

JGR Solid Earth

RESEARCH ARTICLE

10.1029/2019JB017519

Key Points:

- We have developed an algorithm that incorporates spatial independent component analysis to automatically detect signs of volcanic unrest
- We are able to automatically detect the increase in inflation that preceded the 2018 eruption of Sierra Negra
- One component of our algorithm separates geophysical and atmospheric signals, and we release this freely to the scientific community

Supporting Information:

- Supporting Information S1

Correspondence to:

M. E. Gaddes,
eemeg@leeds.ac.uk

Citation:

Gaddes, M. E., Hooper, A., & Bagnardi, M. (2019). Using machine learning to automatically detect volcanic unrest in a time series of interferograms. *Journal of Geophysical Research: Solid Earth*, 124, 12,304–12,322. <https://doi.org/10.1029/2019JB017519>

Received 8 FEB 2019

Accepted 7 SEP 2019

Accepted article online 13 SEP 2019

Published online 28 NOV 2019

Using Machine Learning to Automatically Detect Volcanic Unrest in a Time Series of Interferograms

M. E. Gaddes¹ , A. Hooper¹ , and M. Bagnardi^{1,2,3} 

¹COMET, University of Leeds, Leeds, UK, ²Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA, ³Now at Cryospheric Sciences Laboratory, NASA Goddard Space Flight Center, Greenbelt, MD, USA

Abstract The latest generation of synthetic aperture radar satellites produce measurements of ground deformation at the majority of the world's subaerial active volcanoes and can be used to detect signs of volcanic unrest. We present an automatic detection algorithm that uses these data to automatically warn when deformation at a volcano departs from the background. We demonstrate our approach on synthetic data sets and the unrest leading to the 2018 eruption of Sierra Negra (Galapagos). Our algorithm encompasses spatial independent component analysis and uses a significantly improved version of the ICASO algorithm, which we term ICASAR, to robustly perform spatial independent component analysis. We use ICASAR to isolate signals of geophysical interest from atmospheric signals, before monitoring the evolution of these signals through time in order to detect the onset of a period of volcanic unrest.

1. Introduction

There are ~1,400 volcanoes globally with the potential to erupt, but with only ~100 volcano observatories, many volcanoes remain unmonitored (Loughlin et al., 2015). However, the routine global acquisition of the latest generation of synthetic aperture radar (SAR) satellites (e.g., The European Space Agency's Sentinel-1 constellation), combined with fast formation of interferograms from newly acquired images (e.g., the LiCSAR processor González et al., 2016), produces measurements of ground deformation that can be used to monitor the majority of the world's volcanoes. Though these measurements of ground deformation may be easily interpretable to the human observer, the sheer volume of data required for routine global monitoring would create an onerous task. Therefore, we seek to develop an algorithm that is able to automatically detect signs of volcanic unrest in a time series of interferograms. In contrast to the algorithm of Anantrasirichai et al. (2018) which aims to identify any deformation, we aim to identify deformation that departs from the background rate/patterns and, through using time series methods, to detect signals that may not be clear in single interferograms. In previous work (Gaddes et al., 2018), we investigated how best to use blind signal separation methods with interferometric synthetic aperture radar (InSAR) data and demonstrated how the components these methods isolated could be used in a simple detection algorithm. Here, we build on these results to produce a complete automatic detection algorithm, and, as our algorithm contains information about the spatial and temporal nature of deformation at a volcano, we also perform an exploratory study into the inflation prior to the 2018 eruption of Sierra Negra (Galapagos Archipelago, Ecuador).

In our simple detection algorithm (Gaddes et al., 2018), we applied a common machine learning approach, spatial independent component analysis (sICA), to Sentinel-1 data that spanned the 2015 Wolf volcano (Galapagos Archipelago, Ecuador) eruption and were able to isolate three signals that we interpreted as being of geophysical interest and three more as being due to changes in atmospheric conditions. A common problem when applying independent component analysis (ICA) is how to evaluate the reliability of the sources recovered, and in this previous study we performed this step by comparing our recovered sources with those found by other authors who applied different methods to similar data sets. When utilizing ICA in an automatic detection algorithm, this approach cannot be relied upon, and we must instead implement other machine learning methods to automatically assess the significance of the sources we recover before we entrain them into later parts of our algorithm.

The need to ascertain how reliably ICA recovers sources stems from two issues. The first is the more common issue that we wish to determine the statistical significance of our results (i.e., whether it is plausible that they were not simply recovered by chance), while the second is termed “computational reliability” by Hyvärinen

(2012) and is a product of the lack of guarantee that most ICA algorithms will find the global minimum (or maximum) of their objective function.

The computational reliability of the FastICA algorithm can be addressed through running the algorithm multiple times from different starting points (Himberg et al., 2004). This is done in the ICASO algorithm (Himberg et al., 2004) by initiating the first guess of the matrix that recovers the latent sources (commonly termed \mathbf{W} in ICA literature) using a set of random numbers and so ensures that a variety of local minima are sampled as this matrix is iteratively updated. Himberg et al. (2004) argue that as some sources are recovered accurately at all local minima, through sampling many minima we can determine which sources are the most robust as these are the ones that are likely to be recovered at the majority of minima. The sources recovered from these multiple runs can then be analyzed using machine learning methods such as clustering methods, in which compact and isolated clusters are deemed to contain robustly estimated sources.

To address the statistical significance of the results, Meinecke et al. (2002) and Himberg et al. (2004) showed that through resampling the data before the FastICA algorithm was applied; the differences in the sources recovered could be used to establish which were the most reliable. The ICASO algorithm performs this randomization using bootstrapping, in which subsets of the baseline data are generated through randomly selecting (with replacement) a certain number of the original baseline data. The multiple realizations of the recovered sources can then be analyzed through the same clustering approach that was discussed in the preceding paragraph. In practice, the two methods can be used in parallel to produce a single suite of sources recovered from multiple FastICA runs that then require clustering.

A subsequent approach is that of the ISCTEST algorithm (Hyvärinen, 2011; Hyvärinen & Ramkumar, 2013), which tests the reliability of the recovered components on separate data sets that are expected to contain the same underlying signals and as a result is able to ascertain which components are statistically significant. The original algorithm performed the comparison through analyzing the similarities of the mixing matrices (Hyvärinen, 2011), but a subsequent version performs this analysis on the recovered sources instead (Hyvärinen & Ramkumar, 2013). The algorithm was originally used with intersession or intersubject medical imaging but has been applied to InSAR data by Ebmeier (2016), who subdivided the data used into two independent groups (i.e., interferograms in Group 1 did not share any acquisition dates with those in Group 2) before performing the analysis.

We choose to implement the ICASO algorithm rather than the ISCTEST algorithm as our goal of creating an automatic detection algorithm centers on characterizing the baseline behavior of a volcano, instead of isolating a signal of geophysical interest for further investigation. Consequently, we are less concerned with specific components that are recovered and instead focus more on the several sources that we require to characterize a volcano's background behavior.

2. Methods: Automatic Detection Algorithm

Our automatic detection algorithm can be divided into three sections. The first uses the FastICA algorithm within our improved version of the ICASO algorithm to isolate signals of geophysical interest from a time series of interferograms. The second uses the components learned in Stage 1 to characterize the baseline data, while the third then ingests new interferograms as they are formed and determines if the signals present have deviated strongly enough from those in the baseline data to warrant flagging the volcano as having entered a period of unrest. These three stages are described in more detail in the following three subsections, while how to apply the FastICA algorithm to InSAR data is discussed in Gaddes et al. (2018).

In the following description of our algorithm, we consider a “daisy chain” (Biggs et al., 2009) of short temporal baseline Sentinel-1 interferograms of 10^5 pixels that are being automatically created by a processor such as LiCSAR (González et al., 2016). We wait until 15–30 interferograms have accrued (around 180–260 days of data when new images are acquired every 12 days) and use these as our baseline data, while we consider interferograms created after this point as testing data.

2.1. Robust Recovery of Latent Sources

The original ICASO algorithm is described fully in Himberg et al. (2004) and has a modular structure that implements several disparate machine learning methods. Since its creation, the methods used for several modules of the algorithm have been surpassed by newly published methods, which has led to our creation

of a modernized Python version that is specialized for use with InSAR data. We term our Python-based algorithm ICASAR and make it freely available via GitHub (Gaddes, 2017).

Figure 1 shows the intermediate steps associated in running the ICASAR algorithm with InSAR data to recover latent sources and address their significance. In a manner similar to the original ICASO algorithm, our ICASAR algorithm initially calls the FastICA algorithm multiple times with either (1) different starting conditions for the unmixing matrix, \mathbf{W} , (2) bootstrapping of the input data, or (3) both of these steps. When the ICA algorithm is set to recover six sources and run 100 times, we would expect to have a suite of 600 sources, of which many are very similar and reflect true latent sources, while others are recovered infrequently and reflect elements such as combinations of two latent sources or noise. The ICASAR algorithm then performs the following two steps in parallel with the aim of differentiating between the most and least robust sources.

The first method uses a clustering algorithm to identify which sources are similar and label these as belonging to a cluster. This is performed in the ICASO algorithm through the use of agglomerative clustering with average linkage criterion (Himberg et al., 2004), while using the absolute value of the correlation between each source as the distance metric (clustering images is discussed more fully in the supporting information). This method considers each source as an individual cluster, before sequentially joining those that are most similar until a single cluster is formed. The resulting tree like structure can then be cut at a level determined by the user selecting how many clusters they wish to recover. However, this requires input from a user and is therefore not suitable for use in an automatic detection algorithm. Consequently, we have exchanged this section of the ICASO algorithm for the newer algorithm, hierarchical density-based spatial clustering of applications with noise, or HDBSCAN (Campello et al., 2015). This algorithm creates the full hierarchy of merges as clusters form but is also able to cut the tree based on the stability of clusters throughout the merging process and so automatically determine the optimal number of clusters. An additional advantage of HDBSCAN over agglomerative clustering is that it is able to determine points that do not belong to any cluster and to label these as noise.

The second method seeks to provide a representation of the similarities between each recovered source in a manner that is easily interpretable to a human. This is achieved through considering each recovered source as a sample within a space with as many dimensions as the sources have pixels and then fitting a 2-D manifold through this space which preserves the distances between pairs of points in the high dimensional space and on the 2-D manifold. In order to avoid the issue of sign flipped versions of a source being treated as dissimilar, we use the same custom distance metric introduced in the previous paragraph when finding the manifold. The original ICASO algorithm utilized curvilinear component analysis (Demartines & Herault, 1997), but in ICASAR we substitute this algorithm for the newer t-distributed stochastic neighbor embedding (t-SNE; Maaten & Hinton, 2008), as this has been shown to produce 2-D maps that reveal structures within the data more robustly than a variety of methods, including curvilinear component analysis (Maaten & Hinton, 2008).

These two methods can be used in a complementary fashion to create the 2-D plot shown in Figure 1f, in which the points that represent each recovered source are colored depending on which cluster they are a member of or if they are noise. While inspection of this plot is not required, it can be used by a human observer to qualitatively ascertain the robustness of recovered sources. Himberg et al. (2004) suggest that compact and isolated clusters are considered to contain robust sources, while those that form more indistinct clusters are considered to contain more spurious sources. In addition, through coloring the points with the label attached by the HDBSCAN algorithm, we can determine the level of agreement between the two methods, and we postulate that this provides another method to ascertain the robustness of each recovered source.

