

Elliptical Monogenic Wavelets for the analysis and processing of color images

Raphaël Souillard and Philippe Carré

Abstract—This paper studies and gives new algorithms for image processing based on monogenic wavelets. Existing greyscale monogenic filterbanks are reviewed and we reveal a lack of discussion about the synthesis part. The monogenic synthesis is therefore defined from the idea of wavelet modulation, and an innovative filterbank is constructed by using the Radon transform. The color extension is then investigated. First, the elliptical Fourier atom model is proposed to generalize the analytic signal representation for vector-valued signals. Then a color Riesz-transform is defined so as to construct color elliptical monogenic wavelets. Our Radon-based monogenic filterbank can be easily extended to color according to this definition. The proposed wavelet representation provides efficient analysis of local features in terms of shape and color, thanks to the concepts of amplitude, phase, orientation, and ellipse parameters. The synthesis from local features is deeply studied. We conclude the article by defining the color local frequency, proposing an estimation algorithm.

Index Terms—wavelet transform, color, analytic signal, monogenic signal, monogenic wavelets.

I. INTRODUCTION

Through the last 15 years, the *analytic signal* and its 2D generalizations have greatly improved wavelet representations [7], [23], [40], [48] by naturally embedding AM-FM analysis in the subband coding framework. This yields an efficient representation of geometric structures in grayscale images thanks to a *local phase* carrying geometric information complementary to an *amplitude envelope* having good invariance properties. The signal is encoded with more sparsity than with standard wavelets, which globally improves most wavelet-based applications. The seemingly most appropriate 2D extension for image analysis is the *monogenic* representation [15], of which a few wavelet counterparts have been proposed [23], [38], [48] and applied [2], [27], [30], [35], [52].

Existing monogenic wavelet transforms are reviewed in section II, and it turns out that from the numerical point of view, the topic is still at its very first steps.

Manuscript received April 2, 2015; September 10, 2015. First published December ??, 2015; current version published December ??, 2015. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Wenwu Wang.

Copyright © 2015 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

Authors are with the XLIM-SIC Dept. UMR CNRS 7252 - University of Poitiers, France.

This paper has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the author. The material includes the MATLAB source code. Contact raphael.souillard@univ-poitiers.fr for further questions about this work.

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

Digital Object Identifier ??/TSP.???/????

Our main concern is that there is no satisfactory synthesis method from the monogenic wavelet analysis, whereas this is required for a large category of image processing applications, such as denoising, enhancement, compression and watermarking. While studying the invertibility and the synthesis monogenic wavelets, the information encoded in coefficients will be highlighted.

In this paper, we define a clear synthesis method from monogenic coefficients through the idea of *wavelet modulation*. An original numerical algorithm of reversible monogenic filterbank is derived in section III. The scheme is based on the Radon transform and 1D undecimated analytic wavelets. The reconstruction is exact and the synthesis wavelets are modulated according to monogenic coefficients.

The second part of the paper investigates the color extension in section IV. Analyzing color data is essential for a lot of applications but the Fourier analysis and wavelet representations have been almost exclusively limited to scalar signals so far. In practice, processing of color images is most often based on a suboptimal use of scalar tools, either by only considering luminance, or by working separately on each color channel (“marginal” schemes [3]).

We propose to mathematically handle the dimensionality of color signals by starting back to the univariate case. The elliptical oscillation model is proposed to generalize the *analytic signal* to vector-valued signals. This method inspired from [32] provides a new clearly interpretable color phase concept.

The color definition is then extended to 2D. We define a directional color Riesz transform - based on the structure tensor formalism - to build the elliptical monogenic signal model. We end up with a color wavelet transform, whose implementation is easily done in parallel by directly using our Radon-based algorithm.

The proposed wavelet representation provides efficient multiscale representation of local color features in terms of amplitude, phase, orientation, and ellipse parameters. Elliptical monogenic features open up good applicative prospects by coherently embedding AM-FM and multiscale representations, as well as differential geometry, into a unified tool for color images. In particular, the color phase feature is a totally new concept.

The synthesis from coefficients is deeply studied, and two reconstruction methods are finally given so that most known wavelet-based processing application are possible.

In section V, the elliptical monogenic wavelets are successfully used to estimate local frequency from color images, by using the multiscale amplitude and phase data. This experi-

mentally validates the color phase concept in particular and the proposed filterbank in general, opening the way towards higher-level image analysis applications for future work.

For making this research reproducible, the source code that produces the figures and the experiments presented in this paper has been made available publicly [43].

Following our previous publication [44] on color extension of monogenic wavelets, the original contributions of the present article are the new filterbank algorithm, the synthesis method, the enriched color signal model (elliptical) and the local frequency estimation.

Notations :

2D coordinates: $\mathbf{x} = (x, y)$ $\boldsymbol{\omega} = (\omega_1, \omega_2) \in \mathbb{R}^2$
 Euclidean norm: $\|\mathbf{x}\| = \sqrt{x^2 + y^2}$
 Complex algebra: $z = \Re\{z\} + \mathbf{j}\Im\{z\} = |z|e^{\mathbf{j}\arg(z)}$
 Fourier transform: $\mathcal{F}\{f\}(\omega) = \int f(t)e^{-\mathbf{j}\omega t} dt$

II. THE GREYSCALE MONOGENIC REPRESENTATION

A. Generalization of the analytic signal

An *analytic signal* s_A is a multi-component signal associated to a real signal s to analyze. The definition is classical for scalar 1D signals [18], [50]: $s_A = s + \mathbf{j}\mathcal{H}s$, where the Hilbert transform is defined by:

$$\{\mathcal{H}s\}(t) = p.v. \int \frac{1}{\pi t} s(t - \tau) d\tau \xleftrightarrow{\mathcal{F}} -\mathbf{j} \operatorname{sgn}(\omega) \hat{s}(\omega) \quad (1)$$

The polar form of the 1D analytic signal provides an AM/FM representation of s with $|s_A|$ being the *amplitude envelope* and $\varphi = \arg(s_A)$ the *instantaneous phase*.

The growing interest on this tool within the image community is due to the possible interpretation of amplitude, phase and frequency in terms of local geometric *shape*.

Among the different existing 2D extensions [5], [16], [21], the *monogenic signal* [15] seems to be the most advanced definition. It is rotation invariant and generalizes the local phase concept intuitively through a directional model. Given a 2D real signal s , the associated monogenic signal s_M is 3-vector valued (instead of complex-valued in the 1D case) and must be taken in spherical coordinates:

$$s_M = \begin{bmatrix} s \\ \Re\{\mathcal{R}s\} \\ \Im\{\mathcal{R}s\} \end{bmatrix} = \begin{bmatrix} A \cos \varphi \\ A \sin \varphi \cos \theta \\ A \sin \varphi \sin \theta \end{bmatrix} \quad (2)$$

where $\mathcal{R}s$ is the complex-valued Riesz transform of s :

$$\{\mathcal{R}s\}(\mathbf{x}) = p.v. \int \frac{\tau_1 + \mathbf{j}\tau_2}{2\pi\|\boldsymbol{\tau}\|^3} s(\mathbf{x} - \boldsymbol{\tau}) d\boldsymbol{\tau} \xleftrightarrow{\mathcal{F}} \frac{\omega_2 - \mathbf{j}\omega_1}{\|\boldsymbol{\omega}\|} \hat{s}(\boldsymbol{\omega}) \quad (3)$$

(see also [29], [37]). From the analysis point of view, the monogenic signal is formed by the three following features:

$$\begin{aligned} \text{Amplitude: } & A = \sqrt{s^2 + |\mathcal{R}s|^2} \\ \text{Orientation: } & \theta = \arg\{\mathcal{R}s\} \in [-\pi; \pi[\\ \text{1D Phase: } & \varphi = \arg\{s + \mathbf{j}|\mathcal{R}s|\} \in [0; \pi] \end{aligned} \quad (4)$$

A direct link between the angles θ and φ and the geometric local structure of s has been shown [15]. The signal is so expressed like an “A-strong” 1D structure with orientation θ . The quantity φ is analogous to the 1D local phase and indicates

if the structure is rather a line or an edge. As a drawback, non-directional structures are not handled. Evolutions of the monogenic model [17], [51] taking into account more complex shapes are out of the scope of this paper.

From a *signal processing* viewpoint, such an AM/FM representation is accordingly well suited for narrow-band signals. It is natural to embed it in a multiscale transform.

B. Monogenic Wavelets: a review

Defining monogenic wavelet transforms is a recent topic. The main idea is to combine parallel filterbanks whose underlying wavelet functions form a Riesz triple.

This is how, in the 1D case, the “dual-tree” algorithm [40] combines good properties of Gabor/Morlet analytic wavelets [20] (invariance of amplitude) with the sparsity of bi-orthogonal filterbanks - achieving small redundancy and reversibility. However, its 2D extension is not monogenic and fails to extract the actual phase of the signal.

The *monogenic curvelet frames* [45] - in the spirit of the 2D Gabor transform [11] - use a set of *directional* subbands. Such a combination of the intrinsic orientation analysis of the monogenic framework with a multi-oriented scheme is out of the scope of this paper.

Continuous monogenic wavelets were proposed in [38], but does not allow any reconstruction of input signal from coefficients. Yet reversibility is necessary for use in applications such as denoising and compression, for which wavelets are known to be efficient. We will rather focus on wavelet *frames* and reversible filterbanks.

The Riesz-Laplace monogenic wavelets [48] are constructed upon a real-valued primary wavelet frame relying on a dyadic pyramid, whose mother wavelet ψ is defined in the context of polyharmonic splines. The *frame* is the extension of a *basis* obtained by adding a specific *subband regression* algorithm at the synthesis side [49]. The “Riesz part” is obtained from the Riesz transform of ψ , generating an other similar frame (complex-valued). The joint consideration of both transforms form monogenic subbands from which can be extracted the amplitude and phase for an overall redundancy of 4:1.

The choice of the authors is to provide the “minimally-redundant wavelet counterpart of Felsberg’s monogenic signal”. This transform is approximately isotropic and has a parameter γ controlling the number of vanishing moments.

Exactly isotropic tight monogenic wavelet frames can be found in [23], [47], where the wavelet is defined in the Fourier domain by a radial profile.

In these works, the algorithm uses FFT implementation of the filterbank. Note that what we call “monogenic filterbank” is in fact a *pair* of filterbanks - each one being reversible - forming a multiscale monogenic analysis as 3-valued coefficients. The synthesis of the input signal from these coefficients is now discussed.

C. The synthesis issue

We propose to clearly define how the input image should be coherently synthesized from the monogenic analysis. This is not trivial because both the *primary* and the *Riesz-part*

filterbanks have been separately defined as redundant and perfectly reversible (see *e.g.* [10], [23], [47], [48]). The *spatial* redundancy of each one is sometimes exploited through the concept of wavelet *frames* and/or ‘subband regression’ [49]. However, in the monogenic setting, an additional redundancy occurs by the use of *two* such filterbanks, each one being basically self-sufficient in terms of reconstruction.

For applications involving wavelet-based image synthesis, the modified monogenic features (A', φ', θ') can be easily converted back to Cartesian terms with equation (2) - providing separately the wavelet coefficients for the *primary* and the *Riesz-part* filterbanks. However, one does not know if it is better to reconstruct either from the *primary* part only, or from the *Riesz-part*, or from a combination of both.

