



HAL
open science

Speckle reduction in matrix-log domain for synthetic aperture radar imaging

Charles-Alban A Deledalle, Loïc Denis, Florence Tupin

► **To cite this version:**

Charles-Alban A Deledalle, Loïc Denis, Florence Tupin. Speckle reduction in matrix-log domain for synthetic aperture radar imaging. *Journal of Mathematical Imaging and Vision*, 2022, 64, pp.298-320. 10.1007/s10851-022-01067-1 . ujm-03275572

HAL Id: ujm-03275572

<https://ujm.hal.science/ujm-03275572>

Submitted on 1 Jul 2021

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.

Speckle reduction in matrix-log domain for synthetic aperture radar imaging

Charles-Alban Deledalle · Loïc Denis · Florence Tupin

Received: date / Accepted: date

Abstract Synthetic aperture radar (SAR) images are widely used for Earth observation to complement optical imaging. By combining information on the polarization and the phase shift of the radar echos, SAR images offer high sensitivity to the geometry and materials that compose a scene. This information richness comes with a drawback inherent to all coherent imaging modalities: a strong signal-dependent noise called “speckle”.

This paper addresses the mathematical issues of performing speckle reduction in a transformed domain: the matrix-log domain. Rather than directly estimating noiseless covariance matrices, recasting the denoising problem in terms of the matrix-log of the covariance matrices stabilizes noise fluctuations and makes it possible to apply off-the-shelf denoising algorithms. We refine the method MuLoG by replacing heuristic procedures with exact expressions and improving the estimation strategy. This corrects a bias of the original method and should facilitate and encourage the adaptation of general-purpose processing methods to SAR imaging.

C.A. Deledalle
IMB, CNRS, Univ Bordeaux, Bordeaux INP, F-33405 Talence, France
E-mail: charles-alban.deledalle@math.u-bordeaux.fr

L. Denis
Univ Lyon, UJM-Saint-Etienne, CNRS, Institut d Optique Graduate School, Laboratoire Hubert Curien UMR 5516, F-42023, Saint-Etienne, France
E-mail: loic.denis@univ-st-etienne.fr

F. Tupin
LTCl, Télécom Paris, Institut Polytechnique de Paris, Palaiseau, France
E-mail: florence.tupin@telecom-paris.fr

Acknowledgements

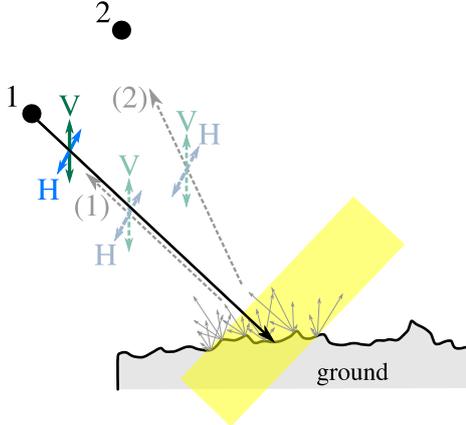
The airborne SAR images processed in this paper were provided by ONERA, the French aerospace lab, within the project ALYS ANR-15-ASTR-0002 funded by the DGA (Direction Générale à l’Armement) and the ANR (Agence Nationale de la Recherche).

1 Introduction

SAR imaging is a key technology in airborne and satellite remote sensing. This active imaging technique based on time-of-flight measurement and coherent processing (the so-called aperture synthesis) has a night and day capability and can produce images through clouds [30]. Beyond intensity images, many SAR systems offer polarimetric information, i.e., they measure how the polarization of the wave is affected by the backscattering mechanisms occurring when the electromagnetic radar wave interacts with the illuminated scene. Backscattered radar waves collected from slightly different points of view can be combined, in a process called interferometry, to perform 3-D reconstructions or to recover very small displacements. In contrast to conventional optical imaging, SAR imaging gives access to both the amplitude *and the phase* of the backscattered wave. Interferometry uses this phase information to relate a phase shift observed between two SAR images to a change in the optical path (i.e., the path length traveled by the wave in its round-trip from the SAR antenna to the scene and back). Figure 1(a) illustrates the geometry of SAR imaging. Depending on the nature of the information available at each pixel, several names are used to describe SAR images:

The physics of SAR imaging:

(a) Geometry of acquisition

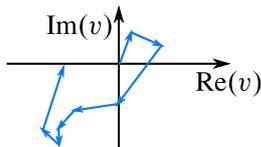


SAR imaging systems use antennas to emit radio waves (with wavelengths ranging from a few centimeters to a meter) and then collect the wave back-scattered by the illuminated scatterers, see (a). When polarimetric information is collected, the radar antenna alternatively emits vertically then horizontally linearly polarized waves. The vertical and horizontal components of the backscattered wave are then measured. In SAR interferometry, the back-scattered wave is collected on two antennas (represented as black dots 1 and 2), either simultaneously (mono-pass interferometry, depicted in the figure) or with a single-antenna system at two different times (multiple-pass interferometry). All the scatterers located within the rectangular region highlighted in the Figure produce radar echoes that are collected simultaneously on the antenna. These echoes interfere and the resulting complex amplitude measured by the antenna is highly fluctuating, depending on the destructive and constructive interferences.

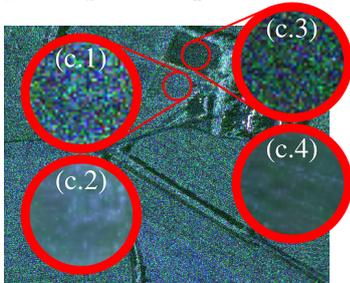
The simplest SAR systems do not record polarimetric or interferometric information, only a single complex amplitude v_k is measured at pixel k . Polarimetric systems collect the back-scattered amplitude in different configurations: for a horizontally linearly polarized emission and horizontally linearly polarized reception (HH), the cross-polarization corresponding to a horizontal emission and vertical reception (HV), or the converse (VH), and for vertically polarized waves in emission and reception (VV). Due to the symmetry of the system, HV and VH measurements are very similar and only one of them is generally considered. A polarimetric SAR image thus contains a vector $\mathbf{v}_k \in \mathbb{C}^3$ at pixel k . SAR interferometry combines echoes collected for two distinct locations of the antenna. Two complex amplitudes are collected per pixel: $\mathbf{v}_k \in \mathbb{C}^2$. Polarimetry and interferometry can be combined, the diffusion vector \mathbf{v}_k then contains 6 complex amplitudes: $\mathbf{v}_k \in \mathbb{C}^6$.

The statistics of speckle:

(b) Coherent summation of radar echoes as a random walk in the complex plane



(c) signal-dependent speckle fluctuations:



The term “speckle” denotes the strong fluctuations of SAR amplitudes from one pixel to the next in homogeneous areas. Due to the coherent illumination, the complex amplitudes of the echoes returned by all the scatterers seen within a given pixel interfere. This corresponds to a random walk in the complex plane, depicted in (b). By the central limit theorem, it follows that, in the limit of a large number of independent elementary scatterers each introducing a uniformly distributed phase shift in $]-\pi, \pi]$, the resulting complex amplitude is distributed according to a circular complex Gaussian [22].

Speckle is a signal-dependent phenomenon: it is stronger in areas with larger reflectivities. This is illustrated in (c) where a polarimetric SAR image is displayed. Regions (c.1) and (c.3) correspond to two homogeneous areas with different reflectivities. Stronger fluctuations can be noted in area (c.1), whose underlying reflectivity is larger (the estimated reflectivity is shown in (c.2)), while region (c.3) displays smaller fluctuations, in line with a lower reflectivity (the estimated reflectivity of this region is shown in (c.4)).

Fig. 1 A short introduction to the physics of SAR imaging and to the speckle phenomenon.

- **SAR** denotes single-channel images: each pixel k contains the complex amplitude $v_k \in \mathbb{C}$ backscattered by the scene. When no interferometric processing involving other single-channel SAR images is to be performed, the phase of v_k can be discarded and only the amplitude $|v_k|$ or the intensity $|v_k|^2$ is considered;
- **PolSAR** denotes multi-channel polarimetric images: each pixel k contains a vector $\mathbf{v}_k = (v_k^1, v_k^2, v_k^3) \in \mathbb{C}^3$ of 3 complex amplitudes backscattered by the scene under different polarizations (horizontally or vertically linearly polarized components at emission or reception, see Figure 1(a));
- **InSAR** denotes multi-channel interferometric images: each pixel k contains a vector $\mathbf{v}_k \in \mathbb{C}^2$ formed by the 2 complex amplitudes backscattered by the scene on the two antenna positions, multi-baseline and tomographic SAR are extensions of SAR interferometry to more than two images;
- **PolInSAR** denotes multi-channel polarimetric and interferometric images: each pixel k contains a vector of 6 complex amplitudes $\mathbf{v}_k \in \mathbb{C}^6$ and corresponds to the combination of polarimetric and interferometric information.

Polarimetric images are generally displayed in false colors using Pauli polarimetric basis, by combining the complex amplitudes collected under the different polarimetric configurations (the red channel corresponds to $\frac{1}{2}|v_k^1 - v_k^3|^2$, the green channel to $2|v_k^2|^2$ and the blue channel to $\frac{1}{2}|v_k^1 + v_k^3|^2$), see [30]. Figure 2 shows airborne images of the same area obtained with each SAR modality (image credit: ONERA). Different structures are visible in the images: cultivated fields, roads (appearing as dark lines in SAR images due to the low reflectivity of such smooth surfaces), trees and several farm buildings (in the bottom left quarter of the image). The SAR illumination comes from the left hand side of the images, shadows are thus visible on the right of all elevated elements (in particular, trees). A striking peculiarity of SAR images is the strong noise observed in all SAR modalities. This noise is unavoidable because it originates from the coherent illumination that is essential to the synthetic aperture processing. The SAR antenna collects at each time sample several echoes that interfere with each other, see Figure 1(b) and (c). The resulting complex amplitude $\mathbf{v}_k \in \mathbb{C}^D$ is distributed, under the well-established speckle model due to Joseph Goodman [22], as a circular complex Gaussian distribution:

$$p_{\mathcal{V}}(\mathbf{v}_k; \boldsymbol{\Sigma}_k) = \frac{1}{\pi^D |\boldsymbol{\Sigma}_k|} \exp(-\mathbf{v}_k^* \boldsymbol{\Sigma}_k^{-1} \mathbf{v}_k), \quad (1)$$

where $*$ denotes the conjugate transpose, and the complex-valued covariance matrix $\boldsymbol{\Sigma}_k$ carries all the

information about the backscattering process: in SAR imaging ($D = 1$), $\boldsymbol{\Sigma}_k$ corresponds to the reflectivity at pixel k , in PolSAR $\boldsymbol{\Sigma}_k \in \mathbb{C}^{3 \times 3}$ characterizes the reflectivity in each polarimetric channel (diagonal of $\boldsymbol{\Sigma}_k$) and the scattering mechanism (matrix $\boldsymbol{\Sigma}_k$ can be decomposed into a sum of matrices corresponding to elementary phenomena such as surface scattering, dihedral scattering and volume scattering), in InSAR $\boldsymbol{\Sigma}_k \in \mathbb{C}^{2 \times 2}$ and diagonal elements correspond to reflectivities while the off-diagonal elements indicate the phase shift $\arg(\Sigma_k^{1,2})$ from one antenna to the other (due to the difference in path length) and the coherence $|\Sigma_k^{1,2}| / \sqrt{\Sigma_k^{1,1} \Sigma_k^{2,2}}$ (i.e., the remaining correlation between the complex amplitudes v_k^1 and v_k^2 : this correlation drops when the two images are captured at two dates that are more separate or when the scene has evolved between the two acquisitions, e.g., due to vegetation growth). Since the covariance matrix $\boldsymbol{\Sigma}_k$ contains all the useful information, it has to be estimated from the diffusion vector \mathbf{v}_k to characterize the radar properties of the scene at each pixel. This is classically done by computing the sample covariance \mathbf{C}_k inside a small window centered at pixel k :

$$\mathbf{C}_k = \frac{1}{L} \sum_{\ell \in \mathcal{N}_k} \mathbf{v}_\ell \mathbf{v}_\ell^*, \quad (2)$$

where \mathcal{N}_k is the set of pixel indices within the window centered at pixel k and $L = \text{Card}(\mathcal{N}_k)$ is the number of pixels in the window. If the speckle is spatially independent and all pixels in the window follow a distribution with a common covariance matrix $\boldsymbol{\Sigma}_k$, the samples \mathbf{v}_ℓ are independent and identically distributed. When $L \geq D$, the sample covariance matrix is then distributed according to a complex Wishart distribution: $\mathbf{C}_k \sim \mathcal{W}(\boldsymbol{\Sigma}_k, L)$, and its multi-variate probability density function is given by [21]

$$p_{\mathcal{C}}(\mathbf{C}_k; \boldsymbol{\Sigma}_k, L) = \frac{L^L |\mathbf{C}_k|^{L-D}}{\Gamma_D(L) |\boldsymbol{\Sigma}_k|^L} \exp(-L \text{tr}(\boldsymbol{\Sigma}_k^{-1} \mathbf{C}_k)), \quad (3)$$

where $\Gamma_D(L) = \pi^{D(D-1)/2} \prod_{k=1}^D \Gamma(L - k + 1)$. In the case of single-channel SAR images ($D = 1$), \mathbf{C}_k corresponds to an intensity I_k and $\boldsymbol{\Sigma}_k$ is the pixel reflectivity R_k . The SAR intensity is then distributed according to a gamma distribution:

$$p_{\mathcal{I}}(I_k; R_k, L) = \frac{L^L I_k^{L-1}}{\Gamma(L) R_k^L} \exp(-L I_k / R_k). \quad (4)$$

The speckle in SAR intensity images is known to be a multiplicative noise in the sense that $\text{Var}[I_k] = R_k^2 / L$, so that the standard deviation of speckle fluctuations is

proportional to the pixel reflectivity, and the speckle-corrupted intensity I_k may be modeled using a generative model of the form: $I_k = R_k S_k$ with S_k a random variable distributed according to a gamma distribution with a mean equal to 1 and a variance equal to $1/L$. In the multi-variate case ($D > 1$), the generative model becomes $\mathbf{C}_k = \boldsymbol{\Sigma}_k^{1/2} \mathbf{S}_k \boldsymbol{\Sigma}_k^{1/2}$ with $\mathbf{S}_k \sim \mathcal{W}(\mathbf{I}, L)$, see [13]. The variance of $C_k^{i,j}$ (the element of matrix \mathbf{C}_k at row i and column j) is equal to $\frac{1}{L} \Sigma_k^{i,i} \Sigma_k^{j,j}$, i.e., the standard deviation of $C_k^{i,j}$ is proportional to the geometric mean of the reflectivities in channels i and j and is thus also signal-dependent.

The ubiquity of speckle noise in SAR images, the multiplicative nature of the fluctuations and heavy-tailed behavior of Wishart distribution has fueled numerous developments of specific image processing methods to reduce speckle. The vast majority of these works considered single-channel SAR images, and adapted techniques based on pixel selection (based on range selection with Lee’s σ -filter [25], window selection [27] or region growing [37]), variational techniques (total variation minimization [1, 18, 35], curvelets [19]), patch-based methods derived from the non-local means [14, 7, 15, 32, 16], or more recently deep-learning based methods (using supervised [8], semi-supervised [11] or self-supervised [28] learning strategies). The adaptation of a novel image denoising technique to the specificities of SAR imaging is a thorough process that includes replacing steps to account for the statistics of speckle and the nature of SAR images where many bright structures reach intensities several orders of magnitude larger than the surrounding area. This slows down the transfer of successful denoising techniques to the field of SAR imaging. In order to circumvent this adaptation process, we recently proposed a generic framework [13], named MuLoG, derived from the general “plug-and-play ADMM” strategy [6] which is related to a wider family of approaches using denoisers to regularize inverse problems [33]. When an image restoration problem is stated in the form of a variational problem and then solved using the alternating directions method of multipliers (ADMM, see for example [3]), or proximal-splitting techniques [9], one step of the algorithm that improves the fidelity to the prior model corresponds to the denoising of an image corrupted by additive white Gaussian noise. The key idea of “plug-and-play ADMM” is then to replace this step by an off-the-shelf Gaussian denoiser.

