Low-rankness transfer for denoising Sentinel-1 SAR images

Nicolas Brodu

HAL Id: hal-01936331

https://hal.science/hal-01936331

Submitted on 27 Nov 2018

Low-rankness transfer for denoising Sentinel-1 SAR images

Nicolas Brodu

Inria Bordeaux Sud-Ouest, 200 Avenue de la vieille tour, 33405 Talence

Email: nicolas.brodu@inria.fr

Abstract—This paper introduces a new algorithm for denoising SAR images. It is directly applicable to Sentinel-1 GRD images, without the need for single-look complex operations. The algorithm builds on the non-local patch matching idea for statistical denoising, similar to the SAR-BM3D and NL-SAR algorithms, but introduces two corrections : 1. A non-uniform prior for the reflectance values is used for the patch matching, thus allowing a better fit to the data, and 2. Denoising is performed in singular values space, with a prior distribution of expected "clean" singular values learned and transfered from optical images. The denoised SAR images show reduced amount of speckle compared to alternative methods.

I. INTRODUCTION

Earth observation programs, such as Copernicus [1] from the European Space Agency (ESA), rely on both optical and radar satellites in order to monitor the environment at short time intervals. The image resolutions are relatively low compared to alternatives, on the order of 10m, but the temporal resolution is excellent, with high repeatability over the same coordinates ($\approx$3-4 days in European latitudes). This configuration is ideal for monitoring land changes, for agricultural purposes, for large scale environmental surveys, etc. When they are available, optical data are best for detecting vegetation and classifying the observed area. However, clouds limit the usefulness of optical data. In tropical regions especially, these data become nearly useless during the rainy season. Synthetic Aperture Radar (SAR) images can then be used as fallback, since these can be acquired even with a high cloud coverage. ESA's Copernicus program thus provides both SAR and optical images, with respectively their Sentinel-1 and Sentinel-2 satellites. However, SAR images are plagued by speckle noise, which makes them not as practical as optical data for land classification. In addition, freely available images for civil applications are rarely in full-polarimetric format, thus limiting the signal information that can be used for classification. In practice, Sentinel-1 images are provided with single or dual polarimetry depending on the region of interest. Thus, the goal of this paper is not to compare with full-polarimetric denoisers and classifiers [2], [3], that may exhibit impressive results but which are of limited use in practice for monitoring the environment with freely available data. Approaches using Deep Learning have recently gained interest [4]. Their problem is similar to that of inpainting: how to ensure that details recovered from the noisy image are actually present in the original data, and not inferred from some external (possibly gigantic) database of similar images. This paper introduces a new method for denoising the mono or dual polarimetry SAR images, making them look more similar to optical images. It is based on non-local patch matching, similar to the popular SAR-BM3D algorithm [6]. Patch matching has been validated [8] not to introduces spurious details, which is ensured by a well-defined statistical test for patch similarity (see the next section). Two key points are introduced in this paper :

- A better prior for (clean) image intensity values. In SAR-BM3D, a uniform prior is used [6]. This is meaningless over a large range of values, and false, since the intensity values in an image are not at all uniformly distributed between their minimal and maximal bounds.

- Denoising operates in singular value space - unlike SAR-BM3D that uses shrinkage in wavelet space - and the method is further improved by allowing to set a prior distribution of singular values. This is accomplished by a low-rank transfer method [9], here adapted to a SAR context, together with using optical Sentinel-2 images as prior information for building statistics.

This combination of using a complete patch-comparison distance together with an improved shrinkage process is the core contribution of this paper. Section II introduces patch-based denoising and the two key improvements on a theoretical basis. Section III shows an application of the method and compares the results to widely used alternatives.

II. PATCH-BASED DENOISING

A. General method

Consider an observed image Z described by its pixels $z(s)$, where $s = (row, column)$ notes the pixel sites. We would like to recover a clean image $X = \{x(s)\}$, that has presumably been corrupted by some noise in order to produce $Z = f(X)$.

