2 \hat{g}(\omega, \varphi) \sin \varphi \cos \varphi, \quad (3.12)$$

$$\mathcal{F}\{\partial_2^2 g\} = -\omega^2 \hat{g}(\omega, \varphi) \sin^2 \varphi, \quad (3.13)$$

where we have used the identities $\omega_1 = \omega \cos \varphi$ and $\omega_2 = \omega \sin \varphi$. Here, ∂_1 and ∂_2 denotes the derivatives of a function with respect to the coordinates x_1 and x_2 . Expanding $\sin \varphi$ and $\cos \varphi$ in terms of $e^{j\varphi}$ and $e^{-j\varphi}$, all of the above can be written as linear combinations of the functions g_n , expressed in the Fourier domain as

$$\hat{g}_n(\omega, \varphi) := \omega^2 e^{-\frac{\omega^2}{2}} e^{jn\varphi}, \quad (3.14)$$

for $n = -2, -1, 0, 1, 2$. We note that, $-2\pi\omega^2 e^{-\omega^2/2}$ is the the Fourier transform of the Laplacian of Gaussian (LoG). All directional second-order derivatives of the Gaussian can be written as linear combinations of g_n .

From the above formula, it is apparent that, for an arbitrary rotation θ_0 ,

$$\hat{g}_n(\omega, \varphi + \theta_0) = e^{jn\theta_0} \hat{g}_n(\omega, \varphi), \quad (3.15)$$

and, using the equivalence of rotations in space and Fourier,

$$g_n(r, \theta + \theta_0) = e^{jn\theta_0} g_n(r, \theta). \quad (3.16)$$

Let f denote a filter corresponding to a directional second-derivative of the Gaussian. Based on above discussion, f can be expanded as

$$f(r, \theta) = \sum_{-2 \leq n \leq 2} c_n g_n(r, \theta). \quad (3.17)$$

The rotation of f by θ_0 is then given by

$$f(r, \theta + \theta_0) = \sum_n c_n g(r, \theta + \theta_0) = \sum_n e^{jn\theta_0} c_n g(r, \theta). \quad (3.18)$$

We use f to determine the orientation of a feature in I_{loc} , by finding its best match against all rotations of f . The response of the filter is computed as

$$\begin{aligned} \langle I_{\text{loc}}, f(\cdot, \cdot + \theta_0) \rangle &= \left\langle I_{\text{loc}}, \sum_n c_n e^{jn\theta_0} g(r, \theta) \right\rangle \\ &= \sum_n e^{jn\theta_0} c_n q_n. \end{aligned} \quad (3.19)$$

where $q_n = \langle I_{\text{loc}}, g_n \rangle$. This means that, once we have the response of the image to $g_n(r, \theta)$, its response to any rotation of the filter $f(r, \theta + \theta_0)$ can be determined analytically using (3.19). This makes it computationally feasible to find the best match of the image against all possible rotations of f .

Other simple examples of steerable filters are based on the gradient or the Hessian. Starting from an isotropic lowpass function $v(x_1, x_2)$, one can create a subspace of steerable derivative-based templates which can serve as basic edge or ridge detectors. In 2D, let $v_{k,l} = \partial_{x_1}^k \partial_{x_2}^l v$ be derivatives of the isotropic function v . By the chain rule of differentiation, for any rotation matrix \mathbf{R}_{θ_0} , the function $\partial_{x_1}^k \partial_{x_2}^l v(\mathbf{R}_{\theta_0} \cdot)$ can be written as a linear combination of $v_{i,j}$ with $i + j = k + l$. Therefore, any anisotropic filter of the form

$$f(x_1, x_2) = \sum_{m=1}^M \sum_{k+l=m} c_{k,l} v_{k,l}(x_1, x_2) \quad (3.20)$$

is steerable. Consequently, for any rotation matrix \mathbf{R}_{θ_0} , an application of the rotated filter to an image I yields

$$(I * f(\mathbf{R}_{\theta_0} \cdot))(x) = \sum_{m=1}^M \sum_{k+l=m} c_{k,l}(\theta) I_{k,l}(x), \quad (3.21)$$

where $I_{k,l} = I * \partial_{x_1}^k \partial_{x_2}^l v$ and $c_{k,l}(\theta)$ is a trigonometric polynomial in $\cos(\theta)$ and $\sin(\theta)$. Once every $I_{k,l}$ is precomputed, the linear combination (3.21) allows us to quickly evaluate the filtering of the image by the anisotropic filter rotated by any angle. We can then “steer” f by changing θ , typically to determine the direction along which the response is maximized and across which it is minimized.

In [101], Jacob and Unser improved the basic steerable filters by imposing Canny-like criteria of optimality [40] on the following properties of the detector: reasonable approximation of the ideal detector, maximum signal-to-noise ratio, good spatial localization, and reduced oscillations. Their formalism boils down to a constrained optimization of the expansion coefficients $c_{k,l}$ using Lagrange multipliers.

In [102], Mühlich *et al.* proposed a combined method for the detection and classification of polar separable patterns (including junctions). The detection phase relies on a classical structure tensor scheme, complemented by further operations to truncate the list of candidate points. The classification step is dedicated to the determination of the precise angular orientation of the “branches”. For such a purpose, multisteerable filters are designed and adjusted to the patterns of interest.

3.2.4 Polar Separable Measurement Functions

Due to the commutativity of rotations on the plane, the matrices $L(\theta_0)$ in (3.5) commute for different values of θ_0 . Consequently, when a diagonalization exists, they are simultaneously diagonalizable. What this means is that we can find a change of basis ξ in (3.5), such that in the new basis, the matrix $L(\theta_0)$ is diagonal for any θ_0 . We can further assume the basis functions to be normalized with $\|\xi_\alpha\|_2^2 = 1$.

In this case, each measurement function ξ_α can be steered independently by any angle θ_0 based on the relation

$$\xi_\alpha(r, \theta + \theta_0) \equiv \frac{c_\alpha(\theta_0)}{c_\alpha(0)} \xi_\alpha(r, \theta). \quad (3.22)$$

It can be shown that functions ξ_α that fulfil the above condition are always polar separable and have the form

$$\xi_\alpha(r, \theta) \equiv \eta_\alpha(r) e^{jn_\alpha \theta} \quad (3.23)$$

for some radial functions η_α and integers $n_\alpha \in \mathbb{Z}$. The integers n_α are called the *harmonics* (they determine the angular periodicity of the measurement function).

If $\xi_\alpha(r, \theta)$ is polar separable as in (3.23), then its Fourier transform is also polar separable as

$$\hat{\xi}_\alpha(\omega, \varphi) = \hat{h}_\alpha(\omega) e^{jn_\alpha \varphi} \quad (3.24)$$

for a function \hat{h}_α related to η_α and n_α . As we shall see in 3.3, ξ_α can also be interpreted as the n th order complex Riesz transform of $h(r)$. This observation leads to a comprehensive theory for the design and analysis of steerable filters and wavelets [98, 103–106].

Conversely, if the $\xi_\alpha(r, \theta)$ are normalized and polar separable as in (3.23), then any linear combination of them such as

$$f = \sum_\alpha c_\alpha(0) \xi_\alpha \quad (3.25)$$

is steerable:

$$\begin{aligned} f(r, \theta + \theta_0) &= \sum_\alpha c_\alpha(0) \eta_\alpha(r) e^{jn_\alpha \theta} e^{jn_\alpha \theta_0} \\ &= \sum_\alpha c_\alpha(\theta_0) \xi_\alpha(r, \theta). \end{aligned} \quad (3.26)$$

From (3.23), the form of the coefficients c_α in (3.4) and (3.26) is given by

$$c_\alpha(\theta_0) \equiv c_\alpha(0) e^{jn_\alpha \theta_0}. \quad (3.27)$$

The functions $\xi_\alpha(r, \theta)$ are orthogonal as long as all n_α 's are different. In this case,

$$\|f\|_2^2 = \sum_\alpha |c_\alpha(0)|^2 \|\xi_\alpha\|_2^2. \quad (3.28)$$

The measurements q_α of (3.3) can be written (in polar coordinates) as

$$\begin{aligned} q_\alpha &= \langle J(\cdot, \cdot + \theta^*), \xi_\alpha(\cdot, \cdot) \rangle + \langle S, \xi_\alpha \rangle \\ &= \langle J(\cdot, \cdot), \xi_\alpha(\cdot, \cdot - \theta^*) \rangle + \langle S, \xi_\alpha \rangle \\ &= \langle J, e^{-jn_\alpha \theta^*} \xi_\alpha \rangle + \langle S, \xi_\alpha \rangle \\ &= e^{-jn_\alpha \theta^*} u_\alpha + s_\alpha, \end{aligned} \quad (3.29)$$

where we have defined $u_\alpha = \langle J, \xi_\alpha \rangle$ and $s_\alpha = \langle S, \xi_\alpha \rangle$.

We make one final assumption regarding measurement functions ξ_α , namely that they have enough vanishing moments to stationarize and cancel the mean of the background signal S .

In cases where the measurement functions ξ_α , as defined in (3.24), all have distinct harmonics n_α , we can drop the index α , and instead, index them directly by their harmonics $n \in H$, where $H \subset \mathbb{Z}$ is the set of all used harmonics. Specifically, we write,

$$\xi_n(r, \theta) = \eta_n(r) e^{jn\theta}, \tag{3.30}$$

and in the Fourier domain,

$$\hat{\xi}_n(\omega, \varphi) = \hat{h}_n(\omega) e^{jn\varphi}. \tag{3.31}$$

We recall that for $n \neq m$, ξ_n and ξ_m are orthogonal, due to the orthogonality of their angular exponential factors.

In practical constructions, one often chooses the same radial pattern $\hat{h}(\omega)$ for all measurement functions [98], that is

$$\hat{\xi}_n(\omega, \varphi) = \hat{h}(\omega) e^{jn\varphi}. \tag{3.32}$$

3.3 Steerable Representations and the Riesz Transform

3.3.1 Definition and Properties of the Riesz Transform

The complex Riesz transform was introduced in the literature by Larkin *et al.* [107, 108] as a multidimensional extension of the Hilbert transform. The Hilbert transform is a 1D shift-invariant operator that maps all cosine functions into sine functions without affecting their amplitude (allpass filter). More precisely, the **Hilbert transform** of a function f , denoted by $\mathcal{H}\{f\}$, is

$$\mathcal{H}\{f\}(x) = \mathcal{F}^{-1} \left\{ -\frac{j\omega}{|\omega|} \hat{f} \right\}(x) = \mathcal{F}^{-1} \{-j \operatorname{sgn}(\omega) \hat{f}\}(x). \tag{3.33}$$

Similarly to the Hilbert transform, the **Riesz transform** of a function on the plane, denoted by $\mathcal{R}\{f\}$ is defined as

$$\mathcal{R}\{f\}(\mathbf{x}) = \mathcal{F}^{-1} \left\{ \frac{(\omega_x + j\omega_y)}{\|\boldsymbol{\omega}\|} \hat{f} \right\}(\mathbf{x}) = \mathcal{F}^{-1} \{e^{j\varphi} \hat{f}\}(\mathbf{x}). \tag{3.34}$$

The transform is a convolution-type operator that also acts as an allpass filter, with a phase response that is completely encoded in the angular part.

The Riesz transform is scale- and shift-invariant,

$$\forall \mathbf{x}_0 \in \mathbb{R}^d, \quad \mathcal{R}\{f(\cdot - \mathbf{x}_0)\}(\mathbf{x}) = \mathcal{R}f(\mathbf{x} - \mathbf{x}_0) \quad (3.35)$$

$$\forall a \in \mathbb{R}^+ \setminus \{0\}, \quad \mathcal{R}\left\{f\left(\frac{\cdot}{a}\right)\right\}(\mathbf{x}) = \mathcal{R}f\left(\frac{\mathbf{x}}{a}\right), \quad (3.36)$$

and provides a unitary mapping from $L_2(\mathbb{R}^2)$ to $L_2(\mathbb{R}^2)$. The Riesz transform commutes with spatial rotations, in the sense that its impulse response is steerable,

$$\mathcal{R}\{\delta\}(\mathbf{R}_{\theta_0}\mathbf{x}) = e^{j\theta_0}\mathcal{R}\{\delta\}(\mathbf{x}), \quad (3.37)$$

where $\mathbf{R}_{\theta_0} = \begin{bmatrix} \cos\theta_0 & -\sin\theta_0 \\ \sin\theta_0 & \cos\theta_0 \end{bmatrix}$ is the matrix that implements a 2D spatial rotation by the angle θ_0 .

The n th-order complex 2D Riesz transform \mathcal{R}^n represents the n -fold iterate of \mathcal{R} . It inherits the invariance properties of the Riesz transform since they are preserved through iteration. This means that we can also use the Riesz transform to map a set of primary wavelets into an augmented one while preserving the scale- and shift-invariant structure.

Comparing the above definition with (3.23) and (3.24), we observe that the polar-separable measurement functions introduced in the previous section can be interpreted as the n_α -th order Riesz transform of an isotropic function $\mathcal{F}^{-1}\{\hat{h}_\alpha\}$.

The relation between the Riesz transform and steerable filters is presented in [103]. The properties of steerable filters using low-order harmonics are analyzed in [104]. In [109], the authors rely on the Riesz transform to analyze predefined junction points in images.

3.3.2 Connection to the Gradient and Directional Derivatives

In this section, we describe the connection between the Riesz transform, the directional Hilbert transform, the gradient, and the directional derivatives. Assuming a function f with sufficient vanishing moments, the Riesz transform is related to the complex gradient operator as

$$\mathcal{R}f(x_1, x_2) = -j \left(\frac{\partial}{\partial x_1} + j \frac{\partial}{\partial x_2} \right) (-\Delta)^{-1/2} f(x_1, x_2). \quad (3.38)$$

Here, $(-\Delta)^{\gamma/2}, \gamma \in \mathbb{R}^+$ is the isotropic fractional differential operator of order γ , whose Fourier representation is $\|\omega\|^\gamma$. The corresponding fractional integrator of order γ is $(-\Delta)^{-\gamma/2}, \gamma/2 \in \mathbb{R}^+$, with Fourier representation $\|\omega\|^{-\gamma}$. The value $\gamma = 1$ is of special interest, providing the link between the Riesz transform and the complex gradient operator. The integral operator acts on all derivative components and has an isotropic smoothing effect, thus, the Riesz transform acts as the smoothed version of the image gradient.

Assuming a function f with sufficient vanishing moments, the higher-order Riesz transforms are related to the partial derivatives of f by

$$\mathcal{R}^n f(x_1, x_2) = (-\Delta)^{-\frac{n}{2}} \sum_{n_1=0}^n \binom{n}{n_1} (-j)^{n_1} \partial_{x_1}^{n_1} \partial_{x_2}^{n-n_1} f(x_1, x_2). \quad (3.39)$$

The fractional integrator acts as an isotropic lowpass filter whose smoothing strength increases with n . The Riesz transform captures the same directional information as derivatives. However, it has the advantage of being better conditioned since, unlike them, it does not amplify the high frequencies.

The directional Hilbert transform is the Hilbert transform along a direction \mathbf{v} . It is related to the Riesz transform by

$$\mathcal{H}_{\mathbf{v}_\theta} f(\mathbf{x}) = \cos \theta f_1(\mathbf{x}) + \sin \theta f_2(\mathbf{x}), \quad (3.40)$$

where $f_1 = \text{Re}(\mathcal{R}f)$ and $f_2 = \text{Im}(\mathcal{R}f)$ are the real and imaginary parts of $\mathcal{R}f$. Assuming again a zero-mean function f , the directional Hilbert transform is related to the derivative in the direction \mathbf{v} by

$$\mathcal{H}_{\mathbf{v}} f(\mathbf{x}) = -(-\Delta)^{-\frac{1}{2}} D_{\mathbf{v}} f(\mathbf{x}). \quad (3.41)$$

Here, the operator $D_{\mathbf{v}}$ is the one defined in (2.1). This result corresponds to the interpretation that the Hilbert transform acts as a lowpass-filtered version of the derivative operator. The n -fold version of the directional Hilbert transform acting on a function f with sufficient vanishing moments along the direction specified by \mathbf{v} can be expressed in term of the partial derivatives of f as

$$\mathcal{H}_{\mathbf{v}}^n f(\mathbf{x}) = (-1)^n (-\Delta)^{-\frac{n}{2}} D_{\mathbf{v}}^n f(\mathbf{x}). \quad (3.42)$$

3.4 Isotropic Steerable Wavelet Frames

In [110], Simoncelli and Freeman proposed a new take on steerable filters: the steerable pyramid. The goal of their design was to combine steerability with a multiscale detection scheme. This pioneering work had many successful applications: contour detection [95], image filtering and denoising [111], orientation analysis [96], and texture analysis and synthesis [112]. In [113] multiscale steerable filters were used in the detection of stellate distortions in mammograms.

Typically, multiresolution directional techniques are motivated by their invariance with respect to primary geometric transformations: translation, dilation, and rotation. In this section, we present a unifying framework for steerable wavelet frames. Their construction is based on the application of the Riesz transform on isotropic wavelet profiles. Hence, the proposed method is translation invariant and multiscale. The rotation invariance is achieved by the Riesz transform that also gives a connection to gradient-like signal analysis [114], as presented in Section 3.3.2. The constructed scheme is flexible in a sense that it is possible to shape the profile of the wavelets. Another advantage is that the method has a natural extension to higher dimensions.

An application of steerable Riesz wavelets for texture learning is presented in [115]. Applications of steerable wavelets on the analysis of CT images are presented in [116, 117]. The extension of the steerable wavelet design based on the Riesz transform for higher dimensions, along with potential biomedical applications, are presented in [35]. Spherical harmonics, which are the 3D counterparts of circular harmonics, have also been used to represent and detect features and shapes in 3D [118], [119]. There, the authors also aim to identify symmetric structures; however, they search for symmetries in an entire 3D shape, not locally.

3.4.1 Radial Wavelet Profiles

Steerable wavelet frames are adapted to capture the local orientation of features (*i.e.*, junctions) within a multiresolution hierarchy. To simplify the notations, we consider wavelets that are centered at the origin.

Proposition 3.1 provides sufficient conditions on an isotropic profile that is used to generate a desired wavelet system.

Proposition 3.1 (cf. [98, Proposition 4.1.]) *Let $\hat{h} : [0, \infty) \rightarrow \mathbb{R}$ be a radial frequency profile satisfying:*

1. $\hat{h}(\omega) = 0$ for $\omega > \pi$
2. $\sum_{i \in \mathbb{Z}} \left| \hat{h}(2^i \omega) \right|^2 = 1$
3. $\left. \frac{d^n \hat{h}}{d\omega^n} \right|_{\omega=0} = 0$ for $n = 0, \dots, N$.

The isotropic mother wavelet ϕ whose Fourier transform is given by

$$\widehat{\phi}(\boldsymbol{\omega}) = \hat{h}(\|\boldsymbol{\omega}\|) \tag{3.43}$$

generates a normalized tight wavelet frame of $L_2(\mathbb{R}^2)$ whose basis functions

$$\phi_{i,\mathbf{k}}(\mathbf{x}) = 2^{-i} \phi\left(\frac{\mathbf{x}}{2^i} - \mathbf{k}\right) \tag{3.44}$$

have vanishing moments up to order N . In particular, any $f \in L_2(\mathbb{R}^d)$ can be represented as

$$f = \sum_{i \in \mathbb{Z}} \sum_{\mathbf{k} \in \mathbb{Z}^d} \langle f, \phi_{i,\mathbf{k}} \rangle \phi_{i,\mathbf{k}}. \tag{3.45}$$

There is a variety of isotropic profiles satisfying Proposition 3.1, typical examples can be found in Table 3.1. For further details, we refer to [120].

We note that, while (3.45) is reminiscent of a decomposition in an orthonormal basis, the main difference is that the basis functions are redundant. The key property that ensures that the transform is self-reversible (tight frame) is assumption 2 in Proposition 3.1.

3.4.2 Circular Harmonic Wavelets

In this section, corresponding to (3.24), we consider applying the multiorder complex Riesz transform to primal isotropic functions that satisfies Proposition 3.1. Compared to (3.3), here the measurement functions and measurements are indexed by the pair $\alpha = (n, i)$ of the harmonic $n \in H$ and the scale i . The wavelet schemes generated in such

Table 3.1: Radial frequency responses of isotropic bandlimited wavelets

Wavelet type	Mother wavelet function
Shannon	$\hat{h}(\omega) = \begin{cases} 1, & \frac{\pi}{2} \leq \omega \leq \pi \\ 0, & \text{otherwise} \end{cases}$
Simoncelli [112]	$\hat{h}(\omega) = \begin{cases} \cos\left(\frac{\pi}{2} \log_2\left(\frac{2\omega}{\pi}\right)\right), & \frac{\pi}{4} < \omega \leq \pi \\ 0, & \text{otherwise} \end{cases}$
Meyer ¹ [51]	$\hat{h}(\omega) = \begin{cases} \sin\left(\frac{\pi}{2} v\left(\frac{4\omega}{\pi} - 1\right)\right), & \frac{\pi}{4} < \omega \leq \frac{\pi}{2} \\ \cos\left(\frac{\pi}{2} v\left(\frac{2\omega}{\pi} - 1\right)\right), & \frac{\pi}{2} < \omega \leq \pi \\ 0, & \text{otherwise} \end{cases}$
Papadakis [124]	$\hat{h}(\omega) = \begin{cases} \sqrt{\frac{1+\sin(5\omega)}{2}}, & \frac{3\pi}{10} < \omega \leq \frac{\pi}{2} \\ 1, & \frac{\pi}{2} < \omega \leq \frac{3\pi}{5} \\ \sqrt{\frac{1-\sin(\frac{5\omega}{2})}{2}}, & \frac{3\pi}{5} < \omega \leq \pi \\ 0, & \text{otherwise} \end{cases}$
Held ² [114]	$\hat{h}(\omega) = \begin{cases} \cos\left(2\pi q\left(\frac{\omega}{2\pi}\right)\right), & \frac{\pi}{4} < \omega \leq \frac{\pi}{2} \\ \sin\left(2\pi q\left(\frac{\omega}{4\pi}\right)\right), & \frac{\pi}{2} < \omega \leq \pi \\ 0, & \text{otherwise} \end{cases}$

a way are often referred to as circular harmonic wavelets [121, 122]. We note that our circular harmonic wavelets are similar to the ones of Jacovitti and Neri [123], with the difference that the latter ones are non-tight.

Using the n th-order Riesz transform, we define the new function $\xi_n := \mathcal{R}^n \phi$, whose translates and dilates at location \mathbf{x}_0 are given by

$$\begin{aligned}
\xi_{n,i}(\mathbf{x} - \mathbf{x}_0) &= 2^{-i} \xi_n\left(\frac{\mathbf{x} - \mathbf{x}_0}{2^i}\right) \\
&= 2^{-i} \mathcal{R}^n \{\phi\}\left(\frac{\mathbf{x} - \mathbf{x}_0}{2^i}\right) \\
&= \mathcal{R}^n \left\{ 2^{-i} \phi\left(\frac{-\mathbf{x}_0}{2^i}\right) \right\}(\mathbf{x}) = \mathcal{R}^n \{\phi_i(\cdot - \mathbf{x}_0)\}(\mathbf{x}), \tag{3.46}
\end{aligned}$$

²For the Meyer wavelet of order N , the auxiliary function v is the polynomial of degree $2N + 1$ such that $v(t) = 0$, if $t \leq 0$, $v(t) = 0$, if $t \geq 1$, and $v(t) + v(1 - t) = 1$, while $v \in C^N([0, 1])$. The auxiliary function that achieves the frequency response with $N = 3$ continuous derivatives is $v(t) = t^4(35 - 84t + 70t^2 - 20t^3)$.

²For the Held wavelet of order N , the auxiliary function q is a polynomial of degree N .

due to the translation- and scale-invariance properties of the transform. More precisely, the circular harmonic wavelet at scale i , location $\mathbf{0}$, and harmonic channel n has the Fourier transform

$$\hat{\xi}_{n,i}(\omega, \varphi) = 2^i \hat{h}(2^i \omega) e^{jn\varphi}. \quad (3.47)$$

Consequently, in the spatial domain, we have

$$\xi_n(r, \theta) = \frac{j^n e^{jn\theta}}{2\pi} \int_0^\infty \omega \hat{h}(\omega) J_n(r\omega) d\omega, \quad (3.48)$$

where J_n denotes the n -th Bessel function of the first kind. We note that

$$\xi_{n,i}(r, \theta) = 2^{-i} \xi_{n,0}(2^{-i} r, \theta). \quad (3.49)$$

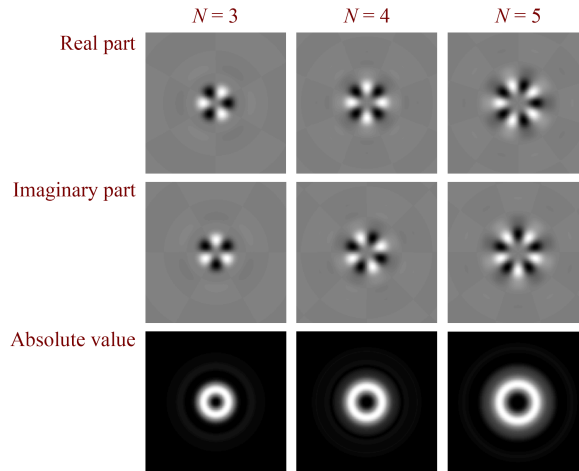
We illustrate the circular harmonic wavelets in Figure 3.1. As can be observed, in the Fourier domain, the support of the wavelets does not depend on n . By contrast, this is no longer so in the space domain, because the spread is determined by J_n . hence, it grows with n . We also note that, the n th-order harmonic wavelet has a rotational symmetry of order n around its center, corresponding to the n th-order rotational symmetry of $e^{jn\varphi}$.

We note that the application of the Riesz transform on $\xi_i(\cdot - \mathbf{x}_0)$ preserves its tight-frame property. Thus, by choosing N distinct values for the integer n (distinct set of harmonics), we form a tight frame of steerable wavelets. Thus, any finite-energy function f can be decomposed as

$$f = \sum_{n,i,\mathbf{k}} \langle f, \xi_{n,i,\mathbf{k}} \rangle \xi_{n,i,\mathbf{k}}, \quad (3.50)$$

where $2^i \mathbf{k}$ is the shift.

Circular Harmonic Wavelets in the Space Domain



Circular Harmonic Wavelets in the Fourier Domain




Figure 3.1: Circular harmonic wavelets at scale 1. From left to right: N th-order harmonic wavelet, for $N = 3$, $N = 4$, and $N = 5$. From top to bottom: real part, imaginary part, and absolute value, visualized in the space and Fourier domains, respectively.

3.5 Steerable Wavelets With Angular Selectivity

Though the circular harmonic wavelets are steerable, they have no directionality. To add directionality to such wavelet schemes, one can generate an orthogonal shaping matrix that accounts for the local directional information.

The new steerable frame functions $\psi_{n,i,\mathbf{k}}$ are obtained by a mapping with the orthogonal shaping matrix \mathbf{U} that is such that

$$\begin{bmatrix} \psi_{n_1,i,\mathbf{k}} \\ \vdots \\ \psi_{n_N,i,\mathbf{k}} \end{bmatrix} = \mathbf{U} \begin{bmatrix} \xi_{n_1,i,\mathbf{k}} \\ \vdots \\ \xi_{n_N,i,\mathbf{k}} \end{bmatrix}. \quad (3.51)$$

The new wavelets span the same space as the wavelet frame $\{\xi_{n,i,\mathbf{k}}\}$. The shaping matrix \mathbf{U} endows the wavelet functions $\psi_{n,i,\mathbf{k}}$ with a desired angular profile.

3.5.1 Monogenic Wavelets

An interesting instance of this design are monogenic wavelets [114, 125–127], which combine the idea of multiscale transforms with the idea of the monogenic signal [103]. We obtain the three-component monogenic wavelet transform by the shaping matrix

$$\mathbf{U}_{\text{mono}} = \frac{1}{\sqrt{2}} \begin{bmatrix} 0 & 1 & 0 \\ -\frac{j}{2} & 0 & -\frac{j}{2} \\ \frac{1}{2} & 0 & -\frac{1}{2} \end{bmatrix}.$$

The monogenic wavelet transform gives access to the local amplitude and the local phase of the image, additionally to the information on the local orientation. Monogenic wavelets have been used, among other things, for the enhancement of anisotropic structures in fluorescence microscopy and multiscale edge detection [35]. The orientation selectivity can be increased by using monogenic curvelets [128] or quasi-monogenic shearlets [129], which are useful for the description of optical filters [130].

Another direction of research related to the monogenic signal is the recovery of directional phase information using complex-valued wavelet transforms such as the monogenic wavelets [103, 114, 125, 127–131]. Preliminary applications include equalization of brightness [114], detection of salient points [132], enhancement of anisotropic

structures in fluorescence microscopy [35], and texture segmentation [133]. Image-analysis tools based on monogenic wavelets are provided by the ImageJ/Fiji plugins Monogenic³ and Monogenic Wavelet Toolbox⁴.

³Available at <http://bigwww.epfl.ch/demo/monogenic/>

⁴Available at <http://www-m6.ma.tum.de/Mamebia/MonogenicWaveletToolbox/>

Chapter 4

Cramér-Rao Lower Bounds for Steerable Detectors

4.1 Overview

Our aim in this chapter is to derive the fundamental Cramér-Rao bounds on the error of estimating local orientations using steerable filters, including multiscale estimates based on steerable wavelet frames. Our framework relies on the model we have built up in Chapter 3.

The Cramér-Rao Lower Bound (CRLB) sets a fundamental limit on the accuracy of unbiased estimators of a parameter from observations. The localization and detection properties of junction detectors have been extensively examined (*e.g.*, in [134]); however their accuracy in orientation estimation has not been fully described yet. In this chapter, we derive the CRLB for estimating the orientation of local features of an image. We study a formulation based on steerable filters that allows us to select the best harmonics for the design of steerable detectors, and identify the optimal radial profile. We also discuss the CRLB for different detector schemes, including multiscale approaches, in particular, wavelets¹.

¹The results in this chapter are based on the manuscript [135]: Z. Püspöki, A. Amini, J.P. Ward, M. Unser, “Cramér-Rao Lower Bounds for the Detection of Steerable Features”, submitted, 2016.

4.2 The Estimation Problem at a Glance

4.2.1 Measurement Framework

We formulate the estimation problem based on our junction model developed in Section 3.2.1. First, we recall the definition of our measurement framework, and then we investigate the elements of the model.

Using the notations of Chapter 3, we denote the general shape of a junction of interest with J , and its unknown rotation with $J_{\theta^*}(\mathbf{x}) = J(\mathbf{R}_{\theta^*}\mathbf{x})$. The estimation problem is to identify the unknown orientation θ^* .

Our local image is defined in polar coordinates, corresponding to (3.2), as

$$I_{\text{loc}}(r, \theta) = J(r, \theta + \theta^*) + S(r, \theta), \quad (4.1)$$

where S denotes the background signal which, for our purposes is considered to be noise.

The observations (measurements) used in estimating θ^* take the form (3.3)

$$q_\alpha = \langle I_{\text{loc}}, \xi_\alpha \rangle = \langle J_{\theta^*}, \xi_\alpha \rangle + \langle S, \xi_\alpha \rangle, \quad (4.2)$$

where $\{\xi_\alpha\}$ is a set of steerable measurement functions. The first term in (4.2) is deterministic, while the second (noise) term is random.

We are focusing on measurement functions of the form

$$\xi_\alpha(r, \theta) = \eta_\alpha(r) e^{jn_\alpha\theta} \quad (4.3)$$

for some radial function η_α and harmonics $n_\alpha \in \mathbb{Z}$.

The problem studied here is to estimate the local orientation angle θ^* from the steerable measurements q_α . We develop our theory in general for steerable detectors. Then, we provide more specialized results on multiscale schemes, in particular, on wavelets.

In practice, discretization affects multiscale steerable analysis in two ways: It limits the practical range of scales and limits the range of harmonics due to angular aliasing. However, as we shall see in the next sections, one can achieve good estimation by selecting an appropriate template with a small amount of polar-separable low order harmonics.

4.2.2 Background Noise Model

The background S is modeled as a realization of a statistically isotropic Gaussian self-similar field with stationary increments. The motivation behind this choice is that the power spectrum of many natural images is isotropic with a $1/|\omega|^\lambda$ type of decay [2, 3], which is consistent with long-range dependencies. Also, this model of background signal fits fluorescence microscopy images well [4], which is relevant to many practical applications of orientation-estimation.

For our background signal S to fulfill the requirements of self-similarity, long-range dependence, and inverse power-law spectrum, we define it as the (non-stationary) solution of a fractional stochastic differential equation. The general theory of such models and their non-Gaussian and sparse extensions is covered in [5]. The most important definitions are given in the Nomenclature.

We assume that our background signal corresponds to a stochastic process S in \mathbb{R}^2 that is defined as the solution of the stochastic equation

$$(-\Delta)^{\frac{\gamma}{2}} S = w, \quad (4.4)$$

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 $\|\omega\|^\gamma$. The intuitive idea here is to shape or “color” the white noise by an appropriately defined inverse fractional Laplacian, which would give the solution an inverse power-law spectrum and the desired (wide sense) self-similarity and long-range dependencies.

The fractional Laplacian operator in (4.4) is inverted in a weak sense; in other words, in the sense of generalized random fields [1]. Formally, for test functions ζ we write

$$\langle S, \zeta \rangle = \left\langle w, I_{\gamma,2}^* \zeta \right\rangle, \quad (4.5)$$

where $I_{\gamma,2}^* = (-\Delta)^{-\frac{\gamma}{2}}$ is the corrected scale-invariant inverse operator of $(-\Delta)^{\frac{\gamma}{2}}$, defined from \mathcal{S} (Schwartz space) to L_2 with an appropriate correction for the singularity at $\omega = \mathbf{0}$. The notation for the inverse operator and the index 2 are related to a special case from a family of operators introduced in [5]. For further details on $I_{\gamma,2}^*$ we refer to the Nomenclature. In the case where w is Gaussian, the field we obtain is the isotropic 2D generalization of a fractional Brownian motion. In the special case of functions ζ with sufficient vanishing moments, which is what we actually need here, $I_{\gamma,2}^*$ acts simply like multiplication by $\|\omega\|^{-\gamma}$ in Fourier.

The correlation form of an isotropic self-similar field as defined in the Nomenclature can be computed as

$$\begin{aligned} \mathcal{B}_S\{f, g\} &= \mathcal{B}_w \left\{ I_{\gamma,2}^* f, I_{\gamma,2}^* g \right\} \\ &= \frac{\sigma_0^2}{(2\pi)^2} \int_0^{2\pi} \int_0^\infty \frac{\widehat{f}(\omega, \varphi) \widehat{g}(\omega, \varphi)}{\omega^{2\gamma}} \omega d\omega d\varphi \end{aligned} \tag{4.6}$$

for functions f and g with sufficiently many vanishing moments. Here, σ_0^2 is the variance of the white noise w .

The above formula can be compared with the correlation form of an isotropic stationary random field X with radial power spectrum $P_X(\omega)$:

$$\mathcal{B}_X\{f, g\} = \frac{1}{(2\pi)^2} \int_0^{2\pi} \int_0^\infty \widehat{f}(\omega, \varphi) \widehat{g}(\omega, \varphi) P_X(\omega) \omega d\omega d\varphi. \tag{4.7}$$

It follows that, in our model, the singular distribution $\omega^{-2\gamma}$ can be interpreted as the generalized power spectrum of our background signal.

4.2.3 Reference Class of Estimators

In this section, we introduce a family of local orientation estimators based on steerable functions that we use for the experimental verification of our results. The performance bounds we derive also apply to estimators outside this family, which use the same steerable measurements.