The literature on monogenic representation almost exclusively relies on *analysis* applications: stereo analysis [30], motion estimation [35], segmentation [2], [27], texture classification [52]. A particular case is image *demodulation* [48], that produces a ‘reconstructed’ image. The result is in fact equal to the amplitude A , which is part of the analysis, so it does not need any synthesis from the filterbanks. In [10], which studies the Riesz-wavelets with examples of inverse problems, the authors present a special monogenic setting requiring to jointly use an additional isotropic filterbank. However, this setting is used for contour detection only, the other applications being based on the Riesz-wavelets.

In all these references, no global *monogenic* synthesis method is given. Exceptions can be found in [47] and [23].

In [47], the monogenic setting is a particular case of multi-component generalized Riesz-based steerable wavelet transform, where the *primary* part is realized with a 0-th order Riesz transform. As a special case of a general inversion scheme, the monogenic synthesis is defined as the sum of the two reconstructions (up to some constant factors). However, this is neither discussed nor applied. A more developed presentation of monogenic synthesis can be found in [23].

1) *Reconstruction from primary part*: In [23], after processing, it is proposed to update the i -th scale *primary* coefficients w_i according to the new amplitude and phase by $w_i(\mathbf{x}) = A_i(\mathbf{x}) \cos(\varphi_i(\mathbf{x}))$ and then to use the synthesis side of the *primary* filterbank only.

Now we give a practical example where this method is not satisfactory. Suppose that an image containing sharp edges is denoised by thresholding the amplitude of its monogenic representation. Whether there is an edge around \mathbf{x}_0 for example, w_i has a zero-crossing structure - since the primary wavelet acts like a Laplacian; whereas $\mathcal{R}w_i$ has a peak-like structure - since the Riesz-wavelet acts like a gradient. As a result, the amplitude A_i shows a peak centered on \mathbf{x}_0 . The problem is that thresholding will probably cancel out the whole neighborhood around \mathbf{x}_0 , leaving only one non-zero coefficient in the Riesz-part $\mathcal{R}w_i(\mathbf{x}_0)$, and nothing in the primary part ($w_i(\mathbf{x}_0) = 0$). When using only the primary part like in [23], the information necessary to retrieve the edge is lost. Actually, the isotropic wavelet is optimal for peaks and lines, while the Riesz wavelet is optimal for edges. Contrary to the method from [23], we believe that *both* filterbanks should be used so as to properly retrieve contours. More precisely, the synthesis

method must be defined such that *one* monogenic coefficient (A, φ, θ) is able to *modulate* the synthesis wavelet, according to its local amplitude, phase and orientation.

2) *Our proposition*: To this end, we propose a direct analogy with the Fourier theory. Recall that the Fourier transform synthesizes any real-valued 2D signal by a sum of plane waves of the form:

$$a_\omega(\mathbf{x}) = A \cos(\omega^\top \mathbf{x} + \varphi) \quad (5)$$

Orientation and frequency of a_ω are determined by the co-ordinate ω , while amplitude and phase are encoded by the Fourier coefficient $\hat{f}(\omega) = Ae^{j\varphi}$. We want to point out that the *modulation* of the plane wave by A and φ is mathematically achieved by the *summation* of odd and even parts:

$$\begin{aligned} a_\omega(\mathbf{x}) &= \Re\{\hat{f}(\omega)\} \cos(\omega^\top \mathbf{x}) - \Im\{\hat{f}(\omega)\} \sin(\omega^\top \mathbf{x}) \\ &= A \cos(\varphi) \cos(\omega^\top \mathbf{x}) - A \sin(\varphi) \sin(\omega^\top \mathbf{x}) \end{aligned} \quad (6)$$

In the monogenic wavelet case, the *primary* filterbank computes the “real part” of the monogenic coefficient, and is responsible for the “even” (isotropic) part of the monogenic wavelet. On the other hand, the *Riesz* filterbank computes the ‘imaginary part’ (as two components) and controls the orientation of the “odd” (directional) part of the wavelet.

This analogy justifies the synthesis method to consist in a *summation* of reconstructed signals from both the *primary* and the *Riesz* filterbanks. This way, the implicit linear combination of even isotropic wavelets and odd directional wavelets will actually produce the desired shape: a wavelet whose amplitude, phase and orientation are encoded in the monogenic coefficient (A, φ, θ) . Note that a strictly analogous summation occurs at the core of the *Riesz* filterbank as well, to *steer* the directional wavelet by a linear combination of its x and y components (although it may be “hidden” by the use of complex algebra like in eq. (3)).

In the 1D case, most works on *analytic* wavelets - including the dual-tree filterbank - intrinsically define the synthesis within the complex algebra, by the same odd/even summation we just proposed. This needed to be clarified in the monogenic case for which complex numbers cannot be used the same way. Note that our proposition turns out to be similar to the definition from [47].

The review of existing monogenic wavelet transforms reveals that the literature is still at its very first steps. The filterbank design in [47] uses all the degrees of freedom for the tuning of the analysis filter to be as isotropic as possible, ending up with highly anisotropic synthesis filters that could not be used for proper monogenic synthesis. We propose to explore an alternative strategy in the next section.

III. RADON TRANSFORM BASED MONOGENIC FILTERBANK

We propose in this section an innovative monogenic wavelet transform algorithm based on the following motivations. The monogenic framework simplifies to a 1D definition in the Radon domain [4]; discrete Radon transform algorithms are available in the literature [6]; and the design of analytic 1D wavelets is well known compared to 2D. In particular, the dual-tree algorithm [40] is an interesting implementation of

mildly redundant analytic wavelets. We propose to investigate a Radon based monogenic filterbank design.

In the first part we present the theoretical relation between Radon transform, analytic and monogenic wavelets. Next we present the numerical algorithm for the Radon transform, and then for the monogenic transform. Note that this section only deals with greyscale images, color will be treated right after.

A. Monogenic wavelets in Radon domain

Given any function $s(x, y)$, its Radon transform is defined by the projection along direction ρ as :

$$s_\rho(t) = \int_{\mathbb{R}} s(\tau \sin \rho + t \cos \rho, -\tau \cos \rho + t \sin \rho) d\tau \quad (7)$$

The Radon domain is known to connect the 1D and 2D worlds, as suggested by the *Fourier slice theorem*:

$$\widehat{s}_\rho(\omega) = \widehat{s}(\omega \cos \rho, \omega \sin \rho) \quad (8)$$

It shows that 1D Fourier transform in the Radon domain is like 2D Fourier transform in the space domain. This property is used in image reconstruction from projection methods. The *convolution property* stems from it. Given two signals $s(x, y)$, $h(x, y)$ and their Radon transforms $s_\theta(t)$, $h_\theta(t)$, we have:

$$(s_\theta * h_\theta)(t) = (s * h)_\theta(t) \quad (9)$$

where $*$ denotes 1D convolution and $**$ denotes 2D convolution. This is convenient to perform isotropic filtering. By applying a 1D symmetric filter h_1 on all Radon projections of a 2D signal, an isotropic 2D filtering is actually performed:

$$s_\theta(t) \leftarrow (s_\theta * h_1)(t) \Leftrightarrow \widehat{s}(\omega) \leftarrow \widehat{s}(\omega) \widehat{h}_1(\|\omega\|) \quad (10)$$

with $h_1(t) = h_1(|t|)$

The resulting 2D filter has a radial frequency response identical to $\widehat{h}_1(\omega)^1$. So Radon domain filtering allows a direct control of the frequency response from 1D filters.

In addition, the Riesz transform turns out to be simplified in the Radon domain, as a weighted Hilbert transform [4]:

$$\{\mathcal{R}s\}_\rho(t) = e^{j\rho} \{\mathcal{H}s_\rho\}(t) \quad (11)$$

Considering that a monogenic wavelet transform essentially consists in some isotropic band-pass filtering combined with the Riesz transform, the properties given above strongly suggest that it can reduce to using 1D band-pass filters and the Hilbert transform *i.e.* an *analytic wavelet* transform. This way we can focus on a simpler 1D filterbank design, while letting the isotropy issue be handled by the Radon transform.

Given any symmetric 1D analytic wavelet:

$$\psi_A(t) = \psi(t) + j\mathcal{H}\psi(t) \quad \text{with} \quad \psi(t) = \psi(|t|) \quad (12)$$

let the monogenic wavelet ψ_M be defined by:

$$\psi_M(\mathbf{x}) = \begin{bmatrix} \psi_{iso}(\mathbf{x}) \\ \Re\{\mathcal{R}\psi_{iso}(\mathbf{x})\} \\ \Im\{\mathcal{R}\psi_{iso}(\mathbf{x})\} \end{bmatrix} \xrightarrow{\text{Radon}} \begin{bmatrix} \psi(t) \\ \cos(\theta)\mathcal{H}\psi(t) \\ \sin(\theta)\mathcal{H}\psi(t) \end{bmatrix} \quad (13)$$

with the following radial frequency response:

$$\widehat{\psi_{iso}}(\omega) = \widehat{\psi}(\|\omega\|) \quad (14)$$

From this construction, let us derive the numerical algorithm.

¹There is no such closed-form relation in the space domain.

B. Discrete Radon transform

The discretization of the Radon transform is difficult to achieve. Most methods in the literature have been devised for computerized tomography or to approximate the continuous formula. None of them, however, were specifically designed to be invertible transforms for discrete images. In previous work [6] we have developed a numerical Radon transform with exact reconstruction based on discrete analytical geometry. This approach has straightforward n -D extension and computational simplicity. The present filterbank is based on it.

Thanks to the Fourier slice theorem (eq. (8)), we know that Radon projections can be obtained from inverse 1D FFT's applied on all *slices* of the 2D FFT of the input signal.

The main idea is to treat each direction θ by a discrete line, having central symmetry and forming a good approximation of the corresponding Euclidean straight line. We want the line to cross both the origin and a given "target point" $\mathbf{p} = (p_1, p_2) \in \mathbb{Z}^2$, tied to orientation $\theta = \arg(p_1 + jp_2)$. The work in [6] proposes to write this line L as a set of discrete 2D coordinates, according to the following definition:

$$L_{\mathbf{p}}^\tau = \left\{ (k_1, k_2) \in \mathbb{Z}^2 \mid |p_2 k_1 - p_1 k_2| \leq \frac{\tau}{2} \right\} \quad (15)$$

where τ is the *arithmetical thickness* parameter. Thickness provides control over the connectivity of the line, and will determine the redundancy factor of our Radon transform. In this paper, we will work with 8-connected "naive" lines, by setting $\tau = \max(|p_1|, |p_2|)$. We actually tried "thicker" lines ("pythagorean", "supercover" [6]) but besides increased redundancy, we experimentally noticed no significant difference.

The algorithm generating the discrete path works iteratively with 1-pixel shifts from \mathbf{p} to $(0, 0)$. At each step, the shift is selected among the 8 possible directions from eq. (15). Path from $(0, 0)$ to $-\mathbf{p}$ is then processed by central symmetry. This way, we ensure that coordinates are well ordered with respect to their radial frequency $\|\omega\|$.

We can now extract all Fourier slices to perform Radon transform. Due to the desired symmetry, the dimension must be odd. We consider an input signal of size $[2K_1 + 1] \times [2K_2 + 1]$, and its 2D FFT $\widehat{s}(\mathbf{k})$ with $\mathbf{k} \in [-K_1..K_1] \times [-K_2..K_2]$.

