

# Automatic Detection of Volcanic Unrest Using Interferometric Synthetic Aperture Radar

Matthew Edward Gaddes

Submitted in accordance with the requirements for the degree of  
Doctor of Philosophy

The University of Leeds  
School of Earth and Environment

May 2019



The candidate confirms that the work submitted is their own, except where work which has formed part of jointly authored publications has been included. The contribution of the candidate and the other authors to this work has been explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others.

The work in Chapter 2 of the thesis has appeared in publication as follows:

**Gaddes, M.**, Hooper, A., Bagnardi, M., Inman, H., and 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, doi:10.1029/2018JB016210.

AH provided discussion on methodology/results/etc. MB processed the Sentinel-1 time series for Wolf volcano. HI calculated the kurtosis of GPS deformation data. FA processed the Sentinel-1 time series for Etna volcano. MG authored the paper and figures (which were improved with comments and suggestions from the authors), developed the software to perform the analysis, and performed the analysis detailed in the paper.

The work in Chapter 3 of the thesis has been accepted for publication pending minor revisions under the following title:

**Gaddes, M.**, Hooper, A., and Bagnardi, M. (2019), *Using machine learning to automatically detect signs of volcanic unrest in a time series of interferograms*, Journal of Geophysical Research: Solid Earth.

AH provided discussion on methodology/results/etc. MB processed the first half of the Sentinel-1 time series for Wolf volcano and the GBIS inversion. MG authored the paper and figures (which were improved with comments and suggestions from the authors), processed the second half of the Sentinel-1 time series, developed the ICASAR algorithm, developed the automatic detection algorithm, and applied the two algorithms to a variety of datasets.

The work in Chapter 4 of the thesis is a manuscript about to be submitted under the following title:

**Gaddes, M.**, Hooper, A., and Albino, F. (in prep), *Classifying and locating deformation in SAR interferograms using deep learning*.

AH provided discussion on methodology/results/etc. FA provided some of the Sentinel-1 data used. MG authored the paper and figures (which were improved with comments and suggestions from the authors), processed some of the Sentinel-1 data, built and trained the convolutional neural network, and applied it to a variety of data.

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement

Copyright © 2019 The University of Leeds and Matthew Edward Gaddes  
The right of Matthew Edward Gaddes to be identified as Author of this work has been asserted by him in accordance with the Copyright, Designs and Patents Act 1988.

# Acknowledgements

The thanks that I owe fall into two categories. The first are traditional, and involve all the people who have helped me reach this stage. Working with my main supervisor, Andy Hooper, has been the greatest pleasure of the PhD process, and I remain grateful for his years of expert supervision. Tim Wright has also provided insightful inputs whenever we met, and Marco Bagnardi's contribution to the Galapagos volcanoes work has been like that of a supervisor. Fabien Albino's help with Sentinel-1 data has also been invaluable, and I've always enjoyed the times that we've met to chat about current work. I'm also grateful to the staff and students of IGT for making Leeds an enjoyable place to work, and to the cohort who started in October 2014.

The second are less traditional, and involve all those who helped Jess and I build our way out of a rather unpleasant boat shaped corner, and ultimately allowed me to focus on this work. My thanks again go to Andy, Tim, and Sebastian Rost for their support of my request for a four month suspension, and their patience when I asked to extend it for another three months. Countless other people also helped with boat jobs, ranging from the pleasurable jobs like sailing the boat to a new location, to the less pleasurable jobs such as the removal of every inch of old paint on it. You are too many to thank here (but most of you live near Skipton), and your help was ultimately invaluable in allowing me to concentrate on my PhD.

I'd also like to thank my parents, Wendy and Mark, for all their help and support over the years, and to my sister Elena for providing some healthy PhD competition, despite starting life three years after me. And finally, I'd like to thank Jess for her years of endless patience and encouragement, and for having to go off to work whilst I kept being a student.



# Abstract

A diverse set of hazards are posed by the world's  $\sim 1500$  subaerial volcanoes, yet the majority of them remain unmonitored. Measurements of deformation provide a way to monitor volcanoes, and synthetic aperture RaDAR (SAR) provides a powerful tool to measure deformation at the majority of the world's subaerial volcanoes. This is due to recent changes in how regularly SAR data are acquired, how they are distributed to the scientific community, and how quickly they can be processed to create time series of interferograms. However, for interferometric SAR (InSAR) to be used to monitor the world's volcanoes, an algorithm is required to automatically detect signs of deformation-generating volcanic unrest in a time series of interferograms, as the volume of new interferograms produced each week precludes this task being achieved by human interpreters. In this thesis, I introduce two complementary methods that can be used to detect signs of volcanic unrest.

The first method centres on the use of blind signal separation (BSS) methods to isolate signals of geophysical interest from nuisance signals, such as those due to changes in the refractive index of the atmosphere between two SAR acquisitions. This is achieved through first comparing which of non-negative matrix factorisation (NMF), principal component analysis (PCA), and independent component analysis (ICA) are best suited for solving BSS problems involving time series of InSAR data, and how InSAR data should best be arranged for its use with these methods. I find that NMF can be used with InSAR data, providing the time series is formatted in a novel way that reduces the likelihood of any pixels having negative values. However, when NMF, PCA, and ICA are applied to a set of synthetic data, I find that the most accurate recovery of signals of interest is achieved when ICA is set to recover spatially independent sources (termed sICA). I find that the best results are produced by sICA when interferograms are ordered as a simple "daisy chain" of short temporal baselines, and when sICA is set to recover around 1 – 3 more sources than were thought to have contributed to the time series. However, I also show that in cases such as deformation centred under a stratovolcano, the overlapping nature of a topographically correlated atmospheric phase screen (APS) signal and a deformation signal produces a pair of signals that are no longer spatially statistically independent, and so cannot be recovered accurately by

sICA.

To validate these results, I apply sICA to a time series of Sentinel-1 interferograms that span the 2015 eruption of Wolf volcano (Galapagos archipelago, Ecuador) and automatically isolate three signals of geophysical interest, which I validate by comparing with the results of other studies. I also apply the sICA algorithm to a time series of interferograms that image Mt Etna, and through isolating signals that are likely to be due to instability of the east flank of the volcano, show that the method can be applied to stratovolcanoes to recover useful signals. Utilising the ability of sICA to isolate signals of interest, I introduce a prototype detection algorithm that tracks changes in the behaviour of a subaerial volcano, and show that it could have been used to detect the onset of the 2015 eruption of Wolf.

However, for use in a detection algorithm that is to be applied globally, the signals recovered by sICA cannot be manually validated through comparison with other studies. Therefore, I seek to incorporate a module into my detection algorithm that is able to quantify the significance of the sources recovered by sICA. I achieve this through extensively modernising the ICASO algorithm to create a new algorithm, ICASAR, that is optimised for use with InSAR time series. This algorithm allows me to assess the significance of signals recovered by sICA at a given volcano, and to then prioritise the tracking of any changes they exhibit when they are used in my detection algorithm.

To further develop the detection algorithm, I create two synthetic time series that characterise the different types of unrest that could occur at a volcanic centre. The first features the introduction of a new signal, and my algorithm is able to detect when this signal enters the time series by tracking how well the baseline sources are able to fit new interferograms. The second features the change in rate of a signal that was present during the baseline stage, and my algorithm is able to detect when this change in rate occurs by tracking how sources recovered from the baseline data are used through time. To further test the algorithm, I extended the Sentinel-1 time series I used to study the 2015 eruption of Wolf to include the 2018 eruption of Sierra Negra, and I find that my algorithm is able to detect the increase in inflation that precedes the eruption, and the eruption itself.

I also perform a small study into the pre-eruptive inflation seen at Sierra Negra using the deformation signal and its time history that are outputted by ICASAR. A Bayesian inversion is performed using the GBIS software package, in which the inflation signal is modelled as a horizontal rectangular dislocation with variable opening and uniform overpressure. Coupled with the time history of the inflation signal provided by ICASAR, this allows me to determine the temporal evolution of the pre-eruptive overpressure since the beginning of the Sentinel-1 time series in 2014. To extend this back to the end of the previous eruption in 2005, I use GPS data that spans the entire

interruptive period. I find that the total interruptive pressure change is  $\sim 13.5$  MPa, which is significantly larger than the values required for tensile failure of an elastic medium overlying an inflating body. I conclude that it is likely that one or more processes occurred to reduce the overpressure within the sill, and that the change in rate of inflation prior to the final failure of the sill is unlikely to be coincidental.

The second method I develop to detect volcanic deformation in a time series of interferograms uses a convolutional neural network (CNN) to classify and locate deformation signals as each new interferogram is added to the time series. I achieve this through building a model that uses the five convolutional blocks of a previously state-of-the-art classification and localisation model, VGG16, but incorporates a classification output/head, and a localisation output/head. In order to train the model, I perform transfer learning and utilise the weights made freely available for the convolutional blocks of a version of VGG16 that was trained to classify natural images. I then synthesise a set of training data, but find that better performance is achieved on a testing set of Sentinel-1 interferograms when the model is trained with a mixture of both synthetic and real data. I conclude that CNNs can be built that are able to differentiate between different styles of volcanic deformation, and that they can perform localisation by globally reasoning with a  $224 \times 224$  pixel interferogram without the need for a sliding window approach.

The results I present in this thesis show that many machine learning methods can be applied to both time series of interferograms, and individual interferograms. sICA provides a powerful tool to separate some geophysical signals from atmospheric ones, and the ICASAR algorithm that I develop allows a user to evaluate the significance of the results provided by sICA. I incorporate these methods into an deformation detection algorithm, and show that this could be used to detect several types of volcanic unrest using data produced by the latest generation of SAR satellites. Additionally, the CNN I develop is able to differentiate between deformation signals in a single interferogram, and provides a complementary way to monitor volcanoes using InSAR.



# Contents

<b>List of Figures</b>	<b>xv</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Volcano monitoring using InSAR . . . . .	1
1.1.1 Volcano monitoring . . . . .	1
1.1.2 InSAR for hazard monitoring . . . . .	6
1.1.3 Characteristics of displacement signals in interferograms of volcanic centres . . . . .	10
1.2 Machine learning . . . . .	12
1.2.1 Unsupervised learning . . . . .	13
1.2.2 Supervised learning . . . . .	15
1.3 Aims and Objectives . . . . .	19
1.4 Thesis outline . . . . .	20
<b>2 Blind Signal Separation Methods for InSAR: The Potential to Automatically Detect and Monitor Signals of Volcanic Deformation</b>	<b>29</b>
2.1 Introduction . . . . .	30
2.2 Blind Signal Separation . . . . .	32
2.2.1 Interferograms as mixtures of signals . . . . .	32
2.2.2 Organising and synthesising data . . . . .	34
2.2.3 Principal Component Analysis . . . . .	37
2.2.4 Independent Component Analysis . . . . .	40
2.2.5 PCA and ICA example . . . . .	43
2.2.6 Non-negative Matrix Factorisation . . . . .	44
2.3 Application of NMF to InSAR time series . . . . .	47
2.4 Comparison of PCA, ICA and NMF . . . . .	47
2.4.1 Single master vs daisy chain time series . . . . .	50
2.4.2 Number of sources sought . . . . .	53
2.4.3 Length of time series . . . . .	53
2.4.4 Strength of turbulent atmosphere contribution . . . . .	55
2.4.5 Spatial independence of sources . . . . .	56

2.4.6	Comparison conclusions and method limitations . . . . .	57
2.5	Application to real data . . . . .	59
2.5.1	Wolf volcano . . . . .	59
2.5.2	Etna volcano . . . . .	60
2.6	Use of sICA in a monitoring algorithm . . . . .	64
2.7	Discussion . . . . .	67
2.8	Conclusion . . . . .	68
2.9	Acknowledgments . . . . .	68
<b>3</b>	<b>Using Machine Learning to Automatically Detect Volcanic Unrest in a Time Series of Interferograms</b>	<b>75</b>
3.1	Introduction . . . . .	76
3.2	Methods: Detection algorithm . . . . .	78
3.2.1	Robust recovery of latent sources . . . . .	79
3.2.2	Characterisation of the baseline data . . . . .	83
3.2.3	Ingestion of new interferograms . . . . .	85
3.3	Application to Sierra Negra . . . . .	86
3.4	Modelling overpressure before the 2018 Sierra Negra eruption . . . . .	93
3.5	Discussion . . . . .	97
3.5.1	Detection Algorithm . . . . .	97
3.5.2	Overpressure before the 2018 Sierra Negra Eruption . . . . .	100
3.6	Conclusion . . . . .	101
3.7	Acknowledgments . . . . .	102
<b>4</b>	<b>Classifying and Locating Deformation in SAR Interferograms Using Deep Learning</b>	<b>107</b>
4.1	Introduction . . . . .	108
4.2	Classification with different data formats . . . . .	110
4.3	Classification and localisation . . . . .	113
4.3.1	Using synthetic data . . . . .	113
4.3.2	Application to real data . . . . .	118
4.3.3	Using Sentinel-1 data . . . . .	121
4.4	Discussion . . . . .	123
4.5	Conclusion . . . . .	127
4.6	Acknowledgments . . . . .	127
<b>5</b>	<b>Discussion and Conclusions</b>	<b>131</b>
5.1	Project aims and key findings . . . . .	131
5.1.1	Characterising baseline signals . . . . .	131
5.1.2	Detecting and displaying deviations from baseline behaviour . . .	133
5.1.3	Complementary monitoring algorithm . . . . .	134

