



RESEARCH ARTICLE

10.1002/2015JB012752

InSAR processing for volcano monitoring and other near-real time applications

Karsten Spaans¹ and Andrew Hooper¹¹COMET, School of Earth and Environment, University of Leeds, Leeds, UK

Key Points:

- The algorithm we present here provides the rapid and flexible processing necessary to use InSAR for monitoring purposes
- Taking variation of scattering mechanism into account improves estimation of coherence
- Tailored point selection for each interferogram improves coverage compared to other time series methods

Correspondence to:

K. Spaans,
eekhs@leeds.ac.uk

Citation:

Spaans, K., and A. Hooper (2016), InSAR processing for volcano monitoring and other near-real time applications, *J. Geophys. Res. Solid Earth*, 121, 2947–2960, doi:10.1002/2015JB012752.

Received 16 DEC 2015

Accepted 12 MAR 2016

Accepted article online 21 MAR 2016

Published online 8 APR 2016

Abstract Radar interferometry (InSAR, interferometric synthetic aperture radar) is routinely used to measure surface deformation prior to, during, and after volcanic events, although not in a monitoring capacity. The improved data availability of some current satellite missions presents us with the opportunity to do just that. We present here a fast and flexible algorithm to estimate coherence and select points on an interferogram-by-interferogram basis, which overcomes limitations of the conventional boxcar ensemble method in areas of marginal coherence. Time series methods, which offer an alternative way to select coherent points, are typically slow, and do not allow for insertion of new data without reprocessing the entire data set. Our new algorithm calculates the coherence for each point based on an ensemble of points with similar amplitude behavior throughout the data set. The points that behave similarly are selected prior to new images being acquired, on the assumption that the behavior of these nearby points does not change rapidly through time. The resulting coherence estimate is superior in resolution and noise level to the boxcar method. In contrast to most other time series methods, we select a different set of coherent points for each interferogram, avoiding the selection compromise inherent to other time series methods. The relative simplicity of this strategy compared to other time series techniques means we can process new images in about 1 h for a typical setup.

1. Introduction

Interferometric synthetic aperture radar (InSAR) uses radar signal interferometry to obtain high-resolution surface deformation measurements with millimeter to centimeter level accuracy, covering areas of hundreds to thousands of square kilometers per interferogram [Bamler and Hartl, 1998; Hanssen, 2001]. Since the first demonstrations of the technique using satellites in the 1980s [e.g., Gabriel et al., 1989; Li and Goldstein, 1990], InSAR has evolved into one of the main geodetic techniques for monitoring surface deformation. Its dense spatial sampling compliments other techniques like GPS and leveling, where the requirement of equipment and manpower in the field limits spatial sampling.

The basic InSAR technique interferes two SAR images to obtain an interferogram, which gives the difference in phase between the two images. Although it contains several nuisance terms, the phase difference largely reflects the surface deformation in the radar line of sight during the period between the two acquisition dates. To form the interferograms, the images have to be in the same geometry. This is achieved by coregistering the slave image to the master and resampling it [Hanssen, 2001]. Each pixel of the master is then multiplied by the complex conjugate of the corresponding pixel in the slave image to obtain the interferogram. Due to the difference in position of the satellite at the times of acquisition, a phase component is introduced. As this geometric component is a function of the known baseline (the difference in position between the satellites) and the topographic height, a digital elevation model (DEM) can be used to remove this phase component from the interferometric phase [Bamler and Hartl, 1998].

One of the main drawbacks of InSAR is that it relies on the scattering properties of the surface, remaining consistent. The signal scattered back toward the satellite by every resolution element is the result of the coherent summation of many individual scatterers. If the scattering properties of these individual scatterers changes, or if the viewing geometry between acquisitions changes, the coherent summation changes as well. This introduces noise into the signal, known as decorrelation [Zebker and Villasenor, 1992].

Coherence is a measure of the amount of correlation and has magnitude between 0 (no correlation) and 1 (full correlation). The coherence of each point in a single interferogram can be estimated from the phase and

©2016. The Authors.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.



Figure 1. Comparison between three different full resolution coherence estimates of the same interferogram (20090618–20090629): (a) A 5×5 window boxcar, (b) 11×11 window boxcar, and (c) sibling-based (25 to 100 siblings per point) 41×41 search window). The boxcar coherence estimates finds many false high coherence points, such as in the field outlined in red.

amplitude statistics of an ensemble of surrounding pixels [Touzi *et al.*, 1999]. The standard approach is to use a two-dimensional boxcar to define the ensemble of pixels. The boxcar method has three main drawbacks. First, it suffers from a resolution problem. Points within the ensemble have different scattering characteristics, and those points with high amplitude (e.g., buildings or natural ridges) dominate the coherence estimation. Neighboring points use nearly the same mix of points within the ensemble, leading to a smearing out of these features. Second, it tends to overestimate the coherence of a large amount of fully decorrelated points because of the random signal being similar for neighboring points by chance. As neighboring points have very similar ensembles, these features also smear out, exacerbating the issue (Figures 1a and 1b). Third, as the method essentially measures the variability of phase within the window, any nonconstant signal biases the coherence estimation, with high-phase gradients leading to low coherence estimates.

Time series analysis techniques were developed partly as a way to deal with the problem of decorrelation. Two broad categories have been developed: the persistent scatterer (PS) and small baselines (SB) methods. The PS techniques focus on pixels which are dominated by a single strong scatterer, which are less sensitive to decorrelation [Ferretti *et al.*, 2001; Hooper *et al.*, 2007]. The SB methods focus on forming interferometric pairs with small perpendicular and temporal (i.e., separation in time of the two images) baselines [Berardino *et al.*, 2002; Hooper, 2008]. They are therefore better able to extract pixels containing many scatterers, known as distributed scatterers, that might decorrelate in longer baseline combinations. Both types of method select a set of points that are deemed coherent throughout all the interferograms used in the time series. By selecting the same set of points in every interferogram, the time domain can be used to assist in unwrapping the phase [Hooper, 2009]. Both the PS and SB methods suffer from long processing times when run at full resolution, due to their complex analysis methods. Also, the fact that they select the same set of points in each interferogram leads to an averaging effect. Coherent points in high-coherence interferograms may not be selected because the same points are decorrelated in too many of the low-coherence interferograms in the data set. Vice versa, decorrelated points may be selected because of good coherence in the majority of the interferograms [Hooper *et al.*, 2011]. This lowers the overall signal-to-noise ratio in the results significantly and inevitably leads to a loss of information. This selection compromise was partly addressed by identifying “semi-PS,” “temporary” PS, or “partial” PS points [Basilico *et al.*, 2004; Hooper *et al.*, 2011], with success heavily dependent on the area [Ferretti *et al.*, 2011].

