



HAL
open science

Just Project! Multi-Channel Despeckling, the Easy Way

Loïc Denis, Emanuele Dalsasso, Florence Tupin

► **To cite this version:**

Loïc Denis, Emanuele Dalsasso, Florence Tupin. Just Project! Multi-Channel Despeckling, the Easy Way. 2024. ujm-04660707

HAL Id: ujm-04660707

<https://ujm.hal.science/ujm-04660707v1>

Preprint submitted on 19 Aug 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Just Project!

Multi-Channel Despeckling, the Easy Way

Loïc Denis, *Senior Member, IEEE*, Emanuele Dalsasso, Florence Tupin, *Senior Member, IEEE*

Abstract—Reducing speckle fluctuations in multi-channel SAR images is essential in many applications of SAR imaging such as polarimetric classification or interferometric height estimation. While single-channel despeckling has widely benefited from the application of deep learning techniques, extensions to multi-channel SAR images are much more challenging.

This paper introduces MuChaPro, a generic framework that exploits existing single-channel despeckling methods. The key idea is to generate numerous single-channel projections, restore these projections, and recombine them into the final multi-channel estimate. This simple approach is shown to be effective in polarimetric and/or interferometric modalities. A special appeal of MuChaPro is the possibility to apply a self-supervised training strategy to learn sensor-specific networks for single-channel despeckling.

Index Terms—SAR polarimetry, SAR interferometry, despeckling, self-supervised learning

I. INTRODUCTION

Synthetic aperture radar (SAR) imaging is an invaluable technique for Earth observation due to its unique ability to see through clouds and its sensitivity to surface roughness and soil moisture. Beyond providing an image of the back-scattered echo intensities, the polarimetric and interferometric modes offer additional information about the scattering mechanisms (single, double, or triple bounces, surface or volume scattering), the topography (elevation, displacement), or the 3D location of scatterers. All these features are achieved thanks to coherent measurement and processing of the SAR signals. Due to the use of coherent illumination, the speckle phenomenon arises: constructive or destructive interferences occur within each resolution cell due to the coherent summation of several back-scattered echoes. Speckle manifests itself in SAR images in the form of strong fluctuations of the intensity. In multi-channel SAR imaging, it corrupts the interferometric phase and the polarimetric covariance matrix, making the analysis of these images and their automatic processing challenging.

Speckle fluctuations can be reduced by averaging pixel values within a small window. To reach a satisfying amount of smoothing, tens of pixels must be averaged, which implies a severe blurring of the image structures as well as notable errors when mixing, within a given window, scatterers with very different backscattering power [1]. More refined filtering

strategies are required to reduce speckle fluctuations while preserving the spatial resolution. Many different approaches have been imagined: some based on the selection of pixels within homogeneous oriented windows [2], clustered by region-growing [3], identified based on patch similarity [4]–[6]; other approaches follow a variational approach [7]–[10]. More recently, deep neural networks have led to very successful despeckling techniques [11]–[13].

The standard way to train a deep neural network is through supervised learning, i.e., by training on pairs formed by a speckled image (provided as input to the network) and the corresponding speckle-free image (the expected output of the network). Generating a training set with such image pairs is challenging, in particular for multi-channel SAR imaging. To obtain the speckle-free image corresponding to an image corrupted by speckle, the easiest way is to start with a speckle-free image and simulate synthetic speckle. Speckle-free images can be obtained by degrading the resolution of very-high-resolution images (i.e., by spatial multi-looking) or by averaging long time series (i.e., by temporal multi-looking). While temporal multi-looking may be applied to SAR polarimetry (PolSAR), it generally does not work in SAR interferometry (InSAR): the availability of time series of interferograms with a constant baseline and negligible coherence between interferograms is generally not a realistic scenario. Simulating speckle while accounting for the spatial correlations due to the SAR transfer function (zero-padding and spectral apodization) and the spatial and temporal decorrelations in InSAR can be challenging. Yet, it is necessary in order to obtain networks that perform well on actual SAR data. The difficulty of producing training sets for supervised learning increases with the dimensionality of multi-channel SAR images, representing a real limit for multi-baseline InSAR, PolInSAR or SAR tomography applications.

To circumvent the difficulty of building training sets for supervised learning, self-supervised approaches have been proposed for SAR despeckling [14]. Self-supervised despeckling strategies share a common principle: the splitting of the noisy observations into two subsets, one provided to the despeckling network, the other only available to evaluate the training loss function. Methods differ in the way this splitting is performed:

- 1) In SAR2SAR [15], images captured at different dates are used: a reference date is processed by the network and compared, during training, to a second noisy date. To account for possible changes occurring between these two dates, a change-compensation step is added.
- 2) Blind-spot approaches [16] consist of designing a specific network architecture such that the receptive field

Loïc Denis is with Université Jean Monnet Saint-Etienne, CNRS, Institut d'Optique Graduate School, Laboratoire Hubert Curien UMR 5516, F-42023, SAINT-ETIENNE, France and also with LTCI, Télécom Paris, Institut Polytechnique de Paris, Palaiseau, France, (e-mail: loic.denis@univ-st-etienne.fr).

Emanuele Dalsasso is with Environmental Computational Science and Earth Observation (ECEO) laboratory, EPFL, 1951 Sion Switzerland.

Florence Tupin is with LTCI, Télécom Paris, Institut Polytechnique de Paris, Palaiseau, France.

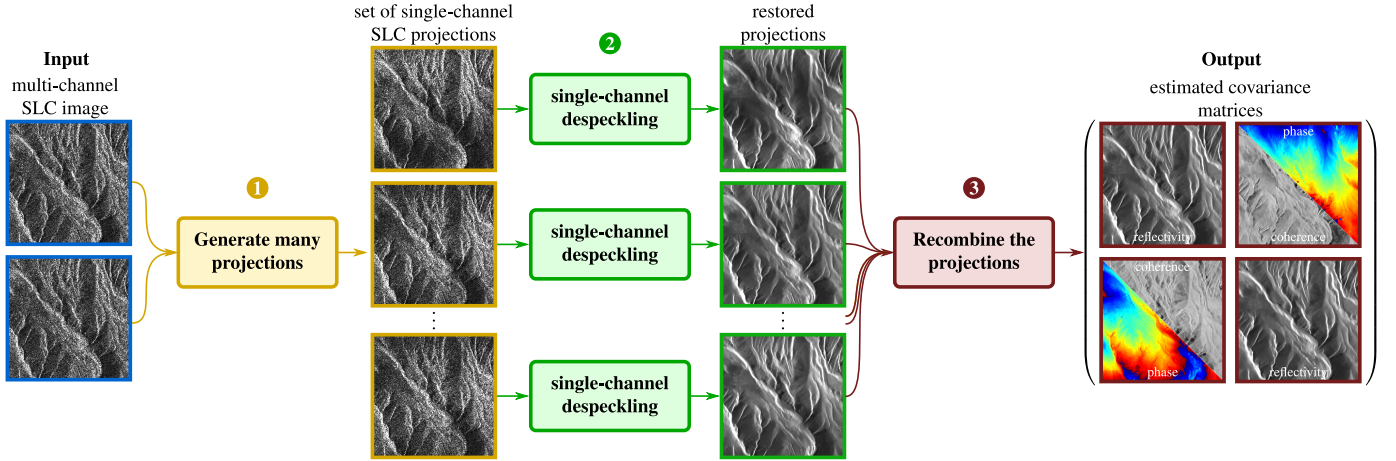


Fig. 1. Scheme of the MuChaPro method introduced in this paper, illustrated here on SAR interferometry: multi-channel SAR images are projected into single-channel SLC images which are, after despeckling, recombined to produce the final estimate.

(i.e., the area in the input image used to predict the value at a given pixel of the output image) excludes the central pixel. The noisy value of that pixel can then be used to evaluate a training loss (the log-likelihood, conditioned by the neighborhood made available to the network). In Speckle2Void [17], Molini *et al.* extend this strategy to SAR despeckling. An alternative is to use more conventional network architectures and apply a preprocessing step to mask out some pixels and use their value to evaluate the training loss, as suggested in [18].

- 3) MERLIN [19] uses the real and imaginary decomposition of a single-look complex SAR image and the property that the speckle in each image of this decomposition is statistically independent of one another (under some assumptions on the SAR transfer function, see [20]).

A radically different approach is the plug-and-play framework MuLoG [9]. Derived from a variational framework and the Alternating Directions Method of Multipliers (ADMM) [21], it consists first of decomposing a multi-channel SAR image into real-valued channels after a matrix-logarithm is applied to the noisy covariance matrices that capture the polarimetric and/or interferometric information at each pixel, then alternately denoising these real-valued channels separately with a conventional graylevel image denoising algorithm (based on a deep neural network or not) and recombining all channels within a non-linear operation. MuLoG presents the advantage of being applicable to SAR images of various dimensionality and to readily include pre-trained networks capable of denoising images corrupted by additive white Gaussian noise. The drawback is the lack of specialization to a specific SAR sensor. In particular, it does not account for the spatial correlations of speckle. Since it is based on the matrix-log decomposition of covariance matrices, some amount of spatial smoothing is necessary, in particular when the dimensionality of the images increases (in multi-baseline interferometry and SAR tomography).

