To be published in the proceedings of CAIP 2015.
The final publication will be available at link.springer.com.

Tensorial Orientation Scores

Jasper J. van de Gronde

Johann Bernoulli Institute for Mathematics and Computer Science,
University of Groningen, P.O. Box 407, 9700 AK Groningen, The Netherlands
j.j.van.de.gronde@rug.nl

Abstract. Orientation scores [7, 10] are representations of images built using filters that only select on orientation (and not on the magnitude of the frequency). Importantly, they allow (easy) reconstruction, making them ideal for use in a filtering pipeline. Traditionally a specific set of orientations has to be chosen, and the response is determined for those orientations. This work introduces an alternative, where a tensorial representation is built that approximates an idealized orientation score in a well-defined way. It is shown that the filter's output can be usefully interpreted in terms of tensor decompositions. Tensorial orientation scores can be considered to fit in a family of filtering schemes that includes not just traditional orientation scores, but also monomial filters [16] and curvelets [3].

Keywords: orientation scores, monomial filters, curvelets, tensors

1 Introduction

Typical wavelet transforms can be used to both analyse and synthesize signals in terms of their content in different frequency bands, but provide only very coarse-grained information on orientation. Orientation scores [7, 10] on the other hand do not concern themselves with frequency magnitude at all, but can give very detailed orientation information, all the while still allowing for both analysis *and* synthesis of signals. This means they are ideally suited as part of a filtering pipeline, as we can convert an image to an orientation score, filter that, and then turn the orientation score into an image again, such that if the filter does not do anything, one simply recovers the original image (rather than a blurred version for example).

In their original formulation, orientation scores only allowed for a discrete set of orientations. More recently, van de Gronde et al. [14] used a tensorial variant of orientation scores. Tensor decompositions were used to extract very precise orientation information, even while using tensors of relatively low degree. This was then used as part of a larger filtering pipeline to extract thin and long (but curved) structures. Unfortunately there was too little room at the time to give a good description and analysis of how these filters worked exactly. This work remedies the situation, by showing how tensorial orientation scores can be developed from first principles.

Having defined idealized tensorial orientation scores it is discussed how such filters can be applied in practice, showing that any singularities present in the idealized filters naturally go away in the spatially bounded and band limited case. Next, the interpretation of the filter output in terms of tensor decompositions is discussed, laying the foundation for a scheme in the spirit of tensor voting [21]. Finally, tensorial orientation scores are compared to traditional orientation scores, monomial filters [16], and curvelets [3].

2 Definitions and notation

Throughout this work, the following definitions for the Fourier transform and its inverse are assumed:

$$\mathcal{F}[f](\boldsymbol{\xi}) = \int_{\mathbb{R}^d} f(\mathbf{x})\, e^{-2\pi\, i\, \mathbf{x}\cdot\boldsymbol{\xi}}\, \mathrm{d}\mathbf{x}, \text{ and}$$
$$\mathcal{F}^{-1}[g](\mathbf{x}) = \int_{\mathbb{R}^d} g(\boldsymbol{\xi})\, e^{2\pi\, i\, \boldsymbol{\xi}\cdot\mathbf{x}}\, \mathrm{d}\boldsymbol{\xi}.$$

In addition $\hat{f}$ will be used to denote the Fourier transform of f.

This work makes extensive use of symmetric tensors. These can be viewed as a generalization of vectors and symmetric matrices. Symmetric tensors can be built from a vector space $\mathbb{C}^d$ using the symmetric tensor product '$\odot$' [18, ch. 4 Proposition 5.7] (which is linear in both arguments). In terms of vectors and matrices, one could say that $\mathbf{a} \odot \mathbf{b}$ ($\mathbf{a}, \mathbf{b} \in \mathbb{C}^d$) is equivalent to the matrix $\frac{1}{2}(\mathbf{a}\,\mathbf{b}^T + \mathbf{b}\,\mathbf{a}^T)$. The result of the n-times repeated tensor multiplication of $\mathbf{a}$ by itself, a degree-n symmetric tensor, will be denoted by $\mathbf{a}^{\odot n}$. It can be shown that a general tensor of degree n is symmetric if and only if it can be written as a sum of such tensor "powers" [5, Lemma 4.2]; the number of such tensor powers needed is the (symmetric) rank of the tensor, so a rank one tensor of degree n can be written as $\mathbf{a}^{\odot n}$, for some $\mathbf{a} \in \mathbb{C}^d$. Note that lower-case letters are used for vectors, while upper-case letters are used for (higher degree) tensors.

The inner product on symmetric tensors can be defined [1, §V.3; 13, §1.25; 19, Eq. (2.1)] based on its linearity and the inner product on the underlying vector space ($\mathbf{a}, \mathbf{b} \in \mathbb{C}^d$):

$$\mathbf{a}^{\odot n} \cdot \mathbf{b}^{\odot n} = (\mathbf{a} \cdot \mathbf{b})^n.$$

Note that this inner product is roughly equivalent to the Frobenius inner product on matrices [11, p. 332].