A different type of time series approach known as SqueeSAR [Ferretti *et al.*, 2011] takes advantage of both PS and distributed scatterer pixels to maximize the amount of information extracted from the interferograms. To achieve this, pixels within a neighborhood that behave in a similar way are identified. In the SqueeSAR method, these pixels are referred to as statistically homogeneous pixels (SHPs). In this text, we will refer to them as siblings, representing their relationship to each other in terms of scattering mechanism. The SqueeSAR method identifies siblings by applying the two-sample Kolmogorov-Smirnov test [Stephens, 1970] on the amplitude vectors (through time) of the current pixel and all neighbors (in turn) contained in a window around it, assuming a level of significance. After identifying the siblings for a given pixel, the method is able to estimate the sample covariance matrix [Ferretti *et al.*, 2011]. For distributed scatterers, which are identified as having a number of siblings greater than a certain threshold, the covariance matrix is used to invert for the

maximum likelihood (ML) phase value of each cluster of siblings in each of the original interferogram. This estimation process also yields a “goodness of fit” measure, akin to a coherence estimate. For a sibling cluster of distributed scatterer pixels, with an estimated coherence above a certain threshold, a single point with the estimated ML phase value replaces the original points, essentially filtering the original interferogram, taking into account different scattering mechanisms. The resulting filtered interferograms are then processed using a conventional PS algorithm, identifying points which stay coherent throughout the time series. The ML phase estimation is nonlinear and therefore requires iterative methods, greatly increasing processing time. In rural areas, the method can, however, significantly increase the amount of information extracted. Similar methods have since appeared, like the CAESAR approach [Fornaro *et al.*, 2015] or the NL-InSAR method [Deledalle *et al.*, 2011], which uses nonlocal ensembles estimated using single interferograms. The idea of SHPs or siblings has been around for much longer, first appearing as part of the sigma filter as a means of despeckling SAR images [Lee, 1983]. Many approaches have been proposed to identify siblings, for example, in Nicolas *et al.* [2001] using a region growing approach, in Parizzi and Brcic [2011] using different statistical tests, and in Vasile *et al.* [2006] exploiting polarimetry information.

Surface deformations in and around volcanoes are an example where InSAR measurements have proven valuable [Amelung *et al.*, 2000; Sigmundsson *et al.*, 2010a]. InSAR typically has a superior spatial measurement density compared to GPS, making it very useful in constraining the source parameters of surface deformation due to magma movements. However, InSAR data have been sparse in time, and to analyze the data often takes days to weeks to complete. Current SAR satellite missions combine a much shorter repeat time with increased data acquisition capabilities, to yield a greatly improved measurement frequency over most areas compared to previous SAR missions. If all potential acquisitions are realized, an average of approximately two images a day is now possible. This high measurement frequency potentially allows us to use InSAR as a near-real time monitoring tool, as is the case for ground-based techniques such as Global Navigation Satellite Systems. But to achieve this, a rethink of our processing strategies is required, shifting away from slow time series techniques toward more flexible, faster processing. This is also facilitated by the shorter repeat time and higher data acquisition volume, which yield far superior overall coherence of the interferograms, reducing the need for complex time series techniques.

Here we present a new InSAR processing algorithm, aimed at fast ingestion of new images, while extracting the maximum amount of information. The algorithm, which we refer to as Rapid Time Series InSAR (RapidSAR), uses sibling information not for preprocessing prior to time series processing like the SqueeSAR method, but to quickly estimate the coherence for each pixel in newly formed interferograms. This avoids many of the problems of boxcar coherence estimation, while retaining much of its speed and flexibility. Our method yields an individual coherence estimate for each interferogram, in contrast to most other time series methods. This allows us to avoid the selection compromise of PS and SB methods and maximize the amount of information extracted. We demonstrate the effectiveness of RapidSAR on data sets covering the Eyjafjallajökull and Bárðarbunga volcanoes, Iceland.

2. Processing Strategy

To allow the fast processing of new images required for effective volcano monitoring, we have split our processing strategy in two. We start from a set of full resolution interferograms coregistered to a common master and use this initial data set to identify for each pixel a set of siblings. This information is stored for use when a new image comes in. A more detailed overview of how siblings are selected can be found in section 2.1.

The second part of the processing strategy occurs when new images are acquired. Upon arrival of a new SAR image, we form a number of interferometric combinations using this image. We then use the stored sibling information to estimate the coherence for the new interferograms. This coherence estimate is completely independent for every interferogram, allowing for great flexibility in parallel processing. Although we assume the sibling information stays valid for several months, we reestimate the siblings based on the most recent data set after a certain amount of new images has been acquired, to ensure the sibling information stays up to date. Section 2.2 describes how we estimate coherence and in what way we use this coherence to select coherent points in each interferogram.

2.1. Sibling Identification

From the initial set of single master interferograms, we can form all possible interferometric combinations. We use this full set of interferograms to estimate a set of siblings for every pixel. For each pixel, we first calculate the mean amplitude over all interferometric combinations:

$$A_{\text{mean},j} = 1/N \sum_{i=1}^N A_{i,j}, \quad (1)$$

where A is the interferometric amplitude, j represents an arbitrary pixel, and N is the number of interferograms. We then compare the mean amplitude of the pixel to the mean amplitude of each pixel in a surrounding window. Points with a mean amplitude within a percentage threshold of the current pixel are accepted as siblings. Both the window size and the mean amplitude percentage threshold are set by the user, with typical values ranging from 20 to 40 pixels for the window size and 5–15% for the threshold.

We also consider the mean amplitude difference between the master and slave over all interferometric combinations. For each pixel, we calculate the mean amplitude difference as

$$\Delta A_{\text{mean},j} = 1/N \sum_{i=1}^N (A_{i,j}^M - A_{i,j}^S), \quad (2)$$

where $\Delta A_{\text{mean},j}$ is the mean amplitude difference for pixel j , $A_{i,j}^M$ represents the master amplitude of interferogram i for pixel j , and $A_{i,j}^S$ the slave amplitude. For each pixel, we ensure that the selected siblings are also within a percentage threshold of the mean amplitude difference, dropping those for which this is not the case. This helps to ensure that points belong to similar scattering surface types. Typical values for the amplitude difference threshold are 10–30%. Contrary to many other techniques that use siblings, we do not enforce connectedness of siblings, as our application does not require it and a wide spread of points can actually improve coherence estimation reliability in many cases.

