Non local image processing

Jean-Michel Morel

Ecole Normale Supérieure de Cachan

Antoni Buades Capó

CNRS, Université Paris Descartes

Thanks for grants: CYCIT, ONR, CNES, DGA

THE EFROS-LEUNG IMAGE AUTOSIMILARITY

A. Efros and Th. K. Leung Texture Synthesis by Non-parametric Sampling, IEEE International Conference on Computer Vision, Corfu, Greece, September 1999

Image autosimilarity

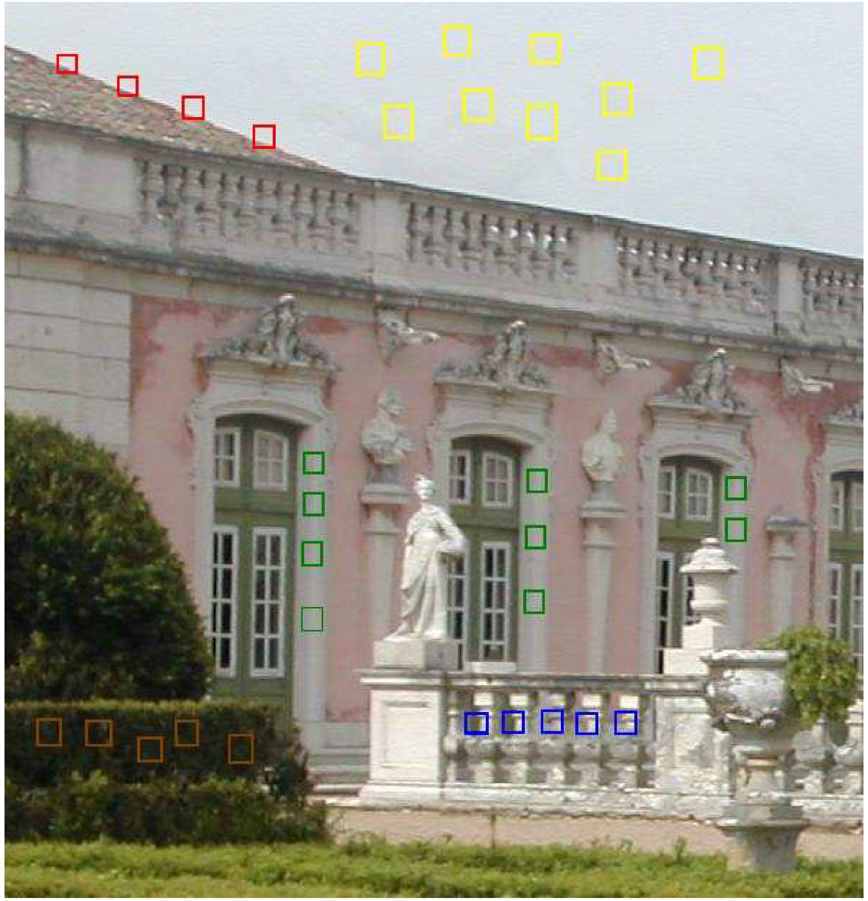

Groups of similar windows in a digital image, long range interaction. First used by Efros and Leung for texture synthesis. Fourier : too global, not geometrically adaptive. Wavelets: not adaptive enough. Main idea: the image generates its own non local model because self-similarities are non local.

• Example 1 : removing text from images by the Efros Leung algorithm ("inpainting")

• Example 2 : changing the content of aerial views

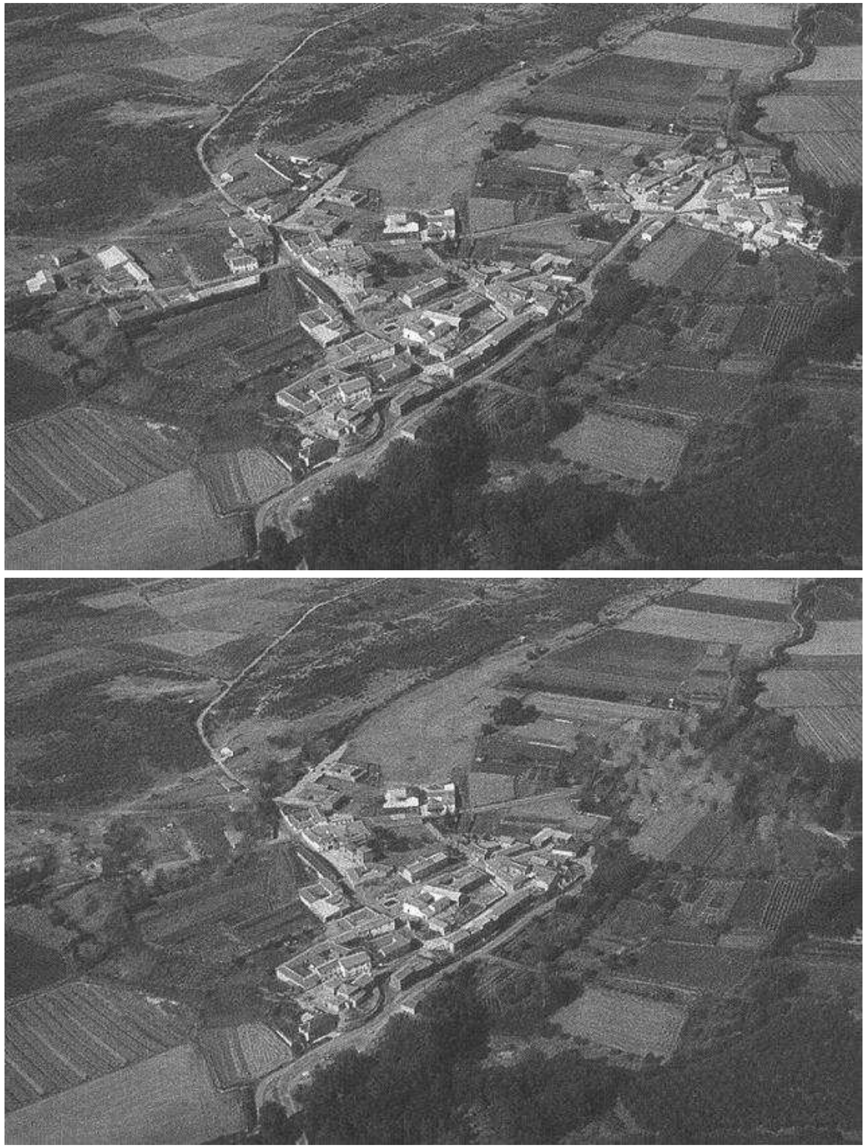

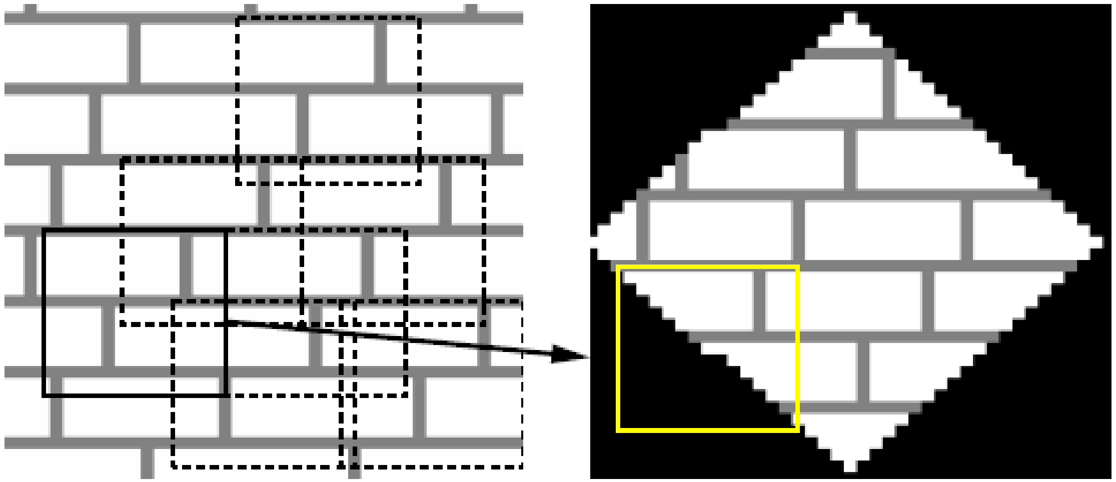

Figure 1. Algorithm Overview. Given a sample texture image (left), a new image is being synthesized one pixel at a time (right). To synthesize a pixel, the algorithm first finds all neighborhoods in the sample image (boxes on the left) that are similar to the pixel's neighborhood (box on the right) and then randomly chooses one neighborhood and takes its center to be the newly synthesized pixel.