Deep neural networks have been designed specifically to estimate the phase and coherence of an interferometric pair.

The Φ -Net [22] is trained in a supervised way. It applies a decorrelation of the real and imaginary components of the complex interferogram to form the two channels provided as input to the network. InSAR-MONet [23] uses a multi-term loss function to restore the interferometric phase. Digital elevation models and empirical coherence maps were used to generate pairs of simulated noisy/noise-free phase images. The loss combines spatial terms, to ensure that the restored phase is close to the ground-truth phase, and a statistical term to enforce that the estimated noise component follows the expected distribution.

Specific networks have also been proposed for SAR polarimetry, generally following a supervised training methodology. In [24], the same matrix-logarithm transform as in MuLoG is applied first, then complex-valued operations are applied throughout the network (i.e., the real and imaginary parts are not extracted and processed like a two-channels image but rather complex-convolutions, complex-activations, and complex-batch normalizations are applied). Tucker and Potter [25] also apply a log-transform but then perform a real-imaginary decomposition and apply a real-valued residual network to estimate the log-transformed speckle component. A different approach is followed in [26] where a network is trained, in a supervised way, to produce weights leading to a weighted combination of the neighboring pixels as close as possible to a ground truth. By considering pairs of polarimetric images acquired at close dates, a network is trained without reference in [27]. Unlike SAR2SAR [15], no compensation for changes is applied, which could represent a serious limitation in quickly evolving areas (e.g., agricultural fields, changes of soil moisture).

Our contributions: This paper introduces a radically different approach to InSAR and PolSAR despeckling. Although the task requires estimating a multi-dimensional and complex-valued covariance matrix at each pixel, we suggest reducing the problem to a series of single-channel real-valued image despeckling by working on projections, see Figure 1. Such an approach, named “MuChaPro” in the following, presents several advantages:

- 1) it greatly simplifies the application of deep neural networks for the estimation of polarimetric and / or interferometric properties,
- 2) a network trained for a given sensor / resolution can be readily applied to various polarimetric or interferometric configurations,
- 3) it prevents some of the issues occurring with increasing covariance matrix dimensions, in particular, the augmentation of the network complexity required by joint processing of numerous channels and the substantial risk of generalization issues (the spatial polarimetric/interferometric patterns become more diverse and the gap between training and inference steps increases with the dimensionality),
- 4) the proposed approach is generic: it can accommodate any despeckling algorithm, including deep neural networks of various architectures as well as algorithms that do not use neural networks,
- 5) various despeckling algorithms can easily be applied in parallel to compare their outputs and get a better sense of possible artifacts / network hallucinations,
- 6) our approach provides a way to train networks for polarimetric or interferometric despeckling in a self-supervised way.

All those features come at a cost: by reducing the multichannel despeckling problem to a series of independent single-channel despeckling tasks, weakly contrasted geometrical features do not benefit from the reinforcement brought by joint processing. We show in the following that the capacity of most recent self-supervised networks to restore details in single-channel SAR images mitigates this limitation and makes our approach competitive.

II. MUCHAPRO: ESTIMATING COVARIANCE MATRICES FROM PROJECTIONS

A. A generic framework for multi-channel restoration using single-channel despeckling techniques

In order to obtain a method that readily generalizes to multi-channel SAR images with an arbitrary number of channels, we propose to project these images into a set of single-channel images and to despeckle these projections. After despeckling, these single-channel images can be recombined to form the final multi-channel covariance estimates. Figure 1 summarizes the principle of our method. Since **Multi-Channel Projections** are at the core of the framework and we perform many projections (“**mucha proyección**”, in Spanish), we call our method MuChaPro.

To explain the rationale behind this approach, it is necessary to recall Goodman’s fully-developed speckle model [28]. A single-look complex (SLC) multi-channel SAR image \mathbf{z} contains at each pixel ℓ a vector $[\mathbf{z}]_\ell \in \mathbb{C}^D$ of D complex amplitudes. Due to speckle, these complex amplitudes are distributed¹ according to a complex circular Gaussian distribution $\mathcal{N}_c(\mathbf{C}_\ell)$ parameterized by the covariance matrix $\mathbf{C}_\ell \in \mathbb{C}^{D \times D}$

¹we neglect here the spatial correlations of speckle due to the SAR transfer function, in section II-C we discuss the impact of these correlations

that contains the polarimetric and/or interferometric information at pixel ℓ .

Let $\{\mathbf{p}_k\}_{k=1..K}$ be a set of K vectors of \mathbb{C}^D . The complex-valued images $\{\mathbf{s}_k\}_{k=1..K}$ obtained by projection of the multi-channel image \mathbf{z} onto the corresponding vectors \mathbf{p}_k are defined by:

$$\forall \ell, [\mathbf{s}_k]_\ell = \mathbf{p}_k^\dagger \cdot [\mathbf{z}]_\ell, \quad (1)$$

where the notation \dagger indicates the conjugate-transpose operation. The value $[\mathbf{s}_k]_\ell$ is a complex number, i.e., the image \mathbf{s}_k corresponds to a single-channel SLC image. Since the projection $[\mathbf{s}_k]_\ell$ at pixel ℓ corresponds to a linear transform of $[\mathbf{z}]_\ell$, it also follows a complex circular Gaussian distribution, with a variance equal to $\mathbf{p}_k^\dagger \cdot \mathbf{C}_\ell \cdot \mathbf{p}_k \equiv [\mathbf{v}_k]_\ell \in \mathbb{R}^+$. The SLC projections can be seen as regular SAR images, corrupted by single-look speckle, and with an underlying reflectivity \mathbf{v}_k . The application of a single-channel despeckling algorithm separately on each of these images produces estimates $\hat{\mathbf{v}}_k$ of the true variance \mathbf{v}_k .

The aim is to recover at each pixel ℓ the $D \times D$ covariance matrix \mathbf{C}_ℓ . The definition $[\mathbf{v}_k]_\ell = \mathbf{p}_k^\dagger \cdot \mathbf{C}_\ell \cdot \mathbf{p}_k$ shows that the variances $[\mathbf{v}_k]_\ell$ can themselves be seen as projections of the covariance matrices \mathbf{C}_ℓ :

$$\forall k, \forall \ell, [\mathbf{v}_k]_\ell = \mathbf{q}_k^\dagger \cdot [\mathbf{c}]_\ell, \quad (2)$$

with $[\mathbf{c}]_\ell = \text{vec}(\mathbf{C}_\ell)$ the vectorized form of \mathbf{C}_ℓ (i.e., the values of matrix \mathbf{C}_ℓ are rearranged to form a column vector of dimension D^2) and vector $\mathbf{q}_k \in \mathbb{C}^{D^2}$ is defined by $\mathbf{q}_k = \text{vec}(\mathbf{p}_k \cdot \mathbf{p}_k^\dagger)$. In other words, by despeckling the projections \mathbf{z}_k , we recovered *projections of the covariances* \mathbf{C}_ℓ . It remains to invert these projections, a step that is discussed in the next paragraph.

B. Recovering the covariance matrix from its projections

The simplest and fastest way to recover, at each pixel ℓ , the covariance matrix \mathbf{C}_ℓ consists of solving the following linear least-squares problem:

$$\hat{\mathbf{C}}_\ell \in \arg \min_{\mathbf{C}_\ell} \sum_{k=1}^K \left(\mathbf{p}_k^\dagger \cdot \mathbf{C}_\ell \cdot \mathbf{p}_k - [\hat{\mathbf{v}}_k]_\ell \right)^2. \quad (3)$$

Note that the inversion of the projections is independent from one pixel to the next and can thus be conducted in parallel. Using the notations introduced in equation (2), the least-squares problem can be rewritten in a more conventional form:

$$[\hat{\mathbf{c}}]_\ell \in \arg \min_{[\mathbf{c}]_\ell} \|\mathbf{Q}^\dagger \cdot [\mathbf{c}]_\ell - [\hat{\mathbf{v}}]_\ell\|_2^2, \quad (4)$$

with \mathbf{Q} the $D^2 \times K$ matrix with the k -th column equal to \mathbf{q}_k and $[\hat{\mathbf{v}}]_\ell$ the column vector of the K restored reflectivities at pixel ℓ obtained by despeckling the K projections $\{\mathbf{s}_k\}_{k=1..K}$.

Provided that there are at least D^2 linearly independent projection directions \mathbf{p}_k , the matrix $\mathbf{Q} \cdot \mathbf{Q}^\dagger$ is invertible and the least squares solution is unique and given by:

$$\forall \ell, \hat{\mathbf{C}}_\ell^{(\text{MuChaPro})} = \text{vec}^{-1} \left((\mathbf{Q} \cdot \mathbf{Q}^\dagger)^{-1} \cdot \mathbf{Q} \cdot [\hat{\mathbf{v}}]_\ell \right), \quad (5)$$

with vec^{-1} the reshaping operation that transforms a column vector of size D^2 back to a $D \times D$ matrix.