Also, contrary to many other techniques, we do not use statistical hypothesis testing on the amplitude vectors of pairs of pixels to select our siblings, opting instead to use the mean amplitude and the mean amplitude difference. We made this choice for several reasons. First, calculating the mean and the mean amplitude difference is less complex than performing the hypothesis testing, therefore taking less time. Equally critical is that during sibling identification, only the mean amplitude and mean amplitude difference have to be kept in memory, as opposed to the full amplitude stack for all combinations, vastly reducing RAM memory and disk read requirements. Finally, the hypothesis tests tend to result in a binary result, either the hypothesis is true or false. We require a minimum number of siblings in the ensemble to obtain a reasonably unbiased estimate for the coherence (see section 2.2) and thus enforce a minimum amount of siblings (typically between 10 and 30). If a pixel does not have sufficient siblings within the threshold, we can add the best sibling candidates outside of the threshold. The same is true if we impose a maximum on the amount of siblings, where we can reject the worst siblings that fall within the threshold.

2.2. Coherence Estimation and Point Selection

When a new image is acquired and coregistered with respect to the common master, we already have for every pixel a preidentified ensemble of sibling pixels. Similar to the “boxcar” method described above, we can now use this sibling ensemble to estimate the coherence:

$$\hat{\gamma} = \frac{\left| \frac{1}{n} \sum_{i=1}^n M_i \cdot \bar{S}_i \right|}{\sqrt{\frac{1}{n} \left(\sum_{i=1}^n M_i \cdot \bar{M}_i \sum_{i=1}^n S_i \cdot \bar{S}_i \right)}}, \quad (3)$$

where n represents the number of points in the ensemble, M represents the master image signal for an arbitrary point, and S represents the slave image signal. The overline indicates the complex conjugate. As siblings are points with similar scattering characteristics, this method avoids or largely mitigates the problems that the boxcar method has; high-coherence targets do not smear out, as they are not part of the sibling ensemble of nearby points with lower coherence, and the chance of erroneously high coherence estimates is reduced, as larger ensembles can be used.

We can use the estimated coherence to select points from the full resolution interferograms directly, and in some cases, this might be preferable. However, if high resolution is not required, it is preferable to select points




Figure 2. Comparison between points selected in (a) full resolution and (b) after multilooking. The variance threshold was set, independently for each interferogram, based on the 99th percentile variance in the ocean (top left corner). The interferogram shown (20090618–20090629) covers the full test area. The area shown in Figures 1 and 4 is indicated by the black box, and the grey box indicates the area shown in Figure 3.

from multilooked versions of the interferograms. Using the sibling-based ensemble to estimate the coherence reduces the spatial correlation between coherence estimates, but it does not completely remove the issues with erroneously high coherence estimates in incoherent areas. However, as these erroneously high estimates no longer smear out, multilooking reduces them to low values, even for small multilook factors.

To select points in the multilooked image, we need to obtain a measure of the quality of each multilooked point. Averaging the coherence is not appropriate; e.g., the coherence of a multilooked point consisting of 100 points with coherence 0.6 is significantly higher than the mean of 0.6. Instead, we use the Cramer-Rao relationship to calculate a variance for each point based on its coherence [Hanssen, 2001]:

$$\sigma_{\text{CR}}^2 = \frac{1 - |\gamma|^2}{2|\gamma|^2} \quad (4)$$

This allows us to weight the interferometric signal by the inverse of the variance before multilooking, as well as calculate a variance for the resulting multilooked point by propagating the uncertainty, which we can then use to select points.

The multilooking is especially effective in creating contrast between incoherent areas and areas with low coherence. For incoherent areas, the occasional, chancy low variance (high coherence) estimate is averaged with the high-variance estimates surrounding it, resulting in an overall high variance. For low-coherence areas, the coherent, low-variance points within a multilook window lower the overall variance within it. This effect is enough to differentiate between incoherent and low-coherence areas (Figure 2). As hinted at above, this averaging works for the sibling-based coherence, since neighboring points do not necessarily use a similar set of points for the coherence estimate. Averaging of the coherence would not work with the boxcar method, as the ensembles used to calculate the coherence for neighboring points share 80–90% of their pixels, thus yielding highly spatially correlated coherence estimates.

After selecting the points with sufficient coherence, we filter them using a Goldstein filter [Goldstein and Werner, 1998]. To allow unwrapping, we use region growing to fill empty points in the grid, similar to the method used in StaMPS [Hooper *et al.*, 2007]. The high-coherence interferograms allow us to use two-dimensional, spatial unwrapping with high success rates. We use the public domain software SNAPHU [Chen and Zebker, 2001], which uses a minimum cost flow approach to solve for the phase ambiguities. The algorithm described above allows us to estimate the coherence for every interferogram individually, providing us with much greater flexibility in terms of parallel processing. We implemented the processing strategy to take full advantage of parallel processing, greatly improving processing efficiency and time.

2.3. Optional Processing Steps

Besides being able to take advantage of parallel processing, another advantage of the ability to process images individually is the flexibility in applying extensions dealing with specific challenges present in individual images. One of the potential pitfalls of using equation (3) is that it assumes no systematic variability in the phase of all the points within the ensemble used to calculate the coherence. However, in the case of




Figure 3. Demonstration of the effect of temporarily changing siblings and the optional sibling rejection proposed to resolve it. (a) The wrapped interferometric phase of interferogram 20090618–20090629. (b) The estimated sibling-based coherence without sibling rejection. (c) The estimated sibling-based coherence, using the same siblings as in Figure 3b but with a third of the siblings rejected. The area covered is a zoom of the riverbed visible in the bottom right side of Figure 1.

high deformation rates, or high-frequency nuisance signals, this is clearly not the case and results in an underestimation of the coherence for these kind of areas. Our method uses points distributed over a larger area than the boxcar method, exacerbating the problem. For images with high-frequency systematic phase, we estimate the spatially correlated phase component of the interferogram by filtering and removing this phase component from the interferometric phase. We then use the residual, “flattened” phase to estimate the coherence. The filtering is done either using multilooking or using a combined low-pass adaptive filtering [Hooper, 2008]. Although removing the spatially correlated phase from interferograms eliminates the underestimation of coherence in areas of high (coherent) phase variability, it comes at the cost of increased processing time and an overestimation of the coherence in incoherent areas (see section 3.4). We, therefore, choose to make the filtering of the phase optional, allowing it to be run on specific images only.

