Lectures on Stochastic Methods for Image Analysis

Agnès Desolneux

CNRS, CMLA, ENS Cachan

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

LECTURE 3

STOCHASTIC GEOMETRY FOR
THE DETECTION OF VANISHING POINTS

PART I : VANISHING POINTS

What’s a vanishing point ?

In a « pinehole » camera model, parallel straight lines in 3D are projected on the image plane as 2D lines that intersect at a single point.

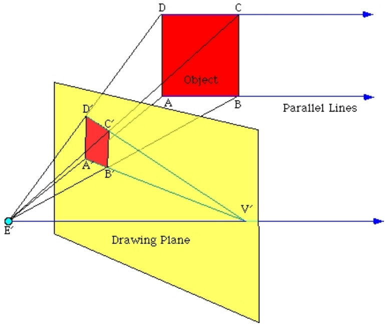

This principle is used to create perspective effects in images. The intersection points are called **vanishing points**.

What's a vanishing point ?

In a « pinehole » camera model, parallel straight lines in 3D are projected on the image plane as 2D lines that intersect at a single point.

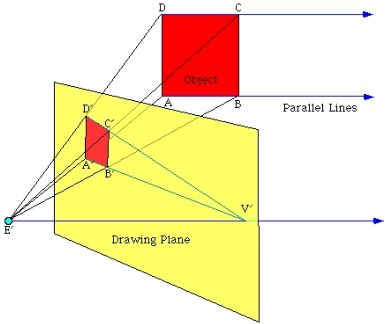

This principle is used to create perspective effects in images. The intersection points are called **vanishing points**.

Why is it useful to detect vanishing points ?

- Camera calibration
- Photogrammetry (the science of making measurements from photographs)
- 3D Reconstruction of a scene

Example 1

Example 2

The Ideal City, Piero della Francesca, 1475.

Example 3

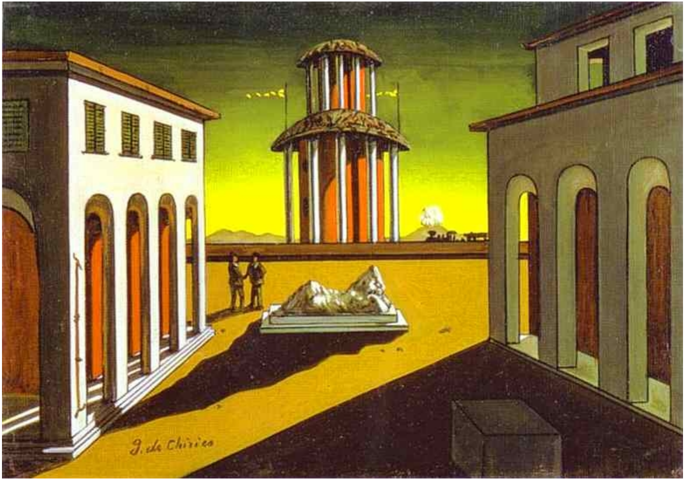

Piazza d'Italia, Giorgio de Chirico, 1913.

A contrario detection of vanishing points

Principle : We start from elementary straight segments detected in the image (by the LSD algorithm for instance), and then we look for regions in the image plane (inside or outside the image domain), such that a « lot of » segments converge towards these regions.

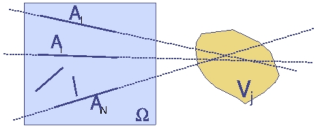

N segments

A contrario detection of vanishing points

Principle : We start from elementary straight segments detected in the image (by the LSD algorithm for instance), and then we look for regions in the image plane (inside or outside the image domain), such that a « lot of » segments converge towards these regions.

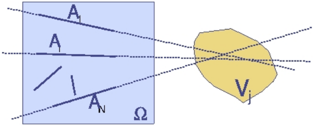

N segments

How to determine what « a lot of » means ?
$\Longrightarrow$ Use the *a contrario* methodology : we look for events that are unlikely to happen by chance.
We start with a noise model $\mathcal{H}_0$: "The N segments are i.i.d. uniform". Then, we look for regions that are intersected by significantly more segments than the number that could be expected under $\mathcal{H}_0$.

A. Almansa, A. Desolneux and S. Vamech, Vanishing Point Detection without Any A Priori Information, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(4), 2003.

Mathematical Formulation

Let Ω be the image domain and let N be the number of elementary segments detected in the image.
Let $D_1, \dots D_N$ be the support lines of the N segments.
We start from a partition of the image plane $\mathbb{R}^2$ into M regions :

$$\mathbb{R}^2 = \bigcup_{j=1}^{M} V_j.$$

Definition

Let V_j be a region and let $k_j = \#\{D_i \text{ s.t. } D_i \cap V_j \neq \emptyset\}$. The Number of False Alarms of V_j is defined by

$$\mathrm{NFA}(V_j) := M \cdot \mathcal{B}(N, k_j, p_j) = M \cdot \sum_{k=k_j}^{N} \binom{N}{k} p_j^k (1-p_j)^{N-k},$$

where p_j is the probability that a random line going through Ω also meets V_j. When $\mathrm{NFA}(V_j) \leqslant \varepsilon$, then we say that the region V_j is ε-meaningful.

Mathematical Formulation

Let Ω be the image domain and let N be the number of elementary segments detected in the image.
Let $D_1, \ldots D_N$ be the support lines of the N segments.
We start from a partition of the image plane $\mathbb{R}^2$ into M regions :

$$\mathbb{R}^2 = \bigcup_{j=1}^{M} V_j.$$

Definition

Let V_j be a region and let $k_j = \#\{D_i \text{ s.t. } D_i \cap V_j \neq \emptyset\}$. The Number of False Alarms of V_j is defined by

$$\mathrm{NFA}(V_j) := M \cdot \mathcal{B}(N, k_j, p_j) = M \cdot \sum_{k=k_j}^{N} \binom{N}{k} p_j^k (1 - p_j)^{N-k},$$

where p_j is the probability that a random line going through Ω also meets V_j. When $\mathrm{NFA}(V_j) \leqslant \varepsilon$, then we say that the region V_j is ε-meaningful.

Proposition

Under $\mathcal{H}_0$, the expected number of ε-meaningful regions is less than ε.

Mathematical Formulation

Let Ω be the image domain and let N be the number of elementary segments detected in the image.
Let $D_1, \ldots D_N$ be the support lines of the N segments.
We start from a partition of the image plane $\mathbb{R}^2$ into M regions :

$$\mathbb{R}^2 = \bigcup_{j=1}^{M} V_j.$$

Definition

Let V_j be a region and let $k_j = \#\{D_i \text{ s.t. } D_i \cap V_j \neq \emptyset\}$. The Number of False Alarms of V_j is defined by

$$\mathrm{NFA}(V_j) := M \cdot \mathcal{B}(N, k_j, p_j) = M \cdot \sum_{k=k_j}^{N} \binom{N}{k} p_j^k (1-p_j)^{N-k},$$

where p_j is the probability that a random line going through Ω also meets V_j. When $\mathrm{NFA}(V_j) \leqslant \varepsilon$, then we say that the region V_j is ε-meaningful.

