Mathematical morphology in non-Euclidean spaces and medical images - Technical report

Samy Blusseau

May, 2017

1 Lattice on the cone of semi positive definite matrices

1.1 Notations

- S_n: the set of $n \times n$ symmetric real valued matrices
- S_n^+: the set of semi-positive definite (SPD) matrices
- S_n^{++}: the set of positive definite matrices
- I_n: the identity matrix
- $tr(M)$: the trace of matrix M
- M^T: transposition of matrix M
- $||M||$: the Euclidean norm of matrix M, for the canonical scalar product $\langle A,B\rangle = tr(A^TB)$
- $\forall \lambda \in \mathbb{R}$, $T_\lambda = \{M \in S_n, tr(M) = \lambda\}$: the set of symmetric matrices with trace λ

1.2 The Loewner ordering is not a lattice ordering

Definition and geometric interpretation The Loewner ordering is defined on S_n as follows: for any $A,B \in S_n$,

$$A \geq B \iff A - B \in S_n^+.$$

If $A \in S_n^+$, it can be geometrically represented by an ellipsoid $\mathscr{E}_A = \{x \in \mathbb{R}^n, x^TAx \leq 1\}$, the lengths of $\mathscr{E}_A$'s semi-axis being $1/\sqrt{\lambda_i(A)}$, if $\lambda_1, \dots, \lambda_n$ denote the eigenvalues of A (the ellipsoid can be degenerated, with infinitely long semi-axis when $\lambda_i = 0$). Then the Loewner ordering corresponds to a *reversed* inclusion ordering on ellipsoids: for $A,B \in S_n^+$,

$$A \geq B \iff \mathscr{E}_A \subseteq \mathscr{E}_B.$$

To get a more intuitive geometrical correspondance where the ellipsoid associated to the smaller element is included in the one associated to the bigger element, we can consider the ellipsoid whose semi-axis's lenghts are $\sqrt{\lambda_i}$. When $A \in S_n^{++}$ this corresponds to $\mathscr{E}_{A^{-1}}$, thus for $A,B \in S_n^{++}$,

$$A \geq B \iff A^{-1} \leq B^{-1} \iff \mathscr{E}_{B^{-1}} \subseteq \mathscr{E}_{A^{-1}},$$

and this can also be extended for non invertible matrices[1].

[1] I prefer not to spend more time on this since the Loewner ordering may not be what we need.

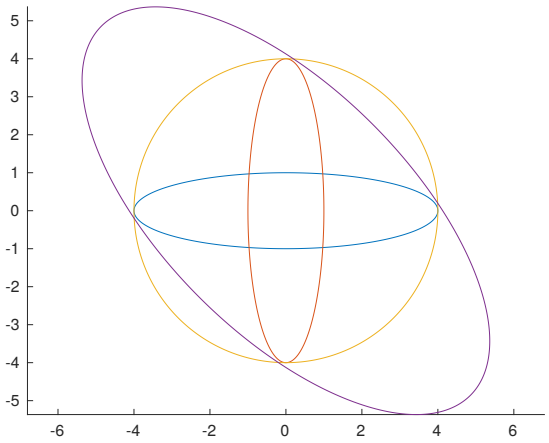

Figure 1: An example of minimal-volume enclosing ellipsoid (yellow), not comparable (in the Loewner sense) to another enclosing ellipsoid (in purple).

S_n^+ with the Loewner ordering is an anti-lattice An anti-lattice [4] is a poset in which, for any two elements A and B, a greatest lower bound exists if and only if A and B are comparable. In [4] Kadison shows that S_n with the Loewner ordering is an anti-lattice. Moreland and Gudder [5] proved that this is also true for S_n^+ (that is, a greatest lower bound exists *in* S_n^+ if and only if $A \geq 0$ and $B \geq 0$ are comparable), and Ando [1] generalizes the result to infinite dimension.

The misleading geometrical interpretation Despite the aforementioned results, one can argue that it is possible to find a unique smallest (largest) ellipsoid containing (contained in) a set of centered ellipsoids. The problem in this definition is that it refers to two different orderings: whereas the set of upper/lower bounds is defined by the Loewner ordering (or equivalently the inclusion ordering on ellipsoids), the smallest/biggest element of this set is defined by a *volume* total ordering. Unfortunately, the volume ordering does not induce the inclusion ordering, and one can find upper/lower bounds in the Loewner order that are not comparable with the "smallest"/"biggest" (in the volume sense) upper/lower bound - see Figure 1.

1.3 The (questionable) proposition of [2]

The authors of [2] acknowledge that the Loewner ordering is not a lattice ordering. However, the paper builds another ordering that is not proved to induce a lattice and that, in fact, seems to suffer from the same problem (the new order seems to be an anti-lattice as well).

In [2], computing the sup of a set of matrices in S_n^+ boils down to finding the smallest sphere covering a set of non centered spheres. As in the case of the ellipsoids, here two orderings are mixed: the inclusion ordering and the volume ordering. Again, we can find examples, as in Figure 2, where the resulting "smallest upper bound" is not comparable to other upper bounds in the ordering defined by the authors. These counter-examples seem to prove that this kind of ordering does not define a lattice. Indeed, if it did, the two blue spheres of Figure 2 would admit a unique sup. Since inclusion implies the ordering on volume, this sup would be the sphere of minimal volume enclosing the two blue spheres, that is the green sphere. The red sphere shows the existence of upper bounds not comparabale to the green one, and contradicts the existence of a unique minimal upper bound.