A patch (noted in this paper in bold face) is defined as a set of pixels $\mathbf{s}(s) = \{l : \|s - l\|_{\mathcal{M}} \le e\}$ in a small neighborhood around a reference pixel s (noted in this paper in regular face), with e the patch extent. The neighborhood norm $\mathcal{M}$ is typically the maximum norm, so the patches become squares of side $2e + 1$ centered around their reference pixel. Let $\lambda = |\mathbf{s}|$ be the number of pixels in a patch.

The probability $p\left(z(\mathbf{s}), z(\mathbf{t})|x(\mathbf{s}) = x(\mathbf{t})\right)$ of observing the (noisy) patches $z(\mathbf{s})$ and $z(\mathbf{t})$, given that they come from the same underlying (clean) patch $x(\mathbf{s}) = x(\mathbf{t})$, can be used as a patch similarity measure. This probability can be calculated if we know the noise process f (additive white noise, multiplicative speckle, camera-specific point-spread function, etc).

The K-nearest patch neighbors are then found using that similarity measure. Assuming all these patches represent a different version of the same underlying clean patch, denoising can be performed by some simple operation on these K patches. For example, in the case of additive noise, non-local means denoising [5] averages all the patches in order to reduce the noise variance. This gives, for each site s, an estimate of the denoised patch $\tilde{\mathbf{s}}(s)$. But due to the patch extent e, each such estimate $\tilde{\mathbf{s}}$ also covers $\lambda - 1$ nearby sites of s. The final denoised image $\tilde{X}$ is computed by aggregating, for each site s, all the corresponding pixels from the set of patches $R = \{\tilde{\mathbf{r}} : s \in \tilde{\mathbf{r}}\}$ covering that site (for example by averaging those pixels).

The next sections describe each of these steps in the case of Sentinel-1 SAR imagery, and improvements made compared to previous work in the domain [6].

B. Application to SAR imagery

For SAR imagery, the additive gaussian noise assumption is invalid and [6] introduces a proper patch similarity measure for speckle noise. Assuming fully developped speckle, the noise is multiplicative and can be modeled as $z(s) = x(s)u(s)$, with u the speckle. Noting $a(s) = \sqrt{z(s)}$ the observed signal amplitude, $p(a(s), a(t)|x(s) = x(t))$ can be used as a pixel similarity measure. Averaging over the range of all possible pixel reflectances $\nu \in \mathcal{D}$:

$$p(a(s), a(t)|x(s) = x(t))$$
$$= \int_{\mathcal{D}} p(a(s)|x(s) = \nu) \, p(a(t)|x(s) = \nu) \, p(\nu) d\nu \quad (1)$$

For each site, the L-look amplitude with a SAR speckle can be modeled as:

$$p(a|\nu) = \frac{2L^L}{\Gamma(L)v^L} a^{2L-1} \exp\left(-\frac{L}{v}a^2\right) \quad (2)$$

For Sentinel-1 satellites, $L = 4.4$ equivalent number of looks for GRD images in IW mode [7]. In order to simplify the expression, [6] assumes a uniform distribution of reflectance values $p(\nu)$ over $\nu \in \mathcal{D}$. However, the reflectances are not at all uniformly distributed in the image. Moreover, the problem is ill-posed as the upper limit for the reflectance is arbitrarily set by the sensor resolution. As an example, the distribution of pixels in Fig. 2, bottom-right, is shown in Fig. 1.

Inspired by the form of Eq. 2, this distribution can be empirically fit with a 2-parameter model:

$$\hat{p}(v) = \frac{C}{v^\alpha} \exp\left(-\frac{\beta}{v}\right) \quad (3)$$

with $C = \frac{\beta^{\alpha-1}}{\Gamma(\alpha-1)}$ the normalisation constant.