The flexibility of MuLoG with respect to various SAR modalities (see the despeckling results obtained with MuLoG in Figure 2) and its ability to benefit from the latest developments in additive Gaussian denoising makes the method very useful for SAR appli-

cations (e.g., multi-temporal filtering [42], deformation analysis [20], height retrieval [41], or despeckling using pre-trained neural networks [12]). The original MuLoG algorithm in [13] is based on approximations that can lead, however, to estimation biases. This paper starts with a brief summary of MuLoG framework in Section 2. We then perform a rigorous analysis of the optimization problem involved and establish the exact closed-form expression for the first and second directional derivatives of the matrix exponential mapping. We discuss the important problem of initialization and regularization of the covariance matrices, in particular in the rank-deficient case $L < D$. We introduce several modifications and show that they suppress the bias of the original method. Beyond their use in MuLoG’s generic framework, these mathematical developments can benefit other variational methods for the restoration or segmentation of multi-channel SAR images, as well as hybrid methods that combine deep learning and an explicit statistical model of speckle by algorithm unrolling [29].

2 An overview of MuLoG framework

In order to give a self-contained presentation of our developments, we recall in this section the principle of MuLoG, as first introduced in [13]. MuLoG’s approach to multi-channel SAR despeckling is built around two key ideas:

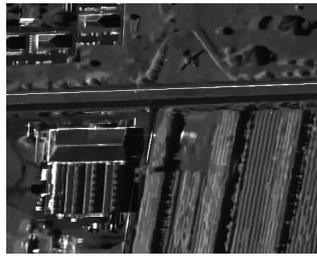
- a non-linear transform that decomposes the field of $D \times D$ noisy covariance matrices $\{\mathbf{C}_k\}_{k=1..n}$ into D^2 real-valued images with n -pixels; this transform approximately stabilizes speckle fluctuations and decorrelates the channels so that each can be denoised independently;
- implicit regularization using a plug-and-play ADMM iterative scheme where the proximal operator associated to the prior term is replaced by an off-the-shelf Gaussian denoiser.

The non-linear transform is defined in three steps: (i) a matrix-log is applied to map each speckled covariance matrix \mathbf{C}_k to a Hermitian matrix with approximately stabilized variance; (ii) the real and imaginary parts of these Hermitian matrices are separated, forming D^2 real-valued channels; (iii) an affine transform, identical for all pixel locations k , whitens these channels. The noisy covariance matrix $\mathbf{C}_k \in \mathbb{C}^{D \times D}$ is then re-parameterized by $\mathbf{y}_k \in \mathbb{R}^{D^2}$, and similarly the covariance matrix of interest $\boldsymbol{\Sigma}_k$ is defined through the real-valued vector $\mathbf{x}_k \in \mathbb{R}^{D^2}$:

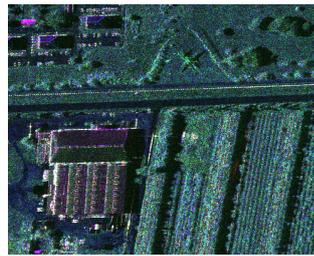
$$\mathbf{C}_k = e^{\Omega(\mathbf{y}_k)} \quad \text{and} \quad \boldsymbol{\Sigma}_k = e^{\Omega(\mathbf{x}_k)} \quad (5)$$

$D = 1$: SAR intensity images

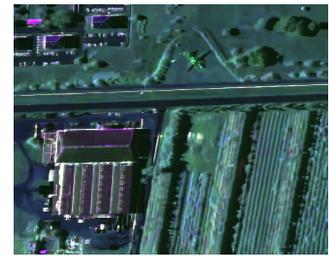

(a.1) noisy intensity



(a.2) restored reflectivity

 $D = 3$: PolSAR polarimetric images


(b.1) noisy polarimetry



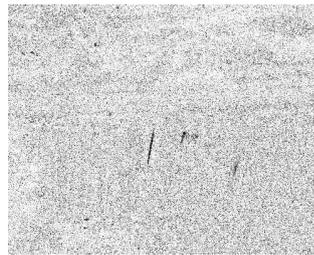
(b.2) restored polarimetry

 $D = 2$: InSAR interferometric images


(c.1) noisy intensity



(c.2) noisy phase



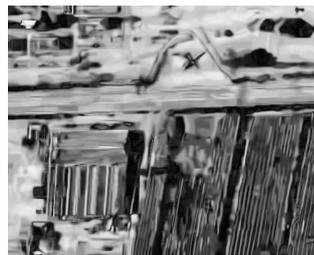
(c.3) noisy coherence



(c.4) restored intensity



(c.5) restored phase



(c.6) restored coherence

 $D = 6$: PolInSAR polarimetric and interferometric images


(d.1) noisy intensity



(d.2) noisy polarimetry



(d.3) noisy phase



(d.4) noisy coherence



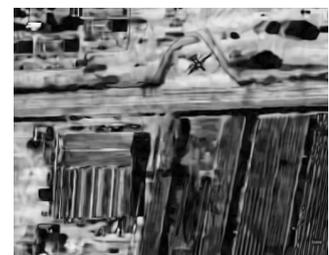
(d.5) restored intensity



(d.6) restored polarimetry



(d.7) restored phase



(d.8) restored coherence

Fig. 2 Synthetic aperture radar imaging offers rich information of a scene but suffers from speckle. The combination of images acquired with slightly different incidence angles (InSAR) or various polarization states (PolSAR) leads to a D -dimensional complex-valued vector per pixel. Reduction of the speckle fluctuations requires appropriate statistical modeling. This paper is devoted to the mathematical analysis of a generic approach for speckle reduction based on matrix-log decompositions. The images shown were obtained with the X-band airborne imaging system SETHI of the French aerospace lab ONERA [2] (after our pre-processing to achieve a trade-off between sidelobe attenuation and speckle decorrelation, the pixel size is $\approx 70\text{cm} \times 70\text{cm}$, the area shown is $\approx 300\text{m} \times 370\text{m}$). The Gaussian denoiser used in the despeckling algorithm is BM3D [10].

where the exponential corresponds to a matrix-exponential and the affine invertible mapping $\Omega: \mathbb{R}^{D^2} \rightarrow \mathbb{C}^{D \times D}$ can be decomposed as $\Omega(\mathbf{x}) = \mathcal{K}(\mathbf{W}\Phi\mathbf{x} + \mathbf{b})$, with $\mathcal{K}: \mathbb{R}^{D^2} \rightarrow \mathbb{C}^{D \times D}$ the linear operator that transforms a vector of D^2 reals into a $D \times D$ Hermitian matrix:

$$\mathcal{K} \begin{pmatrix} \alpha^1 \\ \vdots \\ \alpha^{D^2} \end{pmatrix} = \begin{pmatrix} \alpha^1 & \frac{\alpha^{D+1} + j\alpha^{D+2}}{\sqrt{2}} & \dots & \frac{\alpha^{D^2-1} + j\alpha^{D^2}}{\sqrt{2}} \\ \frac{\alpha^{D+1} - j\alpha^{D+2}}{\sqrt{2}} & \alpha^2 & & \\ \vdots & & \ddots & \vdots \\ \frac{\alpha^{D^2-1} - j\alpha^{D^2}}{\sqrt{2}} & & \dots & \alpha^D \end{pmatrix}, \quad (6)$$

$\mathbf{W} \in \mathbb{R}^{D^2 \times D^2}$ a (whitening) unitary matrix, $\Phi \in \mathbb{R}^{D^2 \times D^2}$ a diagonal positive definite matrix (used to weight each channel) and $\mathbf{b} \in \mathbb{R}^{D^2}$ a real vector (for centering). To compute the real-valued decomposition \mathbf{x} of a covariance matrix Σ , the inverse transform can be applied: $\mathbf{x} = \Phi^{-1}\mathbf{W}^{-1}(\mathcal{K}^{-1}(\log(\Sigma)) - \mathbf{b})$. A principal component analysis is used to compute matrix \mathbf{W} and vector \mathbf{b} :

$$\mathbf{b}_i = \frac{1}{n} \sum_{k=1}^n \alpha_k^i, \quad (7)$$

with $\alpha_k^i = \mathcal{K}^{-1}(\log(\mathbf{C}_k))$, the i -th real-value extracted from the log-transformed covariance $\tilde{\mathbf{C}}_k = \log(\mathbf{C}_k)$, and the columns of matrix \mathbf{W} are unit-norm eigenvectors of the Gram matrix:

$$\frac{1}{n} \sum_{k=1}^n (\alpha_k^i - \mathbf{b}_i)(\alpha_k^j - \mathbf{b}_j). \quad (8)$$

The i -th diagonal element of Φ corresponds to the noise standard deviation of the i -th channel estimated via the median absolute deviation (MAD) estimator.

Figure 3 illustrates the channels obtained with the transform Ω^{-1} . Fluctuations due to the speckle noise have a variance that is approximately stabilized in the channels of \mathbf{y} . Due to the principal component analysis, most of the information is contained in the first channels.

The neg-log-likelihood of the re-parameterized covariance matrix can be derived from the Wishart distribution given in Eq.(3):

$$-\log p_{\mathbf{y}}(\mathbf{y}_k | \mathbf{x}_k) = L \operatorname{tr}[\Omega(\mathbf{x}_k) + e^{-\Omega(\mathbf{x}_k)} e^{\Omega(\mathbf{y}_k)}] + \text{Cst}, \quad (9)$$

where \mathbf{x}_k and \mathbf{y}_k are the D^2 -dimensional real-valued vectors corresponding respectively to the re-parameterization of the noiseless and noisy covariance matrices, while Cst is a constant independent of \mathbf{x}_k .

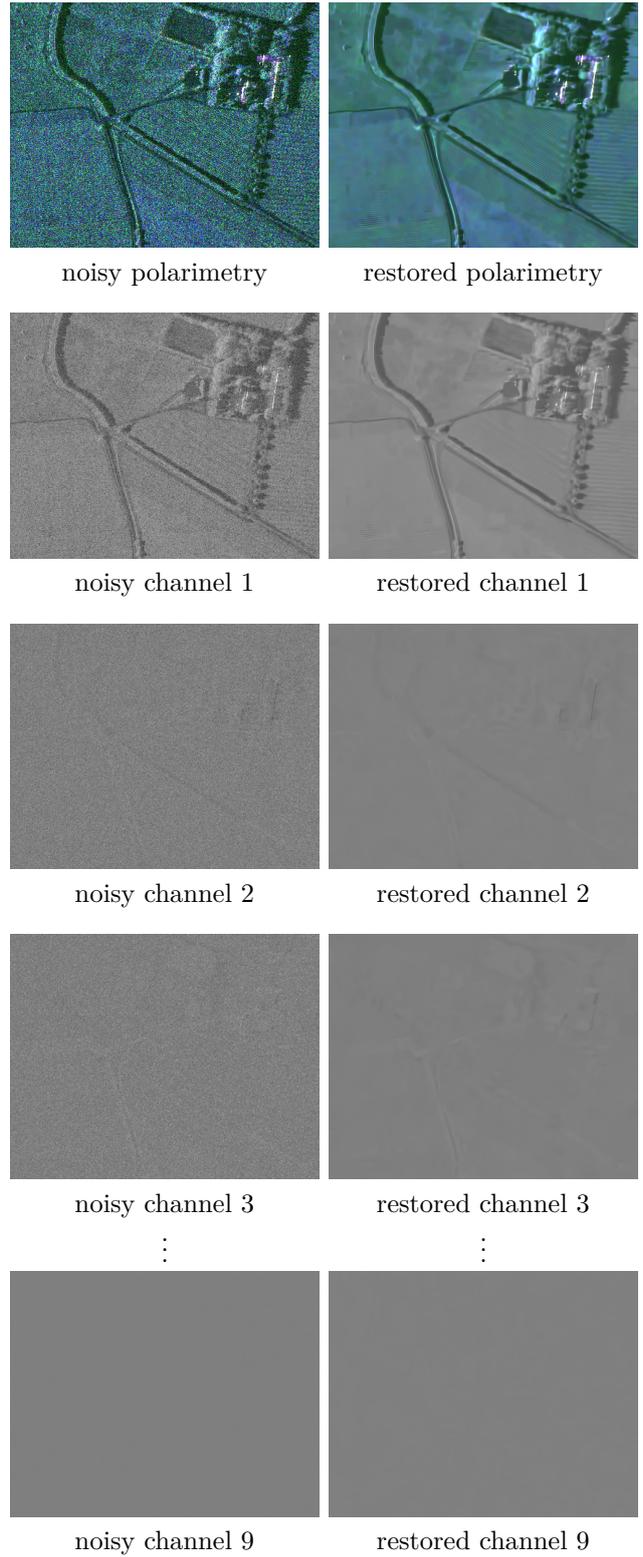


Fig. 3 Illustration of the matrix-log decomposition of a polarimetric SAR image into 9 real-valued channels (airborne image from ONERA-SETHI, see the caption of Figure 2).

MuLoG is detailed in Algorithm 1. In practice, a fixed number of steps is generally used to stop the main loop (typically 6 steps provide a good trade-off between computational cost and restoration quality). Parameter β is increased within the loop, line 11, to ensure the convergence (we use the adaptive update rule given in equation (15) of paper [6] which consists of increasing β whenever the norm of the primal and dual residuals decreases too slowly between two successive steps). Several steps playing an important role in MuLoG algorithm are revisited in this paper:

- the initial estimate used for $\hat{\mathbf{x}}_k$, line 4, computed from the matrices $\{\mathbf{C}_k^{(\text{init})}\}_{k=1..n}$ defined line 1: an adequate initialization speeds up the convergence as discussed in Section 4.1;
- the regularization of the noisy covariance matrices, set line 2, that impacts the minimization problem line 10: regularization strategies can avoid badly behaved cost functions that are hard to minimize but can also lead to estimation biases, see Section 4.2;
- the choice of the denoiser used line 8: as illustrated in the discussion (Section 6), each denoiser suffers from specific artifacts;
- the minimization technique to compute, at each pixel k , the D^2 -dimensional vector $\hat{\mathbf{x}}_k$ line 10: this paper develops in Section 3 a new algorithm based on exact derivatives to improve this crucial step.

3 Improved computation of the data-fidelity proximal operator

Line 10 of Algorithm 1 involves solving n independent D^2 -dimensional minimization problems of an objective function of the form

$$F(\mathbf{x}) = \frac{\beta}{2} \|\mathbf{x} - \mathbf{u}\|^2 + L \text{tr}(\Omega(\mathbf{x}) + e^{\Omega(\mathbf{y})} e^{-\Omega(\mathbf{x})}). \quad (10)$$

This corresponds to computing the data-fidelity proximal operator [9]:

$$\text{prox}_{\text{data}}(\mathbf{u}) = \arg \min_{\mathbf{x}} \underbrace{\frac{\beta}{2} \|\mathbf{x} - \mathbf{u}\|^2 - \log p_{\mathbf{y}}(\mathbf{y}|\mathbf{x})}_{F(\mathbf{x})}, \quad (11)$$

where \mathbf{x} , \mathbf{y} and \mathbf{u} are vectors in \mathbb{R}^{D^2} . We recall that \mathbf{y} corresponds to the noisy log-channels defined in equation (5) and shown in the first column of Figure 3, while $\mathbf{u} = \hat{\mathbf{z}}_k + \hat{\mathbf{d}}_k$ in line 10 of Algorithm 1. Efficient minimization methods require the computation of the gradient $\mathbf{g} = \nabla F(\mathbf{x})$ of F at \mathbf{x} .