---

5.1.4	Study of a volcanic process . . . . .	135
5.1.5	Further Remarks . . . . .	136
5.1.6	Key Findings . . . . .	137
5.2	Future work . . . . .	137
5.3	Concluding remarks . . . . .	139



# List of Figures

1.1	Distribution of human population and Holocene volcanoes . . . . .	3
1.2	Deformation measured using GPS . . . . .	4
1.3	Deformation measured using InSAR . . . . .	5
1.4	Deformation as an indicator of eruption . . . . .	6
1.5	SAR amplitude image . . . . .	7
1.6	LiCSAR processing areas . . . . .	9
1.7	Example of data used . . . . .	11
1.8	Pre-eruptive inflation at Sinabung . . . . .	12
1.9	Example of labelled data . . . . .	13
1.10	Blind signal separation with faces . . . . .	15
1.11	Non-negative matrix factorisation with faces . . . . .	16
1.12	Perceptron and neural network . . . . .	17
1.13	Linearly separable functions . . . . .	18
1.14	Neuron activation functions . . . . .	18
1.15	VGG16 . . . . .	19
2.1	Linear mixing for spatially and temporally organised data . . . . .	35
2.2	Sources and estimated PDFs . . . . .	38
2.3	Spatial sources and BSS . . . . .	45
2.4	Temporal sources and BSS . . . . .	46
2.5	NMF scenarios . . . . .	48
2.6	NMF applied to spatial data . . . . .	49
2.7	Daisy chain vs. single master data . . . . .	51
2.8	PCA and ICA applied to daisy chain and single master data . . . . .	52
2.9	PCA and ICA applied to recover different numbers of sources . . . . .	54
2.10	PCA and ICA applied to time series of different lengths . . . . .	55
2.11	PCA and ICA applied to time series with different atmospheric strengths . . . . .	56
2.12	PCA and ICA applied to data with non-spatially independent sources . . . . .	58
2.13	Galapagos overview . . . . .	60
2.14	Wolf volcano interferograms . . . . .	61
2.15	sICA application to Wolf volcano interferograms . . . . .	62

2.16	sICA application to Etna interferograms . . . . .	65
2.17	Prototype automatic detection algorithm . . . . .	66
3.1	ICASAR Algorithm . . . . .	82
3.2	Automatic detection with synthetic data: 1 . . . . .	87
3.3	Automatic detection with synthetic data: 2 . . . . .	88
3.4	Galapagos overview . . . . .	91
3.5	ICASAR algorithm results . . . . .	92
3.6	ICASAR algorithm cluster centrotypes . . . . .	93
3.7	Automatic detection with Sierra Negra data . . . . .	94
3.8	Modelled deformation signal . . . . .	97
3.9	Temporal evolution of the modelled source . . . . .	98
4.1	Hierarchy of deformation signals . . . . .	110
4.2	Synthetic training data . . . . .	114
4.3	Three channel InSAR data . . . . .	115
4.4	Three channel classification accuracy . . . . .	116
4.5	Classification and localisation CNN . . . . .	119
4.6	Two stage CNN training . . . . .	120
4.7	Results of CNN 1: synthetic data . . . . .	120
4.8	Results of CNN 1: real data . . . . .	122
4.9	Results of CNN 2: real data . . . . .	124
5.1	YOLO9000 . . . . .	140
5.2	GPS time series . . . . .	147
5.3	Mount Etna interferograms . . . . .	148
5.4	Sierra Negra interferograms . . . . .	152
5.5	Sierra Negra GPS data . . . . .	153
5.6	Automatic detection algorithm: Low magnitude signal . . . . .	154
5.7	Bayesian inversion parameters . . . . .	155

# Nomenclature

## List of acronyms

BSS	Blind Signal Separation
CNN	Convolutional Neural Network
DEM	Digital Elevation Model
ESA	European Space Agency
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
ICA	Independent component analysis
InSAR	Interferometric synthetic aperture RaDAR
NN	Neural Network
NMF	Non-negative matrix factorisation
PCA	Principal component analysis
PDF	Probability density function
SAR	Synthetic aperture RaDAR
SRTM	Shuttle RaDAR Topography Mission
SVM	Support vector machines

## List of symbols

<b>W</b>	Unmixing matrix
<b>A</b>	Mixing matrix
<b>S</b>	Sources (as row vectors)
<b>X</b>	Mixtures (as row vectors)



# Chapter 1

## Introduction

This thesis seeks to develop machine learning methods that can be used to automatically detect signs of deformation-generating volcanic unrest in a time series of interferograms. Consequently, the remainder of this chapter is divided into Section 1.1, which introduces volcano monitoring using InSAR, and Section 1.2, which introduces machine learning.

### 1.1 Volcano monitoring using InSAR

#### 1.1.1 Volcano monitoring

The world's  $\sim 1500$  subaerial volcanoes (Siebert and Simkin, 2013) pose a variety of hazards to both those living in the immediate vicinity of a volcano, and those living further afield. The hazards posed by volcanoes are diverse, and whilst they are commonly associated with eruptions, also include events such as highly destructive flank collapses (Ramalho et al., 2015), and secondary lahars (Machado et al., 2015). Eruptive hazards have a wide range of spatial impacts, with processes such as lahars (Smith and Lowe, 1991), pyroclastic density currents (Druitt, 1998), ash falls (Bond and Sparks, 1976), lava flows (Bagnardi et al., 2016), and gas clouds (Othman-Chandev, 1987) causing loss of life and damage to property in the area surrounding a volcano. At larger scales, ash produced by explosive eruptions can close airspace to aviation (Sigmundsson et al., 2010), and gas clouds can drift across continents (Carboni et al., 2012), whilst large eruptions can have impacts on global scales (Schmidt et al., 2012).

However, volcanoes also bring benefits to humanity, such as the nutrient rich soils which are produced through weathering of the edifices that subaerial volcanoes commonly form (Small and Naumann, 2001). As a result, humans can be drawn to living on, or near, volcanic centres, and there are currently 457 Holocene volcanoes with over one million people living within a 100km radius of them (Small and Naumann, 2001). For a more general overview, Figure 1.1 shows the global distribution of Holocene volca-

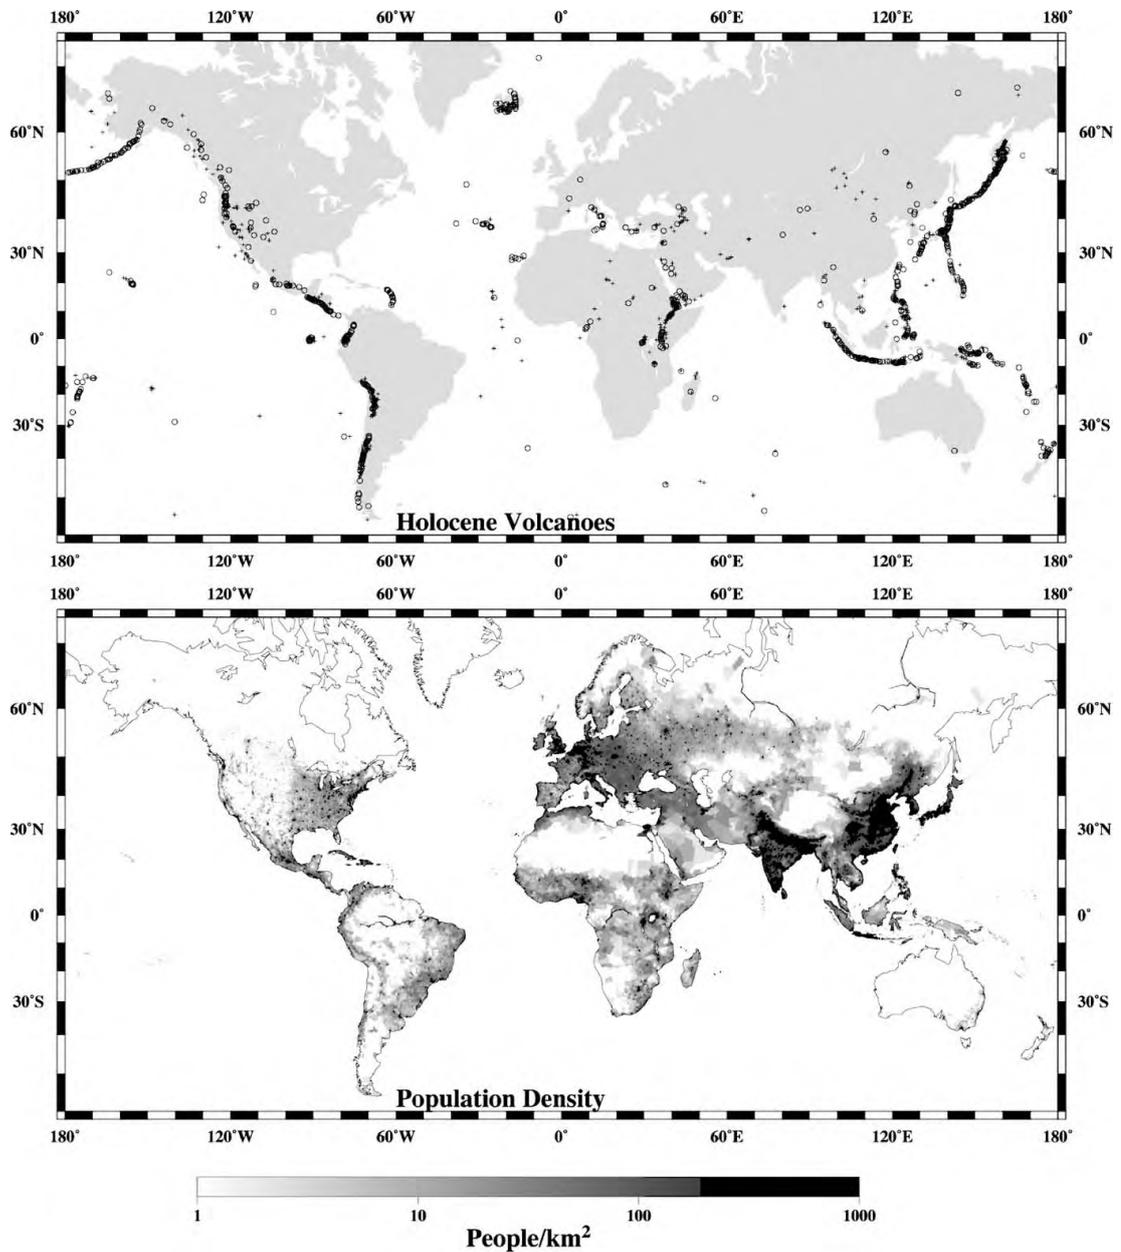
noes and human population. Since the 17th century, this proximity has resulted in over 300,000 deaths due to disasters caused by volcanoes (Small and Naumann, 2001; Tilling, 2008). Eruptions have also created significant economic costs, such as the airspace closure associated with the Eyjafjallajökull eruption which created a multitude of costs, including those to the aviation industry which reached approximately \$250 million per day (Gudmundsson, 2012).

Due to the combination of human and economic costs from volcanic hazards, humans have sought to constrain the likeliness of damage causing events. A requisite for this has been systematic monitoring of volcanoes, and some have now been monitored for over 100 years (Sparks et al., 2012). In addition to improving forecasts of events such as volcanic eruptions, systematic monitoring has also provided data for humanity's more philosophic quest to understand the processes occurring at volcanoes.