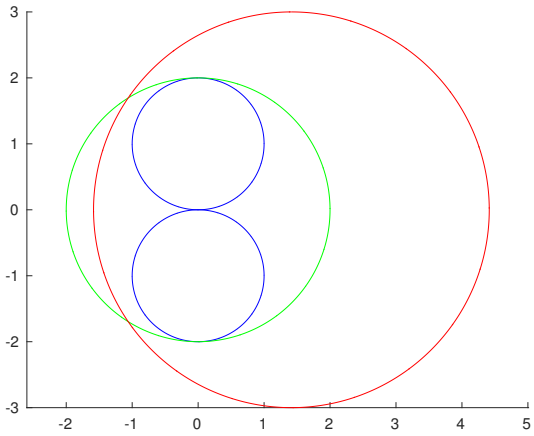

Figure 2: An example of minimal-volume enclosing ball (green), not comparable in the inclusion sense to another enclosing ball (red).

In the following I give a recap and reformulation of the construction presented in[2].

Bases of S_n^+ and their extreme points Figure 3 provides an illustration of the structure of the cone S_n^+, with the main features used for the construction of a new ordering in [2].

For any $\lambda \in \mathbb{R}$, we note $T_\lambda = \{M \in S_n, tr(M) = \lambda\}$. Then [2] recalls that $\mathscr{B}_1 := S_n^+ \cap T_1$, the set of positive matrices with trace 1, is a base of the cone S_n^+, that is to say: for any $M \in S_n^+, M \neq 0$, there exists a unique $\tilde{M} \in \mathscr{B}_1$ and a unique $\mu > 0$ such that $M = \mu\tilde{M}$. Indeed, here $\mu = tr(M)$ and $\tilde{M} = \frac{M}{tr(M)}$. Moreover, $\mathscr{B}_1$ is convex and its set of extreme points is known: it is the set of matrices $E_1 := \{vv^T, v \in \mathbb{R}^n, ||v|| = 1\}$. For a convex set $\mathscr{B}$, extreme points $ext(\mathscr{B})$ are defined as the points such that for any $x \in ext(\mathscr{B})$, $\mathscr{B} \setminus \{x\}$ is still convex, and they have the property to belong to the boundary of the convex set.

In this precise case, E_1 has another interesting property: its elements lie on the Euclidean sphere of centre $\frac{1}{n}I_n$ and radius $\sqrt{1-\frac{1}{n}}$: for any $M = vv^T \in E_1$, $||M - \frac{1}{n}I_n||^2 = tr((M - \frac{1}{n}I_n)^2) = tr(M^2) - \frac{2}{n}tr(M) + \frac{1}{n^2}tr(I_n) = tr(M) - \frac{1}{n} = 1 - \frac{1}{n}$.

Since extreme points are always included in the boundary of the convex set, and since $\mathscr{B}_1$ is a base of S_n^+, knowing $E_1 = ext(\mathscr{B}_1)$ gives a good idea of the general shape of the cone. Indeed, we can view $\mathscr{B}_1$ as some kind of convex polygone inscribed in the latter sphere, and then any "slice" $\mathscr{B}_\lambda = S_n^+ \cap T_\lambda$, $\lambda > 0$ can be deduced from $\mathscr{B}_1$ by homothety. We get that $\mathscr{B}_\lambda$ is also a base of S_n^+, its extremal points are exactly $\lambda \cdot E_1$, and they lie on the Euclidean sphere $\mathscr{S}_\lambda$ of center $\frac{\lambda}{n}I_n$ and radius $\lambda\sqrt{1-\frac{1}{n}}$. We note $\mathscr{C}_\lambda = \mathscr{S}_\lambda \cap T_\lambda$ the intersection of that sphere with T_λ.

The new cone [2] Just as $S_n^+ = \cup_{\lambda \geq 0}\mathscr{B}_\lambda$, a new cone C_n can be defined as $C_n = \cup_{\lambda \geq 0}\mathscr{C}_\lambda$. Although this is not explicit in the paper, the new ordering defined in [2] is the one induced by C_n:

$$A \geq B \iff A - B \in C_n.$$

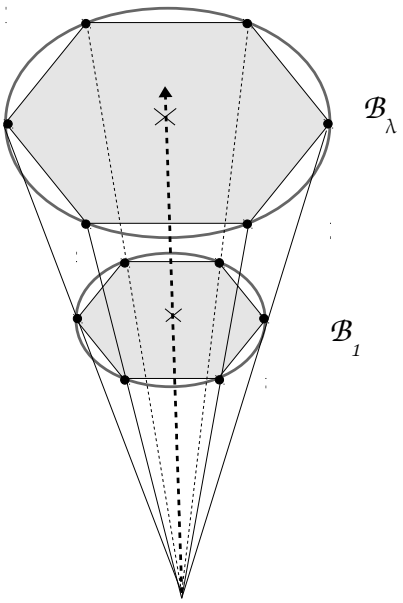

Figure 3: Illustration of the cone S_n^+. The gray polygones represent two bases of the cone, $\mathscr{B}_1$ and an arbitrary $\mathscr{B}_\lambda$ with $\lambda > 1$, intersections of the cone with the hyperplanes T_1 and T_λ. The extreme points of the bases are marked by black dots (there are infinitely many of them in reality), and the central dashed line indicates the span of the identity matrix I_n, orthogonal to each T_μ. On each base $\mathscr{B}_\mu$, the extreme points lie on a sphere $\mathscr{S}_\mu$ centered on $\frac{\mu}{n} I_n$ (marked by a black cross) and with radius $r_\mu = \mu\sqrt{1-\frac{1}{n}}$. Whereas S_n^+ is the union of all the gray polygones $\mathscr{B}_\mu$, the cone that defines the new ordering in [2] is the union of the $\mathscr{C}_\mu = \mathscr{S}_\mu \cap T_\mu$, $\mu \geq 0$.

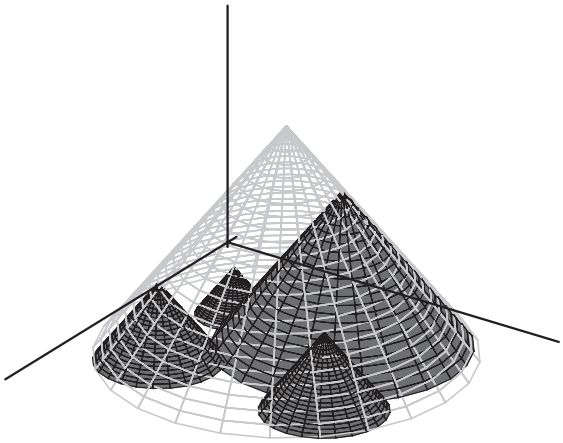

Figure 4: Examples of translated and reversed cones $M - C_n$ for several M. The light gray cone corresponds to the smallest ball $(M - C_n) \cap T_0$ containing all the others. Figure extracted from [2].

Readily, we have $S_n^+ \subseteq C_n$, which means that the Loewner ordering implies the new ordering. Furthermore, we get an easy characterization of C_n:

$$\begin{array}{rcl} M \in C_n & \iff & M \in \mathscr{C}_{tr(M)} \\ & \iff & ||M - \frac{tr(M)}{n} I_n|| \leq tr(M)\sqrt{1-\frac{1}{n}} \\ & \iff & tr(M^2) - \frac{1}{n} tr(M)^2 \leq tr(M)^2(1-\frac{1}{n}) \quad \text{and} \quad tr(M) \geq 0 \end{array}$$

$$M \in C_n \iff ||M|| \leq tr(M).$$

Recalling that a matrix $M = \begin{bmatrix} a & c \\ c & b \end{bmatrix}$ is in S_2^+ if and only if $a \geq 0, b \geq 0$ and $det(M) \geq 0$, we can see that the Loewner ordering and the new ordering are equivalent for $n = 2$: $S_2^+ = C_2$.

In practice: inclusion ordering on spheres The ordering in [2] is directly defined in terms of inclusion of spheres. Using the notations introduced above, it can be written as

$$A \geq B \iff (B - C_n) \cap T_0 \subseteq (A - C_n) \cap T_0 \iff B - \mathscr{C}_{tr(B)} \subseteq A - \mathscr{C}_{tr(A)}.$$

$(B - C_n) \cap T_0$ is the intersection of the "ground" plane T_0, and the reversed cone $-C_n$ whose vertex has been placed in B (see Figure 4). It is therefore the sphere of T_0 centered in $m_B = B - \frac{tr(B)}{n} I_n$ and of radius $r_B = tr(B)\sqrt{1-\frac{1}{n}}$. Similarly, $(A - C_n) \cap T_0$ is the sphere of T_0 centered in $m_A = A - \frac{tr(A)}{n} I_n$ and of radius $r_A = tr(A)\sqrt{1-\frac{1}{n}}$. By applying a criterion on the distance between centres, we can check that this definition of the new ordering, based on the inclusion of spheres, is equivalent to the one we gave earlier:

$$\begin{array}{rcl} A \geq B & \iff & ||m_A - m_B|| \leq r_A - r_B \\ & \iff & ||A - B - \frac{tr(A-B)}{n} I_n|| \leq tr(A-B)\sqrt{1-\frac{1}{n}} \\ & \iff & ||A - B|| \leq tr(A-B) \\ & \iff & A - B \in C_n. \end{array}$$