Algorithm 1: MuChaPro, Multi-channel despeckling

- input** : an N -pixels D -channels SLC image provided as a matrix $\mathbf{Z} \in \mathbb{C}^{D \times N}$
- input** : a single-channel despeckling function $f : \mathbb{C}^N \rightarrow \mathbb{R}_+^N$
- output**: the restored covariance matrices stored as an $D^2 \times N$ matrix $\widehat{\mathbf{C}}$ such that the ℓ -th column contains the D^2 values of $\text{vec}(\widehat{\mathbf{C}}_\ell^{(\text{MuChaPro})})$, i.e., the covariance matrix at pixel ℓ in vectorized form.
- Step 1: generate K SLC images by projection**
- 1 generate a matrix $\mathbf{P} \in \mathbb{C}^{D \times K}$ whose columns define K projection directions $\{\mathbf{p}_k\}_{k=1..K}$
 - 2 $\mathbf{S} \leftarrow \mathbf{P}^\dagger \cdot \mathbf{Z}$ (compute the K projections)
- Step 2: despeckle the projections**
- for** $k = 1$ to K **do**
- 3 $\widehat{\mathbf{V}}_{k,\bullet} \leftarrow f(\mathbf{S}_{k,\bullet})$ (despeckle k -th single-channel image)
- end**
- Step 3: recombine restored projections**
- 4 build matrix $\mathbf{Q} \in \mathbb{R}^{D^2 \times K}$ such that $\mathbf{Q}_{\bullet,k} = \text{vec}(\mathbf{P}_{\bullet,k} \mathbf{P}_{\bullet,k}^\dagger)$
 - 5 $\widehat{\mathbf{C}} \leftarrow (\mathbf{Q} \cdot \mathbf{Q}^\dagger)^{-1} \cdot \mathbf{Q}^\dagger \cdot \widehat{\mathbf{V}}$ (recover all covariance matrices)

The algorithm MuChaPro given at the top of page 4 summarizes the 3 steps of the proposed method by using matrix notation to define each step using linear algebra for an efficient implementation in languages such as Python or Matlab.

The formulation of the least-squares problem (4) does not exploit the Hermitian symmetry of the matrix \mathbf{C}_ℓ . Due to this symmetry, there are not D^2 complex unknowns but actually D^2 real unknowns:

$$\mathbf{C}_\ell = \begin{pmatrix} C_{11} & C_{12} & \cdots & C_{1D} \\ C_{12}^* & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & C_{D-1 D} \\ C_{1D}^* & \cdots & C_{D-1 D}^* & C_{DD} \end{pmatrix} \quad (6)$$

so the vector of all real-valued unknowns $[c]_\ell$ at pixel ℓ can be structured as follows:

$$[c]_\ell = \left(C_{11} \cdots C_{DD} \quad \text{Re}(C_{12}) \cdots \text{Re}(C_{D-1 D}) \quad \text{Im}(C_{12}) \cdots \text{Im}(C_{D-1 D}) \right)^t \quad (7)$$

and the corresponding matrix \mathbf{Q} in equation (4) is structured

Algorithm 2: Enforcing positive-definiteness

- input** : a $D \times D$ symmetrical matrix \mathbf{C}
- input** : thermal noise equivalent reflectivity R_{thml}
- input** : maximum coherence value $\rho_{\text{max}} < 1$
- output**: a $D \times D$ positive definite matrix $\widehat{\mathbf{C}}$
- Step 1: enforce reflectivities at least equal to R_{thml}**
- for** $d = 1$ to D **do**
- 1 $[\widehat{\mathbf{C}}]_{d,d} \leftarrow \max(R_{\text{thml}}, [\mathbf{C}]_{d,d})$ (clip low/negative values)
- end**
- Step 2: enforce coherence values below ρ_{max}**
- for** $d_1 = 1$ to D **do**
- for** $d_2 = d_1 + 1$ to D **do**
 - 2 $\rho^{-1} \leftarrow \max\left(\rho_{\text{max}}, \frac{|[\mathbf{C}]_{d_1, d_2}|}{\sqrt{[\widehat{\mathbf{C}}]_{d_1, d_1} [\widehat{\mathbf{C}}]_{d_2, d_2}}}\right)^{-1}$
 - 3 $[\widehat{\mathbf{C}}]_{d_1, d_2} \leftarrow \rho_{\text{max}} \cdot \rho^{-1} \cdot [\mathbf{C}]_{d_1, d_2}$
 - 4 $[\widehat{\mathbf{C}}]_{d_2, d_1} \leftarrow \rho_{\text{max}} \cdot \rho^{-1} \cdot [\mathbf{C}]_{d_2, d_1}$
- end**
- end**

into 3 blocs:

$$\mathbf{Q} = \begin{pmatrix} | & & | \\ | \mathbf{p}_1|^2 & \cdots & | \mathbf{p}_K|^2 \\ | & & | \\ \hline | & & | \\ | 2\text{Re}(\mathbf{U}\mathbf{q}_1) & \cdots & 2\text{Re}(\mathbf{U}\mathbf{q}_K) \\ | & & | \\ \hline | & & | \\ | -2\text{Im}(\mathbf{U}\mathbf{q}_1) & \cdots & -2\text{Im}(\mathbf{U}\mathbf{q}_K) \\ | & & | \end{pmatrix}, \quad (8)$$

where \mathbf{U} is a $D(D-1) \times D^2$ matrix obtained from the identity matrix of size $D^2 \times D^2$ by keeping only the rows corresponding to elements above the diagonal, i.e., corresponding to the upper triangular part of the $D \times D$ matrix $\mathbf{p}_k \cdot \mathbf{p}_k^\dagger$.

The expression of the least-squares solution remains unchanged provided that this new definition for matrix \mathbf{Q} is used. As discussed in paragraph II-D, this parameterization using a real-valued vector \mathbf{c} leads to better-conditioned matrices $\mathbf{Q} \cdot \mathbf{Q}^\dagger$ and more stable estimates. Note that the least squares solution does not guarantee that the reconstructed matrix $\widehat{\mathbf{C}}_\ell^{(\text{MuChaPro})}$ is positive definite. A post-processing step can be applied to clip negative reflectivity estimates and coherence values greater or equal to 1, see algorithm 2.

C. Theoretical properties of MuChaPro

In this paragraph we establish two theoretical results that have strong implications for the application of MuChaPro to multi-channel SAR data. The first one indicates that, for the

class of linear filters, multi-channel filtering and MuChaPro are equivalent. There is no loss of statistical efficiency when processing single-channel projections rather than performing the standard multi-channel linear processing. The second one shows that the real and imaginary parts of the single-look-complex projections are statistically independent and, therefore, lend themselves to the application of the self-supervised learning method MERLIN [19], in contrast to the case of multi-channel SAR images for which the real and imaginary parts are in general correlated [29].

Proposition 1. *Let $\widehat{\mathbf{C}}^{(\text{lin})}$ be a linear multi-channel estimator defined by:*

$$\widehat{\mathbf{C}}^{(\text{lin})} = \sum_{\ell=1}^L w_{\ell} \cdot [\mathbf{z}]_{\ell} \cdot [\mathbf{z}]_{\ell}^{\dagger},$$

where $[\mathbf{z}]_{\ell} \in \mathbb{C}^D$ is the vector of complex amplitudes at pixel ℓ of the multi-channel SLC image \mathbf{z} .

Let $\{\widehat{v}_k^{(\text{lin})}\}_{k=1..K}$ be the corresponding linear single-channel estimators defined by:

$$\forall k, 1 \leq k \leq K, \widehat{v}_k^{(\text{lin})} = \sum_{\ell=1}^L w_{\ell} \cdot |[\mathbf{s}_k]_{\ell}|^2$$

with the same weights w_{ℓ} and $\{\mathbf{s}_k\}_{k=1..K}$ the K projections of \mathbf{z} defined in equation (1).

If the set of projections $\{\mathbf{p}_k\}_{k=1..K}$ is such that the matrix $\mathbf{Q} \cdot \mathbf{Q}^{\dagger}$ in equation (5) is invertible, then the estimators $\widehat{\mathbf{C}}^{(\text{lin})}$ and $\widehat{\mathbf{C}}^{(\text{MuChaPro})}$ applied to the linearly despeckled reflectivities $\{\widehat{v}_k^{(\text{lin})}\}_{k=1..K}$ coincide.