Example 3: texture synthesis from samples

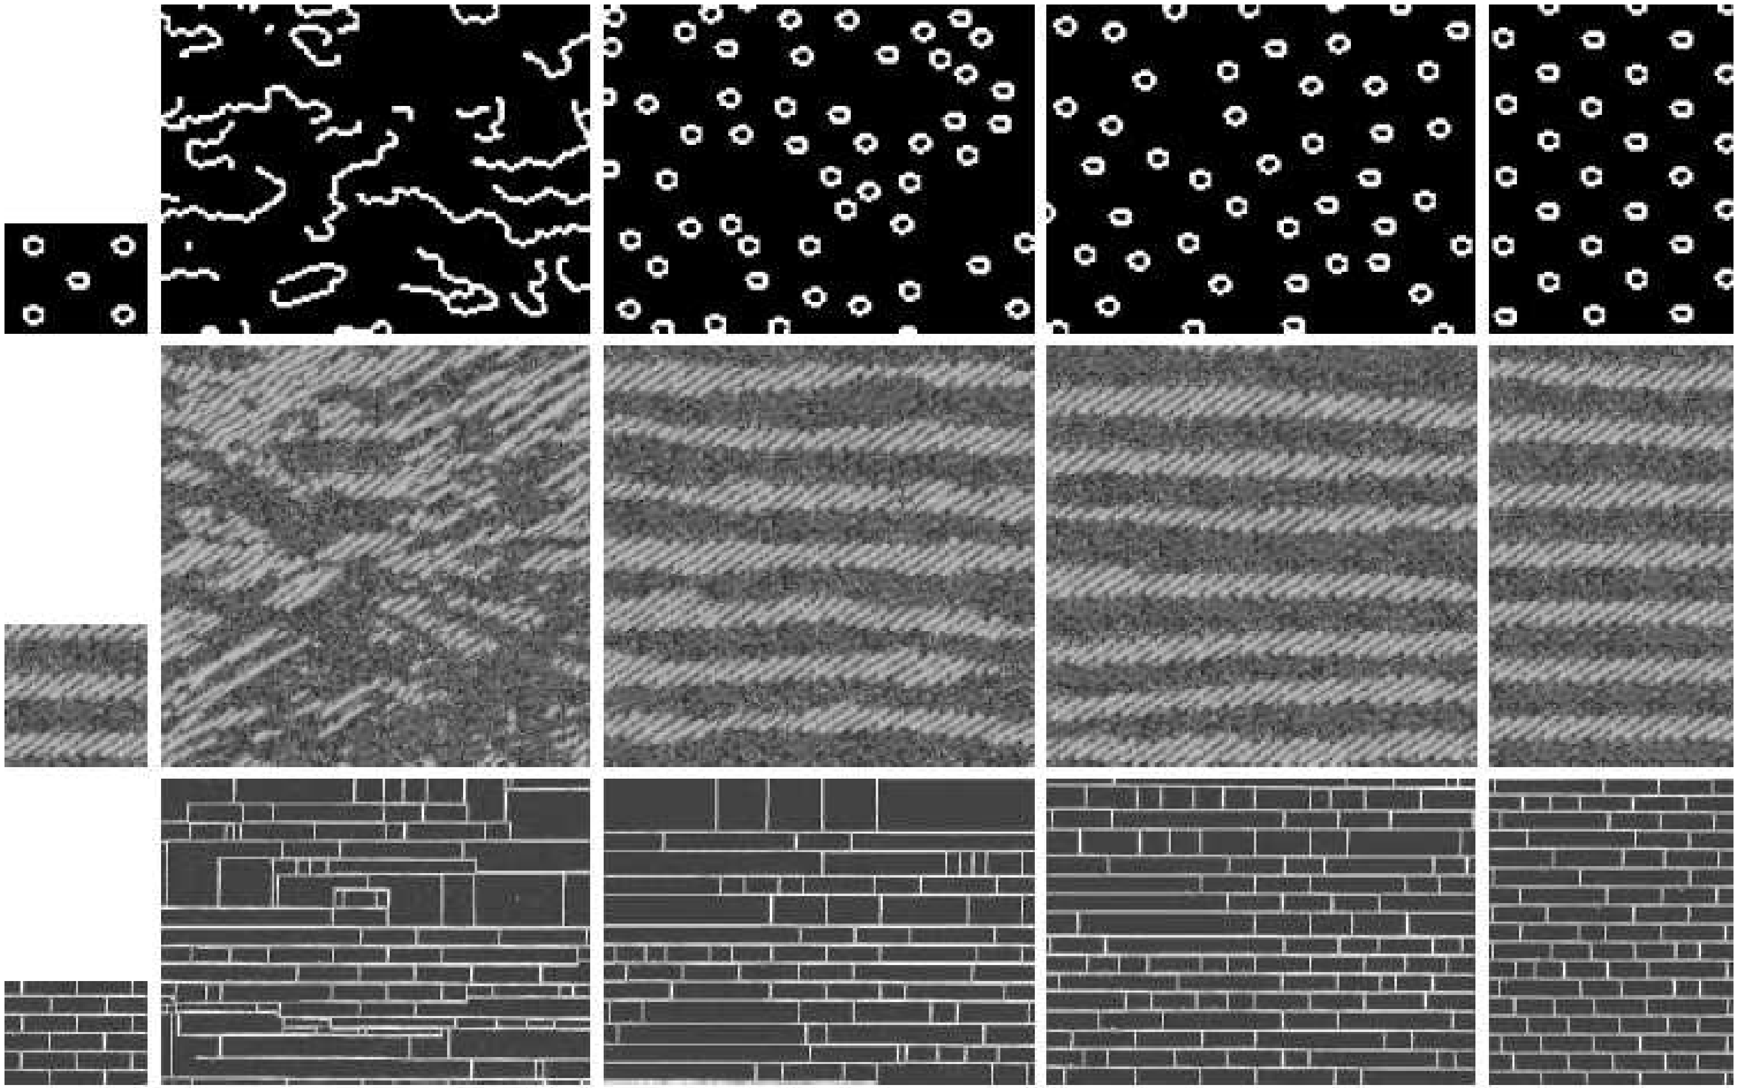

NOISE

APPLICATIONS OF NONLOCAL HEAT EQUATION

- A. Buades, B. Coll, and J.M M., A review of image denoising methods with a new one, Multiscale Modeling and Simulation, 4 (2), 2005

A noisy image...

