



HAL
open science

MuLoG, or How to apply Gaussian denoisers to multi-channel SAR speckle reduction?

Charles-Alban Deledalle, Loïc Denis, Sonia Tabti, Florence Tupin

► **To cite this version:**

Charles-Alban Deledalle, Loïc Denis, Sonia Tabti, Florence Tupin. MuLoG, or How to apply Gaussian denoisers to multi-channel SAR speckle reduction?. 2016. hal-01388858v1

HAL Id: hal-01388858

<https://hal.science/hal-01388858v1>

Preprint submitted on 27 Oct 2016 (v1), last revised 12 Apr 2017 (v2)

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.

MuLoG, or How to apply Gaussian denoisers to multi-channel SAR speckle reduction?

Charles-Alban Deledalle, Loïc Denis, Sonia Tabti, Florence Tupin

Abstract—Speckle reduction is a longstanding topic in synthetic aperture radar (SAR) imaging. Since most current and planned SAR imaging satellites operate in polarimetric, interferometric or tomographic modes, SAR images are multi-channel and speckle reduction techniques must jointly process all channels to recover polarimetric and interferometric information. The distinctive nature of SAR signal (complex-valued, corrupted by multiplicative fluctuations) calls for the development of specialized methods for speckle reduction. Image denoising is a very active topic in image processing with a wide variety of approaches and many denoising algorithms available, almost always designed for additive Gaussian noise suppression. This paper proposes a general scheme, called MuLoG (MULTi-channel LOfarithm with Gaussian denoising), to include such Gaussian denoisers within a multi-channel SAR speckle reduction technique. A new family of speckle reduction algorithms can thus be obtained, benefiting from the ongoing progress in Gaussian denoising, and offering several speckle reduction results often displaying method-specific artifacts that can be dismissed by comparison between results.

Index Terms—SAR, speckle, variance stabilization, ADMM, Wishart distribution

I. INTRODUCTION

Synthetic aperture radar (SAR) imaging is a widely used technique for earth observation. It offers complementary information to the more common optical imaging. Among its distinctive features, one may cite its all-weather and day-and-night imaging capabilities [43], the interferometric configurations that give access to 3-D reconstructions for digital elevation models [48] and displacement estimation [42], [20] or the polarimetric and tomographic modes that give access to estimates of the biomass in forested areas [34].

Coherent combination of several radar echoes within each resolution cell results in interferences and the well-known *speckle phenomenon* [26]. Due to speckle, regions with homogeneous radar properties display strong fluctuations in SAR images. Direct estimation of the reflectivity, the interferometric phase or polarimetric properties is unusable given its prohibitively large variance. Speckle reduction is thus a longstanding topic in SAR imagery.

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

L. Denis is with 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

S. Tabti is with DIETI, University Federico II of Naples, Italy. This work has been done while S. Tabti was with LTCI, Telecom ParisTech, Université Paris Saclay, 46 rue Barrault, 75013 Paris, France, and next with Université Paris Descartes, MAP5, CNRS UMR 8145, Paris, France. email: sonia.tabti@telecom-paristech.fr

F. Tupin is with LTCI, Telecom ParisTech, Université Paris Saclay, 46 rue Barrault, 75013 Paris, France, email: florence.tupin@telecom-paristech.fr

Since speckle statistics departs from the additive Gaussian noise model widely used for processing optical images, a whole line of denoising methods have been designed specifically for speckle reduction. Restoration of intensity images (i.e., single-channel SAR images) remains the first and most studied case. Many different schemes have been proposed in this context, see the recent reviews [3] and [12]. Among the different possible strategies, *selection-based methods* use various criteria to identify homogeneous collections of neighboring pixels. Pixels may be selected by locally choosing the best neighborhood within a fixed set of oriented windows [36]. Another approach consists of selecting connected pixels by region growing [54]. Patch-comparison has been shown to provide a robust way to get a (weighted) selection of similar pixels [15], [12]. *Variational methods* formulate the estimation problem as an optimization problem depending on the whole image. The objective function that is optimized is composed of two terms: a data-fitting term and a regularization term. Due to the non-Gaussian distribution of the intensity of SAR images, the data-fitting term differs from the usual sum of squared differences that arises from a Gaussian assumption. Several regularization terms have been considered in the literature, e.g., total variation (TV) [49], [51], [17], [52], [7], curvelets [18], Gaussian mixture models [53]. Yet another strategy is to transform the data so the noise becomes additive and approximately Gaussian and stationary by using a *homomorphic transform*. As further discussed in Sec.II, the log-transformed intensity is approximately Gaussian distributed and can thus be restored using a Gaussian denoiser, e.g., based on wavelet thresholding [56], [1], [6] or on patch redundancy [41].

Not all these speckle reduction methods generalize well to multi-channel SAR images. *Selection-based methods* require to extend the homogeneity or similarity criteria to multi-variate data. This can be done by using only part of the information, for example the span [54] or the scattering properties [37], or by exploiting the whole covariance matrices [8], [16]. Once relevant pixels have been selected, estimation of polarimetric/interferometric information is straightforward, e.g., by using a linear mean square error approach [35], [16] or a weighted maximum likelihood estimator [13], [14]. Extension of *variational methods* to multi-channel SAR data is more challenging. Direct formulation of the objective function on the covariance matrices raises several problems: (i) computational complexity due to the non-convexity of the data-fitting term; (ii) difficulty to express regularity properties of the complex-valued terms of the covariance matrices; (iii) non-stationary variance of speckle that leads to over/under-smoothing in some areas. These difficulties explain that very

few works were conducted in this direction, with the exception of recent works on multi-channel TV regularization approaches [39], [45], [46], [47]. Finally, while the *homomorphic transform* approach is well understood on intensity images, no variance stabilization transform is known for multi-channel images.

A generic methodology to apply denoising methods from the “additive Gaussian noise world” to multi-channel SAR data is definitely lacking. This paper attempts to fill this gap. The contributions of this paper are the following:

- 1) to provide a generic method, called MuLoG, to embed a Gaussian denoiser in a speckle reduction filter;
- 2) to apply as well to single channel or to multi-channel SAR images;
- 3) to produce estimates of the complex covariance matrix that capture all polarimetric and/or interferometric information,
- 4) to better account for speckle statistics than other methods based on a homomorphic transform¹;
- 5) to require no parameter tuning other than possibly within the Gaussian denoiser.

We introduce our generic methodology by first considering (single-channel) intensity images (Sec. II), then the extension to multi-channel SAR images (Sec. III). We discuss implementation issues (Sec. IV) before illustrating the proposed methodology with several Gaussian denoisers and different types of multi-channel SAR images (Sec. V).

II. SPECKLE REDUCTION FOR SAR INTENSITY IMAGES

A. Statistics of univariate SAR images

a) Intensity: Univariate SAR images are by nature complex-valued and only the modulus (*a.k.a.*, the amplitude) is informative. The square of the modulus (*a.k.a.*, the intensity) is nevertheless easier to manipulate, and according to Goodman’s model [26], it follows a gamma distribution $\mathcal{G}(R; L)$ with a probability density given by

$$p_I(I|R) = \frac{L^L I^{L-1}}{\Gamma(L) R^L} \exp\left(-L \frac{I}{R}\right), \quad (1)$$

where $I \in \mathbb{R}_+$ is the observed intensity, $R \in \mathbb{R}_+$ is the underlying reflectivity (related to the radar cross-section), and $L > 0$ is the number of looks. The intensity I can be decomposed as a product of the reflectivity R and of a speckle component S distributed under a standard gamma distribution ($S \sim \mathcal{G}(1; L)$):

$$I = R \times S, \quad \mathbb{E}[I] = R \quad \text{and} \quad \text{Var}[I] = \frac{R^2}{L}. \quad (2)$$

As the variance depends on the expectation, fluctuations are said to be signal dependent. The top graph in Fig.2(a) illustrates how a simple rectangle signal (solid curve) gets corrupted by speckle: the gray area shows values between the first and third quartiles, the dots represent the expectation

¹note that these other methods are only applicable to intensity SAR images (i.e., single-channel)

and the dashed line a single noisy realization. The signal-dependent nature of the fluctuations can be observed: the difference between first and third quartiles is larger when the underlying signal values are high. Last but not least, the gamma distribution has a heavier right-tail characterizing the typical bright outliers observed in SAR intensity images.

b) Logarithm: The log-transform $y = \log I \in \mathbb{R}$ is often employed to convert multiplicative fluctuations to additive ones. From the gamma distribution (1), y follows the Fisher-Tippett distribution defined as

$$p_y(y|x) = \frac{L^L}{\Gamma(L)} e^{L(y-x)} \exp(-Le^{y-x}), \quad (3)$$

where $x = \log R \in \mathbb{R}$. The Fisher-Tippett distribution, denoted by $\mathcal{FT}(x; L)$, models additive corruptions as [57]

$$y = x + s, \quad (4)$$

$$\mathbb{E}[y] = x - \log L + \Psi(L), \quad (5)$$

$$\text{and} \quad \text{Var}[y] = \Psi(1, L), \quad (6)$$

where $s \sim \mathcal{FT}(0; L)$. The log transform stabilizes the variance, *i.e.*, the fluctuations are made signal independent. Equation (5) shows that the noise has a non-zero mean. If noise is assumed zero-mean during a speckle-reduction step performed on log-transformed data, a subsequent debiasing step is then necessary.

The bottom graph in Fig.2(a) displays the intensities in a log scale. By comparing with the top curve, variance stabilization can be observed, as well as the bias between the expectation of log-transformed intensities (black dots) and ground truth signal (solid curve).

