Graph-Based Mathematical Morphology (a survey)

ICPR 2016 tutorial

Sunday, December, 4rd, 2016

Laurent Najman   Hugues Talbot

Université Paris-Est - LIGM - A3SI - ESIEE Paris -France

Motivation

Mathematical morphology

- First consistent non-linear theory for image analysis
- Born in the late 60's
- Rely on the algebraic structure of complete lattice
- Adaptable to graphs

Graphs

- Generic data structures
- Independent of the dimension of the images
- A long history in all engineering fields
- Intrinsecally discrete, hence adapted to the digital world
- Current trend: adapt signal processing tools to graphs

Link between Graph and Mathemathical Morphology

A major inshight

The datum of a graph $G = (V, E)$ is equivalent to the one of a dilation. Precisely, the neighborhood operator Γ is a dilation

$$\forall v \in V, \Gamma(v) = \{u \in V \mid (v, u) \in E\}$$

- Vincent, 1989 - Signal Processing
- Heijmans & Vincent, 1992, (in a book)
- Cousty, Najman, Dias, F., Serra, J., 2013 - CVIU

Outline

Outline

Outline

Outline

Outline

Outline

Outline

What is graph

- **A graph is a representation of a set of data where sc pairs of data are connected by links**
  - **The data are called vertices or nodes**
  - **The links are called edges**

- **A graph can be weighted**
  - **on vertices**
  - **on edges**
  - **on both edges and vertices**

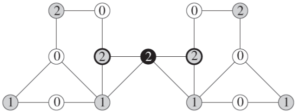

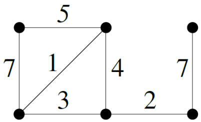

How to build a graph

- Regular grids

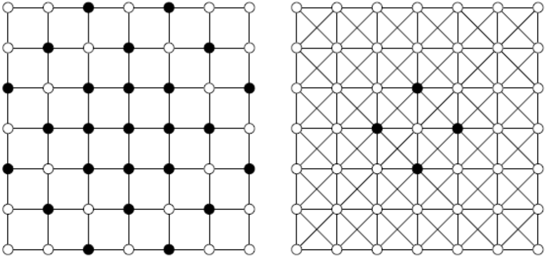

(a) 4-adjacency grid (b) 8-adjacency grid

- Topological issues
  - Jordan curve theorem
  - Thickness of boundaries
  - Merging of adjacent regions

How to build a graph

- Regular grids

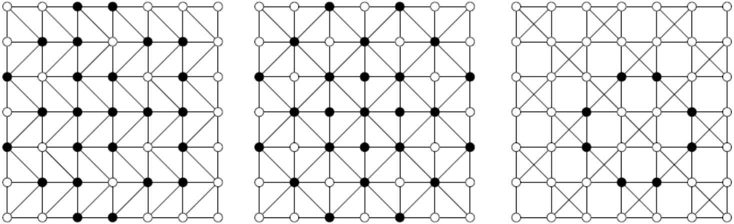

(c) 6-adjacency grid (d) Khalimsky grid (e) Perfect fusion grid

- Topological issues
  - Jordan curve theorem
  - Thickness of boundaries
  - Merging of adjacent regions

How to build a graph

- Spatially-variant grids

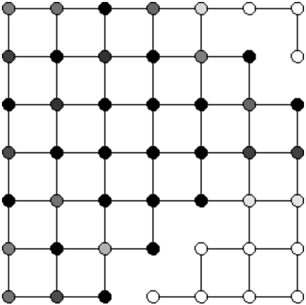

Removing high-gradient edges

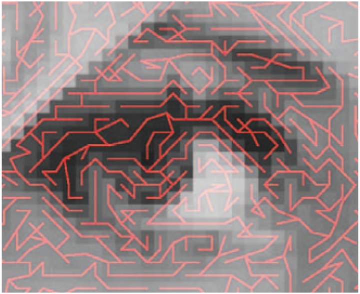

Keeping the edges of a minimum spanning tree

How to build a graph

- Spatially-variant grids

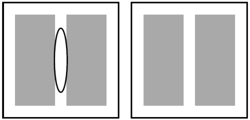

Fig. 3. Closing of an image by an amoeba. The amoeba does not cross the contour and as such preserves even the small canals.

Amoebas: neighborhoods as balls of geodesic distance

Nearest neighbors in a feature space

How to build a graph

- Spatially-variant grids

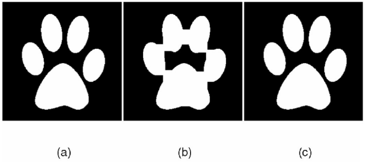

Fig. 2. Clustering Sets: (a) The original image X illustrates five separate objects which expanded by ψ yield (b) the sets making up the cluster. (c) By intersecting the connected components of $\psi(X)$ with X, the operator $\Gamma_x^{\psi}(X)$ extracts the cluster of the previously disconnected objects.

Second order connectivity

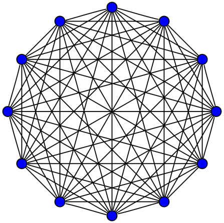

Complete graph (as in non-local means)

How to build a graph

- Region adjacency graphs

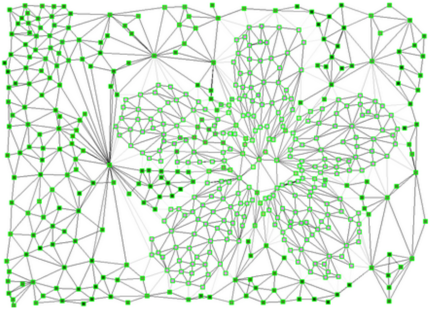

How to build a graph

- Triangular meshes and alpha-shapes

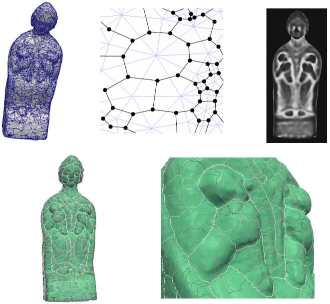

3D surface

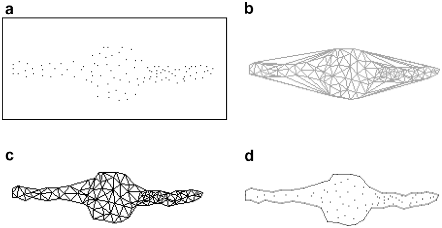

Fig. 2. (a) 2D point set S; (b) Delaunay triangulation *Del(S)* or infinite-complex*(S)*; (c) *opt*-complex (a a simplicial complex) *C_opt(S)* triangulated by *Del_opt(S)*; (d) *opt*-shape as a polytope *S_opt(S)*.

Cloud points

Discrete calculus

Important idea

- *Differential calculus can be written in an algebraic form.*
- *Exact solution on a graph can be computed (no approximation)*

Combinatorial Continuous Max-Flow [Couprie *et al.*, SIAM 2011]

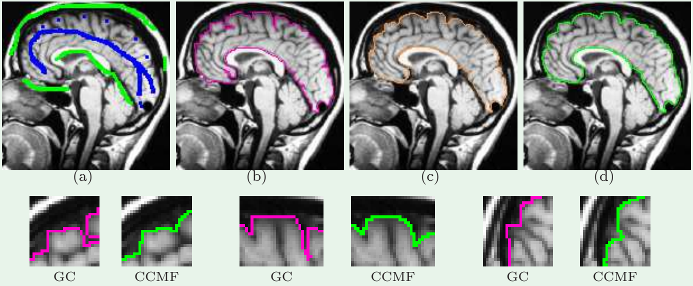

Discrete calculus

Important idea

- *Differential calculus can be written in an algebraic form.*
- *Exact solution on a graph can be computed (no approximation)*

Combinatorial Continuous Max-Flow [Couprie *et al.*, SIAM 2011]

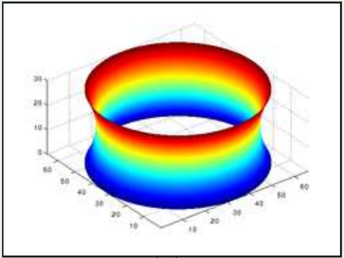

(a)

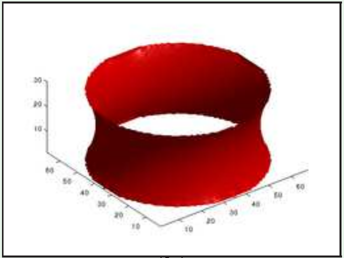

(b)

Discrete calculus

Important idea

- *Differential calculus can be written in an algebraic form.*
- *Exact solution on a graph can be computed (no approximation)*

Combinatorial Continuous Max-Flow [Couprie *et al.*, SIAM 2011]

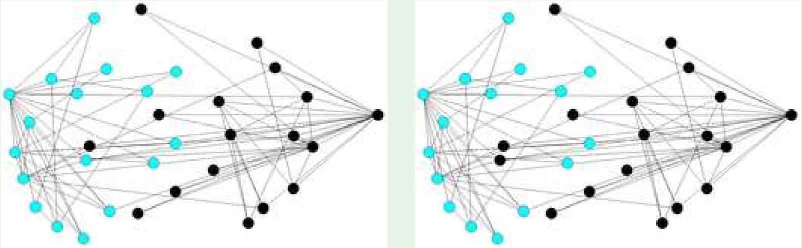

(a) Network and true membership after split

(b) Classification using CCMF

Discrete calculus

Important idea

- *Differential calculus can be written in an algebraic form.*
- *Exact solution on a graph can be computed (no approximation)*

Dual-Constrained Total-Variation [Couprie *et al.*, SIAM 2013]

Noisy

NonLocal TV

NonLocal DCTV

Discrete calculus

Important idea

- *Differential calculus can be written in an algebraic form.*
- *Exact solution on a graph can be computed (no approximation)*

Dual-Constrained Total-Variation [Couprie *et al.*, SIAM 2013]

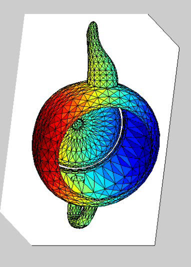

(a) Original mesh

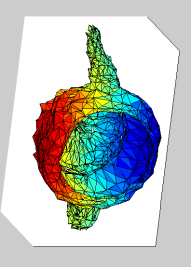

(b) Noisy mesh

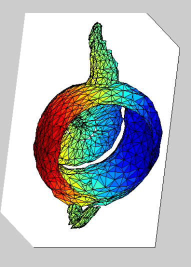

(c) DCTV regularization

Outline

Thin objects filtering

Thin objects are fragile

- Thin objects are much longer than they are wide
- In 3D, this include line-like and plane-like objects
- No isotropic neighborhood fits everywhere in them
- This means that spatially-invariant filtering cannot work in this case.

Vessel reconnection

Vessel connectivity

- Connectivity is crucial for vessel identification and classification (i.e. vein, artery). We need this information for instance for pre-op planning.
- However noise causes disconnections and denoising typically is not enough to reconnect.
- So we need to be more "forceful".

Spatially-variant morphology (SVMM)

Reconnection must be spatially variant

- A natural idea for a morphologist might be to use openings or closings for reconnecting disconnected vessels.
- However, using standard morphology with a spatially invariant structuring element will not work

(a) Noisy vessels (b) SI closing (c) SV closing

SVMM definition

SV and adjunction

- Defining a SV erosion or dilation is *easy*
- Defining their SV adjunct dilation or erosion (resp.) is *not so easy*

$$\forall x \in \mathcal{L}, \forall y \in \mathcal{M},\ \delta(x) \leq y \Longleftrightarrow x \leq \varepsilon(y)$$

- We still have $\delta_B(I) = \bigvee_{p \in B} I_p$ adjunct to $\varepsilon_B(I) = \bigwedge_{p \in \check{B}} I_p$
- However, we need to consider the full definition of the transpose of a SE

$$\check{B}(x) = \{y \mid x \in B(y)\}, \qquad (1)$$

- In the SI case, $\check{B}(x) = -B(x)$, but not in the SV case.
- It is possible to *compute it* but inefficient in the SVMM case.
- Fortunately we have an alternative.

Adjunct erosion computation

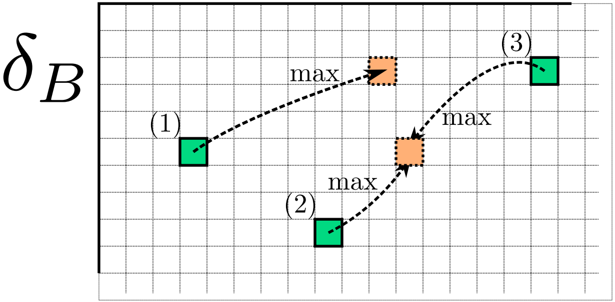

Figure: Spatially variant dilation

This is easy to compute.

Adjunct erosion computation

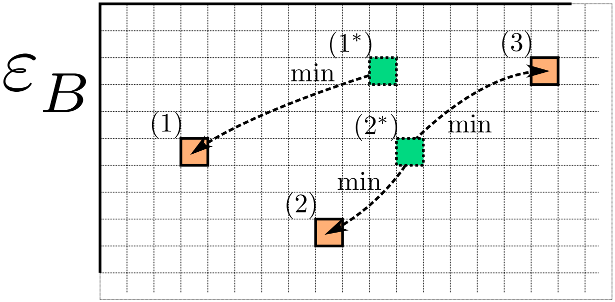

Figure: Spatially variant adjunct erosion with computed SE

This can be expensive to compute.

Adjunct erosion computation

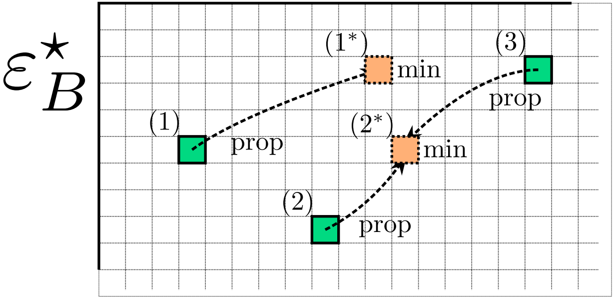

Figure: Spatially variant adjunct erosion alternative definition

This operator is equivalent to the adjunct erosion, and is as efficient to compute as the initial dilation.

More formally: morphology on graphs

- Let (E, Γ) be a graph with vertices E and oriented edges Γ (a.k.a arcs). if x is a vertex, we denote $\Gamma(x)$ its successors in the graph.
- Let $S \in E$ be a subset of E, then

$$\varepsilon_\Gamma(S) = \{\Gamma(x), x \in S\} \tag{2}$$

- Let ψ be an operator on (E, Γ), then we define the dual of ψ for any subset S of E, as $\psi^\star(S) = \overline{\psi(\overline{S})}$, where $\overline{S}$ is the set complement of S.
- Then, the adjunct of ε_Γ is:

$$\delta_\Gamma = \varepsilon^\star_{\Gamma^{-1}}, \tag{3}$$

where (E, Γ^{-1}) is the symmetric graph of (E, Γ), i.e. where all the edge orientations have been reversed.

End of formalities

- This extends the standard erosions and dilations, which correspond to Γ being a regular, reflexive, symmetric graph.
- E.g, with E arranged in a regular square grid, Γ the 4-connected reflexive connectivity, this defines the standard erosion / dilation pair with the 4-connected neighborhood.
- Arbitrary structuring elements are defined by the equivalent graph connectivity.
- Openings and closings are defined as usual:

$$\gamma_\Gamma = \delta_\Gamma \circ \varepsilon_\Gamma \quad \text{(opening)} \tag{4}$$

$$\phi_\Gamma = \varepsilon_\Gamma \circ \delta_\Gamma \quad \text{(closing)} \tag{5}$$

- Grey-level operators are formally defined by threshold decomposition, but implemented efficiently with a max or min operator.

A tubular objects filtering procedure

Filtering pipeline

- Filter the image to eliminate noise with an efficient NLM implementation (MPI + GPU , 5s for a $200 \times 200 \times 200$ image).
- Detect tubular objects using Frangi's vesselness
- Reconnect vessels with a spatially variant closing.

No problem in theory, however to reconnect vessels we require a dense direction field.