...enhanced with DxO Noise

DxO Noise is an original, low cost implementation of the "*Non Local Means*" noise removal approach, recently introduced by Professor Jean Michel Morel's team[1]. T

The main assumption on noise

- **Main Hypothesis** In a digital image, the noise $n(i)$ at each pixel i only depends on the original pixel value $\Phi(i)$ and is additive, i.i.d. for all pixels $j \in J(i)$ with the same original value as i.
- $J(i)$ is the **neighborhood** of i. The challenge is finding $J(i)$ for every i. The simplest idea to do so is to *assume that all pixels with the same observed value $u(i)$ have the same noise model*: neighborhood filters.
- A more sophisticated use of the Hypothesis : *for a given pixel in an image, detect all pixels which have the same underlying model.*
- By the Hypothesis each j in $J(i)$ obeys a model $u(j) = v(i) + n(j)$ where $n(j)$ are i.i.d. It is then licit to perform a denoising of $u(i)$ by replacing it by

$$NFu(i) =: \frac{1}{|J(i)|} \sum_{j \in J(i)} u(j).$$

- By the variance formula for independent variables one then obtains $NFu(i) = v(i) + \tilde{n}(i)$ where

$$\mathrm{Var}(\tilde{n}(i)) = \frac{1}{|J(i)|} \mathrm{Var}(n(i)).$$

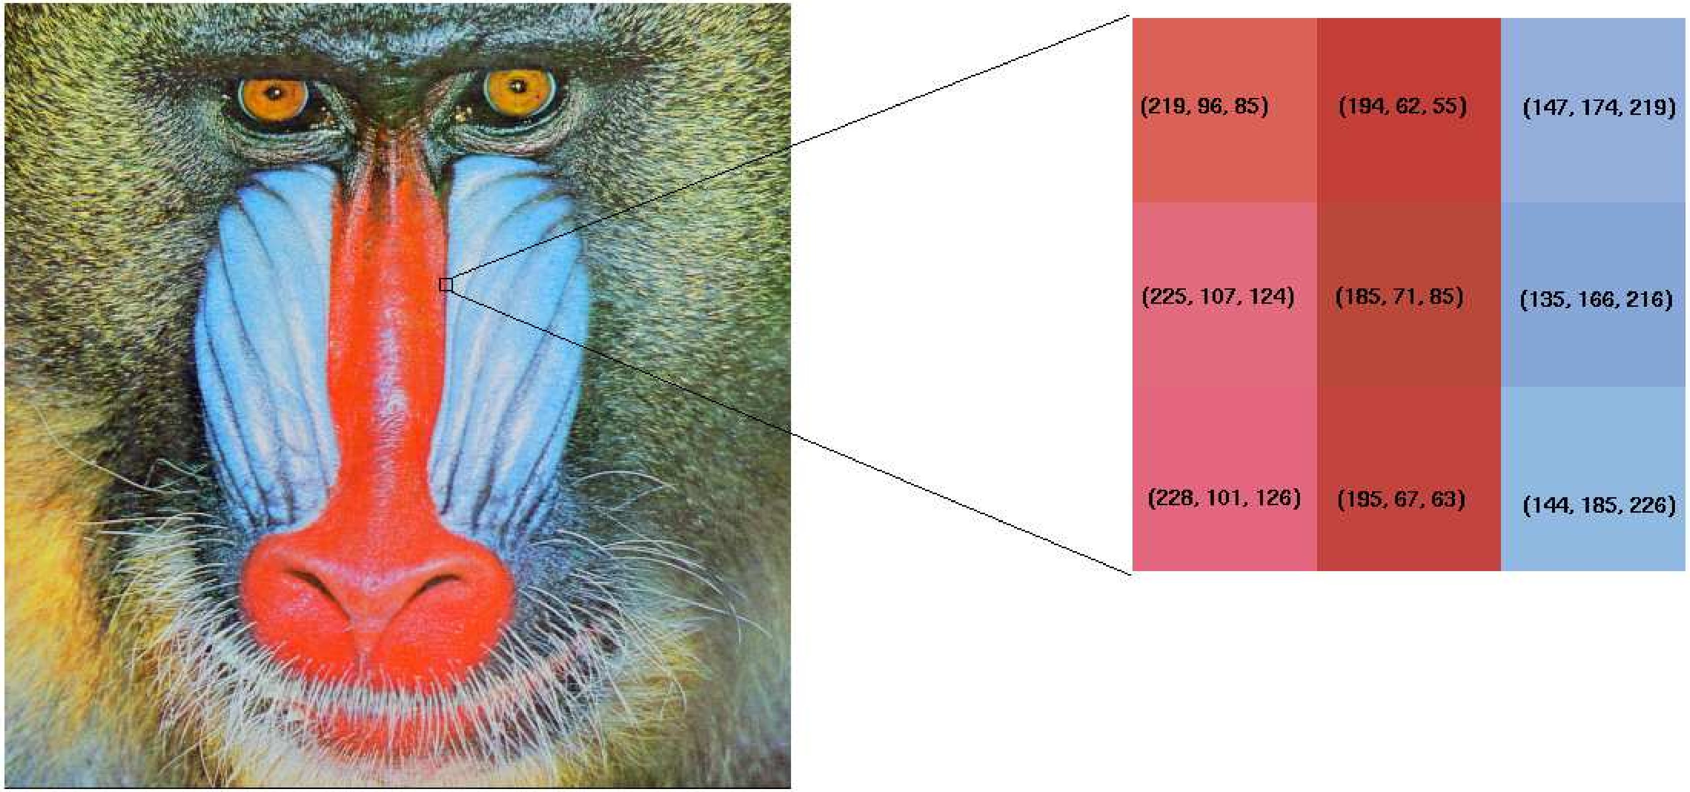

Figure 1: A. Buades, B. Coll, and J.M Morel, "Neighborhood filters and PDE's", Numerische Mathematik, 105 (1), 2006.

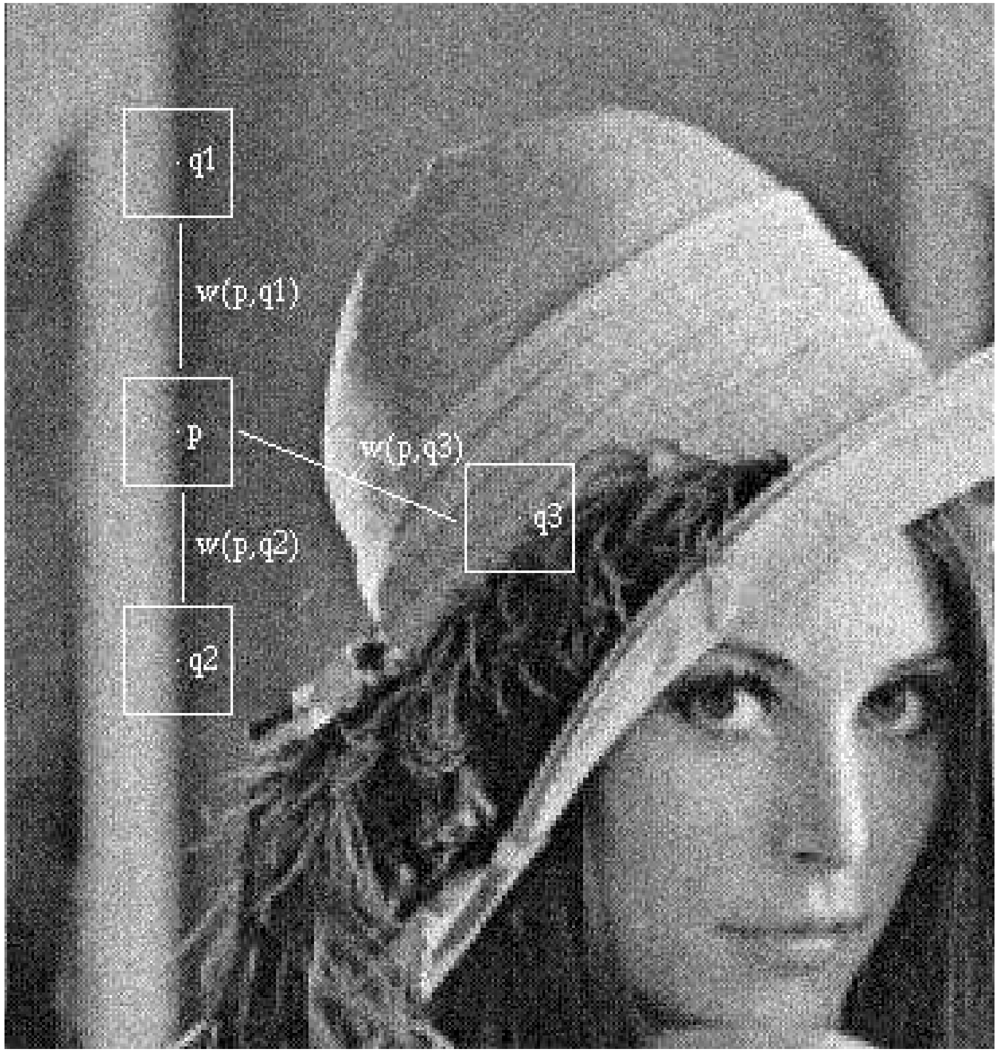
q1
w(p,q1)
p
w(p,q3)
q3
w(p,q2)
q2

FOUR NEIGHBORHOOD FILTERS FROM LOCAL TO NONLOCAL

- Gaussian mean : average of pixels in a whole Gaussian neighborhood.

$$G_\rho[u](\mathbf{x}) = \frac{1}{C(\mathbf{x})}\int_\Omega e^{-\frac{|\mathbf{x}-\mathbf{y}|^2}{\rho^2}} u(\mathbf{y})d\mathbf{y}$$

  where $C(\mathbf{x}) = C$ is the normalization parameter of the Gaussian parameter.

