



Image Analysis and Reconstruction using a Wavelet Transform Constructed from a Reducible Representation of the Euclidean Motion Group*

REMCO DUITS

*Eindhoven University of Technology, Department of Biomedical Engineering,
P.O. Box 513, NL-5600 MB Eindhoven, The Netherlands*
R.Duits@tue.nl

MICHAEL FELSBURG AND GÖSTA GRANLUND

*Computer Vision Laboratory, Department of Electrical Engineering, Linköping University
S-58183 Linköping, Sweden*
mfe@isy.liu.se
gosta@isy.liu.se

BART TER HAAR ROMENY

*Eindhoven University of Technology, Department of Biomedical Engineering,
P.O. Box 513, NL-5600 MB Eindhoven, The Netherlands*
B.M.terHaarRomeny@tue.nl

Received December 13, 2004; Revised October 26, 2005; Accepted November 8, 2005

First online version published in July, 2006

Abstract. Inspired by the early visual system of many mammals we consider the construction of-and reconstruction from- an orientation score $U_f : \mathbb{R}^2 \times S^1 \rightarrow \mathbb{C}$ as a local orientation representation of an image, $f : \mathbb{R}^2 \rightarrow \mathbb{R}$. The mapping $f \mapsto U_f$ is a wavelet transform \mathcal{W}_ψ corresponding to a reducible representation of the Euclidean motion group onto $\mathbb{L}_2(\mathbb{R}^2)$ and oriented wavelet $\psi \in \mathbb{L}_2(\mathbb{R}^2)$. This wavelet transform is a special case of a recently developed generalization of the standard wavelet theory and has the practical advantage over the usual wavelet approaches in image analysis (constructed by irreducible representations of the similitude group) that it allows a stable reconstruction from one (single scale) orientation score. Since our wavelet transform is a unitary mapping with stable inverse, we directly relate operations on orientation scores to operations on images in a robust manner.

Furthermore, by geometrical examination of the Euclidean motion group $G = \mathbb{R}^2 \rtimes \mathbb{T}$, which is the domain of our orientation scores, we deduce that an operator Φ on orientation scores must be left invariant to ensure that the corresponding operator $\mathcal{W}_\psi^{-1} \Phi \mathcal{W}_\psi$ on images is Euclidean invariant. As an example we consider all linear second order left invariant evolutions on orientation scores corresponding to stochastic processes on G . As an application we detect elongated structures in (medical) images and automatically close the gaps between them.

Finally, we consider robust orientation estimates by means of channel representations, where we combine robust orientation estimation and learning of wavelets resulting in an auto-associative processing of orientation features. Here linear averaging of the channel representation is equivalent to robust orientation estimation and an adaptation of the wavelet to the statistics of the considered image class leads to an auto-associative behavior of the system.

* The Netherlands Organization for Scientific Research is gratefully acknowledged for financial support. This work has been supported by EC Grant IST-2003-004176 COSPAL.

1. Introduction

In many medical image applications it is desirable to construct a local orientation-score of a grey-value image. In the case of 2D images $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ such an orientation score $U_f : \mathbb{R}^2 \times \mathbb{T} \rightarrow \mathbb{C}$ depends on 3 variables $(b_1, b_2, e^{i\theta})$, where $\mathbf{b} = (b_1, b_2) \in \mathbb{R}^2$ denote position and where $e^{i\theta} \in \mathbb{T} \leftrightarrow (\cos \theta, \sin \theta) \in S^1$ is a local orientation variable.¹

Mostly, such an **orientation score** is obtained by means of a convolution with some anisotropic wavelet $\psi \in \mathbb{L}_2(\mathbb{R}^2) \cap \mathbb{L}_2(\mathbb{R}^2)$, cf. Kalitzin et al. (1999), van Ginkel (2002):

$$U_f(\mathbf{b}, e^{i\theta}) = \int_{\mathbb{R}^2} \overline{\psi(R_\theta^{-1}(\mathbf{x}' - \mathbf{b}))} f(\mathbf{x}') \, d\mathbf{x}',$$

with $R_\theta = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}$. (1)

This idea is inspired by the early visual system of many mammals, in which receptive fields exist that are tuned to various locations and orientations. Thereby a simple cell receptive field can be parameterized by its position and orientation. Assemblies of oriented receptive fields are grouped together on the surface of the primary visual cortex in a pinwheel like structure, known as the orientation preference structure. The orientation preference structure is an almost everywhere smooth mapping of the Euclidean motion group space $\mathbb{R}^2 \times \mathbb{T}$ onto the 2D surface. Due to the difference in dimensionality, the orientation preference structure is punctuated by so-called pinwheels, which are singularities in this mapping, see Figure 1. Perceptual organization (or image enhancement) on the basis of orientation similarity on images f can be done via their orientation scores U_f , as there exists a linear *well-posed invertible* transformation \mathcal{W}_ψ from the image f to the orientation score U_f and vice versa. This invertibility ensures that no information is lost in the decomposition of an input image into various orientations. As a model for the orientation preference structure in the visual system this implies that the orientation score may serve as an alternative format to the input luminance function, since there is no loss of data evidence.² As a tool for image processing, however, the inverse mapping from orientation score to original image is a very useful one as we will see later.

The domain of an orientation score U_f is the well-known Euclidean motion group $G = \mathbb{R}^2 \times \mathbb{T}$, with

group product

$$g g' = (\mathbf{b}, e^{i\theta})(\mathbf{b}', e^{i\theta'}) = (\mathbf{b} + R_\theta \mathbf{b}', e^{i(\theta+\theta')}),$$

$$g = (\mathbf{b}, e^{i\theta}), g' = (\mathbf{b}', e^{i\theta'}) \in \mathbb{R}^2 \times \mathbb{T}.$$

and the mapping $f \mapsto U_f$ is a wavelet transformation

$$U_f(\mathbf{b}, e^{i\theta}) = (\mathcal{W}_\psi[f])(g) = (\mathcal{U}_g \psi, f)_{\mathbb{L}_2(\mathbb{R}^2)}$$

$$= (\mathcal{T}_\mathbf{b} \mathcal{R}_{e^{i\theta}} \psi, f)_{\mathbb{L}_2(\mathbb{R}^2)}, \quad g = (\mathbf{b}, e^{i\theta}), \quad (2)$$

where $\mathcal{T}_\mathbf{b} \mathcal{R}_{e^{i\theta}} \psi$ is the translated and rotated wavelet and the representation $g \mapsto \mathcal{U}_g$ is given by

$$\mathcal{U}_g \psi(\mathbf{x}) = (\mathcal{T}_\mathbf{b} \mathcal{R}_{e^{i\theta}} \psi)(\mathbf{x}) = \psi(R_\theta^{-1}(\mathbf{x} - \mathbf{b})),$$

$$g = (\mathbf{b}, e^{i\theta}) \in G, \mathbf{x} \in \mathbb{R}^2, \quad (3)$$

where $(\mathcal{T}_\mathbf{b} \psi)(\mathbf{x}) = \psi(\mathbf{x} - \mathbf{b})$, $\mathbf{x} \in \mathbb{R}^d$ and $(\mathcal{R}_{e^{i\theta}} \psi)(\mathbf{x}) = \psi(R_\theta^{-1} \mathbf{x})$, with $R_\theta, \theta \in [0, 2\pi)$, the counter clock-wise rotation given in (1).

Because the local orientation is explicitly encoded in the orientation score, it is much easier to do (enhancement or perceptual organization) operations based on local orientations on the score. In section 3 we quantify the stability of \mathcal{W}_ψ and its inverse, by means of a functional Hilbert space approach to the theory of wavelets. In fact, this approach leads to a generalization of standard wavelet theory, which is *necessary* for our application. Here we will only give a brief explanation of this theory and restrict ourselves to the practical consequences. For a more in-depth mathematical treatment we refer to earlier work, cf. (Duits, 2005), (Duits and Duits, 2004), (Duits, 2004) and (Duits et al., 2004).

In section 4 we discuss explicit constructions of wavelets (so-called proper wavelets) that allow a stable (re)construction. Using these proper wavelets we can do image processing via processing on its score, which is discussed in section 5. This kind of image processing, which is useful for detection and completion of elongated structures in (medical) imaging, is also generalized to image processing via invertible orientation scores of 3D images. This is briefly discussed in section 6, where we do not deal with all technical details, but just show some preliminary results.

Finally, in section 7 we focus on robust orientation estimation rather than image enhancement. For this purpose we will describe a different paradigm, based on channel representations. Here we will also discuss

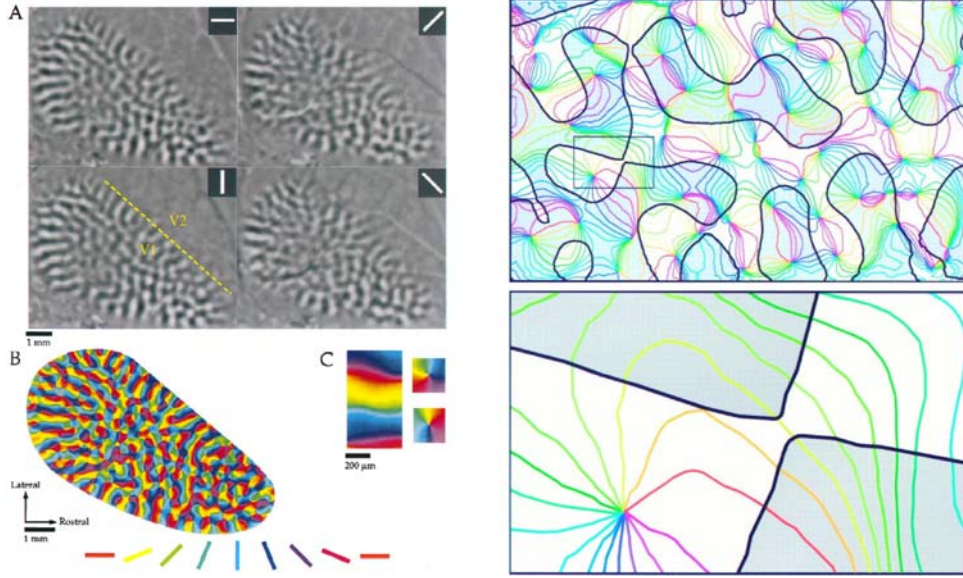


Figure 1. Left: A: Parts of visual cortex active under different orientation stimuli. B: Orientation preference map obtained by vector summation of data obtained for each angle. Orientation preference is color coded according to the key shown below, replicated with permission from Bosking et al. (1997), copyright 1997 Society of Neuroscience. Right: enlarged section of the rectangular area in the upper Figure. Shaded and unshaded areas denote the left and right eye resp. Colored lines connect cells with equal orientation sensitivity, replicated with permission from Ts'o et al. (1990).

the similarity and complementarity of channel representations, with the foregoing theory of orientation scores.

2. Preliminaries and Notation

- The Fourier transform $\mathcal{F} : \mathbb{L}_2(\mathbb{R}^d) \rightarrow \mathbb{L}_2(\mathbb{R}^d)$, is almost everywhere defined by

$$[\mathcal{F}(f)](\omega) = \hat{f}(\omega) = \frac{1}{(2\pi)^{d/2}} \int_{\mathbb{R}^d} f(\mathbf{x}) e^{-i\omega \cdot \mathbf{x}} d\mathbf{x}.$$

Notice that $\|\mathcal{F}[f]\|^2 = \|f\|^2$ and $\mathcal{F}[f * g] = (2\pi)^{\frac{d}{2}} \mathcal{F}[f]\mathcal{F}[g]$, for all $f, g \in \mathbb{L}_2(\mathbb{R}^d)$.

- We use the following notation for Euclidean/polar coordinates in spatial and Fourier domain, respectively: $\mathbf{x} = (x, y) = (r \cos \phi, r \sin \phi)$, $\omega = (\omega_x, \omega_y) = (\rho \cos \varphi, \rho \sin \varphi)$, with $\phi, \varphi \in [0, 2\pi)$, $r, \rho > 0$. The corresponding complex variables will be denoted by $z = x + iy = r e^{i\phi}$ and $w = \omega_x + i\omega_y = \rho e^{i\varphi}$.
- Images are assumed to be within $\mathbb{L}_2(\mathbb{R}^d)$. We mainly consider $d = 2$, unless explicitly stated otherwise. The space of bandlimited (by $\varrho > 0$) images is given

by

$$\mathbb{L}_2^{\varrho}(\mathbb{R}^d) = \{f \in \mathbb{L}_2(\mathbb{R}^2) \mid \text{supp}(\mathcal{F}[f]) \subset B_{\mathbf{0}, \varrho}\}, \quad \varrho > 0, \quad (4)$$

where $B_{\mathbf{0}, \varrho} = \{\omega \in \mathbb{R}^d \mid \|\omega\| < \varrho\}$.

- We will use short notation for the following groups:
 - $\text{Aut}(\mathbb{R}^d) = \{A : \mathbb{R}^d \rightarrow \mathbb{R}^d \mid A \text{ linear and } A^{-1} \text{ exists}\}$
 - dilation group $D(d) = \{A \in \text{Aut}(\mathbb{R}^d) \mid A = aI, a > 0\}$.
 - orthogonal group $O(d) = \{X \in \text{Aut}(\mathbb{R}^d) \mid X^T = X^{-1}\}$
 - rotation group $SO(d) = \{R \in O(d) \mid \det(R) = 1\}$.
 - circle group $\mathbb{T} = \{z \in \mathbb{C} \mid |z| = 1\}$, $z = e^{i\theta}$, $\theta = \arg z$ with group homomorphism $\tau : \mathbb{T} \rightarrow SO(2) \subset \text{Aut}(\mathbb{R}^2)$, given by $\tau(z) = R_\theta$, recall (1)
- With $\mathcal{B}(H)$ we denote the space of bounded operators on H . The range of a linear operator A will be denoted by $\mathcal{R}(A)$ and its nilspace will be denoted by $\mathcal{N}(A)$.

- A representation \mathcal{R} of a group G onto a Hilbert space H is a homomorphism \mathcal{R} between G and $\mathcal{B}(H)$, the space of bounded linear operators on H . It satisfies $\mathcal{R}_{gh} = \mathcal{R}_g \mathcal{R}_h$ for all $g \in G, h \in G$ and $\mathcal{R}_e = I$. A representation \mathcal{R} is irreducible if the only invariant closed subspaces of H are H and $\{0\}$ and otherwise reducible. We mainly consider unitary representations (i.e. $\|\mathcal{U}_g \psi\|_H = \|\psi\|_H$ for all $g \in G$ and $\psi \in H$), which will be denoted by \mathcal{U} rather than \mathcal{R} .

In this article we mainly consider the left regular representation of the Euclidean motion group on $\mathbb{L}_2(\mathbb{R}^d)$, which are given by (3).

- Let $\mathbf{b} \in \mathbb{R}^d, a > 0$ and $g \in G$ with corresponding $\tau(g) \in \text{Aut}(\mathbb{R}^d)$. Then the unitary operators $f \mapsto \check{f}, \tau_{\mathbf{b}}, \mathcal{D}_a$ and \mathcal{R}_g on $\mathbb{L}_2(\mathbb{R}^d)$ are defined by

$$\begin{aligned} \check{f}(\mathbf{x}) &= f(-\mathbf{x}) & \mathcal{T}_{\mathbf{b}}\psi(\mathbf{x}) &= \psi(\mathbf{x} - \mathbf{b}) \\ \mathcal{D}_a\psi(\mathbf{x}) &= \frac{1}{a^2}\psi\left(\frac{\mathbf{x}}{a}\right) & \mathcal{R}_g\psi(\mathbf{x}) &= \frac{1}{\sqrt{\det\tau(g)}}\psi((\tau(g))^{-1}\mathbf{x}), \end{aligned} \quad (5)$$

which are left regular actions of $O(1), \mathbb{R}^d, D(d)$ respectively G in $\mathbb{L}_2(\mathbb{R}^d)$.