The optimization problem that we solve for the estimation is similar to the formulation we introduce in Chapter 6. Specifically, the estimator works by steering a pre-defined pattern to find the maximal response (*i.e.*, best match), against the junction of interest. Let us represent our steerable detector by $f = \sum_\alpha c_\alpha \xi_\alpha$, where the vector $\mathbf{c} = (c_\alpha)$ parametrizes the family of estimators. The coefficients c_α may be obtained by optimizing the shape of the detector to match certain classes of patterns, as in Chapter 6 or, in the simplest case, by projecting a single directional pattern J on the measurement functions ξ_α . We then want to maximize $\langle I_{loc}, f(\cdot + \theta_0) \rangle$ over θ_0 . This means that

we define the estimate as

$$\begin{aligned}\tilde{\theta} &= \arg \max_{\theta_0 \in [0, 2\pi)} \left| \left\langle I_{\text{loc}}, \sum_{\alpha} e^{jn_{\alpha}\theta_0} c_{\alpha} \xi_{\alpha} \right\rangle \right|^2 \\ &= \arg \max_{\theta_0 \in [0, 2\pi)} \left| \sum_{\alpha} q_{\alpha} c_{\alpha} e^{jn_{\alpha}\theta_0} \right|^2,\end{aligned}\quad (4.8)$$

noting that the inner product is conjugate-linear in the first argument.

4.2.4 Cramér-Rao Lower Bound

The CRLB for estimating the parameter θ^* given a vector of measurements \mathbf{q} is defined as

$$\text{CRLB} = \frac{1}{\text{FI}(\theta^*)}, \quad (4.9)$$

where FI denotes the Fisher information. It is calculated from the conditional probability of the observations given θ^* .

Based on (3.29), we recall that the measurements q_{α} of (4.2) are expressed (in polar coordinates) as

$$q_{\alpha} = e^{-jn_{\alpha}\theta^*} u_{\alpha} + s_{\alpha}, \quad (4.10)$$

where $u_{\alpha} := \langle J, \xi_{\alpha} \rangle$ and $s_{\alpha} := \langle S, \xi_{\alpha} \rangle$.

Considering our measurement vector $\mathbf{q} = (q_{\alpha})$ in (4.2), we say that the first term (junction model) is general and deterministic. It depends only on u_{α} , plus the θ^* angle that we are estimating. The second term is random and depends on the properties of the self-similar field defined in (4.4).

Since the real and imaginary parts of $\mathbf{q} = (q_{\alpha})$ are jointly Gaussian random variables, $\mathbf{q} = q_{\alpha}$ is a complex Gaussian random vector. In general, a complex Gaussian vector \mathbf{z} can be characterized by three parameters: the mean vector $\boldsymbol{\mu}$; the covariance matrix \mathbf{C} with entries $\mathbb{E}\{\bar{z}_i z_j\}$; and the pseudo-covariance matrix \mathbf{P} with entries $\mathbb{E}\{z_i z_j\}$. If the pseudo-covariance matrix is zero, following the terminology of [136], we call the complex Gaussian distribution *circular*.

We will later show that, due to the structure of our problem, the CRLB always depends on a circular set of measurements. This property is favorable for us, because it

simplifies the calculation of the Fisher information. For circular complex Gaussian variables, this is given by the following version of the Slepian-Bangs formula [136, B.3.25]:

$$\text{FI}(\theta^*) = 2 \operatorname{Re} \left(\frac{d\boldsymbol{\mu}^H}{d\theta^*} \mathbf{C}^{-1} \frac{d\boldsymbol{\mu}}{d\theta^*} + \operatorname{Tr} \left(\mathbf{C}^{-1} \frac{d\mathbf{C}}{d\theta^*} \mathbf{C}^{-1} \frac{d\mathbf{C}}{d\theta^*} \right) \right). \quad (4.11)$$

There, $\boldsymbol{\mu}$ denotes the mean and \mathbf{C} the covariance matrix of the (Gaussian) observation vector \mathbf{q} .

Due to the vanishing moments of the measurement functions, solely the contribution to the mean of the observation vector comes from the deterministic terms u_α . Also, from (4.2), the random part does not depend on θ^* , meaning the second term of (4.11) vanishes.

In the following sections, we derive the CRLB for single- and multiscale orientation estimators based on steerable functions and study their behavior. We address the question of how to achieve the best bounds given a finite number of measurements. This is done by carefully choosing the harmonics and the radial profile of the measurement functions ξ_α .

4.3 Cramér-Rao Lower Bound for an Estimation with Distinct Harmonics

4.3.1 Finite Number of Distinct Harmonics

In this section, we derive the CRLB in cases where the measurement functions defined in (4.3) have distinct harmonics n_α . In those cases, we can drop the index α . Instead, we index the measurement functions directly by their harmonics $n \in H$, where $H \subset \mathbb{Z}$ is the set of all used harmonics. Specifically, we write that

$$\xi_n(r, \theta) = \eta_n(r) e^{jn\theta} \quad (4.12)$$

and, in the Fourier domain,

$$\hat{\xi}_n(\omega, \varphi) = \hat{h}_n(\omega) e^{jn\varphi}. \quad (4.13)$$

We recall that for $n \neq m$, ξ_n and ξ_m are orthogonal, due to the orthogonality of their angular exponential factors. Note also that $\overline{\xi_n} = \xi_{-n}$. This in particular implies that for a real image I_{loc} ,

$$q_{-n} = \langle I_{\text{loc}}, \xi_{-n} \rangle = \langle I_{\text{loc}}, \overline{\xi_n} \rangle = \overline{\langle I_{\text{loc}}, \xi_n \rangle} = \overline{q_n}. \quad (4.14)$$

Thus, q_{-n} and q_n essentially carry the same information, so that the CRLB based on q_n , $n \in H$, is the same as the CRLB based on q_n with harmonics n in the set

$$H_+ = \{|n| : n \in H\}. \quad (4.15)$$

We further exclude $n = 0$ from consideration, since the corresponding measurement does not depend on the rotation angle θ^* .

We can now calculate the CRLB for estimating the angle θ^* of the junction in (4.1):

Theorem 4.1 *For measurements using distinct harmonics $n \in H_+$, the exact form of the Fisher information is*

$$FI(\theta^*) = \frac{4\pi}{\sigma_0^2} \sum_{n \in H_+} \frac{n^2 |u_n|^2}{\int_0^\infty \omega^{1-2\gamma} |\hat{h}_n(\omega)|^2 d\omega} \quad (4.16)$$

and the CRLB for the estimation problem is

$$CRLB = \frac{\sigma_0^2 / 4\pi}{\sum_{n \in H_+} \frac{n^2 |u_n|^2}{\int_0^\infty \omega^{1-2\gamma} |\hat{h}_n(\omega)|^2 d\omega}}. \quad (4.17)$$

The proof of Theorem 4.1 is given in Section 4.7.

We note that, given its vanishing moments and finite energy, ξ_n is essentially band-pass. Consequently, the integral

$$K_n = \int_0^\infty \omega^{1-2\gamma} |\hat{h}_n(\omega)|^2 d\omega, \quad (4.18)$$

with its homogeneous kernel, which measures the power of background signal seen by ξ_n , is effectively an indication of the scale of ξ_n . This is due to the self-similar nature of the background signal. To make this dependence on scale explicit, take an arbitrary scaling of ξ_n by a parameter a (with proper normalization): $\xi_n(\cdot/a)/a$. The corresponding integral will be

$$\int_0^\infty \omega^{1-2\gamma} |a \hat{h}_n(a\omega)|^2 d\omega = a^{2\gamma} K_n. \quad (4.19)$$

We are going to account for the effect of scaling on the CRLB later, when we look at multiscale estimators. For now, we focus on single-scale estimators, where one often chooses the same radial pattern \hat{h} for all measurement functions [98], which results in

$$\hat{\xi}_n(\omega, \varphi) = \hat{h}(\omega) e^{in\varphi}. \quad (4.20)$$

4.3.2 Estimation from Best N Measurements

The next question we address is, if we can only use a finite number N of harmonics (*i.e.*, a finite number of measurements), which ones to take to have the lowest CRLB. Here, as just noted, we fix the scale-dependent part of ξ_n (the radial frequency pattern \hat{h}) and choose the set H_+ of harmonics to have cardinality N .

We have the following immediate corollary of Theorem 4.1.

Corollary 4.1 *The exact form of the Fisher information in the case of a fixed scale is*

$$FI(\theta^*) = \frac{4\pi/\sigma_0^2}{\int_0^\infty \omega^{1-2\gamma} |\hat{h}(\omega)|^2 d\omega} \sum_{n \in H_+} n^2 |u_n|^2 \quad (4.21)$$

and the CRLB for the estimation problem is

$$CRLB = \frac{1}{\sum_{n \in H_+} n^2 |u_n|^2} \frac{\sigma_0^2}{4\pi} \int_0^\infty \omega^{1-2\gamma} |\hat{h}(\omega)|^2 d\omega. \quad (4.22)$$

Corollary 4.1 provides the CRLB for the estimation of the local orientation of features by steerable detectors at a particular scale.

Considering (4.22), we see that choosing the best measurements is equivalent to identifying the set $H_+ \subset \mathbb{Z}_+$ with N members, for which the sum

$$\sum_{n \in H_+} n^2 |u_n|^2 \quad (4.23)$$

is maximized.

We rearrange u_n in the Fourier domain as

$$\begin{aligned} u_n &= \langle J, \xi_n \rangle = \frac{1}{(2\pi)^2} \langle \hat{J}, \hat{\xi}_n \rangle \\ &= \frac{1}{(2\pi)^2} \int_0^{2\pi} \int_0^\infty \overline{\hat{J}(\omega, \varphi)} \hat{h}(\omega) e^{jn\varphi} \omega d\omega d\varphi \\ &= \frac{1}{2\pi} \int_0^{2\pi} G(\varphi) e^{jn\varphi} d\varphi, \end{aligned} \quad (4.24)$$

where we have defined

$$G(\varphi) = \frac{1}{2\pi} \int_0^\infty \overline{\hat{J}(\omega, \varphi)} \hat{h}(\omega) \omega d\omega. \quad (4.25)$$

It is seen that u_n is the n th Fourier-series coefficient of $G(\varphi)$.

For a continuous and sufficiently differentiable G , we invoke integration by parts to write

$$\begin{aligned} \frac{1}{2\pi} \int_0^{2\pi} \frac{d}{d\varphi} G(\varphi) e^{jn\varphi} d\varphi &= -\frac{1}{2\pi} \int_0^{2\pi} G(\varphi) \frac{d}{d\varphi} e^{jn\varphi} d\varphi \\ &= -\frac{jn}{2\pi} \int_0^{2\pi} G(\varphi) e^{jn\varphi} d\varphi \\ &= -jnu_n. \end{aligned} \tag{4.26}$$

This shows that $-jnu_n$ are the Fourier coefficients of $\frac{d}{d\varphi} G(\varphi)$. The choice of harmonics n that minimizes the CRLB is therefore given by the N largest Fourier coefficients of $\frac{d}{d\varphi} G(\varphi)$, excluding the Hermitian symmetric half with $n < 0$. If \hat{J} is sufficiently smooth, this derivative can be calculated from the partial derivative of \hat{J} as

$$\begin{aligned} \frac{d}{d\varphi} G(\varphi) &= \frac{d}{d\varphi} \int_0^\infty \overline{\hat{J}(\omega, \varphi)} \hat{h}(\omega) \omega d\omega \\ &= \int_0^\infty \overline{\partial_\varphi \hat{J}(\omega, \varphi)} \hat{h}(\omega) \omega d\omega. \end{aligned} \tag{4.27}$$

This result also covers the specific case of having just one measurement at our disposal.

4.3.3 Asymptotic Behavior in a Single Scale

Similarly to Section 4.3.2, we fix the scale-dependent part of ξ_n (the radial frequency pattern \hat{h}) and look at the asymptotic behavior in terms of the number of harmonics, by choosing the set H_+ of harmonics as $\{1, \dots, N\}$, and letting N tend to infinity. Based on (4.22), we conclude that the asymptotic behavior of the CRLB depends on the asymptotic (decay) properties of the coefficients $u_n = \langle J, \xi_n \rangle$ of the directional pattern (junction).

The question is whether the CRLB vanishes asymptotically as $N \rightarrow \infty$, which would suggest the theoretical possibility of perfect estimation with infinitely many measurements. This can only happen if the series $\sum_n n^2 |u_n|^2$ diverges. We now study this question.

We note that, the junction $J \in L_1 \cap L_2$ is a finite image with finite energy. Since the

measurement functions $\{\xi_n\}_{n \in \mathbb{Z}}$ are orthonormal, by the Bessel inequality, we have that

$$\sum_{n \in \mathbb{Z}} |u_n|^2 \leq \|J\|_2^2 < \infty \Rightarrow \{u_n\}_{n \in \mathbb{Z}} \in l_2, \tag{4.28}$$

so the series without the n^2 factors converges.

Next, note that by integration by parts, the equality

$$\langle \partial_\theta J, \xi_n \rangle = -\langle J, \partial_\theta \xi_n \rangle \tag{4.29}$$

holds for continuously differentiable J, ξ_n , which we can then extend by duality to any J . If $\partial_\theta J \in L_2$, we then have

$$\begin{aligned} \|\partial_\theta J\|_2^2 &\geq \sum_{n \in \mathbb{Z}} |\langle \partial_\theta J, \xi_n \rangle|_2^2 = \sum_{n \in \mathbb{Z}} |\langle J, \partial_\theta \xi_n \rangle|_2^2 \\ &= \sum_{n \in \mathbb{Z}} n^2 |\langle J, \xi_n \rangle|_2^2 = \sum_{n \in \mathbb{Z}} n^2 |u_n|_2^2 \\ &= 2 \sum_{n \geq 0} n^2 |u_n|_2^2, \end{aligned} \tag{4.30}$$

using $u_{-n} = \bar{u}_n$ for real J . So in this case the series converges and, consequently, the CRLB does not vanish, even for infinitely many measurements. A similar result can also be obtained in the Fourier domain. This is due to

$$\partial_\varphi \mathcal{F}\{J\}(\omega, \varphi) = \mathcal{F}\{\partial_\theta J\}(\omega, \varphi) \tag{4.31}$$

when J is sufficiently smooth. To see this, differentiate under the integral sign on the left and use integration by parts on the right. Or, observe that rotations in space correspond to the same rotations in Fourier and then use the limit definition of the derivative.

We now examine the conditions for the series to diverge and the CRLB to go to zero asymptotically. We recall that u_n are the Fourier-series coefficients of $G(\varphi)$ in (4.24). The convergence of the series in (4.22) therefore depends on the decay properties of the Fourier-series of $G(\varphi)$, which in turn is related to its differentiability. In particular, if \hat{J} has angular jump discontinuities that are inherited by G , u_n will decay slowly like $1/n$, and the series will diverge. In this case, the CRLB will asymptotically vanish. This for instance happens in many cases where \hat{J} has jump discontinuities along infinite radial lines, which typically goes along with a similar discontinuity in J (see [137]).

4.3.4 Results on the Radial Profile

So far, we mostly focused on the angular properties of the pattern and the detector, and their effect on the CRLB. In this subsection, we consider how the radial distribution of the pattern can be matched by the detector functions to lower the CRLB. We provide results on the general case of non-polar-separable \hat{J} , where the optimal radial profile is different for each n , as we shall see below. Then we simplify the results for the polar-separable case.

Theorem 4.2 gives the shape of the radial profile $\hat{h}_n(\omega)$ for each n that minimizes the CRLB.

Theorem 4.2 *For a general, non-polar-separable directional pattern $\hat{J}(\omega, \varphi)$, the shape of the radial profile $\hat{h}_n(\omega)$ for each n that minimizes the CRLB is*

$$\hat{h}_n(\omega) \propto \omega^{2\gamma} \frac{\frac{1}{2\pi} \int_0^{2\pi} \partial_\varphi \hat{J}(\omega, \varphi) e^{jn\varphi} d\varphi}{\sigma_0^2}. \quad (4.32)$$

The proof of Theorem 4.2 is given in Section 4.7. We have the following corollary of Theorem 4.2.

Corollary 4.2 *If $\hat{J}(\omega, \varphi)$ factorizes as $\hat{J}_{Rad}(\omega) \hat{J}_{Ang}(\varphi)$, the CRLB is minimized by taking all $\hat{h}_n(\omega)$ to be the same functions $\hat{h}(\omega)$ satisfying*

$$\hat{h}(\omega) \propto \omega^{2\gamma} \frac{\overline{\hat{J}_{Rad}(\omega)}}{\sigma_0^2}. \quad (4.33)$$

The proof of Corollary 4.2 is given in Section 4.7.

The optimal shape of h given in the above results depends on the pattern that is being matched, as well as the spectrum of the background signal. The result can be interpreted as stating that the optimal radial profile is the prewhitened version of the radial profile of J , up to some angular derivation/averaging. This implies that the optimal detector embeds the whitening operator Δ^γ for the background process. This is reminiscent of the estimation problem for sparse stochastic processes, where detectors embedding the whitening operator also feature prominently [5].

4.4 Experiments

4.4.1 CRLB for Analytical Junctions

In this section, we compute the CRLB associated with a few directional patterns for which explicit formulas are provided. We study three different types of junctions, one that is angularly smooth, one that is angularly discontinuous, and one that is both angularly and radially smooth. Specifically,

$$J_1(\omega, \varphi) = \cos(1.5\varphi)^\beta \quad (4.34)$$

$$J_2(\omega, \varphi) = \begin{cases} 1, & \text{if } \cos(1.5\varphi)^\beta > 0.8 \\ 0, & \text{otherwise} \end{cases} \quad (4.35)$$

$$J_3(\omega, \varphi) = \left(\frac{1}{1 + \omega^\lambda} \right) \cos(1.5\varphi)^\beta \quad (4.36)$$

within the support of $\hat{h}(\omega)$, with $\lambda = 2.1$ and $\beta = 28$.

We are interested in the following quantities: $|u_n|, n|u_n|$ (as it determines the decay rate of the CRLB), and the CRLB for a fixed number of harmonics. For computing the CRLB, we apply three different strategies: “First N ”; “Best N ”; and “ M -fold”. In the first case, we use the first N coefficients; in the Best N case, we select the harmonics that maximize (4.23). In the case of M -fold symmetric junctions one could alternatively choose the first N multiples of M as harmonics. This latter choice accounts for the name “ M -fold”.

For the experiments, we have chosen the first scale of the Meyer-type profile (cf. Table 3.1) with an auxiliary function $v \in C^3([0, 1])$. Related to (4.17), we have chosen the value $\sigma_0 = 1$ since it provides only a scaling factor and does not influence the decay rate of the curve. Figure 4.1 contains an illustration of the results.

In the smooth angular case, we observe that the junction can be expressed with a finite number of harmonics, as expected from its definition. In that case, the CRLB converges to a theoretical positive value. In the case of sharp edges, the rate of decay of the circular harmonic coefficients permits a theoretical vanishing limit for the CRLB.

Finally, as expected, for these three-fold junctions, only every third component plays a significant role in the estimation of the orientation. This can be seen in the almost flat CRLB curve between multiples of three in the “First N ” strategy. We also observe a difference in performance between generically choosing the first N M -fold symmet-

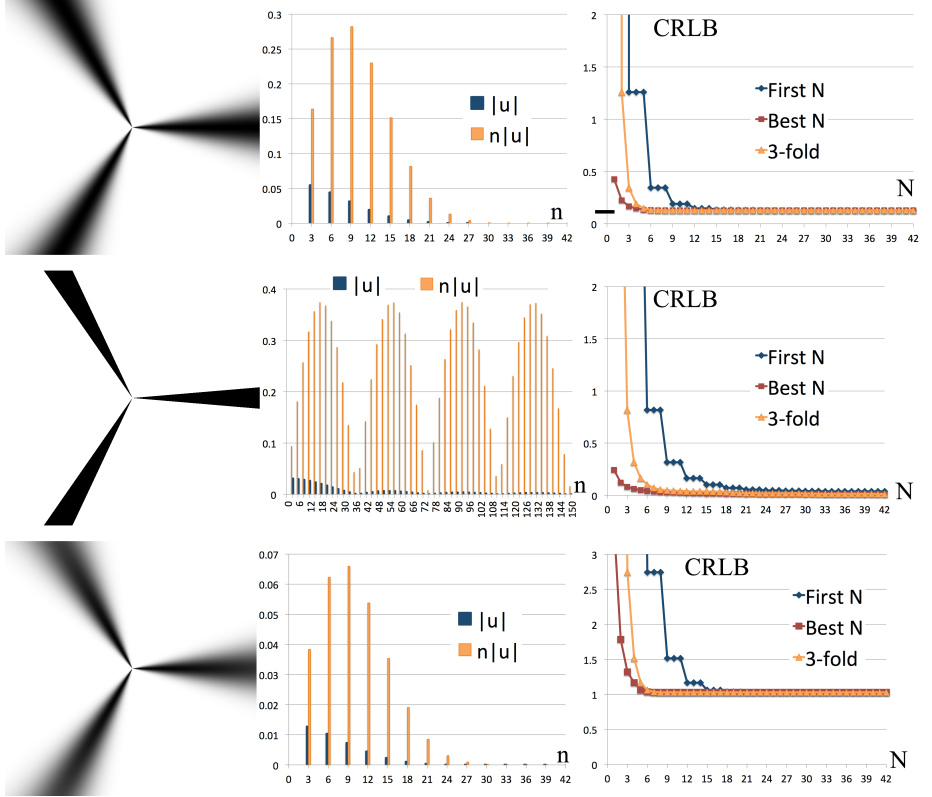


Figure 4.1: First column from top to bottom: illustration of the analytically defined junctions J_1 ; J_2 ; and J_3 in the Fourier domain. Second column: $|u|$ and $n|u|$ as a function of the harmonics n . Third column: the CRLB as a function of the number of harmonics N .

ric coefficients (as a strategy for unknown M -fold junctions) and making our choice of harmonics based on maximizing (4.23).

By looking at Figure 4.1, we observe that, with the right choice of harmonics, the CRLB can be much reduced, even with a small number of harmonics.




Figure 4.2: Examples of non-separable junctions. From left to right: three-fold symmetric, T-shape, arrow-type and four-fold symmetric junctions.

4.4.2 CRLB for Non-Polar-Separable Junctions

In this section, we compute the CRLB of sharp junctions that are not polar separable. Similarly to the previous experiment, we are interested in the following quantities: $|u_n|$; $n|u_n|$; and the CRLB for a fixed number of harmonics. For computing the CRLB, we apply the same three different strategies, namely “First N ”, “Best N ”, and “ M -fold”, as described in Section 4.4.1.

For the experiments, we have chosen the first scale of four different radial profiles: Simoncelli; Meyer; Papadakis; and Held. For further information, we refer to Table 3.1. We expect similar results, since all of those functions approximate the indicator of $[\pi/4, \pi/2]$. To introduce some variety, this time we use a four-fold junction instead of a 3-fold (illustrated in Figure 4.2).

We illustrate the results in Figure 4.3. As before, we observe that the CRLB can be quite small when the harmonics are suitably chosen, even with a small number of harmonics.

4.4.3 A Proposed Estimator and the CRLB

In this section, we test the accuracy of the estimation of the proposed estimator (4.8) on three symmetric (one separable and two non-separable) and two asymmetric junctions. The non-separable junctions are illustrated in Figure 4.2. For symmetric junctions, we have chosen: a sharp, non-polar separable three-fold junction; the analytically defined junction J_2 (4.35); and the analytically defined junction J_3 (4.36) with $\lambda = 2.5$.

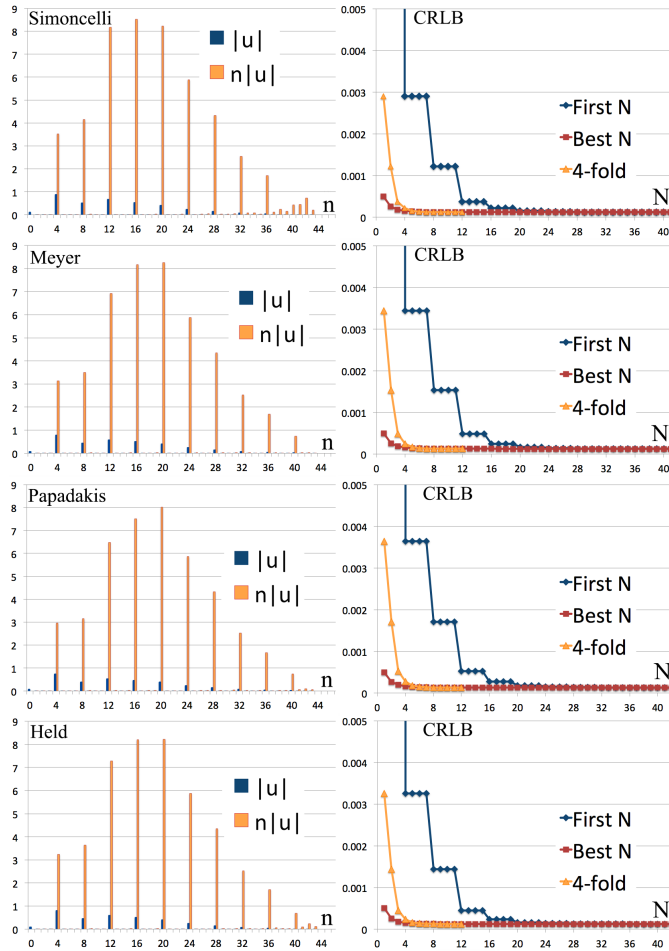


Figure 4.3: From top to bottom: wavelet coefficients of a sharp, non-separable four-fold symmetric junction and the corresponding CRLB of four different radial profiles: Simoncelli; Meyer; Papadakis; and Held. First column: $|u|$ and $n|u|$ as a function of harmonics n . Second column: the CRLB as a function of the number of harmonics N .

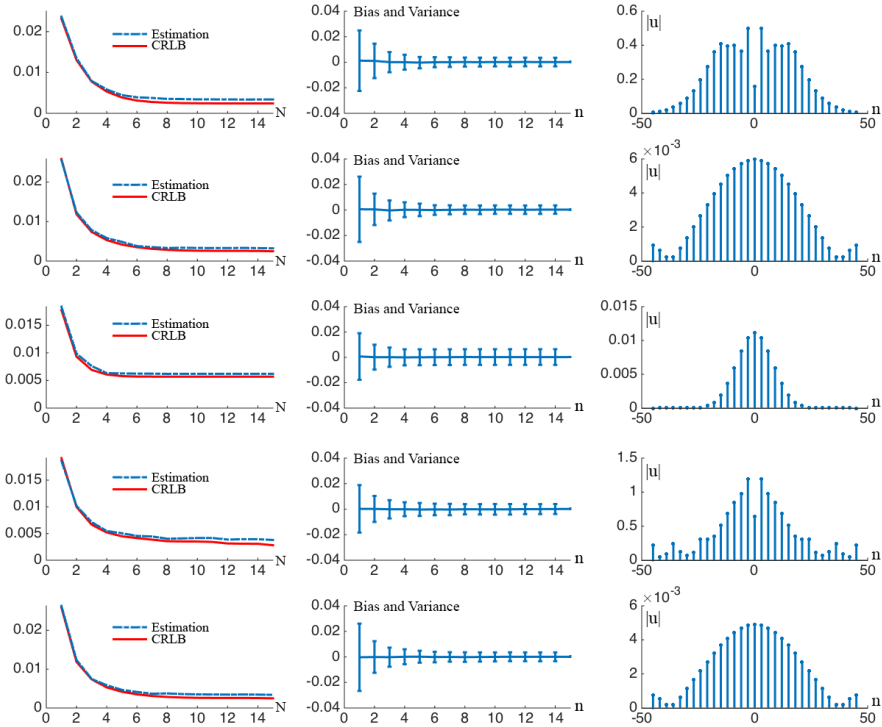


Figure 4.4: First row: results corresponding to a sharp, three-fold symmetric, non-separable junction, with Meyer radial detector profile. Second row: results on the analytically defined junction J_2 , with Meyer radial detector profile. Third row: results on the analytically defined junction J_3 , with $\lambda = 2.5$ and Meyer radial detector profile. Fourth row: results corresponding to a sharp, three-fold symmetric, non-separable junction, with LoG radial detector profile. Fifth row: results on the analytically defined junction J_2 , with LoG radial detector profile. First column: accuracy of the orientation estimation of the proposed estimator (dashed blue line) compared to the CRLB (continuous red line). Second column: bias and variance of the proposed estimator. Third column: $|u|$ as a function of harmonics n .

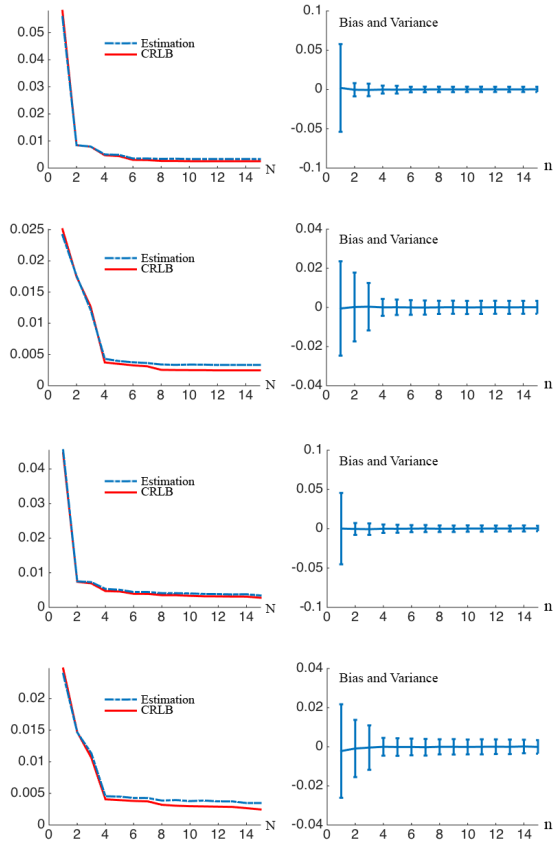


Figure 4.5: First row: results corresponding to a sharp, non-separable, arrow-type junction, with Meyer radial detector profile. Second row: results corresponding to a sharp, non-separable, T-shaped junction, with Meyer radial detector profile. Third row: results corresponding to a sharp, non-separable, arrow-type junction, with LoG radial detector profile. Fourth row: results corresponding to a sharp, non-separable, T-shaped junction, with LoG radial detector profile. First column: accuracy of the orientation estimation of the proposed estimator (dashed blue line) compared to the CRLB (continuous red line). Second column: Bias and variance of the proposed estimator.

The coefficients c_α in (4.8) are obtained in each case by projecting the junction of interest on the measurement functions.

We choose the variance of the noise such that it corresponds to an SNR of 17.22 dB. The typical SNR range of real images where the quality is still acceptable is 15-35 dB, so our experimental conditions are representative of a noisy image. We build 1,000 different realizations to make the experiments statistically reliable.

Since we are detecting three-fold symmetric junctions, only every third component plays a role in the estimation of the orientation. Thus, for the estimator, we use the “ M -fold” strategy with $M = 3$. For the radial part of our detector, we have chosen the first scale of the Meyer wavelet profile and the LoG defined in Section 3.2.3.

The results for the symmetric junctions are illustrated in Figure 4.4. We can observe that the accuracy of the estimator follows closely the CRLB curve, while staying above, as expected. We also show experimentally that the estimator is unbiased with more than one harmonic and that the variance of the estimator is small.

The results for the asymmetric junctions are illustrated in Figure 4.5. Like in the symmetric cases, we can observe that the accuracy of the estimator follows closely the CRLB curve, while staying above as expected. We also show experimentally that the estimator is unbiased with more than one harmonic and that the variance of the estimator is small.

4.5 Extension to Wavelets

Steerable wavelet frames are adapted to capture the local orientation of features, within a multiresolution hierarchy. In this section, we extend our results on the CRLB to such schemes.

To simplify the notations, we consider wavelets that are centered at the origin. In particular, we are interested in circular harmonic wavelets since they correspond to our steerable measurement framework (4.3). We recall that, in the Fourier domain, the circular harmonic wavelet ξ at scale i and harmonic channel n has the form

$$\hat{\xi}_{n,i}(\omega, \varphi) = 2^i \hat{h}(2^i \omega) e^{jn\varphi}, \quad (4.37)$$

where \hat{h} is a radial wavelet profile satisfying Proposition 3.1. Typical examples for \hat{h} are given in Table 3.1. For further details, we refer to Section 3.4.

Similarly to the single-case formulation (4.15), we take the values of n from the pre-defined set of harmonics $H_+ = \{|n| : n \in H\}$. In an extension of (4.2), the measurement

functions and measurements are indexed here by the pair $\alpha = (n, i)$ of the harmonic $n \in H_+$ and the scale i . We note that, in the case of wavelet measurements, we can have multiple measurement functions (at different scales) for the same harmonic n .

We make the additional assumption that \hat{h} is real-valued, which holds true for every radial designs given in Table 3.1. Finally, given that the radial patterns \hat{h} of interest are bandpass, we assume that $\hat{h}(\omega) = 0$ for $\omega \notin (\pi/4, \pi]$ to set the fundamental scale.

4.5.1 The Cramér-Rao Lower Bound

To simplify future formulas, we introduce the notations

$$\tilde{q}_{n,i} = 2^{-i\gamma} q_{n,i} = e^{jn\theta^*} (2^{-i\gamma} u_{n,i}) + 2^{-i\gamma} s_{n,i}, \quad (4.38)$$

$$\tilde{s}_{n,i} = 2^{-i\gamma} s_{n,i}, \quad (4.39)$$

where, we recall, $q_{n,i} = \langle I_{\text{loc}}, \xi_{n,i} \rangle$, $u_{n,i} = \langle J, \xi_{n,i} \rangle$, and $s_{n,i} = \langle S, \xi_{n,i} \rangle$.

The angle θ^* is estimated at each point by steering the whole template, which can be seen as a sum of templates. The task of estimating θ^* based on $q_{n,i}$ is essentially the same as its estimation based on $\tilde{q}_{n,i}$. In particular, the CRLB is the same.

We introduce the following notations:

$$b_z = \frac{1}{2\pi} \int_{\mathbb{R}} \omega^z \hat{h}(\omega)^2 \omega d\omega, \quad (4.40)$$

$$d_z = \frac{1}{2\pi} \int_{\mathbb{R}} \omega^z \hat{h}(\omega) \hat{h}(2\omega) \omega d\omega, \quad (4.41)$$

and

$$B = \sigma_0^2 b_{-2\gamma}, \quad D = \sigma_0^2 2^{1-\gamma} d_{-2\gamma}. \quad (4.42)$$

The covariance of the modified wavelet measurements is

$$\text{Cov}(\tilde{q}_{n,i}, \tilde{q}_{m,k}) = \begin{cases} B, & \text{if } m = n, k = i \\ D, & \text{if } m = n, |k - i| = 1 \\ 0, & \text{otherwise.} \end{cases} \quad (4.43)$$

The last case comes from the fact that, for the radial profiles defined in Table 3.1, the overlap between $\hat{h}(\omega)$ and $\hat{h}(2^{|k-i|}\omega)$ is 0 for $|k - i| > 1$. The calculations are given in Section 4.7. We can find an upper bound for d_z with respect to b_z and, as a consequence, for D with respect to B .

Lemma 4.1 *Assume that \hat{h} is real-valued and that $\hat{h}(\omega) = 0$ for $\omega \notin (\pi/4, \pi]$. Then,*

$$b_z > 2^{z/2+2} |d_z|. \quad (4.44)$$

Consequently,

$$B > 2|D|. \quad (4.45)$$

The proof of Lemma 4.1 is given in the Appendix of [135].

The main challenge in computing the Fisher information from (4.11) for wavelet measurements is that, unlike in context of Section 4.3, the covariance matrix of measurements $\{\tilde{q}_{n,i}\}$ is not diagonal, so it is not as easy to invert. But this challenge is still surmountable because, as we see from (4.43), we can rearrange the measurements such that the covariance matrix is (at most) tridiagonal. We can form such a rearrangement of the measurements as follows: Our starting point is an arbitrary set of unique measurement indices of the form (n, i) , where n is the harmonic and i is the scale of the corresponding measurement. As before, since (n, i) and $(-n, i)$ carry the same information for real-valued patterns, with $q_{-n,i} = \overline{q_{n,i}}$ and $u_{-n,i} = \overline{u_{n,i}}$, we can assume that all harmonics n are positive without loss of generality. We divide the set of measurement indices into as few disjoint sets G_r as possible, subject to three conditions ($r = 1, \dots, g$ is an index for the sets).

- Each set G_r consists of indices with a single fixed harmonic n_r .
- Each set G_r only contains indices with consecutive scales—it can also contain only a single element.
- If the two sets G_r, G_s share the same harmonic $n_r = n_s$, then the scales in G_r, G_s differ by a minimum of 2.

The last condition is a consequence of the first two and requires one to have a minimal number of sets. Here, g denotes the total number of such sets.

