Gaussian Models for Image Patches

Agnès Desolneux

CNRS and Ecole Normale Supérieure Paris-Saclay

Cours Master MVA

Modeling images as samples of Gaussian random fields

Let $u : \Omega \to \mathbb{R}$ be a grey-level image. It is seen as a sample of a multivariate Gaussian distribution $\mathcal{N}(\mu, \Sigma)$ of density

$$p(u) = \frac{\det(\Sigma)^{-\frac{1}{2}}}{(\sqrt{2\pi})^{|\Omega|}} e^{-(u-\mu)^t \Sigma^{-1} (u-\mu)/2},$$

where u is seen as the vector $\{u_j \, ; \, j \in \Omega\}$.

Here $\mu : \Omega \to \mathbb{R}$ is the mean (size $|\Omega|$) and Σ is the covariance matrix, it is a $|\Omega| \times |\Omega|$ symmetric positive matrix.

PREVIOUSLY

In the previous lectures : we have considered **stationary periodic Gaussian** models, implying that

the vector (image) μ is constant

and $\quad \forall i,j \in \Omega, \quad \Sigma(i,j) = \Sigma(0, j-i).$

Number of **parameters** (degrees of freedom) $\simeq 1 + |\Omega|/2$.

Number of **samples** : one image, and all its periodic translations.

We will consider **general Gaussian models** :

Number of **parameters** (degrees of freedom) $= |\Omega| + |\Omega|(|\Omega| + 1)/2$

Ex : When $|\Omega| = 256 \times 256$, nb of parameters to estimate is of the order of $256^4/2$, which is quite huge !

Number of **samples** : just one image

BUT one can consider image patches (*imagettes*) and use the redundancy of natural images.

PATCH-BASED RESTORATION

Patch-based denoising Buades et al. ('05), Awate Whitaker ('06), Dabov et al. ('07), Lebrun et al. ('12), etc.

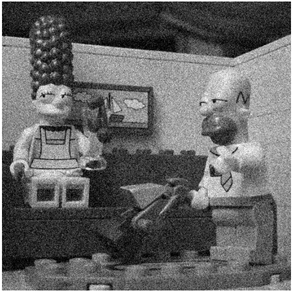

PATCH-BASED RESTORATION

Patch-based denoising Buades et al. ('05), Awate Whitaker ('06), Dabov et al. ('07), Lebrun et al. ('12), etc.

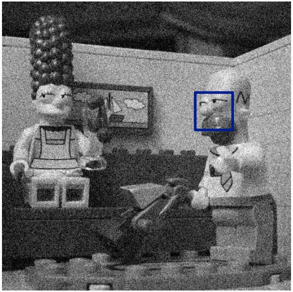

PATCH-BASED RESTORATION

Patch-based denoising Buades et al. ('05), Awate Whitaker ('06), Dabov et al. ('07), Lebrun et al. ('12), etc.

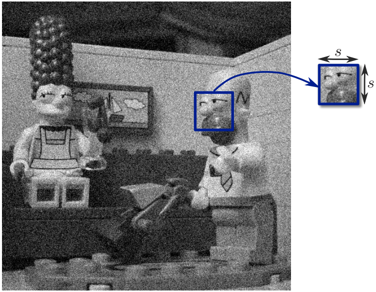

Patch-based restoration

Patch-based denoising Buades et al. ('05), Awate Whitaker ('06), Dabov et al. ('07), Lebrun et al. ('12), etc.

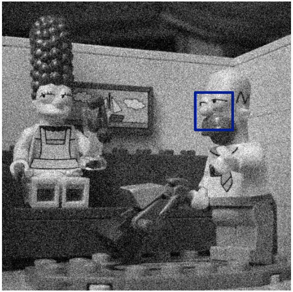

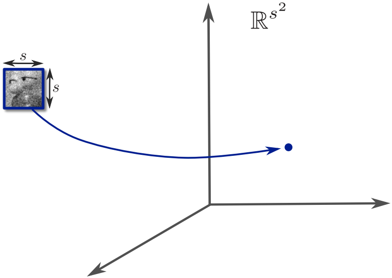

PATCH-BASED RESTORATION

Patch-based denoising Buades et al. ('05), Awate Whitaker ('06), Dabov et al. ('07), Lebrun et al. ('12), etc.

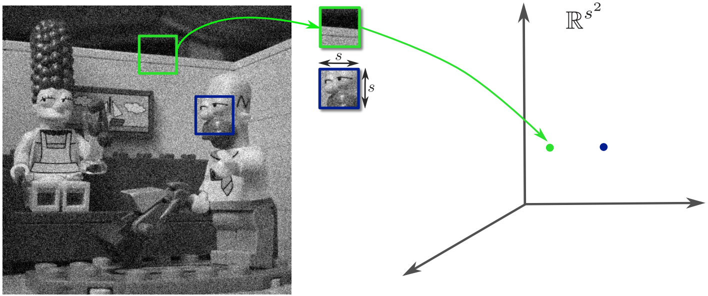

Patch-based restoration

Patch-based denoising Buades et al. ('05), Awate Whitaker ('06), Dabov et al. ('07), Lebrun et al. ('12), etc.