- The 2D-Gaussian kernel G_s at scale s is given by $G_s(\mathbf{x}) = \frac{1}{4\pi s} e^{-\frac{\|\mathbf{x}\|^2}{4s}}$.

3. Quantification of Wellposedness of Transformations between Image and Orientation Score

Because the local orientation is explicitly encoded in the orientation score, it is much easier to do (enhancement or perceptual organization) operations based on local orientations on the score. However, well-posed image enhancement on the basis of orientation similarity in an image f (without loss of data evidence) can be done via its orientation score U_f iff there exists a stable transformation from image f to U_f and vice versa. Stability entails that a small³ perturbation in the image, must correspond to a small perturbation on the orientation score and vice versa. For instance in the case of Fourier transformation, the stability is ensured by the Plancherel theorem, which states that $\|\mathcal{F}(f)\|_{\mathbb{L}_2(\mathbb{R}^d)}^2 = \|f\|_{\mathbb{L}_2(\mathbb{R}^d)}^2$ for all images $f \in \mathbb{L}_2(\mathbb{R}^d)$. In standard wavelet theory there also exists such a theorem, but this is not applicable to our case, which brings us to the challenge of generalizing standard wavelet theory.

3.1. Why Standard Wavelet Theory is not Applicable to our Application

In this subsection we first explain why standard wavelet theory, summarized in Theorem 1, can not be applied to the framework of orientation scores. Then we give a short summary (only as far as relevant for our purpose) of the results from a new wavelet theory which we developed in earlier work. For more mathematical background and generalizations we refer to cf. (Duits, 2005), (Duits and Duits, 2004), (Duits, 2004) and (Duits et al., 2004).

Let H be a Hilbert space, e.g. the space of images $\mathbb{L}_2(\mathbb{R}^d)$. Let \mathcal{U} be an irreducible unitary representation of the locally compact group G , with left-invariant Haar measure μ_G . Recall the definitions in the preliminaries. Let $\psi \in H$ be an admissible wavelet, which means that

$$C_\psi = \int_G \frac{|(\mathcal{U}_g \psi, \psi)|^2}{(\psi, \psi)} d\mu_G(g) < \infty$$

then the wavelet transform $\mathcal{W}_\psi : H \rightarrow \mathbb{L}_2(G)$ is defined by

$$\mathcal{W}_\psi[f](g) = (\mathcal{U}_g \psi, f)_H.$$

The next theorem is well-known in mathematical physics (Twareque Ali, 1998), and is first formulated and proven in Grossmann et al. (1985). For a simple alternative and self-contained proof, see (Duits and Duits, 2004) p. 20, which uses a topological version of Schur's lemma, for proof see (Duits, 2004) p. 86.

Theorem 1 The Wavelet Reconstruction Theorem.

The wavelet transform is a linear isometry (up to a constant) from the Hilbert space H onto a closed subspace $\mathbb{C}_{K_\psi}^G$ of $\mathbb{L}_2(G, d\mu)$:

$$\|\mathcal{W}_\psi[f]\|_{\mathbb{L}_2(G)}^2 = C_\psi \|f\|^2 \quad (6)$$

The space $\mathbb{C}_{K_\psi}^G$ is the unique functional Hilbert space with reproducing kernel $K_\psi(g, g') = \frac{1}{C_\psi} (\mathcal{U}_g \psi, \mathcal{U}_{g'} \psi)$. The corresponding orthogonal projection $\mathbb{P}_\psi : \mathbb{L}_2(G, d\mu) \rightarrow \mathbb{C}_{K_\psi}^G$ is given by

$$\begin{aligned} (\mathbb{P}_\psi \Phi)(g) &= \int_G K_\psi(g, g') \Phi(g') d\mu_G(g') \\ \Phi &\in \mathbb{L}_2(G, d\mu). \end{aligned} \quad (7)$$

Furthermore, \mathcal{W}_ψ intertwines \mathcal{U} and the left regular representation \mathcal{L} —i.e. \mathcal{L}_g is given by $\mathcal{L}_g(\Phi) = (h \mapsto \Phi(g^{-1}h))$ —on $\mathbb{L}_2(G) : \mathcal{W}_\psi \mathcal{U}_g = \mathcal{L}_g \mathcal{W}_\psi$

We notice that if \mathcal{U} is the left regular representation of $G = \mathbb{R} \rtimes D(1) \rtimes O(1)$ (the group consisting of translations, dilations and polarity) onto $H = \mathbb{L}_2(\mathbb{R})$ one obtains the more familiar framework of wavelet theory in 1D signal processing, (Duits et al., 2004) p. 5–6.

Of course, we would like to apply Theorem 1 to the wavelet transformation that maps an image to its orientation score, recall (2), since it would imply that reconstruction of an image from its orientation score is perfectly well-posed in the sense that (just like in Fourier transform) the quadratic norm is preserved. Unfortunately, the next lemma shows us that we are not allowed to do so in our case. Therefore in earlier work, cf. Duits et al. (2004), (Duits and Duits, 2004), (Duits, 2005) we generalized the standard wavelet theory in such a way that irreducibility is neither a requirement nor replaced by a requirement.

Lemma 2. *The left-regular action \mathcal{U} of the Euclidean motion group in $\mathbb{L}_2(\mathbb{R}^2)$, given by (2), is a reducible representation.*

Proof: Consider the subspace consisting of \mathbb{L}_2 -functions whose Fourier transform have a support inside a given disk around the origin with radius, say $a > 0$, i.e. $\mathbb{L}_2^a(\mathbb{R}^2) = \{f \in \mathbb{L}_2(\mathbb{R}^2) \mid \text{supp}(\mathcal{F}[f]) \subset B_{0,a}\}$, then this is a closed proper subspace of $\mathbb{L}_2(\mathbb{R}^2)$ which is invariant under \mathcal{U} , which directly follows by $\mathcal{F}[\mathcal{U}_g \psi] = e^{i\omega \cdot \mathbf{b}} \mathcal{R}_{e^{i\theta}} \mathcal{F}[\psi]$, for all $\psi \in \mathbb{L}_2^a(\mathbb{R}^2)$. \square

We could consider the similitude group $SIM(2) = \mathbb{R}^2 \rtimes \mathbb{T} \times D(1)$ with representation

$$\mathcal{V}_{\mathbf{b}, e^{i\theta}, a} \psi(\mathbf{x}) = \frac{1}{\sqrt{a}} \psi \left(\frac{R_\theta^{-1}(\mathbf{x} - \mathbf{b})}{a} \right),$$

$$a > 0, \theta \in [0, 2\pi), \mathbf{b} \in \mathbb{R}^2$$

which is irreducible, for proof see Louis et al. (1997) p. 51–52. This brings us within the standard wavelet frameworks in 2D image analysis (in particular to 2D Gabor wavelets, (Lee 1996), or Cauchy-wavelets (Antoine, 1999)). But from an implementation/practical point of view we do not want to consider multiple scales, but stick to a single scale. This pertains

to the so-called *Euclidean coherent states* from mathematical physics, (Isham and Klauder, 1991), which are not to be confused with the more familiar Euclidean coherent states constructed from the irreducible⁴ representations of the Euclidean motion group onto $\mathbb{L}_2(S^1)$, cf. (Ali et al., 1999) p. 219–220.

Omitting the dilation group poses an important question. For example in scale space theory, (Duits et al., 2004), it is a well-known problem that reconstruction of a sharp image f , from its (say Gaussian) blurred version $f * G_s$ is extremely ill-posed: Is it possible to get around this ill-posedness by considering all rotated versions of linear combinations of Gaussian derivatives $f * (\partial_x)^p (\partial_y)^q G_s$? Before we give an affirmative answer to this question and deal with the issue of well-posed reconstruction of images from orientation scores, we give an illustration by means of an extremely simplified discrete example, where reconstruction is done by integration over discrete orientations, rather than inverse convolution.

Example: Suppose we construct a discrete orientation score with only 4 orientations, up, down left and right, constructed with the following discrete oriented wavelet $\psi : \mathbb{Z} \times \mathbb{Z} \rightarrow \mathbb{R}$, given by

$$\psi[x^1, x^2] = \begin{cases} 1 & \text{if } (x^1, x^2) \in \{(0, 0), (1, 0)\} \\ -1/3 & \text{if } (x^1, x^2) \in \{(0, 1), (0, -1), (-1, 0)\} \\ 0 & \text{else,} \end{cases}$$

see Figure 2. Then reconstruction of the original discrete image $f : \mathbb{Z} \times \mathbb{Z} \rightarrow \mathbb{R}$ from its orientation score is done by integration over all directions.

$$f[x^1, x^2] = \frac{1}{4} \sum_{k=1}^4 U_f^k[x^1, x^2, e^{ik\pi/2}]. \quad (8)$$

3.2. Generalization of Standard Wavelet Theory by Reproducing Kernel Theory

Before we formulate the main theorem (Theorem 4) from which we can quantify the stability of the transformations between image and orientation score, we give some short explanation on reproducing kernel Hilbert spaces (also called functional Hilbert spaces), which is necessary to read and understand the theorem.

A reproducing kernel Hilbert space is a Hilbert space consisting of complex valued functions on an index set \mathbb{I} on which the point evaluation $\delta_{\mathbf{a}}$, given by $\delta_{\mathbf{a}}(f) = f(\mathbf{a})$ is a continuous linear functional for all $\mathbf{a} \in \mathbb{I}$. This means that $\delta_{\mathbf{a}}(f_n) = f_n(\mathbf{a}) \rightarrow \delta_{\mathbf{a}}(f) = f(\mathbf{a})$ for every

$$\begin{array}{|c|c|c|} \hline 0 & -\frac{1}{3} & 0 \\ \hline -\frac{1}{3} & 1 & 1 \\ \hline 0 & -\frac{1}{3} & 0 \\ \hline \end{array} +
\begin{array}{|c|c|c|} \hline 0 & 1 & 0 \\ \hline -\frac{1}{3} & 1 & -\frac{1}{3} \\ \hline 0 & -\frac{1}{3} & 0 \\ \hline \end{array} +
\begin{array}{|c|c|c|} \hline 0 & -\frac{1}{3} & 0 \\ \hline 1 & 1 & -\frac{1}{3} \\ \hline 0 & -\frac{1}{3} & 0 \\ \hline \end{array} +
\begin{array}{|c|c|c|} \hline 0 & -\frac{1}{3} & 0 \\ \hline -\frac{1}{3} & 1 & -\frac{1}{3} \\ \hline 0 & 1 & 0 \\ \hline \end{array} =
\begin{array}{|c|c|c|} \hline 0 & 0 & 0 \\ \hline 0 & 4 & 0 \\ \hline 0 & 0 & 0 \\ \hline \end{array}$$

Figure 2. Integrating the discrete orientation score U_f^4 over its 4 discrete orientations (8), boils down to convolution with the discrete spike δ .

sequence $\{f_n\}$ in H which converges to f , $f_n \rightarrow f$. It is not difficult to show that a linear functional is continuous if and only if it is bounded. So $\delta_{\mathbf{a}}$ is a continuous linear functional if and only if there exists a constant $C_{\mathbf{a}}$ such that $|f(\mathbf{a})| \leq C_{\mathbf{a}}\|f\|_H$. For example, the spaces $\mathbb{L}_2(\mathbb{R}^d)$ are *not* functional Hilbert spaces, but the well known first order Sobolev space $\mathbb{H}_1(\mathbb{R})$ is such a functional Hilbert space. Another example which is related to image processing is the space of bandlimited images (on a square $[0, a] \times [0, a]$), where the reproducing kernel⁵ is given by $K(\mathbf{x}, \mathbf{x}') = \frac{1}{4a^2\pi^2} \text{sinc}(a(x - x'))\text{sinc}(a(y - y')) = \mathcal{F}^{-1}[1_{[-a, a] \times [-a, a]}](\mathbf{x} - \mathbf{x}')$ and as is pointed out in Duits (2005) p. 121, p. 126–127 and (Duits, 2004) the Nyquist theorem is a direct consequence of Theorem 3 below.

If H is a functional Hilbert space, then $\delta_{\mathbf{a}}$ is a continuous linear functional, so that by the Riesz representation theorem it has a Riesz representant $K_{\mathbf{a}} \in H$ such that

$$f(\mathbf{a}) = \delta_{\mathbf{a}}(f) = (K_{\mathbf{a}}, f)_H,$$

for every $\mathbf{a} \in \mathbb{I}$. The function $K : \mathbb{I} \times \mathbb{I} \rightarrow \mathbb{C}$ given by $K(\mathbf{a}, \mathbf{b}) = (K_{\mathbf{a}}, K_{\mathbf{b}})_H = K_{\mathbf{b}}(\mathbf{a})$ is called reproducing kernel. This reproducing kernel is a function of positive type on \mathbb{I} , i.e.

$$\sum_{i=1}^n \sum_{j=1}^n K(m_i, m_j) \bar{c}_i c_j \geq 0, \text{ for all } n \in \mathbb{N}, \\ c_1, \dots, c_n \in \mathbb{C}, m_1, \dots, m_n \in \mathbb{I}.$$

Conversely, as Aronszajn pointed out in his paper, cf. (Aronszajn, 1950), a function K of positive type on a set \mathbb{I} uniquely induces a reproducing kernel Hilbert space consisting of functions on \mathbb{I} with reproducing kernel K . We denote this unique reproducing kernel Hilbert space by $\mathbb{C}_K^{\mathbb{I}}$. For the explicit construction of this space, we refer to Duits (2005) p. 120, p. 221–222 or (Martens, 2004), (Aronszajn, 1950). We notice that this (abstract) construction, is somewhat disappointing from the engineering point of view as it does not lead to a simple tangible description of the inner product:

Only in some cases it is possible to obtain explicit tangible descriptions of these inner products. Nevertheless, there is the following very general result⁶:

Theorem 3. *Let $V = \{\phi_m \mid m \in \mathbb{I}\}$ be a subset of H such that its linear span is dense in H . Define the function $K : \mathbb{I} \times \mathbb{I} \rightarrow \mathbb{C}$ by $K(m, m') = (\phi_m, \phi_{m'})_H$. Then the transform $\mathcal{W} : H \mapsto \mathbb{C}_K^{\mathbb{I}}$ defined by*

$$(\mathcal{W}[f])(m) = (\phi_m, f)_H \quad (9)$$

is a unitary mapping, i.e. $\|\mathcal{W}[f]\|_{\mathbb{C}_K^{\mathbb{I}}} = \|f\|_H$ for all $f \in H$.

Proof: see Duits and Duits (2004) p. 8 or (Duits, 2005) p. 221–222.