Another optional processing step deals with variability in the siblings with time. An example of this can be seen in Figures 3a and 3b, where an incoherent river can be seen on the right side of the scene in the wrapped phase, but does not show up in the sibling-based coherence estimate. The river is a glacial outlet river that tends to shift due to deposit and flow rate changes. This means that siblings in the otherwise very coherent riverbed will not belong to the same type of scatterer for certain images. In the case of the river, a point in the river will have a few siblings within the incoherent river but also many siblings on the coherent riverbed. This will result in a high-coherence estimate, even though the point itself is clearly incoherent. The random phase within these incoherent areas may cause unwrapping errors, which makes it important to solve this issue. We deal with it by (optionally) evaluating the validity of the sibling information. We commence by estimating the boxcar coherence for the interferogram in question using a small window (typically 5×5 pixels). We then compare the boxcar coherence value for every point to all its siblings and reject the one third of siblings with the largest difference in coherence compared to the current point. We only use the remaining siblings to estimate the coherence for the point in question. This naturally comes at the cost of additional processing time and potentially the requirement of raising the number of siblings per point to be estimated in the first place.

3. Eyjafjallajökull Case Study

Around the time of the 2010 Eyjafjallajökull eruption [Sigmundsson *et al.*, 2010a], a large set of TerraSAR-X SAR images were obtained covering the volcano. This data set represents one of the first demonstrations of the possibilities offered by the high acquisition rates of current satellites. The high data acquisition frequency and the resulting high coherence form a good example of the expected data quality in the years to come.

3.1. Coherence Estimate and Multilooking

A comparison between the boxcar and the sibling coherence estimates for a region covering the south flank and coast of Eyjafjallajökull (5000 by 5000 pixels) is given in Figure 1. The advantages of the sibling coherence estimate over the boxcar method are clearly visible. As described in section 1, the boxcar method tends to smear out coherence, leading to smaller roads becoming invisible in the coherence estimate and high-coherence buildings and wider roads influencing the coherence of the surrounding pixels. This effect

Table 1. Estimated Mean Phase Variance as a Proxy for the Quality of the Three Coherence Estimates Shown in Figure 1, as Well as a Boxcar Coherence Estimate With a 17 by 17 Window

Method	Mean Phase Variance (rad ²)
Boxcar 5 × 5	1.725
Boxcar 11 × 11	1.517
Boxcar 17 × 17	1.550
RapidSAR	1.308

noisy pixels having similar phase by chance or by one pixel in the ensemble having a high amplitude compared to the remaining pixels and dominating the coherence estimation. It is, therefore, more prevalent when a smaller window is used, as the coherence is calculated using fewer points, increasing the chance of a biased coherence estimate due to either of the two aforementioned reasons. However, even the very large 11 × 11 window is not sufficient to overcome this problem. The sibling coherence estimate, on the other hand, appears visually to perform much better (Figure 1). To evaluate quantitatively whether our new approach is actually estimating coherence more reliably, we use as a proxy the phase variance of points with estimated coherence above a certain threshold. For each of the three coherence estimates shown in Figure 1, we select all points with a coherence over 0.5. For each point, we calculate the variance of the phase values of selected points in a 21 by 21 window surrounding the point. To ensure a reasonable variance estimate, we only take into account points with at least 10 selected points in the search window. We also calculated the phase variance for a 17 by 17 window boxcar estimate. Table 1 shows the mean phase variance for the four coherence estimates. The lower mean phase variance demonstrates that the RapidSAR coherence estimate is superior to the three boxcar estimates.

One of the main reasons why the sibling coherence estimate suffers less from the two problems described becomes apparent when looking at the number of siblings selected for each pixel (Figure 4). It shows that objects with stable scattering properties, especially those surrounded by areas of lower overall coherence, tend to have fewer siblings. These pixels use only the few pixels that contain the same or similar objects (e.g., buildings, roads, and cliff faces), not allowing points with different scattering mechanism surrounding these objects to influence the coherence and smear it out. Vice versa, good scatterers do not raise the coherence of surrounding points as well. Furthermore, in fields and other low-coherence areas, pixels tend to have a high number of siblings, often the maximum number allowed. This, combined with the fact that neighboring points can have very different siblings, drastically reduces the noise in the coherence estimate seen in the boxcar coherence.

The coherence estimate obtained using the sibling information is converted to a variance using equation (4) and subsequently used to weight the pixels during multilooking. The benefit of weighted multilooking




Figure 4. Overview of the number of siblings for every pixel. A minimum of 25 siblings per pixel was imposed to ensure a reasonably unbiased estimate for the coherence and a maximum of 100 for efficiency. The area shown covers the same area as in Figure 1.




Figure 5. Comparison between normal and weighted multilooking for a very low-coherence interferogram due to snow cover. The interferogram (20100204–20100331) is different from that shown in previous figures but covers the same area as shown in Figure 1 and is part of the same data set used to identify siblings. The (a) standard and (b) weighted multilooked phase. A multilooking window of 5×5 was used. Also shown are the filtered phases of (c) Figure 5a and (d) Figure 5b. Although the difference between Figures 5a and 5b are subtle, the difference between Figures 5c and 5d clearly shows the value of using weighted multilooking on low-coherence images. The black boxes highlight some areas of improvement.

compared to normal multilooking is most apparent in low-coherence areas. Figure 5 shows a comparison of the effect of normal and weighted multilooking on an interferogram affected by snow cover. Especially after filtering, it becomes clear that the amount of signal retrieved is much higher for weighted multilooking, greatly aiding in the unwrapping of these interferograms.

3.2. Point Selection and Unwrapping

Although we estimate the coherence of every point at full resolution, we typically select points on a (slightly) multilooked version of the interferogram, as described in section 2.2. The reason for this becomes clear when looking at Figure 2. The selection method after multilooking is indeed better able to distinguish between low-coherence fields and incoherent areas like the ocean (top left corner).

To evaluate the performance of our coherence estimation and point selection routine, we compare in Figure 6 points selected using the RapidSAR method (middle column) to an SB time series method (right column) for three different interferometric combinations. For reference, we also show the full, nonmultilooked wrapped phase values. For the RapidSAR coherence estimate, each pixel has between 25 and 100 siblings selected, using all possible combinations of 17 images between July 2009 and March 2010. A search window of 40 pixels was used to find siblings, and the amplitude percentage threshold was set to 10%. The small baseline processing was achieved using the small baseline module in StaMPS [Hooper, 2008], which uses phase stability through time to estimate a temporal coherence measure. A selection threshold of 2% random points was set, and only interferograms with high expected coherence based on the temporal and perpendicular baselines were used in the network [Hooper, 2008].