- Neighborhood filter. Average of pixels with a similar configuration in a whole Gaussian neighborhood.

$$NF_{\rho,h}[u](\mathbf{x}) = \frac{1}{C(\mathbf{x})}\int_\Omega e^{-\frac{|\mathbf{x}-\mathbf{y}|^2}{\rho^2}} e^{-\frac{|u(\mathbf{x})-u(\mathbf{y})|^2}{h^2}} \ u(\mathbf{y})\ d\mathbf{y};$$

- Anisotropic filter (mean curvature motion): Average of spatially close pixels in the direction of the level line

$$AF_h u(\mathbf{x}) = G_h * u_{|l(\vec{\xi})} = \int_{\mathbb{R}} G_h(t)u(\mathbf{x}+t\frac{Du^\perp}{|Du|})dt,$$

- NL-means filter. Average of pixels with a similar configuration in a whole Gaussian neighborhood.

$$NL_{h,a}[u](\mathbf{x}) = \frac{1}{C(\mathbf{x})}\int_\Omega e^{-\frac{1}{h^2}\int_{\mathbb{R}^2} G_a(t)|u(\mathbf{x}+t)-u(\mathbf{y}+t)|^2 dt} \ u(\mathbf{y})\ d\mathbf{y},$$

  where G_a is a Gaussian kernel of standard deviation a and h acts as a filtering parameter. Markovian hypothesis: Pixels with a similar neighborhood have a similar grey level value.

NONLOCAL HEAT KERNELS and EIGENVECTORS OF LAPLACIAN

•Diffusion neighborhoods

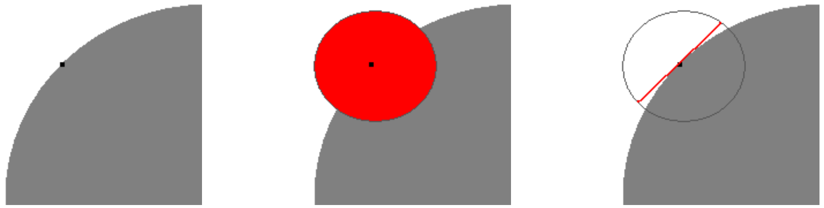

Image Gaussian Anisotropic filt.

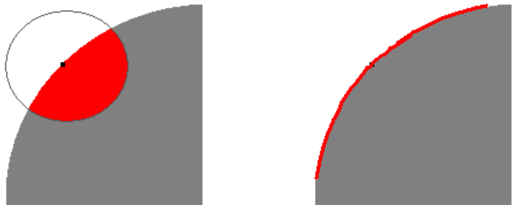

Neighborhood filter NLmeans

The NL-means: An extension of all previous methods or Gestalt Grouping

- Flat region. The large coefficients are spread out like a convolution.

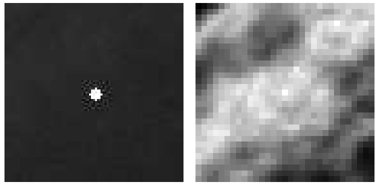

- Straight edge. The large coefficients are aligned like in a anisotropic filter.

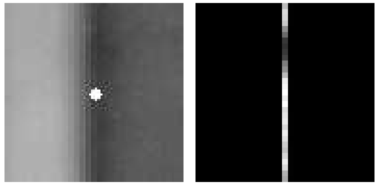

- Curved edge. The weights favor pixels belonging to the same contour.

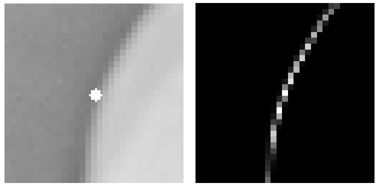

- Flat neighborhood. The average is made in the grey level neighborhood as the neighborhood filter.

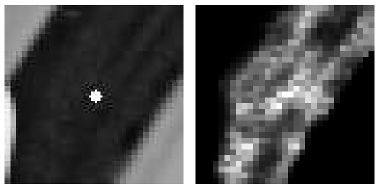

- Periodic case. The large coefficients are distributed across the texture (non local).

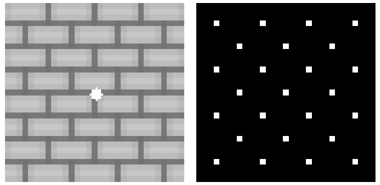

- Repetitive structures. The weights favor similar configurations even they are far away (non local).