By applying this theorem to the case $H = \mathbb{L}_2(\mathbb{R}^2)$, $\mathbb{I} = G$, where $G = \mathbb{R}^2 \rtimes \mathbb{T} \equiv \mathbb{R}^2 \rtimes SO(2)$ and $V = \{\mathcal{U}_g \psi \mid g \in G\}$, which is dense in $\mathbb{L}_2(\mathbb{R}^2)$ iff

$$0 < M_{\psi}(\omega) = (2\pi) \int_0^{2\pi} |\mathcal{F}(\psi)(\rho \cos \theta, \rho \sin \theta)|^2 d\theta \\ < \infty, \quad \rho = \|\omega\|, \quad (10)$$

almost everywhere on \mathbb{R}^2 , and by characterizing the inner product on the space $\mathbb{C}_K^{G=\mathbb{R}^2 \rtimes \mathbb{T}}$ we obtain the following result:

Theorem 4. *The space of orientation scores is a reproducing kernel Hilbert space $\mathbb{C}_K^{\mathbb{R}^2 \rtimes \mathbb{T}}$ which is a closed subspace of $\mathbb{H}_{\psi} \otimes \mathbb{L}_2(\mathbb{T}; \frac{\mu_{\mathbb{T}}}{\det(\tau(i))})$ which is a vector subspace⁷ of $\mathbb{L}_2(G)$, where $\mathbb{H}_{\psi} = \{f \in \mathbb{L}_2(\mathbb{R}^d) \mid M_{\psi}^{-\frac{1}{2}} \mathcal{F}[f] \in \mathbb{L}_2(\mathbb{R}^d)\}$. The inner product on $\mathbb{C}_K^{\mathbb{R}^2 \rtimes \mathbb{T}}$ is given by $(\Phi, \Psi)_{M_{\psi}} = (T_{M_{\psi}}[\Phi], T_{M_{\psi}}[\Psi])_{\mathbb{L}_2(G)}$ where*

$$[T_{M_{\psi}}[\Phi]](\mathbf{b}, \theta) = \mathcal{F}^{-1} \left[\omega \mapsto M_{\psi}^{-\frac{1}{2}}(\omega) \mathcal{F}[\Phi(\cdot, e^{i\theta})](\omega) \right](\mathbf{b}),$$

which is thereby explicitly characterized by means of the function M_{ψ} given in (10). The wavelet transformation which maps an image $f \in \mathbb{L}_2(\mathbb{R}^2)$ onto its

orientation score $U_f \in \mathbb{C}_K^{\mathbb{R}^2 \times \mathbb{T}}$ is a unitary mapping:
 $\|f\|_{\mathbb{L}_2(\mathbb{R}^2)}^2 = \|U_f\|_{M_\psi}^2 = (U_f, U_f)_{M_\psi}$.
 As a result the image f can be exactly reconstructed from its orientation score $U_f = \mathcal{W}_\psi[f]$ by means of the adjoint wavelet transformation \mathcal{W}_ψ^* :

$$f = \mathcal{W}_\psi^* \mathcal{W}_\psi[f]$$

$$= \mathcal{F}^{-1} \left[\omega \mapsto \int_0^{2\pi} \mathcal{F}[U_f(\cdot, e^{i\theta})](\omega) \mathcal{F}[\mathcal{R}_{e^{i\theta}} \psi](\omega) d\theta M_\psi^{-1}(\omega) \right] \quad (11)$$

Proof: Take $d = 2$, $T = \mathbb{T}$ and $\tau : \mathbb{T} \rightarrow \text{Aut}(\mathbb{R}^2)$ given by $\tau(e^{i\theta})\mathbf{x} = R_\theta \mathbf{x}$ in Theorem 4.4 formulated in Duits et al. (2004), which is a summary of results proved in (Duits and Duits, 2004) p. 27–30 (here one must set $S = \mathbb{R}^2$). \square

Consequences and Remarks:

1. This theorem easily generalizes to d dimensional images, i.e. $f \in \mathbb{L}_2(\mathbb{R}^d)$, $d = 2, 3, \dots$. The only thing that changes is that integration now takes place over $SO(d)$ and the function M_ψ becomes

$$M_\psi(\omega) = (2\pi)^{d/2} \int_{SO(d)} |\mathcal{F}(\mathcal{R}_t \psi)(\omega)|^2 d\mu_T(t),$$

where $d\mu_T(t)$ is the normalized left-invariant Haar-measure of $SO(d)$, which is the Fourier transform of

$$\tilde{\psi}(\mathbf{x}) = \int_{SO(d)} \overline{(\mathcal{R}_t \tilde{\psi} * \mathcal{R}_t \psi)(\mathbf{x})} d\mu_T(t).$$

It can be shown that if $\psi \in \mathbb{L}_1(\mathbb{R}^2)$, then M_ψ and $\tilde{\psi}$ are continuous functions in $\mathbb{L}_1(\mathbb{R})$ and thereby vanishing at infinity. As a result the ideal case $M_\psi = (2\pi)^{d/2}$ (in which case we would have $\mathbb{H}_\psi = \mathbb{L}_2(\mathbb{R}^d \times SO(d))$ and thereby (quadratic norm preservation between image and orientation score) cannot be obtained *unless* one uses a Gelfand triple structure (just like Fourier transform) constructed by means of the Laplace operator⁸ but this goes beyond the scope of this paper, for details see Duits (2004).

2. Theorem 4 easily generalizes to the discrete orientation group, i.e. $G = \mathbb{T}_N \times \mathbb{R}^2$, where

$$\mathbb{T}_N = \left\{ e^{ik\Delta\theta} \mid k \in \{0, 1, \dots, N-1\}, \right.$$

$$\left. \Delta\theta = \frac{2\pi}{N} \right\}, \text{ for } N \in \mathbb{N}, \quad (12)$$

by replacing integrations by discrete summation. Notice that the discrete orientation score $U_f^N(\mathbf{b}, e^{ik\Delta\theta})$ of an image $f \in \mathbb{L}_2(\mathbb{R}^2)$ is given by

$$U_f^N(\mathbf{b}, e^{ik\Delta\theta}) = (\mathcal{I}_\mathbf{b} \mathcal{R}_{e^{ik\Delta\theta}} \psi, f)_{\mathbb{L}_2(\mathbb{R}^2)},$$

$$k \in \{0, 1, \dots, N-1\}, \Delta\theta = \frac{2\pi}{N}.$$

and the discrete version of the function M_ψ is

$$M_\psi(\omega) = \frac{1}{N} \sum_{k=0}^{N-1} |\mathcal{F}(\mathcal{R}_{e^{ik\Delta\theta}} \psi)(\omega)|^2.$$

3. The function M_ψ completely determines the stability of the forward and backward transformation. In practice (due to sampling) we work with band-limited images. If we restrict the wavelet transformation to the space of bandlimited images $\mathbb{L}_2^{\mathcal{B}}(\mathbb{R}^2)$ we can define a condition number (with respect to quadratic norms on the space of images and the space of orientation scores), (Duits, 2005). This condition number tends to 1 when M_ψ tends to a constant function on the relevant part of the spectrum say $1_{B_{0,\varrho}}$. We will call wavelets with the property that $M_\psi|_{B_{0,\varrho}} \approx 1$ **proper wavelets** as they guarantee a stable reconstruction. For these type of wavelets we could as well use the approximative reconstruction formula⁹

$$\tilde{\psi} * f = \mathcal{F}^{-1} \left[\omega \mapsto \frac{1}{2\pi} \int_0^{2\pi} \mathcal{F}[U_f(\cdot, e^{i\theta})](\omega) \mathcal{F}[\mathcal{R}_{e^{i\theta}} \psi](\omega) d\theta \right]. \quad (13)$$

4. Construction of Proper Wavelets

Wavelets ψ with $M_\psi = 1_{B_{0,\varrho}}$ induce optimal stability of the (inverse) wavelet transform, but because of the discontinuity at $\rho = \|\omega\| = \varrho$ this choice causes numerical problems with the discrete inverse Fourier transform in practice. To avoid this practical problem we mainly focus on wavelets ψ , with either $M_\psi(\omega) = \mathcal{M}_N(\sigma^2 \rho^2)$, $N \in \mathbb{N}$, $\sigma > 0$, $\rho = \|\omega\|$ where

$$\mathcal{M}_N(\rho^2) = e^{-\rho^2} \sum_{k=0}^N \frac{\rho^{2k}}{k!} \leq 1. \quad (14)$$

In both cases the function M_ψ smoothly approximates $1_{B_{0,\varrho}}$, see Figure 3, and thereby guarantees a stable reconstruction. In the sequel we will call a wavelet $\psi \in \mathbb{L}_2(\mathbb{R}^2) \cap \mathbb{L}_1(\mathbb{R}^2)$, with such a M_ψ , a **proper wavelet**. Within the class of proper wavelets, we are looking for wavelets that are also good practical line detectors, since in the score we want to see a clear




Figure 3. Plots of $\lambda \mapsto \mathcal{M}_N(\rho^2)$, with $\rho^2 = \lambda(N+1)$ for $N = 20, 40, 60, 80, 100$.

distinction between elongated structures and non elongated structures.

To this end we notice that it is hard to give a strict mathematical answer to the question which elongated wavelet to choose given an application where it is required to enhance, detect and complete elongated structures? In practice it is mostly sufficient to consider wavelets that are similar to the local elongated patch one would like to detect and orthogonal to structures of local patches you do not want to detect, in other words employing the basic principles of template matching.

4.1. Construction of Proper Wavelets by Expansion in Eigen Functions of the Harmonic Oscillator

A nice basis to expand wavelet ψ are the eigenfunctions (represented in polar coordinates) of the harmonic oscillator operator $\|\mathbf{x}\|^2 - \Delta$ from quantum mechanics. To this end we notice that this basis is a complete basis in $\mathbb{L}_2(\mathbb{R}^d)$, $d = 2, 3, \dots$ Moreover, it is readily verified that the harmonic oscillator commutes both with the rotation operator \mathcal{R}_R given by $\mathcal{R}_R \psi(\mathbf{x}) = \psi(R^{-1}\mathbf{x})$, $\mathbf{x} \in \mathbb{R}^d$, $\psi \in \mathbb{L}_2(\mathbb{R}^d)$ and with Fourier transformation. As a result they have a common set of eigenfunctions. Consequently, these eigen functions are steerable. Moreover they are (up to a phase factor) Fourier invariant,

which enables us to control the shape of the wavelet in the Fourier domain (in order to get a proper wavelet, i.e. stable (re)construction) and simultaneously in the spatial domain (in order to get a good line detector). *In this respect*, we stress that any other choice of a complete polar-separable base is inferior to the complete base of eigen functions of the harmonic oscillator, due to the Bochner-Hecke Theorem, see Appendix A. In this section we only consider 2D images (so $d = 2$), so $\mathbb{L}_2(\mathbb{R}^2) = \mathbb{L}_2(S^1) \otimes \mathbb{L}_2((0, \infty); r dr)$ and the eigen functions¹⁰ are given by $(h_n^m \otimes Y_m)(\mathbf{x}) = h_n^m(r)Y_m(\phi)$, $\mathbf{x} = (r \cos \phi, r \sin \phi)$, where

$$\begin{aligned} h_n^m(r) &= \left(\frac{2n!}{(|m|+n)!} \right)^{1/2} r^{|m|} e^{-r^2/2} L_n^{(|m|)}(r^2), \\ r &> 0, n \in \mathbb{N} \cup \{0\}, m \in \mathbb{Z}, \\ Y_m(\phi) &= \frac{1}{\sqrt{2\pi}} e^{-im\phi}, \phi \in (0, 2\pi), \end{aligned} \quad (15)$$

where $L_n^{(|m|)}(r)$ is the n -th generalized Laguerre polynomial of type $|m|$, see Figure 4. Considering numerical expansions of local patches of elongated structures into this basis (in order to see which components are important to detect) we notice that for each $|m|$ a soft (linear) cut-off in n is required, see Figure 5. By expanding the wavelet ψ in the complete basis of eigen functions of the Harmonic oscillator we get:

$$\begin{aligned} \psi(\mathbf{x}) &= \sum_{m \in \mathbb{Z}} \sum_{n=0}^{\infty} \alpha_m^n (Y_m \otimes h_n^m)(\phi, r), \\ \mathcal{F}[\psi](\omega) &= \sum_{m \in \mathbb{Z}} \sum_{n=0}^{\infty} (i)^{|m|} (-1)^{n+m} \alpha_m^n (Y_m \otimes h_n^m)(\phi, \rho), \\ (\mathcal{R}_{e^{i\theta}} \psi)(\mathbf{x}) &= \sum_{m \in \mathbb{Z}} \sum_{n=0}^{\infty} \alpha_m^n e^{+im\theta} (Y_m \otimes h_n^m)(\phi, r), \\ U_f(\mathbf{b}, e^{i\theta}) &= (\mathcal{R}_{e^{i(\theta+\pi)}} \bar{\psi} * f)(\mathbf{b}) \\ &= \sum_{m \in \mathbb{Z}} \sum_{n=0}^{\infty} (-1)^m \overline{\alpha_m^n} e^{-im\theta} (Y_{-m} \otimes h_n^m * f)(\mathbf{b}), \\ M_\psi(\omega) &= \sum_{m=-\infty}^{\infty} \left| \sum_{n=0}^{\infty} (-1)^n \alpha_m^n h_n^m(\rho) \right|^2 \\ &= \sum_{m=-\infty}^{\infty} \sum_{n=0}^{\infty} \sum_{n'=0}^{\infty} (-1)^{n+n'} \overline{\alpha_m^n} \alpha_m^{n'} h_n^m(\rho) h_{n'}^m(\rho). \end{aligned} \quad (16)$$




Figure 4. Radial basis functions h_n^m , left for $m = 0$ and middle for $m = 1$ and lighter gray for $n = 0, 1, 2, \dots$ Right, the basis functions are effectively active on $[0, R_{mn})$, where $R_{mn} = \sqrt{2(2n + |m| + 1)}$, as this equals the radius where the total energy $E_{mn} = 2(2n + |m| + 1)$ equals the potential energy given by $V(\mathbf{x}) = r^2$. This is illustrated by joining the graphs of h_n^m , $m = 0, 1, n = 0, 1, 2$ together with their corresponding energy levels and the graph of the potential V .



Figure 5. Local expansion of a 33×33 pixel patch of an MRI-image showing a bifurcating bloodvessel in the retina. The original image is left, the reconstruction with basis function up to $|m| = 32$ and $n = 12$ is in the middle, and the same reconstruction linearly dampening higher m and n components is depicted on the right.

For details on derivations of proper wavelets with $M_\psi(\omega) = \mathcal{M}_N(\sigma^2 \rho^2)$, we refer to earlier work (Duits et al., 2004). Here we only consider a specific case which leads to a nice line detecting proper wavelet, which corresponds to the wavelet proposed by Kalitzin et al. (1999).

Example: The special case $\alpha_m^n = \alpha_m \delta_{n0}$ In this case $M_\psi = \mathcal{M}_N$, recall (14) and (16) simplifies to

$$M_\psi(\omega) = \sum_{m=0}^N |\alpha_m|^2 (h_n^m(\rho))^2 = \mathcal{M}_N(\rho^2), \quad \rho = \|\omega\|, \quad (17)$$

The (up to phase factors unique) solution ψ_N^0 of (17) is now given by ($\alpha_m = 1$ for all m)

$$\begin{aligned} \psi_N^0(\mathbf{x}) &= \sum_{m=0}^N \frac{1}{\sqrt{m!}} r^m e^{-\frac{\rho^2}{2}} \frac{e^{-im\phi}}{\sqrt{2\pi}} = \frac{1}{\sqrt{2\pi}} \sum_{m=0}^N \frac{(\bar{z})^m}{\sqrt{m!}} e^{-\frac{|z|^2}{2}} \\ &= \frac{1}{\sqrt{2\pi}} \sum_{m=0}^N \left(-\frac{1}{2}\right)^m \frac{\left(\frac{a}{\partial z}\right)^m}{\sqrt{m!}} e^{-\frac{|z|^2}{2}} \quad z = r e^{i\phi}. \end{aligned} \quad (18)$$

This series converges uniformly on compacta, but not in \mathbb{L}_2 -sense. The real part of this wavelet corresponds to the wavelet first proposed by Kalitzin et al. (1999) cf. as a line detector in medical images. The imaginary part is a good edge detector.

Practical Aspects: The cutoff index N has a practical upper bound because of sampling. If N increases the reconstruction will become better, but if we choose N too large the wavelet behaves badly along $\phi = 0$, see Figures 6 and 7.

We notice that $\tilde{\psi}_N^0 = \mathcal{F}^{-1}[\omega \mapsto \mathcal{M}_N(\rho^2)]$, which equals the integration over the auto-correlations of all rotated kernels, is an approximation of the identity¹¹, whereas the wavelets ψ_N^0 are not.