Proposition

Under $\mathcal{H}_0$, the expected number of ε-meaningful regions is less than ε.

Question : Determine a “good partition” of the plane, and compute p_j.

Computing p_j : a question of stochastic geometry

What means « a random line » ?

Not an obvious question, as shown by Bertrand's Paradox.

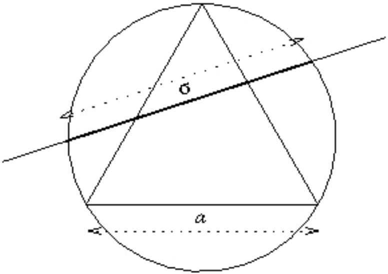

FIGURE: Bertrand's Paradox. Let us consider an equilateral triangle of side length a and let $\mathcal{C}$ be its circumscribed circle. What is the probability that a random line interesecting the circle defines a chord of length $\geqslant a$?

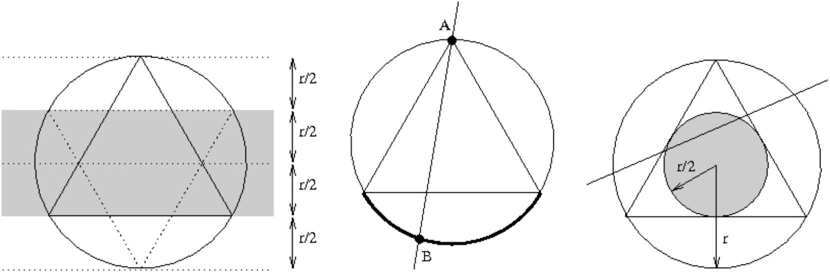

FIGURE: From left to right, depending on the definition of « random line meeting the circle », the answer is respectively $1/2$, $1/3$ or $1/4$.

$\longrightarrow$ What is the « right » answer ?

Measure on the set of lines

A line G of the plane is parametrized by its two polar coordinates : $\rho \geqslant 0$ and $\theta \in [0, 2\pi)$.

$$G(\rho, \theta) = \{(x, y) \in \mathbb{R}^2 \text{ s.t. } x\cos\theta + y\sin\theta = \rho\}.$$

Then, there exists a unique (up to a positive multiplicative constant) measure on the set of lines that is invariant under translations and rotations. It is the Poincaré measure given by

$$d\mu = d\rho\, d\theta.$$

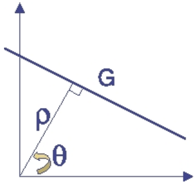

L. Santalo, *Integral Geometry and Geometric Probability*, Cambridge University Press, Second Edition, 2004.

Properties

- Let $K \subset \mathbb{R}^2$ be a bounded closed convex set with non-empty interior. Then

$$\mu\left(\{G \text{ s.t. } G \cap K \neq \emptyset\}\right) = \mathrm{Per}(K),$$

where $\mathrm{Per}(K)$ is the perimeter of K (= length of the boundary).

- Let K_1 and K_2 be two bounded closed convex sets with non-empty interior. Then

$$\mu\left(\{G \text{ s. t. } G \cap K_1 \neq \emptyset \text{ and } G \cap K_2 \neq \emptyset\}\right) = \begin{cases} \mathrm{Per}(K_1) & \text{if } K_1 \subset K_2, \\ L_i - L_e & \text{if } K_1 \cap K_2 = \emptyset, \\ \mathrm{Per}(K_1) + \mathrm{Per}(K_2) - L_e & \text{otherwise.} \end{cases}$$

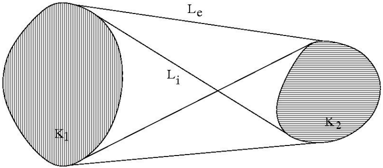

Choice of the plane partition

We want to satisfy two constraints :

1. $p_j =$ constant. All regions are « equally detectable », meaning that they require the same minimal number of lines passing through them in order to become meaningful.
2. Constraint of angular precision :

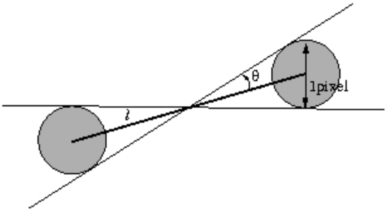

For a segment of length l, its support line is rather a « cone » of angle

$$d\theta = \arcsin\frac{1}{l} \simeq \frac{1}{l}.$$

Construction of the regions

We consider that the image domain is given by $\Omega = D(0, R)$. We fix an angle θ. The exterior regions are defined as portions of circular sectors with angle 2θ.

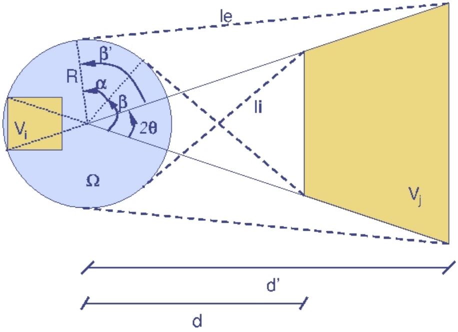

For a region $V_{d,d'}$, portion of a circular sector between the distances d and d', then

$$p_{V_{d,d'}} = \frac{L_i - L_e}{\mathrm{Per}(\Omega)} = \frac{1}{\pi}\left(\tan\beta - \tan\beta' + \frac{1}{\cos\beta'} - \frac{1}{\cos\beta} + \beta' - \beta + 2\theta\right),$$

where $\beta = \arccos\left(\frac{R}{d}\cos\theta\right)$ and $\beta' = \arccos\left(\frac{R}{d'}\cos\theta\right)$

The interior regions are simply chosen as squares of side length $2R\sin\theta$. This implies that

$$p_j = \frac{\mathrm{Per}(V_j)}{\mathrm{Per}(\Omega)} = \frac{4\sin\theta}{\pi} := p_\theta.$$

We can now exactly determine the exterior regions :
- we start with $d_1 = R$,
- then set $d_2 > d_1$ such that $p_{V_{d_1,d_2}} = p_\theta$,
- then $d_3 > d_2$ such that $p_{V_{d_2,d_3}} = p_\theta$,
- and so on, until being larger than d_∞, that is finite and characterized by

$$\forall d' > d_\infty, \quad p_{V_{d_\infty,d'}} < p_\theta.$$

The last region is thus infinite and its probability is $< p_\theta$.

