Building invariance properties for dictionaries of SAR image patches

Sonia Tabti[1], Charles-Alban Deledalle[2], Loïc Denis[3], Florence Tupin[1]

[1]Institut Mines-Télécom, Télécom ParisTech CNRS-LTCI, France
[2] IMB, CNRS-Univ. Bordeaux, France
[3] Laboratoire Hubert Curien, CNRS-Univ. Saint-Etienne, France

July 18, 2014

How to represent information in SAR images?

Different levels of representation

- Pixels.
- Local descriptors.
- Patches.

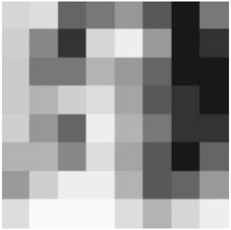
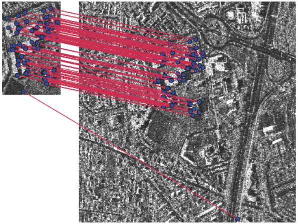
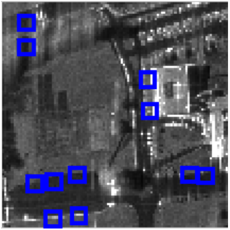

Our motivations

- Elaborate patch-based dictionaries approaches for SAR images.
- Adapt to diversity of SAR data content.
- Provide compact dictionaries.

  $\rightarrow$ Include invariance properties:
  - Shift-invariance.
  - Affine radiometric transform invariance.

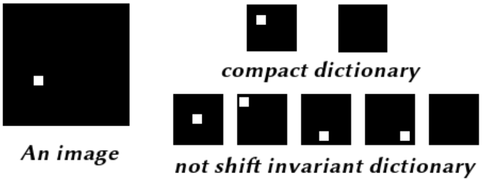

Outline

- Signal $\simeq$ a linear combination of a few number of dictionary elements (the **atoms**).
- **1st problem:** Learn a dictionary of patches.
- **2nd problem:** Find the best coefficients (**sparse coding**).

An image

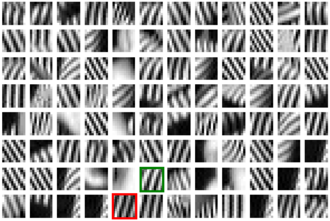

A trained dictionary

 = Coeff 1  + Coeff 2 

An illustration of the sparse coding principle (sparsity degree=2)

- **Fixed**, eg: wavelets, Discrete Cosine Transform (DCT).
- **Adaptive**
  eg: Principal Component Analysis (PCA),
  K-SVD [Aharon et al., 2006].
- **Shift invariant**
  eg: MoTIF (Matching of Time Invariant Features)
  [P. Jost and Gribonval, 2006],
  Epitomes.

128×128 image 50×50 epitome [1]

[1] Image courtesy: Epitomic analysis of appearance and shape, N. Jojic and al.

Outline

Numerical experiment:

- A K-SVD dictionary learned on the log of a 100-looks SAR image (256 atoms of width 11×11)
- Sparse-coding algorithm: OMP (Orthogonal Matching Pursuit)
- Logarithm of patches corrupted with 1-look gamma noise (speckle) extracted from a SAR image.
- Criterion: MSE (Mean Square error) computed over 100 noisy realizations.

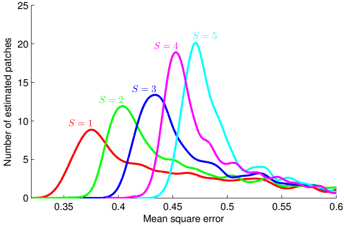

MSE histograms with different sparsity degrees.

Outline

- Context
- Dictionaries of patches
  - Main principle of sparse representations
  - Some types of patch-based dictionaries
- Invariance properties in patch-based dictionaries for SAR images
  - Optimal sparsity degree for 1-look SAR images
  - 1-sparse image approximation method with invariance properties
  - Proposed dictionary learning procedure
- Conclusions and perspectives

Introducing radiometric invariance [2]

The square difference, $\mathrm{SSD}(\boldsymbol{a}, \boldsymbol{p}) = \|\boldsymbol{p} - \boldsymbol{a}\|^2$, is not robust against radiometric changes. Invariant critera should answer positively to the question:

Enforcing affine radiometric invariance into the SSD leads to:

$$\mathrm{SSD}^{CI}(\boldsymbol{a}, \boldsymbol{p}) = \begin{cases} (1 - \mathcal{C}(\boldsymbol{p}, \boldsymbol{a})^2)\|\boldsymbol{p} - \bar{\boldsymbol{p}}\|^2 & \text{if} \quad \boldsymbol{a} \neq \bar{\boldsymbol{a}} \quad \text{and} \quad \boldsymbol{p} \neq \bar{\boldsymbol{p}} \\ \|\boldsymbol{p} - \bar{\boldsymbol{p}}\|^2 & \text{else} \end{cases}$$

with $\mathcal{C}(\boldsymbol{p}, \boldsymbol{a}) = \dfrac{\langle \boldsymbol{p} - \bar{\boldsymbol{p}}, \boldsymbol{a} - \bar{\boldsymbol{a}} \rangle}{||\boldsymbol{p} - \bar{\boldsymbol{p}}||_2 ||\boldsymbol{a} - \bar{\boldsymbol{a}}||_2}$ the normalized correlation,

- $\boldsymbol{p}$, $\boldsymbol{a}$: $\sqrt{m} \times \sqrt{m}$: noisy patches and atom resp.
- $\bar{\boldsymbol{a}}$ and $\bar{\boldsymbol{p}}$: patches with constant values equal to the means of $\boldsymbol{a}$ and $\boldsymbol{p}$ resp.

[2] Template matching with noisy patches: a contrast-invariant GLR test, EUSIPCO 2013
Charles-Alban Deledalle, Loïc Denis, Florence Tupin

Approximation-error of all the patches with a particular atom

Measuring the quality of an atom $\boldsymbol{a}$ with respect to an image $\boldsymbol{y}$ requires evaluating $\mathrm{SSD}^{CI}(\boldsymbol{a}, \boldsymbol{p})$ for all patches $\boldsymbol{p}$ $\rightarrow$ **time consuming**

Proposed extension: fast computation of the map of the criterion between a given atom and all the patches of a **whole image** $\boldsymbol{y}$.

Contrast-invariant and 1-sparse approximation error map of $\boldsymbol{y}$ by $\boldsymbol{a}$

$$\mathrm{SSD}^{\mathrm{CI}}_{\mathrm{map}}(\boldsymbol{a}, \boldsymbol{y}) = \begin{cases} \boldsymbol{y}^2 \star \iota - m(\boldsymbol{y} \star \iota / m)^2 - \dfrac{(\boldsymbol{y} \star (\boldsymbol{a} - \bar{\boldsymbol{a}}))^2}{\|\boldsymbol{a} - \bar{\boldsymbol{a}}\|^2} & \text{if} \quad \boldsymbol{a} \neq \bar{\boldsymbol{a}} \\ \boldsymbol{y}^2 \star \iota - m(\boldsymbol{y} \star \iota / m)^2 & \text{else} \end{cases}$$

where ι is the support of a patch (an $\sqrt{m} \times \sqrt{m}$ window), $\star$ the 2D-discrete convolution, and $\cdot^2$ the element-wise square function.

$\Rightarrow$ Fast computation in the Fourier domain with $O(N \log N)$ operations where N is the size of the image.

Including the shift-invariance

Key idea: Selecting the neighboring patch around a pixel with the smallest error boils down to applying an erosion to $\mathrm{SSD}^{\mathrm{CI}}_{\mathrm{map}}$ (image tiling).

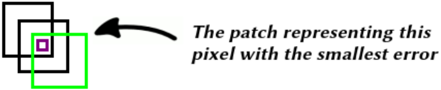

$\mathrm{SSD}^{\mathrm{CI}}_{\mathrm{map}}$ with shift-invariance

$$\mathrm{SSD}^{\mathrm{CI+SI}}_{\mathrm{map}}(\boldsymbol{a}, \boldsymbol{y}) = \mathrm{erode}(\ \mathrm{SSD}^{\mathrm{CI}}_{\mathrm{map}}(\boldsymbol{a}, \boldsymbol{y}), \iota)$$

where $\mathrm{erode}(\cdot, \iota)$ is the erosion operator with structural element ι.

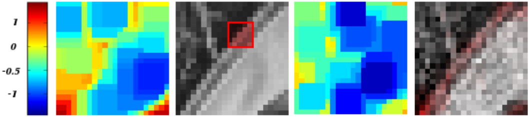

Outline

Proposed dictionary learning procedure

- $\mathcal{C} = \{\boldsymbol{c}_t\}$: set of candidate patches to be atoms or not (e.g., extracted from 100-looks SAR images).