Proof. The expression of the estimator $\widehat{v}_k^{(\text{lin})}$ can be expanded using equation (1):

$$\widehat{v}_k^{(\text{lin})} = \sum_{\ell=1}^L w_{\ell} \cdot |\mathbf{p}_k^{\dagger} \cdot [\mathbf{z}]_{\ell}|^2 = \sum_{\ell=1}^L w_{\ell} \cdot \mathbf{p}_k^{\dagger} \cdot [\mathbf{z}]_{\ell} \cdot [\mathbf{z}]_{\ell}^{\dagger} \cdot \mathbf{p}_k$$

By substituting \mathbf{C} with the expression of the linear multi-channel estimator $\widehat{\mathbf{C}}^{(\text{lin})}$ and $\widehat{v}_k^{(\text{lin})}$ with its expansion, the sum of squares in equation (3) becomes:

$$\begin{aligned} \sum_{k=1}^K \left(\mathbf{p}_k^{\dagger} \cdot \mathbf{C} \cdot \mathbf{p}_k - \widehat{v}_k^{(\text{lin})} \right)^2 = \\ \sum_{k=1}^K \left(\mathbf{p}_k^{\dagger} \cdot \left(\sum_{\ell=1}^L w_{\ell} \cdot [\mathbf{z}]_{\ell} \cdot [\mathbf{z}]_{\ell}^{\dagger} \right) \cdot \mathbf{p}_k \right. \\ \left. - \sum_{\ell=1}^L w_{\ell} \cdot \mathbf{p}_k^{\dagger} \cdot [\mathbf{z}]_{\ell} \cdot [\mathbf{z}]_{\ell}^{\dagger} \cdot \mathbf{p}_k \right)^2 = 0 \end{aligned}$$

By definition (see equation (3)), the estimator $\widehat{\mathbf{C}}^{(\text{MuChaPro})}$ minimizes this sum of squares. When the matrix $\mathbf{Q} \cdot \mathbf{Q}^{\dagger}$ is invertible, this estimator is uniquely defined. Since $\widehat{\mathbf{C}}^{(\text{lin})}$ reaches the lowest possible value for the sum of squares, it corresponds to the least squares estimator, i.e., it coincides with MuChaPro's estimator. \square

Note that, since $\widehat{\mathbf{C}}^{(\text{lin})}$ is linear, weights w_{ℓ} can not depend on the data \mathbf{z} . This result applies to filters like spatial multi-looking (all weights w_{ℓ} are then equal), convolutions

(i.e., shift-invariant weights), spatially-adaptive filtering with weights driven solely by an external image (i.e., guided-filtering, see for example [30] when only the optical image is used to compute the weights) or even given by an oracle. It does not apply, though, to patch-based methods such as [5] that compute weights w_{ℓ} based on the similarity between multi-channel SAR patches.

Proposition 1 implies that, provided that the weights w_{ℓ} match for multi-channel and single-channel filtering, the despeckling result is the same. In a constant region, using identical weights for all pixels ℓ of the region in the estimator $\widehat{\mathbf{C}}^{(\text{lin})}$ leads to the maximum likelihood estimator, which is asymptotically efficient. MuChaPro also is asymptotically efficient in this case and there is no increased variance or bias in the estimation due to processing only projections. Yet, the cost of providing only single-channel images to the despeckling algorithm comes from the impossibility to exploit jointly all channels (i.e., use the co-location of spatial structures in the different channels). A drop of performance is to be expected for non-linear estimators due to the independent processing of each projection. The very good restoration capability of modern single-channel despeckling techniques limits this impact, as shown in section III. We believe that the advantages of MuChaPro outweigh this drawback, in particular the possibility to perform a self-supervised training, as shown by the next proposition:

Proposition 2. *If the SAR system response is real-valued then the projections $\{\mathbf{s}_k\}_{k=1..K}$ defined in equation (1) have statistically independent real and imaginary parts. The self-supervised learning framework MERLIN [19] is then applicable.*

Proof. In section II-A we considered an ideal SAR system response. Yet, in practice, due to zero-padding and spectral apodization, a low-pass filtering correlates the speckle. Let $\mathbf{H} \in \mathbb{R}^{N \times N}$ be the linear operator, in the spatial domain, that models this low-pass response. Assuming that \mathbf{H} is real-valued amounts to a transfer function with Hermitian symmetry for shift-invariant systems. Paragraph II-E discusses how to ensure that \mathbf{H} is real-valued. Let us further assume that this SAR system response is applied identically and separately to all channels of the ideal multi-channel SAR image \mathbf{z} :

$$\forall \ell, [\tilde{\mathbf{z}}]_{\ell} = \sum_{i=1}^N \mathbf{H}_{\ell i} [\mathbf{z}]_i,$$

where $[\tilde{\mathbf{z}}]_{\ell} \in \mathbb{C}^D$ corresponds to the vector of complex amplitudes in the actual SAR image (the image when accounting for the SAR system response, with spatially correlated speckle). The projections $\{\mathbf{s}_k\}_{k=1..K}$ are computed from the multi-channel image $\tilde{\mathbf{z}}$:

$$\forall \ell, [\mathbf{s}_k]_{\ell} = \mathbf{p}_k^{\dagger} \cdot [\tilde{\mathbf{z}}]_{\ell}.$$

Since \mathbf{s}_k is obtained by applying linear transforms to complex circular Gaussian random vectors, the real and imaginary parts $\Re(\mathbf{s}_k)$ and $\Im(\mathbf{s}_k)$ are jointly Gaussian. Therefore, they are independent if and only if $\text{Cov}[\Re([\mathbf{s}_k]_{\ell}), \Im([\mathbf{s}_k]_m)] = 0$ for

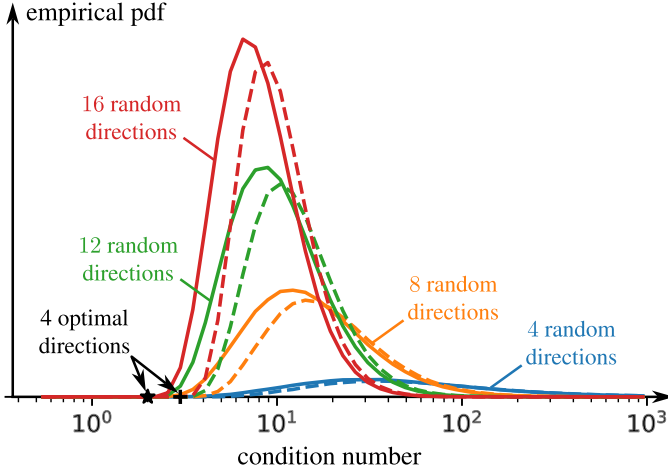


Fig. 2. Empirical probability density function of the condition number of matrix $\mathbf{Q} \cdot \mathbf{Q}^\dagger$ for projection directions \mathbf{p}_k independently drawn according to a standard Gaussian distribution. The case where matrix \mathbf{Q} is built without any constraint on matrix \mathbf{C} is displayed with dashed lines. The case where the Hermitian symmetry is leveraged and matrix \mathbf{Q} is defined by (8) is represented with solid lines. In this experiment, $D = 2$ and K varies between D^2 and $4D^2$.

all ℓ and all m . Since the real and imaginary parts are centered, this covariance is equal to

$$\mathbb{E}[\operatorname{Re}(\mathbf{s}_k)_\ell \cdot \operatorname{Im}(\mathbf{s}_k)_m] = \mathbb{E}\left[\operatorname{Re}\left(\mathbf{p}_k^\dagger \sum_{i=1}^N \mathbf{H}_{\ell i}[\mathbf{z}_i]\right) \cdot \operatorname{Im}\left(\mathbf{p}_k^\dagger \sum_{i=1}^N \mathbf{H}_{m i}[\mathbf{z}_i]\right)\right]$$

Since $\mathbf{H}_{\ell i}$ and $\mathbf{H}_{m i}$ are both real, this expectation is equal to

$$\sum_{i=1}^N \sum_{j=1}^N \mathbf{H}_{\ell i} \mathbf{H}_{m j} \mathbb{E}\left[\operatorname{Re}(\mathbf{p}_k^\dagger[\mathbf{z}_i]) \operatorname{Im}(\mathbf{p}_k^\dagger[\mathbf{z}_j])\right]$$

The D -dimensional random vectors $[\mathbf{z}]_i$ and $[\mathbf{z}]_j$ are independent when $i \neq j$, therefore the expectation is zero in that case. Since $\mathbf{p}_k^\dagger[\mathbf{z}]_\ell \sim \mathcal{N}_c([\mathbf{v}_k]_\ell)$ and $[\mathbf{v}_k]_\ell$ is a scalar, the real and imaginary parts of $\mathbf{p}_k^\dagger[\mathbf{z}]_\ell$ are independent and the expectation is also equal to zero when $i = j$. We have shown that $\operatorname{Re}(\mathbf{s}_k)$ and $\operatorname{Im}(\mathbf{s}_k)$ are jointly Gaussian and decorrelated, they are therefore independent which is the requirement to apply MERLIN self-supervised training, as shown in [19] and further discussed in section II-E. \square

D. Selecting the projections

The directions $\{\mathbf{p}_k\}_{k=1..K}$ onto which the multi-channel image is projected can be chosen arbitrarily provided that the covariance matrices can be recovered by least squares inversion. In particular, a necessary condition for matrix $\mathbf{Q} \cdot \mathbf{Q}^\dagger$ to be invertible is that $K \geq D^2$. The stability of the inversion depends on the condition number of this matrix. The collection of directions should therefore be chosen such that the condition number be as low as possible.