B. Homomorphic approach

We now consider I and R as images, such that $I \in \mathbb{R}_+^n$, and $R \in \mathbb{R}_+^n$, where n is the number of pixels. Let $y \in \mathbb{R}^n$ and $x \in \mathbb{R}^n$ be the entry-wise logarithm of I and R respectively, *i.e.*, such that $y_k = \log I_k$ and $x_k = \log R_k$. A classical approach to estimate R is thus to approach the Fisher-Tippett distribution by a non-zero mean Gaussian distribution, which leads to the following estimation procedure

$$\hat{x} = f_{\Psi(1,L)}(y) + \underbrace{(\log L - \Psi(L)) \mathbf{1}_n}_{\text{debiasing}}, \quad (7)$$

where $f_{\sigma^2} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a denoiser for images contaminated by zero-mean additive white Gaussian noise $\mathcal{N}(0; \sigma^2 \text{Id})$. Typically f can be a regularized least-square solver expressed as

$$f_{\sigma^2}(y) \in \underset{x \in \mathbb{R}^n}{\text{argmin}} \frac{1}{2\sigma^2} \|y - x\|^2 + \mathcal{R}(x), \quad (8)$$

where, within the Maximum A Posteriori (MAP) framework, $\mathcal{R}(x) = -\log p_x(x)$ is a prior term enforcing some regularity on the solution. Finally, the estimate \hat{R} is defined as $\hat{R}_k = \exp \hat{x}_k$ for all pixel index k . This process is summarized on Fig.1, top row. We illustrate the restored 1D signals obtained by applying L2+TV minimization [50] within a homomorphic procedure on Fig.2(b).

The homomorphic approach has been extensively used, *e.g.*, for wavelet prior in [56], for patch-based priors as in KSVD




Fig. 1. Several approaches for speckle reduction discussed in this paper. The proposed scheme generalizes the variational approaches based on a log-transform to multi-channel SAR images. Similarly to homomorphic or variational approaches on log-transformed data, it embeds a Gaussian denoising step.

[21] or for non-local filtering [41]. While for large values of L the Fisher-Tippett distribution approaches the Gaussian distribution, the two distributions differ for low values of L . In particular, the Fisher-Tippett is asymmetrical with a heavier left-tail. Hence, this approximation often leads to remaining dark stains on the resulting images.

C. Variational approach

Rather than applying the homomorphic approach in order to reduce the original problem to the Gaussian case, an alternative proposed in [4], [17] is to consider directly the gamma distribution of $I \in \mathbb{R}_+^n$ leading to a MAP solver of the form (see also Fig.1 second row):

$$\hat{R} \in \underset{R \in \mathbb{R}_+^n}{\operatorname{argmin}} -\log p_I(I|R) + \mathcal{R}(R), \quad (9)$$

$$\text{where } -\log p_I(I|R) = L \sum_{k=1}^n \log R_k + \frac{I_k}{R_k} + \text{Cst.}$$

Not only is this minimization constrained to positive-valued images, but the objective function is also non convex. This

non convexity enforces some robustness to the bright outliers originating from the right tail of the gamma distribution. As a consequence, when standard iterative solvers are used (*e.g.*, gradient descents [4], or the forward backward algorithm [9]), the solution depends on the initialization and the internal parameters of the solver, even if \mathcal{R} is chosen convex. Convexification by replacing the original objective function by its convex hull [31] simplifies the minimization at the cost of a loss of accuracy of the statistical model. In particular, the convex hull does not capture the right tail of the gamma distribution, leading to remaining bright outliers on the resulting images. For some Markov prior regularization, namely convex pairwise regularizations, the global optimum can be obtained in finite time [30]. This optimization method has been applied in the 1D illustration of Fig.2(c). It can be observed that, since large fluctuations in bright areas are more strongly penalized than the small fluctuations in low-level areas, speckle is reduced predominantly in the brightest areas. In the restored signal, the remaining noise variance is made uniform, *i.e.*, signal-independent, in linear scale. Thus,

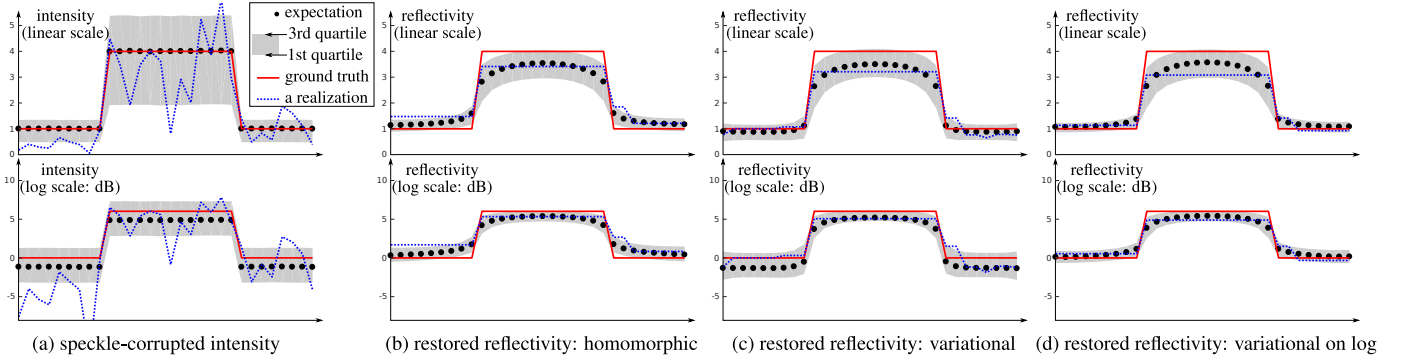


Fig. 2. Illustration of speckle reduction by 1D total variation minimization: (a) speckle-corrupted rectangle signal ($L = 2$), (b) restored reflectivity using the homomorphic approach and “L2 + TV” minimization in log-domain, (c) restored reflectivity using a variational approach “gamma + TV” in the linear domain, (d) restored reflectivity using a variational approach “Fisher-Tippett + TV” in the log domain. Expectation and quartiles are computed from 50 000 noisy realizations. Restorations computed by exact discrete optimization [30] with 2 000 quantization levels.

small details with identical signal to noise ratio will be more likely suppressed if placed in areas of large average value. This phenomenon is also observed on images: speckle noise is reduced more strongly in brighter areas [17], [52], and bright punctual targets are found to be spread out.

D. Variational approach on log-transformed data

Another alternative proposed in [51], [52], [7], [18] consists of formulating the estimation problem in the log domain. Rather than approximating the likelihood of log-transformed reflectivities by a Gaussian distribution (as done by homomorphic approaches discussed in Sec.II-B), the Fisher-Tippett distribution (3) is considered. The regularization is also expressed on the log-transformed reflectivities $x \in \mathbb{R}^n$, leading to a MAP optimization problem of the form:

$$\hat{x} \in \operatorname{argmin}_{\hat{x} \in \mathbb{R}^n} -\log p_y(y|x) + \mathcal{R}(x), \quad (10)$$

$$\text{where } -\log p_y(y|x) = L \sum_{k=1}^n x_k + e^{y_k - x_k} + \text{Cst.}$$

The final estimate is defined as $\hat{R}_k = \exp \hat{x}_k$ for all pixel index k . Remark that $\log p_y(y|x) = \log p_I(\exp(y)|\exp(x)) + \text{Cst.}$, so that problems (9) and (10) only differ in terms of the prior regularization. Nevertheless, this change of variable leads to several advantages compared to (9). First, the optimization is unconstrained as x can be any vector of \mathbb{R}^n . More importantly, the data fidelity becomes convex. As a consequence, if \mathcal{R} is also chosen convex, solutions depend neither on the initialization or the choice of the internal parameters of the solver. Such internal parameters only affect the speed of convergence and thus the number of iterations to perform in practice.

The multiplicative image denoising by augmented Lagrangian (MIDAL) algorithm [7] considers a convex regularization, TV, and minimizes (10) using the alternating direction method of multipliers (ADMM) algorithm [25], [23], [19] that

repeats, for an internal parameter $\beta > 0$, the updates

$$\hat{z} \leftarrow \operatorname{argmin}_z \frac{\beta}{2} \|z - \hat{x} + \hat{d}\|^2 + \mathcal{R}(z), \quad (11)$$

$$\hat{d} \leftarrow \hat{d} + \hat{z} - \hat{x}, \quad (12)$$

$$\hat{x} \leftarrow \operatorname{argmin}_x \frac{\beta}{2} \|x - \hat{z} - \hat{d}\|^2 - \log p_y(y|x). \quad (13)$$

Clearly the minimization for z in (11) depends on the choice of \mathcal{R} and can be solved by a dedicated regularized least square solver (corresponding to a Gaussian denoiser). This is schematized on Fig.1, third row. The minimization for x in (13) amounts to solving n separable convex problems of the form

$$\operatorname{argmin}_{x_k} \frac{\beta}{2} (x_k - a_k)^2 + L (x_k + e^{y_k - x_k}), \quad (14)$$

where $a_k = \hat{z}_k + \hat{d}_k$. The explicit solution is given by the Lambert W functions [7], or can be computed more efficiently with a few iterations of Newton’s method as

$$\hat{x}_k \leftarrow \hat{x}_k - \frac{\beta(\hat{x}_k - a_k) + L(1 - e^{y_k - \hat{x}_k})}{\beta + L e^{y_k - \hat{x}_k}}. \quad (15)$$

Using ten iterations is usually enough to offer good performances within a reasonable computational time, see, e.g., [53].

As already mentioned, when \mathcal{R} is convex, the parameter β only acts on the speed of convergence. Interestingly, as the noise variance is independent from the signal, the convergence is in practice uniform meaning that for a finite number of iterations the same amount of smoothing will be performed both in dark and bright regions. Small details with identical signal to noise ratio will be identitically smoothed whatever the average value of the area. If moreover β is chosen proportional to $\Psi(1, L)^{-1}$, a similar smoothing will be reached for a fixed number of iterations whatever the initial number of looks L .

When \mathcal{R} is non-convex, but satisfies some weak conditions and β is chosen large enough, ADMM still converges to a local minimum [29]. In this case, the solution depends on both the initialization and the choice of β . Again, even in this case, the fact that the variance is stabilized helps at approaching local solutions in finite time with a uniform

amount of smoothing, independently of L when β is chosen proportional to $\Psi(1, L)^{-1}$.

This shows that taking the logarithm not only makes the data fidelity term convex, but also ensures a uniform speed of convergence of a non-constrained optimization problem. These are two key practical ingredients that challenges non-convex strategies that directly deal with gamma distributed values.

We illustrate on Fig.2(d) that minimization of the total variation of log-transformed intensities produces a smaller loss of contrast of bright structures and reduces speckle equally in all regions.