Algorithm 1: MuLoG algorithm

input : a bidimensional field of speckle-corrupted covariance matrices \mathbf{C}

input : a denoiser function $f_{\sigma} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ to remove additive white Gaussian noise with standard deviation σ in n -pixels single-channel images

output: a bidimensional field of estimated covariance matrices $\hat{\Sigma}$

initialization:

- 1 estimate $\{\mathbf{C}_k^{(\text{init})}\}_{k=1..n}$ *(initial estimation)*
- 2 estimate $\{\mathbf{C}_k^{(\text{reg})}\}_{k=1..n}$ *(enforces PSD matrices)*
- 3 compute \mathbf{b} and \mathbf{W} *(Eqs.(7) and (8))*
- 4 $\{\hat{\mathbf{x}}_k \leftarrow \Omega^{-1}(\log \mathbf{C}_k^{(\text{init})})\}_{k=1..n}$ *(initial estimation)*
- 5 $\{\mathbf{y}_k \leftarrow \Omega^{-1}(\log \mathbf{C}_k^{(\text{reg})})\}_{k=1..n}$ *(noisy channels)*
- 6 $\{\mathbf{d}_k \leftarrow \mathbf{0}_{D^2}\}_{k=1..n}$ *(initial scaled multipliers)*
- 7 $\beta \leftarrow 1 + \frac{2}{L}$ *(initial value)*

repeat

- 8 $\{\hat{\mathbf{z}}^i \leftarrow f_{\beta^{-1/2}}(\hat{\mathbf{x}}^i - \hat{\mathbf{d}}^i)\}_{i=1..D^2}$ *(denoise channels)*
- 9 $\hat{\mathbf{d}} \leftarrow \hat{\mathbf{d}} + \hat{\mathbf{z}} - \hat{\mathbf{x}}$ *(update multipliers)*
- 10 $\{\hat{\mathbf{x}}_k \leftarrow \underset{\mathbf{x}_k}{\text{argmin}} \frac{\beta}{2} \|\mathbf{x}_k - \hat{\mathbf{z}}_k - \hat{\mathbf{d}}_k\|^2 - \log p_{\mathbf{y}}(\mathbf{y}_k|\mathbf{x}_k)\}_{k=1..n}$ *(data-fidelity prox)*
- 11 update β

until convergence;

- 12 $\{\hat{\Sigma}_k \leftarrow e^{\Omega(\hat{\mathbf{x}}_k)}\}_{k=1..n}$ *(final estimation)*

As shown in [13] and recalled in Appendix A for the sake of completeness, this gradient is given by

$$\mathbf{g} = \beta(\mathbf{x} - \mathbf{z}) + L\Theta \left(\mathbf{Id}_D - \frac{\partial e^{\Sigma}}{\partial \Sigma} \Big|_{-\Omega(\mathbf{x})}^* [e^{\Omega(\mathbf{y})}] \right), \quad (12)$$

where the linear operator Θ is defined by

$$\Theta(\cdot) = \Phi \mathbf{W}^* \mathcal{K}^*(\cdot) \quad (13)$$

and $\frac{\partial e^{\Sigma}}{\partial \Sigma} \Big|_{-\Omega(\mathbf{x})}^* [e^{\Omega(\mathbf{y})}]$ is the adjoint of the directional derivative of the matrix exponential in the direction $e^{\Omega(\mathbf{y})}$ taken at $-\Omega(\mathbf{x})$. We recall that the directional derivative of a differentiable function f at \mathbf{X} in the direction \mathbf{A} is defined as

$$\frac{\partial f(\mathbf{X})}{\partial \mathbf{X}} \Big|_{\mathbf{X}} [\mathbf{A}] = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} (f(\mathbf{X} + \varepsilon \mathbf{A}) - f(\mathbf{X})). \quad (14)$$

The directional derivative is a linear mapping with respect to the direction \mathbf{A} . The computation of the gradient thus requires computing the adjoint of such linear mapping for the matrix exponential.

In the original derivation of MuLoG algorithm [13], we used an integral expression for the adjoint of the directional derivative of the matrix exponential

$$\left. \frac{\partial e^{\Sigma}}{\partial \Sigma} \right|_{\Sigma}^* [\mathbf{A}] = \int_0^1 e^{u\Sigma} \mathbf{A} e^{(1-u)\Sigma} du \quad (15)$$

to derive an approximation based on a Riemann sum:

$$\left. \frac{\partial e^{\Sigma}}{\partial \Sigma} \right|_{\Sigma}^* [\mathbf{A}] \approx \frac{1}{Q} \sum_{q=1}^Q e^{u_q \Sigma} \mathbf{A} e^{(1-u_q)\Sigma} \quad (16)$$

where $u_q = (q - 1/2)/Q$. In practice, we were using only $Q = 1$ rectangle to get a fast-enough approximation. In the next paragraph, we derive a closed-form expression of the gradient to avoid this approximation, then we obtain the closed-form expression of second-order directional derivative of F in order to obtain an improved minimization method for the computation of $\text{prox}_{\text{data}}$ in equation (11).

3.1 Closed-form expression of the gradient

We leverage studies from [26, 17, 5] regarding the derivatives of matrix spectral functions. As shown in Appendix B, these leads to a closed-form expression for this directional derivative of the matrix exponential as described in the next proposition.

Proposition 1 *Let Σ be a $D \times D$ Hermitian matrix with distinct eigenvalues. Let $\Sigma = \mathbf{E} \text{diag}(\mathbf{A}) \mathbf{E}^*$ be its eigendecomposition where \mathbf{E} is a unitary matrix of eigenvectors and $\mathbf{A} = (\lambda_1, \dots, \lambda_D)$ the vector of corresponding eigenvalues. Then, for any $D \times D$ Hermitian matrix \mathbf{A} , denoting $\bar{\mathbf{A}} = \mathbf{E}^* \mathbf{A} \mathbf{E}$, we have*

$$\left. \frac{\partial e^{\Sigma}}{\partial \Sigma} \right|_{\Sigma} [\mathbf{A}] = \mathbf{E} [\mathbf{G} \odot \bar{\mathbf{A}}] \mathbf{E}^* \quad (17)$$

where \odot is the element-wise (a.k.a, Hadamar) product, and, for all $1 \leq i, j \leq D$, we have defined

$$\mathbf{G}_{i,j} = \begin{cases} \frac{e^{\lambda_i} - e^{\lambda_j}}{\lambda_i - \lambda_j} & \text{if } i \neq j, \\ e^{\lambda_i} & \text{otherwise.} \end{cases} \quad (18)$$

Note that Proposition 1 assumes the eigenvalues of Σ to be distinct. In practice, we observe instabilities when λ_i and λ_j are close to each other. To solve this numerical issue we use that, for $\lambda_i > \lambda_j$,

$$e^{\lambda_j} \leq \frac{e^{\lambda_i} - e^{\lambda_j}}{\lambda_i - \lambda_j} \leq e^{\lambda_i} \quad (19)$$

since exp is convex and increasing. Then, whenever an off-diagonal element $\mathbf{G}_{i,j}$ is out of this constraint we

Algorithm 2: Exact evaluation of the gradient ∇F

input : a vector $\mathbf{x} \in \mathbb{R}^{D^2}$

input : a vector $\mathbf{y} \in \mathbb{R}^{D^2}$

input : a vector $\mathbf{u} \in \mathbb{R}^{D^2}$

input : the affine map $\Omega : \mathbf{x} \mapsto \mathcal{K}(\mathbf{W}\Phi\mathbf{x} + \mathbf{b})$

output: the gradient \mathbf{g} of $F(\mathbf{x})$ defined in (10)

- 1 compute the eigenvalue decomposition:
 $-\Omega(\mathbf{x}) = \mathbf{E} \text{diag}(\mathbf{A}) \mathbf{E}^*$
 - 2 compute \mathbf{G} (Eq.(18))
 - 3 $\{\mathbf{G}_{i,j} \leftarrow \text{Proj}(\mathbf{G}_{i,j})\}_{i,j=1..D^2}$
(projection on $[e^{\lambda_j}, e^{\lambda_i}]$)
 - 4 $\bar{\mathbf{A}} \leftarrow e^{\Omega(\mathbf{y})}$
 - 5 $\bar{\mathbf{A}} \leftarrow \mathbf{E}^* \bar{\mathbf{A}} \mathbf{E}$
 - 6 $\mathbf{g} \leftarrow \beta(\mathbf{x} - \mathbf{u}) + L\Theta(\mathbf{Id}_D - \mathbf{E}[\mathbf{G} \odot \bar{\mathbf{A}}]\mathbf{E}^*)$
-

project its value on the feasible range $[e^{\lambda_j}, e^{\lambda_i}]$. Equation (19) shows that $\lim_{\lambda_j \rightarrow \lambda_i} \frac{e^{\lambda_i} - e^{\lambda_j}}{\lambda_i - \lambda_j} = e^{\lambda_i}$. In the case of duplicate eigenvalues, we use the continuous expansion obtained by replacing the condition $i \neq j$ with $\lambda_i \neq \lambda_j$. We checked numerically that this continuous expansion is working for matrices with repeated eigenvalues.

Proposition 1 gives us an exact closed-form formula for the directional derivative. The next corollary shows that this formula is also valid for its adjoint (see proof in Appendix C).

Corollary 1 *Let Σ be a Hermitian matrix with distinct eigenvalues. The Jacobian of the matrix exponential is a self-adjoint operator*

$$\left. \frac{\partial e^{\Sigma}}{\partial \Sigma} \right|_{\Sigma}^* = \left. \frac{\partial e^{\Sigma}}{\partial \Sigma} \right|_{\Sigma} . \quad (20)$$

Based on the closed-form expressions provided by Proposition 1 and Corollary 1, we define Algorithm 2 for an exact evaluation of the gradient of F .

3.2 A refined quasi-Newton scheme

The computation of the proximal operator requires the minimization of function $F(\mathbf{x})$. This can be performed using quasi-Newton iterations:

$$\mathbf{x} \leftarrow \mathbf{x} - \hat{\mathbf{H}}^{-1} \mathbf{g} \quad (21)$$

where $\hat{\mathbf{H}}$ is an approximation of the Hessian \mathbf{H} of F at \mathbf{x} , i.e., the real symmetric matrix defined by

$$\mathbf{H} = \frac{\partial^2 F(\mathbf{x})}{\partial \mathbf{x}^2} = \frac{\partial}{\partial \mathbf{x}} \nabla_{\mathbf{x}} F(\mathbf{x}) . \quad (22)$$

While in [13] a heuristic was used to define a diagonal approximation to the matrix \mathbf{H} , we consider here the following approximation:

$$\hat{\mathbf{H}} = \gamma \mathbf{I}_{D^2} \quad \text{with} \quad \gamma = \mathbf{d}^* \mathbf{H} \mathbf{d} \quad \text{and} \quad \mathbf{d} = \frac{\mathbf{g}}{\|\mathbf{g}\|} \quad (23)$$

where γ corresponds to the *exact* second derivative of F at \mathbf{x} in the direction \mathbf{d} of the gradient \mathbf{g} .

As proved in Appendix D, this second order derivative is given by:

$$\gamma = \beta + L \left\langle \Theta^*(\mathbf{d}), \left. \frac{\partial^2 e^{\Sigma}}{\partial \Sigma^2} \right|_{-\Omega(\mathbf{x})} [e^{\Omega(\mathbf{y})}, \Theta^*(\mathbf{d})] \right\rangle, \quad (24)$$

where $\Theta^*(\cdot) = \mathcal{K}(\mathbf{W}\Phi \cdot)$ is the adjoint of linear operator Θ and for any matrices \mathbf{X} and \mathbf{Y} , the matrix dot product is defined as $\langle \mathbf{X}, \mathbf{Y} \rangle = \text{tr}(\mathbf{X}^* \mathbf{Y})$. This choice for approximating the Hessian leads to a quasi-Newton step that is exact when restricted to the direction of the gradient. As F has some regions of non-convexity, we consider in practice half the absolute value of the scalar product in (24) in order to avoid some local minima and ensure that quasi-Newton follows a descent direction. Note that, like with the original formulation of MuLoG, we recover the exact Newton method in the mono-channel case ($D = 1$).

The computation of γ thus requires computing the second-order directional derivative of the matrix exponential. We recall that the second-order directional derivative of a function $f : \mathbb{C}^{D \times D} \rightarrow \mathbb{C}^{D \times D}$ at \mathbf{X} in the directions \mathbf{A} and \mathbf{B} is the $D \times D$ complex matrix defined as

$$\begin{aligned} \frac{\partial^2 f(\mathbf{X})}{\partial \mathbf{X}^2} [\mathbf{A}, \mathbf{B}] &= \\ \lim_{\varepsilon \rightarrow 0} \frac{1}{4\varepsilon^2} & [f(\mathbf{X} + \varepsilon \mathbf{A} + \varepsilon \mathbf{B}) - f(\mathbf{X} - \varepsilon \mathbf{A} + \varepsilon \mathbf{B}) \\ & - f(\mathbf{X} + \varepsilon \mathbf{A} - \varepsilon \mathbf{B}) + f(\mathbf{X} - \varepsilon \mathbf{A} - \varepsilon \mathbf{B})]. \end{aligned} \quad (25)$$

The closed form expression for the second-order directional derivative of the matrix exponential is given in the following Proposition (see proof in Appendix E):

Proposition 2 *Let Σ be a $D \times D$ Hermitian matrix with distinct eigenvalues. Let $\Sigma = \mathbf{E} \text{diag}(\mathbf{A}) \mathbf{E}^*$ be its eigendecomposition where \mathbf{E} is a unitary matrix of eigenvectors and $\mathbf{A} = (\lambda_1, \dots, \lambda_D)$ the vector of corresponding eigenvalues. For any $D \times D$ Hermitian matrices \mathbf{A} and \mathbf{B} , denoting $\bar{\mathbf{A}} = \mathbf{E}^* \mathbf{A} \mathbf{E}$ and $\bar{\mathbf{B}} = \mathbf{E}^* \mathbf{B} \mathbf{E}$, we have*

$$\frac{\partial^2 e^{\Sigma}}{\partial \Sigma^2} [\mathbf{A}, \mathbf{B}] = \mathbf{E} [F(\bar{\mathbf{A}}, \bar{\mathbf{B}})] \mathbf{E}^* \quad (26)$$

Algorithm 3: Iterative method to compute $\text{prox}_{\text{data}}$

input : a noisy vector $\mathbf{y} \in \mathbb{R}^{D^2}$
input : a vector $\mathbf{u} \in \mathbb{R}^{D^2}$
input : the affine map $\Omega : \mathbf{x} \mapsto \mathcal{K}(\mathbf{W}\Phi \mathbf{x} + \mathbf{b})$
input : the linear map $\Theta : \boldsymbol{\alpha} \mapsto \Phi \mathbf{W}^* \mathcal{K}^* \boldsymbol{\alpha}$
output: the approximate minimizer of $F(\mathbf{x})$

- 1 $\mathbf{x} \leftarrow \frac{\beta}{\beta+1} \mathbf{u} + \frac{1}{\beta+1} \mathbf{y}$ *(initialization)*
- repeat**
- 2 $\mathbf{g} \leftarrow \text{Algorithm 2}(\mathbf{x}, \mathbf{y}, \mathbf{u}, \Omega)$ *(gradient)*
- 3 $\mathbf{d} \leftarrow \mathbf{g} / \|\mathbf{g}\|$ *(gradient direction)*
- 4 $\bar{\mathbf{B}} \leftarrow \mathbf{E}^* \Theta^*(\mathbf{d}) \mathbf{E}$
- 5 compute $F(\bar{\mathbf{A}}, \bar{\mathbf{B}})$ *(Eq.(27))*
- 6 $\gamma \leftarrow \beta + L |\langle \bar{\mathbf{B}}, F(\bar{\mathbf{A}}, \bar{\mathbf{B}}) \rangle|$ *(inverse step size)*
- 7 $\mathbf{x} \leftarrow \mathbf{x} - \frac{1}{\gamma} \mathbf{g}$ *(quasi-Newton step)*

until *convergence*;

where, for all $1 \leq i, j \leq D$, we have

$$F(\bar{\mathbf{A}}, \bar{\mathbf{B}})_{i,j} = \sum_{k=1}^D \varphi_{i,j,k} (\bar{\mathbf{A}}_{ik} \bar{\mathbf{B}}_{jk}^* + \bar{\mathbf{B}}_{ik} \bar{\mathbf{A}}_{jk}^*) \quad (27)$$

$$\text{with} \quad \varphi_{i,j,k} = \begin{cases} \frac{\mathbf{G}_{ik} - \mathbf{G}_{jk}}{\lambda_i - \lambda_j} & \text{if } i \neq j, \\ \frac{\mathbf{G}_{ii} - \mathbf{G}_{ik}}{\lambda_i - \lambda_k} & \text{if } i = j \text{ and } k \neq i, \\ \frac{\mathbf{G}_{ii}}{2} & \text{if } i = j = k. \end{cases} \quad (28)$$

As in Proposition 1, Proposition 2 assumes the eigenvalues to be distinct. We checked on numerical simulations that the result is still valid for Hermitian matrices with duplicate eigenvalues by simply defining \mathbf{G}_{ij} in Proposition 1 using the condition $\lambda_i \neq \lambda_j$ instead of $i \neq j$.