The size of the wavelet ψ_0^N can be controlled by dilation, $\mathbf{x} \mapsto (\mathcal{D}_\sigma \psi_0^N)(\mathbf{x}) = \frac{1}{\sigma} \psi_0^N(\mathbf{x}/\sigma)$. This does

effect M_ψ , since $\mathcal{F}\mathcal{D}_\sigma = \mathcal{D}_{1/\sigma}\mathcal{F}$, but, for N sufficiently large, the stability remains guaranteed. Moreover, for large N , the wavelet can be smoothed by convolving the wavelet with a relatively small Gaussian kernel:

$$\begin{aligned} \psi_{N,s}^0(\mathbf{x}) &= (G_s * \psi_N^0)(\mathbf{x}) \\ &= e^{-(1/2)\left|\frac{z}{\varrho}\right|^2} \beta(\varrho, s) (\mathcal{D}_{\alpha(\varrho, s)} \psi_N^0)(z), \\ z &= x + iy \end{aligned} \quad (19)$$

where we recall that $\varrho > 0$ equals the maximum frequency radius and where $\alpha(\varrho, s) = 1 + \frac{s}{\varrho^2} \approx 1$ and $\beta(\varrho, s) = \frac{1}{\frac{\varrho^2}{s} + 2 + \frac{s}{\varrho^2}} = O\left(\frac{s}{\varrho^2}\right)$ and $\frac{s}{\varrho^2} \ll 1$. It is easily verified that $\psi \mapsto G_s * \psi$ implies $\tilde{\psi} \mapsto G_{2s} * \tilde{\psi}$, so as long as the scale s is relatively small, the Gaussian convolution of the wavelet is harmless for the stability of the (re)construction.

4.2. Simple Approach to Construction of Proper Wavelets

In Duits (2005) we also developed a more simple approach to obtain proper wavelets, which we will again illustrate with an example. In this more heuristic approach we do not insist on having full control over the analytic expansions in spatial and Fourier domain by means of the Harmonic oscillator basis. Here we take the condition $M_\psi \approx 1$ as a starting point and consider line detector wavelets only in the Fourier domain. Although we do not have full control over the shape of the wavelet in the spatial domain, we will rely on the basic but clear intuition that an elongated wavelet in the Fourier domain (say along axis \mathbf{e}) corresponds to a double sided elongated wavelet in the spatial domain (along axis $R_{\pi/2}\mathbf{e}$). The following example shows that it is possible to obtain




Figure 6. Top row: Left: two different views of graphs $\psi_{N=15}^0$, Right: two different views of graphs $\psi_{N=30}^0$. Bottom row: Left: 3 Plots of graph $\Re(\psi_{N=100}^0)$. Right: 100×100 -pixel grey-value plots of $\Re(\psi_{N=100}^0)$ and Gaussian blurred (with $\sigma = 0.8$ pixels). Notice that the kernel becomes sharper and the oscillatory-ring vanishes as N increases. The locally highly oscillatory behavior within $\Re(\psi_{N=100}^0)$ may seem awkward, but is not really harmful since it disappears immediately by convolution with Gaussian kernel with tiny scale, also see (19).

proper wavelets, which are nice line detectors and which allow a simple and fast reconstruction scheme.¹²

Example : Consider $N = 18$ discrete orientations, so $\Delta = \frac{2\pi}{N}$. The idea is to “fill a cake by pieces of cake” in the Fourier domain. In order to avoid high frequencies in the spatial domain, these pieces must be smooth and thereby they must overlap.

Let $\psi = \mathcal{F}^{-1}[\omega \mapsto \sqrt{A(\varphi)}G_s(\rho)]$ and let $A : S^1 \rightarrow \mathbb{R}^+$ be given by

$$A(\varphi) = \begin{cases} -\frac{2}{\Delta^2} \varphi^2 + 1 & \text{if } |\varphi| \leq \frac{\Delta}{2} \\ \frac{2}{\Delta^2} \varphi^2 - \frac{4}{\Delta} |\varphi| + 2 & \text{if } \frac{\Delta}{2} \leq |\varphi| \leq \Delta, \\ 0 & \text{else} \end{cases} \quad \varphi \in [-\pi, \pi), \quad (20)$$

then it is easily checked that $M_\psi(\omega) = G_s(\rho)$, $\rho = \|\omega\|$. In our experiments we took $s = \frac{1}{2}\sigma^2$ large, which




Figure 7. Left: The graphs of the kernel $\psi_N^0(\mathbf{x})$ cut off at $N = 10, 20, 30, 40, 50$ with $\sigma = 1/8$, restricted to its main direction $\phi = 0$. Notice that the peaks move out as m increases. Notice that the asymptotic formula derived for $\psi_\infty^0(r, \phi = 0) = (8\pi)^{1/4} \sqrt{r} (1 - \frac{1}{16r^2} + O(\frac{1}{r^4})) = (8\pi)^{1/4} \sqrt{r} + O(r^{-\frac{3}{2}})$ see Duits (2005) section 7.3, is a very good approximation (we included the graph of $r \mapsto (8\pi)^{1/4} \sqrt{r}$). Right: The corresponding functions $M_{\psi_N^0} = M_N$ which indeed approximate 1 as $N \rightarrow \infty$.

allowed us to use the approximate reconstruction (13), but this gave similar results than a fast reconstruction by integration over the angles only:

$$\tilde{f}(\mathbf{x}) = \sum_{k=0}^{N-1} U_f(\mathbf{x}, e^{i(k\Delta)}), \quad (21)$$

see Figure 8. In Duits (2005) p. 134 we applied elementary processing on the orientation score of the retinal image in Figure 8 and compared it with a standard line detection technique.

5. Image Enhancement via Left Invariant Operations on Orientation Scores

Now that we have constructed an orientation score U_f from image f , such that it allows a well-posed reconstruction of f from U_f , one can think of suitable operations on the orientation scores.

Let ψ be a proper wavelet, then there exists a 1-to-1 correspondence between bounded operators $\Phi \in \mathcal{B}(\mathbb{C}_G^K)$ on orientation scores and bounded operators $\Upsilon \in \mathcal{B}(\mathbb{L}_2^q(\mathbb{R}^d))$ on band limited images:

$$\Upsilon[f] = (\mathcal{W}_\psi^q)^*[\Phi[\mathcal{W}_\psi^q[f]]], \quad f \in \mathbb{L}_2^q(\mathbb{R}^d). \quad (22)$$

This allows us to relate operations on orientation scores to operations on images in a robust manner. For proper wavelets, we have that the space of orientation scores \mathbb{C}_G^K can be considered as a linear subspace of $\mathbb{L}_2(G)$.




Figure 8. Upper row: Plots of the graph of the real and imaginary part of the proper wavelet ψ given by (20), determined by the discrete inverse Fourier transform of $\omega \mapsto \sqrt{A(\varphi)G_s(\|\omega\|)}$ (we set $s = 800$) sampled on a 256×256 equidistant grid. From left to right: Density plots of $\Re(\psi)$ at true size, 3D views of the graphs of $\Re(\psi)$ and $\Im(\psi)$. Bottom row: MRI-image of the retina, three slices $U_f^N(e^{ik\Delta}, \cdot)$, $k = 0, 2, 4$ of the discrete orientation score and fast approximative reconstruction, which is close to exact reconstruction.

Let $\Phi : \mathbb{L}_2(G) \rightarrow \mathbb{L}_2(G)$ be some bounded operator on $\mathbb{L}_2(G)$, then the range of the restriction of this operator to the subspace \mathbb{C}_K^G of orientation scores need not be contained in \mathbb{C}_K^G , i.e. $\Phi(U_f)$ need not be the orientation score of an image. The adjoint mapping of $\mathcal{W}_\psi^e : \mathbb{L}_2^e(\mathbb{R}^d) \rightarrow \mathbb{L}_2(G)$, given by $\mathcal{W}_\psi^e[f] = \mathcal{W}_\psi[f]$, $f \in \mathbb{L}_2^e(\mathbb{R}^d)$ is given by¹³

$$(\mathcal{W}_\psi^e)^*(V) = \int_G \mathcal{U}_g \psi V(g) d\mu_G(g), \quad V \in \mathbb{L}_2(G).$$

The operator $\mathbb{P}_\psi = \mathcal{W}_\psi^e (\mathcal{W}_\psi^e)^*$ is the orthogonal projection on the space of orientation scores \mathbb{C}_K^G . This projection can be used to decompose the manipulated orientation score:

$$\Phi(U_f) = \mathbb{P}_\psi(\Phi(U_f)) + (I - \mathbb{P}_\psi)(\Phi(U_f)).$$

Definition 5. An operator $\Phi : \mathbb{L}_2(G) \rightarrow \mathbb{L}_2(G)$ is left invariant iff

$$\Phi \circ \mathcal{L}_g = \mathcal{L}_g \circ \Phi, \quad \text{for all } g \in G, \quad (23)$$

where the left regular action \mathcal{L}_g (also known as shift-twist transformation, cf. Williams and Zweck (2003)) of $g \in G$ onto $\mathbb{L}_2(G)$ is given by

$$\mathcal{L}_g \psi(h) = \psi(g^{-1}h) = \psi(R_\theta^{-1}(\mathbf{b}' - \mathbf{b}), e^{i(\theta' - \theta)}), \quad (24)$$

with $g = (\mathbf{b}, e^{i\theta}) \in G$, $h = (\mathbf{b}', e^{i\theta'}) \in G$.

Theorem 6. Let Φ be a bounded operator on \mathbb{C}_K^G . Then the unique corresponding operator Υ on $\mathbb{L}_2^e(\mathbb{R}^d)$, which is given by $\Upsilon[f] = (\mathcal{W}_\psi^e)^* \circ \Phi \circ \mathcal{W}_\psi^e[f]$ is Euclidean invariant, i.e. $\mathcal{U}_g \Upsilon = \Upsilon \mathcal{U}_g$ for all $g \in G$ if and only if Φ is left invariant, i.e. $\mathcal{L}_g \Phi = \Phi \mathcal{L}_g$ for all $g \in G$.

Here we omit the proof. It follows by $\mathcal{W}_\psi^e \mathcal{U}_g = \mathcal{L}_g \mathcal{W}_\psi^e$, for all $g \in G$, cf. Duits et al. (2004) p. 38.

Practical Consequence: Euclidean invariance of Υ is of great practical importance, since the result should not be essentially different if the original image is rotated or translated. So by Theorem 6 the only reasonable operations on orientation scores are left invariant. It is not a problem when the mapping $\Phi : \mathbb{C}_K^G \rightarrow \mathbb{L}_2(G)$ maps an orientation score to an element in $\mathbb{L}_2(G) \setminus \mathbb{C}_K^G$, but be aware that $\mathbb{P}_\psi \Phi : \mathbb{C}_K^G \rightarrow \mathbb{C}_K^G$ yields the same result.¹⁴

All linear left invariant kernel operators $\Phi : \mathbb{L}_2(G) \rightarrow \mathbb{L}_2(G)$ are G -convolution operators. They are given by

$$\begin{aligned} [\Phi(U)](g) &= \int_G K(h^{-1}g)U(h) dh, \quad g = (\mathbf{b}, e^{i\theta}) \\ &= \int_{\mathbb{R}^2} \int_0^{2\pi} K(R_{\theta'}^{-1}(\mathbf{b} - \mathbf{b}'), e^{i(\theta - \theta')}) U(\mathbf{b}', e^{i\theta'}) d\theta' db'_1 db'_2, \end{aligned} \quad (25)$$

for almost every $g = (\mathbf{b}, e^{i\theta}) \in G$. From the practical point of view(speed) these can be implemented via

impuls response and then taking the G -convolution. Before we propose left invariant operators on orientation scores, we give a brief overview of the interesting geometry within the domain G of orientation scores, which is the Euclidean Motion Group.

For any Lie-group G the tangent space $T_e(G)$ at the unity element equipped with the product

$$[A, B] = \lim_{t \downarrow 0} \frac{a(t)b(t)(a(t))^{-1}(b(t))^{-1} - e}{t^2},$$

where $t \mapsto a(t)$ resp. $t \mapsto b(t)$ are any smooth curves in G with $a(0) = b(0) = e$ and $a'(0) = A$ and $b'(0) = B$, is isomorphic to $\mathcal{L}(G)$. $\mathcal{L}(G)$ is the Lie-algebra of left invariant vector fields on G , i.e. all vector fields \tilde{A} on G such that

$$\tilde{A}_g f = \tilde{A}_e(f \circ L_g) = \tilde{A}_e(h \mapsto f(g h)),$$

equipped with product $[\tilde{A}, \tilde{B}] = \tilde{A}\tilde{B} - \tilde{B}\tilde{A}$. The isomorphism is given by $A \leftrightarrow \tilde{A} \leftrightarrow \tilde{A}_g(\phi) = A(\phi \circ L_g) = A(h \mapsto \phi(g h))$ for all smooth $\phi : G \supset O_g \rightarrow \mathbb{R}$ and all g, h in G . In our case of the Euclidean motion group we have that $T_e(G)$ is spanned by $\{A_1 = \mathbf{e}_\theta, A_2 = \mathbf{e}_\xi, A_3 = \mathbf{e}_\eta\}$ with

$$\xi = b_1 \cos \theta + b_2 \sin \theta, \mathbf{e}_\xi = \cos \theta \mathbf{e}_{b_1} + \sin \theta \mathbf{e}_{b_2} \quad (26)$$

(in the spatial plane along the measured orientation) and

$$\eta = -b_1 \sin \theta + b_2 \cos \theta, \mathbf{e}_\eta = -\sin \theta \mathbf{e}_{b_1} + \cos \theta \mathbf{e}_{b_2} \quad (27)$$

(in the spatial plane orthogonal to the measured orientation). The corresponding left (or shift-twist) invariant vector fields are given by

$$\{\tilde{A}_1 = \partial_\theta, \tilde{A}_2 = \partial_\xi = \cos \theta \partial_{b_1} + \sin \theta \partial_{b_2}, \tilde{A}_3 = \partial_\eta = -\sin \theta \partial_{b_1} + \cos \theta \partial_{b_2}\}. \quad (28)$$

It is easily verified that

$$\begin{cases} \tilde{A}_3 = [\tilde{A}_1, \tilde{A}_2] \\ -\tilde{A}_2 = [\tilde{A}_1, \tilde{A}_3] \\ 0 = [\tilde{A}_2, \tilde{A}_3] \end{cases} \text{ and } \begin{cases} A_3 = [A_1, A_2] \\ -A_2 = [A_1, A_3] \\ 0 = [A_2, A_3] \end{cases},$$

which coincides with $A \longleftrightarrow \tilde{A}$. For dimensional consistency define $X_1 = (\frac{2\pi}{Z})A_1$, $X_2 = A_2$ and $X_3 = A_3$, where Z is the width of the image domain (so $b_1, b_2 \in [0, Z]$, assuming a square image domain). A group element $g = (\mathbf{b}, \theta)$ can then be parameterized using either so-called coordinates of the first kind

$\{\alpha_i\}_{i=1,2,3}$, see Duits (2005) section 7.6 p. 228–229, or by the coordinates of the second kind $\{\beta_i\}_{i=1,2,3}$:

$$\begin{aligned} g &= (\mathbf{b}, e^{i\theta}) = \exp\left(\sum_{i=1}^3 \alpha_i X_i\right) \\ &= \left(\frac{Z\alpha_3}{2\pi\alpha_1} \left(\cos\left(\frac{2\pi\alpha_1}{Z}\right) - 1\right) + \frac{Z\alpha_2}{2\pi\alpha_1} \sin\left(\frac{2\pi\alpha_1}{Z}\right), \right. \\ &\quad \left. \frac{Z\alpha_3}{2\pi\alpha_1} \sin\left(\frac{2\pi\alpha_1}{Z}\right) - \frac{Z\alpha_2}{2\pi\alpha_1} \left(\cos\left(\frac{2\pi\alpha_1}{Z}\right) - 1\right), e^{i\frac{2\pi\alpha_1}{Z}}\right), \quad (29) \\ g &= (\mathbf{b}, e^{i\theta}) = \prod_{i=1}^3 \exp(\beta_i X_i) \\ &= \left(\beta_2 \cos\left(\frac{2\pi\beta_1}{Z}\right) - \beta_3 \sin\left(\frac{2\pi\beta_1}{Z}\right), \right. \\ &\quad \left. \beta_2 \sin\left(\frac{2\pi\beta_1}{Z}\right) + \beta_3 \cos\left(\frac{2\pi\beta_1}{Z}\right), e^{i\frac{2\pi\beta_1}{Z}}\right), \end{aligned}$$

The coordinates of the second kind correspond to $(e^{i\eta}, \theta, \xi)$, since by (29): $(\frac{2\pi\beta^1}{Z}, \beta^2, \beta^3) = (\theta, R_{-\theta}\mathbf{b}) = (\theta, \xi, \eta)$.