From the above conditions, it follows that each set G_r has some l_r elements of the form $(n_r, i_r), \dots, (n_r, i_r + l_r - 1)$, for some harmonic $n_r \in H_+$ and minimum scale i_r . We introduce the notation $(r|e) := (n_r, i_r + e - 1)$ for $r = 1, \dots, g$ and $e = 1, \dots, l_r$. In this way, we can re-index all measurements as $\tilde{q}_{r|e}$. The idea behind this re-indexing is that,

following (4.43), the covariance matrix of the measurements $\{\tilde{q}_{r|e}\}$ for each fixed r is Toeplitz-tridiagonal of the form

$$\mathbf{T}_r = \begin{pmatrix} B & D & 0 & 0 & \dots & 0 \\ D & B & D & 0 & \dots & 0 \\ 0 & D & B & D & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & D & B & D \\ 0 & 0 & \dots & 0 & D & B \end{pmatrix}_{l_r \times l_r}. \quad (4.46)$$

Moreover, the overall covariance matrix is block-diagonal, with \mathbf{T}_r , $r = 1, \dots, g$, as its diagonal blocks.

Using this reformulation of the problem, we are then able to compute the CRLB explicitly.

Theorem 4.3 *For wavelet measurements $q_{r|e}$, $r = 1, \dots, g$, $e = 1, \dots, l_r$, the Fisher information is given by*

$$FI(\theta^*) = 2 \sum_{r=1}^g n_r^2 \sum_{t=1}^{l_r} \left(B + 2D \cos\left(\frac{t\pi}{l_r+1}\right) \right)^{-1} \times \left| \sum_{e=1}^{l_r} 2^{-(i_r+e-1)\gamma} u_{r|e} \sin\left(e \frac{t\pi}{l_r+1}\right) \right|^2. \quad (4.47)$$

The CRLB for the estimation problem is given by $1/FI(\theta^*)$. It satisfies

$$\frac{B + |2D|}{2 \sum_{n,i} n^2 4^{-i\gamma} |u_{n,i}|^2} \geq CRLB \geq \frac{B - 2|D|}{2 \sum_{n,i} n^2 4^{-i\gamma} |u_{n,i}|^2}. \quad (4.48)$$

The proof of Theorem 4.3 is given in the Appendix of [135].

After this general result, we now look at the asymptotic behavior in number of harmonics and scales.

4.5.2 Asymptotic Behavior for All Scales and Harmonics

The correlation pattern of the background noise, given in (4.6), reveals that the high-frequency components of the junction are affected by noise of smaller magnitude. Thus, it is tempting to consider fine wavelet scales to estimate the angle, as they tend to become rather noiseless. It is, however, important that the wavelet coefficients of the junction decay slower than the predicted rate for the noise.

Table 4.1: Excluded γ values

Wavelet type	Excluded γ values
Shannon	∞
Simoncelli	0.736307866729164
Meyer	1.978000155843049
Papadakis	1.476205252555835
Held	1.442557587435832

We now consider the case where we have access to each wavelet coefficient $q_{n,i} = \langle I_{loc}, \xi_{n,i} \rangle$ for $n \in \mathbb{Z}$, $i \in \mathbb{Z}$, assuming that $J(r, \theta)$ is expanded in terms of $\{\xi_{n,i}\}_{n \in \mathbb{Z}, i \in \mathbb{Z}}$. Once again, for real-valued data, we can limit ourselves to $n \in \mathbb{Z}_+$ for computing the CRLB, due to Hermitian symmetry.

Theorem 4.4 *Assume γ is such that $\frac{b_0}{|d_0|} \neq 2^{1-\gamma} + 2^{1+\gamma}$ and such that J has an expansion in terms of $\{\xi_{n,i}\}$. Then, by observing $\langle I_{loc}, \xi_{n,i} \rangle$, the CRLB can vanish only if*

$$(-\Delta)^{\gamma/2} \frac{\partial}{\partial \theta} J(r, \theta) \notin L_2. \quad (4.49)$$

The proof of Theorem 4.4 is given in the Appendix of [135].

The γ values to exclude for specific isotropic bandlimited wavelets are collected in Table 4.1.

4.5.3 Experiments with Wavelets

In this section, we compute the CRLB of Junction J_3 with $\lambda = 2.1$ and 4.5. These choices of λ correspond to two cases: first, where the background decays faster than the junction (in Fourier); and second, where it is the junction that decays faster. The value of γ is fixed as 2.5.

For the experiments, we have chosen the Meyer-type wavelet with an auxiliary function $v \in C^3([0, 1])$. We again choose $\sigma_0 = 1$ in (4.48) since it provides a scaling factor that does not influence the rate of decay of the curve.

Based on the graphs of Figure 4.6, we can observe the following. For $\lambda = 2.1$, one wavelet scale was not enough for the CRLB to converge to a theoretical perfect estimation (like in Figure 4.1, third experiment). However, by increasing the number of scales,




Figure 4.6: First column: Illustration of the analytically defined junction J_3 with $\lambda = 2.1$ and 4.5 in the Fourier domain. Second column: The CRLB as a function of the largest wavelet scale.

the 0 bound is asymptotically achievable. This illustrates the main result of Theorem 4.4, which shows that, by using finer scales, wavelets can improve the estimate, even in the case of junctions of a fixed size, provided the radial profile of the junction fulfills the conditions of the theorem. For $\lambda = 4.5$, as expected, the CRLB converges to a positive theoretical value.

4.6 Conclusion

In this chapter, we derived the Cramér-Rao lower bound (CRLB) on the error of estimation of the local orientations of directional patterns in images using steerable filters. Based on those bounds, we provided solutions to improve and optimize the choice of the detector functions. In particular, we addressed the problem of choosing a finite number of harmonics to achieve the lowest CRLB, as well as finding the optimal radial profile of the detector function. We also studied the bounds on the performance of multiscale steerable estimators and contrasted them with the single-scale case. We proposed an estimator for the estimation of orientations and compared its estimation error to the theoretical bounds. We provided several experiments on different realistic junctions and directional patterns.

4.7 Notes

Proof of Theorem 4.1

As noted in the introduction to Theorem 4.1, the CRLBs based on the sets H and H_+ are the same. Moreover, for $m, n \in H_+$, $\mathbb{E}\{s_m s_n\} = \mathbb{E}\{\overline{s_{-m}} s_n\} = 0$ (see the calculation of the correlation form (4.51)). This shows that the pseudo-covariance matrix of q_n , $n \in H_+$ is $\mathbf{0}$. In other words, the measurements q_n , $n \in H_+$, are *circularly* Gaussian [138]. (We have excluded q_0 , which does not depend on the parameter being estimated.) Consequently, we can use (4.11) to compute the Fisher information.

We need to compute the mean and covariance of $q_n = e^{-jn\theta^*} u_n + s_n$, $n \in H$, where $u_n = \langle J, \xi_n \rangle$ and $s_n = \langle S, \xi_n \rangle$. Due to the assumptions on the vanishing moments of ξ_n , all contributions to the mean come from the first term, so that

$$\mu_n = e^{-jn\theta^*} u_n. \quad (4.50)$$

The covariance can be written as

$$c_{mn} = \mathbb{E}\{\overline{s_m} s_n\} = \mathbb{E}\left\{\overline{\langle S, \xi_m \rangle} \langle S, \xi_n \rangle\right\}, \quad (4.51)$$

which is the same as the correlation form $\mathcal{B}_S\{\xi_m, \xi_n\}$ and can be computed by the formula given in (4.6) like

$$c_{mn} = \frac{\sigma_0^2}{(2\pi)^2} \int_0^{2\pi} e^{j(n-m)\varphi} d\varphi \int_0^\infty \rho^{1-2\gamma} \overline{\hat{h}_m(\rho)} \hat{h}_n(\rho) d\rho, \quad (4.52)$$

where we have used the Fourier expressions for ξ_m, ξ_n from (4.13). We note that the integral converges despite the singularity at the origin, due to the vanishing moments of ξ_m and ξ_n . The first integral vanishes for $m \neq n$, and takes the value 2π otherwise, which means the covariance matrix is diagonal with diagonal variance terms

$$c_{nn} = \frac{\sigma_0^2}{2\pi} \int_0^\infty \rho^{1-2\gamma} |\hat{h}_n(\rho)|^2 d\rho. \quad (4.53)$$

Using (4.11), we get

$$\text{FI}(\theta^*) = \frac{4\pi}{\sigma_0^2} \sum_{n \in H_+} \frac{n^2 |u_n|^2}{\int_0^\infty \rho^{1-2\gamma} |\hat{h}_n(\rho)|^2 d\rho}. \quad (4.54)$$

The CRLB of the estimation problem is the reciprocal of the Fisher information.

Proof of Theorem 4.2

We start our derivation from equation (4.22).

We define the quantity

$$M_n := \frac{n^2 |u_n|^2}{\frac{1}{2\pi} \int_0^\infty |\hat{h}_n(\omega)|^2 P(\omega) \omega d\omega}, \quad (4.55)$$

where $P(\omega) = \sigma_0^2 \omega^{-2\gamma}$, and, for convenience, we introduce the notation

$$P_n := \frac{1}{2\pi} \int_0^\infty |\hat{h}_n(\omega)|^2 P(\omega) \omega d\omega. \quad (4.56)$$

Then, the CRLB can be written, in general, as

$$\text{CRLB} = \frac{1}{2 \sum_{n \in H_+} M_n}. \quad (4.57)$$

We optimize the CRLB by maximizing M_n term-by-term.

We recall that for sufficiently regular ξ_n and \hat{J} ,

$$n u_n = \langle J, \partial_\theta \xi_n \rangle = -\langle \partial_\theta J, \xi_n \rangle = -\frac{1}{(2\pi)^2} \langle \partial_\varphi \hat{J}, \hat{\xi}_n \rangle. \quad (4.58)$$

Therefore,

$$\begin{aligned}
 M_n &= \frac{(2\pi)^{-4} \left| \int_0^\infty \int_0^{2\pi} \overline{\partial_\varphi \hat{J}(\omega, \varphi)} \hat{h}_n(\omega) e^{jn\varphi} d\varphi \omega d\omega \right|^2}{P_n} \\
 &= \frac{(2\pi)^{-2} \left| \int_0^\infty \frac{\frac{1}{2\pi} \int_0^{2\pi} \overline{\partial_\varphi \hat{J}(\omega, \varphi)} e^{jn\varphi} d\varphi}{\sqrt{P(\omega)}} \sqrt{P(\omega)} \hat{h}_n(\omega) \omega d\omega \right|^2}{P_n} \\
 &\leq \frac{(2\pi)^{-2} \int_0^\infty \frac{\left| \frac{1}{2\pi} \int_0^{2\pi} \overline{\partial_\varphi \hat{J}(\omega, \varphi)} e^{jn\varphi} d\varphi \right|^2}{P(\omega)} \omega d\omega P_n}{P_n} \\
 &= (2\pi)^{-2} \int_0^\infty \frac{\left| \frac{1}{2\pi} \int_0^{2\pi} \overline{\partial_\varphi \hat{J}(\omega, \varphi)} e^{jn\varphi} d\varphi \right|^2}{\sigma_0^2 \omega^{-2\gamma}} \omega d\omega, \tag{4.59}
 \end{aligned}$$

using the Cauchy-Schwartz inequality, with equality iff

$$\hat{h}_n(\omega) \propto \frac{\frac{1}{2\pi} \int_0^{2\pi} \overline{\partial_\varphi \hat{J}(\omega, \varphi)} e^{jn\varphi} d\varphi}{\sigma_0^2 \omega^{-2\gamma}}. \tag{4.60}$$

This gives the shape of the radial profile $\hat{h}_n(\omega)$ for each n that minimizes the CRLB.

Proof of Corollary 4.2

Based on the assumption, $\hat{J}(\omega, \varphi)$ factorizes as

$$\hat{J}(\omega, \varphi) = \hat{J}_{\text{Rad}}(\omega) \hat{J}_{\text{Ang}}(\varphi). \tag{4.61}$$

Using the elements of the proof of Theorem 4.2, we optimize (minimize) the CRLB by maximizing M_n (4.55) term-by-term. Then,

$$M_n \leq (2\pi)^{-2} \int_0^\infty \frac{\left| \frac{1}{2\pi} \int_0^{2\pi} \overline{\partial_\varphi \hat{J}(\omega, \varphi)} e^{jn\varphi} d\varphi \right|^2}{P(\omega)} \omega d\omega, \tag{4.62}$$

$$= (2\pi)^{-2} \int_0^\infty \frac{\left| \frac{1}{2\pi} \hat{J}_{\text{Rad}}(\omega) \int_0^{2\pi} \overline{\partial_\varphi \hat{J}_{\text{Ang}}(\varphi)} e^{jn\varphi} d\varphi \right|^2}{\sigma_0^2 \omega^{-2\gamma}} \omega d\omega. \tag{4.63}$$

Here, $\int_0^{2\pi} \overline{\partial_\varphi \hat{J}_{\text{Ang}}(\varphi)} e^{jn\varphi} d\varphi$ is a constant c_n , more precisely,

$$c_n = \mathcal{F}_n \left\{ \overline{\partial_\varphi \hat{J}_{\text{Ang}}(\varphi)} \right\} \quad (4.64)$$

We obtain equality, and hence minimize the CRLB iff we choose $\hat{h}_n(\omega)$ such that

$$\hat{h}_n(\omega) \equiv \hat{h}(\omega) \propto \frac{\overline{\hat{J}_{\text{Rad}}(\omega)}}{\sigma_0^2 \omega^{-2\gamma}}. \quad (4.65)$$

Covariance of the Wavelet Measurements

The covariance of the modified wavelet measurements $\tilde{q}_{n,i}, \tilde{q}_{m,k}$ is equal to the covariance of their random parts $\tilde{s}_{n,i}, \tilde{s}_{m,k}$. To compute it, we note that it is proportional to the covariance of $\langle S, \xi_{n,i} \rangle$ and $\langle S, \xi_{m,k} \rangle$, which can be computed from (4.6). Thus,

$$\begin{aligned} & \text{Cov}(\tilde{q}_{n,i}, \tilde{q}_{m,k}) \\ &= \mathbb{E} \left\{ \overline{\tilde{s}_{n,i}} \tilde{s}_{m,k} \right\} \\ &= 2^{-(i+k)\gamma} \mathbb{E} \left\{ \overline{s_{n,i}} s_{m,k} \right\} = 2^{-(i+k)\gamma} \mathcal{B}_S \left\{ \xi_{n,i}, \xi_{m,k} \right\} \\ &= \frac{\sigma_0^2 2^{-(i+k)\gamma}}{(2\pi)^2} \int_0^{2\pi} \int_0^\infty \frac{2^{(i+k)} \hat{h}(2^i \omega) \hat{h}(2^k \omega)}{\omega^{2\gamma}} e^{j(m-n)\varphi} \omega d\omega d\varphi \\ &= \frac{\sigma_0^2}{(2\pi)^2} 2^{(i+k)(1-\gamma)} \int_0^{2\pi} e^{j(m-n)\varphi} d\varphi \int_0^\infty \frac{\hat{h}(2^i \omega) \hat{h}(2^k \omega)}{\omega^{2\gamma}} \omega d\omega \\ &= \frac{\sigma_0^2}{2\pi} 2^{|k-i|(1-\gamma)} \delta[n-m] \int_0^\infty \frac{\hat{h}(\omega) \hat{h}(2^{|k-i|} \omega)}{\omega^{2\gamma}} \omega d\omega. \end{aligned} \quad (4.66)$$

Chapter 5

Template-free Detection of Symmetries

5.1 Overview

In this chapter, we propose an approach to detect and group different kinds of local symmetries in images in a multiscale, and rotation-invariant way. Our method is also template-free, in the sense that there is need neither to design a specific junction or ridge-shaped template nor to align it with the pattern of interest. Instead, we propose an efficient wavelet-based method to determine the order of local symmetry at each location. Our algorithm relies on circular harmonic wavelets (Chapter 3 and [98, 123]), which distribute the energy of the signal among a set of angular harmonics. Based on this angular distribution, we propose a measure of symmetry on the wavelet coefficients and a hypothesis test for local symmetry at each pixel. In particular, we use the F-test to examine the distribution of the energy across different channels. 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 for matching the template to image features of different orientations. Finally, as a result of the statistical formu-

lation, 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 chapter is organized as follows: In Section 5.2, we recall the important properties of the circular harmonic wavelets that define the analytical framework used in this chapter. 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. Based on the definitions given in the Nomenclature, we lay down the mathematical background of our null hypothesis and examine the distribution of the wavelet coefficients. Based on the results, we formulate a hypothesis test in Section 5.4 for the detection of local symmetries of any given order. Next, in Section 5.5, we deal with the problem of classifying local symmetries of different orders. Finally, in Section 5.6, we provide experimental results on synthetic images, biological micrographs, and electron-microscopy images to demonstrate the performance of the algorithm¹.

5.2 Concept

5.2.1 Local Wavelet Energies

We rely on the circular harmonic wavelets that have been presented in Section 3.4. We recall that, based on a predefined set $H \subset \mathbb{Z}$ of harmonics, the N -channel tight frame of wavelets is generated as $\{\xi_{n,i,\mathbf{k}} = \mathcal{F}^{-1}\{\hat{\xi}_{n,i,\mathbf{k}}\}\}_{n \in H}$. There, ξ_n is defined as $\xi_n := \mathcal{R}^n \phi$, where ϕ denotes the isotropic mother wavelet. The translates and dilates of ξ_n are given by

$$\begin{aligned} \xi_{n,i}(\mathbf{x} - \mathbf{x}_0) &= 2^{-i} \xi_n \left(\frac{\mathbf{x} - \mathbf{x}_0}{2^i} \right) \\ &= 2^{-i} \mathcal{R}^n \left\{ \phi \right\} \left(\frac{\mathbf{x} - \mathbf{x}_0}{2^i} \right) \\ &= \mathcal{R}^n \left\{ 2^{-i} \phi \left(\frac{\cdot - \mathbf{x}_0}{2^i} \right) \right\} (\mathbf{x}) = \mathcal{R}^n \left\{ \phi_i(\cdot - \mathbf{x}_0) \right\} (\mathbf{x}), \end{aligned} \quad (5.1)$$

due to the translation- and scale-invariance properties of the transform. The wavelets $\xi_{n,i,\mathbf{k}}$ form a tight wavelet frame, thus any finite-energy function f can be decomposed

¹The results in this chapter are based on the article [121]: Z. Püspöki, M. Unser, "Template-Free Wavelet-Based Detection of Local Symmetries", *IEEE Transactions on Image Processing*, vol. 24, no. 10, pp. 3009 - 3018, October 2015.

as

$$f = \sum_{n,i,\mathbf{k}} \langle f, \xi_{n,i,\mathbf{k}} \rangle \xi_{n,i,\mathbf{k}}. \quad (5.2)$$

We note that $\overline{\xi_n} = \xi_{-n}$. This implies that for a real image I_{loc} ,

$$q_{-n} = \langle I_{\text{loc}}, \xi_{-n} \rangle = \langle I_{\text{loc}}, \overline{\xi_n} \rangle = \overline{\langle I_{\text{loc}}, \xi_n \rangle} = \overline{q_n}. \quad (5.3)$$

It means that q_{-n} and q_n essentially carry the same information. We can also exclude $n = 0$ from consideration, since the corresponding measurement is not sensitive to symmetries. Thus, we define the set of all used harmonics $H_+ \subset H$ as

$$H_+ := \{|n| : n \in H \setminus 0\}. \quad (5.4)$$

Let H_{M_+} be the subset of H_+ consisting of harmonics that are integer multiples of M with $H_{M_+} = \{n \in H_+ : n = Mk, k \in \mathbb{Z}\} \subset H_+$, and let $q_{n,i} = \langle \xi_{n,i}(\cdot - \mathbf{x}_0), f \rangle$ be the n th-channel coefficient at the scale and location of interest.

In the following, we take advantage of the fact that an n th-order harmonic wavelet has a rotational symmetry of order n around its center, corresponding to the n th-order rotational symmetry of $e^{in\varphi}$.

We propose to base our analysis on the quantity

$$E_{n,i} = \frac{|q_{n,i}|^2}{\sum_{m \in H_+} |q_{m,i}|^2} \quad (5.5)$$

for each $n \in H_+$ at a given scale and location, which we refer to as normalized wavelet energies. We present in Figure 5.1 the distribution of normalized wavelet energies at a particular scale ($i = 1$) at three different locations: three-fold junction (a), ridge (b), and non-symmetric point (c). We observe that the “energy” of the wavelet is highly concentrated in M -fold periodic harmonics in cases 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.

5.2.2 Null Hypothesis

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

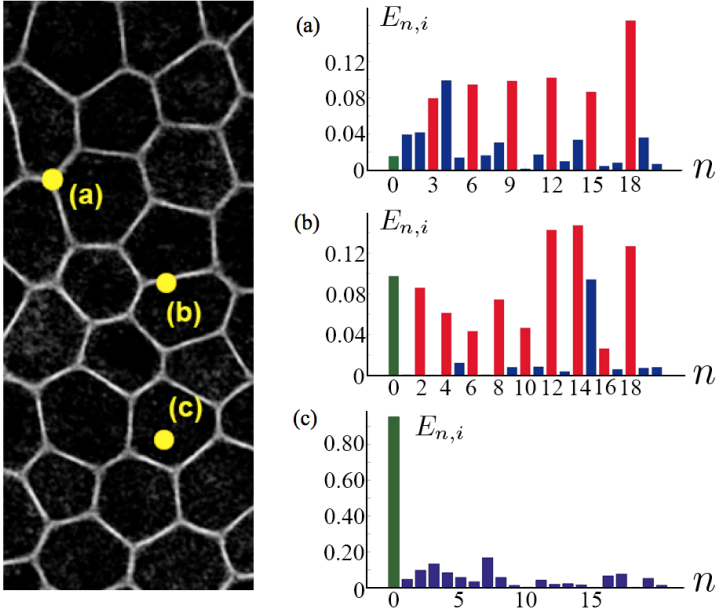


Figure 5.1: Illustration of local symmetries in an image at three different locations: (a), (b) and (c). The original image (hexagonal embryonic stem cells in light microscopy [31]) is on the left. Subfigure (a) is associated with the normalized wavelet energy vector $E_{n,1}$ (5.5) at location (a) and scale 1. Location (a) corresponds to a three-fold symmetric junction, hence every third component has a high amplitude. Subfigure (b) is associated with the vector at location (b) and scale 1. Location (b) corresponds to a two-fold junction (ridge), hence every second component has a high amplitude. Subfigure (c) is associated with 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.

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 H_{M_+} must be the same or lower than the average over its complement $H_+ \setminus H_{M_+}$. 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 $q_{n,i}$. In the following, we demonstrate that, the circular harmonic wavelet coefficients of images that are realizations of isotropic random fields are uncorrelated on a scale-by-scale basis. Moreover, when the isotropic random field is Gaussian, the wavelet coefficients are i.i.d. Gaussian within each wavelet scale. This result serves as one of the theoretical contribution of this chapter.

5.3 Wavelet-Domain Statistics

We shall consider two types of fields: stationary ones, and self-similar ones, which also have a convenient expansion in terms of wavelets. Hence, for the analysis, we use elements of the theory of generalized stochastic fields to model statistically isotropic images. For the basic definitions and mathematical preliminaries we refer to the Nomenclature. For a thorough overview on stochastic fields we refer to [5].

5.3.1 Stationary Isotropic Fields

We are interested in wide-sense stationary fields with zero mean [cf. 0.18]. Our first result is as follows.

Theorem 5.1 *Let S be a realization of a second-order w.s.s. field with isotropic power spectrum $P_S(\boldsymbol{\omega}) = P_S(\omega)$. Then, the wavelet coefficients $q_{n,i} = \{\langle S, \xi_{n,i}(\cdot - \mathbf{x}_0) \rangle\}_n$ at a given scale i and position \mathbf{x}_0 are uncorrelated with zero mean and constant variance*

$$\frac{1}{2\pi} \int_0^\infty P_S\left(\frac{\omega}{2^i}\right) |\hat{h}(\omega)|^2 \omega d\omega. \quad (5.6)$$

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

The proof of Theorem 5.1 is given in Section 5.8.

5.3.2 Self-similar Isotropic Fields

Our formalism allows us to extend our theory to isotropic self-similar fields whitened by $L = (-\Delta)^{\frac{\gamma}{2}}$. This particular choice of background model corresponds to the one presented in Section 4.2.2.

Theorem 5.2 *Let S be a realization of an isotropic self-similar field of order γ and variance σ_0^2 . Then, the wavelet coefficients $q_{n,i} = \{\langle S, \xi_{n,i}(\cdot - \mathbf{x}_0) \rangle\}_n$ at a given scale i and position \mathbf{x}_0 are uncorrelated with zero mean and constant variance*

$$\frac{\sigma_0^2}{2\pi} 4^{i\gamma} \int_0^\infty \frac{1}{\omega^{2\gamma}} |\hat{h}(\omega)|^2 \omega d\omega. \quad (5.7)$$

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

This result is an immediate corollary of (4.66) of Section 4.5.

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.

Theorem 5.1 and Theorem 5.2 establish the validity of our null hypothesis for a broad class of isotropic processes.

5.4 F -Test for M -Fold Symmetric Junctions

5.4.1 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

$$F(\mathbf{q}; M) = \frac{\frac{1}{|H_{M_+}|} \sum_{n \in H_{M_+}} |q_n|^2}{\frac{1}{|H_+| - |H_{M_+}|} \sum_{n \in H_+ \setminus H_{M_+}} |q_n|^2} \quad (5.8)$$

at each point \mathbf{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 $2|H_{M_+}|$ and $2(|H_+| - |H_{M_+}|)$ degrees of freedom, respectively. It means that F in (5.8)

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

$$p(F \geq F_{\text{obs}}) < \alpha, \quad (5.9)$$

where *p* is the *F* probability distribution with $2|H_{M_+}|$ and $2(|H_+| - |H_{M_+}|)$ degrees of freedom. Equivalently, we may compute F_α , the value of the statistic such that

$$p(F \geq F_\alpha) = \alpha. \quad (5.10)$$

Then, if $F_{\text{obs}} > F_\alpha$, we conclude that (5.9) 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.

5.4.2 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

$$F_i(\mathbf{q}; M) = \frac{\frac{1}{|H_{M_+}|} \sum_{n \in H_{M_+}} |q_{n,i}|^2}{\frac{1}{|H_+| - |H_{M_+}|} \sum_{n \in H_+ \setminus H_{M_+}} |q_{n,i}|^2} \quad (5.11)$$

at each point \mathbf{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 5.4.1.

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

$$p = \text{Prob} \bigcup_l \left\{ F_l \geq \max_i F_{i,\text{obs}} \right\}, \quad (5.12)$$

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

$$\sum_l \text{Prob} \left\{ F_l \geq \max_i F_{i,\text{obs}} \right\} = I \times p(F \geq \max_i F_{i,\text{obs}}), \quad (5.13)$$

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

$$p(F \geq \max_i F_{i,\text{obs}}) < \frac{\alpha}{I}, \quad (5.14)$$

or, equivalently, that

$$\max_i F_{i,\text{obs}} > F_{\alpha/I}. \quad (5.15)$$

5.5 Classification of Symmetry Orders

The F -score defined earlier can also be used to classify keypoints in 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 $p(m|\mathbf{q})$ of having an m -fold symmetry at the point of interest, given the wavelet coefficients \mathbf{q} . The MAP estimator of the local order of symmetry at each point would then be

$$\hat{M} = \underset{m}{\text{argmax}} p(m|\mathbf{q}), \quad (5.16)$$

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

As in (5.8), let us denote by $F(\mathbf{q}; m)$ the observed value of the F -score for the m -fold test. Recall that, by definition, a higher value of $F(\mathbf{q}; m)$ indicates a higher concentration of energy in m -fold symmetric wavelets. It is therefore natural to assume that, for fixed m , the probability $p(m|\mathbf{q})$ increases monotonically with $F(\mathbf{q}; m)$; that is,

$$p(m|\mathbf{q}) = g_m(F(\mathbf{q}; m)), \quad (5.17)$$

where g_m is some monotonically increasing function.

We can thus construct different classification rules inspired by the MAP approach. To do so, we choose different families of monotonically increasing functions g_m . 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 (ignoring normalization):

1. The identity function $g_m(u) \equiv u$ for all m . The corresponding classifier is defined by

$$\hat{M} = \operatorname{argmax}_m F(\mathbf{q}; m). \quad (5.18)$$

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 $g_m(u) \equiv -\log p_m(F^{(m)} \geq u)$, where p_m is the probability associated with the m -fold score $F^{(m)}$ under the null hypothesis of the earlier sections. The motivation for this choice is that, in this case, the value

$$\exp(-g_m(F(\mathbf{q}; m))) = p_m(F^{(m)} > F(\mathbf{q}; m)) = p_m \quad (5.19)$$

is the p -value associated with rejecting the null hypothesis of the earlier sections in favor of m -fold symmetry. Consequently, this choice leads to the classification rule

$$\hat{M} = \operatorname{argmax}_m (-\log p_m) = \operatorname{argmin}_m p_m, \quad (5.20)$$

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 5.6.

5.6 Experimental Results

Our algorithm to detect symmetries has been programmed as a plug-in for ImageJ [139]. 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, ..., 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` library². To evaluate the performance of 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 two-fold symmetries or junctions.

5.6.1 Robustness Against Background Signal

To make the detection task more realistic, we have used the model described in Section 4.2.2 to generate a background signal that is added to the patterns of interest to simulate the autofluorescence that is present in real micrographs.

First, we generate a series of (1024×1024) test images populated by three-, four-, and six-fold junctions at fixed locations. The ground-truth data contains 610 three-fold, 921 four-fold, and 304 six-fold symmetric junctions, respectively. We corrupt the images with isotropic Brownian motion, with a mean of zero and σ of 0 to 100. An illustration of the test images (a cropped version of them) can be seen in Figure 5.2. 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 (RMS). The RMS error is computed for the matched detections. The definition of the Jaccard index is given in the Nomenclature. An illustration of the results (crop of one of the test images) can be seen in Figure 5.2, and the Jaccard index in Figure 5.3. Based on the graphs, we claim that our algorithm performs well under a wide range of background noise.

²<http://www.jsc.nildram.co.uk/api/jsc/distributions/Beta.html>




Figure 5.2: From left to right and from top to bottom: test image corrupted with isotropic Brownian motion (mean 0, $\sigma = 10$); corresponding detections; image corrupted with isotropic Brownian motion (mean 0, $\sigma = 50$); corresponding detections. The junctions are at least 25 pixels away from each other.

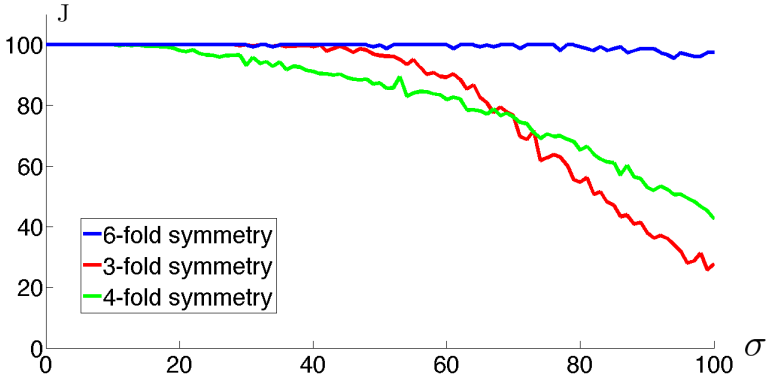


Figure 5.3: Jaccard index under isotropic Brownian motion, as a function of σ .

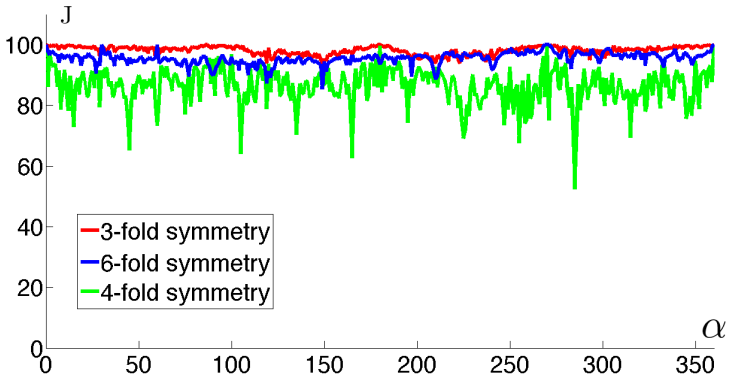


Figure 5.4: Jaccard index under isotropic Brownian motion of zero mean and $\sigma = 10$ for rotations α from 0 to 360 degrees.

5.6.2 Rotation Invariance

We generate another series of (1024×1024) 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 $\sigma = 10$, as illustrated on Figure 5.2) on these images. The ground-truth data contains approximately 600 three-fold, 900 four-fold, and 300 six-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 5.4) of the detections. Based on the graph, we confirm that our method is essentially rotation invariant and performs well independently of the orientation of the junctions.

5.6.3 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 5.5. 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 5.6 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.

5.6.4 Micrographs

Figure 5.7 features hexagonal embryonic stem cells imaged by fluorescence microscopy with three-fold symmetric junctions. 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 5.8 contains an atomic-resolution ADF-STEM image of hexagonal graphene crystals with three-fold symmetric junctions. The detection task is more challenging than in the embryonic stem-cell image due to blurring effects. Here the image quality

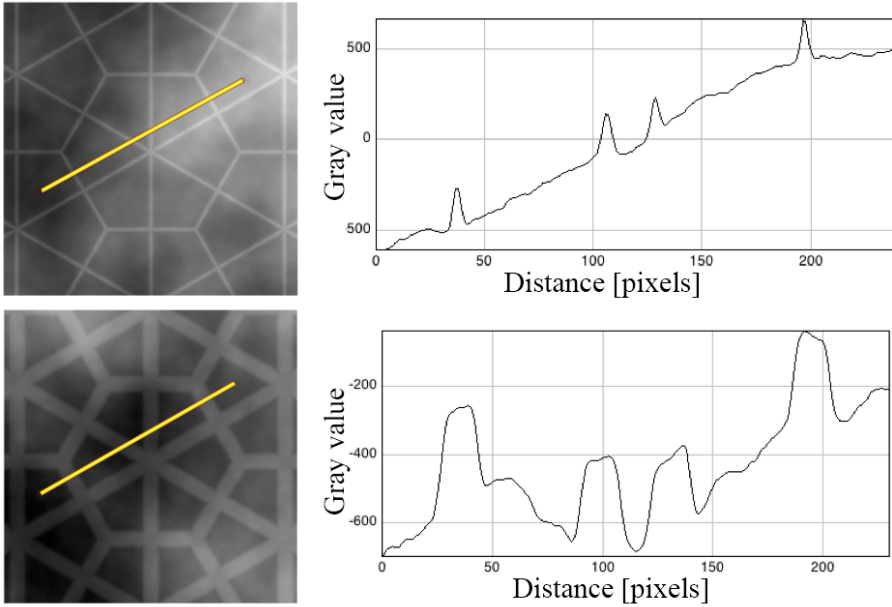


Figure 5.5: Test images of different line thicknesses, corrupted with isotropic Brownian motion (mean 0, $\sigma = 10$). From left to right: test image; profile plot along the corresponding bright line. From top to bottom: line thickness $l = 1.0$, line thickness $l = 10.0$, respectively.

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.

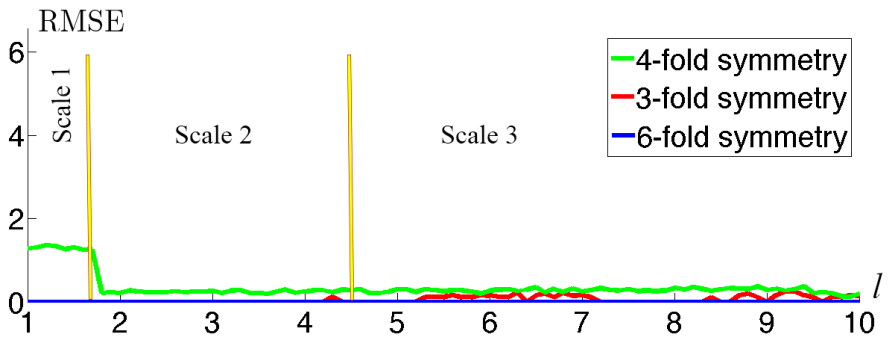


Figure 5.6: Position error in the sense of root-mean-square error (RMSE) under isotropic Brownian motion of zero mean and $\sigma = 10$ for line thicknesses 1 to 15 with steps 0.1. Scale 1, 2, and 3 indicates the scale at which the detections are made.