Algorithm 3 details the quasi-Newton optimization scheme to solve the minimization problem (11). The algorithm starts with an initial value for \mathbf{x} , line 1, obtained by approximating the data likelihood with a Gaussian (i.e., Wiener).

3.3 Numerical validation

In order to validate the correctness of the derived closed-form expressions of the gradient and the directional second derivative, we leverage that for $\mathbf{g} = \nabla F(\mathbf{x})$ and $\gamma = \mathbf{d}^* \frac{\partial^2 F(\mathbf{x})}{\partial \mathbf{x}^2} \mathbf{d}$, the following two identities regarding the first and second directional derivatives hold true

$$\langle \mathbf{g}, \mathbf{u} \rangle = \lim_{\varepsilon \rightarrow 0} \frac{F(\mathbf{x} + \varepsilon \mathbf{u}) - F(\mathbf{x})}{\varepsilon}, \quad \forall \mathbf{u} \in \mathbb{R}^{D^2}, \quad (29)$$

and

$$\gamma = \lim_{\varepsilon \rightarrow 0} \frac{F(\mathbf{x} + \varepsilon \mathbf{d}) - 2F(\mathbf{x}) + F(\mathbf{x} - \varepsilon \mathbf{d})}{\varepsilon^2}. \quad (30)$$

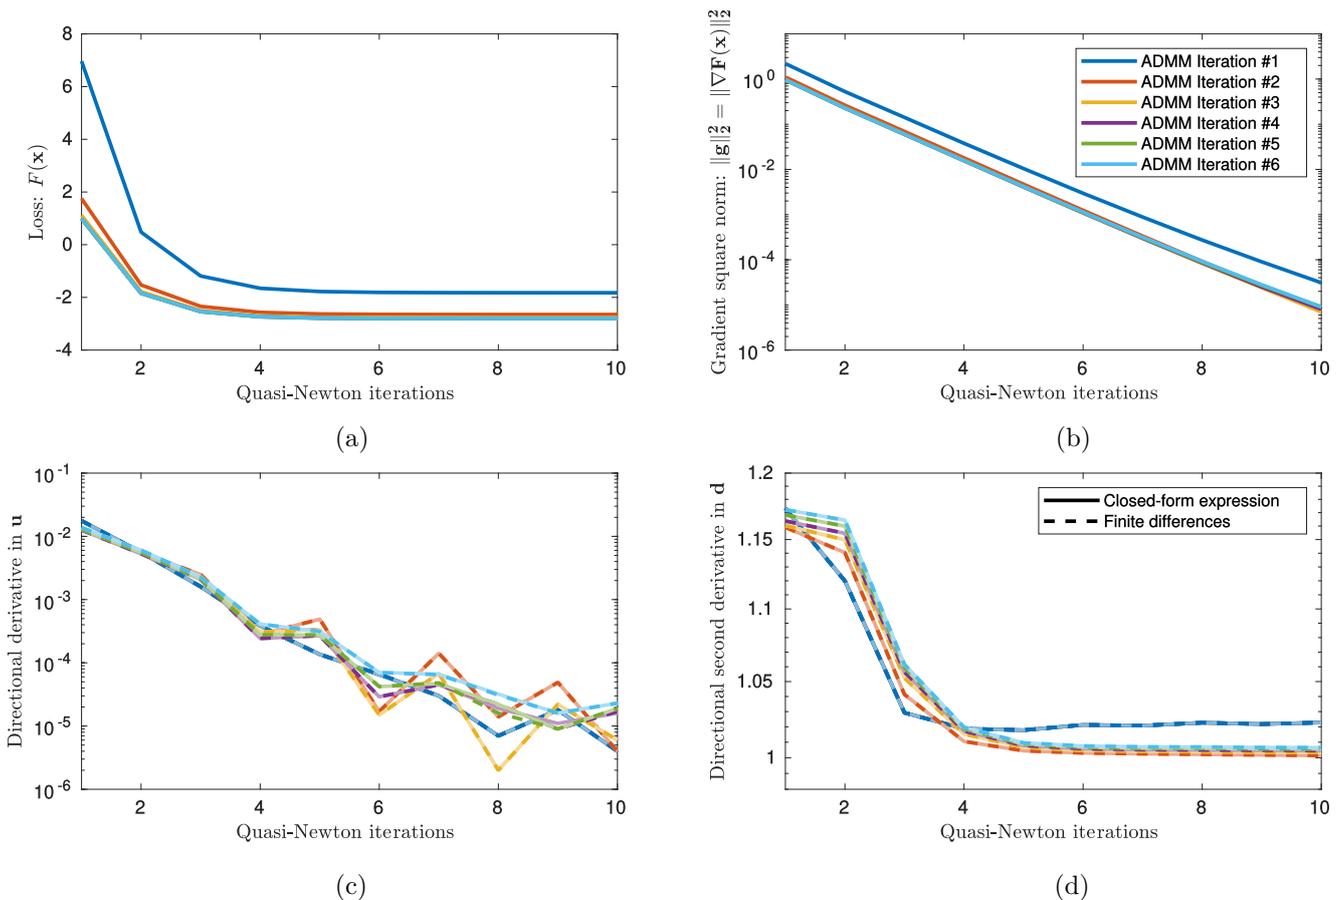


Fig. 4 Evolution with respect to the quasi-Newton iterations of (a) the loss function $F(\mathbf{x})$, (b) the square norm of its gradient $\mathbf{g} = \nabla F(\mathbf{x})$, (c) the directional derivative $\langle \mathbf{g}, \mathbf{u} \rangle$ for a random direction \mathbf{u} evaluated in closed-form and by finite differences, (d) the second directional derivative $\gamma = \mathbf{d}^* \frac{\partial^2 F(\mathbf{x})}{\partial \mathbf{x}^2} \mathbf{d}$ for the direction $\mathbf{d} = \mathbf{g} / \|\mathbf{g}\|_2$ evaluated in closed-form and by finite differences. The curves are provided for the six iterations of Plug-and-Play ADMM.

During the iterations of the proposed quasi-Newton scheme, we evaluate the left-hand sides of these equalities by running Algorithm 2 and 3 to compute \mathbf{g} and γ . Independently, we evaluate the right-hand sides of these equalities by finite difference with a small non-zero value of ε . The direction \mathbf{u} was chosen as a fixed white standardized normal vector.

Figure 4 shows the evolution of these four quantities during the iterations of the proposed quasi-Newton scheme for an arbitrary choice of the image \mathbf{y} , initializations and constant L and β . In addition, the evolution of $F(\mathbf{x})$ and $\|\mathbf{g}\|_2^2 = \|\nabla F(\mathbf{x})\|_2^2$ are also provided. On the one hand, we can notice that the proposed first and second directional derivatives are very close to the ones estimated by finite differences, which shows the validity of our formula. On the other hand, the objective $F(\mathbf{x})$ is decreasing and its gradient converges to 0, showing that the obtained stationary point is likely to be a local minimum. Furthermore, one can notice that the second directional derivative varies less than 20% showing that

the loss F is nearly quadratic in the vicinity of its minimizers.

4 Initialization and regularization of rank-deficient matrix fields

Multi-channel SAR images with n pixels are provided in the form of a bidimensional field of either n diffusion vectors $\mathbf{v}_k \in \mathbb{C}^D$ (single-look data) or n covariance matrices $\mathbf{C}_k \in \mathbb{C}^{D \times D}$ (multi-look data), see Figure 1. The statistical distribution of speckle is defined with respect to full-rank covariance matrices Σ_k . The initial guess for this matrix must be positive-definite, which requires an adapted strategy. We discuss different strategies for this initialization in Paragraph 4.1. Then, we describe in Paragraph 4.2 how the noisy covariance matrices \mathbf{C}_k can be regularized so that the neg-log-likelihood is better behaved.

4.1 Initialization

Under the assumption that the radar properties of the scene vary slowly with the spatial location, a spatial averaging can be used to estimate a first guess of the covariance matrix $\Sigma_k \stackrel{\text{def}}{=} \mathbb{E}[\mathbf{v}_k \mathbf{v}_k^*]$. At pixel k , the so-called *radar coherence* between channels i and j of the D -channels SAR image is defined by $\rho_{k,i,j} = |[\Sigma_k]_{i,j}| / \sqrt{[\Sigma_k]_{i,i}[\Sigma_k]_{j,j}}$, i.e., the correlation coefficient between the two channels. This coherence is notably over-estimated by the sample covariance estimator when L is small, see [36]. In the extreme case of $L = 1$, $\mathbf{C}_k = \mathbf{v}_k \mathbf{v}_k^*$ and the coherence equals 1, see Figure 5. Some amount of smoothing is necessary to obtain more meaningful coherence values and to guarantee that the covariance matrix $\mathbf{C}_k^{(\text{init})}$ be positive definite. The coherence is estimated by weighted averaging with Gaussian weights:

$$\hat{\rho}_{k,i,j} = \frac{|\sum_{\ell} w_{k,\ell} [\mathbf{C}_{\ell}]_{i,j}|}{\sqrt{(\sum_{\ell} w_{k,\ell} [\mathbf{C}_{\ell}]_{i,i})(\sum_{\ell} w_{k,\ell} [\mathbf{C}_{\ell}]_{j,j})}}, \quad (31)$$

where the weights $w_{k,\ell}$ are defined, based on the spatial distance between pixels k and ℓ , using a Gaussian kernel with spatial variance $\tau/2\pi$. Spatial parameter τ is chosen in order to achieve a trade-off between bias reduction and resolution loss. Figure 5 shows coherence images $\hat{\rho}$ of a single-look interferometric pair ($L = 1$ and $D = 2$) for different values of τ . We use the empirical rule $\tau = D / \min(L, D)$ to increase the filtering strength with the dimensionality of SAR images and reduce it for multi-look images.

The estimated coherence values $\hat{\rho}_{k,i,j}$ are then used to define a first estimate $\hat{\mathbf{C}}_k$ of the covariance matrix at pixel k :

$$[\hat{\mathbf{C}}_k]_{i,j} = \hat{\rho}_{k,i,j} \frac{\sqrt{[\mathbf{C}_{\ell}]_{i,i}[\mathbf{C}_{\ell}]_{j,j}}}{|[\mathbf{C}_{\ell}]_{i,j}|} [\mathbf{C}_k]_{i,j}. \quad (32)$$

Note that on-diagonal entries are unchanged: $\forall i, [\hat{\mathbf{C}}_k]_{i,i} = [\mathbf{C}_k]_{i,i}$, while the magnitude of off-diagonal entries becomes equal to $\hat{\rho}_{k,i,j}$ and the phase of off-diagonal entries is preserved. In practice, the coherence $\hat{\rho}_{k,i,j}$ is smaller than 1 for all values of τ larger than 0. In order to guarantee that the covariance matrix $\hat{\mathbf{C}}_k$ actually is positive definite, off-diagonal entries can additionally be shrunk toward 0 by a small factor (by multiplication by a factor .99 for example). The filtering step applied in equations (31) and (32) largely improves the conditioning of covariance matrix $\hat{\mathbf{C}}_k$, which helps performing the principal component analysis required to define transform Ω (by producing a more stable analysis).

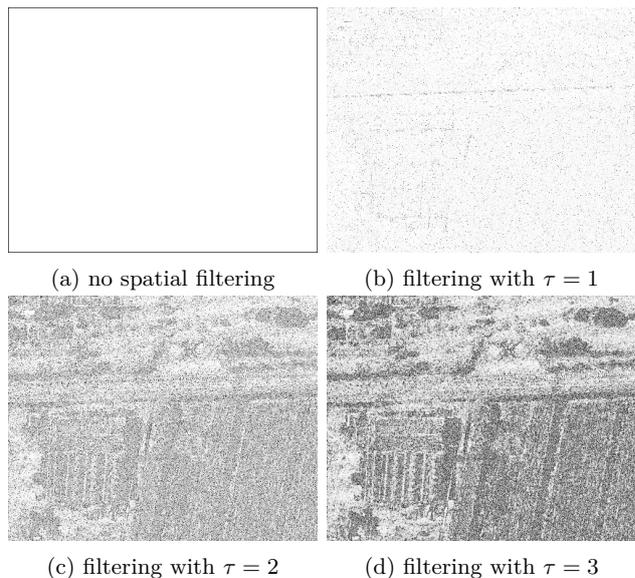


Fig. 5 Estimation of the interferometric coherence by local Gaussian filtering, for different filtering strengths. In the absence of spatial filtering, the coherence is equal to 1 everywhere. All images are displayed using the same graylevel scale with a zero coherence displayed in black and a unit coherence displayed in white. The area shown corresponds to the images of Figure 2.

In contrast to intensity data for which $\mathbb{E}\{[\mathbf{C}_k]_{ij}\} = [\Sigma_k]_{ij}$ for all channels i and j , the average of log-transformed data is known to suffer from a systematic bias [38] that can be quantified on diagonal elements, for all $1 \leq k \leq D$, by

$$\mathbb{E}[\log \mathbf{C}_{kk}] - \log \Sigma_{kk} = \psi(L) - \log L, \quad (33)$$

where ψ is the di-gamma function. MuLoG does not estimate $\log \Sigma$ by averaging, but by iterating the lines 8 to 11 in Algorithm 1. The sequence of these iterations leads to an unbiased estimate. However, at the first iteration, the Gaussian denoiser function f_{σ} is applied to the initial guess $\mathbf{C}_k^{(\text{init})}$. This denoiser performs an averaging of the values in homogeneous areas. Convergence to the final values is faster if the bias is pre-compensated in the initial guess $\mathbf{C}^{(\text{init})}$, as done in line 5 of Algorithm 4.

4.2 Regularization

The previous paragraph described the strategy to build an initial guess $\mathbf{C}_k^{(\text{init})}$ of the covariance matrix at pixel k . This guess serves as a starting point for the iterative estimation procedure conducted by Algorithm 4. Line 2 of Algorithm 4, a regularized version $\mathbf{C}_k^{(\text{reg})}$ of the covariance matrix is defined to compute vector \mathbf{y}

Algorithm 4: Initial estimation of the covariances

input : a bidimensional field of speckle-corrupted covariance matrices \mathbf{C}

output: the initial guess $\mathbf{C}^{(\text{init})}$
(to be used line 1 of Algorithm 1)

- 1 $\tau \leftarrow D / \min(L, D)$ *(filtering strength)*
- 2 $\mathbf{C}^{(\text{filt})} \leftarrow \mathbf{Gaussian_filter}(\mathbf{C}, \tau/2\pi)$
(filter with a Gaussian kernel of variance $\tau/2\pi$)
- 3 $\left\{ \hat{\rho}_{k,i,j} \leftarrow \frac{|[\mathbf{C}_\ell^{(\text{filt})}]_{i,i}|}{\sqrt{[\mathbf{C}_\ell^{(\text{filt})}]_{i,i}[\mathbf{C}_\ell^{(\text{filt})}]_{j,j}}} \right\}_{k=1..n, i,j=1..D}$
(estimate coherences)
- 4 $\left\{ [\hat{\mathbf{C}}_k]_{i,j} \leftarrow \hat{\rho}_{k,i,j} \frac{\sqrt{[\mathbf{C}_\ell]_{i,i}[\mathbf{C}_\ell]_{j,j}}}{|[\mathbf{C}_\ell]_{i,j}|} [\mathbf{C}_k]_{i,j} \right\}_{k=1..n, i,j=1..D}$
(impose the estimated coherences)
- 5 $\left\{ \mathbf{C}_k^{(\text{init})} \leftarrow \exp(\log L - \psi(L)) \hat{\mathbf{C}}_k \right\}_{k=1..n}$
(pre-compensate bias)

(line 5) and then used to define the data-fidelity proximal operator (line 10). Although a rank-deficient matrix \mathbf{C}_k such as $\mathbf{v}_k \mathbf{v}_k^*$ could possibly be used to define the proximal operator (by replacing $\exp(\Omega(\mathbf{y}))$ by this rank-deficient matrix in the definition of $F(\mathbf{x})$ in equation (10)), rank-deficient or ill-conditioned covariance matrices lead to cost functions $F(\mathbf{x})$ that are harder to minimize. On the other hand, we show in the sequel that the spatial smoothing strategy used in Algorithm 4 to build our initial guess $\mathbf{C}_k^{(\text{init})}$ should not be used to compute $\mathbf{C}_k^{(\text{reg})}$ (as done in the original algorithm MuLoG [13]) since this may lead to significant biases.