To ensure exact reconstruction, the set of lines must cover the whole lattice. Thus the set of orientations is defined so that every boundary point

$$\mathbf{p} \in [-K_1..K_1] \times [-K_2..K_2] \text{ with } |p_1| = K_1 \text{ or } |p_2| = K_2 \quad (16)$$

together with its symmetric $-\mathbf{p}$, is the target point for one line $L_{\mathbf{p}}^\tau$. We end up with $2K_1 + 2K_2 + 1$ Fourier slices covering the whole 2D FFT and generating a discretization of θ in $[0; \pi[$.

After line extraction, a 1D inverse FFT provides the final Radon projection.

Finally, the reconstruction has to be defined. After having transformed the Radon projections back in the 1D Fourier domain, the 1D Fourier coefficients must be set back into the 2D spectrum. The same discrete 2D paths used at the extraction stage are used again to match Radon domain coordinate to Fourier domain. A choice has to be made since most 2D Fourier samples are represented through *several* Fourier slices (redundancy). The solution proposed in [6] is

a pointwise averaging of the multiple contributions from the different slices.

Although this method is satisfactory in the general case, it is not adapted for the monogenic filterbank because the Riesz transform needs to be implemented by a function depending on θ . In this case, for *one* Fourier sample, most of the contributions are tied to values of θ that do not fit well the angular coordinate of the Fourier sample, resulting in wrong values. Therefore the contributions must not be averaged.

We propose an alternative reconstruction method in this paper. Every Fourier sample $\hat{s}(\mathbf{k})$ will be extracted only once, from the *best* possible slice $\hat{s}_\theta(\mathbf{k})$. This slice will be chosen to have θ value as close as possible to the exact orientation of the Fourier sample: $\theta_{ref} = \arg(k_1 + \mathbf{j}k_2)$. Note that when the Radon transform is not modified, both reconstruction methods are identical, but the present one is required for the monogenic filterbank. The last inversion step is simply the inverse 2D FFT. The whole reconstruction algorithm is numerically exact up to machine's precision.

C. Discrete Monogenic wavelets

Now we have to choose a 1D analytic wavelet transform to be applied in the Radon domain. The first idea is to use the well established *dual-tree* wavelet transform [40], that combines good approximation of Hilbert pairs of filters, and efficient filterbank implementation. However, this method would downsample the Radon projections at each scale, which would produce half-length projections without reducing the number of orientations. The discrete Radon algorithm is not compatible with such resizing of the data, so the inverse Radon transform could not be used. We will then rather focus on *undecimated* wavelet transforms.

This choice presents the advantage of releasing strong constraints on filter design, and potentially improving the signal analysis. This also implies increased redundancy, non-orthogonal decomposition, and that we found no existing work considering a Hilbert pair of *undecimated* filterbanks. Literature on analytic wavelets is either focused on subsampled IIR filterbanks (Phaselets [19], B-Splines [8]), or on continuous definitions, like the generalized Morse wavelets [31] and the well known Gabor/Morlet wavelets [20]. A known decimated method may be adapted by simply removing the sub-sampling steps but it would limit us to less optimal filters.

Therefore, we propose the customized filterbank design outlined on figure 1.




Fig. 1. Undecimated analytic 1D filterbank iterated for $i \in \{1..J\}$. Input is $s = s_0$, wavelet coefficients are $w_i + \mathbf{j}w_{h_i}$ and the approximation is s_J .

The perfect reconstruction of input signal is ensured by:

$$s_{i-1} = s'_{i-1} \Leftrightarrow \hat{f}_i \hat{f}_i + (\hat{g}_i \hat{g}_i + \hat{g}_{h_i} \hat{g}_{h_i})/2 = 1 \quad (17)$$

To achieve the Radon domain monogenic analysis explained in section III-A, the primary low-pass (resp. high-pass) filters f_i and \tilde{f}_i (resp. g_i and \tilde{g}_i) must be symmetric in the space domain, and the Hilbert part (anti-symmetric) must verify $\tilde{g}_{h_i} = \mathcal{H}g_{h_i}$. For convenient interpretation of wavelet coefficients, we choose to use the same filters on the analysis and synthesis stages. With these constraints, the perfect reconstruction condition (17) reduces to $|\hat{f}_i|^2 + |\hat{g}_i|^2 = 1$. We propose the following definition of the filters in the Fourier domain, based on a Gaussian low-pass filter:

$$\begin{aligned} \tilde{f}_i(t) = f_i(t) &\xleftrightarrow{\mathcal{F}} \hat{f}_i(\omega) = \exp(-\omega^2 \sigma_i^2/2) \\ \tilde{g}_i(t) = g_i(t) &\xleftrightarrow{\mathcal{F}} \hat{g}_i(\omega) = \sqrt{1 - \hat{f}_i(\omega)^2} \\ \tilde{g}_{h_i}(t) = -g_{h_i}(t) &\xleftrightarrow{\mathcal{F}} -\mathbf{j} \operatorname{sgn}(\omega) \hat{g}_i(\omega) \end{aligned} \quad (18)$$

Like in [8], we opt for Fourier domain definition and FFT implementation, allowing for high order IIR filters. In this case the Hilbert transform is approximated by the truncation $\omega \in [-\pi; \pi]$. Note that odd-length signals will be preferred (this is always verified with our symmetric discrete Radon projections) to avoid singularity at Nyquist frequency $\omega = \pm\pi$.

We choose the dyadic dilation rule $\sigma_i = 2^{i-2}$ so that $f_i(t) = \frac{1}{2} f_{i-1}(\frac{t}{2})$ (*idem* for g_i and g_{h_i}). Subbands are normalized by the energy of equivalent filters at each scale. We refer the reader to the source code for further details [43].

Now we can define the final monogenic filterbank. The combination of discrete Radon transform, analytic filterbank and θ weighting according to eq. (11) is outlined on the figure 2. Note that FFT-based 1D filtering allows to work directly on Fourier slices and skip all the 1D FFTs.

The analysis stage basically consists in Radon domain analytic wavelet transform followed by $\cos(\theta)$ - $\sin(\theta)$ weighting and inverse Radon transform.

At the synthesis stage, back in the Radon domain, the two Riesz components are multiplied again by $\cos(\theta)$ - $\sin(\theta)$ and summed so that we get back the imaginary part of the analytic wavelet transform. With respect to discussion at section II-C, where we suggest to sum monogenic components *after* filtering, we must point out that the orientation information of each monogenic coefficient will here still properly modulate the orientation of the synthesis wavelet. The reason is that the *steering* will be actually done in the last inverse Radon transform. It is shown on figure 3 that the values of the monogenic coefficient actually modulate the wavelet shape in terms of phase and orientation. It is also visible that our




Fig. 3. Wavelets reconstructed from one unique monogenic coefficient with different values of φ and θ ($A = 1$, $i = 6$).

Radon-based monogenic wavelet is not exactly isotropic. This is due to the Fourier slices having varying Nyquist frequencies from π at x - and y - axes to $\pi/\sqrt{2}$ at diagonals. The 1D filters are applied regardless of this bias, which intrinsically *dilates*




Fig. 2. Radon based monogenic 2D filterbank (iterated for $i \in \{1..J\}$). Input image is $s(\mathbf{x}) = s_0(\mathbf{x})$. Monogenic wavelet coefficients at i -th scale are $(w_i(\mathbf{x}), w_{x,i}(\mathbf{x}), w_{y,i}(\mathbf{x}))$. The low-frequency approximation is $s_J(\mathbf{x})$.

their frequency response around diagonals, in the 2D FFT, resulting in a slightly cross-shaped wavelet. This points out the tight link between the Radon algorithm and isotropy - which has always been a crucial point in the discrete world. Improvements may be possible but out of the scope of this paper, while the present result is perceptually coherent with respect to phase and orientation, which is satisfactory.

D. Discussion

Note that our filterbank is close to the algorithm of *complex ridgelets* [9]. In the ridgelet analysis, the goal is to construct oriented wavelets. Its complex extension is aimed at gaining translation invariance, by using analytic 1D wavelets in the Radon domain, resulting in Radon domain wavelet coefficients. In the present work, the constructed wavelets are the multiscale extension of an AM-FM model. Their definition requires the subsequent use of the inverse Radon transform, to set the coefficients back in the space domain. This particularity is necessary for the interpretation of the data as a monogenic analysis. For example, the 2D phase concept does not exist with the *complex ridgelets*, whereas it is natural in the monogenic case.

Compared to the existing monogenic pyramids [23], [48], our undecimated algorithm benefits from theoretical simplicity, and especially from a clear synthesis method, achieving coherent modulation of the wavelets. Let us recall that the representation is translation and rotation invariant. The n -D extension can be straightforwardly done by extending the discrete line concept like in [22] and considering $n - 1$ orientation angles.

Now our greyscale monogenic filterbank is defined, let us extend the whole concept to the tri-valued bi-dimensional signals that are color images.

IV. ELLIPTICAL MODEL FOR COLOR MONOGENIC REPRESENTATION

In a previous work [41], following the lead from [12], we proposed a color monogenic analysis based on Cauchy-Riemann equations in the Clifford algebra. However, from the “image processing” point of view, this construction does not provide a clear physical interpretation of the data.

We have next carried out an alternative construction [44] based on differential geometry to handle the *color orientation* concept. This time, the amplitude and phase were clearly identified, and a “color axis” concept emerged. Yet, the synthesis

method has to be improved, since only the primary part of the filterbank was used (see discussion in II-C).

This paper goes one step further toward color signal processing, starting from the fundamental concept of *Fourier atom* to construct a full color signal model. The interpretation of the features is improved thanks to a true vector approach of the basic signal processing concepts, and the synthesis method in terms of modulated color wavelets is achieved.

Note that we would think premature to consider advanced color spaces at this stage, so we restrict ourselves to Euclidean ones - like the usual RGB space. The study of perceptually meaningful color spaces and of the dependencies between colour channels is an important prospect.

A. Fourier atom

Redefining *Fourier analysis* for multidimensional signals $s : \mathbb{R}^p \mapsto \mathbb{R}^q$ is an open mathematical topic, usually involving quaternions or Clifford algebra (see [1], [24] for example). From the physical point of view, this literature does not yet provide sufficiently intuitive interpretation of Fourier coefficients. We propose a pragmatic approach.

As a construction basis, it seems reasonable to consider separate Fourier analyzes on R, G and B channels. The color signal is then expressed like a sum of sine waves that we will call *Fourier atoms*. In the 1D case, we define the *color Fourier atom* $\mathbf{a}(t)$ as:

$$\mathbf{a}(t) = \begin{bmatrix} a^R(t) \\ a^G(t) \\ a^B(t) \end{bmatrix} = \begin{bmatrix} A^R \cos(\omega t + \varphi^R) \\ A^G \cos(\omega t + \varphi^G) \\ A^B \cos(\omega t + \varphi^B) \end{bmatrix} \quad (19)$$

For any fixed frequency ω , the parameters of $\mathbf{a}(t)$ are given by three *Fourier coefficients*:

$$\hat{\mathbf{a}} = \left[A^R e^{j\varphi^R} \ A^G e^{j\varphi^G} \ A^B e^{j\varphi^B} \right]^T \quad (20)$$

An example of color Fourier atom is given on figure 4a.

Here is our starting point for a non-marginal analysis of color signals. As is, those three amplitudes and three phases do not intuitively convey the oscillation’s behavior in the “color world”. We may prefer some unified amplitude and phase concepts, as well as color features. The atom $\mathbf{a}(t)$ must then be studied as a point-trajectory within the 3D color space - a “color path” - as illustrated on figure 4b.