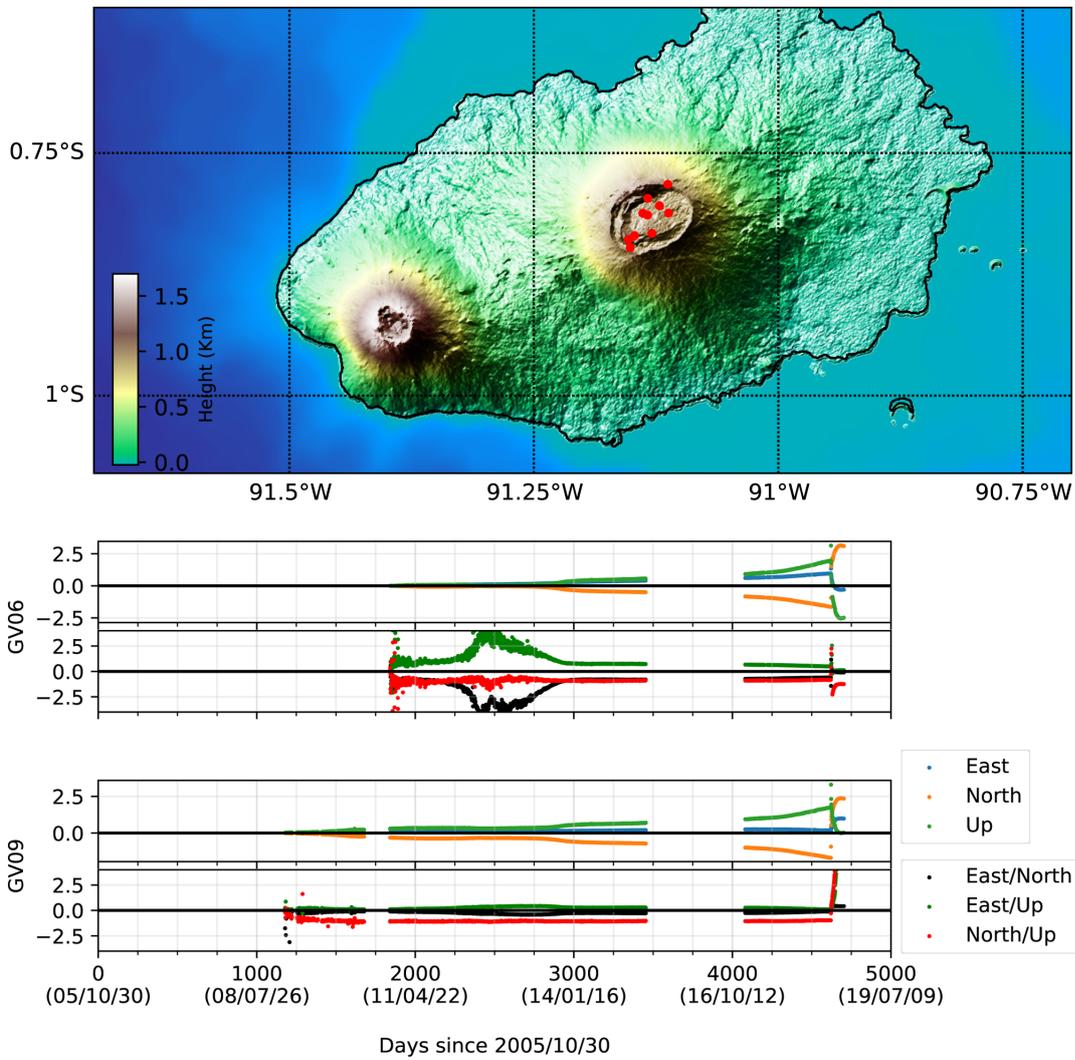
Monitoring is commonly performed using data gathered from approaches such as levelling, tiltmeters, seismology, strain gauges, satellite derived thermal imagery, satellite derived gas emissions, satellite derived deformation measurements, and geochemical analysis of fumarolic gases (Sparks et al. (2012), and references therein). Changes that can be detected using these methods or instruments can be indicative of a volcano entering a period of unrest, and eruptions are almost always preceded by a period of volcanic unrest (Tilling, 2008). Therefore, through monitoring a volcano, we hope to ascertain what the baseline behaviour of a volcano is, and to then detect if it has entered a period of unrest. Knowledge of the changes defining the period of unrest can then be used to forecast eruptions, and hopefully reduce the human and economic losses that may result from one. Additionally, data recorded during periods of unrest can also be used to test hypotheses on the physical processes occurring at volcanoes.

However, despite the motivation to forecast volcanic hazards creating a requirement for systematic monitoring of subaerial volcanoes, the majority remain unmonitored (Sparks et al., 2012). This is likely to be due to a combination of factors, of which the cost of extending the monitoring methods mentioned in the previous paragraphs is likely to be key. At present, most systematic measurement is performed using seismometers, Global Positioning System (GPS) receivers, and tiltmeters (Sparks et al., 2012), and the extension of these monitoring techniques requires the procurement of new instruments, their installation and upkeep, and the expertise required to interpret the data they produce. Figure 1.2 shows a GPS installation at Sierra Negra (Galapagos Archipelago, Ecuador), which produces relatively temporally dense but spatially sparse data.

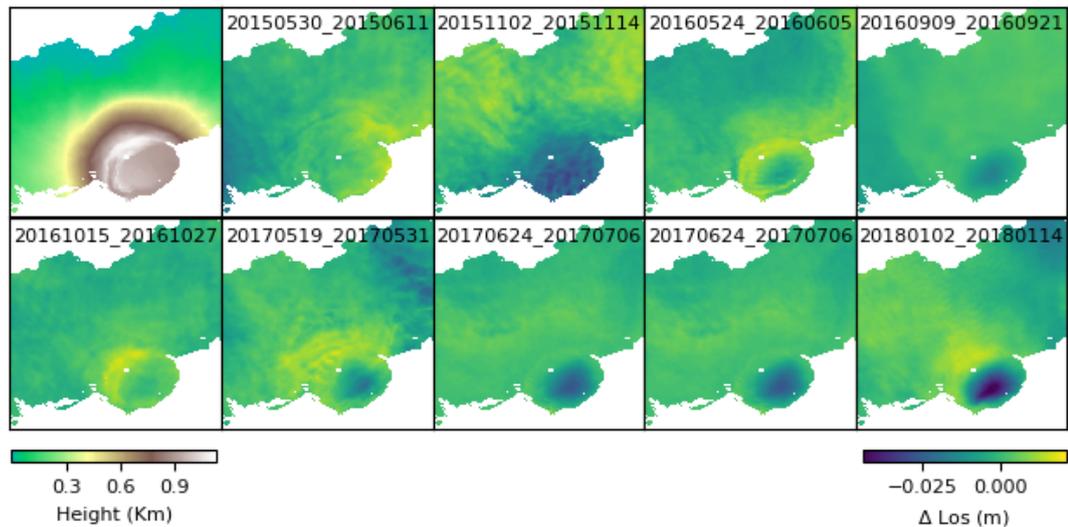
In contrast to the incremental costs of expanding current common monitoring methods to new volcanoes, deformation measurements produced by the latest generation of synthetic aperture RaDAR (SAR) satellites now provide a novel opportunity to perform near-global monitoring of subaerial volcanoes. Interferometric SAR (InSAR) has



**Figure 1.1:** Distribution of human population and Holocene volcanoes from Small and Naumann (2001). Volcanoes that have been active during historical times are marked as circles, whilst the remaining volcanoes are marked as crosses. A logarithmic scale is used for population. Note the coexistence of high population densities and active volcanoes in areas such as Indonesia and Japan.



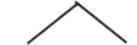
**Figure 1.2:** Deformation measured using GPS at Sierra Negra (Galapagos Archipelago, Ecuador). Red dots indicate the locations of the 10 GPS stations, whilst the lower subplots (GV06 and GV09) show the East/North/Up components of the deformation measured, and the ratios of each combination of components. Despite several times during which there are no data, the majority of the data is temporally dense and shows deformation that can be interpreted as pre-eruptive uplift.



**Figure 1.3:** Deformation measured using InSAR at Sierra Negra (Galapagos Archipelago, Ecuador). The first tile shows the digital elevation model (DEM), with water masked in the top left corner, and vegetated areas with poor coherence masked in the lower right corner. The remaining nine tiles show the change in line of sight (LOS) for a random selection of 12 day interferograms, with master and slave dates shown in the format YYYYMMDD. The signal contained in the later interferograms can be interpreted as pre-eruptive uplift (i.e. ground to satellite shortening).

been used as a tool for retrospective analysis of volcano deformation since the mid 1990s (Massonnet et al., 1995), and has been used to measure signals that have been attributed to processes such as pre-eruptive inflation of a magma chamber (e.g. Mt. Sinabung - Chaussard et al. (2013)), subsidence due to flank loading by new material (e.g. Arenal and Santiaguito - Ebmeier et al. (2013)), and subsidence due to cooling of a magma body below a volcano (e.g. Torfajökull - Ofeigsson et al. (2011)). InSAR also provides other advantages over methods that require the installation of sensors on a volcano, such as its ability to measure deformation in areas that are remote or potentially dangerous (e.g. Nyiragongo, in The Democratic Republic of the Congo (Wauthier et al., 2012)). Figure 1.3 shows a selection of interferograms that image Sierra Negra and, in contrast to the GPS data shown in Figure 1.2, provides data that is relatively spatially dense but temporally sparse. The expansion of these measurements to cover all the volcanoes of the Galapagos Archipelago is relatively straightforward, as a single Sentinel-1 track images the majority of the main island, and contrasts with performing the same expansion using GPS, which would require the installation of tens of receivers.

An insight into the expected results of monitoring a large number of volcanoes with InSAR is provided by Biggs et al. (2014). Here, the authors completed an extensive literature search of measurements of volcanic deformation made using InSAR and compared this with whether the volcano had erupted or not. The truth table in Figure 1.4 shows that 46% of volcanoes that deformed also erupted. However, of the 34 that

Systematic Coverage	Erupted	Non-Erupted
Deformed 	$DE$ 25 True positive	$DE^{\bar{}}$ 29 False positive
Non-deformed 	$\bar{DE}$ 9 False negative	$\bar{DE}^{\bar{}}$ 135 True negative

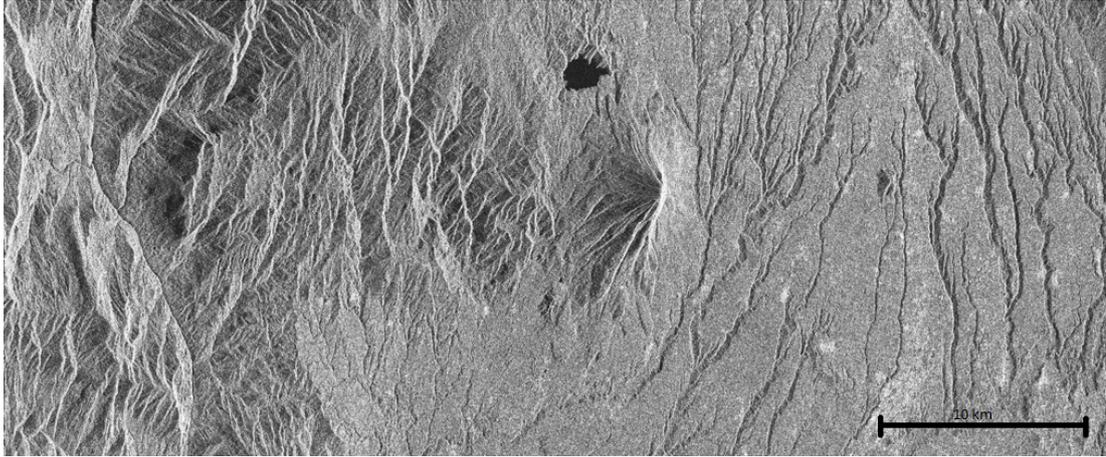
**Figure 1.4:** “Contingency table linking volcanoes that deformed and erupted”, reproduced from Biggs et al. (2014). Two important cases from the table are that very few volcanoes that do not deform also erupted ( $9/(9 + 135)$ ), and that approximately half ( $25/(25 + 29)$ ) of volcanoes that deformed also erupted.

erupted, only 25 also deformed, leaving 9 (or  $\sim 25\%$ ) volcanoes that did not show deformation recorded with InSAR but did also go on to erupt (i.e. false negative results). False negatives are of particular concern as these are cases in which measurements made with InSAR did not record any deformation and would not have flagged the volcano as having entered a period of unrest, yet the volcanoes did erupt. However, when using the lack of deformation as a predictor (in the sense of predicting a result, rather than implying eruption followed deformation) for a lack of eruption, these results show that it is correct in 94% of cases, for a specific time interval. Considering these statistics, it can be argued that if a volcano is monitored solely using InSAR based deformation measurements, the presence of a deformation signal is not a “strong diagnostic of imminent eruption” (Biggs et al., 2014). However, this study was performed before Sentinel-1 data were available, and the reduced revisit time that these satellites provide may strengthen the link between deformation observed with InSAR and eruptions. Additionally, as many of the world’s volcanoes are currently not monitored, the rate of false positives suggests that monitoring using deformation measurements derived from InSAR may still provide valuable constraints on the behaviour of many volcanoes.

### 1.1.2 InSAR for hazard monitoring

In the previous section, the arguments for monitoring all of the world’s subaerial volcanoes were described, and the potential for InSAR to be used to achieve this introduced. In this section, the recent changes that have allowed InSAR to evolve from a retrospective tool are discussed, and a brief introduction to the method given.

Individual SAR images consist of complex numbers, detailing the phase and amplitude of the return from a given area of land (Hanssen, 2001). Whilst the phase appears random, the amplitude information can be viewed as a grayscale image of the Earth, though the oblique look angle of SAR satellites causes features to be slightly distorted.