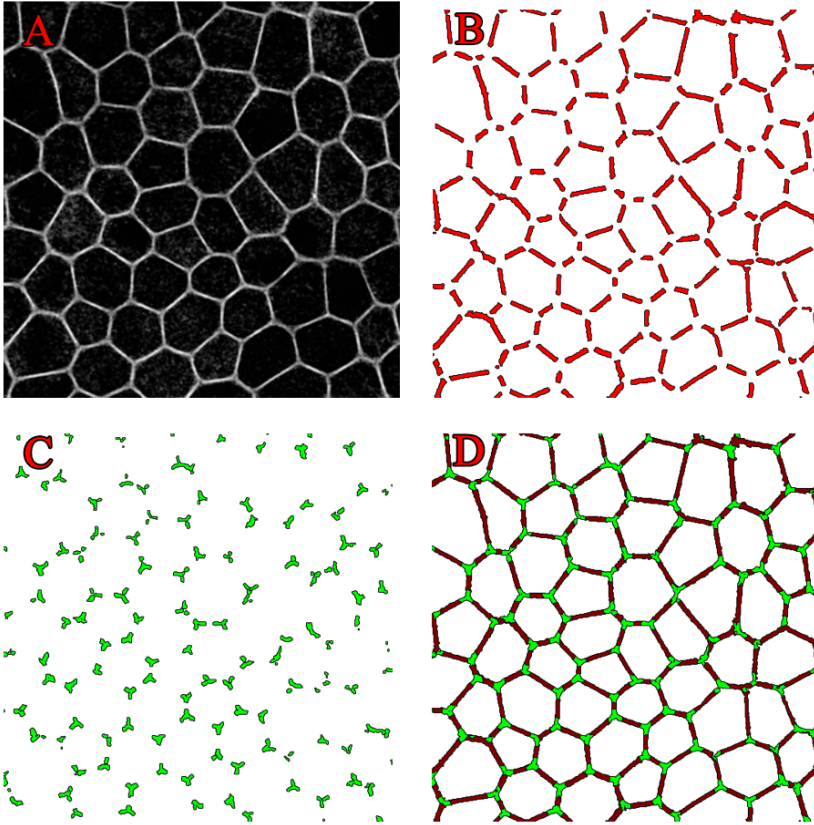


Figure 5.7: Hexagonal embryonic stem cells in light microscopy [31]. Detection of two-fold symmetries or ridges and three-fold junctions. Subfigures: A: original image, B: detected ridges (red), C: detected three-fold junctions (green), D: combined image. Size of the original image: 512×512 pixels; the junctions are at least 20 pixels away from each other.




Figure 5.8: Atomic-resolution ADF-STEM image of hexagonal graphene crystals [140]. Detection of two-fold symmetries or ridges (red) and three-fold junctions (green). Subfigures: original image, combined RGB image. Size of the original image: 256×256 pixels; the junctions are at least 15 pixels away from each other.

5.7 Conclusion

In this chapter 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.

5.8 Notes

Proof of Theorem 5.1

For a given scale i and position \mathbf{x}_0 , we consider the two harmonics n_1, n_2 . We use (18) 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

$$\begin{aligned}
 & \mathcal{B}_S \{ \psi_{n_1, i}(\cdot - \mathbf{x}_0), \psi_{n_2, i}(\cdot - \mathbf{x}_0) \} \\
 &= \frac{2^{2i}}{(2\pi)^2} \int \overline{\hat{\psi}_{n_1}(2^i \boldsymbol{\omega})} \hat{\psi}_{n_2}(2^i \boldsymbol{\omega}) P_S(\boldsymbol{\omega}) d\boldsymbol{\omega} \\
 &= \frac{2^{2i}}{(2\pi)^2} \int_0^\infty \int_{-\pi}^\pi P_S(\boldsymbol{\omega}) \left| \hat{h}(2^i \boldsymbol{\omega}) \right|^2 e^{-jn_1 \varphi} e^{jn_2 \varphi} \boldsymbol{\omega} d\boldsymbol{\omega} d\varphi \\
 &= \frac{1}{(2\pi)^2} \int_0^\infty \int_{-\pi}^\pi P_S\left(\frac{\boldsymbol{\omega}}{2^i}\right) \left| \hat{h}(\boldsymbol{\omega}) \right|^2 e^{-jn_1 \varphi} e^{jn_2 \varphi} \boldsymbol{\omega} d\boldsymbol{\omega} d\varphi \\
 &= \frac{1}{(2\pi)^2} \int_0^\infty P_S\left(\frac{\boldsymbol{\omega}}{2^i}\right) \left| \hat{h}(\boldsymbol{\omega}) \right|^2 \boldsymbol{\omega} d\boldsymbol{\omega} \int_{-\pi}^\pi e^{j(n_2 - n_1) \varphi} d\varphi \\
 &= C_i \int_{-\pi}^\pi e^{j(n_2 - n_1) \varphi} d\varphi, \tag{5.21}
 \end{aligned}$$

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 2π if $n_1 = n_2$ and vanishes otherwise. This means that the wavelet coefficients $\{ \langle S, \psi_{n, i}(\cdot - \mathbf{x}_0) \rangle \}_n$ are uncorrelated with common variance

$$\frac{1}{2\pi} \int_0^\infty P_S\left(\frac{\boldsymbol{\omega}}{2^i}\right) \left| \hat{h}(\boldsymbol{\omega}) \right|^2 \boldsymbol{\omega} d\boldsymbol{\omega}. \tag{5.22}$$

In the Gaussian case, uncorrelation is equivalent to independence.

Chapter 6

Steerable Wavelet Design

6.1 Overview

In Chapter 5, within the framework of 2D steerable wavelet frames, we proposed a novel method for the detection of local symmetries in a template-free way. By template-free we meant that there was no need to design a specific junction, or ridge-shaped template, nor to align it with the pattern of interest. Our detection scheme was multiscale due to the wavelet-type representation. In this chapter, we follow a complementary approach that serves to address two objectives not covered by the methods of the previous chapter. First, we target the detection of junctions and patterns without constraining ourselves to symmetric patterns only. Secondly, we aim not only to detect these local patterns, but also to estimate their orientation. These dual objectives are achieved by designing steerable wavelets that are shaped to match a desired angular profile. With such wavelet schemes we are not only able to detect local symmetry centers and junction points, but also to determine the local orientation of the patterns of interest.

Our methodology introduces a novel wavelet design scheme along with an innovative detection algorithm based on analytical optimization. Within the framework of 2D steerable wavelet frames, we define wavelets that are polar-separable in the Fourier domain. The decomposition involves a pyramid characterized by a radial wavelet function and a set of directional components encoded using circular harmonics. The tightness

of the generated frame ensures the perfect reconstruction property ¹.

In Section 6.2, we recall the principle of steerability and the construction of steerable wavelet frames, which provide the analytical framework used in this chapter. In Section 6.3, we provide a brief summary of the work of Slepian *et al.* [141–143] that we will adapt to our framework. Then, in Section 6.4, we examine the design of M -fold symmetric steerable wavelet frames both in the space and the Fourier domains. In Section 6.5, we show that our symmetric steerable wavelet design can be generalized to other classes of local structures and junctions. In Section 6.6, we evaluate the different design methods based on the ease of generation of symmetric patterns. We then present in Section 6.7 an algorithm to detect centers of symmetry. We finally apply our method to synthetic and real biological micrographs in Section 6.8.

6.2 Concept

Once more, we rely on the circular harmonic wavelets that have been presented in Section 3. To put our approach in context, we summarize the basic ideas of the construction of such wavelet frames.

We start our design from a bandlimited isotropic mother wavelet over \mathbb{R}^2 , whose shifts and dilations form a wavelet frame. Proposition 3.1 in Chapter 2 states sufficient conditions for such primal profiles, specific examples are given in Table 3.1.

We apply the multiorder complex Riesz transform on our isotropic wavelet frame to obtain circular harmonic wavelets. We recall that, the circular harmonic wavelet at scale i , location $\mathbf{0}$, and harmonic channel n has the Fourier transform

$$\hat{\xi}_{n,i}(\omega, \varphi) = 2^i \hat{h}(2^i \omega) e^{jn\varphi}. \quad (6.1)$$

Consequently, in the spatial domain, we have

$$\xi_n(r, \theta) = \frac{j^n e^{jn\theta}}{2\pi} \int_0^\infty \omega \hat{h}(\omega) J_n(r\omega) d\omega, \quad (6.2)$$

where J_n denotes the n -th Bessel function of the first kind.

¹The results in this chapter are based on the article [105]: Z. Püspöki, V. Uhlmann, C. Vonesch, M. Unser, “Design of Steerable Wavelets to Detect Multifold Junctions”, *IEEE Transactions on Image Processing*, vol. 25, no. 2, pp. 643-657, February 2016.

The N -channel tight frame of wavelets are generated as $\{\xi_{n,i,\mathbf{k}} = \mathcal{F}^{-1}\{\hat{\xi}_{n,i,\mathbf{k}}\}\}_{n \in H}$, where the values of n are distinct and taken from a predefined set $H = \{n_1, \dots, n_N\} \subset \mathbb{Z}$. We take advantage of these wavelets to construct new steerable representations that detect features of interest in a multiscale way.

The new steerable frame functions are obtained by a mapping with an orthogonal shaping matrix \mathbf{U}

$$\begin{bmatrix} \psi_{n_1,i,\mathbf{k}} \\ \vdots \\ \psi_{n_N,i,\mathbf{k}} \end{bmatrix} = \mathbf{U} \begin{bmatrix} \xi_{n_1,i,\mathbf{k}} \\ \vdots \\ \xi_{n_N,i,\mathbf{k}} \end{bmatrix}. \quad (6.3)$$

The new wavelets span the same space as the wavelet frame $\{\xi_{n,i,\mathbf{k}}\}$. The shaping matrix \mathbf{U} endows the wavelet functions $\psi_{n,i,\mathbf{k}}$ with a desired angular profile. We note that, once wavelets with the designed angular profile have been designed, they can be rotated arbitrarily and continuously as they inherit the steerability property of the circular harmonic wavelets.

There are several ways to design the shaping matrix. We propose to define an angular energy concentration like $\sum_n u_n \xi_n$ (corresponding to a single row of \mathbf{U} , thus to a single wavelet channel) and to optimize it with respect to the angular weights u_n that control the concentration of the energy along the desired pattern, with $\mathbf{u} = (u_1, u_2, \dots, u_N)$ and $\mathbf{u}^H \mathbf{u} = 1$. This formulation is reminiscent of what appears in the works of Slepian *et al.* [141–143]. He investigated the connection between the concentration of functions in the time and frequency domains in several settings (continuous, multi-dimensional, and discrete), and considered solutions that were optimal in different ways: by being optimally concentrated in one domain and band-limited in the other, or being optimally concentrated in both domains. In Section 6.3 we recall some of his results, since they provide constructive ideas for our steerable wavelet design.

6.3 Optimal Concentrations of Energy

Let us denote an index-limited (finite) sequence of length N by $\{a_n\}$, such that $a_n = 0$ for all $n < N_0$ and $n > N_0 + N - 1$. The question to answer is, what is the maximal concentration it can achieve in frequency within a bandwidth of $2B$?

Slepian defines bandwidth concentration as

$$\tilde{\mu} = \frac{\int_{-B}^B |A(\rho)|^2 d\rho}{\int_{-\frac{1}{2}}^{\frac{1}{2}} |A(\rho)|^2 d\rho}, \quad (6.4)$$

where $A(\rho) = \sum_n a_n e^{2\pi j n \rho}$ is the discrete-time Fourier transform of the sequence $\{a_n\}$ (parametrized here by the frequency $\rho \in [-\frac{1}{2}, \frac{1}{2})$ instead of ω). This formula quantifies the ratio of energy of $A(\rho)$ contained in the interval $[-B, B]$ compared to its total energy. To maximize the bandwidth concentration defined in equation (6.4), we note two things. First, the denominator of the above expression is equal to $\sum_n |a_n|^2$ by Parseval. Second, the numerator can be re-expressed as

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} |A(\rho)|^2 \text{rect}\left(\frac{\rho}{2B}\right) d\rho. \quad (6.5)$$

We replace $|A(\rho)|^2$ by the product

$$A(\rho) \overline{A(\rho)} = \left(\sum_n a_n e^{2\pi j n \rho} \right) \overline{\left(\sum_m a_m e^{2\pi j m \rho} \right)} = \sum_{n,m} a_n \overline{a_m} e^{2\pi j (n-m)\rho}, \quad (6.6)$$

multiply by $w(\rho)$ and integrate from $-\frac{1}{2}$ to $\frac{1}{2}$, and use the rect-sinc Fourier transform formula

$$\int_{-\frac{1}{2}}^{\frac{1}{2}} e^{2\pi j (n-m)\rho} \text{rect}\left(\frac{\rho}{2B}\right) d\rho = \frac{\sin(2\pi B(n-m))}{\pi(n-m)}. \quad (6.7)$$

These steps allow us to write the numerator as a quadratic form in the finite sequence (vector) $\{a_n\}$, namely

$$\sum_{m,n} \frac{\sin(2\pi B(n-m))}{\pi(n-m)} \overline{a_m} a_n. \quad (6.8)$$

The numerator can thus be written in matrix form as $\mathbf{a}^H \mathbf{P} \mathbf{a}$ for a matrix \mathbf{P} (defined by the sinc above), and the denominator as $\|\mathbf{a}\|_2^2$. The problem is then to maximize the ratio

$$\tilde{\mu} = \frac{\mathbf{a}^H \mathbf{P} \mathbf{a}}{\|\mathbf{a}\|_2^2}. \quad (6.9)$$

We rewrite the problem in terms of the unit vector $\mathbf{u} = \frac{\mathbf{a}}{\|\mathbf{a}\|}$. Dividing the numerator and denominator by $\|\mathbf{a}\|^2$ we get the equivalent problem of maximizing

$$\mathbf{u}^H \mathbf{P} \mathbf{u} \quad \text{subject to} \quad \|\mathbf{u}\|^2 = 1. \quad (6.10)$$

The solution of the above problem is then the maximal eigenvector of \mathbf{P} . Slepian *et al.* arrive at this solution from an eigenvalue problem for a degenerate integral equation instead of a matrix one, but we can show it equivalently as follows.

Since \mathbf{P} is symmetric, its eigenvectors form an orthonormal basis $\{\mathbf{v}_n\}$ for \mathbb{R}^N . Expanding \mathbf{u} in this basis, we get for the objective function,

$$\begin{aligned} \mathbf{u}^H \mathbf{P} \mathbf{u} &= \left(\sum_m c_m \mathbf{v}_m \right)^H \mathbf{P} \left(\sum_n c_n \mathbf{v}_n \right) = \left(\sum_m c_m \mathbf{v}_m \right)^H \left(\sum_n c_n \lambda_n \mathbf{v}_n \right) \\ &= \sum_{m,n} \lambda_n \overline{c_m} c_n \mathbf{v}_m^H \mathbf{v}_n = \sum_n \lambda_n |c_n|^2, \end{aligned} \quad (6.11)$$

where in the last step we have used the orthonormality of the basis, and in two steps before it we have used the definition of an eigenvector (the λ_n s being corresponding eigenvalues).

The above derivation shows that to maximize $\mathbf{u}^H \mathbf{P} \mathbf{u}$, we should concentrate the coefficient vector $\{c_n\}$ on the largest eigenvalue (let's assume this is λ_1), whereby we get the solution

$$\mathbf{u} = \sum_n c_n \mathbf{v}_n = c_1 \mathbf{v}_1. \quad (6.12)$$

Furthermore, since \mathbf{u} is a unit vector, $|c_1| = 1$ and we can simply take $\mathbf{u} = \mathbf{v}$ (maximal eigenvector) as our solution.

Our formulation to define wavelets whose energy is concentrated along specific angular directions is similar to this formulation. In particular, its solution is found similarly, through an eigen decomposition.

6.4 Design of M -fold-symmetric Wavelets

We can formulate our energy optimization either in the spatial or in the Fourier domain. As we shall see shortly, both designs are appropriate to generate symmetric junction templates; however, each one has its own advantages and disadvantages. We examine and discuss both cases in Sections 6.4.1 and 6.4.2, respectively.

6.4.1 Design of Steerable Wavelets in the Space Domain

To design a wavelet that responds to M -fold symmetries, we propose to impose the same M -fold-symmetric pattern on ψ . We achieve this by minimizing a well chosen energy functional. As an illustration, in the case of a unimodal detector, we are looking for the angular profile that is most concentrated around the angle 0. This means that we are targeting a minimum-variance solution, which results in an energy that concentrates around a chosen axis.

For a wavelet centered at the origin, the spatial quadratic energy term we propose to minimise has the form

$$E\{\psi\} = \frac{1}{2\pi} \int_0^\infty \int_{-\pi}^\pi |\psi(r, \theta)|^2 w(\theta) d\theta r dr, \quad (6.13)$$

where w is an M -fold-symmetric nonnegative weighting function with M equidistant minima on the unit circle that favors M -fold-symmetric solutions. Minimizing E thus forces the solution ψ to be localized symmetrically around the M -fold minima. Once the mother wavelet ψ is found, its translates and dilates will naturally share the optimal angular profile around their center. We note that, the formulation here is reminiscent of the one of Slepian, however, we use a more general window function instead of a rect (see Section 6.3).

By expanding ψ as $\sum_n u_n \xi_n$ and imposing a unit norm on \mathbf{u} , this formalism boils down to a quadratic optimization problem with quadratic constraints similar to the one discussed in Section 6.3, which can be solved through an eigen decomposition. During the minimization process, the radial part of the wavelet function results in Hankel-like integrals. We take advantage of (6.2) to determine the quadratic optimization problem.

By replacing $\sum_{n \in H} u_n \xi_n(r, \theta) = \sum_{n \in H} u_n \eta_n(r) e^{jn\theta}$ for $\psi(r, \theta)$ in the energy functional, we arrive at the form

$$E\{\psi\} = \mathbf{u}^H \mathbf{W} \mathbf{u}, \quad (6.14)$$

where the entries of \mathbf{W} are given by

$$W_{n, n'} = \frac{1}{2\pi} \int_{-\pi}^\pi e^{j(n-n')\theta} w(\theta) d\theta \int_0^\infty \eta_n(r) \overline{\eta_{n'}(r)} r dr. \quad (6.15)$$

The radial and angular factors are computed separately. The angular factor

$$W(n - n') := \frac{1}{2\pi} \int_{-\pi}^\pi e^{j(n-n')\theta} w(\theta) d\theta \quad (6.16)$$

is the $(n - n')$ th Fourier coefficient of $w(\theta)$, while the factor

$$\begin{aligned} H_{\text{sp.}}(n, n') &:= \int_0^\infty \eta_n(r) \overline{\eta_{n'}(r)} r dr \\ &= \frac{j^{n-n'}}{(2\pi)^2} \int_0^\infty \left(\int_0^\infty \omega \hat{h}(\omega) J_n(r\omega) d\omega \right) \left(\int_0^\infty \overline{\omega \hat{h}(\omega)} J_{n'}(r\omega) d\omega \right) r dr \end{aligned} \quad (6.17)$$

is a multiplicative term that corresponds to the effect of the radial part. Finally, the entries of \mathbf{W} are computed as

$$\mathbf{W}_{n,n'} = H_{\text{sp.}}(n, n') W(n - n'). \quad (6.18)$$

It turns out that \mathbf{W} is a Hermitian-symmetric matrix with entries indexed by $n, n' \in H$.

With the previous derivation, we end up with a quadratic optimization problem (minimization of $E = \mathbf{u}^H \mathbf{W} \mathbf{u}$) with quadratic constraints ($\mathbf{u}^H \mathbf{u} = 1$), which is the same as the problem we studied in 6.3. Minimizing the energy functional is equivalent to finding the eigenvector that corresponds to the smallest eigenvalue of \mathbf{W} . We note that the matrix \mathbf{W} is positive-semi-definite since it is derived from a nonnegative weight function, which itself defines a nonnegative energy (Bochner's theorem).

A particular example of a three-fold-symmetric wavelet obtained by the proposed method can be seen in Figure 6.1. The weight function w is $\frac{2\pi}{3}$ -periodic, with $w(\theta) = \theta^2$ for $|\theta| \leq \frac{\pi}{3}$ to achieve maximal energy concentration about the minima in the sense of variance. An example of a complex three-fold-symmetric wavelet obtained by the proposed method is presented in Figure 6.1.

6.4.2 Design of Steerable Wavelets in the Fourier Domain

Since rotations and symmetries in the space domain carry over to the frequency domain, we can formulate our minimization problem in either one. The advantage of the frequency-domain formulation is to avoid the computation of Hankel-like integrals. This happens because, in the Fourier domain, the radial part of the wavelet has no effect on the optimization.

The Fourier-domain energy functional E is defined as

$$E\{\hat{\psi}\} = \frac{1}{2\pi} \int_0^\infty \int_{-\pi}^{+\pi} |\hat{\psi}(\omega, \varphi)|^2 w(\varphi) d\varphi \omega d\omega, \quad (6.19)$$

where $w \geq 0$ is an M -fold-symmetric weighting function with M equidistant minima on the unit circle. As stated in the Section 6.4.1, minimizing E will force the solution $\hat{\psi}$ to be localized symmetrically near the M -fold minima.

The polar factorization of $\hat{\psi}$ can be written as

$$\hat{\psi}(\omega, \varphi) = \sum_{n \in H} u_n \hat{\xi}_n(\omega, \varphi) = \hat{h}(\omega) \sum_{n \in H} u_n e^{jn\varphi}. \quad (6.20)$$

The energy functional (6.19) therefore factorizes as

$$E\{\hat{\psi}\} = \int_0^\infty |\hat{h}(\omega)|^2 \omega d\omega \cdot \frac{1}{2\pi} \int_{-\pi}^{\pi} \left| \sum_{n \in H} u_n e^{jn\varphi} \right|^2 w(\varphi) d\varphi. \quad (6.21)$$

This means that the radial part of the wavelet function has no role in the optimization problem since it is factored out of the integral. As a result, we can optimize the angular profile of the wavelet by defining the Fourier-based energy for the angular factor alone as

$$E_{\text{ang}}\{\hat{\psi}\} = \frac{1}{2\pi} \int_{-\pi}^{\pi} \left| \sum_{n \in H} u_n e^{jn\varphi} \right|^2 w(\varphi) d\varphi. \quad (6.22)$$

Consequently, formulating the energy term in the Fourier domain results in a substantial simplification of the problem. After replacing $\hat{\psi}(\varphi)$ by its expansion $\sum_{n \in H} u_n e^{jn\varphi}$, we have that

$$\begin{aligned} E_{\text{ang}}\{\hat{\psi}\} &= \sum_{n' \in H} \sum_{n \in H} \overline{u_{n'}} u_n \frac{1}{2\pi} \int_{-\pi}^{+\pi} e^{j(n-n')\varphi} w(\varphi) d\varphi \\ &= \sum_{n' \in H} \sum_{n \in H} \overline{u_{n'}} W(n-n') u_n = \mathbf{u}^H \mathbf{W} \mathbf{u}, \end{aligned} \quad (6.23)$$

where

$$W(n-n') = \frac{1}{2\pi} \int_{-\pi}^{+\pi} e^{j(n-n')\varphi} w(\varphi) d\varphi \quad (6.24)$$

is the $(n' - n)$ th Fourier-series coefficient of w (as in (6.16)) and \mathbf{W} is the corresponding Hermitian-symmetric matrix with entries indexed by $n, n' \in H$. Finally, and similarly to the space-domain design, we end up with an eigenvalue problem, but for a matrix that is much easier to compute.

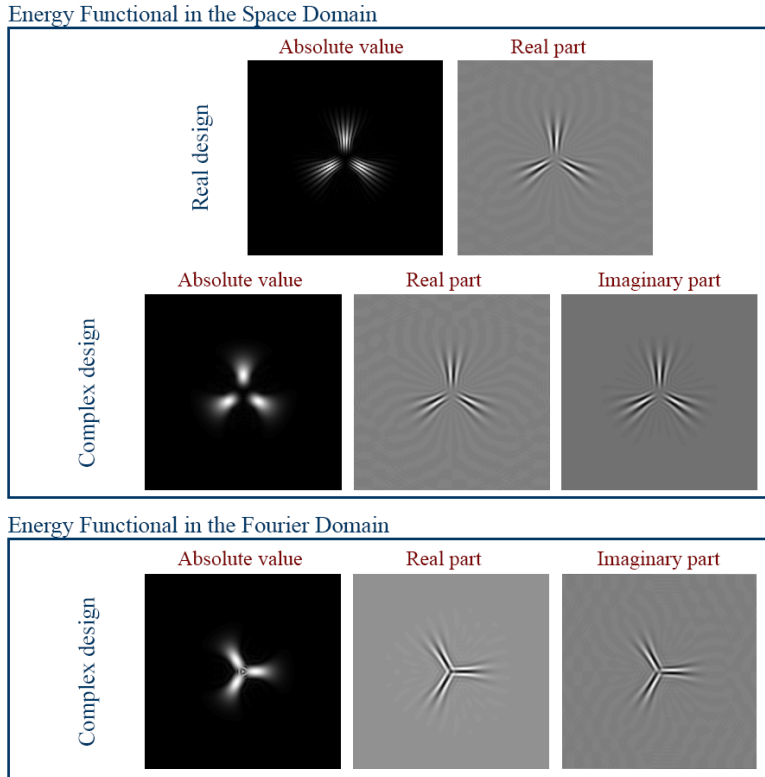


Figure 6.1: First row: three-fold-symmetric wavelet designed by the energy functional in the space domain and visualized in the space domain. From left to right: absolute value and real part. The imaginary part is 0. Harmonics: $H = \{3n : n = -8, \dots, 8\}$. Second row: three-fold-symmetric complex wavelet designed by the energy functional in the space domain and visualized in the space domain. From left to right: absolute value, real, and imaginary parts. Harmonics: $H = \{3n : n = 0, \dots, 9\}$. Third row: Three-fold-symmetric wavelet designed by the energy functional in the Fourier domain and visualized in the space domain. From left to right: magnitude, real, and imaginary parts. Harmonics: $H = \{3n : n = 0, \dots, 9\}$.




Figure 6.2: Magnitude of four-fold-, five-fold-, and six-fold-symmetric wavelets designed by the energy functional in the Fourier domain and visualized in the space domain. Harmonics: $H = \{Mn : n = 0, \dots, 9\}$.

A particular example of a three-fold-symmetric wavelet obtained by the proposed method can be seen in Figure 6.1. The weight function w is $\frac{2\pi}{M}$ -periodic, with $w(\varphi) = \varphi^2$ for $|\varphi| \leq \frac{\pi}{M}$ to achieve maximal energy concentration about the minima in the sense of variance. Examples of higher-order symmetric wavelets can be seen in Figure 6.2.

For implementation purposes, the wavelets are projected on bandlimited basis functions that have a certain cutoff frequency. This suppresses all patterns above that frequency and explains why the center of our wavelets vanish. This effect can be seen in Figures 6.1 and 6.2. When the number of harmonics is high, we do not have enough bandwidth to represent the directional patterns within a disk of radius R . (Based on the Shannon-Nyquist sampling theorem, this occurs when $R < \frac{\sqrt{2}}{\pi}N$, where N is the highest harmonic.) This phenomenon can be seen in Figure 6.1 at the center of the wavelet. However, we note that i) this phenomenon does not reduce the detection property of the wavelets; ii) one typically does not apply a large number of harmonics, also due to computational cost.

6.5 Angular Weights and Asymmetric Wavelets

Our framework is flexible and easy to generalize to many classes of local structures and junctions. Here, we take another look at the angular weighting functions and propose

to design asymmetric wavelets, targeting familiar shapes like corners or T-junctions. The advantage, as before, is that we can “steer” these wavelets to any arbitrary angle by a systematic complex rescaling of the wavelet coefficients. The design of asymmetric wavelets is possible only in the space domain, since the rotational equivalence of the space and Fourier domains is no longer relevant for asymmetric profiles.

We begin our formulation from (6.18), by recalling the nature of the angular weighting function for the symmetric case. Then we adapt the angular factor W and note that it corresponds to the (time-reversed) Fourier series of the angular weight function w . Thus, we achieve the desired shape by choosing w such that its local extrema are placed at angles that correspond to the desired branches of the wavelet.

In the case of the symmetric three-fold example in Sections 6.4.1, we had identified a profile that, after proper steering, was simultaneously concentrated around zero, $\frac{2\pi}{3}$, and $-\frac{2\pi}{3}$. This function is given by the parabola θ^2 over the interval $[-\pi/3, \pi/3]$, and its translates to the intervals $[\pi/3, \pi]$ and $[-\pi/3, -\pi]$. The function w_{S3} on the interval $[-\pi, \pi]$ is illustrated in Figure 6.3. It has the formula

$$w_{S3}(\theta) = C_{S3} (\theta^2 \mathbf{1}_{[-\pi/3, \pi/3]}(\theta) + (\theta - 2\pi/3)^2 \mathbf{1}_{[\pi/3, \pi]}(\theta) + (\theta + 2\pi/3)^2 \mathbf{1}_{[-\pi, -\pi/3]}(\theta)). \quad (6.25)$$

Here, C_{S3} is the normalization constant $(3/\pi)^2$. The corresponding kernel is

$$W_{S3}(l) = C_{S3} \frac{(1 + 2 \cos(\frac{2\pi}{3}l)) ((\pi^2 l^2 - 18) \sin(\frac{\pi}{3}l) + 6\pi l \cos(\frac{\pi}{3}l))}{9l^3 \pi}. \quad (6.26)$$

In the case of a T junction, we intend to identify profiles that, after proper steering, are simultaneously concentrated around zero, $\pi/2$, and $-\pi/2$. Similarly to the symmetric case, using the energy minimization paradigm, we take w to have three identical minima at 0, $\pi/2$, and $-\pi/2$, given by the parabola θ^2 over the interval $[-\pi/4, \pi/4]$ and its translates to the intervals $[\pi/4, 3\pi/4]$ and $[-3\pi/4, -\pi/4]$. We keep w_T constant over the intervals $[3\pi/4, \pi]$ and $[-\pi, -3\pi/4]$. The function w_T on the interval $[-\pi, \pi]$ is illustrated in Figure 6.3. Specifically,

$$\begin{aligned} w_T(\theta) = C_T & (\theta^2 \mathbf{1}_{[-\pi/4, \pi/4]}(\theta) + (\theta - \pi/2)^2 \mathbf{1}_{[\pi/4, 3\pi/4]}(\theta) \\ & + (\theta + \pi/2)^2 \mathbf{1}_{[-3\pi/4, -\pi/4]}(\theta) + \mathbf{1}_{[-\pi, -3\pi/4]}(\theta) + \mathbf{1}_{[3\pi/4, \pi]}(\theta), \end{aligned} \quad (6.27)$$

where C_T is the normalization constant $(4/\pi)^2$. The corresponding kernel is

$$W_T(l) = \frac{\pi^2 l^2 \sin(\pi l) + 8\pi l (\cos(\frac{3\pi}{4}l) + 2 \cos(\frac{\pi}{4}l)) - 32 \sin(\frac{3\pi}{4}l)}{\pi^3 l^3}. \quad (6.28)$$

The wavelet obtained by the proposed method can be seen in Figure 6.4.

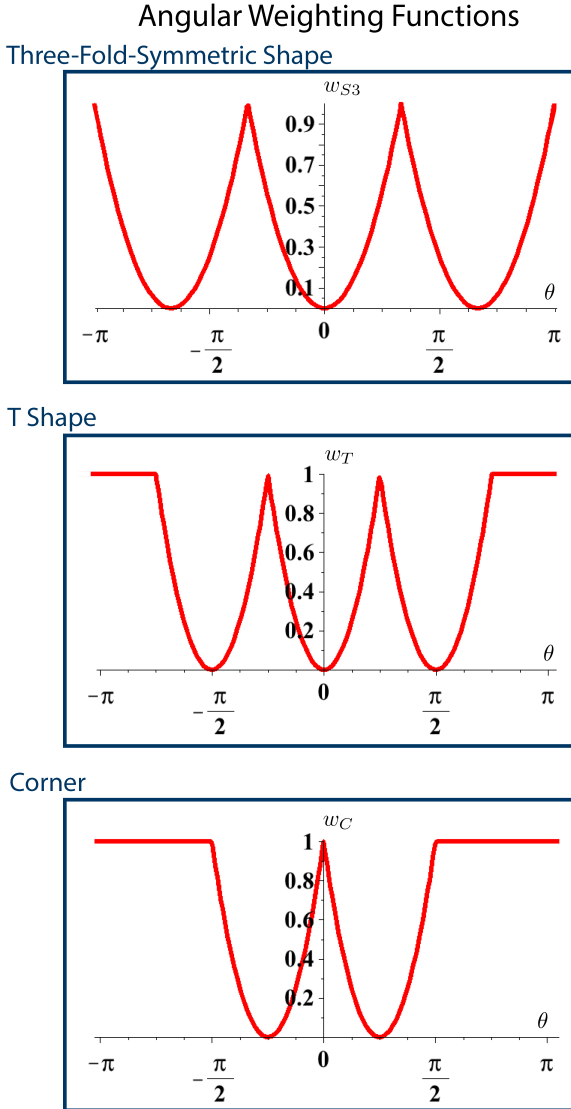
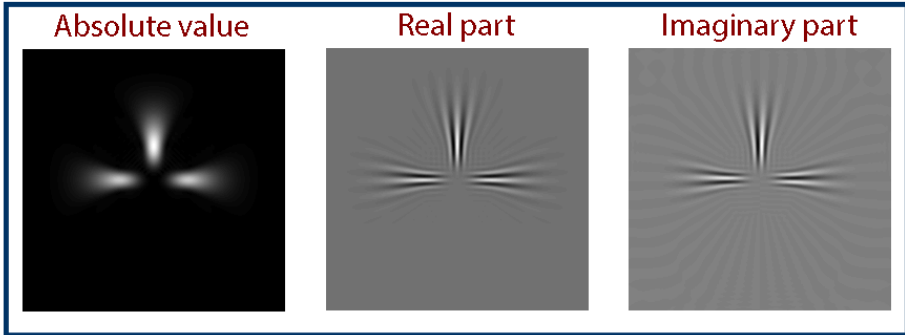


Figure 6.3: From top to bottom: angular weighting function for the design of three-fold-symmetric wavelets; angular weighting function for the design of T-shaped wavelets; angular weighting function for the design of corner-shaped wavelets.

T Shape



Corner

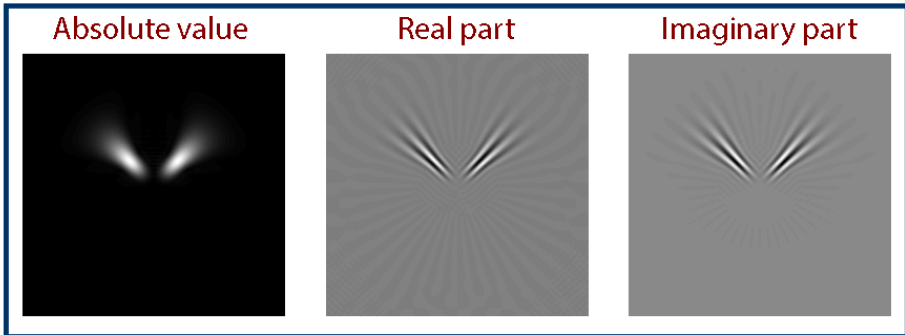


Figure 6.4: First row: T-shaped asymmetric wavelet in the space domain. From left to right: absolute value, real, and imaginary parts, respectively. Harmonics: $H = \{0, \dots, 40\}$. Second row: corner-shaped asymmetric wavelet in the space domain. From left to right: absolute value, real, and imaginary parts, respectively. Harmonics: $H = \{0, \dots, 40\}$.

Similarly, for the perpendicular corner, the weight function consists of two identical parabolas placed at $-\pi/4$ and $\pi/4$. We keep w_C constant over the intervals $[\pi/2, \pi]$ and

Table 6.1: Comparison of designs.