The symmetric degree-n identity tensor $\mathbf{I}_n$, with n even, is defined here as the unique symmetric tensor that satisfies $\mathbf{a}^{\odot n} \cdot \mathbf{I}_n = \|\mathbf{a}\|^n$ for all non-zero $\mathbf{a} \in \mathbb{R}^d$ (this characterization does not hold for non-real vectors). It can be given explicitly be either of the following equivalent constructions:

$$\mathbf{I}_n = D_{n,d} \int_{S_d} \mathbf{s}^{\odot n}\, \mathrm{d}\mathbf{s} \qquad \text{or} \qquad \mathbf{I}_n = \mathbf{I}_2^{\odot n/2}.$$

Here S_d is the sphere of all unit length vectors in $\mathbb{R}^d$ and $D_{n,d}$ is a (positive real) constant fully determined by $\mathbf{a}^{\odot n} \cdot \mathbf{I}_n = \|\mathbf{a}\|^n$ (for any non-zero $\mathbf{a} \in \mathbb{R}^d$), while $\mathbf{I}_2$ can be seen to be $\sum_{k \in \mathcal{K}} \mathbf{e}_k^{\odot 2}$ for any orthonormal basis $\{\mathbf{e}_k\}_{k \in \mathcal{K}}$ of $\mathbb{R}^d$.

The fractional anisotropy of a degree-2 tensor on a 2D vector space is defined here as the ratio between the absolute difference of the largest and smallest eigenvalue, and the square root of the sum of the squares of the eigenvalues. This is one for a rank one tensor, and zero for the identity tensor.

3 Tensorial Orientation Scores

To find local orientations, traditional orientation scores use filters that divide the Fourier domain into several sectors, resembling the way a pie (or cake) is cut. If we increase the number of sectors and take the limit as the number of sectors goes to infinity, then instead of getting values for a discrete number of orientations per point, we would – per point – get a function whose domain is the unit sphere. So any method to represent such a function can be used for computing orientation scores, and here tensors are used instead of sampling.

The continuous orientation score can be considered to take infinitesimal slices from the spectrum. Since at each point the cross-section of such a slice is proportional to the $(d-1)$-th power of the radius, we have ($\mathbf{x} \in \mathbb{R}^d$, $\mathbf{s} \in S_d$):

$$U_f(\mathbf{x}, \mathbf{s}) = \int_0^\infty r^{d-1}\, \mathcal{F}[f](r\,\mathbf{s})\, e^{2\pi\, i\, r\, \mathbf{s}\cdot\mathbf{x}}\, \mathrm{d}r. \tag{1}$$

Note that the integral over the unit sphere of $U_f(\mathbf{x})$ recovers $f(\mathbf{x})$.

Now, instead of sampling the continuous orientation score, here it is approximated using symmetric tensors. In particular, if $U_f : \mathbb{R}^d \times S_d \to \mathbb{C}$ is the continuous orientation score, then we compute a symmetric tensor field $\mathbf{U}_{n,f} : \mathbb{R}^d \to (\mathbb{C}^d)^{\odot n}$ such that for any (symmetric) degree-n tensor $\mathbf{A}$ and position $x \in \mathbb{R}^d$

$$\int_{S_d} (\mathbf{A} \cdot \mathbf{s}^{\odot n})\, U_f(\mathbf{x}, \mathbf{s})\, \mathrm{d}\mathbf{s} = \mathbf{A} \cdot \mathbf{U}_{n,f}(\mathbf{x}). \tag{2}$$

An explicit form for $\mathbf{U}_{n,f}$ can be found by plugging in Eq. (1) in Eq. (2):

$$\begin{aligned}
\int_{S_d} (\mathbf{A} \cdot \mathbf{s}^{\odot n})\, U_f(\mathbf{x}, \mathbf{s})\, \mathrm{d}\mathbf{s} &= \mathbf{A} \cdot \int_{S_d} \mathbf{s}^{\odot n} \int_0^\infty r^{d-1}\, \mathcal{F}[f](r\,\mathbf{s})\, e^{2\pi\, i\, r\, \mathbf{s}\cdot\mathbf{x}}\, \mathrm{d}r\, \mathrm{d}\mathbf{s} \\
&= \mathbf{A} \cdot \int_{S_d} \int_0^\infty \frac{\mathbf{s}^{\odot n}}{r^d}\, r^{d-1}\, \mathcal{F}[f](r\,\mathbf{s})\, e^{2\pi\, i\, r\, \mathbf{s}\cdot\mathbf{x}}\, \mathrm{d}r\, \mathrm{d}\mathbf{s} \\
&= \mathbf{A} \cdot \int_{\mathbb{R}^d} \frac{\boldsymbol{\xi}^{\odot n}}{\|\boldsymbol{\xi}\|^n}\, \mathcal{F}[f](\boldsymbol{\xi})\, e^{2\pi\, i\, \boldsymbol{\xi}\cdot\mathbf{x}}\, \mathrm{d}\boldsymbol{\xi}.
\end{aligned}$$

The tensor-valued filter can thus be defined by

$$\hat{\mathbf{H}}_n(\boldsymbol{\xi}) = \frac{\boldsymbol{\xi}^{\odot n}}{\|\boldsymbol{\xi}\|^n}. \tag{3}$$

Here $\hat{\mathbf{H}}_n$ is the Fourier transform of $\mathbf{H}_n$.