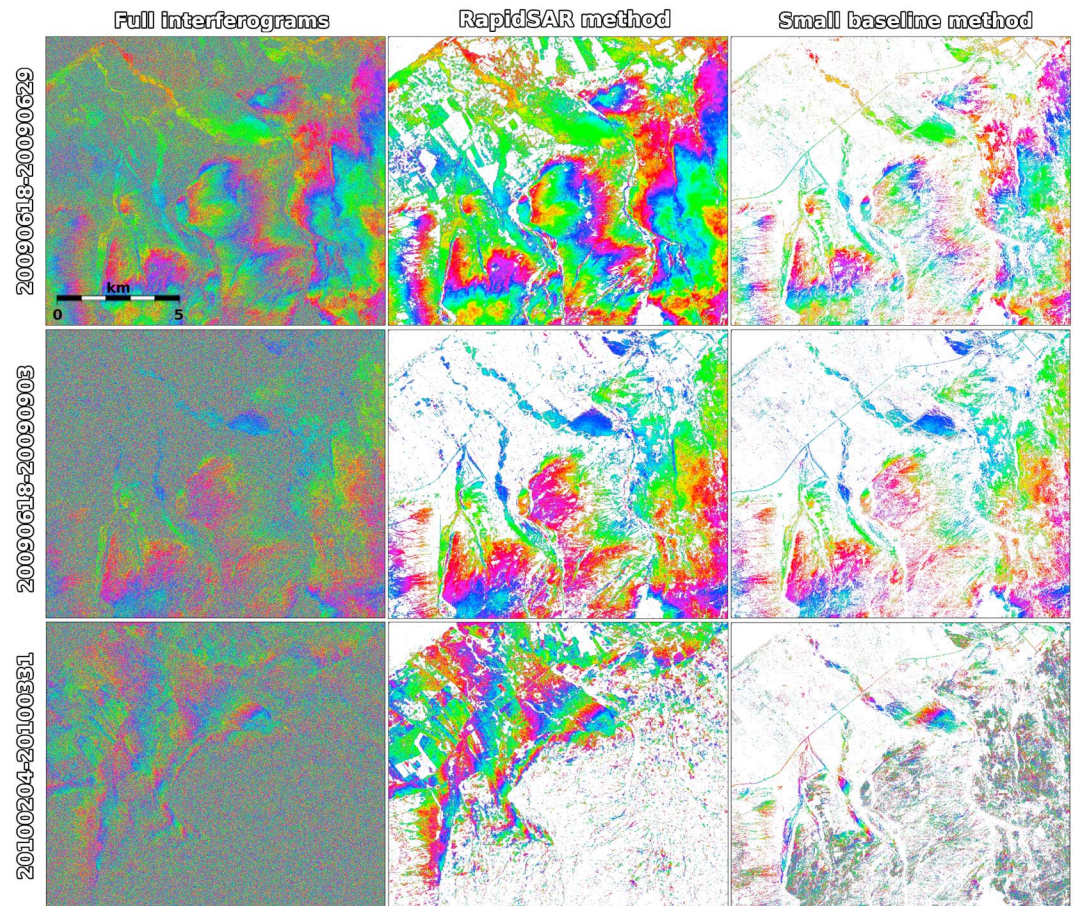


Figure 6. Comparison of selected points using the (middle column) RapidSAR method and (right column) small baseline processing. Both selections are plotted in the same resolution. (left column) For comparison, the full, nonmultilooked interferograms are shown. Three different interferometric combinations are shown: A short temporal baseline, highly coherent summer acquisition (20090618–20090629, top row), a combination with a longer, 3 month baseline (20090618–20090903, middle row) and a short temporal baseline, winter acquisition (20100204–20100331, bottom row). Area shown is the same as that in Figure 2.

The three combinations shown are chosen for their different levels of coherence and their different decorrelation mechanisms. Figure 6 (top row) shows a high coherence, small baseline combination, as is typically expected for current short repeat cycle satellites like TerraSAR-X, Cosmo-SkyMed, and Sentinel 1. The RapidSAR method clearly selects more points compared to the SB method. The main reason for this is that the SB method selects a single set of points for all interferograms, only selecting points that stay consistently coherent. This leads to a significant loss in information that is extracted from the interferograms. The reason that the SB method selects less points in the top left half of the scene becomes apparent when looking at the longer temporal baseline combination shown in Figure 6 (middle row). The top left area of the scene consists mainly of agricultural fields, which decorrelate rapidly over time. Due to combinations like this one, with images separated in time by 3 months, the SB method selects fewer points in all interferogram for this area. The RapidSAR method selects few points in the area for the long temporal baseline combinations only, without affecting other combinations.

Figure 6 (bottom row) shows another problem caused by the limitation of selecting one set of points. Inherently, this selection becomes a trade-off between retaining signal in high-quality interferograms and reducing noise in lower quality interferograms [Hooper *et al.*, 2011]. The scene in Figure 6 (bottom row) is a winter acquisition, heavily decorrelated due to snow cover on the higher altitude area in the bottom right half of the scene.

Table 2. Processing Time to Calculate the Coherence of 10 Interferometric Combinations of a 5000 by 5000 TerraSAR-X Scene, for Varying Number of Processing Cores^a

Number of CPUs	Processing Time (min)
2	84.6
4	58.4
8	45.9
16	36.5

^aApproximately 8 Gb of RAM memory was used during all processing runs.

As there are few winter acquisitions in the time series, the small baseline method selects many points in the affected area, the majority of which contain no signal. Our method selects few points in the area, most of which are on rocky outcrops. Figure 6 clearly shows the advantage of individual point selection for each interferogram and demonstrates the effectiveness of the RapidSAR method for revisit times typical in current satellites.

The processing time necessary to process the scene in Figure 6 depends heavily on the computational facilities used. One of the most critical variables for performance is the amount of processor cores. Table 2 gives a representative example of processing times of the coherence estimate for the same set of 5000 by 5000 pixel interferograms. A clear diminishing return can be seen on the amount of processing cores used. Partly, this is due to overhead for the parallel computations, but mostly this is caused by certain operations not allowing parallelization. The processing times given here are meant as an indication of expected processing times and could be improved in several ways (e.g., increased internal memory or optimized implementation).

The high density of coherent selected points allows for effective filtering and unwrapping of the phase. Figure 7 shows the filtered and unwrapped phase of the points selected using our method for the high-quality, short baseline acquisition of Figure 6. The smooth, filtered phase is unwrapped effectively, even given the complex nature of the wrapped phase pattern and the many discontinuities. Figure 8 shows an incremental time series processed using the RapidSAR method.

3.3. Optional Sibling Rejection

As described in section 2.3, our method is less effective where sibling relationships temporarily change. This problem mostly occurs when a thin band of points decorrelate with respect to their surrounding siblings, such as along a river with highly varying water levels. Figure 3c shows the result of the optional processing step designed to deal with this issue. We rejected one third of the siblings for every point, based on the 5×5 boxcar coherence estimate for every point. For temporarily decorrelated points in the river, this resulted in many of the siblings that were in the coherent riverbed being rejected, greatly reducing the coherence estimated for those points. The coherence estimate in the low-coherence fields also changes slightly. It is difficult to evaluate if this slight change in the fields is an improvement. However, due to the nature in which we select points, after multilooking, this difference largely disappears.