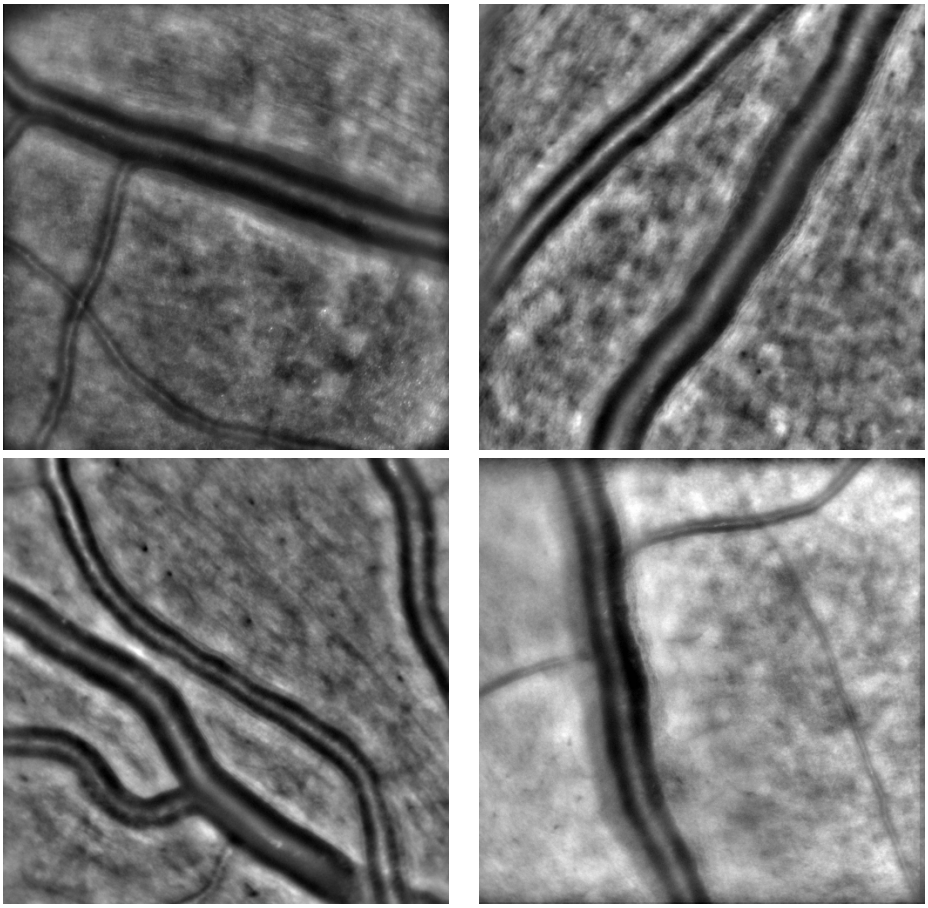

Figure 5: Example of input images.

Conclusion on [2] The paper proposes a new ordering but does not check whether it induces a lattice or not. From simple geometrical considerations it seems clear that it does not, and [6] asserts this ordering does not produce a lattice. I wonder if it is worth trying to show it is actually an anti-lattice.

This ordering and the definition of smallest upper bound may still be used as an approximation for morphological methods on S_n^+. The thesis [6] gives a quality measure of this approximation, I still need to have a look at it. However, it is not clear how it is better than the original Loewner ordering and the inclusion of centered ellipsoids. Therefore, if we decide to work with approximations of dilations and erosion, I would stay with the Loewner ordering.

2 Analysis of images of vessels based on structure tensors

2.1 Example of images

See Figure 5. In the following, we note f the image and suppose it is of size $N \times N$.

2.2 Structure tensor

The definition and computation of the structure tensor follows G. Peyré's numerical tour: `http://www.numerical-tours.com/matlab/pde_3_diffusion_tensor/`.

Gradient The following centered finite difference approximation of ∇f is used:$\nabla f(x,y) = (f_X(x,y), f_Y(x,y))^T$, where $f_X(x,y) = \frac{1}{2}(f(x+1,y) - f(x-1,y))$ for any y and $2 \leq x \leq N-1$, $f_X(1,y) = \frac{1}{2}(f(2,y) - f(N,y))$ and $f_X(N,y) = \frac{1}{2}(f(1,y) - f(N-1,y))$ (and similarly for $f_Y(x,y)$).

Tensor at scale zero The scale zero tensor T_0 maps each (x,y) to the symmetric matrix with rank ≤ 1

$$T_0(x,y) = \nabla f(x,y) \cdot \nabla f(x,y)^T = \begin{bmatrix} f_X(x,y)^2 & f_X(x,y)f_Y(x,y) \\ f_X(x,y)f_Y(x,y) & f_Y(x,y)^2 \end{bmatrix}.$$

It is straightforward that $\nabla f(x,y)$ is eigenvector of $T_0(x,y)$ with $||\nabla f(x,y)||^2$ as corresponding eigenvalue, and 0 is the other eigenvalue. Hence for any (x,y), $T_0(x,y)$ is an extreme point of the cone of positive semi-definite matrices S_2^+.

Tensor T_σ at scale $\sigma > 0$ It is a smoothed version of T_0, obtained by convolving the latter with the σ-scale Gaussian kernel

$$T_\sigma = G_\sigma T_0 = \begin{bmatrix} G_\sigma(f_X^2) & G_\sigma(f_X f_Y) \\ G_\sigma(f_X f_Y) & G_\sigma(f_Y^2) \end{bmatrix}$$

where G_σ is the smoothing operator. Note that, for any $u = (x,y)$, $T_\sigma(u)$ is a weighted sum of tensors $T_0(v)$ with positive weights, and is therefore a positive semi-definite matrix.