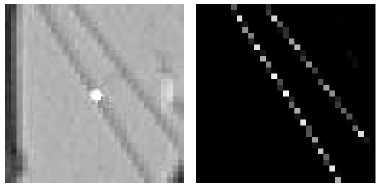

6 G. PEYRÉ

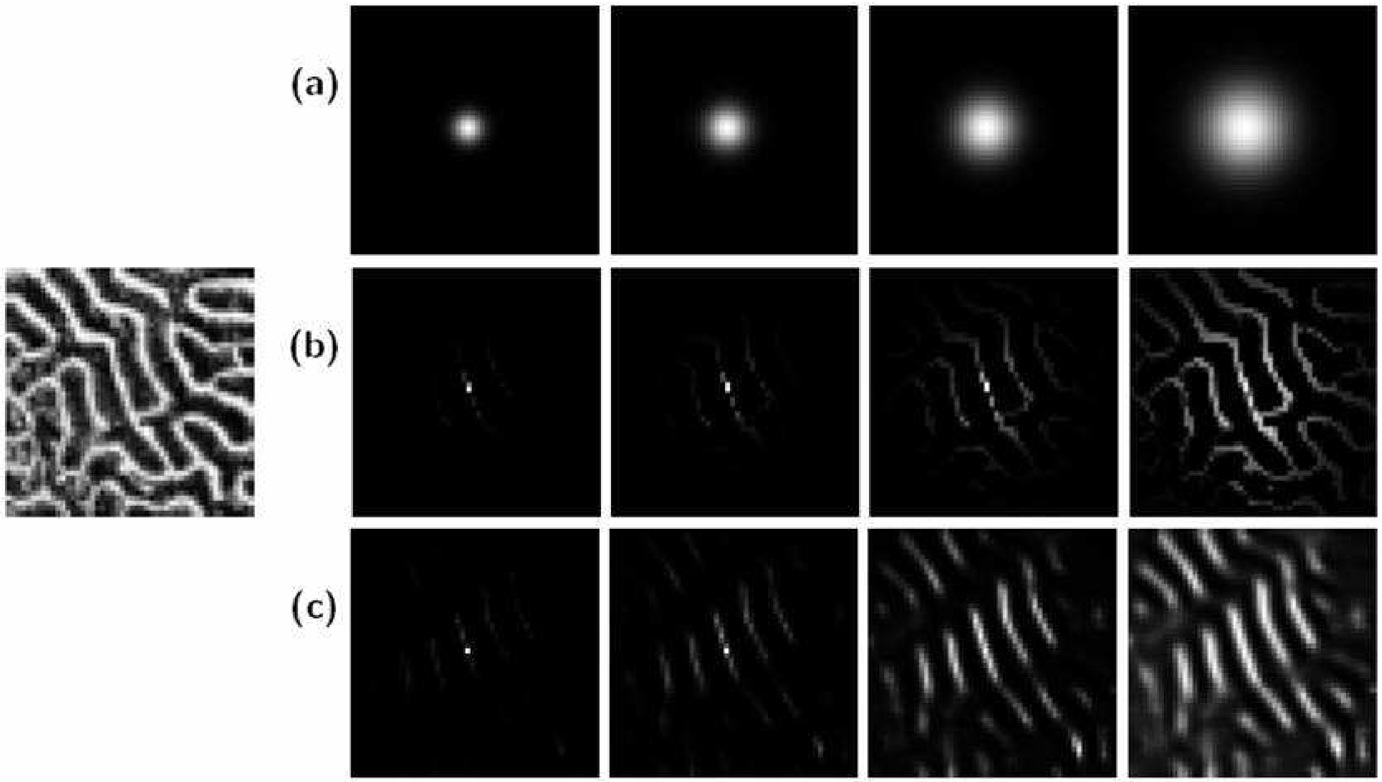

FIGURE 2.1. *Left: original image f. Right: heat diffusions with an increasing time for: (a) local embedding $x \mapsto x$, (b) semi-local embedding $x \mapsto (x, \lambda f(x))$, (c) non-local embedding $x \mapsto p_x(f)$.*

Figure 2: G. Peyré, Manifold models for signals and images, preprint, (2007).

Comparison: Visual quality: Input, Gauss, Neighborhood classic, Non local means

- Comparison on a periodic image

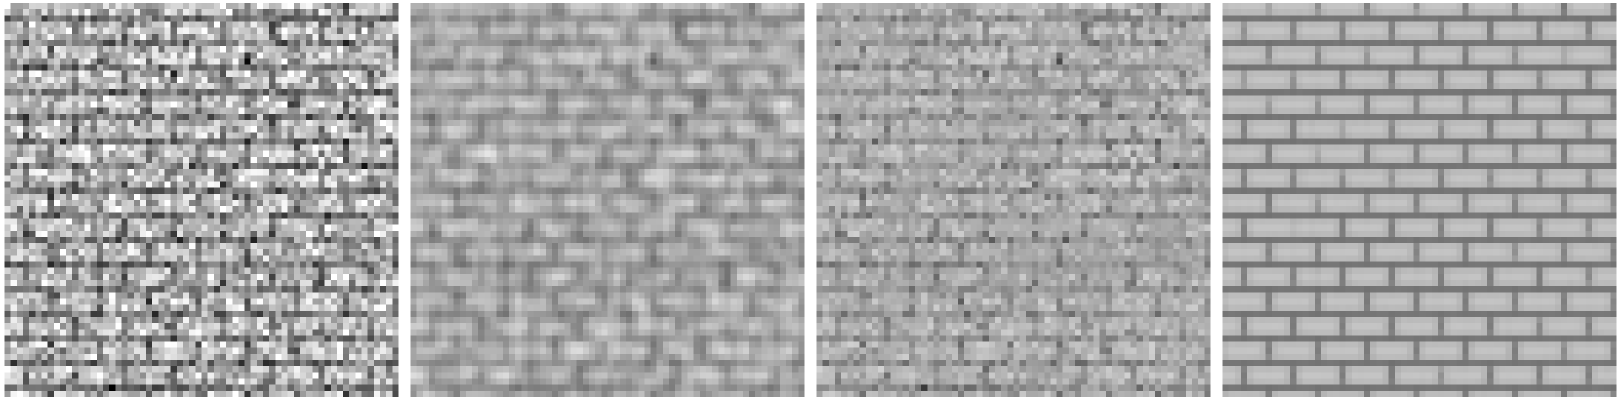

Visual Comparison

- Restored images and removed noise by Gaussian convolution, sigma filter and NL-means.

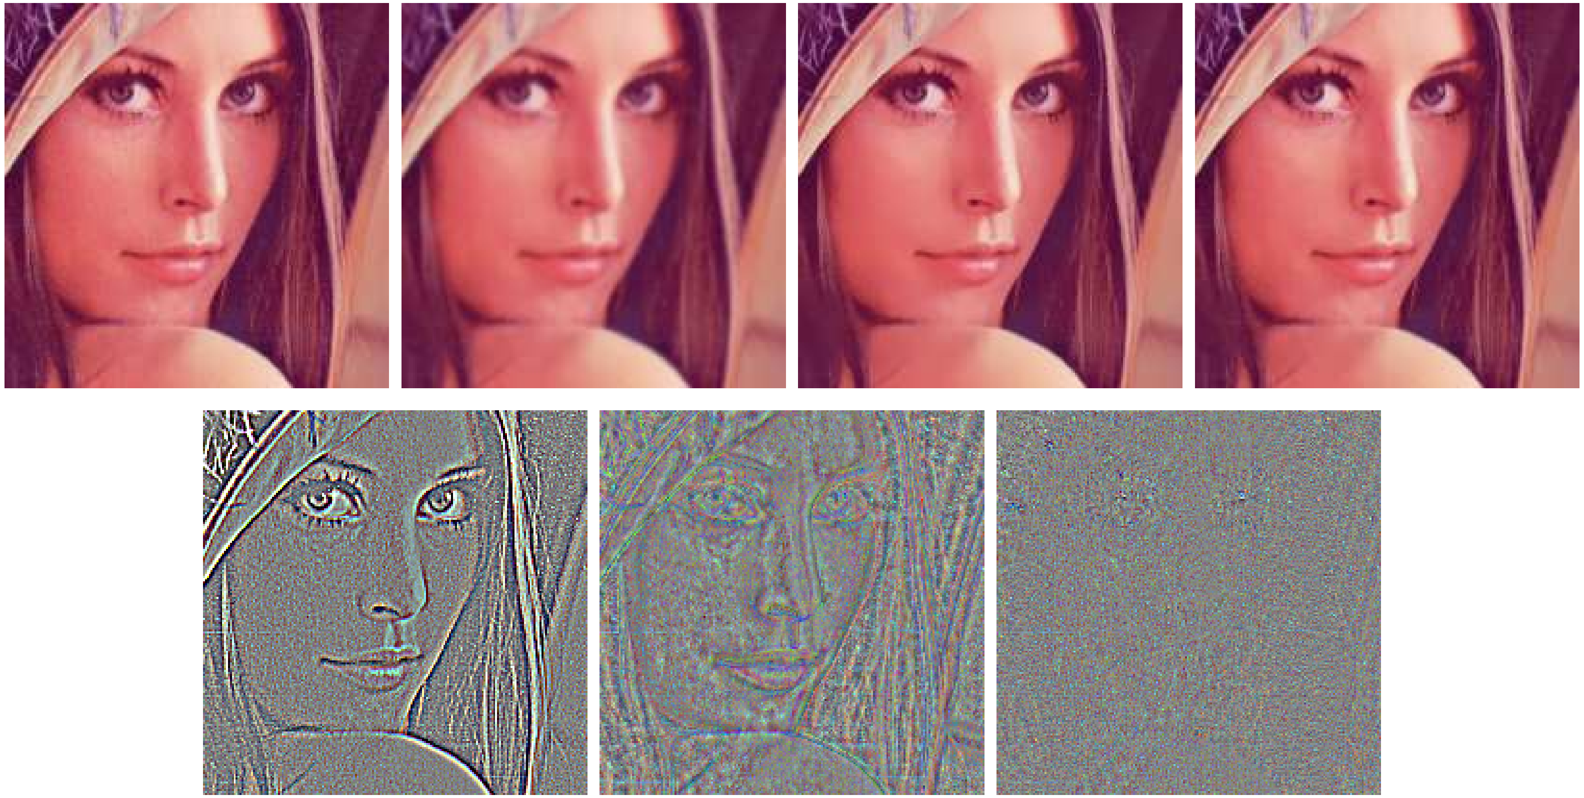

- Restored images and removed noise by the anisotropic filter, the sigma filter and the NL-means.

