

InSAR-BM3D: A Nonlocal Filter for SAR Interferometric Phase Restoration

Francescopaolo Sica, Davide Cozzolino, Xiao Xiang Zhu¹, *Senior Member, IEEE*,
Luisa Verdoliva², *Senior Member, IEEE*, and Giovanni Poggi³

Abstract—The block-matching 3-D (BM3D) algorithm, based on the nonlocal approach, is one of the most effective methods to date for additive white Gaussian noise image denoising. Likewise, its extension to synthetic aperture radar (SAR) amplitude images, SAR-BM3D, is a state-of-the-art SAR despeckling algorithm. In this paper, we further extend BM3D to address the restoration of SAR interferometric phase images. While keeping the general structure of BM3D, its processing steps are modified to take into account the peculiarities of the SAR interferometry signal. Experiments on simulated and real-world Tandem-X SAR interferometric pairs prove the effectiveness of the proposed method.

Index Terms—Nonlocal filtering, synthetic aperture radar (SAR), SAR interferometry (InSAR).

I. INTRODUCTION

SYNTHETIC aperture radar interferometry (InSAR) is one of the most successful and widespread applications of active microwave remote sensing. By unwrapping the interferometric phase, one can generate a digital elevation map of the surface which, depending on the system parameters, may have extremely high precision. However, because of the strong speckle noise affecting SAR images, the interferometric phase may be quite noisy itself, calling for the use of some phase restoration methods before unwrapping. A straightforward solution, still widespread among practitioners, consists in performing a simple moving average on a rectangular window (boxcar filtering). Indeed, Seymour and Cumming [1] have shown that, when all involved processes are wide-sense stationarity, the sample average is the maximum-likelihood (ML) estimator for interferometric phase and coherence. However, these images are inherently nonstationary, due to the topography variations or line-of-sight displacements of the observed scenes which, on the other hand, represent the most valuable information for InSAR systems. Boxcar filtering, with its

strong smoothing action, causes a marked loss of spatial resolution, with large phase and coherence estimation errors near signal discontinuities and, ultimately, inconsistent results.

The first attempts to face the problem of signal nonstationarity in phase estimation date back to the filters proposed by Lee *et al.* [2] and by Goldstein and Werner [3], working in the spatial and frequency domains, respectively. Both filters estimate the local fringe morphology, and take advantage of it, as also done explicitly in [4]. The Lee filter adapts the shape of the estimation window to the local fringe direction in order to average similar phase values. Likewise, the Goldstein filter estimates the local power spectrum of the signal and keeps only the dominant component, corresponding again to adapting to the local direction of fringes. Moreover, if we consider the modified Goldstein filter of Baran *et al.* [5], both methods adapt to the local noise intensity, preserving the signal in case of areas that show high coherence (i.e., low noise level).

In the last 20 years, on the wake of these pioneering papers, a large number of filters have been proposed, following the main approaches in image restoration. Several papers propose improvements to the basic Lee filter, such as the local adaptive filter [6], or directional filters [7]–[9] which adapt shape and size of the filtering window to better follow the fringe morphology. Other researchers, instead, refine the Goldstein and Baran filters through a better estimate of the coherence which prevents underfiltering incoherent areas [10], [11]. Besides the frequency domain, also the wavelet domain has been considered for phase noise filtering. Lopez-Martinez and Fabregas [12] propose a wavelet filter based on the suitable modeling of the interferometric phase noise in the complex domain. Undecimated wavelet transform and wavelet packets have also been used in [13] and [14], respectively. In general, wavelet-domain filters guarantee a good preservation of spatial resolution.

Other approaches rely on suitable local models, such as [15], based on polynomial approximation, or [16], relying on sparse coding. In [17], SAR phase filtering is recast as a Bayesian estimation problem where the image prior is modeled as a Markov random field (MRF) and the filtered phase field is obtained as the configuration with maximum posterior probability. Through a suitable choice of the image prior, one can force the solution to possess some desired properties, such as a small number of residues or a smooth phase field. However, this requires modeling the prior distribution, which is still an open problem. MRFs have also been used in [18] to perform joint SAR phase and amplitude filtering.

Manuscript received April 22, 2017; revised October 25, 2017 and December 19, 2017; accepted January 23, 2018. The work of X. X. Zhu was supported in part by the European Research Council through the European Unions Horizon 2020 Research and Innovation Programme under Grant ERC-2016-StG-714087 and in part by the Helmholtz Association under the framework of the Young Investigators Group SiPEO under Grant VH-NG-1018. (Corresponding author: Giovanni Poggi.)

F. Sica and X. X. Zhu are with the German Aerospace Center (DLR), 82234 Oberpfaffenhofen, Germany (e-mail: francescopaolo.sica@dlr.de; xiao.zhu@dlr.de).

D. Cozzolino, L. Verdoliva, and G. Poggi are with the Department of Electrical Engineering and Information Technology, University of Naples Federico II, 80125 Naples, Italy (e-mail: davide.cozzolino@unina.it; verdoliv@unina.it; poggi@unina.it).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TGRS.2018.2800087

In recent years, nonlocal filtering has gained huge popularity for image restoration, including SAR data filtering. Nonlocal (NL) filtering, proposed originally [19] for natural image denoising, has been readily applied to the restoration of all kinds of SAR imagery [20], from amplitude images [21]–[23] to interferometric phase fields [24]–[28], to polarimetric images [28]–[30], and multitemporal stacks [31], [32], always with state-of-the-art results. Like previous approaches, nonlocal filtering deals with image nonstationarities by adapting estimation to the local signal behavior. However, it looks for estimator pixels not only in a local neighborhood of the target pixel but in the whole image (in practice, a large search area), relying on a suitable measure of patch similarity to select the best ones. This ensures that a large number of candidate estimators are considered. Among them, for the self-similarity property of images, it is very likely to find good estimators, even for uncommon signals.

The nonlocal approach was first applied to interferometric phase filtering by Deledalle *et al.* [24]. An iterative version of the nonlocal means algorithm was proposed, with estimator pixels selected through a probabilistic criterion based on both image intensities and interferometric phase values. Further improvements were then proposed in [25], where a pyramidal representation is used in order to better preserve textural details, and in [27], where nonlocal filtering is based on higher order singular value decomposition. In [28], instead, a unified framework is proposed where an adaptive procedure is carried out to select the best nonlocal estimation, ensuring a good preservation of radar structures and discontinuities.

In this paper, we propose a new filter for SAR interferometric phase, inspired to the nonlocal block-matching 3-D (BM3D) algorithm [33]. BM3D was proposed originally for the restoration of images corrupted by additive white Gaussian noise (AWGN). It combines nonlocal patch-based estimation with other powerful tools, such as wavelet transforms and Wiener filtering, thereby exploiting effectively both local and nonlocal redundancies. It grants major improvements with respect to its predecessors, and provides a still unsurpassed performance for AWGN denoising. Its adaptation to multiplicative noise, the SAR-BM3D algorithm [22], follows the same general principles and architecture, but uses models and tools more fit to SAR amplitude images. Like BM3D for AWGN denoising, it is arguably the state-of-the-art for SAR despeckling.

These success stories motivated us to propose, here, a nonlocal filter for SAR interferometric phase restoration, called InSAR-BM3D in the following, based on the same architecture as BM3D, but relying on suitable phase-oriented solutions. More in detail, we keep the two-pass strategy of BM3D, with the second pass driven by the pilot image generated in the first pass, and keep the grouping, collaborative filtering, and aggregation steps. However, to adapt to the peculiar statistics of interferometric signal and noise, a new *ad hoc* similarity measure is defined. Moreover, for each stack of similar blocks collected from the complex interferogram, filtering is carried out separately on the real and imaginary parts after a local decorrelating transform, using wavelet thresholding and empirical Wiener filtering. The usual aggregation of results is




Fig. 1. Two-pass filtering procedure of BM3D.

then applied before inverting all steps to recover the filtered phase.

Experiments on both simulated and real-world interferometric phase images prove the effectiveness of the proposed method in terms of both objective distortion measures and perceived image quality.

In the following, we recall briefly the structure of the BM3D algorithm (see Section II) then describe in detail the proposed InSAR-BM3D algorithm (see Section III), present the results of a number of experiments to assess performance (see Section IV), and finally draw conclusions (see Section V).

II. BM3D ALGORITHM

The BM3D denoising algorithm, originally proposed for AWGN [33], combines the nonlocal approach with wavelet-domain shrinkage and Wiener filtering. The algorithm comprises two passes, as shown in Fig. 1. In the first pass, the noisy image z is filtered to produce a basic estimate \hat{x}_1 . Then, the very same image is filtered again, but this second filtering process is driven by the first-pass estimate, leading to a largely improved result, \hat{x}_2 . Both passes consist of three steps: grouping, collaborative filtering, and aggregation, described graphically in Fig. 2.

In the *grouping* step, for each image block, a number of similar blocks are singled out all over the image and collected in a 3-D stack. The selection is based on a block-similarity measure, which takes into account the statistics of noise. By filtering jointly all blocks in the stack, both local (intra-block) and nonlocal (inter-block) dependences are considered.