Multi-scale approach

How to choose θ ?
$\longrightarrow$ Use several values !
This is also necessary in order to have a good balance between detectability and localization.

We choose n angular values $\theta_s = \frac{2\pi}{2^s}$ with $s = 4, 5, \ldots, n+3$.

For each θ_s, we have a “partition” of the image plane : $\cup_{j=1}^{M_s} V_{j,s}$. The number of false alarms of a region is then defined as

$$\mathrm{NFA}(V_{j,s}) = n \cdot M_s \cdot B(N_s, k, p_{\theta_s}),$$

where N_s is the number of segments having a precision at least θ_s.

The region $V_{j,s}$ is said ε-meaningful iff $\mathrm{NFA}(V_{j,s}) \leqslant \varepsilon$.

Maximality

When « a lot of « lines pass through a region $V_{j,s}$, then the neighbouring regions are also intersected by « a lot of » lines.
$\longrightarrow$ Need to select the « best regions ».

Definition

The region $V_{j,s}$ is said maximal ε-meaningful if it is ε-meaningful and if

$$\forall s' \in [s_1, \ldots, s_n], \quad \forall j' \in [1, \ldots, M_{s'}], \quad \overline{V_{j',s'}} \cap \overline{V_{j,s}} = \emptyset \Longrightarrow \mathrm{NFA}(V_{j,s}) \leqslant \mathrm{NFA}(V_{j',s'}).$$

Example of results (1)

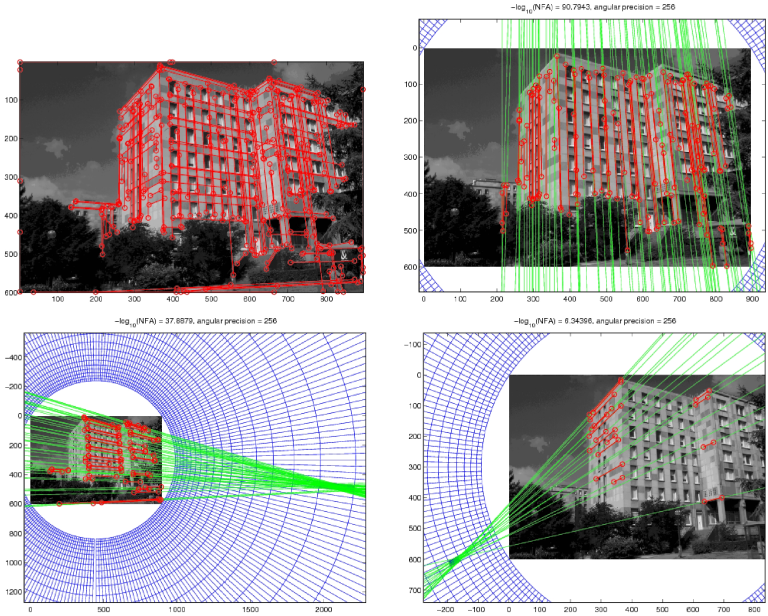

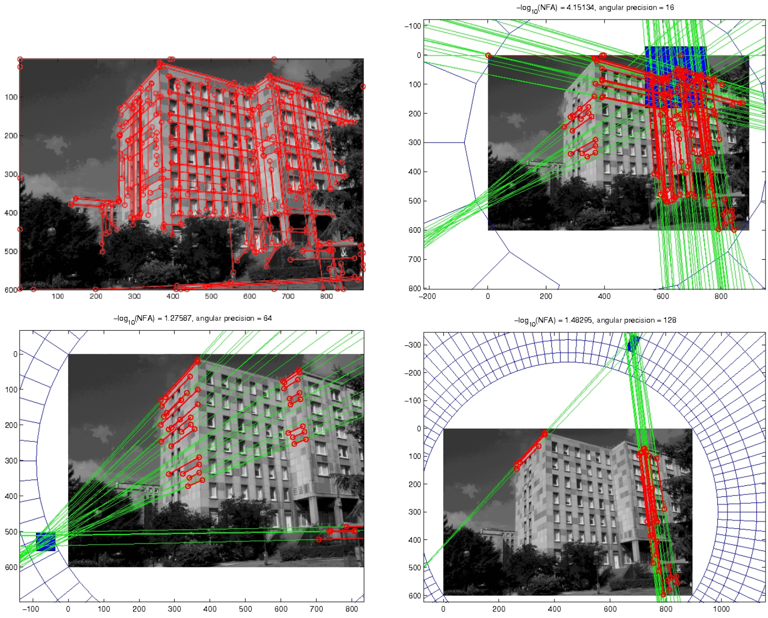
$-\log_{10}(\text{NFA}) = 4.15134$, angular precision = 16
$-\log_{10}(\text{NFA}) = 1.27587$, angular precision = 64
$-\log_{10}(\text{NFA}) = 1.48295$, angular precision = 128

Exclusion Principle

The problem : Some maximal meaningful regions are the « mixture » of two sets of lines that converge to two different vanishing points.

Proposed solution : an exclusion principle

- Each line votes only for one maximal meaningful region : the one that has the minimal NFA.
- We compute again the NFA of maximal meaningful regions by counting only the number of lines that have voted for them.
- If the maximal meaningful regions still satisfy the test "$\mathrm{NFA} \leqslant \varepsilon$", then they are said EP-meaningful.

In the following we will only show the EP-meaningful regions.

Example of results (2)