For even n the inner product between $\mathbf{U}_{n,f}(x)$ and the identity tensor gives $f(x)$, as $\frac{\boldsymbol{\xi}^{\odot n}}{\|\boldsymbol{\xi}\|^n} \cdot \mathbf{I}_n = 1$, away from the origin at least. This ensures that the per-position inner product between a filtered signal and $\mathbf{I}_n$ recovers the original signal. One way to make sense of what happens at the origin is to say that $\hat{\mathbf{H}}_n(0) = \frac{\mathbf{I}_n}{\mathbf{I}_n \cdot \mathbf{I}_n}$. This is compatible with convolving (in the frequency domain) $\hat{\mathbf{H}}_n$ with a rotationally invariant Gaussian that integrates to one, and taking the limit as the standard deviation goes to zero. Similarly, for odd n the response at the origin can be considered zero. Note that it can be shown that $\mathbf{H}_n$ is homogeneous of degree $-d$: $\mathbf{H}_n(\mathbf{x}) = \|\mathbf{x}\|^{-d} \mathbf{H}_n(\mathbf{x}/\|\mathbf{x}\|)$.

It is useful to note that $\mathbf{H}_n$ has the form of the n-th (tensorial) derivative of the fundamental solution to the $n/2$-Laplace equation [2, 4, 12]. For example, $\mathbf{H}_n$ is the Hessian of the fundamental solution to the ordinary Laplace equation. In general, for all n and d except when $n = d + 2m$ for some $m \in \mathbb{N}$,

$$\mathbf{H}_n(\mathbf{x}) = \nabla^{\odot n} \left[\frac{\Gamma(d/2 - n/2)\, \|\mathbf{x}\|^{n-d}}{(2i)^n\, \pi^{d/2}\, \Gamma(n/2)} \right]. \tag{4}$$

Here $\nabla^{\odot n} f(x)$ is taken to be the symmetric tensor corresponding to the n-fold application of the gradient operator. When n does equal $d + 2m$, there unfortunately seems to be some disagreement in the literature about the correct general form. The interested reader is referred to Gel'fand and Shilov [12] and Boyling [2].

The observant reader may have deduced from Eq. (4) that for odd n, the filters are purely imaginary in the spatial domain, while for even n the filters are purely real. This stems from the fact that the odd numbered filters capture "odd" structures, while the even numbered filters capture "even" structures, and that a real-valued signal has conjugate symmetry in the Fourier domain. This also means one needs both an odd and an even filter to give a proper approximation of U_f (although this will be ignored for now). Knutsson and Westin [15] discuss how to combine even and odd filters to get phase-invariant filters.

3.1 Computation

Although the completely continuous and unbounded filters defined in Eq. (3) are useful for reasoning about the idealized form of the filters, some care is needed in applying them in real life situations. In particular: the filters have a singularity at the origin, and the integral of the filter is problematic (considering the filter in the spatial domain). That is, although above we have seen that one can in fact give meaning to the overall integral of the entire filter, this is a fairly fragile trick, and it breaks down as soon as we are interested in integrating the absolute value of the filter or in integrating over a wedge of the filter for example. Luckily, these problems go away as soon as the filters are made band-limited and spatially bounded.

Making the filters band limited is (conceptually) easy: just convolve the spatial representation with a sampling kernel k (or multiply its frequency domain

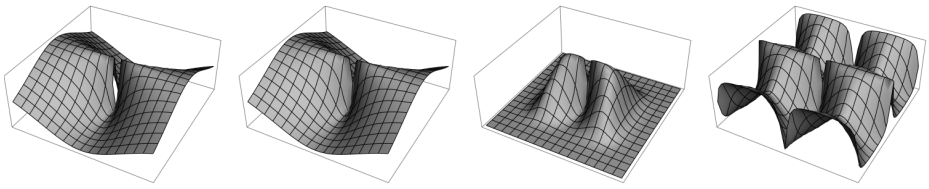

Fig. 1. The different stages in constructing a properly sampled and bounded kernel (in the Fourier domain). From left to right: the idealized filter ($\hat{\mathbf{H}}_n$, undefined at the origin), the windowed filter ($\hat{w} \star \hat{\mathbf{H}}_n$), the filter after applying the sampling kernel ($\hat{k}\,(\hat{w} \star \hat{\mathbf{H}}_n)$), and the filter after convolving with a Dirac comb to implement spatial sampling. Here, k was a third order cardinal spline, while w was a Gaussian with $\sigma = 20$. The graphs show the Fourier transforms on the domain $[-1, 1]^2$.

representation with $\hat{k}$). As a result, the inner product with the identity tensor no longer gives the Dirac delta, but rather the kernel used for limiting the bandwidth. Similarly, it can be seen that the response at the origin is now finite and well-defined, since the integral (in the Fourier domain) of the original filter multiplied by the Fourier transform of the sampling kernel will be bounded. The integral of the filter in the spatial domain is still ill-defined though.