5.1. Basic Left Invariant Operations on Orientation Scores

In image analysis it is well known that differential operators used for corner/line/edge/blob detection must be Euclidean invariant. Mostly, such differential invariants are easily expressed in a local coordinate system (gauge coordinates) where in 2D one coordinate axis (say v) is along the isophote and the other runs along the gradient direction (say w), cf. (Florack, 1997).

Rather than putting these gauge coordinates along isophotes we propose a local coordinate system along the measured orientation. Note to this end that in some medical image applications the elongated structures are not along isophotes. So in our orientation scores ξ and η play the role of v and w . Moreover we can differentiate along the direction θ and obtain directional frequencies.

Besides these local left invariant operators we can think of more global left-invariant operators, like normalization

$$[\Phi(U_f)](\mathbf{b}, e^{i\theta}) = U_f(\mathbf{b}, e^{i\theta}) / \left(\int_{\mathbb{R}^2} |U_f(\mathbf{x}, e^{i\theta})|^p d\mathbf{x} \right)^{\frac{1}{p}}, p > 1, \quad (30)$$

cf. Figure 9, or grey-value transformations

$$\Phi(U_f) = (U_f - \min_g \{U_f(g)\})^q, q > 0. \quad (31)$$

to enhance elongated structures.




Figure 9. Illustration of image processing via elementary operation on orientation score. Visual illusion. Normalization (see (30)) of the orientation layers in the orientation scores reveals the most contrasting lines in the triangles.

5.2. Evolution Equations Corresponding to Left Invariant Stochastic Processes on the Euclidean Motion Group

Just like the well-known Gaussian scale space satisfies the translation and rotation invariance axiom, (Duits et al., 2004), the following linear evolutions on orientation scores are left invariant:

$$\begin{cases} \partial_s W = A W \\ \lim_{s \downarrow 0} W(\cdot, s) = U_f(\cdot), \end{cases} \quad (32)$$

where the generator A acting on $\mathbb{L}_2(G)$ is given by (the closure of)

$$A = \left[\frac{a_1}{Z} \partial_\theta + a_2 \partial_\xi + a_3 \partial_\eta + \frac{D_{11}}{Z^2} (\partial_\theta)^2 + D_{22} (\partial_\xi)^2 + D_{33} (\partial_\eta)^2 \right], \quad (33)$$

$$a_i, D_{ii} \in \mathbb{C}, i = 1, \dots, 3.$$

The first order derivatives take care of transport (convection) and the second order derivatives give diffusion. We first consider the case where all D_{ii} 's are zero and the initial condition is a spike-bundle $\delta_{\theta_0, \mathbf{b}_0}$ (i.e. one ‘‘oriented particle’’ so to speak). This spike will move over time along exponential curves, which are straight-

lines in a spatial plane, spirals through G and straight lines along θ -direction. By introducing the variables, $t = s \frac{a_1}{Z}$, $\lambda_2 = \frac{a_2}{a_1} Z$, $\lambda_3 = \frac{a_3}{a_1} Z$ Equations (32–33) reduces to

$$\begin{cases} \partial_t W = [\partial_\theta + \lambda_2 \partial_\xi + \lambda_3 \partial_\eta] W & \lambda_2, \lambda_3 \in \mathbb{C} \\ \lim_{t \downarrow 0} W(\cdot, t) = \delta_{\theta_0, \mathbf{b}_0}. \end{cases}$$

Notice that indeed $[s] = 1 \leftrightarrow [t] = 1$ and $[a_1] = [a_2] = [a_3] = [\text{length}] \leftrightarrow [\lambda_2] = [\lambda_3] = [\text{length}]$. It follows by equality (29) that the orbit of the Dirac distribution at initial position $(\mathbf{b}_0, e^{i\theta_0})$ is given by

$$\begin{aligned} & (b_0^1 + \lambda_3(\cos(t + \theta_0) - \cos(\theta_0)) + \lambda_2(\sin(t + \theta_0) - \sin(\theta_0)), \\ & b_0^2 + \lambda_3(\sin(t + \theta_0) - \sin(\theta_0)) - \lambda_2(\cos(t + \theta_0) \\ & - \cos(\theta_0)), e^{i(t+\theta_0)}), \end{aligned}$$

which is a circular spiral with radius $\sqrt{\lambda_2^2 + \lambda_3^2}$ around central point

$$(-\lambda_3 \cos \theta_0 - \lambda_2 \sin \theta_0 + b_0^1, \lambda_2 \cos \theta_0 - \lambda_3 \sin \theta_0 + b_0^2),$$

which *exactly* corresponds to the results from our numerical implementation. The solution of the pure diffusion problem i.e. $a_1 = a_2 = a_3 = 0$ in (33) is a G -convolution kernel operator with some positive kernel $K_s \in \mathbb{L}_1(G)$, which can be sharply estimated from above and below by Gaussian kernels on G , viz.

$$c'(V(t))^{-1/2} e^{-b' \frac{|g|^2}{t}} \leq K_s(g) \leq c(V(t))^{-1/2} e^{-b \frac{|g|^2}{t}},$$

with $c, c', b, b' > 0$. For details see Dungey et al. (2003). In the degenerate case $a_1 = a_2 = a_3 = D_{11} = 0$, the diffusion boils down to an ordinary spatial convolution for each fixed θ with an anisotropic Gaussian kernel where $\frac{D_{22}}{D_{33}}$ gives the anisotropy factor of Gaussian convolution along \mathbf{e}_ξ and \mathbf{e}_η .

The evolution equations given by (32) correspond to stochastic processes. For example the case $a_1 = a_3 =$




Figure 10. Illustration of image processing via elementary operation on orientation score. From left to right: 1: Noisy medical image with guide wire, 2–3: Small oriented wavelet ψ with corresponding processed image $\tilde{f} = \mathcal{W}_\psi^* [\tilde{U}_f]$, with $\tilde{U}_f = \left(\frac{\Re(U_f) - \min(\Re(U_f))}{\max(\Re(U_f)) - \min(\Re(U_f))} \right)^2$. 4–5: Same as 2–3 with relatively larger kernel $\frac{1}{\sigma} \psi_N^0(\frac{x}{\sigma})$, where we recall that ψ_N^0 was given by (18). For the sake of clarity the wavelet plots (2,4) are zoomed in with a factor of 2.




Figure 11. Example of perceptual organization. From left to right: 1. Original image; 2. detection of elongated structures via orientation scores U_f : $\tilde{f} = \mathcal{W}_\psi^*[\tilde{U}_f]$ with $\tilde{U}_f = \left(\frac{\Re U_f - \min \Re(U_f)}{\max \{\Re(U_f) - \min \Re(U_f)\}}\right)^2$; 3. Inverse transformation of evolved orientation score $\mathcal{W}_\psi^*[\Phi[\tilde{U}_f]]$, where Φ denotes the shooting process by maintaining curvature and direction; 4. Inverse transformation of probability density of collision of forward and backward process on orientation score, see (34). In contrast to related work (Williams and Zweck, 2003) we do not put sources and sinks by hand, but use our orientation scores instead. The only parameters involved are, range of wavelet t , decay time α and the stochastic process parameters D_{11} , D_{22} , D_{33} in (33).

0, $D_{11} = \frac{1}{2}\sigma^2$ and $D_{22} = D_{33} = 0$ is the forward Kolmogorov equation corresponding to the stochastic process known as the so-called direction process, cf. Mumford's work (Mumford, 1994):

$$\begin{cases} \bar{\kappa} = \frac{1}{L} \int_0^L k(s) ds = \frac{1}{L} \int_0^L |\dot{\theta}(s)| ds \sim \mathcal{N}(0, \sigma^2) \\ \mathbf{x}(s) = \int_0^s \begin{pmatrix} \cos \theta(\tau) \\ \sin \theta(\tau) \end{pmatrix} d\tau + \mathbf{x}(0), \\ L \sim NE(\alpha). \end{cases}$$

which is the limit of the following *discrete* stochastic process

$$\begin{cases} \theta(s_i + \Delta s) = \theta(s_i) + \Delta s \varepsilon, & \varepsilon \sim \mathcal{N}(0, N\sigma^2) \\ \mathbf{x}(s_i + \Delta s) = \mathbf{x}(s_i) + \Delta s \begin{pmatrix} \cos \theta(s_i) \\ \sin \theta(s_i) \end{pmatrix} \\ \Delta s = \frac{L}{N}, \text{ with } L \sim NE(\alpha). \end{cases}$$

Just like scale space theory¹⁵, a scale space representation $u(\mathbf{x}, s) = (G_s * f)(\mathbf{x})$ can be regarded as an isotropic stochastic process, where the distribution of positions of particles evolve over time (the well-known Wiener process), evolutions on orientation scores can be considered as stochastic processes where the distribution of positions of *oriented* particles evolve over time.

The life time T of a particle travelling with unit speed (so $T = L$) through G is assumed to be negative exponentially distributed, $T \sim NE(\alpha)$, i.e. $p(T = t) = \alpha e^{-\alpha t}$, with expected life time $E(T) = \frac{1}{\alpha}$, because it is memoryless, which must be the case in a Markov-process. The probability density of finding a

$e^{i\theta}$ -oriented particle at position \mathbf{b} is given by

$$\begin{aligned} p(g) = p(\mathbf{b}, \theta) &= \int_0^\infty p(\mathbf{b}, \theta | T = t) p(T = t) dt \\ &= \alpha \int_0^\infty [e^{tA} U_f](\mathbf{b}, \theta) e^{-t\alpha} dt \\ &= -\alpha [(A - \alpha I)^{-1} U_f](\mathbf{b}, \theta). \end{aligned}$$

Consider two independent stochastic processes generated by $A = \text{Conv} + \text{Diff}$, where Conv, resp. Diff stands for the convection resp. diffusion part of A , given by (33), and its adjoint $A^* = -\text{Conv} + \text{Diff}$. So the direction of shooting particles is opposite and the stochastic behavior is similar in the two processes. The probability-density of collision of particles from these 2 processes yields the following left invariant operation, see Figure 11,

$$(\Phi[U_f])(g) = [(A - \alpha I)^{-1} U_f](g) [(A^* - \alpha I)^{-1} U_f](g). \quad (34)$$

For detailed numerical algorithms and analytic approximations¹⁶ (By deriving exact Green's functions of a nil-potent group of Heisenberg type) of the Green's functions of the involved left invariant evolution equations we refer to Duits (2005) section 4.9, p. 163–177 and (van Almsick et al., 2005).

6. Invertible Orientation Scores of 3D Images

We generalized all of our results for 2D-image processing via left invariant operations on invertible orientation scores of 2D images to the more complicated case



Figure 12. Density plots of through XOZ -plane and YOZ -plane of the 3D-equivalent (for details see Duits et al., 2004) of the 2D-wavelet obtained in the simple approach framework, see (20), a joint contour plot of iso-surfaces of this 3D-equivalent ψ at $\psi(\mathbf{x}) = +0.02$ and $\psi(\mathbf{x}) = -0.02$, show that it is rather a surface patch detector than a line detector.

of 3D-image processing via left invariant operations on invertible orientation scores of 3D images. In this section we will not deal with all technical details, but restrict ourselves to the definition of a 3D orientation score and some preliminary results of line detection in 3D by means of left invariant operations on invertible orientation scores.

Although some generalizations are straightforward, some difficulties arise that did not arise in the 2D-case. First of all $SO(3)$ is not commutative, so the $SO(3)$ -irreducible representations are not one-dimensional. Secondly, in practice one is mainly interested in constructing orientation scores by “sigar-shaped” wavelets, i.e. wavelets that are invariant under the stabilizer of the north pole, which brings us to the 2-sphere $S^2 = SO(3)/SO(2)$. Thirdly, it is not obvious which discrete subgroup of $SO(3)$ to take and thereby the question arises of how to store the orientation score, since an equidistant sampling in spherical coordinates does not make sense.

Let $f \in \mathbb{L}_2^G(\mathbb{R}^3)$ be a bandlimited 3D image, then we define its wavelet transform $\mathcal{W}_\psi[f] \in \mathbb{C}_K^G$ by

$$\mathcal{W}_\psi[f](g) = \int_{\mathbb{R}^3} \overline{\psi(R^{-1}(\mathbf{x} - \mathbf{b}))} f(\mathbf{x}) d\mathbf{x},$$

$$g = (\mathbf{b}, R) \in G = \mathbb{R}^3 \rtimes SO(3).$$

We restrict ourselves to the case where the wavelet ψ is invariant under the stabilizer of the north-pole \mathbf{e}_z , which is the subgroup of $SO(3)$ consisting of all rotations around the z -axis. So we assume

$$\psi(R\mathbf{x}) = \psi(\mathbf{x}), \text{ for all } R \in \text{Stab}(\mathbf{e}_z). \quad (35)$$

On $SO(3)$ we define the following equivalence relation:

$$R_1 \sim R_2 \Leftrightarrow (R_2)^{-1}R_1 \in \text{Stab}(\mathbf{e}_z) \equiv SO(2).$$

The equivalence classes are the left cosets $[R] = R \text{Stab}(\mathbf{e}_z)$, $R \in SO(3)$. The partition of all equivalence classes will be denoted by $SO(3)/SO(2)$, which is isomorphic to S^2 and thereby not a group. Rather than using the canonical parameterization given by

$$R_{\mathbf{a},\phi}(\mathbf{x}) = (\cos \phi)\mathbf{x} + (1 - \cos \phi)(\mathbf{a}, \mathbf{x})\mathbf{a} \\ + \sin \phi(\mathbf{a} \times \mathbf{x}),$$

$$\mathbf{x}, \mathbf{a} \in \mathbb{R}^3, \phi \in [0, 2\pi) \quad (36)$$

of $SO(3)$ we will use the well-known Euler angle parameterization $Y : B_{0,2\pi} \rightarrow SO(3)$:

$$Y(\mathbf{x}) = R_{\mathbf{e}_z,\gamma} R_{\mathbf{e}_y,\beta} R_{\mathbf{e}_z,\alpha},$$

$$\mathbf{x} = (\alpha \cos \gamma \sin \beta, \alpha \sin \gamma \sin \beta, \alpha \cos \beta)^T, \quad (37)$$

which gives us directly an explicit isomorphism between S^2 and $\frac{SO(3)}{SO(2)}$:

$$S^2 \ni \mathbf{n}(\beta, \gamma) = (\cos \gamma \sin \beta, \sin \gamma \sin \beta, \cos \beta)^T \\ \Leftrightarrow [R_{\mathbf{e}_z,\gamma} R_{\mathbf{e}_y,\beta}] \in \frac{SO(3)}{SO(2)}.$$

Because of our assumption (35) we can define the orientation score $U_f : \mathbb{R}^3 \times S^2 \rightarrow \mathbb{C}$ corresponding to image $f \in \mathbb{L}_2(\mathbb{R}^3)$ by means of

$$U_f(\mathbf{b}, \mathbf{n}(\beta, \gamma)) = \mathcal{W}_\psi[f](\mathbf{b}, R_{\mathbf{e}_z,\gamma} R_{\mathbf{e}_y,\beta}).$$

We can again expand the wavelet in eigen functions of the harmonic oscillator (which are invariant under rotations around the z -axis, i.e. $m = 0$), see Appendix A, and thereby we obtain (for details see Duits (2005))

par. 4.7.2, p. 146–151):

$$\begin{aligned}\psi(\mathbf{x}) &= \sum_{n=0}^{\infty} \sum_{l=0}^{\infty} \alpha_l^n g_n^l(r) Y_l^0(\theta, \phi), \\ &\quad \text{where } \alpha_l^n = \alpha_{0l}^n, \\ \psi((R_{\mathbf{e}_z, \gamma} R_{\mathbf{e}_y, \beta})^{-1} \mathbf{x}) &= \sum_{n=0}^{\infty} \sum_{l=0}^{\infty} \sum_{m'=-l}^l \alpha_l^n [\mathcal{D}^{(l)}(R_{\mathbf{e}_z, \gamma} R_{\mathbf{e}_y, \beta})]_{m'}^0 \\ &\quad g_n^l(r) Y_l^{m'}(\theta, \phi), \\ \mathcal{F}[\psi](\boldsymbol{\omega}) &= \sum_{n=0}^{\infty} \sum_{l=0}^{\infty} \alpha_l^n (-1)^{n+l} (i)^l g_n^l(\rho) Y_l^0(\vartheta, \varphi), \\ \mathcal{M}_\psi(\boldsymbol{\omega}) &= \sum_{l=0}^{\infty} \sum_{n=0}^{\infty} \sum_{\bar{n}=0}^{\infty} (-1)^{n+\bar{n}} \alpha_l^n \alpha_{\bar{l}}^{\bar{n}} g_n^l(\rho) g_{\bar{n}}^{\bar{l}}(\rho), \\ &\quad \text{for all } \boldsymbol{\omega} \in \mathbb{R}^3,\end{aligned}$$

where

$$[\mathcal{D}^{(l)}(R_{\mathbf{e}_z, \gamma} R_{\mathbf{e}_y, \beta})]_{m'}^0 = \frac{(-1)^{m'} Y_l^{m'}(\beta, \gamma)}{\sqrt{\frac{(2l+1)}{4\pi}}}.$$

Completely analogous to the 2D case the case $\alpha_l^n = \alpha_l \delta^{n0}$ establishes the 3D-equivalent of wavelet (18), which is again a proper wavelet with $\mathcal{M}_\psi(\boldsymbol{\omega}) = \mathcal{M}_N(\rho^2)$, $N \in \mathbb{N}$, $\rho = \|\boldsymbol{\omega}\|$, see Figure 13.