A quantitative approach to the robustness of each recovered source is also presented in Himberg et al. (2004). They consider the most significant sources to be the ones that form compact and isolated clusters, which they measure as the difference between the mean intracluster similarity (ideally 1) and the mean similarity between members of the cluster and all other sources (ideally 0). They term this the cluster quality index, I_q , and we use it to rank sources in importance before their use in subsequent stages of the algorithm. The ICASO algorithm also allows for selection of a source that is most representative of each cluster, which Himberg et al. (2004) term the centroid. This is calculated as the point within a cluster that minimizes the distance between it and the remaining points within that cluster. Combining the centroids with the cluster quality index allows us to rank some sources as more robust and significant than others and to take these sources and their associated confidences into the subsequent stages of our automatic detection algorithm.

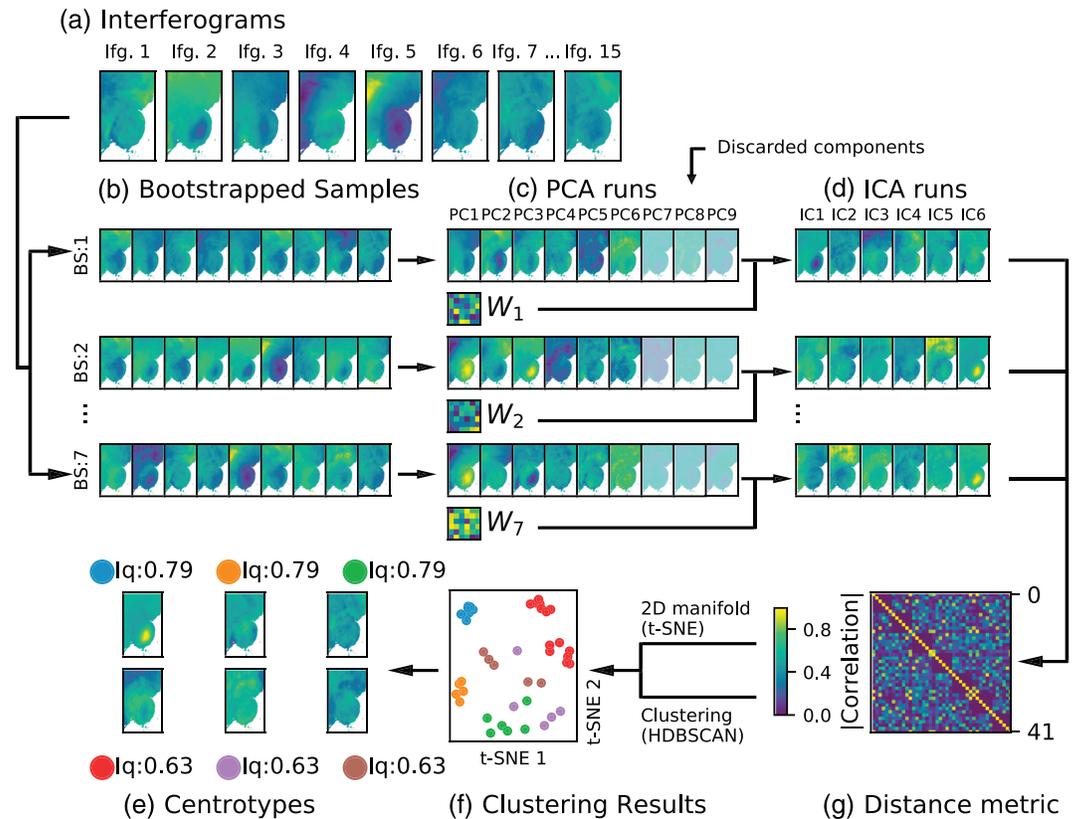


Figure 1. Depiction of the ICASAR algorithm when applied to a time series of interferometric synthetic aperture radar data that cover Sierra Negra (Galapagos Archipelago, Ecuador). This time series features uplift of the caldera floor, which is clearest in Interferogram 2, and poor coherence requiring the masking of pixels on the southeastern flank of the volcano (i.e., the white area). Interferograms are chosen from the mean centered baseline data (a) randomly to create multiple bootstrapped samples (b). PCA is performed on each of these samples (c), and lower-order components are discarded (shown with reduced opacity) to reduce the dimensionality to the required level. Spatial ICA is then performed on each whitened sample (d), before the similarity/distance between each recovered source is estimated (e). The distance matrix is then used in parallel by the HDBSCAN clustering algorithm and the t-SNE manifold learning algorithm to produce a 2-D visualization of the recovered sources (f). The centrotypes of each cluster are then recovered (g) and ordered by their cluster quality index, Iq. PCA = principal components analysis; ICA = independent component analysis; t-SNE = t-distributed stochastic neighbor embedding.

In the first part of our automatic detection algorithm, we utilize the ICASAR algorithm with both bootstrapping and random initiation of the unmixing matrix and seek 200 runs that converge. The choice of the number of components to recover with ICA when the number of latent sources is unknown remains a difficult problem. In our previous study, we found that when applying the algorithm to noisy data, the best results are generated when the algorithm seeks around two sources more than present in the data, as this ensures that the majority of the signals of interest that may exist in lower-importance principal components are not discarded (Gaddes et al., 2018). Consequently, we set the FastICA algorithm to recover around six sources, as we postulate that majority of the world's volcanic centers are likely to contain several consistent atmospheric signals (e.g., a topographically correlated atmospheric phase screen [APS] or consistently orientated pressure gradients) and possibly persistent deformation (e.g., subsidence, such as that measured by InSAR at Askja, Iceland, since 1995; Pagli et al., 2006). However, in the simpler synthetic cases that we present in Figures 2 and 3, our data were generated from only three and two latent sources, respectively, and consequently we set the FastICA algorithm to recover only four sources (i.e., approximately two more than were thought to have generated the data). When applying the algorithm globally, the choice of the number of sources to recover does not impact the number of signals our algorithm could isolate in total but does impact the number it could isolate at a given volcano. Therefore, if a single volcano contains an

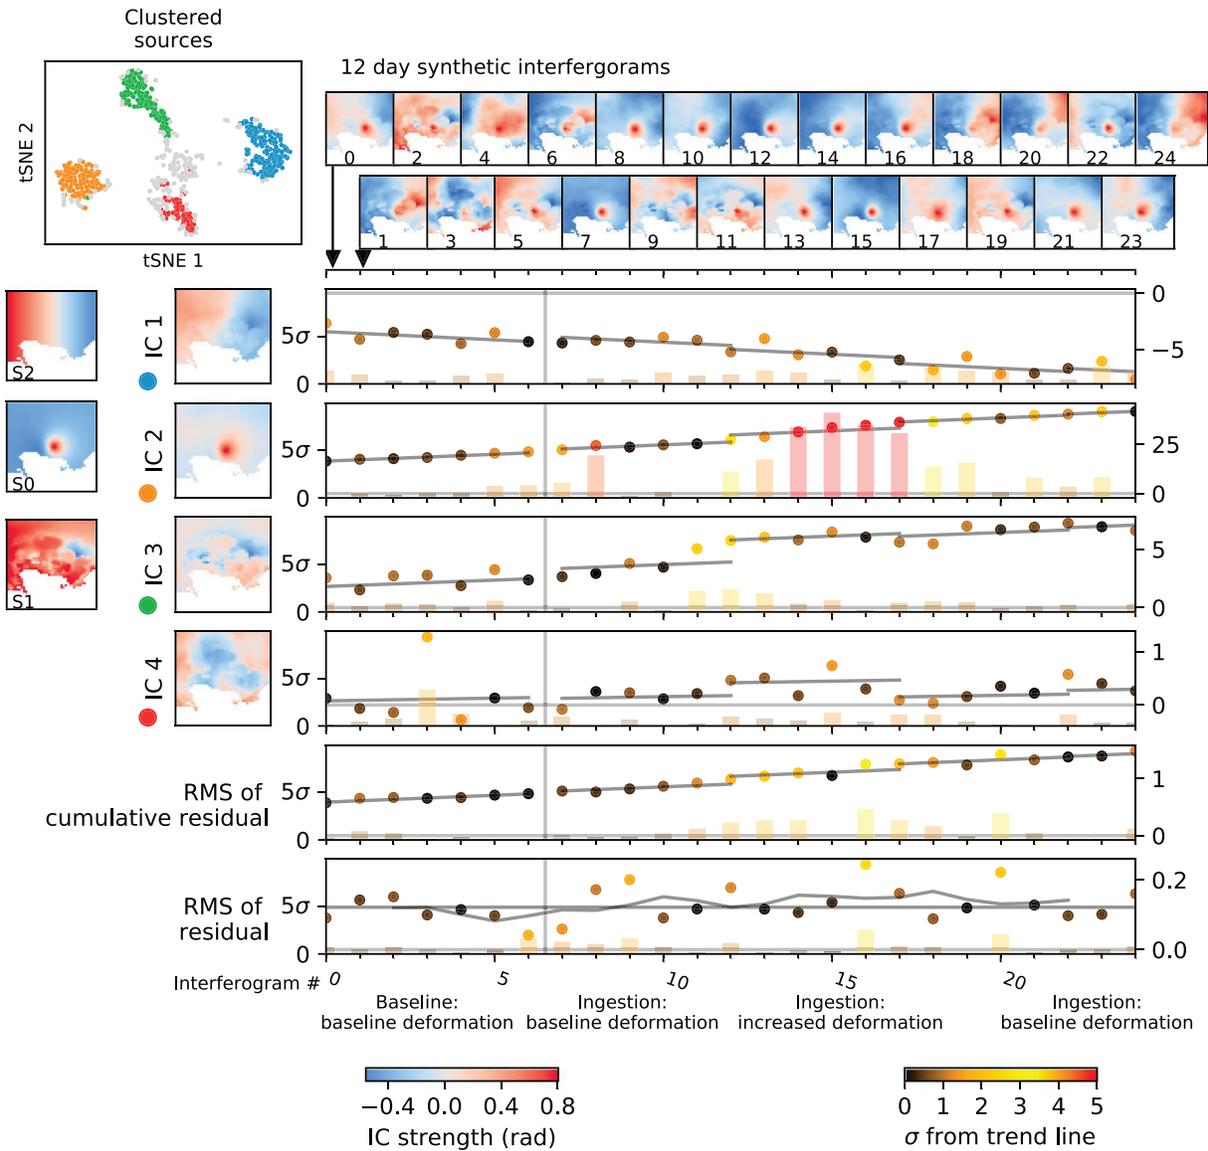


Figure 2. The results of applying the automatic detection algorithm to a synthetic time series of 25 interferograms over the Campi Flegrei caldera complex (Italy), which features a change in rate of the previously constant uplift to simulate a period of volcanic unrest. The interferograms are shown on the top two rows, with the arrows highlighting that the lower left corner of each interferogram is taken as its x value. The centropypes of the four clusters are shown as IC1–IC4, of which three can be seen to correspond to synthetic sources (S0 = deformation source; S1 = topographically correlated atmospheric phase screen; S2 = east-west phase gradient). The cumulative time course of each recovered source is shown to its right, along with every fifth line of best fit drawn at the gradient learned during the training stage. The deviations between each point and the line of best fit are shown as both the color of the data point and as the height and color of the low opacity bars (black to red). Interferograms 12–16 contain an increased contribution from the deformation source (S0), and this synthetic period of unrest is flagged in the time course for IC2 (purple to orange labeling of points). t-SNE = t-distributed stochastic neighbor embedding; RMS = root-mean-square.

unusually high number of different types of deformation, increasing the number of sources to be recovered slightly may be necessary to isolate them all successfully.