A simple strategy to set the projection directions could be to use random directions. As shown in Fig. 2, random

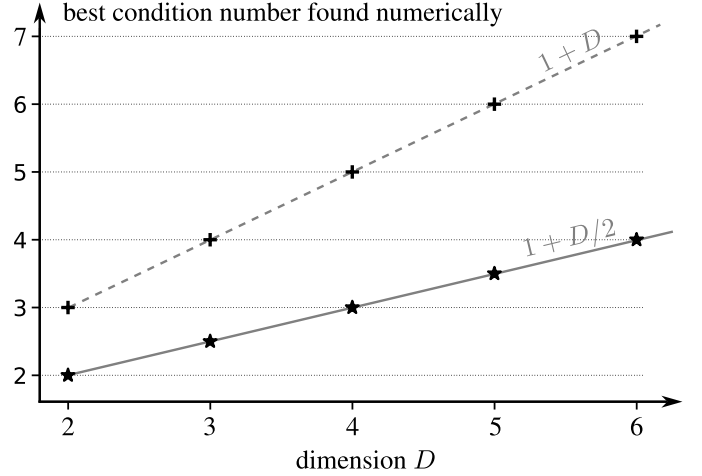


Fig. 3. Evolution of the optimal condition number found numerically, as a function of the dimension D , and expression conjectured from the first 5 points. Solid line: inversion using the Hermitian symmetry of the covariance matrix; dashed line: without symmetry constraint.

directions can lead to very poor condition numbers. The set of random directions leading to the best condition number among one million trials is still about 10 to 15% worse than the best directions identified by optimization when $D = 2$ and two orders of magnitude larger when $D = 6$. Exploiting the Hermitian symmetry offers an improvement of the condition number both for random directions (the distributions are shifted to the left in Fig. 2) and for the optimal directions found by optimization using the technique described below (the best condition number drops from 3 to 2 in the 2-dimensional case illustrated in Fig. 2, the gap between the two approaches increases for higher dimensions, as shown in Fig.3).

Optimizing the condition number of a Gram matrix is a non-smooth and non-convex problem and is thus non-trivial to solve, see [31] for an approach based on the Clarke generalized gradient and an exponential smoothing. We tested several optimization strategies to find directions leading to the best possible condition number. We found experimentally that a quasi-Newton method (L-BFGS [32]) combined with the smoothing approach of Chen *et al.* [31] and Pytorch automatic differentiation was quite efficient when combined with multiple random initializations (at least for the small values of D shown in Fig. 3). The smoothing consists of replacing the condition number of a Hermitian matrix \mathbf{A} with the following approximation

$$\operatorname{cond}_\mu(\mathbf{A}) = -\frac{\log\left(\sum_{i=1}^n e^{\lambda_i(\mathbf{A})/\mu}\right)}{\log\left(\sum_{i=1}^n e^{-\lambda_i(\mathbf{A})/\mu}\right)} \quad (9)$$

where $\lambda_1(\mathbf{A}) \geq \dots \geq \lambda_N(\mathbf{A})$ are the eigenvalues of matrix \mathbf{A} and μ is a smoothing parameter.

Interestingly, while increasing the number of random directions helps to improve the condition number (see the evolution from 4 random directions to 16 random directions in Fig. 2), the optimal condition number found by numerical optimization is reached even for the lowest number of directions, i.e., D^2 . This means that the computational overhead required by processing projections rather than directly handling the multi-

channel images remains limited and is in fact similar to that of some multi-channel approaches (for example, it involves the same number of 2D image denoising operations as a single iteration of MuLoG algorithm [9], [33]).

The best condition numbers obtained by numerical optimization in the cases $D = 1, 2, 3 \dots$ are very close to integer or half-integer values, which is a strong indication that particular structures for the optimal matrices \mathbf{P} and \mathbf{Q} may be found. We could identify such structures in the case without constraint, the more interesting case of the Hermitian symmetry seems more intricate, though, and we could not figure out a simple expression for the optimal projections. We have to rely on the numerical search to identify these projections. Fig. 3 reports the evolution of the best condition number obtained for different dimensions D . We conjecture that these condition numbers follow a simple affine trend: $1+D$ (formulation without constraint) or $1+D/2$ (formulation leveraging the Hermitian symmetry). Note that, while it is numerically more stable to use well-conditioned matrices, it is not mandatory to identify the best overall projection directions for the method to work well.

E. Self-supervised learning for multi-channel despeckling

We introduced in [19], under the name MERLIN, a self-supervised training strategy to learn despeckling networks based on the decomposition of SLC images into real and imaginary parts. Due to coherence between the channels of a polarimetric or interferometric SAR images, the real and imaginary values of multi-channel SAR images are not independent and a direct extension of the MERLIN framework is not possible for these images. Yet, once projected, we have shown in Proposition 2 that the real and imaginary parts of each projection are statistically independent. It is then possible to apply self-supervised learning to these images. In the proof of Proposition 2 we used the fact that the SAR system response was real-valued. If this is not the case, for instance, due to a zero-Doppler shift, then correlations appear and the training does not lead to satisfactory results [20]: speckle fluctuations remain because the network could guess the values of the imaginary component based solely on the values of the real component. SAR images captured by a satellite in stripmap mode require a simple centering of the Doppler centroid. Other modes like spotlight or TOPS and the use of a large bandwidth in airborne imaging require additional processing to ensure that the response is real-valued [20]. The power spectrum density of the SLC image is expected to be symmetrical when the response is real. Spectral shifts, demodulation, and restriction to a symmetrical support can be used to enforce this symmetry.

Once the symmetry of the spectrum has been ensured, a single-channel deep neural network can be trained by feeding the real (or imaginary) component to the network and evaluating the likelihood of the reflectivities estimated by the network with respect to the other component (imaginary or real), as described in [19]. Since projections are similar irrespective of the number of channels, a single network can be trained for a given sensor and then applied to restore polarimetric and

interferometric acquisitions from the same sensor. The ability to vary the number of channels without the need to retrain the network combined with the possibility to apply self-supervised learning represent very appealing features of the MuChaPro framework.

III. APPLICATION TO POLSAR, INSAR, AND MULTI-BASELINE INSAR

In this section, we illustrate the application of MuChaPro to polarimetric, and interferometric SAR images.

We first illustrate the applications of the method in polarimetry and interferometry. Figure 4 illustrates the application of MuChaPro on a TerraSAR-X dual-polarization image (area north of Belleville, Ontario, Canada, captured on June 3, 2016 ©DLR/ESA). Projections along 4 optimal directions (optimal in the sense of the condition number, as discussed in paragraph II-D) are shown in Fig.4(a). Application of the MERLIN despeckling algorithm² leads to the restored images in Fig.4(a). Application of the inversion formulae (5) and (8) give the restored polarimetric image shown in Fig.4(c), bottom row. A qualitative comparison with the multilooked polarimetric estimation shows that polarimetric behaviors are in agreement and that the spatial resolution is better preserved when applying MuChaPro with MERLIN. Note that we have checked numerically that applying multilooking to the polarimetric image gives the exact same result as multilooking the projections first and then inverting the projections to reconstruct the covariance matrices, as stated in Proposition 1.

In figure 5, we illustrate the application of MuChaPro to an interferometric couple of TerraSAR-X Stripmap images captured close to the Jungfrau summit in Switzerland on July 4, 2020, and July 15, 2020 (©DLR/ESA). The same optimal projection directions are used as in Fig.4. The combination of the two images leads to visible fringes in some of the projections, see 5(a) and 5(b). After inversion, the interferometric fringes are recovered. In 5(c), we show the recovered interferogram in the form of a color image with the hue encoding the interferometric phase, the saturation the coherence, and the value the reflectivity of the reference image. In this mountainous area at the confluence of several glaciers, there are low-coherence rapidly moving areas and higher-coherence areas where topographic fringes are recovered. The qualitative comparison of the interferogram obtained by 5×5 multilooking and with MuChaPro shows that the fringes have similar locations and orientations but that thin fringes are better recovered by MuChaPro in conjunction with MERLIN and appear much smoother. Note that to slightly improve the interferogram estimation, we performed a despeckling of each image of the interferometric pair and substituted the reflectivities recovered in MuChaPro's interferogram with those estimates.