Property	Fourier-domain design	Space-domain design
Real odd symmetric wavelets	Not possible	Possible
Design of asymmetric wavelets	Not possible	Possible
Computational complexity	Analytical results	Numerical integration needed
Ease of design	Fast	Slower
Design of arbitrary orders of symmetry	By periodization of a unifold design	By solving an eigen decomposition for each order

$[-\pi, -\pi/2]$. This leads to the kernel illustrated in Figure 6.3 with

$$w_C(\theta) = C_C \left((\theta - \pi/4)^2 \mathbf{1}_{[0, \pi/2]}(\theta) + (\theta + \pi/4)^2 \mathbf{1}_{[-\pi/2, 0]}(\theta) \right) + \mathbf{1}_{[-\pi, -\pi/2]}(\theta) + \mathbf{1}_{[\pi/2, \pi]}(\theta), \quad (6.29)$$

where C_C is the normalization constant $(4/\pi)^2$. The corresponding kernel is

$$W_C(l) = \frac{\pi^2 l^2 \sin(\pi l) - 32 \sin\left(\frac{\pi}{2} l\right) + 8\pi l \left(1 + \cos\left(\frac{\pi}{2} l\right)\right)}{\pi^3 l^3}. \quad (6.30)$$

The wavelet obtained by the proposed method and its weight function can be seen in Figures 6.3 and 6.4.

6.6 Comparison of Different Designs

The main features of the two designs of Section 6.4 are summarized in Table 6.1. We stress that the characteristics that address speed pertain only to the initial design of the wavelets (calculation of the design matrix \mathbf{U}). Once the wavelets have been designed, the analysis and processing of images runs at the same speed for a given number of harmonics, independent of the design.

Both approaches provide a comprehensive basis to detect symmetric junctions. The Fourier-domain design is the easiest, since the radial part of the wavelet has no effect on the optimization process. The space-domain design is more cumbersome because it requires one to handle the Hankel-like integrals associated with the radial part of the wavelet. In return, it sidesteps the limitations of the Fourier domain that forbid wavelets that are asymmetric and real. Also, it allows the design of asymmetric shapes.

Proposition 6.1 *No real-valued wavelets with odd symmetries can result from minimizing (6.19) in the Fourier domain.*

The proof of Proposition 6.1 is given in Section 6.10.

While the Fourier design is not appropriate when odd-symmetric real wavelets are desired, this does not impair the detection of junctions as complex odd designs are still possible (using positive harmonics alone). Also, by defining the energy functional $E\{\psi\}$ in the space domain, it is possible to obtain real wavelets.

The Fourier techniques are not suitable for the design of asymmetric wavelets (e.g., corners, T junctions), since the Fourier design relies on the equivalence of rotational symmetries between the spatial and Fourier domains. On the contrary, the space design provides a free parametric framework.

6.7 Application to Image Analysis

Our junction- and structure-detection algorithm consists of a main part, namely, the detection of key points (i.e., centers of symmetry) and their orientation, followed by an optional connecting of the detected junctions, which is formulated as an optimal-path-finding problem.

6.7.1 Summary of the Algorithm

Our algorithm to detect junctions of a given multiplicity M along with the orientation of the corresponding edges, and the connections between these junctions, can be decomposed in the steps described below. It uses an M -fold steerable wavelet as introduced in the earlier sections.

(1) Wavelet analysis with optimally steered wavelets

We decompose the image with the steerable wavelet. The local orientation θ_0 is determined as the one that maximizes the “detector response” $|\langle I, \psi_{i,k}(\cdot, \cdot + \theta_0) \rangle|$, for each

scale i and location \mathbf{k} . The output of this first stage is a map of maximal steerable wavelet responses $Q(i, \mathbf{k})$ and orientations at every scale i and location \mathbf{k} .

(2) Maximal projection across scales

We perform a maximum-intensity projection to aggregate the key points detected at different scales. The detectors have normalized energy across scales, thus, a maximal projection across scales is a meaningful choice. For each position, we keep the coefficient corresponding to the largest detector response across existing scales, like

$$Q(\mathbf{k}) = \max_i(Q(i, \mathbf{k})). \quad (6.31)$$

(3) Thresholding and the detection of local maxima

We assume that the pixels that correspond to key points are sparse in the image. Points with a detector response smaller than the mean of the whole image are considered to belong to the background and are removed. We additionally apply local non-maximum suppression over a user-predefined window in order to prevent multiple detections of the same junction [40].

(4) Edge computation using dynamic programming (Optional)

Based on the key points detected after Step (3), we use dynamic programming to establish paths between all possible pairs of symmetry centers that are closer than a given threshold. “True” junctions are finally extracted from the set of possible edges by favoring configurations that involve low-cost paths that are coherent with each type of center of symmetry.

The details of Steps (1) and (4) are described below in more detail.

6.7.2 Steering

In this section, we present a novel analytical method that applies efficiently to all detectors. We determine the optimal steering angle by maximizing a trigonometric polynomial in θ_0 , which can be achieved by computing the roots of its first-order derivative.

The key points in the image correspond to maxima in the response of the wavelet detector. The wavelets therefore have to be “steered” to look for the angle that elicits the largest response. Let $\{q_1, \dots, q_N\}$ denote the coefficients of the N channels computed by analyzing the input image at scale i and position \mathbf{k} using the original N -channel wavelets $\{\xi_{n_1, i, \mathbf{k}}(r, \theta), \dots, \xi_{n_N, i, \mathbf{k}}(r, \theta)\}$. For the steering, we want to rotate the wavelet ψ at each (i, \mathbf{k}) and find the maximum of $|\langle I, \psi_{i, \mathbf{k}}(\cdot, \cdot + \theta_0) \rangle|$ as a function of θ_0 . For a given

(i, \mathbf{k}) , we thus rewrite the function to maximize as

$$Q(i, \mathbf{k}, \theta_0) = \left| \left\langle \sum_{n \in H} I_n e^{jn\theta_0} u_n \xi_{n,i,\mathbf{k}} \right\rangle \right|^2 = \left| \sum_{n \in H} q_{n,i,\mathbf{k}} u_n e^{jn\theta_0} \right|^2, \quad (6.32)$$

noting that the inner product is conjugate-linear in the first argument.

We propose to maximize (6.32) by testing the roots of its derivative with respect to θ_0 . The order of the polynomial to solve can be reduced substantially by making a change of variable, relying on the fact that we use harmonics that are multiples of M for an M -fold pattern. We introduce the variable $z = e^{jM\theta_0}$ for the given set of harmonics $H = \{Mk : k = k_0, k_0 + 1, \dots, k_0 + N - 1\}$. The polynomial to be maximized on the unit circle hence becomes

$$Q(z) = \left(\sum_{k=k_0}^{k_0+N-1} r_k z^k \right) \overline{\left(\sum_{k=k_0}^{k_0+N-1} r_k z^k \right)}, \quad (6.33)$$

where $r_k = q_{Mk} u_{Mk}$. Extrema of this expression are found by looking for the zeros of its derivative with respect to θ_0 . After finding the points that satisfy this condition, we evaluate $Q(z)$ at all of them to find the one corresponding to a maximum.

It is important to note that, on the unit circle, Q can be rewritten as

$$Q(z) = \sum_k s_k z^k, \quad (6.34)$$

where $s = r[\cdot] * \bar{r}[-\cdot]$ is the discrete auto correlation of r . The derivative then yields

$$\frac{d}{d\theta_0} Q(z) = (jM) \sum_k k s_k z^k. \quad (6.35)$$

This form greatly simplifies the processing as the complex roots of $\frac{d}{d\theta_0} Q(z)$ can be easily computed, now that an analytical expression for the coefficients of this polynomial is known.

Let \tilde{z} be the root on the unit circle for which $Q(\tilde{z})$ is maximal. From the definition $z = e^{jM\theta_0}$, the optimal steering angle is given by

$$\tilde{\theta} = \frac{\angle \tilde{z}}{M}. \quad (6.36)$$

Due to the M -fold symmetry, $\tilde{\theta} + m \frac{2\pi}{M}$, $m \in \mathbb{Z}$, are equivalent solutions.

6.7.3 Connecting Junctions

We propose an approach based on dynamic programming in order to link the detected centers of symmetry and segment the objects of interest [144]. The best path between two centers of symmetry is sought on a discrete grid built around axes indexed by k and v , the former corresponding to the straight line joining the starting and end points and the latter to the perpendicular line to k going through the starting point (*i.e.*, $k = 0$). We refer to v_k as the value on the v axis of the decision node at location k . We define the cost c of the path connecting two junctions up to node v_{k+1} as

$$c(v_{k+1}) = c(v_k) + \frac{\lambda}{L_\zeta} \left(\sum_{(x,y) \in \zeta} I(x,y) \right) + (1 - \lambda) |v_{k+1} - v_k|, \quad (6.37)$$

where ζ is the segment that joins v_k and v_{k+1} , L_ζ its length, and I the input image. The cost is therefore equal to the cost up to the previous node v_k , plus two terms weighted by the data-dependent parameter $\lambda \in [0, 1]$. The first term ensures fidelity to data as it is composed of the integral on the input image over the straight line between v_k to v_{k+1} . This formulation favors paths that explore dark pixels; it can be negated when bright paths are searched for. The second term enforces smoothness by penalizing large displacements along the v axis. The solution found at any step k is guaranteed to be optimal as it is built by propagating the optimal solution up to this step.

As we have no prior information on which other points each center of symmetry links to, we apply a global approach to connect the junctions in the whole image. For a given center, we therefore compute all possible paths linking it to its neighbors within a certain range of distance and save their cost. We proceed in this fashion for all points and therefore obtain a list of costs for each possible connection between two centers of symmetry. After ordering this list in a decreasing order with respect to the cost value, we start with the first entry, then go through each of the following and draw the valid connections. In order to determine which connections are most likely to be correct, we rely on the knowledge of the order of symmetry of each center. We define a link as valid under two conditions. First, the centers involved have to be “free”: they have to be connected to less than M other points if they have M -fold symmetry. Then, as connections in an M -fold center are assumed to be spaced by an angle $\frac{\pi}{M}$, the angle between the new link and the already validated ones is checked for coherency. We proceed in this way until all the centers of symmetry are connected.

6.8 Experimental Results

Our symmetry-detection algorithm has been implemented as a plug-in for the open-source image-processing software ImageJ [139]. In our implementation, we use the isotropic mother wavelets of Simoncelli [112] with dyadic scale progressions. To evaluate the performance of the algorithm, we test it on a variety of synthetic images in the presence of noise and, finally, on real microscopic images. In addition to demonstrating the effectiveness of the approach in detecting junctions and segmenting structured images, the aim of this section is also to measure to which degree our methods honor translation, rotation, and scale invariance, as well as their robustness against noise.

We evaluate two versions of the space design: a real version with positive and negative harmonic pairs; and a complex version with positive harmonics only. We compare these results with the ones given by the Fourier template. The number of harmonics are fixed for all cases: $H = \{3n : n = 0, 1, \dots, 9\}$ for the complex wavelets and $H = \{3n : n = -8, -7, \dots, 7, 8\}$ for the real wavelets.

To generate our test images, we assume that we can represent the background signal (autofluorescence) present in biological images with the model described in Section 4.2.2. We consider the background signal as noise to the detection problem.

In general, to make a quantitative evaluation, we compute the Jaccard index, and root mean-square errors for both position and angle. The definition of the Jaccard index is given in the Nomenclature.

6.8.1 Robustness Against Noise

First, we generate a series of 1024×1024 test images, where we control the location of the junctions and the angles. To separate distinct effects, in this experiment we use only the first wavelet scale to make detections. We corrupt the images with isotropic Brownian motion, with a mean of zero and a σ of $1, 2, \dots, 100$. We intend to detect 643 junctions, with 2 different angle values: 0 and 60 degrees.

To make a quantitative evaluation we compute the Jaccard index, and both position and angle root mean-square errors. 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 closer than 5 pixels to the original nodes. Otherwise, they are counted as false positives. The RMS error is computed for the matched detections.

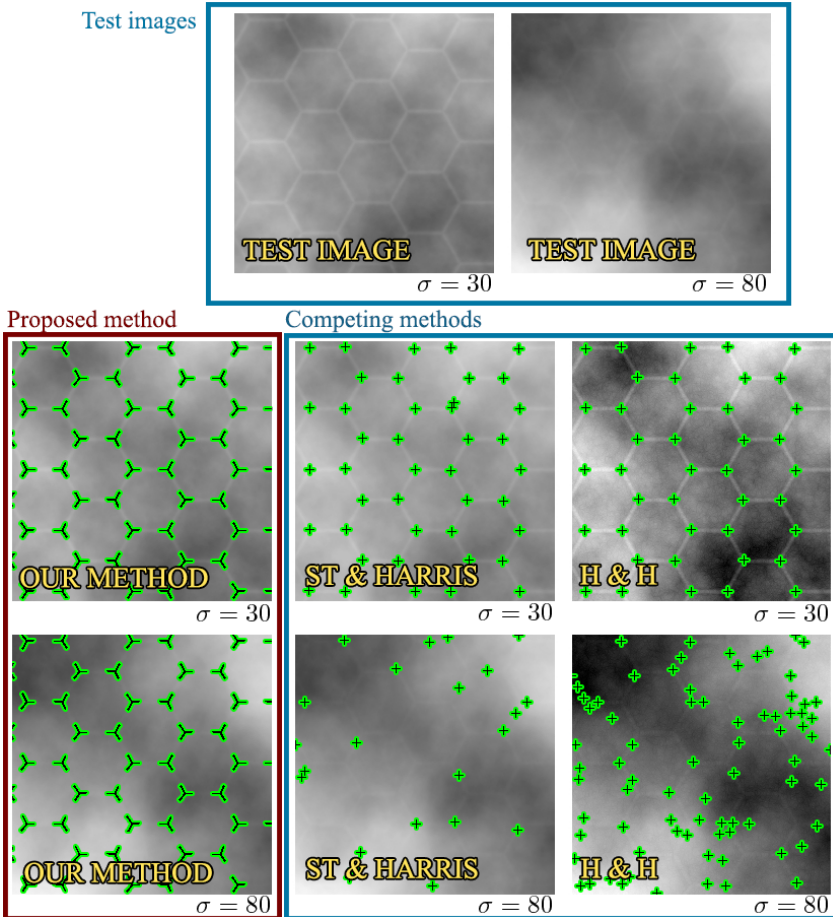


Figure 6.5: First row: test images corrupted with isotropic Brownian motion (mean 0, $\sigma = 30$ and mean 0, $\sigma = 80$). Second row: detections on the first test image with our method; with Harris' corner detector based on the structure tensor ("ST & Harris"); and with Harris' corner detector based on the largest eigenvalues of the Hessian ("H & H"). Third row: detections on the second test image with the three approaches: our method; "ST & Harris"; and "H & H".

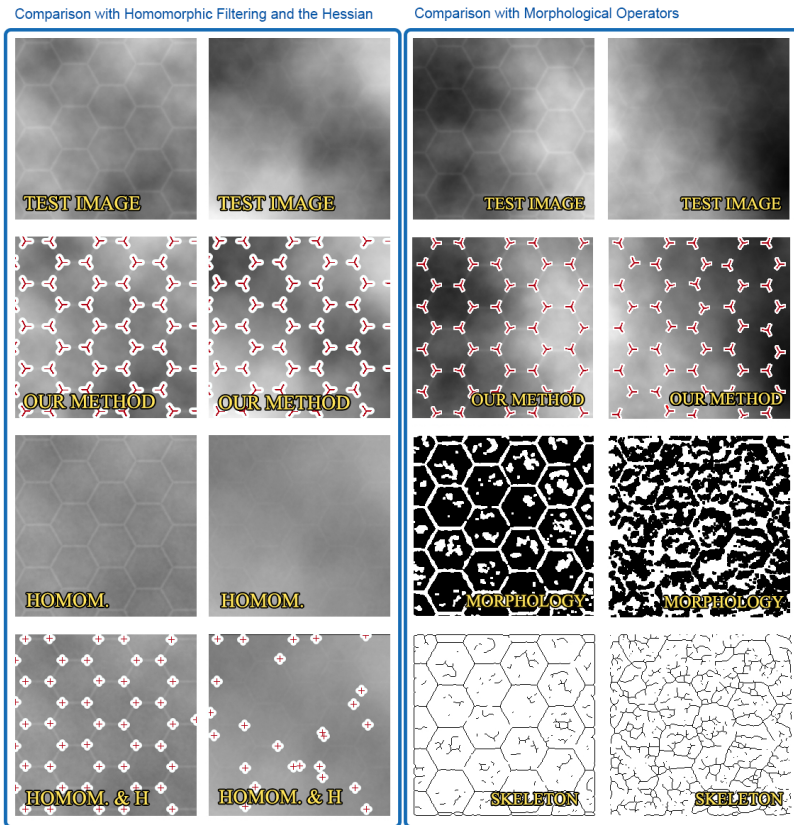


Figure 6.6: First row, left rectangle: test images corrupted with isotropic Brownian motion (mean 0, $\sigma = 30$ and mean 0, $\sigma = 80$). Second row, left rectangle: detections with our method on both images (results directly obtained on test images). Third row, left rectangle: output of a homomorphic filter, after careful tuning. Fourth row, left rectangle: detections with Harris' corner detector based on the output of the homomorphic filter. First row, right rectangle: test images corrupted with isotropic Brownian motion (mean 0, $\sigma = 30$ and mean 0, $\sigma = 80$). Second row, right rectangle: detections with our method on both images. Third row, right rectangle: output of morphological filtering, after careful tuning. Fourth row, right rectangle: skeleton of the underlying structure, obtained by skeletonizing the output of the morphological filter. To provide a fair comparison, in all cases, we used only the first wavelet scale of our method to make the detections.

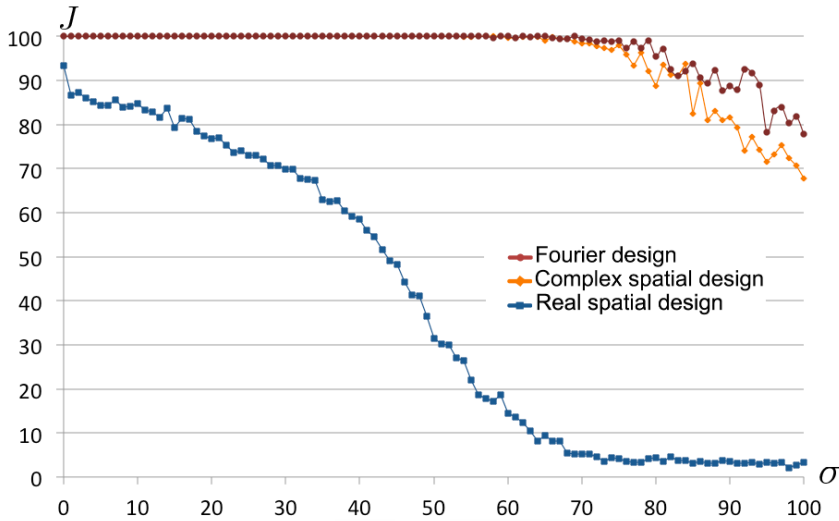


Figure 6.7: Jaccard index J under isotropic Brownian motion, as a function of σ of the noise.

An illustration of the results is given in Figure 6.5. To provide a fair comparison, we only took advantage of the first wavelet scale of our method to make detections. To generate the results for the competing methods, we used the ImageJ/Fiji built-in Hessian plug-in. As visible from these images, our algorithm performs much better than structure tensor and Hessian-based methods in case of strong noise.

Figure 6.6 provides further comparisons. First, we apply a homomorphic filter, as a nonlinear contour enhancement technique on the original test images to improve the output of the Harris' corner detector. We note that, if there are objects of interest at different scales, finding a single filter matching them all is often not possible. In this sense, our method shows a clear advantage, both in terms of the quality of the results and in a sense that we have a result equivalent to a non-linear filter at different scales and the detection in one shot. Then, we apply morphological operators on the original image. They typically work well in the absence of noise, or in case of clear structures. As mentioned in the Nomenclature, the spectral power density of fluorescence

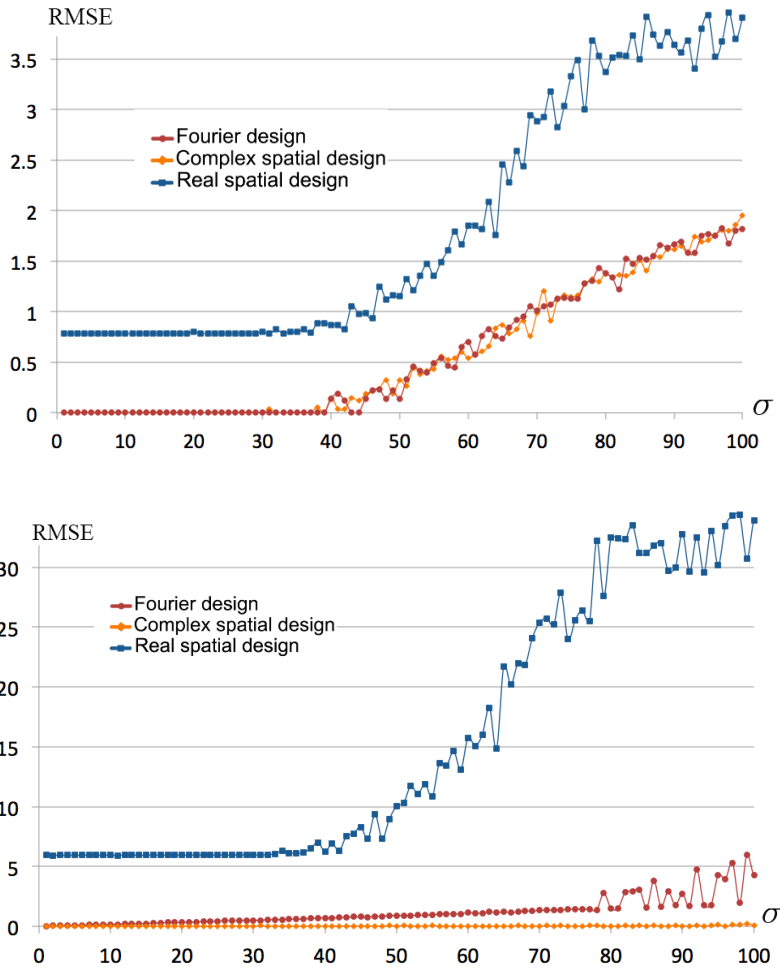


Figure 6.8: Position and angle error in the root-mean-square (RMSE) sense under isotropic Brownian motion, as a function of σ of the noise.

microscopy images is isotropic and well represented by a power law corresponding to fractional Brownian motion. When applying morphological operators such as opening, closing and skeletonization (provided by the morphological plugins of Fiji/ImageJ), we observe that they are mostly ineffective in such images.

The Jaccard index is presented in Figure 6.7, and the RMS errors in Figure 6.8. The noise level of $\sigma = 30$, which corresponds to the first test image of Figure 6.5, is in the range when our method performs the best (Jaccard index 100%, negligible RMS errors). The noise level of $\sigma = 80$, which corresponds to the second test image of Figure 6.5, is related to the point where the Jaccard index curves cut off from 100%.

Based on the graphs, we can say that the Fourier and the complex spatial design performed best. The real spatial design had difficulties under heavy noise; they produced a bias even in the noise-free case. This has two explanations. First, the shape of the wavelet (Figure 6.1) contains oscillations related to the Hankel functions. When matching, the algorithm tries to fit the junction with the oscillations, such that the maxima overlap. Second, when generating the wavelet we had to use more harmonics (17 channels in total) to achieve the same blade size, hence, the same detection range. When steering the wavelets, these extra harmonics mean that we have to find the roots of a polynomial of a much higher order (17 instead of 10), thus increasing numerical errors.

6.8.2 Multiscale Properties

We present results challenging the multiscale aspect of our approach using the test image shown in Figure 6.9. This synthetic image shows a uniform tiling of the hyperbolic plane and exhibits three- and four-fold local symmetries of various size. The goal is to identify these junctions at different scales.

Next, we generate yet another series of 1024×1024 test images. First, we build a synthetic grid as in the experiments of Sections 6.8.1-6.8.3, and then change the profile of the edges from thickness 1 to 5.1 with steps of 0.1 (see Figure 6.10). We add isotropic Brownian motion of zero mean and $\sigma = 10$. There are approximately 85 junctions with specific angle values to detect.

To make a quantitative evaluation, we computed the RMS error for position and angle (as seen in Figure 6.11). We use the Jaccard index to decide whether to make the decisions at higher scale or not: when the Jaccard index is getting lower, we jump to the next scale.

Based on Figure 6.11, we claim that our method is multiscale and performs well, for patterns within a wide range of scales.




Figure 6.9: Detection results on synthetic data exhibiting three- (blue) and four-fold (red) local symmetries. Size of the original image: 512×512 pixels.




Figure 6.10: Test images with different line thickness, corrupted with isotropic Brownian motion (mean 0, $\sigma = 10$). From top to bottom: line thickness 1.0 and 5.0, respectively.




Figure 6.11: Position and angle error in the sense of root mean square under isotropic Brownian motion for scaling from 1 to 5.1.




Figure 6.12: Illustration of the rotation invariance of our method. From top to bottom: A: original test image, B: test image corrupted with additional Brownian motion (PSNR: 20.8), C: test image corrupted with additional white Gaussian noise (PSNR: 18.26). The junctions were rotated with 10 degree steps. The centers of the junctions were correctly detected in all cases. The numbers next to the junctions correspond to the detected rotation angle of the junctions with respect to their original state.




Figure 6.13: Jaccard index J under isotropic Brownian motion of zero mean and $\sigma = 30$ for rotations α from 0 to 360 degrees.

6.8.3 Rotation Invariance

The rotation invariance property of our method is illustrated in Figure 6.12. The junctions in the figure are rotated with 10 degree steps with respect to their previous state. The centers of the junctions are correctly detected in all cases. The numbers next to the junctions correspond to the detected rotation angle of the junctions with respect to their original state.

We also generate another series of 1024×1024 test images. First, we build a synthetic grid as in the previous experiments, and then rotate it around its center with steps of 1 degree, from 0 to 360 degrees. We finally apply isotropic Brownian motion of zero mean and $\sigma = 30$ on these images. This amount of noise corresponds to the one illustrated in Figure 6.5. As in Section 6.8.1, we use only the first wavelet scale to make detections, and we fix the window size to 30×30 pixels. There are approximately 640 junctions (some nodes may “fall out” of the image when rotated) to detect with specific angle values.

To make a quantitative evaluation, we compute the Jaccard index (Figure 6.13). The graph confirms that our method is essentially rotation-invariant and performs well, independently of the orientation of the junctions.




Figure 6.14: Detection of three-fold junctions and segmentation in the bright-field microscopic image of a honeycomb. Left: original image. Right: detection results using the proposed approach. Size of the original image: 256×256 pixels; the junctions are at least 30 pixels away from each other.

6.8.4 Practical Applications

Figure 6.14 features a microscopic view of a honeycomb. This biological structure naturally tend to exhibit a close-to-perfect hexagonal structure. The detection results of our method are shown on the right. They are accurate, both concerning the detection of junctions and their linking, which was expected from the regularity of the structure.

Figure 6.15 features a picture of endothelial cells of the cornea with the detection results. The detection task is more challenging then in the honeycomb image due to the variety of orientations and sizes of the junctions. In this case, the method once again delivered good results, which suggests that our algorithm works well even if the structures are only semi-regular.




Figure 6.15: Detection of three-fold junctions in an image of endothelial cells of the cornea. Image courtesy of Jorge Fischbarg [145]. Size of the original image: 539×666 pixels; the junctions are at least 16 pixels away from each other.

6.9 Conclusion

In this chapter, we presented a general wavelet-based framework to detect various types of junctions and their orientation in images. Our approach is rotation-invariant, multi-scale, robust to noise, and can be tuned to identify arbitrary symmetries as well as other types of intersections. It is hence possible to obtain a precise estimate of the location of junctions across a large range of scales and with any orientation. We proposed an algorithm that combines detectors of local symmetry and dynamic programming to extract the structure of, and segment images featuring multifold junctions. The effectiveness of our approach in practical applications was demonstrated on synthetic and real data in the presence and absence of noise.

6.10 Notes

Proof of Proposition 6.1

The quadratic Fourier-domain energy term to minimize takes the form

$$\begin{aligned}
 E\{\hat{\psi}\} &= \frac{1}{2\pi} \int_{-\pi}^{\pi} |\hat{\psi}(\varphi)|^2 w(\varphi) d\varphi \\
 &= \frac{1}{2\pi} \int_{-\pi}^{\pi} \left(\sum_{n=-N}^N u_n e^{jn\varphi} \right) \overline{\left(\sum_{n'=-N}^N u_{n'} e^{jn'\varphi} \right)} w(\varphi) d\varphi \\
 &= \sum_{n=-N}^N \sum_{n'=-N}^N u_n \bar{u}_{n'} \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{j(n-n')\varphi} w(\varphi) d\varphi \\
 &= \frac{1}{2} \sum_{n=-N}^N \sum_{n'=-N}^N u_n \bar{u}_{n'} \frac{1}{2\pi} \int_{-\pi}^{\pi} \left(e^{j(n-n')\varphi} w(\varphi) + e^{j(n'-n)\varphi} w(-\varphi) \right) d\varphi \\
 &= \frac{1}{2} \left(\sum_{n=-N}^N \sum_{n'=-N}^N u_n \bar{u}_{n'} W(n-n') + \sum_{n=-N}^N \sum_{n'=-N}^N u_n \bar{u}_{n'} W(n'-n) \right) \\
 &= \mathbf{u}^H \mathbf{W} \mathbf{u}, \tag{6.38}
 \end{aligned}$$

where we have made the assumption that w is even, with $w(\varphi) = w(-\varphi)$.

The Fourier coefficients of w are computed as

$$W(n-n') = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{j(n-n')\theta} w(\theta) d\theta. \tag{6.39}$$

These calculations show that the eigenvalue problem involve a real and symmetric Toeplitz matrix \mathbf{W} , which necessarily has real eigenvectors. In addition, the eigenvectors of a symmetric Toeplitz matrix are either symmetric or anti-symmetric around their center [146].

1. Symmetric case: $\hat{\psi}$ can be written as a sum of cosines with purely real coefficients like

$$\sum_{n=-N}^N u_n e^{jn\varphi} = u_0 + \sum_{n=1}^N 2u_n \cos(n\varphi). \tag{6.40}$$

Thus, $\hat{\psi}$ is purely real, so ψ is Hermitian-symmetric. As a result, $|\psi|$ is symmetric with respect to the origin.

2. Anti-symmetric case: $\hat{\psi}$ can be written as a sum of sines with purely imaginary coefficients like

$$\sum_{n=-N}^N u_n e^{jn\varphi} = \sum_{n=1}^N 2j u_n \sin(n\varphi). \quad (6.41)$$

Thus, $\hat{\psi}$ is purely imaginary, so ψ is Hermitian-antisymmetric. As a result, $|\psi|$ is symmetric with respect to the origin.

In conclusion, $|\psi|$ can have only even orders of symmetry.

Now, we shall attempt to observe how this property influences the detection of odd symmetries. We make the natural assumption that the image f is real, so that

$$|\langle \psi, f \rangle|^2 = (\langle \text{Re}(\psi), f \rangle)^2 + (\langle \text{Im}(\psi), f \rangle)^2. \quad (6.42)$$

If we rotate ψ by π , the magnitude of the wavelet coefficients doesn't change (only the sign of $\langle \text{Re}(\psi), f \rangle$ or $\langle \text{Im}(\psi), f \rangle$ can be different). It means that the magnitude of the wavelet coefficients cannot distinguish a pattern from the rotation the of same pattern by π . Thus, it is not suitable for detecting odd symmetries.

We note that the same derivation is valid if the harmonics are multiples of some integer κ .

Part II

Transformability: Beyond Rotations

Chapter 7

Transformability: Beyond Rotations

7.1 Overview

In the first part of this thesis, we presented the construction and application of steerable representations of functions that are adapted to local rotations of image features. In the second part, our goal is to generalize these ideas to other local transformations, with a special focus on local scaling.

First, in the present chapter, we reinterpret and extend the notion of steerability by taking the broader perspective of Lie groups. We identify classes of functions that fulfill this generalized definition for relevant geometric transformations beside rotation (translation and scaling). Then, in Chapter 8, we develop a framework to combine these functions with tight wavelet frames. This allows us to capture, for instance, the local scale of features, with continuous scaling in between the usual dyadic scales of discrete wavelets.

Theories addressing local geometric transformations have appeared in the literature since the early 1990's. The origins of steerable wavelets can be traced back to the work of Freeman and Adelson [92] on steerable filters. A general approach to construct scalable functions was first given by Simoncelli *et al.* [147]. Perona generalized this to arbitrary compact transformations without requiring the group properties [148]. Teo, in his doc-

toral dissertation [149] and a series of papers with Hel-Or [150, 151], unified the existing theories and provided a solid mathematical framework for the study of transformability based on the theory of Lie groups and Lie algebras. We recall some of his definitions later in this chapter.

Teo's definition of general transformability and the corresponding function families are similar to ours. However, we focus on "local" transformations such as scaling, both in the sense of localization in the space or frequency domains, as well as limiting the transformation parameter (*e.g.*, scale) to a certain range. This aspect distinguishes mainly our approach from previous work; the difference is apparent in Chapter 8.

As discussed later on, the global solutions for scalable functions are non-local, thus a tradeoff for the approximation properties is needed. In Teo's work, this compromise consists of finding an approximate representation of the function of interest, which is scalable within a limited range of scales. We, on the other hand, are interested in deriving an exact representation of the function, that is approximately (but very closely) scalable at each location. Furthermore, we develop tight wavelet frames based on this definition, which are completely novel, while Teo never discusses wavelets.

In the following, we summarize the general definitions and concepts that serve as a basis for the new constructions and results in Chapter 8.

7.2 Lie groups and transformation groups

Definition 7.1 (Lie group) *Assume a group $A = \{a\}$ with operation $*$ and inverse \sim . Based on that, we define the two maps*

$$\begin{aligned} \iota : a &\mapsto \sim a && \text{(inverse map)} \\ \nu : (a_1, a_2) &\mapsto a_1 * a_2. && \text{(composition map)} \end{aligned}$$

If the maps ι and ν are infinitely differentiable, then A is called a Lie group.

This definition assumes that the notion of differentiation with respect to a is defined (more precisely, that A is a smooth manifold).

Example 7.1 *Let us consider the group $A = \mathbb{R}^2$ with addition. Then*

$$\iota(a) = -a \quad \text{and} \quad \nu(a_1, a_2) = a_1 + a_2 \tag{7.1}$$

are both smooth functions. This defines a Lie group.

Definition 7.2 (Transformation group) A transformation on space X (e.g., the space of images) is an invertible map $X \rightarrow X$. Transformations can be composed to construct new transformations. A transformation group is a group of transformations, with the group operation of composition.

Definition 7.3 (Lie transformation group) A Lie transformation group is a transformation group G on some space X that can be parametrized by a Lie group A , such that for $a_1, a_2 \in A$:

$$T(a_1) \circ T(a_2) = T(a_1 * a_2) = T(v(a_1, a_2)), \quad (7.2)$$

where \circ denotes the composition of operators (we later drop \circ to simplify notation).

It follows from the definition that

$$T^{-1}(a) = T(\sim a) = T(i(a)), \quad (7.3)$$

and $T(a)$ is identity if a is the identity of A .

Example 7.2 Let us consider $A = (0, \infty)$ under multiplication, and let X be the space of real images on \mathbb{R}^2 : $X = \{f : \mathbb{R}^2 \rightarrow \mathbb{R}\}$. We define $T(a)$, $a \in A$, as

$$T(a) : f(\cdot) \mapsto f(\cdot / a). \quad (\text{dilation by } a)$$

Since dilation by $a_1 a_2$ is equivalent to separate dilations by a_2 followed by a_1 , we have

$$T(a_1)T(a_2) = T(a_1 a_2). \quad (7.4)$$

Thus $\{T(a) : a \in A\}$ is a transformation group, and specifically the transformation group of dilations.

We are interested in transformation groups G on signals and images, where the parameter group A is a Lie group. More specifically, our interest is in rotation ($A = \mathbb{T}$, the circle group); shifts or translation (in 1D $A = \mathbb{R}$ and in 2D $A = \mathbb{R}^2$ under addition); and scaling ($A = (0, \infty)$ under multiplication).