It can be shown [32] that $\mathbf{a}(t)$ always draws an *ellipse* in the color space. We propose to study the ellipse parameters to physically handle the atom’s behavior in a non-marginal way.

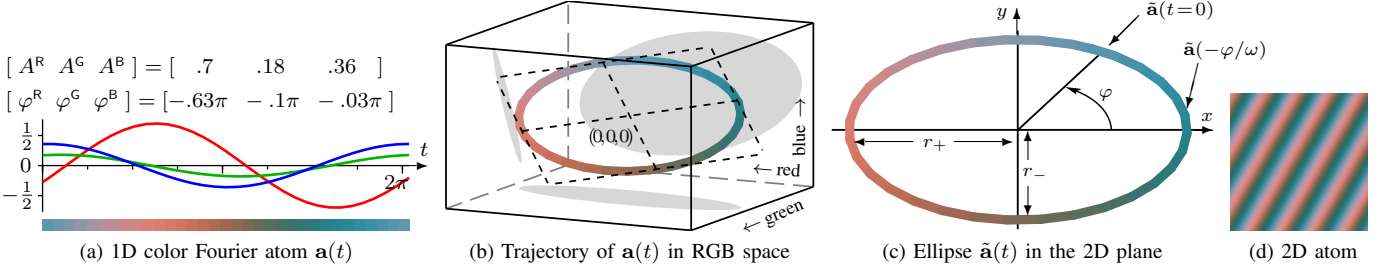


Fig. 4. Color Fourier atom. Colors of $\mathbf{a}(t)$ are generated by centering around middle grey. $\tilde{\mathbf{a}}(t)$ is the rotated version of $\mathbf{a}(t)$ by $(\alpha_1, \alpha_2, \alpha_3)$.

B. The elliptical model

New “Fourier-like features” can be carried out by studying the color Fourier atom as an ellipse. First, the *amplitude* concept can intuitively be extended to the ellipse’s *size*. By considering the *apogees* as sinusoid’s *maxima*, the *phase* can then be defined as the time-relative position of the apogees, with respect to $\mathbf{a}(0)$, as illustrated in figure 4c.

The apogees can be quickly identified analytically from the analysis of $\|\mathbf{a}(t)\|$, whose minimum and maximum values must be equal to the least radius r_- and the greatest radius r_+ (see figure 4c). It can be verified that:

$$\|\mathbf{a}(t)\|^2 = \frac{A^2 - |\hat{\mathbf{a}}^T \hat{\mathbf{a}}|}{2} + |\hat{\mathbf{a}}^T \hat{\mathbf{a}}| \cos^2 \left(\omega t + \frac{\arg(\hat{\mathbf{a}}^T \hat{\mathbf{a}})}{2} \right) \quad (21)$$

where the amplitude A is defined by

$$A = \sqrt{A^R^2 + A^G^2 + A^B^2} \quad (\text{Amplitude}) \quad (22)$$

This writing explicitly provides the extrema of $\|\mathbf{a}(t)\|^2$:

$$r_-^2 = (A^2 - |\hat{\mathbf{a}}^T \hat{\mathbf{a}}|)/2 \quad r_+^2 = (A^2 + |\hat{\mathbf{a}}^T \hat{\mathbf{a}}|)/2 \quad (23)$$

By defining the phase as

$$\varphi = \arg(\hat{\mathbf{a}}^T \hat{\mathbf{a}})/2 \quad (\text{Phase}) \quad (24)$$

the point $\mathbf{a}(-\varphi/\omega)$ is one of the two apogees, as expected, since its norm equals r_+ . The classical scalar case is retrieved with $\hat{\mathbf{a}} = A e^{j\varphi}$, $r_- = 0$, $r_+ = A$.

The quantity $\hat{\mathbf{a}}^T \hat{\mathbf{a}}$ can be viewed as a combination of the three Fourier coefficients, so that φ results as a weighted average of the separate phases. The modulus $|\hat{\mathbf{a}}^T \hat{\mathbf{a}}|$ is an ellipse feature in itself, when expressed through the “linearity”:

$$\lambda = \frac{|\hat{\mathbf{a}}^T \hat{\mathbf{a}}|}{A^2} = \frac{r_+^2 - r_-^2}{r_+^2 + r_-^2} \in [0; 1] \quad (\text{Linearity}) \quad (25)$$

The linearity λ is as close to 1 as the three separate phases are similar (modulo π). In this case the ellipse has a *thin* shape, and the Fourier atom oscillates between two main colors by crossing the origin. When λ is small, the ellipse is *round* and draws a richer color path.

Now the ellipse’s position must be identified. We propose to define the apogee by:

$$\mathbf{a}_+ = \mathbf{a}\left(\frac{-\varphi}{\omega}\right) = \begin{bmatrix} A^R \cos(\varphi^R - \varphi) \\ A^G \cos(\varphi^G - \varphi) \\ A^B \cos(\varphi^B - \varphi) \end{bmatrix} = r_+ \begin{bmatrix} \cos \alpha_1 \cos \alpha_2 \\ \sin \alpha_1 \cos \alpha_2 \\ \sin \alpha_2 \end{bmatrix} \quad (26)$$

We consider the direction of \mathbf{a}_+ as the main *color axis*, encoded by α_1 and α_2 . The angles are computed by appropriate arc-tangent operations. Then the perigee is defined by:

$$\mathbf{a}_- = \mathbf{a}\left(\frac{-\varphi - \frac{\pi}{2}}{\omega}\right) = \begin{bmatrix} A^R \sin(\varphi^R - \varphi) \\ A^G \sin(\varphi^G - \varphi) \\ A^B \sin(\varphi^B - \varphi) \end{bmatrix} = R_{\alpha_1} R_{\alpha_2} \begin{bmatrix} 0 \\ r_- \cos \alpha_3 \\ r_- \sin \alpha_3 \end{bmatrix} \quad (27)$$

with the rotation matrices:

$$R_{\alpha_1} = \begin{bmatrix} \cos \alpha_1 & -\sin \alpha_1 & 0 \\ \sin \alpha_1 & \cos \alpha_1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad R_{\alpha_2} = \begin{bmatrix} \cos \alpha_2 & 0 & -\sin \alpha_2 \\ 0 & 1 & 0 \\ \sin \alpha_2 & 0 & \cos \alpha_2 \end{bmatrix} \quad (28)$$

The perigee’s direction is encoded by the last parameter α_3 . Its computation requires to use the rotation by $R_{\alpha_2}^T R_{\alpha_1}^T$ that sets \mathbf{a}_+ on the x -axis and \mathbf{a}_- in the y - z plane. α_3 is the angle between the y and z coordinates of the rotated perigee $R_{\alpha_2}^T R_{\alpha_1}^T \mathbf{a}_-$.

The back conversion from $(A, \lambda, \varphi, \alpha_1, \alpha_2, \alpha_3)$ first consists in getting back r_+ and r_- :

$$r_+ = A\sqrt{(1+\lambda)/2} \quad r_- = A\sqrt{(1-\lambda)/2} \quad (29)$$

and rotating them to retrieve \mathbf{a}_+ and \mathbf{a}_- from equations (26)-(27). The original separate Fourier coefficients are equal to:

$$\hat{\mathbf{a}} = (\mathbf{a}_+ + \mathbf{j} \cdot \mathbf{a}_-) e^{j\varphi} \quad (30)$$

Let us finally give angle wrapping details. At this stage, two equivalent ellipse feature sets are possible for one given Fourier atom, according to the equivalence:

$$\begin{pmatrix} A & \lambda & \varphi & \alpha_1 & \alpha_2 & \alpha_3 \end{pmatrix} \equiv \begin{pmatrix} A & \lambda & \varphi + \pi & \alpha_1 + \pi & -\alpha_2 & -\alpha_3 \end{pmatrix} \quad (31)$$

In order to have a unique encoding, we propose to do

$$\text{if } \cos(\varphi^R - \varphi) < 0 \text{ then } \varphi \leftarrow \varphi + \pi \quad (32)$$

before the computation of \mathbf{a}_+ and \mathbf{a}_- . This corresponds to arbitrarily choosing the apogee being in the half of the color space where the red channel is positive. The advantage is that two oscillations having similar color properties will be quickly identified as having similar values of α_i , even if they are out of phase each other. We finally have $(\alpha_1, \alpha_2) \in [-\pi/2; \pi/2]^2$ and $(\varphi, \alpha_3) \in [-\pi; \pi]^2$.

The proposed ellipse parameters are strongly inspired from [32], where the elliptical model has been studied to analyze trajectories of particles in the field of seismology, and our definitions of A , λ and φ are in fact identical.

However, the approach is different with respect to the encoding of the ellipse's position as well as the calculus. In [32], the method relies on the computation of a *normal vector*, orthogonal to the plane containing $\mathbf{a}(t)$. Then two specific rotation angles ("azimuth" and "zenith") allow to study the ellipse in the x - y plane, so that r_+ , r_- , φ and a last "precession" angle are defined. We believe that when working in a color space, the choice of these particular angles is not justified. In contrast, our definition provides a notion of *color axis* which is analogous to the so-called *color axis* proposed in [44], as pointing toward two opposite main colors involved in the oscillation. When $\lambda = 1$, the *color axis* from present paper and this from [44] are equal. When the oscillation is more complex ($\lambda < 1$), the *color axis* of present paper encodes a still meaningful *main color direction*, while this from [44] becomes less relevant because it is based on a simpler model. It should also be noted that the ellipse parameters computation from [32] is unstable when $\lambda \approx 1$. The reason is that the normal vector is ill-defined when $\lambda = 1$. In contrast, our method relying on the computation of the apogee and perigee, is numerically more efficient. The only natural consequence when $\lambda = 1$ is that α_3 is irrelevant because the perigee is null ($r_- = 0$). Finally, we believe that our approach provides a simpler computation of the conversion between the Fourier coefficients and the ellipse features.

Thanks to the elliptical model, the three Fourier analyzes can be *converted* into intuitive amplitude and phase completed by 4 colorimetric features λ , α_1 , α_2 , α_3 in a non-marginal way. We can now define the corresponding color *analytic signal*.

C. Color analytic signal

The analytic signal (see II-A) makes a pointwise estimation of amplitude and phase, under the hypothesis that the signal to analyze locally resembles a *Fourier atom* $a(t) = A \cos(\omega t + \varphi)$. The extracted local information has the form of a *Fourier coefficient* $Ae^{j\varphi}$. In the case of an input 3-valued signal

$$\mathbf{s}(t) = [s^R(t) \ s^G(t) \ s^B(t)]^T \quad (33)$$

we define the color extension by the use of independent Hilbert transforms on R, G and B channels, followed by the conversion to elliptical parameters. The color analytic signal is then

$$\begin{aligned} \mathbf{s}_A(t) &= [s^R(t) \ s^G(t) \ s^B(t) \ \mathcal{H}s^R(t) \ \mathcal{H}s^G(t) \ \mathcal{H}s^B(t)]^T \\ &\Rightarrow [A(t) \ \varphi(t) \ \lambda(t) \ \alpha_1(t) \ \alpha_2(t) \ \alpha_3(t)]^T \end{aligned} \quad (34)$$

Now the ellipse parameters are *time varying*. We plot an example of oscillating color 1D signal with varying properties, together with the estimated amplitude and phase on figure 5. The synthetic color path is shown in the RGB space as a 3d




Fig. 5. Color 1D signal $\mathbf{s}(t)$ and its estimated amplitude A and phase φ .

plot on the left-hand side, as three red, green and blue 1d plots on the upper graph, and as a pixel map on the middle. Colors are generated by centering the values of $\mathbf{s}(t)$ in the RGB space around the middle grey.

