Published in the proceedings of GSI 2013.
The final publication is available at link.springer.com.

Frames for Tensor Field Morphology

Jasper J. van de Gronde and Jos B.T.M. Roerdink

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. We propose to apply our recently developed frame-based framework for group-invariant morphology to the problem of tensor field morphology. Group invariance (and particularly rotation invariance) have been, and are, motivated to be relevant for filtering tensor fields. This leads to the development of a rotation-invariant frame for tensors, which can be used to easily define rotation-invariant morphological operators on tensor fields. We also show how our method can be used to filter structure tensor fields.

1 Introduction

An image can be described as a function from positions to values. A value at a position can be thought of as a grey level, concentration, colour, depth, etc. In any case, the values can be considered separate from the position space, and the image is the only thing that links the two. However, it is becoming increasingly popular to work with tensor fields. In this case, the values are non-scalar, and intimately linked to the space of positions.

Typically, tensor fields describe things like flow, diffusion, and other physical processes. In other cases, like the structure tensor [7], they might describe the gradient magnitude or edge strength in an image. Unfortunately, it is not always straightforward to apply traditional image processing methods to tensor fields. In particular, just applying (non-linear) operators to the image “channels” corresponding to the tensor components often results in nonsensical results.

Burgeth et al. [2] already gave several conditions for morphological operations on tensor fields. In particular, they motivate that morphological operations on tensor fields should be invariant to rotations. We demonstrate how our recent framework [4] for group-invariant morphology can be applied to tensor fields.

2 Filtering tensor fields

Tensor fields present a challenge when it comes to applying morphological filters. But what exactly are the problems? In principle, we can easily filter the tensor components separately, but in practice the results are “weird”, as illustrated in Fig. 1. So what exactly goes wrong?

* This research is funded by the Dutch National Science Foundation (NWO), project no. 612.001.001.

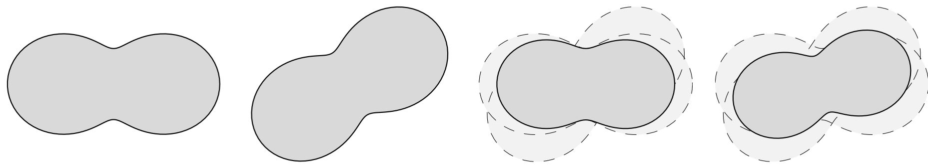

Fig. 1. Glyphs of the matrix (3 0 ; 0 1) and its 30°rotation $(5/2\ \sqrt{3}/2\,;\sqrt{3}/2\ 3/2)$ (the semicolons separate rows), followed by the component-wise minimum and the pseudo-meet derived from our method. The dashed lines indicate the shapes of the original two glyphs. The glyphs are polar plots of vMv^T, with $v = \big(\cos(\alpha)\ \sin(\alpha)\big)$, and M the matrix being plotted. While the component-wise minimum just shrinks the axis-aligned matrix a bit (not even within the confines of the other), our method results in a close fit to what is effectively the intersection of the two original glyphs.

A tensor field typically describes a process or effect. For example, it may describe how a fluid flows, or how a signal changes when moving around in space. When we rotate the tensor field, the tensors should thus be transformed in concert to still describe the same (but rotated) situation. Also, a filter should still act in the exact same way on the rotated signal. After all, our choice of orientation is (typically) arbitrary, and we could just as easily have chosen a different one. In other words, the filter should be invariant to rotations.

The importance of rotation invariance was already observed by Burgeth et al. [2]. The root of the problem is that mathematical morphology is based on lattice theory, and it can be shown [1, thm. XV.1] that one *cannot* define an appropriate lattice on tensors. Burgeth et al. chose to essentially forego the lattice-theoretic foundation, and define operators that have a qualitatively similar result. In contrast, by constructing a different representation, we *can* stay within the confines of established lattice theory. Only when it is necessary to go back to the original representation do we have to let go of the lattice-theoretic framework.