We do not cover the general theory of Lie algebras and exponential maps here, but simply note a definition that is used in one of the results.

Definition 7.4 (Generator) We consider a Lie transformation group G with scalar parameter $a \in A$. The generator of G , denoted as L , is defined as the operator

$$L: f \mapsto \frac{d}{da} \left(T(a)f \right)_{a=0}. \quad (7.5)$$

If the parameters are vectors $\mathbf{a} = (a_1, \dots, a_d)$, there are d generators:

$$L_k: f \mapsto \frac{\partial}{\partial a_k} \left(T(\mathbf{a})f \right)_{\mathbf{a}=0}. \quad (7.6)$$

Example 7.3 For the Lie transformation group of x -translations, $T(a): f \mapsto f(\cdot - a, \cdot)$. The generator is given by

$$Lf = \frac{d}{da} \left(T(a)f \right)_{a=0} = \frac{d}{da} \left(f(\cdot - a, \cdot) \right)_{a=0} = -\partial_x f, \quad (7.7)$$

that is, $L = -\partial_x$.

7.3 Transformability and Eigenfunctions

Similar to Definition 3.1 of Section 3.2.2, we define a transformable family of functions. For this we replace rotations by more general transformation groups.

Definition 7.5 Let $G = \{T(a)\}_{a \in A}$ be a transformation group of bounded linear operators on $L_2(\mathbb{R}^d)$, indexed by a set A . A finite family of functions $\mathcal{M} = \{m_n\}$ is **transformable** by G if the span of \mathcal{M} is invariant under the action of any operator from G .

This definition is a generalization of the idea of eigenfunctions. In the simplest case, we have a single function $\mathcal{M} = \{m\}$ that is an eigenfunction for each transform $T(a)$ and there exist $\lambda_a \in \mathbb{R}$ such that

$$T(a)m = \lambda_a m. \quad (7.8)$$

For $\mathcal{M} = \{m_n\}_{n=1}^N$ containing several linearly independent functions, this condition is generalized, as the eigenvalues are replaced by matrices. For any $a \in A$, there must be a matrix $\mathbf{L}(a)$ such that

$$\begin{pmatrix} T(a)m_1 \\ \vdots \\ T(a)m_N \end{pmatrix} = \mathbf{L}(a) \begin{pmatrix} m_1 \\ \vdots \\ m_N \end{pmatrix}. \quad (7.9)$$

Given a linearly independent G -transformable set \mathcal{M} and an expansion

$$f = \sum_n c_n m_n = \mathbf{c}^T \begin{pmatrix} m_1 \\ \vdots \\ m_N \end{pmatrix}, \quad (7.10)$$

the transformation of f is determined by $T(a) \in G$, by modifying the expansion coefficients:

$$T(a)f = \mathbf{c}^T \mathbf{L}(a) \begin{pmatrix} m_1 \\ \vdots \\ m_N \end{pmatrix}. \quad (7.11)$$

Such a function or representation is **transformable**. We mention three important examples in dimension 1.

Example 7.4 For the family of shifts in \mathbb{R} , the eigenfunctions are exponential functions, and in particular complex exponentials $m_\omega(x) = e^{j\omega x}$ for all $\omega \in \mathbb{R}$. The related transformable representation is the continuous-time Fourier transform (or, more generally, the two-sided Laplace transform):

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{f}(\omega) e^{j\omega x} d\omega, \quad (7.12)$$

where \hat{f} is the forward Fourier transform of f (the uncountable sum becomes an integral).

Example 7.5 For the family of circular (periodic) shifts with period 2π on the unit circle, the eigenfunctions are complex exponentials with period $2\pi/n$: $m_n(x) = e^{jn x}$, $n = 0, 1, 2, \dots$. The related transformable representation is the Fourier series:

$$f(x) = \sum_{n \in \mathbb{Z}} c_n e^{jn x}, \quad (7.13)$$

where c_n , $n \in \mathbb{Z}$, are Fourier series coefficients of f .

Example 7.6 For the family of dilations, the eigenfunctions are homogeneous functions $m_\alpha(x) = x^\alpha$. The related transformable representation is the inverse Mellin transform:

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} x^{-\frac{1}{2}-js} \tilde{f}(s) ds, \quad (7.14)$$

Table 7.1: Comparison of the different nomenclature.

Teo	Our work
steerable	transformable
basis for equivariant space	transformable family of functions
equivariant space	(the span of a transformable family)

where $\tilde{f}(s)$ is the forward Mellin transform of f and is given by

$$\tilde{f}(s) = \int_0^{+\infty} x^{-\frac{1}{2}+js} f(x) dx. \quad (7.15)$$

In Table 7.1, we compare our definitions to those of Teo. In particular, our transformable families form a basis for equivariant spaces in Teo's terminology. A main result in characterizing these spaces is the following.

Theorem 7.1 *Given a Lie transformation group G with generator(s) L_k , and a basis $\mathbf{m} = (m_1, \dots, m_N)$, $\text{span}\{\mathbf{m}\}$ is equivariant under G iff for each L_k there is an $N \times N$ (numerical) matrix \mathbf{B}_k , such that*

$$L_k \mathbf{m} = \mathbf{B}_k \mathbf{m}. \quad (7.16)$$

Since the L_k , as seen from the example of x -translations, are differential operators, fulfilling the above identities involves finding general solutions of linear systems of differential equations.

7.3.1 Characterization of shiftable and scalable families

Here, we are interested in studying scaling properties, so we consider the transformation group $G = \{T(a)\}_{a \in (0, \infty)}$, where $T(a)f = f(\cdot/a)$. Thus, we require

$$\text{span}\{g_n : n = 1, \dots, N\} = \text{span}\{g_n(\cdot/a) : n = 1, \dots, N\} \quad (7.17)$$

for any dilation $a \in (0, \infty)$. Notice that

$$g_n(\cdot) = g_n\left(2^{\log_2(\cdot)}\right). \quad (7.18)$$

This defines the new functions $m_n := g_n(2^{(\cdot)})$ that should satisfy

$$\text{span}\{m_n : n = 1, \dots, N\} = \text{span}\{m_n(\cdot + b) : n = 1, \dots, N\} \quad (7.19)$$

for any real b . Hence the span of $\{m_n\}$ should be translation-invariant. Conversely, any translation-invariant function can be used to define a dilation-invariant one. This means that, once we identify all shiftable functions, we can derive scalable families. Finite-dimensional translation-invariant sets of functions have been characterized; the following results follow from [152].

Proposition 7.1 *Let $\{m_1, \dots, m_{n_{\max}}\}$ be a family of differentiable functions on \mathbb{R} , and let \mathbf{m} denote the vector with m_n as its n th component. Further, suppose that the collection is shiftable so that, for any $b \in \mathbb{R}$, there is a matrix $\mathbf{L}(b)$ such that*

$$\mathbf{m}(\cdot + b) = \mathbf{L}(b)\mathbf{m}. \quad (7.20)$$

If the components of \mathbf{L} are differentiable functions, then each m_n is a linear combination of functions of the form

$$x^k 2^{\alpha x} \quad (7.21)$$

for $x \in \mathbb{R}^+ \setminus \{0\}$, $k \in \mathbb{N}$, and $\alpha \in \mathbb{C}$.

Based on the above result, we characterize scalable families.

Proposition 7.2 *Let $\{g_1, \dots, g_{n_{\max}}\}$ be a differentiable collection of functions, and let \mathbf{g} denote the vector with g_n as its n th component. Further, suppose that the collection is scalable such that, for any $a \in (0, \infty)$, there is a matrix $\mathbf{L}(a)$ that satisfies*

$$\mathbf{g}(\cdot / a) = \mathbf{L}(a)\mathbf{g}(\cdot). \quad (7.22)$$

If the components of \mathbf{L} are differentiable functions, then each g_n is a linear combination of functions of the form

$$(\log_2(x))^k 2^{\alpha \log_2(x)} = (\log_2(x))^k x^\alpha \quad (7.23)$$

for $x \in \mathbb{R}^+ \setminus \{0\}$, $k \in \mathbb{N}$, and $\alpha \in \mathbb{C}$.

This gives a complete characterization of (globally) scalable functions, that is one of the main components of the construction of local scalable representations of Chapter 8.

Chapter 8

Scalable Filters for Tight Wavelet Frames

8.1 Overview

In Chapter 3, we presented how the multi-order Riesz transform is used to equip an isotropic tight wavelet frame with the property of “steerability”. The relevance of the Riesz transform in this context was due to the fact that, in the Fourier domain, it is represented by complex exponentials on the circle, which are the eigen functions of the rotation group.

In this chapter, in analogy with steerable wavelets, we present a general construction of adaptable tight wavelet frames, with an emphasis on scaling operations. In particular, the wavelets that we derive can be “dilated” by a procedure comparable to the operation of steering steerable wavelets. This permits us to detect and estimate the scale (or frequency localization) of features that fall between the dyadic wavelet scales of classical wavelet schemes. Thus, we incorporate several advantages of continuous wavelet transforms, while benefiting from the computational efficiency of discrete wavelets.

The fundamental aspects of the construction are comparable to steerable wavelets: an admissible collection of Fourier multipliers is used to extend a tight wavelet frame (paralleling the Riesz transform), and the “scale” of the wavelets is adapted by scaling the multipliers. As applications, we show how these new constructions can be used for

frequency localization, and for estimating the shape and location of cellular shapes that appear frequently in problems in biology and material science.

Despite the similarity of the main concept, our formulation here slightly differs from the steerable framework we developed earlier. This deviation follows from the fact that, unlike for rotations, where we were limited to a finite range of angles, for other transformations, in particular for scaling, the range of parameters is unbounded. To overcome this challenge, we limit the range within which our wavelets can be adapted (for instance to the scales in between two wavelet bands), by scaling the wavelet within a certain parameter “window”. We combine this “finite” version of scaling with the dyadic wavelet scales to cover the entire range of scales ¹.

In Section 8.2, we provide the mathematical formulation of the problem. More precisely, in Section 8.2.1, we introduce admissible families of Fourier multipliers, and we connect them with transformable families of functions in Sections 8.2.2 and 8.2.3. Section 8.2.4 contains their application to wavelet frames to define transformable wavelet families. Next, in Section 8.3, we use this approach to define wavelets with adaptable frequency (or scale) localization. This translates to a close approximation of continuous scaling in-between wavelet bands, which we call pseudo-scaling. In Section 8.4, we present an example of the construction, with a straightforward generalization to higher dimensions. In Section 8.5, we describe our software implementation and provide experimental validation of the approach in practical applications.

8.2 Mathematical Formulation

8.2.1 Admissible Multipliers and Extended Frames

The two components of our construction are a primal wavelet frame and a collection of bounded linear operators on L_2 . Since we are working with a wavelet system, we require these operators to commute with translations. Every operator of this form can be described as a Fourier multiplier in $L_\infty(\mathbb{R}^d)$ [155, Theorem 3.18]. Specifically, for each such operator T , there is a symbol $\widehat{T} \in L_\infty(\mathbb{R}^d)$ such that

$$\mathcal{F}\{Tf\} = \widehat{T}\widehat{f}, \tag{8.1}$$

¹The results in this chapter are based on the manuscript [153]: Z. Püspöki, J.P. Ward, D. Sage, M. Unser, “On The Continuous Steering of the Scale of Tight Wavelet Frames”, in press and on the article [154]: Z. Püspöki, D. Sage, J.P. Ward, M. Unser, “SpotCaliper: Fast Wavelet-based Spot Detection with Accurate Size Estimation”, *Bioinformatics*, in press, 10.1093/bioinformatics/btv728, 2015.

for every $f \in L_2(\mathbb{R}^d)$.

We have previously defined an admissible collection of Fourier multipliers as one that satisfies the partition-of-unity property below.

Definition 8.1 (cf. [99, Definition 2.1]) Consider a collection of complex-valued functions $\mathcal{M} = \{M_n\}_{n=1}^{n_{\max}}$. They are admissible if

1. Each M_n is Lebesgue-measurable;
2. The squared moduli of the elements of \mathcal{M} form the partition of unity

$$\sum_{n=1}^{n_{\max}} |M_n(\boldsymbol{\omega})|^2 = 1 \quad (8.2)$$

for every $\boldsymbol{\omega} \in \mathbb{R}^d \setminus \{\mathbf{0}\}$.

The partition-of-unity property is important as it allows us to use an admissible collection to extend a tight frame to a new tight frame.

Theorem 8.1 (cf. [99, Theorem 2.4]) Suppose $\{\xi_k\}_{k \in \mathbb{Z}}$ is a tight frame of $L_2(\mathbb{R}^d)$, with

$$f = \sum_k \langle f, \xi_k \rangle \xi_k \quad (8.3)$$

and

$$\|f\|_{L_2}^2 = \sum_k |\langle f, \xi_k \rangle|^2 \quad (8.4)$$

for every $f \in L_2(\mathbb{R}^d)$. Also, let $\mathcal{M} = \{M_n\}_{n=1}^{n_{\max}}$ be admissible. Then, the collection

$$\{\mathcal{F}^{-1}\{M_n \widehat{\xi}_k\}\}_{n=1, \dots, n_{\max}; k \in \mathbb{Z}} \quad (8.5)$$

is also a tight frame.

A tight frame that is extended in this way is called an **extended frame**.

Similarly to steerable wavelets, we adapt (or shape) the multipliers by applying an isometry to them. This operation maintains the the tight-frame property.

Proposition 8.1 Let $\{M_n\}_{n=1}^{n_{\max}}$ be an admissible collection and define the vector \mathbf{M} to have entries $[\mathbf{M}]_n = M_n$. If the $(n'_{\max} \times n_{\max})$ matrix \mathbf{U} is an isometry (i.e., $\mathbf{U}^* \mathbf{U}$ is the $(n_{\max} \times n_{\max})$ identity matrix), then the elements of $\mathbf{U}\mathbf{M}$ form an admissible collection of size n'_{\max} .

The proof of Proposition 8.1 is given in Section 8.8.

In Section 8.2.3, we take advantage of Proposition 8.1 to combine a collection of individual trigonometric functions to build trigonometric polynomials.

8.2.2 Admissible Multipliers and Transformable Families

Here, we briefly review the idea of transformability in the context of rotations by recalling an illustrative example from the first part of the thesis. Then, we extend our formulation to the general case.

In 2D, an example of an admissible collection of homogeneous multipliers is $\mathcal{M} = \{M_1, M_2\}$, where

$$M_1(\boldsymbol{\omega}) = \sin(v) \quad (8.6)$$

$$M_2(\boldsymbol{\omega}) = \cos(v), \quad (8.7)$$

and v denotes the angle that $\boldsymbol{\omega}$ makes with a fixed direction. Any rotation of these functions can be written as a weighted sum of the unrotated functions, as in

$$\begin{pmatrix} \cos(v + v_0) \\ \sin(v + v_0) \end{pmatrix} = \begin{pmatrix} \cos(v_0) & -\sin(v_0) \\ \sin(v_0) & \cos(v_0) \end{pmatrix} \begin{pmatrix} \cos(v) \\ \sin(v) \end{pmatrix}. \quad (8.8)$$

Now suppose that we have a tight frame $\{\xi_k\}$, which is extended by \mathcal{M} as described in Theorem 8.1, and let $f \in L_2(\mathbb{R}^2)$. Then the rotated multipliers

$$\mathcal{M}_{v_0} = \{\cos(\cdot + v_0), \sin(\cdot + v_0)\} \quad (8.9)$$

are also admissible and can be used to extend the frame. We use the frame coefficients from one frame to compute the coefficients for the other by

$$\begin{pmatrix} \langle \widehat{f}, \widehat{\xi}_k \cos(v + v_0) \rangle \\ \langle \widehat{f}, \widehat{\xi}_k \sin(v + v_0) \rangle \end{pmatrix} = \begin{pmatrix} \cos(v_0) & -\sin(v_0) \\ \sin(v_0) & \cos(v_0) \end{pmatrix} \begin{pmatrix} \langle \widehat{f}, \widehat{\xi}_k \cos(v) \rangle \\ \langle \widehat{f}, \widehat{\xi}_k \sin(v) \rangle \end{pmatrix}. \quad (8.10)$$

Our steerable wavelets are constructed using a tight wavelet frame $\{\xi_k\}$, where the mother wavelet is bandlimited and isotropic. We note that, the rotation invariance of the primal wavelets means that the derived wavelets themselves (not only the multipliers) are rotated by matrix multiplication. In the first part of the thesis, this property is used to capture the local orientation of features in images.

The steerable-wavelet construction is based on the group of rotation operators on $L_2(\mathbb{R}^d)$; however, an analogous construction can be performed in a more general setting, using families of multipliers that are transformable in the sense of 7.3.

In the general case of a transformable (*e.g.*, scalable) family of multipliers, the span of the multipliers is invariant to the action of the operator. Thus, the wavelet coefficients of the transformed tight frame can be derived from the wavelet coefficients of the untransformed frame.

8.2.3 Families of Dilation Multipliers

We consider the transformation group of dilations. We restrict our attention to admissible collections of multipliers $\{M_n\}$, where each multiplier is radial. This means that there exists $m_n : (0, \infty) \rightarrow \mathbb{C}$ such that $M_n(\boldsymbol{\omega}) = m_n(|\boldsymbol{\omega}|)$ (using $|\boldsymbol{\omega}|$ to denote absolute value in 1D and magnitude of a vector in higher dimensions).

Families of functions that are transformable by dilations are fully characterized by Proposition 7.2. They all have the form of

$$g(x) = m(\log_2(x)) = (\log_2(x))^k 2^{\alpha \log_2(x)} = (\log_2(x))^k e^{\alpha' \log_2(x)} \quad (8.11)$$

for $x \in \mathbb{R}^+ \setminus \{0\}$, $k \in \mathbb{N}$, and $\alpha, \alpha' \in \mathbb{C}$.

Multipliers are required to be bounded to form admissible collections. Therefore, functions of the form identified in (8.11) with $k > 0$, or with the real part of α non-zero cannot be included. Consequently, for our scalable multipliers, we consider admissible collections composed of trigonometric polynomials in $\log_2(x)$. A particular example is given by the collection of trigonometric functions

$$\mathcal{M} = \{m_n(\log_2|\cdot|)\}_{n=1}^{n_{\max}}, \quad (8.12)$$

where n_{\max} is some fixed degree and where

$$\{m_n\}_{n=1}^{n_{\max}} = \{\alpha_0\} \bigcup_{l=1}^{l_{\max}} \left\{ \alpha_l \cos\left(\frac{2\pi l}{\sigma} \cdot\right), \alpha_l \sin\left(\frac{2\pi l}{\sigma} \cdot\right) \right\}, \quad (8.13)$$

$\sum_{l=0}^{l_{\max}} |\alpha_l|^2 = 1$, and $\sigma > 0$. Note that this collection is scalable due to the angle-addition formulas

$$\begin{aligned} \sin(\omega_1 + \omega_2) &= \sin(\omega_1) \cos(\omega_2) + \cos(\omega_1) \sin(\omega_2) \\ \cos(\omega_1 + \omega_2) &= \cos(\omega_1) \cos(\omega_2) - \sin(\omega_1) \sin(\omega_2). \end{aligned} \quad (8.14)$$

Admissibility follows immediately from the fact that $\sin^2 + \cos^2 = 1$.

We define a general class of admissible collections of multipliers by combining trigonometric functions into trigonometric polynomials; a criterion for admissibility is given in [98].

Proposition 8.2 *For $\sigma > 0$, the collection $\mathcal{M} = \{m(\log_2(\omega_n |\cdot|))\}_{n=1}^{n_{\max}}$, with*

$$m(\omega) = \frac{\alpha_0}{\sqrt{n_{\max}}} + \sum_{l=1}^{l_{\max}} \sqrt{\frac{2}{n_{\max}}} \alpha_l \cos\left(\frac{2\pi l}{\sigma} \omega\right), \quad (8.15)$$

$$\omega_n = 2^{\sigma n / n_{\max}}, \quad \text{and} \quad \sum_{l=0}^{l_{\max}} |\alpha_l|^2 = 1, \quad (8.16)$$

is admissible if $n_{\max} \geq 2l_{\max} + 1$.

The proof of Proposition 8.2 is given in Section 8.8.

8.2.4 Adapting Extended Frames

The scale-invariance property (7.17) is important because it allows us to scale the multipliers in an admissible collection \mathcal{M} using matrix multiplications. We combine our multipliers with a tight frame to form an extended frame. In this case, the scale-invariance property means that, once we have computed the wavelet coefficients for the system derived from $\{M_n(a \cdot)\}$, matrix multiplications are used to determine the coefficients of the wavelet system from $\{M_n(a' \cdot)\}$ for $a \neq a'$. Moreover, we a can be chosen independently at each point.

As an illustration, suppose that we have at our disposal a normalized tight frame $\{\xi_k\}_{k=1}^{\infty}$ and the scalable, admissible family of dilation multipliers

$$\{M_n\}_{n=1}^{n_{\max}} = \{m_n(\log_2(|\cdot|))\}_{n=1}^{n_{\max}}. \quad (8.17)$$

Thus, the collection $\mathcal{M}_a = \{M_n(a \cdot)\}$ is also admissible for any $a > 0$. Indeed,

$$\{\mathcal{F}^{-1}\{M_n(a \cdot)\widehat{\xi}_k\} : k \in \mathbb{Z}, n = 1, \dots, n_{\max}\} \quad (8.18)$$

is a normalized tight frame.

Now, let \mathbf{M}_a be the matrix whose n th entry is $M_n(a \cdot)$. As these collections are scalable, there exist matrices $\mathbf{L}(a)$ such that

$$\mathbf{M}_a = \mathbf{L}(a^{-1})\mathbf{M}. \quad (8.19)$$

Therefore, based on the frame coefficients of $f \in L_2(\mathbb{R}^d)$, the coefficients for $a' \neq a$ and $a > 0$ are computed as

$$\begin{pmatrix} \langle \widehat{f}, \widehat{\xi}_k M_1(a' \cdot) \rangle \\ \vdots \\ \langle \widehat{f}, \widehat{\xi}_k M_{n_{\max}}(a' \cdot) \rangle \end{pmatrix} = \mathbf{L}(a'^{-1})\mathbf{L}(a) \begin{pmatrix} \langle \widehat{f}, \widehat{\xi}_k M_1(a \cdot) \rangle \\ \vdots \\ \langle \widehat{f}, \widehat{\xi}_k M_{n_{\max}}(a \cdot) \rangle \end{pmatrix}. \quad (8.20)$$

As a particular example, suppose all multipliers are dilations of a single multiplier. Let \mathcal{M} be defined as in Proposition 8.2 with $n_{\max} \geq 2l_{\max} + 1$ and $\sigma = 2$. In this case,

$$\mathbf{M}_a = \mathbf{U}\mathbf{D}_a\mathbf{b}, \quad (8.21)$$

where \mathbf{U} is the $(n_{\max} \times (2l_{\max} + 1))$ matrix with entries

$$[\mathbf{U}]_{n,l} = \frac{1}{\sqrt{n_{\max}}} e^{j\pi l \log_2(\omega_n)} \quad (8.22)$$

for $l = -l_{\max}, \dots, l_{\max}$ and $n = 1, \dots, n_{\max}$; \mathbf{b} is the vector of functions with entries

$$[\mathbf{b}]_l = \begin{cases} \frac{\alpha|l|}{\sqrt{2}} e^{j\pi l \log_2(|\omega|)}, & l \neq 0 \\ \alpha_0, & l = 0 \end{cases} \quad (8.23)$$

and \mathbf{D}_a is the diagonal matrix with entries

$$[\mathbf{D}_a]_{l,l} = e^{j\pi l \log_2(a)}. \quad (8.24)$$

The matrix \mathbf{U} is an isometry, thus for any $a' > 0$, we have that

$$\mathbf{M}_{a'} = \mathbf{U}\mathbf{D}_{a'}\mathbf{b} \quad (8.25)$$

$$= \mathbf{U}\mathbf{D}_{a'}(\mathbf{D}_{a^{-1}}\mathbf{U}^T\mathbf{M}_a) \quad (8.26)$$

$$= \mathbf{U}\mathbf{D}_{a'a^{-1}}\mathbf{U}^T\mathbf{M}_a. \quad (8.27)$$

Therefore, $\mathbf{T}_{a,a'} = \mathbf{U}\mathbf{D}_{a'a^{-1}}\mathbf{U}^T$ is the matrix used to transform \mathcal{M}_a into $\mathcal{M}_{a'}$.

8.2.5 Isotropic Wavelets and Extended Frames

Similarly to Chapter 3, our construction is initialized with a tight wavelet frame of $L_2(\mathbb{R}^d)$, whose basis functions are generated by dilations and translations of the single mother wavelet ξ . Proposition 3.1 exhibits sufficient conditions for such a wavelet system.

In Section 8.3, we use Meyer-type wavelets, constructed using the techniques of [99]. We want our final wavelets to take the shape of the multipliers in an admissible family within a finite bandwidth (or range of scales). Therefore, we would ideally like to use a primal wavelet where h is a constant multiple of the characteristic function of $[\pi/4, \pi]$. However, discontinuities in the Fourier domain correspond to slow decay in the spatial domain. Therefore, as a tradeoff, we propose a profile h_ϵ that is a smooth approximation to the characteristic function of $[\pi/4, \pi]$, where ϵ is an approximation parameter. For this construction, we define a smooth, non-decreasing function G that satisfies $G(\gamma) = 0$ for $\gamma < -1$ and $G(\gamma) = \pi/2$ for $\gamma > 1$. Then, for an approximation parameter ϵ , we define

$$H_\epsilon(\gamma) = G\left(\frac{\gamma+1}{\epsilon}\right) - \frac{\pi}{2} + G\left(\frac{\gamma-1}{\epsilon}\right) \quad (8.28)$$

and

$$h_\epsilon(\gamma) = 2^{-1/2} \cos\left(H_\epsilon\left(\log\left(\frac{2^{1+\epsilon}}{\pi}\gamma\right)\right)\right), \quad (8.29)$$

with support in $[4^{-1-\epsilon}\pi, \pi]$.

For

$$G(\gamma) = \begin{cases} 0, & \gamma < -1 \\ \frac{35\pi}{64} \left(\frac{-1}{7}\gamma^7 + \frac{3}{5}\gamma^5 - \gamma^3 + \gamma + \frac{16}{35}\right), & -1 \leq \gamma < 1 \\ \frac{\pi}{2}, & \gamma \geq 1, \end{cases} \quad (8.30)$$

the resulting function h_ϵ and a dilated version are plotted in Figure 8.1.

8.3 Smooth Scaling of Wavelets Using Localized Frequency Multipliers

Our aim in this section is to demonstrate the construction of an extended frame. Here, we adapt locally our wavelets within each dyadic scale to have a continuous transition in scale to the adjacent dyadic scales. Hence, we define scalable multipliers that are




Figure 8.1: Plot of $h_\epsilon(\omega)$ and $h_\epsilon(\omega/2)$ for the Meyer wavelet system.

localized in the frequency domain, so that one cycle of these multipliers is neatly captured by the Meyer-type wavelets of the previous section. This means that, we require the energy of the multipliers to be localized within the support of the wavelets, such that scaling the multipliers provides a close approximation to scaling the wavelets.

8.3.1 Example of a Localized Frequency Multiplier

In this section, we provide an example of the proposed construction. We define the radial mother wavelet ξ in the Fourier domain by $\widehat{\xi}(\omega) = h_\epsilon(\omega)$ where h_ϵ is given as in (8.28), (8.29), and (8.30). Fourier-domain representations of $h_\epsilon(\omega)$ and $h_\epsilon(\omega/2)$ are shown in Figure 8.1. We extend this wavelet frame using the admissible collection of Proposition 8.2 with $n_{\max} = 2l_{\max} + 1$. The trigonometric polynomial m is defined by the vector

$$\alpha = \frac{\sqrt{4685}}{14055} (125, 101\sqrt{2}, 53\sqrt{2}, 16\sqrt{2}, 2\sqrt{2}), \quad (8.31)$$




Figure 8.2: Plot of the Fourier multiplier $m(\log_2(\cdot))$ defined by (8.31) on the interval $[\pi/32, 2\pi]$.

which is determined by sampling a polynomial B-spline of degree 3. The corresponding multiplier is illustrated in Figure 8.2. We note that, the choice of the multiplier depends on the application. The one used in this example corresponds to a wavelet that has good localization in the Fourier domain. For α defined in this way, we see that the collection is admissible by Theorem 8.2, and the extended wavelet frame is tight.

8.3.2 Pseudo-Scaling

Our construction allows us to perform pseudo-dilations of the wavelets. The procedure is as follows. We define the scaled multipliers

$$\mathcal{M}^a := \{m(\log_2(\omega_n |a \cdot|))\}_{n=1}^{n_{\max}}. \quad (8.32)$$

We recall that h_ϵ is the Fourier profile of the primal mother wavelet ξ . We define the interval

$$(4^{-1-\epsilon}\pi + \epsilon', 2(4^{-1-\epsilon}\pi + \epsilon')] \tag{8.33}$$

within the support $[4^{-1-\epsilon}\pi, \pi]$ of the profile h_ϵ for a given parameter

$$0 < \epsilon' < \frac{\pi}{2} \left(1 - \frac{2}{4^{1+\epsilon}} \right). \tag{8.34}$$

Note that the dyadic dilations of this interval are non-overlapping.

Here, we formally define the operation of pseudo-scaling wavelets as scaling of the corresponding multipliers as in (8.20). The pseudo-scaled wavelet system is formed by applying \mathcal{M}^a to the primal-wavelet system. Within this system, we can very closely approximate true scaling.

We consider the multiplier defined in Section 8.3.1. Let M_0 be an element of \mathcal{M}^1 such that the profile of $\widehat{\xi}M_0$ attains its maximal value in the interval (8.33). Let p_0 denote the location of the maximum, and let $p(a) = p_0a$, as illustrated in Figure 8.3. We define the wavelet $\psi := \mathcal{F}^{-1}\{M_0\widehat{\xi}\}$. The profile of the dilated version of the Fourier transform of ψ , $\widehat{\psi}(a^{-1}\cdot)$, attains a maximum at $p(a)$. The pseudo-dilated version of ψ is defined as

$$\psi_a := \mathcal{F}^{-1}\{M_0(a\cdot)\widehat{\xi}(2^{q_a}\cdot)\}, \tag{8.35}$$

where q_a satisfies

$$p(a) \in (2^{q_a}(4^{-1-\epsilon}\pi + \epsilon'), 2^{q_a}2(4^{-1-\epsilon}\pi + \epsilon')). \tag{8.36}$$

The underlying idea is that we scale the multiplier M_0 by a , and at the same time, find the dyadic wavelet band (defined by $\widehat{\xi}(2^{q_a}\cdot)$) that best captures the localized peak of $M_0\widehat{\xi}$. In other words, $p(a)$ belongs to the 2^{q_a} dilation of the interval (8.33). The Fourier transform of ψ and a sequence of pseudo-dilations are shown in Figure 8.3. We note that, the transition as $p(a)$ crosses the point $2(4^{-1-\epsilon}\pi + \epsilon')$, where the dyadic scale changes. The wavelets maintain their shape due to the overlapping primal windows.

To quantify the approximation of scaling by pseudo-scaling, we measure the correlation between the truly scaled ψ and its pseudo-dilated counterpart ψ_a as

$$\varrho_\psi(a) := \max \left(\left\langle \frac{\psi_a}{\|\psi_a\|_{L_2}}, \frac{\psi(a\cdot)}{\|\psi\|_{L_2}} \right\rangle, 0 \right). \tag{8.37}$$




Figure 8.3: Left: Profile plot of $\hat{\psi}(a^{-1}\cdot)$ for $a = 1.0$ (solid) and $a = 1.5$ (dotted). The location of the peak is $p(a)$, which is indicated by the dashed vertical line. Right: Profile plot of pseudo-dilations $\hat{\psi}_a$, for $a = 1, 1.3, 1.5, 1.7$. The dashed vertical line is located at $2(4^{-1-\epsilon}\pi + \epsilon')$.




Figure 8.4: Quality metric ρ_ψ as a function of the dilation a .

This is similar to the quality metric defined in [156] to measure the approximation of shiftable wavelets. Values close to one indicate that the pseudo-dilated wavelets provide a close match to true dilation, while smaller values indicate a poor match. We show in Figure 8.4 the evolution of ρ_ψ as a function of dilation. There, we illustrate the quality of pseudo-scaling using the empirically defined value $\epsilon' = 0.45$. Due to the log-periodic nature of pseudo-scaling, the minimum of ρ_ψ is found by analyzing a single period of the trigonometric function. We compute this minimum to be 0.998.

8.3.3 Comments on the Construction

In Proposition 8.2, the periodicity of the trigonometric functions means that the multipliers are periodic in scale. For example,

$$\cos(\pi l \log_2(2^2 \cdot)) = \cos(\pi l \log_2(\cdot)), \quad (8.38)$$

since l is an integer.

Essentially, our tight-frame construction consists of the application of an admissible collection $\{M_n\}$ to bandlimited wavelets ξ , where $\hat{\xi}$ is zero in a neighborhood of the origin. Therefore, the supports of the scaled wavelets $\hat{\xi}(2^q \cdot)$ segment the Fourier domain, such that, for each scale, the support of $\hat{\xi}(2^q \cdot)$ covers a finite number of scaling periods of the multipliers. The fact that cyclic scaling can be combined naturally with a geometric division of the Fourier domain motivates our interest in cyclically scalable families for use with wavelets. We note that, there should exist synchronization between the cyclic periodicity and the discrete scales of the wavelets. As a result, one could get the rough discrete scale from the wavelet hierarchy, while the local scale is being refined by adapting the multipliers. This is the approach we follow with defining the intervals in (8.36) and the corresponding coarse dyadic scale given by q_a . Furthermore, we can adapt the frame independently at each point, by using a pointwise independent scaling parameter.

8.4 Special Case: One Complex Multiplier

In the localized-frequency construction, arbitrary localization can be obtained if one is willing to work with a highly redundant wavelet frame. In this section, we look at a simple case, where the admissible collection contains only two functions: sine and cosine.

We propose to use the mother wavelets ψ^{\cos} and ψ^{\sin} that are defined by their Fourier domain profiles

$$\widehat{\psi}^{\cos}(\boldsymbol{\omega}) = h(|\boldsymbol{\omega}|) \cos(\omega_0 \log_2(\kappa |\boldsymbol{\omega}|)) \quad (8.39)$$

$$\widehat{\psi}^{\sin}(\boldsymbol{\omega}) = h(|\boldsymbol{\omega}|) \sin(\omega_0 \log_2(\kappa |\boldsymbol{\omega}|)), \quad (8.40)$$

which have two parameters: κ and ω_0 . Note that the sin and cos functions are special cases of the trigonometric functions identified in (8.12) and (8.13). The parameter κ specifies a local shift of the cosine and sine under the window h in the frequency domain. The dependence of this frequency shift on κ is cyclic: all values of κ that differ by $2^{n2\pi/\omega_0}$ produce equivalent frequency shifts. The second parameter is ω_0 . It determines the size of these local-frequency cycles as well as the number of oscillations of cosine and sine that fall within the support of h . In practical designs, one can tune these parameters to select the number of oscillations (ω_0) and the phase shift (κ) of the sinusoids within the support of h . In our publication [157], we have used $\omega_0 = 4\pi$ and $\kappa = 2^5/\pi$. The function h can be defined as $2^{-1/2}h_\epsilon$, where h_ϵ is from (8.29).

The two real mother wavelets are combined in the single complex wavelet

$$\psi(\mathbf{x}) = \psi^{\cos}(\mathbf{x}) + j\psi^{\sin}(\mathbf{x}) \quad (8.41)$$

$$= \mathcal{F}^{-1} \left\{ h(|\cdot|) e^{j\omega_0 \log_2(\kappa|\cdot|)} \right\}(\mathbf{x}). \quad (8.42)$$

The polar form of the analysis coefficients of the complex wavelets $\psi_{i,\mathbf{k}}$ is

$$A_{i,\mathbf{k}} e^{j\beta_{i,\mathbf{k}}} = \langle f, \psi_{i,\mathbf{k}} \rangle, \quad (8.43)$$

where $i \in \mathbb{Z}$ denotes the scale and $\mathbf{k} \in \mathbb{Z}^d$ the position.