7. Channel Representations

Orientation scores are a nice tool for *well-posed image enhancement* and stochastic completion of “gaps” by means of left-invariant processing on the orientation scores. However, for some practical applications where detection of oriented structures is required it is more relevant to obtain a fast robust estimation for a *single* orientation per position, rather than image enhancement via orientation scores where the columns $U_f(\mathbf{x}, \cdot)$ represent a *distribution of the local orientation* per position \mathbf{x} .



Figure 13. The 3D-equivalent of the 2D proper wavelet given by (18) is also a proper wavelet and is given by $\psi_N^{0,3D}(\mathbf{x}) = \sum_{l=0}^N \frac{1}{\sqrt{l!}} r^l e^{-\frac{r^2}{2}} Y_l^0(\theta, \phi)$,

where Y_l^0 is the well-known surface harmonic $Y_l^0(\theta, \phi) = P_l^0(\cos \theta) \sqrt{\frac{2l+1}{4\pi}}$. From left to right, plots of $\psi_N^{0,3D}(x, 0, z)$, for $N = 10, 20, 40$ and $G_s * \psi_{N=40}^{0,3D}(x, 0, z)$, for tiny scale s and finally a joined 3D plot of the iso-intensity contours of the rotated 3D kernel $\psi_i((R_{\mathbf{e}_z, \phi} R_{\mathbf{e}_x, \theta})^{-1} \mathbf{x}) = 0.5, -0.5$, ($\phi \approx -\frac{\pi}{4}$ and $\theta \approx \frac{3\pi}{4}$).

A simple way to get a first orientation estimate $\hat{\theta}(\mathbf{x})$ would be by storing the angle where the response is maximum within each orientation column of the score. To each orientation estimate $\hat{\theta}(\mathbf{x})$ we attach a measure of reliability of the estimate $r(\mathbf{x})$. This can be obtained by computing the value of the maximum relative to other response values in the orientation column $U_f(\mathbf{x}, \cdot)$. This would give a rough estimate of the local orientation (with its corresponding confidence). For example, if the score is constructed by a second order derivative of the 2D-Gaussian kernel with respect to y (so oriented along x) this boils down to an orientation estimate by means of the angle of the eigenvector $\mathbf{v}_1(\mathbf{x})$ in the Hessian matrix (consisting of Gaussian derivatives) with smallest absolute eigen value, so $\hat{\theta}(\mathbf{x}) = \angle \mathbf{v}_1(\mathbf{x})$ and we can attach the confidence $r(\mathbf{x}) = \frac{|\lambda_2(\mathbf{x})| - |\lambda_1(\mathbf{x})|}{|\lambda_1(\mathbf{x})| + |\lambda_2(\mathbf{x})|}$. Typically these kind of rough orientation estimates are noisy and unreliable, likewise the example in the middle image in the bottom row of Figure 17. Therefore, the main goal in this section is to obtain a more robust orientation estimates from the rough orientation estimate $\mathbf{x} \mapsto \hat{\theta}(\mathbf{x})$ with confidence $\mathbf{x} \mapsto r(\mathbf{x})$. This brings us to the framework of channel representations, (Granlund, 2000). This approach has quite some analogy with the orientation score framework of the previous sections in the sense that

- they both provide a new object which is an orientation decomposition of the image and which is a function on the Euclidean motion group $\mathbb{R}^2 \times \mathbb{T}$ (or rather $\mathbb{R}^2 \times \mathbb{T}_N$), and therefore all theory and algorithms concerning left-invariant operations on orientation scores may as well be applied to (the smoothing of) channel representations,
- they pose the requirement for exact invertibility,




Figure 14. Illustration of robust line detection in 3D images via orientation scores. 3D-images ($64 \times 64 \times 64$) illustrated only by 3 different 2D-cuts (along xy -plane, on $z = 2, 12, 22$). First column original image of a straight line and 2 circular spirals. Parameterized by respectively $(10 + (0.2)t, 20 + (0.2)t, t)$, $(32 + 10 \cos(\frac{2\pi t}{128}), 32 + 10 \sin(\frac{2\pi t}{128}), t)$ and $(20 + 12 \cos(\frac{2\pi t}{80}), 20 + 12 \sin(\frac{2\pi t}{80}), t)$. In the second column we added other geometrical structures, some spots and a cube. In the third column we obtained f_1 by adding strong correlated Gaussian distributed noise on the grey-values. 4th and 5th column: two layers of the orientation score $U_{f_1}(\cdot, \tilde{\mathbf{c}}_i) = \mathcal{W}_\psi[f_1]$, with wavelet $\psi_{N=16}^{3D}$ illustrated in Figure 13 and with $\tilde{\mathbf{c}}_2 \approx (-0.19, 0.58, 0.79)$, with $\tilde{\mathbf{c}}_7 \approx (-0.30, 0.93, 0.19)$, 6th column: the approximative reconstruction $f_1 * \tilde{\psi}$ and last column the enhanced/processed image after simple power enhancement (see (31)) in the orientation score. We did **not** use any thresholding on the grey-values (which is by definition an ill-posed operation). These experiments show us that we can enhance lines (and distinguish them from other geometrical structures such as lines and planes) already by extremely simple operation on the 3D-scores in a robust manner. For application of this technique in medical imaging (detection of Adam Kiewicz-vessel, responsible for blood supply to spinal court), we refer to van der Put (2005).

- admit an operational scheme for reconstruction from the orientation decomposition.

Nevertheless we stress that the starting point with respect to the invertibility of the frameworks is different: In the orientation score framework the reconstructable input is the original image $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, whereas the reconstructable input in the framework of channel representations is a first rough orientation estimate $\theta : \mathbb{R}^2 \rightarrow \mathbb{T}$, obtained¹⁷ from original image $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, rather than the original image itself.

Next we explain the method of robust orientation by channel representations in more detail: Let $L > 0$, then the channel representation $\mathbf{u} : [0, L] \times \mathbb{R}^+ \rightarrow \mathbb{R}^L$ is an encoding of a *signal value* $f(\mathbf{x}) \in [0, L]$ obtained by a modular¹⁸ *signal* $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ at position \mathbf{x} , and an associated confidence $r(\mathbf{x}) > 0$, given by

$$\mathbf{u}(f(\mathbf{x}), r(\mathbf{x})) = r(\mathbf{x})(B^0(f(\mathbf{x})), \dots, B^{L-1}(f(\mathbf{x})))^T,$$

where $B^n : [0, L] \rightarrow \mathbb{R}^+$, $n = 0, \dots, L-1$ are continuous functions, such that the mapping $(y, r) \mapsto \mathbf{u}(y, r)$ is injective. The vector $\mathbf{u}(f(\mathbf{x}), r(\mathbf{x}))$ is usually called

channel vector (at position $\mathbf{x} \in \mathbb{R}^2$) and its components are called channels. The injectivity assumption allows us to decode $(f(\mathbf{x}), r(\mathbf{x}))$ from the channel vector $\mathbf{u}(f(\mathbf{x}), r(\mathbf{x}))$.

Here we only consider the case where $f(\mathbf{x}) = \frac{L\hat{\theta}(\mathbf{x})}{2\pi}$, where $\hat{\theta}(\mathbf{x})$ is a first rough orientation estimation (for example by means of a gradient, Riesz transform, or Hessian of the original image). Thereby we have $f(\mathbf{x}) \in [0, L]$ for all $\mathbf{x} \in \mathbb{R}^2$ and the interval $[0, L]$ will be equipped with group product $f(\mathbf{x}) + g(\mathbf{x}) = (f(\mathbf{x}) + g(\mathbf{x})) \bmod L$ and modular distance

$$d_L(f(\mathbf{x}), g(\mathbf{x})) = \min(\text{mod}(f(\mathbf{x}) - g(\mathbf{x}), L), \text{mod}(g(\mathbf{x}) - f(\mathbf{x}), L)),$$

which makes $[0, L]$ isomorphic to the circle group \mathbb{T} . Moreover, we will only consider the special case

$$B^n(f(\mathbf{x})) = B^0(d_L(f(\mathbf{x}), n)),$$

i.e. the basis functions B^n are obtained by a modular shift over n from smooth basis function B^0 . The

support of the symmetric function B_0 will be a connected bounded interval $[-\frac{W}{2}, \frac{W}{2}]$, where W is called the kernel width. As a result the support of B^n equals $[n - \frac{W}{2}, n + \frac{W}{2}]$. Consequently, we have a nice practical localized smooth basis. Further we notice that the number N_A of active (non-zero) channels in the channel representation is limited as it equals $1 + 2\lfloor \frac{W}{2} \rfloor$. This is a major practical advantage over for example the discrete Fourier basis where all discrete frequencies are needed to represent a discrete δ -spike.

Example: B-spline channels The B-spline channel representation, (Felsberg et al., 2006), with parameter $K \in \mathbb{N}$ is obtained by setting

$$B_K^0(f(\mathbf{x})) = (\text{rect} *^{(K)} \text{rect})(f(\mathbf{x})),$$

where we used K -fold periodic convolution (so that the channel width equals $K + 1$) and where $\text{rect}(y) = 1$ if $|y| < \frac{1}{2}$ and $\text{rect}(y) = \frac{1}{2}$ if $|y| = \frac{1}{2}$ and $\text{rect}(y) = 0$ elsewhere. The decoding is linear and is given by

$$(f(\mathbf{x}), r(\mathbf{x})) = \left(\frac{1}{r(\mathbf{x})} \sum_{n=0}^{L-1} n u_K^n(f(\mathbf{x}), r(\mathbf{x})), \sum_{n=0}^{L-1} u_K^n(f(\mathbf{x}), r(\mathbf{x})) \right), \quad (38)$$

where $u_K^n(f(\mathbf{x}), r(\mathbf{x})) = r(\mathbf{x}) B_K^0(f(\mathbf{x}) - n)$. It is not difficult to give a formal proof of (38):

1. The first part of (38) can easily be proved by induction: The case $K = 1$ is trivial, moreover if we assume that it holds for $K - 1$ then we have

$$\begin{aligned} \sum_{n=0}^{L-1} n B_K^n(y) &= \sum_{n=0}^{L-1} n B_{K-1}^0(y - n) \\ &= \sum_{n=0}^{L-1} n \int_{n=0}^{L-1} B_0^0(u) B_{K-1}^0(y - n - u) du \\ &= \int_{n=0}^{L-1} B_0^0(u) \left(\sum_{n=0}^{L-1} n B_{K-1}^0(y - n - u) \right) du \\ &= \int_{n=0}^{L-1} B_0^0(u) (y - u) du = 1 \cdot y + 0 = y \end{aligned}$$

2. With respect to the second argument we notice that

$$\begin{aligned} \sum_{n=0}^{L-1} B_0^n(y) &= 1 \text{ for all } y \in [0, L] \Rightarrow \sum_{n=0}^{L-1} B_K^n(y) \\ &= 1 \text{ for all } y \in [0, L], \end{aligned}$$

from which we deduce that $\sum_{n=0}^{L-1} u^n(y, r) =$

$$r \sum_{n=0}^{L-1} B_K^n(y) = r.$$

In practice one uses decoding after manipulation and then one would like to use a decoding window (of size say $2p + 1$) and obtain the following fast decoding estimate

$$\tilde{f}(\mathbf{x}) = \frac{\sum_{n=n_0-p}^{n_0+p} n u^n(f(\mathbf{x}), r(\mathbf{x}))}{\sum_{n=n_0-p}^{n_0+p} u^n(f(\mathbf{x}), r(\mathbf{x}))}. \quad (39)$$

In the experiments discussed in subsection 7.1 we set $p = 1$ and $K = 2$, i.e. we used quadratic B-splines with a decoding window size of 3, where the decoding estimate (39) is given by

$$\tilde{f}(\mathbf{x}) = \frac{u^{n_0+1}(f(\mathbf{x}), r(\mathbf{x})) - u^{n_0-1}(f(\mathbf{x}), r(\mathbf{x}))}{u^{n_0+1}(f(\mathbf{x}), r(\mathbf{x})) + u^{n_0}(f(\mathbf{x}), r(\mathbf{x})) + u^{n_0-1}(f(\mathbf{x}), r(\mathbf{x}))} + n_0, \quad (40)$$

the window center n_0 can be chosen in various (somewhat ad-hoc) ways. For example n_0 could be chosen such that $u^{n_0+1}(f(\mathbf{x}), r(\mathbf{x})) + u^{n_0}(f(\mathbf{x}), r(\mathbf{x})) + u^{n_0-1}(f(\mathbf{x}), r(\mathbf{x}))$ is maximal.

Example: Cosine square channels The cosine square channel representation, cf. Forssén and Granlund, (2000), with parameter $\omega > 0$ is obtained by setting

$$B^\omega(f(\mathbf{x})) = C_\omega^0(f(\mathbf{x})) := \begin{cases} \cos^2(\omega f(\mathbf{x})) & \text{If } \omega |f(\mathbf{x})| \leq \frac{\pi}{2}, \\ 0 & \text{else.} \end{cases} \quad (41)$$

Note that $\frac{\pi}{\omega}$ equals the channel width.

The decoding algorithm is more complicated in this case compared to the B-spline case. It is non-linear and involves \mathbb{L}_2 -approximation for $\omega > \frac{\pi}{N}$, but it is possible to give an exact decoding for the case $\omega = \frac{\pi}{N}$ (i.e. channel width $W = N$), $N \geq 3$:

$$\begin{aligned} f(\mathbf{x}) &= l + \frac{N}{2\pi} \arg \left[\sum_{n=l}^{l+N-1} u_{\omega=\frac{\pi}{N}}^n(f(\mathbf{x}), r(\mathbf{x})) e^{i \frac{2(n-l)\pi}{N}} \right], \\ &= l + \frac{N}{2\pi} \arg \left[\sum_{n=0}^{N-1} u_{\omega=\frac{\pi}{N}}^{n+l}(f(\mathbf{x}), 1) e^{i \frac{2n\pi}{N}} \right] \end{aligned} \quad (42)$$

where l is the index of the first channel in the decoding interval, see Forssén (2004) p. 17–19 for more details.

7.1. Channel Smoothing for Robust Orientation Estimation

In this section we will consider robust orientation estimation by means of channel smoothing. The idea is to




Figure 15. Illustration of robust orientation estimation by means of channel smoothing.

construct a channel representation, to smooth the channel representation and to decode from the smoothed channels to obtain an orientation estimate, see Figure 15.