In figure 6, we consider another TerraSAR-X Stripmap interferogram over an area of the Grand Canyon, Arizona,

²we used the pretrained network made available by Telecom Paris and Hi! PARIS at <https://github.com/hi-paris/deepdespeckling/>, after a resampling step to obtain the same transfer function as single-polarization TerraSAR-X images used during the training phase of the network

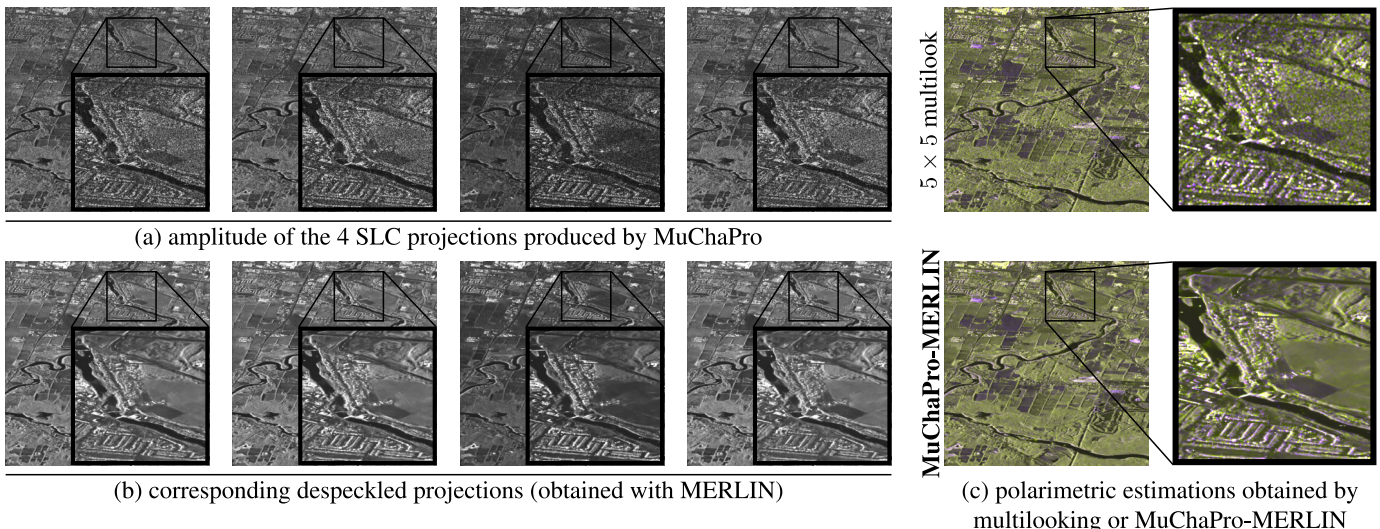


Fig. 4. Application of MuChaPro to a dual-polarization TerraSAR-X image (©ESA/DLR). To estimate the 2×2 polarimetric covariance matrices at each pixel, 4 SLC projections are generated in MuChaPro algorithm (shown in (a)). These projections are despeckled independently using a deep neural network trained beforehand with MERLIN (despeckling results are given in (b)). These despeckled projections are then recombined using Algorithm 1, equation (8), and Algorithm 2 to obtain the estimation shown in (c) bottom row, to be compared with a 5×5 multilooking in the top row.

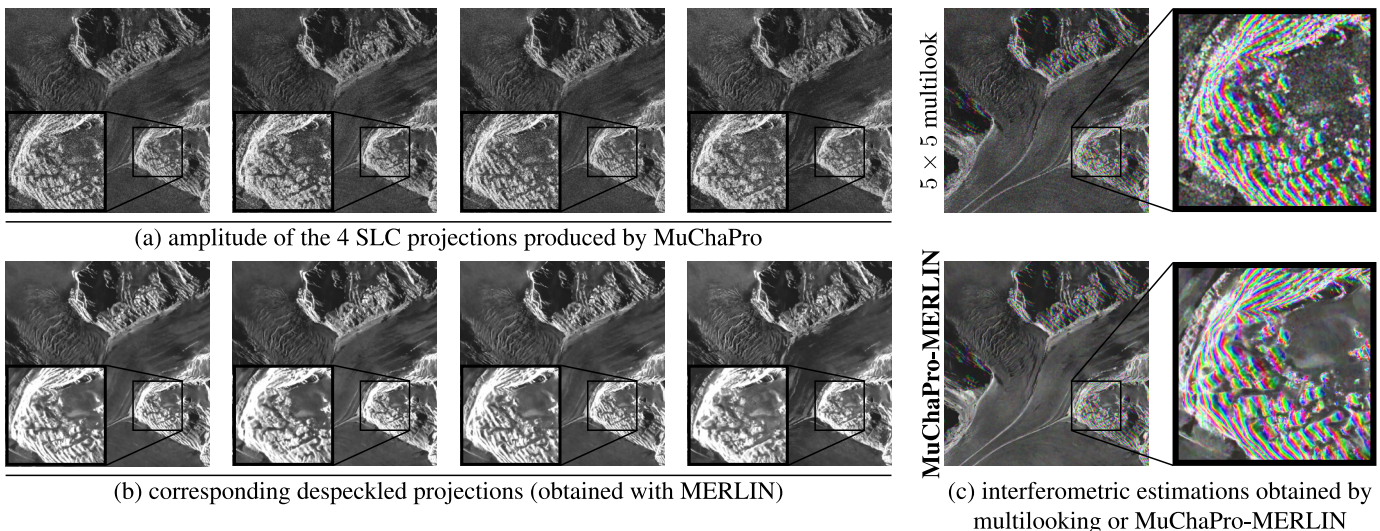


Fig. 5. Application of MuChaPro to an interferometric pair of TerraSAR-X images (©ESA/DLR) of the confluence of glaciers named *Konkordiaplatz*, close to the Jungfrau summit in Switzerland. As in Fig. 4, we show in (a) the noisy projections and in (b) the despeckling results. Combining these restored projections leads to MuChaPro estimates shown in (c). The use of MERLIN despeckling network gives a better preservation of interferometric fringes than by multilooking, with an improved spatial smoothing. In this color composition, the hue corresponds to the interferometric phase, the saturation to the coherence, and the value is given by the restored reflectivity of the reference date.

USA (©Airbus Defence and Space), where the coherence is much larger. We compare the interferograms recovered by several methods: (a) multilooking, (b) multi-channel estimation with MuLoG [9], [33], (c) the application of MuLoG to the projections produced by MuChaPro followed by the inversion of the projections to recover the interferogram, (d) the deep-learning method InSAR-MONET [23] (that only produces an estimate of the interferometric phase), (e) the deep-learning method Φ -NET [22] (that gives both an estimation of the phase and coherence), and finally (f) MuChaPro with MERLIN as a single-channel image despeckling method. Before processing the images with MuLoG, we performed a downsampling by a factor 2 to reduce the spatial correlations of the speckle. The images have been resampled to remove all zero-padding and

reduce speckle correlations before applying InSAR-MONET and Φ -NET, we found this preprocessing step to improve the results with these networks. When applying MERLIN, we did not resample the images since the network was already trained to handle the spatially correlated speckle in TerraSAR-X single-polarization Stripmap images. Figure 7 offers zoomed-in views to analyze the details recovered by each method. The interferometric phases are in agreement with all methods but are much more noisy with multilooking. Interestingly, MuChaPro seems to better transfer the fine spatial details observed in the reflectivity to the coherence and interferometric phase channels.

Finally, we show in Figure 8 that MuChaPro can also be applied to multibaseline interferometry. We processed 3