We note the following scaling relationship for ψ :

$$\operatorname{Re} \left(e^{j\beta} \psi \right) = \cos(\beta) \psi^{\cos} - \sin(\beta) \psi^{\sin} \quad (8.44)$$

$$= \mathcal{F}^{-1} \left\{ h(|\cdot|) \cos(\omega_0 \log_2(2^{\beta/\omega_0} \kappa |\cdot|)) \right\}. \quad (8.45)$$

A similar relationship holds between ψ , ψ^{\cos} , and ψ^{\sin} upon taking the imaginary part. When analyzing a real-valued function f , we use this property to scale the complex wavelets. We arrive at an expansion of f as a sum of locally scaled versions of ψ^{\cos} ,

where ψ^{\cos} is adapted to the signal at each scale i and position \mathbf{k} as in

$$f = \operatorname{Re} \left(\sum_{i,\mathbf{k}} \langle f, \psi_{i,\mathbf{k}} \rangle \psi_{i,\mathbf{k}} \right) \quad (8.46)$$

$$= \operatorname{Re} \left(\sum_{i,\mathbf{k}} \left(A_{i,\mathbf{k}} e^{j\beta_{i,\mathbf{k}}} \right) \psi_{i,\mathbf{k}} \right) \quad (8.47)$$

$$= \operatorname{Re} \left(\sum_{i,\mathbf{k}} A_{i,\mathbf{k}} \left(e^{j\beta_{i,\mathbf{k}}} \psi_{i,\mathbf{k}} \right) \right) \quad (8.48)$$

$$= \sum_{i,\mathbf{k}} A_{i,\mathbf{k}} \mathcal{F}^{-1} \left\{ 2^{id} e^{-j2^i \mathbf{k} \cdot \boldsymbol{\omega}} h(2^i |\cdot|) \cos \left(\omega_0 \log_2 \left(2^{\beta_{i,\mathbf{k}} / \omega_0} \kappa |\cdot| \right) \right) \right\}, \quad (8.49)$$

where d denotes the dimension.

Definition 8.2 Let $\{\psi_{i,\mathbf{k}}\}$ be a wavelet system as defined above. For a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, let the wavelet coefficients be defined as in (8.43). Then, we define the scale-adapted wavelets $\psi_{i,\mathbf{k}}^f$ by

$$\psi_{i,\mathbf{k}}^f = \mathcal{F}^{-1} \left\{ 2^{id} e^{-j2^i \mathbf{k} \cdot \boldsymbol{\omega}} h(2^i |\cdot|) \cos \left(\omega_0 \log_2 \left(\kappa_{i,\mathbf{k}} |\cdot| \right) \right) \right\} \quad (8.50)$$

$$\kappa_{i,\mathbf{k}} = 2^{\beta_{i,\mathbf{k}} / \omega_0} \kappa, \quad (8.51)$$

for $0 \leq \beta_{i,\mathbf{k}} < 2\pi$.

Using the scale-adapted wavelets, we view $A_{i,\mathbf{k}}$ as the coefficient associated to the analysis of f with a version of ψ^{\cos} that is optimally scaled locally. Scaling $\psi_{i,\mathbf{k}}$ by $2^{(-\beta_{i,\mathbf{k}})/\omega_0}$ creates the best match between f and a locally scaled version of ψ^{\cos} .

An application of the single complex multiplier case was presented in our publication [157]. To compare the estimation properties of the single complex multiplier with the more general multichannel case (with α corresponding to (8.31)), we generated a series of test images. The spots range in size from 8 pixels to 11 pixels with step size 0.2. The reference images (left) and the correspondence between the measured spot radius and the actual spot size is illustrated in Figure 8.5. Since the multichannel case provides a better estimate of size (with a relatively small computational overhead), from now on we focus on that case.




Figure 8.5: Reference images (left) and radius estimations (right). The measured spot size as a function of the actual spot size. The straight red curve represents the theoretical values. The yellow curve shows the estimated radius extracted from the phase of the wavelet coefficients in the complex case. The blue curve visualizes the estimated radius determined by scaling the multipliers in the multichannel case.

8.5 SpotCaliper

Based on the mathematical tools we developed in Sections 8.2 and 8.3, we created a novel wavelet-based image analysis software that provides a fast automatic detection scheme for circular patterns (spots), combined with the precise estimation of their size. The software is implemented as an ImageJ plugin under the name SpotCaliper. Our software can be applied to a broad class of images; it is not linked to one specific biological problem. However, we note two particular applications, where the precise detection of the spots and their radius estimation is of high interest. First, in antimicrobial susceptibility testing of significant bacterial isolates, the size of the spots describing the zones of the inhibition areas are measured and used as an indicator of antibiotic-resistance [33]. Second, detection and size measurements are also important for understanding *Enzyme-Linked ImmunoSpot* images [32]. As discussed in Chapter 1, those assays are used to monitor the adaptive immune responses of humans and animals by visualizing antigen-antibody reactions. Accompanying the interaction, the naturally transparent cells are releasing cytokines that color the surrounding substance, typically

resulting in blue or black spots. The size of the spots is proportional to the strength of the reaction.

8.5.1 Steps of the Algorithm

The software is able to automatically find spots of varying sizes, and it gives precise measurements of their radii. Here, we give a brief summary of the algorithm. We suppose that a scalable wavelet with the proper Fourier multipliers is at our disposal.

(1) **(Wavelet Analysis With Scalable Wavelets)**

We decompose the image with the scalable wavelet. At each location \mathbf{k} and dyadic scale i , we have n_{\max} channels, corresponding to the number of Fourier multipliers. The first stage outputs a map of wavelet coefficients $\{w_n(i, \mathbf{k})\}_{n=1}^{n_{\max}}$.

(2) **(Local Maxima Detection)**

We apply a local non-maximum suppression in order to prevent multiple detections of the same spot.

(3) **(Fine Estimation of the Radius)**

We perform the adaptation step (steering of the scale) at each location that was retained by the previous step. The size of an object corresponds to a maximum in the response of the wavelet detector. The wavelets therefore have to be “scaled” to look for the precise scale that elicits the largest response. Since our wavelets are scaled on a quasi-continuum, we obtain precise results.

(4) **(Selection)**

We rank or filter the candidates based on a particular measure (*e.g.*, strength of the response of the wavelet, contrast with respect to the neighboring background, signal-to-noise ratio (SNR)).

(4) **(Visualization of the Detections)**

The user is able to edit the results and their visualization in several ways, or save the results to a comma-separated text file (.csv).

8.5.2 Modes and Parameter Settings

The software has two user interfaces: regular (SpotCaliper) and plain (SpotCaliper Plain). In the regular setup, the SpotCaliper user interface visualizes the detection results on an image, and detections are collected in a table. When the computations are finished, the software switches to post-processing mode, enabling the user to edit the results: modify the measurements, extract data, and do further analysis. The plain interface differs

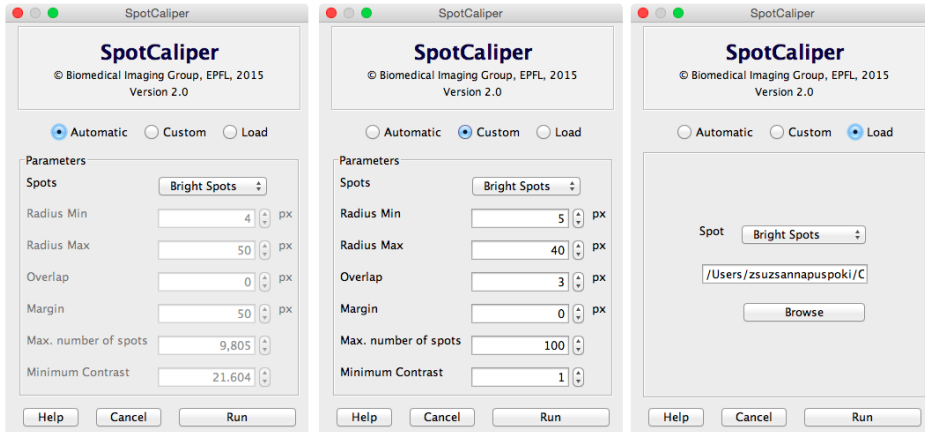


Figure 8.6: Automatic, custom, and load modes of SportCaliper.

from the regular version in the sense that the detection results are directly saved to a comma-separated text file (.csv), facilitating batch processing.

The SpotCaliper graphical user interface has three different modes corresponding to three different ways of parameter configuration: automatic, custom, and load. The automatic mode accounts for a completely automatic detection process, the custom mode permits the user to set several parameters, and the load option enables recovering previously saved measurements. The three different modes of SpotCaliper are shown in Figure 8.6.

Once the detection phase is complete, the user interface automatically switches to the selection panel (Figure 8.7) and visualizes the detections. Initially, all of the detected spots are visualized on the image and collected into a table.

8.5.3 Post-processing and Visualization

The detections are collected in a table, displaying the following parameters: ID of a spot, its location (x and y coordinates), radius, confidence, contrast, SNR and type. Each spot has its own unique identification number. This helps the user find the corresponding spot on the image. As described in Section 8.5.1, the detection algorithm is based on

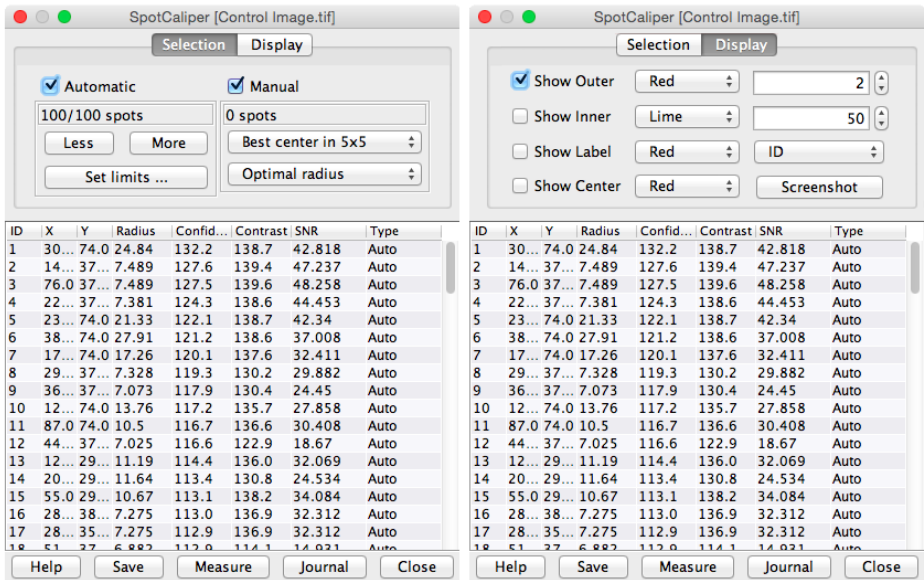


Figure 8.7: The selection and display panels of SpotCaliper.

the strength of wavelet coefficients. When the coefficients have high values, one can be more certain that the corresponding detections are correct. The confidence level reflects this property. The contrast measures the distinction between the object and its background: the higher the difference between the average intensity of the spot and the surrounding concentric neighborhood, the higher the probability that the detection is correct. Similarly to the contrast, the SNR characterizes the “correctness” of the spots. It gives the ratio between the average of the intensity of the spot (inside the surrounding circle) and the standard deviation of the background (surrounding concentric neighborhood).

It is possible to fine tune the given measurements, and the post-processing of the data is also an available option. Depending on the area of application of the software, there are two possible approaches to treat the number of false positive/negative results. In some cases, one might prefer to detect all of the spots, accepting some false detec-

tions; in other instances, it may be favorable to keep fewer detections, with a higher level of certainty. The above-mentioned confidence level, contrast level, and SNR measures address this issue by enabling the user to filter out the least reliable detections. Higher confidence levels, contrast levels, and SNR correspond to a higher level of certainty that the detections are correct.

It is also possible to interact with the data directly: the user can manually add a spot with an automatically computed radius, drag and drop or remove it. The drag and drop option includes an optimization of the center in an $M \times M$ window facilitating the editing of the results. When adding or dragging a spot, the software automatically searches for the highest detection response in an $M \times M$ window around the cursor. It means that there is no need to exactly click to the center point of the spot, the software finds it automatically. The size of the search window can be adjusted by the user.

The detected objects are visualized by their outline, covered with a variable opacity disk, or identified by both methods. There is a possibility to label the spots either by their ID, radius, confidence level, contrast, or SNR. The centers of the spots are optionally shown as well with crosses. The display panel controlling these settings is presented in Figure 8.7.

8.6 Experimental Results

To evaluate the performance of our algorithm, we use a variety of test images, including synthetic ones and real micrographs. The aim of our experiments is to measure the speed of our method, its accuracy, and its robustness against background signal. We also want to compare our method to other popular spot detection methods in the literature.

The artificial test images we generate aim at mimicking the circular structures of biological substances mentioned in the introduction of SpotCaliper. Our basic geometric forms are randomly distributed circular disk patterns. However, we add substantial noise corresponding to the non-stationary background signal of microscopy images in order to make the detection task challenging (*e.g.*, in Figure 8.8 the disks are barely visible).




Figure 8.8: First row: ground truth; detections obtained by our method on the test image corrupted with additive isotropic Brownian motion (mean 0, $\sigma = 2.0$); detections obtained by our method on the test image corrupted with additive isotropic Brownian motion (mean 0, $\sigma = 7.0$). Second row: detections obtained by competing algorithms on the first test image. Third row: detections obtained by competing algorithms on the second test image.

8.6.1 Reference Methods

There exist several approaches dedicated to the detection of circular objects and measuring their radii. We can separate these approaches into three main categories: classical global methods (*e.g.*, morphology and adaptive thresholding); methods based on the detection and analysis of edges and gradients (*e.g.*, the Circle Hough transform [158] and active contours); and approaches based on filtering (*e.g.*, Laplacian of Gaussian (LoG), determinant of Hessian (DoH); and wavelet-based techniques, [159]).

Global methods are mostly used for the evaluation of clear structures without background signal or noise. Their accuracy drops significantly for complex structures that appear in biological experiments, due to their sensitivity to noise and variable backgrounds. Edge-based methods are highly demanding in computational time and capacity. Filter-based methods include detectors with parametric templates that correspond to a specific range of sizes or scales.

In [4], it was shown that the LoG filter can be likened to a whitened matched filter and offers optimal properties for detection in a broad category of images. In [160], Lindeberg proposed a multiscale extension of the LoG-based detection scheme to overcome the limitation of the single scale. His method uses numerous passes of the LoG filter to capture the location and the size of spots. The computational efficiency of the algorithm highly depends on the diversity of the spot sizes.

For the comparison, we have chosen the following methods: Circle Hough transform, Multiscale LoG, Thresholding and connected components.

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 from the original nodes. Otherwise, they are counted as false positives. To make a quantitative evaluation, we compute the Jaccard index and the root-mean-square error (RMSE) for the estimation of the position and radius. We note that the RMSE is computed for the matched detections only. The definition of the Jaccard index is given in the Nomenclature.

8.6.2 Running Time

When evaluating different methods, it is often desirable to compare their performance with respect to their speed or running time. To do so, we perform the following experiment. We generate a series of test images (of size $1,000 \times 1,000$) where we control the location of the spots and their radius. The ground-truth data contain 200 disks, with

radii varying between 8 and 40 pixels. We allow overlap between neighboring spots, by at most 10 pixels.

In case of the multiscale Laplacian of Gaussian and the Circle Hough transform, the running time is influenced by the parameter settings of the method. In general, there is a trade-off between speed and performance: with generous settings they provide better results; however, their computation time increases dramatically.

The multiscale LoG method uses numerous passes of the LoG filter. It means that the range of the scales (depending on the size of the blobs) and the sampling step between the scales influences the running time. The bigger the range of the scales, or the finer the sampling steps are the more dramatic is the increase in computation time.

In the case of the Circle Hough transform, one can set parameters that determine the precision of the algorithm (corresponding to its “search space”).

Our method relies on the computation of wavelet coefficients. We apply a thresholding and local maxima search on the coefficients, and do the precise estimation of the radius only for a selected list of candidates. The list of the candidate points can be adjusted. The algorithm reaches its slowest computation time when we process all the pixels equivalently without any pre-selection step (*i.e.*, the list of candidates contains all the points).

To test the speed of the algorithm we set the parameters in such a way that they achieve a specific running time. Then, we measure how well can those methods detect and estimate the radius in that amount of time. We compute the RMS error of the radius and position estimation to quantify the results. In this case, we are using noise-free images, as described above. The results are summarized in Figure 8.9.

Due to the noise-free nature of the dataset the “thresholding and connected components” method (built-in plugin of ImageJ) performs the best. It is fast and accurate, which we note with a dot close to (but not in) the origin. However, we highlight that, the application area of the method “thresholding and connected components” is limited to essentially noise-free data. Our method performs better than the Circle Hough transform, even when allowing the latter one run for a much (4-4.5 times) longer time. Our method also outperforms the multiscale LoG in the sense that we reach the same quality of results in less time (2.5 times faster).




Figure 8.9: Position and radius estimation error (in pixels) in the sense of root mean square error (RMSE) as a function of running time.

8.6.3 Robustness Against Background Signal

To generate our test images for this experiment, we assume that we can represent the background signal (autofluorescence) present in biological images with the model described in Section 4.2.2.

As in the previous case, we generate a series of test images (of size $1,000 \times 1,000$), where we exactly control the location of the spots and their radii. We intend to detect different spots, with radii varying between 8 and 40. We allow overlap between neighboring spots, by at most 10 pixels. To make the detection more challenging, we add some isotropic background signal (fBm), with a mean of zero and σ ranging from 0 to 10. An illustration of typical test results obtained by our method and the competing algorithms can be seen in Figure 8.8, along with the ground truth. We set the running time of the corresponding algorithms to the same order of magnitude (favoring the competing ones).

The Jaccard index and the RMS errors are presented in Figure 8.10. We plot the RMSE curves for each method for a Jaccard index higher than 60%. We also plot the distribution of the radius estimation error (in pixels) in the sense of EMS error as a function of the Jaccard index. One can observe that the results of our method are concentrated in the lower left portion of the diagram, which is favourable for applications (higher Jaccard index and higher accuracy). Based on the graphs, we confirm that our algorithm has a clear advantage in terms of accuracy and detection power in the presence of significant background.

We can see that the noise level of $\sigma = 2.0$ corresponds to the second line of the test images in Figure 8.8. In this range of noise parameters we obtain with our method a Jaccard index of 100%. The noise level of $\sigma = 7.0$, which corresponds to the third line of test image Figure 8.8 reflects the point where the Jaccard index curve of our method cuts off from 100%. Figure 8.8 allows a visual comparison of the different methods at those points. One can observe that, when our method is still working perfectly, the others are already making false positive and negative detections, as well as misalignments.

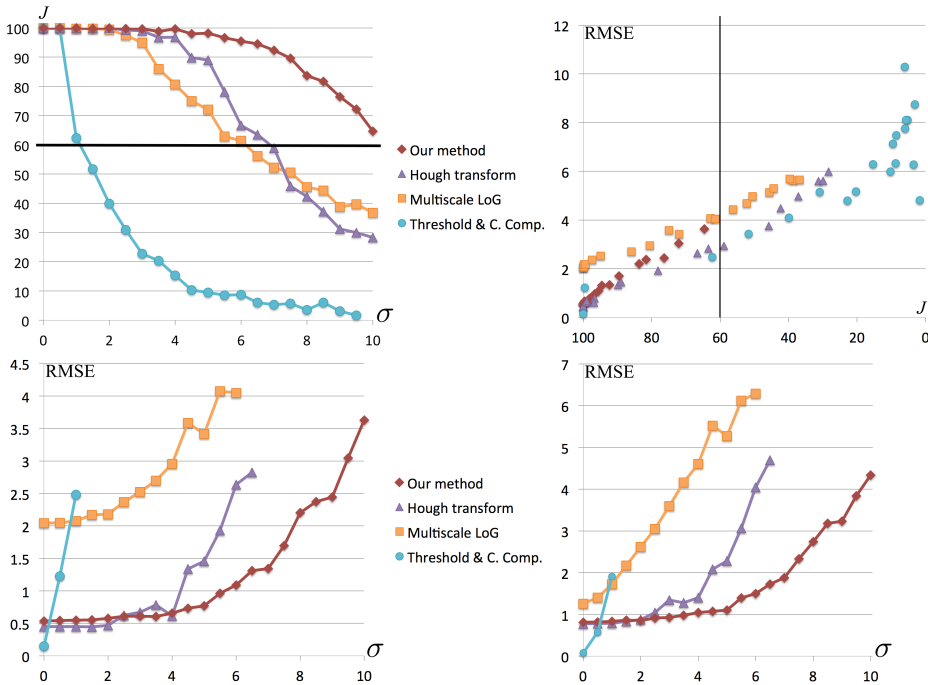


Figure 8.10: First row from left to right: Jaccard index under isotropic Brownian motion (background signal), as a function of σ of the noise; distribution of the radius estimation error (in pixels) in the sense of root mean square error (RMSE) as a function of the Jaccard index. Second row: Position and radius estimation error (in pixels) in the sense of root mean square error (RMSE) under isotropic Brownian motion (background signal), as a function of σ .




Figure 8.11: From left to right: cells in fluorescence microscopy; corresponding detections.

8.6.4 Practical Applications

In this section we illustrate the performance of our algorithm on biological micrographs. Figure 8.11 features cells in fluorescence microscopy. The detections and radius measurements are accurate, despite the heavy background signal and the fact that the intensity of the cells are varying. Figure 8.12 visualizes human HT29 colon-cancer cells². Our algorithm works well in this case as well. Figure 8.13 visualizes detection results on ELISpot data. Figure 8.14 illustrates results on the detection of cell nuclei labelled with DAPI staining³.

²Source: Broad Bioimage Benchmark Collection, <http://www.broadinstitute.org/bbbc/BBBC008/>

³Source: Online image database of the Perelman School of Medicine, University of Pennsylvania, www.med.upenn.edu/gtp/morphology_gallery.shtml.




Figure 8.12: From left to right: human colon-cancer cells; corresponding detections.




Figure 8.13: Detection results on ELISpot data. From left to right: original test image; automatic detections; semi-automatic detections.




Figure 8.14: Detection of the nucleus of cells labelled with DAPI staining. From left to right: original test image; corresponding detections.

8.7 Conclusion

In this chapter, we presented a general construction of adaptable tight wavelet frames, focusing on scaling operations. We applied our wavelet-based framework to detect and estimate the scale of circular structures in images. The attractive features of our algorithm are (i) our wavelets can be scaled on a quasi continuum without significant computational overhead; (ii) robustness; and (iii), speed. The effectiveness of our approach in practical applications was demonstrated on synthetic and real biological data.

8.8 Notes

Proof of Proposition 8.1

The result follows immediately from the definitions of \mathbf{M} and \mathbf{U} .

Proof of Proposition 8.2

This is a log mapping of Theorem 5.2 from [98].

Chapter 9

Conclusions

In this thesis, we presented a comprehensive approach and general frameworks for image representations that can be adapted to local geometric transformations. In addition to a multitude of theoretical results, including fundamental bounds and optimal solutions in many cases, we also verified the practical applicability of our approach in a wide range of experiments.

In the first section below, we summarize the main contributions of our work. In the second section, we discuss the potential areas of interest for future research related to our work.

9.1 Summary of Results

Transforms and Operators for Directional Bioimage Analysis

We presented a survey on the directional analysis of bioimages. We discussed the benefits and drawbacks of classical gradient- and derivative-based methods, directional multiscale transforms, and (possibly multiscale) steerable filters. From a user perspective, we identified steerable filters and wavelets among the most attractive methods. The advantage of steerable wavelets is that they combine high directional selectivity and multiscale analysis, which allows the efficient processing of oriented patterns at different scales.

Design Criteria on Optimal Detectors

We derived the Cramér-Rao Lower Bounds (CRLB) on the error of estimating local orientations of directional patterns in images using steerable filters. Based on those bounds we provided solutions for improving and optimizing the choice of the detector functions. In particular, we addressed the problem of choosing the best harmonics to achieve the lowest CRLB, as well as finding the optimal radial profile of the detector function. We also studied the bounds on the performance of multi-scale steerable estimators, and the differences with the single-scale case. We proposed an estimator for the estimation of orientation, and compared its estimation error to the theoretical bounds. Our proposed estimator closely followed these bounds in experiments. We provided several experiments on different realistic junctions and directional patterns.

Analysis of Local Symmetry Points

We presented a framework for the detection and classification of the centers of local symmetry in images. Our algorithm relied 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 used the F-test to examine the distribution of the energy across different channels. 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. We provided experimental results on synthetic images, biological micrographs, and electron-microscopy images to demonstrate the performance of the algorithm.

Design of Steerable Filters for Junction Detection

We studied junction detectors based on steerable wavelets that are polar-separable in the Fourier domain. We defined a general framework to design steerable wavelets that optimally match desired profiles. This permitted us to obtain precise estimates of junctions locations and orientations across a wide range of scales and at any orientation. The multiscale selectivity of the method was provided by the radial part of the detector, while the angular profile was represented by circular harmonics, and encoded by a shaping matrix. The shaping matrix was obtained by an optimization problem of maximally localizing the wavelets along some particular directions. The focus of the derivation was on the design of multi-fold symmetric wavelet profiles, however we provided

solutions for the general case as well. Based on the estimation of the junctions, we proposed an algorithm combining local symmetry detectors and dynamic programming to detect and segment structures in images featuring multifold junctions. The effectiveness of our approach in practical applications was demonstrated on synthetic and real biology data in the presence of noise.

General Construction of Adaptable Tight Wavelet Frames

In analogy with steerable wavelets, we presented a general construction of adaptable tight wavelet frames, with an emphasis on scaling operations. The fundamental aspects of the construction were reminiscent of steerable wavelet design: an admissible collection of Fourier multipliers were used to extend a tight wavelet frame, and the “scale” of the wavelets was adapted continuously in between dyadic scales, by scaling the multipliers. As applications, the proposed wavelets were used for frequency localization, and for detecting the location and size of cellular shapes.

Based on the theory discussed above, we have created a novel wavelet-based image-analysis software, called SpotCaliper, providing a fast automatic detection scheme for circular patterns (spots), combined with the precise estimation of their size. It is implemented as an ImageJ plugin with a friendly user interface. The user can additionally edit the results by modifying the measurements (in a semi-automated way), and extract data for further analysis. The fine tuning of the detections includes the possibility of adjusting or removing the original detections, as well as adding further spots.

9.2 Outlook

Integrated Design of an Optimal Steerable Pattern

In Chapter 4, we identified the harmonics and radial profiles for steerable detectors that achieve the best Cramér-Rao bounds. We demonstrated some of the significant practical implications of this in the experiments in that section (especially with regard to the optimal choice of the harmonics). However, it would be of interest to approach this question from a more practical perspective, and provide a step-by-step guide to designing the optimal steerable detector in applications. Based on the CRLB results for steerable detectors, the optimal radial profile and the number of harmonics could be easily determined. The performance of such schemes could be predicted and tested. We are

currently actively working towards a paper on this topic.

Steerable Wavelets in 3D

Instead of the circular harmonic wavelet frames used in Chapters 5 and 6, we could adapt our framework to 3D with spherical or zonal harmonics. However, the design and the shaping of the harmonics is not straightforward. While the circular harmonic components correspond to specific angular directions in the 2D plane, this is no longer true in the spherical harmonic case; the 3D basis functions can take various shapes. Also, the steering itself is computationally more expensive, since one has to optimize the rotation with respect to horizontal and vertical directions that cannot be decoupled.

Another interesting point to address comes from the imaging properties of microscopes. Typically, the axial and lateral resolution is not the same, so one cannot consider the data to be sampled on an isotropic grid. Thus, the implementation of the proposed algorithms raises several future questions related to steering on non-uniform lattices. Also, the proper choice of the radial profile should be revisited, possibly favoring elliptical solutions over isotropic ones.

SWM: Steerable Wavelet Machines

It could be of interest to combine the proposed methods with machine learning, by using our framework as descriptors for image regions. Our methods could serve to bridge the gap between handcrafted moving frame based approaches and learned representations. We are currently conducting active research on this topic, and a paper on this subject will be forthcoming.

Scaling With Non-Isotropic Wavelet Profiles

In the current framework, for scalable wavelet analysis, we focus on isotropic wavelet frames and their adaptation with the family of dilation multipliers. Theoretically, we could apply a particular basis of scalable filters to a non-isotropic (elliptical) dyadic tight wavelet frame to create wavelet channels, and the resulting wavelet frame would be also arbitrarily scalable. However, in case of ellipsoidal patterns, the orientation of the ellipses would raise several further questions. Due to the directionality of the patterns one would need to do steering not only with respect to dilations, but also jointly, with

respect to rotations. The joint optimization process would result in a more complicated scenario, but could serve as a target for future research.

Scaling in 3D

It could be of interest to extend our “scalable” framework for 3D data. Note that our theoretical results and constructions for scalable design are not restricted to 2D. Assuming isotropy, we could take the 3D counterpart of our primal wavelet and apply the same classes of Fourier multipliers that we have defined for the 2D case. However, the proper handling of anisotropic data (which is typical in micrographs, as discussed above) requires more care concerning the implementation, and possibly non-isotropic wavelet profiles need to be considered, as mentioned in the previous point.

Bibliography

- [1] I.M. Gelfand and N.Y. Vilenkin, *Generalized Functions, Vol. 4: Applications of Harmonic Analysis*, Academic Press, New York, USA, 1964.
- [2] B. Pesquet-Popescu and J.L. Veheil, “Stochastic fractal models for image processing,” *IEEE Signal Processing Magazine*, vol. 19, no. 5, pp. 48–62, September 2002.
- [3] A.P. Pentland, “Fractal-based description of natural scenes,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-6, no. 6, pp. 661–674, November 1984.
- [4] 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 Transactions on Image Processing*, vol. 14, no. 9, pp. 1372–1383, September 2005.
- [5] M. Unser and P.D. Tafti, *An Introduction to Sparse Stochastic Processes*, Cambridge University Press, Cambridge, UK, 2014.
- [6] J. Fageot, E. Bostan, and M. Unser, “Wavelet statistics of sparse and self-similar images,” *SIAM Journal on Imaging Sciences*, vol. 8, no. 4, pp. 2951–2975, 2015.
- [7] P.N. Bernasconi, D.M. Rust, and D. Hakim, “Advanced automated solar filament detection and characterization code: Description, performance, and results,” *Solar Physics*, vol. 228, no. 1-2, pp. 97–117, May 2005.
- [8] Y. Yuan, F.Y. Shih, J. Jing, H. Wang, and J. Chae, “Automatic solar filament segmentation and characterization,” *Solar Physics*, vol. 272, no. 1, pp. 101–117, August 2011.

- [9] M.A. Schuh, J.M. Banda, P.N. Bernasconi, R.A. Angryk, and P.C.H. Martens, "A comparative evaluation of automated solar filament detection," *Solar Physics*, vol. 289, no. 7, pp. 2503–2524, July 2014.
- [10] F. Tupin, H. Maitre, J.-F. Mangin, J.-M. Nicolas, and E. Pechersky, "Detection of linear features in SAR images: application to road network extraction," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 36, no. 2, pp. 434–453, March 1998.
- [11] H. Jiuxiang, A. Razdan, J.C. Femiani, C. Ming, and P. Wonka, "Road network extraction and intersection detection from aerial images by tracking road footprints," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 45, no. 12, pp. 4144–4157, December 2007.
- [12] L. Wu and Y. Hu, "Vision-aided navigation for aircrafts based on road junction detection," in *Proceedings of the IEEE International Conference on Intelligent Computing and Intelligent Systems ICIS'09*, Shanghai, China, November 20–22 2009, vol. 4, pp. 164–169.
- [13] B. Dan, A. W. K. Ma, E. H. Hároz, J. Kono, and M. Pasquali, "Nematic-like alignment in SWNT thin films from aqueous colloidal suspensions," *Industrial & Engineering Chemistry Research*, vol. 51, no. 30, pp. 10232–10237, July 2012.
- [14] G. Agam, S.G. Armato, and C. Wu, "Vessel tree reconstruction in thoracic CT scans with application to nodule detection," *IEEE Transactions on Medical Imaging*, vol. 24, no. 4, pp. 486–499, April 2005.
- [15] B.S.Y. Lam, Y. Gao, and A.W.-C. Liew, "General retinal vessel segmentation using regularization-based multiconcavity modeling," *IEEE Transactions on Medical Imaging*, vol. 29, no. 7, pp. 1369–1381, July 2010.
- [16] N. Patton, T.M. Aslam, T. MacGillivray, I.J. Deary, B. Dhillon, R.H. Eikelboom, K. Yogesana, and I.J. Constable, "Retinal image analysis: concepts, applications and potential," *Progress in Retinal and Eye Research*, vol. 25, no. 1, pp. 99–127, January 2006.
- [17] N. Honnorat, R. Vaillant, J.S. Duncan, and N. Paragios, "Curvilinear structures extraction in cluttered bioimaging data with discrete optimization methods," in *Proceedings of the Eighth IEEE International Symposium on Biomedical Imaging*:

- From Nano to Macro ISBI'11*, Chicago IL, USA, March 30-April 2, 2011, pp. 1353–1357.
- [18] J.-F. Le Garrec, C.V. Ragni, S. Pop, A. Dufour, J.-C. Olivo-Marin, M.E. Buckingham, and S.M. Meilhac, “Quantitative analysis of polarity in 3D reveals local cell coordination in the embryonic mouse heart,” *Development*, vol. 140, no. 2, pp. 395–404, January 2012.
- [19] S. Pop, A.C. Dufour, J.-F. Le Garrec, C.V. Ragni, C. Cimper, S.M. Meilhac, and J.-C. Olivo-Marin, “Extracting 3D cell parameters from dense tissue environments: application to the development of the mouse heart,” *Bioinformatics*, vol. 29, no. 6, pp. 772–779, March 2013.
- [20] E. Meijering, M. Jacob, J.-C.F. Sarria, P. Steiner, H. Hirling, and M. Unser, “Design and validation of a tool for neurite tracing and analysis in fluorescence microscopy images,” *Cytometry Part A*, vol. 58A, no. 2, pp. 167–176, April 2004.
- [21] G. González, F. Fleurety, and P. Fua, “Learning rotational features for filament detection,” in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition CVPR'09*, Miami Beach, FL, USA, June 20–25 2009, IEEE, pp. 1582–1589.
- [22] R. Rezakhaniha, A. Agianniotis, J.T.C. Schrauwen, A. Griffa, D. Sage, C.V.C. Bouten, F.N. van de Vosse, M. Unser, and N. Stergiopoulos, “Experimental investigation of collagen waviness and orientation in the arterial adventitia using confocal laser scanning microscopy,” *Biomechanics and Modeling in Mechanobiology*, vol. 11, no. 3–4, pp. 461–473, March 2012.
- [23] E. Meijering, “Neuron tracing in perspective,” *Cytometry A*, vol. 77, no. 7, pp. 693–704, July 2010.
- [24] M. Jacob, T. Blu, C. Vaillant, J.H. Maddocks, and M. Unser, “3-D shape estimation of DNA molecules from stereo cryo-electron micro-graphs using a projection-steerable snake,” *IEEE Transactions on Image Processing*, vol. 15, no. 1, pp. 214–227, January 2006.
- [25] F. Aguet, S. Geissbühler, I. Märki, T. Lasser, and M. Unser, “Super-resolution orientation estimation and localization of fluorescent dipoles using 3-D steerable filters,” *Optics Express*, vol. 17, no. 8, pp. 6829–6848, April 2009.