The estimated amplitude $A(t)$ is plotted as a dotted curve in the upper graph. It visually well conveys the local energy of the oscillation. The estimated instantaneous phase $\varphi(t)$ is plotted in the lower graph. Because the frequency of input signal is supposed to evolve smoothly, the phase is expected to be locally linear, which is confirmed here. The slope of φ is locally according with the frequency of \mathbf{s} . Note the singularity of φ where amplitude equals 0, which is usual with polar representations, and has no consequence as long as A is used as some *significance* information.

The linearity $\lambda(t)$ (not plotted) also gives the right values:

- $\lambda \approx 0.5$ for the first oscillation (blue-yellow), whose trajectory in the RGB space is visibly "round", as can be seen on the left-hand 3D plot;
- $\lambda \approx 1$ for the second oscillation (pink-green), whose trajectory is far much linear.

The other colorimetric parameters α_1 , α_2 and α_3 are constant in the first part, and slowly varying in the second part, according to the color patterns visible on the pixel-map. The interpretation of (α_1, α_2) is similar to this of the *color axis* proposed in [44]. This axis points towards the two opposite main colors involved in the oscillation. More generally, the interpretation of color angles is still an open issue, and we expect to get more intuition about it by working on the 2D extension.

The reader should notice that the phase estimation is well behaved through the four quadrants, and not altered by the variation of the other parameters. The quality of the estimation of φ is a direct benefit from the use of elliptical modeling. With respect to the phase estimation from our previous work [44], this phase data is numerically more stable.

Let us now combine this new formalism with the Riesz transform to carry out the elliptical monogenic representation.

D. Color Riesz transform

We have studied above how to analyse a color signal by combining it with a phase-shifted version of its color channels (its Hilbert transform) in the 1D case, in perfect analogy with the analytic signal formalism. The extension to images must consider the *monogenic* formalism, according to which the signal is combined with its Riesz transform ("RT"). For the extension to be well driven, we have to consider the *directional* RT, defined for any scalar 2D signal $\mathbf{s}(\mathbf{x})$ by:

$$\mathcal{R}_\theta \mathbf{s}(\mathbf{x}) = \cos(\theta) \Re\{\mathcal{R}\mathbf{s}(\mathbf{x})\} + \sin(\theta) \Im\{\mathcal{R}\mathbf{s}(\mathbf{x})\} \quad (35)$$

$$= |\mathcal{R}\mathbf{s}(\mathbf{x})| \cos(\theta - \arg(\mathcal{R}\mathbf{s}(\mathbf{x}))) \quad (36)$$

Here the RT acts like a pair of *steerable* filters, of which a linear combination can provide any rotation of them. The RT is retrieved by combining them along x - and y - axes² $\mathcal{R}\mathbf{s} = \mathcal{R}_0\mathbf{s} + \mathbf{j}\mathcal{R}_{\frac{\pi}{2}}\mathbf{s}$. The *directional* RT is equivalent to a

²The exact same analysis is well known about the *gradient* operator, when studied as the solution of a maximum variation problem [26].

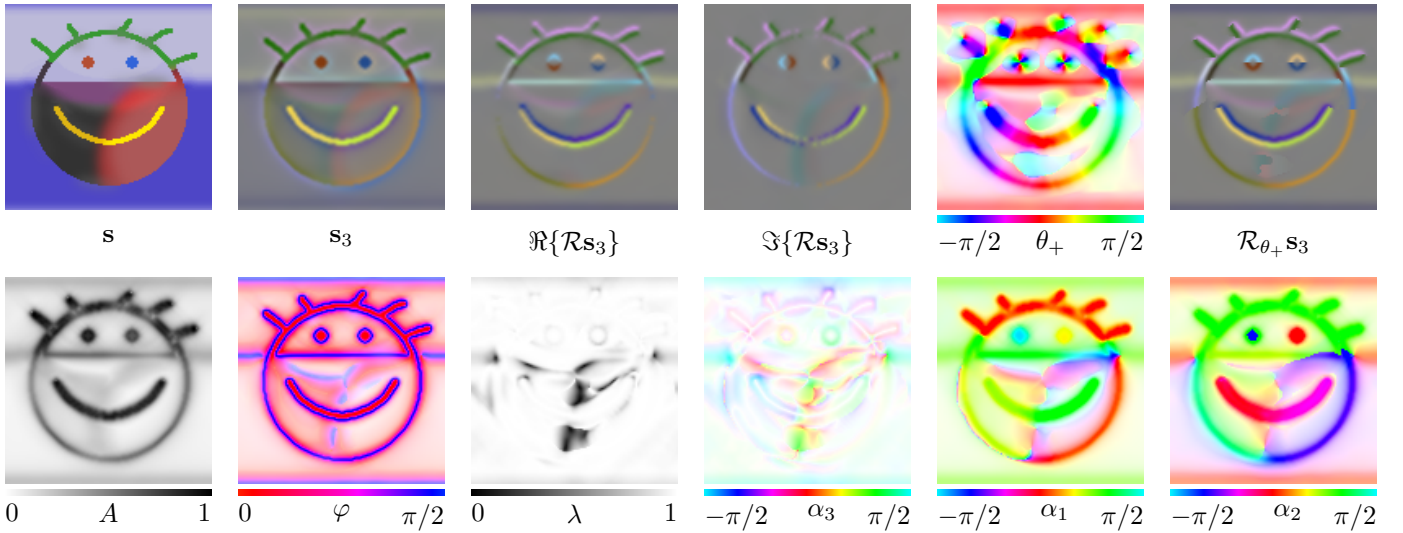


Fig. 6. Color monogenic wavelet transform at third scale (Subband s_3). Oscillating color signals are centered and normalized in the RGB cube. Amplitude A and linearity λ are shown as normalized greyscale bitmaps. Angles θ_+ , φ , α_1 , α_2 and α_3 are displayed as *hue* in the HSV color space, while *saturation* is controlled by A so as to whiten irrelevant values (The linearity is also used for α_3). The angles φ and α_3 are wrapped. The smoothing kernel h for calculus of θ_+ is a Gaussian filter with $\sigma = 2$.

Hilbert transform in the case of strictly directional 2D signals (a.k.a. “simple signal” or “intrinsically 1D”):

$$s_{\theta_1}(\mathbf{x}) = s_1(\cos(\theta_1)x_1 + \sin(\theta_1)x_2) \quad (37)$$

$$\Rightarrow \mathcal{R}_{\theta_1} s_{\theta_1}(\mathbf{x}) = \{\mathcal{H}s_1\}(\cos(\theta_1)x_1 + \sin(\theta_1)x_2) \quad (38)$$

This expresses well that the monogenic signal is like the 1D analytic signal, locally steered toward its main direction. We can now feed the elliptical model with the *directional* RT as phase-shifted versions of s . The last step is to choose a direction, toward which the RT will be steered.

In the grayscale case, the direction is naturally $\arg(\mathcal{R}s)$. Here, we propose to use the one that maximizes the *color variation*. This mathematical problem is solved in terms of Euclidean color distances, by the well known color gradient from Di Zenzo [13], [39]. Thanks to a strong theoretical link between the RT and the gradient operator [28], [48], we have already proposed in [44] a Riesz-based color structure tensor, analogous to Di Zenzo’s gradient. It is defined by:

$$T_R(s) = T_R^R + T_R^G + T_R^B = \begin{bmatrix} T_{11} & T_{12} \\ T_{12} & T_{22} \end{bmatrix} \quad (39)$$

$$\text{with } T_R^c = h * \begin{bmatrix} \Re\{\mathcal{R}s^c\}^2 & \Re\{\mathcal{R}s^c\}\Im\{\mathcal{R}s^c\} \\ \Re\{\mathcal{R}s^c\}\Im\{\mathcal{R}s^c\} & \Im\{\mathcal{R}s^c\}^2 \end{bmatrix} \quad (40)$$

and h is a smoothing filter. The main orientation of color local structures is equal to the orientation of the eigenvector tied to the largest eigenvalue of T_R :

$$\theta_+ = \arg(T_{11} - T_{22} + 2jT_{12})/2 \quad (41)$$

This method is used in the grayscale case in [48] to regularize the orientation with h . Now we use θ_+ to steer the directional RT so as to perform amplitude/phase analysis in a *unique* local direction. The phase-shifted signal is defined by:

$$\mathcal{R}_T s(\mathbf{x}) = \begin{bmatrix} \mathcal{R}_{\theta_+(\mathbf{x})} s^R(\mathbf{x}) & \mathcal{R}_{\theta_+(\mathbf{x})} s^G(\mathbf{x}) & \mathcal{R}_{\theta_+(\mathbf{x})} s^B(\mathbf{x}) \end{bmatrix} \quad (42)$$

The equation (34) can now be fed by the three channels of the signal s on the one hand, and by those of the phase-shifted version $\mathcal{R}_T s$ on the other hand. Let us finally use it to define the elliptical monogenic representation.

E. Elliptical monogenic wavelets

The elliptical monogenic representation of any color input 2D signal s is given by:

$$s_M = \begin{bmatrix} s^R(\mathbf{x}) \\ s^G(\mathbf{x}) \\ s^B(\mathbf{x}) \\ \mathcal{R}s^R(\mathbf{x}) \\ \mathcal{R}s^G(\mathbf{x}) \\ \mathcal{R}s^B(\mathbf{x}) \end{bmatrix} \xrightarrow{T_R \rightarrow \theta_+(\mathbf{x})} \begin{bmatrix} \mathcal{R}_{\theta_+} s^R \\ \mathcal{R}_{\theta_+} s^G \\ \mathcal{R}_{\theta_+} s^B \end{bmatrix} \xrightarrow{\text{eqs. (22) to (27)}} \begin{bmatrix} A(\mathbf{x}) \\ \varphi(\mathbf{x}) \\ \lambda(\mathbf{x}) \\ \alpha_1(\mathbf{x}) \\ \alpha_2(\mathbf{x}) \\ \alpha_3(\mathbf{x}) \end{bmatrix} \quad (43)$$

The conversion to ellipse parameters is possible thanks to the steering of RTs toward the local main orientation θ_+ based on the Riesz-part tensor constructed from s_M . Note that the construction basically uses three parallel Riesz transforms, before the non-marginal conversion of all the components. The wavelet representation can therefore be based on three parallel monogenic filterbanks. The color decomposition is defined by the use of grayscale monogenic wavelets on every color channel followed by the conversion of output coefficients according to eq. (43). In this paper we use the filterbank defined in section III.

Let us comment the illustration of the wavelet transform at the third scale given on figure 6. We first comment the upper row. The subband s_3 seems well filtered in an isotropic way, with no visibly favored direction. The Riesz part actually shows x - and y - phase-shifting, and the tensor orientation θ_+ is according to our perception, and equally efficient on lines and edges thanks to the smoothing kernel. The Riesz-part is well combined as \mathcal{R}_{θ_+} , where directional structures are clearly

phase-shifted in the proper direction (See how the blue/yellow edge is curved to follow the input yellow line).