These dependences are exploited in the *collaborative filtering* step to produce a denoised version of the whole stack. As the name suggests, each block of the stack collaborates to the restoration of all the others. Filtering is performed by transform-domain shrinkage as follows.

- 1) A separable 3-D wavelet transform is applied to the stack.
- 2) Shrinkage is applied to the transform coefficients to attenuate the noise component.
- 3) Inverse transform produces the denoised stack.

All filtered blocks in the stack are then returned to their original positions in the image. Since several blocks are filtered in each stack, a large number of overlapping filtered blocks are eventually generated. As a consequence, multiple estimates are available for each image pixel, which are all taken into account in the final *aggregation* step through a weighted average.

Although both passes of BM3D perform the same three steps, the second pass can exploit some precious information provided by the first one. In particular, the partially denoised pilot image can be used to improve the selection of similar

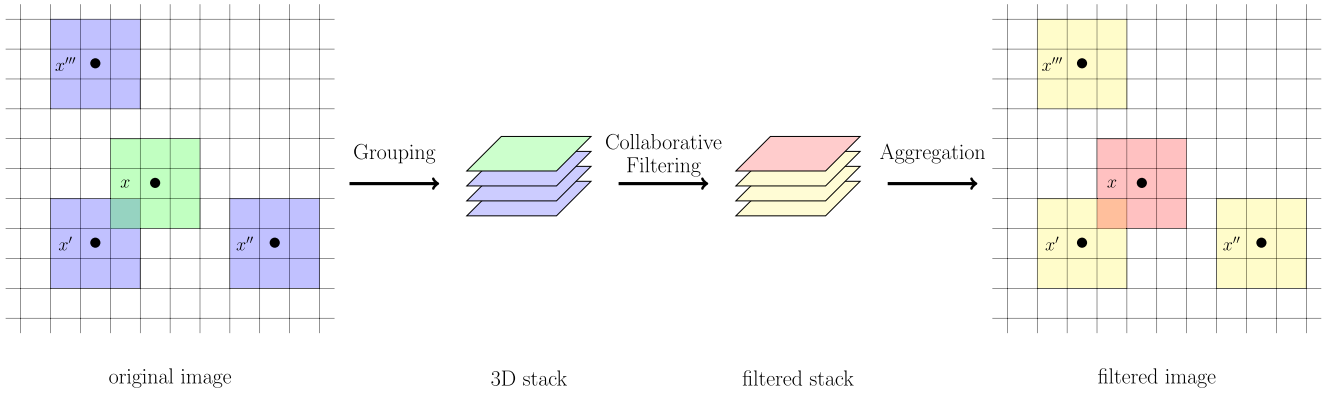


Fig. 2. Main conceptual steps of BM3D filtering.

blocks in the grouping step. Moreover, while in the first pass, a simple hard thresholding (HT) is used to perform collaborative filtering, in the second pass, a more accurate empirical Wiener filtering can be carried out, based on the statistics estimated on the pilot image. These modifications significantly improve the quality of the final denoised image with respect to the basic estimate.

III. INSAR-BM3D

In this section, we describe in detail the proposed method. Since our aim is to extend the successful ideas and tools of BM3D to the case of interferometric phase, we retain as far as possible the filtering structure of the original algorithm, introducing only the modifications necessary to adapt the filtering process to the new domain. Therefore, we keep BM3D's two-pass strategy, with the second pass driven by the pilot image output in the first pass, as well as its grouping, collaborative filtering, and aggregation steps. However, rather than working on the phase signal, which exhibits a large number of annoying 2π phase jumps, we filter (after a suitable decorrelating transform) the real and imaginary components of the interferogram, and reconstruct the denoised phase from these filtered versions.

In the following, we formalize our signal model, and then describe in detail all the steps of the algorithm.

A. Signal Model

We adopt the model proposed in [24] wherein the interferometric pair $(u_1, u_2)^T$ is expressed¹ as the product of two standard circular Gaussian random variables $(v_1, v_2)^T$ with a suitable matrix T

$$\begin{pmatrix} u_1 \\ u_2 \end{pmatrix} = T \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} \quad (1)$$

where

$$T = \begin{pmatrix} A_1 & 0 \\ A_2 \rho e^{-j\psi} & A_2 \sqrt{1 - \rho^2} \end{pmatrix} \quad (2)$$

is the Cholesky decomposition of the desired covariance matrix, depending only on the true amplitudes, A_1 and A_2 ,

¹For notational simplicity, here and in the following, we drop spatial dependences unless necessary.

interferometric phase ψ , and coherence ρ . As a consequence, the complex interferogram

$$I = u_1 u_2^* = I_0 + n \quad (3)$$

can be regarded as the sum of its true value

$$I_0 = A_1 A_2 \rho e^{j\psi} \quad (4)$$

and a zero-mean signal-dependent noise term

$$n = A_1 A_2 \rho e^{j\psi} (|v|^2 - 1) + A_1 A_2 \sqrt{1 - \rho^2} v_1 v_2^* \quad (5)$$

with covariance matrix

$$C_n = \frac{1}{2} A_1^2 A_2^2 \begin{pmatrix} 1 + \rho^2 \cos(2\psi) & \rho^2 \sin(2\psi) \\ \rho^2 \sin(2\psi) & 1 - \rho^2 \cos(2\psi) \end{pmatrix}. \quad (6)$$

The terms on the main diagonal are the variances of the real and imaginary parts of the noise. Note that, in case of multilooked data, the covariance matrix becomes simply C_n/L , with L the number of looks.

B. Grouping

The selection of the blocks, which are most similar to the target block, is based on a suitable similarity measure. In general, this measure could take into account not only the observed phase but also the observed amplitudes and coherence, as done, for example, in [24]. However, here we are only interested in estimating the true phase. The possible improvements deriving from the use of other quantities are arguably balanced by the uncertainty in their estimation and by the increased complexity. Hence, we resort to a phase-only similarity measure defined as the cosine of the phase difference

$$S(\psi_1, \psi_2) = \cos(\psi_1 - \psi_2). \quad (7)$$

The similarity is maximum when the two phases coincide (obviously, mod 2π , here and in all subsequent developments), while it is minimum when they are opposite. In addition, near the origin, $1 - S(\psi_1, \psi_2)$ approximates the Euclidean distance between the phases.

Let $\Psi_p = \{\psi_{p_1}, \dots, \psi_{p_N}\}$ be the size- N interferometric phase block centered on the target pixel p , and $\Psi_q = \{\psi_{q_1}, \dots, \psi_{q_N}\}$ be the candidate estimator block centered on pixel q . The similarity between these blocks computes as

$$S(\Psi_p, \Psi_q) = \frac{1}{N} \sum_{i=1}^N \cos(\psi_{p_i} - \psi_{q_i}) \quad (8)$$

varying from a maximum of 1, when the phase blocks coincide, to a minimum of -1 , when all phases are pairwise opposite.

In the first pass of the algorithm, lacking any side information, this basic similarity measure is used to assess the suitability of the block centered on q as a predictor for the target block

$$S_1(p, q) = S(\Psi_p, \Psi_q). \quad (9)$$

In the second pass, however, a basic estimate of the interferometric phase field is already available, and can be used to improve grouping accuracy. Replacing altogether the noisy phases ψ_p with their first-pass estimates $\hat{\psi}_p$ is not advisable, as this would reinforce the artifacts introduced in the first pass. Instead, following [33] and all subsequent multipass nonlocal filters, in the second pass, we use a weighted average of the similarity measures computed over the original noisy phases and over first-pass estimates:

$$S_2(p, q) = \tau S(\hat{\Psi}_p, \hat{\Psi}_q) + (1 - \tau)S(\Psi_p, \Psi_q). \quad (10)$$

The weight $\tau \in [0, 1]$, which establishes the relative importance of the two contributions, can be set based on theoretical considerations, when possible, or on preliminary experiments.

C. Stack-Based Parameter Estimation

Assuming that the grouping process succeeds in finding similar signal patches, the stack formed by these blocks can be loosely regarded as a collection of independent observations of the same signal patch, corrupted by different noise realizations. This is the central idea on the nonlocal approach and the ultimate reason for its success. Although only approximately true, this multilook interpretation justifies the use of nonlocal estimates, performed over the stack, which are usually much more accurate than local averages, since more samples are used. In particular, with this approach, we can easily compute the ML estimates of coherence and phase through stackwise averages.