The maximum likelihood fit of (α, β) is also shown in Fig. 1. Although not perfect, it is clearly more adequate than a uniform distribution to serve as a prior for estimating the clean image. The maximum likelihood estimator for (α, β) is derived in Appendix.
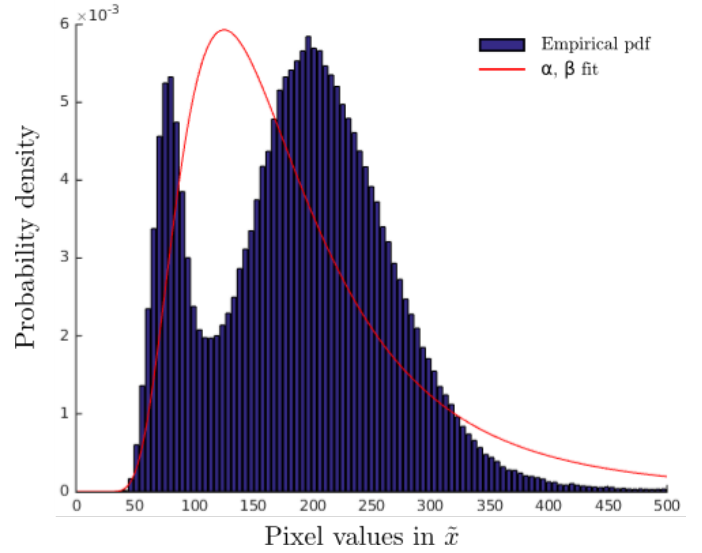

Figure 1. Distribution of pixel values in the region of interest. The histogram is rescaled into an empirical probability distribution. It has two modes. The maximum likelihood fit for α, β in Eq. 3 is overlayed. By definition it has only one mode, but using this fit is still better than using an uniform prior.

Inserting (2) and (3) into (1), the similarity between pixels becomes

$$p[a(s), a(t)|x(s) = x(t)] = \int_0^\infty \frac{A}{v^n} \exp\left(-\frac{B}{v}\right) dv$$

with:

$$n = 2L + \alpha$$

$$A = \frac{\beta^{\alpha-1}}{\Gamma(\alpha-1)} \frac{4L^{2L}}{\Gamma^2(L)} [a(s)a(t)]^{2L-1}$$

$$B = L\left[a(s)^2 + a(t)^2\right] + \beta$$

Using the identity from [6], the pixel similarity becomes:

$$p[a(s), a(t)|x(s) = x(t)] = AB^{1-n}\Gamma(n-1)$$

Assuming that, given the clean pixel values, the noise is independent from site to site, the similarity for whole patches is

$$p(a(\mathbf{s}), a(\mathbf{t})|x(\mathbf{s}) = x(\mathbf{t})) =$$
$$\prod_j p(a(s+j), a(t+j)|x(s+j) = x(t+j))$$

where j is an offset running over all the pixels in the patch. The patch matching search for K nearest neighbors is best done in log space, which does not change the relative ordering. A dissimilarity measure is then given by $d(\mathbf{s}, \mathbf{t}) = -\log p(a(\mathbf{s}), a(\mathbf{t})|x(\mathbf{s}) = x(\mathbf{t}))$:

$$d(\mathbf{s}, \mathbf{t}) = (1 - 2L) \sum_j \log \left[a(s+j) a(t+j) \right]$$

$$+ (2L + \alpha - 1) \sum_j \log \left(L \left[a(s+j)^2 + a(t+j)^2 \right] + \beta \right) + c_{\alpha\beta}(s)$$

$$\tag{4}$$

with

$$c_{\alpha\beta}(s) = -N \log \frac{4 L^{2L}}{\Gamma^2(L)} \frac{\beta^{\alpha-1}}{\Gamma(\alpha-1)} \Gamma(2L + \alpha - 1)$$

The term $c_{\alpha\beta}$ only depends on α and β, but these values are estimated by a maximum likelihood fit at each site s (see the next section). When looking for the K neighbors of patch $\mathbf{s}$, we use the α, β values defined at site s (thus, $d(\mathbf{s}, \mathbf{t}) \neq d(\mathbf{t}, \mathbf{s})$ in general).