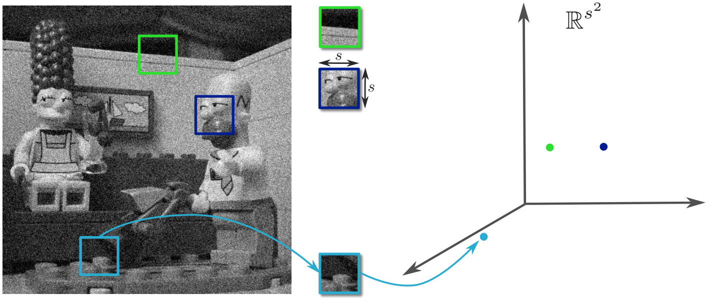

PATCH-BASED RESTORATION

Patch-based denoising Buades et al. ('05), Awate Whitaker ('06), Dabov et al. ('07), Lebrun et al. ('12), etc.

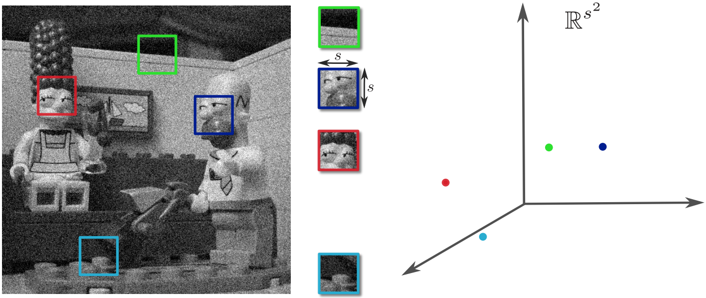

PATCH-BASED RESTORATION

Patch-based denoising Buades et al. ('05), Awate Whitaker ('06), Dabov et al. ('07), Lebrun et al. ('12), etc.

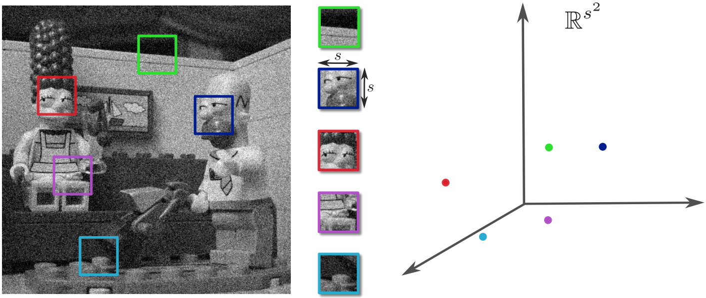

PATCH-BASED RESTORATION

Patch-based denoising Buades et al. ('05), Awate Whitaker ('06), Dabov et al. ('07), Lebrun et al. ('12), etc.

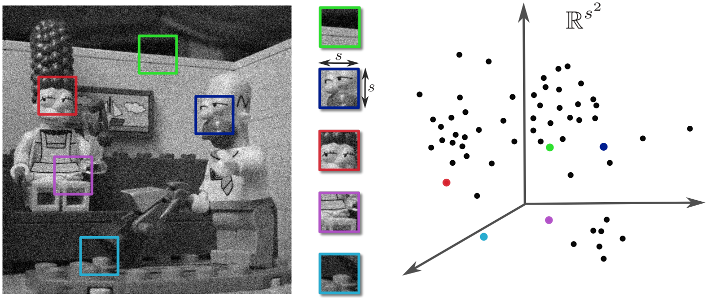

Image edition and synthesis

- texture synthesis Efros-Leung ('99)
- image retargeting or reshuffling Barnes et al. ('09)
- style transfer Frigo et al. ('16)

Image restoration

- Gaussian denoising Buades et al. ('05), Awate Whitaker ('06), Dabov et al. ('08), Lebrun et al. ('12),
- non Gaussian denoising : Poisson, Speckle Deledalle et al. ('10), ('12), impulse noise Delon Desolneux ('13)
- inpainting Wexler et al. ('04), Criminisi Perez ('04), Newson et al. ('14)
- interpolation Yu et al. ('12), demosaicing Buades et al. ('07)
- high dynamic range images (HDR) Aguerrebere et al. ('17), decompression

Image edition and synthesis

- texture synthesis Efros-Leung ('99)
- image retargeting or reshuffling Barnes et al. ('09)
- style transfer Frigo et al. ('16)

Image restoration

- Gaussian denoising Buades et al. ('05), Awate Whitaker ('06), Dabov et al. ('08), Lebrun et al. ('12),
- non Gaussian denoising : Poisson, Speckle Deledalle et al. ('10), ('12), impulse noise Delon Desolneux ('13)
- inpainting Wexler et al. ('04), Criminisi Perez ('04), Newson et al. ('14)
- interpolation Yu et al. ('12), demosaicing Buades et al. ('07)
- high dynamic range images (HDR) Aguerrebere et al. ('17), decompression

IMAGE DENOISING : NLMEANS (2005)

We observe

$$\tilde{u} = u + n$$

with n_j i.i.d. and $\sim \mathcal{N}(0, \sigma^2)$. The aim is to recover u.

IMAGE DENOISING : NLMEANS (2005)

We observe

$$\tilde{u} = u + n$$

with n_j i.i.d. and $\sim \mathcal{N}(0, \sigma^2)$. The aim is to recover u.