It should be noted that when bootstrapping our baseline data, we choose a subset of the interferograms at random with replacement. However, we must have at least as many independent interferograms in our baseline data as sources that we seek to recover. To demonstrate this case, we consider n different interferograms that are chosen with replacement to create a sample containing m interferograms (e.g., from a baseline set of five interferograms, we choose a sample of 10, and so some interferograms must feature multiple times). The data now lie within a nD hyperplane in the mD space, and only n principal components can be found. If we

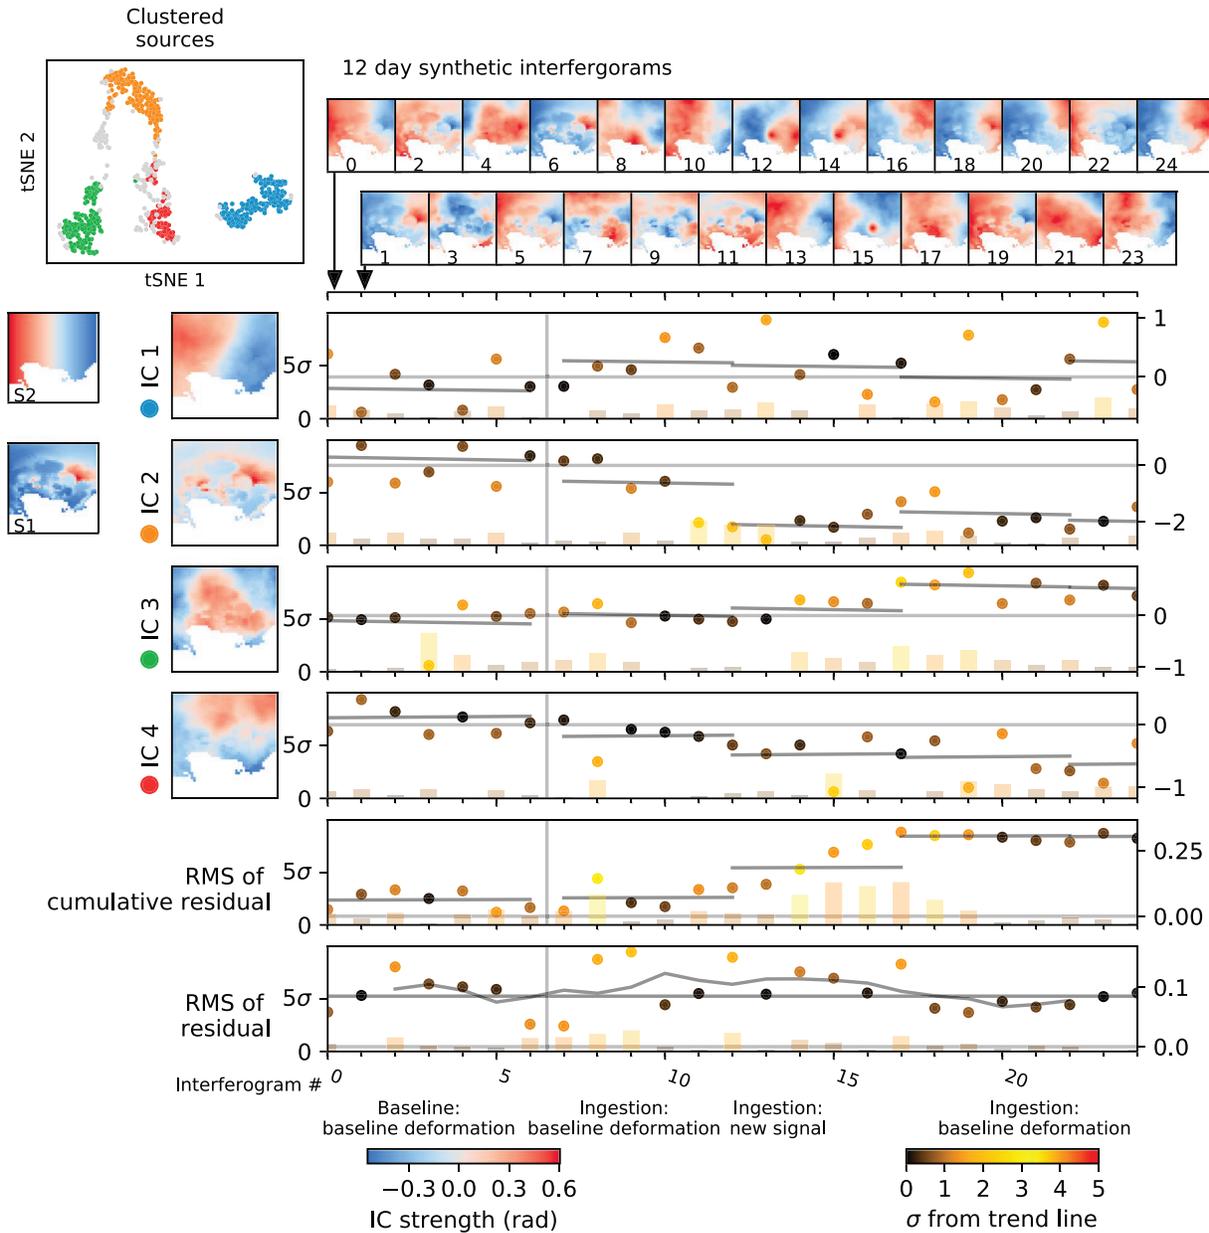


Figure 3. The results of applying the automatic detection algorithm to a synthetic time series of 25 interferograms over the Campi Flegrei caldera complex (Italy), which features the emergence of a new signal to simulate a period of volcanic unrest. The majority of the features of this figure are consistent with Figure 2, with Interferograms 12–16 again containing the synthetic period of unrest. As these interferograms contain a new deformation signal, they cannot be fit well by the background sources, and so the RMS cumulative residual is seen to deviate from the baseline rate, causing the points to be flagged as showing significant deviations from the baseline phase. t-SNE = t-distributed stochastic neighbor embedding; RMS = root-mean-square.

seek s sources and $s > n$, the preprocessing step would fail and we would not be able to perform ICA. Consequently, we reject any bootstrapped samples which contain less than s independent interferograms.

As our goal is to automate the detection of periods of volcanic unrest, we avoid any manual inspection of the results of the ICASAR algorithm through using the number of clusters that is automatically selected by the HDBSCAN algorithm and rank these from most to least confident using the cluster quality index, I_q . At the conclusion of our first step, we have recovered around six spatial maps that express the spatial nature of both atmospheric signals and any deformation signals present at a volcanic center and can rank these signals in terms of our confidence of their significance. It should be noted that in the case that more than six sources

generated the data, it is still possible that all of them may be recovered despite the FastICA algorithm only seeking six sources, as the number of clusters that HDBSCAN detects determines the number of sources that are selected, and this is not constrained to be less than the number of sources that the ICA algorithm recovers. In the case that eight sources are present in the data (e.g., multiple atmospheric signals and multiple deformation signals), but the FastICA algorithm was only set to recover six sources, we would expect a different selection of the eight true sources to be recovered each time, which would still result in eight clusters being identified by HDBSCAN.

2.2. Characterization of the Baseline Data

The second stage of our algorithm seeks to characterize the temporal nature of the sources recovered in the previous stage. Figures 2 and 3 depict the application of our algorithm to synthetic data, and the points detailed in the remainder of this section are illustrated in these figures. To determine the temporal nature of the spatial signals, we perform a simple least squares inversion to fit each of the training interferograms using the recovered spatial sources. The time history for each spatial pattern is commonly termed a “time course” in ICA literature, and through summing these we can ascertain the cumulative use of a given spatial pattern throughout the baseline data.

To characterize the use of these spatial patterns, we first fit a linear trend line through the cumulative time courses that span the baseline stage and calculate the residuals between each data point and the line of best fit. For each source, we can compute the standard deviations of the residuals, before classifying the number of standard deviations each point is from the line. This is shown in Figures 2 and 3 as the colors of each point, ranging from black for points lying on the trend line to orange for points lying over five standard deviations from the trend line.

Our ability to fit the baseline data using our learned sources can be also characterized through measuring the residual between each interferogram and the results of the inversion to fit it with the recovered sources. The introduction of a new signal will decrease our ability to fit new interferograms using the existing recovered sources and so lead to a detectable increase in the residual. We measure the residual in two complementary ways; in the first, we record the root-mean-square (RMS) of the residual between each interferogram and the weighted sum of sources from the inversion, which we term “RMS residual.” In the second, we sum the residual for each pixel throughout time, before calculating the RMS of the residual at each time step, and term this “RMS cumulative residual.” This second method can be used to discard false positives that may be caused by strong nonstandard atmospheric signals in a single SAR image, as the interferograms before and after this image will show the signal with the opposite sign and cancel in the cumulative residual. Note that only nonconsistent signals, such as turbulence, can lead to a strong residual; signals that appear in multiple acquisitions, such as the correlation with topography, are captured in modeled components. As is the case for the time courses, the onset of slow deformation will show as a change in rate for the cumulative residual, although always as an increase rather than a decrease. A strong atmospheric signal in a single acquisition would also cause a significant jump in both residual terms, but the cumulative residual of the subsequent acquisition will then drop down again. In a manner similar to that described for each spatial pattern's baseline cumulative time course, we fit lines to each type of residual, calculate the standard deviation of the line to point misfits, and then use this to determine if future deviations from the line of best fit are significant.

2.3. Ingestion of New Interferograms

When a new interferogram is formed, it is ingested into the algorithm and a simple inversion is performed to fit it using the learned spatial components (i.e., in a manner similar to the baseline data). For the cumulative time courses, RMS residual, and RMS cumulative residual, the line of best fit is then extrapolated to the date that the new interferogram spans and the residuals calculated. This is then compared to the standard deviation of the previous residuals, in order to determine how significant any changes are. In the case of an individual component contributing significantly more or less than before, the gradient of the cumulative time course will change, causing the new point to lie a large distance from the line of best fit.

Following a sustained period of changed rate of either a cumulative time course or RMS cumulative residual, the behavior of the signals present in the time series may return to that seen during the ingestion phase. Through periodically redrawing the lines of best fit at the same gradient as learned during the training

phase, but shifted appropriately vertically, the data points after the period of unrest again lie close to the line of best fit and are no longer flagged as expressing a significant deviation from the baseline stage, provided that the physical processes operating during the baseline stage are again active in the same manner. This is demonstrated in the time course of IC2 in Figure 2 and the RMS cumulative residual in Figure 3.

In the following two sections, we present two synthetic time series that cause either a recovered sources' cumulative time courses or the RMS cumulative residual to change significantly. Both time series contain a deformation signal (modeled as the inflation of a point source in an elastic half space; Mogi, 1958), a turbulent APS, an east-west phase gradient, and a topographically correlated APS (discussed in Gaddes et al., 2018).

2.3.1. Acceleration of an Existing Signal

This scenario aims to demonstrate our algorithm's ability to detect a period of volcanic unrest created when a relatively steady geophysical process changes in rate during the time series. As this is a synthetic data set, we know that only three signals were used to generate the time series, while the turbulent APS contributes only noise. Consequently, we set the FastICA algorithm to recover five sources, as this follows the previously discussed point of searching for several more sources than we think were used to generate the data. Figure 2 shows the results of applying first the ICASAR algorithm and then our automatic detection algorithm to these data, with the 2-D plot showing the results of these repeated ICA runs. All of the intermediate sources are seen to form distinct clusters which are identified by both t-SNE (the 2-D manifold learning method) and HDBSCAN (the clustering algorithm), which we interpret as an indicator of the robust nature of each of the sources. Visual inspection of IC1, IC2, and IC3 suggests that these are accurate reconstructions of the three synthetic sources (S1, S2, and S3), while the fourth appears to capture an aspect of the turbulent atmosphere.