Direction field

Dense direction field

We need to:

- Estimate vessel directions from the Hessian eigenvectors
- Robustify these directions by sampling them near the center of the vessels
- Dilate the direction field
- Perform the SV closing with a segment oriented along these directions

SV closing illustration

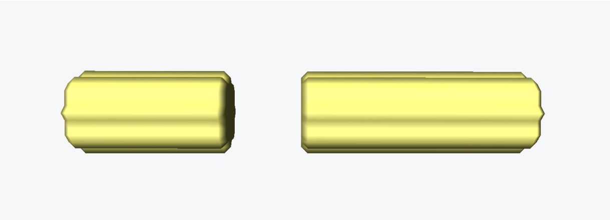

SV closing illustration

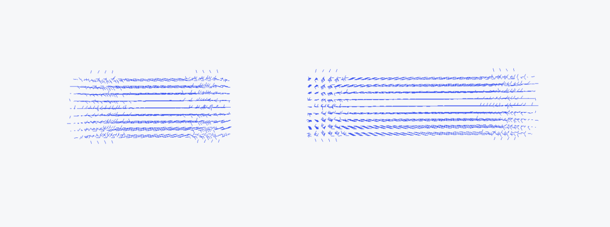

SV closing illustration

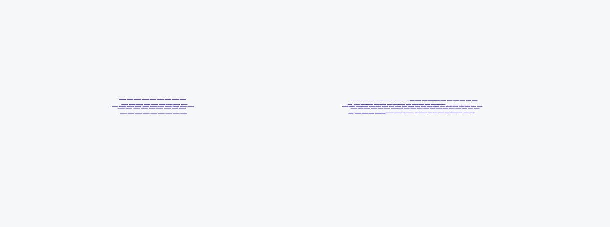

SV closing illustration

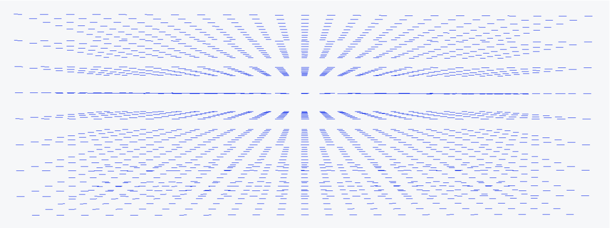

SV closing illustration

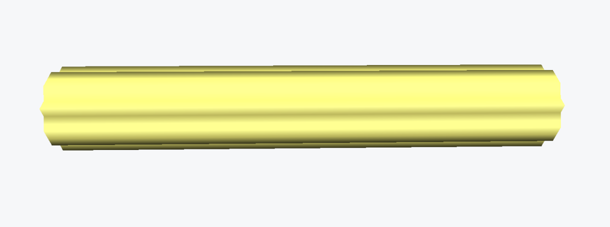

Visual results

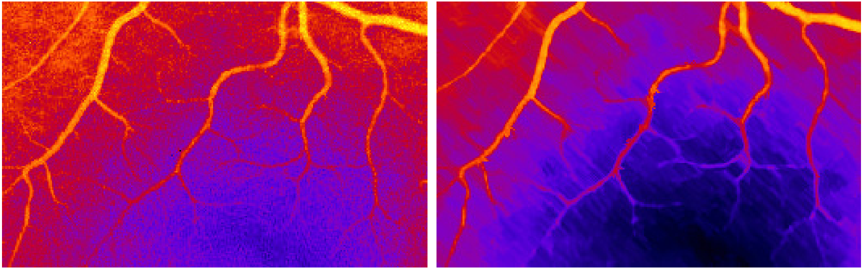

Figure: Eye fundus filtering

Visual results

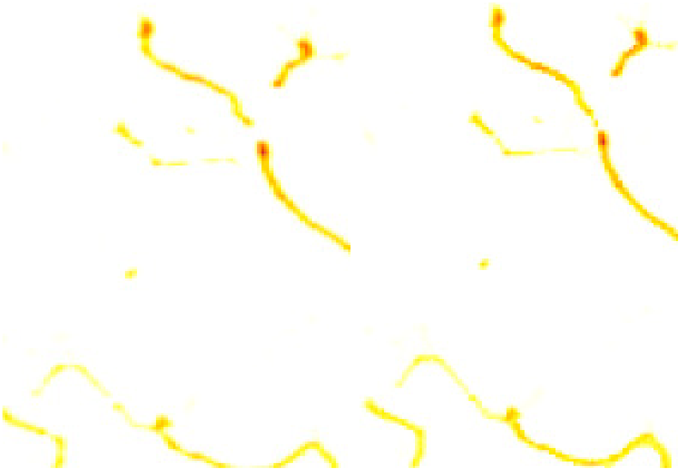

Figure: Neurite filtering

Visual results

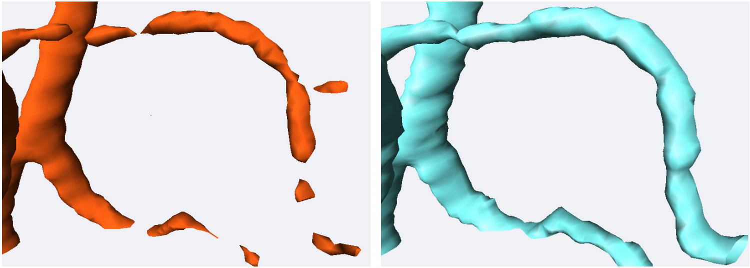

Figure: 3D image of vessels in the brain

Phantom for validation

Description and origin

- We use a phantom from [8], which is a $100 \times 100 \times 100$ image used in a MICCAI workshop.
- It is tortuous and vessel-like
- grey-level with a parabolic intensity from 200 at the center to 150 at the edges. The background is 100, consistent e.g. with TOF MRA.
- In the following ROC analyses, the triangle indicates best fully-connected result.

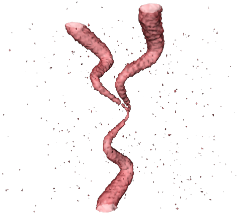

Validation

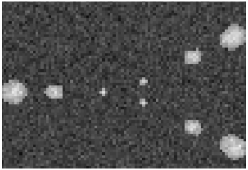

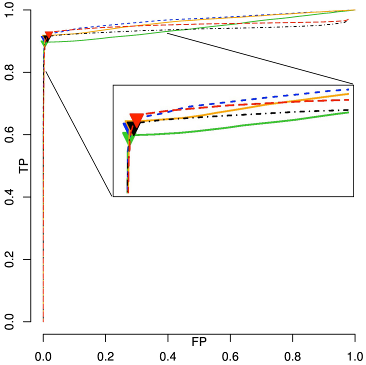

Figure: Level of noise standard deviation $\sigma = 10$

Validation

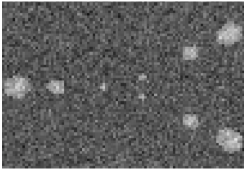

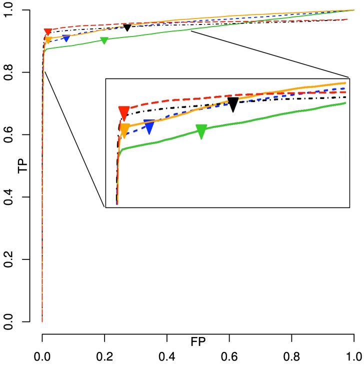

Figure: Level of noise standard deviation $\sigma = 20$

Validation

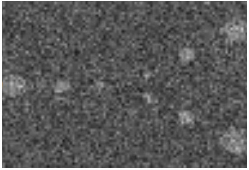

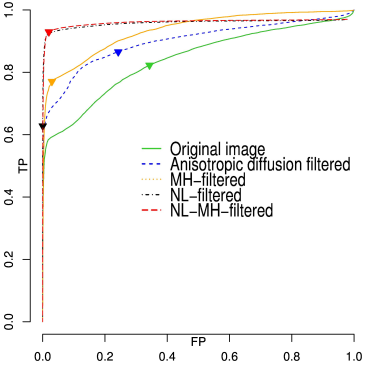

Figure: Level of noise standard deviation $\sigma = 40$

Validation

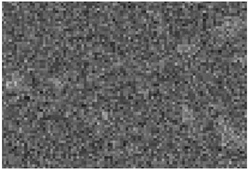

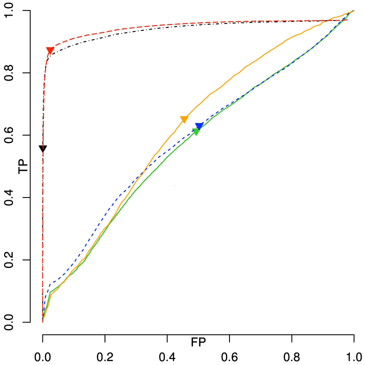

Figure: Level of noise standard deviation $\sigma = 80$

Notice that the filtered phantom remains connected even at very high noise levels.

Discussion

Remarks

- Noise reduction achieved with non-local approaches, orientation measured by vesselness, reconnection achieved by Spatially Variant morphology.
- Combining noise reduction techniques with morphology allows us to achieve extremely robust results for thin object detection

Publications

This work is described in greater detail in [25], as well as [26, 22, 23, 24, 6, 7, 17].

Discussion and challenges

- This pipeline is effective but requires the tuning of a number of parameters ;
- It requires significant hardware to be sufficiently fast
- Vessel detection is limited by the vesselness measure, which is not very effective
- It still needs to be evaluated on larger dataset, e.g. full brain vascular network, but annotated data is difficult to obtain.

VIVABRAIN project

Objective: Vascular network analysis from brain MRA data

- **Filtering**
  Improve images (Denoising, contrast enhancement)
- **Segmentation**
  Detecting the vascular network
- **Post-processing**
  Reconnexion, quantitative data analysis: directions, diameter, vessel density ...)

Filtering aspect

A new filtering method to improve existing segmentation pipeline

2 complementary axes :

- Noise reduction
- Vascular network contrast enhancement

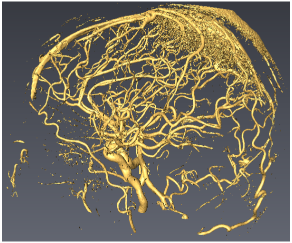

3D MRA data surface rendering

Maximum intensity projection

Classical tubular segmentation

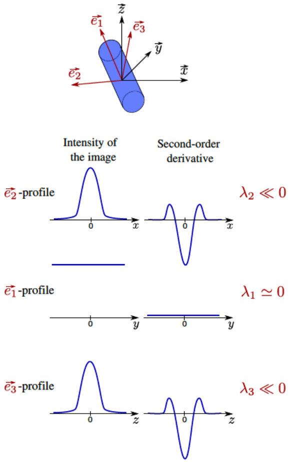

Figure: Classical approach using the Hessian

Errors due to scale-space

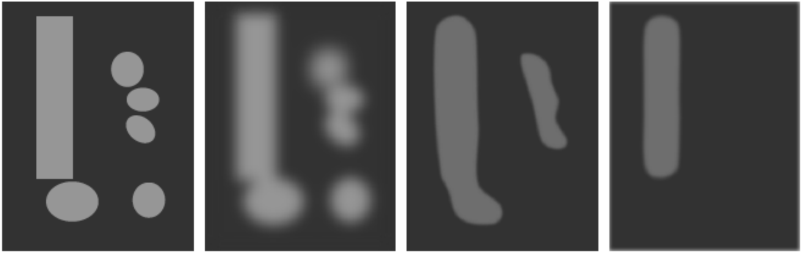

(a) Initial image (b) Gaussian filter result (c) Hessian-based filter result (d) Expected result

Figure: Scale-space filtering problem

Errors due to locality

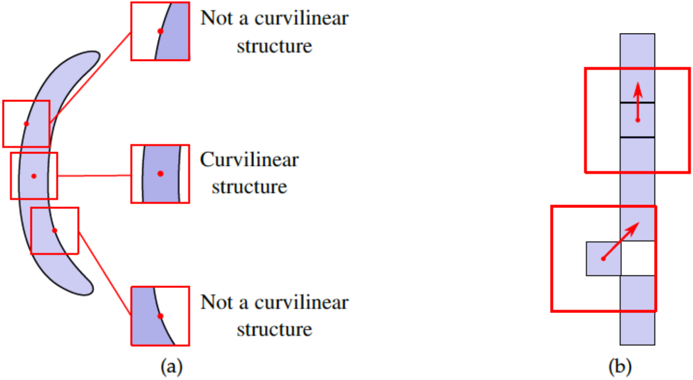

Scale space methods use local neighborhoods and are susceptible to misinterpretation at some scales.

Proposed solution

- Scale selection and combination is a challenging problem in traditional scale-space methods.
- One solution is to use semi-local neighborhoods, i.e. that gather information over long distances at all scales.
- We propose the use of paths.

Adjacency graph

A path, **a**, is a set of neighboring pixels on a graph defining an adjacency relation $x \rightarrow y$:

$$\mathbf{a} = (a_1, a_2, ..., a_L) \quad si \ a_k \rightarrow a_{k+1}$$

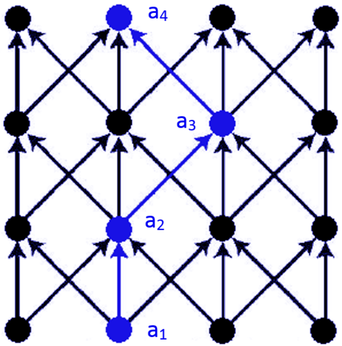

Adjacency graph (black) and vertical path **a** *of length 4 (blue).*

Multiple orientations

Filtering of an image by a path opening

Preserving thin structures in arbitrary orientations imposes to filter the image by several paths each using a particular adjacency graph.

The 2D space is discretized in 4 different orientations :

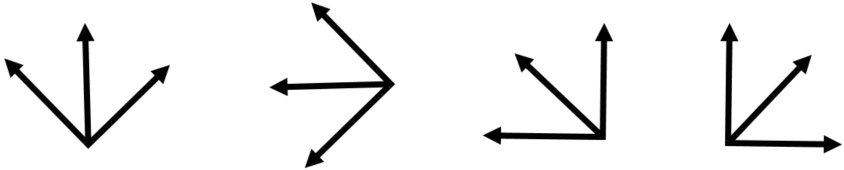

Multiple orientations in 3D

In 3D, the discrete space is discretized in 7 different orientations :

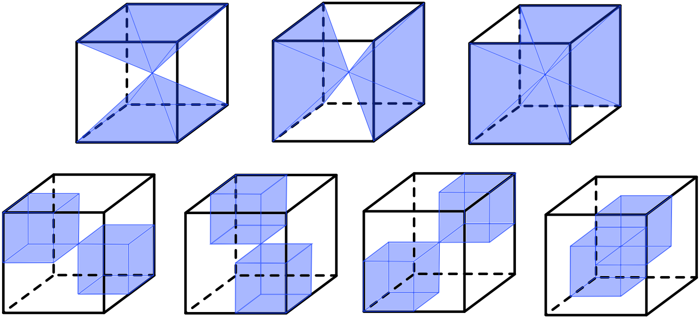

Path filtering

Example binary path opening

$$\alpha_L = \bigvee \{\sigma(\mathbf{a}), \mathbf{a} \in \Pi_L(X)\}$$

σ_L : Set of all pixels belonging to path **a**.
$\Pi_L(X)$: Set of all paths of length L.

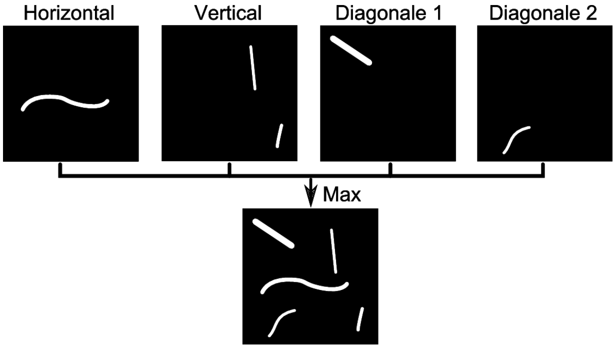

Principle

Path definition relaxation

A path can now admit K consecutive noise pixels between path pixels

This makes it possible to preserve partially disconnected thin/tubular structures :

Path with $L = 10$ *and* $K = 1$ *noise pixel*

This notion is different from that of *path incompleteness* by Heijman et al, it was proposed by F. Cokelaer [5] and is simpler to implement.

Example

RPO Example on a synthetic, noisy 2D image (centered AWGN $\sigma = 20$)

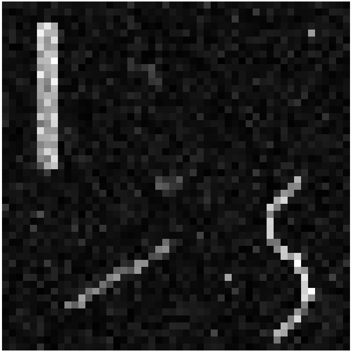

Initial image 50x50px

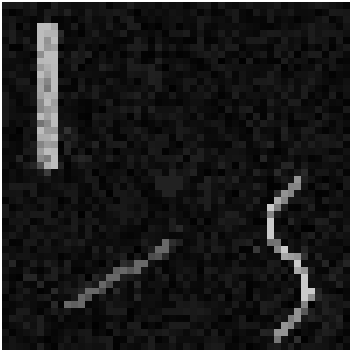

RPO L=10, K=1

The 3D case is more complicated than 2D

2D Case

2 Types of structures :

Fibres and Blobs

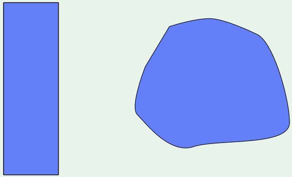

RPO preserves only fibres if blobs are not too big.

3D Case

3 Types of structures :

Tubes, Planes and Blobs

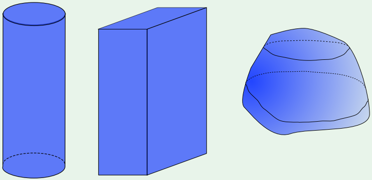

RPO preserves both tubes and planes.

An RPO by itself preserves more than just tubes in 3D images.

Principle

Hypothesis

Planar structures should be detected in at least one more orientation than tubular structures

Test of this hypothesis on 3 synthetic structures :

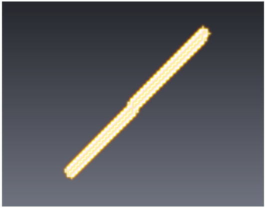

Tube

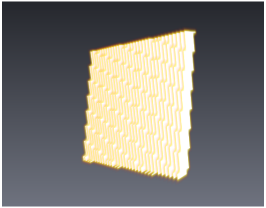

Plane

Half-ellipsoid surface

Hypothesis testing

Test :

Filtering 100 3D images of each structure and measuring the number of RPO orientations still containing the structure after filtering

Histogram of the number of orientations preserving the structure:

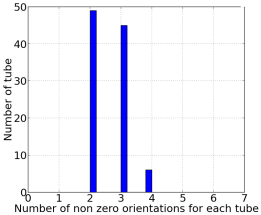

Tubes

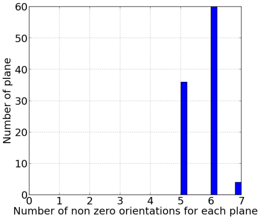

Planes

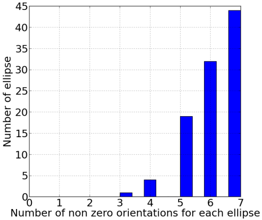

Half-ellipsoids

Methodology

New operator

We order the result of each RPO orientation pixelwise and compute

$$\text{RORPO} = RPO_1 - RPO_i$$

RPO_1 : Result of standard RPO (max of all RPOs)
RPO_i : The $i-th$ rank of the RPO.

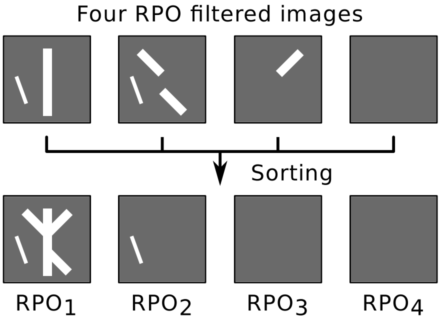

Robustness test

We compute the RORPO error rate on 100 random synthetic structure of each type.

$$\%error = \frac{nb_{error}}{nb_{pixels}} \times 100$$

nb_{error}: number of false negative pixels for the tubes and of false positifs for the planes and half-ellipsoids.