Non-local Means, Buades, Coll, Morel, ('05)

$$\forall i \in \Omega, \quad NLu_i = \frac{\sum_{j \in \Omega} w_{i,j} \tilde{u}_j}{\sum_j w_{i,j}}.$$

where the $w_{i,j}$ are weights that measure the similarity between patches centered at i and j, typically

$$w_{i,j} = e^{-\|x_i^{\tilde{u}} - x_j^{\tilde{u}}\|_2^2 / 2h^2}.$$

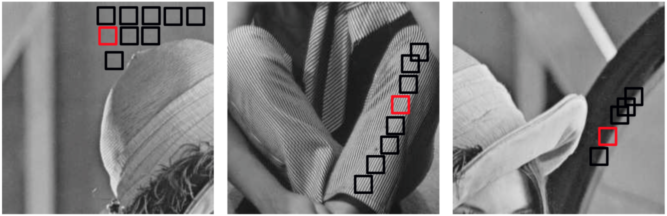

Région uniforme | Région texturée | Contour géométrique

IMAGE DENOISING : NLMEANS (2005)

Non-local means, Buades, Coll, Morel, ('05)

Noisy image, $\sigma = 20$

NL-means

Style Transfer with patches (2016)

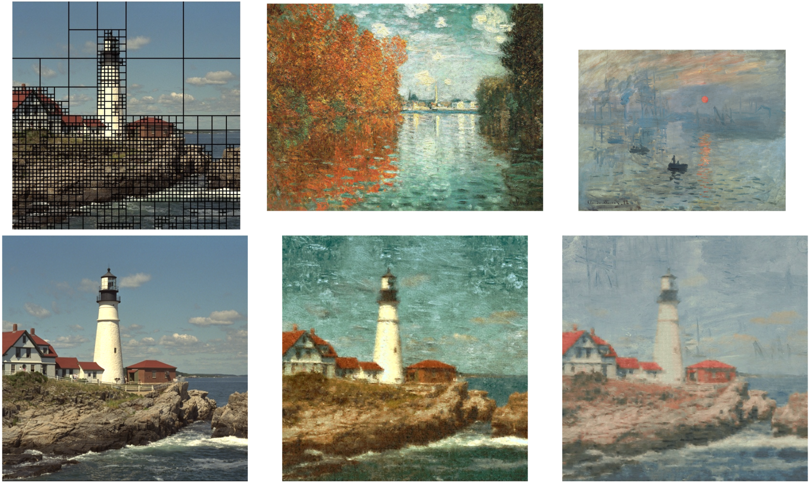

technicolor

Frigo et al. ('16) Markov random field, optimization

Inverse Problems

Model : we observe a corrupted image $\widetilde{u}$

$$\underset{\text{observation}}{\widetilde{u}} = \underset{\text{operator}}{A} \quad \underset{\text{unknown}}{u} + \underset{\text{noise}}{n}$$

Aim : recover u from $\tilde{u}$

Inverse Problems

Model : we observe a corrupted image $\widetilde{u}$

$$\underset{\text{observation}}{\widetilde{u}} = \underset{\text{operator}}{A} \; \underset{\text{unknown}}{u} + \underset{\text{noise}}{n}$$

Aim : recover u from $\tilde{u}$

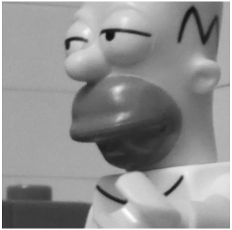

image u

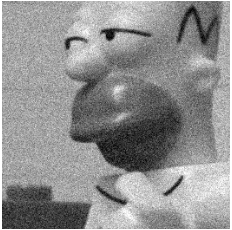

noisy

blurry

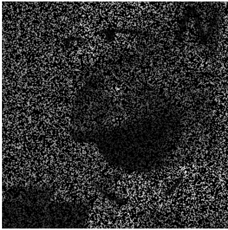

missing pixels

Inverse Problems

Model : for all patches y_i in the image $\tilde{u}$

$$\underset{\text{observation}}{y_i} = \underset{\text{operator}}{A_i} \; \underset{\text{unknown}}{x_i} + \underset{\text{noise}}{n_i}$$

Aim : estimate the « clean » patches $x_i \in \mathbb{R}^p$ from the observed $\{y_i\}_i$

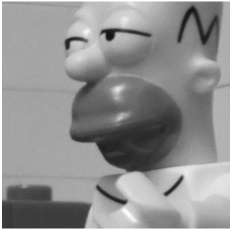

image u

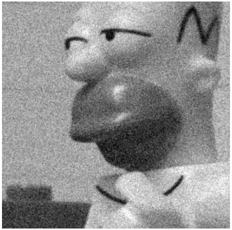

noisy

blurry

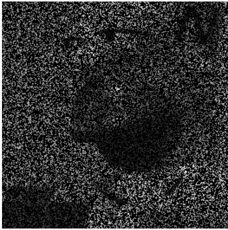

missing pixels

RESTORATION STRATEGIES

Under the hypothesis of an additive Gaussian noise and a prior distribution $p(x)$ on the patches, the **posterior distribution** is

$$p(x|y) \propto p(y|x)p(x) \propto e^{-\frac{\|Ax-y\|^2}{2\sigma^2}} p(x).$$

Restoration strategies

Under the hypothesis of an additive Gaussian noise and a prior distribution $p(x)$ on the patches, the **posterior distribution** is

$$p(x|y) \propto p(y|x)p(x) \propto e^{-\frac{\|Ax-y\|^2}{2\sigma^2}} p(x).$$

Some restoration strategies

- $\widehat{x} = \mathbb{E}[X|Y=y]$ the **minimum mean square error (MMSE)** estimator