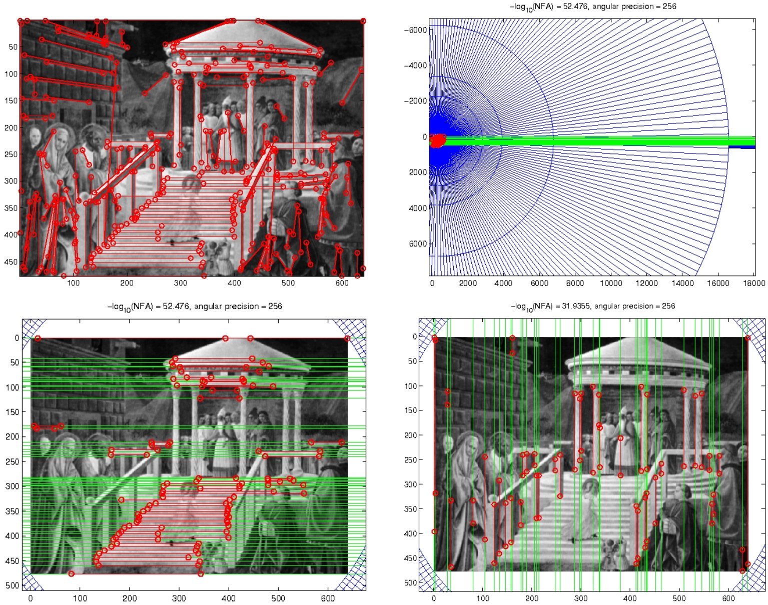

The masking phenomenon

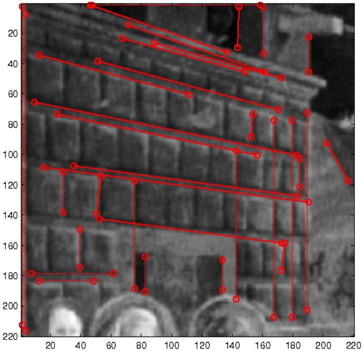

$-\log_{10}(\mathrm{NFA}) = 4.75545$, angular precision = 256

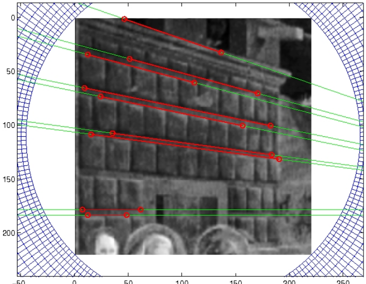

$-\log_{10}(\mathrm{NFA}) = 12.4332$, angular precision = 256

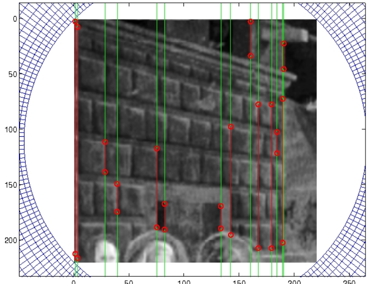

Example of results (3)

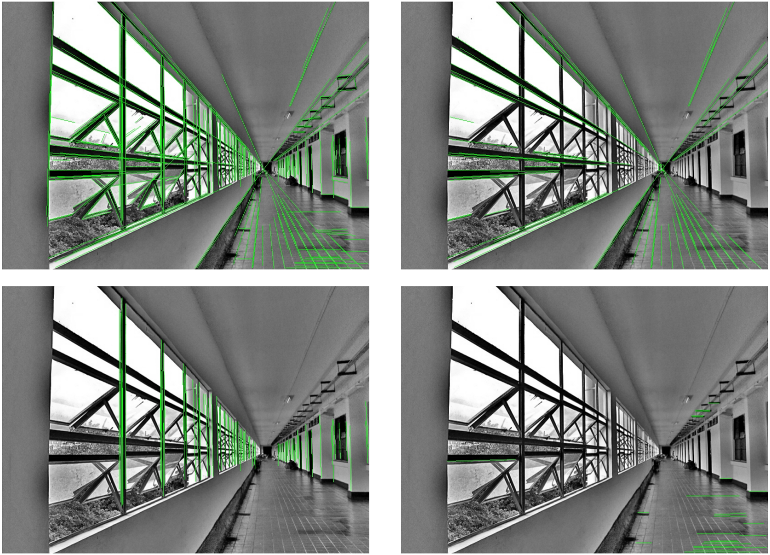

Example of results (4)

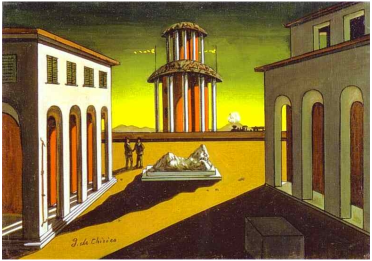
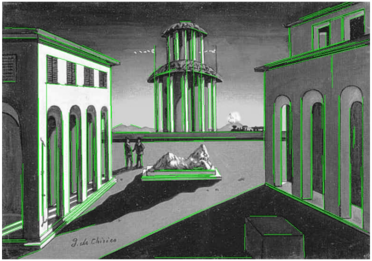
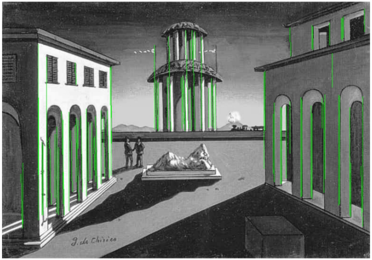
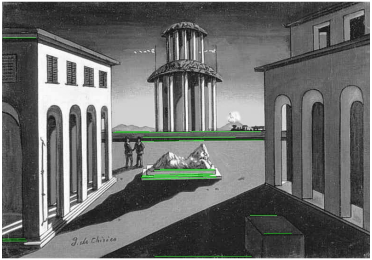

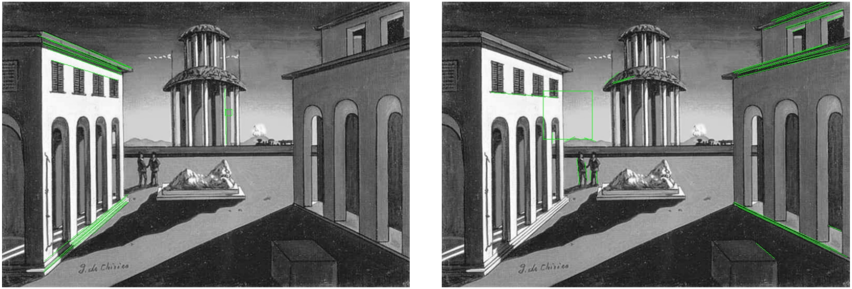

FIGURE: Left : the 3rd and last EP-meaningful region. Right : masked region (it is not meaningful but becomes meaningful when we remove the lines that have votes for the EP-meaningful regions).

PART II : CHANGING THE A CONTRARIO NOISE MODEL

Anisotropic detection of convergence points

In some images, a main convergence point is « normal », and we don't want to detect it as meaningul, we rather want to include it in the a contrario noise model.

Example : a mammogram

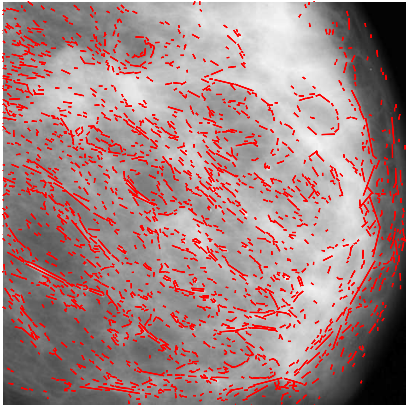

An anisotropic distribution on lines

Given a main point of convergence M, and a « width » σ, we can define a Gaussian law on lines by

$$
\begin{aligned}
d\mu_g &= \frac{1}{\pi\sqrt{2\pi}\sigma} e^{-(\rho - x_M cos\theta + y_M \sin\theta)^2/2\sigma^2} \, d\rho \, d\theta \\
&= \frac{1}{\pi\sqrt{2\pi}\sigma} e^{-(\rho - r_M \cos(\theta - \theta_M))^2/2\sigma^2} \, d\rho \, d\theta
\end{aligned}
$$

This is equivalent to say that the signed distance from a random line to the point M, that is equal to $\rho - r_M \cos(\theta - \theta_M)$, follows a $\mathcal{N}(0, \sigma^2)$ distribution.