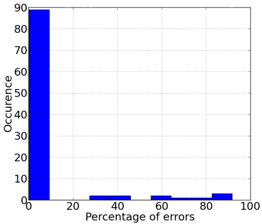

Tubes $(m = 4\%)$

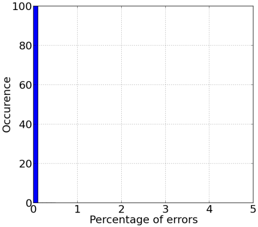

Planes $(m = 0\%)$

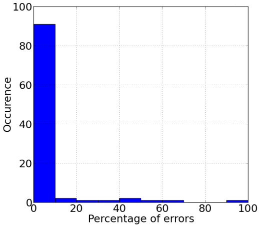

Half ellipsoids $(m = 4\%)$

Multi-scale approach

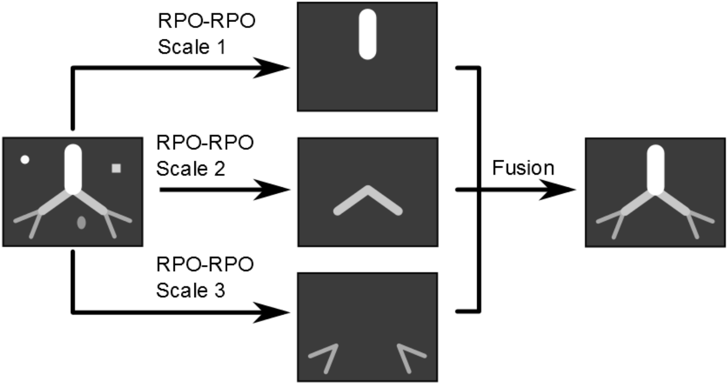

Multiscale Principle

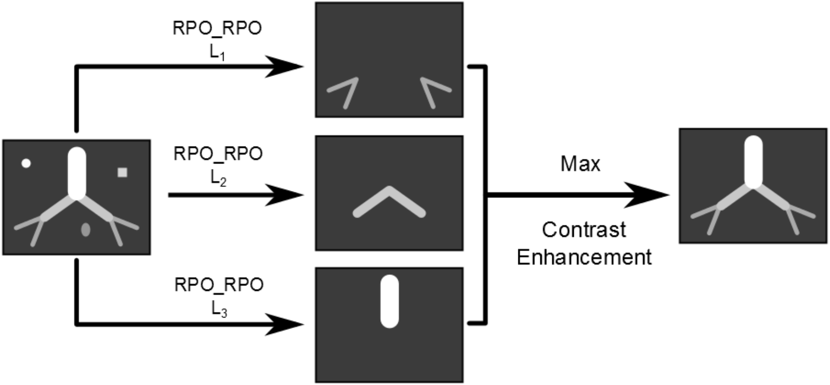

Comparisons

We performed qualitative comparisons of various methods according to four criteria on a full cerebral MRA

- Capacity to reduce background noise
- Capacity to detect large blood vessels
- Capacity to detect small blood vessels
- Presence of artifacts

RORPO with classical adjacencies and a multiscale approach based on path lengths seems to provide the best compromise.

Computing directions from RORPO

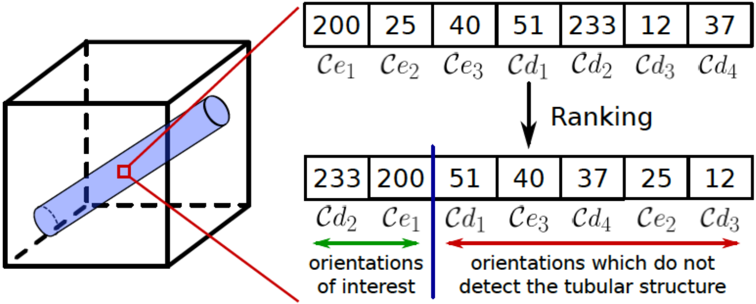

Figure: Computing directions from RORPO can be done by averaging the directions of high response.

Orientation results (in 2D)

(a)

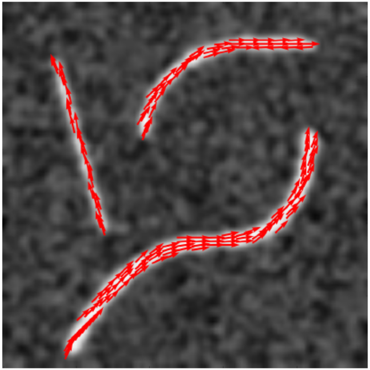

(b)

Figure: Orientation feature in 2D

Brain MRA result

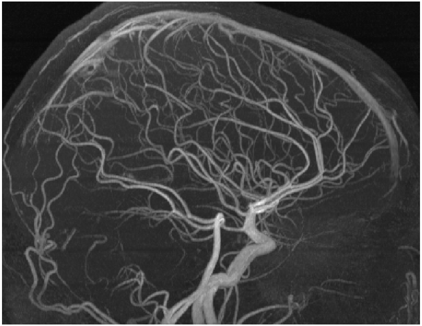

Initial image MIP

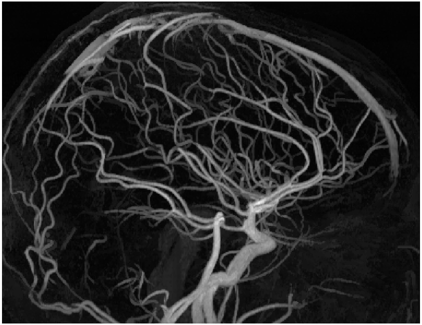

Length-based multiscale RORPO

MRA Result

Initial image

Multiscale RORPO

Comparison with Frangi vesselness

Proposed method

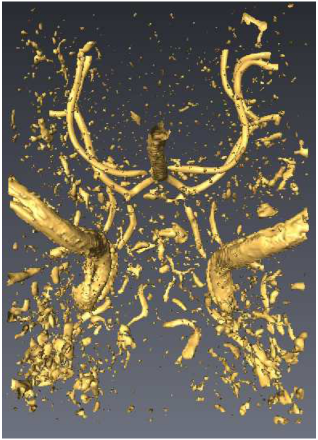

Optimized Frangi vesselness

Quantitative comparison: input data

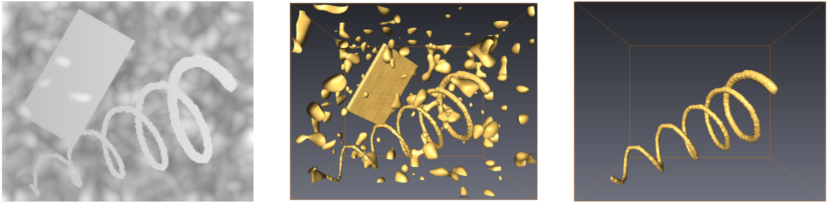

(a) CCM=0.605, Dice=0.634 (b) (c)

Figure: Synthetic image: (a) maximum intensity projection and (b) isosurface. (c) Ground truth.

Quantitative comparison: filtering response

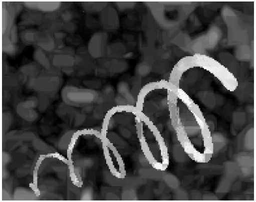

(a) CCM=0.88,Dice=0.89

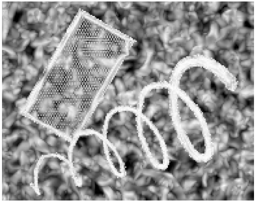

(b) CCM=0.71,Dice=0.73

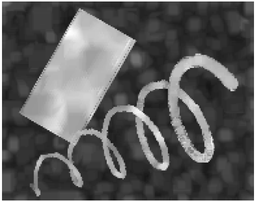

(c) CCM=0.66,Dice=0.65

Figure: Filtered synthetic image: maximum intensity projection. (a) RORPO. (b) Frangi's vesselness. (c) and RPO-top-hat.

Quantitative comparison - ROC analysis

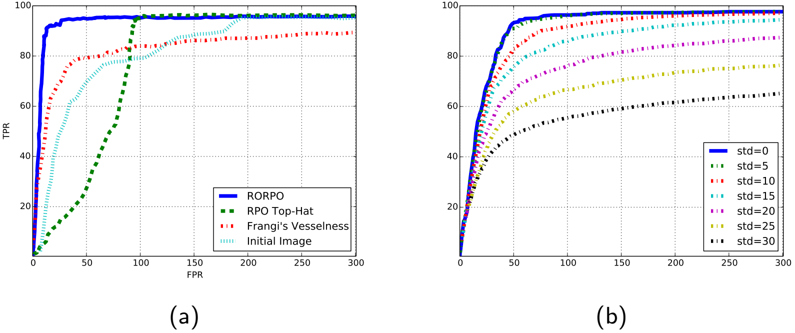

Figure: ROC curves on synthetic data. (a) Comparison of the three filters, plus the native image. (b) Noise robustness of the RORPO filter.

Quantitative Comparison

MICCAI Rotterdam Coronaries Database

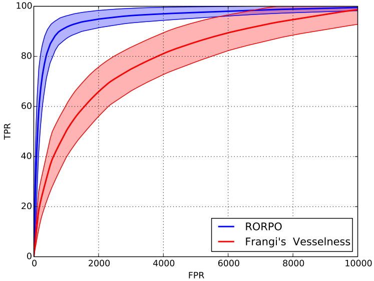

Figure: ROC curves of RORPO and Frangi's Vesselness on the Rotterdam repository. For both filtering the central curve is the mean ROC curve and the two others are the mean plus or minus one standard deviation ROC curve.

Quantitative comparison, synthetic data

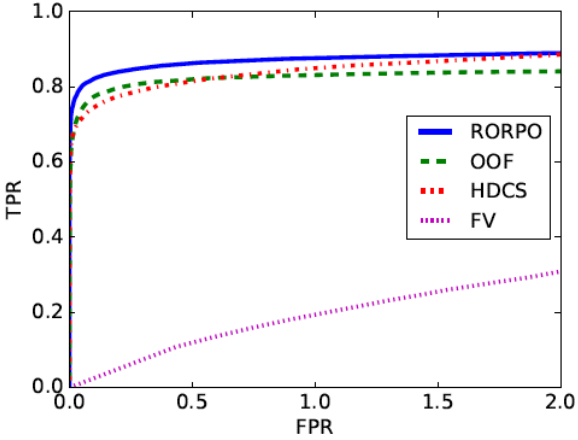

Three-way ROC analysis RORPO vs. OOF [11], HDCS [13] and Frangi Vesselness (FV)

Quantitative comp., Heart Coronaries

(a)

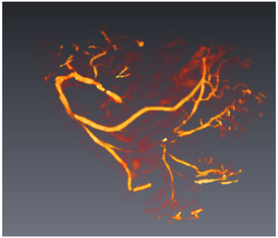

(b) MCC = 0.541

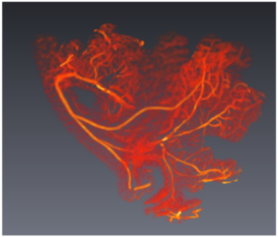

(c) MCC = 0.529

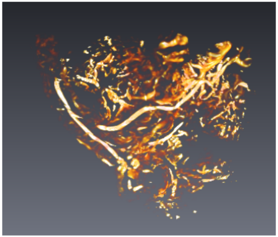

(d) MCC = 0.405

Ground truth (a) : RORPO (b), OOF [11] (c), Frangi (d)

Orientation feature 3D

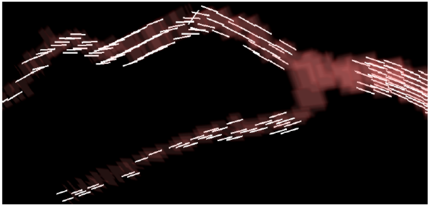

(a)

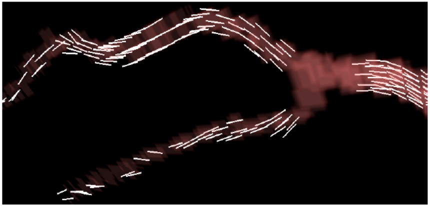

(b)

Orientation feature in 3D, Heart data: RORPO (a) vs Frangi Vesselness (

Tubular structures segmentation

- So far we have proposed a solution for curvilinear structure filtering.
- Segmentation of more complex structures that include tubes/cylinders can be built from this.
- We propose to use a variational framework by improving Total Variation (TV) to include a directional component.

Variational framework

We consider a convex variational framework

$$\min_{u \in X} F(u, f) + \lambda G(u). \tag{6}$$

- Here F is a data fidelity term and G a regularization.
- f is the input data and u the desired result.
- Typically F is associated to a noise model and G to an image model.
- A common image model is the Total Variation (TV)

Standard TV

This is isotropic standard TV regularization term (in 2D):

$$\mathrm{TV}(u) = \|\nabla u\|_{2,1} = \sum_{0 \le i,j < N} \sqrt{((\nabla u)^x_{i,j})^2 + ((\nabla u)^y_{i,j})^2} \tag{7}$$

where $\nabla u = \big((\nabla u)^x, (\nabla u)^y\big)$ is the 2D gradient.
It is classical in mathematics, and was proposed for image regularization in [18] (ROF model).

Directional TV

We define a directional gradient, $\nabla_D \in X^p$:

- We first define a *span* $(\mathrm{v}_1, \dots, \mathrm{v}_p)$ of p unitary vectors.
- This span contains all the discrete undirected orientations in a $k \times k$ neighborhood
- then:

$$\nabla_D u = \left(D^1 \circ (\nabla_d u)^1, \cdots, D^p \circ (\nabla_d u)^1 \right) \tag{8}$$

$$(\nabla_D u)_{i,j} = D^1_{i,j} (\nabla_d u)^1_{i,j} \mathrm{v}_1 + \cdots + D^p_{i,j} (\nabla_d u)^p_{i,j} \mathrm{v}_p \tag{9}$$

- with $D^q \in X,\ 1 \leq q \leq p$ a weight image such that:

$$D^q_{i,j} = d^q_{i,j} \Phi_{i,j} + (1 - \Phi_{i,j}) \tag{10}$$

Link with RORPO

- $\Phi \in X$ is a vesselness-like intensity feature normalized to the interval $[0, 1]$
- $(d^i)_{i \in [\![1,p]\!]}$, are computed from an orientation field

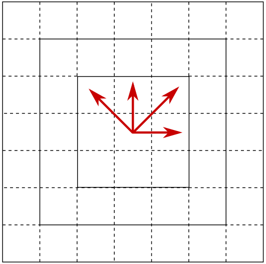

Span of vectors in a 3×3 neighborhood.

We used the RORPO response as the vesselness feature and the orientation field computed from RORPO.

Directional TV idea

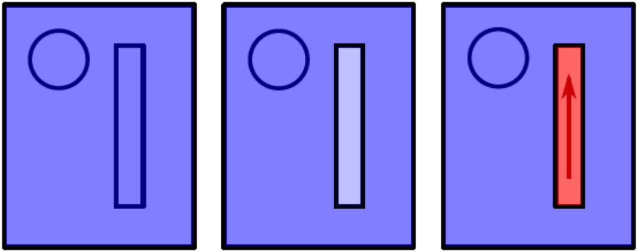

(a) TV (b) Weighted TV [60] (c) Directional TV

Directional TV idea

- Standard TV will penalize all edges identically; weighted TV may attempt to penalize edges less in a curvilinear object.
- A directional feature will not penalize edges inside a curvilinear object

Directional TV edges

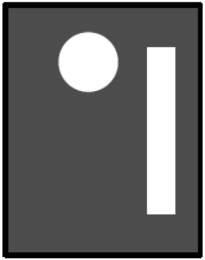
(a) u

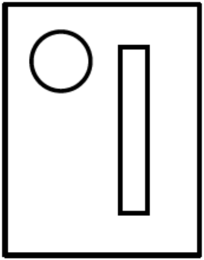
(b) ∇u

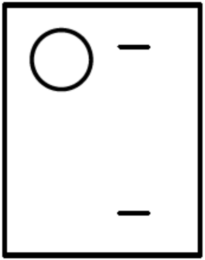
(c) $\nabla_d u$

Theoretical edge weights

- Thin, curvilinear object are all edges, and so are easily filtered out in standard/weighted TV.
- A directional TV will filter only along the direction of the curvilinear object, preserving it.

Segmentation model

Our model is base on TV_D, as follows:

$$\underset{u\in[0,1]^{N\times N}}{\text{minimize}} \quad \langle c_f, u\rangle_F + \lambda\|\nabla_D u\|_{2,1} \tag{11}$$

- $\|\nabla_D u\|_{2,1}$ is the directional Total Variation
- $\langle c_f, u\rangle_F$ is the Chan et al. data fidelity term [3] where $(c_f)_{i,j} = (c_1 - f_{i,j})^2 - (c_2 - f_{i,j})^2$ and $\langle u, v\rangle_F$ is the Frobenius product.
- The scalars c_1 and c_2 are respectively the foreground and background constant and f is the initial image.

DRIVE segmentation result

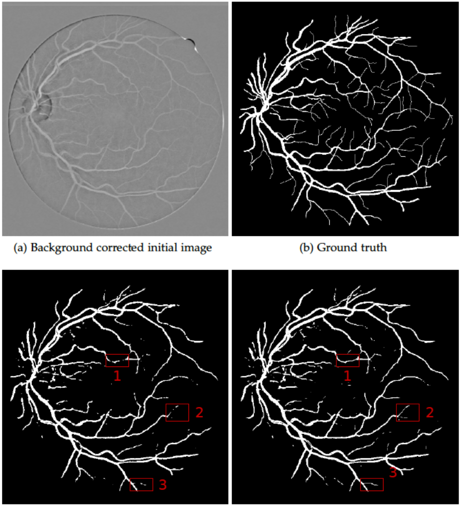

(a) Background corrected initial image

(b) Ground truth

(c) Chan model

(d) Proposed model

2D Results on DRIVE

DRIVE result details

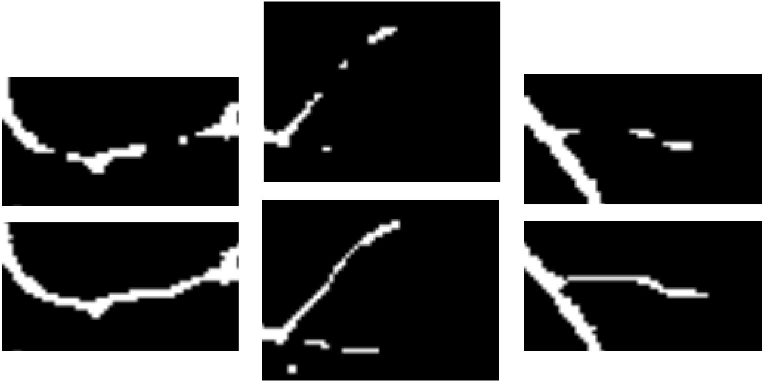

(e) box 1 (f) box 2 (g) box 3

2D Result on DRIVE (details) top: standard TV ; bottom: directional TV

DRIVE numerical results

Our segmentation result compare favourably with the state of the art. Note that some learning-based approaches can still outperform these results.

	TP	TN	Acc
Standard TV	0.656	0.985	0.9421
Directional TV	0.690	0.981	0.9434
Staal [20]	-	-	0.9442
Human observer	-	-	0.9470

Figure: Quantitative segmentation results on the DRIVE database.