- $\boldsymbol{y}$ is the noisy image to approximate.

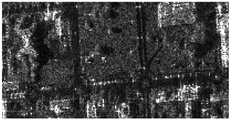

- **Goal:** learn a dictionary $\mathcal{D}$ with a fixed (low) number of atoms from $\mathcal{C}$ which approximate the content of $\boldsymbol{y}$ with respect to the noise level.

- **Method:** replace iteratively one atom by a candidate from $\mathcal{C} = \{\boldsymbol{c}_t\}$.

- **Mean:** minimize $\mathrm{SSD}_{\mathrm{map}}^{\mathrm{CI+SI}}$ at each step to improve the dictionary.

Computation of the approximation error map

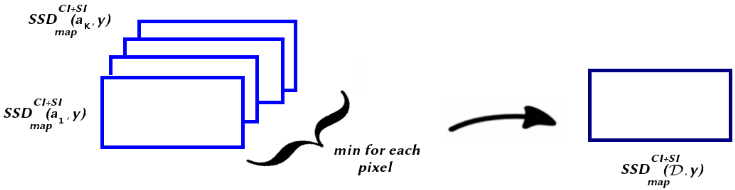

Approximation error map of $\boldsymbol{y}$ with a dictionary $\mathcal{D}$ at pixel i

$$\mathrm{SSD}_{\mathrm{map}}^{\mathrm{CI+SI}}(\mathcal{D}, \boldsymbol{y})_i = \min_k \ \mathrm{SSD}_{\mathrm{map}}^{\mathrm{CI+SI}}(a_k, \boldsymbol{y})_i$$

Dictionary update: is $\boldsymbol{c}_t$ replacing $\boldsymbol{a}_k$?

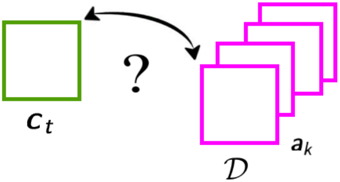

Gain map at pixel i if candidate $\boldsymbol{c}_t$ replaces an atom $\boldsymbol{a}_k$

$$\mathrm{SSD}_{\mathrm{map}}^{\mathrm{CI+SI}}(\mathcal{D}, \boldsymbol{y})_i - \mathrm{SSD}_{\mathrm{map}}^{\mathrm{CI+SI}}(\mathcal{D} \cup \{\boldsymbol{c}_t\} \setminus \{\boldsymbol{a}_k\}, \boldsymbol{y})_i$$

Condition to add a candidate patch c_t, to the dictionary: find an atom a_k such that the average gain is positive (ie the average approximation error is reduced).

Results

Noisy image $\boldsymbol{y}$

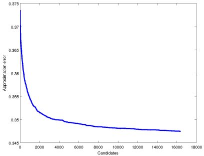

Energy decay wrt atoms updated

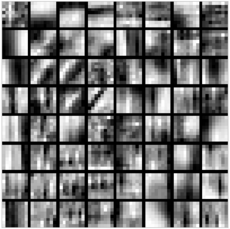

K-SVD dictionary

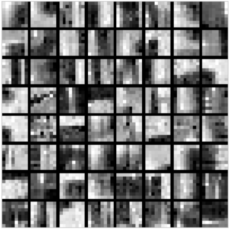

Our dictionary

Results: Illustration of the dictionnary

Which atom is used to approximate patches from the noisy image?

▷ all features

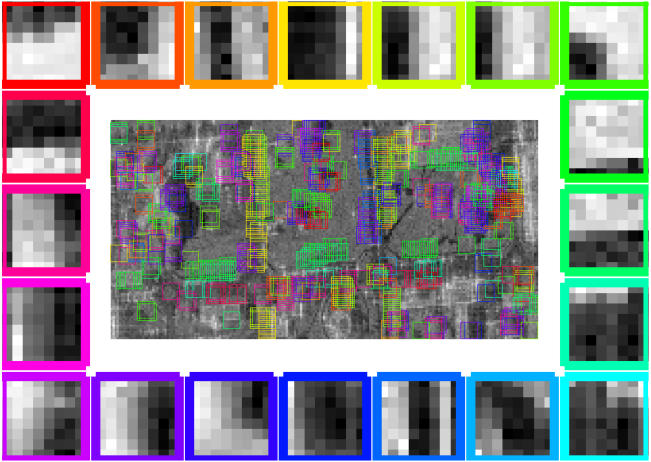

Results: Illustration of the dictionnary

Which atom is used to approximate patches from the noisy image?