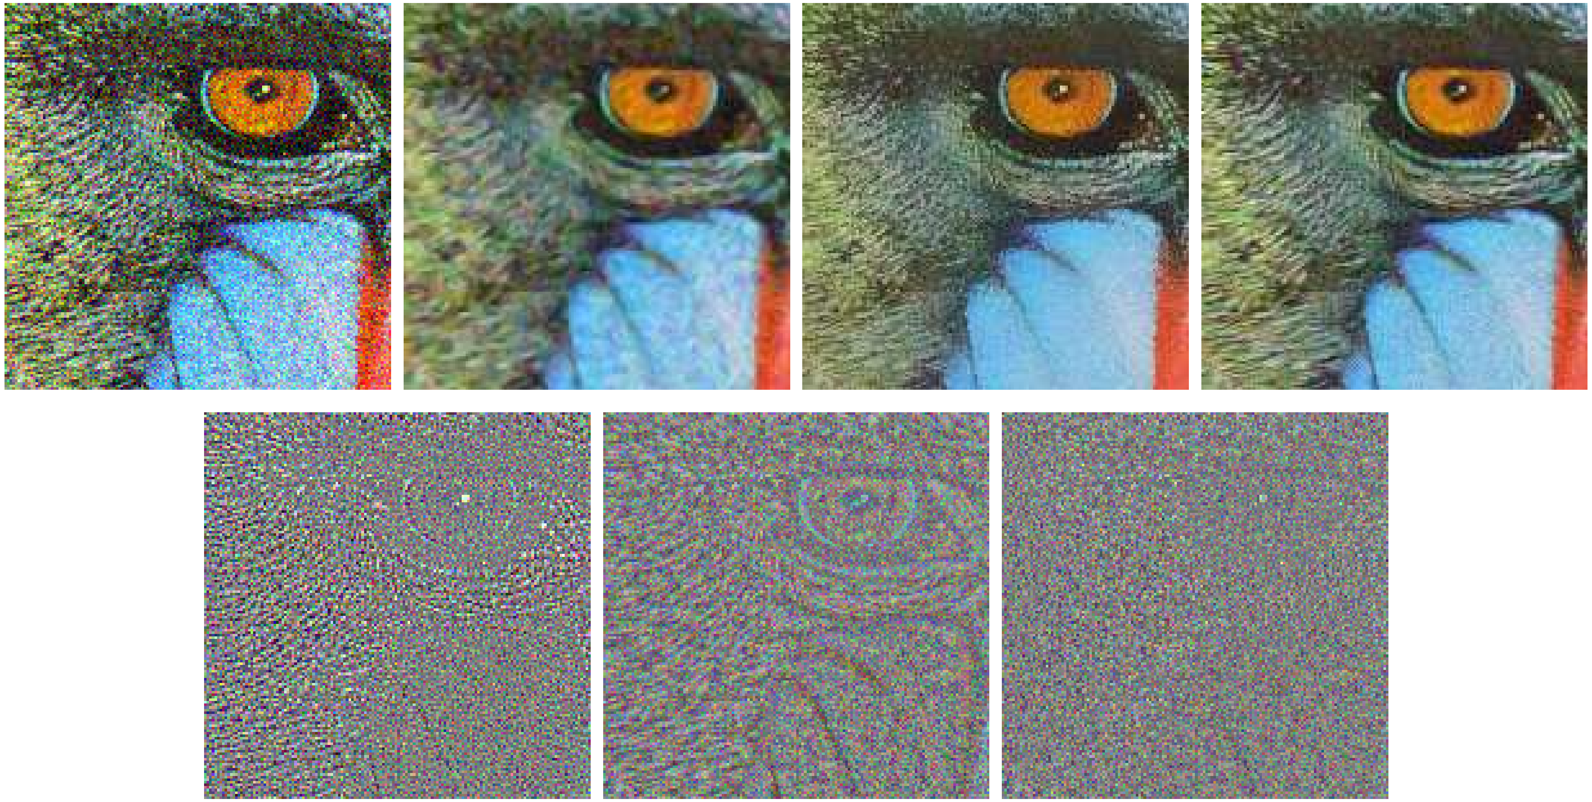

TWO MAIN DENOISING PRINCIPLES

Denoising

We define a denoising method D_h as a decomposition

$$v = D_h v + n(D_h, v), \tag{1}$$

where h is a filtering parameter which usually depends on the standard deviation of the noise σ.

- **Preservation of original information.** Features in $n(D_h, v) = v - D_h v$ are removed from v. We call this difference *method noise* when v is non or slightly noisy.

  Principle 1 *For every denoising algorithm, the method noise must be zero if the image contains no noise and should be in general an image of independent zero-mean random variables.*

- **No artifacts** The transformation of a white noise into any correlated signal creates structure and artifacts.

  Principle 2 *A denoising algorithm must transform a white noise image into a white noise image (with lower variance).*

Other classic algorithms to be compared

- Minimization of the total variation

$$TVF_\lambda(v) = \arg\min_u \int_\Omega |Du| + \lambda \int |v-u|^2 \tag{1}$$

- Wavelet thresholding

$$HWT = \sum_{\{(j,k) | |\langle v, \psi_{j,k}\rangle| > \tau\}} \langle v, \psi_{j,k}\rangle \psi_{j,k}$$

where $\mathcal{B} = \{\psi_{j,k}\}_{(j,k)}$ is a wavelet basis and τ the threshold.

Preservation of original information. We recall that we defined the method noise as the image difference $n(D_h, v) = v - D_h v$ when v is non or slightly noisy.

> **Principle 1** *For every denoising algorithm, the method noise must be zero if the image contains no noise and should be in general an image of independent zero-mean random variables.*

Theorem 1 *The convolution with a gaussian kernel G_h is such that*

$$u - G_h * u = -h^2 \Delta u + o(h^2),$$

for h small enough.

Theorem 2 *The image method noise of an anisotropic filter AF_h is*

$$u(\boldsymbol{x}) - AF_h u(\boldsymbol{x}) = -\frac{1}{2} h^2 |Du| curv(u)(\boldsymbol{x}) + o(h^2),$$

Theorem 3 *The method noise of the Total Variation minimization is*

$$u(\boldsymbol{x}) - TVF_\lambda(u)(\boldsymbol{x}) = -\frac{1}{2\lambda} curv(TVF_\lambda(u))(\boldsymbol{x}).$$

Theorem 4 *The method noise of a hard thresholding $HWT_\mu(u)$ is*

$$u - HWT_\mu(u) = \sum_{\{(j,k) || \langle u, \psi_{j,k} \rangle | < \tau\}} \langle u, \psi_{j,k} \rangle \psi_{j,k}$$

Method Noise Method noise of six denoising methods

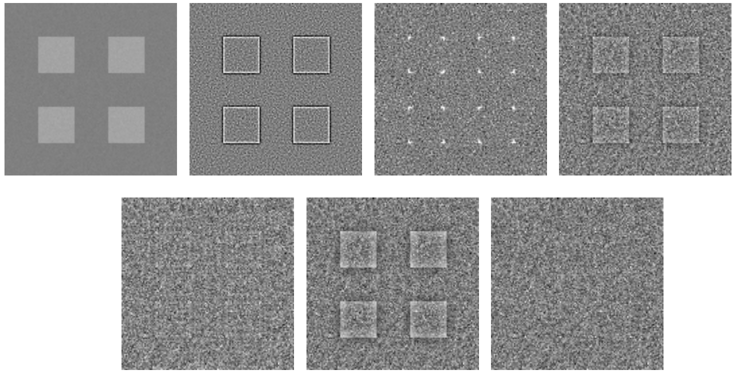

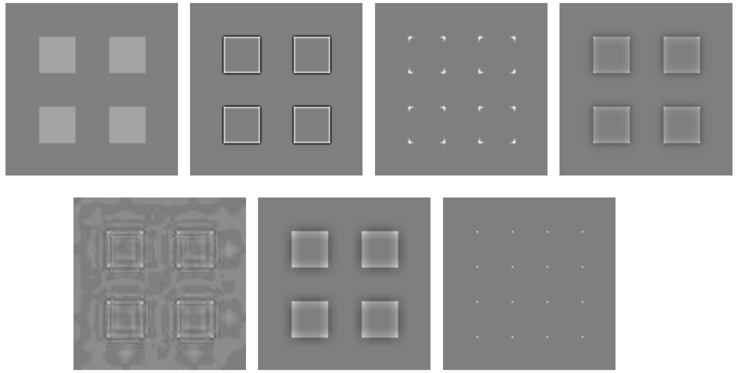

Gaussian, Anisotropic filtering, TV, stationary wavelet, neighborhood filter, NL-means.
In the noisy case, parameters are fixed in order to remove exactly an energy σ^2 ($\sigma = 2.5$). The same parameters have been used in the second experiment on the real image.