**Figure 1.5:** SAR amplitude image of Mt. Sinabung (Indonesia) acquired by the Radarsat-2 satellite. Note the asymmetric nature of the volcano’s edifice is a feature of the right looking nature of the SAR antenna used to create this image during a descending (approximately north to south) portion of the orbit. Data provided as part of the “Sinabung Supersite” by UNAVCO.

As SAR images are produced through active illumination using electromagnetic waves that can penetrate cloud, SAR images can be acquired at night and in weather conditions that preclude the use of optical imagery. Figure 1.5 shows an example of a SAR amplitude image of a stratovolcano, and images like this can be used to image dome growth (Wang et al., 2015), and estimate the volume of lava flows (Arnold et al., 2017).

Pairs of SAR images can be combined through differencing the phase values of each pixel to create an interferogram, which commonly contains a combination of signals:

$$\varphi = W \{ \varphi_{def} + \varphi_{orb} + \varphi_{atm} + \Delta\varphi_{\theta} + \varphi_N \} \quad (1.1)$$

where  $\varphi_{def}$  is the phase change due to deformation of the ground surface,  $\varphi_{orb}$  is the phase due to errors in the location of the satellite at each acquisition,  $\varphi_{atm}$  is the phase change due to changes in the atmospheric delay,  $\Delta\varphi_{\theta}$  is the phase due to misestimation of the look angle,  $\varphi_N$  is the phase noise, and  $W$  is a wrapping operator that results in the phase lying between  $-\pi$  and  $\pi$  (Hooper et al., 2012). Deformation is generally the signal of interest for geophysical applications, and a variety of methods exist to reduce the contribution from other signals, such as the TRAIN toolbox that can be used to reduce the contribution from changes in the atmosphere (Bekaert et al., 2015). Additionally, the combined signal expressed in an interferogram is *modulo*  $2\pi$ , or more commonly termed “wrapped”. Unwrapping solves for the phase ambiguity created each time wrapped phase progresses above  $\pi$ , and algorithms exist to do this (Chen and Lapusta, 2009; Hooper, 2010). The deformation signals within interferograms are also relative measurements, requiring a reference area to be set outside the deforming region of interest, and due to the look angle of the SAR satellites presently operating,

they remain relatively insensitive to north-south movement (Wright, 2004).

When two SAR images fall either side of an event of geophysical interest, the resulting interferogram that spans these dates can be used to infer details about the event. This was originally applied to tectonic events (e.g. Massonet et al. (1993)), but has also been applied to volcanic events (e.g. González et al. (2015)). However, through the creation of multiple interferograms, time series can be produced to study additional processes, such as continued uplift prior to an eruption (Chaussard and Amelung, 2012), or those such as interseismic strain which feature low magnitude signals (Wright et al., 2001). Approaches to the creation of time series fall into two broad categories. Persistent scatter (PS) methods (Ferretti et al., 2001; Hooper et al., 2007) seek pixels that are have a RaDAR return dominated by a single scatterer that remains stable on a time scale of years, and create interferograms relative to a single master. In contrast, short baseline methods (Berardino et al., 2002; Hooper, 2008) seek to minimise the decorrelation of pixels that do not feature a dominant scatterer through producing pairs of interferograms with either short temporal baselines, or small offsets between the satellites at the time of each acquisition. Signal processing methods have also been successfully applied to the resulting time series in order to isolate signals of geophysical interest, such as principal component analysis (Chaussard et al., 2014), and independent component analysis (Ebmeier, 2016).

Whilst these methods have allowed InSAR to flourish as a tool for retrospective analysis, it has not been widely suitable for hazard monitoring. Obstacles included the cost of acquiring data, the latency in the availability of data, the lack of data over some regions of geophysical interest, the long revisit times between acquisitions, and the time consuming nature of the creation of interferograms. However, the latest generation of SAR satellites such as the European Space Agency’s Sentinel-1 constellation provide free and open data that are made available quickly due to the creation of the European Data Relay System, and the 250km wide swaths produced by “Terrain observation with progressive scans” (TOPS, De Zan and Guarnieri (2006)) allows for the revisit time to be reduced to around 12 days. These advances have been mirrored by the creation of facilities to automatically create interferograms from the SAR images (e.g. González et al. (2016) and Meyer et al. (2016)), and in refinements to the algorithms used to create interferograms (e.g. Spaans and Hooper (2016)). Figure 1.6 shows the areas that are covered by interferograms created automatically by the LiCSAR processor (González et al., 2016), and the length of the time series at each of these locations. In addition to producing time series of unwrapped interferograms, these facilities also produce associated products, such as coherence images, that have previously been used to perform tasks such as mapping fresh lava flows (Ebmeier et al., 2018).

Of the signals that combine to form an interferogram (Equation 2.1), the atmospheric component is commonly challenging to separate from deformation signals



**Figure 1.6:** Areas for which Sentinel-1 interferograms are automatically created by the LiC-SAR processor. Cool colours indicate areas with relatively short time series, whilst warmer colours indicate longer time series, such as over Europe and Turkey. Figure modified from the COMET-LiCS Sentinel-1 InSAR portal (González et al., 2016).

(Hooper et al., 2012). When producing interferograms, it is initially assumed that the electromagnetic waves propagate at uniform velocity between the satellite and ground. However, for most geodetic InSAR applications, microwaves must travel through the Earth’s atmosphere and, consequently, the assumption of uniform wave velocity is broken. For C-band satellites, the effect of the ionosphere can be regarded as minimal (Hooper et al., 2012), and the delay to the electromagnetic waves depends on their angle of inclination, and the integral of the refractive index with height, which in turn depends on pressure (P), temperature (T), and the partial water vapour pressure (e):

$$N = k_1 \frac{P}{T} + \left( k_2' \frac{e}{T} + k_3 \frac{e}{T^2} \right) \quad (1.2)$$

and three coefficients,  $k_1$ ,  $k_2'$  and  $k_3$  (Hanssen, 2001). As the first term is not influenced by water vapour, it is often termed the dry or hydrostatic delay, whilst the later two combine to be termed the wet delay. The hydrostatic delay is several times larger than the wet, but, as temperature and pressure vary smoothly throughout the atmosphere (in contrast to partial water vapour pressure), differential interferograms often contain signals with length scales of the order of 10 km due to changes in the wet delay. Additionally, changes in the dry delay or changes in vertically stratified water vapour can produce signals in differential interferograms that are correlated with topography (Pavez et al., 2006).

Over the last 20 years, a suite of methods have evolved to correct for atmospheric delays. GNSS receivers have been used to measure the total delay above them (Williams et al., 1998), weather models have been used to calculate the parameters required to cal-

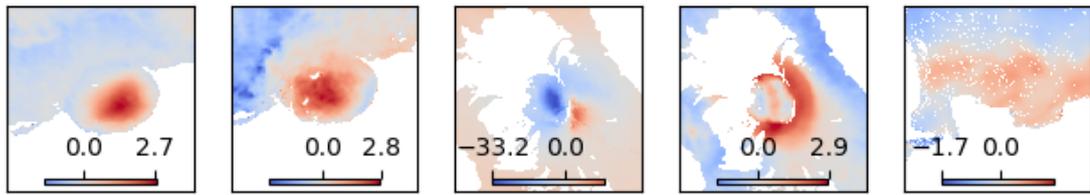
culate  $N$  (Wadge et al., 2002), and other spectrometers (such as MERIS and MODIS) have been used to measure the wet component of the delay (Walters et al., 2013). Methods have also been designed that do not require additional data and rely solely on the interferometric phase. Topographically correlated delays have been corrected using linear methods (Wicks, 2002; Elliott et al., 2008) and, more recently, with a variable power law relationship (Bekaert et al., 2015; Shen et al., 2019). Of particular importance to volcanic studies are topographically correlated delays as, due to the conic shape of many volcanoes, the delay may appear in a pattern very similar to that expected from a change in volume of a magma chamber below the volcano. Consequently, they are likely to be one of the largest challenges in designing an algorithm to detect deformation-generating volcanic unrest.

To summarise, the latest generation of SAR satellites produce imagery that spans the majority of the world’s active volcanoes, and is made freely available to the scientific community in a timely manner. These amplitude images can be useful for volcano monitoring, but the subtle deformation signals that can be measured through the creation of time series of interferograms are most applicable for volcano monitoring due to the preliminary results on the link between deformation measured using InSAR and volcanic eruptions (Biggs et al., 2014). However, future work may also be able to utilise other products associated with SAR imagery, such as coherence maps, and changes in SAR amplitude. In the remainder of this chapter, I provide a brief introduction to machine learning, as this field provides many methods that may be suitable for use in an detection algorithm.

### 1.1.3 Characteristics of displacement signals in interferograms of volcanic centres

Measurements of ground deformation provided by satellite based InSAR are traditionally spatially dense but temporally poor. The reduction in revisit time provided by the Sentinel-1 satellites has increased the temporal density of measurements, with 6 or 12 day revisit times for areas of geophysical interest. Deformation signals can be divided based on whether they are contained within a single revisit pair, or if they span multiple pairs. In the case that they span multiple pairs, time series methods can be used to determine rates of deformation, but in the case that the deformation event is contained within a single revisit period, the calculated deformation rate is likely to be an underestimate.

Additionally, for a signal to be detectable, it must be visible above noise contained within an interferogram, which was previously described in Equation 2.1 and the accompanying text. Signals due to the turbulent atmosphere commonly have magnitudes of the order 1 – 2 cm (Lohman and Simons, 2005), but due to their spatially random nature, the stacking of multiple interferograms to create time series can minimise their

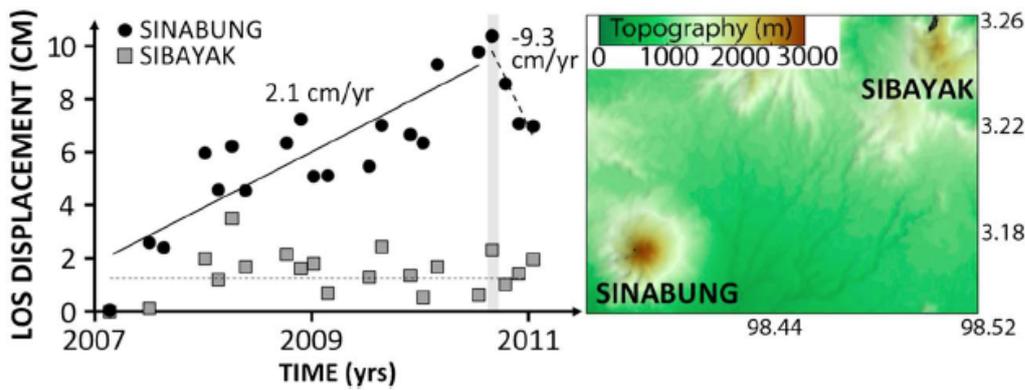


**Figure 1.7:** Examples of Sentinel-1 data used in this study, showing line of sight displacement in cm. Left to right, (1) Sierra Negra (Galapagos Archipelago) with the vegetated southern side of the volcano masked showing pre-eruptive uplift of  $\sim 3$  cm in 12 days. (2) Sierra Negra, as per (1), interpreted as showing a similar deformation signal to (1), but due to a stronger atmospheric signal, the deformation is less clear. (3) Wolf (Galapagos Archipelago) with the vegetated western side of the volcano masked, showing co-eruptive deformation captured by a 12 day interferogram spanning the start of the 2015 eruption. (4) Wolf, as per 3, showing no obvious deformation, but a signal that can be interpreted as a topographically correlated APS can be seen to show the “upturned soup bowl” shape of the volcano. (5) Campi Flegrei (Italy) with areas of low coherence masked and a signal that can be interpreted as a turbulent APS visible, but deformation is not visible in this 12 day interferogram.

impact (e.g. Wright et al. (2001)). However, of particular importance at volcanic centres are atmospheric signals that are topographically correlated, which, due to their persistent location, cannot be removed through a simple stacking approach, yet may not be independent of deformation. In the case of stratovolcanoes, delays of  $\sim 9$  rad/km over  $\sim 4000$  m of topography can create an apparent deformation signal of 17 cm (Pinel et al., 2011).

For a deformation signal to be detectable in a single 12 day interferogram, a signal of at least a magnitude greater than the turbulent APS is required. Examples of these include the signal associated with the 2015 eruption of Calbuco which featured  $\sim 12$  cm of deformation (Delgado et al., 2017), and the 2014/15 eruption of Fogo, which featured  $\sim 25$  cm of deformation associated with the intrusion of a dyke (González et al., 2015). Figure 1.7 also shows examples of single interferogram pairs that are used within this thesis. Subplots (1) and (2) show 12 day interferograms of Sierra Negra (Galapagos Archipelago), in which a deformation signal of magnitude  $\sim 3$  cm can be seen both clearly (1), and more subtly in (2) due to what is likely to be a strong turbulent APS. Subplots (3) and (4) show 12 day interferograms of Wolf (Galapagos Archipelago), in which (3) shows a larger signal ( $\sim 30$  cm of motion away from the satellite) associated with the 2015 eruption, whilst (4) shows what is likely to be a strong topographically correlated APS, producing  $\sim 3$  cm of apparent deformation.

