Lectures on
Stochastic Methods for Image Analysis

Agnès Desolneux

CNRS, CMLA, ENS Cachan

4th Stochastic Geometry Days,
Poitiers, 24-26 August 2015

LECTURE 2

A CONTRARIO DETECTION
OF GEOMETRIC STRUCTURES IN IMAGES

PART I : DETECTION OF ALIGNMENTS IN AN IMAGE

Meaningful Alignments : Introduction

Let $\{u(i,j)\}_{1 \leqslant i,j \leqslant N}$ be a grey level image of size $N \times N$ pixels.

A discrete oriented **segment** S of length l is a sequence of l pixels, determined by its starting point (pixel) x_1 and its ending point (pixel) x_l.

The image domain contains a **finite** number of segments.

Let $m(l)$ be the number of discrete oriented segments of length $l \geqslant 1$, then

$$\sum_{l=1}^{l_{max}} m(l) = N^2(N^2 - 1) \simeq N^4.$$

A. Desolneux, L. Moisan and J.-M. Morel, Meaningful Alignements, *International Journal of Computer Vision*, 2000.

The direction at a point (i, j) is

$$\vec{d}(i,j) = \frac{\nabla u^{\perp}}{\|\nabla u\|},$$

where

X1	X2
X3	X4

$$\nabla u = \frac{1}{2} \begin{pmatrix} X_2 - X_1 + X_4 - X_3 \\ X_3 - X_1 + X_4 - X_2 \end{pmatrix}.$$

Exercise : show that it is the gradient of the second-order interpolation at the center of the 2×2 window.

We say that a point M is aligned with a direction $\vec{v}$ up to precision p if

$$|Angle(\vec{d}(M), \vec{v})| \leqslant p\pi.$$

Null Hypothesis

Hypothesis $\mathcal{H}_0$

The directions at the pixels of the image are independent and uniformly distributed on $[0, 2\pi)$*.*

This hypothesis is satisfied in a white noise image if we only consider pixels at distance $\geqslant 2$.

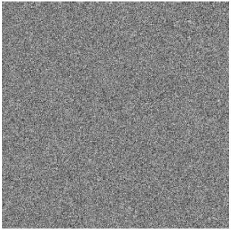

Null Hypothesis

Hypothesis $\mathcal{H}_0$

The directions at the pixels of the image are independent and uniformly distributed on $[0, 2\pi)$.

This hypothesis is satisfied in a white noise image if we only consider pixels at distance $\geqslant 2$.

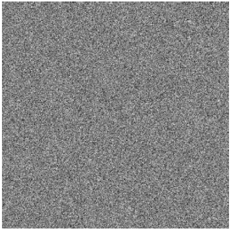

$$\mathbb{P}_{\mathcal{H}_0}[\text{dir. at a point is aligned with dir. } \vec{v} \text{ up to precision } p] = \frac{2\pi p}{2\pi} = p.$$

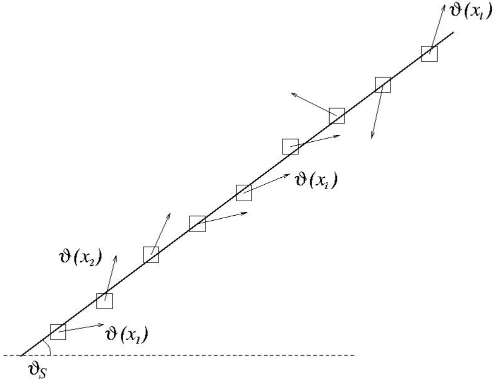

ϑ(x_l)
ϑ(x_i)
ϑ(x_2)
ϑ(x_1)
ϑ_S

Grouping principle

Principle

On a given segment S *of length* l*, we count the number of points of* S *that are aligned with* S *up to precision* p*. Let* k *be this number. If this number is high enough to be very unlikely under the null hypothesis* $\mathcal{H}_0$*, we keep the segment.*

Grouping principle

Principle

On a given segment S of length l, we count the number of points of S that are aligned with S up to precision p. Let k be this number. If this number is high enough to be very unlikely under the null hypothesis $\mathcal{H}_0$, we keep the segment.

Let $\mathbf{S} = \{x_1, \dots x_l\}$ be a discrete segment of length l (counted in independent points, i.e. at distance 2) in an image of Gaussian white noise of size $N \times N$.

Let X_i be the random variable that has value 1 if x_i is aligned with $\mathbf{S}$ up to precision p, and 0 otherwise. The random variable X_i follows a Bernoulli distribution of parameter p.

Let $S_l = \sum_{i=1}^{l} X_i$ = number of aligned points with the direction of $\mathbf{S}$, then :

$$\mathbb{P}_{\mathcal{H}_0}[S_l = k] = \binom{l}{k} p^k (1-p)^{l-k}$$

and then $\mathbb{P}_{\mathcal{H}_0}[S_l \geqslant k] = \mathcal{B}(l,k,p) := \sum_{j \geqslant k} \binom{l}{j} p^j (1-p)^{l-j}.$

ε-meaningful segment

Definition

A segment $\mathbf{S}$ of length l in a $N \times N$ image is said ε-meaningful if it contains at least $k(l)$ aligned points, where

$$k_{\min}(l) := \min\{k \in \mathbb{N},\ \mathcal{B}(l,k,p) \leqslant \frac{\varepsilon}{N^4}\}.$$

Rk : $\mathcal{B}(l,k,p)$ is the tail of the binomial distribution of parameters l and p, and it is a decreasing function of k.