Restoration strategies

Under the hypothesis of an additive Gaussian noise and a prior distribution $p(x)$ on the patches, the **posterior distribution** is

$$p(x|y) \propto p(y|x)p(x) \propto e^{-\frac{\|Ax-y\|^2}{2\sigma^2}} p(x).$$

Some restoration strategies

- $\widehat{x} = \mathbb{E}[X|Y=y]$ the **minimum mean square error (MMSE)** estimator
- $\widehat{x} = Dy + \alpha$ such that D and α minimize $\mathbb{E}(\|DY + \alpha - X\|^2)$ (linear MMSE or **Wiener estimator**)

RESTORATION STRATEGIES

Under the hypothesis of an additive Gaussian noise and a prior distribution $p(x)$ on the patches, the **posterior distribution** is

$$p(x|y) \propto p(y|x)p(x) \propto e^{-\frac{\|Ax-y\|^2}{2\sigma^2}} p(x).$$

Some restoration strategies

- $\widehat{x} = \mathbb{E}[X|Y=y]$ the **minimum mean square error (MMSE)** estimator
- $\widehat{x} = Dy + \alpha$ such that D and α minimize $\mathbb{E}(\|DY + \alpha - X\|^2)$ (linear MMSE or **Wiener estimator**)
- **Maximum *a posteriori* (MAP)**

$$\widehat{x} = \arg\max_{x \in \mathbb{R}^p} p(x|y) = \arg\min_{x} \underbrace{E_{data}(x,y)}_{\text{fidelity term}} + \underbrace{E_{smooth}(x)}_{\text{regularity term}}$$

Restoration strategies

Under the hypothesis of an additive Gaussian noise and a prior distribution $p(x)$ on the patches, the **posterior distribution** is

$$p(x|y) \propto p(y|x)p(x) \propto e^{-\frac{\|Ax-y\|^2}{2\sigma^2}} p(x).$$

Some restoration strategies

- $\widehat{x} = \mathbb{E}[X|Y = y]$ the **minimum mean square error (MMSE)** estimator
- $\widehat{x} = Dy + \alpha$ such that D and α minimize $\mathbb{E}(\|DY + \alpha - X\|^2)$ (linear MMSE or **Wiener estimator**)
- **Maximum *a posteriori* (MAP)**

$$\widehat{x} = \arg\max_{x \in \mathbb{R}^p} p(x|y) = \arg\min_{x} \underbrace{E_{data}(x, y)}_{\text{fidelity term}} + \underbrace{E_{smooth}(x)}_{\text{regularity term}}$$

Link with local variational approaches

Ex : *local TV denoising*, Louchet-Moisan ('11), $\widehat{x} = \arg\min_x \|x - y\|_2^2 + \lambda \mathrm{TV}(x)$.

Gaussian prior

For Gaussian priors, MAP = MMSE = Linear MMSE.

If $X \sim \mathcal{N}(\mu, \Sigma)$ and $N \sim \mathcal{N}(0, \sigma^2 I_p)$ are independant,

$$\hat{x} = \psi(y) := \text{Argmax}_x \; \log p[x|y] = \text{Argmin}_x \frac{1}{\sigma^2}(Ax - y)^t(Ax - y) + (x - \mu)^t \Sigma^{-1}(x - \mu)$$
$$= \mu + \Sigma A^t (A \Sigma A^t + \sigma^2 I_p)^{-1}(y - A\mu)$$

GAUSSIAN PRIOR

For Gaussian priors, MAP = MMSE = Linear MMSE.

If $X \sim \mathcal{N}(\mu, \Sigma)$ and $N \sim \mathcal{N}(0, \sigma^2 I_p)$ are independant,

$$\hat{x} = \psi(y) := \text{Argmax}_x \ \log p[x|y] = \text{Argmin}_x \frac{1}{\sigma^2}(Ax - y)^t(Ax - y) + (x - \mu)^t \Sigma^{-1}(x - \mu)$$
$$= \mu + \Sigma A^t (A\Sigma A^t + \sigma^2 I_p)^{-1}(y - A\mu)$$

Pure denoising case ($A = I_p$). Let $\Sigma = Q\text{diag}(\lambda_1, \dots, \lambda_p)Q^t$, then

$$\widehat{x} = \mu + Q \begin{pmatrix} \frac{\lambda_1}{\lambda_1+\sigma^2} & 0 & \cdots & 0 \\ 0 & \frac{\lambda_2}{\lambda_2+\sigma^2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \frac{\lambda_p}{\lambda_p+\sigma^2} \end{pmatrix} Q^t(y - \mu)$$

$\longrightarrow$ link with the diagonal estimation (Q well chosen basis, dictionary, PCA, etc.), Wiener or thresholding. Ex : Deledalle et al. ('11), Zhang et al ('10).