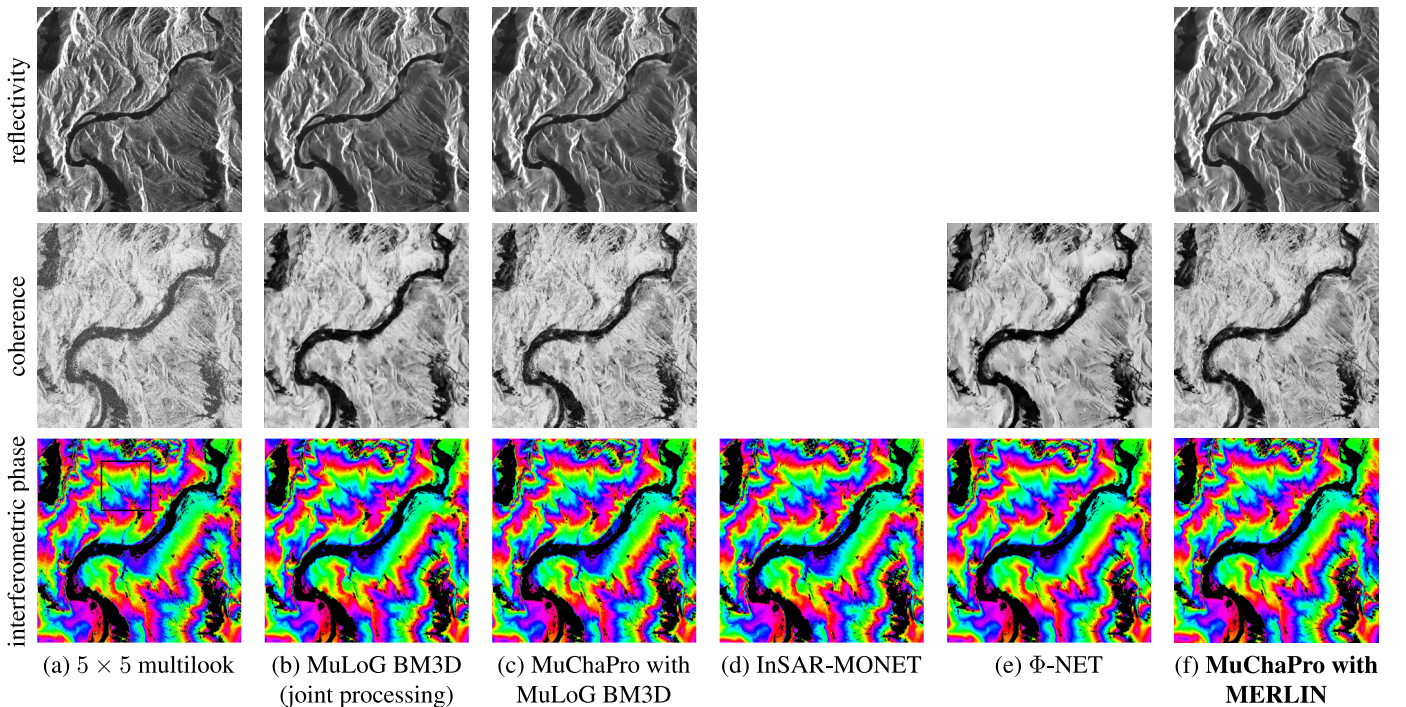


Fig. 6. Computation of an interferogram from a pair of TerraSAR-X images over the Colorado river ©Airbus Defence and Space. Images are best viewed on-screen with a strong magnification. Fig.7 gives a zoomed-in version of the area depicted by a black square in column (a), last row.

TerraSAR-X Stripmap images of the city of Domancy, France (©DLR) captured in September and October 2009. Parts of the images covered by vegetation display a low coherence. This coherence tends to be over-estimated by multi-looking. The interferometric phase estimated by MuChaPro and multilooked are in agreement at a large scale. It is less noisy, in particular in low coherence areas corresponding to fields on the lower-diagonal part of the image.

IV. CONCLUSION

This paper introduced a new generic framework to perform multi-channel SAR despeckling using only single-channel despeckling algorithms. The idea is to produce a set of single-channel SAR images by projecting the multi-channel SAR image onto different directions. Once despeckled, these images give access to projections of the underlying covariance matrices which can be recovered thanks to a linear least-squares estimator.

This approach is equivalent to multi-channel processing when linear filters are applied. It also makes the self-supervised learning strategy MERLIN possible. To the best of our knowledge, this is the first time that a self-supervised training strategy has been proposed for interferometric and polarimetric data. A key advantage of the method is that, once a network is trained for a given sensor, it can be readily applied to multi-channel images of the same sensor whatever their dimensionality: dual-pol, full-pol, InSAR, multi-baseline InSAR, ...

This flexibility comes at a price: by despeckling separately single-channel images, MuChaPro does not benefit from the redundancy present in multi-channel SAR imaging when

restoring weakly contrasted spatial structures. The ability to train networks in a self-supervised way, a possibility that is not (yet) offered by joint despeckling techniques, tends to mitigate this drawback, though.

The application of MuChaPro over different TerraSAR-X images in polarimetric or interferometric configurations confirms the potential of the proposed approach to recover the information with a high spatial resolution. Further work is necessary to perform a more in-depth analysis of the performance and limits of this approach.

ACKNOWLEDGMENT

The authors would like to thank the French National Research Agency (ANR) and the Direction Générale de l'Armement (DGA) under ASTRAL project ANR-21-ASTR-0011 for funding this research. They are grateful to ESA for making available TerraSAR-X images through the TerraSAR-X ESA archive collection, disseminated under ESA's Third Party Missions. The SAR image of Domancy was obtained from the DLR under project LAN-1706. The SAR image of the Colorado is part of sample data freely distributed by Airbus Industry.

REFERENCES

- [1] P. Formont, F. Pascal, G. Vasile, J.-P. Ovarlez, and L. Ferro-Famil, "Statistical classification for heterogeneous polarimetric SAR images," *IEEE Journal of selected topics in Signal Processing*, vol. 5, no. 3, pp. 567–576, 2010.
- [2] J.-S. Lee, M. R. Grunes, and G. De Grandi, "Polarimetric SAR speckle filtering and its implication for classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 37, no. 5, pp. 2363–2373, 1999.

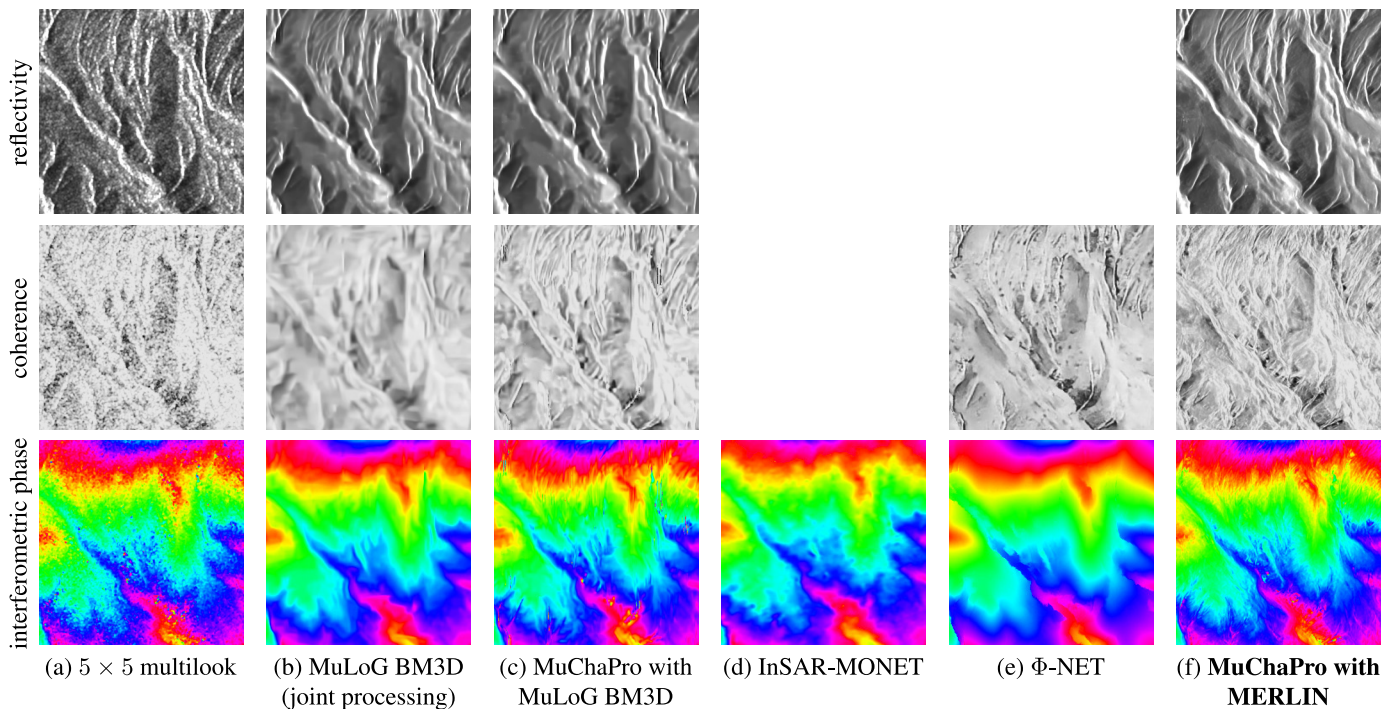


Fig. 7. Close-up view of the area delimited by a black rectangle in Fig.6(a) (SAR image ©Airbus Defence and Space).