Expectation of the number of detections under $\mathcal{H}_0$

Proposition

The expectation of the number of ε-meaningful segments in a random image of size $N \times N$ pixels following the null hypothesis $\mathcal{H}_0$, is less than ε.

$\longrightarrow$ On the average, the number of detected ε-meaningful segments in a Gaussian white noise image is less than ε.

Proof :

Let $e_i = 1$ if the i-th segment of the image is ε-meaningful, and 0 otherwise.

Let R be the number of ε-meaningful segments in the image. Then

$$\mathbb{E}_{\mathcal{H}_0}[R] = \sum_{i=1}^{Nseg} \mathbb{E}[e_i] = \sum_{i=1}^{Nseg} \mathbb{P}_{\mathcal{H}_0}[S_{l_i} \geqslant k_{\min}(l_i)] = \sum_{i=1}^{Nseg} \mathcal{B}(l_i, k_{\min}(l_i), p) \leqslant Nseg \times \frac{\varepsilon}{N^4} \leqslant \varepsilon.$$

Number of false alarms

Definition

Let S be a segment of length $l(S)$ containing $k(S)$ aligned points. The number of false alarms (NFA) of S is defined by

$$\mathrm{NFA}(S) = \mathrm{NFA}(l(S), k(S)) := N^4 \mathcal{B}(l(S), k(S), p) = N^4 \sum_{k=k(S)}^{l(S)} \binom{l(S)}{k} p^k (1-p)^{l(S)-k}.$$

Number of false alarms

Definition

Let S be a segment of length $l(S)$ containing $k(S)$ aligned points. The number of false alarms (NFA) of S is defined by

$$\mathrm{NFA}(S) = \mathrm{NFA}(l(S), k(S)) := N^4 \mathcal{B}(l(S), k(S), p) = N^4 \sum_{k=k(S)}^{l(S)} \binom{l(S)}{k} p^k (1-p)^{l(S)-k}.$$

$\rightarrow$ it measures the degree of confidence, or the “meaningfulness” of an observed alignment.

The smaller $\mathrm{NFA}(S)$ is, the more meaningful the segment S is.

Number of false alarms

Definition

Let S be a segment of length $l(S)$ containing $k(S)$ aligned points. The number of false alarms (NFA) of S is defined by

$$\mathrm{NFA}(S) = \mathrm{NFA}(l(S), k(S)) := N^4 \mathcal{B}(l(S), k(S), p) = N^4 \sum_{k=k(S)}^{l(S)} \binom{l(S)}{k} p^k (1-p)^{l(S)-k}.$$

$\rightarrow$ it measures the degree of confidence, or the “meaningfulness” of an observed alignment.

The smaller $\mathrm{NFA}(S)$ is, the more meaningful the segment S is.

Link between NFA and meaningfulness :

$$S \text{ is } \varepsilon\text{-meaningful} \iff k(S) \geqslant k_{\min}(l(S)) \iff \mathrm{NFA}(S) \leqslant \varepsilon.$$

Some properties of the NFA

Using elementary properties of the binomial distribution, we have :

1. $NFA(l,0) = N^4$
2. $NFA(l,l) = N^4 p^l$
3. $NFA(l,k+1) < NFA(l,k)$
4. $NFA(l,k) < NFA(l+1,k)$ $\quad \rightarrow \rightarrow \bullet \rightarrow\rightarrow \bullet$
5. $NFA(l+1,k+1) < NFA(l,k)$ $\quad \rightarrow \rightarrow \bullet \rightarrow\rightarrow\rightarrow$

Proof : exercise !

Remark : If a segment S of length $l(S)$ is ε-meaningful, then

$$\varepsilon \geqslant \mathrm{NFA}(S) \geqslant N^4 p^{l(S)},$$

and therefore

$$l(S) \geqslant \frac{4 \log N - \log \varepsilon}{-\log p}.$$

For instance for $N = 512$, $p = 1/16$ and $\varepsilon = 1$, we have $l(S) \geqslant 9$.

Sufficient Conditions of meaningfulness

We assume that $p < 0.5$.

Proposition (sufficient condition)

Let S be a segment of length l, containing k aligned points. If

$$k \geqslant pl + \sqrt{\frac{4\log N - \log\varepsilon}{h(p)}}\sqrt{l},$$

where $p \mapsto h(p)$ is the function given by

$$h(p) = \frac{1}{1-2p}\log\frac{1-p}{p}$$

then S is ε-meaningful.

Proof : Use Hoeffding Inequality, to get for any $1 \geqslant r \geqslant p > 0$,

$$\mathcal{B}(l, rl, p) \leqslant e^{-l[r\log\frac{r}{p} + (1-r)\log\frac{1-r}{1-p}]} \leqslant e^{-l(r-p)^2 h(p)}.$$

Necessary Conditions of meaningfulness

Proposition (Necessary condition 1)

Assume that $pN^4 \geqslant 1$ (Rk : not very restrictive !). Let S be a segment of length l, containing k aligned points. If S is 1-meaningful then

$$k \geqslant pl + (1-p).$$

Proposition (Necessary condition 2)

Assume that $p \leqslant \frac{1}{4}$ and $pN^4 \geqslant 1$. Let S be a segment of length l, containing k aligned points. If S is ε-meaningful then

$$k \geqslant pl + \alpha(N,\varepsilon)\sqrt{lp(1-p)},$$

where $\alpha(N,\varepsilon)$ is defined by $\frac{1}{\sqrt{2\pi}}\int_{\alpha(N,\varepsilon)}^{+\infty} e^{-x^2/2}dx = \frac{\varepsilon}{N^4}$.