Conclusion

We have studied a thin object filtering methods called RORPO
Associated with a multiscale approaches based on path length
Our method is effective at significantly reducing background noise while simultaneously suppressing non-tubular structures.

Perspectives

Quantitative evaluation of our results :

- Use phantoms produced by VascuSynth
- Use ground truth from heart MRA data.

Perspectives

- Produce images of scales
- Adapt the path operator slidework to the max-tree/min-tree slidework
    - This would allow discriminating objects on more complex measures than mere length
    - Think about incorporating robustness to max-trees / min-trees

Literature on path operators

- Definitions and early algorithms [2, 9, 10]
- Faster algorithms [1, 21]
- Extension to 3D and regularisation [12]
- RPO and 3D [5], [4]
- Applications [27, 28, 29, 19]
- RORPO [15, 16, 14]

Outline

Image representations

Decomposition into primitive or fundamental elements that can be more easily interpreted:

- Functional decompositions;
- Multiresolution decompositions;
- Multi-scale representations;
- Threshold decompositions;
- Hierarchical representations.

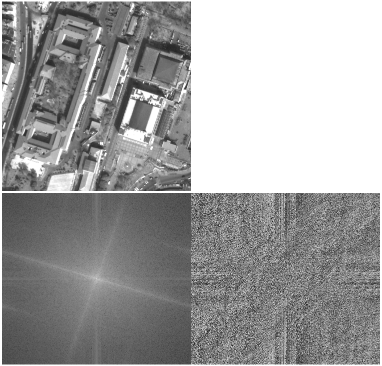

Amplitude Phase

Image representations

Decomposition into primitive or fundamental elements that can be more easily interpreted:

- Functional decompositions;
- Multiresolution decompositions;
- Multi-scale representations;
- Threshold decompositions;
- Hierarchical representations.

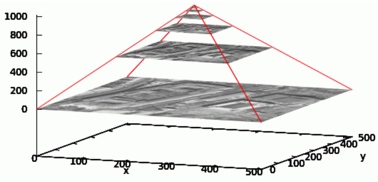

Image representations

Decomposition into primitive or fundamental elements that can be more easily interpreted:

- Functional decompositions;
- Multiresolution decompositions;
- Multi-scale representations;
- Threshold decompositions;
- Hierarchical representations.

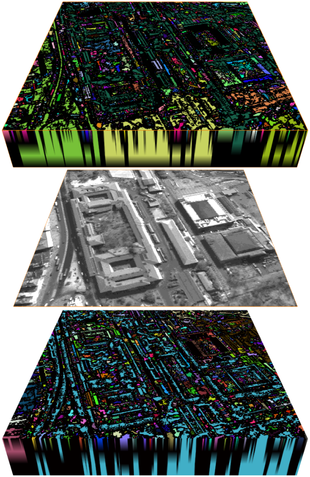

Image representations

Decomposition into primitive or fundamental elements that can be more easily interpreted:

- Functional decompositions;
- Multiresolution decompositions;
- Multi-scale representations;
- Threshold decompositions;
- Hierarchical representations.

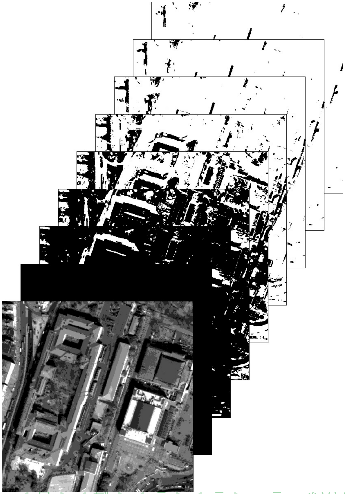

Image representations

Decomposition into primitive or fundamental elements that can be more easily interpreted:

- Functional decompositions;
- Multiresolution decompositions;
- Multi-scale representations;
- Threshold decompositions;
- Hierarchical representations.

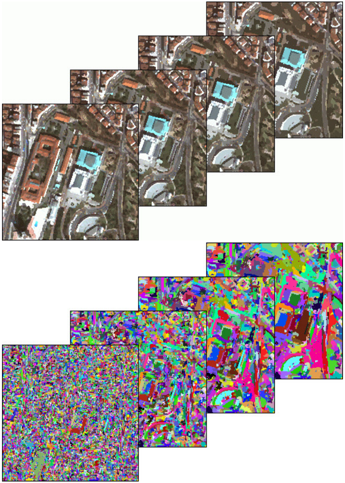

Image representations

Decomposition into primitive or fundamental elements that can be more easily interpreted:

- Functional decompositions;
- Multiresolution decompositions;
- Multi-scale representations;
- Threshold decompositions;
- Hierarchical representations.

Not mutually exclusive.
Properties inherited from those of underlying operations.
Choice driven by the application needs.

Connected operators

What's connected operators ?

Filtering tools that merge flat zones.

Properties

- No new contours,
- Keep contours' position.

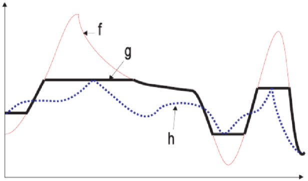

Leveling with marker.
f : input,
h : marker,
g : result.

One popular implementation [Salembier & Wilkinson, SPM, 2009]

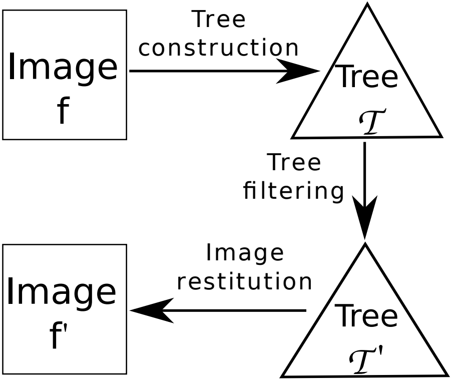

Level sets and components

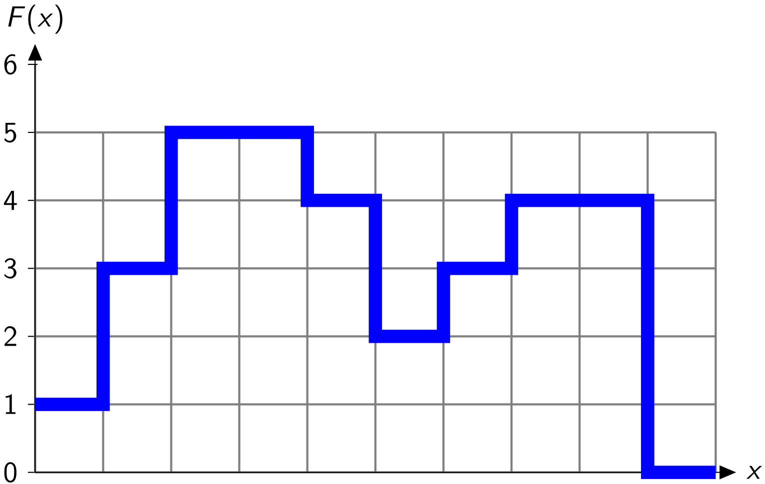

$$y = F(x).$$

Level sets and components

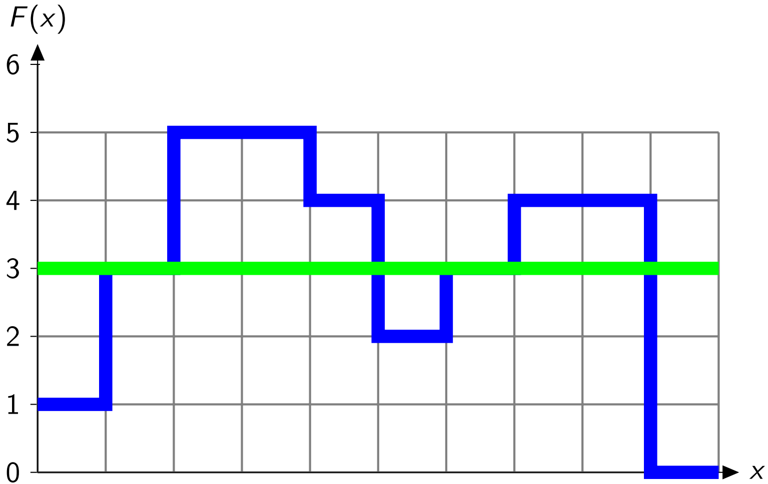

$F(x) \geq 3.$

Level sets and components

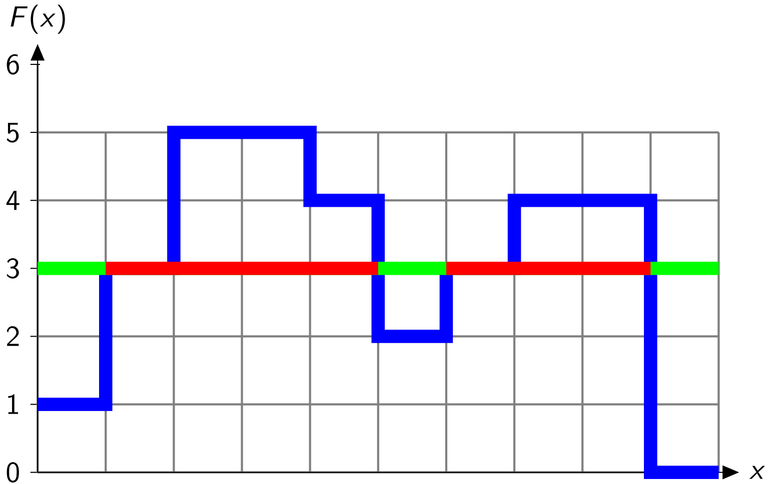

$$F_3 = \{x \mid F(x) \geq 3\}.$$

Level sets and components

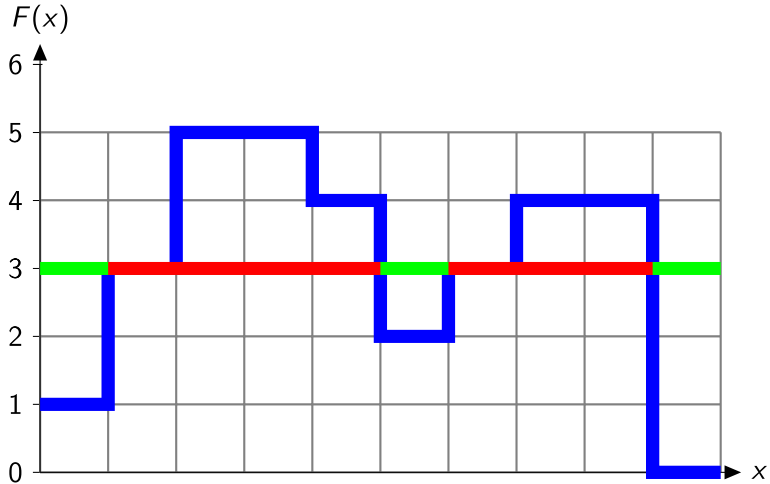

$$F_k = \{x \mid F(x) \geq k\}.$$

(Max) component tree

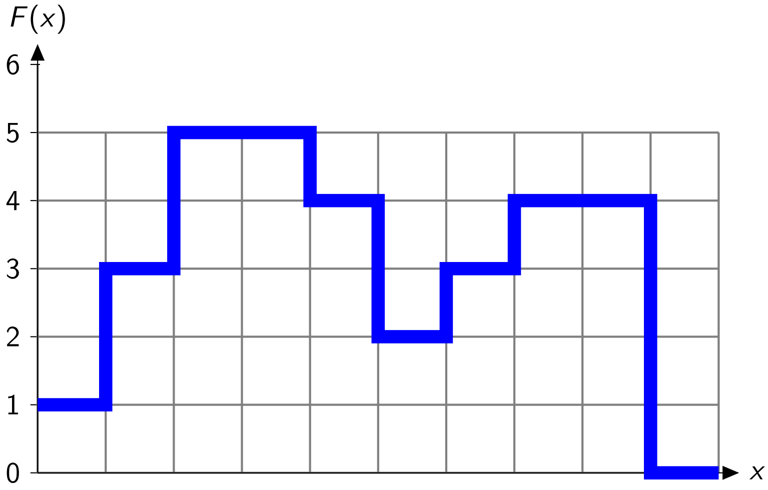

(Max) component tree

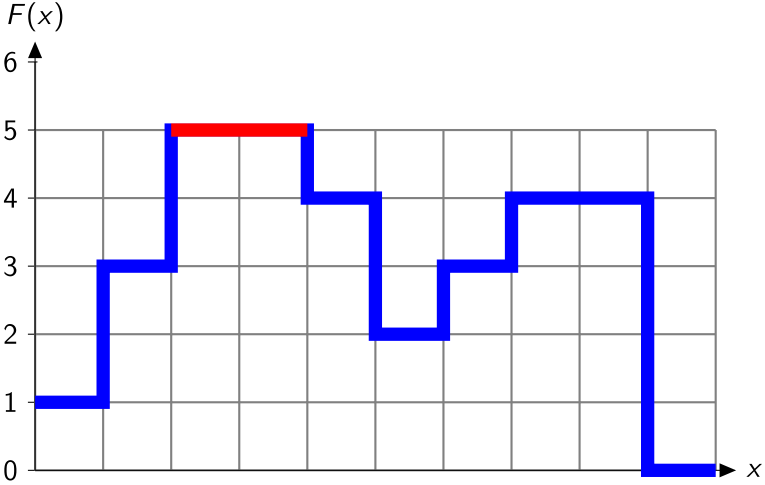

$F_5 = \{x \mid F(x) \geq 5\}.$

(Max) component tree

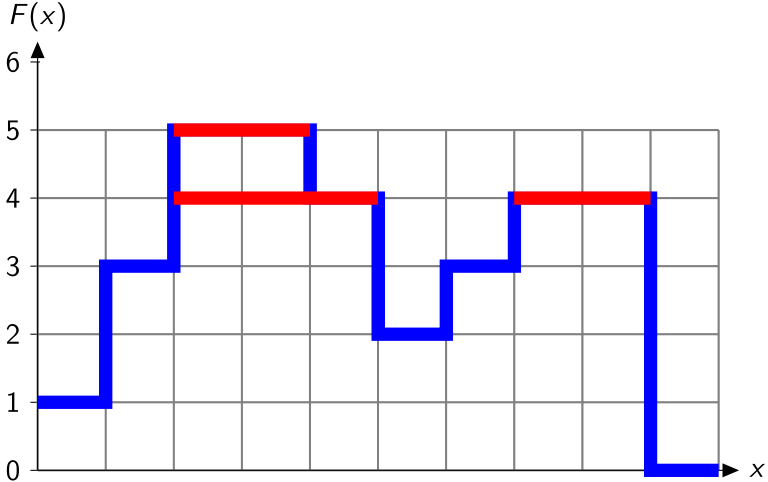

$F_4 = \{x \mid F(x) \geq 4\}.$

(Max) component tree

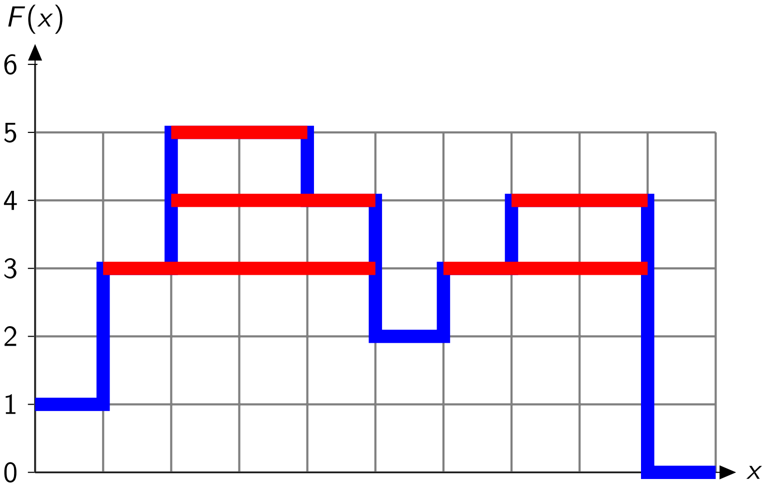

$F_3 = \{x \mid F(x) \geq 3\}$.