C. Denoising and low-rank transfer

At this point, for each site s and associated patch $\mathbf{s}$, a set $\mathcal{N}$ of K nearest neighbors is found using the patchmatch algorithm [10] and the above dissimilarity function. The goal of this section is to build a denoised patch estimate $\tilde{\mathbf{s}}$.

When noise is additive, a simple operation [5] is to average the patches: $\tilde{\mathbf{s}}_{additive} = \langle \mathbf{t} \rangle_{\mathbf{t} \in \mathcal{N}}$. A more elaborated scheme is to work in a sparse representation of $\mathcal{N}$, and perform shrinkage of the values: all small coefficients are discarded, retaining only the largest ones. This is more robust to small errors, that creep in the form of the small coefficients. Working in wavelet space, the shrinkage operation is equivalent to some Wiener filtering as shown by [6]. More recent works [11] improve this approach by working in the singular values space. The logarithm of the pixel values in each patches in $\mathcal{N}$ are stored as the K column vectors of a matrix M. If the patches were exactly the same, that matrix M would be of rank one. Due to the noise, some small non-zero singular values of M appear:

$$M = USV^T$$

with S a diagonal matrix of singular values $(\sigma_1, \ldots \sigma_q)$ of which $\sigma_i \ll \sigma_1 \ \forall i > 1$, with U and V matrices built with the left and right singular vectors, and with $q = \min(\lambda, K)$.

Low-rank denoising [ref from Hicham] amounts to setting all singular values of M to zero except the largest one, and then reconstructing $\tilde{M} = U \begin{pmatrix} \sigma_1 & \\ & 0 \end{pmatrix} V^T$. This operation forces $\tilde{M}$ to be of rank one. The clean estimate patches are then extracted from $\tilde{M}$ and aggregated (see the next Section).

The rank-1 assumption is however unrealistic. Comparing patches with $p\left(z(\mathbf{s}), z(\mathbf{t}) | x(\mathbf{s}) = x(\mathbf{t}) \right)$ induces type-II statistical errors: clean patches $x(\mathbf{s})$ and $x(\mathbf{t})$ may differ at locations $\mathbf{s}$ and $\mathbf{t}$, but the noisy patch similarity may be wrongly accepted. Hence, some natural variability in the set of K nearest neighbors is observed, even in the case of a clean image, so long as a statistical patch comparison is used for finding similar patches. The matrix M is thus not of rank 1, even for clean images. The idea of low-rank transfer is to quantify that natural variability and preserve it while denoising. A reference clean image is used and sets of K nearest neighbors are built with the same method. The singular values are computed on these clean sets, yielding an expected distribution of singular values $(\sigma_1^c, \ldots \sigma_q^c)$ in the clean case. Then, denoising is perform by setting the values of $\tilde{M}$ according to that distribution instead of truncating all values to 0 except the first. Since clean SAR data are by definition not available, we use infrared data from optical Sentinel-2 images as a reference.

In this paper, the average clean distribution is represented as a set of ratios $(\gamma_1, \ldots \gamma_q) = \left(1, \langle \sigma_2^c / \sigma_1^c \rangle, \ldots \langle \sigma_q^c / \sigma_1^c \rangle \right)$, where averages are computed over the reference clean image sites. Reconstruction is then performed with $\tilde{M} = U \begin{pmatrix} s_1 & \\ & \gamma_i s_1 \end{pmatrix} V^T$, scaling the ratios to the leading singular value in order to preserve the patch reflectance. This method is similar to [12] in that it reweights the singular values, albeit with a different weighting scheme.

The ratios can be easily stored in a lookup table, allowing to perform denoising even when the user does not provide a clean image. In this paper, the reference region is around Bordeaux, France, comprising fields, coastal areas, urban areas, rivers, forests, etc. Even in the case of denoising a region with a completely different landscape (tropical area, see Section III), accounting for some variability in the singular values is still better than truncating them to 0 (results shown in Fig. 2, middle-right).