Proof of Condition 2 : Consequence of Slud's Theorem (1977), stating that if $p \leqslant \frac{1}{4}$ and $pl \leqslant k \leqslant l$, then

$$\mathcal{B}(l,k,p) \geqslant \int_{\frac{k-pl}{\sqrt{lp(1-p)}}}^{+\infty} e^{-x^2/2}dx.$$

Asymptotique behavior of $k(l)$

We have $k_{\min}(l) \simeq pl + \sqrt{Cl \log \frac{N^4}{\varepsilon}}$, with $2p(1-p) \leqslant C \leqslant \frac{1}{2}$.

$\rightarrow$ Notice that $k_{\min}(l)$ depends on $\log(\varepsilon)$ and $\log(N)$. Choice $\varepsilon = 1$ is generally satisfying in experiments.

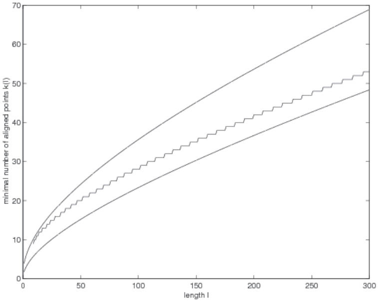

Middle curve : $k_{\min}(l)$ for $N = 512$, $\varepsilon = 1$ and $p = \frac{1}{16}$. Top curve : Sufficient condition. Bottom curve : Necessary condition 2.

Dependence in ε

$\varepsilon = 0.1.$

Dependence in ε

$\varepsilon = 0.01.$

Dependence in ε

$\varepsilon = 0.001.$

Maximality

Generally, S meaningful $\Longrightarrow$ Many segments containing S or being contained in it are also meaningful.

Maximality

Generally, S meaningful $\Longrightarrow$ Many segments containing S or being contained in it are also meaningful.

Definition

A segment S is said maximal meaningful if it is meaningful and if

$$\forall B \subset S, \quad \mathrm{NFA}(B) \geqslant \mathrm{NFA}(S),$$
$$\forall B \supset S, \quad \mathrm{NFA}(B) > \mathrm{NFA}(S).$$

Properties of maximal meaningful segments :

- The two ending points of S are aligned with S,
- The two points, one before and one after S, are not aligned with S.

Maximality

Conjecture

If S_1 *and* S_2 *are two distinct maximal meaningful segments lying on the same straight line, and such that* $S_1 \cap S_2 \neq \emptyset$*, the*

$$\min(\mathrm{NFA}(S_1 \cup S_2), \mathrm{NFA}(S_1 \cap S_2)) < \max(\mathrm{NFA}(S_1), \mathrm{NFA}(S_2)).$$

Consequence :

Two maximal meaningful segments lying on the same straight line cannot meet.

Remark : The above conjecture is equivalent to say that

$$\min(\mathcal{B}(l_1 + l_2 - l_\cap, k_1 + k_2 - k_\cap, p), \mathcal{B}(l_\cap, k_\cap, p)) < \max(\mathcal{B}(l_1, k_1, p), \mathcal{B}(l_2, k_2, p)).$$

Maximality

$\varepsilon = 10^{-3}$.

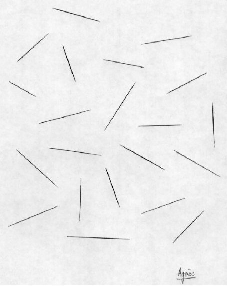 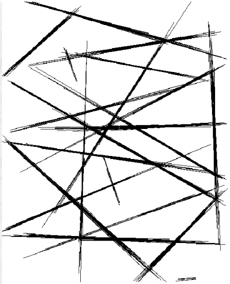

Maximality

$\varepsilon = 10^{-3}$.

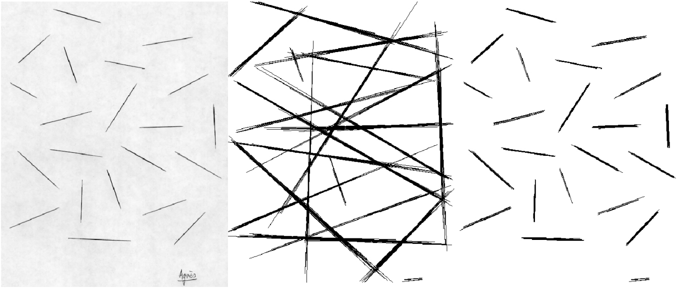

Example of maximal meaningful segments

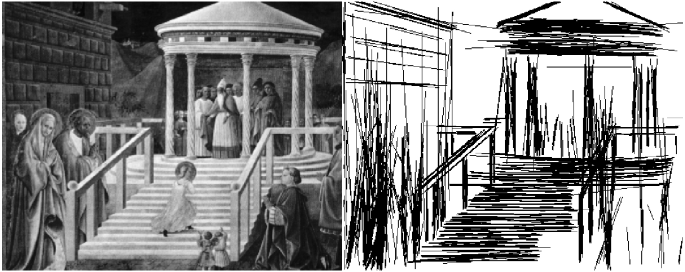

Meaningful Segments : defaults of the method

The main defaults of the method are :

- It is very slow !
- We often get « bundles « of segments.