III. EXTENSION TO MULTI-VARIATE SAR IMAGERY

A. Statistics of covariance matrices

Multi-variate complex SAR images carry much more information than univariate SAR images since inter-channel cross-correlations capture geophysical features (*e.g.*, height is related to the interferometric phase, geometrical configuration with the polarimetric properties). All relevant information of these complex valued images can be gathered together at each pixel by forming a $D \times D$ complex covariance matrix \mathbf{C} . Goodman's model [26] describes speckle in this matrix as being circular complex Wishart distributed, for $L \geq D$:

$$p_{\mathbf{C}}(\mathbf{C}|\mathbf{\Sigma}) = \frac{L^L D |\mathbf{C}|^{L-D}}{\Gamma_D(L) |\mathbf{\Sigma}|^L} \exp(-L \operatorname{tr}(\mathbf{\Sigma}^{-1} \mathbf{C})) , \quad (16)$$

where $\mathbf{\Sigma}$ is the underlying covariance matrix encoding reflectivities and complex correlations, and $L \geq D$ is the number of looks. Both \mathbf{C} and $\mathbf{\Sigma}$ belong to the open cone of complex hermitian positive definite matrices. While fluctuations in univariate SAR images are multiplicative, the Wishart distribution, denoted by $\mathcal{W}(\mathbf{\Sigma}; L)$ models fluctuations in multivariate SAR images as

$$\mathbf{C} = \mathbf{\Sigma}^{1/2} \mathbf{S} \mathbf{\Sigma}^{1/2} , \quad (17)$$

$$\mathbb{E}[\mathbf{C}] = \mathbf{\Sigma} , \quad (18)$$

$$\operatorname{Var}[\mathbf{C}_{ij}] = \frac{1}{L} \mathbf{\Sigma}_{ii} \mathbf{\Sigma}_{jj} , \quad (19)$$

where $\mathbf{S} \sim \mathcal{W}(\operatorname{Id}; L)$ (see, *e.g.*, [40])². Note that the variance for off-diagonal elements does not depend on $\mathbf{\Sigma}_{ij}$ but on $\mathbf{\Sigma}_{ii}$ and $\mathbf{\Sigma}_{jj}$, which indicates that the fluctuations are not only intra channel signal dependent but inter channel signal dependent. Interestingly, according to [27], [40], we have the following relations regarding the determinant and the trace

$$|\mathbf{C}| = |\mathbf{\Sigma}| |\mathbf{S}| , \quad (20)$$

$$\operatorname{tr} \mathbf{C} = \operatorname{tr}(\mathbf{\Sigma} \mathbf{S}) , \quad (21)$$

$$\mathbb{E}[\operatorname{tr} \mathbf{C}] = \operatorname{tr} \mathbf{\Sigma} , \quad (22)$$

$$\operatorname{Var}[\operatorname{tr} \mathbf{C}] = \frac{1}{L} \operatorname{tr} \mathbf{\Sigma}^2 , \quad (23)$$

$$\text{and } \mathbb{E}[\operatorname{tr}(\mathbf{C}^2 - \mathbf{\Sigma}^2)] = \frac{1}{L} (\operatorname{tr} \mathbf{\Sigma})^2 . \quad (24)$$

²for complex random variables $\operatorname{Var}[\mathbf{C}_{ij}] = \mathbb{E}[|\mathbf{C}_{ij}|^2] - |\mathbb{E}[\mathbf{C}_{ij}]|^2$

B. Limit of a direct variational approach

As for the univariate case, a variational approach considering the Wishart distribution of \mathbf{C} can be expressed as

$$\hat{\mathbf{\Sigma}} \in \operatorname{argmin}_{\mathbf{\Sigma} \succ 0} -\log p_{\mathbf{\Sigma}}(\mathbf{\Sigma}|\mathbf{C}) + \mathcal{R}(\mathbf{\Sigma}) , \quad (25)$$

where $-\log p_{\mathbf{\Sigma}}(\mathbf{\Sigma}|\mathbf{C}) = L \sum_{k=1}^n \log |\mathbf{\Sigma}_k| + \operatorname{tr}(\mathbf{\Sigma}_k^{-1} \mathbf{C}_k) + \text{Cst.}$

where $\mathbf{\Sigma} \succ 0$ reads as “ $\mathbf{\Sigma}$ is hermitian positive definite”. While estimating directly the reflectivity in a variational framework in the case of univariate data requires to enforce a non-negativity constraint, optimizing for $\mathbf{\Sigma}$ requires optimizing on the open cone of complex hermitian positive definite matrices, which is much more challenging. Moreover, as noted for the univariate case, the neg-log-likelihood associated to the Wishart distribution is highly non-convex, so that finding a good quality local optimum is very difficult. This approach is summarized on the second row of Fig.1.

To circumvent the difficulty arising from the non convexity of the neg-log-likelihood, Nie *et al.* approximated it by its convex envelope [45], [46]. In their first algorithm, WisTV [45], \mathcal{R} was chosen as a matricial total-variation regularizer, and next in [46] it was replaced by a matricial non-local total-variation regularizer [24] to better preserve textures. Replacing the neg-log-likelihood by its convex envelope is a rather crude approximation since it leads to underestimating the right tail of the distribution, and thus prevents from being robust against bright outliers.

The noise components in \mathbf{C} are intensively signal dependent, so that speckle suppression is not as strong in all regions. Last but not least, as for the univariate case, it has been recently observed that the performance of WisTV can drop significantly when the dynamic range of the input SAR image becomes very high [47]. To cope with these different issues we suggest mimicking the univariate case by making use of the matrix logarithm, thereby extending MIDAL approach [7] to multi-channel SAR images.

C. The Wishart-Fisher-Tippett distribution

Our objective is to generalize the use of the log transform in univariate SAR images to multi-variate SAR images. To that end, we resort to the matrix logarithm defined as

$$\mathbf{\Sigma} \mapsto \tilde{\mathbf{\Sigma}} = \log \mathbf{\Sigma} = \mathbf{E} \operatorname{diag}(\tilde{\mathbf{\Lambda}}) \mathbf{E}^{-1} \quad \text{where} \quad \tilde{\mathbf{\Lambda}}_i = \log \mathbf{\Lambda}_i , \quad (27)$$

$\mathbf{E} \in \mathbb{C}^{D \times D}$ is the matrix whose column vectors are eigenvectors (with unit norm) of $\mathbf{\Sigma}$ and $\mathbf{\Lambda} \in \mathbb{R}_+^D$ is the vector of corresponding eigenvalues, such that $\mathbf{\Sigma} = \mathbf{E} \operatorname{diag}(\mathbf{\Lambda}) \mathbf{E}^{-1}$, and $\tilde{\mathbf{\Lambda}} \in \mathbb{R}^D$. Its inverse transform is the matrix exponential defined similarly as

$$\tilde{\mathbf{\Sigma}} \mapsto \mathbf{\Sigma} = e^{\tilde{\mathbf{\Sigma}}} = \mathbf{E} \operatorname{diag}(\mathbf{\Lambda}) \mathbf{E}^{-1} \quad \text{where} \quad \mathbf{\Lambda}_i = \exp \tilde{\mathbf{\Lambda}}_i . \quad (28)$$

While $\mathbf{\Sigma}$ lies in the open cone of complex hermitian positive definite matrices, $\tilde{\mathbf{\Sigma}}$ lies in the vector space of complex hermitian matrices which is isomorphic to \mathbb{R}^{D^2} . The change

$$\mathcal{K} : \begin{pmatrix} \alpha^1 \\ \vdots \\ \alpha^{D^2} \end{pmatrix} \mapsto \begin{pmatrix} \alpha^1 & (\alpha^{D+1} + j\alpha^{D+2})/\sqrt{2} & \dots & (\alpha^{D^2-1} + j\alpha^{D^2})/\sqrt{2} \\ (\alpha^{D+1} - j\alpha^{D+2})/\sqrt{2} & \alpha^2 & & \\ \vdots & & \ddots & \vdots \\ (\alpha^{D^2-1} - j\alpha^{D^2})/\sqrt{2} & \dots & & \alpha^D \end{pmatrix} \quad (26)$$




Fig. 3. Three two dimensional sections of the Wishart-Fisher-Tippett negative log likelihood with respect to (x^1, x^2) , (x^1, x^3) and (x^3, x^4) .

of variables $\tilde{\mathbf{C}} = \log \mathbf{C}$ and $\tilde{\Sigma} = \log \Sigma$ lead to the distribution of log-transformed matrices $\tilde{\mathbf{C}}$:

$$p_{\tilde{\mathbf{C}}}(\tilde{\mathbf{C}}|\tilde{\Sigma}) = \frac{L^{LD} f(\tilde{\mathbf{C}})}{\Gamma_D(L)} e^{L \text{tr}(\tilde{\mathbf{C}} - \tilde{\Sigma})} \exp\left(-L \text{tr}(e^{\tilde{\mathbf{C}}} e^{-\tilde{\Sigma}})\right) \quad (29)$$

with $f(\tilde{\mathbf{C}}) = |\mathbf{J}_{\exp}(\tilde{\mathbf{C}})| / \exp[D \text{tr}(\tilde{\mathbf{C}})]$ a normalization factor that involves the jacobian of log transform $|\mathbf{J}_{\exp}(\tilde{\mathbf{C}})|$ and that is equal to 1 when $D = 1$. We call such a distribution the Wishart-Fisher-Tippett distribution, denoted as $\mathcal{WFT}(\tilde{\Sigma}; L)$, as it generalizes the Fisher-Tippett distribution to the case where $D > 1$. The expectation and variance of $\tilde{\mathbf{C}}$ do not seem to be known in closed form in the literature. Nevertheless, according to [2] its trace (which coincides with the logarithm determinant of \mathbf{C} : $\text{tr} \tilde{\mathbf{C}} = \log |\mathbf{C}|$) has the following statistics

$$\text{tr} \tilde{\mathbf{C}} = \text{tr} \tilde{\Sigma} + \text{tr} \tilde{\mathbf{S}}, \quad (30)$$

$$\mathbb{E}[\text{tr} \tilde{\mathbf{C}}] = \text{tr} \tilde{\Sigma} + \sum_{i=1}^D \Psi(0, L - i + 1) - D \log L, \quad (31)$$

$$\text{and } \text{Var}[\text{tr} \tilde{\mathbf{C}}] = \sum_{i=1}^D \Psi(1, L - i + 1), \quad (32)$$

where $\tilde{\mathbf{S}} \sim \mathcal{WFT}(\mathbf{0}; L)$. This shows that the trace of the matrix logarithm suffers from additive non-zero mean signal independent corruptions. Note that (30) is a direct consequence of (20). It follows that the D^2 channels of $\tilde{\mathbf{C}}$ can reasonably be assumed to have approximately a stabilized variance (see Section III-G for numerical evidences), which opens the door to regularization with iterative schemes. Note that we cannot use the matrix log transform directly as a variance stabilization procedure, as done in the univariate homomorphic case, because we do not have an inversion formula (i.e., a bias correction formula). We will thus adopt a variational strategy.