No artifacts The transformation of a white noise into any correlated signal creates structure and artifacts.

> **Principle 2** *A denoising algorithm must transform a white noise image into a white noise image (with lower variance).*

Application of the denoising algorithms to a noise sample and its Fourier transforms. Gauss, TV, wavelets, neighborhood filter, NL-means

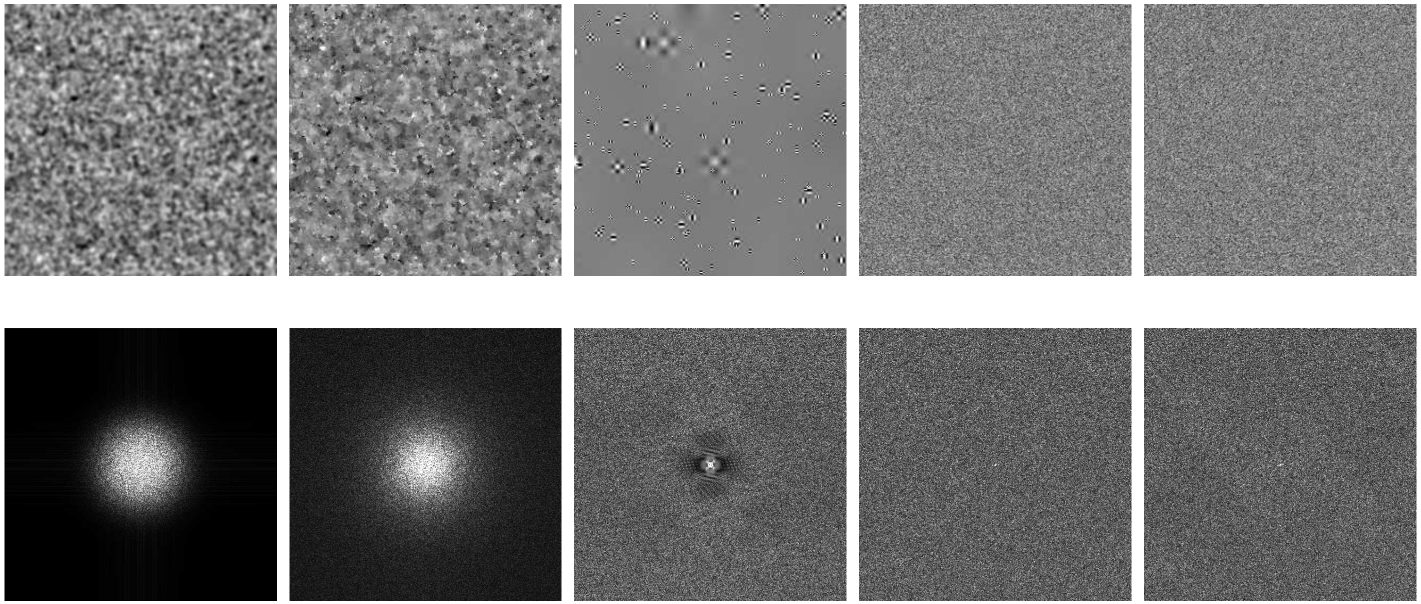

Films and NL-means

- The NL-means algorithm does not need to calculate the trajectories. It simply looks for the resembling pixels, no matter where they lie in the movie.

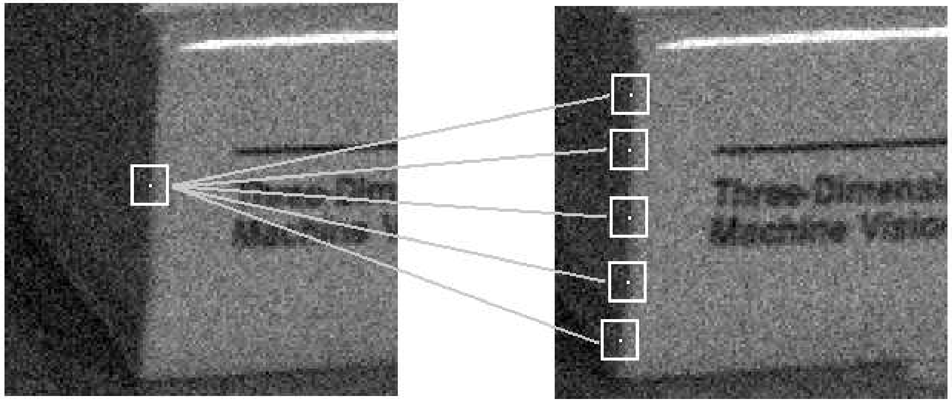

Why do not average all?

- Straightforward extension as a spatiotemporal filter.

Probability distributions in movement

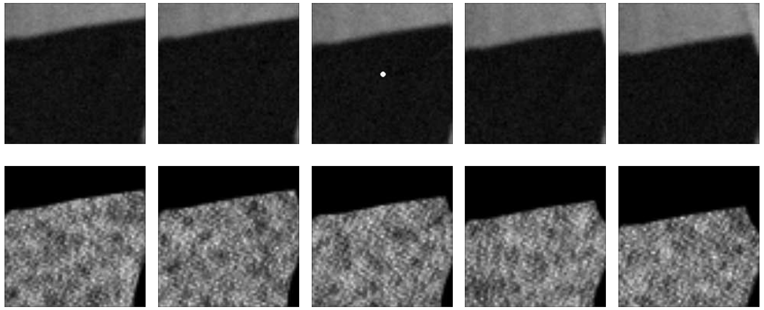

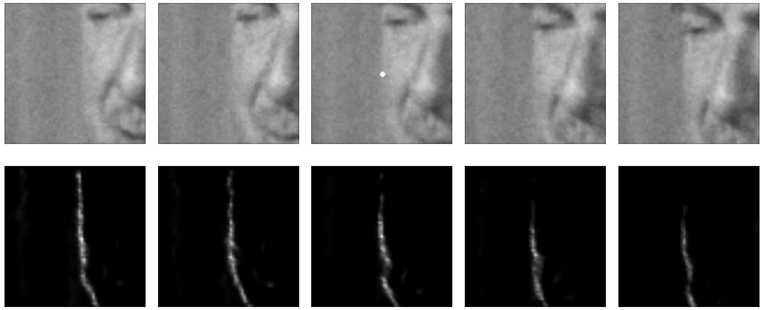

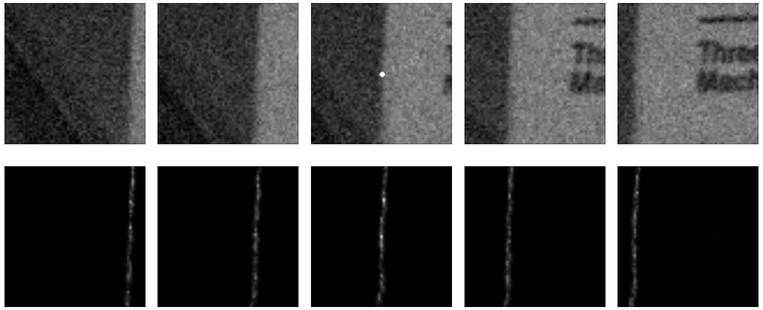

The algorithm looks for the pixels with a more similar configuration even they have moved. This algorithm is adapted to moving pictures without the need of an explicit motion estimation.