Making the filters spatially bounded is equally easy: just multiply with the windowing function w in the spatial domain (or convolve with $\hat{w}$ in the frequency domain). Analogous to the usual requirement that the integral of a sampling kernel is one, the spatial windowing function must be one at the origin. As a result, the corresponding kernel in the frequency domain can be considered to be a smoothing kernel (effectively a sampling kernel for the frequency domain). If the kernel is rotationally symmetric, we recover the previously stated result that the integral of the spatial response of the filter is $\frac{\mathbf{I}_n}{\mathbf{I}_n \cdot \mathbf{I}_n}$ (the scaling follows from the fact that $\frac{\boldsymbol{\xi}^{\odot n}}{\|\boldsymbol{\xi}\|^n} \cdot \mathbf{I}_n = 1$ away from the origin). In general, if the signal is already band limited, and a suitable windowing function is chosen, the integral of the windowed filter (in the spatial domain) becomes perfectly well-defined.

The trick is now in combining the above two procedures. For example, $w(\mathbf{x})\,(k \star \mathbf{H}_n)(\mathbf{x}) \cdot \mathbf{I}_n = w(\mathbf{x})\,(k \star (\mathbf{H}_n \cdot \mathbf{I}_n))(\mathbf{x}) = w(\mathbf{x})\,(k \star \delta)(\mathbf{x}) = w(\mathbf{x})\,k(\mathbf{x})$, which may not be what we want (as it breaks our model of the signal). On the other hand, $k \star (w\,\mathbf{H}_n) \cdot \mathbf{I}_n = k \star (w\,\mathbf{H}_n \cdot \mathbf{I}_n) = k \star (w\,\delta) = k$, as long as $w(0) = 1$. Of course the the convolution with k now interacts with the windowing function w, but this is expected to be less objectionable. Either procedure yields a kernel that is both band limited as well as spatially bounded. Note that spatially bounded should be taken to mean that the bulk of the kernel is concentrated in a bounded region (around the origin). In this work the convention $\mathbf{H}_n^* = k \star (w\,\mathbf{H}_n)$ will be used.

Note that when convolving the filter $\mathbf{H}_n$ with a sampled signal f represented using splines, the ideal procedure would involve computing $\mathring{\varphi} \star (w\,\mathbf{H}_n) \star \varphi \star f$ and sampling the result. Here φ is a reconstruction filter, while $\mathring{\varphi}$ is the comple-

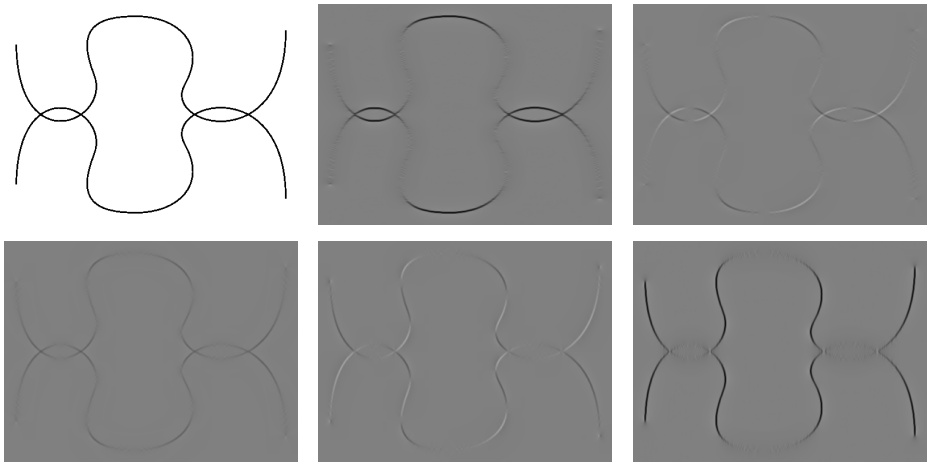

Fig. 2. Original image and images showing the corresponding tensor component images corresponding to: 1111, 1112, 1122, 1222, 2222. The tensor component images' intensities have been scaled, and subsequently adjusted so the background is grey, all images are rotated 90 degrees anti-clockwise. The inner product with the degree-four identity tensor corresponds to summing the 1111, 1122, and 2222 images, weighted by 1, 2, and 1, respectively; this recovers the original image.

mentary prefilter, as described by Unser [22]. Because of the commutativity of convolution it is equivalent to compute $\mathbf{H}_n^* \star f = k \star (w\,\mathbf{H}_n) \star f$, where $k = \mathring{\varphi} \star \varphi$. We now just need to sample the result on the same grid as f.