D. Log-channel decomposition

As mentioned in the previous paragraph, $\tilde{\Sigma}$ lies in the vector space of complex Hermitian matrices that is isomorphic to \mathbb{R}^{D^2} . In this section we describe a re-parameterization of the log-transformed covariance matrix $\tilde{\Sigma}$ as a vector of D^2 reals denoted \mathbf{x} . We first define in (26) a transform \mathcal{K} that maps real-valued vectors α of \mathbb{R}^{D^2} to Hermitian $D \times D$ matrices. We also introduce a whitening affine map (\mathbf{A}, \mathbf{b}) so that the D^2 channels of \mathbf{x} are better decorrelated, and a scaling transform Φ (i.e., a diagonal matrix) to balance the variance of noise between channels. The transform Ω between (log-transformed) covariance matrices $\tilde{\Sigma}$ and the vector of parameter \mathbf{x} is:

$$\tilde{\Sigma} = \Omega(\mathbf{x}) = \mathcal{K}(\mathbf{A}\Phi\mathbf{x} + \mathbf{b}). \quad (33)$$

We denote its inverse $\Omega^{-1}(\tilde{\Sigma}) = \Phi^{-1}\mathbf{A}^{-1}(\mathcal{K}^{-1}(\tilde{\Sigma}) - \mathbf{b})$ and introduce the linear operator $\Omega^* = \Phi\mathbf{A}^*\mathcal{K}^*$, where $*$ denotes the adjoint. The re-parameterized log-transformed data are noted $\mathbf{y} \in \mathbb{R}^{D^2}$ and defined by $\mathbf{y} = \Omega^{-1}(\tilde{\mathbf{C}})$, see Fig. 1.

The affine transform (\mathbf{A}, \mathbf{b}) is chosen such that the channels x^i of the re-parameterized vector are approximately independent:

$$p_{\tilde{\Sigma}}(\tilde{\Sigma}) = p_{\mathbf{x}}(\mathbf{x}) \approx \prod_{i=1}^{D^2} p_{x^i}(x^i). \quad (34)$$

The affine transform (\mathbf{A}, \mathbf{b}) is obtained by principal component analysis, as described in Section III-G. Moreover, the scaling matrix Φ is chosen such that the variance of y^i is equal to 1 for all channels i , see Section III-G.

E. Proposed variational approach

We now extend \mathbf{x} and \mathbf{y} into images of n real-valued vectors each of dimension D^2 . Writing $-\log p_{x^i}(x^i) = \mathcal{R}(x^i)$ and

$$\Delta_{\infty}^{k,i} = \frac{\beta(x_k^i - a_k^i) + L \left[\Omega^* \left(\mathbf{Id} - \int_0^1 e^{(u-1)\Omega(\mathbf{x}_k)} e^{\Omega(\mathbf{y}_k)} e^{-u\Omega(\mathbf{x}_k)} du \right) \right]^i}{\beta + L \left| \Omega^* \int_0^1 e^{(u-1)\Omega(\mathbf{x}_k)} e^{\Omega(\mathbf{y}_k)} e^{-u\Omega(\mathbf{x}_k)} du \right|^i}. \quad (35)$$

$$\Delta_Q^{k,i} = \frac{\beta(x_k^i - a_k^i) + L \left[\Omega^* \left(\mathbf{Id} - \frac{1}{Q} \sum_{q=1}^Q e^{(u_q-1)\Omega(\mathbf{x}_k)} e^{\Omega(\mathbf{y}_k)} e^{-u_q\Omega(\mathbf{x}_k)} \right) \right]^i}{\beta + L \left| \Omega^* \frac{1}{Q} \sum_{q=1}^Q e^{(u_q-1)\Omega(\mathbf{x}_k)} e^{\Omega(\mathbf{y}_k)} e^{-u_q\Omega(\mathbf{x}_k)} \right|^i}, \quad u_q = \frac{q - \frac{1}{2}}{Q}. \quad (36)$$

using the relations (34) within the MAP framework leads to the following minimization problem

$$\hat{\mathbf{x}} \in \underset{\mathbf{x}}{\operatorname{argmin}} -\log p_{\mathbf{y}}(\mathbf{y}|\mathbf{x}) + \sum_{i=1}^{D^2} \mathcal{R}(x^i), \quad (37)$$

where, from (29), we have

$$-\log p_{\mathbf{y}}(\mathbf{y}|\mathbf{x}) = L \sum_{k=1}^n \operatorname{tr} \left(\Omega(\mathbf{x}_k) + e^{\Omega(\mathbf{y}_k)} e^{-\Omega(\mathbf{x}_k)} \right) + \text{Cst.}, \quad (38)$$

and the final estimate $\hat{\Sigma}_k$ at pixel index k is defined as $\hat{\Sigma}_k = \exp \Omega(\hat{\mathbf{x}}_k)$. Note the difference between \mathbf{x}_k a vector of D^2 coefficients at pixel k and x^i a scalar image of n pixels corresponding to the i -th channel of \mathbf{x} .

While the direct multivariate variational approach (25) requires optimizing on the cone of complex hermitian matrices, a first major advantage is that Problem (37) is unconstrained on \mathbb{R}^{D^2} . Even though the likelihood term (38) is still non convex, it appears to be much more suitable for optimization. Figure 3 illustrates the evolution of this term (with $D = 2$, and $n = 1$) on three two-dimensional cross sections of \mathbb{R}^4 showing that for each of such sections the energy appears to be convex. Unlike the direct multivariate variational approach (25), (38) appears to be convex in many scenarios, *e.g.*, in the univariate case (10), and more generally when $\hat{\mathbf{x}}_k$ is restricted to the vectorial subspace where $e^{-\Omega(\hat{\mathbf{x}}_k)}$ commutes with $e^{\Omega(\mathbf{y}_k)}$, hence satisfying $e^{\Omega(\mathbf{y}_k)} e^{-\Omega(\hat{\mathbf{x}}_k)} = e^{\Omega(\mathbf{y}_k - \mathbf{x}_k)}$. Convexity in this case follows from the fact that the trace is linear, hence convex, the function $\mathbf{x}_k \rightarrow \operatorname{tr} e^{\Omega(-\mathbf{x}_k + \mathbf{y}_k)}$ is a convex spectral function [22], [38], and Ω is affine.

As for the mono-dimensional case, we will thus consider the ADMM algorithm which iteratively performs the updates

$$\hat{\mathbf{z}} \leftarrow \underset{\mathbf{z}}{\operatorname{argmin}} \frac{\beta}{2} \|\mathbf{z} - \hat{\mathbf{x}} + \hat{\mathbf{d}}\|^2 + \sum_{i=1}^{D^2} \mathcal{R}(z^i), \quad (39)$$

$$\hat{\mathbf{d}} \leftarrow \hat{\mathbf{d}} + \hat{\mathbf{z}} - \hat{\mathbf{x}}, \quad (40)$$

$$\hat{\mathbf{x}} \leftarrow \underset{\mathbf{x}}{\operatorname{argmin}} \frac{\beta}{2} \|\mathbf{x} - \hat{\mathbf{z}} - \hat{\mathbf{d}}\|^2 - \log p_{\mathbf{y}}(\mathbf{y}|\mathbf{x}). \quad (41)$$

Since the noise variance is approximately stabilized and equal to about 1 on each channel (see, III-G), a single value for β can be chosen for all channels and all possible number of

looks L . With such a choice, the same amount of smoothing is reached in practice for a fixed number of iterations, whatever L and D .

Equation (39) is separable on the different channels, and can then be solved by applying D^2 regularized least-square solvers. Equation (41) amounts to solving n separable problems of the form

$$\underset{\mathbf{x}_k}{\operatorname{argmin}} \frac{\beta}{2} \|\mathbf{x}_k - \mathbf{a}_k\|^2 + L \operatorname{tr} \left(\Omega(\mathbf{x}_k) + e^{\Omega(\mathbf{y}_k)} e^{-\Omega(\mathbf{x}_k)} \right) \quad (42)$$

where $\mathbf{a}_k = \hat{\mathbf{z}}_k + \hat{\mathbf{d}}_k$. As for the univariate case, we will consider Newton's method to solve this optimization problem. Its gradient is given by

$$\beta(\mathbf{x}_k - \mathbf{a}_k) + L \Omega^* \left(\mathbf{Id} - \int_0^1 e^{(u-1)\Omega(\mathbf{x}_k)} e^{\Omega(\mathbf{y}_k)} e^{-u\Omega(\mathbf{x}_k)} du \right), \quad (43)$$

see the development in Appendix A. By mimicking the univariate case, we consider the following approximation for the Hessian

$$\operatorname{diag} \left(\beta \mathbf{Id} + L \Omega^* \int_0^1 e^{(u-1)\Omega(\mathbf{x}_k)} e^{\Omega(\mathbf{y}_k)} e^{-u\Omega(\mathbf{x}_k)} du \right), \quad (44)$$

leading to quasi Newton's iterations given by

$$\hat{\mathbf{x}}_k^i \leftarrow \hat{\mathbf{x}}_k^i - \Delta_{\infty}^{k,i}, \quad (45)$$

where $\Delta_{\infty}^{k,i}$ is defined in equation (35). As for the univariate case, we notice that ten iterations is enough to offer good performance with a reasonable computational time.