(Max) component tree

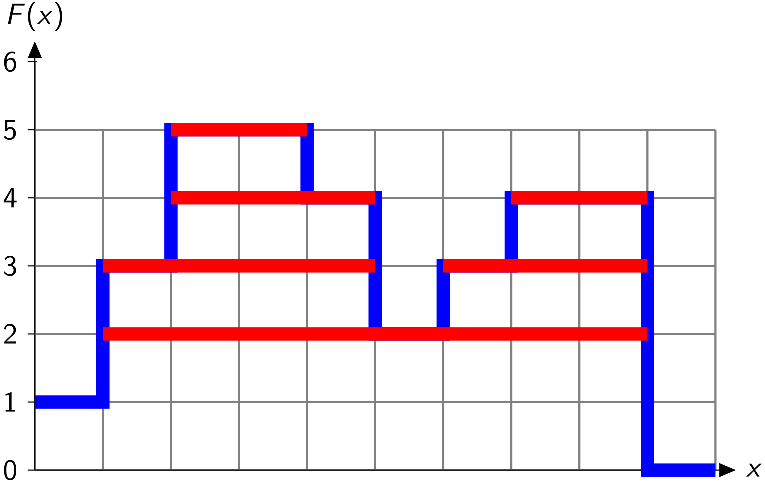

$F_2 = \{x \mid F(x) \geq 2\}.$

(Max) component tree

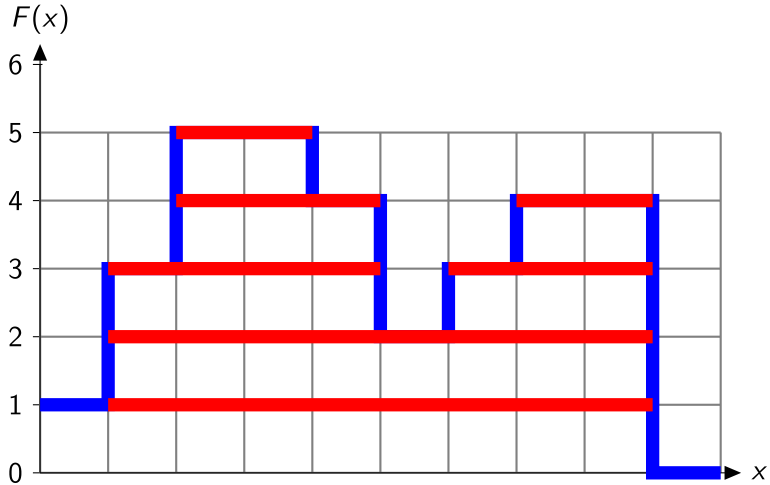

$F_1 = \{x \mid F(x) \geq 1\}.$

(Max) component tree

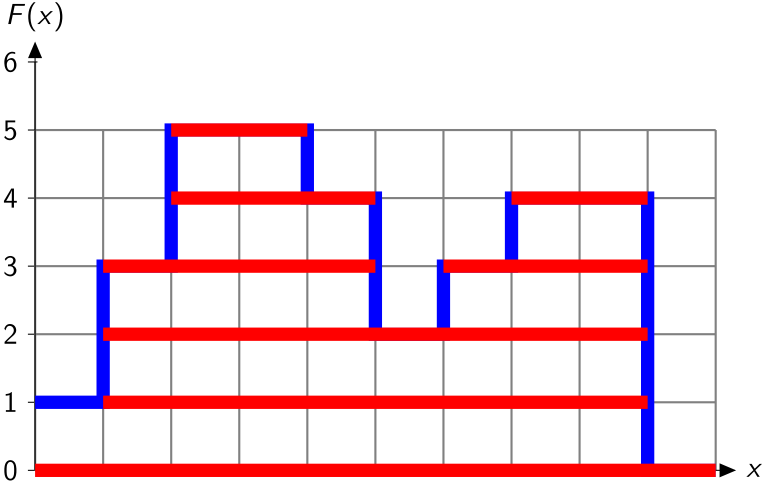

$$F_0 = \{x \mid F(x) \geq 0\}.$$

(Max) component tree

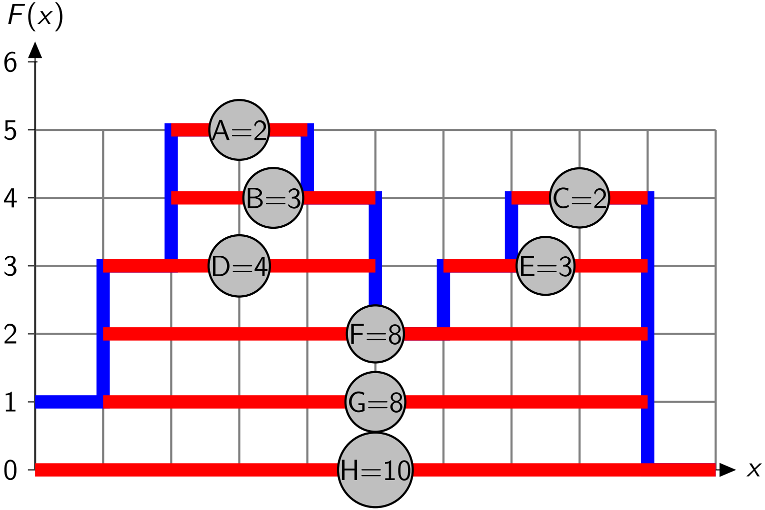

(Max) component tree

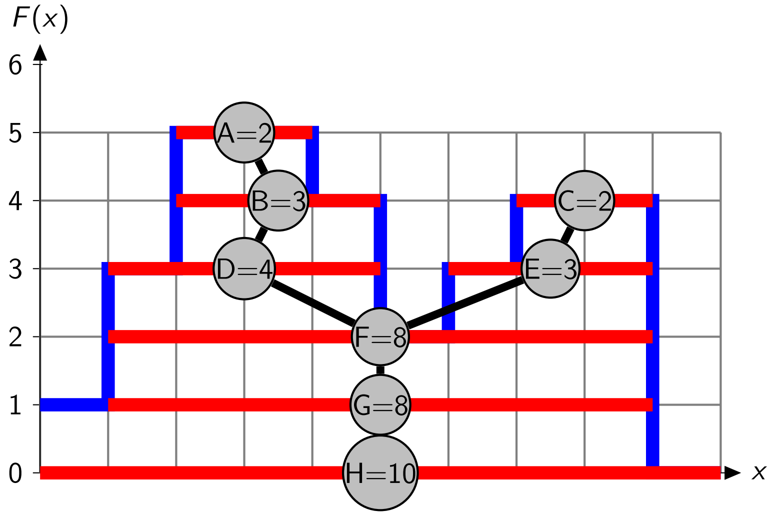

Components + inclusion relationship = component tree.

Some of the many possible trees

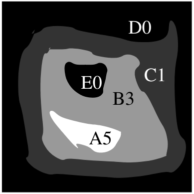

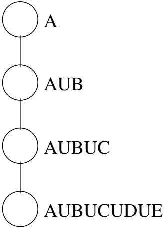

Max–tree

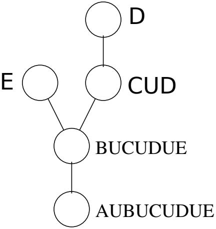

Min–tree

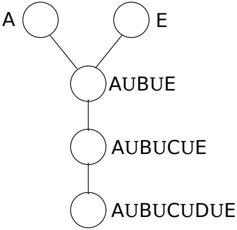

Tree of shapes

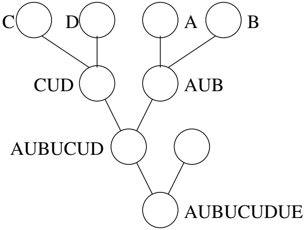

Binary Partition Tree

Some link with Morse's Theory

Important idea

Some nodes are more important than others

- *Leaves* $\Leftrightarrow$ *"Extrema"*
- *Nodes with more than one child* $\Leftrightarrow$ *"Saddle points"*
- *Hence, filtering is linked with* topological persistence

A Topological Approach to Local Feature detection

Y. Xu *et al.*, ITIP 2014

Interest point detection

- Important for 3D reconstruction, image registration, . . .
- Many methods exist: DoG, Corners, MSER, . . .
- None of them is invariant to contrast change

A Topological Approach to Local Feature detection

Y. Xu *et al.*, ITIP 2014

Tree-Based Morse Regions (TBMR)

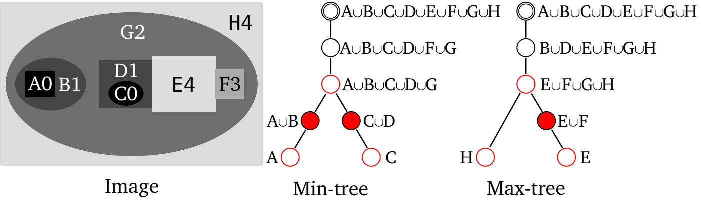

- Select critical nodes (leaves and nodes with several children)
- The scale of a critical node is the largest region containing it and topologically equivalent in its tree.

A Topological Approach to Local Feature detection

Y. Xu *et al.*, ITIP 2014

Attributes

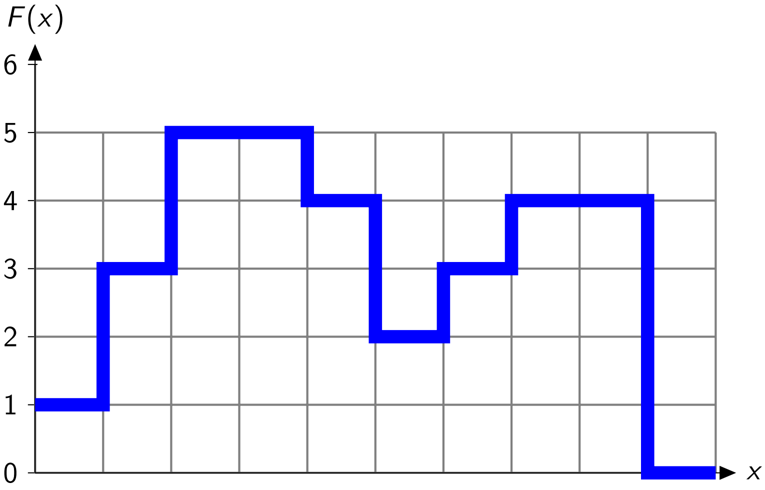

Attributes

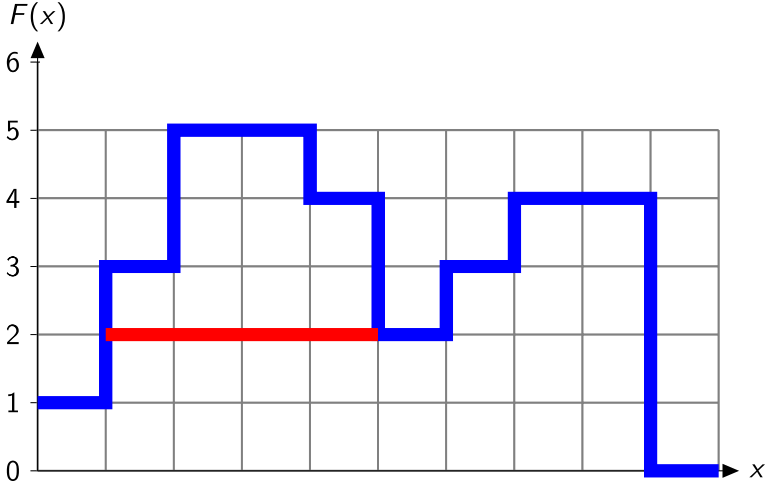

A connected component.

Attributes

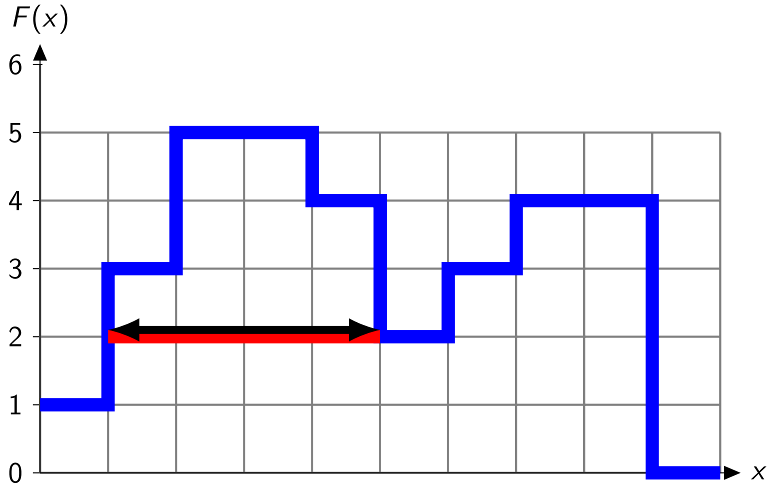

Area.

Attributes

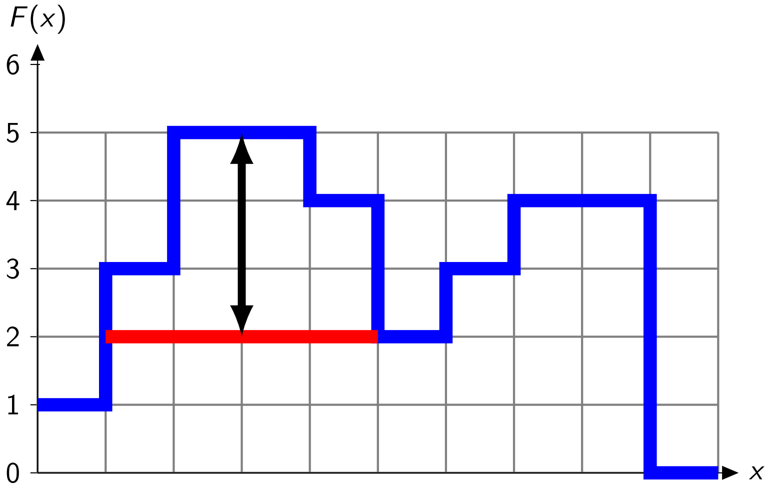

Height.

Attributes

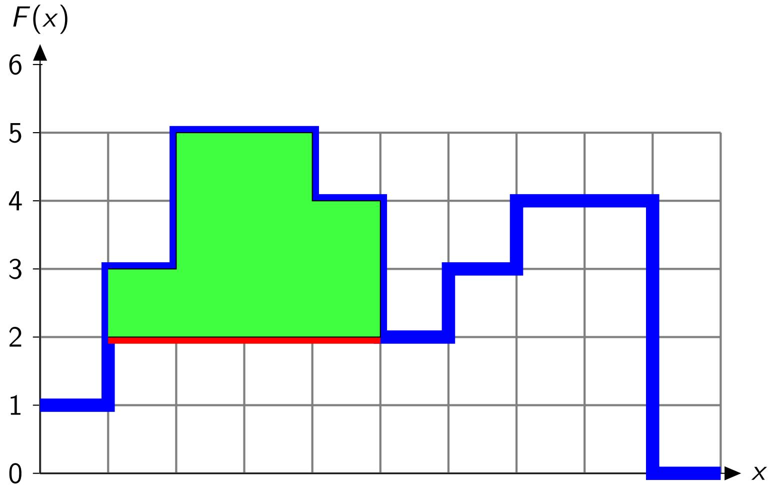

Volume.

Attributes

Increasing attributes

Increasing attributes : $A \subseteq B \Rightarrow \mathcal{A}(A) \leq \mathcal{A}(B)$.
Examples : Area, height, volume.

Non-increasing attributes

Shape attributes.

- I/A^2 minimum for a round object,
- Circularity : $area/(\pi \times l_{max}^2)$,
- Elongation : L_{max}/L_{min}.

L_{min} and L_{max} : Length of the two main axes of the best fitting ellipse.

Filtering with increasing attributes

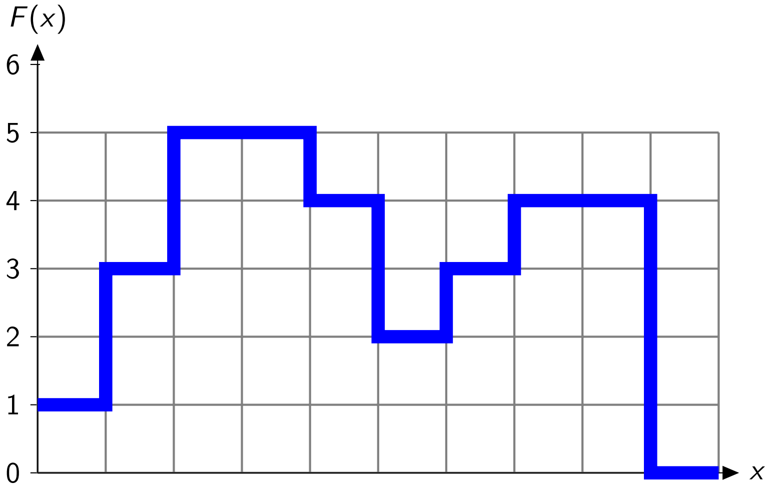

Filtering with increasing attributes

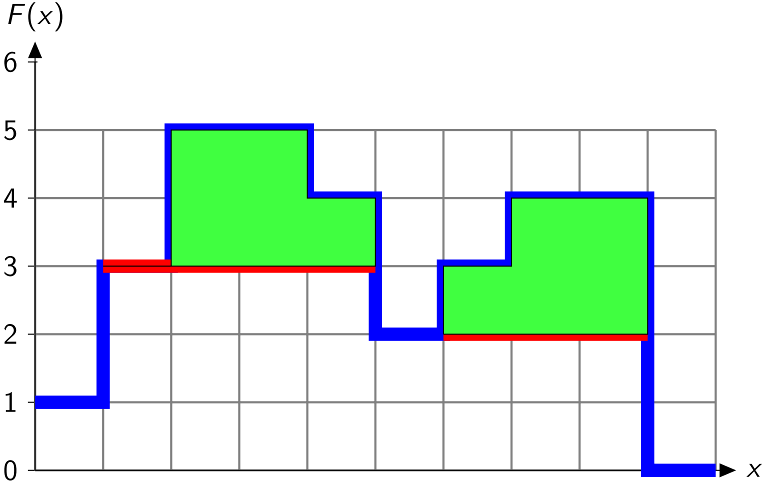

Volume ≤ 5.

Filtering with increasing attributes

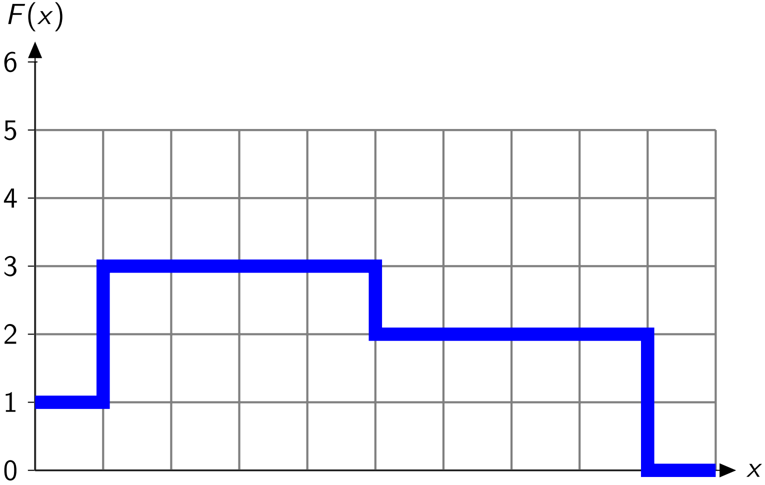

Filtered function.

Applications: filtering with increasing attribute

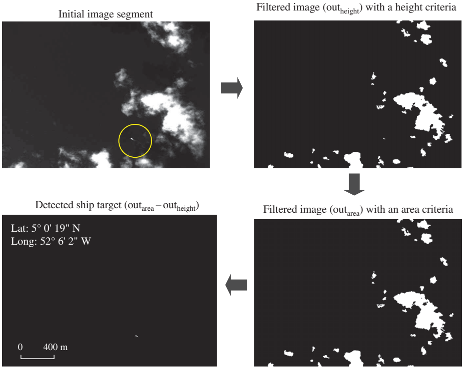

Ship detection on optical satellite image

C. Corbane *et al.*, International Journal of Remote Sensing 31 (22), 5837-5854

Applications: filtering with increasing attribute

FIRST

future. According to Dealogic, 87 SPACs have begun trading in the U.S. since the end of 2003, buying some well-known companies like smoothie purveyor Jamba Juice. Last year alone, 40 SPACs worth $3.4 billion were announced, up from $484 million two years earlier.

SPACs are essentially shell companies. They go public with little more to show investors than a management team and an agreement that the money raised will be used to fund an acquisition in a particular sector, such as retail, or in an emerging market like China. Ergo, a SPAC is a roll of the dice. "Imagine paying $50 to go to a Broadway show, and you have no idea what's behind the curtain," says lawyer Mitchell Littman of Littman Krooks LLP, who works with a lot of SPACs. "You are relying on the fortune and integrity of the management team." Once a deal is completed, the SPAC's managers (who typically receive 20% of the public shares as compensation) are free to sell their holdings, usually after a lockup period.