A. Desolneux and F. Doré, An anisotropic a contrario framework for the detection of convergences in images, *submitted*, 2015.

Uniform versus Gaussian law on lines

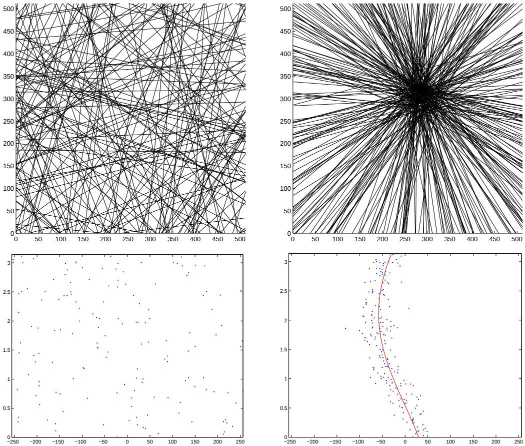

FIGURE: Left column : a set of 200 lines sampled from the uniform measure conditioned to meet the image domain Ω. Right column : a set of 200 lines sampled from the Gaussian law on lines conditioned to intersect Ω with $\sigma = 20$ and $(x_M, y_M) = (286, 306)$.

Support function of a convex set

Definition

Let K be a closed bounded convex set. The support function of K is defined for all $\varphi \in [0, 2\pi)$ by

$$s_K(\varphi) = \sup_{\mathbf{x} \in K} \langle \mathbf{x}, \mathbf{e}_\varphi \rangle,$$

where $\mathbf{e}_\varphi$ is the unit vector having an angle φ with the horizontal axis, and $\langle \cdot, \cdot \rangle$ is the usual Euclidean scalar product in $\mathbb{R}^2$.

Link with the Perimeter :

$$\mathrm{Per}\, K = \int_0^{2\pi} s_K(\varphi) d\varphi.$$

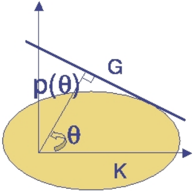

Useful Formulas

Proposition

Let K be a closed bounded convex set and let s_K denote its support function. The measure under the law μ_g of the set of lines meeting K is given by

$$\mu_g(D \cap K \neq \emptyset) = \frac{1}{\pi}\int_{\varphi=0}^{\pi}\left[\Phi\left(\frac{s_K(\varphi) - r_M\cos(\theta_M - \varphi)}{\sigma}\right) - \Phi\left(\frac{-s_K(\varphi+\pi) - r_M\cos(\theta_M - \varphi)}{\sigma}\right)\right]d\varphi,$$

where Φ is the cumulative distribution function of the standard normal distribution, i.e.

$$\forall t \in \mathbb{R}, \quad \Phi(t) := \int_{-\infty}^{t} \frac{1}{\sqrt{2\pi}} e^{-u^2/2}\, du.$$

We can also compute a similar formula for two convex sets, i.e. for

$$\mu_g(D \cap K_1 \neq \emptyset \text{ and } D \cap K_2 \neq \emptyset).$$

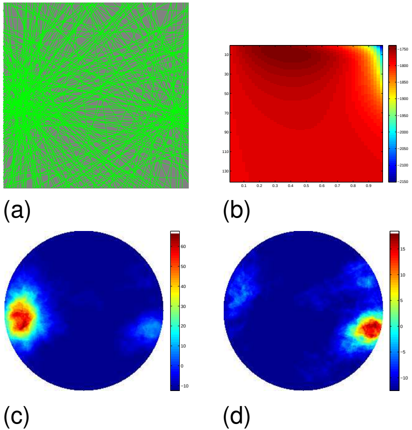

FIGURE: (a) We simulated $N = 300$ straight lines with a parametric mixture model with two Gaussian terms and one uniform term : the main convergence is centered in $(x_M, y_M) = (-100, 14)$ with a standard variation $\sigma = 20$ and a weight $p = 0.4$, and the secondary convergence has center $(x_{M'}, y_{M'}) = (114, 28)$, $p' = 0.2$ and $\sigma' = 10$. The uniform term has thus weight $1 - p - p' = 0.4$. (b) Image of the log-likelihood as a function of p and σ. It is maximal for $(\hat{p}, \hat{\sigma}) = (0.37, 17)$. (c) Image of $-\log \mathrm{NFA}_1$. (d) Image of $-\log \mathrm{NFA}_{f_{\hat{p},M,\hat{\sigma}}}$.

PART III : WHEN THE A CONTRARIO APPROACH BECOMES GENERATIVE

Some questions about the a contrario methodology

In short, the **a contrario methodology** works like this :

- Observe N_{test} geometric events $E_1, \dots E_{N_{test}}$ (special arrangements of elementary objects).
- Define a noise model H_0 on the elementary objects (generally, the i.i.d. uniform distribution).
- Compute for each event E_j, $\mathrm{NFA}(E_j) := N_{test} \cdot \mathbb{P}_{H_0}(E_j)$ and declare the event E_j as being ε-meaningful when $\mathrm{NFA}(E_j) \leqslant \varepsilon$.

Some questions about the a contrario methodology

In short, the **a contrario methodology** works like this :

- Observe N_{test} geometric events $E_1, \ldots E_{N_{test}}$ (special arrangements of elementary objects).
- Define a noise model H_0 on the elementary objects (generally, the i.i.d. uniform distribution).
- Compute for each event E_j, $\mathrm{NFA}(E_j) := N_{test} \cdot \mathbb{P}_{H_0}(E_j)$ and declare the event E_j as being ε-meaningful when $\mathrm{NFA}(E_j) \leqslant \varepsilon$.

Some questions :

- What happens if we change H_0 ?
- Given the observed events, are there distributions H_0 under which no events are meaningful ?
- If yes, which one is at the same time « as random » as possible ?
- What does a sample from this distribution look like ?

The Line Segment Detector framework

The framework of the **LSD (Line Segment Detector)** Algorithm of Grompone et. al is the following :

Let us consider a grey level image I^0 defined on a discrete domain $\Omega = \{1, \dots, M\} \times \{1, \dots, N\}$. We compute its orientation field $\theta^0 : \Omega \to S^1$ by

$$\forall x \in \Omega, \quad \theta^0(x) = \frac{\pi}{2} + \operatorname{Arg} \frac{\nabla I^0(x)}{\|\nabla I^0(x)\|},$$

where ∇I^0 is the gradient computed on a 2×2 window.