3 Definitions

3.1 Hilbert space

A Hilbert space is a vector space with a positive-definite inner product '$\cdot$'.[1] Any Hilbert space has an indexed set of vectors $\{\mathbf{e}_k\}_{k\in\mathcal{K}}$ that spans it, such that if you remove any of the vectors, the set no longer spans it. Such an indexed set is called a basis, and is associated with a dual basis $\{\mathbf{e}^k\}_{k\in\mathcal{K}}$ such that $\mathbf{e}^k \cdot \mathbf{e}_m$ equals one if $k = m$ and zero otherwise (for all $k, m \in \mathcal{K}$). Note that we will always work with Hilbert spaces over the reals, and that we denote vectors in a Hilbert space using bold face.

[1] Technically, a Hilbert space must also be complete with respect to the metric induced by the inner product, but this is not an issue here.

3.2 Transformation groups

A transformation on a (Hilbert) space V is a bijection between V and itself. A transformation group is a set of transformations that is closed under function composition '$\circ$', is closed under taking the inverse, and contains the identity mapping. Typically, we will consider groups of linear transformations and will not write '$\circ$' explicitly. Thus, given two elements τ_1 and τ_2 of a group $\mathbb{T}$ acting on V, and an element $\mathbf{a}$ of a Hilbert space V, $(\tau_1 \circ \tau_2)(\mathbf{a}) = \tau_1\,\tau_2\,\mathbf{a}$. A function f on V is considered invariant to $\mathbb{T}$ if for all $\tau \in \mathbb{T}$ and $\mathbf{a} \in V$, $f(\tau\,\mathbf{a}) = \tau\, f(\mathbf{a})$.

3.3 Tensors

Tensors can be viewed as a generalization of vectors and matrices. In general, tensors can be built from a vector space V using the associative tensor product '$\otimes$' (which is linear in both arguments). In terms of (column) vectors and matrices, one could say that $\mathbf{A} = \mathbf{a} \otimes \mathbf{b}$ $(\mathbf{a}, \mathbf{b} \in V)$ is equivalent to the matrix $\mathbf{a}\,\mathbf{b}^T$. We will use $\mathbf{a}^{\otimes n} \in V^{\otimes n}$ to denote the result of n-times repeated tensor multiplication of $\mathbf{a}$ by itself. Note that lower-case letters are used for vectors, while upper-case letters are used for (higher degree) tensors.

Tensors can be classified by their *degree* (or rank). A degree-zero tensor is a scalar (here a real), a degree-one tensor is a vector, and a degree-two tensor is a sum of elements of the form $\mathbf{a} \otimes \mathbf{b}$ (with $\mathbf{a}, \mathbf{b} \in V$). In general, a degree-$n$ tensor (with $n \in \mathbb{N}$) is a sum of tensor products of n vectors. A degree-n *symmetric* tensor is a sum of tensors of the form $\mathbf{a}^{\otimes n}$. The space of all degree-n tensors is denoted by $V^{\otimes n}$, the space of all degree-n symmetric tensors by $\mathrm{Sym}^n(V)$.

If V is a Hilbert space, then every $V^{\otimes n}$ can be considered a Hilbert space as well. The degree-zero and degree-one tensors trivially constitute a Hilbert space of course. However, higher degree tensors also form a Hilbert space, by making use of the inner product on V. Note that any degree-n tensor $\mathbf{A}$ in $V^{\otimes n}$ can be written as a sum of tensors of the form $\mathbf{a}_1 \otimes \mathbf{a}_2 \otimes \cdots \otimes \mathbf{a}_n$. The inner product of two such tensors can be computed as follows (with parentheses for clarity):

$$(\mathbf{a}_1 \otimes \mathbf{a}_2 \otimes \cdots \otimes \mathbf{a}_n) \cdot (\mathbf{b}_1 \otimes \mathbf{b}_2 \otimes \cdots \otimes \mathbf{b}_n) = \prod_{i=1}^{n} \mathbf{a}_i \cdot \mathbf{b}_i.$$

This is roughly equivalent to the Frobenius inner product on matrices.