The deformation of  $\sim 3$  cm/12 days measured in the Sentinel-1 interferograms of Sierra Negra is equivalent to  $\sim 90$  cm/year, and demonstrates that for persistent signals to be visible in individual interferograms, large deformation rates are required (see Figure 2b of Ebmeier et al. (2018) for a summary of deformation rates measured with InSAR). Subplot (5) of Figure 1.7 shows Campi Flegrei (Italy), and lower rate defor-



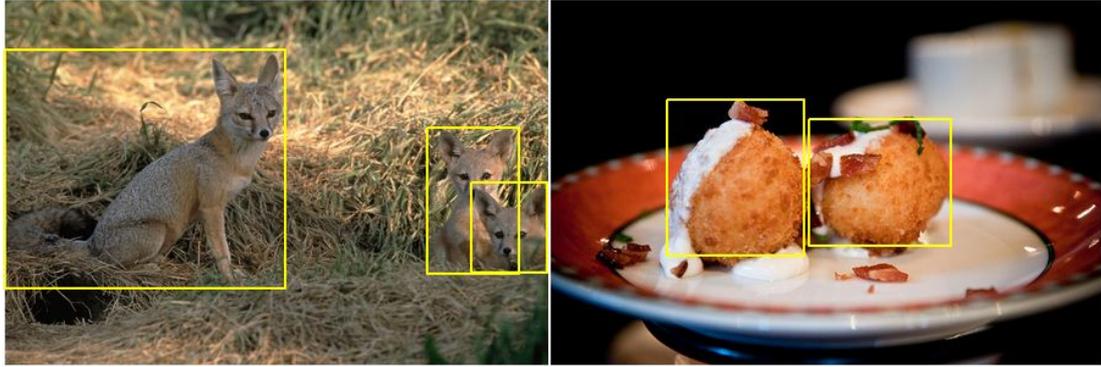
**Figure 1.8:** Pre-eruptive inflation at Sinabung (Indonesia), reproduced from Chaussard et al. (2013). The black points show the deformation for Sinabung, with inflation before and deflation after the 2010 eruption visible. In contrast, no deformation is observed at the inactive volcano Sibayak (grey squares).

mation that is not visible in a single 12 day interferogram above what is likely to be a turbulent APS with magnitude of 1 – 2 cm.

However, lower deformation rates that produce signals that may not be visible in single interferogram pairs have been measured using time series approaches. In Indonesia, the use of L-band ALOS SAR has allowed pre-eruptive inflation of 2.1 cm/year at Sinabung prior to the 2010 eruption to be measured (Chaussard et al., 2013), which is shown in Figure 1.8. Additionally, deformation of 2.1 cm/year has been measured at Kernici, and 7.8 cm/year at Agung (Chaussard et al., 2013). In other regions, examples of the use of L-band SAR include the measurement of a subsidence signal of  $\sim 5.5$  cm/year seen on a patch of size  $\sim 5.5$  km<sup>2</sup> in the Parícutin lava fields (Chaussard, 2016). Using C-band SAR, examples include deformation of  $\sim 3$  cm/year that was measured at lava flows on the flanks of Sierra Negra produced by the 1979 eruption (Amelung et al., 2000), and of  $\sim 28$  cm/year at Paka (East African Rift) that was measured over a nine month period (Biggs et al., 2009). Consequently, some of the rates of deformation that will be encountered by a detection algorithm are likely to be low enough that the signals they produce are not visible above the noise within a single 12 day interferogram, and a detection algorithm that uses a time series of data may be required to detect them.

## 1.2 Machine learning

The data used to train machine learning algorithms can be either labelled or unlabelled. Unlabelled data consists of a packet of information (such as an image of a handwritten letter), but this can be converted to labelled data through the application of an annotation detailing a point of interest in the data (such as what the letter is). Algorithms that use labelled data are termed supervised, whilst those that use unlabelled data



**Figure 1.9:** An example of labelled data from the ImageNet database (Deng et al., 2009). In addition to the bounding boxes visible in the image, the left hand image has also been labelled as “Kit fox”, and the right hand as “croquette”.

are termed unsupervised and, more recently, a third class of algorithms that use both labelled and unlabelled data has been developed, which are termed semi-supervised.

Interferograms are unlabelled data. To perform supervised (or semi-supervised) learning, labels are required for deformation patterns due to a variety of geophysical processes (such as flank instability, or uplift), and for cases where no deformation is present and an interferogram is dominated by atmospheric signals. Labelling interferograms must be performed by a domain expert as features due to tectonic processes are often subtle, which contrasts with other machine learning datasets such as ImageNet (Deng et al., 2009), in which the labelling of everyday images is relatively straightforward. Figure 1.9 shows an example of ImageNet data that has both class labels, and location labels detailing where in the scene the item of interest resides.

Consequently, the creation of a large scale database of labelled deformation patterns for use with machine learning methods is likely to be prohibitively expensive. However, interferograms spanning periods of volcanic unrest are likely to be the focus of scientific research and so it is possible that the data could be labelled using published findings at very little extra effort. Therefore, whilst the majority of data used to train an algorithm would be unlabelled, it is possible that some cases of labelled data will exist. The remainder of this section introduces both unsupervised and supervised learning.

### 1.2.1 Unsupervised learning

The two most common approaches to unsupervised learning are cluster analysis and latent variable models. Of latent variable models, the several methods to perform blind signal separation (BSS) are often cited as ways to extract features from large databases of images (e.g. Delac et al. (2005)).

BSS is the process of using several observations of a mixture of signals to reconstruct what the original signals were. A traditional blind signal separation problem is the

“cocktail party problem”. In this problem, a listener at a cocktail party is tasked with separating a linear mixture of sounds from within the party (e.g music, background conversation, and an immediate conversation) in order to listen to the one of interest (Cherry, 1953).

Mathematically, this is commonly written in the form:

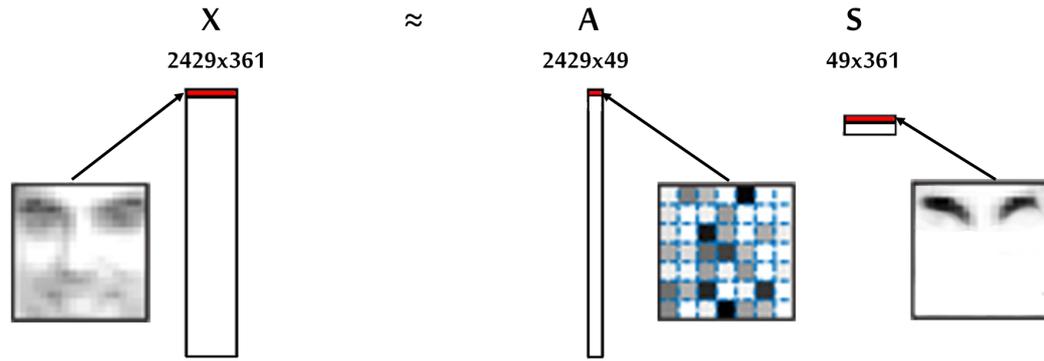
$$\mathbf{X} = \mathbf{A}\mathbf{S} \quad (1.3)$$

Where  $\mathbf{X}$  is a matrix containing the mixtures as row vectors,  $\mathbf{A}$  is a mixing matrix, and  $\mathbf{S}$  is a matrix containing the latent sources as row vectors. Given that only the mixtures,  $\mathbf{X}$ , are available to the BSS algorithm, the problem is also commonly expressed as:

$$\mathbf{S} = \mathbf{W}\mathbf{X} \quad (1.4)$$

Where  $\mathbf{W}$  is the unmixing matrix, or inverse of  $\mathbf{A}$ . When performing matrix multiplication in the mixing case, the first column of  $\mathbf{A}$  controls the contribution of the first latent source (row one of  $\mathbf{S}$ ) to each of the mixtures, and is often termed a “time course” in BSS literature. Figure 1.11 shows an example of BSS from Lee and Seung (1999) in which a database of faces are considered as linear mixtures of the constituent parts of a face using non-negative matrix factorisation (NMF), whilst Figure 1.11 shows my results from applying BSS methods to a database of faces to recover the underlying parts/sources. Other studies have also used different BSS methods on different types of data, such as McKeown et al. (1998) who used spatial independent component analysis (sICA) to show that fMRI data can be considered as a mixture of spatially independent brain regions, Frappart et al. (2011) who used sICA to show that Gravity Recovery and Climate Experiment (GRACE) data are made up from geophysical data in addition to a more prominent striping pattern, Chaussard et al. (2014) who used principal component analysis (PCA) to separate long term and season deformation in interferograms, Ebmeier (2016) who used both spatial and temporal ICA (sICA and tICA) to isolate signals of geophysical interest in interferograms featuring volcanoes, Amato et al. (2008) who used PCA and ICA to mask clouds in optical imagery, Barnie and Oppenheimer (2015) who used ICA to extract the radiance of high temperature events imaged by geostationary satellites, and Liu et al. (2015) who used PCA and ICA to isolate signal of geophysical interest in GPS time series.

Interferograms of volcanoes are usually composed of signals from several sources, such as deformation, the atmosphere, and orbital inaccuracies. An algorithm to automatically detect volcanic unrest needs to separate the tectonic deformation signal from



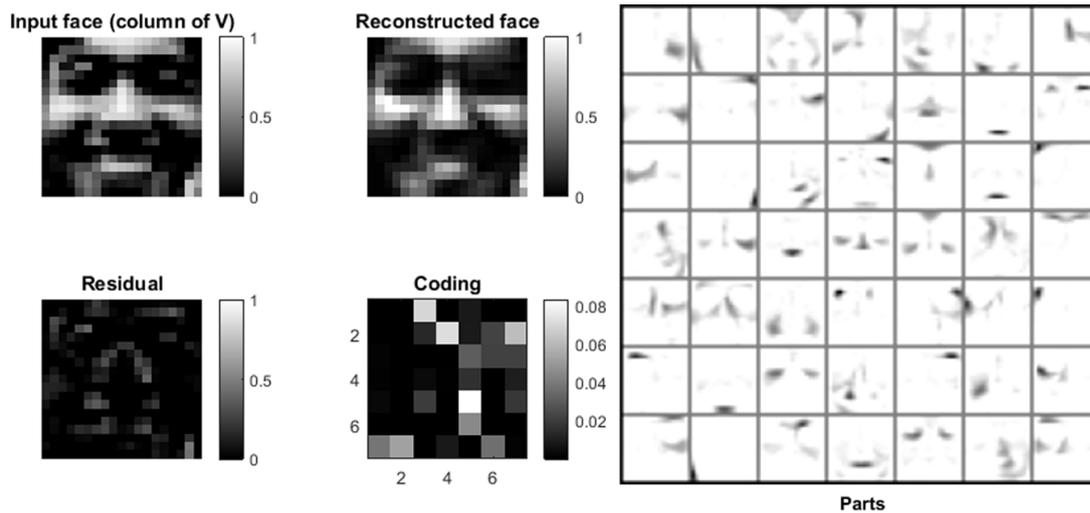
**Figure 1.10:** Schematic depicting the separation of mixtures into sources and a mixing matrix. A database of  $2429 \times 19 \times 19$  pixel images are flattened to row vectors in  $\mathbf{X}$ , which is decomposed into a mixing matrix,  $\mathbf{A}$ , and a matrix of sources or parts,  $\mathbf{S}$ . In NMF literature, it is common to visualise the row vector of the mixing matrix as a square, with each element corresponding to one of the sources/parts. Modified from Lee and Seung (1999).

the composite signal (the interferogram) in order to interpret it accurately. This problem falls within the category of “blind signal separation”, where several observations are made of a signal that is the product of several component signals, of which there is no information of what they are, or how they are mixed. Consequently, my work on unsupervised machine learning has focussed on three blind signal separation methods - non-negative matrix factorisation (NMF), principal component analysis (PCA), and independent component analysis (ICA).

### 1.2.2 Supervised learning

In contrast to unsupervised learning, supervised learning algorithms must first be trained using a selection of labelled data, before they can then be used to perform tasks such as predicting the labels of new data. In order for the algorithm to train, it must incorporate some system that allows it to adjust parameters within it in order to correctly label the training data, and this adjustment is commonly referred to as learning. Currently, there exist a suite of supervised learning algorithms, such as support vector machines (SVMs), neural networks (NNs), and decision trees. However, of this suite of methods, convolutional neural networks are currently the state-of-the-art for working with images (Zhou and Kainz, 2018), and their introduction forms the remainder of this section.