Denoising

Most of the recent patch-based denoising methods use Gaussian or GMM priors, EPLL ('11), NL-Bayes ('12), PLE ('12), S-PLE ('13), DA3D ('15).

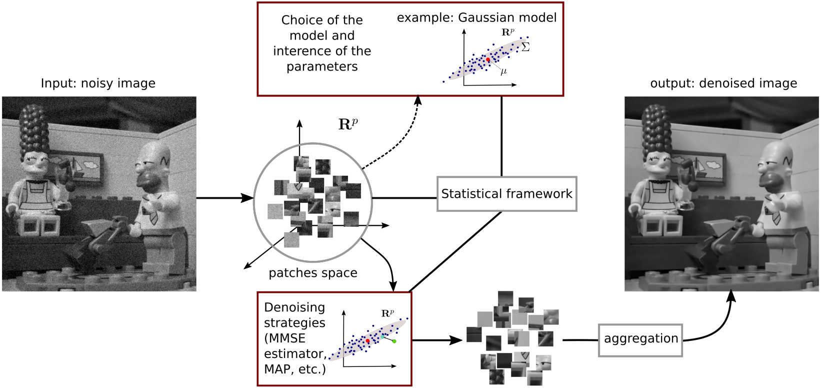

TWO MAIN ISSUES

1. **How to reconstruct an image $\hat{u}$ from the set of denoised patches ?**

2. **How to estimate (μ, Σ) from a set of noisy patches $\{y_i\}$?**

RECONSTRUCTION OF $\hat{u}$ FROM THE DENOISED PATCHES

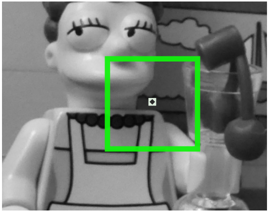

Central value

Reconstruction of $\hat{u}$ from the denoised patches

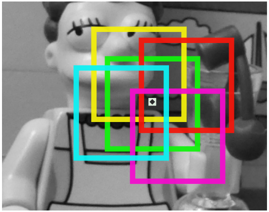

Central value

Aggregation of the estimators

Reconstruction of $\hat{u}$ from the denoised patches

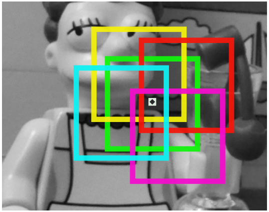

Central value

Aggregation of the estimators

Global optimization in u, EPLL, Zoran-Weiss ('11) ; Teodoro et al ('16),

$$\mathrm{Argmin}_u \frac{\lambda}{2}\|Au-\tilde{u}\|_2^2 - \sum_i \log\ p(x_i^u).$$

Numerical scheme : half quadratic splitting with auxiliary variables $\{z_i\}$:

$$\mathrm{Argmin}_{u,\{z_i\}_i} \frac{\lambda}{2}\|Au-\tilde{u}\|_2^2 + \sum_i \frac{\beta}{2}\|x_i^u - z_i\|^2 - \sum_i \log\ p(z_i)$$

Alternate :

1. update the patches with the MAP, with data fidelity with the patches of the previous step.
2. update u = linear combination of the previous image and the aggregated actual patches.

Estimation of the Gaussian prior on the patches

A first solution : local estimation

Patch y_1 given $\longrightarrow$ set of patches $y_2, \ldots, y_M$ similar to y_1 (for a certain « distance »). Empirical estimation of (μ, Σ) from this set.

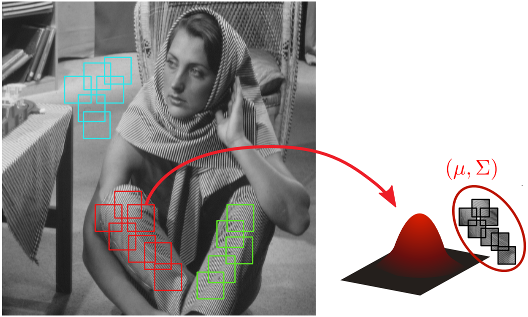

Estimation of the Gaussian prior on the patches

A first solution : local estimation

Patch y_1 given $\longrightarrow$ set of patches $y_2, \ldots, y_M$ similar to y_1 (for a certain « distance »). Empirical estimation of (μ, Σ) from this set.

NL-Bayes, Lebrun et al. ('13)

$$A = \mathrm{I}_p, \quad \widehat{\mu} = \frac{1}{M}\sum_{i=1}^{M} y_i$$

$$\widehat{\Sigma} = \frac{1}{M-1}\sum_{i=1}^{M}[y_i - \widehat{\mu}][y_i - \widehat{\mu}]^T - \sigma^2 \mathrm{I}_p$$

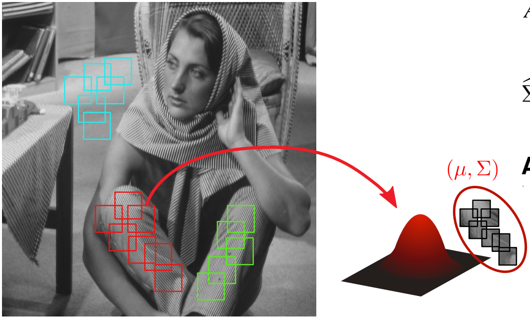

ALGO

1. $\forall i$, estimate (μ_i, Σ_i), restore $\hat{x}_i = \mathbb{E}[X_i | Y_i = y_i, \mu_i, \Sigma_i]$ and aggregate.
2. $\forall i$, re-estimate (μ_i, Σ_i) on the image of 1., and restore/aggregate again.