Let U_1 and U_2 be a couple of paired stacks formed by blocks drawn from images u_1 and u_2 , respectively. Each paired stack collects the target block and the $M - 1$ most similar blocks drawn from the corresponding image and selected according to the phase similarity measures. Since we collect blocks with similar phase patterns, and of small size, we can assume the phase to be stationary over the whole stack, namely, both in space and across the blocks. Amplitude patterns, instead, are not required to be similar to one another, so we assume the amplitude to be only spatially stationary, but variable across the blocks. Accordingly, we estimate phase and coherence over the whole stack, but compute a separate amplitude estimate for each block. In detail, assuming the same amplitude over the two images, $A_1 = A_2 = A$, the ML estimate of the (square) amplitude is

$$\hat{A}_{\text{ML}}^2(m) = \frac{1}{N} \sum_{i=1}^N \frac{|U_1(m, i)|^2 + |U_2(m, i)|^2}{2} \quad (11)$$

where m and i span the M blocks of the stack and the N pixels of a block, respectively.

Then, the quantity

$$\hat{\Gamma} = \frac{1}{M} \sum_{m=1}^M \frac{\frac{1}{N} \sum_{i=1}^N U_1(m, i) U_2^*(m, i)}{\hat{A}_{\text{ML}}^2(m)} \quad (12)$$

gives us the ML estimates of coherence

$$\hat{\rho}_{\text{ML}} = |\hat{\Gamma}| \quad (13)$$

and phase

$$\hat{\psi}_{\text{ML}} = \arg(\hat{\Gamma}). \quad (14)$$

Note that by performing stackwise estimates, we renounce some resolution in favor of reliability, but this makes perfect sense considering that in the following developments, we will perform stackwise denoising.

D. Noise Decorrelating Transform

As said before, the proposed algorithm filters the real and imaginary parts of the interferogram to recover the interferometric phase only afterward. Since these components are affected by *correlated* noise, as shown by (6), they should be filtered jointly, taking into account noise correlation. As an alternative, one can decorrelate the components through an appropriate rotation, and then filter them separately. While joint processing is theoretically superior, we follow the approach based on decorrelating transforms, which is much simpler and provides excellent results in a large number of image processing problems.

For element (m, i) of the stack, the covariance matrix between the real and imaginary parts is

$$\frac{1}{2} A^4(m) \begin{pmatrix} 1 + \rho^2 \cos(2\psi) & \rho^2 \sin(2\psi) \\ \rho^2 \sin(2\psi) & 1 - \rho^2 \cos(2\psi) \end{pmatrix} \quad (15)$$

with ρ and ψ stationary over the whole stack but A depending on the block index m . Therefore, to remove this latter dependence, we consider the normalized interferogram

$$\frac{U_1(m, i) U_2^*(m, i)}{A^2(m)} = I_R + i I_I \quad (16)$$

with covariance matrix

$$C'_n = \frac{1}{2} \begin{pmatrix} 1 + \rho^2 \cos(2\psi) & \rho^2 \sin(2\psi) \\ \rho^2 \sin(2\psi) & 1 - \rho^2 \cos(2\psi) \end{pmatrix}. \quad (17)$$

The latter can be written as

$$C'_n = Q^T C_w Q \quad (18)$$

(T indicates transpose) with Q a unitary matrix whose columns are the eigenvectors of C'_n , and C_w a diagonal matrix whose entries are the eigenvalues of C'_n . Hence, to decorrelate the data, we must only premultiply them by Q (Karhunen–Loève transform)

$$z = x + w = \begin{pmatrix} z' \\ z'' \end{pmatrix} = Q \begin{pmatrix} I_R \\ I_I \end{pmatrix} \quad (19)$$

so that z' and z'' (no more real and imaginary parts of the data) have diagonal covariance matrix

$$C_z = C_w = Q C'_n Q^T = \frac{1}{2} \begin{pmatrix} 1 + \rho^2 & 0 \\ 0 & 1 - \rho^2 \end{pmatrix} \quad (20)$$

with unequal variances depending on the local coherence.

The rotation matrix

$$Q = \begin{pmatrix} \cos(\psi) & \sin(\psi) \\ -\sin(\psi) & \cos(\psi) \end{pmatrix} \quad (21)$$

can be computed exactly given the true phase value ψ , namely, the very same quantity we want to estimate. However, for our needs, we do not need perfect decorrelation in this preliminary step, and can resort to the ML estimates of Section III-C.

E. Collaborative Filtering

We can now proceed with the actual filtering step, which is performed by wavelet-domain shrinkage applied separately to both components

$$z' = x' + w' \text{ and } z'' = x'' + w'' \quad (22)$$

of each rotated stack. Let us indicate with capital letters the corresponding transforms

$$Z' = X' + W' \text{ and } Z'' = X'' + W''. \quad (23)$$

Since the two components are filtered exactly in the same way, in the following, we will discard the apexes.

In the first pass of the algorithm, we use a 2-D biorthogonal wavelet for spatial transform and a 1-D Haar wavelet along the blocks. Then, a simple HT is applied, as in BM3D [33]. Accordingly, the first-pass estimate of the signal coefficient is

$$\hat{X}_1 = \begin{cases} Z & \text{if } |Z| > Z_{\text{th}} \\ 0 & \text{otherwise.} \end{cases} \quad (24)$$

The threshold is given by the noise standard deviation σ_w multiplied by a constant λ

$$Z_{\text{th}} = \lambda \sigma_w. \quad (25)$$

The standard deviation σ_w , instead, can be derived by the blockwise estimates of coherence

$$\hat{\sigma}_w = \sqrt{(1 \pm \hat{\rho}^2)/2} \quad (26)$$

with the plus[minus] sign used for the first[second] component. In BM3D, the λ parameter is fixed to 2.7 through preliminary experiments, and we found this value satisfactory for our case.

In the second pass, instead, discrete cosine transform (DCT) and Haar wavelet transform are used for 2-D and 1-D transforms, respectively. The clean coefficient \hat{X}_2 is estimated through Wiener shrinkage as

$$\hat{X}_2 = \max \left(0, \frac{\hat{X}_1^2 - \sigma_w^2}{\hat{X}_1^2} \right) \cdot Z. \quad (27)$$

In this case, since an estimate of the true signal is already available, this can be used to estimate the noise variance based on the difference between the noisy quantity and its basic estimate.

F. Aggregation

In this phase, the algorithm follows closely the original BM3D, except for the need to work in parallel on two images. All patches of a filtered stack are returned to their original location, as shown in Fig. 2. Since for each pixel, multiple estimates are obtained, due to overlapping patches, they are all averaged, with suitable weights based on the estimate reliability, to produce the final filtered image.

IV. EXPERIMENTAL ANALYSIS

In this section, we describe the results of a number of experiments carried out to assess the performance of the proposed method. Experiments are performed on both simulated and real SAR interferometric data, and results are compared with those of several widespread reference methods.

Simulated data allow us to investigate the performance in controlled conditions by objective measures and to study the dependence on the main signal features, especially coherence. On the other hand, simulation methods cannot replicate the complex features of real SAR data; Hence, we will also conduct visual inspection on the output of the filters applied to some real-world images.

As reference methods, we will consider the following:

- 1) boxcar filter with a 5×5 window;
- 2) Lee filter [2] with 16 25/27-pixel directional windows;
- 3) Goldstein filter [3] with $\text{dimWindow} = 21$ and $\alpha = 0.9$;
- 4) NL-SAR [28] with $\text{dimWindowMax} = 25$ and $\text{dimPatchMax} = 11$;
- 5) NL-InSAR [24] with $\text{dimPatch} = 7$ and $\text{dimWindow} = 21$.

Note that all parameters are set, when applicable, as suggested by the authors of the original papers, or else chosen to optimize the performance. Turning to the proposed method, all parameters are fixed in advance, with no need of user intervention. Most of them are inherited from BM3D, where an extensive experimental analysis has been carried out to optimize them. Others, more critical and task-specific, have been set only after some preliminary experiments. In particular, we studied how the performance depends on the threshold used for shrinkage λ and on the spatial and stackwise transforms. Overall, the performance appears to depend only weakly on these choices, so we eventually used the original setting, since it has been field-tested on a much wider set of sources, providing good guarantees of robustness. In the grouping phase, the image is scanned with subsampling step 3. For each 8×8 pixel target patch, the 64 most similar patches in a 21×21 pixel search window are collected, using the distances of (9) and (10), respectively, in the two passes. The stack is then transformed along spatial and stackwise coordinates for the collaborative filtering. Spatial transforms are wavelet with bior1.5 basis (three levels) in the first pass and DCT in the second pass, while the Haar wavelet (four levels) is always used along the stack. Finally, in the aggregation phase, a Kaiser window with $\beta = 0$ is used.

To ensure reproducibility of results, the executable code of the proposed method and the simulated

images used in the experiments are available online at <http://www.grip.unina.it/web-download.html>.