The Perceptron was introduced by Rosenblatt (1958) and built on the mathematical description of biological neurons described in McCulloch and Pitts (1943) to produce a machine that was able to learn to classify images of letters and numbers. Figure 1.12 shows an example of a perceptron, which considers a set of binary values,  $\mathbf{x}$ , as its inputs. In the case of the Mark 1 Perceptron built using custom hardware, the input was a  $20 \times 20$  array of photocells which each produced a binary output. These inputs



**Figure 1.11:** Results of applying NMF to a database of faces. The right hand square shows the 49 sources/parts that are recovered, whilst the “coding” matrix shows how these can be used to reconstruct a face. Note that some of the sources/parts can be interpreted as intuitive elements of a face, such as the eye of part 47, and the mouth of part 46.

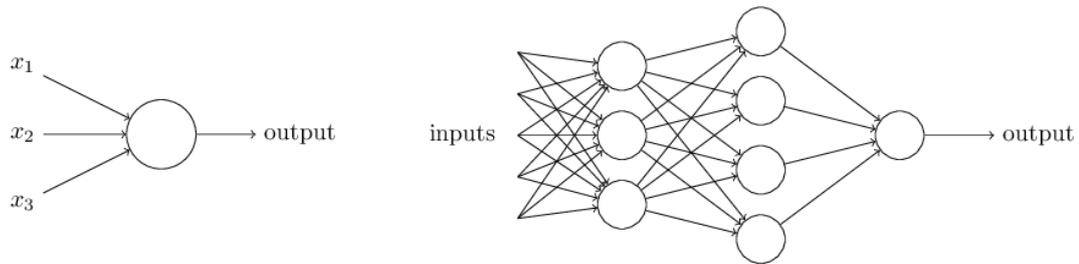
are each scaled by one of a set of weights,  $\mathbf{w}$ , before testing to see if the result exceeds a threshold which produces an output that can be used for binary classification problems. Mathematically, this can be written as:

$$\text{Binary output} = \begin{cases} 1 & \text{if } \mathbf{w} \cdot \mathbf{x} > \text{threshold} \\ 0 & \text{if } \mathbf{w} \cdot \mathbf{x} \leq \text{threshold} \end{cases} \quad (1.5)$$

However, this can be simplified by considering an extra parameter termed bias,  $b$ , and removing the *threshold* term:

$$\text{Binary output} = \begin{cases} 1 & \text{if } \mathbf{w} \cdot \mathbf{x} + b > 0 \\ 0 & \text{if } \mathbf{w} \cdot \mathbf{x} + b \leq 0 \end{cases} \quad (1.6)$$

Therefore, for a perceptron with a high bias,  $b$ , only a small input,  $\mathbf{w} \cdot \mathbf{x}$ , is required for it to produce a high output, or in biological parlance, to “fire”. Through using a set of labelled training data, the perceptron was able to use a simple update rule to change the weights, and learn to approximate a function that mapped the  $20 \times 20$  inputs to the correct label. However, despite successes in pattern recognition provided by the Perceptron (Minsky et al., 1969), the threshold that divides the two outputs is  $\mathbf{w} \cdot \mathbf{x} + b$ , which can be thought of as the equation of a hyperplane. Additionally, the change from one binary output to another which is triggered at a certain threshold can also be considered as a step activation function, which is shown Figure 1.14. Consequently, the perceptron is only a linear binary classifier, and was therefore unable to approxi-

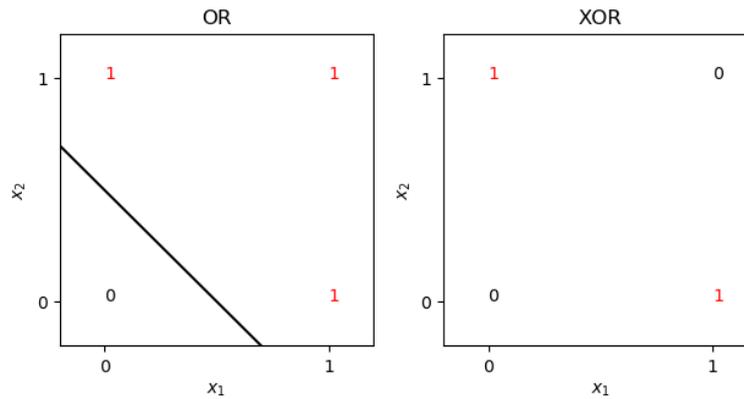


**Figure 1.12:** Left: Simple model of a perceptron in which the “neuron” (circular element) multiplies the three binary inputs,  $x_1, x_2, x_3$  with three real weights (not shown), before passing the output through a step function displaced in the  $x$  direction by a certain *threshold*. Right: Perceptrons can be stacked to produce multi-layer perceptrons, which are commonly termed artificial neural networks due to the differences in the neurons used to construct them and the original perceptron/neuron proposed by Rosenblatt (1958). Figure modified from Nielsen (2015)

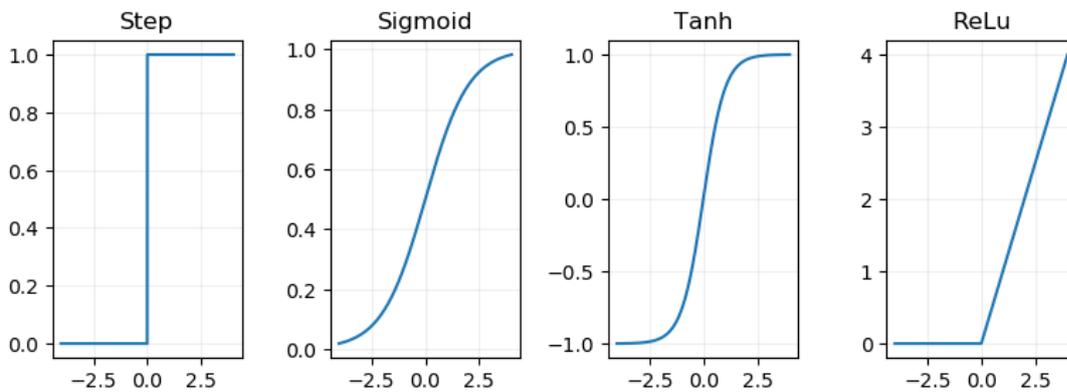
mate some simple but non-linear functions, such as exclusive OR (XOR). Figure 1.13 demonstrates the OR function, and the decision boundary that the perceptron could learn to approximate this function, and the XOR function, which it cannot.

This shortcoming of an individual perceptron can be overcome by building multi-layer networks of individual perceptrons (Minsky et al., 1969), an example of which is shown in Figure 1.12. Whilst these multilayer perceptrons can learn more complex functions such as XOR, they cannot be trained in the manner originally described in Rosenblatt (1958) (Nielsen, 2015), due to the step activation functions contained within individual perceptrons causing training to be highly unstable, as small changes in the weights of early layers produce either no or drastic changes in the subsequent layers, depending on if the threshold of the step activation function is exceeded. However, this limitation was overcome with the development of backpropagation (Rumelhart et al., 1986), which substituted the step activation function used in the original perceptrons for a differentiable sigmoid function (shown in Figure 1.14), and allowed for multi-layer models to be trained using differentiation. Due to the multi-layer nature of these models, they are commonly termed “deep models”, and their utilisation termed “deep learning”.

The performance of multilayer models was further improved through the use of convolutional neural networks (CNNs), in which the  $2D$  structure of an image can be utilised by a network through its learning filters that can be convolved across an image. This contrasts with earlier approaches in which an image is flattened to be a  $1D$  vector, before each input (i.e. pixel) is connected to subsequent neurons. When larger images are used, the large number of weights to be learned for each of these connections can be challenging, but through using only a small number of weights that are convolved across all the pixels of an image, CNNs greatly reduce the complexity of training a model, for a given size of input image. This approach was used by LeCun et al. (1998) to create a



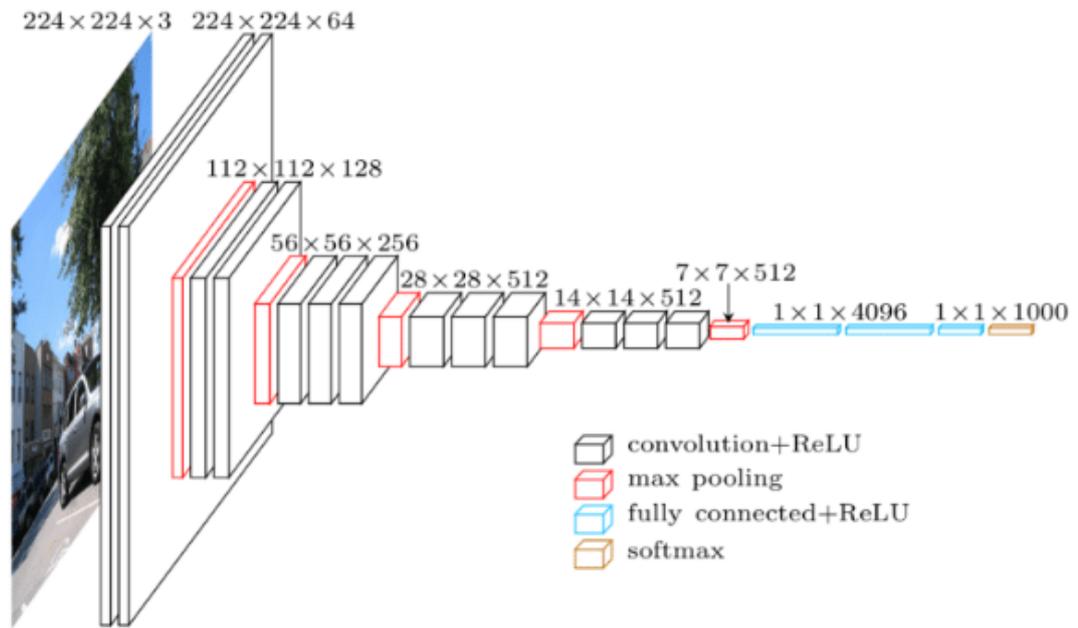
**Figure 1.13:** Comparison of a linearly separable function, OR, and a more complex function that cannot be separated linearly, XOR (except OR). A linear decision boundary has been added to the OR function, but cannot be added to the XOR function.



**Figure 1.14:** A comparison of several common activation functions. The step function was used in perceptrons, before the differentiable sigmoid and hyperbolic tangent (tanh) functions were used in multilayer networks. The rectified linear unit (ReLU) is also a common activation function due to its low computational cost.

model with state-of-the-art performance for classifying hand written digits, but it was not until AlexNet (Krizhevsky et al., 2012) that this approach was successfully scaled to larger problems involving more classes and larger input data. Figure 1.15 shows an overview of a model similar to AlexNet termed VGG16 (Simonyan and Zisserman, 2014), in which convolutional layers produce successively deeper representations of an image, before it is assigned a class label.

Subsequently, deep learning has permeated through many aspects of society, ranging from search engines to content filtering on social media to scientific applications (LeCun et al., 2015). The introduction of graphical processing units (GPUs) and computationally inexpensive activation functions such as rectified linear units (ReLUs, shown in Figure 1.14) has reduced the time required to train models, whilst easy to use optimisers such as NADAM (Dozat, 2016) do not require the time consuming choice



**Figure 1.15:** An overview of the VGG16 (Simonyan and Zisserman, 2014), reproduced from Sugata and Yang (2017). A three channel image of size  $224 \times 224$  is converted into a series of successively spatially smaller but deeper feature maps, culminating in a  $7 \times 7 \times 512$  representation. This is flattened to become a vector, which is connected to a series of fully connected layers, which are similar to the traditional neural network shown in Figure 1.12.

of a learning rate. Additionally, high level application programming interfaces (APIs), such as Keras, allow for different model architectures to be explored easily, and the availability of weights learned in successful models allows for new and complex models to be constructed using relatively little data (e.g. Anantrasirichai et al. (2018)).

The application of unsupervised learning is central to the work featured in chapters 2 and 3, whilst supervised learning is used in 4. The remainder of this chapter introduces the aims and objectives of this thesis.

### 1.3 Aims and Objectives

With many of the requirements for InSAR to evolve from a tool used for retrospective analysis into one used for hazard monitoring having been met (e.g. the automatic creation of Sentinel-1 interferograms (González et al., 2016)), I aim to develop an algorithm to detect deformation-generating volcanic unrest that produces deformation. I envisage that this algorithm will utilise a time series of interferograms, and so be able to detect small signals that are not apparent in single interferograms. As the separation of nuisance atmospheric signals from deformation signals has traditionally been important for geophysical studies, I envisage that a key part of the algorithm will be its ability to determine if signals present in new interferograms are atmospheric, or

due to deformation.

The objectives of this thesis are:

1. Determine how to characterise the baseline signals present in a time series of interferograms of a volcanic centre.
2. Develop a monitoring algorithm to both detect when the signals present in new interferograms deviate from the baseline signals, and to display this information in a form that is easily interpretable to a domain expert.
3. Develop a separate monitoring algorithm that can be used to ascertain if individual interferograms contain deformation signals.
4. Perform a small study into a volcanic process that is initially investigated during the construction of a monitoring algorithm.