It should be noted that, perhaps counter intuitively, the width of w does not affect the integral of $w\,\mathbf{H}_n$. Instead it couples scale with orientation sensitivity: the narrower w is, the worse the orientation sensitivity for large features. It can be determined[1] that for a Gaussian windowing function $w(\mathbf{x}) = e^{-\|\mathbf{x}\|^2/(2\sigma^2)}$ – in a 2D context with degree-two tensors – setting σ to the width of an infinite line (box profile) results in a fractional anisotropy (FA) of roughly 0.55, while setting it to twice the width results in an FA of roughly 0.78, and three times the width corresponds to an FA of 0.86.

The filters used here have been generated by working in the Fourier domain[1], using numerical integration to incorporate the effects of the windowing function w (a Gaussian in our case, with $\sigma = 20$), and summation to implement (approximate) the effect of sampling in the spatial domain. The end result – which can be reused over and over – is an array representing $\hat{\mathbf{H}}_n^*$. Note that there should be some spatial aliasing, due to sampling the frequency domain (and a windowing function with infinite support), but this appears to be negligible. The process is illustrated in Fig. 1, some examples of filtered results are shown in Figs. 2 and 3.

A faster method to compute (not apply!) an analogue of $\hat{\mathbf{H}}_n^*$ is to go back to the observation that $\mathbf{H}_n$ can be built by differentiating the fundamental solu-

[1] Code available at `http://bit.ly/1xc3OC3`.

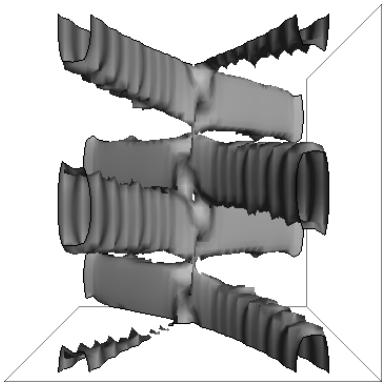

Fig. 3. Showing a contour of $\mathbf{U}_{4,f}$ for a section of the image in Fig. 2 (around the left-most crossing), with angle along the vertical axis. Note that each curve is present at two angles, corresponding to the even symmetry of a degree-four tensor. Also, the crossing can be seen to be less present at this contour, which is a result of the intensity of the curve being spread over both angles. It also interesting to note that since the curves are *curved*, their angle varies along the curve (as evidenced by the tapered sections at the top and bottom of the plot).

tion to the Laplace operator. Both differentiation and the Laplace operator have standard discrete approximations (using finite differences) that can be used instead of their continuous counterparts. This directly gives a closed expression for a spatially discrete kernel. So far this seems to give a visibly lower quality result than directly windowing and sampling $\hat{\mathbf{H}}_n$ though, so it has not been used here. It would, however, be interesting to look at this option further in the future, in the interest of simplifying the generation of tensorial orientation score kernels.

3.2 Interpretation of output

An examination of the (idealised) behaviour of the filters developed above can provide additional insight, as well as help in developing post-processing tools. In particular, it turns out tensor decompositions can play an important role in analysing the filter output.

First of all, if the original signal is constant, then the Fourier transform of this signal is clearly a Dirac delta at the origin, whose weight is equal to the constant. This can be easily verified by looking at the *inverse* Fourier transform of such a Dirac delta:

$$\mathcal{F}^{-1}[\boldsymbol{\xi} \mapsto c\,\delta(\boldsymbol{\xi})](\mathbf{x}) = \int_{\mathbb{R}^d} c\,\delta(\boldsymbol{\xi})\, e^{2\pi\, i\, \boldsymbol{\xi}\cdot\mathbf{x}}\, \mathrm{d}\boldsymbol{\xi} = c\, e^{2\pi\, i\, 0\cdot\mathbf{x}} = c.$$

The corresponding symmetric tensor field can be seen to be $\frac{c}{\mathbf{I}_n \cdot \mathbf{I}_n}\, \mathbf{I}_n$ for even n and zero otherwise.

If the input is zero except for a single infinite line through the origin, then the Fourier transform is zero except for a hyperplane through the origin, oriented perpendicularly to the original line. From Eq. (3) it is then clear that all tensors in the resulting tensor field are orthogonal to $\mathbf{v}^{\odot n}$, where $\mathbf{v}$ is the direction of the line. Furthermore, we can see that cross-sections of the tensor field that are perpendicular to the line correspond to a lower-dimensional version of the filter.