The update matrix $\Delta_{\infty}^{k,i}$, defined in equation (35), requires numerical integration. We use Riemann integral approximation with Q rectangles as defined in equation (36). Of course, the smaller Q , the faster the iteration (45). Figure 4 shows the evolution of $\|\Delta_{100}\|/\|\hat{\mathbf{x}}\|$ with respect to D (with $L = D$ and $\beta = 10L$) after 10 iterations of Newton's method where Δ_{∞} is substituted by Δ_Q , for $Q = \{1, 2, 4, 8, 16\}$. We furthermore define Δ_0 by substituting the integral by $\exp(\Omega(\mathbf{y}_k - \mathbf{x}_k))$ as a crude approximation of the gradient (and exact in the case where $e^{-\Omega(\hat{\mathbf{x}}_k)}$ commutes with $e^{\Omega(\mathbf{y}_k)}$). Ideally, the algorithm would reach the optimum solution such that $\Delta_{\infty} \approx \Delta_{100} = 0$ (as for the univariate case). This error is relatively small whatever Q and D . For low values of D , the error reduces by about a factor 10 each time Q is multiplied by 4. The




Fig. 4. Evolution with respect to the matrix dimension D of the error $\|\Delta_{100}\|/\|\mathbf{x}\|$ after 10 iterations of Newton's method for which the update Δ_∞ is replaced by Δ_Q , for $Q = \{0, 1, 2, 4, 8, 16\}$. For each D , the results are averaged over 100 realizations of \mathbf{C} (with $L = D$), each realization are obtained from different versions of Σ (also randomly generated).

increase with D is due to accumulation of numerical errors of the exponential matrix function, even more important when Q gets larger. This experiment reveals that $Q = 1$ is a good practical choice reaching a relative error of about $2 \cdot 10^{-3}$ whatever $D \in [2, 16]$ with reasonable computation time.

F. Adaptation to advanced filters

As in the univariate case, we notice that ADMM converges when the regularizer \mathcal{R} is chosen as non convex. More remarkably, as observed in deblurring contexts, *e.g.*, [55], [11], replacing the minimization problem in (39) by the solution of any Gaussian denoiser leads to appealing results. The resulting algorithm, coined MULTI-channel LOGarithm with Gaussian denoising (MuLoG), is given as

$$\hat{z}^i \leftarrow f_{\beta-1/2}(\hat{x}^i - \hat{d}^i), \quad \text{for } i = 1, \dots, D^2, \quad (46)$$

$$\hat{\mathbf{d}} \leftarrow \hat{\mathbf{d}} + \hat{\mathbf{z}} - \hat{\mathbf{x}}, \quad (47)$$

$$\hat{\mathbf{x}} \leftarrow \underset{\mathbf{x}}{\operatorname{argmin}} \frac{\beta}{2} \|\mathbf{x} - \hat{\mathbf{z}} - \hat{\mathbf{d}}\|^2 - \log p_{\mathbf{y}}(\mathbf{y}|\mathbf{x}), \quad (48)$$

where f_{β}^2 is again a denoiser for images contaminated by zero-mean white Gaussian noise $\mathcal{N}(0; \sigma^2 \text{Id})$, and (48) is solved as mentioned in the previous paragraph. While it is difficult to make a general statement, we believe that such an algorithm converges as soon as the denoiser $f_{\beta-1/2}$ is at least non-expansive. We depict this algorithm on the last row of Fig.1.

G. Calibration

In practice, the operators \mathbf{A} , \mathbf{b} and Φ are obtained from the log matrix of the input image \mathbf{C} with n pixels. Let $\{\alpha_1, \dots, \alpha_n\}$ be the collection of n vectors extracted from the matrix logarithm of the input image as $\alpha_k = \mathcal{K}^{-1}(\tilde{\mathbf{C}}_k)$. We thus define $\mathbf{b} = \frac{1}{n} \sum_{k=1}^n \alpha_k$ and \mathbf{A} the matrix whose columns are the eigenvectors (with unit norm) of $\frac{1}{n} \sum_{k=1}^n (\alpha_k - \mathbf{b})(\alpha_k - \mathbf{b})^t$. As a result, the vector $\mathbf{A}^{-1}(\alpha_k - \mathbf{b})$ is the representation of α_k in the principal component analysis of

$\{\alpha_1, \dots, \alpha_n\}$ known to maximize channel de-correlation, and thus leads to (34). By assuming the noise variance is stabilized on each channel of $\mathbf{A}^{-1}(\alpha - \mathbf{b})$, we define the diagonal matrix Φ as $\text{diag}(\hat{\sigma}_1^2, \dots, \hat{\sigma}_D^2)$, where $\hat{\sigma}_1^2, \dots, \hat{\sigma}_D^2$ are obtained by estimating separately the variance of the noise on each channel of $\mathbf{A}^{-1}(\alpha - \mathbf{b})$. This procedure ensures that the noise variance is about 1 on each channel of \mathbf{y} . In practice, we resort to the median absolute deviation estimate (MAD) [28] for this task, but other strategies could be used.

Figure 5 displays the nine components of \mathbf{y} for an image of 3×3 matrices \mathbf{C} . While the noise variance in \mathbf{C} is clearly signal dependent, the amplitudes of the fluctuations in the different channels of \mathbf{y} appear fairly constant whatever the underlying signal Σ . Nevertheless, the noise is not signal independent, we can observe that some regions are a smidgen noisier than others. Remark also that the noise variance is also about the same for all channels. Then, we can reasonably claim from this experiment that our log channel decomposition approximately stabilizes the noise variance.

Thanks to this approximate stabilization, we notice in practice that choosing $\beta = 4$ and 6 iterations provide satisfying solutions in all tested situations, whatever the dimension D , the number of looks L and the embedded Gaussian denoiser. We always choose $Q = 1$ according to Sec. III-E.

When $L < D$, the matrix \mathbf{C}_k is rank deficient, and the likelihood term (38), though convex when restricted to commutative matrices, is however no longer strictly convex and has an infinite number of minimizers. Worse, its logarithm $\tilde{\mathbf{C}}$ is undefined which prevents the computation of \mathbf{A} , \mathbf{b} and Φ used in our log-channel decomposition. To deal with these issues a practical solution is to perform a small diagonal loading of the input matrices \mathbf{C}_k . We perform a spatially varying re-scaling of the off-diagonal elements as

$$(\mathbf{C}_k)_{i,j}^{\text{new}} \leftarrow \frac{|\sum_l w_{k,l}(\mathbf{C}_l)_{i,j}|}{\sqrt{\sum_l w_{k,l}(\mathbf{C}_l)_{ii} \sum_l w_{k,l}(\mathbf{C}_l)_{jj}}} (\mathbf{C}_k)_{i,j}, \quad (49)$$

where $w_{k,l}$ are the weights of a Gaussian kernel with bandwidth 1 pixel. This procedure inevitably introduces some bias in the solution, but this bias does not seem to be significant.

IV. EFFICIENT MATRIX LOG AND EXP TRANSFORMS

Our algorithm requires several times to compute n matrix logarithms, matrix exponentials and matrix-by-matrix products. It is thus important to make these operations as fast as possible. Our first attempt to call n times Matlab's functions `logm`, `expm` and `mtimes` leads to very slow computations. It is in fact important to use a vectorial implementation of these functions. When $D = 2$, we can write the matrix \mathbf{C} and $\tilde{\mathbf{C}}$ as

$$\mathbf{C} = \begin{pmatrix} a & c^* \\ c & b \end{pmatrix} \quad \text{and} \quad \tilde{\mathbf{C}} = \begin{pmatrix} \tilde{a} & \tilde{c}^* \\ \tilde{c} & \tilde{b} \end{pmatrix}, \quad (50)$$

and the relation between both are given, for $\tilde{\mathbf{C}} = \log \mathbf{C}$, as

$$\begin{cases} \tilde{a} &= \frac{1}{\delta} [(a - b + \delta)l_1 - (a - b - \delta)l_2], \\ \tilde{b} &= \frac{1}{\delta} [-(a - b - \delta)l_1 - (a - b + \delta)l_2], \\ \tilde{c} &= \frac{1}{\delta} [cl_1 - cl_2], \end{cases} \quad (51)$$




Fig. 5. (a) the nine log channels $\mathbf{x}^1, \dots, \mathbf{x}^9$, with approximate unit stabilized variance, of (c) the image \mathbf{C} generated from (b) the image $\mathbf{\Sigma}$ displayed with a RGB representation based on its decomposition in Pauli's basis.

and, for $\mathbf{C} = \exp \tilde{\mathbf{C}}$, as

$$\begin{cases} a = \frac{1}{\delta} \left[(\tilde{a} - \tilde{b} + \tilde{\delta})e_1 - (\tilde{a} - \tilde{b} - \tilde{\delta})e_2 \right], \\ b = \frac{1}{\delta} \left[-(\tilde{a} - \tilde{b} - \tilde{\delta})e_1 - (\tilde{a} - \tilde{b} + \tilde{\delta})e_2 \right], \\ c = \frac{1}{\delta} [\tilde{c}e_1 - \tilde{c}e_2], \end{cases} \quad (52)$$

where $\delta = \sqrt{4|c| - (a-b)^2}$, $l_1 = \log\left(\frac{a+b+\delta}{2}\right)$, $l_2 = \log\left(\frac{a+b-\delta}{2}\right)$, $\tilde{\delta} = \sqrt{4|\tilde{c}| - (\tilde{a}-\tilde{b})^2}$, $e_1 = \exp\left(\frac{\tilde{a}+\tilde{b}+\tilde{\delta}}{2}\right)$ and $e_2 = \exp\left(\frac{\tilde{a}+\tilde{b}-\tilde{\delta}}{2}\right)$. With these established relations, we can compute the matrix logarithm in a vectorial way, starting by computing for the n pixels the coefficient \tilde{a} , and next \tilde{b} and finally \tilde{c} . We have implemented this vectorial procedure in Matlab (using the element wise operators $\cdot *$ and $\cdot /$). and observed an acceleration by a factor larger than 600 compared to calling n times the dedicated Matlab's function. We were also able to derive closed form formula for the case where $D = 3$ (the formulas are too long to be inserted in this paper), leading in this case to an acceleration by a factor 75. The same workaround is applied for the matrix exponential and the matrix-by-matrix product.