During the "Ingestion: baseline deformation" phase, the time series continues in a manner similar to that of the baseline data and no significant deviations from the extrapolated trend lines are seen for either the cumulative time courses and cumulative residual. However, in Interferograms 12 to 16, the strength of the deformation signal in the synthetic time series is approximately doubled, to mimic a period of volcanic unrest due to increased inflation. This causes IC2 to be used more strongly in the inversions to fit each of these algorithms, which has the effect of increasing the gradient of the cumulative time course of IC2. As this deviates from the extrapolated line of best fit, the points become further from the line and are flagged as showing more significant deviations (orange to yellow colors). Insignificant deviations are seen in the remaining cumulative time courses and RMS cumulative residual, as by increasing the strength of IC2, we continue to fit the new interferograms well. After Interferogram 16 the time series returns to the behavior seen during baseline stage, and the algorithm returns to assigning insignificant deviations to each data point.

2.3.2. Emergence of a New Signal

This scenario aims to demonstrate our algorithm's ability to detect a period of volcanic unrest created when a new deformation signal enters the time series. Figure 3 shows the results of applying first the ICASAR algorithm and then our automatic detection algorithm to this data. We observe that the two most significant sources recovered (IC1 and IC2) correspond to the two synthetic sources, while we interpret the remaining two as capturing aspects of the turbulent APS. Interferograms 12–16 contain a deformation signal that was not present in the baseline data, and consequently our algorithm is unable to fit these interferograms well, which produces an increase in the RMS cumulative residual. This deviation is flagged as the orange/yellow points and ceases when the new deformation signal disappears from the time series.

3. Application to Sierra Negra

To demonstrate our automatic detection algorithm's ability to detect signs of volcanic unrest in real data, we present results from Sierra Negra, a shield volcano in the Galapagos Archipelago (Ecuador). This example was chosen as we expect it to contain the two types of volcanic unrest we aim to detect: Acceleration of uplift occurs before the June–August 2018 eruption, and the eruption itself produces deformation signals not previously seen in the baseline data.

Sierra Negra is a basaltic shield volcano located on the southern half of Isabela island in the Galapagos Archipelago (Reynolds et al., 1995). Figure 4 shows an overview of the area described, with the caldera of Sierra Negra visible within the box delimiting the area shown by the following figures. The asymmetry of the area used in the following figures is due to the difference in vegetation between the northern rocky

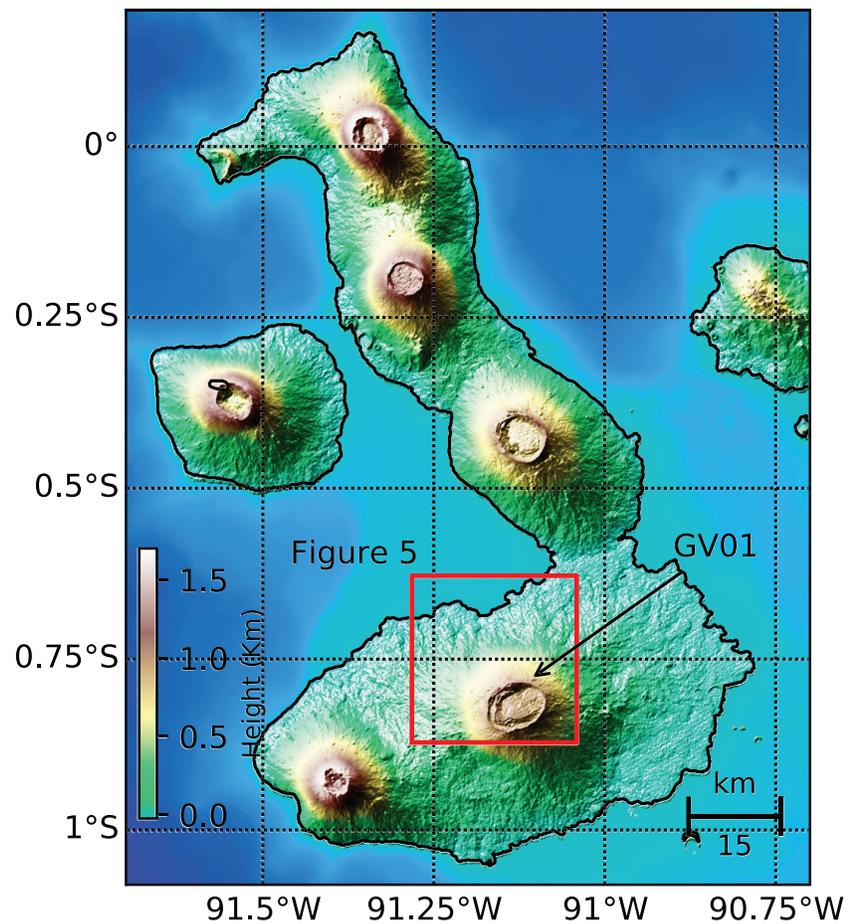


Figure 4. Overview of the Western Galapagos Islands of Fernandina and Isabella (Ecuador). Sierra Negra is visible as the area of high topography delineated by the red box, which shows the extent of the interferograms presented in the remaining figures. An arcuate area of high topography is visible on the southern and western edges of the caldera floor, which has been interpreted as the results of repeated trapdoor faulting events (Chadwick et al., 2006). The GPS station used in section 4 is labeled as GV01.

slopes which provide good InSAR coherence and the densely vegetated southern slopes which do not. On the southern and western edges of the caldera floor, a second area of raised topography is seen, which is attributed to the interplay of cycles of uplift and trapdoor faulting (Jónsson et al., 2005).

Prior to the 2018 eruption, Sierra Negra last erupted in 2005. Inflation preceding this eruption was imaged using both InSAR (Amelung et al., 2000; Jónsson, 2009; Jónsson et al., 2005) and GPS (Chadwick et al., 2006), with total uplift between 1992 and 2005 reaching just below 5 m prior to the eruption. During this period, uplift was in part accommodated by trapdoor faulting, which may have acted to reduce the buildup of strain and delay the eventual eruption (Chadwick et al., 2006). Modeling of the observed deformation by different studies suggested that it was caused by the inflation of a sill at a depth of ~ 2 km below the caldera floor, or ~ 1 km below sea level (Amelung et al., 2000; Chadwick et al., 2006; Jónsson, 2009; Yun et al., 2006).

Between the 2005 and 2018 eruptions, the deployment of a temporary broadband seismic network termed the SIGNET array (Tepp et al., 2014) has provided additional insights into the structure of the crust beneath Sierra Negra. Body wave tomography has been used to infer the existence of large low velocity zone at depths of 8–15.5 km below sea level (BSL) coupled with smaller areas of high and low velocities at depths of 3–15.5 km BSL (Tepp et al., 2014), while attenuation tomography has been used to infer the existence of separate low velocity zones from 0.5–3 and 3–10.5 km BSL (Rodd et al., 2016). Combining these measurements with the geodetic studies of deformation before the 2005 eruption has led to the conclusion that both a shallow magma chamber and a deeper magma chamber embedded in a larger mush zone exist below

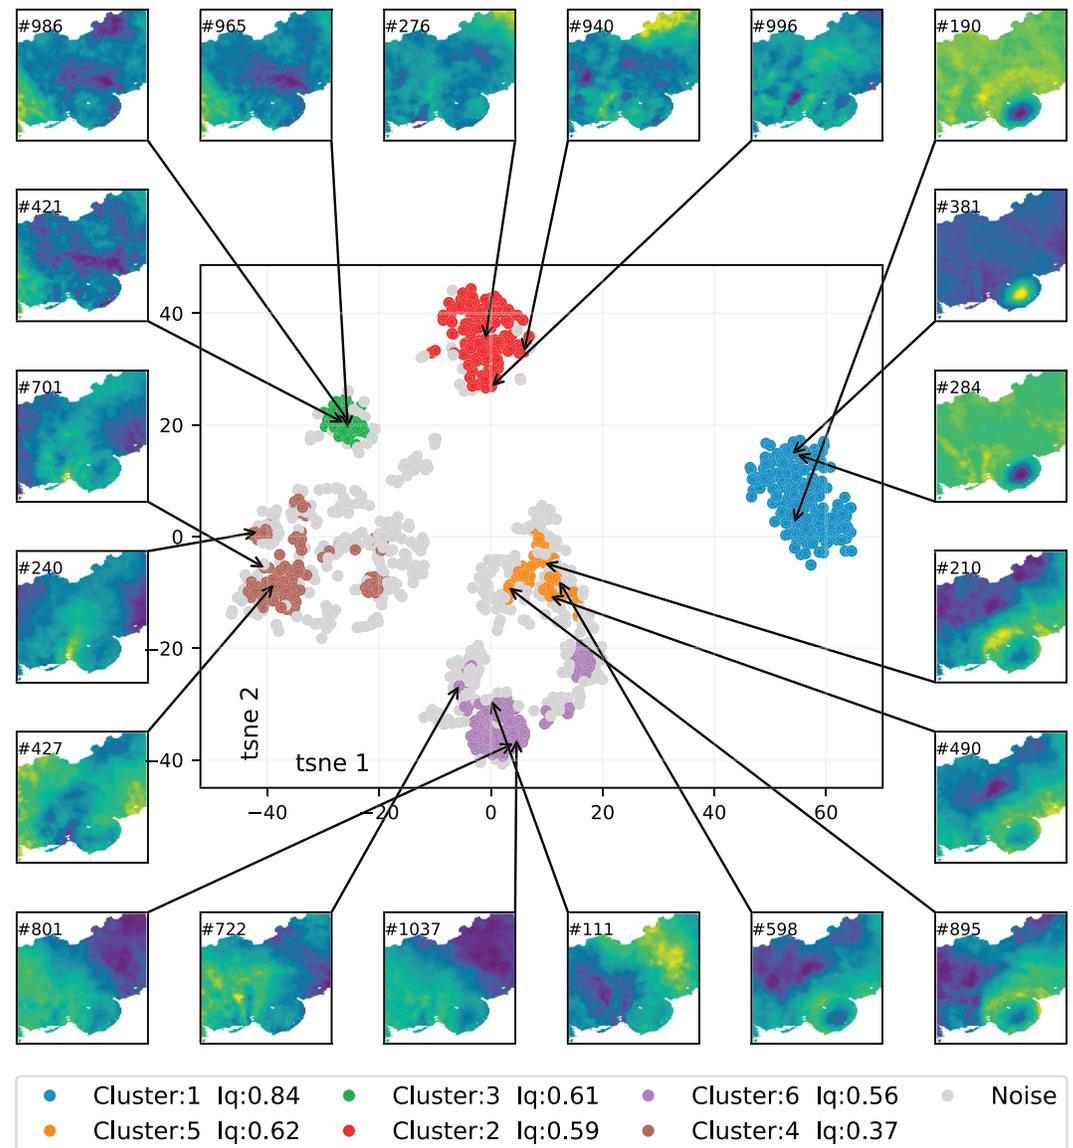


Figure 5. Results from application of the ICASAR algorithm to the time series of Sentinel-1 data at Sierra Negra Volcano (Galapagos Archipelago). The sources that are recovered by the multiple runs of the FastICA algorithm are expressed as dots in the central scatter plot, while a random subset of these are shown in full around the edge of the plot. The position of each point within the 2-D scatter plot is found by the manifold learning method t-SNE, while the labels and associated colors are found by the clustering algorithm, HDBSCAN. The two methods can be seen to be broadly in agreement, with isolated clusters that are formed by t-SNE being labeled homogeneously by HDBSCAN. The clusters are ranked by their cluster quality index, I_q , with the highest value attained by the cluster that contains the caldera floor deformation signal. t-SNE = t-distributed stochastic neighbor embedding.