A refinement of the method [9] is to use separate distribution of singular values for each land occupation type. In practice, the effects of this refinement are small, yet not null. In this paper, the moments of a patch $\mathbf{s}$ are used as a characteristic descriptor of that patch. A kernel ridge regressor is trained on the clean image to predict, from the moment vectors, the distribution of singular value ratios. During denoising, the patch moments of a first denoised image estimate $\tilde{y}$ are computed for each location s (see Section II-E). The trained regressor is used for producing a refined distribution of singular value ratios, per site, hence accounting for the land type variability. These ratios are used in the second pass for refining the denoised result (Section II-E, Step 7). In addition to refining the singular value ratios $(\gamma_1, \ldots \gamma_q)$, the K nearest neighbors sets $\mathcal{N}$ of the first denoised estimate $\tilde{y}$ are also used to fit a couple (α, β) for each site, in order to refine the patch comparison for the next denoising stage (see the previous section).

D. Patch aggregation

For each patch $\mathbf{s}$ and similar patch $\mathbf{t}$, $d(\mathbf{s}, \mathbf{t})$ can be seen as a proxy for how confident we are that $x(\mathbf{s}) = x(\mathbf{t})$, i.e. that the patch $\mathbf{t}$ actually represents the same underlying clean image as the patch $\mathbf{s}$. The set $\mathcal{N}_s$ of the K nearest patch neighbors of $\mathbf{s}$ is thus augmented with a weight w_k reflecting that confidence : $\mathcal{N}_s = \{(\mathbf{t}_k, w_k)\}_{k \leq K}$. We set $w_k = \exp\left(d(\mathbf{s}, \mathbf{s}) - d(\mathbf{s}, \mathbf{t}_k) \right)$, so that the largest weight 1 is attributed to $\mathbf{s}$ itself and all other weights are decreasingly lower with k.

For each site s, consider the set of patches $R = \{\mathbf{r} : s \in \mathbf{r}\}$ covering that site. Each of these $\mathbf{r}$ patches is at least in the neighbor set $\mathcal{N}_r$ of itself, with weight 1, but may also belong to an arbitrarily large number U of other neighbor sets $\mathbf{r} \in \mathcal{N}_u$, with associated weights $w_{u \leq U}$. Each of these lead to a denoised estimate of the pixel $\tilde{x}_u(s)$, which we also associate with weight w_u.

Due to the expected distribution of observed noisy pixels (Eq. 2), we aggregate the U patches using a weighted geometric average: $\tilde{x}(s) = \exp\left[\left(\sum_u w_u \log \tilde{x}_u(s)\right) / \left(\sum_u w_u\right)\right]$.

E. Summary of the algorithm

The full algorithm is iterative, similar to [6] and bayesian patchmatch denoising [13]. Extra steps are needed for fitting α, β and applying the low-rank transfer functions. The full algorithm can be summarized by :

Step 1: Find an initial set of K-nearest neighors using the simple patch similarity from [13]. This step is the same as in SAR-BM3D, using a uniform prior for $p(\nu)$, which yields $d(\mathbf{s}, \mathbf{t}) = (2L-1) \sum_j \log\left[\frac{a(s+j)}{a(t+j)} + \frac{a(t+j)}{a(s+j)}\right]$.

Step 2: Do a low-rank shrinkage in SVD space, using the average singular value ratios (see Section II-C).

Step 3: Reconstruct a first estimate of the clean image $\tilde{y}$ from the patches (aggregation phase, Section II-D).

Step 4: Using the same nearest neighors locations, use K patches from $\tilde{y}$ to fit of couple (α, β) at each site.

Step 5: Recompute the set of K-nearest neighbors using the full form patch comparison function in Eq. 4.

Step 6: Use the moments of each patch in $\tilde{y}$ in order to refine the singular value ratios for each site, using the transfer functions learned in Section II-C.