Comparison

- Three consecutive frames of a noisy image sequence. The noisy sequence has been obtained by the addition of a Gaussian additive white noise ($\sigma = 15$) to the original sequence.

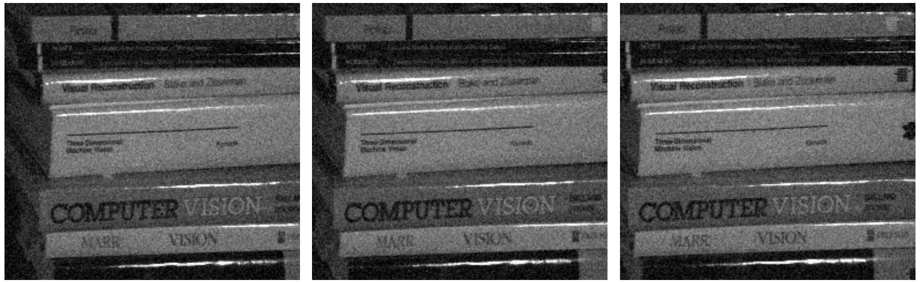

- Comparison experiment between the motion compensated neighborhood filter and the NL-means. The motion estimation has been obtained by the block matching algorithm.

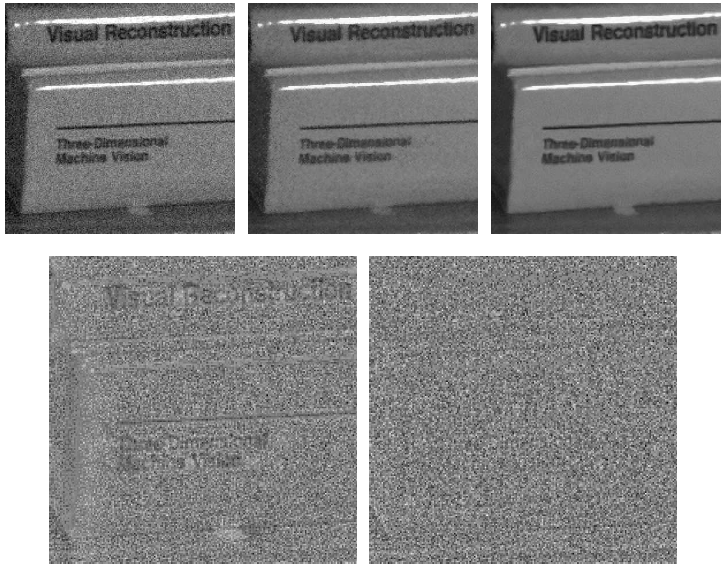

COLORIZATION AND SEGMENTATION BY DIFFUSION

Colorization and segmentation

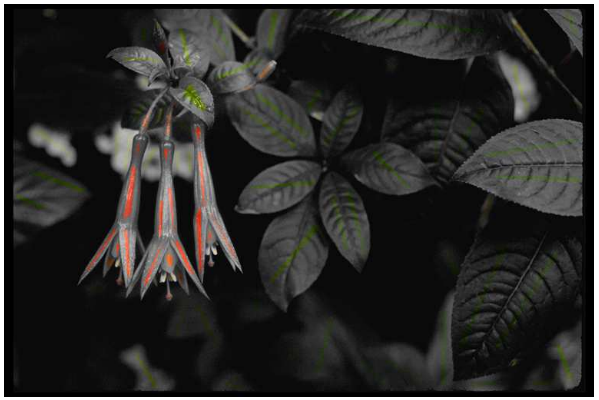

which is decomposed as in Y, U, V as (seeds in color).

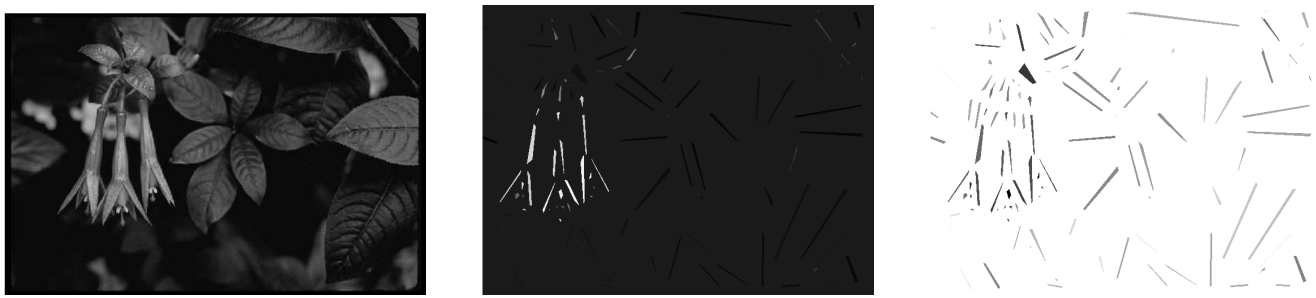

- Non local heat equation with neighborhood fixed by the grey level image and color seeds Dirichlet condition.

- Replace each pixel (x, y) with no initial color by the weighted mean

$$\hat{u}(x, y) = \frac{1}{|B|} \sum_{(v,w) \in B} u(v, w).$$

where B is the set of pixels with grey level similar to (x, y).

- Iterate

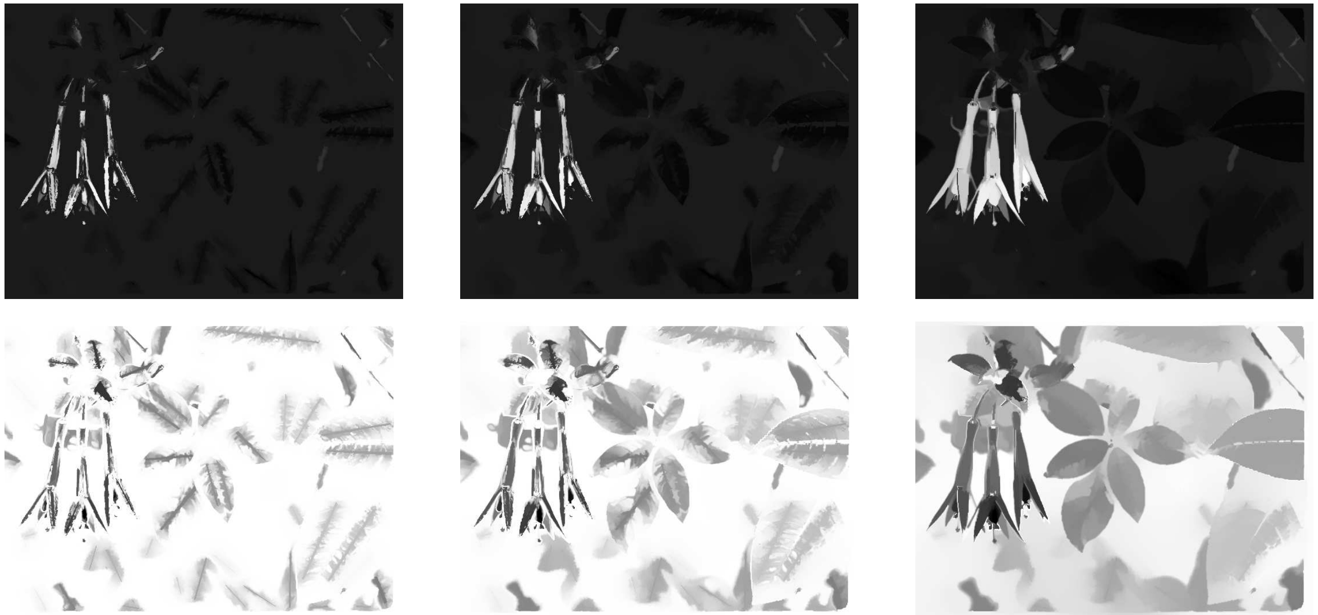

The single point comparison of the neighborhood filter can lead to the mixing of objects color. Middle: neighborhhod filter (Grady), right : NL-means.