For a rectangle r in Ω with principal orientation $\varphi(r)$, we define the number of aligned points it contains, up to a precision p, by

$$k(r; \theta^0) := \sum_{x \in r} \mathbb{1}_{|\theta^0(x) - \varphi(r)| \leqslant p\pi}.$$

We denote $n(r) = \#r$ the total number of pixels in the rectangle r.

- Let us define a « noise model », that is a probability distribution P on orientation fields $\Theta : \Omega \to S^1$.

- We define the Number of False Alarms of the rectangle r in θ_0, under the noise model P by

$$\mathrm{NFA}_P(r; \theta^0) = N_{tests} \cdot \mathbb{P}_P[k(r; \Theta) \geqslant k(r; \theta^0)],$$

where N_{tests} is the number of tests, that is the number of rectangles in an image of size $M \times N$ (it is of the order of $(MN)^{5/2}$).

- In this definition, Θ is random orientation field following the distribution P and $k(r; \Theta)$ is therefore a random variable given by

$$k(r; \Theta) := \sum_{x \in r} 1\!\!\mathrm{I}_{|\Theta(x) - \varphi(r)| \leqslant p\pi}.$$

- When $\mathrm{NFA}_P(r; \theta^0) < \varepsilon$, we say that the rectangle r is (ε)-P-meaningful in θ^0.

In the LSD algorithm :

- The noise model is $P = U$, the i.i.d. uniform distribution (*i.e.* we assume that the $\Theta(x)$ are independent and uniformly distributed on S^1).
- The result of the LSD algorithm is a list of U-rectangles, denoted by $r_1, \ldots, r_m$.
- Several precision values p are tested in the LSD algorithm, but one can notice that in most images, at least 95% of the meaningful rectangles are obtained for the precision $p = 1/8$. In all the following, we will use only this fixed value for p.
- The obtained meaningful rectangles $r_1, \ldots, r_m$ are disjoint (it is almost true, since in practice less than 1% of the pixels belong to two meaningful rectangles at the same time).

Changing the noise model

Definition

Let $\theta^0 : \Omega \to S^1$ be an orientation field. Let $r_1, \dots, r_m$ be the (disjoint) U-meaningful rectangles, results of the LSD Algorithm on θ^0. We define the 3 following sets of distributions :

- Let $\mathcal{P}$ be the set of distributions P on Θ such that none of the regions $r_1, \dots, r_m$ are P-meaningful in θ^0. That is :

$$P \in \mathcal{P} \iff \forall 1 \leqslant j \leqslant m, \quad \mathrm{NFA}_P(r_j; \theta^0) \geqslant \varepsilon.$$

- Let $\mathcal{Q}$ be the set of distributions Q on Θ such that if an orientation field θ is sampled from Q, then, « in most cases, we have the same detections as in θ_0 ». That is :

$$Q \in \mathcal{Q} \iff \forall 1 \leqslant j \leqslant m, \quad \mathrm{Med}_Q(\mathrm{NFA}_U(r_j; \Theta)) \leqslant \mathrm{NFA}_U(r_j; \theta^0).$$

- Finally, let $\mathcal{I}$ be the set of distributions on Θ such that the $\Theta(x)$ are independent (but not necessarily identically distributed).

Proposition

1. *The distribution* $P_0 \in \mathcal{P} \cap \mathcal{I}$ *that has maximal entropy is given by the following probability density :* $f_{P_0}(\theta) = \prod_x f_{P_0}^{(x)}(\theta_x)$ *with*

$$f_{P_0}^{(x)}(\theta_x) = \begin{cases} \frac{1}{2\pi} & \text{if } x \notin \cup_{j=1}^m r_j \\ \frac{1}{2p\pi} B^{-1}_{n(r_j),k(r_j;\theta^0)}\left(\frac{\varepsilon}{N_{tests}}\right) & \text{if } x \in r_j \text{ and } |\theta_x - \varphi(r_j)| \leqslant p\pi \\ \frac{1}{2(1-p)\pi}\left(1 - B^{-1}_{n(r_j),k(r_j;\theta^0)}\left(\frac{\varepsilon}{N_{tests}}\right)\right) & \text{if } x \in r_j \text{ and } |\theta_x - \varphi(r_j)| > p\pi \end{cases}$$

2. *The distribution* $Q_0 \in \mathcal{Q} \cap \mathcal{I}$ *that has maximal entropy is given by a formula analoguous to the one of* P_0*, where*

$$\text{we just replace } B^{-1}_{n(r_j),k(r_j;\theta^0)}\left(\frac{\varepsilon}{N_{tests}}\right) \text{ by } B^{-1}_{n(r_j),k(r_j;\theta^0)}\left(\frac{1}{2}\right) \simeq \frac{k(r_j;\theta^0)}{n(r_j)}.$$