Sierra Negra (Rodd et al., 2016). This theory is in broad agreement with studies of other Galapagos volcanoes, such as Bagnardi and Amelung (2012), which identified at least two areas of magma storage below Fernandina Volcano, and Stock et al. (2018), which identified two magma reservoirs beneath Wolf Volcano.

The 2018 eruption began in the afternoon of 26 June, with lava emitted from four fissures (Sennert, 2018a), which primarily flowed down the northern flank to reach the sea (Vasconez et al., 2018) but also flowed into the caldera (Sennert, 2018a). Lava flows continued to be active throughout July and August, but by 23 August activity had ceased (Sennert, 2018b). During the eruption, SO_2 emissions were visible to the Deep Space Climate Observatory satellite (Carn et al., 2018).

To test our automatic detection algorithm on Sierra Negra, we first create a daisy chain of the shortest possible temporal baseline interferograms from 98 Sentinel-1 SAR images (shown in Figure S1) using the LiCSAR processor (González et al., 2016), which includes filtering with a Goldstein filter (Goldstein & Werner, 1998) and unwrapping using SNAPHU (Chen & Zebker, 2001). The data span 13 December 2014 to 1 July 2018, with the last 12-day interferogram capturing the co-ruptive signal associated with the start of the 2018 eruption. The average coherence for each pixel is calculated after filtering, and any pixels with a mean coherence below 0.7 are removed. As the majority of the southern flank is densely vegetated, these pixels are removed, but the remainder of the volcano exhibits high coherence and the majority of the pixels are retained.

We apply the ICASAR algorithm to the first 35 interferograms (December 2014 to April 2015) and recover the suite of sources shown in Figure 5. Visual inspection suggests that the clustering performed by HDBSCAN and the manifold learning performed by t-SNE are broadly in agreement, as the distinct clusters found by t-SNE are similar to the clusters found by HDBSCAN. The order of the clusters is random, but the cluster quality index (I_q) is displayed in the legend and can be used as a metric to rank the sources based on their robustness, as discussed in section 2.1. While further human analysis of these components is not required for our automatic detection algorithm, we present a brief discussion here, as stand-alone use of the ICASAR algorithm to isolate signals of geophysical interest may be useful for motives other than volcano monitoring (e.g., the extraction of small signals associated with the cooling of lava flows for further study; Chaussard, 2016). Figure 6 shows the centrotypes of each cluster ordered by their cluster quality index and compares them against the Shuttle Radar Topography Mission Digital Elevation Model (DEM, Farr et al., 2007). Visual inspection of IC1 suggests that this signal contains the uplift signal centered at the caldera, and the IC strength versus height plot shows that the majority of the signal is located above ~ 800 -m height, which can be seen from the DEM to be the approximate height of the caldera floor. A combination of visual inspection of the sources and computation of the correlation coefficient between each recovered source and the DEM suggests that as IC2 is very similar to topography, it is likely that this component is capturing a topographically correlated APS. However, above ~ 800 -m height, some pixels lie significant distances from the linear trend, suggesting that a small amount of the uplift signal may be present in this component. The remaining components show no easily interpretable patterns and are likely to correspond to atmospheric signals.

Figure 7 shows the results of applying our automatic detection algorithm to the time series. The most striking feature is the flagging of the time course of IC1, as indicated by the orange coloring of the points, due to the rate of inflation increasing. Closer to the eruption, other time courses also exhibit unusual behavior which is flagged as a sign of unrest (e.g., the time course of IC3 from Interferogram 88 onward) and may be due to processes such as slip on the intracaldera faults causing slight changes in the shape of the uplift pattern, which then requires different use of the baseline components during the inversion step. Automatic detection of the new large signals associated with the onset of the eruption captured in Interferogram 97 is achieved through the inability of the learned components to fit these new signals, which causes both measures of the residual to increase rapidly.

Both the RMS residual and RMS cumulative residual also increase transiently during Interferograms 55 and 56. Inspection of these interferograms (shown in Figure 7) shows that the lower left quadrant of each interferogram contains a strong signal of opposite sign. We conclude that this is due to a strong APS in the SAR image that the two interferograms share and highlights the ability of our two measures of the residual to determine this. While the RMS residual rises for these two interferograms as each cannot be fit well, the RMS cumulative residual rises for a single image, before falling back to a level that does not indicate unrest. This is because the opposite sign of the strong atmospheric contribution to each interferogram causes it to cancel when summing the residual for each pixel. Other than waiting for the next acquisition, setting a higher threshold at which deviations from the baseline data are flagged to a user could avoid events such as Interferogram 55 being flagged as indicating unrest.

For a user seeking to apply our algorithm, maximizing the length of the baseline data is also likely to improve the accuracy of the components recovered by allowing more subtle signals to become apparent above background noise. Additionally, increasing the length of the baseline data used at each volcano is likely to allow the ICs selected to characterize more atmospheric signals, and therefore calibrate the expected level of residual better. However, in the case presented here, the baseline stage could not be extended significantly without shortening the ingestion phase before the change in rate observed in mid-2017.

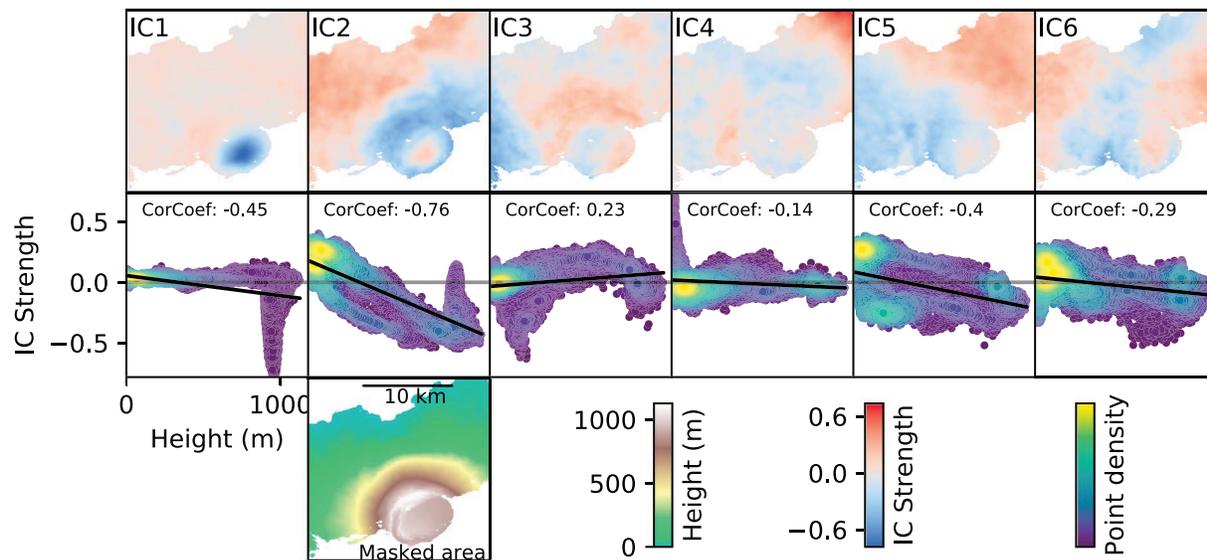


Figure 6. The centrotypes of the six clusters from Figure 5 (top row), the topography as described by the SRTM-1 DEM (Farr et al., 2007, bottom row), and 2-D Gaussian kernel density estimates between the pixels of the DEM and each source (middle row). Visual comparison of the scatter plots of each IC versus DEM height indicates that IC2 most closely resembles the DEM, while the IC2 to DEM scatter plot also shows a clear trend and has the highest correlation coefficient (-0.76). SRTM DEM = Shuttle Radar Topography Mission Digital Elevation Model.

4. Modeling Overpressure Before the 2018 Sierra Negra Eruption

The pressure change in a magma chamber is of geophysical interest as an increase in the pressure can overcome the tensile failure stress of the elastic crust overlying the chamber and so create an opening that allows magma to rise to the surface (Pinel et al., 2010). As our InSAR time series contains significant deformation located within the caldera, we postulate that the 2018 eruption was caused by an increase in magma chamber overpressure and seek to constrain how large this increase was. We expect that future applications of our automatic detection algorithm will identify volcanoes deforming in a similar manner, and modeling of the overpressure could then allow for forecasts of the timing of an eruption (Bato et al., 2017).

Through modeling IC1 as the signal of the caldera preeruptive uplift, we can infer information about the source of this deformation, such as the pressure change. This is likely the same source as the 2005 eruption, although we solve for the depth and geometry independently. Expected values for the overpressure in a magma chamber required for both a dyke to form and magma within it not to freeze before reaching the surface vary widely, from 10–100 MPa for silicic magmas to 1 Mpa for basaltic magmas (Manga & Brodsky, 2006). This wide range is due to variations of viscosity of magmas of different compositions, which impacts the ability of a dyke to continue to propagate without being arrested due to freezing of its walls (Jellinek & DePaolo, 2003). In both endmember cases, the overpressure is lower than values derived in laboratory experiments for the stress required to fracture the elastic crust (8.6 ± 1.4 MPa for basalt, to 13.8 ± 2.1 MPa for granite; Albino et al., 2018; Pinel et al., 2010; Touloukain, 1989). However, these values are likely to be upper bounds, as volcanic processes such as the formation of three-dimensional fracture networks due to contraction during cooling are likely to significantly weaken the crust (Schultz, 1995). As Sierra Negra erupts iron-rich tholeiitic basalts (Reynolds & Geist, 1995), we expect an overpressure required for eruption at the lower end of the reported values. Through measuring the preeruptive pressure change using geodetic methods, we seek to determine if the pressure change is comparable to the stresses required to rupture the elastic crust and allow a dyke to propagate to the surface.

We first modeled the deformation using a kinematic approach, with a horizontal rectangular dislocation embedded within an elastic half space with uniform opening (Okada, 1985), and performed a Bayesian inversion using the GBIS software (Bagnardi & Hooper, 2018). We chose to ignore the effects of topography as the majority of the deformation is located within the flat caldera, and the caldera rim is only ~ 100 m in height.