V. NUMERICAL EXPERIMENTS

Figures 6, 7 and 8 give three illustrations of the proposed technique. Figure 6 corresponds to an airborne single-look SAR interferometric image ($D = 2, L = 1$) of a building in Toulouse, France (sensed by RAMSES). Figure 7 corresponds to a spaceborne single-look SAR interferometric image ($D = 2, L = 1$) of the dam of Serre-Ponçon, France (sensed by TerraSAR-X). Figure 8 corresponds to an airborne SAR

polarimetric image ($D = 3, L = 1$) of the city of Kaufbeuren, Germany (sensed by F-SAR).

In the sub-figures (b) are displayed the results for the isotropic total-variation (TV) regularization, i.e., $\mathcal{R}(x^i) = \lambda \sum_{k=1}^n \|(\nabla x^i)_k\|$ which corresponds to the multivariate extension of [7]. The parameter λ has been tuned by hand once and kept the same for all data-sets, whatever the dimension D or the number of looks L . In this case, ADMM find a local minimum of (37). We can observe that the solution in each channel inherits from the well-known behavior of the univariate TV regularization: the solution appears piece-wise constant, small details and low contrasted features are lost, and a slight bias can be observed.

In the sub-figures (c) and (d) are displayed the results for DDID [33] and BM3D [10]. The inner parameters of these two algorithms have been kept the same as the one provided in the authors' implementation. In these cases, our proposed algorithm does not explicitly minimize a given energy but we can observe that it converges in practice to relevant solutions. As for TV, we can observe that the solution in each channel inherits from the behavior of the original method: in this case the small details and low contrasted features are well restored but some known typical artifacts (small oscillating features) of these two methods can be observed as well.

Table I reports the computation time of our proposed approach for different Gaussian denoisers, different image sizes and different covariance matrix sizes. Our implementation uses parallelization obtained by running eq. (46) in parallel on the D^2 channels, and running eq. (48) on different subsets of the n pixels. Our experiments were conducted on a processor with 4 cores. In this experiment, we observe that the total computation time is always faster than running successively $6 \times D^2$ times the original algorithm.

This paper only deals with the problem of how to apply Gaussian filters to multi-variate SAR data. The study of comparing their relative performance or comparing them to state-of-the-art multi-variate SAR filters is out of the scope of this paper. We however believe that this study should be performed in a future work and is of main interest for the SAR community. For this reason, we have released a Matlab script³, under CECILL license, that takes as input a multivariate SAR image, its number of looks and a Gaussian denoiser provided by the user, and outputs the filtered SAR image.

VI. CONCLUSION

While a large diversity of speckle reduction methods exist for intensity images, only few can be extended to process multi-channel SAR images. In particular, extension of variational methods lead to challenging optimization problems due to the non-convexity of the neg-log-likelihood and the positive definiteness constraint on the covariance matrices. Furthermore, signal and channel dependent variance lead to restoration results with an uneven suppression of speckle. This paper introduced a general scheme based on a matrix logarithm

³<http://www.math.u-bordeaux.fr/~cdeledal/mulog>

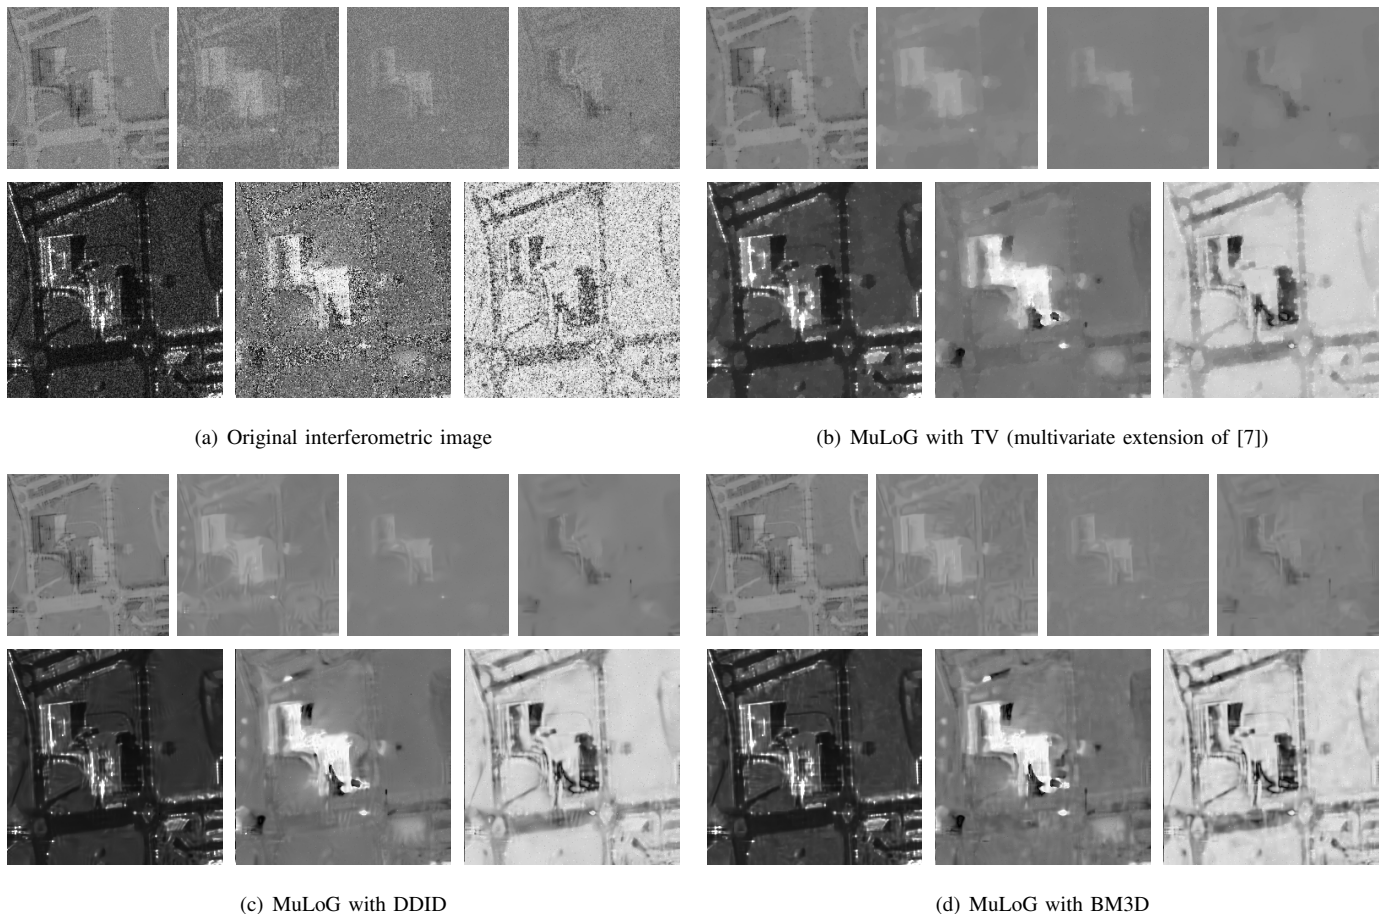


Fig. 6. (a) A single-look interferometric image ($D = 2, L = 1$) of Toulouse (France) sensed by RAMSES (©ONERA). Estimation with our proposed multivariate framework MuLoG for (b) Total-Variation (TV), (c) DDID and (d) BM3D. For each image, the four log channels are displayed on top x^1, x^2, x^3, x^4 , below are the trace $\text{tr} \mathbf{C}$, the phase $\arg \mathbf{C}_{12}$ and the coherence $|\mathbf{C}_{12}|/\sqrt{|\mathbf{C}_{11}\mathbf{C}_{22}|}$ respectively.

TABLE I

COMPUTATION TIME ON AN “INTEL(R) CORE(TM) i5-4590S CPU @ 3.00GHz”. TIME FOR GD IS THE RUNNING TIME OF THE GAUSSIAN DENOISER ON ONE SCALAR CHANNEL. TIME FOR NM IS FOR THE 10 ITERATIONS OF NEWTON’S METHOD WITH $Q = 1$. TOTAL TIME IS THE ACTUAL TIME OF OUR IMPLEMENTATION WITH 6 ITERATIONS OF THE ADMM SCHEME WHERE SOME PARTS OF THE CODE ARE RUNNED IN PARALLEL ON 4 DIFFERENT CORES.

Image size	D^2	File size	Gaussian denoiser	Time for GD	Time for NM	Total time
256×256	2×2	2Mb	BM3D	0.87s		8.97s
			TV	1.43s	0.35s	17.08s
			DDID	38.4s		4min 29s
256×256	3×3	4.5Mb	BM3D	0.87s		29.21s
			TV	1.43s	2.20s	48.18s
			DDID	38.4s		12min 6s
512×512	3×3	18Mb	BM3D	3.6s		2min 09s
			TV	6.3s	9.97s	4min 4s
			DDID	2min 37s		48s 52s

transform to approximately stabilize speckle variance and produce close to independent channels. Each channel can then be processed with a user-defined Gaussian denoiser. Upon re-iterating, a good fit of the restored multi-channel image with the Wishart distribution of input covariance matrices is enforced.

Special care is paid to ensure that the method requires no parameter tuning other than possibly within the Gaussian denoiser, and that these parameters, if any, can be tuned once for all. The resulting generic method can then include Gaussian denoisers selected by the user and tremendously extends

the set of available speckle reduction methods applicable to multi-channel SAR images. We believe that this offers several notable advantages: (i) the SAR imaging community will directly benefit from upcoming progress made in the field of image denoising; (ii) several images with reduced speckle can be produced by very different denoising algorithms, and these images can be compared to discard artifacts and confirm weak structures; (iii) this family of speckle reduction methods can serve as a reference to benchmark future specialized speckle reduction algorithms. This motivated the release of an open-source code implementing our method³.