Example 1

Original image I^0

Results of the LSD

Example 1

Original image I^0

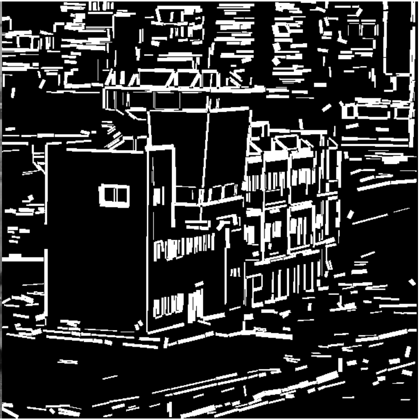

Rectangles of the LSD

Example 1

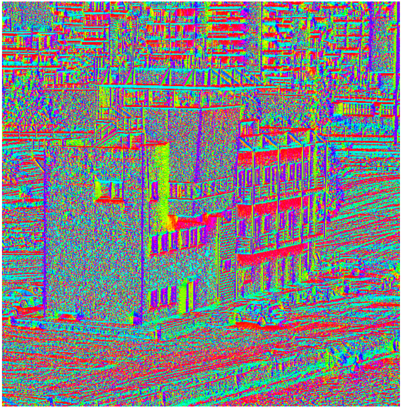

Orientation field θ^0

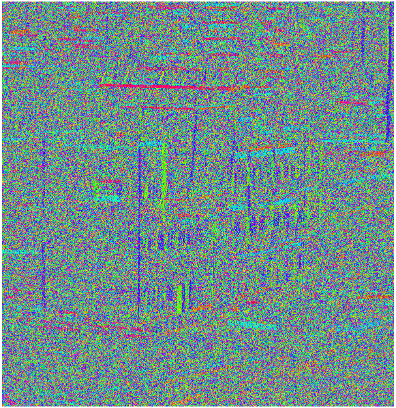

Sample from P_0

Example 1

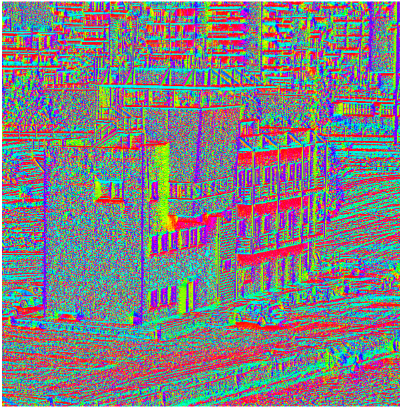

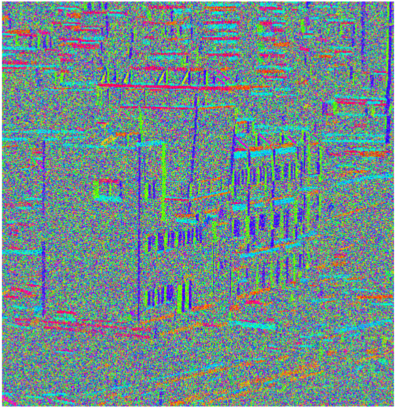

Orientation field θ^0

Sample from Q_0

Example 2

Original image I^0

Results of the LSD

Example 2

Original image I^0

Rectangles of the LSD

Example 2

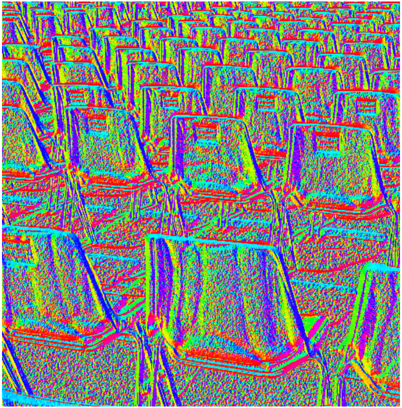

Orientation field θ^0

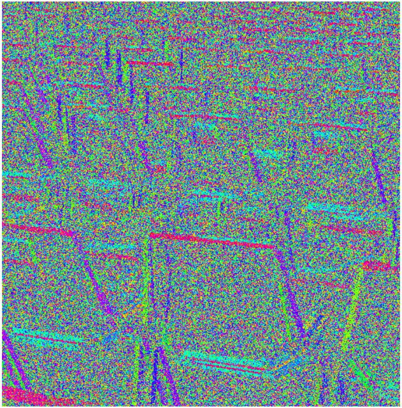

Sample from P_0

Example 2

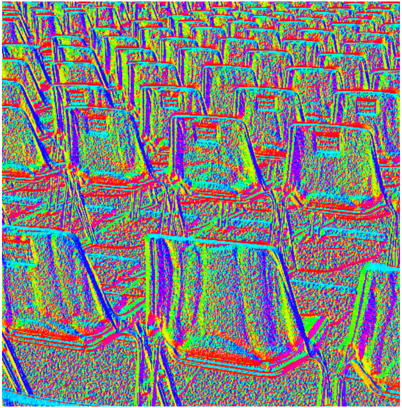

Orientation field θ^0

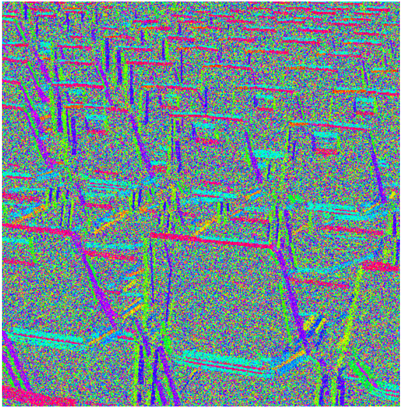

Sample from Q_0

Example 3

Original image I^0

Results of the LSD

Example 3

Original image I^0

Rectangles of the LSD

Example 3

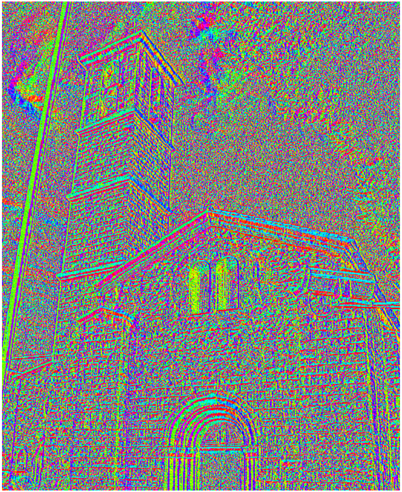

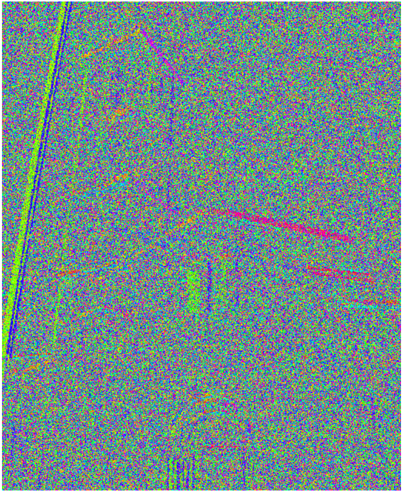

Orientation θ^0

Sample from P_0

Example 3

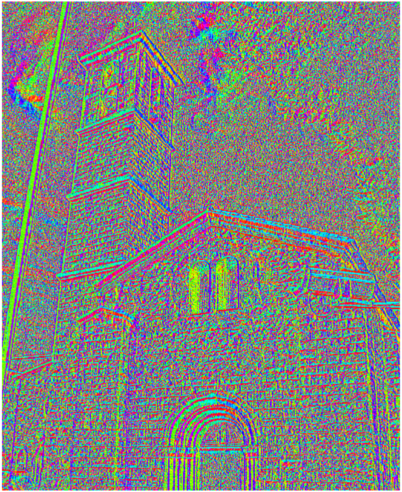

Orientation θ^0

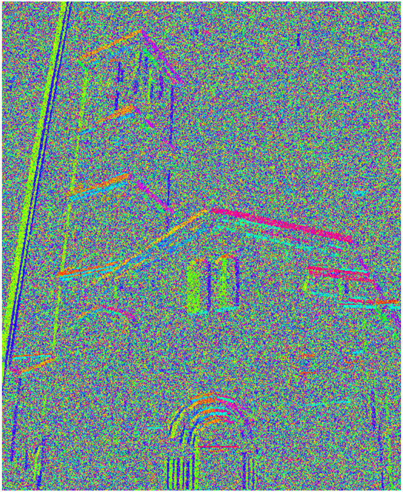

Sample from Q_0

Reconstructing an image from an orientation field

Question : How to reconstruct an image from an orientation field θ ?