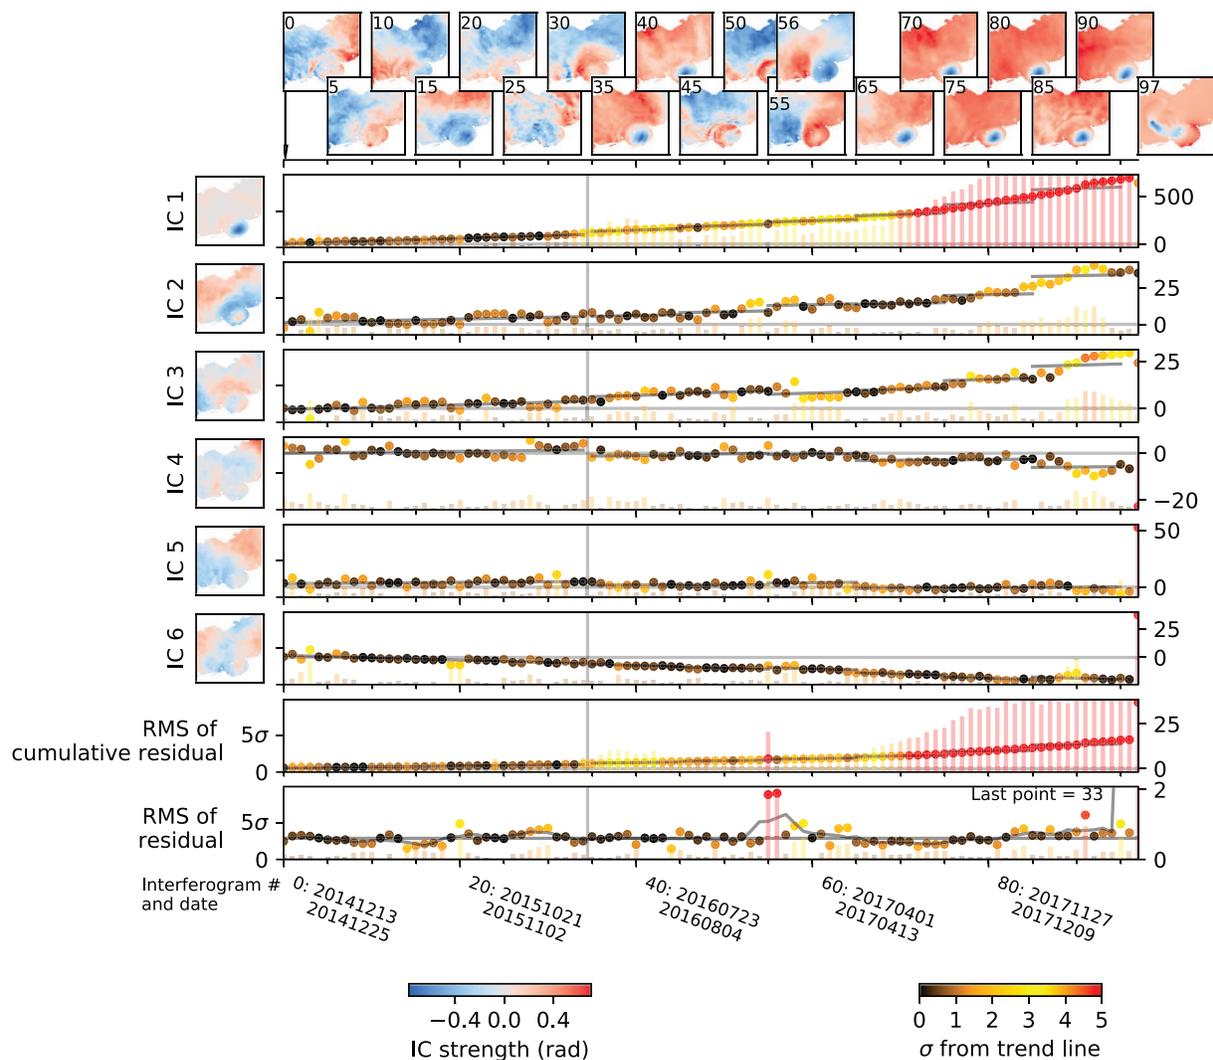


Figure 7. The results of applying our automatic detection algorithm to a time series of Sentinel-1 interferograms which cover the final ~3.5 years of inflation before the 2018 eruption of Sierra Negra (the anomalous points associated with the coeruptive interferogram can be seen on the right-hand y axis). Roughly every five interferograms are shown, but some liberty is taken to ensure those of particular interest are visible (e.g., 55, 56, and 97). The components are initially used in a similar fashion before and after the switch to the ingestion phase (marked by the black vertical line), before more pronounced deformation from around Interferogram 65 causes IC1 to be flagged as having deviated significantly from the baseline data. The residual when the final coeruptive interferogram is fitted is an order of magnitude larger than seen previously and is omitted from the RMS residual plot for clarity. The two high values of RMS residual for Interferograms 55 and 56 are due to a strong atmospheric signal in the acquisition common to the two. RMS = root-mean-square.

However, our model does not take into account the effects of the caldera bounding faults, and the possible implications of this are discussed further in section 5.2. While the data can be fit using this approach, it does not constrain the overpressure within the sill, which required the use of a physical model. We initially tried to fit the deformation using a penny shaped crack (Fialko et al., 2001), but this model's requirement of radial symmetry is not suitable for the rectangular deformation pattern, and we instead used the boundary element approach detailed in Hooper et al. (2011) to solve for spatially variable opening with uniform overpressure for a horizontal sill divided into discrete rectangular patches implemented in GBIS. Using this approach, we solve for the location (x, y, z), length, width, and pressure change. We note that given the longevity of the deformation observed at Sierra Negra, a relatively thick magma chamber is required for it to remain unfrozen. However, as we are able to fit the data well using a modeled sill, we conclude that a diapir-shaped magma chamber is most likely, as these have been shown to produce surface deformation patterns almost identical to sills (Yun et al., 2006).

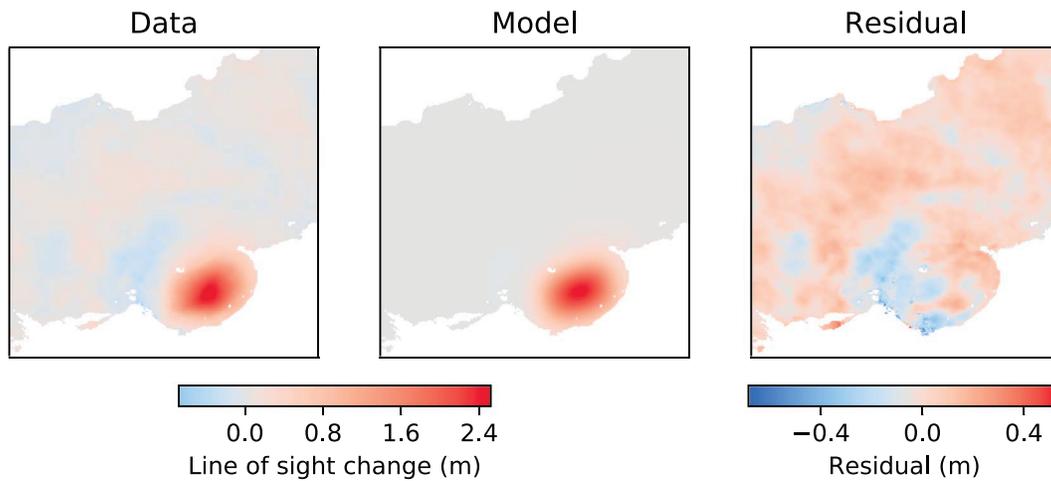


Figure 8. Data: The signal contained in IC1 throughout the Sentinel-1 time series, showing ~ 2.4 m of motion toward the satellite during the Sentinel-1 time series. Model: The result of our optimal forward model, which treats the magma chamber as a 6.2×3.7 -km² rectangular dislocation at a depth of ~ 2.0 km. Residual: The misfit between our model and the data, which is dominated by a mottled pattern across the majority of the scene which independent component analysis is unable to remove from IC1 and our model is unable to fit.

We use the total inflation due to IC1 as the input to our inversion, which we reconstruct using its spatial pattern and time course. The mean centered interferograms for the time spanned by the Sentinel-1 time series can be reconstructed using

$$\mathbf{X}_{ic1} = \mathbf{A}_{ic1} \times \mathbf{S}_{ic1} \quad (1)$$

Where \mathbf{X}_{ic1} are the mean centered interferograms (i.e., the mixtures, in ICA terminology), \mathbf{A}_{ic1} is the first column of the matrix of time courses, and \mathbf{S}_{ic1} is the first row of the matrix of recovered sources. As the interferograms were originally created relative to a reference pixel outside the deforming region (located at 91.2 west, 0.7 south) but mean centered for use with ICA, we return the signal contained in \mathbf{X}_{ic1} to be relative to this reference pixel through performing the reverse of the mean centering processes. We then sum the phase change for each pixel in our reconstructed time series to calculate the cumulative motion of IC1 through the time series (shown in Figure 8), which we use as the input for our modeling.

Figure 8 shows the modeled deformation and the residual between it and the input data. Our best fitting model places a 6.2×3.7 km² rectangular dislocation within the caldera at a depth of ~ 2.0 km which, when the shear strength of the crust is set as 10 GPa (Jónsson, 2009), has undergone a pressure change of 10.4 MPa (probability density functions are provided for model parameters in the supporting information). However, this modeling finds only the pressure change within the modeled sill between the 2018 eruption and the first Sentinel-1 acquisition (13 December 2014), and not since the end of the 2005 eruption. While we could extrapolate the linear inflation seen in the early part of the Sentinel-1 time series (shown as IC1's cumulative time course in Figure 7, and in Figure 9) back to the end of the 2005 eruption, we instead utilize GPS data to first investigate if the spatial patterns remains constant, before using it to extrapolate the InSAR time series.

To determine if the spatial pattern of the deformation source remains unchanged prior to the Sentinel-1 data, we examine the ratios of the east, north, and up components for a selection of GPS stations across the caldera complex. We find that the ratios remain approximately constant from 2005 to 2018 (shown in the supporting information) and conclude that the spatial pattern of the deformation is unlikely to have changed significantly. Consequently, we are satisfied that we can extrapolate our model of an inflating sill from the Sentinel-1 time series to the entire intereruptive time series.

To determine if the temporal nature remains comparable to that measured throughout the Sentinel-1 time series, we use data from one of 10 continuous GPS stations located on Sierra Negra (Blewitt et al., 2018). Station GV01 is not ideally placed in that it lies on the edge of the caldera but, unlike the other nine GPS stations, contains near-daily data between 2005 and 2018. To combine this displacement time series with

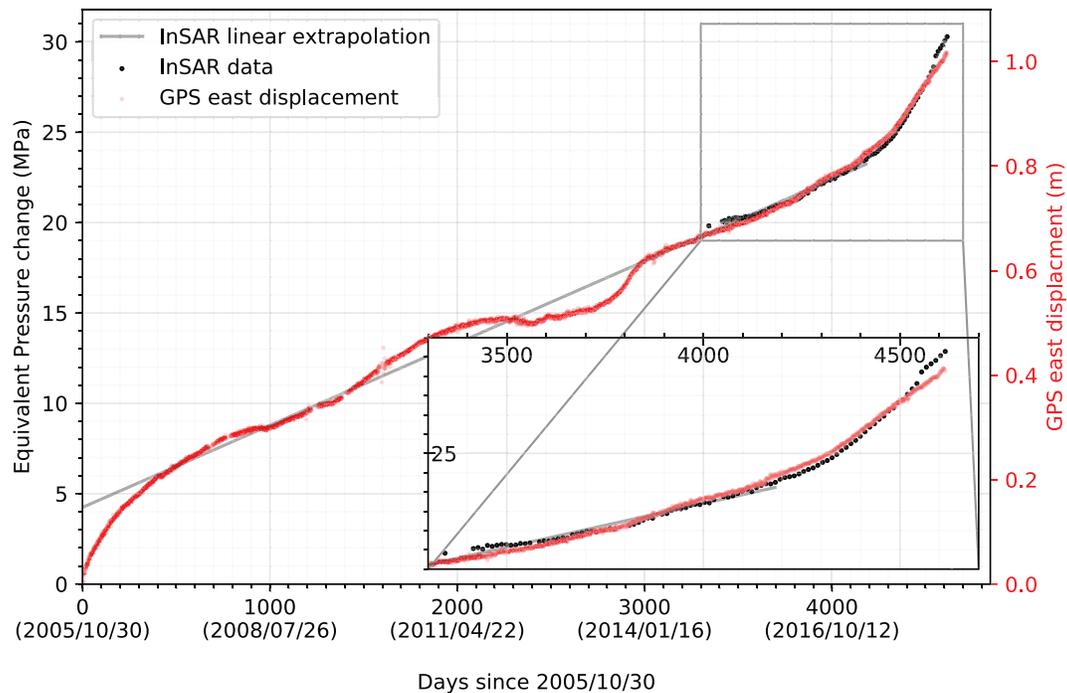


Figure 9. Temporal evolution of the modeled source. The cumulative time course for component IC1 is shown as black points and is rescaled in terms of the pressure change in the modeled rectangular dislocation, which can be seen to attain a maximum value of ~ 30 MPa before the 2018 eruption. The preacceleration section of this time course (prior to Day ~ 4200) can be fit with a linear trend (gray line) and extrapolated back to the 2005 eruption. However, using GPS station GV01 to constrain the temporal evolution of the deformation shows that while the majority of the 2005–2016 time series can be fit with a linear function, there is an initial period of faster uplift, which increases the total pressure change by ~ 4 MPa. InSAR = interferometric synthetic aperture radar.