A. Experiments on Simulated Data

We generated four 256×256 pixel images of interest for typical interferometric application: a *Cone* with circular symmetry and constant phase gradient along the radius, a richer image, *Peaks*, obtained by summing some Gaussian profiles, which simulates a more complex scenery with a varied geometry, a straight *Ramp* with variable phase gradient, with fringe separation going from 28 to 8 pixels, and finally an image, *Squares*, with a large number of phase jumps as might occur in some urban scenarios. Data are simulated by generating two circular complex standard Gaussian processes and correlating them through the matrix T in (2). In order to gain insight into the dependence of filtering quality on phase noise intensity, coherence grows linearly from left to right, from a minimum of 0.1 (very noisy) to a maximum of 0.9 (almost clean). The amplitude does not impact on noise intensity, but enters the similarity measure of NL-InSAR and NL-SAR and therefore might affect their performance. Therefore, in the first two images, the amplitude grows from 21 to 255 digital numbers along the y -axis, allowing us to study its effect on the output independently of the coherence. In the third image, it is kept constant in order not to interfere with fringe frequency variations. In the fourth image, it follows the clean phase (plus an offset to keep positive values) to account for amplitude-phase correlations in urban scenarios. Fig. 3 shows, for each of these images, the noisy original phase image (left) and the noiseless reference (right) with the usual false-color representation in $[-\pi, \pi]$ as shown by the color bar on the right.

On these images, we will compute two objective measures of performance, the root mean square error (RMSE) between the filtered phase image and the corresponding clean reference and the number of residues [34] remaining after filtering.

Figs. 4–7 show, for all simulated scenes, the filtered images (top) and the corresponding error images (bottom) for both the proposed and the reference techniques. Numerical results, RMSE and number of residues, are instead collected in Tables I and II, respectively, where we report the average over ten repetitions of the experiment (standard deviations, always quite low, are not shown for the sake of clarity). Moreover, in order to gain insight into the dependence on coherence, in Fig. 8, we also show curves obtained by computing RMSE and number of residues in sliding-window modality over vertical strips of width 32 pixels. No similar figures are devoted to explore the dependence on amplitude, which seems very limited.

All methods work fairly well on the right of the images (high coherence) although the boxcar filter shows a clear loss of resolution and the Lee filter produces some sparse residues on the flat-phase region of *Peaks*. On the contrary, moving toward the medium- and low-coherence areas, in the center and left of the images, large differences in performance emerge. In the absence of phase jumps, InSAR-BM3D provides by far the best results, preserving the structure of the original phase

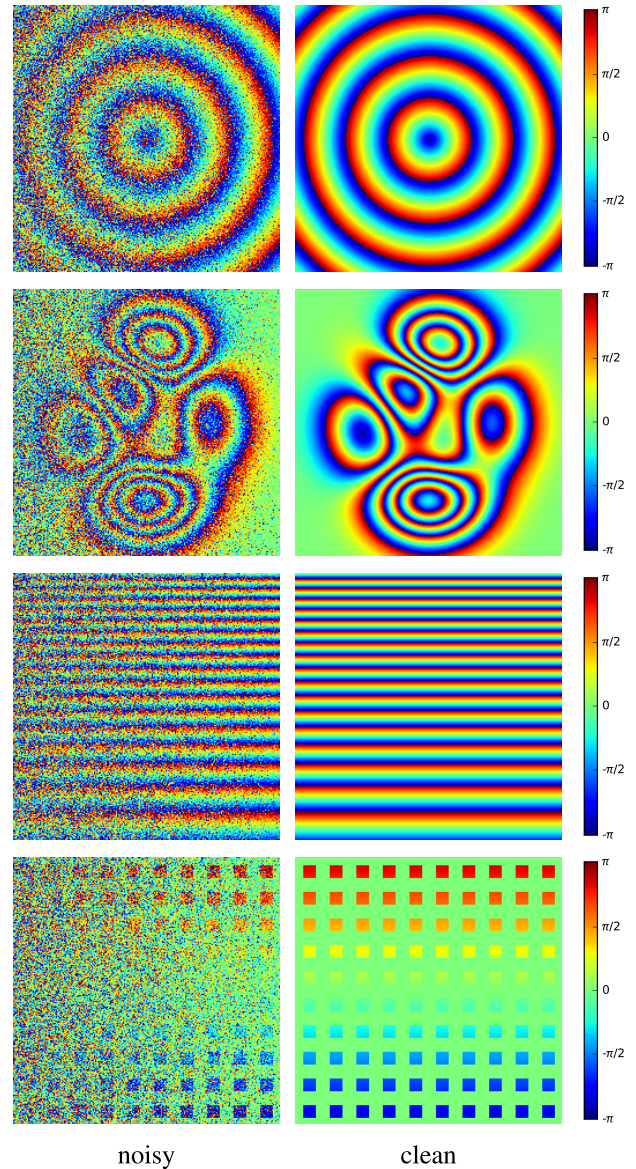


Fig. 3. (Left) Noisy simulated images used in the experiments and (Right) corresponding noiseless references. (From top to bottom) Cone, Peaks, Ramp, and Squares. In all images, the coherence increases from 0.1 to 0.9 going from left to right.

fringes and giving rise to very few residues. NL-InSAR also performs generally well, except in regions characterized by low coherence and rapid phase variations, like the center of *Peaks* or the left of *Ramp*. In these regions, in fact, only a small number of patches similar to the target fall in the search area, and these have low quality. This interpretation is confirmed by the fact that the results improve again when fringe frequency increases beyond a critical level (image top) because multiple fringes now fit in the search area, increasing the number of patches similar to the target. The other reference methods, and especially the Goldstein filter, work rather poorly in low-coherence areas, which is clear from visual inspection and confirmed by the sliding-window performance curves. Note that the amplitude variations in *Cone* and *Peaks* do not produce appreciable effects on any filter. In the presence of phase jumps, that is, in the *Squares* image, all methods work rather

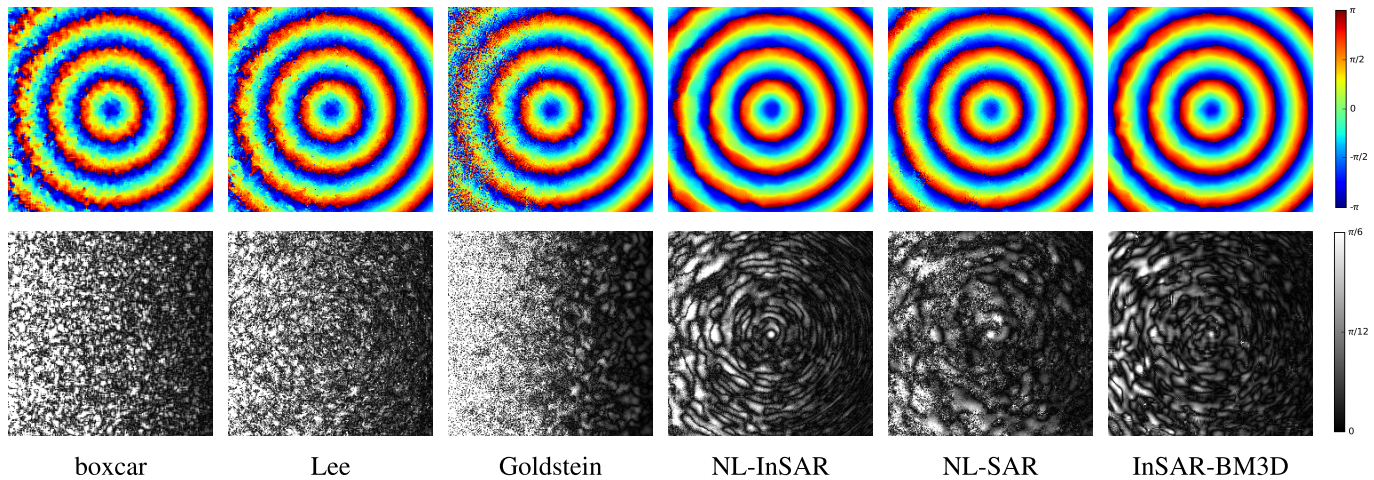


Fig. 4. (Top) Filtered images and (Bottom) error images generated by reference methods and InSAR-BM3D for the Cone noisy image. The error images of NL-InSAR and InSAR-BM3D, with patches oriented along phase fringes, prove these filters' ability to restore the general structure of the clean phase image.