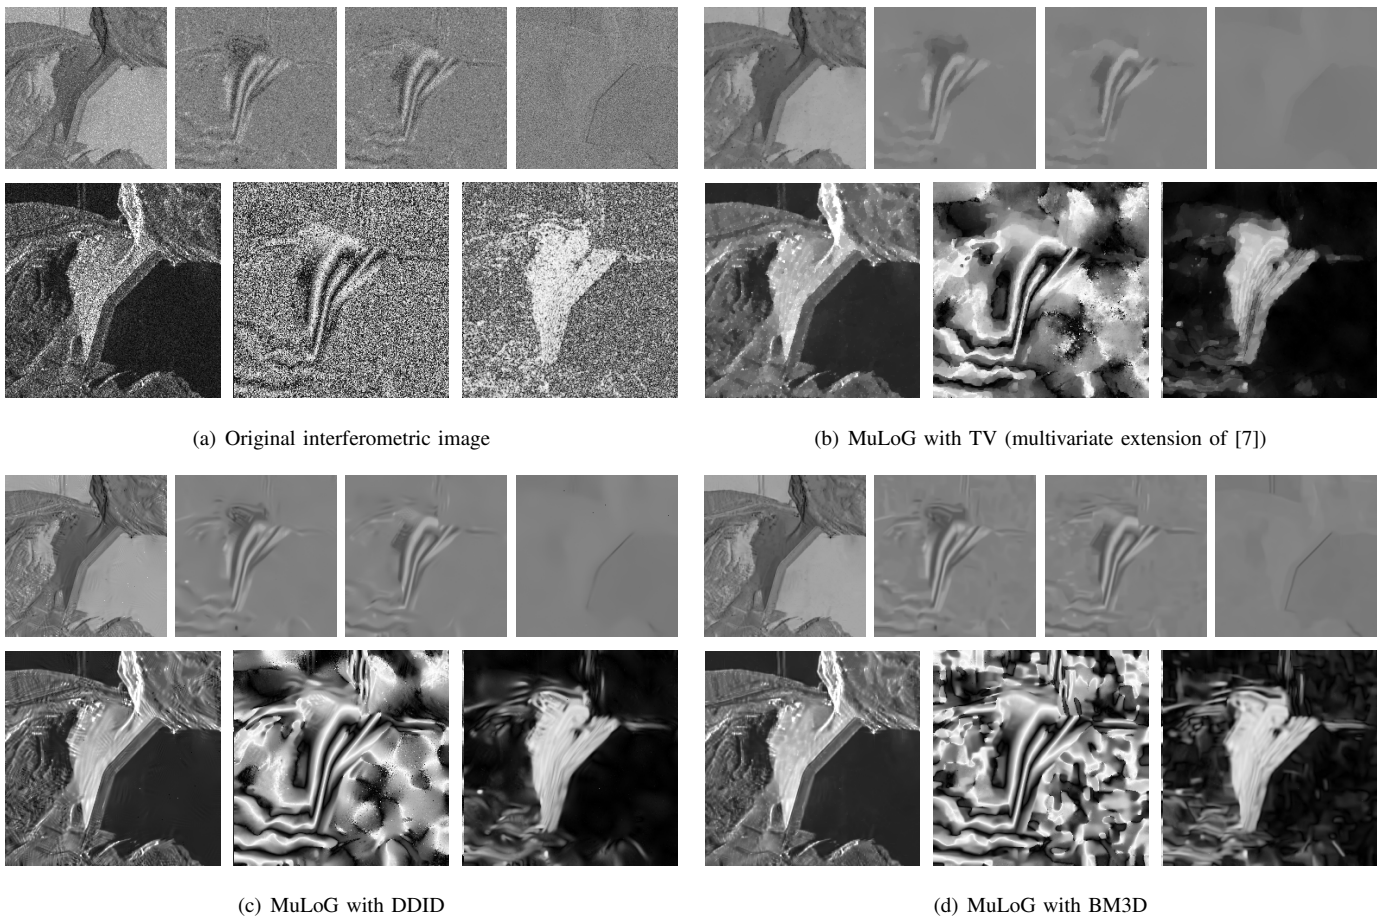


Fig. 7. (a) A single-look interferometric image ($D = 2, L = 1$) of Serre-Ponçon (France) sensed by TerraSar-X (image courtesy to Airbus Defence and Space). Estimation with our proposed multi-variate framework MuLoG for (b) Total-Variation (TV), (c) DDID and (d) BM3D. For each image, the four log channels are displayed on top $\mathbf{x}^1, \mathbf{x}^2, \mathbf{x}^3, \mathbf{x}^4$, below are the trace $\text{tr } \mathbf{C}$, the phase $\arg \mathbf{C}_{12}$ and the coherence $|\mathbf{C}_{12}|/\sqrt{|\mathbf{C}_{11}\mathbf{C}_{22}|}$ respectively.

ACKNOWLEDGMENTS

The authors would like to thank the Centre National d'Études Spatiales, the Office Nationale d'Études et de Recherches Aérospatiales and the Délégation Générale pour l'Armement for providing the RAMSES data, the Airbus Defence and Space for the TerraSAR-X data and the German Aerospace Center (DLR) for the F-SAR data. S. Tabti also thanks the Délégation Générale pour l'Armement and the Centre National de la Recherche Scientifique for funding her PhD program.

APPENDIX A

GRADIENT OF THE NEG LOG LIKELIHOOD

In this section, we establish the part of formula (43) corresponding to the gradient of $\text{tr}(\Omega(\mathbf{x}) + e^{\Omega(\mathbf{y})} e^{-\Omega(\mathbf{x})})$. To this end, recall that for a real valued differentiable function f ,

$$df(\mathbf{x}) = \text{tr}[g(\mathbf{x})d\mathbf{x}] \Leftrightarrow \nabla_{\mathbf{x}} f(\mathbf{x}) = g(\mathbf{x})^* . \quad (53)$$

where g is a matrix valued function whose dimension depends on that of \mathbf{x} , and $*$ denotes the adjoint operator. Since $\Omega(\mathbf{x}) = \mathcal{K}(\mathbf{A}\Phi\mathbf{x} + \mathbf{b})$ is affine, this directly implies that

$$\nabla_{\mathbf{x}} \text{tr} \Omega(\mathbf{x}) = \Phi^* \mathbf{A}^* \mathcal{K}^* = \Omega^* . \quad (54)$$

We now use that for any matrix valued function h , we have (proven in [32], [5] according to [44])

$$de^{h(\mathbf{x})} = \int_0^1 e^{uh(\mathbf{x})} (dh(\mathbf{x})) e^{(1-u)h(\mathbf{x})} du . \quad (55)$$

It follows that, for any matrix \mathbf{L} , we have

$$d \text{tr} [\mathbf{L} e^{h(\mathbf{x})}] = \text{tr} [\mathbf{L} (de^{h(\mathbf{x})})] , \quad (56)$$

$$= \text{tr} \left[\mathbf{L} \int_0^1 e^{uh(\mathbf{x})} (dh(\mathbf{x})) e^{(1-u)h(\mathbf{x})} du \right] , \quad (57)$$

$$= \int_0^1 \text{tr} \left[\mathbf{L} e^{uh(\mathbf{x})} (dh(\mathbf{x})) e^{(1-u)h(\mathbf{x})} \right] du , \quad (58)$$

$$= \int_0^1 \text{tr} \left[e^{(1-u)h(\mathbf{x})} \mathbf{L} e^{uh(\mathbf{x})} (dh(\mathbf{x})) \right] du , \quad (59)$$

$$= \text{tr} \left[\left(\int_0^1 e^{(1-u)h(\mathbf{x})} \mathbf{L} e^{uh(\mathbf{x})} du \right) (dh(\mathbf{x})) \right] . \quad (60)$$

Choosing $h(\mathbf{x}) = -\Omega(\mathbf{x})$ and $\mathbf{L} = e^{\Omega(\mathbf{y})}$ leads to

$$\nabla_{\mathbf{x}} \text{tr}(e^{\Omega(\mathbf{y})} e^{-\Omega(\mathbf{x})}) = -\Omega^* \int_0^1 e^{(u-1)\Omega(\mathbf{x})} e^{\Omega(\mathbf{y})} e^{-u\Omega(\mathbf{x})} du , \quad (61)$$




Fig. 8. (a) A single-look interferometric image ($D = 3, L = 1$) of Kaufbeuren (Germany) sensed by F-SAR (©DLR). Estimation with our proposed multivariate framework MuLoG for (b) Total-Variation (TV), (c) DDID and (d) BM3D. For each image, the nine log channels are displayed on top $\mathbf{x}^1, \dots, \mathbf{x}^9$, below are a RGB representation of Σ based on its decomposition in Pauli's basis.

the integral term being Hermitian since it reads as the integral of Hermitian matrices

$$\int_0^{\frac{1}{2}} \underbrace{e^{(u-1)\Omega(\mathbf{x})} e^{\Omega(\mathbf{y})} e^{-u\Omega(\mathbf{x})} + e^{-u\Omega(\mathbf{x})} e^{\Omega(\mathbf{y})} e^{(u-1)\Omega(\mathbf{x})}}_{\text{Hermitian}} du.$$