our InSAR results, we perform a simple inversion to find the two parameters required to rescale and translate the cumulative eastward component of the displacement to match the cumulative time course of IC1. As we solve for only one rescaling parameter, the GPS derived deformation and InSAR derived deformation can only be fit both before and after the rate change if they feature the same proportional increase. We then apply this operation to the entire GPS eastward cumulative displacement time series and set the displacement to 0 after the 2005 eruption. Figure 9 shows the results of this process, and visual inspection shows that GPS data are in broad agreement with IC1's time course, with features such as the change in rate in early 2017 occurring in both time series. The results indicate that the rate before and after the rate change cannot both be fit exactly with only two scaling parameters, and if our inversion were performed before the change in rate observed on Day ~ 4100 to ensure that both had a more closely agreeing gradient (at the cost of the fit after the change in rate), the total pressure change would be slightly decreased, but we estimate that this would produce only a minor change. For comparative purposes, we also fit a linear trend to the initial part of IC1's cumulative time course, and while this fits the majority of the data well, it can be seen to underestimate the deformation due to period of faster motion immediately after the 2005 eruption.

To calculate the total preeruptive pressure change, we rescale the pressure change calculated during the Sentinel-1 time series by the ratio of the total GPS derived displacement, which we show as the y axis of Figure 9. We find that between the 2005 and 2018 eruptions, the observed surface deformation would translate to a change in magma chamber pressure of ~ 30 MPa.

5. Discussion

5.1. Automatic Detection Algorithm

The ICASAR algorithm provides a method to apply sICA to InSAR data and to assess how robust the recovered sources are. The ICASAR algorithm differs from the older ICASO algorithm through incorporating a newer manifold learning algorithm (t-SNE) and a newer clustering algorithm (HDBSCAN) and is optimized for the application of sICA to InSAR data.

When using ICASAR to recover signals of geophysical interest at a given study site, the use of a manifold learning algorithm, a clustering algorithm, and a measure of cluster quality (termed I_q) provides three independent measures of the significance and robustness of the recovered sources. Those deemed of geophysical interest can then be utilized for further investigation, in a manner similar to our use of the deformation source in section 4. We utilize ICASAR as an integral part of our automatic detection algorithm, and the ability of the HDBSCAN algorithm to automatically detect the number of clusters allows for minimal intervention from a human user. The cluster quality index, I_q , also provides a measure to automatically determine which sources come from the most robust clusters.

The application of our automatic detection algorithm to Sierra Negra's preruptive time series has demonstrated its ability to detect changes in established signals; when the existing preruptive inflation changes in rate, this is detected by our algorithm through the marked increase in the number of standard deviations each new data point of IC1 and the "RMS of cumulative residual" lie from the extrapolated lines of best fit (seen more clearly as a lengthening and change in color of the bar graphs from Interferogram ~70 onward in Figure 7). Additionally, our algorithm is able to detect the emergence of new signals, such as those associated with the movement of magma to the surface during the 2018 eruption, through the flagging of the last data points of the time courses of IC3–IC6 and the RMS residual as lying over 5 standard deviations from the extrapolated lines of best fit (seen as their change from black/orange to red at the extreme right of Figure 7).

Our algorithm is also computationally inexpensive, as the baseline data for a given volcano can be converted to a selection of spatially independent sources and their associated time courses on a single-core ~2.4-GHz desktop computer in several minutes, and ingesting new interferograms to be fit using these components can be completed in several seconds. We therefore envisage that our algorithm will be suitable for application to Sentinel-1 data as they are acquired in order to begin the monitoring of subaerial volcanoes using InSAR. Our method differs from that of Anantrasirichai et al. (2018) as instead of using only the most recently acquired interferogram, we use the full time series and can therefore detect subtle signals that are not clear in individual images. Additionally, our algorithm's central tenet of seeking deviations from the baseline behavior also allows it to avoid flagging interferograms that contain deformation of the type that is normal for that volcano, unless there is a change in their strength. False positives generated by strong atmospheric turbulence can be identified as false when the signal reverses in the subsequent image. Inspection of other acquisitions from a different geometry or satellite, prior to the subsequent image, could also be used to rule out deformation. Whether any action is taken based on a single interferogram is a decision for monitoring agencies, who will also use other data that are available to them.

During routine monitoring, our algorithm will be faced with both more subtle signals and noisier interferograms than those presented in this study. We envisage that our primary tool to combat this will be adjustment of the frequency with which our extrapolated lines of best fit are redrawn from the default of every 10 interferograms. As discussed in section 2.3, the redrawing of the lines of best fit ensure that any small misestimations in the baseline gradient do not cause the cumulative time courses/residual to gradually diverge from the line of best fit. However, in the case that existing deformation changes in strength only slightly or a new signal is of small magnitude, the change in rate of cumulative time course/residual increase will only be slight. If lines of best fit are recalculated infrequently (e.g., every 60 interferograms), this slight change in rate will become significant and be flagged. However, in the case that lines of best fit are recalculated frequently (e.g., every five interferograms), the deviation between the cumulative time course/residual and the line of best fit is likely to remain small, and the unrest may not be flagged. An example similar to those shown in section 2 but featuring a more subtle unrest signal that is not visible in a single 12-day interferogram is provided in the supporting information.

5.2. Overpressure Before the 2018 Sierra Negra Eruption

Our estimation of a preruptive overpressure of ~30 MPa in our modeled sill below the Sierra Negra caldera is significantly larger than the overpressure values of ~1 MPa suggested by other authors required for the propagation of a dyke to the surface when a magma chamber contains mafic magma (as discussed in section 4). This value is sensitive to the shear modulus used during the inversion, but even if we reduce this from the value of 10 GPa used by Jónsson (2009) to the lower value of 4.5 GPa suggested by Hooper et al. (2002), the preruptive overpressure reduces to ~14 MPa. While this value does decrease the difference

between our calculated overpressure and existing values of the overpressure required for an eruption, the disparity remains significant and we believe it is unlikely that the 2018 eruption can be thought of simply as the result of an inflating sill that has increased in pressure until the tensile failure stress of the overlying elastic crust has been overcome.

Previous studies of Sierra Negra found that a trapdoor faulting event prior to the 2005 eruption reduced the pressure within the magma chamber by 3 MPa, which has been postulated as a method to delay eruption (Jónsson, 2009). Interferograms that both capture preeruptive inflation and span episodes of faulting within the caldera would feature different spatial patterns to those that feature only inflation, and we therefore expect that when our automatic detection algorithm fits an interferogram containing both signals, it would use the baseline components in a new manner. This is seen in the time courses of IC3 and IC6 (Figure 7), as both deviate from their baseline rates in the latter stages and, while they are unlikely to represent faulting, may be being used by the algorithm to partially fit any new signals that may be present in later interferograms. However, these signals could be due to other processes, and further investigation of the type described in Jónsson (2009) would be required to determine the reduction in overpressure caused by any potential slip events.

Application of our automatic detection algorithm also reveals an increase in the rate of inflation through the change in slope of IC1's time course at the end of 2016 (Figure 7). We postulate that this is due to an increase in the influx of magma to the sill, which would in turn cause the overpressure within it to increase at a rate that outpaced potential pressure-reducing processes such as viscoelastic relaxation and eventually cause the failure in the crust required for the movement of magma to the surface seen during the 2018 eruption. However, the further analysis of the InSAR data required to further explore this hypothesis remains beyond the scope of this paper, which seeks to primarily address volcano monitoring using InSAR.

6. Conclusion

Our study demonstrates that the reliability of latent sources recovered when sICA is applied to InSAR data can be assessed through use of an updated ICASO algorithm. We term this new algorithm ICASAR and note that it is particularly suited for application to automation projects as, unlike the original ICASO algorithm, it is able to automatically determine the number of sources present in the recovered data.

We use our ICASAR algorithm as a constituent of an automatic detection algorithm, which we demonstrate is able to detect signs of volcanic unrest due to both a change in rate of a preexisting signal and the emergence of a new signal. We apply this algorithm to a time series of Sentinel-1 data that span the run-up to the 2018 eruption of Sierra Negra and show that we would have been able to promptly flag this volcano as entering a period of increased activity when the rate of inflation increased approximately 1 year before the eventual eruption and so have provided some warning that the volcano had entered a period of heightened unrest. While Sierra Negra is already closely monitored, the algorithm could be applied globally and hence provide warning for less closely watched volcanoes.

Combining the results of the ICASAR algorithm with a time series of GPS data spanning the period between the 2005 and 2018 eruptions of Sierra Negra, we attribute the intereruptive inflation equivalent to an increase in pressure of 14 MPa in a sill 2.0 km below the caldera floor. This value is significantly larger than the values suggested for a mafic dyke to propagate to the surface, and we postulate that either viscoelastic relaxation is occurring and/or that the spatial pattern of the measured deformation may have changed due to processes such as slip on the intracaldera faults acting to reduce the overpressure within the sill.

References

- Albino, F., Amelung, F., & Gregg, P. (2018). The role of pore fluid pressure on the failure of magma reservoirs: Insights from Indonesian and Aleutian Arc Volcanoes. *Journal of Geophysical Research: Solid Earth*, *123*, 1328–1349. <https://doi.org/10.1002/2017JB014523>
- Amelung, F., Jónsson, S., Zebker, H., & Segall, P. (2000). Widespread uplift and 'trapdoor' faulting on Galápagos volcanoes observed with radar interferometry. *Nature*, *407*(6807), 993–996. <https://doi.org/10.1038/35039604>
- Anantrasirichai, N., Biggs, J., Albino, F., Hill, P., & Bull, D. (2018). Application of machine learning to classification of volcanic deformation in routinely generated InSAR data. *Journal of Geophysical Research: Solid Earth*, *123*, 6592–6606. <https://doi.org/10.1029/2018JB015911>
- Bagnardi, M., & Amelung, F. (2012). Space-geodetic evidence for multiple magma reservoirs and subvolcanic lateral intrusions at Fernandina Volcano, Galápagos Islands. *Journal of Geophysical Research*, *117*, B10406. <https://doi.org/10.1029/2012JB009465>
- Bagnardi, M., & Hooper, A. (2018). Inversion of surface deformation data for rapid estimates of source parameters and uncertainties: A Bayesian approach. *Geochemistry, Geophysics, Geosystems*, *19*, 2194–2211. <https://doi.org/10.1029/2018GC007585>

Acknowledgments

The Copernicus Sentinel data (covering Sierra Negra) were acquired by the European Space Agency (ESA) and were obtained by the authors of this study from the Alaska Satellite Facility archive (<https://www.asf.alaska.edu/>). This work was supported by the NERC Centre for the Observation and Modelling of Earthquakes, Volcanoes and Tectonics (COMET). M. E. G. is supported by the LiCS grant (NE/K011006/1). M. B. was supported by COMET and by an appointment to the NASA Postdoctoral Program at the Jet Propulsion Laboratory, administered by the Universities Space and Research Administration (USRA) through a contract with NASA. The authors are grateful for reviews from Benoit Taisne and an anonymous reviewer and for comments from Susanna Ebmeier and Juliet Biggs made during the examination of MEG's thesis. Figures were prepared in Matplotlib (Hunter, 2007). The Python implementation of HDBSCAN used was from McInnes et al. (2017), and the Python implementation of t-SNE used was from Pedregosa et al. (2011).