We will also consider the tensor product of linear transformations. If τ_1 and τ_2 are linear transformations on a vector space V, then $\tau_1 \otimes \tau_2$ is a linear transformation on $V \otimes V$, such that $(\tau_1 \otimes \tau_2)(\mathbf{a} \otimes \mathbf{b}) = \tau_1(\mathbf{a}) \otimes \tau_2(\mathbf{b})$. It can be seen that the adjoint and inverse operations distribute over taking the tensor product, so $(\tau_1 \otimes \tau_2)^* = \tau_1^* \otimes \tau_2^*$ and $(\tau_1 \otimes \tau_2)^{-1} = \tau_1^{-1} \otimes \tau_2^{-1}$.

3.4 Tensor fields

In our context, a tensor field is a map $f : V \to V^{\otimes n}$. The idea is that the tensor field describes something that is happening in the underlying space. For example, the tensors in diffusion tensor imaging describe how water diffuses in different directions, and the structure tensor describes how much the image changes (locally) when we move in a particular direction.

If the underlying space is transformed by some linear transformation $\tau : V \to V$, so that $f(x)$ corresponds to $f'(\tau\, x)$, then the tensors should be transformed as well, in such a way that they still describe the same situation. So-called *contravariant* tensors simply transform like positions, thus $f'(\tau\, x) = \tau^{\otimes n} f(x)$. For *covariant* tensors we want to ensure that (for all $x, \mathbf{a}_i \in V$):

$$f'(\tau\, x) \cdot \big(\tau^{\otimes n}(\mathbf{a}_1 \otimes \mathbf{a}_2 \otimes \cdots \otimes \mathbf{a}_n)\big) = f(x) \cdot (\mathbf{a}_1 \otimes \mathbf{a}_2 \otimes \cdots \otimes \mathbf{a}_n).$$

Thus, we should have $f'(\tau\, x) = (\tau^{\otimes n})^{-*} f(x)$, which uses the transpose of the inverse of $\tau^{\otimes n}$. Note that if τ is orthogonal (like a rotation), the same transform is applied to both contra- and covariant tensors. Some examples: a velocity vector and a diffusion tensor are contravariant, the gradient of a function and the structure tensor [7] are covariant.

3.5 Frames

Generalizing the concept of a basis, a frame [3] is a set of vectors $\{\mathbf{f}_i\}_{i \in \mathcal{I}}$ (not necessarily finite or even countable) spanning a Hilbert space V, for which there are finite, positive constants A and B such that for any $\mathbf{a} \in V$

$$A\,\|\mathbf{a}\|^2 \leq \|F\mathbf{a}\|^2 \leq B\,\|\mathbf{a}\|^2.$$

Here the linear operator $F : V \to \mathbb{R}^{\mathcal{I}}$ is called the analysis operator, and is defined by $(F\mathbf{a})_i = \mathbf{f}_i \cdot \mathbf{a}$ for all $i \in \mathcal{I}$. The squared norm of $\mathbf{a} \in V$ is given by $\|\mathbf{a}\|^2 = \mathbf{a} \cdot \mathbf{a}$. Similarly, we take $\|F\mathbf{a}\|^2 = F\mathbf{a} \cdot F\mathbf{a}$. Obviously this requires the definition of an inner product on $\mathbb{R}^{\mathcal{I}}$. For simplicity, just assume that such an inner product exists, we will explicitly give the inner product where necessary.

The condition above is sufficient to ensure that there is at least one linear operator that acts as a left-inverse for F. In particular, there is a linear operator F^+ that minimizes $\mathbf{u} - F\hat{F}^*\mathbf{u}$ in a least-squares sense (with $\mathbf{u} \in \mathbb{R}^{\mathcal{I}}$).

4 A rotation-invariant frame

Since Euclidean space is invariant to rotation, we will first look at rotating the tensors in isolation, without rotating the underlying space. Afterwards, we will consider what can be done about rotating the underlying space.

Our applications only deal with symmetric[2] tensors. Traditionally, a $d \times d$ matrix is used for a degree-two symmetric tensor associated with $\mathbb{R}^d$, with the coefficients in the matrix being symmetric about the diagonal. In our approach, a basis is built using "tensor squares" of vectors. Different options are possible, for example (in 2D) $\{\mathbf{e}_1^{\otimes 2}, \mathbf{e}_2^{\otimes 2}, (\mathbf{e}_1 + \mathbf{e}_2)^{\otimes 2}\}$.