Let us now observe the lower row, containing the ellipse parameters. The amplitude and phase are visibly coherent. Color lines and edges are equally detected by the amplitude A , that shows a maximum at their centers. The phase is wrapped in one quadrant of the trigonometric circle to highlight the line-edge information. We actually find 0 (red) and $\pi/2$ (blue) on line- and edge- like structures respectively, which confirms that φ encodes contour type information. The linearity of local color ellipse λ is near 1 almost everywhere. This means that the chosen test image mainly contains “simple” color oscillations, involving two main colors at their apogees in the color space. For example, the yellow line in test image is encoded in the subband by a simple variation between blue and yellow. The information conveyed by all these features is further explored below in the context of synthesis. Note that this result confirms the advantage of having defined a numerically stable computation of ellipse parameters, instead of using the method in [32], which is ill-defined when $\lambda = 1$. Another advantage is that the angles (α_1, α_2) can be interpreted as the main direction in the color space, pointing towards the two opposite colors around a contour (see also [44]). Note how the two angles show locally constant values around contours of a same color.

Note that features appear independent of each other, suggesting a *complementary* encoding of the visual information. Before going into the synthesis method for this wavelet transform, let us briefly discuss some related work.

Related work: The color quaternion Fourier transform

It should be noted that elliptical color paths had already been noticed in the context of color quaternion Fourier transform by Ell and Sangwine [14]. Yet, those ellipses are different from these of present paper, because they arise from a very different combination of real and imaginary parts of red, green and blue Fourier transforms. The “modulus” and “phase” from [14] are then different quantities than those presented here. We think that the quaternion framework - intrinsically 4D - is impeding to handle well the required 6 ellipse parameters. Moreover, the ellipses from [14] are parameterized by a globally constant color “axis”, responsible for the particular combination of color channels, whereas in this paper, the analysis is non-parametric.

The prior work by McCabe *et al.* [36] should also be cited, because it revealed elliptical paths from the use of a classical Fourier transform on image’s chrominance.

Let us now study image synthesis from our monogenic coefficients. We expect monogenic synthesis to provide more interpretative elements about the features.

F. Synthesis

Because this research field is motivated by the global idea of improving existing wavelet transforms, we believe that the famous applications involving reconstruction - compression, contrast enhancement, denoising, watermarking... - must not

be put aside, contrary to most works about analytic/monogenic wavelets. The synthesis wavelet modulation process proposed in section II-C must now be extended to the elliptical feature set, so that color and shape information can be transferred to the synthesis atoms.

Thanks to the parallel structure of the model, reconstruction from Cartesian components is straightforwardly defined by the parallel use of synthesis monogenic filterbanks. As shown on figure 7, the final synthesis method relies on the way we choose to coherently retrieve the Cartesian coordinates after having processed the elliptical monogenic features. At




Fig. 7. Color monogenic analysis and synthesis filterbank. The ? symbol stands for “How to get coherent Cartesian coordinates back from the monogenic features”.

this stage of the work, the color monogenic signal does not perform a “split of identity” *i.e.* we have no reversible polar-like representation of the whole 9D feature set s_M . Actually, we are close to it, since the ellipse parameters provide a 6D split of identity of primary and *directional* Riesz part. But the orientation information of the Riesz part - reduced to θ_+ - is not fully included in the model.

Fortunately, coherent synthesis is possible in some way. We here propose two particular methods. The first one is direct and exact but limited to amplitude-based processing, the second one allows modification of all ellipse parameters but with approximate reconstruction.

1) *Amplitude-based processing*: Many wavelet-based image processing methods such as wavelet shrinkage (denoising) and contrast enhancement only modify the *amplitude* of coefficients. Since the amplitude A is defined as the Euclidean norm of the monogenic vector, applying any weighting k to A is equivalent to applying this weighting on every of its Cartesian components:

$$A \leftarrow kA \quad \Leftrightarrow \quad s_M \leftarrow k[s^R \ s^G \ s^B \ R s^R \ R s^G \ R s^B]^T \quad (44)$$

This way, the effect of any processing of amplitude is properly propagated to the Cartesian monogenic representation and then to the synthesis image.

For example, a wavelet-shrinkage color image denoising scheme can be carried out with:

$$k(x) = \begin{cases} 1 & \text{if } A(x) > T \\ 0 & \text{else} \end{cases} \quad (45)$$

where T is a given threshold. An experimental result is shown on figure 8, featuring a noisy version of the “sailing” picture, and the denoised result. This experiment is the clear illustration that the color monogenic wavelet transform proposed in this paper is able to coherently synthesize images. More precisely, at those coordinates x_i where amplitude is significant, the 9 Cartesian components of $s_M(x_i)$ are preserved. The parallel

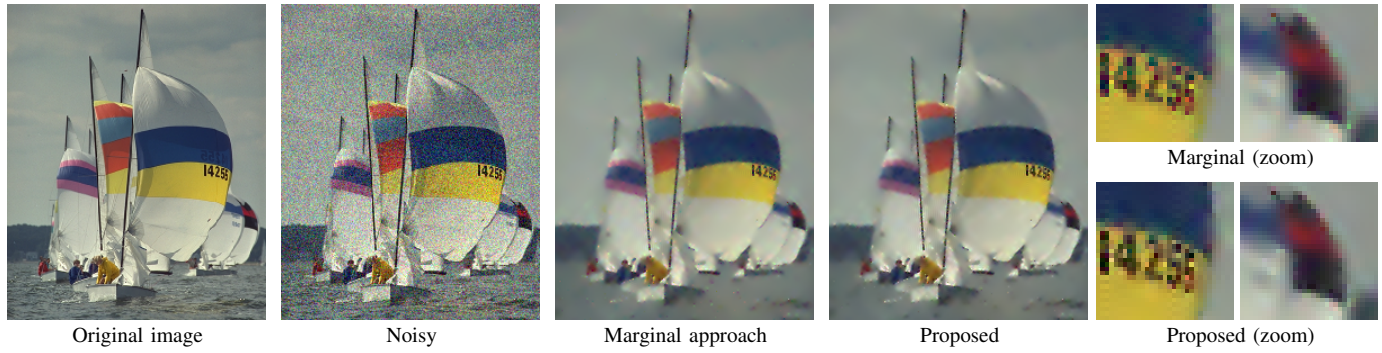


Fig. 8. Example of color monogenic wavelet shrinkage based on amplitude. The noise is Gaussian with $\sigma = 0.1$ and hard thresholding is done through 7 scales. Marginal approach: A^R , A^G and A^B are processed separately with $T = 0.4$. Proposed approach: the elliptical amplitude A is processed with $T = 0.58$.

monogenic filterbanks are then fed with the full information to properly combine odd and even wavelet shapes and produce the corresponding color modulation. To highlight the contribution of embedding the color dimension in our model, the corresponding marginal approach is illustrated for comparison. It consists in applying an amplitude-based greyscale monogenic thresholding on every color channel separately (The threshold has been experimentally adapted to compare as fairly as possible). It can be seen that the marginal approach is less stable with respect to false colors. See the red artifacts around the sail's number, and the green artifacts around the contour of the black-red sail. They are produced by some large noisy wavelet coefficients, which happen to be better regularized in a true color approach. With the marginal method, the unbalanced processing of the color channels generates a strong quantification-like bias in the encoding of the color information. Note how this rounding of color features in the RGB space produces “primary color” artifacts, here red and green, and visually annoying. In contrast, the proposed approach either discards or preserves all channels as a whole, which tends to preserve accurate local color information.

Note that the development of a complete denoising experiment would require a probabilistic model aimed at discriminating the information from the noise, which is out of the scope of this paper.

This wavelet transform is in the spirit of *sparse representations*, because we expect the visual information to be succinctly encoded by few significant coefficients, and closely to the human visual system. Given the rich color and shape description contained in one only monogenic coefficient, we may even think that a good synthesis could still be obtained by keeping only those whose amplitude is a local maxima, like in [34] (This topic is a future work).

In order to use the elliptical monogenic wavelet representation, we have to further understand the link between features and visual elements. Let us now study the shape of synthesis wavelets and their modulation by the elliptical features, through an alternative synthesis method.

2) *The pure directional case:* The full color monogenic filterbank proposed in this paper is tied to a very large dictionary of color wavelet shapes, parameterized in geometry and color by a 9D feature set. Nevertheless, according to the directional model upon which it is based, its proper interpretation requires

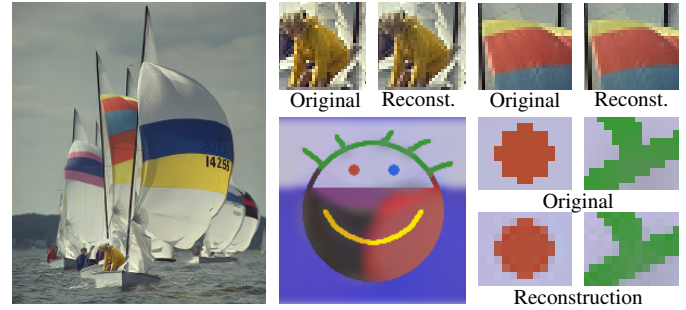


Fig. 9. Pure directional reconstruction of the sailing image (original on figure 8) and the face image (original on figure 6), and zoomed parts.

a unique local orientation, encoded by θ_+ . In some cases, s_M can be such that the Riesz parts have different orientations across color channels, resulting in a biased interpretation of ellipse parameters.

We propose to ignore the multi-oriented cases by assuming:

$$\theta_+(\mathbf{x}) = \arg(\mathcal{R}s^R(\mathbf{x})) = \arg(\mathcal{R}s^G(\mathbf{x})) = \arg(\mathcal{R}s^B(\mathbf{x})) \quad (46)$$

Such assumption is usual in vector differential geometry, for example with Di Zenzo's gradient. This restriction provides a new synthesis method where the directional Riesz part is steered back to θ_+ with $\mathcal{R}s^c = e^{j\theta_+} \mathcal{R}_{\theta_+} s^c$ for $c \in \{R, G, B\}$. We obtain:

$$\begin{bmatrix} A(\mathbf{x}) \\ \varphi(\mathbf{x}) \\ \lambda(\mathbf{x}) \\ \alpha_1(\mathbf{x}) \\ \alpha_2(\mathbf{x}) \\ \alpha_3(\mathbf{x}) \end{bmatrix} \xrightarrow[\theta_+(\mathbf{x})]{\text{eqs. (22) to (27)}} \begin{bmatrix} \mathcal{R}_{\theta_+} s^R \\ \mathcal{R}_{\theta_+} s^G \\ \mathcal{R}_{\theta_+} s^B \end{bmatrix} \xrightarrow{\times e^{j\theta_+}} \begin{bmatrix} s^R(\mathbf{x}) \\ s^G(\mathbf{x}) \\ s^B(\mathbf{x}) \\ \mathcal{R}s^R(\mathbf{x}) \\ \mathcal{R}s^G(\mathbf{x}) \\ \mathcal{R}s^B(\mathbf{x}) \end{bmatrix} \quad (47)$$

The dictionary is reduced to the subset of *pure directional* wavelets, so the reconstruction is not exact in general. It is still a good approximation, since we experimentally noticed almost no visible artifact with our test images, as shown on figure 9. The only visible differences are very mild and occur around the sharp synthetic edges from image *face*. It turns out that most local structures actually fit the directional model, so they can be properly retrieved.

The advantage is that the synthesis wavelets can now be




Fig. 10. Examples of modulation of color monogenic wavelets; synthesized from the Radon based filterbank at the 6th scale. Note that the contrast of the first column has been enhanced.

coherently modulated by the 6 ellipse parameters and θ_+ . Examples of modulated synthesis wavelets are given on figure 10. This experiment confirms that the modulation is properly achieved in terms of phase, orientation and color content.