Figure 7. Demonstration of the phase unwrapping for interferogram 20090618–20090629 (same as Figure 6, top row). (a) Filtered phase values. (b) Unwrapped phase values.

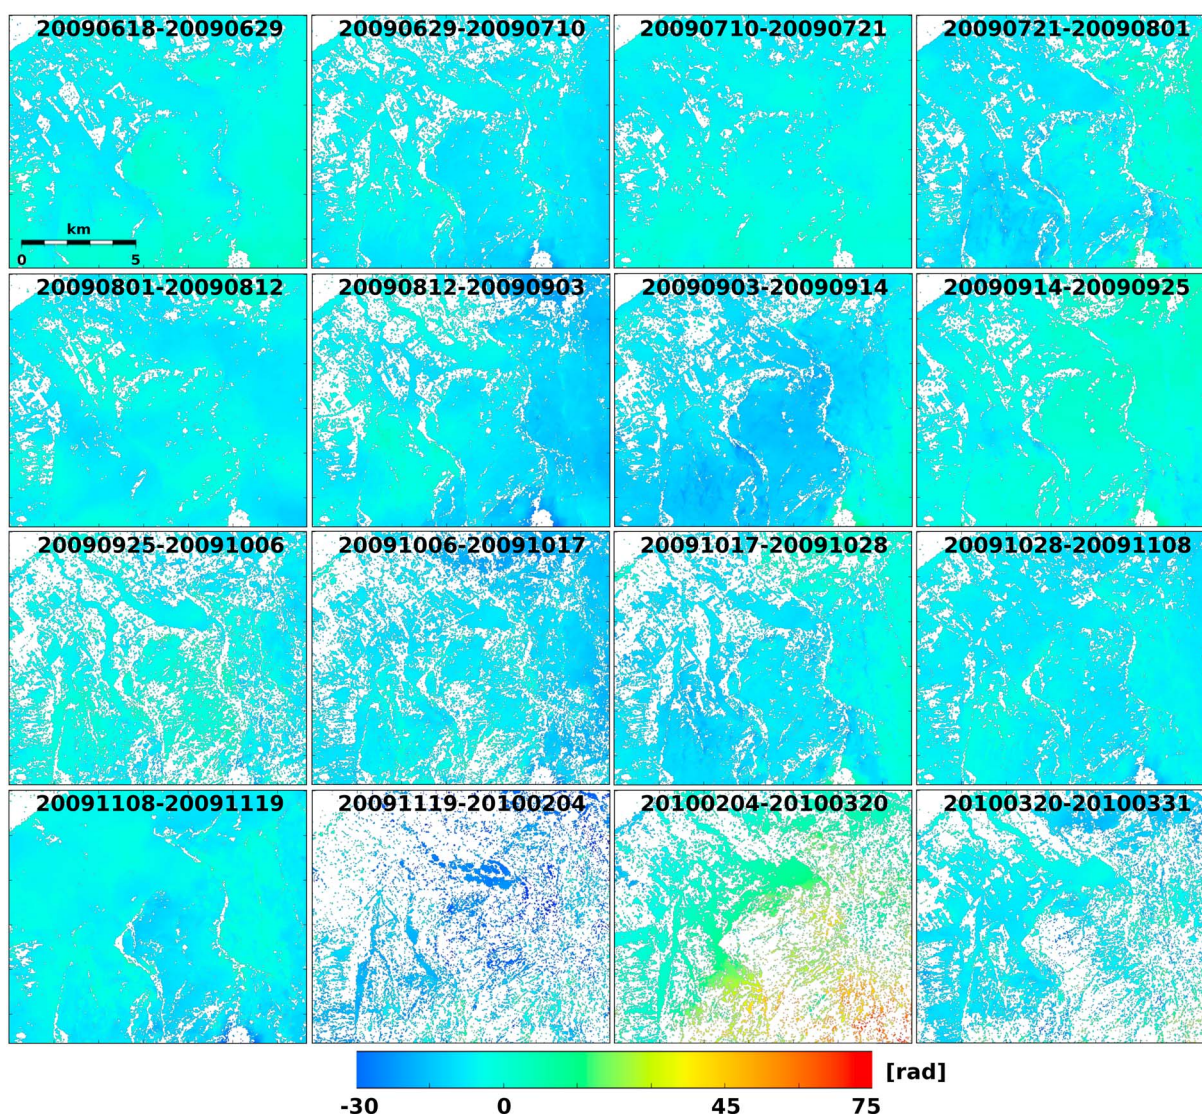


Figure 8. Incremental time series of unwrapped phase values of selected points for the same scene as that covered by Figure 2. Unwrapped phase values are referenced to an area on the coast (top left corner). The Eyjafjallajökull volcano is to the bottom right of the scene. Deformation associated with the first eruptive episode of the 2010 Eyjafjallajökull eruption, which commenced on 20 March 2010 [Sigmundsson *et al.*, 2010b], is clearly visible in the interval 4 February 2010 and 20 March 2010, even though it is one of the lowest coherence pairs in the time series.

3.4. Optional Phase Flattening

Figure 9 shows an interferogram (19,500 by 19,000 pixels) covering the initial weeks of 2014–2015 Bárðarbunga-Holuhraun eruption [Sigmundsson *et al.*, 2015]. In the week leading up to and the weeks after the eruption, relative deformation of over 2 m was measured over less than 20 km [Sigmundsson *et al.*, 2015]. Such high-deformation gradients lead to a high variability in the interferometric phase, as is evident in the interferogram shown in Figure 9. As discussed before, the coherence estimation using equation (3) is essentially a measure of the phase variability of the points in the ensemble. Thus, a high systematic phase gradient will lead to an erroneously low coherence estimate. Figures 9b and 9d show the effectiveness of removing the spatially correlated phase from the interferogram before estimating coherence. This makes it clear that in interferograms with high fringe rates, it is essential to remove the spatially correlated phase before estimating coherence. However, filtering the phase also removes part of the phase variability in completely decorrelated areas, slightly raising the coherence there. This results in the need for a lower variance threshold, leading to a loss of signal.