$\longrightarrow$ Solution : make « thick » segments.

LSD : Line Segment Detector [Grompone et al. 2008]

LSD Algorithm

- Partition of the image into *Line-Support Regions* (connected sets of pixels sharing the same orientation up to precision $p\pi$) ;
- Approximation of these sets by rectangular regions ;
- Computation of the NFA of each region : for a region (rectangle) r containing l points with k of them aligned with it, define

$$\mathrm{NFA}(r) = N^5 \sum_{j=k}^{l} \binom{l}{j} p^j (1-p)^{l-j}.$$

- The regions such that $\mathrm{NFA} < 1$ are kept.

R. Grompone von Gioi, J. Jakubowicz, J.-M. Morel, G. Randall, A Fast Line Segment Detector with a False Detection Control, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2010.

Examples

Examples

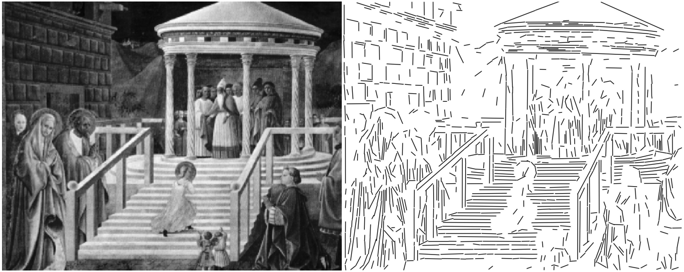

PART II : DETECTION OF CONTRASTED OR SMOOTH CURVES

Meaningful Boundaries : Introduction

- Classical problem in Image Processing : find the boundaries (contours) in an image (also called edge detection problem).
- Dual problem : image segmentation (= segment the image into homogeneous regions).
- Principle of color (or grey level) constancy in Gestalt Theory.
- Perception is invariant under contrast changes.
- What are the « interesting curves » in an image ?
  - All possible curves ? (pb of computation time !)
  - Natural candidates : the level lines of the image (or pieces of them).

A. Desolneux, L. Moisan and J.-M. Morel, Edge Detection by Helmholtz Principle, *Journal of Mathematical Imaging and Vision*, 2001.

Level sets

Definition

The (upper) level sets of an image $u : \Omega \to \mathbb{R}$ are the sets

$$\chi_\lambda(u) = \{x \in \mathbb{R}^2; u(x) \geqslant \lambda\}, \text{ where } \lambda \in \mathbb{R}.$$

These sets are « decreasing » : $\forall \lambda \leqslant \mu, \quad \chi_\mu \subset \chi_\lambda$.
The knowledge of all these sets is enough to reconstruct u by

$$u(x) = \sup\{\lambda;\ x \in \chi_\lambda\}.$$

Definition

The lower level sets of an image $u : \Omega \to \mathbb{R}$ are the sets

$$\chi^\lambda(u) = \{x \in \mathbb{R}^2; u(x) \leqslant \lambda\}, \text{ where } \lambda \in \mathbb{R}.$$

These sets are « increasing » : $\forall \lambda \leqslant \mu, \quad \chi^\lambda \subset \chi^\mu$.

If g is a change of contrast (increasing function), then u and $g(u)$ globally have the same level sets ($\forall \lambda\ \exists \mu$ s.t. $\chi_\lambda(u) = \chi_\mu(g(u))$).

Level sets

Definition

The (upper) level sets of an image $u : \Omega \to \mathbb{R}$ are the sets

$$\chi_\lambda(u) = \{x \in \mathbb{R}^2; u(x) \geqslant \lambda\}, \text{ where } \lambda \in \mathbb{R}.$$

These sets are « decreasing » : $\forall \lambda \leqslant \mu, \quad \chi_\mu \subset \chi_\lambda$.
The knowledge of all these sets is enough to reconstruct u by

$$u(x) = \sup\{\lambda;\ x \in \chi_\lambda\}.$$

Definition

The lower level sets of an image $u : \Omega \to \mathbb{R}$ are the sets

$$\chi^\lambda(u) = \{x \in \mathbb{R}^2; u(x) \leqslant \lambda\}, \text{ where } \lambda \in \mathbb{R}.$$

These sets are « increasing » : $\forall \lambda \leqslant \mu, \quad \chi^\lambda \subset \chi^\mu$.

If g is a change of contrast (increasing function), then u and $g(u)$ globally have the same level sets ($\forall \lambda\ \exists \mu$ s.t. $\chi_\lambda(u) = \chi_\mu(g(u))$).

$\Rightarrow$ **Def. level lines of u = topological boundaries of its level sets**

Examples

Level lines for levels λ that are multiples of 5.

Examples

Level lines for levels λ that are multiples of 20.

Examples

Level lines for levels λ that are multiples of 50.

Contrasted Boundaries

Let u be a discrete grey level image of size $N \times N$ and let N_{ll} be the (finite) number of level lines it contains.

Let L be a level line of u with (discrete) length l measured with points at distance 2, that is $L = \{x_1, \dots x_l\}$.

We define the **contrast of u at a point** x by $c(x) = |\nabla u(x)|$.

Definition

A meaningful boundary is a level line that is « long enough » and « contrasted enough » not to appear just by chance.

Contrasted Boundaries

Let u be a discrete grey level image of size $N \times N$ and let N_{ll} be the (finite) number of level lines it contains.

Let L be a level line of u with (discrete) length l measured with points at distance 2, that is $L = \{x_1, \dots x_l\}$.