If the decoding scheme $f(\mathbf{x}) = \Phi^{-1}[\mathbf{u}(f(\mathbf{x}), r(\mathbf{x}))]$ is linear we have that smoothing of f (by means of a continuous convolution kernel $K : \mathbb{R}^2 \rightarrow \mathbb{R}^+$, with $\int_{\mathbb{R}^2} K(\mathbf{x})d\mathbf{x} = 1$) leads to the same result as decoding its smoothed channel representation $\tilde{\mathbf{u}}(f(\mathbf{x}), r(\mathbf{x})) = (K * \mathbf{u}(f(\cdot), r(\cdot)))(\mathbf{x})$. To this end we notice that

$$\begin{aligned} (K * f)(\mathbf{x}) &= \Phi^{-1}[\tilde{\mathbf{u}}(f(\mathbf{x}), r(\mathbf{x}))] \Leftrightarrow \\ \int_{\mathbb{R}^d} f(\mathbf{y})K(\mathbf{x}-\mathbf{y})d\mathbf{y} &= \int_{\mathbb{R}^d} \Phi^{-1}[\mathbf{u}(f(\mathbf{x}), r(\mathbf{x}))]K(\mathbf{x}-\mathbf{y})d\mathbf{y} \\ &= \Phi^{-1}\left[\int_{\mathbb{R}^d} \mathbf{u}(f(\mathbf{x}), r)K(\mathbf{x}-\mathbf{y})d\mathbf{y}\right], \end{aligned} \quad (43)$$

for all $\mathbf{x} \in \mathbb{R}^d$. The windowed decoding in the B -spline case and the windowed decoding in the cosine square case (42), or similarly for general ω in Forssén (2004) p. 18 are non-linear (even in case $r(\mathbf{x}) = 1$). As a result, in general, the left hand side and right hand side (43) do not coincide. In fact, they only coincide if the orientation measurement $\mathbf{x} \mapsto f(\mathbf{x})$ is such that the decoding becomes linear. For example, in the B -spline case, it requires that $f(\mathbf{x})$ is such that the active part of the smoothed channels is a subset of the decoding window, since then we have $\tilde{u}_{n_0-1}(f) + \tilde{u}_{n_0}(f) + \tilde{u}_{n_0+1}(f) = 1$ with the consequence that (40) and (38) coincide. This will only be the case if the image is locally 1D at \mathbf{x} . Also at discontinuities the channel will be smoothed, but here the contributions of the smoothing falls outside the decoding window, i.e. the smoothing contributions lie in the nil-space of the decoding operation. So, effectively the channel smoothing does not smooth over discontinuities, which is desirable in a robust orientation estimation.

In the cosine square channel case a similar observation can be made. Here we notice that if

$$\begin{aligned} &\arg\left[\sum_{k=0}^{N-1} \tilde{u}^{k+l}(f(\mathbf{x}), r(\mathbf{x}))e^{\frac{2\pi ik}{N}}\right] \\ &= \frac{1}{N} \arg\left[\prod_{k=0}^{N-1} u^{k+l}e^{\frac{2\pi ik}{N}}\right] \\ &= \frac{1}{N} \sum_{k=0}^{N-1} \arg(u^{k+l}e^{\frac{2\pi ik}{N}}) = \frac{\pi(N-1)}{N}, \end{aligned}$$

which is the case¹⁹ if $\tilde{\mathbf{u}}(f(\mathbf{x}), r(\mathbf{x}))$ is symmetric (for the sake of simplicity we assume N is odd):

$$\begin{aligned} \tilde{u}^{l+\frac{N-1}{2}+j}(f(\mathbf{x}), r(\mathbf{x})) &= \tilde{u}^{l+\frac{N-1}{2}-j}(f(\mathbf{x}), r(\mathbf{x})), \\ j &= 1, \dots, \frac{N-1}{2}, \end{aligned}$$

which again will be the case if the image is locally 1D at \mathbf{x} . Moreover we again have that effectively the channel smoothing does not smooth over discontinuities.

7.1.1. Robust Estimation In this section we will draw some parallels from channel smoothing to non-parametric methods in random variable estimation. Assume that $f : \mathbf{x} \mapsto f(\mathbf{x})$ is a realization of a stochastic process \mathcal{P} that is ergodic for all $\mathbf{x} \in \Omega$. This implies that we can replace averaging realizations of \mathcal{P} with averaging f in $f(\Omega) \subset [0, L]$. We denote the probability density function over f by pdf.

The orientation estimate due to decoding is robust in the sense that it minimizes the energy

$$\mathcal{E}(g) = \int_0^L \rho(f-g)\text{pdf}(f)df = (\rho * \text{pdf})(g).$$

If we use B -spline channels, it follows by (40) that the influence function is given by

$$\psi(f) = \rho'(f) = B_2(f-1) + B_2(f+1).$$

This influence function is smooth and compactly supported at $[-2.5, 2.5]$ and locally linear near the origin. The influence function tells us how much an outlier affects the orientation estimate. To this end we notice that $\frac{\partial \mathcal{E}(f)}{\partial f} \Big|_{f=f_0} = -(\psi * \text{pdf})(f)$. For example, if the outlier is outside the compact support $[-2.5, 2.5]$ of the influence function it does not affect the estimate and if the outlier equals ± 1 then puts maximal damage to the estimate. Clearly, this influence function (see




Figure 16. The influence function in case of robust orientation estimation by means of channel smoothing on B -splines.

Figure 16) is highly preferable over the linear influence function in the non-robust case of a least square estimate $\rho(f) = f^2$.

The channel vector $\mathbf{u}(\mathbf{x})$ at position \mathbf{x} is related to the probability density function (pdf) of the generating stochastic process \mathcal{P} at position \mathbf{x} . By definition a kernel density estimation with kernel B_0 is given by $\tilde{p}_{f_n}(f) = \frac{1}{N} \sum_{n=0}^N B_0(f - f_n)$, where $\{f_n\}$ is a set of N independent samples of the stochastic variable f .

By the ergodicity assumption averaging of the elements of the channel representation $\tilde{u}^n(f(\mathbf{x}), 1) = \frac{1}{|\Omega|} \int_{\Omega} u^n(f(\mathbf{x}), 1) d\mathbf{x}$, is equivalent to sampling of a kernel density estimate with the symmetric positive kernel B^0 . The expectation value of the kernel density estimate equals $(B^0 * \text{pdf})(f)$ and thereby the expectation of the channel averaging is

$$E \left\{ \int_{\Omega} u^n(f(\mathbf{x}), 1) \right\} = (B^0 * \text{pdf})(f) \Big|_{f=f_n}. \quad (44)$$

7.1.2. Explicit Example of Channel Smoothing

Analogue to the orientation score framework where we ensured Euclidean invariance of the processing on the original image $\mathbf{x} \mapsto f(\mathbf{x})$ we would like to ensure Euclidean invariant processing on the original orientation map $\mathbf{x} \mapsto \hat{\theta}_f(\mathbf{x})$. This means that operations on the channel representation should be left invariant. Therefore channel smoothing should be done by means of a G_L -convolution, recall (12), the discrete version of the G -convolution given by (25):

$$\begin{aligned} \tilde{u}^n(\mathbf{x}) &= (\tilde{K} *_{G_L} u)(\mathbf{x}, e^{in\Delta\theta}) \\ &= \sum_{k=0}^{L-1} \int_{\mathbb{R}^2} \tilde{K}(R_{k\Delta\theta}^{-1}(\mathbf{x} - \mathbf{b}'), e^{i(n-k)\Delta\theta}) U(\mathbf{b}', e^{ik\Delta\theta}) d\mathbf{b}', \\ n &= 0, \dots, L-1, \end{aligned}$$

where L denotes the total number of channels, where $G_L = \mathbb{R}^2 \rtimes \mathbb{T}_L$ and with $U : G_L \rightarrow \mathbb{R}$ is given by $U(\mathbf{x}, e^{in\theta}) = u^n(\mathbf{x})$, where $\tilde{K} : G_L \rightarrow \mathbb{R}$ represents the smoothing kernel and where $\{u^n\}_{n=0}^{L-1}$ respectively $\{\tilde{u}^n\}_{n=0}^{L-1}$ represents the input and output channels. In particular if \tilde{K} is a singular with respect to the angle²⁰, that is $\tilde{K}(\mathbf{x}, e^{in\Delta\theta}) = \delta_n K(\mathbf{x})$, we obtain:

$$\tilde{u}^n(\mathbf{x}) = (K(R_{n\Delta\theta}^{-1}) *_{\mathbb{R}^2} u^n)(\mathbf{x}).$$

To obtain *adaptive* channel smoothing we quadratically fitted the convolution kernels $K \geq 0$ of the type $\psi_{s,\omega}(\mathbf{x}) = e^{-rs} C_{\omega}^0(\phi)$, $\mathbf{x} = (r \cos \phi, r \sin \phi)$, where we recall that C_{ω}^0 is given by (41), to the auto-correlations

$$A_n(\mathbf{x}) = \mathcal{F}^{-1}[\omega \mapsto |\mathcal{F}(u^n(\cdot, r=1))(\omega)|^2](\mathbf{x})$$

of the channel vectors $u^n(\cdot, r=1)$. More precisely, we minimized²¹

$$\min_{s,\omega} \sum_{n=0}^{N-1} \int_{\mathbb{R}^2} \left| \frac{\|A_n\|_{\mathbb{L}_1(\mathbb{R}^2)}}{\|\psi_{s,\omega}\|_{\mathbb{L}_1(\mathbb{R}^2)}} \psi_{s,\omega}(\mathbf{x}) - A_n(R_n^{-1}\mathbf{x}) \right|^2 w(\mathbf{x}) d\mathbf{x},$$

where $\|A_n\|_{\mathbb{L}_1(\mathbb{R}^2)} = \int_{\mathbb{R}^2} |A_n(\mathbf{x})| d\mathbf{x}$ and where $R_n \in SO(2)$ is a rotation such that the direction of the n -th channel is mapped to the direction of the 0-th channel, so that $\mathbf{x} \mapsto A_n(R_n^{-1}\mathbf{x})$ is aligned to $A_0(\mathbf{x})$. In the implementation we replaced the integral by a discrete summation and using a steepest descent algorithm. We took $w(\mathbf{x}) = \frac{1}{\|\mathbf{x}\| + \epsilon}$ ($1 \gg \epsilon > 0$, to avoid singularity problems at the origin) to compensate for the Jacobian of a polar coordinate transformation. Notice to this end that local orientation means locally 1D signals, so we wanted to have a 1D subspace projection along the orientation.

Some results are illustrated in Figure 17. This method showed very good results compared to other existing orientation estimation methods. For a detailed comparison to other orientation estimation methods, we refer to earlier work (Felsberg et al., 2004) p. 22–25.

7.1.3. The Number of Channels L and the Channel Width W

The shape of the B -spline and the cosine square channel representation are pretty similar if the corresponding channel widths coincide, i.e. $\frac{\pi}{\omega} = K + 1$.

In practice there is not much difference between the two representations. It is rather the number of channels L and the channel width W that makes the difference. Both quantities lead to a trade-off situation.




Figure 17. Robust orientation estimation by means of adaptive channel smoothing. Top row: from left to right, test pattern, $t =$ test pattern with 1 percent salt and pepper noise and Gaussian noise of variance 0.01, ground truth of the orientation. Bottom row: from left to right, we considered orientation estimation only inside the white disk area, and the rough orientation estimation computed from gradient, the result after channel smoothing of $f(x) = 2 * \arccos\left(\frac{\text{gradr}(x) \cdot \mathbf{e}}{\|\text{gradr}(x)\|}\right)$, where $\mathbf{e} \in \mathbb{R}^2$ is a fixed normalized direction.

On the one hand the channel width W should not be too narrow, since then the estimated pdf becomes spiky, and consequently finding the maxima of the pdf is ill-posed. On the other hand W should not be too broad, as this leads to interference of the local maxima.

On the one hand the number of channels L should not be too large, as it requires a lot of samples of the pdf and there is only a limited number of samples of the measurement. On the other hand L should not be too small as the accuracy of the maxima of the pdf becomes better if there are more channels.

8. Conclusion

Given an image $f \in \mathbb{L}_2(\mathbb{R}^d)$, we construct a local orientation score U_f , which is a complex-valued function on the Euclidean motion group G . The corresponding transformation $f \mapsto U_f$ is a wavelet transform constructed from an oriented wavelet and a representation of the Euclidean motion group onto $\mathbb{L}_2(\mathbb{R}^2)$. Since this representation \mathcal{U} is *reducible* the well-known wavelet reconstruction theorem, which allows perfectly well-posed reconstruction, does not apply.

Therefore we generalized standard wavelet theory by means of reproducing kernel theory. From this new generalization it follows that our wavelet transformation is a *unitary* mapping between the space of bandlimited images, modelled by $\mathbb{L}_2^g(\mathbb{R}^2)$ and the functional

Hilbert space \mathbb{C}_K^G of orientation scores which is explicitly characterized. The norm on \mathbb{C}_K^G explicitly depends on the oriented wavelet ψ via a function M_ψ , which thereby characterizes the stability of the explicitly described inverse wavelet transformation. As a result, by proper choice of the wavelet ψ , the image f can be reconstructed from U_f in a robust way. We developed and implemented several approaches to obtain proper wavelets (which are also good line detectors in practice) in the 2 dimensional case, $d = 2$. These results are generalized to (and implemented in) the 3 dimensional case, $d = 3$.

These proper wavelets give rise to a stable transformation from an image to a *single scale* orientation score and vice versa, which allows us to relate operations Φ on orientation scores to operations Υ on images in a robust way.

We showed that operations Φ must be left invariant in order to obtain a Euclidean invariant transformation Υ on images. As a result the only left-invariant kernel operators on the orientation scores are G -convolutions. As an example we observe the probability of collision of particles from two stochastic processes on the Euclidean motion group, which is used to automatically detect elongated structures and to close gaps between them.

Finally, we focussed on robust orientation estimation rather than image enhancement. For this purpose we

used the framework of (invertible) channel representations, which has a lot of analogy with the framework of invertible orientation scores, where a *single* orientation estimate for each position is encoded into channel vectors. We obtained robust orientation estimates via the decoding algorithm after channel smoothing (which is done by a discrete G -convolution on the orientation channels). This is analogue to image processing via left invariant stochastic processes on invertible orientation scores in the sense that the processing is done by left-invariant processing on the Euclidean motion group followed by an inverse transformation to the space of input. The difference though is that in the framework of invertible orientation channels the input is given by the original orientation estimates $\hat{\theta}_f : \mathbb{R}^2 \rightarrow \mathbb{T}$ (obtained from image $f : \mathbb{R}^2 \rightarrow \mathbb{R}$), whereas in the framework of invertible orientation scores the input is the original image f .

