The 22nd British Machine Vision Conference
University of Dundee, 29 August – 2 September 2011

Image denoising with patch-based PCA: local versus global

Charles-Alban Deledalle[1], Joseph Salmon[2], Arnak Dalalyan[3]

[1] Institut Telecom, Telecom ParisTech, CNRS LTCI, Paris, France
[2] LPMA, Université Paris Diderot-Paris 7, Paris, France
[3] Imagine / LIGM, Université Paris-Est, Champs-sur-Marne, France

August 31, 2011

Motivations

Image denoising: find an estimation of the true image from the noisy image.

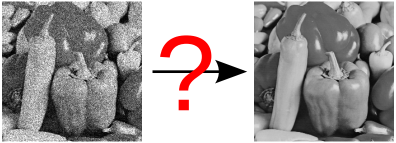

Noisy image

True image

Motivations

Image denoising: find an estimation of the true image from the noisy image.

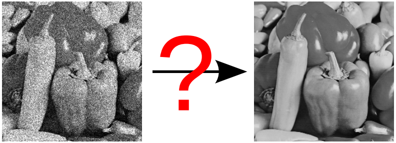

Noisy image   True image

How to denoise an image?

- Assuming sparsity,
- Assuming regularity,
- Assuming self-similarity,
- With hybrid models

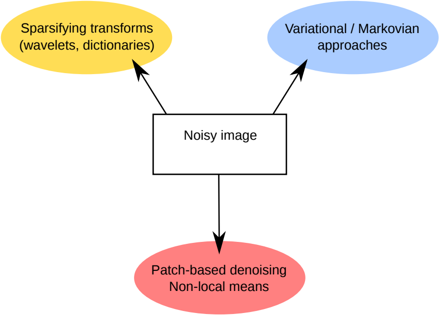

Motivations

Image denoising: find an estimation of the true image from the noisy image.

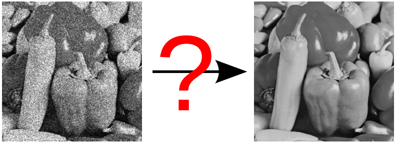

Noisy image True image

How to denoise an image?

- Assuming sparsity,
- Assuming regularity,
- Assuming self-similarity,
- With hybrid models

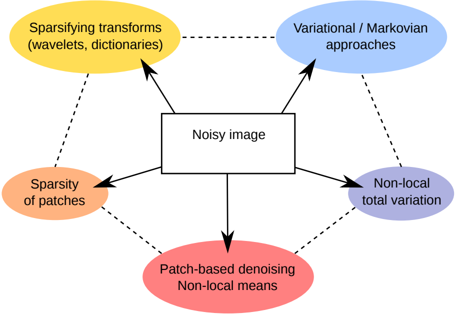

Image denoising: find an estimation of the true image from the noisy image.

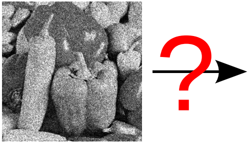

Noisy image

True image

How to denoise an image?

- Assuming sparsity,
- Assuming regularity,
- Assuming self-similarity,
- With hybrid models

Our solution:

- Patch decomposition,
- Principal component analysis (PCA),
- Sparse reconstruction.

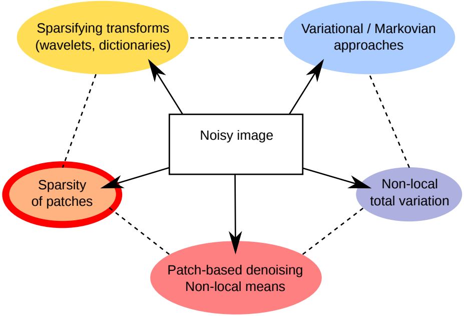

Image denoising: find an estimation of the true image from the noisy image.

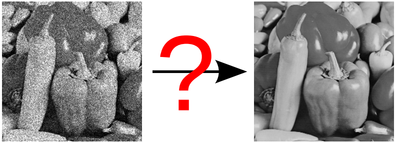

Noisy image

True image

How to denoise an image?

- Assuming sparsity,
- Assuming regularity,
- Assuming self-similarity,
- With hybrid models

Our solution:

- Patch decomposition,
- Principal component analysis (PCA),
- Sparse reconstruction.

Main advantages:

- Easy to design,
- Good performance,
- Few parameters.