Step 7: Do a low-rank shrinkage in SVD space, using the refined ratios at each site.

Step 8: Reconstruct the final estimate $\tilde{x}$ of the clean image from the patches (aggregation phase, Section II-D).

Optionally, loop to Step 4 until a desired number of iterations is reached. In practice, the algorithm converges very fast and the second iteration does not bring any noticeable difference.

III. RESULTS

In order to demonstrate the effect of the algorithm, a tropical region was selected containing a mixture of urban areas, fields, forests, river banks and mountains. The selected region covers some suburbs of Guayaquil, Equador, and was captured by Sentinel-1 on 2015, July 23. This GRD image contains a single band of VV polarisation, in satellite referential, and can be freely downloaded from ESA's SciHub service [1]. The raw (noisy) amplitude image of the region of interest is shown in Fig. 2. Results of the Sigma-Lee filtering, the SAR-BM3D algorithm and of the proposed method are shown for comparison.

The amount of details is controlled by number of retained singular values. When only 1 singular value is retained, only shrinkage is performed and the result (Fig. 2 middle-right) has lost some of the fine details. Smooth regions appear blurred, yet the borders (e.g. the river banks) are crisp. This method could be used when smoothing the image is desired, for example for enhancing the spatial consistency of pixel-wise classifiers. The result is much less blurry than the widely used Sigma-Lee filter, shown above. Low-rank transfer with 5 singular values (bottom-left) restores the details. Compared to the SAR-BM3D result (middle-left), there is a net gain of spatial consistency, with less intensity variations in zones of uniform nature. Adding more singular values restores the fine structures, but at the risk of also restoring some intensity variations due to noise. The result with 10 singular values is shown in bottom-right of Fig. 2. Adding more singular values do not visually change the results. $K = 25$ nearest patch neighors were used in all pictures. Images using 50 instead of 25 neighbors with the same number of singular values are visually undistintinguishable. Fig. 3 shows the denoising of a second region of interest with urban, coastal, forest and agriculture zones. Noise is successfully suppressed while details are preserved, which is especially apparent in the watershed zone.

IV. DISCUSSION

The results presented above are only applicable to each channel individually. The full-polarimetric equivalent of Eq. 2 has been associated with a patch matching algorithm in [3] with some success. The same improvements over SAR-BM3D [6] that were presented in this paper could be also incorporated in this full-polarimetric context. Conversely [3] uses supposedly homogenous regions in the image to build empirical noise statistics. These, in turn, are used in order to improve the patch matching statistical test $p\left[a(s), a(t)|x(s) = x(t)\right]$. These ideas could also be incorporated to the algorithm presented above in future works.

APPENDIX: MAXIMUM LIKELIHOOD ESTIMATOR FOR α, β

We assume in Section II-B that the distribution of pixel values takes the form:

$$p(v) = \frac{\beta^{\alpha-1}}{v^\alpha \Gamma(\alpha-1)} \exp\left(-\frac{\beta}{v}\right)$$

We have a set of observations v_i, $i = 1 \ldots N$, gathered from a set of similar patches in Step 4 of the algorithm (see Section II-E). The joint patch density is $p(v_1, \ldots v_N | \alpha, \beta)$. The likelihood is the same function, but the v_i are fixed and the α, β are the unknowns. We take the log likelihood for maximization: $L(\alpha, \beta; v_1, \ldots v_N) = \log p(v_1, \ldots v_N | \alpha, \beta)$. Assuming pixel independance:

$$L(\alpha, \beta; v_1, \ldots v_N) = \log \prod_i p(v_i | \alpha, \beta) = \sum_i \log p(v_i | \alpha, \beta)$$

$$L(\alpha, \beta; v_1, \ldots v_N) = \sum_i \log\left(\frac{\beta^{\alpha-1}}{v_i^\alpha \Gamma(\alpha-1)} \exp\left(-\frac{\beta}{v_i}\right)\right)$$