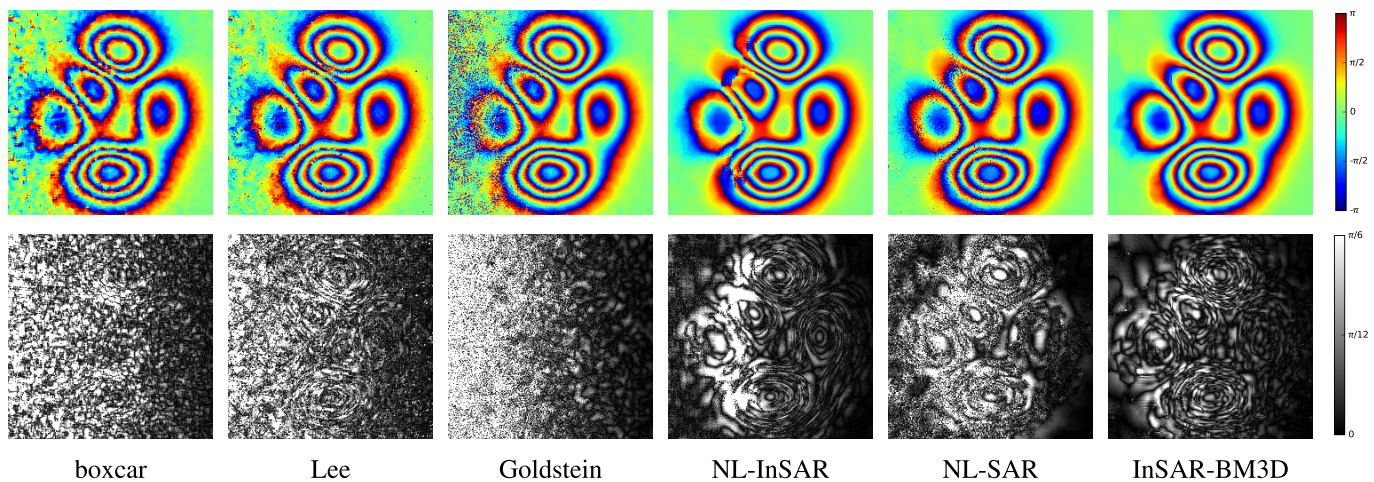


Fig. 5. (Top) Filtered images and (Bottom) error images generated by reference methods and InSAR-BM3D for the Peak noisy image. NL-InSAR outputs a very smooth phase image but does not recover well high-frequency phase fringes. In this case, InSAR-BM3D outperforms all reference methods.

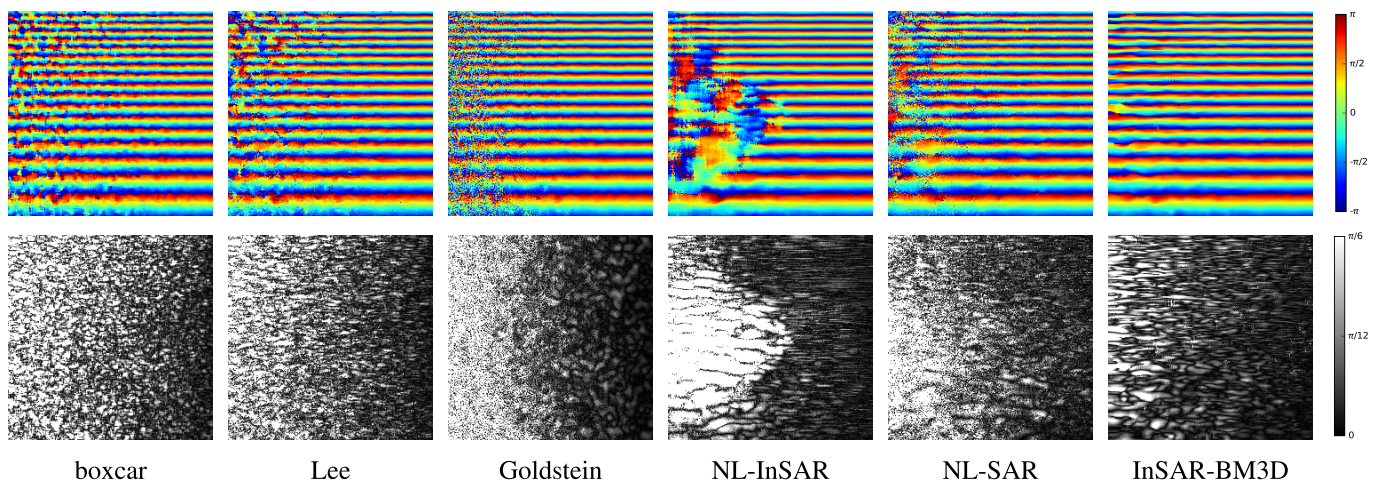


Fig. 6. (Top) Filtered images and (Bottom) error images generated by reference methods and InSAR-BM3D for the Ramp noisy image. (Right) All methods work well on the high coherence, (Left) but only InSAR-BM3D provides satisfactory results also on the low coherence. The behavior of NL-InSAR and NL-SAR at low coherence also depends on the local fringe frequency, which impacts on the number of good patches available for filtering.

poorly, with a sharp quality impairment especially, but not exclusively, in the low-coherence left region. On this image, NL-InSAR and NL-SAR could be expected to outperform the

other methods since they use a similarity measure that takes into account jointly phase and amplitude, strongly correlated by construction. However, this does not seem to be the case.

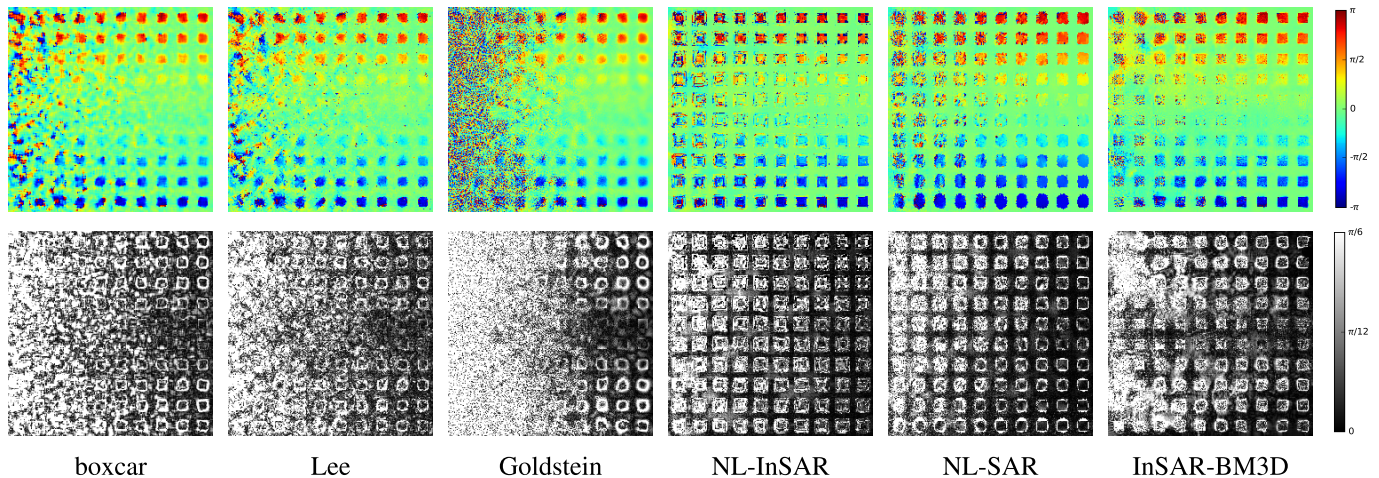


Fig. 7. (Top) Filtered images and (Bottom) error images generated by reference methods and InSAR-BM3D for the Squares noisy image. This is a very challenging case and all methods perform rather poorly, even in the high-coherence region, due to the large number of phase jumps. InSAR-BM3D provides the lowest RMSE, while Boxcar and Lee guarantee the least residuals, due to their strong smoothing.

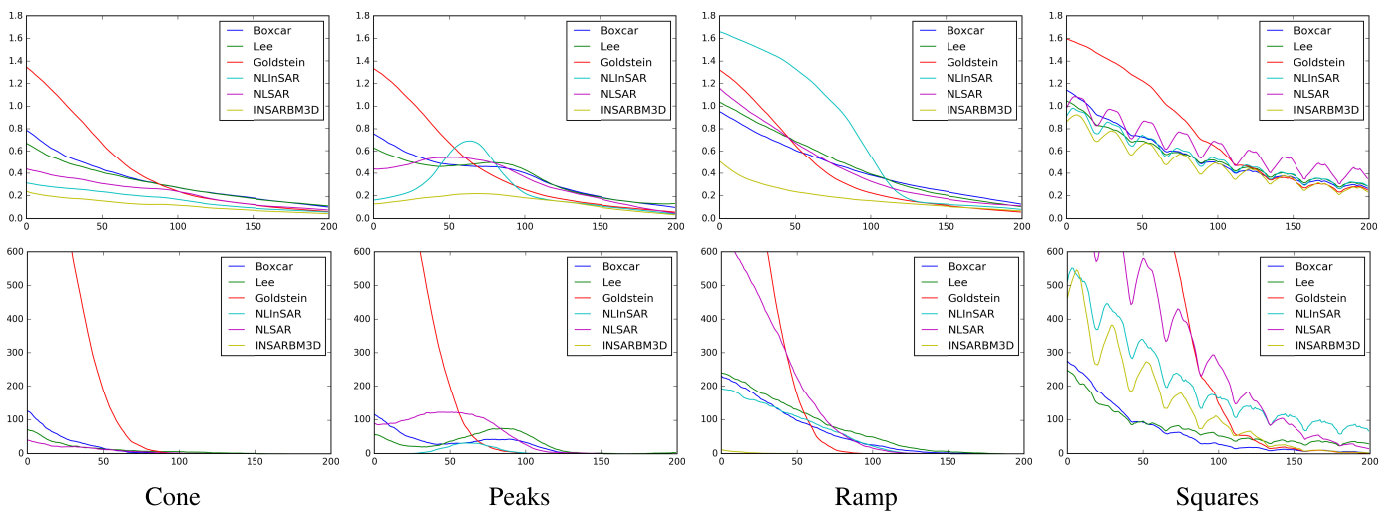


Fig. 8. Sliding-window performance indicators. (Top) RMSE and (Bottom) number of residues for all simulated images. The indicators are computed on a vertical slice of width 32 pixels sliding over the image. Performance improves going from left (low coherence) to right (high coherence). At low coherence, InSAR-BM3D provides consistently the lowest RMSE, while at high coherence, all methods have a similar performance. A similar behavior is observed for the number of residues, with InSAR-BM3D generating a very small number of residues even in low-coherence areas, in the absence of phase jumps. The performance of NL-InSAR, otherwise competitive, is clearly impaired by the scarcity of good patches at some phase slopes (see Ramp). Note also the bimodal behavior of the Goldstein filter, which is almost useless at low coherence, and among the best methods at medium-to-high coherence. On Squares, all methods exhibit a much worse performance, with no large differences among one another.