In general, an even feature of dimension e corresponds to an even degree tensor that has a zero response along the feature itself, and has a uniform response (like an identity tensor) along the other $d - e$ dimensions. This is illustrated in

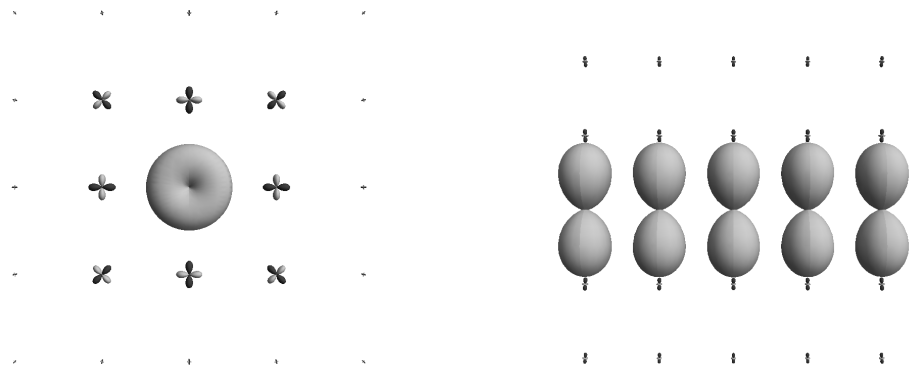

Fig. 4. Slices of two 3D orientation score tensor fields, visualizing the degree-2 tensors as spherical plots of $\mathbf{s}^{\odot 2} \cdot \mathbf{U}_2(\mathbf{x})$. The original image for the left tensor field had a single line (of finite length for this example), perpendicular to and through the middle of the slice. One can see that the middle tensor is effectively an identity tensor constrained to a 2D plane perpendicular to the line. The dark lobes of the other tensors correspond to negative responses (the inner product with the identity tensor should give the original value: zero). On the right we see the effect of a plane, resulting in rank one tensors on the plane (oriented perpendicularly to the plane).

Fig. 4. This makes it interesting to compute decompositions of such tensors into sums of rank one tensors, as well as components corresponding to (scaled) identity tensors on (higher dimensional) linear subspaces of the underlying vector space, as these components correspond directly to features of a certain dimension (similar in spirit to what was done by Tang and Medioni [21]). This strategy was used by van de Gronde et al. [14] to convert tensorial orientation scores into a sparse set of tangent vectors for further processing. For odd tensors it is a little harder to give a good characterization, but a similar pattern occurs (see Fig. 5).

3.3 Related methods

The above filters were inspired by the orientation scores developed by Duits [6], Duits and Franken [7, 8]. A fairly accessible discussion can be found in Franken and Duits [10], where the following kernel is proposed (in 2D):

$$\psi(\mathbf{x}) = \frac{1}{N}\mathcal{F}^{-1}\left[\omega \mapsto B^k\left(\frac{(\varphi \bmod 2\pi) - \pi/2}{s_\theta}\right) f(\rho)\right](\mathbf{x})\, G_s(\mathbf{x}).$$

Here N is a normalization constant, B^k denotes the k-th order B-spline[2], $\boldsymbol{\omega} = (\rho\cos(\varphi), \rho\sin(\varphi))$ (note that it is not entirely clear what convention was used for the Fourier transform), $s_\theta = \frac{2\pi}{n_\theta}$ (n_θ is the number of orientations), f ensures

[2] This convention does not guarantee that the different orientations sum to one. As an alternative, cardinal splines [22] could be used.

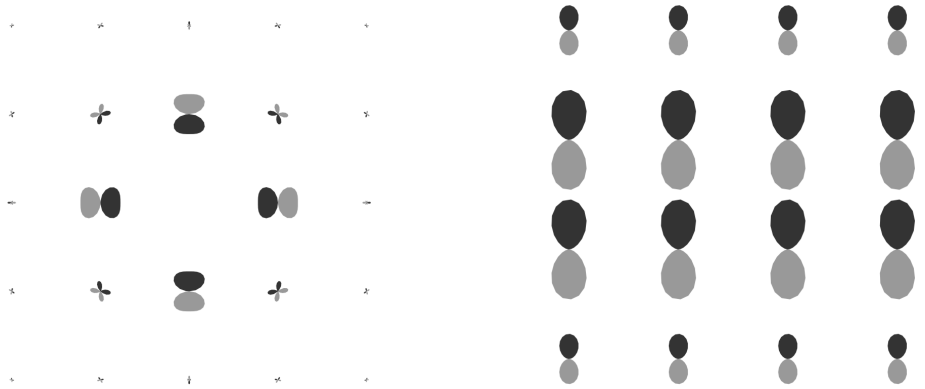

Fig. 5. Details of two 2D orientation score tensor fields, showing the degree-3 tensors as glyphs corresponding to a polar plot of $|\mathbf{s}^{\odot 2} \cdot \mathbf{U}_2(\mathbf{x})|$ (with dark lobes corresponding to negative values). The original image for the left tensor field had a single dot in the middle, while on the right the original had ones on the top half and zeroes on the bottom. Note how the "width" of the response is inversely related to the dimensionality of the feature.

the signal is band limited (a Gaussian divided by its Taylor series up to order q), and G_s is a Gaussian with standard deviation s. Clearly, f plays the role of the Fourier transformed sampling kernel $\hat{k}$, while G_s plays the role of the windowing function w (note though that here $w\,(k \star \mathbf{H}_n)$ is used). In essence, this filter is very similar to $\mathbf{H}_n^* \cdot \mathbf{s}^{\odot n}$, but sampling specific directions rather than building a tensor-based representation. Note that in 3D it is non-trivial to find a good set of orientations, and that the tensor-based representation in principle makes it easier to finely pinpoint a particular orientation (in any dimension).