Today's SPAC boom harkens back to the flurry of "blank check" companies that sprang up in the 1980s. Those were widely discredited by a wave of scams in which fraudsters would take a shell company public, announce a merger, pump the stock, and then dump it before everyone realized that the hot target company was anything but. Investors have a lot more protection with SPACs, which must hold almost all the money in escrow until a deal is done.

From Charney's perspective, there is no downside to the arrangement. It lets him raise cash relatively quickly, remain the largest shareholder, and avoid having to sell other investors on his vision until after his company is already publicly traded. "I'm in a hurry to get the money," he says, citing his plan to expand beyond his current 145 stores. "What's great about the SPAC is, it's kind of in reverse. First you get publicly traded, then you do the compliance stuff. And if you brought in private equity, they might want to run the company in a more [cautious] manner," he says, referring to American Apparel's risqué image.

Keeping control is particularly important for a guy who is the creative force behind the brand, as well as one whose flamboyant management style—there is one sexual-harassment suit outstanding against him by a former employee; two others have been dismissed—might scare off many investors. "He is very talented," says former J.C. Penney CEO Allen Questrom, a partner in private-equity firm Lee Equity Partners, "but he is a *meshugana* [read: nutty] type guy. It's very difficult to see him in a public sphere." (Charney, for his part, dismisses the talk about his reputation as "all tabloid.")

FIRST YOU GET PUBLICLY TRADED, THEN YOU DO COMPLIANCE.

Charney's new bedfellows don't appear concerned, even if they're more likely to wear dark suits than American Apparel's colorful leggings. Endeavor's president, Jonathan Ledecky, is familiar with unusual financing vehicles; he founded U.S. Office Products, the once-hot roll-up that went bankrupt in 2001 (long after Ledecky cashed out). Boldfaced names on the board include Kerry Kennedy, Bobby's daughter, and Edward Mathias, a managing director at the Carlyle Group. On Jan. 31 the same group formed a new $250 million SPAC, Victory Acquisition, underwritten by Citi. The filing notes that it could compete directly with Endeavor. And as with Endeavor, once Victory goes public it will pay a Ledecky "affiliate" $7,500 monthly to run its Manhattan office.

Hedge funds are drawn to these investments. Since SPACs typically must announce a deal within 18 months of their IPOs or return the money raised, with interest, hedgies see them as safe places to park cash while potentially profiting from the price fluctuations between a deal's announcement and its close—a sort of arbitrage play. That may be why Steven Cohen's $12 billion hedge fund SAC Capital snapped up 8.7% of Endeavor's stock when the deal was announced.

Even if investors are SPAC-happy, not all regulators are: So far only the American Stock Exchange lists SPACs, and many bankers and lawyers say the SEC is deliberately taking its time vetting deals, hoping to slow down the pace of offerings. SPACs are probably best left to those, like Charney, who like to live life on the edge.

SPAC Attack!

The volume of special purpose acquisition companies going public in the U.S. is rising sharply.

Number of deals
Value of deals (in billions)

13 $0.5 — 2004
29 $2.1 — '05
40 $3.4 — '06

UNHEALTHY RETURN Since smoothie chain Jamba Juice went public in November via a SPAC, its shares have fallen 19%.

FORTUNE CHART / SOURCE: DEALOGIC; PHOTO: DAVID MCNEW—GETTY

12 • FORTUNE February 19, 2007

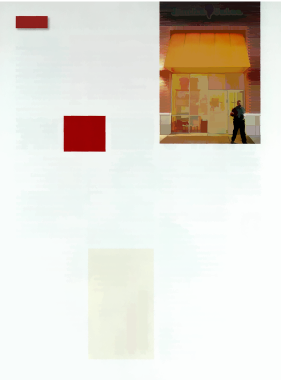

Box detection on a document image

Attribute: (width, height) of the component

G. Lazzara *et al.*, DAS 2014

Some functions and spaces on which to compute the trees

- An image
- A gradient
  - Especially usefull with the watershed
- A node-weighted graph
- An edge-weighted graph
- A weighted mesh
- A density function
- and more...

Filtering with increasing attributes

Pruning the trees

$\mathcal{A} \uparrow$, Pruning the leaves = Attribute thresholding.

Non-increasing attributes

How to process the filtering?

A popular implementation of connected filters

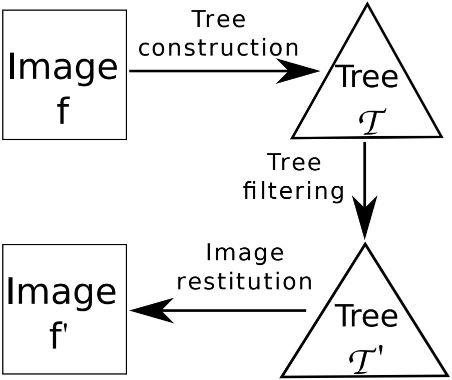

Shape-space filtering [Y. Xu *et al.*, ICPR, 2012]

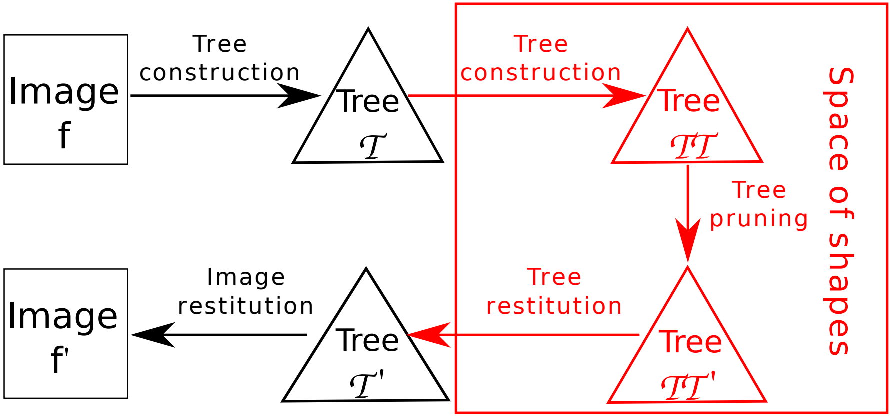

Outline

Construction of second tree representation

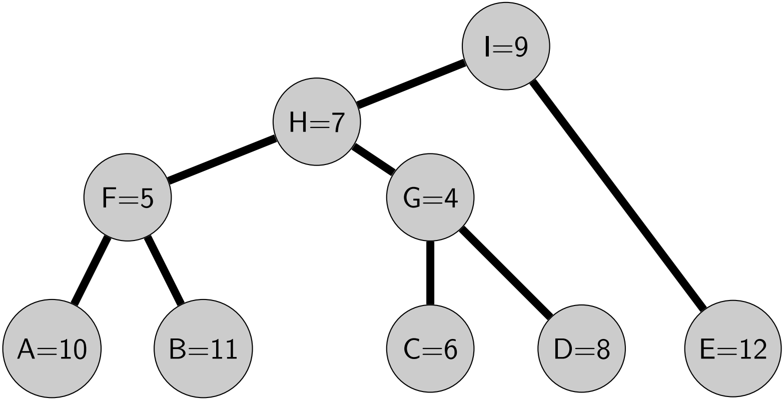

Construction of second tree representation

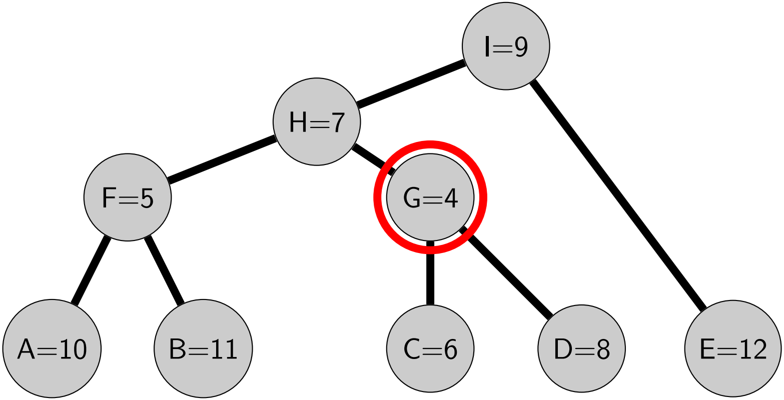

Level $\{x|A(x) \leq 4\}$.

Construction of second tree representation

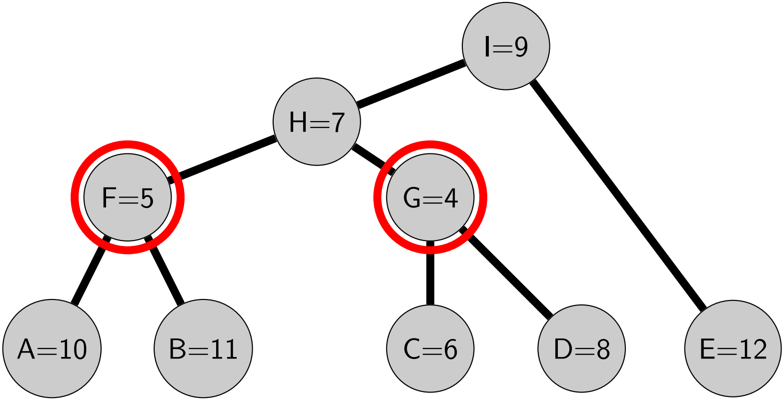

Level $\{x|A(x) \leq 5\}$.

Construction of second tree representation

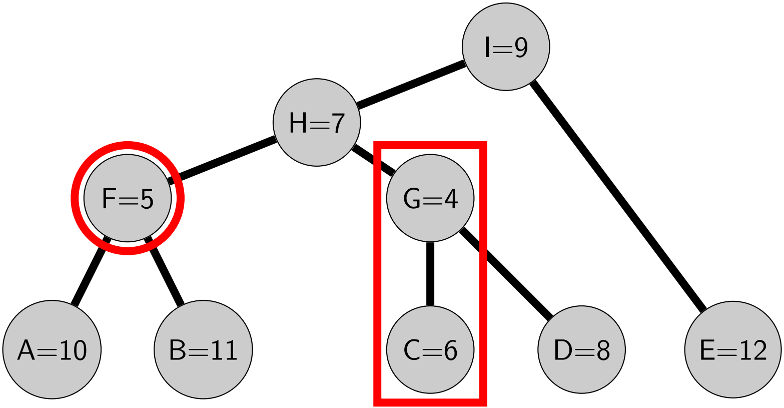

Level $\{x|A(x) \leq 6\}$.

Construction of second tree representation

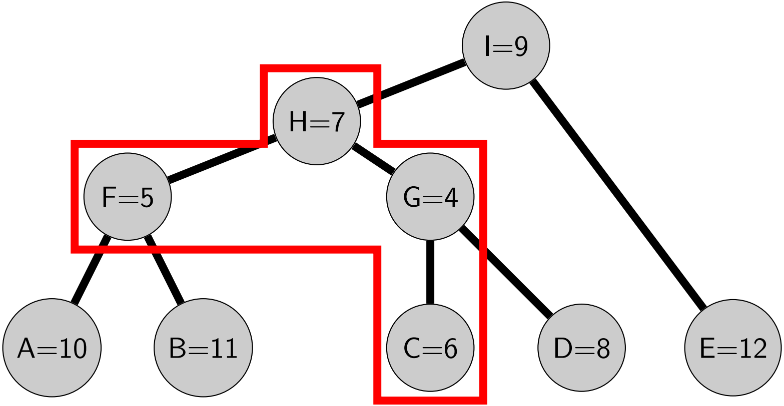

Level $\{x | A(x) \leq 7\}$.

Construction of second tree representation

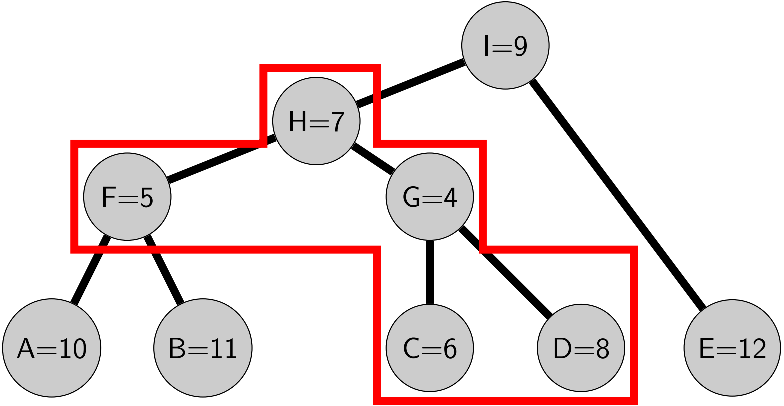

Level $\{x|A(x) \leq 8\}$.

Min-tree of a tree-based image representation

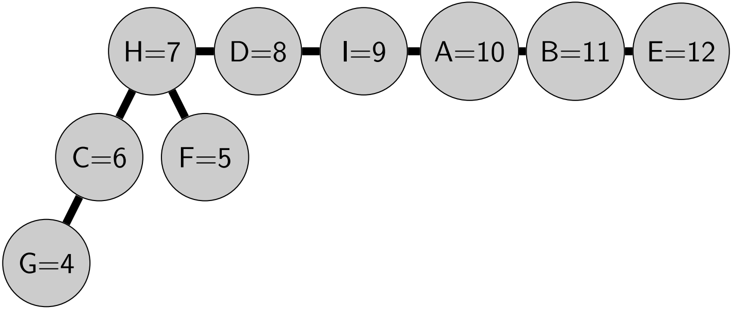

Important idea

Computing a Min-Tree on a node-weighted graph instead of a matrix image.
Easy thanks to Olena [Levillain & Géraud & Najman, ICIP, 2010], the generic image processing platform `http://olena.lrde.epita.fr`.

Morphological shapings

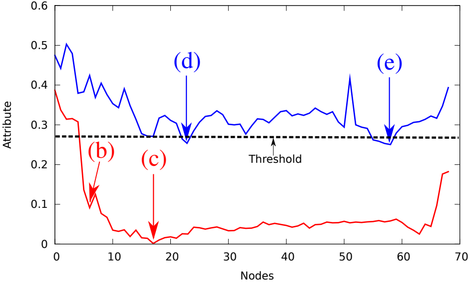

Evolution of circularity on two branches.

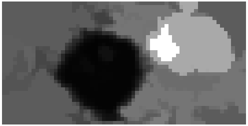

Thresholding.

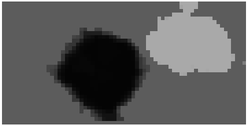

Our shaping.

(b)

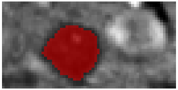

(c)

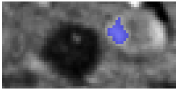

(d)

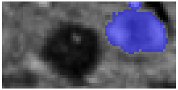

(e)

Morphological shapings

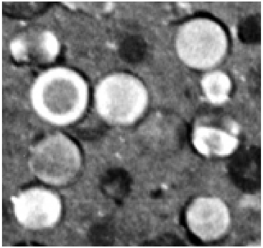

Input image.

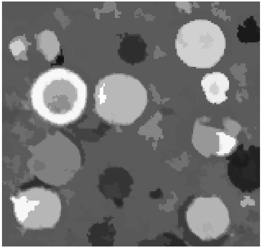

Shaping based on $\mathcal{A}$

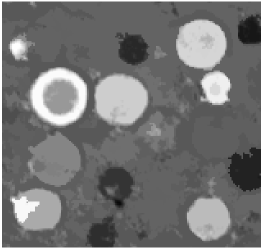

Low threshold of $\mathcal{A}$.

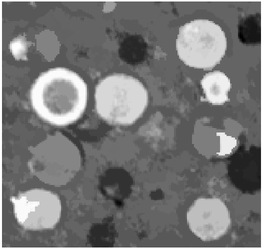

Higher threshold of $\mathcal{A}$.

Morphological shapings

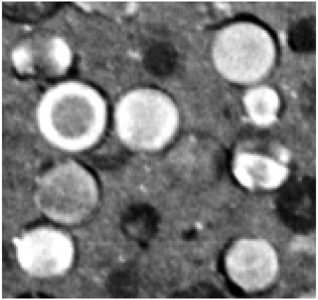

Input image.

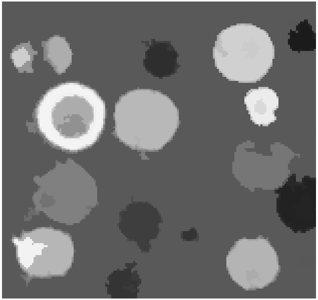

Our shaping 2.

Using a combination of attributes $\mathcal{A}$.

Application to object segmentation

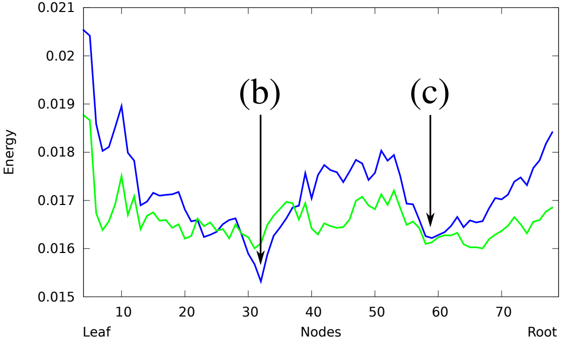

Energy in a branch of the tree;
blue : ICIP 2012 energy; green : snake energy.

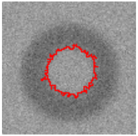

(b)

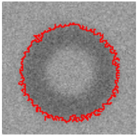

(c)

Object detection principle

Significant minima $\Leftrightarrow$ Objects.

Object detection results

Context-based energy estimator

Input image.

Objects detected.

Object detection results

Shape attribute

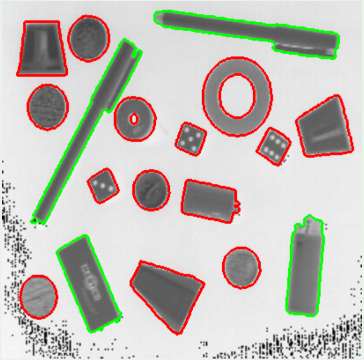

Objects detected using shape attribute.
Red ones : circularity-based; Green ones : Inverse elongation-based.

Shape-based lower/upper levelings

Input image.

Round objects based upper-leveling.

Blood vessels segmentation in retinal images


(a) Input image.


(b) Input image.


(c) Input image.


(d) Input image.