- [26] Berlemont S. and J.-C. Olivo-Marin, "Combining local filtering and multiscale analysis for edge, ridge, and curvilinear objects detection," *IEEE Transactions on Image Processing*, vol. 19, no. 1, pp. 74–84, January 2010.
- [27] F. Zhao, P.R.S. Mendonca, R. Bhotika, and J.V. Miller, "Model-based junction detection algorithm with applications to lung nodule detection," in *Proceedings of the Fourth IEEE International Symposium on Biomedical Imaging: From Nano to Macro ISBI'07*, Arlington VA, USA, April 12-15 2007, pp. 504–507.
- [28] G-S. Xia, J. Delon, and Y. Gousseau, "Accurate junction detection and characterization in natural images," *International Journal of Computer Vision*, vol. 106, no. 1, pp. 31–56, January 2014.
- [29] M. Maire, P. Arbelaez, C. Fowlkes, and J. Malik, "Using contours to detect and localize junctions in natural images," in *Proceedings of the 2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition CVPR'08*, Anchorage, Alaska, USA, June 24-26 2008, pp. 1–8.
- [30] T-A. Pham, M. Delalandre, S. Barrat, and J.-Y. Ramel, "Accurate junction detection and reconstruction in line-drawing images," in *Proceedings of the 21st International Conference on Pattern Recognition ICPR'12*, Tsukuba Science City, Japan, November 11-15 2012, pp. 693–696.
- [31] F. Osakada, H. Ikeda, M. Mandai, T. Wataya, K. Watanabe, N. Yoshimura, A. Akaike, Y. Sasai, and M. Takahashi, "Toward the generation of rod and cone photoreceptors from mouse, monkey and human embryonic stem cells," *Nature Biotechnology*, vol. 26, no. 2, pp. 215–224, February 2008.
- [32] P.V. Lehmann, "Image analysis and data management of elispot assay results," in *Handbook of ELISPOT*, A.E. Kalyuzhny, Ed., vol. 302 of *Methods in Molecular Biology*, pp. 117–131. Humana Press, 2005.
- [33] M.J. Ferraro J.H. Jorgensen, "Antimicrobial susceptibility testing: A review of general principles and contemporary practices," *Clinical Infectious Diseases*, vol. 49, no. 11, pp. 1749–1755, December 2009.
- [34] C. Damerval and S. Meignen, "Blob detection with wavelet maxima lines," *IEEE Signal Processing Letters*, vol. 14, no. 1, pp. 39–42, January 2007.

-
- [35] N. Chenouard and M. Unser, “3D steerable wavelets in practice,” *IEEE Transactions on Image Processing*, vol. 21, no. 11, pp. 4522–4533, November 2012.
- [36] D.H. Hubel and T.N. Wiesel, “Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex,” *The Journal of Physiology*, vol. 160, no. 1, pp. 106–154, January 1962.
- [37] B.A. Olshausen and D.J. Field, “Emergence of simple-cell receptive field properties by learning a sparse code for natural images,” *Nature*, vol. 381, pp. 607–609, June 1996.
- [38] D. Marr and E. Hildreth, “Theory of edge detection,” *Proceedings of the Royal Society of London. Series B, Biological Sciences*, vol. 207, no. 1167, pp. 187–217, February 1980.
- [39] Z. Püspöki, M. Storath, D. Sage, and M. Unser, “Transforms and operators for directional bioimage analysis: A survey,” *Advances in Anatomy, Embryology and Cell Biology*, 2016, in press.
- [40] J. Canny, “A computational approach to edge detection,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 6, pp. 679–698, June 1986.
- [41] B. Jähne, *Digital Image Processing: Concepts, Algorithms, and Scientific Applications*, Springer-Verlag New York, Inc., Secaucus, NJ, USA, 4th edition, 1997.
- [42] W. Förstner, “A feature based correspondence algorithm for image matching,” *International Archives of Photogrammetry and Remote Sensing*, vol. 26, no. 3, pp. 150–166, August 1986.
- [43] J. Bigün and G.H. Granlund, “Optimal orientation detection of linear symmetry,” in *Proceedings of the First IEEE International Conference on Computer Vision ICCV’87*, London, Great Britain, June 8-11 1987, pp. 433–438.
- [44] C. Harris and M. Stephens, “A combined corner and edge detector,” in *Proceedings of the Fourth Alvey Vision Conference*, Manchester, Great Britain, August 31-September 2 1988, pp. 147–151.
- [45] E. Fonck, G.G. Feigl, J. Fasel, D. Sage, M. Unser, D.A. Rüfenacht, and N. Stergiopoulos, “Effect of ageing on elastin functionality in human cerebral arteries,” in

Proceedings of the ASME 2008 Summer Bioengineering Conference SBC'08, Marco Island FL, USA, June 25-29, 2008, pp. SBC2008–192727–1/2.

- [46] U. Köthe, “Edge and junction detection with an improved structure tensor,” in *Pattern Recognition*. 2003, vol. 2781 of *Lecture Notes in Computer Science*, pp. 25–32, Springer.
- [47] J. Bigun, T. Bigun, and K. Nilsson, “Recognition by symmetry derivatives and the generalized structure tensor,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 12, pp. 1590–1605, December 2004.
- [48] A.F. Frangi, W.J. Niessen, K.L. Vincken, and M.A. Viergever, “Multiscale vessel enhancement filtering,” in *Medical Image Computing and Computer-Assisted Intervention*, 1998, vol. 1496 of *Lecture Notes in Computer Science*, pp. 130–137.
- [49] C. Lorenz, I.-C. Carlsen, T.M. Buzug, C. Fassnacht, and J. Weese, “Multi-scale line segmentation with automatic estimation of width, contrast and tangential direction in 2D and 3D medical images,” in *Proceedings of the First Joint Conference Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery CVRMed - MRCAS'97*, Grenoble, France, March 19-22 1997, pp. 233–242, Springer.
- [50] Y. Sato, S. Nakajima, N. Shiraga, H. Atsumi, S. Yoshida, T. Koller, G. Gerig, and R. Kikinis, “Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images,” *Medical Image Analysis*, vol. 2, no. 2, pp. 143–168, June 1998.
- [51] I. Daubechies, *Ten Lectures on Wavelets*, vol. 61, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1992.
- [52] D.S.G. Pollock and I. Lo Cascio, “Non-dyadic wavelet analysis,” in *Optimisation, Econometric and Financial Analysis*, vol. 9 of *Advances in Computational Management Science*, pp. 167–203. Springer Berlin Heidelberg, 2007.
- [53] C.K. Chui and X. Shi, “Orthonormal wavelets and tight frames with arbitrary real dilations,” *Applied and Computational Harmonic Analysis*, vol. 9, no. 3, pp. 243 – 264, October 2000.

-
- [54] P. Auscher, "Wavelet bases for $L^2(\mathbb{R})$ with rational dilation factor," in *Wavelets and Their Applications*, M. et al. Ruskai, Ed., pp. 439–451. Jones & Bartlett, Boston, 1992.
- [55] P. Auscher, *Ondelettes fractales et applications*, Ph.D. thesis, Université Paris, 1998.
- [56] M. Feilner, D. Van De Ville, and M. Unser, "An orthogonal family of quincunx wavelets with continuously adjustable order," *IEEE Transactions on Image Processing*, vol. 14, no. 4, pp. 499–510, April 2005.
- [57] D. Van De Ville, T. Blu, and M. Unser, "On the multidimensional extension of the quincunx subsampling matrix," *IEEE Signal Processing Letters*, vol. 12, no. 2, pp. 112–115, February 2005.
- [58] N. Kingsbury, "The dual-tree complex wavelet transform: a new technique for shift invariance and directional filters," in *IEEE Digital Signal Processing Workshop*, 1998, vol. 86, pp. 120–131.
- [59] I.W. Selesnick, R.G. Baraniuk, and N.C. Kingsbury, "The dual-tree complex wavelet transform," *IEEE Signal Processing Magazine*, vol. 22, no. 6, pp. 123–151, November 2005.
- [60] J. Ma and G. Plonka, "The curvelet transform," *IEEE Signal Processing Magazine*, vol. 27, no. 2, pp. 118–133, March 2010.
- [61] K.N. Chaudhury, A. Muñoz Barrutia, and M. Unser, "Fast space-variant elliptical filtering using box splines," *IEEE Transactions on Image Processing*, vol. 19, no. 9, pp. 2290–2306, September 2010.
- [62] B.G. Sherlock, D.M. Monro, and K. Millard, "Fingerprint enhancement by directional Fourier filtering," in *IEE Proceedings on Vision, Image and Signal Processing*, April 1994, vol. 141, pp. 87–94.
- [63] E. Franken, R. Duits, and B. ter Haar Romeny, "Nonlinear diffusion on the 2D Euclidean motion group," in *Proceedings of the First International Conference on Scale Space and Variational Methods in Computer Vision (SSVM'07)*, Ischia, Italy, May 30–June 2, 2007, pp. 461–472, Springer.

- [64] E. Franken and R. Duits, "Crossing-preserving coherence-enhancing diffusion on invertible orientation scores," *International Journal of Computer Vision*, vol. 85, no. 3, pp. 253–278, December 2009.
- [65] S. Mallat, *A wavelet tour of signal processing: The sparse way*, Academic Press, 2008.
- [66] T.S. Lee, "Image representation using 2D Gabor wavelets," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 10, pp. 959–971, October 1996.
- [67] J.G. Daugman, "Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters," *Journal of the Optical Society of America A*, vol. 2, no. 7, pp. 1160–1169, July 1985.
- [68] J.G. Daugman, "Complete discrete 2D Gabor transforms by neural networks for image analysis and compression," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 36, no. 7, pp. 1169–1179, July 1988.
- [69] J.-P. Antoine, R. Murenzi, and P. Vandergheynst, "Directional wavelets revisited: Cauchy wavelets and symmetry detection in patterns," *Applied and Computational Harmonic Analysis*, vol. 6, no. 3, pp. 314–345, May 1999.
- [70] P. Vandergheynst and J.-F. Gobbers, "Directional dyadic wavelet transforms: Design and algorithms," *IEEE Transactions on Image Processing*, vol. 11, no. 4, pp. 363–372, April 2002.
- [71] L. Jacques, L. Duval, C. Chaux, and G. Peyré, "A panorama on multiscale geometric representations, intertwining spatial, directional and frequency selectivity," *Signal Processing*, vol. 91, no. 12, pp. 2699–2730, December 2011.
- [72] D.K. Hammond, P. Vandergheynst, and R. Gribonval, "Wavelets on graphs via spectral graph theory," *Applied and Computational Harmonic Analysis*, vol. 30, no. 2, pp. 129–150, March 2011.
- [73] D.I. Shuman, S.K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Processing Magazine*, vol. 30, no. 3, pp. 83–98, May 2013.

- [74] A. Sandryhaila and J.M.F. Moura, "Discrete signal processing on graphs," *IEEE Transactions on Signal Processing*, vol. 61, no. 7, pp. 1644–1656, April 2013.
- [75] P. Kovesei, "Symmetry and asymmetry from local phase," in *Proceedings of the Tenth Australian Joint Conference on Artificial Intelligence*, Perth, Australia, November 30 - December 4 1997, Lecture Notes in Computer Science, pp. 2–4, Springer.
- [76] E. Candès and D. Donoho, "New tight frames of curvelets and optimal representations of objects with piecewise C2 singularities," *Communications on Pure and Applied Mathematics*, vol. 57, no. 2, pp. 219–266, February 2004.
- [77] M.N. Do and M. Vetterli, "The contourlet transform: an efficient directional multiresolution image representation," *IEEE Transactions on Image Processing*, vol. 14, no. 12, pp. 2091–2106, December 2005.
- [78] D. Labate, W.-Q. Lim, G. Kutyniok, and G. Weiss, "Sparse multidimensional representation using shearlets," in *Optics & Photonics 2005*, 2005, pp. 59140U–59140U.
- [79] S. Yi, D. Labate, G.R. Easley, and H. Krim, "A shearlet approach to edge analysis and detection," *IEEE Transactions on Image Processing*, vol. 18, no. 5, pp. 929–941, May 2009.
- [80] G. Kutyniok and D. Labate, *Shearlets: Multiscale analysis for multivariate data*, Birkhauser Basel, 2012.
- [81] L. Semler and L. Dettori, "Curvelet-based texture classification of tissues in computed tomography," in *Proceedings of the Thirteenth IEEE International Conference on Image Processing ICIP'06*, Atlanta, Georgia, USA, October 8-11 2006, IEEE, pp. 2165–2168.
- [82] Y. Dong, D. Tao, X. Li, J. Ma, and J. Pu, "Texture classification and retrieval using shearlets and linear regression," *IEEE Transactions on Cybernetics*, vol. 45, no. 3, pp. 358–369, March 2015.
- [83] J.-L. Starck, E. Candès, and D. Donoho, "The curvelet transform for image denoising," *IEEE Transactions on Image Processing*, vol. 11, no. 6, pp. 670–684, June 2002.

- [84] J.-L. Starck, F. Murtagh, E. Candès, and D. Donoho, “Gray and color image contrast enhancement by the curvelet transform,” *IEEE Transactions on Image Processing*, vol. 12, no. 6, pp. 706–717, June 2003.
- [85] J. Friel, “Sparse regularization in limited angle tomography,” *Applied and Computational Harmonic Analysis*, vol. 34, no. 1, pp. 117–141, January 2013.
- [86] E. Candès and D. Donoho, “Continuous curvelet transform: I. Resolution of the wavefront set,” *Applied and Computational Harmonic Analysis*, vol. 19, no. 2, pp. 162–197, September 2005.
- [87] G. Kutyniok and D. Labate, “Resolution of the wavefront set using continuous shearlets,” *Transactions of the American Mathematical Society*, vol. 361, no. 5, pp. 2719–2754, May 2009.
- [88] G. Kutyniok, “Data separation by sparse representations,” in *Compressed Sensing: Theory and Applications*. Cambridge University Press, 2012.
- [89] K. Guo, D. Labate, and W.-Q. Lim, “Edge analysis and identification using the continuous shearlet transform,” *Applied and Computational Harmonic Analysis*, vol. 27, no. 1, pp. 24–46, July 2009.
- [90] M. Storath, “Separation of edges in X-Ray images by microlocal analysis,” *Proceedings in Applied Mathematics and Mechanics PAMM*, vol. 11, no. 1, pp. 867–868, December 2011.
- [91] S.T. Ali, J.-P. Antoine, and J.-P. Gazeau, “Multidimensional wavelets and generalizations,” in *Coherent States, Wavelets, and Their Generalizations*, Theoretical and Mathematical Physics, pp. 411–456. Springer New York, 2014.
- [92] W.T. Freeman and E.H. Adelson, “Steerable filters for early vision, image analysis, and wavelet decomposition,” in *Proceedings of the Third International Conference on Computer Vision*, Osaka, Japan, December 4-7 1990, IEEE, pp. 406–415.
- [93] W. T. Freeman, *Steerable Filters and Local Analysis of Image Structure*, Ph.D. thesis, Massachusetts Institute of Technology, 1992.
- [94] W.T. Freeman and E.H. Adelson, “The design and use of steerable filters,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 9, pp. 891–906, September 1991.

-
- [95] P. Perona, “Steerable-scalable kernels for edge detection and junction analysis,” *Image and Vision Computing*, vol. 10, no. 10, pp. 663–672, December 1992.
- [96] E.P. Simoncelli and H. Farid, “Steerable wedge filters for local orientation analysis,” *IEEE Transactions on Image Processing*, vol. 5, no. 9, pp. 1377–1382, September 1996.
- [97] M. Unser and D. Van De Ville, “Wavelet steerability and the higher-order Riesz transform,” *IEEE Transactions on Image Processing*, vol. 19, no. 3, pp. 636–652, March 2010.
- [98] M. Unser and N. Chenouard, “A unifying parametric framework for 2D steerable wavelet transforms,” *SIAM Journal on Imaging Sciences*, vol. 6, no. 1, pp. 102–135, 2013.
- [99] J. P. Ward, K. N. Chaudhury, and M. Unser, “Decay properties of Riesz transforms and steerable wavelets,” *SIAM Journal on Imaging Sciences*, vol. 6, no. 2, pp. 984–998, 2013.
- [100] J.P. Ward and M. Unser, “Harmonic singular integrals and steerable wavelets in $l_2(\mathbb{R}^d)$,” *Applied and Computational Harmonic Analysis*, vol. 36, no. 2, pp. 183–197, March 2014.
- [101] M. Jacob and M. Unser, “Design of steerable filters for feature detection using Canny-like criteria,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 8, pp. 1007–1019, August 2004.
- [102] M. Mühlich, D. Friedrich, and T. Aach, “Design and implementation of multi-steerable matched filters,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 2, pp. 279–291, February 2012.
- [103] M. Felsberg and G. Sommer, “The monogenic signal,” *IEEE Transactions on Signal Processing*, vol. 49, no. 12, pp. 3136–3144, December 2001.
- [104] U. Köthe, “Low-level feature detection using the boundary tensor,” in *Visualization and Processing of Tensor Fields*, J. Weickert and H. Hagen, Eds., Mathematics and Visualization, pp. 63–79. Springer Berlin Heidelberg, 2006.

- [105] Z. Püspöki, V. Uhlmann, C. Vonesch, and M. Unser, "Design of steerable wavelets to detect multifold junctions," *IEEE Transactions on Image Processing*, vol. 25, no. 2, pp. 643–657, February 2016.
- [106] Z. Püspöki, C. Vonesch, and M. Unser, "Detection of symmetric junctions in biological images using 2-D steerable wavelet transforms," in *Proceedings of the Tenth IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI'13)*, San Francisco CA, USA, April 7-11, 2013, pp. 1488–1491.
- [107] K.G. Larkin, D.J. Bone, and M.A. Oldfield, "Natural demodulation of two-dimensional fringe patterns. I. General background of the spiral phase quadrature transform," *Journal of the Optical Society of America A*, vol. 18, pp. 1862–1870, August 2001.
- [108] K.G. Larkin, "Natural demodulation of two-dimensional fringe patterns. II. Stationary phase analysis of the spiral phase quadrature transform," *Journal of the Optical Society of America A*, vol. 18, pp. 1871–1881, August 2001.
- [109] R. Marchant and P. Jackway, "Modelling line and edge features using higher-order Riesz transforms," in *Advanced Concepts for Intelligent Vision Systems*, Jacques Blanc-Talon, Andrzej Kasinski, Wilfried Philips, Dan Popescu, and Paul Scheunders, Eds., vol. 8192 of *Lecture Notes in Computer Science*, pp. 438–449. Springer International Publishing, 2013.
- [110] E.P. Simoncelli and Freeman W.T., "The steerable pyramid: A flexible architecture for multi-scale derivative computation," in *Proceedings of the Second IEEE International Conference on Image Processing ICIP'95*, Washington, DC, USA, October 23-26 1995, IEEE, vol. 3, pp. 444–447.
- [111] A.A. Bharath and J. Ng, "A steerable complex wavelet construction and its application to image denoising," *IEEE Transactions on Image Processing*, vol. 14, no. 7, pp. 948–959, July 2005.
- [112] J. Portilla and E.P. Simoncelli, "A parametric texture model based on joint statistics of complex wavelet coefficients," *International Journal of Computer Vision*, vol. 40, no. 1, pp. 49–70, October 2000.
- [113] N. Karssemeijer and G.M. te Brake, "Detection of stellate distortions in mammograms," *IEEE Transactions on Medical Imaging*, vol. 15, no. 5, pp. 611–619, October 1996.

- [114] S. Held, M. Storath, P. Massopust, and B. Forster, “Steerable wavelet frames based on the Riesz transform,” *IEEE Transactions on Image Processing*, vol. 19, no. 3, pp. 653–667, March 2010.
- [115] A. Depeursinge, A. Foncubierta-Rodriguez, D. Van De Ville, and H. Müller, “Rotation-covariant texture learning using steerable Riesz wavelets,” *IEEE Transactions on Image Processing*, vol. 23, no. 2, pp. 898–908, February 2014.
- [116] A. Depeursinge, M. Yanagawa, A.N. Leung, and D.L. Rubin, “Predicting adenocarcinoma recurrence using computational texture models of nodule components in lung ct,” *Medical Physics*, vol. 42, no. 4, pp. 2054–2063, 2015.
- [117] A. Depeursinge, C. Kurtz, C.F. Beaulieu, S. Napel, and D.L. Rubin, “Predicting visual semantic descriptive terms from radiological image data: Preliminary results with liver lesions in CT,” *IEEE Transactions on Medical Imaging*, vol. 33, no. 8, pp. 1669–1676, August 2014.
- [118] M. Reisert and H. Burkhardt, “Harmonic filters for generic feature detection in 3D,” in *Proceedings of the Thirty-first Annual Symposium of the German Association for Pattern Recognition DAGM’09*, Joachim Denzler, Gunther Notni, and Herbert Susse, Eds. September 9-11 2009, vol. 5748 of *Lecture Notes in Computer Science*, pp. 131–140, Springer.
- [119] H. Skibbe, Q. Wang, O. Ronneberger, H. Burkhardt, and M. Reisert, “Fast computation of 3D spherical Fourier harmonic descriptors—A complete orthonormal basis for a rotational invariant representation of three-dimensional objects,” in *Proceedings of the Twelfth IEEE International Conference on Computer Vision Workshops ICCV’09*, Kyoto, Japan, September 29 - October 2 2009, pp. 1863–1869, IEEE.
- [120] M. Unser, N. Chenouard, and D. Van De Ville, “Steerable pyramids and tight wavelet frames in $l_2(\mathbb{R}^d)$,” *IEEE Transactions on Image Processing*, vol. 20, no. 10, pp. 2705–2721, October 2011.
- [121] Z. Püspöki and M. Unser, “Template-free wavelet-based detection of local symmetries,” *IEEE Transactions on Image Processing*, vol. 24, no. 10, pp. 3009–3018, October 2015.

- [122] Z. Püspöki and M. Unser, "Wavelet-based identification and classification of local symmetries in microscopy images," in *Proceedings of the Eleventh IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI'14)*, Beijing, People's Republic of China, April 29-May 2, 2014, pp. 1035–1038.
- [123] G. Jacovitti and A. Neri, "Multiresolution circular harmonic decomposition," *IEEE Transactions on Signal Processing*, vol. 48, no. 11, pp. 3242–3247, November 2000.
- [124] J.R. Romero, S.K. Alexander, S. Baid, S. Jain, and M. Papadakis, "The geometry and the analytic properties of isotropic multiresolution analysis," *Advances in Computational Mathematics*, vol. 31, no. 1, pp. 283–328, October 2009.
- [125] S.C. Olhede and G. Metikas, "The monogenic wavelet transform," *IEEE Transactions on Signal Processing*, vol. 57, no. 9, pp. 3426–3441, September 2009.
- [126] M. Unser, K. Balać, and D. Van De Ville, "The monogenic Riesz-Laplace wavelet transform," in *Proceedings of the Sixteenth European Signal Processing Conference EUSIPCO'08*, Lausanne VD, Swiss Confederation, August 25-29, 2008.
- [127] M. Unser, D. Sage, and D. Van De Ville, "Multiresolution monogenic signal analysis using the Riesz-Laplace wavelet transform," *IEEE Transactions on Image Processing*, vol. 18, no. 11, pp. 2402–2418, November 2009.
- [128] M. Storath, "Directional multiscale amplitude and phase decomposition by the monogenic curvelet transform," *SIAM Journal on Imaging Sciences*, vol. 4, no. 1, pp. 57–78, January 2011.
- [129] S. Häuser, B. Heise, and G. Steidl, "Linearized Riesz transform and quasis-monogenic shearlets," *International Journal of Wavelets, Multiresolution and Information Processing*, vol. 12, no. 3, February 2014.
- [130] B. Heise, M. Reinhardt, S.E. Schausberger, S. Häuser, S. Bernstein, and D. Stifter, "Fourier plane filtering revisited – analogies in optics and mathematics," *Sampling Theory in Signal and Image Processing*, vol. 13, no. 3, pp. 231–248, April 2014.
- [131] R. Souillard, P. Carré, and C. Fernandez-Maloigne, "Vector extension of monogenic wavelets for geometric representation of color images," *IEEE Transactions on Image Processing*, vol. 22, no. 3, pp. 1070–1083, March 2013.

-
- [132] M. Storath, L. Demaret, and P. Massopust, "Signal analysis based on complex wavelet signs," *Applied and Computational Harmonic Analysis*, 2016, In press, preprint arXiv:1208.4578.
- [133] M. Storath, A. Weinmann, and M. Unser, "Unsupervised texture segmentation using monogenic curvelets and the Potts model," in *Proceedings of the 2014 IEEE International Conference on Image Processing (ICIP'14)*, Paris, France, October 27–30, 2014, pp. 4348–4352.
- [134] K. Rohr, *Landmark-based image analysis : Using geometric and intensity models*, Computational imaging and vision. Kluwer Academic Publ. cop., Dordrecht, Boston, London, 2001.
- [135] Z. Püspöki, A. Amini, J.P. Ward, and M. Unser, "Cramér-Rao lower bounds for the detection of steerable features," *submitted*, 2016.
- [136] P. Stoica and R. Moses, *Spectral Analysis of Signals*, Prentice-Hall, Upper Saddle River, NJ, USA, 2005.
- [137] P. R. Smith, T. M. Peters, and R. H. T. Bates, "Image reconstruction from finite numbers of projections," *Journal of Physics A: Mathematical, Nuclear and General*, vol. 6, no. 3, pp. 361–382, March 1973.
- [138] B. Picinbono, "Second-order complex random vectors and normal distributions," *IEEE Transactions on Signal Processing*, vol. 44, no. 10, pp. 2637–2640, October 1996.
- [139] M.D. Abràmoff, P.J. Magalhães, and S.J. Ram, "Image processing with ImageJ," *Biophotonics International*, vol. 11, no. 7, pp. 36–42, July 2004.
- [140] P.Y. Huang, C.S. Ruiz-Vargas, A.M. van der Zande, W.S. Whitney, M.P. Levendorf, J.W. Kevek, S. Garg, J.S. Alden, C.J. Hustedt, Y. Zhu, J. Park, P.L. McEuen, and D. A. Muller, "Grains and grain boundaries in single-layer graphene atomic patchwork quilts," *Nature*, vol. 469, no. 7330, pp. 389–392, January 2011.
- [141] D. Slepian and H.O. Pollack, "Prolate spheroidal wave-functions, Fourier-analysis, and uncertainty—I," *Bell System Technical Journal*, vol. 40, pp. 43–63, January 1961.

- [142] D. Slepian, "Prolate spheroidal wave-functions, Fourier-analysis, and uncertainty—V: The discrete case," *Bell System Technical Journal*, vol. 57, pp. 1371–1430, May-June 1978.
- [143] D. Slepian, "Some asymptotic expansions for prolate spheroidal wave functions," *Journal of Mathematics and Physics*, vol. 44, no. 2, pp. 99–140, June 1965.
- [144] R. Bellman, "Dynamic programming and Lagrange multipliers," *The Bellman Continuum: A Collection of the Works of Richard E. Bellman*, p. 49, 1986.
- [145] J. Fischbarg, "Fluid transport across leaky epithelia: Central role of the tight junction and supporting role of aquaporins," *Physiological Reviews*, vol. 90, no. 4, pp. 1271–1290, October 2010.
- [146] J. Makhoul, "On the eigenvectors of symmetric Toeplitz matrices," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 29, no. 4, pp. 868–872, August 1981.
- [147] E.P. Simoncelli, W.T. Freeman, E.H. Adelson, and D.J. Heeger, "Shiftable multiscale transforms," *IEEE Transactions on Information Theory*, vol. 38, no. 2, pp. 587–607, March 1992.
- [148] P. Perona, "Deformable kernels for early vision," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 17, no. 5, pp. 488–499, May 1995.
- [149] P.C. Teo, *Theory and Applications of Steerable Functions*, Ph.D. thesis, Stanford University, Stanford, CA, USA, 1998.
- [150] P.C. Teo and Y. Hel-Or, "Lie generators for computing steerable functions," *Pattern Recognition Letters*, vol. 19, no. 1, pp. 7–17, January 1998.
- [151] P.C. Teo and Y. Hel-Or, "Design of multiparameter steerable functions using cascade basis reduction," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 21, no. 6, pp. 552–556, June 1999.
- [152] P. M. Anselone and J. Korevaar, "Translation invariant subspaces of finite dimension," *Proceedings of the American Mathematical Society*, vol. 15, pp. 747–752, 1964.

-
- [153] Z. Püspöki, J.P. Ward, D. Sage, and M. Unser, “On the continuous steering of the scale of tight wavelet frames,” *in press*, 2015.
- [154] Z. Püspöki, D. Sage, J.P. Ward, and M. Unser, “Spotcaliper: Fast wavelet-based spot detection with accurate size estimation,” *Bioinformatics*, 2015, *in press*.
- [155] E. M. Stein and G. Weiss, *Introduction to Fourier analysis on Euclidean spaces*, vol. 32 of *Princeton Mathematical Series*, Princeton University Press, Princeton, N.J., 1971.
- [156] K. N. Chaudhury and M. Unser, “On the shiftability of dual-tree complex wavelet transforms,” *IEEE Transactions on Signal Processing*, vol. 58, no. 1, pp. 221–232, January 2010.
- [157] Z. Püspöki, J.P. Ward, D. Sage, and M. Unser, “Fast detection and refined scale estimation using complex isotropic wavelets,” in *Proceedings of the Twelfth IEEE International Symposium on Biomedical Imaging: From Nano to Macro ISBI'15*, Brooklyn NY, USA, April 16-19, 2015, pp. 512–515.
- [158] P. Hough, “Method and Means for Recognizing Complex Patterns,” U.S. Patent 3.069.654, December 1962.
- [159] J.-C. Olivo-Marin, “Extraction of spots in biological images using multiscale products,” *Pattern Recognition*, vol. 35, no. 9, pp. 1989 – 1996, September 2002.
- [160] T. Lindeberg, “Feature detection with automatic scale selection,” *International Journal of Computer Vision*, vol. 30, pp. 79–116, November 1998.

Curriculum Vitæ



Zsuzsanna Püspöki

e-mail: zsuzsanna.puspoki@gmail.com

Education

PhD in Electrical Engineering (2016)

Ecole polytechnique fédérale de Lausanne (EPFL), Lausanne, Switzerland (2012 – 2016)

Thesis: «Local Geometric Transformations in Image Processing»

Advisor: Prof. Michael Unser

MSc in Electrical Engineering (Excellent with Honors, 2011)

Budapest University of Technology and Economics, Budapest, Hungary (2009 – 2011)

Major: «Control and Robotic Systems»

BSc in Electrical Engineering (Excellent with Honors, 2009)

Budapest University of Technology and Economics, Budapest, Hungary (2005 – 2009)

Major: «Embedded Information and Control Systems»

Experience

Visiting Scientist (2015)

Heidelberg Collaboratory for Image Processing, Heidelberg, Germany

Developed texture descriptors for Ilastik in Prof. Hamprecht's lab.

2nd ICY Coding Party (2015)

Institut Pasteur, Paris, France

Implemented a wavelet-based tool (SpotCaliper) for the analysis of circular shapes in biological micrographs in ICY.

11th IEEE EMBS International Summer School on Biomedical Imaging (2014)

Saint-Jacut-de-la-Mer, France

Attended lecture series from leading experts of different branches of biomedical imaging. Interacted with other researchers working on current bioimage analysis challenges.

R&D Engineer (Summer 2010)

Biomedical Imaging Group, EPFL, Lausanne, Switzerland

Developed algorithms and software implementations for stereo matching for 3D depth perception.

R&D Engineer (Summer 2009)

Biomedical Imaging Group, EPFL, Lausanne, Switzerland

Developed a software implementation for box-spline-based object detection in biomedical images.

Outcome: conference paper (ISBI'13).

Publications

JOURNAL

Z. Püspöki, A. Amini, J.P. Ward, M. Unser, «Cramér-Rao Lower Bounds for Steerable Feature Detection,» submitted.

A. Depeursinge, **Z. Püspöki**, J.P. Ward, M. Unser, «Steerable Wavelet Machines (SWM): Learning Moving Frames for Texture Classification,» submitted.

Z. Püspöki, J.P. Ward, D. Sage, M. Unser, «On the Continuous Steering of Tight Wavelet Frames,» SIAM Journal on Imaging Sciences, in press.

Z. Püspöki, D. Sage, J.P. Ward, M. Unser, «SpotCaliper: Fast Wavelet-Based Spot Detection with Accurate Size Estimation,» Bioinformatics, in press.

Z. Püspöki, M. Storath, D. Sage, M. Unser, «Transforms and Operators for Directional Bioimage Analysis: A Survey,» Advances in Anatomy, Embryology and Cell Biology, in press.

Z. Püspöki, V. Uhlmann, C. Vonesch, M. Unser, «Design of Steerable Wavelets to Detect Multifold Junctions,» IEEE Transactions on Image Processing, vol. 25, no. 2, pp. 643-657, February 2016.

Z. Püspöki, M. Unser, «Template-Free Wavelet-Based Detection of Local Symmetries,» IEEE Transactions on Image Processing, vol. 24, no. 10, pp. 3009-3018, October 2015.

CONFERENCE PAPERS

Z. Püspöki, J.P. Ward, D. Sage, M. Unser, «Fast Detection and Refined Scale Estimation Using Complex Isotropic Wavelets,» Proceedings of the Twelfth IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI'15), Brooklyn NY, USA, April 16-19, 2015, pp. 512-515.

Z. Püspöki, M. Unser, «Wavelet-Based Identification and Classification of Local Symmetries in Microscopy Images,» Proceedings of the Eleventh IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI'14), Beijing, People's Republic of China, April 29-May 2, 2014, pp. 1035-1038.

Z. Püspöki, C. Vonesch, M. Unser, «Detection of Symmetric Junctions in Biological Images Using 2-D Steerable Wavelet Transforms,» Proceedings of the Tenth IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI'13), San Francisco CA, USA, April 7-11, 2013, pp. 1488-1491.

K.N. Chaudhury, **Z. Püspöki**, A. Muñoz Barrutia, D. Sage, M. Unser, «Fast Detection of Cells Using a Continuously Scalable Mexican-Hat-Like Template,» Proceedings of the Seventh IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI'10), Rotterdam, the Netherlands, April 14-17, 2010, pp. 1277-1280.

CONFERENCE ABSTRACTS

Z. Püspöki, M. Unser, «Local Symmetry Assessment in Bioimages,» Eleventh IEEE EMBS International Summer School on Biomedical Imaging, Saint-Jacut-de-la-Mer, France, June 19-27, 2014.

Z. Püspöki, C. Vonesch, M. Unser, «Symmetric Junctions in Biological Micrographs,» Proceedings of the Second International Symposium in Applied Bioimaging: From Molecule to Man (ISAB'13), Porto, Portugal, October 3-4, 2013.

Awards and Honors

The Google Anita Borg Memorial Scholarship (2010, 2008)

Europe, the Middle East and Africa - finalist

Award of the Hungarian Republic for Students (2010, 2009, 2008, 2007)

General Electric (GE) Foundation Scholar Leaders Award (2007)

Teaching and Supervision

Teaching Assistant (TA)

Image Processing I (2012 - 2014)

Image Processing II (2013 - 2015)

Supervised the Following Students at EPFL:

Sahar Hanna (co-advised with Julien Fageot and Prof. Michael Unser), 2016

Project: «Java Optimization Package Development»

Robin Weiskopf (co-advised with Dr. Daniel Sage and Prof. Michael Unser), 2016

Project: «Spot Detection Plug-In Development for Image Processing Applications»

Lucio Genini (co-advised with Dr. John Paul Ward, Dr. Daniel Sage and Prof. Michael Unser), 2014

Master Thesis: «Scalable Wavelets on Biological Images»

Kilian Thomas (co-advised with Julien Fageot and Prof. Michael Unser), 2014

Project: «Junction Detection in Biomedical Micrographs»

Christopher Finelli (co-advised with Virginie Uhlmann and Prof. Michael Unser), 2014

Project: «Classifying Galaxy Type Using the ESA Euclid Space Telescope»

Pritish Chakravarty (co-advised with Virginie Uhlmann and Prof. Michael Unser), 2013

Project: «Designing Tip Detectors»

Emmanuel Girardin (co-advised with Prof. Michael Unser), 2013

Project: «Finding Strong Gravitational Lenses Using the ESA Euclid Space Telescope»

Professional Activities

Reviewer for the Following Journals

IEEE Transactions on Signal Processing, IEEE Transactions on Image Processing, IEEE Journal of Selected Topics in Signal Processing, Advances in Anatomy, Embryology and Cell Biology, Journal of Microscopy (Wiley), and Journal of Mathematical Imaging and Vision

Reviewer for the Following Conferences

IEEE International Symposium on Biomedical Imaging 2014, 2015

IEEE Student Member (2012 - ongoing)

Technical Skills

Coding and Computational Programming – Java, C/C++, MATLAB, Python, Mathematica, Comsol, HTML, CSS

Technologies and Libraries – ImageJ, Fiji, ICY, SVN, Git

Environments – Mac OS X, Windows, LaTeX

Languages

Hungarian – Native

English – C2 level certificate (IELTS)

French – Fluent both spoken and written

German – Limited working proficiency

Personal

Running – 30 km of Budapest (2015)

– Thyon-Dixence Swiss Alpine trail run (2015)