In order to control the condition number of the covariance matrices $\mathbf{C}_k^{(\text{reg})}$, we adjust the eigenvalues by applying an affine map that rescales the eigenvalues from the range $[\lambda_{\min}, \lambda_{\max}]$ to the range $[\lambda_{\max}/\bar{c}, \lambda_{\max}]$, with $\bar{c} = \min(c, \lambda_{\max}/\lambda_{\min})$. This transformation ensures that the resulting covariance matrix has a condition number at most equal to c . Moreover, the largest eigenvalue λ_{\max} is left unchanged and the ordering of the eigenvalues is preserved by this strictly increasing mapping (provided that $c > 1$). It seems that this latter property is beneficial to limit the bias introduced by the covariance regularization scheme. If the condition number is larger than the actual condition number of \mathbf{C}_k , the affine map corresponds to the identity map. The computation of the regularized covariance matrices is summarized in Algorithm 5. We use in the following the value $c = 10^3$.

Algorithm 5: Regularization of covariance matrices

input : a covariance matrix $\mathbf{C}_k \in \mathbb{C}^{D \times D}$

input : a condition number c

output: a regularized matrix $\mathbf{C}_k^{(\text{reg})} \in \mathbb{C}^{D \times D}$

- 1 compute the eigenvalue decomposition:
 $\mathbf{C}_k = \mathbf{E} \text{diag}(\boldsymbol{\Lambda}) \mathbf{E}^*$
- 2 $\lambda_{\min} \leftarrow \min(\boldsymbol{\Lambda})$ *(smallest eigenvalue)*
- 3 $\lambda_{\max} \leftarrow \max(\boldsymbol{\Lambda})$ *(largest eigenvalue)*
- 4 $\bar{c} \leftarrow \min(c, \lambda_{\max}/\lambda_{\min})$ *(targeted condition number)*
- 5 $\left\{ \lambda_k \leftarrow \frac{\lambda_k - \lambda_{\min}}{\lambda_{\max} - \lambda_{\min}} \lambda_{\max} (1 - 1/\bar{c}) + \lambda_{\max}/\bar{c} \right\}_{k=1..D}$
(affine mapping of the eigenvalues)
- 6 $\mathbf{C}_k^{(\text{reg})} \leftarrow \mathbf{E} \text{diag}(\boldsymbol{\Lambda}) \mathbf{E}^*$
(recomposition with the modified eigenvalues)

5 Numerical experiments**5.1 Simulated data**

In a first experiment, we compare the impact of the modifications introduced in Sections 3 and 4 with respect to the original MuLoG algorithm in [13] on a simulated PolSAR image generated from optical satellite images by building at each pixel index k the following covariance matrix

$$\boldsymbol{\Sigma}_k = \frac{1}{4} \begin{pmatrix} (G_k + R_k)^2 & 0 & (G_k^2 - R_k^2)(1+j) \\ 0 & 4B_k^2 & 0 \\ (G_k^2 - R_k^2)(1-j) & 0 & (G_k - R_k)^2 \end{pmatrix}, \quad (34)$$

where R_k , G_k , and B_k are the Red-Green-Blue (RGB) channels of the optical image, and the polarimetric channels of the covariance matrix are organized in the following order HH , HV , and VV . This way the optical image coincides with the RGB representation of $\boldsymbol{\Sigma}$ when represented by fake colors in the Pauli basis, as described page 3. This model considers that channel HV is decorrelated from channels HH and VV , while channels HH and VV have a correlation of $1/\sqrt{2} \approx 0.71$. Given such a ground truth image $\boldsymbol{\Sigma}$, we next simulated noisy versions \mathbf{C} with L looks by random sampling, at each pixel index k ,

$$\mathbf{C}_k = \boldsymbol{\Sigma}_k^{1/2} \mathbf{Z}_k \boldsymbol{\Sigma}_k^{1/2} \text{ with } \mathbf{Z}_k = \frac{1}{L} \sum_{l=1}^L (\mathbf{v}_k^{(l)})(\mathbf{v}_k^{(l)})^*, \quad (35)$$

where $\boldsymbol{\Sigma}_k^{1/2}$ denotes the Hermitian square root of $\boldsymbol{\Sigma}_k^{1/2}$, and $\mathbf{v}_k^{(l)}$ are independent complex random vectors with real and imaginary parts drawn according to a Gaussian white noise with standard deviation $1/\sqrt{2}$. By construction, this gives $\mathbf{Z}_k \sim \mathcal{W}(\mathbf{Id}_D, L)$ and $\mathbf{C}_k \sim \mathcal{W}(\boldsymbol{\Sigma}, L)$.

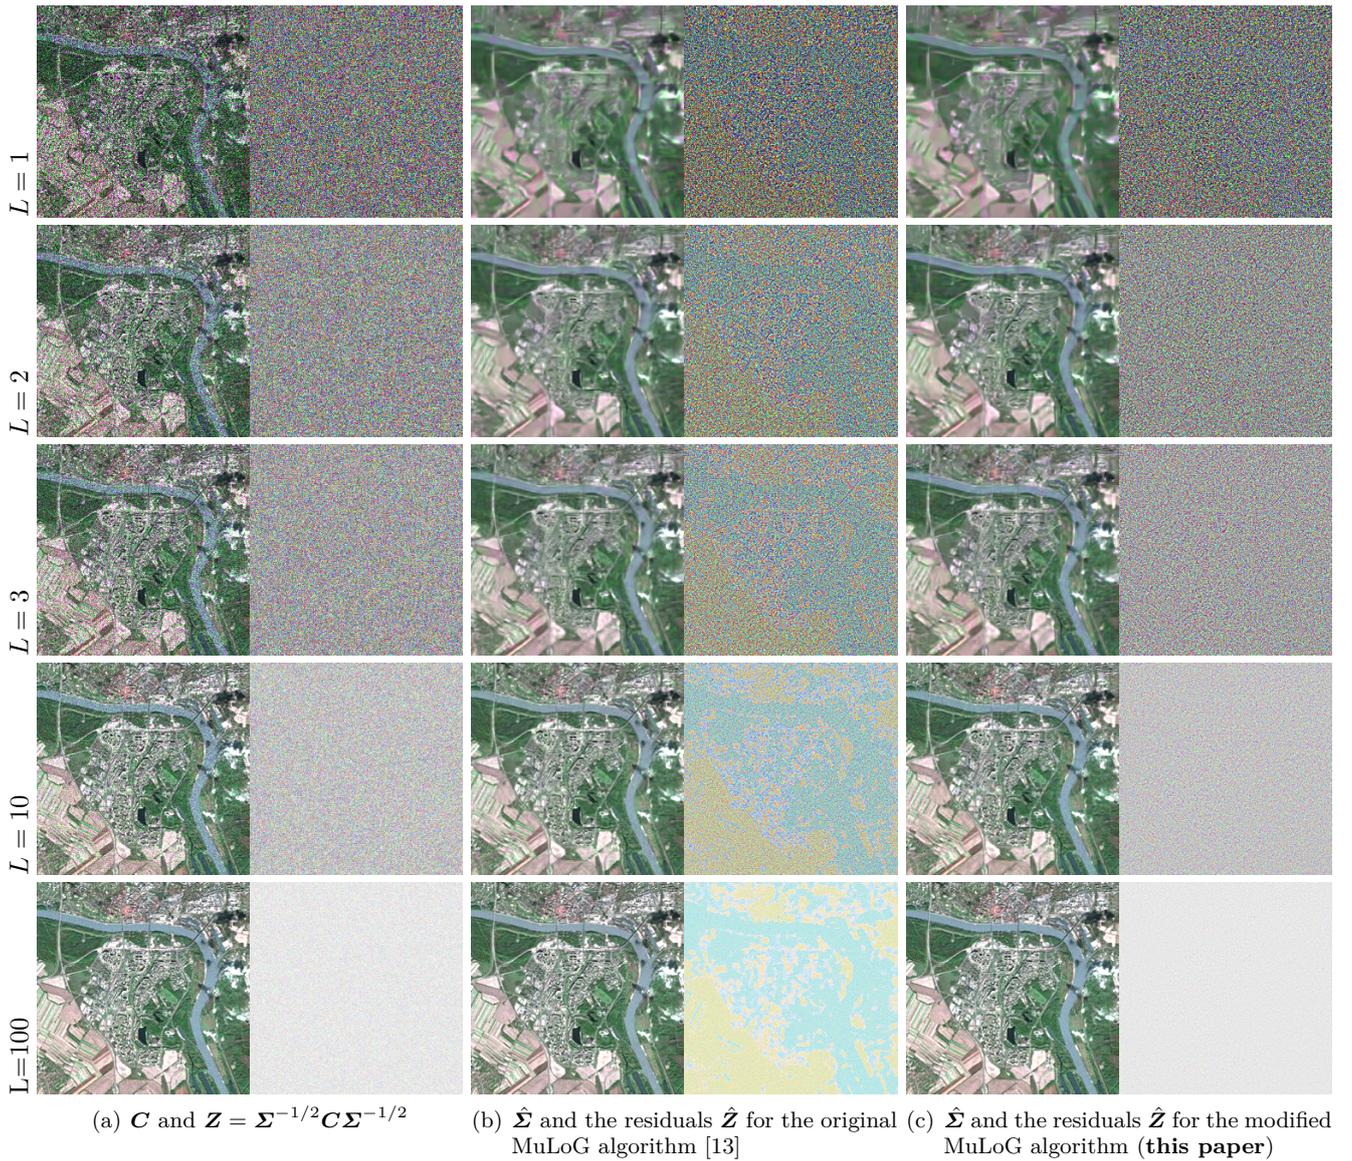


Fig. 6 Comparisons of the original MuLoG algorithm [13] and the modified version introduced in this paper, for different values of L . (a) Noisy image \mathbf{C} and its noise component $\mathbf{Z} = \boldsymbol{\Sigma}^{-1/2} \mathbf{C} \boldsymbol{\Sigma}^{-1/2}$ simulated from the 2003 SPOT Satellite image of Bratislava. (b-c) Estimates $\hat{\boldsymbol{\Sigma}}$ and the corresponding residuals $\hat{\mathbf{Z}} = \hat{\boldsymbol{\Sigma}}^{-1/2} \mathbf{C} \hat{\boldsymbol{\Sigma}}^{-1/2}$ (method noise) for the original and modified versions of MuLoG. The residuals \mathbf{Z} should ideally be as close as possible to the actual noise component \mathbf{Z} . The ground truth is not displayed but is extremely close to \mathbf{C} when $L = 100$ (last row).

5.2 Evaluation with Simulations

Using the procedure described in the previous paragraph, we simulated a PolSAR data $\boldsymbol{\Sigma}$ from a 2003 SPOT Satellite image of Bratislava (Slovakia)¹. We suggest performing first a visual comparison of the estimated $\hat{\boldsymbol{\Sigma}}$ respectively obtained by the original version of MuLoG and by the modified version introduced in this paper, on noisy versions \mathbf{C} obtained for five differ-

ent numbers of looks: $L = 1, 2, 3, 10$ and 100 . In addition, in order to get more insight into the behavior of each estimator, we display the residuals (aka, method noise [4]), defined by $\hat{\mathbf{Z}} = \hat{\boldsymbol{\Sigma}}^{-1/2} \mathbf{C} \hat{\boldsymbol{\Sigma}}^{-1/2}$ where $\hat{\boldsymbol{\Sigma}}^{-1/2}$ is the inverse of the Hermitian square root $\hat{\boldsymbol{\Sigma}}^{1/2}$. Should the estimation be perfect, $\hat{\boldsymbol{\Sigma}}$ would exactly be equal to $\boldsymbol{\Sigma}$, and $\hat{\mathbf{Z}}$ would perfectly match the white speckle component \mathbf{Z} . Comparing $\hat{\mathbf{Z}}$ to \mathbf{Z} is thus an efficient way to assess the bias/variance trade-off achieved by the estimator. Areas of $\hat{\mathbf{Z}}$ that appear less noisy than \mathbf{Z} indicate an under-smoothing. If $\hat{\mathbf{Z}}$ contains structures not present in \mathbf{Z} , this indicates an over-smoothing. Where-

¹ Provided by CNES under Creative Commons Attribution-Share Alike 3.0 Unported license. See: https://commons.wikimedia.org/wiki/File:Bratislava_SPOT_1027.jpg

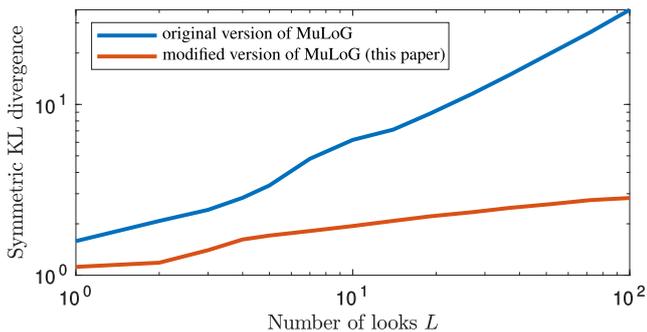


Fig. 7 Evolution of the symmetric KL divergence between the ground-truth Σ and the estimated residuals $\hat{\Sigma}$ obtained with the original or the modified version of MuLoG. The divergence is computed on the simulated PolSAR data obtained from the 2003 SPOT Satellite image of Bratislava shown on Figure 6, for various number of looks L .

ever \hat{Z} has a different color than Z , this is a sign of bias.

Results are provided in Figure 6. The estimated images $\hat{\Sigma}$ obtained with the original version of MuLoG and the modified version introduced in this paper are visually very close. As expected, the quality of the restored images improves with the number of looks along with the signal-to-noise ratio in the speckle-corrupted images. The residuals \hat{Z} clearly indicate a bias with the original version of MuLoG. This bias is smaller with respect to the speckle variance when the number of looks L is lower.

The residuals \hat{Z} obtained with the modified MuLoG algorithm are comparable to the true residuals Z : no geometrical structure from the image is noticeable in the residuals, which indicates that contrasted features were not removed from the image by the despeckling processing. The levels of fluctuations in \hat{Z} and Z seem similar. In order to perform a more quantitative comparison of the residuals, we report in Figure 7 the symmetrical Kullback-Leibler divergence (KLD) between the distribution of the residuals Σ and $\hat{\Sigma}$ for different numbers of looks ranging from $L = 1$ to 100. The KLD, averaged over the pixel index k , between the distributions $\mathcal{W}(\Sigma_k; L)$ and $\mathcal{W}(\hat{\Sigma}_k; L)$ is defined by

$$SD_{\mathcal{W}}(\hat{\Sigma}_k \| \Sigma_k) = L \operatorname{tr}(\Sigma_k \hat{\Sigma}_k^{-1} + \Sigma_k^{-1} \hat{\Sigma}_k) - 2LD. \quad (36)$$

A KLD of 0 indicates a perfect match while a larger value indicates a discrepancy. The divergence increases with the number of looks, which is expected because the KLD is a measure of divergence relative to the signal-to-noise ratio: the larger the signal-to-noise ratio, the more conserving is the KLD. We observe that for all values of L , the modified version of MuLoG leads to residuals closer to the theoretical distribution of speckle residuals. This is in agreement with the behaviour observed

on Figure 6 where the bias is seen to become prominent for large numbers of looks with the original version of MuLoG.

We checked by comparing the average running time on several images that the modifications introduced in this paper do not slow down MuLoG: a slight speedup was even observed in our experiments.

5.3 Evaluation with Real data

Figures 8, 9 and 10 compare the restoration performance of the original MuLoG algorithm and of the modified version introduced in this paper on PolSAR images from 3 different sensors (AIRSAR from NASA, PISAR from JAXA and SETHI from ONERA). The equivalent numbers of looks, estimated by maximum likelihood, are respectively equal to 2.7, 1.7 and 1 for each image. From the 1 look image of SETHI, we build a 4 looks image by spatial averaging and downsampling by a factor 2 in the horizontal and vertical directions.