Knutsson et al. [16] introduced monomial filters consisting of a directional part and a (bandpass) radial part, for the purpose of *analysing* local orientation in images. The directional part, given by $D_n(\omega/\|\omega\|) = (\omega/\|\omega\|)^{\odot n}$ can be seen to be exactly equal to $\hat{\mathbf{H}}_n$. However, this work shows that these filters can in fact also be used without applying a bandpass filter, and that doing so gives us an invertible transformation (like with orientation scores). Also, some more insight into the filter response is given: the filter can be seen as a (higher order) derivative of the fundamental solution to the p-Laplace equation, and it is shown how the filter responds to features of various dimensionalities, suggesting the potential of tensor decompositions to aid in extracting features from the output of the filter.

Finally, it is interesting to note that curvelets [3] also divide the spectrum into segments similar to an orientation score's "cake kernels", except that the segments are *also* divided radially. More precisely, given a suitable angular window ν and radial window w, for all integers j and l, such that $j \geq 0$ and $\ell = 0, 1, \ldots, 2^j - 1$, [3, Eq. 2.1]

$$\chi_{j,\ell}(\boldsymbol{\xi}) = w(2^{-2j}\,\|\boldsymbol{\xi}\|)\,(\nu(2^j\,\theta - \pi\,\ell) + \nu(2^j\,\theta - \pi\,(\ell - 2^j))).$$

	All-pass	Band-pass
Sampled	Orientation scores	Curvelets
Tensorial	Tensorial orientation scores	Monomial filters

Table 1. A summary of methods related to tensorial orientation scores. All-pass means that the filters have (virtually) no radial selectivity in the frequency domain, while band-pass filters select a specific band (or bands, in the case of curvelets). Sampled filters select a specific set of orientations to filter on, while tensorial filters build a tensor-based representation that describes the response for all orientations. Note that both types of orientation scores allow local (per-position) reconstruction. Monomial filters do not allow reconstruction, while curvelets form a tight frame.

Here $\boldsymbol{\xi}$ presumably equals $(\|\boldsymbol{\xi}\| \cos(\theta), \|\boldsymbol{\xi}\| \sin(\theta))$. Note that the amplitude of the filters defined by $\chi_{j,\ell}$ in the frequency domain does not depend on j (the radial level), giving overall scale invariant behaviour similar to orientation scores. Also, [3, Eq. 2.6]

$$|\chi_0(\boldsymbol{\xi})|^2 + \sum_{j \geq 1} \sum_{\ell=0}^{2^j - 1} |\chi_{j,\ell}(\boldsymbol{\xi})|^2 = 1,$$

illustrating that the entire spectrum is filled uniformly. Note that this equation does imply that curvelets have a different method of reconstruction from orientation scores: they form a tight frame rather than allowing purely local reconstruction. Also rather than selecting a single band like monomial filters, a complete representation is built from the responses in different bands. Table 1 summarizes the properties of the methods discussed here.

4 Conclusion and future work

A tensorial analogue of orientation scores has been presented. It is shown that the resulting filter of degree n can be interpreted in the spatial domain as a (tensorial) derivative of order n of the fundamental solution to the $n/2$-Laplace equation. Furthermore, despite this filter having a singularity at the origin and having an ill-defined integral, it is demonstrated that it can in fact be applied to real world signals, detailing the steps required in computing the discretized filter. Unfortunately this method is relatively complicated and slow (although once one has generated the filter it can of course be used many times). A different technique is suggested that is much simpler. However, it still has to be seen how far this simpler technique can come in terms of the quality of the result.

Tensorial orientation scores are compared to traditional orientation scores, monomial filters, and curvelets. It is argued that these four methods all share the same underlying idea, but with different combinations of features. Tensorial orientation scores represent a novel combination of features, in that they do not filter on frequency magnitude (like traditional orientation scores), but are based

on tensors rather than sampling specific directions (like monomial filters). Compared to curvelets, the method of reconstruction is different though: curvelets form a tight frame, while orientation scores allow local reconstruction. It is not immediately obvious which option is to be preferred, but it would definitely be interesting to see if a variant of tensorial orientation scores can be built that forms a tight frame.

One question about the currently proposed method (and to some extent monomial filters as well) is whether the discretization of the filter introduces bias (so features at certain orientations are better detected than features at other orientations). Similarly, a more extensive evaluation of the effect of making the filters spatially bounded would be quite interesting as well. Also, this work only considers using the developed filters to analyse an image and then possibly reconstructing an image (after having filtered the orientation score for example), but what about using the tensor components as basis functions to represent an image? Given the nature of the functions this might have some relation to the work done by Elder [9] and Orzan et al. [20].