- [3] G. Vasile, E. Trouvé, J.-S. Lee, and V. Buzuloiu, “Intensity-driven adaptive-neighborhood technique for polarimetric and interferometric SAR parameters estimation,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 44, no. 6, pp. 1609–1621, 2006.
- [4] J. Chen, Y. Chen, W. An, Y. Cui, and J. Yang, “Nonlocal filtering for polarimetric SAR data: A pretest approach,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 5, pp. 1744–1754, 2010.
- [5] C.-A. Deledalle, L. Denis, F. Tupin, A. Reigber, and M. Jäger, “NL-SAR: A unified nonlocal framework for resolution-preserving (Pol)(In) SAR denoising,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 53, no. 4, pp. 2021–2038, 2014.
- [6] F. Sica, D. Cozzolino, X. X. Zhu, L. Verdoliva, and G. Poggi, “InSAR-BM3D: A nonlocal filter for SAR interferometric phase restoration,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 56, no. 6, pp. 3456–3467, 2018.
- [7] H. Hongxing, J. M. Bioucas-Dias, and V. Katkovnik, “Interferometric phase image estimation via sparse coding in the complex domain,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 53, no. 5, pp. 2587–2602, 2014.
- [8] X. Nie, H. Qiao, B. Zhang, and X. Huang, “A nonlocal TV-based variational method for PolSAR data speckle reduction,” *IEEE Transactions on Image Processing*, vol. 25, no. 6, pp. 2620–2634, 2016.
- [9] C.-A. Deledalle, L. Denis, S. Tabti, and F. Tupin, “MuLoG, or how to apply Gaussian denoisers to multi-channel SAR speckle reduction?” *IEEE Transactions on Image Processing*, vol. 26, no. 9, pp. 4389–4403, 2017.
- [10] C.-A. Deledalle, L. Denis, G. Ferraioli, V. Pascazio, G. Schirinzì, and F. Tupin, “Very-high-resolution and interferometric SAR: Markovian and patch-based non-local mathematical models,” *Mathematical Models for Remote Sensing Image Processing: Models and Methods for the Analysis of 2D Satellite and Aerial Images*, pp. 137–189, 2018.
- [11] G. Fracastoro, E. Magli, G. Poggi, G. Scarpa, D. Valsesia, and L. Verdoliva, “Deep learning methods for synthetic aperture radar image despeckling: An overview of trends and perspectives,” *IEEE Geoscience and Remote Sensing Magazine*, vol. 9, no. 2, pp. 29–51, 2021.
- [12] B. Rasti, Y. Chang, E. Dalsasso, L. Denis, and P. Ghamisi, “Image restoration for remote sensing: Overview and toolbox,” *IEEE Geoscience and Remote Sensing Magazine*, vol. 10, no. 2, pp. 201–230, 2021.
- [13] X. X. Zhu, S. Montazeri, M. Ali, Y. Hua, Y. Wang, L. Mou, Y. Shi, F. Xu, and R. Bamler, “Deep learning meets SAR: Concepts, models, pitfalls, and perspectives,” *IEEE Geoscience and Remote Sensing Magazine*, vol. 9, no. 4, pp. 143–172, 2021.
- [14] E. Dalsasso, L. Denis, M. Muzeau, and F. Tupin, “Self-supervised training strategies for SAR image despeckling with deep neural networks,” in *EUSAR 2022; 14th European Conference on Synthetic Aperture Radar*. VDE, 2022, pp. 1–6.
- [15] E. Dalsasso, L. Denis, and F. Tupin, “SAR2SAR: A semi-supervised despeckling algorithm for SAR images,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 14, pp. 4321–4329, 2021.
- [16] S. Laine, T. Karras, J. Lehtinen, and T. Aila, “High-quality self-supervised deep image denoising,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [17] A. B. Molini, D. Valsesia, G. Fracastoro, and E. Magli, “Speckle2Void: Deep self-supervised SAR despeckling with blind-spot convolutional neural networks,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–17, 2021.
- [18] F. Sica and M. Schmitt, “On the estimation of InSAR phase and coherence through self-supervised learning,” in *EUSAR 2024; 15th European Conference on Synthetic Aperture Radar*, 2024, pp. 736–739.
- [19] E. Dalsasso, L. Denis, and F. Tupin, “As if by magic: self-supervised training of deep despeckling networks with MERLIN,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–13, 2021.
- [20] E. Dalsasso, F. Brigui, L. Denis, R. Abergel, and F. Tupin, “Self-supervised learning of deep despeckling networks with MERLIN: ensuring the statistical independence of the real and imaginary parts,” 2023.
- [21] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein *et al.*, “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [22] F. Sica, G. Gobbi, P. Rizzoli, L. Bruzzone *et al.*, “Φ-Net: Deep Residual Learning for InSAR Parameters Estimation,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 59, no. 5, pp. 3917–3941, 2021.
- [23] S. Vitale, G. Ferraioli, V. Pascazio, and G. Schirinzì, “InSAR-MONet: Interferometric SAR phase denoising using a multiobjective neural network,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–14, 2022.
- [24] A. G. Mullissa, C. Persello, and J. Reiche, “Despeckling Polarimetric SAR Data Using a Multistream Complex-Valued Fully Convolutional Network,” *IEEE Geoscience and Remote Sensing Letters*, vol. 19, pp. 1–5, 2022.
- [25] D. Tucker and L. C. Potter, “Polarimetric SAR Despeckling With Convolutional Neural Networks,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–12, 2022.

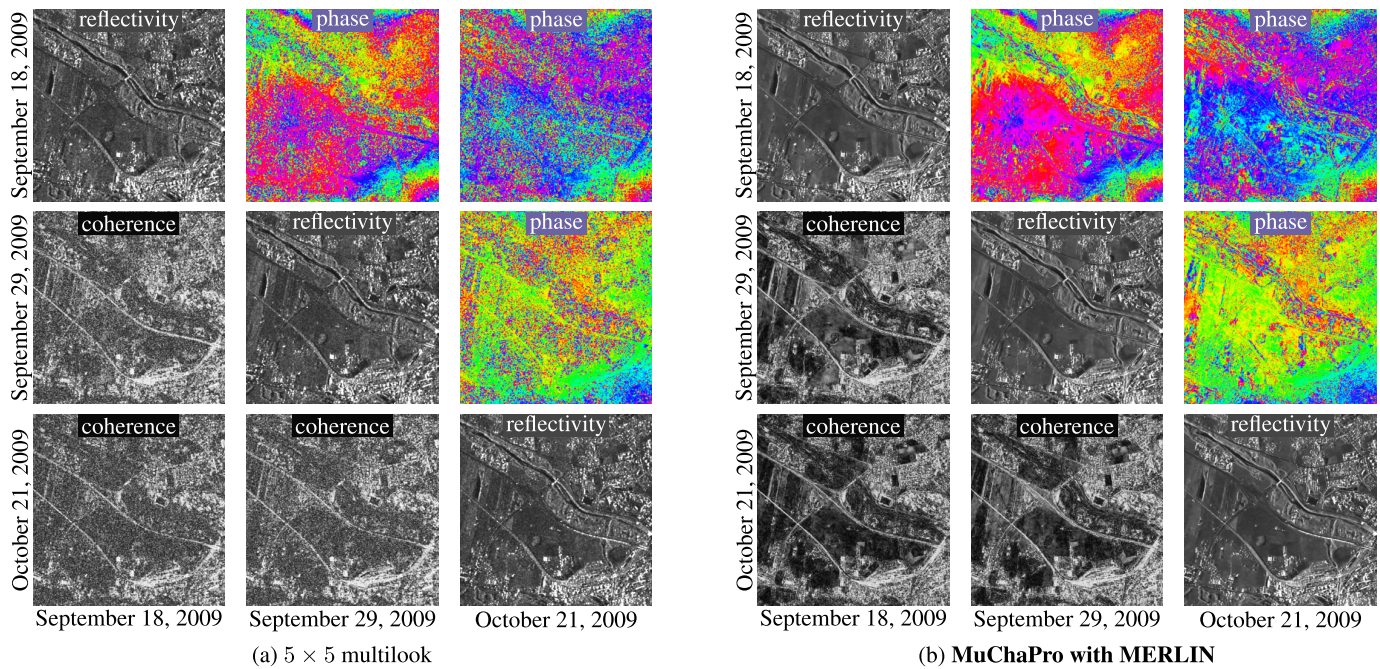


Fig. 8. Application of MuChaPro to multi-baseline interferometry: TerraSAR-X images captured over Domancy, France, at 3 different dates are combined in a multi-baseline interferometric covariance matrix.

- [26] H. Lin, K. Jin, J. Yin, J. Yang, T. Zhang, F. Xu, and Y.-Q. Jin, "Residual in residual scaling networks for polarimetric sar image despeckling," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–17, 2023.
- [27] J. Li, L. Lin, M. He, J. He, Q. Yuan, and H. Shen, "Sentinel-1 Dual-polarization SAR Images Despeckling Network Based on Unsupervised Learning," *IEEE Transactions on Geoscience and Remote Sensing*, 2024.
- [28] J. W. Goodman, *Laser Speckle and Related Phenomena*. Springer Berlin, Heidelberg, 1984, ch. Statistical properties of laser speckle patterns, pp. 9–75.
- [29] L. Denis, "Comment on the paper "PolMERLIN: Self-Supervised Polarimetric Complex SAR Image Despeckling with Masked Networks"," *IEEE Geoscience and Remote Sensing Letters*, pp. 1–1, 2024.
- [30] L. Verdoliva, R. Gaetano, G. Ruello, and G. Poggi, "Optical-driven nonlocal SAR despeckling," *IEEE Geoscience and Remote Sensing Letters*, vol. 12, no. 2, pp. 314–318, 2014.
- [31] X. Chen, R. S. Womersley, and J. J. Ye, "Minimizing the condition number of a Gram matrix," *SIAM Journal on optimization*, vol. 21, no. 1, pp. 127–148, 2011.
- [32] J. Nocedal and S. J. Wright, *Numerical optimization*. Springer, 1999.
- [33] C.-A. Deledalle, L. Denis, and F. Tupin, "Speckle reduction in matrix-log domain for synthetic aperture radar imaging," *Journal of Mathematical Imaging and Vision*, vol. 64, no. 3, pp. 298–320, 2022.