As previously observed on simulated data, while the results of the two versions of the algorithm are similar when the number of looks is small, a bias is visible in the residuals with the original MuLoG algorithm for larger equivalent numbers of looks.

6 Discussion

MuLoG is a generic framework that offers the possibility of a straightforward application of denoisers developed for additive white Gaussian noise (i.e., optical imaging). It suppresses the need for a time-consuming adaptation of these algorithms to the specifics of SAR imagery. Beyond a much faster transfer of state-of-the-art denoising methods, it makes it easier to run several denoisers in parallel and compare the restored images obtained by each. Figure 11 illustrates such restoration results obtained with 6 different denoising techniques. Method-specific artifacts can be identified in the results: the total-variation denoiser [34] suppresses oscillating patterns and creates artificial edges in smoothly varying regions; BM3D [10], based on wavelets transforms, is very good at restoring oscillating patterns but it may also produce oscillating artifacts in some low signal-to-noise ratio areas; DDID [23] creates oscillating artifacts and some ripples around edges; FEPLL [31] tends to suppress low signal-to-noise ratio oscillating patterns; DnCNN [39] produces smooth images with some point-like artifacts and a suppression of the low signal-to-noise ratio oscillating patterns; IRNCNN [40] better restores some of the oscillating patterns but introduces many artifacts in the form of fake-edges in

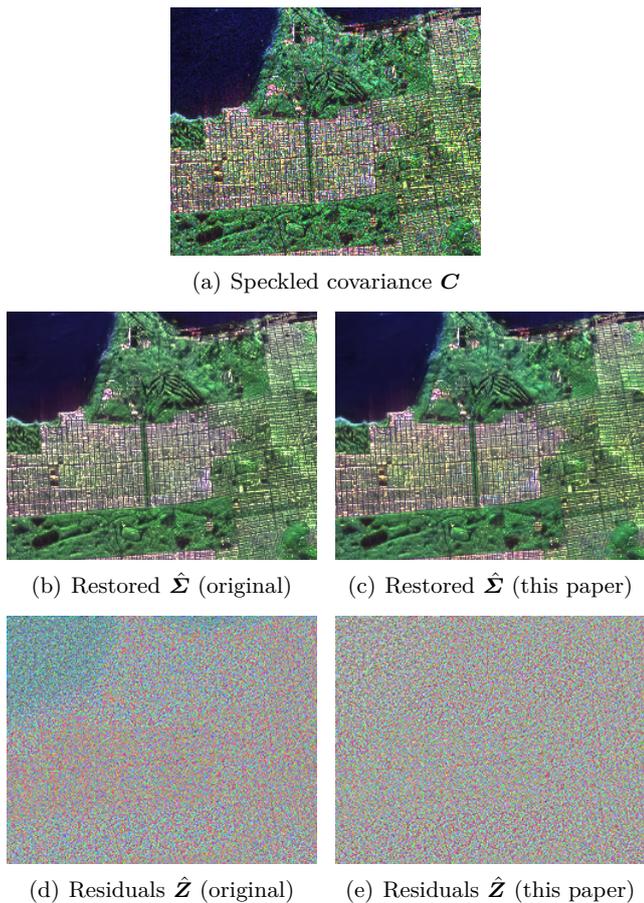


Fig. 8 Despeckling performance comparison of the original algorithm MuLoG and the modified version introduced in this paper. AIRSAR polarimetric image of San-Francisco, USA (©NASA): (a) The original image; (b-c) the despeckling results; (d-e) the residuals.

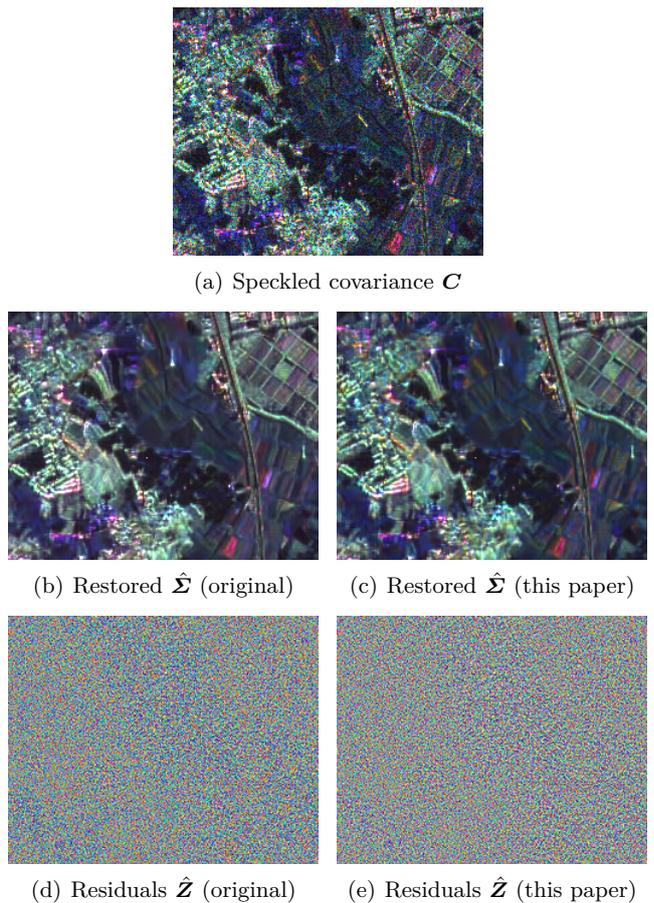


Fig. 9 Despeckling performance comparison of the original algorithm MuLoG and the modified version introduced in this paper. PISAR polarimetric image of Tsukuba, Japan (©JAXA): (a) The original image; (b-c) the despeckling results; (d-e) the residuals.

homogeneous areas. We believe that the possibility to include deep learning techniques with MuLoG is particularly useful to multi-dimensional SAR imaging for which the supervised training of dedicated networks is very difficult to achieve due to the lack of ground-truth data and the dimensionality of the patterns (spatial patterns that extend in the D^2 real-valued dimensions of the covariance matrices).

This paper introduced several modifications to the MuLoG algorithm [13] based on the closed-form derivation of first and second order derivatives of data-fidelity proximal operator. These mathematical developments can benefit other algorithms than MuLoG, in particular the large body of methods derived from the plug-in-ADMM strategy [6], algorithm unrolling [29] or stochastic sampling techniques such as those based on Langevin dynamics [24].

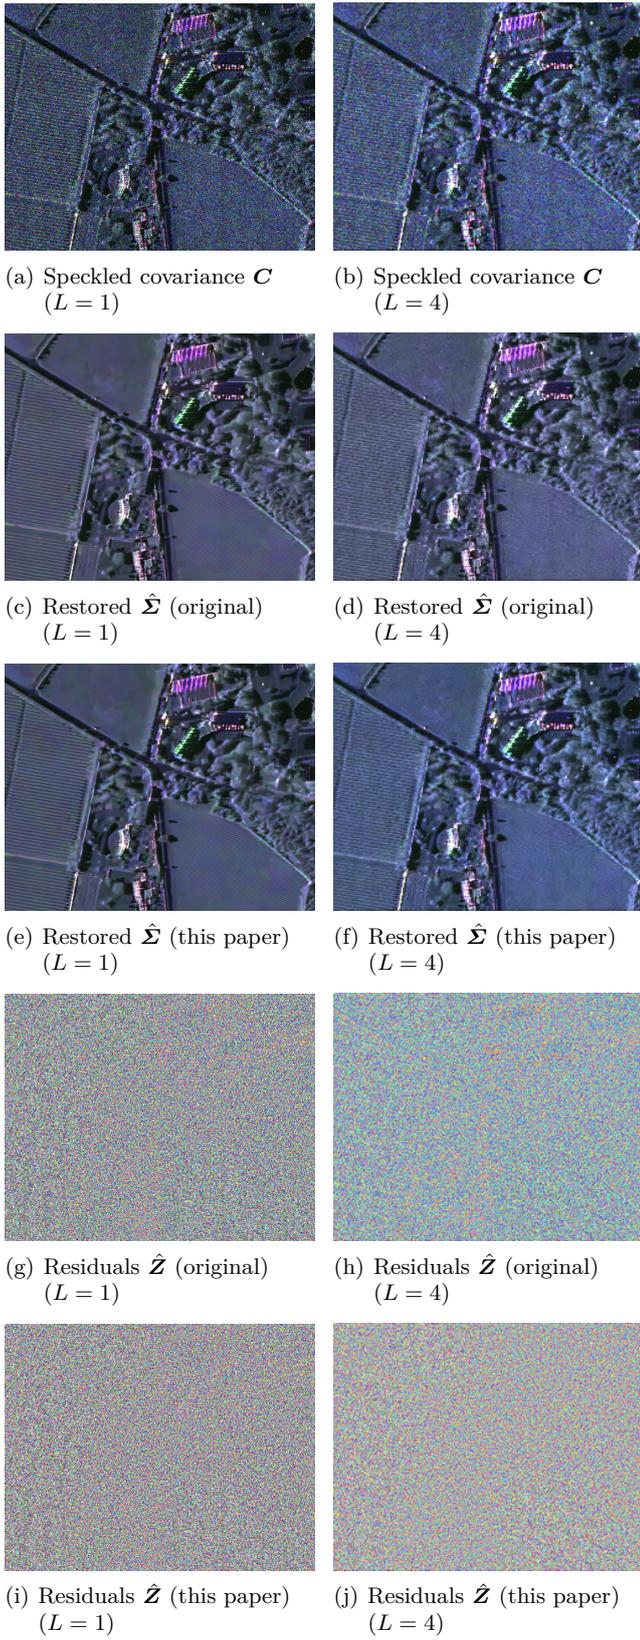


Fig. 10 Despeckling performance comparison of the original algorithm MuLoG and the modified version introduced in this paper. SETHI polarimetric image of Nîmes area, France (©ONERA): first column, single-look image; second column, 4-looks image. (a,b) The original images; (c-f) the despeckling results; (g-j) the residuals.

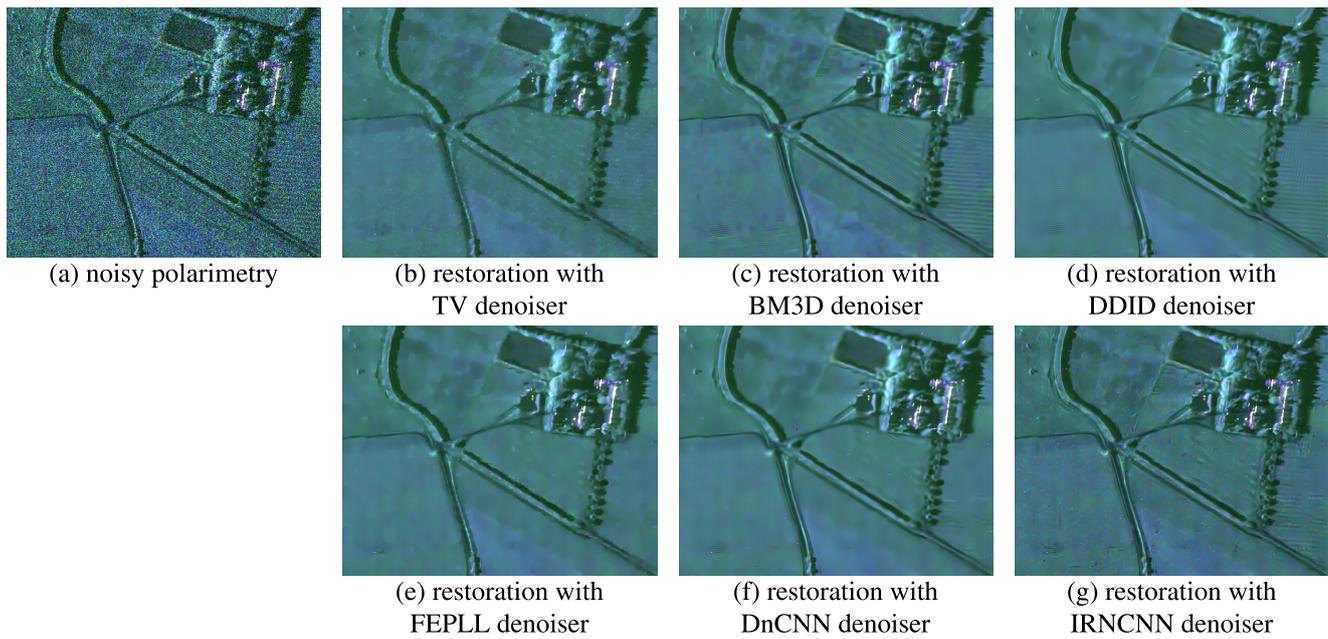


Fig. 11 Many different Gaussian denoisers can be used in conjunction with MuLoG. A speckled polarimetric image (a) is restored with MuLoG and conventional Gaussian denoisers: (b) total variation minimization [34]; (c) the block-matching collaborative 3D filtering algorithm BM3D [10]; (d) the dual domain image denoising method DDID [23]; or more recent learning-based techniques: (e) a fast algorithm based on a model of image patches as a mixture of Gaussians FEPLL [31]; and two deep neural networks: (f) DnCNN [39] and (g) IRNCNN [40].

A Gradient of the objective F

Proof (Proof of eq. (12)) Applying the chain rule to eq. (10) leads to the following decomposition

$$\mathbf{g} = \nabla_{\mathbf{x}} \left[\frac{\beta}{2} \|\mathbf{x} - \mathbf{z}\|^2 + L \operatorname{tr}(\Omega(\mathbf{x}) + e^{\Omega(\mathbf{y})} e^{-\Omega(\mathbf{x})}) \right] \quad (37)$$

$$= \beta(\mathbf{x} - \mathbf{z}) + L \left. \frac{\partial \Omega(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}}^* \left(\nabla_{\mathbf{X}} \operatorname{tr} \mathbf{X} \Big|_{\Omega(\mathbf{x})} \right. \quad (38)$$

$$\left. - \left. \frac{\partial e^{\mathbf{X}}}{\partial \mathbf{X}} \right|_{-\Omega(\mathbf{x})}^* \left. \frac{\partial e^{\Omega(\mathbf{y})} \mathbf{X}}{\partial \mathbf{X}} \right|_{e^{-\Omega(\mathbf{x})}}^* \nabla_{\mathbf{X}} \operatorname{tr} \mathbf{X} \Big|_{e^{\Omega(\mathbf{y})} e^{-\Omega(\mathbf{x})}} \right).$$

Using that for any matrix \mathbf{A} and \mathbf{B}

$$\left. \frac{\partial \mathbf{A} \mathbf{X}}{\partial \mathbf{X}} \right|_{e^{-\Omega(\mathbf{x})}}^* [\mathbf{B}] = \mathbf{B} \mathbf{A}^*, \quad (39)$$

$$\left. \frac{\partial \Omega(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}}^* [\mathbf{B}] = \Theta(\mathbf{B}), \quad (40)$$

$$\nabla_{\mathbf{X}} \operatorname{tr} \mathbf{X} \Big|_{e^{\Omega(\mathbf{y})} e^{-\Omega(\mathbf{x})}} = \mathbf{Id}_D. \quad (41)$$

concludes the proof.

B Proof of Proposition 1

Let us start by recalling the following Lemma whose proof can be found in [26, 17, 5].