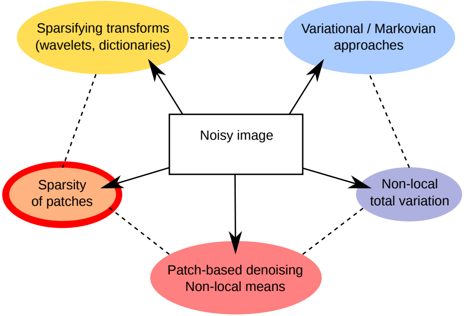

Outline

Outline

1. Patch dictionary with PCA
2. Local adaptive dictionaries
3. Denoising in the patch PCA domain
4. Experiments and results

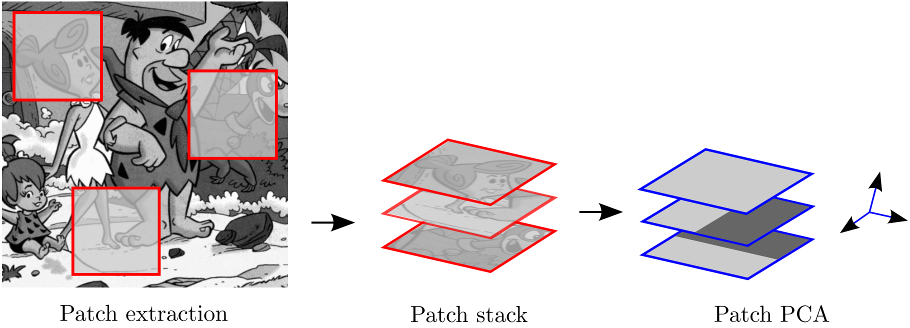

Patch-based image model

- Extract patches: small sub-images extracted in the neighbourhood of a pixel
  $\rightarrow$ represented by vectors of size n,
  the size n is usually between $9 = 3 \times 3$ and $64 = 8 \times 8$,
- We are now interested in analysing the collection of patches.

Patch-based global PCA

(a) Input image

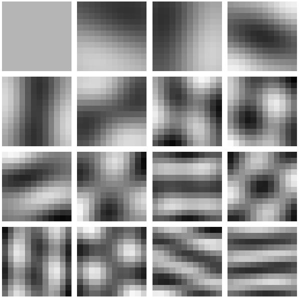

(b) 16 first axes

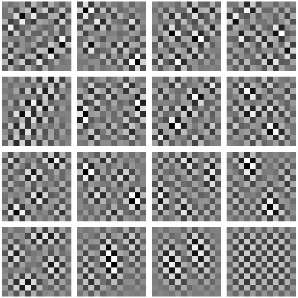

(c) 16 last axes

Principal component analysis

- Compute the $n \times n$ empirical covariance matrix (independent of the image size),
- Size n relatively small $\Rightarrow$ simple extraction of the n eigenvalues and eigenvectors,
- The eigenvectors form an orthogonal basis.

√ They encode the main patterns that occur in the image,

× They are unable to adequately represent a significant proportion of patches.

Outline

1. Patch dictionary with PCA
2. Local adaptive dictionaries
3. Denoising in the patch PCA domain
4. Experiments and results

Patch-based local PCA

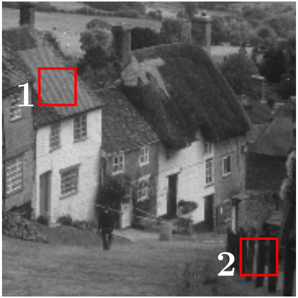

(d) Local search windows

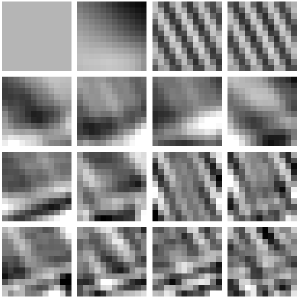

(e) 16 first axes in window 1

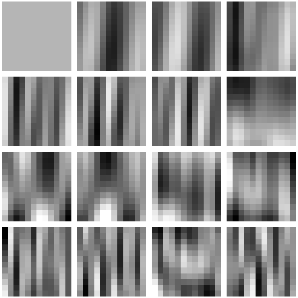

(f) 16 first axes in window 2

- Compute the PCA in a sliding local window,
  - → Similar idea in [Muresan and Parks, 2003, Zhang et al., 2010],
- Local dictionaries can be quite different.

√ This decomposition suitably describes the local features of the image,

× It leads to a longer computing time.