▷ vertical edges

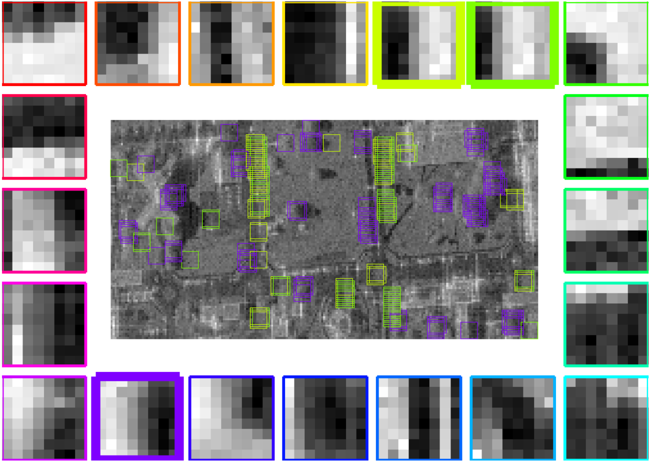

Results: Illustration of the dictionnary

Which atom is used to approximate patches from the noisy image?

▷ horizontal edges

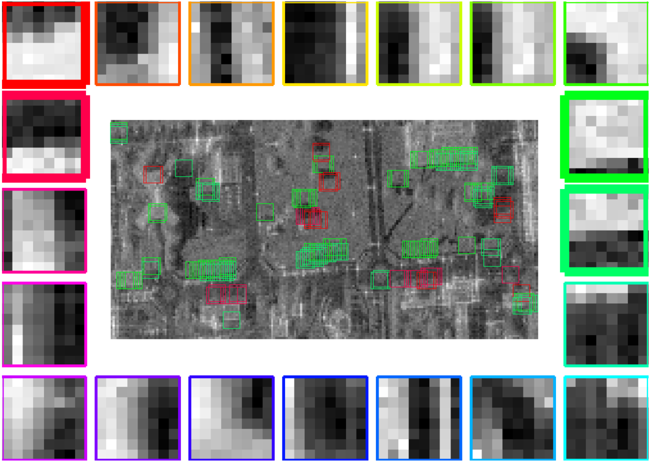

Results: Illustration of the dictionnary

Which atom is used to approximate patches from the noisy image?

▷ corners

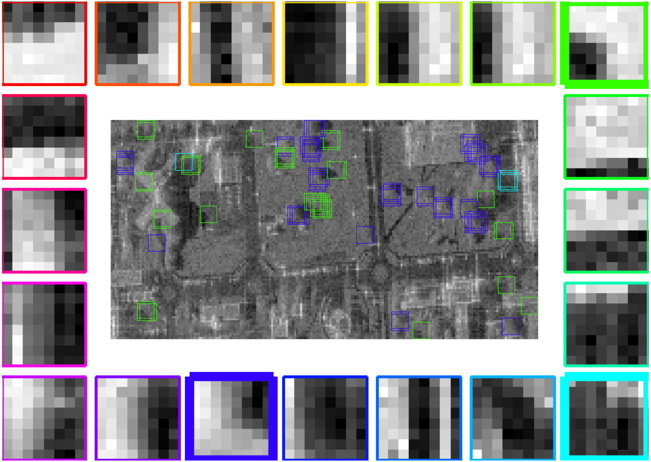

Results: Illustration of the dictionnary

Which atom is used to approximate patches from the noisy image?

▷ low-radiometry linear features

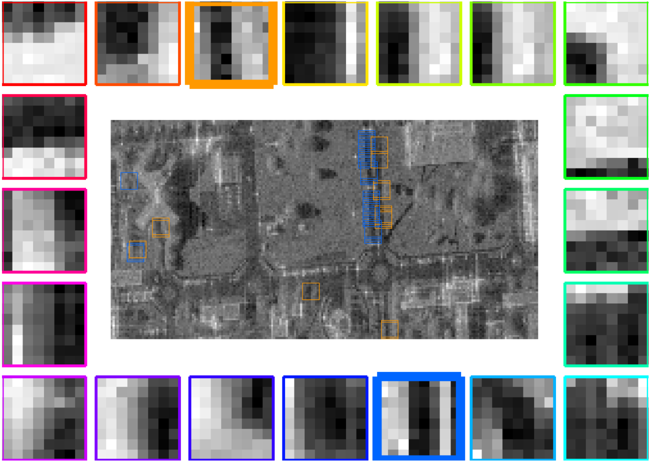

Results: Illustration of the dictionnary

Which atom is used to approximate patches from the noisy image?