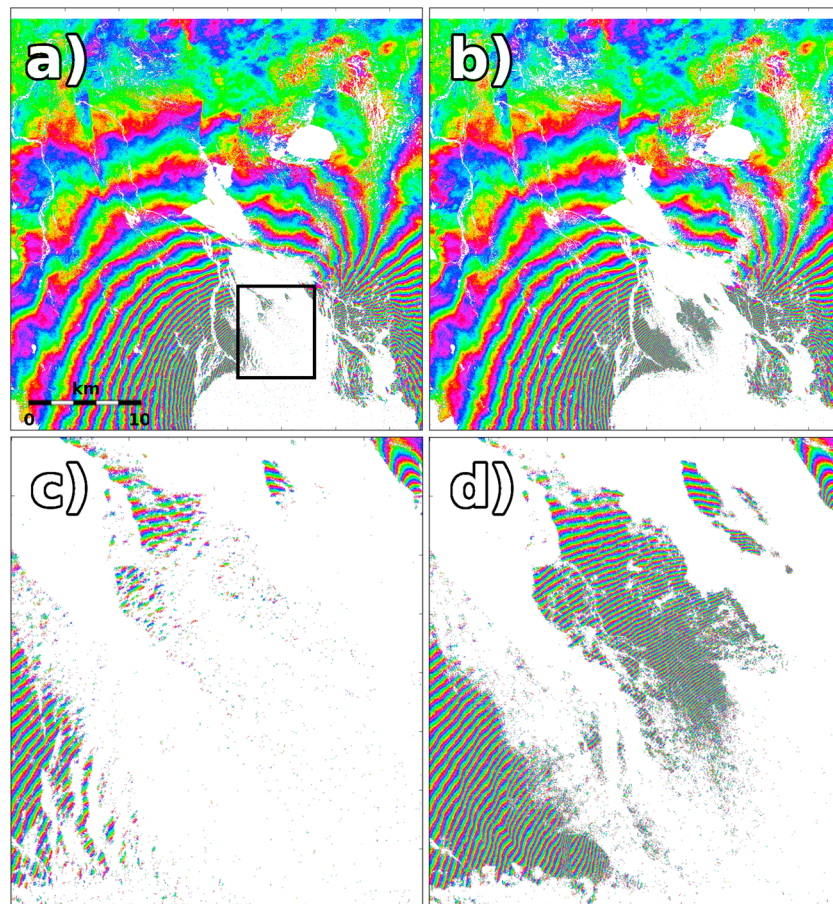


Figure 9. The effect of optionally removing the spatially correlated phase from interferograms before estimating coherence. The interferogram covers the 2014 Bárðarbunga-Holuhraun eruption and was generated using data obtained by the Cosmo-SkyMed satellite constellation. (a) The points selected using coherence estimated with the original phase. (b) Points selected using coherence estimated after removing the filtered phase. A closer look at the fast deforming areas of (c) Figure 9a and (d) Figure 9b as outlined by the black box in Figure 9a.

4. Conclusions and Outlook

We present RapidSAR, a new algorithm that is able to handle the high volumes of data that current generation SAR satellites generate and produce high signal-to-noise ratio deformation maps in a timely fashion. RapidSAR is developed for volcano monitoring, where timely processing of the data is key. Besides the relatively fast processing time, the individual coherence estimate and subsequent point selection allows us to avoid the selection compromise inherent in most other time series methods. The efficiency of the algorithm in extracting coherent points from high-coherence data sets also makes it suitable for other surface deformation applications. The algorithm will be used operationally to monitor volcanic systems in Iceland as part of the FutureVolc project. With the successful operation of Sentinel 1A since late 2014, a freely available, nearly constant stream of SAR data has become available. Sentinel 1A acquisitions cover large areas, albeit at a lower resolution compared to TerraSAR-X and Cosmo-SkyMED. This means that Sentinel 1A data contain fewer points per unit area. Combined with the natural data partitioning resulting from the TOPS mode bursts [De Zan and Guarnieri, 2006], our method is fully scalable to handle Sentinel 1A data.

RapidSAR is well suited to relatively short baseline time series, as obtained over most volcanic systems by current satellite systems. There may however be some situations where methods like the PS and small baseline techniques may perform better, and the techniques should therefore be seen as complimentary. Especially for steady, small-scale signals, PS and small baseline may perform better due to the way in which they select points. In this case, the fact that they select the same points in every image will also aid in unwrapping and estimating nuisance signals, even though it might lead to a loss of information. The loss of information is

mitigated by using methods like SqueeSAR or CAESAR, but these techniques require even more processing power, making them less suited to routine, automated processing of data over numerous areas of interest.

The focus of the RapidSAR methodology is on near-real time monitoring, where the main importance lies in what has happened in the period between the latest two acquisitions. It is here where our method really outperforms other time series methods, which excel at extracting long-term steady state deformations, but tend to struggle with sudden changes in deformation rates and/or patterns. This fact, combined with the point selection compromise in other time series methods, makes RapidSAR well suited to monitoring volcanoes and other deforming processes. Although not currently a focus of the method, the increase in information extracted from the data set should also profit longer-term time series generation. The variable set of points selected in each interferogram will make time series generation more complex, however, and will need to be addressed in the future.

Further improvements to our method are certainly possible. There are other ways in which siblings could be identified, especially with an eye to the current Sentinel 1 mission, which captures dual-polarization data in its standard operation mode. Using polarimetric information could prove very useful in identifying siblings and also in rejecting siblings on an image by image basis (section 2.3). Polarimetry can further be used to find the optimal combination of polarimetric channels, which will yield interferograms with higher coherence [e.g., Navarro-Sanchez and Lopez-Sanchez, 2012]. Furthermore, although the overall good coherence of the short temporal baseline interferograms of current missions greatly aids the phase unwrapping, a more advanced unwrapping algorithm that could incorporate the sparse information in the time domain would improve reliability.

Acknowledgments

K.S. is funded through the FutureVolc project, which received funding through the European Union's Seventh Programme for research, technological development, and demonstration under grant agreement 308377. This work is also supported by the UK Natural Environmental Research Council (NERC) the Centre for the Observation and Modelling of Earthquakes, Volcanoes and Tectonics (COMET). COSMO-SkyMed data were provided by the Italian Space Agency (ASI) and TerraSAR-X data by the German Space Agency (DLR) through the Icelandic Volcanoes Supersite project supported by the Committee on Earth Observing Satellites (CEOS). Information on how to access the data for this paper can be obtained by contacting the corresponding author at eekhs@leeds.ac.uk.