We define the **contrast of u at a point** x by $c(x) = |\nabla u(x)|$.

Definition

A meaningful boundary is a level line that is « long enough » and « contrasted enough » not to appear just by chance.

What *a contrario* noise model (hypothesis $\mathcal{H}_0$) on the contrast ?

- uniform ?
- empirical : $H(\mu) = \frac{1}{N^2} \#\{x\,;\ |\nabla u_0|(x) \geqslant \mu\}$
- To avoid flat regions artefacts, take $H(\mu) = \frac{\#\{x\,;\ |\nabla u_0|(x) \geqslant \mu\}}{\#\{x\,;\ |\nabla u_0|(x) \neq 0\}}$

Contrasted Boundaries

Let X_i be the contrast at a point x_i.
The hypothesis $\mathcal{H}_0$ is that the X_i are i.i.d.
Under $\mathcal{H}_0$, we have

$$\mathbb{P}_{\mathcal{H}_0}[\forall i,\ X_i \geqslant \mu] = \mathbb{P}_{\mathcal{H}_0}[X_1 \geqslant \mu]^l.$$

Definition

Let H be the empirical distribution of $\mathbb{P}[|\nabla u| \geqslant \mu]$. We define the Number of False Alarms of a level line L with discrete length l and minimal contrast $\mu = \min_{x \in L} |\nabla u(x)|$ by

$$\mathrm{NFA}(L) = N_{ll}\, H(\mu)^l.$$

The level line L is said to be an ε-meaningful boundary iff $\mathrm{NFA}(L) \leqslant \varepsilon$.

Proposition

*The ε-meaningful level lines are **invariant** under affine contrast changes.*

Examples

Image and the distribution of the norm of the gradient.

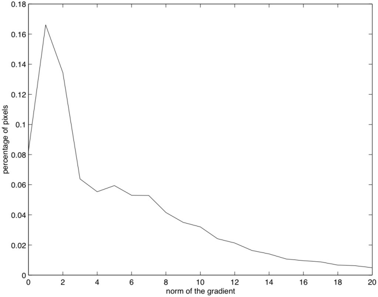

Properties

Let $F(\mu, l) = H(\mu)^l$.

- If $l \leqslant l'$ and μ fixed, then $F(\mu, l) \geqslant F(\mu, l')$ (because $H(\mu) \leqslant 1$).
- If $\mu \leqslant \mu'$ and l fixed, then $F(\mu, l) \geqslant F(\mu', l)$ (because H is decreasing).
- A level line with minimal contrast μ is $\varepsilon-$meaningful iff its length is larger than
$$l_{\min}(\mu) = \frac{\log \varepsilon - \log N_{ll}}{\log H(\mu)}.$$
- A level line with length l is $\varepsilon-$meaningful iff its minimal contrast μ is larger than
$$\mu_{\min}(l) = H^{-1}\left(\left(\frac{\varepsilon}{N_{ll}}\right)^{\frac{1}{l}}\right).$$

Maximality

The set of all meaningful level lines of an image is organized in a tree structure.

Definition

A **monotonic branch** in the tree of level lines is a branch along which the grey level is monotonic and such that each level line has a unique child. A monotonic branch is maximal if it is not contained in another monotonic branch.

Definition

A level line is a **maximal meaningful boundary** if it is meaningful and if its NFA is minimal in its maximal monotonic branch of the tree of level lines.

Examples

On the left, the original image. In the middle, all meaningful boundaries with $\varepsilon = 1$. On the right, all maximal meaningful boundaries with $\varepsilon = 1$.

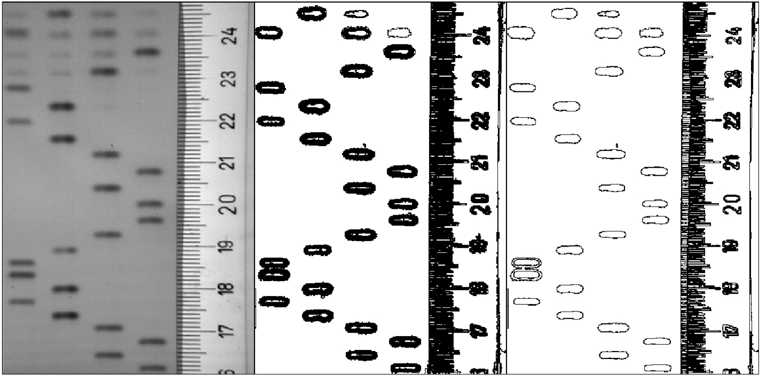

Meaningful good continuations

Goal : Look for « smooth » curves in the image, without considering the contrast.

Let $\Gamma = (p_0, \dots, p_{l+1})$ be a discrete curve of length l, and let κ be its maximal discrete curvature :

$$\kappa = \max_{1 \leqslant i \leqslant l} |\text{Angle}(p_{i+1} - p_i, p_i - p_{i-1})|.$$

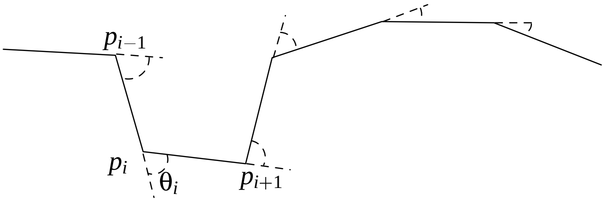

From : F.Cao. “Application of the Gestalt principles to the detection of good continuations and corners in image level lines”, *Computing and Visualization in Science*, 2004.