Figure 2. Denoising with the proposed method. Top-left: original amplitude signal. Top-right: Sigma-Lee filtering. Middle-Left: SAR-BM3D. Middle-Right: Pure shrinkage in SVD space (only one singular value retained). Bottom-left: Low-rank transfer with 5 singular values. Bottom-right: Low-rank transfer with 10 singular values.
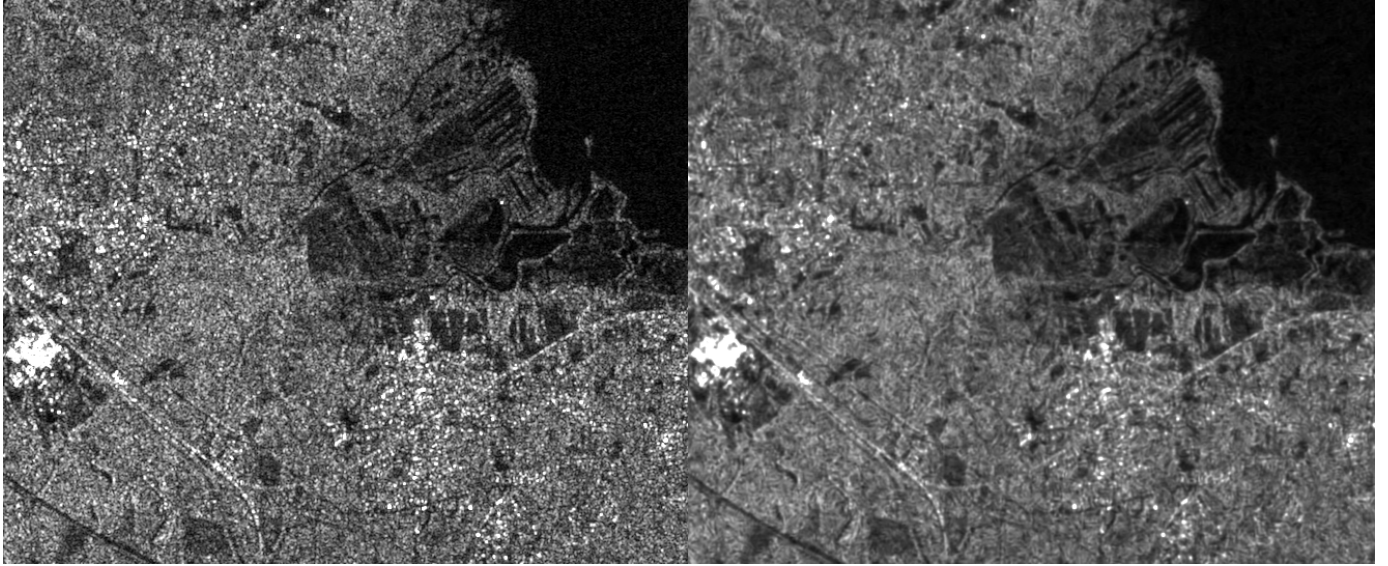
Figure 3. Leyre watershed, Sentinel 1A, 2018/09/27, in the original satellite reference frame. Left: original amplitude signal. Right: Low-rank transfer with 5 singular values and 25 neighbors.

The independance assumption is not true, but for the purpose of this algorithm we just seek an imperfect fit that is still better than a uniform prior (see Section II-B and Fig. 1).

$$L(\alpha, \beta; v_1, \ldots v_N) = N \log \frac{\beta^{\alpha-1}}{\Gamma(\alpha-1)} - \alpha \sum_i \log v_i - \beta \sum_i \frac{1}{v_i}$$

$$= N(\alpha-1) \log \beta - N \log \Gamma(\alpha-1) - \alpha \sum_i \log v_i - \beta \sum_i \frac{1}{v_i}$$

Seeking the maximum likelihood $\frac{\partial L}{\partial \beta} = 0$:

$$\frac{N(\alpha-1)}{\beta} - \sum_i \frac{1}{v_i} = 0 \implies \beta = \frac{N(\alpha-1)}{\sum_i \frac{1}{v_i}} \quad (5)$$