Lemma 1 *Let Σ be an Hermitian matrix with distinct eigenvalues. Let $\Sigma = \mathbf{E} \operatorname{diag}(\mathbf{A}) \mathbf{E}^*$ be its eigendecomposition where \mathbf{E} is a unitary matrix of eigenvectors and $\mathbf{A} = (\lambda_1, \dots, \lambda_D)$ the vector of corresponding eigenvalues. We have for a Hermitian matrix \mathbf{A}*

$$\frac{\partial \mathbf{A}}{\partial \Sigma} [\mathbf{A}] = \operatorname{diag}(\bar{\mathbf{A}}) \quad \text{and} \quad \frac{\partial \mathbf{E}}{\partial \Sigma} [\mathbf{A}] = \mathbf{E} (\mathbf{J} \odot \bar{\mathbf{A}}) \quad (42)$$

where $\bar{\mathbf{A}} = \mathbf{E}^* \mathbf{A} \mathbf{E}$ and \mathbf{J} is the skew-symmetric matrix

$$\mathbf{J}_{ij} = \begin{cases} \frac{1}{\lambda_j - \lambda_i} & \text{if } i \neq j, \\ 0 & \text{otherwise.} \end{cases} \quad (43)$$

Proof (Proof of Proposition 1) Recall that $\exp \Sigma = \mathbf{E} \operatorname{diag}(e^{\mathbf{A}}) \mathbf{E}^*$. From Lemma 1, by applying chain rule, we have

$$\frac{\partial e^{\Sigma}}{\partial \Sigma} [\mathbf{A}] = \mathbf{E} (\mathbf{J} \odot \bar{\mathbf{A}}) \operatorname{diag}(e^{\mathbf{A}}) \mathbf{E}^* \quad (44)$$

$$+ \mathbf{E} \operatorname{diag}(e^{\mathbf{A}}) (\mathbf{J} \odot \bar{\mathbf{A}})^* \mathbf{E}^* \quad (45)$$

$$+ \mathbf{E} \operatorname{diag}(e^{\mathbf{A}}) \operatorname{diag}(\bar{\mathbf{A}}) \mathbf{E}^*. \quad (46)$$

We have for $i \neq j$

$$\left[(\mathbf{J} \odot \bar{\mathbf{A}}) \operatorname{diag}(e^{\mathbf{A}}) + \operatorname{diag}(e^{\mathbf{A}}) (\mathbf{J} \odot \bar{\mathbf{A}})^* \right]_{ij} \quad (47)$$

$$= \mathbf{J}_{ij} \bar{\mathbf{A}}_{ij} e^{\lambda_j} + e^{\lambda_i} \mathbf{J}_{ji} \bar{\mathbf{A}}_{ji} = \mathbf{J}_{ij} \bar{\mathbf{A}}_{ij} e^{\lambda_j} - e^{\lambda_i} \mathbf{J}_{ij} \bar{\mathbf{A}}_{ij} \quad (48)$$

$$= \mathbf{J}_{ij} \bar{\mathbf{A}}_{ij} (e^{\lambda_j} - e^{\lambda_i}) = \bar{\mathbf{A}}_{ij} \mathbf{G}_{ij} = [\mathbf{G} \odot \bar{\mathbf{A}}]_{ij}. \quad (49)$$

For $i = j$, since $\mathbf{J}_{ii} = 0$, we conclude the proof.

C Proof of Corollary 1

Proof We need to prove that for any two $D \times D$ Hermitian matrices \mathbf{A} and \mathbf{B} , we have

$$\left\langle \frac{\partial e^{\Sigma}}{\partial \Sigma} [\mathbf{A}], \mathbf{B} \right\rangle = \left\langle \mathbf{A}, \frac{\partial e^{\Sigma}}{\partial \Sigma} [\mathbf{B}] \right\rangle \quad (50)$$

for the matrix dot product $\langle \mathbf{X}, \mathbf{Y} \rangle = \operatorname{tr}[\mathbf{X} \mathbf{Y}^*]$. According to Proposition 1, this amounts to show

$$\operatorname{tr}(\mathbf{E}[\mathbf{G} \odot (\mathbf{E}^* \mathbf{A} \mathbf{E})] \mathbf{E}^* \mathbf{B}^*) = \operatorname{tr}(\mathbf{E}[\mathbf{G} \odot (\mathbf{E}^* \mathbf{B} \mathbf{E})] \mathbf{E}^* \mathbf{A}^*) \quad (51)$$

where \mathbf{E} and \mathbf{G} are defined from Σ as in Proposition 1. Denoting $\bar{\mathbf{A}} = \mathbf{E}^* \mathbf{A} \mathbf{E}$ and $\bar{\mathbf{B}} = \mathbf{E}^* \mathbf{B} \mathbf{E}$, this can be recast as

$$\operatorname{tr}((\mathbf{G} \odot \bar{\mathbf{A}}) \bar{\mathbf{B}}^*) = \operatorname{tr}((\mathbf{G} \odot \bar{\mathbf{B}}) \bar{\mathbf{A}}^*) \quad (52)$$

Expanding the left hand side allows us to conclude the proof as follows

$$\operatorname{tr}((\mathbf{G} \odot \bar{\mathbf{A}}) \bar{\mathbf{B}}^*) = \sum_{k=1}^D \sum_{l=1}^D (\mathbf{G} \odot \bar{\mathbf{A}})_{kl} (\bar{\mathbf{B}}^*)_{lk} \quad (53)$$

$$= \sum_{k=1}^D \sum_{l=1}^D \mathbf{G}_{kl} \bar{\mathbf{A}}_{kl} \bar{\mathbf{B}}_{kl} \quad (54)$$

$$= \sum_{k=1}^D \sum_{l=1}^D \mathbf{G}_{kl} \bar{\mathbf{B}}_{kl} (\bar{\mathbf{A}}^*)_{lk} \quad (55)$$

$$= \operatorname{tr}((\mathbf{G} \odot \bar{\mathbf{B}}) \bar{\mathbf{A}}^*). \quad (56)$$

D Hessian of the objective F

Proof (Proof of eq. (24)) Applying the chain rule to eq. (12) leads to

$$\mathbf{H} = \frac{\partial}{\partial \mathbf{x}} \left[\beta(\mathbf{x} - \mathbf{z}) + L \Theta \left(\mathbf{Id}_D - \left. \frac{\partial e^{\Sigma}}{\partial \Sigma} \right|_{-\Omega(\mathbf{x})}^* [e^{\Omega(\mathbf{y})}] \right) \right] \quad (57)$$

$$= \beta \mathbf{Id}_D - L \Theta \left(\frac{\partial}{\partial \mathbf{x}} \left(\left. \frac{\partial e^{\Sigma}}{\partial \Sigma} \right|_{-\Omega(\mathbf{x})}^* [e^{\Omega(\mathbf{y})}] \right) [\cdot] \right) \quad (58)$$

$$= \beta \mathbf{Id}_D - L \Theta \left(\left. \frac{\partial^2 e^{\Sigma}}{\partial \Sigma^2} \right|_{-\Omega(\mathbf{x})} [e^{\Omega(\mathbf{y})}, \frac{\partial -\Omega(\mathbf{x})}{\partial \mathbf{x}} [\cdot]] \right) \quad (59)$$

$$= \beta \mathbf{Id}_D + L \Theta \left(\left. \frac{\partial^2 e^{\Sigma}}{\partial \Sigma^2} \right|_{-\Omega(\mathbf{x})} [e^{\Omega(\mathbf{y})}, \Theta^* [\cdot]] \right). \quad (60)$$

It follows that

$$\mathbf{H} \mathbf{d} = \beta \mathbf{d} + L \Theta \left(\left. \frac{\partial^2 e^{\Sigma}}{\partial \Sigma^2} \right|_{-\Omega(\mathbf{x})} [e^{\Omega(\mathbf{y})}, \Theta^* [\mathbf{d}]] \right) \quad (61)$$

and thus, as $\|\mathbf{d}\|_2 = 1$, it follows that

$$\mathbf{d}^* \mathbf{H} \mathbf{d} = \beta + L \mathbf{d}^* \Theta \left(\left. \frac{\partial^2 e^{\Sigma}}{\partial \Sigma^2} \right|_{-\Omega(\mathbf{x})} [e^{\Omega(\mathbf{y})}, \Theta^* [\mathbf{d}]] \right) \quad (62)$$

$$= \beta + L \left\langle \mathbf{d}, \Theta \left(\left. \frac{\partial^2 e^{\Sigma}}{\partial \Sigma^2} \right|_{-\Omega(\mathbf{x})} [e^{\Omega(\mathbf{y})}, \Theta^* [\mathbf{d}]] \right) \right\rangle \quad (63)$$

$$= \beta + L \left\langle \Theta^*[\mathbf{d}], \frac{\partial^2 e^{\Sigma}}{\partial \Sigma^2} \Big|_{-\Omega(\mathbf{x})} [e^{\Omega(\mathbf{y})}, \Theta^*[\mathbf{d}]] \right\rangle \quad (64)$$

which concludes the proof.

E Proof of Proposition 2

Proof (Proof of Proposition 2) The second directional derivative can be defined from the adjoint of the directional derivative as

$$\frac{\partial^2 e^{\Sigma}}{\partial \Sigma^2} [\mathbf{A}, \mathbf{B}] = \frac{\partial}{\partial \Sigma} \left(\frac{\partial e^{\Sigma}}{\partial \Sigma} \Big|_{\Sigma}^* [\mathbf{A}] \right) [\mathbf{B}]. \quad (65)$$

By virtue of Corollary 1, we have $\frac{\partial e^{\Sigma}}{\partial \Sigma} \Big|_{\Sigma}^* = \frac{\partial e^{\Sigma}}{\partial \Sigma} \Big|_{\Sigma}$, and then from Proposition 1 it follows that

$$\frac{\partial^2 e^{\Sigma}}{\partial \Sigma^2} [\mathbf{A}, \mathbf{B}] = \frac{\partial}{\partial \Sigma} (\mathbf{E}[\mathbf{G} \odot \bar{\mathbf{A}}] \mathbf{E}^*) [\mathbf{B}]. \quad (66)$$

In order to apply the chain rule on $\mathbf{E}[\mathbf{G} \odot \bar{\mathbf{A}}] \mathbf{E}^*$, let us first rewrite \mathbf{J} and \mathbf{G} in Lemma 1 and Proposition 1 as

$$\mathbf{J} = (\mathbf{1}_D \mathbf{1}_D^* - \mathbf{Id}_D) \odot (\mathbf{Id}_D + \mathbf{1}_D \mathbf{A}^* - \mathbf{A} \mathbf{1}_D^*) \quad (67)$$

$$\mathbf{G} = \text{diag}(e^{\Lambda}) - (e^{\Lambda} \mathbf{1}_D^* - \mathbf{1}_D e^{\Lambda^*}) \odot \mathbf{J} \quad (68)$$

where $\mathbf{1}_D$ is a D dimensional column vector of ones, \odot denotes the element-wise division and e^{Λ^*} must be understood as the row vector $(e^{\Lambda})^*$. From Lemma 1, we have for a Hermitian matrix \mathbf{B}

$$\frac{\partial}{\partial \Sigma} \text{diag}(e^{\Lambda}) [\mathbf{B}] = \text{diag}(e^{\Lambda}) \odot \bar{\mathbf{B}} \quad (69)$$

$$\begin{aligned} \frac{\partial}{\partial \Sigma} (e^{\Lambda} \mathbf{1}_D^* - \mathbf{1}_D e^{\Lambda^*}) [\mathbf{B}] \\ = \text{diag}(e^{\Lambda}) \text{diag}(\bar{\mathbf{B}}) \mathbf{1}_D^* - \mathbf{1}_D \text{diag}(\bar{\mathbf{B}})^* \text{diag}(e^{\Lambda}) \end{aligned} \quad (70)$$

$$\frac{\partial}{\partial \Sigma} (\mathbf{Id}_D + \mathbf{A} \mathbf{1}_D^* - \mathbf{1}_D \mathbf{A}^*) [\mathbf{B}] = \text{diag}(\bar{\mathbf{B}}) \mathbf{1}_D^* - \mathbf{1}_D \text{diag}(\bar{\mathbf{B}})^* \quad (71)$$

where $\bar{\mathbf{B}} = \mathbf{E}^* \mathbf{B} \mathbf{E}$. By application of the chain rule, we get

$$\frac{\partial \mathbf{J}}{\partial \Sigma} [\mathbf{B}] = -(\text{diag}(\bar{\mathbf{B}}) \mathbf{1}_D^* - \mathbf{1}_D \text{diag}(\bar{\mathbf{B}})^*) \odot \mathbf{J} \odot \mathbf{J} \quad (72)$$

$$\begin{aligned} \frac{\partial \mathbf{G}}{\partial \Sigma} [\mathbf{B}] &= \text{diag}(e^{\Lambda}) \odot \bar{\mathbf{B}} \\ &- \left[\text{diag}(e^{\Lambda}) \text{diag}(\bar{\mathbf{B}}) \mathbf{1}_D^* - \mathbf{1}_D \text{diag}(\bar{\mathbf{B}})^* \text{diag}(e^{\Lambda}) \right. \\ &\quad \left. - \mathbf{G} \odot (\text{diag}(\bar{\mathbf{B}}) \mathbf{1}_D^* - \mathbf{1}_D \text{diag}(\bar{\mathbf{B}})^*) \right] \odot \mathbf{J} \end{aligned} \quad (73)$$

where we used that $(e^{\Lambda} \mathbf{1}_D^* - \mathbf{1}_D e^{\Lambda^*}) \odot \mathbf{J} = -\mathbf{G} \odot \mathbf{J}$. Let \mathbf{A} be a Hermitian matrix and $\bar{\mathbf{A}} = \mathbf{E}^* \mathbf{A} \mathbf{E}$, by Lemma 1, we have

$$\frac{\partial \bar{\mathbf{A}}}{\partial \Sigma} [\mathbf{B}] = \bar{\mathbf{A}} (\mathbf{J} \odot \bar{\mathbf{B}}) - (\mathbf{J} \odot \bar{\mathbf{B}}) \bar{\mathbf{A}}. \quad (74)$$

We are now equipped to apply the chain rule to $\mathbf{E}[\mathbf{G} \odot \bar{\mathbf{A}}] \mathbf{E}^*$ in the direction \mathbf{B} , which leads us to

$$\frac{\partial^2 e^{\Sigma}}{\partial \Sigma^2} [\mathbf{A}, \mathbf{B}] = \mathbf{E}[\mathbf{F}(\bar{\mathbf{A}}, \bar{\mathbf{B}})] \mathbf{E}^* \quad (75)$$

$$\text{with } \mathbf{F}(\bar{\mathbf{A}}, \bar{\mathbf{B}}) = \mathbf{F}_1 + \mathbf{F}_2 + \mathbf{F}_3 \quad (76)$$

$$\text{and } \mathbf{F}_1 = (\mathbf{J} \odot \bar{\mathbf{B}}) (\mathbf{G} \odot \bar{\mathbf{A}}) - (\mathbf{G} \odot \bar{\mathbf{A}}) (\mathbf{J} \odot \bar{\mathbf{B}}) \quad (77)$$

$$\mathbf{F}_2 = \mathbf{G} \odot [\bar{\mathbf{A}} (\mathbf{J} \odot \bar{\mathbf{B}}) - (\mathbf{J} \odot \bar{\mathbf{B}}) \bar{\mathbf{A}}] \quad (78)$$

$$\mathbf{F}_3 = \frac{\partial \mathbf{G}}{\partial \Sigma} [\mathbf{B}] \odot \bar{\mathbf{A}}. \quad (79)$$

We have for all $1 \leq i \leq D$ and $1 \leq j \leq D$

$$[\mathbf{F}_1]_{ij} = \sum_{k=1}^D \mathbf{J}_{jk} \mathbf{G}_{ik} \bar{\mathbf{A}}_{ik} \bar{\mathbf{B}}_{jk}^* + \mathbf{J}_{ik} \mathbf{G}_{jk} \bar{\mathbf{B}}_{ik} \bar{\mathbf{A}}_{jk}^* \quad (80)$$

$$\text{and } [\mathbf{F}_2]_{ij} = -\mathbf{G}_{ij} \sum_{k=1}^D \mathbf{J}_{jk} \bar{\mathbf{A}}_{ik} \bar{\mathbf{B}}_{jk}^* + \mathbf{J}_{ik} \bar{\mathbf{B}}_{ik} \bar{\mathbf{A}}_{jk}^*. \quad (81)$$

Hence, we get

$$\begin{aligned} [\mathbf{F}_1 + \mathbf{F}_2]_{ij} &= \sum_{k=1}^D \mathbf{J}_{jk} (\mathbf{G}_{ik} - \mathbf{G}_{ij}) \bar{\mathbf{A}}_{ik} \bar{\mathbf{B}}_{jk}^* \\ &+ \sum_{k=1}^D \mathbf{J}_{ik} (\mathbf{G}_{jk} - \mathbf{G}_{ij}) \bar{\mathbf{B}}_{ik} \bar{\mathbf{A}}_{jk}^*. \end{aligned} \quad (82)$$