Knutsson et al. [17, §3] suggested a scheme based on the gradient operator that could generate a family of monomial filters that is superficially similar to the family of curvelets, in that the bandwidths follow a geometric progression, and higher bandwidths allow for better orientation selectivity. It would be interesting to see if such a scheme could be used to develop a true tensorial analogue of curvelets.

References

[1] Bourbaki, N.: Topological vector spaces. Elements of Mathematics, Springer-Verlag (1987)

[2] Boyling, J.B.: Green's functions for polynomials in the Laplacian. Z. angew. Math. Phys. 47(3), 485–492 (1996)

[3] Candès, E.J., Donoho, D.L.: New tight frames of curvelets and optimal representations of objects with piecewise C2 singularities. Comm. Pure Appl. Math. 57(2), 219–266 (2004)

[4] Cheng, Antes, H., Ortner, N.: Fundamental solutions of products of Helmholtz and polyharmonic operators. Eng. Anal. Bound. Elem. 14(2), 187–191 (1994)

[5] Comon, P., Golub, G., Lim, L.H., Mourrain, B.: Symmetric Tensors and Symmetric Tensor Rank. SIAM J. Matrix Anal. Appl. 30(3), 1254–1279 (2008)

[6] Duits, R.: Perceptual organization in image analysis: a mathematical approach based on scale, orientation and curvature. Ph.D. thesis, Eindhoven University of Technology (2005)

[7] Duits, R., Franken, E.: Left-invariant parabolic evolutions on SE(2) and contour enhancement via invertible orientation scores Part I: Linear left-invariant diffusion equations on SE(2). Quart. Appl. Math. 68(2), 255–292 (2010)

[8] Duits, R., Franken, E.: Left-invariant parabolic evolutions on SE(2) and contour enhancement via invertible orientation scores Part II: Nonlinear left-invariant diffusions on invertible orientation scores. Quart. Appl. Math. 68(2), 293–331 (2010)
[9] Elder, J.H.: Are Edges Incomplete? Int. J. Comput. Vis. 34(2), 97–122 (1999)
[10] Franken, E., Duits, R.: Crossing-Preserving Coherence-Enhancing Diffusion on Invertible Orientation Scores. Int. J. Comput. Vis. 85(3), 253–278 (2009)
[11] Friedberg, S.H., Insel, A.J., Spence, L.E.: Linear algebra. Pearson Education (2003)
[12] Gel'fand, I.M., Shilov, G.E.: Generalized Functions, vol. 1. Academic Press, London (1964)
[13] Greub, W.H.: Multilinear algebra. Springer-Verlag (1978)
[14] van de Gronde, J.J., Lysenko, M., Roerdink, J.B.T.M.: Path-based mathematical morphology on tensor fields. In: Visualization and Processing of Higher Order Descriptors for Multi-Valued Data. Dagstuhl Follow-Ups, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2014), to appear
[15] Knutsson, H., Westin, C.F.: Monomial Phase: A Matrix Representation of Local Phase. In: Westin, C.F., Vilanova, A., Burgeth, B. (eds.) Visualization and Processing of Tensors and Higher Order Descriptors for Multi-Valued Data, pp. 37–73. Math. Vis., Springer Berlin Heidelberg (2014)
[16] Knutsson, H., Westin, C.F., Andersson, M.: Representing Local Structure Using Tensors II. In: Heyden, A., Kahl, F. (eds.) Image Analysis, LNCS, vol. 6688, pp. 545–556. Springer Berlin Heidelberg (2011)
[17] Knutsson, H., Westin, C.F., Andersson, M.: Structure Tensor Estimation: Introducing Monomial Quadrature Filter Sets. In: Laidlaw, D.H., Vilanova, A. (eds.) New Developments in the Visualization and Processing of Tensor Fields, pp. 3–28. Math. Vis., Springer Berlin Heidelberg (2012)
[18] Kostrikin, A.I., Manin, I.I.: Linear algebra and geometry, Algebra, Logic and Applications, vol. 1. Gordon and Breach (1997)
[19] Minc, H.: Permanents, Encyclopedia of mathematics and its applications, vol. 6. Addison-Wesley (1978)
[20] Orzan, A., Bousseau, A., Winnemöller, H., Barla, P., Thollot, J., Salesin, D.: Diffusion Curves: A Vector Representation for Smooth-Shaded Images. ACM Transactions on Graphics 27(3), 92:1–92:8 (2008)
[21] Tang, C.K., Medioni, G.: Inference of integrated surface, curve and junction descriptions from sparse 3D data. IEEE Trans. Pattern Anal. Mach. Intell. 20(11), 1206–1223 (1998)
[22] Unser, M.: Splines: a perfect fit for signal and image processing. IEEE Signal Processing Magazine 16(6), 22–38 (1999)