Eigen-decomposition of T_σ For each $u = (x,y)$, $T_\sigma(u)$ can be decomposed as

$$T_\sigma(u) = \lambda_1(u) e_1(u) \cdot e_1(u)^T + \lambda_2(u) e_2(u) \cdot e_2(u)^T$$

where $0 \leq \lambda_2(u) \leq \lambda_1(u)$ are $T_\sigma(u)$'s eigenvalues and $(e_1(u), e_2(u))$ its basis of orthogonal eigenvectors. If we represent $T_\sigma(u)$ by its corresponding ellipse

$$\mathscr{E}_u = \{X \in \mathbb{R}^2, X^T T_\sigma(u) X \leq 1\}$$

then the main direction of $\mathscr{E}_u$ is given by the second eigenvector $e_2(u)$. For $\sigma = 0$, this vector is orthogonal to the gradient $\nabla f(u)$; more generally, for small σ, Peyré points out the Taylor expansion

$$T_\sigma(u) = T_0(u) + \sigma^2 Hf(u)^2 + O(\sigma^3),$$

where Hf^2 is the Hessian of f^2. It is not clear what $e2(u)$ represents for larger σ, but it seems that inside a vessel the ellipse's main direction roughly coincides with the vessel's main direction.

Trace and anisotropy images (Figures 7 and 8) From the tensor field T_σ we can build two scalar images: the anisotropy image I_{an} and the trace image I_{tr}, respectively defined as

$$I_{an}(u) = 1 - \frac{2\lambda_2(u)}{tr(T_\sigma(u))} \quad \text{and} \quad I_{tr}(u) = tr(T_\sigma(u)) = \lambda_1(u) + \lambda_2(u).$$

We have $0 \leq I_{an} \leq 1$, and the closer $I_{an}(u)$ to 1, the more $T_\sigma(u)$ is anisotropic. I_{tr} can be seen as a smoothed version of the square norm of the ∇f.

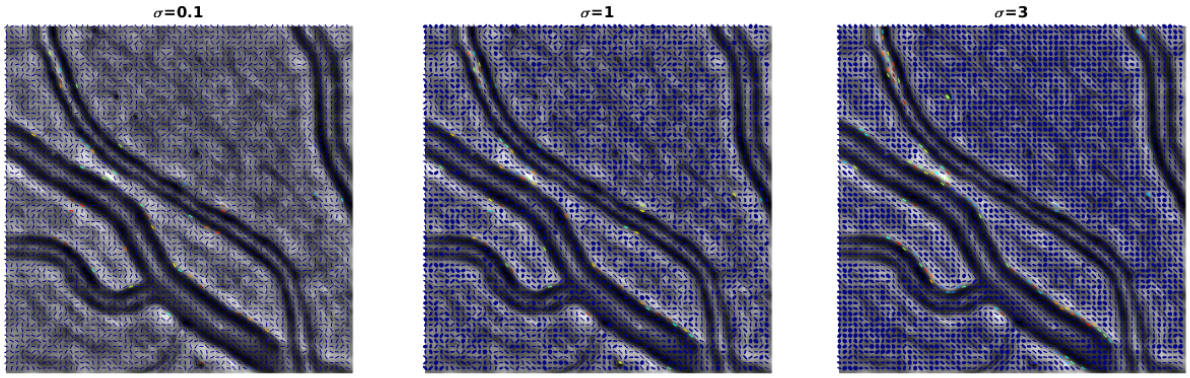

Figure 6: Structure tensors T_σ for the bottom left hand image of Figure, with $\sigma = 0.1$, $\sigma = 1$ and $\sigma = 3$. The shape of the represented ellipses show their anisotropy and main direction, whereas their colors must encode their trace (this is not clear yet).