[2] Our method does not rely on positive definiteness, but typically does preserve it.

To create rotation-invariant operators for tensor fields we select a suitable basis and create a rotation-invariant frame using rotated copies of that basis. We then construct a rotation-invariant morphological operator on the frame representation (which is fairly easy). If desired we can also examine the, often optional, least-squares projection back to the original tensor space. Based on our previous work [4], we know that the overall result will be rotation invariant as long as the inner product on the frame representation is invariant to the transformations induced by rotations of the original vector space.

What kind of basis makes sense for symmetric tensors? In many cases (diffusion tensors, stress tensors, structure tensors, etc.), it is meaningful to compute the (tensor) dot product between a symmetric tensor and the tensor square of a direction vector. For example, this gives the apparent diffusion coefficient in the case of diffusion tensors, and the squared norm of the directional derivative in the case of structure tensors. A natural *dual* basis[3] for such symmetric tensors could thus be a set of tensor squares of (uniformly distributed) direction vectors.

If we start with a (dual) basis $\{\mathbf{E}^k\}_{k\in\mathcal{K}_n}$ for $\mathrm{Sym}^n(V)$ that contains only tensor powers of unit vectors ($\mathbf{a}^{\otimes n}$ with $\mathbf{a} \in V$), then making it invariant to rotations of the form $r^{\otimes n}$ (with r a rotation on V) results in a frame consisting of tensor powers of all vectors on the unit (hyper)sphere. The frame representation can thus be interpreted as giving a distribution over orientation (picture the glyphs in Figs. 1 and 3). More explicitly, there is a unit vector $\mathbf{v} \in V$ such that (with $\mathbf{A} \in \mathrm{Sym}^n(V)$ and r a rotation on V):

$$(F\mathbf{A})_{r,k} = (r^{\otimes n}\,\mathbf{A})_k = (r^{*\otimes n}\,\mathbf{E}^k) \cdot \mathbf{A} = \mathbf{v}^{\otimes n} \cdot \mathbf{A}.$$

Group-invariant operators on the frame representation can be defined by lifting an operator on the original space or by using group morphology [6, 8]. When lifting an operator we view the frame representation as consisting of many transformed copies of the original, and we apply the operator to each copy.

For projecting back in a least-squares manner, there are two choices for defining "least-squares" (and hence the projection): lift the original inner product on $\mathrm{Sym}^n(V)$, or define an inner product directly on the frame coefficients. The former results in a particularly easy back projection; if the analysis operator corresponds to taking all rotated versions of a tensor, then the backprojection operator consists of rotating all those tensors back to their original orientation and taking the average. The other option is slightly more involved, but one can always compute the Moore-Penrose pseudo-inverse of the analysis operator (using a finite number of rotations/vectors). *Not* projecting back is typically preferable, as it better preserves structure (and aids further processing).

We now have all the tools to construct a rotation-invariant frame and an appropriate backprojection technique, allowing easy definition of rotation-invariant operators on tensors. However, we are interested in tensor *fields*, so we should not just rotate the tensors, the grid has to be rotated in concert. Alternatively, the operator acting on it must be changed so that it acts like it was applied to the

[3] We are talking about a dual basis because it is the inner product with these vectors that is meaningful, not necessarily a weighted sum of these vectors.

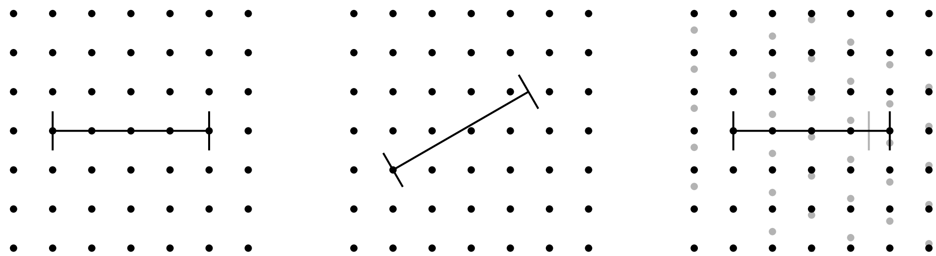

Fig. 2. The original structuring element has length five, and fits horizontally like in the left-most figure. At a thirty degree angle it is no longer aligned with the grid (middle figure) and interpolation is needed. Rather than rotating the image, we skew it (right). The image is skewed until the structuring element becomes horizontal again, and the structuring element length is rounded to an integer (the grey elements show the original grid and structuring element length).