All these observations find further confirmation in the synthetic results of Tables I and II. In particular, InSAR-BM3D is characterized by the smallest RMSE (followed closely by NL-InSAR only for Cone) and only a few occasional residues, except for Squares.

The phase error images, obtained as the difference between filtered images and clean reference² besides confirming the above observations, provide some more information on the noise-rejection behavior of the algorithms. Patch-based non-local methods generate errors that are structured in patches as well, generally aligned with the phase fringes of the

²It is worth underlining that these error images are not “method noise” images. The former are computed with respect to the noiseless original and, ideally, should be as smooth as possible, the latter are computed with respect to the noisy original and, ideally, should contain all and only the image noise.

TABLE I
RMSE FOR PROPOSED AND REFERENCE METHODS

	Cone	Peaks	Ramp	Squares
boxcar	0.413	0.438	0.534	0.660
Lee	0.371	0.408	0.580	0.630
Goldstein	0.653	0.654	0.642	0.944
NL-InSAR	0.194	0.323	0.976	0.650
NL-SAR	0.269	0.376	0.606	0.767
InSAR-BM3D	0.145	0.155	0.254	0.589

clean reference. This indicates the ability to correctly reproduce the correct phase structures, accounting for noise spikes through small coherent spatial shifts of the fringes themselves.




Fig. 9. Phase and amplitude of real images used in the experiments. (From left to right) Wei Hai, Java Surabaya, and Euskirchen (crop highlighted by black box).




Fig. 10. (Top) Filtered images and (Bottom) method noise images generated by reference methods and InSAR-BM3D for the Wei Hai image. Several reference methods (Goldstein, NL-InSAR, and NL-SAR) provide smoother and more pleasant images than InSAR-BM3D. This may come at the cost of detail loss (not measurable without a ground truth) and staircase artifacts (NL-InSAR and NL-SAR). Ideally, the method noise should be a white-noise field. NL-InSAR and NL-SAR depart markedly from this behavior, indicating possible artifacts.

TABLE II
RESIDUES FOR PROPOSED AND REFERENCE METHODS

	Cone	Peaks	Ramp	Squares
boxcar	167.6	224.0	488.1	538.1
Lee	110.8	205.1	580.2	612.0
Goldstein	2049.0	2025.8	2016.3	4557.4
NL-InSAR	0.0	37.0	445.2	1611.0
NL-SAR	69.1	356.9	1209.9	2351.9
InSAR-BM3D	0.0	0.2	13.6	1114.9

For other methods, instead, phase errors have a fine-grain and even white-noise structure (especially in the more noisy right part) showing a tendency of the filtered images to “follow” the local noise pattern rather than recovering the large-scale structure of the noiseless reference.

B. Experiments on Real-World SAR Data

To assess the performance of the proposed method on real-world SAR images, a set of three Tandem-X pairs has been used. The noisy phase images are shown in Fig. 9, together with the corresponding amplitude images, limited to the 256×256 pixel section used in the experiments. In the following, they will be referred to as Wei Hai, Java Surabaya,

and Euskirchen, after the regions of acquisition, in China, Indonesia, and Germany, respectively. All images have the same spatial resolution while showing different phase patterns and noise levels. Figs. 10–12 show, for each of these scenes, the filtered images (top) and the corresponding method noise images (bottom) for both the proposed and the reference techniques. Since for real-world images, there is no clean reference, we cannot compute phase error images. However, the method noise image, obtained as the difference between the filtered image and the noisy original (rather than clean reference), also provides some valuable information on a method behavior. In the ideal case of perfect filtering, signal and noise are perfectly separated; hence, the filtered image corresponds to the ideal noise-free phase image, and the method noise image contains only the noise component. On the contrary, with imperfect filtering, some residual noise will affect the filtered image, while some signal components leak in the method noise image. Therefore, method noise images showing traces of signal structures or departing from the expected noiselike behavior are good indicators of filtering imperfections.

The visual inspection of filtered images confirms, in general, the observations made for the synthetic data. The boxcar filter generates basically a low-pass version of the noisy interferogram, showing blurred fringe edges as it is reasonable to expect. By adapting the shape of the estimation window to the fringe direction, the Lee filter preserves the details

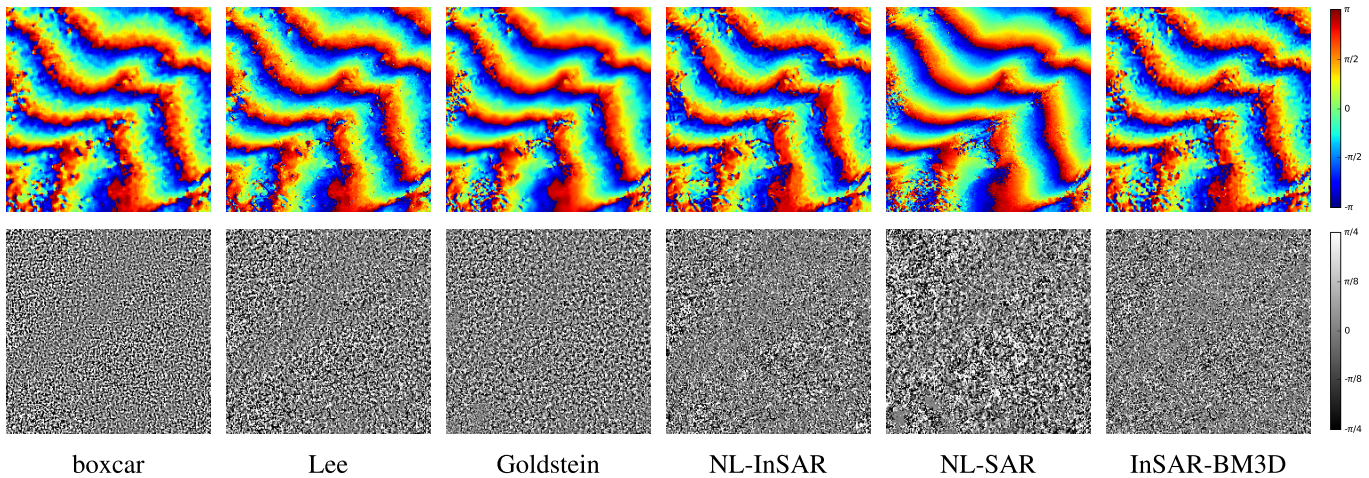


Fig. 11. (Top) Filtered images and (Bottom) method noise images generated by reference methods and InSAR-BM3D for the Java image.

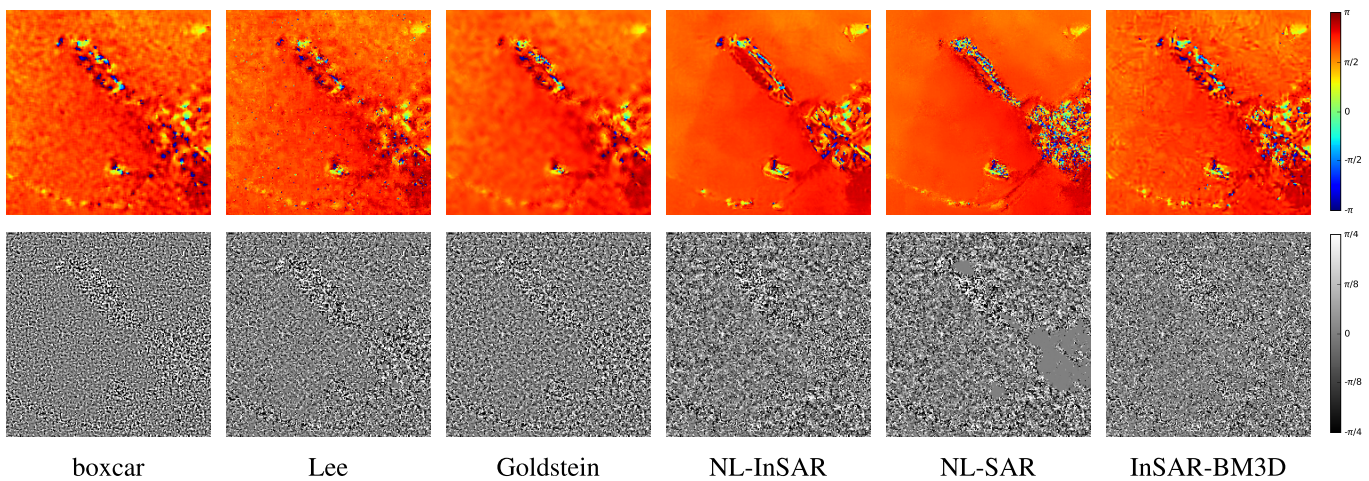


Fig. 12. (Top) Filtered images and (Bottom) method noise images generated by reference methods and InSAR-BM3D for the Euskirchen image. NL-InSAR apparently loses some details due to oversmoothing (see also Fig. 13). NL-SAR does not filter at all some high-activity areas, as obvious from the flat areas in the method noise image. Boxcar, Goldstein, and InSAR-BM3D produce similar results that are semantically similar but with increasing spatial resolution.