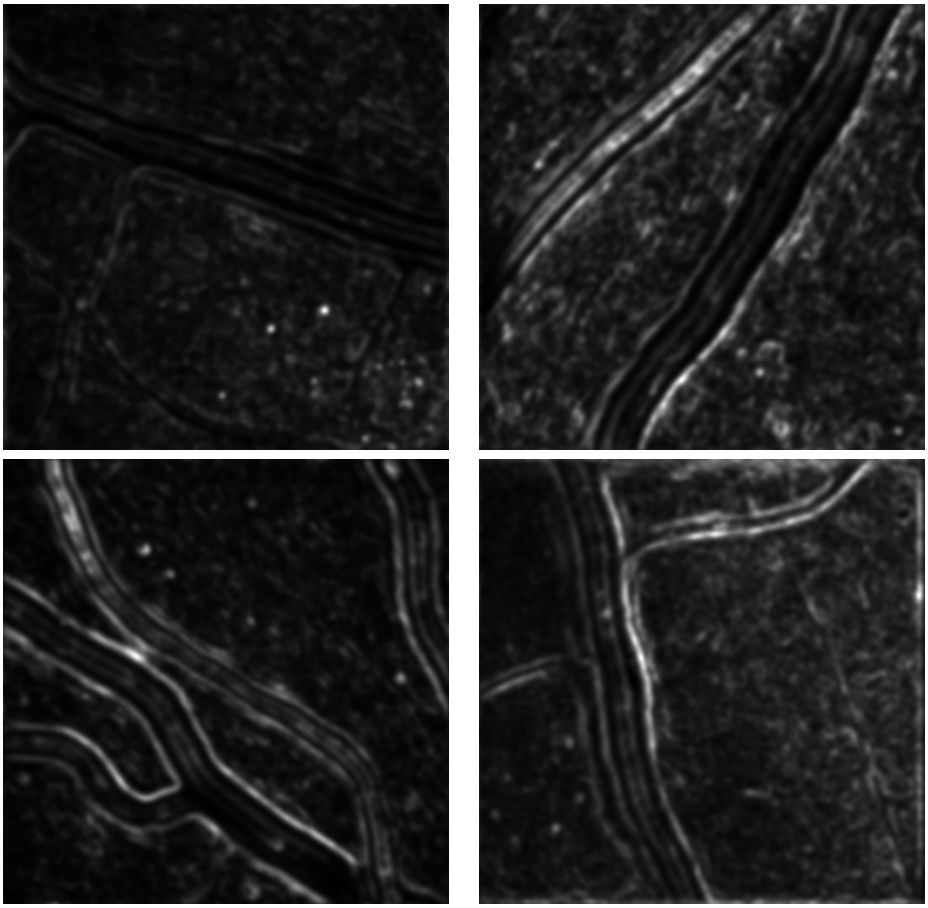

Figure 7: Trace images corresponding to the input of Figure 5, and the tensor field T_σ with $\sigma = 3$. They can be seen as smoothed version of the square norm of the ∇f.

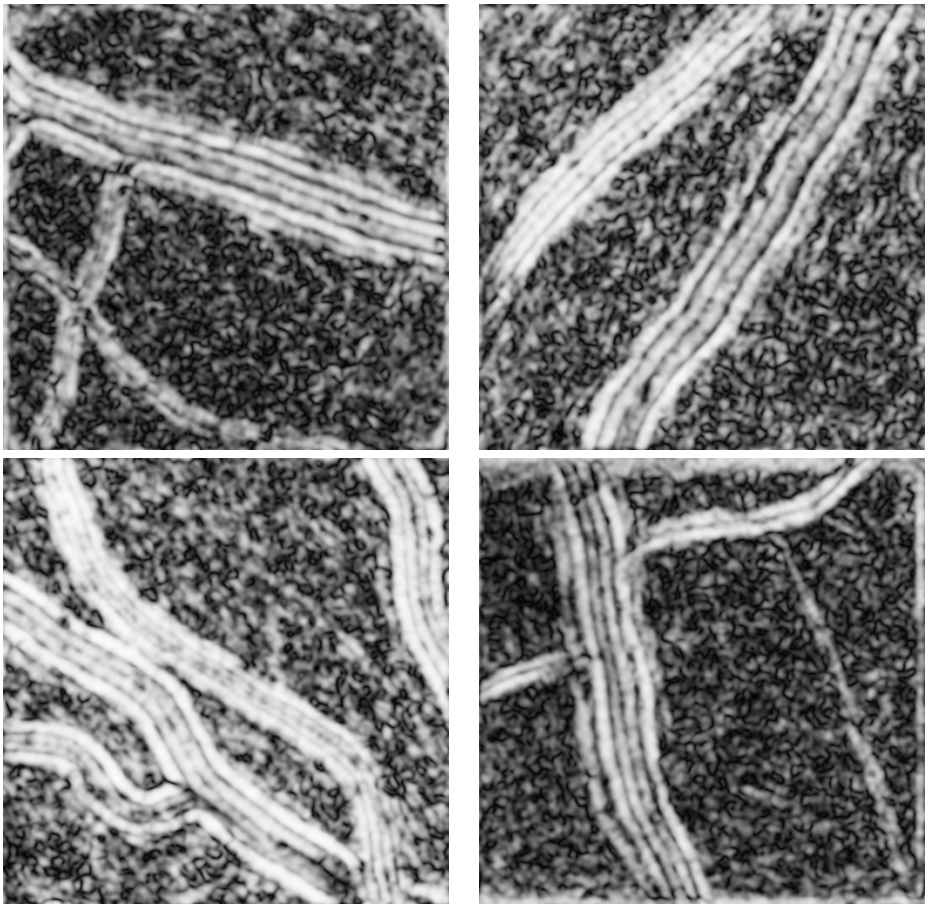

Figure 8: Anisotropy images corresponding to the input of Figure 5, and the tensor field T_σ with $\sigma = 3$. The closer $I_{an}(u)$ to 1, the more $T_\sigma(u)$ is anisotropic.