## 1.4 Thesis outline

The subsequent chapters in this thesis are organised as follows:

- Chapter 2 details determining which of a suite of blind signal separation methods are best suited to recovering signals of geophysical interest from a time series of interferograms, and how time series should be organised for the best results. Spatial independent component analysis (sICA) is found to be the most suitable method, and it is applied to a time series of Sentinel-1 interferograms that image Wolf volcano (Galapagos Archipelago, Ecuador), and Mount Etna (Italy). sICA is then used in a prototype detection algorithm, which we show would have been able to detect the 2015 eruption of Wolf volcano. This chapter has been published in the *Journal of Geophysical Research: Solid Earth*.
- Chapter 3 details the development of an algorithm, ICASAR, to robustly apply sICA to time series of interferograms in order to identify the spatial patterns and time histories of the signals that formed them. The spatial and temporal nature of the signals that this algorithm is able to extract at a volcanic centre is used as a component of a more complete detection algorithm, which we apply to a time series of Sentinel-1 interferograms that image the lead up to the 2018 eruption of Sierra Negra (Galapagos Archipelago, Ecuador). Additionally, one of the signals extracted by our algorithm is attributed to the inflation of a sill, and modelling to calculate the pre-eruptive overpressure within this sill is performed. This chapter is in review with the *Journal of Geophysical Research: Solid Earth*.
- Chapter 4 details the development of a convolutional neural network that is able to classify several types of deformation that may be present at a volcanic centre, and to locate the deformation within the interferogram. Additionally, the

optimal format for InSAR data to be used with models trained on natural images is investigated. This chapter is pending submission to *Geophysical Journal International*.

- Chapter 5 discusses the work contained within the preceding chapters in relation to the goal of using InSAR to automatically monitor the majority of the world's subaerial volcanoes.

## References

- Amato, Umberto, Anestis Antoniadis, Vicenze Cuomo, Luisa Cutillo, M Franzese, Loredana Murino, and Carmine Serio (2008). "Statistical cloud detection from SEVIRI multispectral images". In: *Remote Sensing of Environment* 112.3, pp. 750–766.
- Amelung, Falk, S Jónsson, Howard Zebker, and Paul Segall (2000). "Widespread uplift and 'trapdoor' faulting on Galápagos volcanoes observed with radar interferometry." In: *Nature* 407.6807, pp. 993–996. ISSN: 0028-0836. DOI: 10.1038/35039604. URL: <http://www.ncbi.nlm.nih.gov/pubmed/11069176>.
- Anantrasirichai, N, Juliet Biggs, F. Albino, P. Hill, and D. Bull (2018). "Application of Machine Learning to Classification of Volcanic Deformation in Routinely Generated InSAR Data". In: *Journal of Geophysical Research : Solid Earth*, pp. 1–15. DOI: 10.1029/2018JB015911.
- Arnold, DWD, J Biggs, K Anderson, S Vallejo Vargas, Geoff Wadge, SK Ebmeier, MF Naranjo, and P Mothes (2017). "Decaying Lava Extrusion Rate at El Reventador Volcano, Ecuador, Measured Using High-Resolution Satellite Radar". In: *Journal of Geophysical Research: Solid Earth* 122.12, pp. 9966–9988.
- Bagnardi, Marco, Pablo J González, and Andrew Hooper (2016). "High-resolution digital elevation model from tri-stereo Pleiades-1 satellite imagery for lava flow volume estimates at Fogo Volcano". In: *Geophysical Research Letters* 43.12, pp. 6267–6275.
- Barnie, Talfan and Clive Oppenheimer (2015). "Extracting High Temperature Event radiance from satellite images and correcting for saturation using Independent Component Analysis". In: *Remote Sensing of Environment* 158, pp. 56–68. ISSN: 00344257. DOI: 10.1016/j.rse.2014.10.023. URL: <http://dx.doi.org/10.1016/j.rse.2014.10.023>.
- Bekaert, D P S, A Hooper, and T J Wright (2015). "A spatially variable power law tropospheric correction technique for InSAR data". In: *Journal of Geophysical Research: Solid Earth* 2, pp. 1–12. ISSN: 2169-9356. DOI: 10.1002/2014JB011557.A.
- Berardino, P., J. Bequignon, S. Borgstrom, G. De Natale, P. Capuano, G. Fornaro, R. Lanari, F. Pingue, G.P. Ricciardi, E. Sansosti, and C. Troise (2002). "Evidence for a peculiar style of ground deformation at Vesuvius volcano revealed by 10 years of

- ERS mission". In: *IEEE International Geoscience and Remote Sensing Symposium* 6.9, pp. 2–5. ISSN: 0094-8276. DOI: 10.1109/IGARSS.2002.1027267.
- Biggs, J., E.Y. Y. Anthony, and C.J. J. Ebinger (Oct. 2009). "Multiple inflation and deflation events at Kenyan volcanoes, East African Rift". In: *Geology* 37.11, pp. 979–982. ISSN: 0091-7613. DOI: 10.1130/G30133A.1. URL: <http://geology.gsapubs.org/cgi/doi/10.1130/G30133A.1>.
- Biggs, J, S K Ebmeier, W P Aspinall, Z Lu, M E Pritchard, R S J Sparks, and T a Mather (Jan. 2014). "Global link between deformation and volcanic eruption quantified by satellite imagery." In: *Nature communications* 5, p. 3471. ISSN: 2041-1723. DOI: 10.1038/ncomms4471.
- Bond, A. and RSJ Sparks (1976). "The Minoan eruption of Santorini, Greece". In: *Journal of the Geological Society* 132.1, pp. 1–16.
- Carboni, Elisa, R Grainger, Joanne Walker, Anu Dudhia, and Richard Siddans (2012). "A new scheme for sulphur dioxide retrieval from IASI measurements: application to the Eyjafjallajökull eruption of April and May 2010". In: *Atmospheric Chemistry and Physics* 12.23, pp. 11417–11434.
- Chaussard, E., F. Amelung, and Y. Aoki (Aug. 2013). "Characterization of open and closed volcanic systems in Indonesia and Mexico using InSAR time series". In: *Journal of Geophysical Research: Solid Earth* 118.8, pp. 3957–3969. ISSN: 21699313. DOI: 10.1002/jgrb.50288. URL: <http://doi.wiley.com/10.1002/jgrb.50288>.
- Chaussard, E, R Bürgmann, M Shirzaei, E J Fielding, and B Baker (2014). "Predictability of hydraulic head changes and characterization of aquifer-system and fault properties from InSAR-derived ground deformation". In: *Journal of Geophysical Research : Solid Earth* 119, pp. 6572–6590. DOI: 10.1002/2014JB011266.
- Chaussard, Estelle (2016). "Subsidence in the Parícutin lava field: Causes and implications for interpretation of deformation fields at volcanoes". In: *Journal of Volcanology and Geothermal Research* 320, pp. 1–11.
- Chaussard, Estelle and F. Amelung (Nov. 2012). "Precursory inflation of shallow magma reservoirs at west Sunda volcanoes detected by InSAR". In: *Geophysical Research Letters* 39.21, n/a–n/a. ISSN: 00948276. DOI: 10.1029/2012GL053817. URL: <http://doi.wiley.com/10.1029/2012GL053817>.
- Chen, Ting and Nadia Lapusta (2009). "Scaling of small repeating earthquakes explained by interaction of seismic and aseismic slip in a rate and state fault model". In: *Journal of Geophysical Research: Solid Earth* 114.1, pp. 1–12. ISSN: 21699356. DOI: 10.1029/2008JB005749.
- Cherry, E (1953). *Some Experiments on the Recognitin of Speeck with one and Two ears*.
- De Zan, Francesco and A Monti Guarnieri (2006). "TOPSAR: Terrain observation by progressive scans". In: *IEEE Transactions on Geoscience and Remote Sensing* 44.9, pp. 2352–2360.

- Delac, Kresimir, Mislav Grgic, and Sonja Grgic (2005). “Independent comparative study of PCA, ICA, and LDA on the FERET data set”. In: *International Journal of Imaging Systems and Technology* 15.5, pp. 252–260. ISSN: 08999457. DOI: 10.1002/ima.20059.
- Delgado, Francisco, Matthew E Pritchard, Susanna Ebmeier, Pablo González, and Luis Lara (2017). “Recent unrest (2002–2015) imaged by space geodesy at the highest risk Chilean volcanoes: Villarrica, Llaima, and Calbuco (Southern Andes)”. In: *Journal of Volcanology and Geothermal Research* 344, pp. 270–288.
- Deng, Jia, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei (2009). “Imagenet: A large-scale hierarchical image database”. In: *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on. Ieee*, pp. 248–255.
- Dozat, Timothy (2016). “Incorporating nesterov momentum into adam”. In:
- Druitt, TH (1998). “Pyroclastic density currents”. In: *Geological Society, London, Special Publications* 145.1, pp. 145–182.
- Ebmeier, S K, B J Andrews, M C Araya, D W D Arnold, J Biggs, C Cooper, E Cottrell, M Furtney, J Hickey, J Jay, R Lloyd, A L Parker, M E Pritchard, E Robertson, E Venzke, and J L Williamson (2018). “Synthesis of global satellite observations of magmatic and volcanic deformation : implications for volcano monitoring & the lateral extent of magmatic domains”. In: *Journal of Applied Volcanology*, pp. 1–26. DOI: 10.1186/s13617-018-0071-3.
- Ebmeier, S. K. (2016). “Application of Independent Component Analysis to multi-temporal InSAR data with volcanic case studies”. In: *Journal of Geophysical Research: Solid Earth*, pp. 8970–8986. ISSN: 21699313. DOI: 10.1002/2016JB013765. URL: <http://doi.wiley.com/10.1002/2016JB013765>.
- Ebmeier, S. K., J. Biggs, T. a. Mather, and F. Amelung (May 2013). “On the lack of InSAR observations of magmatic deformation at Central American volcanoes”. In: *Journal of Geophysical Research: Solid Earth* 118.5, pp. 2571–2585. ISSN: 21699313. DOI: 10.1002/jgrb.50195. URL: <http://doi.wiley.com/10.1002/jgrb.50195>.
- Elliott, John R., J. Biggs, B. Parsons, and T. J. Wright (2008). “InSAR slip rate determination on the Altyn Tagh Fault, northern Tibet, in the presence of topographically correlated atmospheric delays”. In: *Geophysical Research Letters* 35.12, pp. 1–5. ISSN: 00948276. DOI: 10.1029/2008GL033659.
- Ferretti, a., C. Prati, and F. Rocca (2001). “Permanent scatterers in SAR interferometry”. In: *IEEE Transactions on Geoscience and Remote Sensing* 39.1, pp. 8–20. ISSN: 01962892. DOI: 10.1109/36.898661. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=898661>.
- Frappart, Frédéric, Guillaume Ramillien, Marc Leblanc, Sarah O. Tweed, Marie Paule Bonnet, and Philippe Maisongrande (2011). “An independent component analysis filtering approach for estimating continental hydrology in the GRACE gravity data”. In: *Remote Sensing of Environment* 115.1, pp. 187–204. ISSN: 00344257. DOI:

- 10.1016/j.rse.2010.08.017. URL: <http://dx.doi.org/10.1016/j.rse.2010.08.017>.
- González, Pablo J., Richard J. Walters, E. L. Hatton, Karsten Spaans, A. McDougall, Andrew Hooper, and Tim J. Wright (2016). “LiCSAR: Tools for automated generation of Sentinel-1 frame interferograms”. In: *AGU Fall meeting*.
- González, Pablo J, Marco Bagnardi, Andrew J Hooper, Yngvar Larsen, Petar Marinkovic, Sergey V Samsonov, and Tim J Wright (2015). “The 2014–2015 eruption of Fogo volcano: Geodetic modeling of Sentinel-1 TOPS interferometry”. In: *Geophysical research letters* 42.21, pp. 9239–9246.
- Gudmundsson, Agust (2012). “Magma chambers: Formation, local stresses, excess pressures, and compartments”. In: *Journal of Volcanology and Geothermal Research* 237-238, pp. 19–41. ISSN: 03770273. DOI: 10.1016/j.jvolgeores.2012.05.015. URL: <http://dx.doi.org/10.1016/j.jvolgeores.2012.05.015>.
- Hanssen, Ramon F (2001). *Radar interferometry: data interpretation and error analysis*. Vol. 2. Springer Science & Business Media.
- Hooper, A., P. Segall, and H. Zebker (2007). “Persistent scatterer interferometric synthetic aperture radar for crustal deformation analysis, with application to Volcan Alcedo, Galapagos”. In: *Journal of Geophysical Research: Solid Earth* 112.7, pp. 1–21. ISSN: 21699356. DOI: 10.1029/2006JB004763.
- Hooper, Andrew (Aug. 2008). “A multi-temporal InSAR method incorporating both persistent scatterer and small baseline approaches”. In: *Geophysical Research Letters* 35.16, p. L16302. ISSN: 0094-8276. DOI: 10.1029/2008GL034654. URL: <http://doi.wiley.com/10.1029/2008GL034654>.
- (2010). “A statistical-cost approach to unwrapping the phase of InSAR time series”. In: *Proceeding of International Workshop on ERS SAR 2009*. March, pp. 1–5. URL: [http://radar.tudelft.nl/%7B~%7Dahooper/Hooper%7B%5C\\_%7DFRINGE%7B%5C\\_%7D2009.pdf](http://radar.tudelft.nl/%7B~%7Dahooper/Hooper%7B%5C_%7DFRINGE%7B%5C_%7D2009.pdf).
- Hooper, Andrew, David Bekaert, Karsten Spaans, and Mahmut Arıkan (Jan. 2012). “Recent advances in SAR interferometry time series analysis for measuring crustal deformation”. In: *Tectonophysics* 514-517, pp. 1–13. ISSN: 00401951. DOI: 10.1016/j.tecto.2011.10.013.
- Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E Hinton (2012). “Imagenet classification with deep convolutional neural networks”. In: *Advances in neural information processing systems*, pp. 1097–1105.
- LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton (2015). “Deep learning”. In: *nature* 521.7553, p. 436.
- LeCun, Yann, Léon Bottou, Yoshua Bengio, and Patrick Haffner (1998). “Gradient-based learning applied to document recognition”. In: *Proceedings of the IEEE* 86.11, pp. 2278–2324.