$\longrightarrow$ We do like in the so-called « Poisson editing » method of Perez et. al. : we look for an image u defined on Ω such that

$$\int_\Omega |\nabla u(x) - R(x)e^{i\theta(x)}|^2 \, dx \quad \text{is minimal ,}$$

where the $R(x)$ are given amplitudes.

Solution : just solve $\Delta u = \mathrm{div}(Re^{i\theta})$ (very easily with the discrete Fourier transform)

Results on example 1

Original image I^0

Reconstruction from P_0 and R_{rand}

Results on example 1

Original image I^0

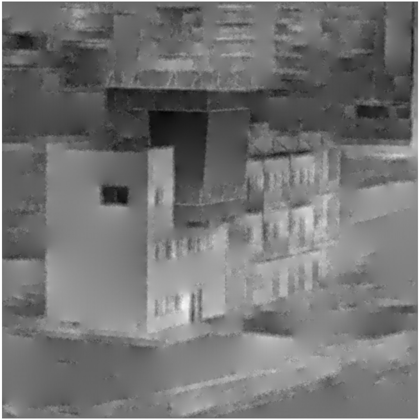

Reconstruction from P_0 and R_{100}

Results on example 1

Original image I^0

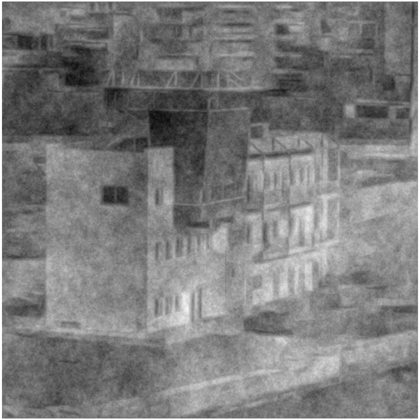

Reconstruction from Q_0 and R_{rand}

Results on example 1

Original image I^0

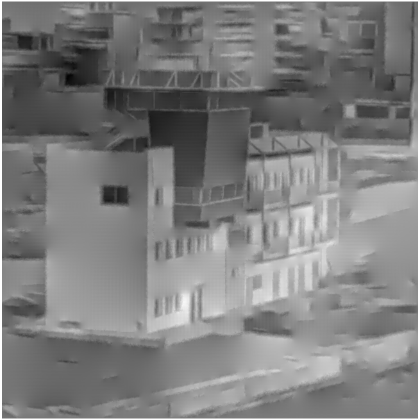

Reconstruction from Q_0 and R_{100}

Results on example 1

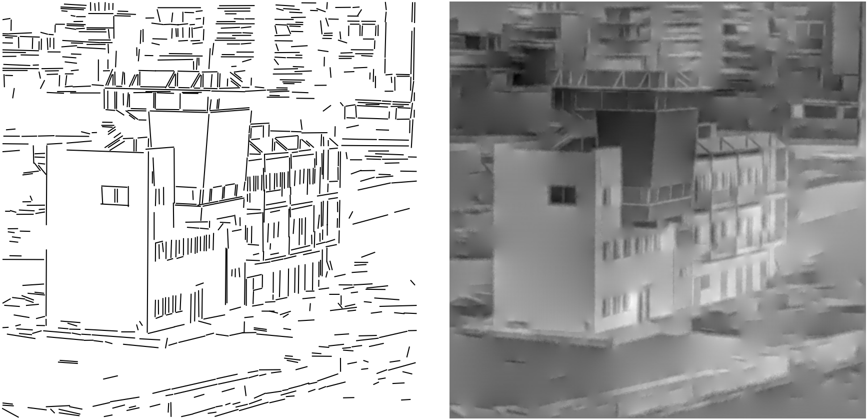

Reconstruction from Q_0 and R_{100}

Results on example 2

Original image I^0

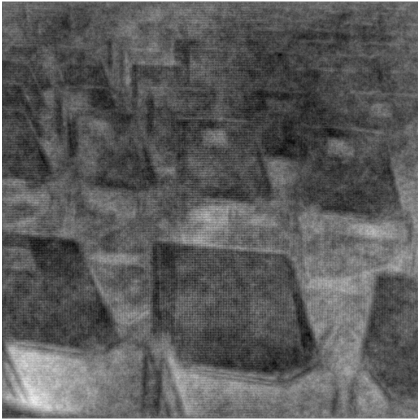

Reconstruction from P_0 and R_{rand}

Results on example 2

Original image I^0

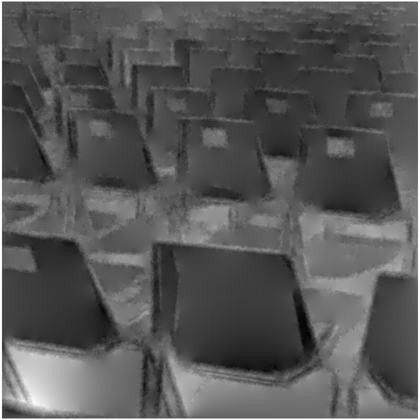

Reconstruction from P_0 and R_{100}

Results on example 2

Original image I^0

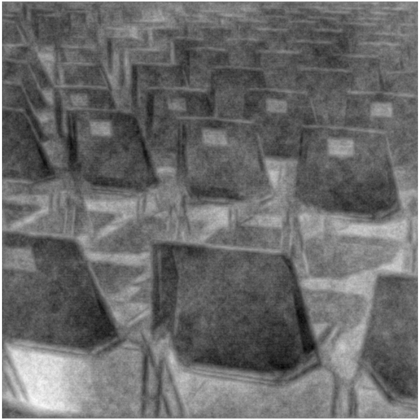

Reconstruction from Q_0 and R_{rand}

Results on example 2

Original image I^0

Reconstruction from Q_0 and R_{100}

Results on example 2

Reconstruction from Q_0 and R_{100}

Results on example 3

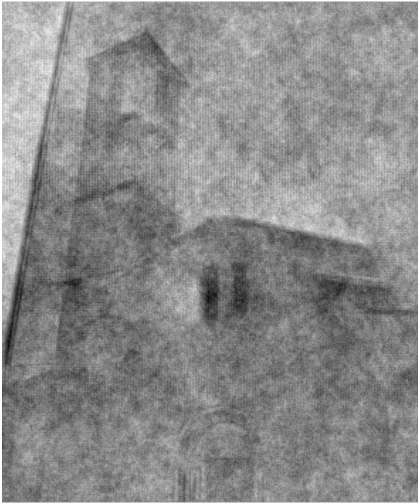

Original image I^0

Reconstruction from P_0 and R_{rand}

Results on example 3

Original image I^0

Reconstruction from P_0 and R_{100}

Results on example 3

Original image I^0

Reconstruction from Q_0 and R_{rand}

Results on example 3

Original image I^0

Reconstruction from Q_0 and R_{100}

Results on example 3

Reconstruction from Q_0 and R_{100}