much better than the boxcar filter while keeping the same denoising capability. On the down side, this mechanism can also lead to underfiltered areas with residual noise that appears as isolated spikes over the image. Turning to the Goldstein filter, this provides quite a good performance in terms of both detail preservation and noise suppression, which is somewhat surprising given the poor results observed on synthetic data. However, remember that the Goldstein filter did work already well on the *right* of the simulated images, those at high coherence, which is the case of our real images, characterized by a relatively high signal-to-noise ratio. In fact, this filter relies on the estimation of a few frequency components of the signal in the target pixel neighborhood. When data are not reliable enough, these estimates impair rapidly, leading to the poor results observed on the *left* of the simulated images.

As for the nonlocal methods, they all guarantee a strong noise suppression. NL-SAR, as already observed on synthetic data, seems to oversmooth the image, in general, losing some details. This can also be observed in the method noise images, where traces of the signal are visible. NL-InSAR, instead, offers both strong noise suppression and good detail preservation, providing filtered images that are very sharp and visually appealing, like in the case of Euskirchen where large

smooth areas are obtained. In some cases, however, there seems to be a conflict between sharpness and fidelity to the original image structures. With reference again to Euskirchen, for example, some “yellow” spots in the bottom of the image are emphasized, while some more, quite similar to the former, are removed. In the center of the image, instead, a whole region characterized by small phase structures is reduced to a uniform region (see also the enlarged crop of Fig. 13). Lacking a ground truth, we cannot speak of errors, but results are dubious in these cases. In addition, in areas characterized by a linearly varying phase, NL-InSAR tends to produce a staircase output signal, with sharp transitions that may be visually appealing but hardly correct. This problem, common to most filters based on nonlocal means, becomes relevant in interferogram filtering, where regions with linear phase variations occur frequently due to constant slope terrains. The staircase effect of NL-InSAR is even more evident in the unwrapped phase image of Fig. 14 (center) while it is absent when boxcar (left) or InSAR-BM3D (right) is used for filtering. Finally, the proposed InSAR-BM3D seems being able to remove a great deal of noise, ensuring also good detail preservation. The output phase fringes are less sharp than those of NL-InSAR, but all details of interest survive, and




Fig. 13. Some filtering results for a crop (see Fig. 9) of the Euskirchen image. At a visual inspection, InSAR-BM3D looks more faithful to the original than NL-InSAR.




Fig. 14. Some unwrapping results for the Weihai image. InSAR-BM3D and NL-InSAR provide a smoother result than boxcar. NL-InSAR, however, introduces some staircase effects.

no staircase effect is observed. Indeed, the structure of the filtered phase is very similar to that of boxcar or Lee filtering, only less noisy. If the boxcar filtered image is regarded simply as a low-pass version of the original, this observation supports our claim about the fidelity of InSAR-BM3D filtering.

An analysis of the method noise images provides further support to the above observations. The method noise images produced by InSAR-BM3D and boxcar filtering are the most similar to white-noise fields, expected with perfect filtering. Accordingly, signal structures are well preserved. For the other techniques, method noise images appear generally more grainy. This holds especially for NL-InSAR and NL-SAR, the latter showing on Euskirchen both a leakage of signal structures (oversmoothing) and a flat area (no filtering).

V. CONCLUSION

We have proposed InSAR-BM3D, a new nonlocal filter for SAR interferometric phase images, which keeps the algorithmic structure of BM3D, but adapts the individual processing steps to the peculiarities of InSAR data. Experiments on simulated images show the proposed filter to outperform all references in terms of objective quality measures, mean square error, and number of residues. The performance gain is especially relevant in areas characterized by low-coherence or rapid phase variations. These observations are mostly confirmed by the visual inspection of filtering results on several real-world high-coherence InSAR images. Clearly, a much larger experimental basis is necessary to draw solid conclusions about the merits of the proposed filter. To this end, and also to guarantee the full reproducibility of research, both the proposed filter and the simulated InSAR images are made available online.

Experiments on real InSAR images show that several filters, and in particular, the Goldstein filter, NL-InSAR, and InSAR-BM3D, provide valuable results, although with very different features. This suggests a path for future research,

which may focus on the development of a new method which keeps the best features of these different approaches, or some form of postfiltering fusion, as proposed in [35]. A further topic of interest for future work is the development and testing of methods based on the joint processing of all data, rather than on the use of decorrelating transforms.

ACKNOWLEDGMENT

The authors would like to thank G. Baier for preparing the data sets. The TanDEM-X data were obtained through Project NTI_INSA0098—Multi-Channel Phase Unwrapping for TanDEM-X.

REFERENCES

- [1] M. S. Seymour and I. G. Cumming, "Maximum likelihood estimation for SAR interferometry," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, Aug. 1994, pp. 2272–2275.
- [2] J.-S. Lee, K. P. Papathanassiou, T. L. Ainsworth, M. R. Grunes, and A. Reigber, "A new technique for noise filtering of SAR interferometric phase images," *IEEE Trans. Geosci. Remote Sens.*, vol. 36, no. 5, pp. 1456–1465, Sep. 1998.
- [3] R. M. Goldstein and C. L. Werner, "Radar interferogram filtering for geophysical applications," *Geophys. Res. Lett.*, vol. 25, no. 21, pp. 4035–4038, Nov. 1998.
- [4] E. Trouvé, J.-M. Nicolas, and H. Maître, "Improving phase unwrapping techniques by the use of local frequency estimates," *IEEE Trans. Geosci. Remote Sens.*, vol. 36, no. 6, pp. 1963–1972, Nov. 1998.
- [5] I. Baran, M. P. Stewart, B. M. Kampes, Z. Perski, and P. Lilly, "A modification to the Goldstein radar interferogram filter," *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 9, pp. 2114–2118, Sep. 2003.
- [6] N. Wu, D.-Z. Feng, and J. Li, "A locally adaptive filter of interferometric phase images," *IEEE Geosci. Remote Sens. Lett.*, vol. 3, no. 1, pp. 73–77, Jan. 2006.
- [7] S. Fu, X. Long, X. Yang, and Q. Yu, "Directionally adaptive filter for synthetic aperture radar interferometric phase images," *IEEE Trans. Geosci. Remote Sens.*, vol. 51, no. 1, pp. 552–559, Jan. 2013.
- [8] C.-F. Chao, K.-S. Chen, and J.-S. Lee, "Refined filtering of interferometric phase from InSAR data," *IEEE Trans. Geosci. Remote Sens.*, vol. 51, no. 12, pp. 5315–5323, Dec. 2013.
- [9] G. Vasile, E. Trouvé, J.-S. Lee, and V. Buzuloiu, "Intensity-driven adaptive-neighborhood technique for polarimetric and interferometric SAR parameters estimation," *IEEE Trans. Geosci. Remote Sens.*, vol. 44, no. 6, pp. 1609–1621, Jun. 2006.
- [10] R. Song, H. Guo, G. Liu, Z. Perski, and J. Fan, "Improved Goldstein SAR interferogram filter based on empirical mode decomposition," *IEEE Geosci. Remote Sens. Lett.*, vol. 11, no. 2, pp. 399–403, Feb. 2014.
- [11] M. Jiang *et al.*, "The improvement for Baran phase filter derived from unbiased InSAR coherence," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 7, no. 7, pp. 3002–3010, Jul. 2014.
- [12] C. Lopez-Martinez and X. Fabregas, "Modeling and reduction of SAR interferometric phase noise in the wavelet domain," *IEEE Trans. Geosci. Remote Sens.*, vol. 40, no. 12, pp. 2553–2566, Dec. 2002.
- [13] Y. Bian and B. Mercer, "Interferometric SAR phase filtering in the wavelet domain using simultaneous detection and estimation," *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 4, pp. 1396–1416, Apr. 2011.
- [14] X. Zha, R. Fu, Z. Dai, and B. Liu, "Noise reduction in interferograms using the wavelet packet transform and Wiener filtering," *IEEE Geosci. Remote Sens. Lett.*, vol. 5, no. 3, pp. 404–408, Jul. 2008.
- [15] J. Bioucas-Dias, V. Katkovnik, J. Astola, and K. Egiazarian, "Absolute phase estimation: Adaptive local denoising and global unwrapping," *Appl. Opt.*, vol. 47, no. 29, pp. 5358–5369, 2008.
- [16] H. Hongxing, J. M. Bioucas-Dias, and V. Katkovnik, "Interferometric phase image estimation via sparse coding in the complex domain," *IEEE Trans. Geosci. Remote Sens.*, vol. 53, no. 5, pp. 2587–2602, May 2015.
- [17] G. Ferraiuolo and G. Poggi, "A Bayesian filtering technique for SAR interferometric phase fields," *IEEE Trans. Image Process.*, vol. 13, no. 10, pp. 1368–1378, Oct. 2004.
- [18] L. Denis, F. Tupin, J. Darbon, and M. Sigelle, "Joint regularization of phase and amplitude of InSAR data: Application to 3-D reconstruction," *IEEE Trans. Geosci. Remote Sens.*, vol. 47, no. 11, pp. 3774–3785, Nov. 2009.