Patch-based hierarchical PCA

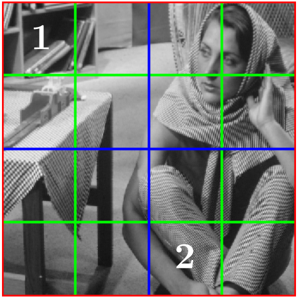

(g) Quad-tree decomposition

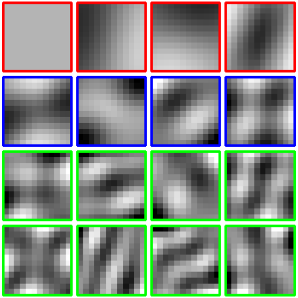

(h) 16 first axes in part 1

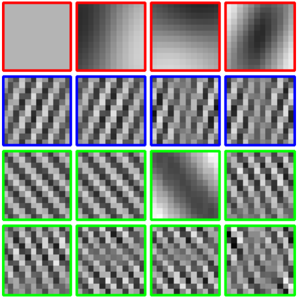

(i) 16 first axes in part 2

- Learn recursively the dictionary with PCA on a quad-tree decomposition,
- Local dictionaries describe more and more the local features.

√ This approach also suitably describes the local features of the image,

√ Compared to patch-based local PCA, it reduces the computing time.

Outline

Keep or kill (KOK)

- Signal is concentrated on the first axes while noise is spread out over all directions,
- Keep the $n' < n$ first axes and kill the remaining axes,

$\surd$ Preserves the maximal variance among all subspaces of dimension n',

$\times$ Cannot represent the many "rare" patches of the image.

Denoising in the patch PCA domain

Keep or kill (KOK)

- Signal is concentrated on the first axes while noise is spread out over all directions,
- Keep the $n' < n$ first axes and kill the remaining axes,

√ Preserves the maximal variance among all subspaces of dimension n',

× Cannot represent the many "rare" patches of the image.

Thresholding

- All axes are relevant to describe the whole of patches,
- Only few of them are required for a given patch,
- *Sparse* representations obtained by shrinking of coefficients:
  - Hard thresholding (set to zero the coefficients smaller than a threshold λ),
  - Soft thresholding (same with rescaling the coefficients bigger than λ),
  - Linear rescaling[a] (Wiener filtering, rescale all coefficients wrt the noise variance).

√ Minimise the squared error under sparsity priors,

√ Can represent both "frequent" and "rare" patches.

[a] choice of [Muresan and Parks, 2003, Zhang et al., 2010].

Denoising in the patch PCA domain

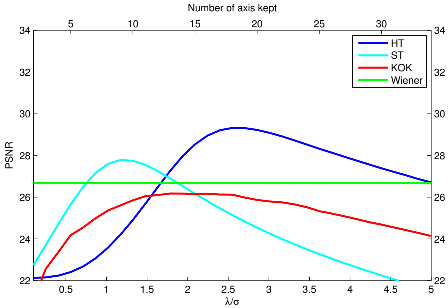

(j) Cameraman (k) House

Comparing various strategies of reconstruction from the projections onto the basis provided by PCA, for House and Cameraman ($\sigma = 20$): Hard Thresholding, Soft Thresholding, "Keep or Kill" and Wiener Filtering. The x axes are different: on the top is the number of coefficients kept for the "Keep or Kill" strategy and on the bottom is the threshold ratio λ/σ.

Denoising in the patch PCA domain

Reprojection [Dabov et al., 2007, Salmon and Strozecki, 2010]

- Each patch is denoised several times in the different stacks,
- Patches naturally overlap,
- Reproject all denoised patches to their original positions,
  - ⇒ Uniformly average the several estimates in each pixel.

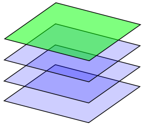

Stack of denoised patches after collaborative filtering

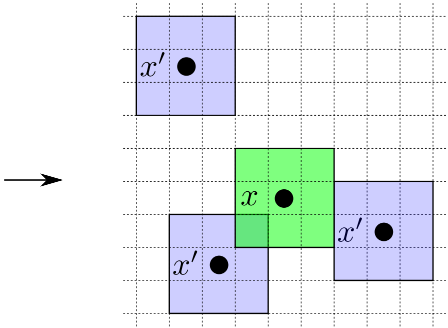

Reprojection

Outline

- Challenge the non-local means (NL means) [Buades et al., 2005]

(a) Noisy image

(b) NL means (PSNR=32.90)

(c) Local PCA (PSNR=33.70)

BM3D versus the patch-based local PCA. Noise level $\sigma = 10$.