• Assume $i \neq j$. We have

$$[\mathbf{F}_3]_{ij} = -[\mathbf{G}_{ii} \bar{\mathbf{B}}_{ii} - \mathbf{G}_{jj} \bar{\mathbf{B}}_{jj}] \mathbf{J}_{ij} \bar{\mathbf{A}}_{ij} \quad (83)$$

$$\begin{aligned} &- \mathbf{G}_{ij} (\bar{\mathbf{B}}_{ii} - \bar{\mathbf{B}}_{jj}) \mathbf{J}_{ij} \bar{\mathbf{A}}_{ij} \\ &= \mathbf{J}_{ij} (\mathbf{G}_{jj} - \mathbf{G}_{ij}) \bar{\mathbf{B}}_{ij} \bar{\mathbf{A}}_{ij} \\ &- \mathbf{J}_{ij} (\mathbf{G}_{ii} - \mathbf{G}_{ij}) \bar{\mathbf{A}}_{ij} \bar{\mathbf{B}}_{ii}. \end{aligned} \quad (84)$$

For $k \neq i$ and $k \neq j$, we have

$$\mathbf{J}_{ik} (\mathbf{G}_{jk} - \mathbf{G}_{ij}) - \mathbf{J}_{ij} (\mathbf{G}_{jk} - \mathbf{G}_{ik}) \quad (85)$$

$$= \frac{e^{\lambda_j} - e^{\lambda_k}}{(\lambda_j - \lambda_k)(\lambda_k - \lambda_i)} - \frac{e^{\lambda_i} - e^{\lambda_j}}{(\lambda_i - \lambda_j)(\lambda_k - \lambda_i)} \quad (86)$$

$$\begin{aligned} &- \frac{e^{\lambda_j} - e^{\lambda_k}}{(\lambda_j - \lambda_k)(\lambda_j - \lambda_i)} + \frac{e^{\lambda_i} - e^{\lambda_k}}{(\lambda_i - \lambda_k)(\lambda_j - \lambda_i)} \\ &= \frac{(\lambda_i - \lambda_j)(e^{\lambda_j} - e^{\lambda_k}) - (\lambda_j - \lambda_k)(e^{\lambda_i} - e^{\lambda_j})}{(\lambda_i - \lambda_j)(\lambda_j - \lambda_k)(\lambda_k - \lambda_i)} \end{aligned} \quad (87)$$

$$\begin{aligned} &+ \frac{(\lambda_k - \lambda_i)(e^{\lambda_j} - e^{\lambda_k}) + (\lambda_j - \lambda_k)(e^{\lambda_i} - e^{\lambda_k})}{(\lambda_i - \lambda_j)(\lambda_j - \lambda_k)(\lambda_k - \lambda_i)} \\ &= \frac{e^{\lambda_i} (\lambda_k - \lambda_j + \lambda_j - \lambda_k)}{(\lambda_i - \lambda_j)(\lambda_j - \lambda_k)(\lambda_k - \lambda_i)} \end{aligned} \quad (88)$$

$$\begin{aligned} &+ \frac{e^{\lambda_j} (\lambda_i - \lambda_j + \lambda_j - \lambda_k + \lambda_k - \lambda_i)}{(\lambda_i - \lambda_j)(\lambda_j - \lambda_k)(\lambda_k - \lambda_i)} \\ &+ \frac{e^{\lambda_k} (\lambda_j - \lambda_i + \lambda_i - \lambda_k + \lambda_k - \lambda_j)}{(\lambda_i - \lambda_j)(\lambda_j - \lambda_k)(\lambda_k - \lambda_i)} \\ &= 0. \end{aligned} \quad (89)$$

Similarly, we have $\mathbf{J}_{jk} (\mathbf{G}_{ik} - \mathbf{G}_{ij}) = \mathbf{J}_{ij} (\mathbf{G}_{jk} - \mathbf{G}_{ik})$. Hence, we get the following

$$\begin{aligned} [\mathbf{F}_1 + \mathbf{F}_2]_{ij} &= \sum_{k \neq i, k \neq j} \mathbf{J}_{ij} (\mathbf{G}_{jk} - \mathbf{G}_{ik}) (\bar{\mathbf{A}}_{ik} \bar{\mathbf{B}}_{jk}^* + \bar{\mathbf{B}}_{ik} \bar{\mathbf{A}}_{jk}^*) \\ &- \mathbf{J}_{ij} (\mathbf{G}_{ii} - \mathbf{G}_{ij}) \bar{\mathbf{A}}_{ii} \bar{\mathbf{B}}_{ji}^* \\ &+ \mathbf{J}_{ij} (\mathbf{G}_{jj} - \mathbf{G}_{ij}) \bar{\mathbf{B}}_{ij} \bar{\mathbf{A}}_{jj}^*. \end{aligned} \quad (90)$$

It follows that

$$\begin{aligned} [\mathbf{F}_1 + \mathbf{F}_2 + \mathbf{F}_3]_{ij} \\ = \sum_{k \neq i, k \neq j} \mathbf{J}_{ij} (\mathbf{G}_{jk} - \mathbf{G}_{ik}) (\bar{\mathbf{A}}_{ik} \bar{\mathbf{B}}_{jk}^* + \bar{\mathbf{B}}_{ik} \bar{\mathbf{A}}_{jk}^*) \end{aligned} \quad (91)$$

$$\begin{aligned}
& + \mathbf{J}_{ij}(\mathbf{G}_{ji} - \mathbf{G}_{ii})(\bar{\mathbf{A}}_{ii}\bar{\mathbf{B}}_{ji}^* + \bar{\mathbf{B}}_{ii}\bar{\mathbf{A}}_{ji}^*) \\
& + \mathbf{J}_{ij}(\mathbf{G}_{jj} - \mathbf{G}_{ij})(\bar{\mathbf{A}}_{ij}\bar{\mathbf{B}}_{jj}^* + \bar{\mathbf{B}}_{ij}\bar{\mathbf{A}}_{jj}^*) \\
= & \sum_{k=1}^D \underbrace{\mathbf{J}_{ij}(\mathbf{G}_{jk} - \mathbf{G}_{ik})}_{\varphi_{i,j,k}}(\bar{\mathbf{A}}_{ik}\bar{\mathbf{B}}_{jk}^* + \bar{\mathbf{B}}_{ik}\bar{\mathbf{A}}_{jk}^*). \quad (92)
\end{aligned}$$

• Now assume that $i = j$. We have $[\mathbf{F}_3]_{ii} = \mathbf{G}_{ii}\bar{\mathbf{A}}_{ii}\bar{\mathbf{B}}_{ii}$. It follows that

$$[\mathbf{F}_1 + \mathbf{F}_2 + \mathbf{F}_3]_{ii} \quad (93)$$

$$= \sum_{k=1}^D \mathbf{J}_{ik}(\mathbf{G}_{ik} - \mathbf{G}_{ii})(\bar{\mathbf{A}}_{ik}\bar{\mathbf{B}}_{ik}^* + \bar{\mathbf{B}}_{ik}\bar{\mathbf{A}}_{ik}^*) \quad (94)$$

$$+ \mathbf{G}_{ii}\bar{\mathbf{A}}_{ii}\bar{\mathbf{B}}_{ii} \quad (95)$$

$$= \sum_{k \neq i} \underbrace{\mathbf{J}_{ik}(\mathbf{G}_{ik} - \mathbf{G}_{ii})}_{\varphi_{i,i,k}}(\bar{\mathbf{A}}_{ik}\bar{\mathbf{B}}_{ik}^* + \bar{\mathbf{B}}_{ik}\bar{\mathbf{A}}_{ik}^*) \quad (96)$$

$$+ \underbrace{\frac{1}{2}\mathbf{G}_{ii}(\bar{\mathbf{A}}_{ii}\bar{\mathbf{B}}_{ii}^* + \bar{\mathbf{B}}_{ii}\bar{\mathbf{A}}_{ii}^*)}_{\varphi_{i,i,i}}$$

which concludes the proof.

References

- Aubert, G., Aujol, J.F.: A variational approach to removing multiplicative noise. *SIAM journal on applied mathematics* **68**(4), 925–946 (2008)
- Baqué, R., du Plessis, O.R., Castet, N., Fromage, P., Martinot-Lagarde, J., Nouvel, J.F., Oriot, H., Angelliaume, S., Brigui, F., Cantalloube, H., et al.: SETHI/RAMSES-NG: New performances of the flexible multi-spectral airborne remote sensing research platform. In: 2017 European Radar Conference (EURAD), pp. 191–194. IEEE (2017)
- Boyd, S., Parikh, N., Chu, E.: Distributed optimization and statistical learning via the alternating direction method of multipliers. Now Publishers Inc (2011)
- Buades, A., Coll, B., Morel, J.M.: A review of image denoising algorithms, with a new one. *Multiscale Modeling & Simulation* **4**(2), 490–530 (2005)
- Candes, E.J., Sing-Long, C.A., Trzasko, J.D.: Unbiased risk estimates for singular value thresholding and spectral estimators. *IEEE Transactions on Signal Processing* **61**(19), 4643–4657 (2013)
- Chan, S.H., Wang, X., Elgendy, O.A.: Plug-and-play ADMM for image restoration: Fixed-point convergence and applications. *IEEE Transactions on Computational Imaging* **3**(1), 84–98 (2016)
- Chen, J., Chen, Y., An, W., Cui, Y., Yang, J.: Nonlocal filtering for polarimetric SAR data: A pretest approach. *IEEE Transactions on Geoscience and Remote Sensing* **49**(5), 1744–1754 (2010)
- Chierchia, G., Cozzolino, D., Poggi, G., Verdoliva, L.: SAR image despeckling through convolutional neural networks. In: 2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS), pp. 5438–5441. IEEE (2017)
- Combettes, P.L., Pesquet, J.C.: Proximal splitting methods in signal processing. In: Fixed-point algorithms for inverse problems in science and engineering, pp. 185–212. Springer (2011)
- Dabov, K., Foi, A., Katkovnik, V., Egiazarian, K.: Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on Image Processing* **16**(8), 2080–2095 (2007)
- Dalsasso, E., Denis, L., Tupin, F.: Sar2sar: a semi-supervised despeckling algorithm for sar images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* (2021)
- Dalsasso, E., Yang, X., Denis, L., Tupin, F., Yang, W.: SAR Image Despeckling by Deep Neural Networks: from a pre-trained model to an end-to-end training strategy. *Remote Sensing* **12**(16), 2636 (2020)
- Deledalle, C.A., Denis, L., Tabti, S., Tupin, F.: MuLoG, or How to apply Gaussian denoisers to multi-channel SAR speckle reduction? *IEEE Transactions on Image Processing* **26**(9), 4389–4403 (2017)
- Deledalle, C.A., Denis, L., Tupin, F.: Iterative weighted maximum likelihood denoising with probabilistic patch-based weights. *IEEE Transactions on Image Processing* **18**(12), 2661–2672 (2009)
- Deledalle, C.A., Denis, L., Tupin, F.: NL-InSAR: Nonlocal interferogram estimation. *IEEE Transactions on Geoscience and Remote Sensing* **49**(4), 1441–1452 (2010)
- Deledalle, C.A., Denis, L., Tupin, F., Reigber, A., Jäger, M.: NL-SAR: A unified nonlocal framework for resolution-preserving (Pol)(In) SAR denoising. *IEEE Transactions on Geoscience and Remote Sensing* **53**(4), 2021–2038 (2014)
- Deledalle, C.A., Vaiter, S., Peyré, G., Fadili, J.M., Dossal, C.: Risk estimation for matrix recovery with spectral regularization. In: ICML’2012 workshop on Sparsity, Dictionaries and Projections in Machine Learning and Signal Processing (2012)
- Denis, L., Tupin, F., Darbon, J., Sigelle, M.: SAR image regularization with fast approximate discrete minimization. *IEEE Transactions on Image Processing* **18**(7), 1588–1600 (2009)
- Durand, S., Fadili, J., Nikolova, M.: Multiplicative noise removal using L1 fidelity on frame coefficients. *Journal of Mathematical Imaging and Vision* **36**(3), 201–226 (2010)
- Even, M., Schulz, K.: InSAR deformation analysis with distributed scatterers: A review complemented by new advances. *Remote Sensing* **10**(5), 744 (2018)
- Goodman, J.: Some fundamental properties of speckle. *Journal Optical Society of America* **66**(11), 1145–1150 (1976)
- Goodman, J.W.: Statistical properties of laser speckle patterns. In: *Laser speckle and related phenomena*, pp. 9–75. Springer (1975)
- Knaus, C., Zwicker, M.: Dual-domain image denoising. In: 2013 IEEE International Conference on Image Processing, pp. 440–444. IEEE (2013)
- Laumont, R., De Bortoli, V., Almansa, A., Delon, J., Durmus, A., Pereyra, M.: Bayesian imaging using Plug & Play priors: when Langevin meets Tweedie. arXiv preprint arXiv:2103.04715 (2021)
- Lee, J.S.: Digital image smoothing and the sigma filter. *Computer vision, graphics, and image processing* **24**(2), 255–269 (1983)
- Lewis, A.S., Sendov, H.S.: Twice differentiable spectral functions. *SIAM Journal on Matrix Analysis and Applications* **23**(2), 368–386 (2001)
- Lopes, A., Touzi, R., Nezry, E.: Adaptive speckle filters and scene heterogeneity. *IEEE transactions on Geoscience and Remote Sensing* **28**(6), 992–1000 (1990)
- Molini, A.B., Valsesia, D., Fracastoro, G., Magli, E.: Speckle2Void: Deep Self-Supervised SAR Despeckling with Blind-Spot Convolutional Neural Networks. *IEEE Transactions on Geoscience and Remote Sensing* (2021)
- Monga, V., Li, Y., Eldar, Y.C.: Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing. *IEEE Signal Processing Magazine* **38**(2), 18–44 (2021)

30. Moreira, A., Prats-Iraola, P., Younis, M., Krieger, G., Hajnsek, I., Papathanassiou, K.P.: A tutorial on synthetic aperture radar. *IEEE Geoscience and remote sensing magazine* **1**(1), 6–43 (2013)
31. Parameswaran, S., Deledalle, C.A., Denis, L., Nguyen, T.Q.: Accelerating GMM-Based Patch Priors for Image Restoration: Three Ingredients for a 100× Speed-Up. *IEEE Transactions on Image Processing* **28**(2), 687–698 (2018)
32. Parrilli, S., Poderico, M., Angelino, C.V., Verdoliva, L.: A nonlocal SAR image denoising algorithm based on LLMMSE wavelet shrinkage. *IEEE Transactions on Geoscience and Remote Sensing* **50**(2), 606–616 (2011)
33. Reehorst, E.T., Schniter, P.: Regularization by denoising: Clarifications and new interpretations. *IEEE transactions on computational imaging* **5**(1), 52–67 (2018)
34. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena* **60**(1-4), 259–268 (1992)
35. Steidl, G., Teuber, T.: Removing multiplicative noise by Douglas-Rachford splitting methods. *Journal of Mathematical Imaging and Vision* **36**(2), 168–184 (2010)
36. Touzi, R., Lopes, A., Bruniquel, J., Vachon, P.W.: Coherence estimation for SAR imagery. *IEEE Transactions on Geoscience and Remote Sensing* **37**(1), 135–149 (1999)
37. Vasile, G., Trouvé, E., Lee, J.S., Buzuloiu, V.: Intensity-driven adaptive-neighborhood technique for polarimetric and interferometric SAR parameters estimation. *IEEE Transactions on Geoscience and Remote Sensing* **44**(6), 1609–1621 (2006)
38. Xie, H., Pierce, L.E., Ulaby, F.T.: Statistical properties of logarithmically transformed speckle. *IEEE Transactions on Geoscience and Remote Sensing* **40**(3), 721–727 (2002)
39. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Transactions on Image Processing* **26**(7), 3142–3155 (2017)
40. Zhang, K., Zuo, W., Gu, S., Zhang, L.: Learning deep CNN denoiser prior for image restoration. In: *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 3929–3938 (2017)
41. Zhang, Y., Zhu, D.: Height retrieval in postprocessing-based VideoSAR image sequence using shadow information. *IEEE Sensors Journal* **18**(19), 8108–8116 (2018)
42. Zhao, W., Deledalle, C.A., Denis, L., Maître, H., Nicolas, J.M., Tupin, F.: Ratio-based multitemporal SAR images denoising: RABASAR. *IEEE Transactions on Geoscience and Remote Sensing* **57**(6), 3552–3565 (2019)