rotated grid. Which of these options is easier depends on the application. In our case, we use structural openings and closings with line segments and adjust the structuring element to the current rotation angle, interpolating while filtering. This is illustrated in Fig. 2.

The basis of our method is that if linear interpolation is used, then extrema are always on grid positions. Thus, to compute the structural dilation/erosion on a rotated grid, it is sufficient to skew the grid and take samples on (skewed) grid positions. Only the end points of the structuring element need special care. We chose to simply round the length to an integer for each angle, which gives a fairly decent approximation (except for the smallest of kernel sizes).

5 Proof of principle

One possible use of morphological filters on tensor fields is in processing structure tensor fields. The structure tensor is derived from the gradient of an image and locally describes (the square of) the magnitude of the directional derivative for all directions. More formally, in the greyscale case, we have:

$$(\nabla f(x) \cdot \mathbf{a})^2 = (\nabla f(x))^{\otimes 2} \cdot \mathbf{a}^{\otimes 2} = \mathbf{T}(x) \cdot \mathbf{a}^{\otimes 2}.$$

The tensor $\mathbf{T}(x)$ is called the structure tensor at x.

Köthe [7] suggested several improvements to computing the structure tensor, including using a non-linear filter for smoothing along edges but not perpendicular to them. This filter essentially has a non-isotropic kernel that is aligned with the gradient at every position. We can do something similar, using a standard 1D morphological filter for every orientation[4], see Figs. 3 and 4.

A closing on the structure tensor field, as in Fig. 4c, helps to avoid a zero response at junctions, which could lead to trouble during further processing (segmentation, corner detection, etc.). A short line segment is used as the structuring

[4] Implementation available at `http://bit.ly/15MoLEI`.

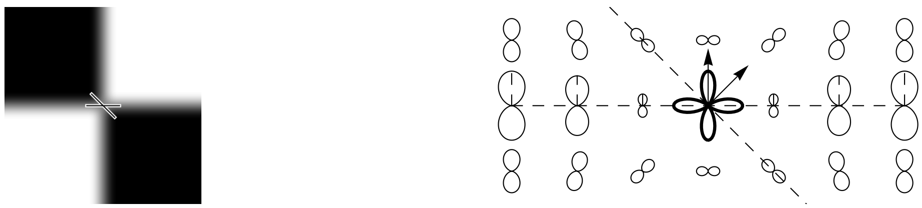

Fig. 3. A dilation of a structure tensor field (right) with a (centered) 1D structuring element. The original image is on the left. The glyphs show the frame representation as polar plots centered on the tensor's position. The thin glyphs show the input (which is zero at the center), the thick glyph shows the output at the center. For two angles, the contributing positions are indicated (corresponding to the two solid line segments in the left image).

element. When applying this filter to the frame coefficient corresponding to a certain orientation, the structuring element is oriented perpendicularly (Fig. 3). This is because the gradient magnitude is highest perpendicular to the edge.

There are several things in favour of our technique. For one thing, the technique used by Köthe [7] adapts to the local orientation using the gradient, and is thus problematic in regions where the gradient has a very small magnitude. Our technique on the other hand does not pick a certain orientation at each point, but simply processes all orientations. Also, our method opens up the possibility of using all sorts of other morphological filters. For example, rather than using openings and closings by line segments, it might make sense to use path openings and closings or attribute filters [5, 10].

6 Conclusion

Our method for constructing group-invariant lattices based on frames allows for the straightforward application of standard tools from mathematical morphology to tensor fields (in a meaningful manner). The practical feasibility and potential for application of the method is illustrated by an example.