A. The Bochner-Hecke Theorem and the Spectral Decomposition of the Hankel Transform

Theorem 7. *Let H be a harmonic homogeneous polynomial of degree m in d variables. Let F be an element of $\mathbb{L}_2((0, r); r^{d+2m-1} dr)$ then the Fourier transform of their direct product $(r, \mathbf{x}) \mapsto F(r)H(\mathbf{x})$, which is in $\mathbb{L}_2(\mathbb{R}^d)$, is given by:*

$$\begin{aligned} (-i)^m H(\omega) \int_0^\infty (r\rho)^{-(\frac{d-2}{2}+m)} J_{\frac{d-2}{2}+m}(\rho r) F(r) r^{d+2m-1} dr = \\ (-i)^m H(\omega) \rho^{-(\frac{d-1}{2}+m)} \left[\mathcal{H}_{\frac{d-2}{2}+m} r^{\frac{d-1}{2}+m} F \right](\rho), \end{aligned}$$

where $\rho = \|\omega\|$ and the $\mathcal{H}_{\frac{d-2}{2}+m}$ the Hankel Transform given by:

$$\begin{aligned} (\mathcal{H}_{\frac{d-2}{2}+m} \phi)(\rho) = \int_0^\infty (\rho r)^{1/2} \phi(r) J_{\frac{d-2}{2}+m}(\rho r) dr, \\ \phi \in \mathbb{L}_2((0, \infty)). \end{aligned} \quad (45)$$

The proof can be found in Faraut and Harzallah (1984) pp. 24–25. The Hankel Transform \mathcal{H}_μ , $\mu = \frac{d-2}{2} + m$, is a unitary map on $\mathbb{L}_2((0, \infty), dr)$ and has a complete set of orthonormal eigen functions $\{E_n^\mu\}$ given by

$$\begin{aligned} r \mapsto E_n^\mu(r) = \left(\frac{2n!}{\Gamma(n + \mu + 1)} \right)^{\frac{1}{2}} r^{\mu + \frac{1}{2}} e^{-\frac{r^2}{2}} L_n^{(\mu)}(r^2), \\ n = 0, 1, 2, \dots, r > 0, \end{aligned} \quad (46)$$

where $L_n^{(\mu)}(r)$ is the n -th generalized Laguerre polynomial of type $\mu > -1$,

$$L_n^{(\mu)}(r) = \frac{r^{-\alpha} e^r}{n!} \left(\frac{d}{dr} \right)^n (e^{-r} r^{n+\mu}), r > 0,$$

with corresponding eigen values $(-1)^n$: $(\mathcal{H}_\mu \phi) = \sum_{n=0}^\infty (-1)^n (E_n^\mu, \phi)_{\mathbb{L}_2((0, \infty), dr)} E_n^\mu$. The functions E_n^μ are also eigen functions of the operator $-\frac{d^2}{dr^2} + r^2 + \frac{\mu^2 - \frac{1}{4}}{r^2} - 2\alpha$, with eigen value $4n + 2$, cf. (Eijndhoven and de Graaf, 1982) p. 79, which coincides with the fact that the Harmonic oscillator $\|\mathbf{x}\|^2 - \Delta$ commutes with Fourier transform. These results give us:

- For $d = 2$, we have $\mathbb{L}_2(\mathbb{R}^2) = \mathbb{L}_2(S^1) \otimes \mathbb{L}_2((0, \infty); rdr)$ and a Fourier invariant orthonormal base is given by $\{Y_m \otimes h_n^m\}_{m \in \mathbb{Z}, n \in \mathbb{N} \cup \{0\}}$, where

$$h_n^m = r^{-1/2} E_n^m(r), \quad (47)$$

and $Y_m(\phi) = \frac{1}{\sqrt{2\pi}} e^{im\phi}$. It now follows by the Bochner-Hecke Theorem that:

$$\begin{aligned} \mathcal{F}(Y_m \otimes h_n^m) &= \mathcal{F}(r^m Y_m \otimes \frac{h_n^m}{r^m}) \\ &= (-1)^n (-i)^{|m|} (Y_m \otimes h_n^m) \\ &= (-1)^{n+m} (i)^{|m|} (Y_m \otimes h_n^m). \end{aligned} \quad (48)$$

- For $d = 3$, we have $\mathbb{L}_2(\mathbb{R}^2) = \mathbb{L}_2(S^1) \otimes \mathbb{L}_2((0, \infty); rdr)$. All l homogeneous harmonic polynomials are spanned by $\{\mathbf{x} \mapsto r^l Y_l^m(\theta, \phi)\}_{l=0, \dots, \infty; m=-l, \dots, l}$. Define

$$g_n^l(r) = r^{-1} E_n^{l+\frac{1}{2}}(r), r > 0 \quad (49)$$

then $g_n^l \in \mathbb{L}_2((0, \infty); rdr)$ are eigen functions of $\rho^{-1} \mathcal{H}_{\frac{1}{2}+m} r^1$ with corresponding eigenvalues $(-1)^n$. Therefore it follows by the Bochner-Hecke Theorem that

$$\mathcal{F}(Y_l^m \otimes g_n^l) = (i)^l (-1)^{n+l} (Y_l^m \otimes g_n^l). \quad (50)$$

Notes

1. The torus \mathbb{T} is the group of elements in the set S^1 with group product $e^{i\theta} e^{i\theta'} = e^{i(\theta+\theta')}$
2. This does by no means imply that the visual system actually runs an inverse process.
3. Notice that this depends on the norm imposed on the set of images and on the norm imposed on the set of orientation scores.

4. They are in fact, up to equivalence, the only irreducible representations of the Euclidean Motion group, cf. (Sugiura, 1990).
5. Also the space $L_2^g(\mathbb{R}^2)$ is a reproducing kernel Hilbert space with reproducing kernel $K(\mathbf{x}, \mathbf{x}') = \mathcal{F}^{-1}[1_{B_{0,e}}](\mathbf{x} - \mathbf{x}') = J_1\left(\frac{\|\mathbf{x}-\mathbf{x}'\|}{e}\right)\frac{e}{\|\mathbf{x}-\mathbf{x}'\|}$.
6. Special cases include frame theory, cf. (Duits, 2005), sampling theorems and wavelet theory cf. (Duits, 2004)
7. i.e. is a subspace as a vector space, but is equipped with a different norm.
8. This commutes with the left regular actions \mathcal{U}_g for all $g \in \mathbb{R}^d \rtimes SO(d)$.
9. We stress that even if $M_\psi \neq 1$ stability is still manifest. The only requirement on M_ψ is that it remains overall finite and non-vanishing. Recall that in general one has to use (11) for exact reconstruction.
10. The eigen-values of $Y_m \otimes h_n^m$, $m \in \mathbb{Z}$, $n \in \mathbb{N} \cup \{0\}$, with respect to \mathcal{F} , $\mathbb{R}_{e^{i\theta}}$, $\|\mathbf{x}\|^2 - \Delta$ are respectively $(-1)^{n+m} |j|^{m|}$, $e^{im\theta}$, $2(2n + |m| + 1)$.
11. This means $f * \psi_N^0 \rightarrow f$ in L_2 -sense for all $f \in L_2(\mathbb{R}^2)$ as $N \rightarrow \infty$
12. Once again exact reconstruction can be obtained by (11), but to avoid any kind of de-blurring, that is divisions in the Fourier domain, we give fast and simple approximative reconstruction schemes (either by (13) or even by (21)) that are sufficient in practical applications.
13. Note that the approximative reconstruction can be written $f * \tilde{\psi} = (\mathcal{W}_\psi^e)^* \mathcal{W}_\psi[f]$, which approximates the original image $f = (\mathcal{W}_\psi)^* \mathcal{W}_\psi[f]$.
14. One can always compute the angle between $\Phi(U_f)$ and \mathbb{C}_K^G to see how effective the operation Φ is (in most of our applications this angle was small).
15. In a scalespace representation $u(\mathbf{x}, s) = (G_s * f)(\mathbf{x})$ the evolution parameter $s = \frac{\sigma^2}{2}$ inherits the physical dimension of the generator Δ of the corresponding evolution equation $u_s = \Delta u$ and thereby in image analysis $s > 0$ is usually considered as the scale of observation of an image f . Scale can be related to time via a diffusion constant D : $Dt = s = \frac{1}{2}\sigma^2$.
16. In Thornber and Williams (1996) it is claimed that the authors found the analytic solution of the direction process. This claim is incorrect: Although their Green's function is for some parameter settings an appropriate approximation of the true solution, it is not the exact Green's function (it is a Green's function on a group of Heisenberg type, not on the Euclidean motion group). This is easily verified by substitution of their solution into the partial differential equation. We did find the analytic solution (in fact we found the solutions of all second order left invariant evolution equations on the Euclidean motion group) in terms of elliptic functions of a special kind, the existence of which was conjectured by Mumford cf (Mumford, 1994) p. 497. This issue is the main topic of our recent work (Duits and van Almsick, 2005).
17. For example obtained by an orientation score.
18. With modular, we mean periodic in its co-domain. For channel representations on non-modular signals, see Forssén (2004).
19. $\arg(e^{\frac{\pi i(N-1)}{N}}(1+z+\bar{z})) = \arg(e^{\frac{\pi i(N-1)}{N}}) = \frac{\pi(N-1)}{N}$ for all $z \in \mathbb{C}$.
20. One may also want to consider non-singular convolution kernels, to model interaction between the orientation channels.
21. Although this minimization works fine, it does not result in the best possible noise suppression. For an optimal minimization

(with the same type of kernel) concerning the signal to noise ratio we refer to Felsberg et al. (2006).

References

- Ali S.T., Antoine J.P. and Gazeau J.P. 1999. *Coherent States, Wavelets and Their Generalizations*. Springer Verlag, New York, Berlin, Heidelberg.
- Antoine J.P. 1999. Directional wavelets revisited: Cauchy wavelets and symmetry detection in patterns. *Applied and Computational Harmonic Analysis*, 6:314–345.
- Aronszajn, N. 1950. Theory of reproducing kernels. *Trans. A.M.S.*, 68:337–404.
- August J. and Zucker S.W. 2003. The curve indicator random field and markov processes. *IEEE-PAMI, Pattern Recognition and Machine Intelligence*, 25.
- Bosking W.H., Zhang Y., Schofield B. and Fitzpatrick D. 1997. Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *The Journal of Neuroscience*, 17(6):2112–2127.
- Duits M. 2004. A functional Hilbert space approach to frame transforms and wavelet transforms. Master thesis in Applied Analysis group at the department of Mathematics and Computer Science at the Eindhoven University of Technology.
- Duits, M. and Duits, R. 2004. A functional Hilbert space approach to the theory of wavelets. Technical report, TUE, Eindhoven, RANA/CASA Report RANA-7-2004, available on the web: <ftp://ftp.win.tue.nl/pub/rana/rana04-07.pdf> Department of Mathematics Eindhoven University of Technology.
- Duits, R. 2005. *Perceptual Organization in Image Analysis*. PhD thesis, Eindhoven University of Technology, Department of Biomedical Engineering, The Netherlands. A digital version is available on the web: URL: <http://www.bmi2.bmt.tue.nl/Image-Analysis/People/RDuits/THESISRDUITS.pdf>.
- Duits, R. and van Almsick, M. 2005. The explicit solutions of the left invariant evolution equations on the Euclidean motion group. Eindhoven University of Technology, Eindhoven, 5–43. Available on the web <http://yp.bmt.tue.nl/pdfs/6321.pdf>. An improved version of which is recently submitted to the Quarterly of Applied Mathematics (journal of American Mathematical Society.)
- Duits, R., Duits, M. and van Almsick, M. 2004. Invertible orientation scores as an application of generalized wavelet theory. Technical report, TUE, Eindhoven. Technical Report 04-04, Biomedical Image and Analysis, Department of Biomedical Engineering, Eindhoven University of Technology.
- Duits, R., Florack, L.M.J., de Graaf, J. and ter Haar Romeny, B. 2004. On the axioms of scale space theory. *Journal of Mathematical Imaging and Vision*, 20:267–298.
- Duits, R., van Almsick, M., Duits, M., Franken, E. and Florack, L.M.J. 2004. Image processing via shift-twist invariant operations on orientation bundle functions. In Niemann Zhuravlev et al. Gepener, Gurevich, editor, *7th International Conference on Pattern Recognition and Image Analysis: New Information Technologies*, 193–196, St. Petersburg. Extended version is to appear in special issue of the International Journal for Pattern Recognition and Image Analysis MAIK.
- Dungey, N., ter Elst, A.F.M. and Robinson, D.W. 2003. *Analysis on Lie groups with polynomial growth*, volume 214. Birkhauser-Progress in Mathematics, Boston.

- Eijndhoven, S.J.L. and de Graaf, J. 1982. Some results on hankel invariant distribution spaces. *Proceedings of the Koninklijke Akademie van Wetenschappen, Series A*, 86(1):77–87.
- Faraut, J. and Harzallah, K. 1984. *Deux cours d'analyse harmonique*. Birkhaeuser, Tunis.
- Felsberg, M., Forssén, P.-E. and Scharf, H. 2004. Efficient robust smoothing of low-level signal features. Technical Report LiTH-ISY-R-2619, SE-581 83 Linköping, Sweden.
- Felsberg, M., Forssén, P.-E. and Scharf, H. 2006. Channel smoothing: Efficient robust smoothing of low-level signal features. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(2):209–222.
- Florack, L.M.J., 1997. *Image Structure*. Kluwer Academic Publishers, Dordrecht, The Netherlands.
- Forssén, P.-E. and Granlund, G. H. 2000. Sparse feature maps in a scale hierarchy. In G. Sommer and Y.Y. Zeevi, editors, *Proc. Int. Workshop on Algebraic Frames for the Perception-Action Cycle*, volume 1888 of *Lecture Notes in Computer Science*, Kiel, Germany, Springer, Heidelberg.
- Forssén, P.E. 2004. *Low and Medium Level Vision using Channel Representations*. PhD thesis, Linköping University, Dept. EE, Linköping, Sweden.
- van Ginkel, M. 2002. *Image Analysis using Orientation Space based on Steerable Filters*. PhD thesis, Delft University of Technology, Delft, Netherlands.
- Granlund, G.H. 2000. An associative perception-action structure using a localized space variant information representation. In *Proceedings of Algebraic Frames for the Perception-Action Cycle (AFPAC)*, Kiel, Germany, Also as Technical Report LiTH-ISY-R-2255.
- Grossmann, A., Morlet, J. and Paul, T. 1985. Integral transforms associated to square integrable representations. *J.Math.Phys.*, 26:2473–2479.
- Isham, C.J. and Klauder J.R. 1991. Coherent states for n -dimensional euclidean groups $e(n)$ and their application. *Journal of Mathematical Physics*, 32(3):607–620.
- Kalitzin, S.N., ter Haar Romeny, B.M. and Viergever, M.A. 1999. Invertible apertured orientation filters in image analysis. *International Journal of Computer Vision*, 31(2/3):145–158.
- Lee T.S., 1996. Image representation using 2d gabor wavelets. *IEEE-Transactions on Pattern Analysis and Machine Intelligence*, 18(10):959–971.
- Louis, A.K., Maass, P. and Rieder, A. 1997. *Wavelets, Theory and Applications*. Wiley, New York.
- Martens, F.J.L. 2004. *Spaces of analytic functions on inductive/projective limits of Hilbert Spaces*. PhD thesis, University of Technology Eindhoven, Department of Mathematics and Computing Science, Eindhoven, The Netherlands, 1988. This PHD thesis is available on the webpages of the Technische Universiteit Eindhoven. Webpage in: <http://alexandria.tue.nl/extra3/proefschrift/PRF6A/8810117.pdf>.
- Mumford, D. 1994. *Elastica and computer vision. Algebraic Geometry and Its Applications*. Springer-Verlag, 491–506.
- van der Put, R.W. 2005. Methods for 3d orientation analysis and their application to the study of arterial remodelling. Master's thesis, Department of Biomedical Engineering Eindhoven University of Technology, Technical Report BMIA-0502.
- Sugiura, M. 1990. *Unitary representations and harmonic analysis*. North-Holland Mathematical Library, 44., Amsterdam, Kodansha, Tokyo, second edition.
- Thorner, K.K. and Williams, L.R. 1996. Analytic solution of stochastic completion fields. *Biological Cybernetics*, 75:141–151.
- Ts'0, D.Y., Frostig, R.D., Lieke, E.E. and Grinvald, A. 1990. Functional organization of primate visual cortex revealed by high resolution optical imaging. *Science*, 249:417–20.
- Twareque Ali, S. 1998. A general theorem on square-integrability: Vector coherent states. *Journal of Mathematical Physics*, 39.
- van Almsick, M.A., Duits, R., Franken E. and ter Haar Romeny, B.M. 2005. From stochastic completion fields to tensor voting. In *Proceedings DSSCC-workshop on Deep Structure Singularities and Computer Vision*, Maastricht the Netherlands, Springer-Verlag.
- Williams, L.R. and Zweck, J.W. 2003. A rotation and translation invariant saliency network. *Biological Cybernetics*, 88: 2–10.