PROGRESS IN IMAGE DENOISING

image u

$\tilde{u}$, $\sigma = 30$
RMSE = 30, PSNR = 18.58

DCT2 (2001)
RMSE = 16.8, PSNR = 23.6

TV-L2 (1992)
RMSE = 8.35, PSNR = 29.7

NLMeans (2005)
RMSE = 7.98, PSNR = 30.1

NLBayes (2012)
RMSE = 6.23, PSNR = 32.24

Estimation of the Gaussian prior on the patches

Second solution : global estimate with a GMM The hypothesis is that the patches are independant (!) samples of a random variable $X \in \mathbb{R}^p$ with probability density

$$f(x) = \sum_{k=1}^{K} \pi_k g(x; \mu_k, \Sigma_k),$$

where $g(x; \mu_k, \Sigma_k)$ is the density of $\mathcal{N}(\mu_k, \Sigma_k)$.

Estimation of the Gaussian prior on the patches

Second solution : global estimate with a GMM The hypothesis is that the patches are independant (!) samples of a random variable $X \in \mathbb{R}^p$ with probability density

$$f(x) = \sum_{k=1}^{K} \pi_k g(x; \mu_k, \Sigma_k),$$

where $g(x; \mu_k, \Sigma_k)$ is the density of $\mathcal{N}(\mu_k, \Sigma_k)$.

- EPLL, Zoran-Weiss ('11). Global model learned on a large dataset of images.
- fast-EPLL, Parameswaran et al. ('17).
- PLE, Yu et al. ('12) ; more K-means rather than GMM,
- SURE-PLE, Wang et Morel ('13), based on Mixture PCA, Tipping & Bishop ('99)
- Single Frame Denoising and Inpainting, Teodoro et al. ('15)
- HDMI (High-Dimensional Mixture Model for Image denoising), Houdard et al. ('17)

BUT...

THE CURSE OF DIMENSIONALITY

- In high dimension, the data are **isolated** : to capture a fraction $a = 1\%$ of the volume of the unit hypercube in dimension $p = 100$, we need a hypercube of side length $a^{\frac{1}{p}} \simeq 0.95$
- If the number n of samples is not large enough, the estimate $\widehat{\Sigma}$ can be **ill-conditioned or singular**

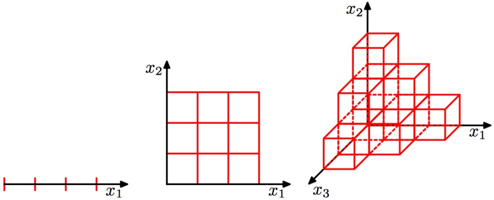

THE CURSE OF DIMENSIONALITY

Some solutions

- **Reduce the dimension**, use small patches (3×3 or 5×5 in NL-Bayes)
- **Regularize** : use hyperpriors HBE, Aguerrebere et al. ('17)
- Add **sparsity hypothesis**
  - covariances of small dimension, given by the model SURE-PLE, Wang and Morel ('13) or
  - small intrinsic dimension, estimated for each Gaussian component HDMI, Houdard et al. ('17)

HBE - Hyperprior on (μ, Σ)

HBE Aguerrebere et al. [17]. Given a set of similar patches $y_1, \ldots, y_M$, we want to estimate the clean patches $x_1, \ldots x_M$ by computing

$$\text{Argmax}_{\{x_i\},\mu,\Sigma} \; p(\{x_i\}_i, \mu, \Sigma \mid \{y_i\}_i) =$$

$$\text{Argmax}_{\{x_i\},\mu,\Sigma} \; \underbrace{p(\{y_i\} \mid \{x_i\}, \mu, \Sigma)}_{\mathcal{N}(A_i x_i, \Sigma_{N_i})} \cdot \underbrace{p(\{x_i\} \mid \mu, \Sigma)}_{\mathcal{N}(\mu, \Sigma)} \cdot \underbrace{p(\mu|\Sigma)}_{\mathcal{N}(\mu_0, \Sigma/\kappa)} \cdot \underbrace{p(\Sigma)}_{\mathcal{IW}(\nu\Sigma_0, \nu)} \; .$$

HBE - Hyperprior on (μ, Σ)