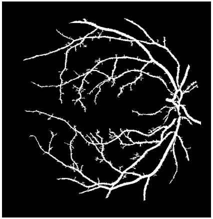
(e) Segmentation.

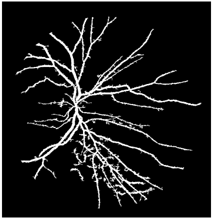
(f) Segmentation.

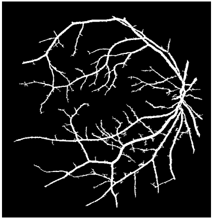
(g) Segmentation.

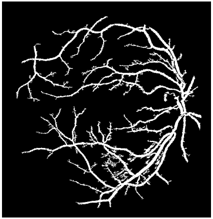
(h) Segmentation.

Cleaning a manuscript with a shape-based lower-leveling

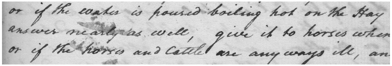

(a) Input image.

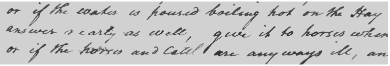

(b) Result of $\psi_{s\downarrow}$.

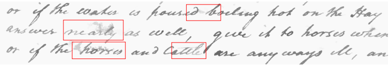

(c) Low threshold of $\mathcal{A}$ (subtractive).

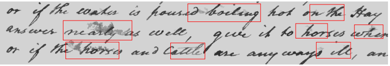

(d) Low threshold of $\mathcal{A}$ (average).

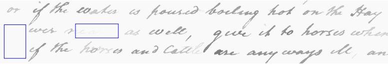

(e) Higher threshold of $\mathcal{A}$ (subtractive).

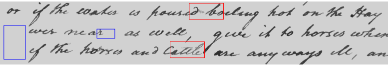

(f) Higher threshold of $\mathcal{A}$ (average).

Saliency maps Najman *et al.*, PAMI 1996 – Najman, JMIV 2011

(a) Original image.

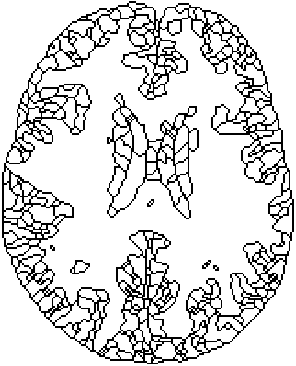

(b) Some contours.

Saliency maps Najman *et al.*, PAMI 1996 – Najman, JMIV 2011

(a) Original image.

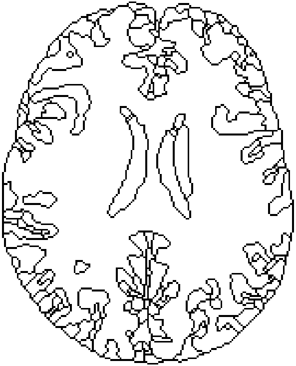

(b) Some contours.

Saliency maps Najman *et al.*, PAMI 1996 – Najman, JMIV 2011

(a) Original image.

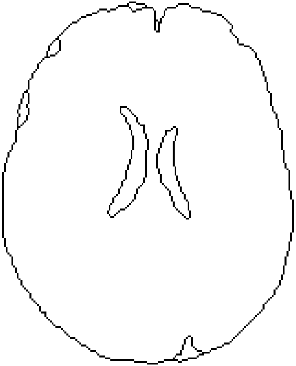

(b) Some contours.

Saliency maps Najman *et al.*, PAMI 1996 – Najman, JMIV 2011

Stacking the contours gives a saliency map [Najman & Schmitt, PAMI, 1996]

(a) Original image.

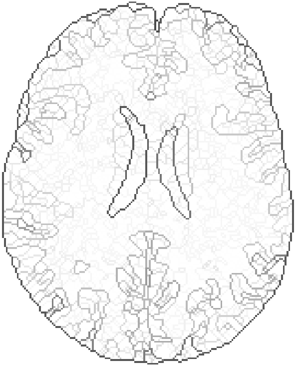

(b) A saliency map.

Different representations

[L. Najman - JMIV - 2011] Mathematical definitions, equivalence between ultrametric watersheds, saliency maps and trees of segmentations

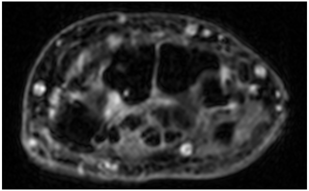

(a) Original image.

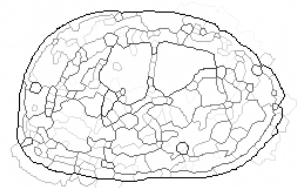

(b) Ultrametric watershed.

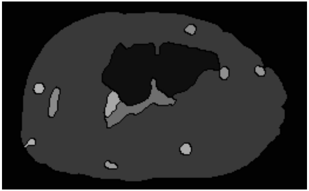

(c) One of the segmentations.

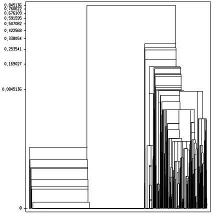

(d) Dendrogram.

Saliency maps can be computed on a mesh

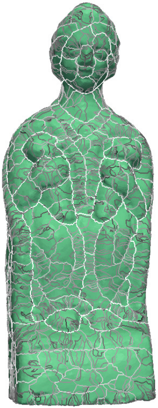

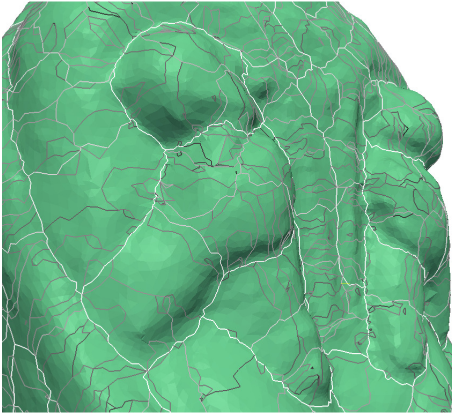

Object spotting and saliency maps

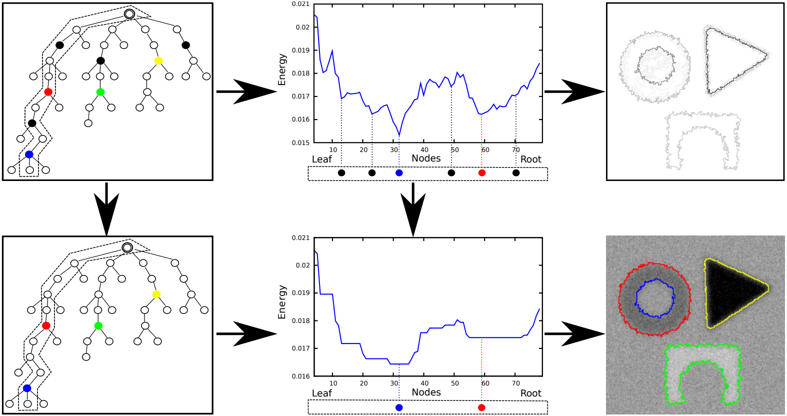

Object spotting and saliency maps

Input image.

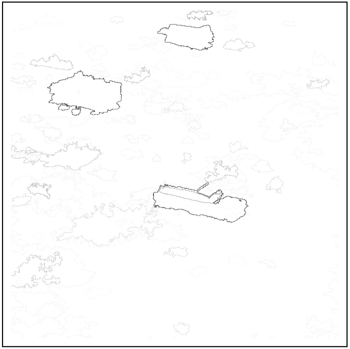

Saliency map.

Hierarchical simplification based on Mumford-Shah

Original.

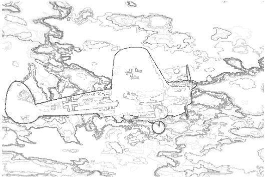

Saliency map.

Simplified.

Hierarchical simplification based on Mumford-Shah

Original.

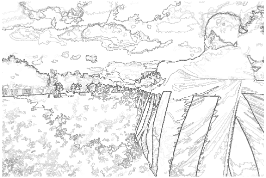

Saliency map.

Simplified.

Hierarchical simplification based on Mumford-Shah

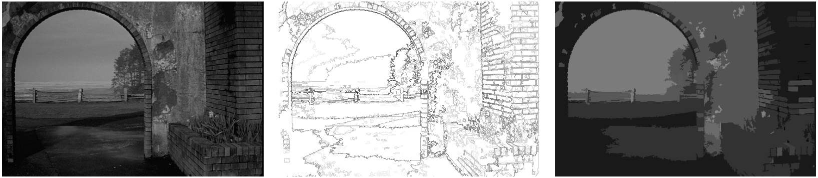

Original. Saliency map. Simplified.

Shape oriented saliency maps on the GTSDB dataset

Original.

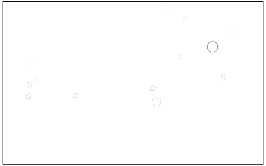

Circular.

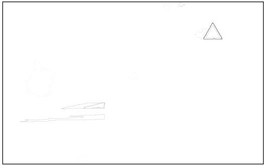

Triangular.

Method	Detection rate		Area under curve	
	Prohibitive	Danger	Prohibitive	Danger
Our	96%	**95%**	**92.16%**	**93.10%**
Viola-Jones	**98.8%**	74.6 %	90.81%	46.26%
HOG+LDA	91.3%	90.7%	70.33%	35.94%
Hough-like	55.3%	65.1%	26.09%	30.41%

Precision of other methods is 10%. Our precision is 59% (resp. 41%)

Shape oriented saliency maps on the GTSDB dataset

Original.

Circular.

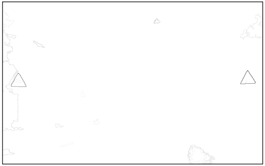

Triangular.

Method	Detection rate		Area under curve	
	Prohibitive	Danger	Prohibitive	Danger
Our	96%	**95%**	**92.16%**	**93.10%**
Viola-Jones	**98.8**%	74.6 %	90.81%	46.26%
HOG+LDA	91.3%	90.7%	70.33%	35.94%
Hough-like	55.3%	65.1%	26.09%	30.41%

Precision of other methods is 10%. Our precision is 59% (resp. 41%)

Shape oriented saliency maps on the GTSDB dataset

Original.

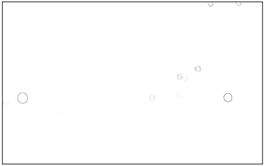

Circular.

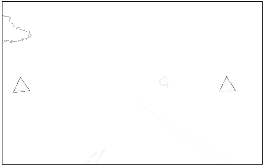

Triangular.

Method	Detection rate		Area under curve	
	Prohibitive	Danger	Prohibitive	Danger
Our	96%	**95%**	**92.16%**	**93.10%**
Viola-Jones	**98.8%**	74.6 %	90.81%	46.26%
HOG+LDA	91.3%	90.7%	70.33%	35.94%
Hough-like	55.3%	65.1%	26.09%	30.41%

Precision of other methods is 10%. Our precision is 59% (resp. 41%)

Smartphone document capture competition (ICDAR 2015)

Input Frame.

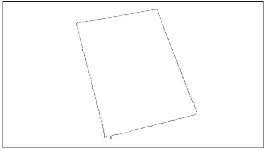

Saliency map.

Extracted document.

Smartphone document capture competition (ICDAR 2015)

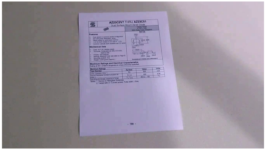

Input Frame.

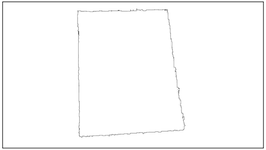

Saliency map.

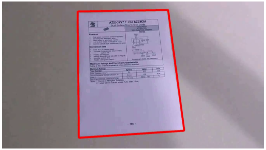

Extracted document.

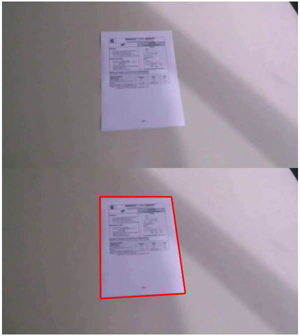

Smartphone document capture competition (ICDAR 2015)

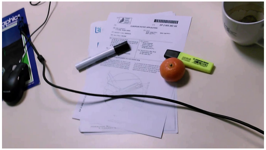

Input Frame.

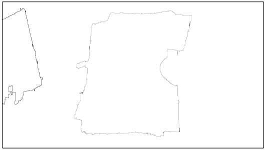

Saliency map.

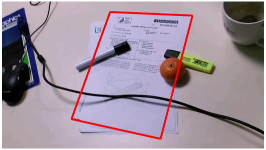

Extracted document.

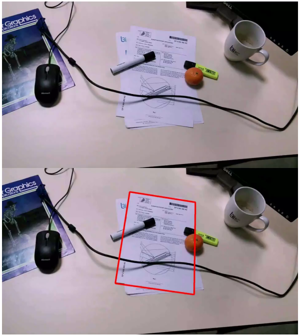

Smartphone document capture competition (ICDAR 2015)

Ranking	Method	Jaccard Index	Confidence Interval
1	Our2	**0.9816**	**[0.9813, 0.9819]**
1	Our	**0.9716**	**[0.9710, 0.9721]**
2	ISPL-CVML	0.9658	[0.9649, 0.9667]
3	SmartEngines	0.9548	[0.9533, 0.9562]
4	NetEase	0.8820	[0.8790, 0.8850]
5	A2iA run 2	0.8090	[0.8049, 0.8132]
6	A2iA run 1	0.7788	[0.7745, 0.7831]
7	RPPDI-UPE	0.7408	[0.7359, 0.7456]
7	SEECS-NUST	0.7393	[0.7353, 0.7432]

Outline

The Power watershed framework C. Couprie *et al.*, PAMI 2011

$$x^*_{p,q} = \arg\min_x \underbrace{\sum_{e_{ij}\in E} {w_{ij}}^p |x_i - x_j|^q}_{\text{Smoothness term}} + \underbrace{\sum_{v_i \in V} {w_i}^p |x_i - l_i|^q}_{\text{Data term}}$$

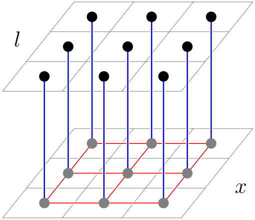

The Power watershed framework C. Couprie *et al.*, PAMI 2011

$$x^*_{p,q} = \arg\min_x \underbrace{\sum_{e_{ij} \in E} {w_{ij}}^p |x_i - x_j|^q}_{\text{Smoothness term}} + \underbrace{\sum_{v_i \in V} {w_i}^p |x_i - l_i|^q}_{\text{Data term}}$$

$$w_{ij} = \exp^{-\alpha d(l_i, l_j)}$$

The Power watershed framework C. Couprie *et al.*, PAMI 2011

$$x^*_{p,q} = \arg\min_x \underbrace{\sum_{e_{ij}\in E} {w_{ij}}^p |x_i - x_j|^q}_{\text{Smoothness term}} + \underbrace{\sum_{v_i\in V} {w_i}^p |x_i - l_i|^q}_{\text{Data term}}$$

$$\boxed{\bar{x} = \lim_{p\to\infty} x^*_{p,q}}$$

The Power watershed framework C. Couprie *et al.*, PAMI 2011

$$x^*_{p,q} = \arg\min_x \underbrace{\sum_{e_{ij} \in E} {w_{ij}}^p |x_i - x_j|^q}_{\text{Smoothness term}} + \underbrace{\sum_{v_i \in V} {w_i}^p |x_i - l_i|^q}_{\text{Data term}}$$

$$x^*_{p,q} = \arg\min_x \left[w_{\max}^p \left(\sum_{e_{ij} \in E_{\max}} |x_i - x_j|^q + \sum_{e_{ij} \notin E_{\max}} \frac{w_{ij}^p}{w_{\max}^p} |x_i - x_j|^q \right) \right]$$

The Power watershed framework C. Couprie *et al.*, PAMI 2011

$$x^*_{p,q} = \arg\min_x \underbrace{\sum_{e_{ij}\in E} {w_{ij}}^p |x_i - x_j|^q}_{\text{Smoothness term}} + \underbrace{\sum_{v_i \in V} {w_i}^p |x_i - l_i|^q}_{\text{Data term}}$$

$$x^*_{p,q} = \arg\min_x \left[w_{\max}^p \left(\sum_{e_{ij}\in E_{\max}} |x_i - x_j|^q + \varepsilon \right) \right]$$

Algorithm for the power watershed

1. Compute a max-tree of the edge-weighted graph
   - Maximum spanning tree
2. Run through all the connected components of the max-tree by decreasing altitude
   - Optimize

$$\min_x \sum_{e_{ij} \in \text{connected component}} |x_i - x_j|^q$$

   on the connected component, using the previous computations as initial conditions

Why PW it a watershed?

Watershed cut Cousty *et al.*, PAMI 2009

- In edge-weighted graphs, it is a cut satisfying the drop of water principle

Why PW it a watershed?

Watershed cut Cousty *et al.*, PAMI 2009

- In edge-weighted graphs, it is a cut satisfying the drop of water principle
- **Consistency property:** two equivalent characterisations by
  - Catchment basins (steepest descent property)
  - Dividing lines (drop of water principle)

Why PW it a watershed?

Watershed cut Cousty *et al.*, PAMI 2009

- In edge-weighted graphs, it is a cut satisfying the drop of water principle
- **Consistency property:** two equivalent characterisations by
  - Catchment basins (steepest descent property)
  - Dividing lines (drop of water principle)
- **Optimality property:**
  - Characterization by minimum/maximum spaning forests/trees

Surface reconstruction using power watershed Couprie *et al.*, ISMM 2011

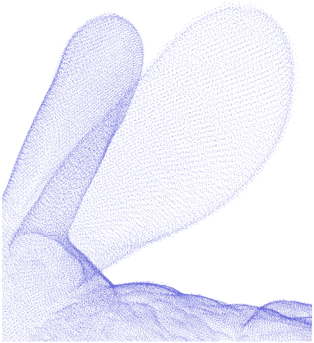

$\Rightarrow$

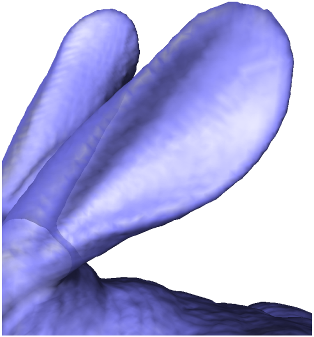

- Goal : given a noisy set of dots, find an explicit surface fitting the dots.