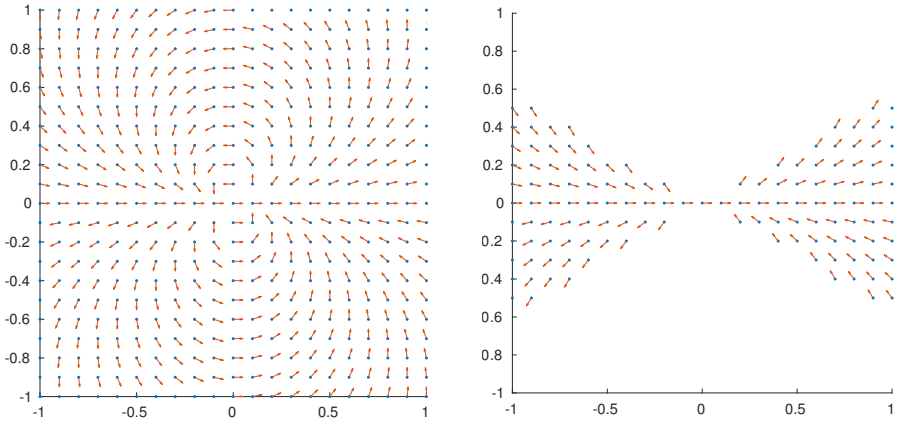

Figure 9: Left: a set of elements $(p, \overrightarrow{v})$ co-circular to the central element $(p_0, \overrightarrow{v_0})$ where p_0 is the origin and $v_0 = (1,0)^T$; Right: the elements form the left hand image that comply with the additional constraint that the (smaller) angle between the directions defined by $\overrightarrow{p_0 p}$ and $\overrightarrow{v}$ is smaller than $\frac{\pi}{6}$.

2.3 Structuring elements

Here we define for each pixel u the neighboorhood to be taken into account in the computation of flat dilation and erosion. The idea is to choose in a spatial window around u, those pixels v with for which the main orientation of $T_\sigma(v)$ is *consistent*, in some way, with $T_\sigma(u)$. So far I have used co-circularity as consistency criterion.

Co-circularity Given two points $p_1, p_2 \in \mathbb{R}^2$ and two vectors $\overrightarrow{v_1}, \overrightarrow{v_2}$, we say that $(p_1, \overrightarrow{v_1})$ and $(p_2, \overrightarrow{v_2})$ are co-circular if there is a circle tangent in p_1 and p_2 to $\overrightarrow{v_1}$ and $\overrightarrow{v_2}$ respectively. An infinite radius for the circle corresponds to the case when $\overrightarrow{p_1 p_2}$, $\overrightarrow{v_1}$ and $\overrightarrow{v_2}$ are colinear, which we consider a particular case of cocircularity.

Additional constraint As shown on the left of Figure 9, co-circularity allows "ladder" configurations, which we may want to discard in the following. If so, one can impose an additional constraint, namely that the angle between $\overrightarrow{p_1 p_2}$ and $\overrightarrow{v_2}$ is below a certain threshold.

Neighbourhood graph From the above we define an (undirected) graph (G,E) as follows: the set of nodes G is the set of pixels in the image $\{u_1, u_2, \ldots, u_{N^2}\}$; $(u_i, u_j) \in E$ iff u_j is in a square window of fixed size $2p+1$, centered on u_i (that is, $||u_i - u_j||_\infty \leq p$), and $(u_i, e_2(u_i))$ and $(u_j, e_2(u_j))$ are co-circular up to a certain angular tolerance, with the additional constraint described earlier to avoid ladder configurations - see Figure 9.

2.4 Morphological filters

Max-plus convolution Building on the graph (G,E) of the previous section, let W the $N^2 \times N^2$ adjacency matrix defined by $W_{ij} = 1 \iff (u_i, u_j) \in E$ or $i = j$, and $W_{ij} = 0$ otherwise. It encodes the structuring elements for any pixel, and we can now use it to compute max-plus convolutions (dilations

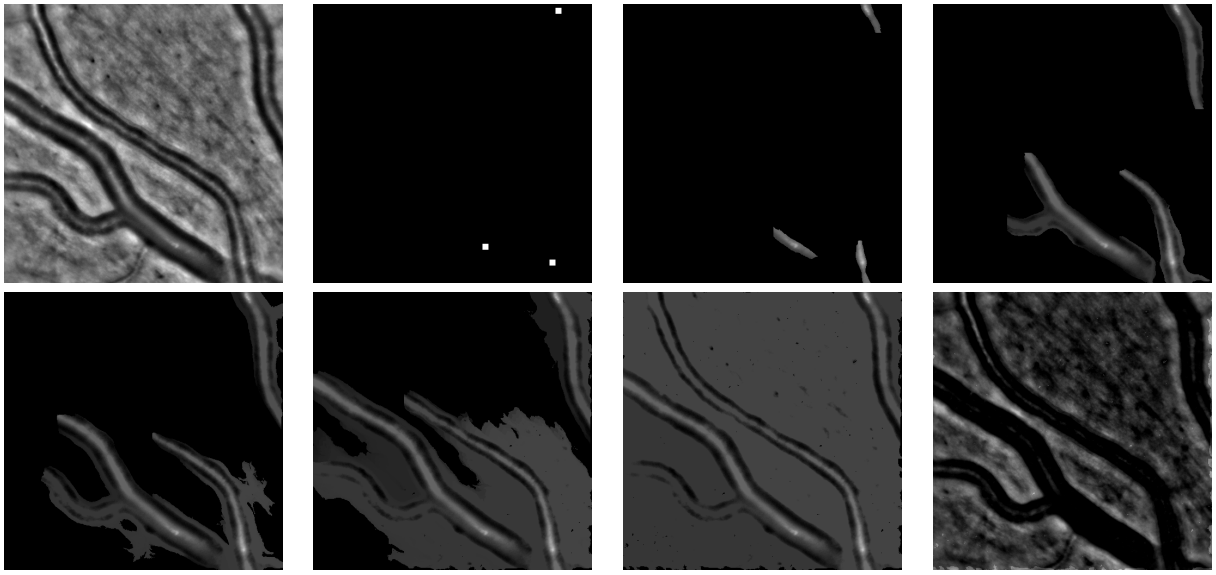

Figure 10: Illustration of the geodesic reconstruction by dilation. From left to right, top to bottom: mask image I (in this case, the original image); marker image R (defined manually); I_5, I_{25}, I_{35}, I_{50}, final reconstruction $\tilde{I} = I_{178}$; residual image $I - \tilde{I}$.