HBE Aguerrebere et al. [17]. Given a set of similar patches $y_1, \ldots, y_M$, we want to estimate the clean patches $x_1, \ldots x_M$ by computing

$$\mathrm{Argmax}_{\{x_i\},\mu,\Sigma} \; p(\{x_i\}_i, \mu, \Sigma \mid \{y_i\}_i) =$$

$$\mathrm{Argmax}_{\{x_i\},\mu,\Sigma} \; \underbrace{p(\{y_i\} \mid \{x_i\}, \mu, \Sigma)}_{\mathcal{N}(A_i x_i, \Sigma_{N_i})} \cdot \underbrace{p(\{x_i\} \mid \mu, \Sigma)}_{\mathcal{N}(\mu,\Sigma)} \cdot \underbrace{p(\mu|\Sigma)}_{\mathcal{N}(\mu_0,\Sigma/\kappa)} \cdot \underbrace{p(\Sigma)}_{\mathcal{IW}(\nu\Sigma_0,\nu)} .$$

Setting $\Lambda = \Sigma^{-1}$, we have

$$\begin{aligned} f(\{x_i\}, \mu, \Lambda) \;&:=\; -\log p(\{x_i\}, \mu, \Lambda \mid \{y_i\}) \\ &=\; \frac{1}{2}(y_i - A_i x_i)^T \Sigma_{N_i}^{-1}(y_i - A_i x_i) - \frac{\nu - p + M}{2} \log|\Lambda| \\ &+\; \frac{1}{2}\sum_{i=1}^{M}(x_i - \mu)^T \Lambda (x_i - \mu) + \frac{\kappa}{2}(\mu - \mu_0)^T \Lambda(\mu - \mu_0) + \frac{1}{2}\mathrm{trace}[\nu \Sigma_0 \Lambda] \end{aligned}$$

on $\mathbb{R}^{pM} \times \mathbb{R}^p \times S_p^{++}(\mathbb{R})$.

The function f is bi-convex in $(\{x_i\}, \mu)$ and $\Lambda \longrightarrow$ numerical scheme with alternate minimization (explicit formulas available).

RESULTS

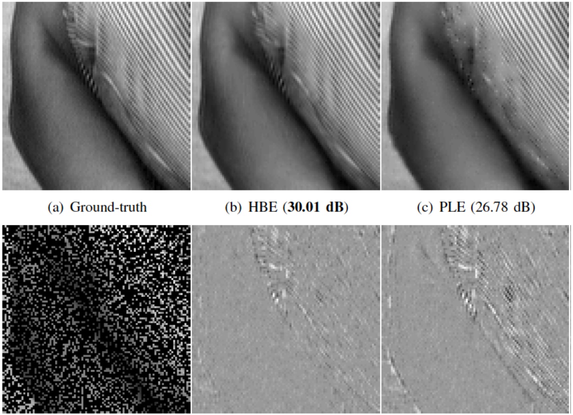

(a) Ground-truth (b) HBE (**30.01 dB**) (c) PLE (26.78 dB)

70% missing pixels.

Results

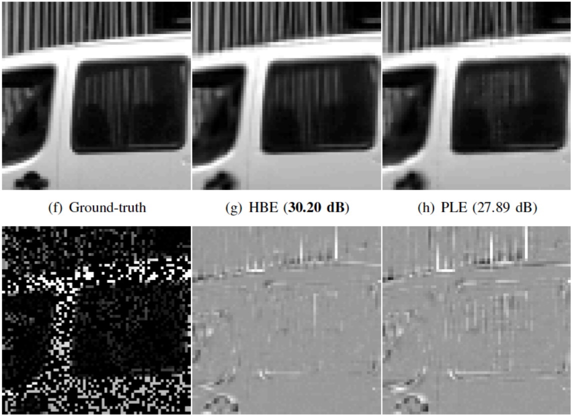

(f) Ground-truth (g) HBE (**30.20 dB**) (h) PLE (27.89 dB)

70% missing pixels.

RESULTS

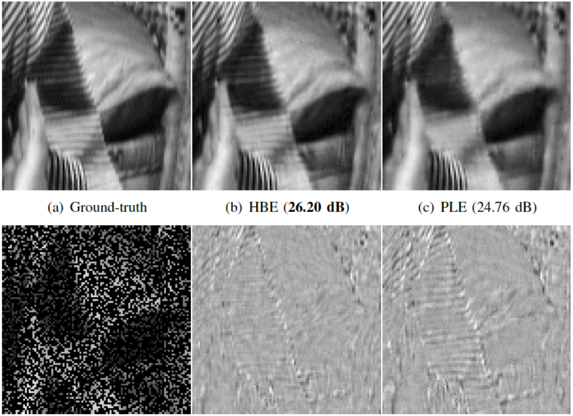

70% missing pixels, Gaussian noise $\sigma = 10$.

Results

(f) Ground-truth (g) HBE (**28.34 dB**) (h) PLE (27.50 dB)

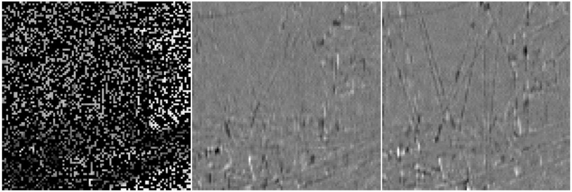

70% missing pixels, Gaussian noise $\sigma = 10$.