- Bato, M. G., Pinel, V., & Yan, Y. (2017). Assimilation of deformation data for eruption forecasting: Potentiality assessment based on synthetic cases. *Frontiers in Earth Science*, 5, 48.
- Biggs, J., Anthony, E. Y., & Ebinger, C. J. (2009). Multiple inflation and deflation events at Kenyan volcanoes, East African Rift. *Geology*, 37(11), 979–982. <https://doi.org/10.1130/G30133A.1>
- Blewitt, G., Hammond, W., & Kreemer, C. (2018). Harnessing the GPS data explosion for interdisciplinary science. *Eos*, 99. <https://doi.org/10.1029/2018EO104623>
- Campello, R., Moulavi, D., Zimek, A., & Sander, J. (2015). Hierarchical density estimates for data clustering, visualization, and outlier detection. *ACM Transactions on Knowledge Discovery from Data*, 10, 1–5.
- Carn, S. A., Krotkov, N. A., Fisher, B. L., Li, C., & Prata, A. J. (2018). First observations of volcanic eruption clouds from the L1 Earth-Sun Lagrange point by DSCOVR / EPIC. *Geophysical Research Letters*, 45, 11,456–11,464. <https://doi.org/10.1029/2018GL079808>
- Chadwick, W. W., Geist, D., Jonsson, S., Poland, M., Johnson, D., & Meertens, C. (2006). A volcano bursting at the seams: Inflation, faulting, and eruption at Sierra Negra volcano, Galapagos. *Geology*, 34(12), 1025–1028. <https://doi.org/10.1130/G22826A.1>
- Chaussard, E. (2016). Subsidence in the parícutin lava field: Causes and implications for interpretation of deformation fields at volcanoes. *Journal of Volcanology and Geothermal Research*, 320, 1–11.
- Chen, C. W., & Zebker, H. A. (2001). Two-dimensional phase unwrapping with use of statistical models for cost functions in nonlinear optimization. *Journal of the Optical Society of America*, 18(2), 338–351.
- Demartines, P., & Herault, J. (1997). Curvilinear component analysis: A self-organizing neural network for nonlinear mapping of data sets. *IEEE Transactions on Neural Networks*, 8(1), 148–154.
- Ebmeier, S. K. (2016). Application of independent component analysis to multi-temporal InSAR data with volcanic case studies. *Journal of Geophysical Research: Solid Earth*, 121, 8970–8986. <https://doi.org/10.1002/2016JB013765>
- Farr, T., Rosen, P., Caro, E., & Crippen, R. (2007). The Shuttle Radar Topography Mission. *Reviews of Geophysics*, 45, RG2004. <https://doi.org/10.1029/2005RG000183.1>
- Fialko, Y., Khazan, Y., & Simons, M. (2001). Deformation due to a pressurized horizontal circular crack in an elastic half-space, with application to volcano geodesy. *Geophysical Journal International*, 146, 181–190.
- Gaddes, M. E. (2017). Icasar. <https://github.com/mattcyp88/ICASAR>
- Gaddes, M., Hooper, A., Bagnardi, M., Inman, H., & Albino, F. (2018). Blind signal separation methods for InSAR: The potential to automatically detect and monitor signals of volcanic deformation. *Journal of Geophysical Research: Solid Earth*, 123, 10,226–10,251. <https://doi.org/10.1029/2018JB016210>
- Goldstein, R. M., & Werner, C. L. (1998). Radar interferogram filtering for geophysical applications. *Geophysical Research Letters*, 25(21), 4035. <https://doi.org/10.1029/1998GL900033>
- González, P. J., Walters, R. J., Hatton, E. L., Spaans, K., McDougall, A., Hooper, A., & Wright, T. J. (2016). LiCSAR: Tools for automated generation of Sentinel-1 frame interferograms. AGU Fall meeting.
- Himberg, J., Hyva, A., & Esposito, F. (2004). Validating the independent components of neuroimaging time series via clustering and visualization. *NeuroImage*, 22, 1214–1222. <https://doi.org/10.1016/j.neuroimage.2004.03.027>
- Hooper, A., Ófeigsson, B., Sigmundsson, F., Lund, B., Einarsson, P., Geirsson, H., & Sturkell, E. (2011). Increased capture of magma in the crust promoted by ice-cap retreat in Iceland. *Nature Geoscience*, 4(11), 783.
- Hooper, A., Segall, P., Johnson, K., & Rubinstein, J. (2002). Reconciling seismic and geodetic models of the 1989 Kilauea south flank earthquake. *Geophysical Research Letters*, 29(22), 2062. <https://doi.org/10.1029/2002GL016156>
- Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing In Science & Engineering*, 9(3), 90–95.
- Hyvärinen, A. (2011). Testing the ICA mixing matrix based on inter-subject or inter-session consistency. *NeuroImage*, 58(1), 122–136. <https://doi.org/10.1016/j.neuroimage.2011.05.086>
- Hyvärinen, A. (2012). Independent component analysis: Recent advances. *Philosophical Transactions of The Royal Society A*, 371, 20110534.
- Hyvärinen, A., & Ramkumar, P. (2013). Testing independent component patterns by inter-subject or inter-session consistency. *Frontiers in Human Neuroscience*, 7(March), 1–15. <https://doi.org/10.3389/fnhum.2013.00094>
- Jellinek, A. M., & DePaolo, D. J. (2003). A model for the origin of large silicic magma chambers: Precursors of caldera-forming eruptions. *Bulletin of Volcanology*, 65(5), 363–381.
- Jónsson, S. (2009). Stress interaction between magma accumulation and trapdoor faulting on Sierra Negra volcano, Galápagos. *Tectonophysics*, 471(1-2), 36–44. <https://doi.org/10.1016/j.tecto.2008.08.005>
- Jónsson, S., Zebker, H., & Amelung, F. (2005). On trapdoor faulting at Sierra Negra volcano, Galapagos. *Journal of Volcanology and Geothermal Research*, 144, 59–71. <https://doi.org/10.1016/j.jvolgeores.2004.11.029>
- Loughlin, S. C., Sparks, R. S. J., Brown, S. K., Jenkins, S. F., & Vye-Brown, C. (2015). *Global volcanic hazards and risk*. Cambridge: Cambridge University Press.
- Maaten, L. V. D., & Hinton, G. (2008). Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9, 2579–2605.
- Manga, M., & Brodsky, E. (2006). Seismic triggering of eruptions in the far field: Volcanoes and geysers. *Annual Review of Earth and Planetary Science*, 34, 263–291. <https://doi.org/10.1146/annurev.earth.34.031405.125125>
- McInnes, L., Healy, J., & Astels, S. (2017). hdbscan: Hierarchical density based clustering. *The Journal of Open Source Software*, 2, 11–12. <https://doi.org/10.21105/joss.00205>
- Meinecke, F., Ziehe, A., & Kawanabe, M. (2002). A resampling approach to estimate the stability of one-dimensional or multidimensional independent components. *IEEE Transactions on Biomedical Engineering*, 49(12), 1514–1525.
- Mogi, K. (1958). Relations between the eruptions of various volcanoes and the deformations of the ground surfaces around them. *Bulletin of the Earthquake Research Institute*, 36, 99–134. <https://doi.org/10.1016/j.eprl.2004.04.016>
- Okada (1985). Surface deformation due to shear and tensile faults in a half-space. *International Journal of Rock Mechanics and Mining Sciences Geomechanics Abstracts*, 75(4), 1135–1154. [https://doi.org/10.1016/0148-9062\(86\)90674-1](https://doi.org/10.1016/0148-9062(86)90674-1)
- Pagli, C., Sigmundsson, F., Arnadóttir, T., Einarsson, P., & Sturkell, E. (2006). Deflation of the askja volcanic system: Constraints on the deformation source from combined inversion of satellite radar interferograms and GPS measurements. *Journal of Volcanology and Geothermal Research*, 152(1-2), 97–108.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.
- Pinel, V., Jaupart, C., & Albino, F. (2010). On the relationship between cycles of eruptive activity and growth of a volcanic edifice. *Journal of Volcanology and Geothermal Research*, 194(4), 150–164. <https://doi.org/10.1016/j.jvolgeores.2010.05.006>
- Reynolds, R. W., & Geist, J. (1995). Petrology of lavas from Sierra Negra volcano, Isabela Island, Galapagos archipelago. *Journal of Geophysical Research*, 100, 24,537–24,553.

- Reynolds, R. W., Geist, D., & Kurz, M. (1995). Physical volcanology and structural development of Sierra Negra volcano, Isabela Island, Galapagos archipelago. *Geological Society of America Bulletin*, *107*(12), 1398–1410.
- Rodd, R. L., Lees, J. M., & Tepp, G. (2016). Three-dimensional attenuation model of Sierra Negra Volcano, Galápagos Archipelago. *Geophysical Research Letters*, *43*, 6259–6266. <https://doi.org/10.1002/2016GL069554>
- Schultz, R. (1995). Limits on strength and deformation properties of jointed basaltic rock masses. *Rock Mechanics and Rock Engineering*, *28*(1), 1–15.
- Sennert, S. K. e (2018a). Global volcanism program, 2018. Report on Sierra Negra (Ecuador). Smithsonian Institution and US Geological Survey, (Weekly Volcanic Activity 886 Report, 18 July-24 July 2018).
- Sennert, S. K. e (2018b). Global volcanism program, 2018. Report on Sierra Negra (Ecuador). Smithsonian Institution and US Geological Survey, (Weekly Volcanic Activity Report, 29 August-4 September 2018).
- Stock, M. J., Bagnardi, M., Neave, D. A., MacLennan, J., Bernard, B., Buisman, I., et al. (2018). Integrated petrological and geophysical constraints on magma system architecture in the western Galápagos Archipelago: Insights from Wolf volcano. *Geochemistry, Geophysics, Geosystems*, *19*, 4722–4743. <https://doi.org/10.1029/2018GC007936>
- Tepp, G., Ebinger, C. J., Ruiz, M., & Belachew, M. (2014). Imaging rapidly deforming ocean island volcanoes in the western Galápagos Archipelago, Ecuador. *Journal of Geophysical Research: Solid Earth*, *119*, 442–463. <https://doi.org/10.1002/2013JB010227>
- Touloukain, Y. S. (1989). *Physical properties of rocks and minerals, CINDAS Data Series on Materials Properties* (pp. 63–76). New York: Hemisphere Pub. Corp.
- Vasconez, B. F., Ramón, P., Hernandez, S., Hidalgo, S., & Bernard, B. (2018). The different character of the recent eruptions of Galápagos shield volcanoes (Ecuador): La Cumbre vs. Sierra Negra. *EarthArXiv Preprints* (pp. 1–11). <https://doi.org/10.31223/osf.io/t7f6k>
- Yun, S., Segall, P., & Zebker, H. (2006). Constraints on magma chamber geometry at Sierra Negra Volcano, Galapagos Islands, based on InSAR observations. *Journal of Volcanology and Geothermal Research*, *150*, 232–243. <https://doi.org/10.1016/j.jvolgeores.2005.07.009>