Surface reconstruction using power watershed

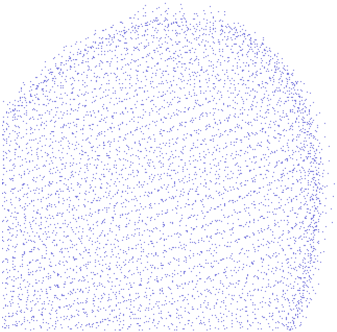

$\Rightarrow$

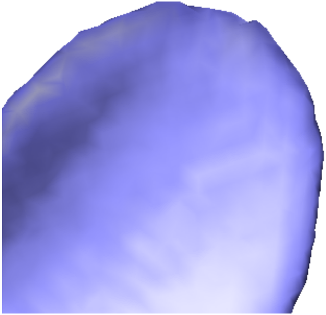

- Goal : given a noisy set of dots, find an explicit surface fitting the dots.

How to solve this problem

- Graph : 3D grid
- Here x represents the object indicator to recover.

$$\lim_{p\to\infty} \arg\min_x \sum_{e_{ij}\in E} {w_{ij}}^p |x_i - x_j|^{q=2}$$

$$\text{s.t. } x(F) = 1,\ x(B) = 0$$

- weights : distance function from the set of dots to fit

Why PW are a good fit for this problem ?

numerous plateaus around the dots to fit $\rightarrow$ smooth isosurface is obtained

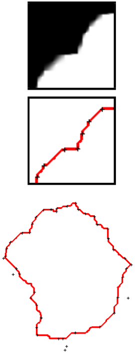

Power watershed solution

Comparisons

Graph cuts
Lemptizky-Boykov,
CVPR 2007

Power watershed
Couprie-Bresson-Najman-Talbot-Grady, ISMM 2011

Comparisons

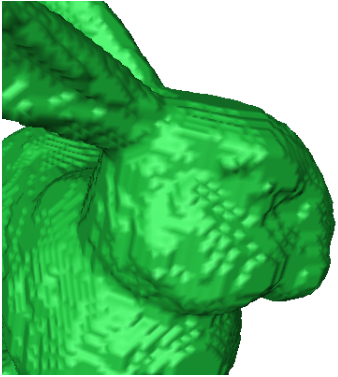

Graph cuts
Lemptizky-Boykov,
CVPR 2007

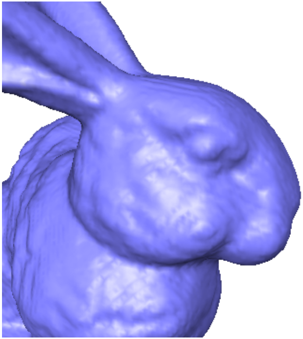

Power watershed
Couprie-Bresson-Najman-Talbot-Grady, ISMM 2011

Comparisons

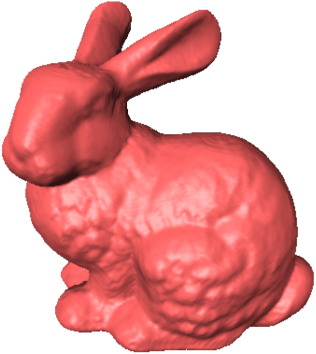

Total variation

Power watershed
Couprie-Bresson-Najman-Talbot-Grady, ISMM 2011

Comparisons

Total variation

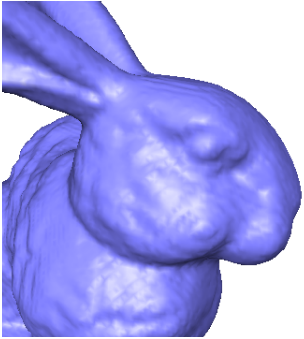

Power watershed
Couprie-Bresson-Najman-Talbot-Grady, ISMM 2011

Non-convex diffusion using power watersheds

- Anisotropic diffusion [Perona-Malik 1990]

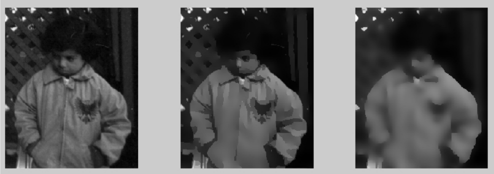

Image　　100 iterations　　200 iterations

Goals of this work:

- perform anisotropic diffusion using an ℓ_0 norm to avoid the blurring effect
- optimize a non convex energy using Power Watershed [Couprie-Grady-Najman-Talbot, ICIP 2010]

Anisotropic diffusion and ℓ_0 norm

$$x^* = \arg\min_x \underbrace{\sum_{e_{ij}\in E} \sigma(x_i - x_j)}_{\text{smoothness term}} + \underbrace{\lambda \sum_{v_i \in V} \sigma(x_i - f_i)}_{\text{data fidelity term}}$$

$$\sigma(x) = 1 - \mathrm{e}^{-\alpha x^2}$$

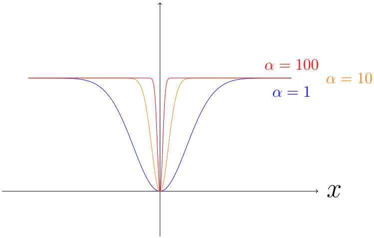

Leads to piecewise constant results

Original image

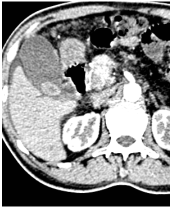

PW result

Power watersheds: more applications coming soon!

- Spectral clustering
- Image filtering

Outline

Conclusion: emerging topics

- Trees for color and multi-modalities images
- A little further with graphs: discrete calculus
  - Discrete version of morphological PDE
  - Hierarchical optimization
  - Links with other computer vision frameworks
- Beyond graphs: other interesting structures
  - Directed graphs
  - Hyper-graphs
  - Simplicial complexes and discrete topological analysis
  - New applications with the emergence of big data
- Morphological tools available online
  - PINK (C library and tools): http://pinkhq.com/
  - MILENA (generic C++ library): https://www.lrde.epita.fr/wiki/Olena/Milena

Learning Hierarchical Features for Scene Labeling

C. Farabet *et al.*, PAMI 2013

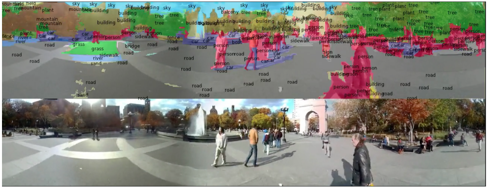

Learning Hierarchical Features for Scene Labeling

C. Farabet *et al.*, PAMI 2013

Thank for your attention !

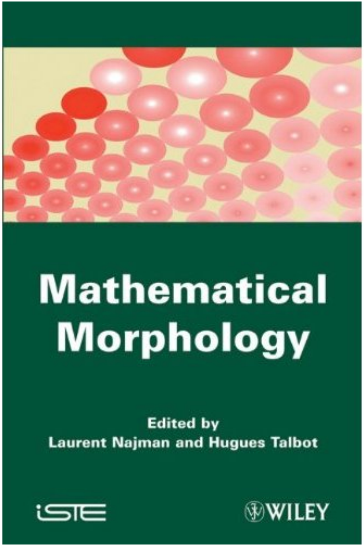

Pink: http://pinkhq.com
Olena: http://www.lrde.epita.fr/cgi-bin/twiki/view/Olena

Some references

Main survey reference (without equations!)

Laurent Najman, Jean Cousty. *A graph-based mathematical morphology reader.* Pattern Recognition Letters, Elsevier, 2014, 47, pp.3-17.

Some references

Optimization

- Camille Couprie, Leo Grady, Laurent Najman, Jean-Christophe Pesquet, Hugues Talbot. *Dual constrained TV-based regularization on graphs.* SIAM Journal on Imaging Sciences, 2013, 6 (3), pp.246-1273.
- Camille Couprie, Leo Grady, Hugues Talbot, Laurent Najman. *Combinatorial Continuous Maximal Flows.* SIAM Journal on Imaging Sciences, 2011, 4, pp.905-930
- Camille Couprie, Leo Grady, Laurent Najman, Hugues Talbot. *Power Watersheds: A Unifying Graph Based Optimization Framework.* IEEE Transactions on Pattern Analysis and Machine Intelligence, 2011, 33 (7), pp.1384-1399.

Some references

On trees

- Yongchao Xu, Pascal Monasse, Thierry Géraud, Laurent Najman. *Tree-Based Morse Regions: A Topological Approach to Local Feature Detection.* IEEE Transaction on Image Processing, 2014, 23 (12), pp 5612-5625.
- Yongchao Xu, Thierry Géraud, Laurent Najman. *Connected Filtering on Tree-Based Shape-Spaces.* IEEE Transaction on PAMI, 2016, 38 (6), pp.1126 - 1140.
- Yongchao Xu, Thierry Géraud, Laurent Najman. *Hierarchical Segmentation Using Tree-Based Shape Spaces.* IEEE Transaction on PAMI, 2016 (to appear).

Some references

On watersheds, hierarchies and topology

- Jean Cousty, Gilles Bertrand, Michel Couprie, Laurent Najman. *Collapses and watersheds in pseudomanifolds of arbitrary dimension.* JMIV, 2014, 50 (3), pp.261-285.
- Laurent Najman. *On the equivalence between hierarchical segmentations and ultrametric watersheds.* JMIV, 2011, 40 (3), pp.231-247.
- Jean Cousty, Gilles Bertrand, Laurent Najman, Michel Couprie. *Watershed Cuts: Minimum Spanning Forests and the Drop of Water Principle.* IEEE PAMI, 2009, 31 (8), pp.1362-1374
- Jean Cousty, Gilles Bertrand, Michel Couprie, Laurent Najman. *Fusion graphs: merging properties and watersheds.* JMIV, 2008, 30 (1), pp.87-104.

Some references

Some recent applications

- G. Lazzara, T. Géraud. and R. Levillain, *Planting, Growing, and Pruning Trees: Connected Filters Applied to Document Image Analysis.*, Document Analysis Systems, 2014, pp. 36 - 40
- Clément Farabet, Camille Couprie, Laurent Najman, Yann Lecun. *Learning Hierarchical Features for Scene Labeling.* IEEE PAMI, 2013, 35 (8), pp.1915 - 1929.
- Sylvie Philipp-Foliguet, Michel Jordan, Laurent Najman, Jean Cousty. *Artwork 3D model database indexing and classification.* Pattern Recognition, Elsevier, 2011, 44 (3), pp.588-597
- Christina Corbane, Laurent Najman, Emilien Pecoul, Laurent Demagistri, Michel Petit. *A complete processing chain for ship detection using optical satellite imagery.* International Journal of Remote Sensing, 2010, 31 (22), pp.5837–5854

Some references

Some algorithms

- Laurent Najman, Michel Couprie. *Building the component tree in quasi-linear time.* IEEE Transactions on Image Processing, 2006, 15 (11), pp.3531-3539.
- Thierry Géraud, Edwin Carlinet, Sébastien Crozet, Laurent Najman. *A quasi-linear algorithm to compute the tree of shapes of n-D images.* ISMM, 2013, LNCS 7883, pp.97-108
- Laurent Najman, Jean Cousty, Benjamin Perret. *Playing with Kruskal: algorithms for morphological trees in edge-weighted graphs.* ISMM, 2013, LNCS 7883, pp.135-146
- Edwin Carlinet and Thierry Géraud. *A comparative review of component tree computation algorithms.* IEEE Trans Image Process. 2014 Sep;23(9):3885-95. Jul 8.

Other references I

Ben Appleton and Hugues Talbot.
Efficient path openings and closings.
In C. Ronse, L. Najman, and E. Decencière, editors, *Mathematical Morphology: 40 Years On*, volume 30 of *Computational Imaging and Vision*, pages 33–42, Dordrecht, 2005. Springer-Verlag.

M. Buckley and H. Talbot.
Flexible linear openings and closings.
In *Mathematical Morphology and its application to image analysis*, pages 109–118, Palo Alto, June 2000. Kluwer.

Tony F. Chan, Selim Esedoglu, and Mila Nikolova.
Algorithms for finding global minimizers of image segmentation and denoising models.
SIAM J. Appl. Math, 66(5):1632–1648, 2006.

Other references II

F. Cokelaer, H. Talbot, and J. Chanussot.
Efficient robust d-dimensional path operators.
IEEE Journal of Selected Topics in Signal Processing, 6(7):830 –839, nov. 2012.

François Cokelaer.
Améliorations des ouvertures par chemins pour l'analyse d'images à N dimensions et implémentations optimisées.
PhD thesis, Université de Grenoble, 2013.

A. Dufour, O. Tankyevych, B. Naegel, H. Talbot, C. Ronse, J. Baruthio, P. Dokladal, and N. Passat.
Filtering and segmentation of 3D angiographic data: Advances based on mathematical morphology.
Medical Image Analysis, 17(2):147–164, 2013.

Other references III

Alice Dufour, Christian Ronse, Joseph Baruthio, Olena Tankyevych, Hugues Talbot, and Nicolas Passat.
Morphology-based cerebrovascular atlas.
In *Biomedical Imaging (ISBI), 2013 IEEE 10th International Symposium on*, pages 1210–1214. IEEE, 2013.

A. Enquobahrie, L. Ibanez, E. Bullitt, and S. Aylward.
Vessel enhancing diffusion filter.
ISC/NA-MIC Workshop on Open Science at MICCAI, 7 2007.

H. Heijmans, M. Buckley, and H. Talbot.
Path-based morphological openings.
In *Proceedings of IEEE ICIP 2004*, pages 3085–3088, Singapore, October 2004.

Other references IV

H. Heijmans, M. Buckley, and H. Talbot.
Path openings and closings.
Journal of Mathematical Imaging and Vision, 22:107–119, 2005.

M. W. K. Law and A. C. S. Chung.
Three dimensional curvilinear structure detection using optimally oriented flux.
In *Proc. Eur. Conf. Comput. Vis.*, volume 5305 of *Lect. Notes Comput. Sci.*, pages 368–382. Springer, 2008.

C.L. Luengo Hendriks.
Constrained and dimensionality-independent path openings.
Image Processing, IEEE Transactions on, 19(6):1587–1595, 2010.

Other references V

Adriënne M Mendrik, Evert-Jan Vonken, Annemarieke Rutten, Max A Viergever, and Bram van Ginneken.
Noise reduction in computed tomography scans using 3-d anisotropic hybrid diffusion with continuous switch.
IEEE Trans. Med. Imag., 28(10):1585–1594, 2009.

Odyssée Merveille, Olivia Miraucourt, Stéphanie Salmon, Nicolas Passat, and Hugues Talbot.
A variational model for thin structure segmentation based on a directional regularization.
In *Image Processing (ICIP), 2016 IEEE International Conference on*, pages 4324–4328. IEEE, 2016.

Other references VI

Odyssee Merveille, Hugues Talbot, Laurent Najman, and Nicolas Passat.
Tubular structure filtering by ranking orientation responses of path operators.
In *Computer Vision–ECCV 2014*, pages 203–218. Springer, 2014.

Odyssée Merveille, Hugues Talbot, Laurent Najman, and Nicolas Passat.
Ranking orientation responses of path operators: Motivations, choices and algorithmics.
In *International Symposium on Mathematical Morphology and Its Applications to Signal and Image Processing*, pages 633–644. Springer, 2015.

Other references VII

Tuan-Anh Nguyen, Alice Dufour, Olena Tankyevych, Amir Nakib, Eric Petit, Hugues Talbot, and Nicolas Passat.
Thin structure filtering framework with non-local means, gaussian derivatives and spatially-variant mathematical morphology.
In *Image Processing (ICIP), 2013 20th IEEE International Conference on*, pages 1237–1241, Melbourne, Australia, September 2013.

Leonid I. Rudin, Stanley Osher, and Emad Fatemi.
Nonlinear total variation based noise removal algorithms.
Phys. D, 60(1-4):259–268, 1992.

Other references VIII

- Eysteinn Mor Sigurosson, Silvia Valero, Jon Atli Benediktsson, Jocelyn Chanussot, Hugues Talbot, and Einar Stefánsson.
  Automatic retinal vessel extraction based on directional mathematical morphology and fuzzy classification.
  Pattern Recognition Letters, 47:164–171, 2014.

- Joes Staal, Michael D Abràmoff, Meindert Niemeijer, Max A Viergever, and Bram van Ginneken.
  Ridge-based vessel segmentation in color images of the retina.
  Medical Imaging, IEEE Transactions on, 23(4):501–509, 2004.

- H. Talbot and B. Appleton.
  Efficient complete and incomplete paths openings and closings.
  Image and Vision Computing, 25(4):416–425, April 2007.

Other references IX

- O. Tankyevych, H. Talbot, P. Dokladal, and N. Passat.
  Direction-adaptive grey-level morphology. application to 3d vascular brain imaging.
  In *Proceedings of ICIP 2009*, pages 2261–2264, Cairo, Egypt, 2009. IEEE.

- O. Tankyevych, H. Talbot, P. Dokladal, and N. Passat.
  Spatially-variant morpho-hessian filter: efficient implementation and application.
  In *Mathematical Morphology and Its Application to Signal and Image Processing, Proceedings of the 9th International Symposium on Mathematical Morphology (ISMM) 2009*, pages 137–148, Groningen, The Netherlands, 2009.

Other references X

O. Tankyevych, H. Talbot, N. Passat, M. Musacchio, and M. Lagneau.
Angiographic image analysis.
In G. Dougherty, editor, *Medical Image Processing: Techniques and Applications*, pages 115–145. Springer, 2011.
ISBN 978-1441997692.

Olena Tankyevych.
Filtering of thin objects, applications to vascular image analysis.
PhD thesis, Université Paris-Est, October 2010.

Olena Tankyevych, Hugues Talbot, and Petr Dokládal.
Curvilinear morpho-hessian filter.
In *proceedings of the International Symposium on Biomedical Imaging (ISBI)*, pages 1011–1014, 2008.

Other references XI

S. Valero, J. Chanussot, J.A. Benediktsson, and H. Talbot.
Détection automatique du réseau vasculaire rétinien basée sur la morphologie directionnelle et la fusion de décision.
In *Proceedings of GRETSI*, Dijon, France, 2009. INIST-CNRS. Paru.

S. Valero, J. Chanussot, J.A. Benediktsson, H. Talbot, and B. Waske.
Directional mathematical morphology for the detection of the road network in very high resolution remote sensing images.
In *Proceedings of ICIP 2009*, pages 3725–3728, Cairo, Egypt, 2009.

Other references XII

S. Valero, J. Chanussot, J.A. Benediktsson, H. Talbot, and B. Waske.
Advanced directional mathematical morphology for the detection of the road network in very high resolution images.
Pattern Recognition Letters, 31(10):1120–1127, July 2010.