REFERENCES

- [1] A. Achim, P. Tsakalides, and A. Bezerianos. SAR image denoising via Bayesian wavelet shrinkage based on heavy-tailed modeling. *IEEE Transactions on Geoscience and Remote Sensing*, 41(8):1773–1784, 2003.
- [2] S. N. Anfinsen, A. P. Doulergis, and T. Eltoft. Estimation of the equivalent number of looks in polarimetric synthetic aperture radar imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 47(11):3795–3809, 2009.
- [3] F. Argenti, A. Lapini, T. Bianchi, and L. Alparone. A tutorial on speckle reduction in synthetic aperture radar images. *IEEE Geoscience and remote sensing magazine*, 1(3):6–35, 2013.
- [4] G. Aubert and J.-F. Aujol. A variational approach to removing multiplicative noise. *SIAM Journal on Applied Mathematics*, 68(4):925–946, 2008.
- [5] R. Bellman. *Introduction to matrix analysis*, volume 960. SIAM, 1970.
- [6] M. I. H. Bhuiyan, M. O. Ahmad, and M. Swamy. Spatially adaptive wavelet-based method using the Cauchy prior for denoising the sar images. *IEEE Transactions on Circuits and Systems for Video Technology*, 17(4):500–507, 2007.
- [7] J. M. Bioucas-Dias and M. A. Figueiredo. Multiplicative noise removal using variable splitting and constrained optimization. *IEEE Transactions on Image Processing*, 19(7):1720–1730, 2010.
- [8] 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*, 49(5):1744–1754, 2011.
- [9] P. L. Combettes and V. R. Wajs. Signal recovery by proximal forward-backward splitting. *Multiscale Modeling & Simulation*, 4(4):1168–1200, 2005.
- [10] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on image processing*, 16(8):2080–2095, 2007.
- [11] A. Danielyan, V. Katkovnik, and K. Egiazarian. Bm3d frames and variational image deblurring. *IEEE Transactions on Image Processing*, 21(4):1715–1728, 2012.
- [12] C.-A. Deledalle, L. Denis, G. Poggi, F. Tupin, and L. Verdoliva. Exploiting patch similarity for SAR image processing: the nonlocal paradigm. *IEEE Signal Processing Magazine*, 31(4):69–78, 2014.
- [13] C.-A. Deledalle, L. Denis, and F. Tupin. Iterative weighted maximum likelihood denoising with probabilistic patch-based weights. *IEEE Transactions on Image Processing*, 18(12):2661–2672, 2009.
- [14] C.-A. Deledalle, L. Denis, and F. Tupin. Nl-insar: Nonlocal interferogram estimation. *IEEE Transactions on Geoscience and Remote Sensing*, 49(4):1441–1452, 2011.
- [15] C.-A. Deledalle, L. Denis, and F. Tupin. How to compare noisy patches? Patch similarity beyond Gaussian noise. *International journal of computer vision*, 99(1):86–102, 2012.
- [16] 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*, 53(4):2021–2038, 2015.
- [17] L. Denis, F. Tupin, J. Darbon, and M. Sigelle. SAR image regularization with fast approximate discrete minimization. *IEEE Transactions on Image Processing*, 18(7):1588–1600, 2009.
- [18] S. Durand, J. Fadili, and M. Nikolova. Multiplicative noise removal using L1 fidelity on frame coefficients. *Journal of Mathematical Imaging and Vision*, 36(3):201–226, 2010.
- [19] J. Eckstein and D. P. Bertsekas. On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55(1-3):293–318, 1992.
- [20] A. Ferretti, C. Prati, and F. Rocca. Nonlinear subsidence rate estimation using permanent scatterers in differential sar interferometry. *IEEE Transactions on Geoscience and Remote Sensing*, 38(5):2202–2212, 2000.
- [21] S. Foucher. SAR image filtering via learned dictionaries and sparse representations. In *IGARSS 2008-2008 IEEE International Geoscience and Remote Sensing Symposium*, volume 1, pages 1–229. IEEE, 2008.

- [22] S. Friedland. Convex spectral functions. *Linear and multilinear algebra*, 9(4):299–316, 1981.
- [23] D. Gabay and B. Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & Mathematics with Applications*, 2(1):17–40, 1976.
- [24] G. Gilboa and S. Osher. Nonlocal operators with applications to image processing. *Multiscale Modeling & Simulation*, 7(3):1005–1028, 2008.
- [25] R. Glowinski and A. Marroco. Sur l’approximation, par éléments finis d’ordre un, et la résolution, par pénalisation-dualité d’une classe de problèmes de Dirichlet non linéaires. *Revue française d’automatique, informatique, recherche opérationnelle. Analyse numérique*, 9(2):41–76, 1975.
- [26] J. Goodman. Some fundamental properties of speckle. *Journal Optical Society of America*, 66(11):1145–1150, 1976.
- [27] N. Goodman. The distribution of the determinant of a complex Wishart distributed matrix. *The Annals of mathematical statistics*, 34(1):178–180, 1963.
- [28] F. R. Hampel. The influence curve and its role in robust estimation. *Journal of the American Statistical Association*, 69(346):383–393, 1974.
- [29] M. Hong, Z.-Q. Luo, and M. Razaviyayn. Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems. *SIAM Journal on Optimization*, 26(1):337–364, 2016.
- [30] H. Ishikawa. Exact optimization for Markov random fields with convex priors. *IEEE transactions on pattern analysis and machine intelligence*, 25(10):1333–1336, 2003.
- [31] M. Kang, S. Yun, and H. Woo. Two-level convex relaxed variational model for multiplicative denoising. *SIAM Journal on Imaging Sciences*, 6(2):875–903, 2013.
- [32] R. Karplus and J. Schwinger. A note on saturation in microwave spectroscopy. *Physical Review*, 73(9):1020, 1948.
- [33] C. Knaus and M. Zwicker. Dual-domain image denoising. In *2013 IEEE International Conference on Image Processing*, pages 440–444. IEEE, 2013.
- [34] T. Le Toan, S. Quegan, M. Davidson, H. Balzter, P. Paillou, K. Papathanassiou, S. Plummer, F. Rocca, S. Saatchi, H. Shugart, et al. The biomass mission: Mapping global forest biomass to better understand the terrestrial carbon cycle. *Remote sensing of environment*, 115(11):2850–2860, 2011.
- [35] J.-S. Lee, T. L. Ainsworth, Y. Wang, and K.-S. Chen. Polarimetric sar speckle filtering and the extended sigma filter. *IEEE Transactions on Geoscience and Remote Sensing*, 53(3):1150–1160, 2015.
- [36] 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*, 37(5):2363–2373, 1999.
- [37] J.-S. Lee, M. R. Grunes, D. L. Schuler, E. Pottier, and L. Ferro-Famil. Scattering-model-based speckle filtering of polarimetric SAR data. *IEEE Transactions on Geoscience and Remote Sensing*, 44(1):176–187, 2006.
- [38] A. S. Lewis. The convex analysis of unitarily invariant matrix functions. *Journal of Convex Analysis*, 2(1):173–183, 1995.
- [39] H. Li and Z. Qin. PolSAR image speckle reduction algorithm based on TV-PDE. In *Seventh International Symposium on Multispectral Image Processing and Pattern Recognition (MIPPR2011)*, pages 80020S–80020S. International Society for Optics and Photonics, 2011.
- [40] D. Maiwald and D. Kraus. Calculation of moments of complex Wishart and complex inverse Wishart distributed matrices. *IEE Proceedings-Radar, Sonar and Navigation*, 147(4):162–168, 2000.
- [41] M. Mäkitalo, A. Foi, D. Fevrale, and V. Lukin. Denoising of single-look SAR images based on variance stabilization and nonlocal filters. In *2010 International Conference on Mathematical Methods in Electromagnetic Theory*, pages 1–4. IEEE, 2010.
- [42] D. Massonnet and K. L. Feigl. Radar interferometry and its application to changes in the earth’s surface. *Reviews of geophysics*, 36(4):441–500, 1998.
- [43] A. Moreira, P. Prats-Iraola, M. Younis, G. Krieger, I. Hajnsek, and K. P. Papathanassiou. A tutorial on synthetic aperture radar. *IEEE Geoscience and Remote Sensing Magazine*, 1(1):6–43, 2013.
- [44] I. Najfeld and T. F. Havel. Derivatives of the matrix exponential and their computation. *Advances in Applied Mathematics*, 16(3):321–375, 1995.
- [45] X. Nie, H. Qiao, and B. Zhang. A variational model for PolSAR data speckle reduction based on the Wishart distribution. *IEEE Transactions on Image Processing*, 24(4):1209–1222, 2015.
- [46] 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*, 25(6):2620–2634, 2016.
- [47] X. Nie, B. Zhang, Y. Chen, and H. Qiao. A new algorithm for optimizing TV-based PolSAR despeckling model. *IEEE Signal Processing Letters*, 23(10):1409, 2016.
- [48] B. Rabus, M. Eineder, A. Roth, and R. Bamler. The shuttle radar topography mission – a new class of digital elevation models acquired by spaceborne radar. *ISPRS journal of photogrammetry and remote sensing*, 57(4):241–262, 2003.
- [49] L. Rudin, P.-L. Lions, and S. Osher. Multiplicative denoising and deblurring: Theory and algorithms. In *Geometric Level Set Methods in Imaging, Vision, and Graphics*, pages 103–119. Springer, 2003.
- [50] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1):259–268, 1992.
- [51] J. Shi and S. Osher. A nonlinear inverse scale space method for a convex multiplicative noise model. *SIAM Journal on Imaging Sciences*, 1(3):294–321, 2008.
- [52] G. Steidl and T. Teuber. Removing multiplicative noise by Douglas-Rachford splitting methods. *Journal of Mathematical Imaging and Vision*, 36(2):168–184, 2010.
- [53] S. Tabti, C.-A. Deledalle, L. Denis, and F. Tupin. Modeling the distribution of patches with shift-invariance: application to SAR image restoration. In *2014 IEEE International Conference on Image Processing (ICIP)*, pages 96–100. IEEE, 2014.
- [54] G. Vasile, E. Trouvé, J. Lee, and V. Buzuloiu. Intensity-driven adaptive-neighborhood technique for polarimetric and interferometric SAR parameters estimation. *IEEE Transactions on Geoscience and Remote Sensing*, 44(6):1609–1621, 2006.
- [55] Y.-W. Wen, M. K. Ng, and W.-K. Ching. Iterative algorithms based on decoupling of deblurring and denoising for image restoration. *SIAM Journal on Scientific Computing*, 30(5):2655–2674, 2008.
- [56] H. Xie, L. E. Pierce, and F. T. Ulaby. SAR speckle reduction using wavelet denoising and Markov random field modeling. *IEEE Transactions on Geoscience and Remote Sensing*, 40(10):2196–2212, 2002.
- [57] H. Xie, L. E. Pierce, and F. T. Ulaby. Statistical properties of logarithmically transformed speckle. *IEEE Transactions on Geoscience and Remote Sensing*, 40(3):721–727, 2002.