Texture synthesis : Efros-Leung (1999)

Efros - Leung ('99)

- Underlying model : Markov random field. The goal is to estimate u_i knowing u in a neighborhood of i.
- First paper using a patch-based approach and exploiting the idea that natural images are redundant.
- Global optimization instead of sequential one Kwatra et al. ('03) , Synthesis patch by patch (instead of pixels) Efros-Freeman ('01) , Mathematical analysis Levina and Bickel ('06)

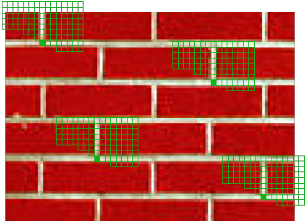

input image

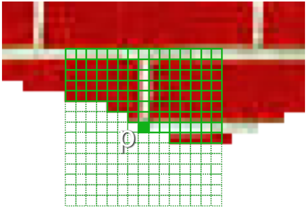

synthesizing a pixel

TEXTURE SYNTHESIS : EFROS-LEUNG (1999)

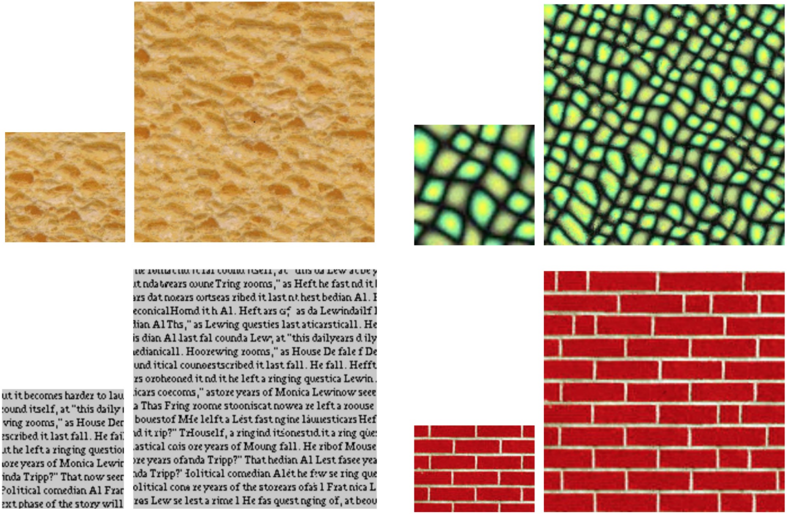

Try online on IPOL : `demo.ipol.im/demo/59/`

Extension : Gaussian modeling of the patches

Proposed by Raad et. al ('14) : to avoid the verbatim effect, iterative generative method with a local Gaussian model and a quilting step (as in Efros-Freeman).

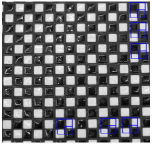

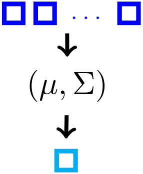

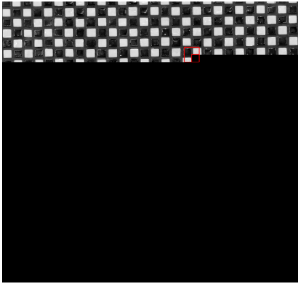

Estimating and sampling a Gaussian patch

Estimating

From a set of m similar patches $y_1, \ldots, y_m$ (seen as vectors in $\mathbb{R}^{s^2}$), estimate the empirical mean and covariance by

$$\mu = \frac{1}{m}\sum_{j=1}^{m} y_j \quad \text{and} \quad \Sigma = \frac{1}{m-1}(Y-\mu)(Y-\mu)^t,$$

where Y is the $s^2 \times m$ matrix containing the patches.

Sampling

How to sample from $\mathcal{N}(\mu, \Sigma)$? Very simple, since

$$z = \mu + \frac{1}{\sqrt{m-1}} \sum_{j=1}^{m} a_j (y_j - \mu),$$

is $\mathcal{N}(\mu, \Sigma)$ distributed when the $a_j \sim \mathcal{N}(0, 1)$ i.i.d.

From patch models to image model

Patch model fusion Saint-Dizier, Delon, Bouveyron ('18)

Model p_i on patch x_i, $i = 1, \ldots, I$. Each patch domain is denoted Ω_i.

Fused image model :

$$p(u) \propto \prod_{i=1}^{I} p_i(u_{|\Omega_i})$$

Rk : Model obtained by conditioning on the equality of patches on their intersections. Generalization of EPLL, Zoran Weiss ('12)

GAUSSIAN CASE

When $p_i = \mathcal{N}(\mu_i, \Sigma_i)$, then $p = \mathcal{N}(\mu, \Sigma)$, with

$$\Sigma^{-1}(k,l) = \sum_{i;k,l \in \Omega_i} \Sigma_i^{-1}(k,l),$$

$$(\Sigma^{-1}\mu)(k) = \sum_{i;k \in \Omega_i} (\Sigma_i^{-1}\mu_i)(k).$$

CONCLUSION

Today : Gaussian models on patches (non stationary, non periodic), many applications

Next lecture (27 February) : Julie Delon, on GMM.

EXERCISE (PRACTICAL)

1. Load the *Barbara* image.
2. Add Gaussian noise to it, with variance $\sigma = 20$ for instance.
3. Crop two parts of the image of size 100×100 for instance : a piecewise smooth part, and a highly textured part.
4. Implement different denoising methods : simple average value in a neighborhood, NL-means and NL-Bayes. What are the key parameters ? And what do you observe ?