The first column illustrates the special subset of isotropic wavelets, for which $\varphi = 0$ (the structure is symmetric) and $\lambda = 1$ *i.e.* the elliptical color path reduces to a line between two opposite colors. For these cases, the orientation is not used and set to an arbitrary value. We can see that the produced wavelets are approximately isotropic and actually feature two opposite colors, explicitly indicated in the color space by the color axis (α_1, α_2) . Since $\lambda = 1$, the angle α_3 has no influence.

The other columns imply directional wavelets, for which $\varphi = \pi/2$ indicates an edge-like shape. For the second and third columns, we can verify the correspondence between θ_+ and the perceptible orientation of the colored wavelets. here again, only two main colors are involved because the ellipse reduces to a line ($\lambda = 1$).

The colorimetric pattern becomes richer when the elliptical color path is rounder ($\lambda < 1$), as can be seen in the last two columns. In these cases, α_3 encode an “orthogonal color axis” containing the information of the additional colors present in the oscillation. The smaller λ , the more present these new colors. In the first row for example, some *brown* has been added to the *lime-purple* oscillation. These wavelets may not be involved for retrieving simple edges but rather for color textures. This suggests using the linearity feature λ to discriminate contours ($\lambda \approx 1$) from textural points ($\lambda \approx 0$).

This alternative synthesis scheme may be used for image regularization, where not only the amplitude could be thresholded, but the other elliptical features could also be regularized. For example, the phase can be simplified under some hypothesis about its local linearity, or the color angles can be made locally constant *etc.*

Finally, let us quickly deal with the open issue of multi-oriented cases. By handling three independent orientations instead of just θ_+ , it is possible to produce multi-oriented synthesis wavelets, that do not fit the directional model, like illustrated in the figure 11. In the greyscale computer vision




Fig. 11. Examples of multi-oriented color wavelets. $\lambda = 1$, $\varphi = \pi/2$, $\alpha_1 = 0.13\pi$, $\alpha_2 = -0.06\pi$. Values of $(\theta_R, \theta_G, \theta_B)$ are given under each one.

literature, these shapes are referred to as “intrinsically 2D” and include *corners*, *junctions*, *crossings* *etc.* A prospect of this work is to extend the model so that it includes these shapes as well. This could involve geometric notions of *isotropy* and *coherency* - well known in the structure tensor literature [46] - and may be linked to the *conformal monogenic* representation [17], [51]. The challenge is to preserve the phase interpretation, which has been so far strictly tied to directional models. Such a research lead could end up with a 9D split of identity associated to a unique modulating reconstruction scheme.

Now the synthesis has been defined and the elliptical monogenic features explained, let us conclude this article in a color image analysis context, through the definition and the estimation of the local color frequency.

V. LOCAL COLOR FREQUENCY ESTIMATION

It is well known in the signal processing literature that the local frequency ν can be basically obtained by differentiating the local phase φ . This topic is numerically challenging and several methods can be used [25]. In the 2D world of images, the topic has a long history in the AM-FM literature [3] where local frequency is interpreted as some *granularity* feature.

The monogenic representation presented in this paper gives direct access to local phase in a multiscale fashion. This is the first time to our knowledge that phase data can be obtained from a color image. This motivates the proposition of a local frequency estimation method based on our filterbank. The first issue is numerical phase differentiation.

A. Differentiating the phase

After having decomposed any color image into elliptical monogenic subbands thanks to the filterbank of section IV-E, we need to compute the spatial derivative of $\varphi(\mathbf{x})$.

An analytic method is given with the greyscale monogenic filterbank from [48]. This consists in using an other version of the filterbank, where the Riesz part filters are tuned to correspond to partial derivatives of the Riesz transform. The coefficients can then be combined to analytically find the local frequency. In our color case, extending the analytical formula of the derivative of $\varphi(\mathbf{x})$ is more intricate, so we adopt another approach: the *finite central difference* method.

The phase is locally understood toward the direction θ_+ , and the differentiation must of course be applied in this same direction. In the discrete world, we can only compute horizontal, vertical or diagonal differences. To steer the calculation properly, our continuous orientation data θ_+ must be quantized




Fig. 12. Estimation of the local frequency $\nu(\mathbf{x})$ in radians/pixel. The noisy version is obtained by adding Gaussian noise with $\sigma = 0.1$.

so as to select the neighbor “left” φ_l and “right” φ_r samples to be compared to the local value $\varphi(\mathbf{x})$. We propose the following quantization formula

$$\theta_q = (\pi/4) \text{round}(\theta_+ / (\pi/4)) \quad (48)$$

giving 4 possible directions: horizontal (0), vertical ($\pi/2$ and $-\pi/2$ are merged) and the two diagonals ($\pm\pi/4$).

For example, if θ_+ lies in $[-\frac{\pi}{8}; \frac{\pi}{8}]$; then the wavefront is approximately parallel to the y -axis, and the phase evolves mainly along the x -axis. So $\varphi(x_1, x_2)$ has to be differentiated with $\varphi_l = \varphi(x_1 - 1, x_2)$ and $\varphi_r = \varphi(x_1 + 1, x_2)$.

The circularity of angle data can be handled with complex exponentials, and the local frequency defined by:

$$\nu = \left| \arg \left((\exp(\varphi_r) / \exp(\varphi_l))^{\frac{1}{2}} \right) / \cos(\theta_+ - \theta_q) \right| \quad (49)$$

The right side term compensates the bias due to orientation quantization, according to a local linear model of φ . Note that for diagonal cases ($\theta_q = \pm\pi/4$), ν must also be divided by $\sqrt{2}$. It shall also be noted that in practice, the phase difference is not directly estimated between φ_r and φ_l , but computed as the average of the phase differences on both sides, which prevents some phase wrapping. For more computation related details, see the source code [43].

B. Multiscale estimation

The advantage of working in the wavelet domain is that we can locally select the “best scale” - which is the one containing the natural frequency of the signal *i.e.* showing the largest energy. The pointwise best scale $k(\mathbf{x})$ is defined by:

$$k(\mathbf{x}) = \arg \max_i \{A_i(\mathbf{x})\} \quad (50)$$

In order to improve the estimation’s robustness, the final algorithm is then as follows:

- Process color monogenic wavelet transform;
- Differentiate the phase $\varphi_i(\mathbf{x})$ at all scales i to compute local frequency $\nu_i(\mathbf{x})$;
- Compute the local best scale $k(\mathbf{x})$;
- Apply 3×3 median filtering to $k(\mathbf{x})$ so that the used scale is locally constant for numerical stability;
- Extract the local frequency $\nu(\mathbf{x}) = \nu_{k(\mathbf{x})}(\mathbf{x})$

The performance of this method is illustrated on figure 12. This first test is done on a synthetic image, featuring explicit local frequencies for various examples of colorimetric and geometric parameters. We expect our method to measure constant values of ν within every spatial sub-square, as

shown on the second graphic. The measure (third graphic) is strikingly close to theoretical values, with some expected and not damaging instability around borders. The estimation method is satisfactory for low frequencies down to $\nu = 0.35$ (period of about 18 pixels, lower-left wave) and up to $\nu = 2.5$ (≈ 2.5 pixels, upper-right wave). Efficiency is regardless of orientation and color features, which experimentally confirms the soundness of the method.

The multiscale property of our method is expected to bring robustness with respect to noise. The figure 12 shows a noisy version of the same test image, and the estimated local frequency (fifth graphic). It is visible that the result is stable, which is confirmed by the good quality of estimation obtained by a simple further regularization of ν with a 5×5 median filter (sixth graphic). The advantage of using a multiscale local frequency estimation is therefore confirmed here.

We finally test it on a natural image as shown on figure 13. This image features a perceptually constant unique frequency,




Fig. 13. From left to right: Natural image, estimated frequency ν with 5×5 median filtering, histogram of ν .

oscillating horizontally. The difficulty is that colorimetric properties are various - blue, green and yellow. We also expect measurements not to be as regular as in the synthetic case, in particular due to richer frequency content, and randomness of real data. The estimated $\nu(\mathbf{x})$ has been regularized by a 11×11 median filtering.

The figure 13 shows the estimated frequency (center graphic), together with a histogram of the measured $\nu(\mathbf{x})$ (right graphic). Apart from a mild random variation due to the original image itself, we can see that ν still lies within a small interval around $\nu = 0.51$ rad./pixel. This value corresponds to a period of around 12 pixels, which is well in accord with the visible periodicity of the pattern in this image.

This experiment shows that our elliptical monogenic representation actually provides sound data, according to an oscillatory model in the spirit of the classical Fourier analysis, with a intuitive extension of the amplitude and phase concepts. This opens the way towards many possible applications.

Texture segmentation of course could benefit from the color frequency feature, as well as the local color linearity λ which may discriminate complex color patterns from simple color contours. The ellipse angles α_1 , α_2 and α_3 can also provide simple discrimination. Phase data has also been shown useful for texture classification in [42]. Note that the development of a full segmentation method based on the monogenic features would require the design and normalization of optimal feature vectors, as well as the modelling of meaningful distances, which is out of the scope of this paper.

The elliptical monogenic features make new invariant descriptors of local color and shape for *keypoint* detection. The invariance properties and the signal flavor of this work could make it a competitive color alternative to the famous SIFT detector [33], probably closer to the human visual system.

Stereo vision and motion estimation are also possible applications, since the phase data can be used to detect sub-pixel accurate shifts (as already done in the greyscale case).

VI. CONCLUSION

This article investigates the numerical representation of color images by monogenic wavelet transforms.

Monogenic representations have been gaining much interest for providing rich physical meaning with respect to usual wavelet transforms. This family of sparse representations suffers from a lack of synthesis algorithms as well as color extensions. Applications have been so far limited to greyscale image analysis only, putting aside the famous wavelet-based denoising and compression as well as any color handling.

To overcome this limit, we have proposed in this paper a clear synthesis method, that we exploited through the definition of a whole new monogenic filterbank. Our numerical algorithm has an explicit synthesis part, and is then able to produce synthesis wavelets being coherently modulated by the amplitude, phase and orientation encoded in one monogenic coefficient. The construction is based on a discrete Radon transform and an original 1D analytic undecimated wavelet transform. This is the first time that the synthesis of the 2D signal from its monogenic wavelet coefficients is presented, thanks to the concept of modulated synthesis wavelets.

The second part of the paper tackles the color extension. The monogenic representation is already multi-dimensional in the scalar case, and its extension to vector-valued signals is not trivial. We start in 1D by defining the *color analytic signal*, revealing the fundamental elliptical oscillation and the *color phase* concept. The representation is then brought to 2D thanks to a directional approach, making use of the color structure tensor to build the *directional Riesz transform*. The *elliptical monogenic representation* is carried out as a non-marginal extension of the monogenic framework, tied to a directional color elliptical oscillation model. It achieves local analysis of images in terms of color and geometry features, including amplitude, phase, elliptical parameters and orientation.

The monogenic filterbank derived in the first part is used to implement the new color wavelet transform. We observe a good separation of geometrical and colorimetric information through the various local features. The synthesis is again

deeply studied, and two methods are proposed. Exact reconstruction is possible in the case of amplitude-based processing, allowing for a variety of applications such as wavelet-based denoising and enhancement. For the price of an approximate reconstruction - which is experimentally not visible - a full color wavelet modulation is achieved from the elliptical feature set, opening the way to new kinds of image processing tasks such as phase based regularization or any automatic local modification of graphical elements, by tuning their elliptical properties.

Finally, the color monogenic representation is applied to estimate from a color image the local frequency information, in a multiscale fashion. The method appears to be robust against additive noise, and is successfully used on synthetic and natural images.