- Challenge the non-local means (NL means) [Buades et al., 2005]
- Simple implementation almost as good as one of the best approaches [Dabov et al., 2007]


(a) Noisy image


(b) BM3D (PSNR=33.90)


(c) Local PCA (PSNR=33.70)

BM3D versus the patch-based local PCA. Noise level $\sigma = 10$.

- Local: +1 dB

(a) Noisy image (b) Global (PSNR=33.6) (c) Local (PSNR=34.8)

Global versus local. Noise level $\sigma = 10$, with their PSNR.

- Local: $+1$ dB, better restoration of local features

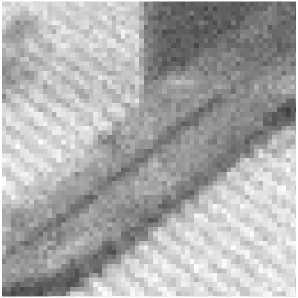

(a) Noisy image

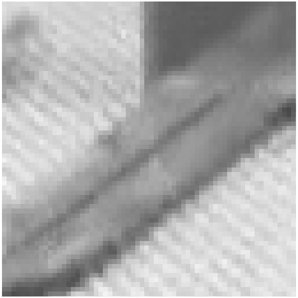

(b) Global (PSNR=33.6)

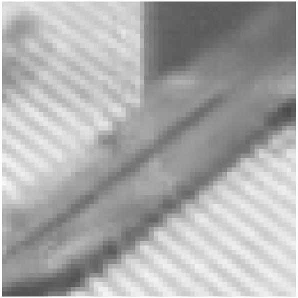

(c) Local (PSNR=34.8)

Global versus local. Noise level $\sigma = 10$, with their PSNR.

- Local: $+1$ dB, better restoration of local features

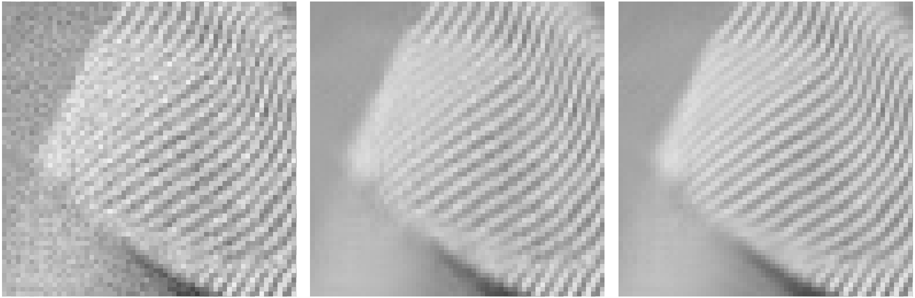

(a) Noisy image (b) Global (PSNR=33.6) (c) Local (PSNR=34.8)

Global versus local. Noise level $\sigma = 10$, with their PSNR.

- Local: +1 dB, better restoration of local features

(a) Noisy image (b) Global (PSNR=29.7) (c) Local (PSNR=31.1)

Global versus local. Noise level $\sigma = 20$, with their PSNR.

- Local: $+1$ dB, better restoration of local features

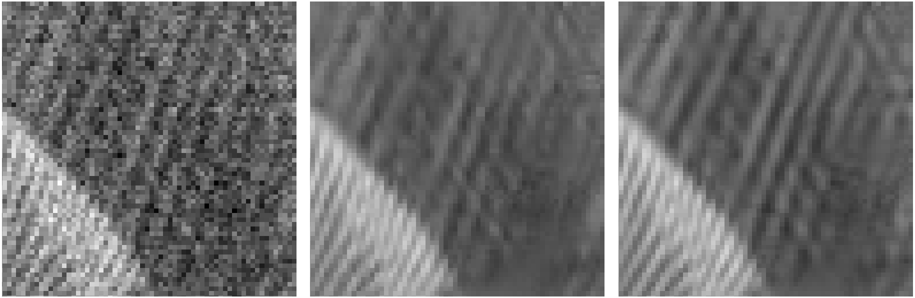

(a) Noisy image (b) Global (PSNR=33.6) (c) Local (PSNR=34.8)

Global versus local. Noise level $\sigma = 20$, with their PSNR.