▷ high-radiometry linear features

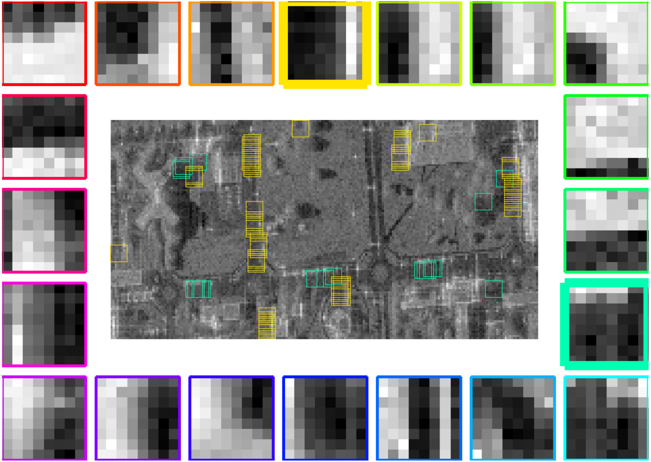

Contributions:

- A 1-sparse similarity criterion with shift and affine radiometric transform invariance.
  - Can be adapted to compare the performances of different dictionaries.
- A dictionary learning algorithm with an updating procedure based on this criterion.
  - Complexity: $O(TN \log N)$ with T the number of candidates.
  - Provides a compact summary of an image (no redundant atoms).
  - Is a good way to initialize a dictionary based method.

Perspectives:

- Replacement of more than one atom by iteration.
- Applications to denoising and clustering.

Thank you for your attention.
Any questions?

sonia.tabti@telecom-paristech.fr

References I

Aharon, M., Elad, M., and Bruckstein, A. (2006).
The K-SVD: An algorithm for designing of overcomplete dictionaries for sparse representation.
IEEE Trans. on Signal Processing, 54(11):4311–4322.

P. Jost, P. Vandergheynst, S. L. and Gribonval, R. (2006).
MoTIF: an efficient algorithm for learning translation invariant dictionaries.
In *IEEE International Conference on Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006*, volume 5, pages V–V. IEEE.

Computation of the gain: different cases

Trick to make it fast: keep track of the two atoms with the best approximation errors in $\mathcal{D}$. Their indices: k_i^{1st}, k_i^{2nd} resp.

- $k \neq k_i^{1st}$: Either $\boldsymbol{c}_t$ improves or does not affect the error at pixel i and the gain is:

$$\max(\ \mathrm{SSD}^{\mathrm{CI+SI}}_{\mathrm{map}}(\mathcal{D}, \boldsymbol{y})_i -\ \mathrm{SSD}^{\mathrm{CI+SI}}_{\mathrm{map}}(\boldsymbol{c}_t, \boldsymbol{y})_i, 0) \quad (1)$$

- $k = k_i^{1st}$:
  - $\boldsymbol{c}_t$ is better than $\boldsymbol{a}_{k_i^{1st}}$ and $\boldsymbol{a}_{k_i^{2nd}}$ so the gain is:

$$\mathrm{SSD}^{\mathrm{CI+SI}}_{\mathrm{map}}(\mathcal{D}, \boldsymbol{y})_i -\ \mathrm{SSD}^{\mathrm{CI+SI}}_{\mathrm{map}}(\boldsymbol{c}_t, \boldsymbol{y})_i \quad (2)$$

  - $\boldsymbol{c}_t$ is worse than $\boldsymbol{a}_{k_i^{2nd}}$ and the gain is:

$$\mathrm{SSD}^{\mathrm{CI+SI}}_{\mathrm{map}}(\mathcal{D}, \boldsymbol{y})_i -\ \mathrm{SSD}^{\mathrm{CI+SI}}_{\mathrm{map}}(\boldsymbol{a}_{k_i^{2nd}}, \boldsymbol{y})_i \quad (3)$$