The elliptical monogenic wavelet transform provides a ready-to-use multiscale pointwise low-level analysis combining classical computer vision methods and new kinds of information, while being efficiently designed for color images. This suggests many prospective applications.

We think that a SIFT-like keypoint detector could be based on this filterbank, and bring better connections with the human visual system, as well as better handling of the color information. Segmentation may also benefit from the advanced encoding of color contours and texture-like information (as linearity and frequency for example). Image reconstruction from wavelet maxima seems an interesting lead, given that significant feature sets are located on amplitude's maxima, and certainly contain a sparse encoding of the neighborhoods. This could be used for denoising. Prospects of this work also include the definition of a full 9D split-of-identity for complete reversibility. Finally, perceptually meaningful color spaces and dependencies between colour channels have to be considered.

VII. THANKS

R. Souillard works with CRESCEN2O. We want to thank the anonymous reviewers for their valuable comments.

REFERENCES

- [1] T. Batard and M. Berthier. Spinor Fourier transform for image processing. *IEEE J. Sel. Topics Signal Process. (Special Issue on Differential Geometry in Signal Processing) IEEE JSTSP*, 7(4):605–613, 2013.
- [2] A. Belaid, D. Boukerroui, Y. Maingourd, and J. F. Lerallut. Phase-based level set segmentation of ultrasound images. *Trans. Info. Tech. Biomed.*, 15(1):138–147, Jan. 2011.
- [3] A. C. Bovik. *Handbook of Image and Video Processing*. Academic Press, May 2000.
- [4] F. Brackx, B. De Knock, and H. De Schepper. On generalized Hilbert transforms and their interaction with the Radon transform in clifford analysis. *Math. Meth. Appl. Sci.* 2007; 30:1071-1092, Wiley InterScience, 2006.
- [5] T. Bülow. Hypercomplex spectral signal representation for the processing and analysis of images. *Thesis*, Aug. 1999.
- [6] P. Carré and E. Andres. Discrete analytical ridgelet transform. *Sig. Proc.*, 84:2165–2173, 2004.
- [7] W. L. Chan, H. H. Choi, and R. G. Baraniuk. Coherent multiscale image processing using dual-tree quaternion wavelets. *IEEE Trans. Image Process.*, 17(7):1069–1082, Jul. 2008.
- [8] K. N. Chaudhury and M. Unser. Construction of Hilbert transform pairs of wavelet bases and Gabor-like transforms. *IEEE Trans. Signal Process.*, 57(9):3411–3425, 2009.
- [9] G. Chen and B. Kégl. Image denoising with complex ridgelets. *Pat. Rec.*, 40(2):578–585, 2007.

- [10] N. Chenouard and M. Unser. 3D steerable wavelets in practice. *IEEE Transactions on Image Processing*, 21(11):4522–4533, November 2012.
- [11] J. G. Daugman. Two-dimensional spectral analysis of cortical receptive field profiles. *Vis. Res.*, 20(10):847–856, 1980.
- [12] G. Demarcq, L. Mascarilla, M. Berthier, and P. Courtellemont. The color monogenic signal: Application to color edge detection and color optical flow. *J. Math. Im. Vis. (JMIV)*, 40(3):269–284, 2011.
- [13] S. Di Zenzo. A note on the gradient of a multi-image. *Computer Vision, Graphics, and Image Processing*, 33(1):116–125, 1986.
- [14] T. A. Ell and S. J. Sangwine. Hypercomplex Fourier transforms of color images. *IEEE Trans. Image Process.*, 16(1):22–35, 2007.
- [15] M. Felsberg. Low-level image processing with the structure multivector. *Thesis*, 2002.
- [16] M. Felsberg and G. Sommer. The monogenic signal. *IEEE Trans. Signal Process.*, 49(12):3136–3144, 2001.
- [17] O. Fleischmann, L. Wietzke, and G. Sommer. Image analysis by conformal embedding. *J. of Math. Im. and Vis.*, 40(3):305–325, 2011.
- [18] D. Gabor. Theory of communication. *J. Inst. Elec. Eng.*, 93(3):429–457, 1946.
- [19] R. A. Gopinath. The phaselet transform-an integral redundancy nearly shift-invariant wavelet transform. *IEEE Trans. Signal Process.*, 51(7):1792–1805, 2003.
- [20] A. Grossman and J. Morlet. Decomposition of Hardy functions into square integrable wavelets of constant shape. *SIAM J. Math. Anal.*, 15(4):723–736, 1984.
- [21] S. L. Hahn. Multidimensional complex signals with single-orthant spectra. *Proc. IEEE*, 80(8):1287–1300, 1992.
- [22] D. Helbert, P. Carré, and E. Andrès. 3-D discrete analytical ridgelet transform. *IEEE Trans. Image Process.*, 15(12):3701–3714, 2006.
- [23] S. Held, M. Storath, P. Massopust, and B. Forster. Steerable wavelet frames based on the Riesz transform. *IEEE Trans. Image Process.*, 19(3):653–667, 2010.
- [24] E. Hitzer and S. J. Sangwine, editors. *Quaternion and Clifford Fourier Transforms and Wavelets*. Trends in Mathematics. Birkhäuser Springer Basel, 2013.
- [25] N. E. Huang, Z. Wu, S. R. Long, K. C. Arnold, X. Chen, and K. Blank. On instantaneous frequency. *Advances in Adaptive Data Analysis*, 1:177–229, 2009.
- [26] B. Jähne. *Digital Image Processing 6th Edition*. Springer, Berlin, 2005.
- [27] N. Joshi, S. Bond, and M. Brady. The segmentation of colorectal MRI images. *Medical Image Analysis*, 14(4):494–509, 2010.
- [28] U. Köthe and M. Felsberg. Riesz-transforms versus derivatives: On the relationship between the boundary tensor and the energy tensor. In J. Weickert R. Kimmel, N. Sochen, editor, *Proc. Scale-Space, LNCS 3459*, pages 179–191. Springer, 2005.
- [29] K. G. Larkin, D. Bone, and M. A. Oldfield. Natural demodulation of two-dimensional fringe patterns: I. general background to the spiral phase quadrature transform. *J. Opt. Soc. Am.*, 18 (8):1862–1870, 2001.
- [30] J. Li, H. Zhao, X. Zhou, and C. Shi. Robust stereo image matching using a two-dimensional monogenic wavelet transform. *Opt. Lett.*, 34(22):3514–3516, 2009.
- [31] J. Lilly and S. Olhede. Generalized morse wavelets as a superfamily of analytic wavelets. *IEEE Trans. Image Process.*, 60(11):6036–6041, 2012.
- [32] J. M. Lilly. Modulated oscillations in three dimensions. *IEEE Trans. Signal Process.*, 59(12):5930–5943, 2011.
- [33] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *Int'l J. of Comp. Vis.*, 60(2):91–110, 2004.
- [34] S. Mallat and S. Zhong. Characterization of signals from multiscale edges. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14(7):710–732, 1992.
- [35] T. Maltaverne, P. Delachartre, and A. Basarab. Motion estimation using the monogenic signal applied to ultrasound elastography. In *Proc. IEEE Eng. Med. Biol. Soc. Conf. (EMBC'10)*, pages 33–36, Buenos Aires, Argentina, 2010.
- [36] A. McCabe, T. Caelli, G. West, and A. Reeves. Theory of spatiochromatic image encoding and feature extraction. *Journal of The Optical Society of America A-optics Image Science and Vision*, 17, 2000.
- [37] M. N. Nabighian. Toward a three-dimensional automatic interpretation of potential field data via generalized Hilbert transforms: Fundamental relations 780786. *Geophysics*, 49(6), 1984.
- [38] S. C. Olhede and G. Metikas. The monogenic wavelet transform. *IEEE Trans. Signal Process.*, 57:3426–3441, 2009.
- [39] G. Sapiro and D. Ringach. Anisotropic diffusion of multivalued images with applications to color filtering. *IEEE Trans. Image Process.*, 5(11):1582–1586, 1996.
- [40] I. Selesnick, R. Baraniuk, and N. Kingsbury. The dual-tree complex wavelet transform - a coherent framework for multiscale signal and image processing. *IEEE Signal Processing Mag.*, 22(6):123–151, 2005.
- [41] R. Souillard and P. Carré. Color monogenic wavelets for image analysis. In *Proc. IEEE Int'l Conf. on Image Processing*, pages 277–280, Brussels, Belgium, Sep. 2011.
- [42] R. Souillard and P. Carré. Quaternionic wavelets for texture classification. *Pat. Rec. Lett.*, 32(13):1669–1678, 2011.
- [43] R. Souillard and P. Carré. MATLAB source code for the implementation and use of Elliptical Monogenic Wavelets. https://hal.archives-ouvertes.fr/hal-01138349/file/EMWT_MATLAB_2015_Revision.zip, Apr. 2015.
- [44] R. Souillard, P. Carré, and C. Fernandez-Maloigne. Vector extension of monogenic wavelets for geometric representation of color images. *IEEE Trans. Image Process.*, 22(3):1070–1083, Mar. 2013.
- [45] M. Storath. Directional multiscale amplitude and phase decomposition by the monogenic curvelet transform. *SIAM J. Imaging Sciences*, 4(1):57–78, 2011.
- [46] D. Tschumperlé and R. Deriche. Vector-valued image regularization with PDEs: A common framework for different applications. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(4):506–517, 2005.
- [47] M. Unser and N. Chenouard. A unifying parametric framework for 2D steerable wavelet transforms. *SIAM Journal on Imaging Sciences*, 6(1):102–135, 2013.
- [48] M. Unser, D. Sage, and D. Van De Ville. Multiresolution monogenic signal analysis using the Riesz-Laplace wavelet transform. *IEEE Trans. Image Process.*, 18(11):2402–2418, 2009.
- [49] M. Unser and D. Van De Ville. The pairing of a wavelet basis with a mildly redundant analysis via subband regression. *IEEE Trans. Image Process.*, 17(11):1–13, 2008.
- [50] J. Ville. Théorie et applications de la notion de signal analytique. *Cables et transmission*, 2A(1):61–74, 1948.
- [51] L. Wietzke and G. Sommer. The signal multi-vector. *Journal of Mathematical Imaging and Vision*, 37:132150, June 2010.
- [52] L. Zhang, L. Zhang, Z. Guo, and D. Zhang. Monogenic-LBP: A new approach for rotation invariant texture classification. In *Proc. IEEE Int'l Conf. on Image Processing*, pages 2677–2680, 2010.

Raphaël Souillard Dr Raphaël Souillard received the PhD degree in 2012 from the University of Poitiers, France. He is currently a postdoctoral research fellow at the XLIM laboratory, France. He received the national prize “Blanc-Lapierre” granted by the SEE (Society of the electricity, the electronics, and the information technologies) for his master's thesis in 2009. His research interests include wavelets and filterbanks, analytic/monogenic representation methods, extension of signal and image processing tools for multidimensional/multivalued signals like color images.

Philippe Carré Professor Philippe Carré received the Engineer degree in Computer Engineering from Compiegne University of Technology, France, in 1995, and the Ph.D. degree in Signal Processing and Computer Science in University of Poitiers, France, in 2000 for his work on the definition of discrete atomic Time-Frequency representations. He is presently Professor at the XLIM laboratory, France. His interests include signal and image processing, wavelets, Time-Frequency and multiscale decompositions theory and their applications to the fields of denoising, segmentation and watermarking.