A contrario noise model $\mathcal{H}_0$: the angles are i.i.d. with uniform law on $[0, 2\pi)$, *i.e.* the curve is a discrete » random walk ».
Let N_c be the number of considered curves (in practice, it will be the number of pieces of level lines of the image).

Definition (Meaningful Good Continuation)

We say that a discrete curve Γ is an ε-meaningful good continuation if

$$\kappa < \frac{\pi}{2} \quad \textit{and} \quad \mathrm{NFA}(\Gamma) = N_c \left(\frac{\kappa}{\pi}\right)^l \leqslant \varepsilon.$$

Definition of maximality : a meaningful good continuation Γ is maximal meaningful if : $\forall \Gamma' \subset \Gamma$, $NFA(\Gamma') \geqslant NFA(\Gamma)$ and $\forall \Gamma' \supsetneq \Gamma$, $NFA(\Gamma') > NFA(\Gamma)$.

Property : if Γ and Γ' are two maximal meaningful good continuations lying on the same level line then $\Gamma \cap \Gamma' = \emptyset$. (Exercise !)

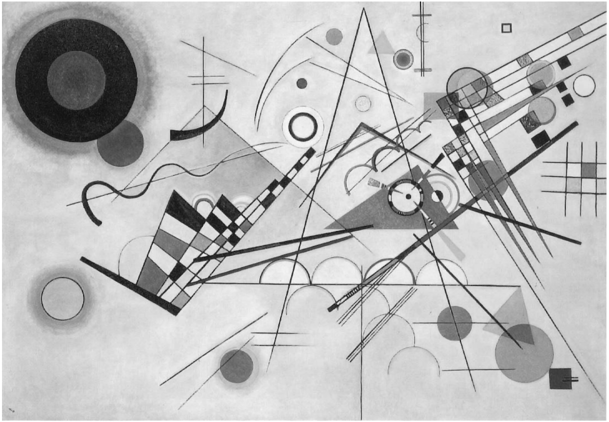

FIGURE: Original image : Kandinsky.

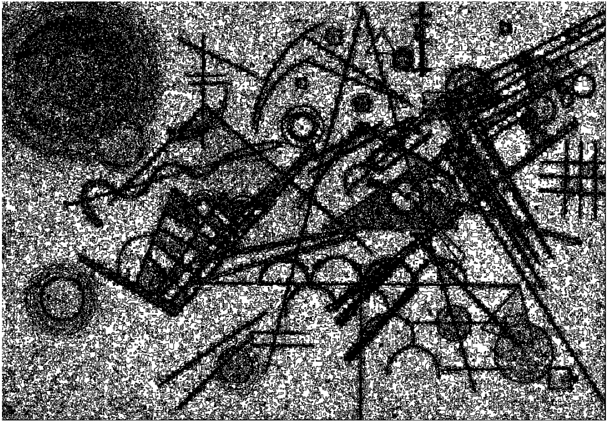

FIGURE: « All » level lines.

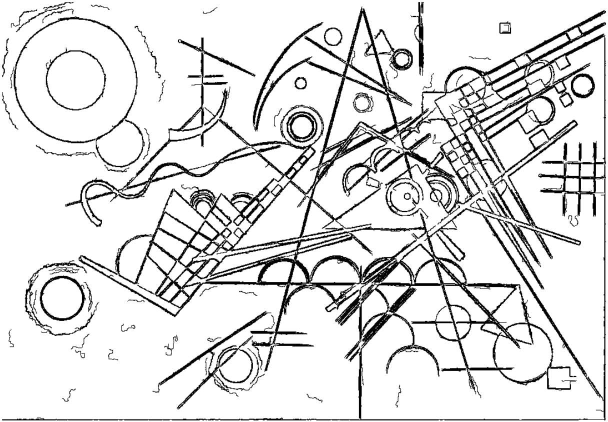

FIGURE: Maximal meaningful good continuations.

FIGURE: Image (INRIA) of the Church of Valbonne.

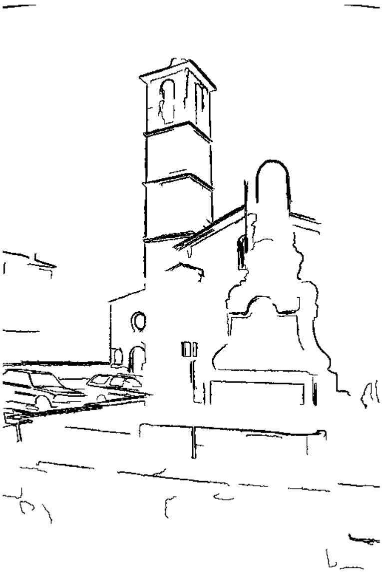

FIGURE: Maximal meaningful good continuations.

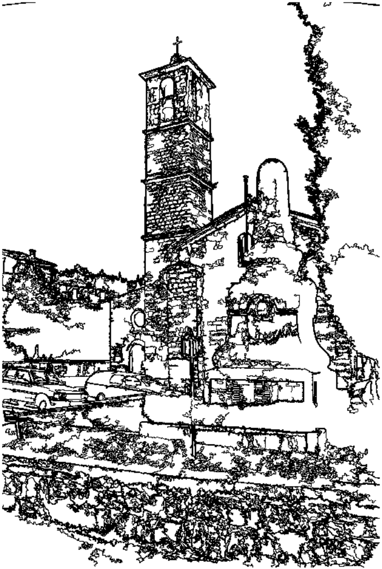

FIGURE: Maximal meaningful contrasted boundaries.