and erosions) on any sclar image I of size $N \times N$ [7]:

$$\delta(I)(u_i) = \bigvee_j \Big(I(u_j) + \log(W_{ij}) \Big) \quad \text{and} \quad \varepsilon(I)(u_i) = \bigwedge_j \Big(I(u_j) - \log(W_{ji}) \Big).$$

In particular, I can be the original image f, the anisotropy image or the trace image (Figures 5, 7 and 8).

Geodesic reconstruction by dilation We want to reconstruct a structure in an image I, e. g. a vessel, and possibly only this structure. If the structure is bright on a darker background, a possible strategy is to start from a marker image R representing a small part of the structure, and dilate it recursively under the constraint that the produced image remains smaller than I. By doing so, we hope to recover the bright structure as it is in I, while the background should remain relatively flat and smaller than it is in I.

More formally, given a mask image I and a marker image R, both $N \times N$, the geodesic reconstruction by dilation consists in building recursively a sequence of images $(I_n)_{n \geq 0}$,

$$I_0 = R \wedge I \quad \text{and} \quad I_{n+1} = \delta(I_n) \wedge I,$$

and finally take the sup $\tilde{I} = \bigvee_{n \geq 0} I_n$. Since in our case $W_{ii} = 1$, the dilation is extensive: $\delta(I)(u_i) \geq I(u_i)$, and the sequence $(I_n)_{n \geq 0}$ is increasing. Furthermore, $I_n \leq I$ for any n, hence the sequence converges to its maximal element $I_{n_{\max}}$ after a finite number $n_{\max}$ of iterations, producing the final reconstruction $\tilde{I} = I_{n_{\max}}$.

Figure 10 shows an example of such a reconstruction, in which the mask image I is an original image of vessels, and the markers have been defined manually to match bright regions in the vessels. The structuring elements (or equivalently the adjacency matrix W) was calculated on the tensor field T_σ associated with the original image, with $\sigma = 3$.

Ideas

- In the geodesic reconstruction, the mask image could be the anisotropy or trace image as well. Interestingly, they are independent on the contrast (it should work for bright vessels on dark background as well as for dark vessels on bright background). I have tried to work with the anisotropy but it is quite a noisy image and the reconstruction is not very accurate with respect to the original shape of the vessel - more work needs to be done on that. Experiments with the trace are on going as well.

- In the definition of the structuring elements, more information can be included than the spatial and angular consistency.I have tried to include the trace information, to avoid the existence of a path between pixels inside a vessel and pixels outside a vessel. So far this has not really worked out.

- The good part of working with the tensor field T_σ is that it provides positive semi-definite matrices, on which we can test methods to be adapted later on diffusion tensors. However, it raises some issues regarding the analysis of vessel images. First, it contains a scale parameter σ, which means that either a multi-scale approach or an automatic scale definition is required. Second, as said earlier, a structure tensor $T_\sigma(u)$ is merely a weighted sum of (rank 1) positive semi-definite matrices. It is not clear how this averaging behaves with respect to orientation information, especially in bifurcations.

- It seems that Frangi's vessel enhancement [3], based on the analysis of the Hessian, gives accurate segmentation of vessels. It may be interesting to figure out how our approach can add to Frangi's.

- As suggested in previous discussions, we should compare the results obtained with our definition of the structuring elements, to more classic ones (non-adaptive structuring elements, sets of differently oriented segments as structuring elements...).

References

[1] Tsuyoshi Ando. *Problem of Infimum in the Positive Cone*, pages 1–12. Springer Netherlands, Dordrecht, 1999.

[2] Bernhard Burgeth, Andres Bruhn, Nils Papenberg, Martin Welk, and Joachim Weickert. Mathematical morphology for matrix fields induced by the Loewner ordering in higher dimensions. *Signal Processing*, 87(2):277 – 290, 2007. Tensor Signal Processing.

[3] Alejandro F Frangi, Wiro J Niessen, Koen L Vincken, and Max A Viergever. Multiscale vessel enhancement filtering. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 130–137. Springer, 1998.

[4] Richard V. Kadison. Order properties of bounded self-adjoint operators. *Proceedings of the American Mathematical Society*, 2(3):505–510, 1951.

[5] T. Moreland and S. Gudder. Infima of hilbert space effects. *Linear Algebra and its Applications*, 286(1):1 – 17, 1999.

[6] Jasper Joris van de Gronde. *Beyond Scalar Morphology: A Theoretical Framework for Mathematical Morphology on Colour and Tensor-valued Images*. University of Groningen, 2015.

[7] Santiago Velasco-Forero and Jesús Angulo. *Nonlinear Operators on Graphs via Stacks*, pages 654–663. Springer International Publishing, Cham, 2015.