Earlier methods by Burgeth et al. [2] already recognized the importance of rotation invariance, but tried to implement this directly on the original tensor space. Since it is not possible to define a rotation-invariant vector lattice on the original tensor space, this resulted in a loss of most properties that are taken for granted in traditional morphology. In contrast, our method relies on constructing a new representation that *does* admit a rotation-invariant vector lattice.

References

[1] Birkhoff, G.: Lattice theory, American Mathematical Society Colloquium Publications, vol. 25. American Mathematical Society (1961)

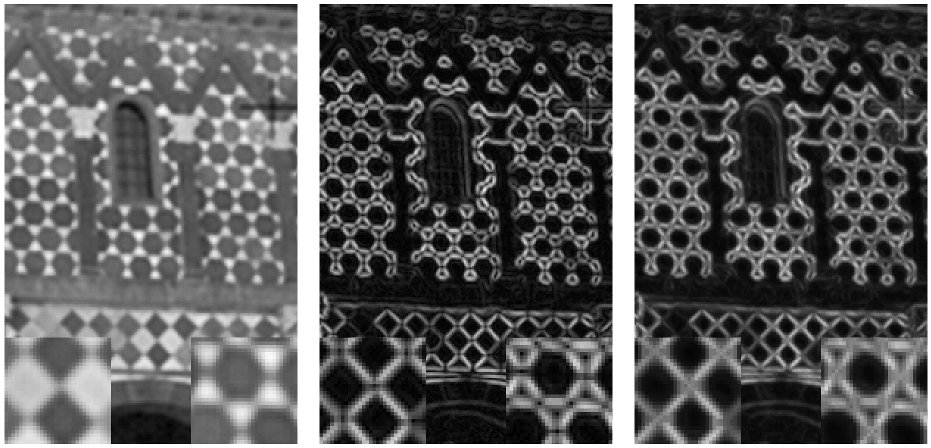

(a) Original (b) Before filtering (c) After filtering

Fig. 4. The original image (a), the gradient magnitude (b) and (c) the square root of the maximum frame coefficient after filtering (before filtering this corresponds to the gradient magnitude). Each image has two insets (on the bottom-left and -right) showing details. The (1D) closing (length 5) used here is applied in the same manner as the dilation in Fig. 3. Notice how the closing fills in the gaps at crossings, while keeping the edges thin.

[2] Burgeth, B., Bruhn, A., Didas, S., Weickert, J., Welk, M.: Morphology for matrix data: Ordering versus PDE-based approach. Image and Vision Computing 25(4), 496–511 (Apr 2007)

[3] Christensen, O.: Frames and Bases: An Introductory Course. Birkhäuser ; Springer e-books (2008)

[4] van de Gronde, J.J., Roerdink, J.B.T.M.: Group-invariant colour morphology based on frames. IEEE Transactions on Image Processing (submitted)

[5] Heijmans, H., Buckley, M., Talbot, H.: Path Openings and Closings. Journal of Mathematical Imaging and Vision 22(2), 107–119 (May 2005)

[6] Heijmans, H.J.A.M.: Morphological image operators. Academic Press (1994)

[7] Köthe, U.: Edge and Junction Detection with an Improved Structure Tensor. In: Michaelis, B., Krell, G. (eds.) Pattern Recognition, Lecture Notes in Computer Science, vol. 2781, pp. 25–32. Springer Berlin Heidelberg (2003)

[8] Roerdink, J.B.T.M.: Group morphology. Pattern Recognition 33(6), 877–895 (Jun 2000)

[9] Serra, J.: Anamorphoses and function lattices. Image Algebra and Morphological Image Processing IV 2030(1), 2–11 (Jun 1993)

[10] Wilkinson, M.H.F.: Hyperconnectivity, Attribute-Space Connectivity and Path Openings: Theoretical Relationships. In: Wilkinson, M.H.F., Roerdink, J.B.T.M. (eds.) Mathematical Morphology and Its Application to Signal and Image Processing, Lecture Notes in Computer Science, vol. 5720, chap. 5, pp. 47–58. Springer, Berlin, Heidelberg (2009)