- Local: $+1$ dB, better restoration of local features

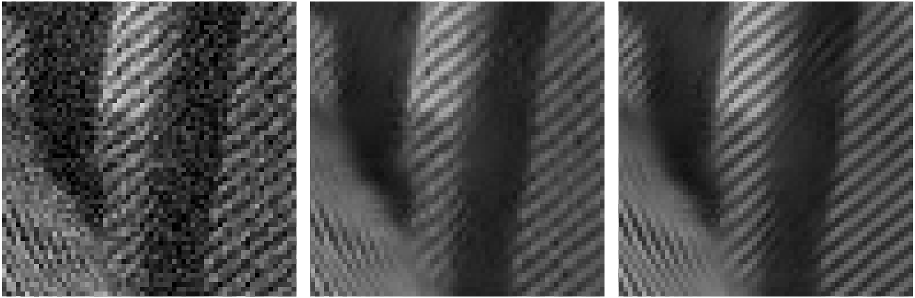

(a) Noisy image (b) Global (PSNR=33.6) (c) Local (PSNR=34.8)

Global versus local. Noise level $\sigma = 20$, with their PSNR.

- Local: $+1$ dB, better restoration of local features, overfitting problem

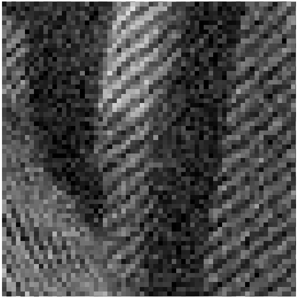

(a) Noisy image

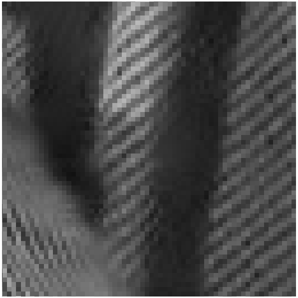

(b) Global (PSNR=33.6)

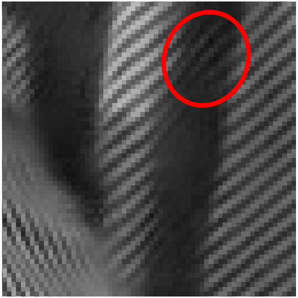

(c) Local (PSNR=34.8)

Global versus local. Noise level $\sigma = 20$, with their PSNR.

Conclusion and perspectives

Conclusions

- Learn an orthogonal dictionary from the data itself with patch-based PCA,
- Simple implementation,
- Provide results close to or challenging other state-of-the-art approaches,
- Relatively fast and require few parameters:
  - Size of the patches,
  - Threshold level,
  - Searching zone or number of recursions (optionals).

Conclusion and perspectives

Conclusions

- Learn an orthogonal dictionary from the data itself with patch-based PCA,
- Simple implementation,
- Provide results close to or challenging other state-of-the-art approaches,
- Relatively fast and require few parameters:
  - Size of the patches,
  - Threshold level,
  - Searching zone or number of recursions (optionals).

Perspectives

- Two-pass filtering [Dabov et al., 2007, Zhang et al., 2010]
- Perform PCA on clusters [Mairal et al., 2009]
- Take into account the correlation between nearby patches:
  - model the statistical correlation of coefficients, or [Portilla et al., 2003],
  - use *grouped*-sparsity. [Mairal et al., 2009].

Questions?

deledalle@telecom-paristech.fr
joseph.salmon@duke.edu
dalalyan@imagine.enpc.fr

http://josephsalmon.org/ → More details and software available.

[Buades et al., 2005] Buades, A., Coll, B., and Morel, J.-M. (2005).
A review of image denoising algorithms, with a new one.
Multiscale Model. Simul., 4(2):490–530.

[Dabov et al., 2007] Dabov, K., Foi, A., Katkovnik, V., and Egiazarian, K. O. (2007).
Image denoising by sparse 3-D transform-domain collaborative filtering.
IEEE Trans. Image Process., 16(8):2080–2095.

[Mairal et al., 2009] Mairal, J., Bach, F., Ponce, J., Sapiro, G., and Zisserman, A. (2009).
Non-local sparse models for image restoration.
ICCV.

[Muresan and Parks, 2003] Muresan, D. D. and Parks, T. W. (2003).
Adaptive principal components and image denoising.
In *ICIP*, pages 101–104.

[Portilla et al., 2003] Portilla, J., Strela, V., Wainwright, M., and Simoncelli, E. P. (2003).
Image denoising using scale mixtures of gaussians in the wavelet domain.
IEEE Trans. Image Process., 12(11):1338–1351.

[Salmon and Strozecki, 2010] Salmon, J. and Strozecki, Y. (2010).
From patches to pixels in non-local methods: Weighted-Average reprojection.
In *ICIP*.

[Zhang et al., 2010] Zhang, L., Dong, W., Zhang, D., and Shi, G. (2010).
Two-stage image denoising by principal component analysis with local pixel grouping.
Pattern Recogn., 43(4):1531–1549.