Another example

Original image

Another example

Meaningful segments (LSD algorithm of Grompone et. al)

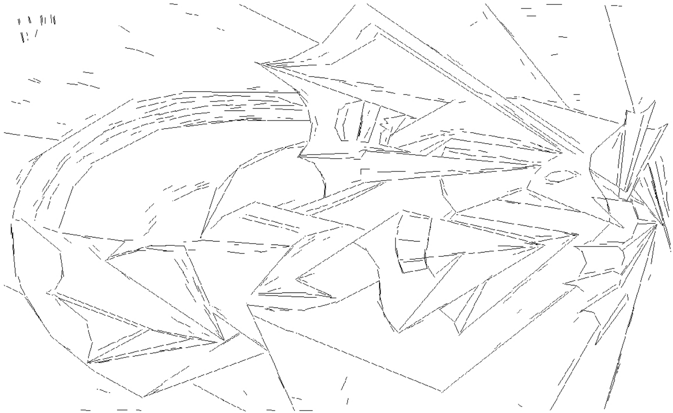

Another example

Maximal meaningful contrasted boundaries

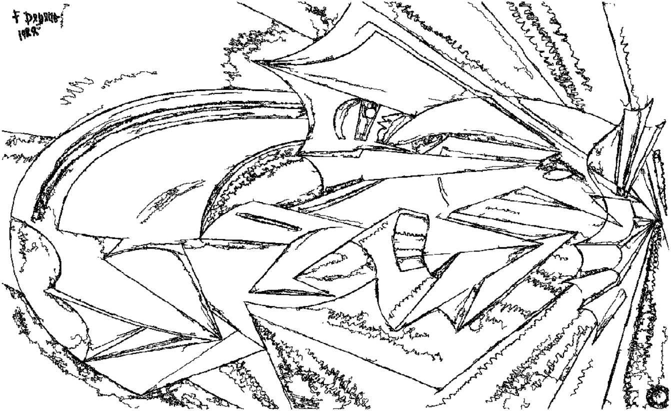

Another example

Maximal meaningful good continuations

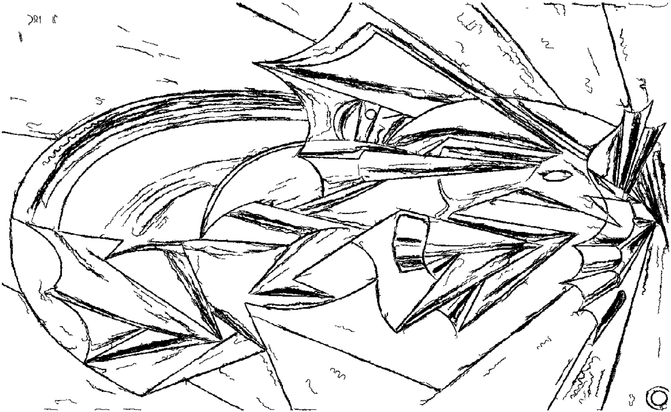

PART III : OTHER A CONTRARIO DETECTIONS AND DISCUSSION

Similarity of a scalar attribute

For a uniform scalar attribute (gray level, orientation, etc.)
Assume we have M "objects", and each of them has an attribute $q \in \{1, 2, \ldots, L\}$. Let a group of k of them has their scalar attribute q such that $a \leqslant q \leqslant b$. Define its Number of False Alarm by

$$\mathrm{NFA}([a,b]) = \frac{L(L+1)}{2} \cdot \mathcal{B}\left(M, k, \frac{b-a+1}{L}\right)$$

For a scalar attribute with decreasing distribution (area, length, etc.)
Define its Number of False Alarm by

$$\mathrm{NFA}([a,b]) = \frac{L(L+1)}{2} \cdot \max_{p \in \mathcal{D}} \mathcal{B}\left(M, k, \sum_{i=a}^{b} p(i)\right)$$

where $\mathcal{D}$ is the set of decreasing probability distributions on $\{1, ..., L\}$.

A first example

A first application is the study of an image grey-level histogram. Looking for the maximal meaningful intervals is a way to obtain an automatic gray level quantization.

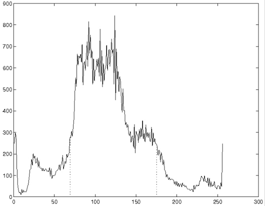

A second example : recursivity

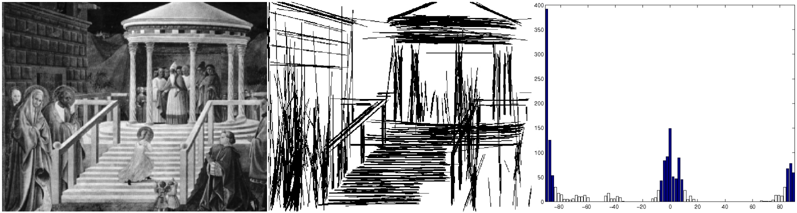

Uccello's painting : maximal meaningful alignments and histogram of orientations. Two maximal meaningful modes are found corresponding respectively to the horizontal and vertical segments.