References

- Amelung, F., S. Jónsson, H. Zebker, and P. Segall (2000), Widespread uplift and "trapdoor" faulting on Galápagos volcanoes observed with radar interferometry, *Nature*, *407*, 993–996, doi:10.1038/35039604.
- Bamler, R., and P. Hartl (1998), Synthetic aperture radar interferometry, *Inverse Prob.*, *14*, 1–54.
- Basilico, M., A. Ferretti, F. Novali, C. Prati, and F. Rocca (2004), Advances in permanent scatterers analysis: Semi and temporary PS, paper presented at European Conference on Synthetic Aperture Radar, Ulm, Germany, pp. 349–350, 25–27 May.
- Berardino, P., G. Fornaro, R. Lanari, and E. Sansosti (2002), A new algorithm for surface deformation monitoring based on small baseline differential SAR interferograms, *IEEE Trans. Geosci. Remote Sens.*, *40*(11), 2375–2383, doi:10.1109/TGRS.2002.803792.
- Chen, C., and H. Zebker (2001), Two-dimensional phase unwrapping with use of statistical models for cost functions in nonlinear optimization, *J. Opt. Soc. Am. A*, *18*(2), 338–351, doi:10.1364/JOSAA.18.000338.
- De Zan, F., and A. Guarnieri (2006), TOPSAR: Terrain observation by progressive scans, *IEEE Trans. Geosci. Remote Sens.*, *44*(9), 2352–2360, doi:10.1109/TGRS.2006.873853.
- Deledalle, C., L. Denis, and F. Tupin (2011), NL-InSAR: Non-local interferogram estimation, *IEEE Trans. Geosci. Remote Sens.*, *49*(4), 1441–1452, doi:10.1109/TGRS.2010.2076376.
- Ferretti, A., C. Prati, and F. Rocca (2001), Permanent scatterers in SAR interferometry, *IEEE Trans. Geosci. Remote Sens.*, *39*(1), 8–20, doi:10.1109/36.898661.
- Ferretti, A., A. Fumagalli, F. Novali, C. Prati, F. Rocca, and A. Rucci (2011), A new algorithm for processing interferometric data-stacks: SqueeSAR, *IEEE Trans. Geosci. Remote Sens.*, *49*(9), 3460–3470, doi:10.1109/TGRS.2011.2124465.
- Fornaro, G., S. Verde, and D. Reale (2015), CAESAR: An approach based on covariance matrix decomposition to improve multibaseline-multitemporal interferometric SAR processing, *IEEE Trans. Geosci. Remote Sens.*, *53*(4), 2050–2065, doi:10.1109/TGRS.2014.2352853.
- Gabriel, A., R. Goldstein, and H. Zebker (1989), Mapping small elevation changes over large areas: Digital radar interferometry, *J. Geophys. Res.*, *94*(B7), 9183–9191.
- Goldstein, R., and C. Werner (1998), Radar interferogram filtering for geophysical applications, *Geophys. Res. Lett.*, *25*(21), 4035–4038.
- Hanssen, R. (2001), Radar interferometry: Data interpretation and error analysis, PhD thesis, Delft Univ. of Technol., Netherlands.
- Hooper, A. (2008), A multi-temporal InSAR method incorporating both persistent scatterer and small baseline approaches, *Geophys. Res. Lett.*, *35*, L16302, doi:10.1029/2008GL034654.
- Hooper, A. (2009), A statistical-cost approach to unwrapping the phase of InSAR time series, paper presented at FRINGE 2009 Workshop, ESA SP-677, Frascati, Italy, 30 Nov.–4 Dec.
- Hooper, A., P. Segall, and H. Zebker (2007), Persistent scatterer interferometric synthetic aperture radar for crustal deformation analysis, with application to Volcán Alcedo, Galápagos, *J. Geophys. Res.*, *112*, B07407, doi:10.1029/2006JB004763.
- Hooper, A., B. Ófeigsson, F. Sigmundsson, B. Lund, P. Einarsson, H. Geirsson, and E. Sturkell (2011), Increased capture of magma in the crust promoted by ice-cap retreat in Iceland, *Nat. Geosci.*, *4*, 783–786, doi:10.1038/ngeo1269.
- Lee, J. (1983), Digital noise smoothing and the sigma filter, *Comput. Vision Graphics Image Process.*, *21*, 255–269.
- Li, F., and R. Goldstein (1990), Studies of multibaseline spaceborne interferometric synthetic aperture radars, *IEEE Trans. Geosci. Remote Sens.*, *28*(1), 88–97.
- Navarro-Sanchez, V., and J. Lopez-Sanchez (2012), Improvement of persistent-scatterer interferometry performance by means of a polarimetric optimization, *IEEE Geosci. Remote Sens. Lett.*, *9*(4), 609–613, doi:10.1109/LGRS.2011.2176715.
- Nicolas, J., F. Tupin, and H. Maitre (2001), Smoothing speckled SAR images by using maximum homogeneous region filters: An improved approach, paper presented at IEEE 2001 International Geoscience and Remote Sensing Symposium (IGRSS'01), pp. 1503–1505, Sydney, Australia, 9–13 Jul.
- Parizzi, A., and R. Brcic (2011), Adaptive InSAR stack multilooking exploiting amplitude statistics: A comparison between different techniques and practical results, *IEEE Geosci. Remote Sens. Lett.*, *8*(3), 441–445, doi:10.1109/LGRS.2010.2083631.

- Sigmundsson, F., et al. (2010a), Intrusion triggering of the 2010 Eyjafjallajökull explosive eruption, *Nature*, 468(7322), 426–430, doi:10.1038/nature09558.
- Sigmundsson, F., V. Pinel, B. Lund, F. Albino, C. Pagli, H. Geirsson, and E. Sturkell (2010b), Climate effects on volcanism: Influence on magmatic systems of loading and unloading from ice mass variations, with examples from Iceland, *Philos. Trans. R. Soc. A*, 368(1919), 2519–2534, doi:10.1098/rsta.2010.0042.
- Sigmundsson, F., et al. (2015), Segmented lateral dyke growth in a rifting event at Bárðarbunga volcanic system, Iceland, *Nature*, 517, 191–195, doi:10.1038/nature14111.
- Stephens, M. (1970), Use of the Kolmogorov-Smirnov, Cramér-Von Mises and related statistics without extensive tables, *J. R. Stat. Soc., Ser. B*, 32(1), 115–122.
- Touzi, R., A. Lopes, J. Bruniquel, and P. Vachon (1999), Coherence estimation for SAR imagery, *IEEE Trans. Geosci. Remote Sens.*, 37(1), 135–149.
- Vasile, G., E. Trouvé, J.-S. Lee, and V. Buzoloiu (2006), Intensity-driven adaptive-neighborhood technique for polarimetric and interferometric SAR parameters estimation, *IEEE Trans. Geosci. Remote Sens.*, 44(6), 1609–1620, doi:10.1109/TGRS.2005.864142.
- Zebker, H., and J. Villasenor (1992), Decorrelation in interferometric radar echoes, *IEEE Trans. Geosci. Remote Sens.*, 30(5), 950–959.