For $\frac{\partial L}{\partial \alpha} = 0$, we get:

$$N \log \beta - \sum_i \log v_i - N\psi(\alpha-1) = 0$$

$$\psi(\alpha-1) = \log \beta - \frac{1}{N} \sum_i \log v_i$$

With ψ the digamma function. Plugin the solution for β:

$$\psi(\alpha-1) = \log \frac{N(\alpha-1)}{\sum_i \frac{1}{v_i}} - \frac{1}{N} \sum_i \log v_i$$

$$\psi(\alpha-1) = \log(\alpha-1) - \log \frac{1}{N} \sum_i \frac{1}{v_i} - \frac{1}{N} \sum_i \log v_i$$

$$\log(\alpha-1) - \psi(\alpha-1) = \log \frac{1}{N} \sum_i \frac{1}{v_i} + \frac{1}{N} \sum_i \log v_i$$

$\log(\alpha-1) - \psi(\alpha-1)$ is strictly decreasing on $]1 \ldots \infty[$ so a solution α can be easily computed or quickly estimated by a lookup table. Once α is known, β is given by Eq. 5.

REFERENCES

[1] https://scihub.copernicus.eu/

[2] J. Lee, M. Grunes, and G. De Grandi, "Polarimetric SAR speckle filtering and its implication for classification," IEEE Transactions on Geoscience and Remote Sensing 37(5):2, pp. 2363–2373, 1999.

[3] C.-A. Deledalle, L. Denis, F. Tupin, A. Reigber, M. Jäger, "NL-SAR: a unified Non-Local framework for resolution-preserving (Pol)(In)SAR denoising". IEEE Transactions on Geoscience and Remote Sensing, 53(4), pp.2021-2038, 2015

[4] P. Wang, H. Zhang and V. M. Patel, "SAR Image Despeckling Using a Convolutional Neural Network". IEEE Signal Processing Letters 24(12), pp. 1763-1767, 2017

[5] A. , C. Bartomeu, J-M. Morel, "A non-local algorithm for image denoising". Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, vol 2, 2005.

[6] S. Parrilli, M. Poderico, C.V. Angelino, L. Verdoliva: "A nonlocal SAR image denoising algorithm based on LLMMSE wavelet shrinkage". IEEE Transactions on Geoscience and Remote Sensing 50(2), pp. 606-616, 2012

[7] Sentinel-1 product definition, ESA reference S1-RS-MDA-52-7440, p7-10, 2016.

[8] A. Lucas, O. Aharonson, C.-A. Deledalle, A.G. Hayes, R. Kirk, E. Howington-Kraus, "Insights into Titan's geology and hydrology based on enhanced image processing of Cassini RADAR data". Journal of Geophysical Research Planets 119, pp. 2149-2166, 2014.

[9] H. Badri, H. Yahia, D. Aboutajdine, "Low-Rankness Transfer for Realistic Denoising". IEEE Transactions on Image Processing 25(12) pp. 5768-5779, 2016.

[10] C. Barnes, E. Shechtman, A. Finkelstein, D.B. Goldman, "PatchMatch: A Randomized Correspondence Algorithm for Structural Image Editing". ACM Transactions on Graphics (Proc. SIGGRAPH) 28(3), 2009.

[11] W. Dong, G. Shi, X. Li, "Nonlocal image restoration with bilateral variance estimation: a low-rank approach". IEEE transactions on image processing 22(2) pp.700-711, 2013.

[12] S. Gu, L. Zhang, W. Zuo, X. Feng, "Weighted nuclear norm minimization with application to image denoising". Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2862-2869, 2014

[13] C. Deledalle, L. Denis, F. Tupin, "Iterative weighted maximum likelihood denoising with probabilistic patch-based weights," IEEE Transactions on Image Processing 18(12), pp. 2661–2672, 2009.