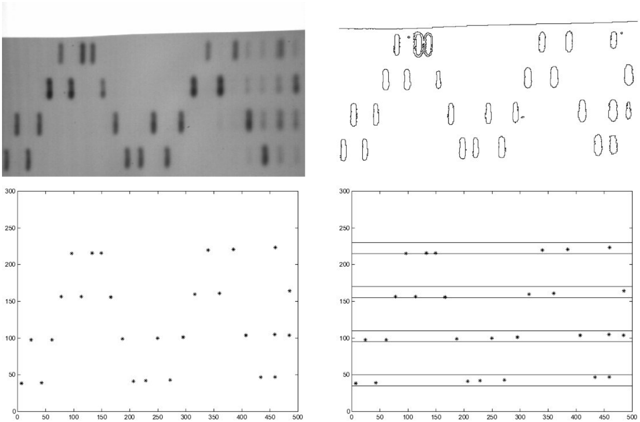

Gestalt grouping principles at work for building an “order 3” gestalt (alignment of blobs of the same size). First row : original DNA image (left) and its maximal meaningful boundaries (right). Second row : left, barycenters of all meaningful regions whose area is inside the only maximal meaningful mode of the histogram of areas ; right, meaningful alignments of these points.

Conclusion

- Helmholtz principle combined with Gestalt grouping laws can become a powerful computational tool.

Conclusion

- Helmholtz principle combined with Gestalt grouping laws can become a powerful computational tool.
- Many applications in detection problems, but also in shape recognition [Musé, Sur, Cao, Gousseau] ; image matching [Rabin, Delon, Gousseau] ; epipolar geometry [Moisan, Stival], motion detection and analysis [Cao, Veit, Bouthemy] ; clustering [Cao et al.] ; stereovision [Sabater et al.] ; image denoising (by grain filters) ; etc.

Conclusion

- Helmholtz principle combined with Gestalt grouping laws can become a powerful computational tool.
- Many applications in detection problems, but also in shape recognition [Musé, Sur, Cao, Gousseau] ; image matching [Rabin, Delon, Gousseau] ; epipolar geometry [Moisan, Stival], motion detection and analysis [Cao, Veit, Bouthemy] ; clustering [Cao et al.] ; stereovision [Sabater et al.] ; image denoising (by grain filters) ; etc.
- Automatic computation of thresholds. Link with human visual perception thresholds ?

Open questions

1. How to deal with the “over-determination” of images (i.e. the fact that visual objects usually have several qualities at the same time) ?

Open questions

1. How to deal with the "over-determination" of images (i.e. the fact that visual objects usually have several qualities at the same time) ?
2. How to define a computational tool (like the NFA) for more than one attribute ?

Open questions

1. How to deal with the "over-determination" of images (i.e. the fact that visual objects usually have several qualities at the same time) ?
2. How to define a computational tool (like the NFA) for more than one attribute ?
3. How to deal with "conflicts" of qualities ?

Examples of conflicts (1)

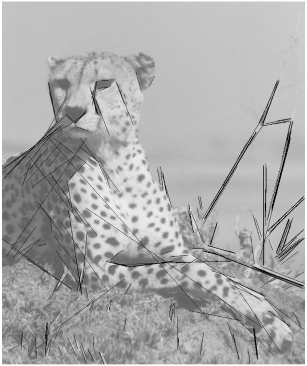

FIGURE: Smooth convex sets or alignments ?

The casual alignments in the Cheetah fur are caused by the presence of many oval shapes. Such alignments are perceptually masked and should be computationally masked !

Examples of conflicts (2)

A dense cluster of points creates a meaningful amount of dots in many strips and the result is the detection of obviously wrong alignments. Again, the detection of a cluster should inhibit such alignment detections.

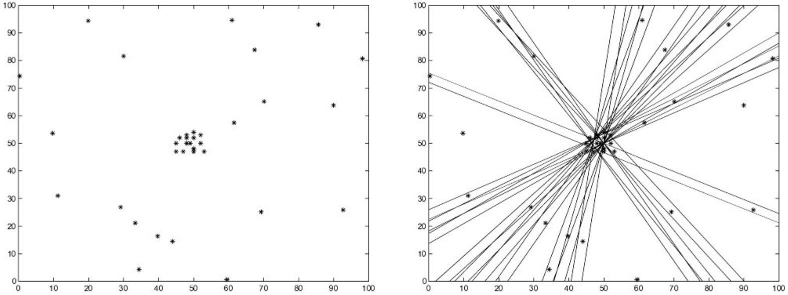

FIGURE: One cluster or several alignments ?

Cooperation of attributes of objects

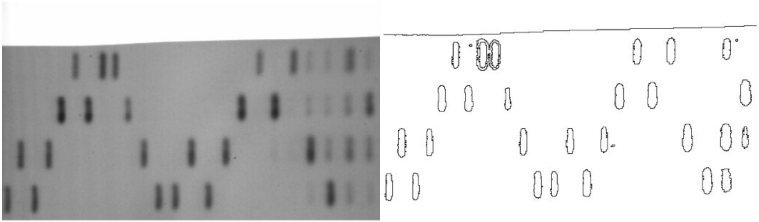

PART IV : MAKE SOME EXPERIMENTS BY YOURSELF

Reproducible Research in Image Processing

IPOL Journal = Image Processing On Line

« IPOL is a research journal of image processing and image analysis which emphasizes the role of mathematics as a source for algorithm design and the reproducibility of the research. Each article contains a text on an algorithm and its source code, with an online demonstration facility and an archive of experiments. Text and source code are peer-reviewed and the demonstration is controlled. IPOL is an Open Science and Reproducible Research journal. »

Let's go to http://www.ipol.im/