- [19] A. Buades, B. Coll, and J.-M. Morel, "A review of image denoising algorithms, with a new one," *Multiscale Model. Simul.*, vol. 4, no. 2, pp. 490–530, 2005.
- [20] C.-A. Deledalle, L. Denis, G. Poggi, F. Tupin, and L. Verdoliva, "Exploiting patch similarity for SAR image processing: The nonlocal paradigm," *IEEE Signal Process. Mag.*, vol. 31, no. 4, pp. 69–78, Jul. 2014.
- [21] C.-A. Deledalle, L. Denis, and F. Tupin, "Iterative weighted maximum likelihood denoising with probabilistic patch-based weights," *IEEE Trans. Image Process.*, vol. 18, no. 12, pp. 2661–2672, Dec. 2009.
- [22] S. Parrilli, M. Poderico, C. V. Angelino, and L. Verdoliva, "A nonlocal SAR image denoising algorithm based on LLMMSE wavelet shrinkage," *IEEE Trans. Geosci. Remote Sens.*, vol. 50, no. 2, pp. 606–616, Feb. 2012.
- [23] D. Cozzolino, S. Parrilli, G. Scarpa, G. Poggi, and L. Verdoliva, "Fast adaptive nonlocal SAR despeckling," *IEEE Geosci. Remote Sens. Lett.*, vol. 11, no. 2, pp. 524–528, Feb. 2014.
- [24] C.-A. Deledalle, L. Denis, and F. Tupin, "NL-InSAR: Nonlocal interferogram estimation," *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 4, pp. 1441–1452, Apr. 2011.
- [25] R. Chen, W. Yu, R. Wang, G. Liu, and Y. Shao, "Interferometric phase denoising by pyramid nonlocal means filter," *IEEE Geosci. Remote Sens. Lett.*, vol. 10, no. 4, pp. 826–830, Jul. 2013.
- [26] X. X. Zhu, R. Bamler, M. Lachaise, F. Adam, Y. Shi, and M. Eineder, "Improving TanDEM-X DEMs by non-local InSAR filtering," in *Proc. EUSAR*, 2014, pp. 1–4.
- [27] X. Lin, F. Li, D. Meng, D. Hu, and C. Ding, "Nonlocal SAR interferometric phase filtering through higher order singular value decomposition," *IEEE Geosci. Remote Sens. Lett.*, vol. 12, no. 4, pp. 806–810, Apr. 2015.
- [28] 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 Trans. Geosci. Remote Sens.*, vol. 53, no. 4, pp. 2021–2038, Apr. 2015.
- [29] J. Chen, Y. Chen, W. An, Y. Cui, and J. Yang, "Nonlocal filtering for polarimetric SAR data: A pretest approach," *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 5, pp. 1744–1754, May 2011.
- [30] J. Hu, R. Guo, X. Zhu, G. Baier, and Y. Wang, "Non-local means filter for polarimetric SAR speckle reduction-experiments using TerraSAR-X data," in *Proc. ISPRS Ann. Photogramm., Remote Sens. Spatial Inf. Sci.*, 2015, pp. 71–77.
- [31] X. Su, C. Deledalle, F. Tupin, and H. Sun, "Two-step multitemporal nonlocal means for synthetic aperture radar images," *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 10, pp. 6181–6196, Nov. 2014.
- [32] F. Sica, D. Reale, G. Poggi, L. Verdoliva, and G. Fornaro, "Nonlocal adaptive multilooking in SAR multipass differential interferometry," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 8, no. 4, pp. 1727–1742, Apr. 2015.
- [33] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.
- [34] R. Bamler and P. Hartl, "Synthetic aperture radar interferometry," *Inv. Probl.*, vol. 14, no. 4, pp. R1–R54, 1998.
- [35] D. Gragnaniello, G. Poggi, G. Scarpa, and L. Verdoliva, "SAR image despeckling by soft classification," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 9, no. 6, pp. 2118–2130, Jun. 2016.



Francesco Paolo Sica received the Laurea (M.S.) degree (*summa cum laude*) in telecommunication engineering and the Ph.D. degree in information engineering from the University of Naples Federico II, Naples, Italy, in 2012 and 2016, respectively.

He is currently a Researcher with the Microwave and Radar Institute, German Aerospace Center (DLR), Oberpfaffenhofen, Germany. From 2014 to 2016, he was a Visiting Student with the Remote Sensing Technology Institute, DLR. His research interests include the processing and classification of synthetic aperture radar images with specific application to single and multibaseline interferometry.

Dr. Sica serves regularly as a Reviewer for the IEEE TRANSACTIONS ON GEOSCIENCE AND REMOTE SENSING, the IEEE GEOSCIENCE AND REMOTE SENSING LETTERS, and the IEEE JOURNAL OF SELECTED TOPICS IN APPLIED EARTH OBSERVATIONS AND REMOTE SENSING.



Davide Cozzolino received the Laurea degree in computer engineering and the Ph.D. degree in information engineering from the University Federico II of Naples, Naples, Italy, in 2011 and 2015, respectively.

He currently holds a post-doctoral position at the University Federico II of Naples. His research interests include image processing and deep learning, with applications to remote sensing (image analysis and restoration) and multimedia forensics (image forgery detection and localization).

Dr. Cozzolino serves as a Reviewer for the IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY, the IEEE GEOSCIENCE AND REMOTE SENSING LETTERS, and *Remote Sensing* (MDPI).



Xiao Xiang Zhu (S'10–M'12–SM'14) received the M.Sc., Dr.-Ing., and Habilitation degrees from the Technical University of Munich (TUM), Munich, Germany, in 2008, 2011, and 2013, respectively, all in signal processing.

She was a Guest Scientist or Visiting Professor with the Institute for Electromagnetic Sensing of the Environment, Italian National Research Council, Naples, Italy, Fudan University, Shanghai, China, the University of Tokyo, Tokyo, Japan, and the University of California at Los Angeles, Los Angeles, CA, USA, in 2009, 2014, 2015, and 2016, respectively. She is currently a Professor with the Signal Processing in Earth Observation, Technical University of Munich (TUM) and German Aerospace Center (DLR), Oberpfaffenhofen, Germany, the Head of the Team Signal Analysis, DLR, and the Head of the Helmholtz Young Investigator Group (SiPEO), DLR and TUM. Her research interests include remote sensing and earth observation, signal processing, machine learning and data science, with a special application focus on global urban mapping.

Dr. Zhu is a member of Young Academy (Junge Akademie/Junges Kolleg) at the Berlin-Brandenburg Academy of Sciences and Humanities, the German National Academy of Sciences Leopoldina, and the Bavarian Academy of Sciences and Humanities. She is an Associate Editor of the IEEE TRANSACTIONS ON GEOSCIENCE AND REMOTE SENSING.



Luisa Verdoliva (M'14–SM'17) is currently an Assistant Professor of telecommunications with the Department of Electrical Engineering and Information Technology, University of Naples Federico II, Naples, Italy. Her research interests include image processing, in particular restoration, fusion, and classification of remote-sensing images, both optical and synthetic aperture radar, and multimedia forensics.

Dr. Verdoliva has been an Elected Member of the IEEE Information Forensics and Security Technical Committee for the 2016–2018 period. She is an

Associate Editor of the IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY.



Giovanni Poggi is currently a Full Professor of telecommunications and the Vice Director of the Department of Electrical Engineering and Information Technology, University of Naples Federico II, Naples, Italy. His research interests include statistical image processing, including compression, restoration, segmentation, and classification, with application to remote-sensing (both optical and synthetic aperture radar images) digital forensics, and biometry.

Dr. Poggi is an Associate Editor of *Remote Sensing* (MDPI). He has been an Associate Editor of the IEEE TRANSACTIONS ON IMAGE PROCESSING and *Signal Processing* (Elsevier).