- Lee, Daniel D and H Sebastian Seung (1999). “Learning the parts of objects by non-negative matrix factorization”. In: *Nature* 401.October 1999, pp. 788–791. ISSN: 0028-0836. DOI: 10.1038/44565. arXiv: arXiv:1408.1149. URL: <http://www.ncbi.nlm.nih.gov/pubmed/10548103>.
- Liu, Bin, Wujiao Dai, Wei Peng, and Xiaolin Meng (2015). “Spatiotemporal analysis of GPS time series in vertical direction using independent component analysis”. In: *Earth, Planets and Space* 67.1, p. 189.
- Lohman, Rowena B. and Mark Simons (2005). “Some thoughts on the use of InSAR data to constrain models of surface deformation: Noise structure and data down-sampling”. In: *Geochemistry, Geophysics, Geosystems* 6.1. ISSN: 15252027. DOI: 10.1029/2004GC000841.
- Machado, Guillermo, Valeria Lupiano, Maria Vittoria Avolio, Francesco Gullace, and Salvatore Di Gregorio (2015). “A cellular model for secondary lahars and simulation of cases in the Vascún Valley, Ecuador”. In: *Journal of Computational Science* 11, pp. 289–299.
- Massonet, Didier, Marc Rossi, Cesar Carmona, Frederic Adragna, Gilles Peltzer, Kurt Feigi, and Thierry Rabaute (1993). “The displacement field of the Landers earthquake mapped by radar interferometry”. In: *Nature* 364, pp. 138–142.
- Massonet, Didier, Pierre Briole, and Alain Arnaud (1995). *Deflation of Mount Etna monitored by spaceborne radar interferometry*. DOI: 10.1038/375567a0.
- McCulloch, Warren S and Walter Pitts (1943). “A logical calculus of the ideas immanent in nervous activity”. In: *The bulletin of mathematical biophysics* 5.4, pp. 115–133.
- McKeown, M J, T P Jung, S Makeig, G Brown, S S Kindermann, T W Lee, and T J Sejnowski (1998). “Spatially independent activity patterns in functional MRI data during the stroop color-naming task.” In: *Proceedings of the National Academy of Sciences of the United States of America* 95.3, pp. 803–810. ISSN: 00278424. DOI: 10.1073/pnas.95.3.803. URL: <http://www.pnas.org/content/95/3/803%7B%5C%7D5Cnhttp://www.ncbi.nlm.nih.gov/pubmed/9448244%7B%5C%7D5Cnhttp://www.pnas.org/content/95/3/803.full.pdf%7B%5C%7D5Cnhttp://www.pnas.org/content/95/3/803.long>.
- Meyer, F. J., P. W. Webley, J. Dehn, S. A. Arko, D. B. McAlpin, and W. Gong (2016). “The SARVIEWS Project: Automated SAR Processing in Support of Operational Near Real-time Volcano Monitoring”. In: *AGU Fall meeting*.
- Minsky, ML, SA Papert, and First Perceptrons (1969). “The MIT Press: Cambridge”. In: *Mass.(Rev. Edition, 1988)*.
- Nielsen, Michael A. (2015). *Neural Networks and Deep Learning*. Determination Press. URL: <http://neuralnetworksanddeeplearning.com/index.html>.
- Ofeigsson, Benedikt G., Andrew Hooper, Freysteinn Sigmundsson, Erik Sturkell, and Ronni Grapenthin (May 2011). “Deep magma storage at Hekla volcano, Iceland, revealed by InSAR time series analysis”. In: *Journal of Geophysical Research* 116.B5,

- B05401. ISSN: 0148-0227. DOI: 10.1029/2010JB007576. URL: <http://doi.wiley.com/10.1029/2010JB007576>.
- Othman-Chandev, M (1987). “The Cameroon volcanic gas disaster. An analysis of a makeshift response”. In: *Disasters* 11.2, pp. 96–101.
- Pavez, A, D Remy, S Bonvalot, M Diament, G Gabalda, Jean-Luc Froger, P Julien, Denis Legrand, and D Moisset (2006). “Insight into ground deformations at Lascar volcano (Chile) from SAR interferometry, photogrammetry and GPS data: Implications on volcano dynamics and future space monitoring”. In: *Remote Sensing of Environment* 100.3, pp. 307–320.
- Pinel, V., A. Hooper, S. De la Cruz-Reyna, G. Reyes-Davila, M. P. Doin, and P. Bascou (2011). “The challenging retrieval of the displacement field from InSAR data for andesitic stratovolcanoes: Case study of Popocatepetl and Colima Volcano, Mexico”. In: *Journal of Volcanology and Geothermal Research* 200.1-2, pp. 49–61. ISSN: 03770273. DOI: 10.1016/j.jvolgeores.2010.12.002.
- Ramalho, Ricardo S, Gisela Winckler, José Madeira, George R Helffrich, Ana Hipólito, Rui Quartau, Katherine Adena, and Joerg M Schaefer (2015). “Hazard potential of volcanic flank collapses raised by new megatsunami evidence”. In: *Science advances* 1.9, e1500456.
- Rosenblatt, Frank (1958). “The perceptron: a probabilistic model for information storage and organization in the brain.” In: *Psychological review* 65.6, p. 386.
- Rumelhart, David E, Geoffrey E Hinton, and Ronald J Williams (1986). “Learning representations by back-propagating errors”. In: *nature* 323.6088, p. 533.
- Schmidt, Anja, Thorvaldur Thordarson, Luke D Oman, Alan Robock, and Stephen Self (2012). “Climatic impact of the long-lasting 1783 Laki eruption: Inapplicability of mass-independent sulfur isotopic composition measurements”. In: *Journal of Geophysical Research: Atmospheres* 117.D23.
- Shen, Lin, Andrew Hooper, and John Elliott (2019). “A spatially-varying scaling method for InSAR tropospheric corrections using a high-resolution weather model”. In: *Journal of Geophysical Research: Solid Earth*.
- Siebert, L. and T. Simkin (2013). *Volcanoes of the World: an Illustrated Catalog of Holocene Volcanoes and their Eruptions*. URL: [www.volcano.si.edu](http://www.volcano.si.edu).
- Sigmundsson, Freysteinn, Sigrún Hreinsdóttir, Andrew Hooper, Thóra Arnadóttir, Rikke Pedersen, Matthew J Roberts, Níels Óskarsson, Amandine Auriac, Judicael Decriem, Páll Einarsson, Halldór Geirsson, Martin Hensch, Benedikt G Ofeigsson, Erik Sturkell, Hjörleifur Sveinbjörnsson, Kurt L Feigl, Thóra Árnadóttir, Rikke Pedersen, Matthew J Roberts, Níels Óskarsson, Amandine Auriac, Judicael Decriem, Páll Einarsson, Halldór Geirsson, Martin Hensch, Benedikt G. Ófeigsson, Erik Sturkell, Hjörleifur Sveinbjörnsson, and Kurt L Feigl (Nov. 2010). “Intrusion triggering of the 2010 Eyjafjallajökull explosive eruption.” In: *Nature* 468.7322, pp. 426–30. ISSN: 1476-4687. DOI: 10.1038/nature09558. URL: <http://dx.doi.org/10.1038/nature09558>.

- org/10.1038/nature09558%7B%5C%7D5Cnhttp://www.nature.com/doi/finder/10.1038/nature09558%20http://www.ncbi.nlm.nih.gov/pubmed/21085177.
- Simonyan, Karen and Andrew Zisserman (2014). “Very deep convolutional networks for large-scale image recognition”. In: *arXiv preprint arXiv:1409.1556*.
- Small, Christopher and Terry Naumann (2001). “The global distribution of human population and recent volcanism”. In: *Global Environmental Change Part B: Environmental Hazards* 3.3, pp. 93–109.
- Smith, Gary A and Donald R Lowe (1991). “Lahars: Volcano hydrologic-events and deposition in the debris flow—Hyperconcentrated flow continuum”. In:
- Spaans, Karsten and Andy Hooper (2016). “InSAR processing for volcano monitoring and other near-real time applications”. In: *Journal of Geophysical Research: Solid Earth*, n/a–n/a. ISSN: 21699313. DOI: 10.1002/2015JB012752.
- Sparks, R. S. J., J. Biggs, and J. W. Neuberg (2012). “Monitoring Volcanoes”. In: *Science* 335.6074, pp. 1310–1311. ISSN: 0036-8075. DOI: 10.1126/science.1219485.
- Sugata, TLI and CK Yang (2017). “Leaf App: Leaf recognition with deep convolutional neural networks”. In: *Materials Science and Engineering Conference Series*. Vol. 273. 1, p. 012004.
- Tilling, RI (2008). “The critical role of volcano monitoring in risk reduction”. In: *Advances in Geosciences* 14, pp. 3–11.
- Wadge, G., P. W. Webley, I. N. James, R. M. Bingley, a. H. Dodson, S. J. Waugh, T. Veneboer, G. Puglisi, M Mattia, D. Baker, S. C. Edwards, Edwards S. J., and P. J. Clarker (2002). “Atmospheric models, GPS and InSAR measurements of the tropospheric water vapour field over Mount Etna”. In: *Geophysical Research Letters* 29.19, pp. 4–7. ISSN: 0094-8276. DOI: 10.1029/2002GL015159.
- Walters, R. J., J. R. Elliott, Z. Li, and B. Parsons (July 2013). “Rapid strain accumulation on the Ashkabad fault (Turkmenistan) from atmosphere-corrected InSAR”. In: *Journal of Geophysical Research: Solid Earth* 118.7, pp. 3674–3690. ISSN: 21699313. DOI: 10.1002/jgrb.50236. URL: <http://doi.wiley.com/10.1002/jgrb.50236>.
- Wang, Teng, Michael P Poland, and Zhong Lu (2015). “Dome growth at Mount Cleveland, Aleutian Arc, quantified by time series TerraSAR-X imagery”. In: *Geophysical Research Letters* 42.24, pp. 10–614.
- Wauthier, C., V. Cayol, F. Kervyn, and N. D’Oreye (2012). “Magma sources involved in the 2002 Nyiragongo eruption, as inferred from an InSAR analysis”. In: *Journal of Geophysical Research: Solid Earth* 117.5, pp. 1–20. ISSN: 21699356. DOI: 10.1029/2011JB008257.
- Wicks, Charles W. (2002). “Magmatic activity beneath the quiescent Three Sisters volcanic center, central Oregon Cascade Range, USA”. In: *Geophysical Research Letters* 29.7, pp. 2–5. ISSN: 0094-8276. DOI: 10.1029/2001GL014205.

- Williams, Simon, Yehuda Bock, Peng Fang, and H Cecil (1998). “Integrated satellite interferometry’ Tropospheric noise, GPS estimates and implications for interferometric synthetic aperture radar products”. In: *Network* 98, pp. 51–67.
- Wright, Tim J. (2004). “Toward mapping surface deformation in three dimensions using InSAR”. In: *Geophysical Research Letters* 31.1, p. L01607. ISSN: 0094-8276. DOI: 10.1029/2003GL018827. URL: <http://doi.wiley.com/10.1029/2003GL018827>.
- Wright, Tim, Barry Parsons, and Eric Fielding (2001). “Measurement of interseismic strain accumulation across the North Anatolian Fault by satellite radar interferometry”. In: *Geophysical Research Letters* 28.10, p. 2117. ISSN: 0094-8276. DOI: 10.1029/2000GL012850.
- Zhou, Keyang and Bernhard Kainz (2018). “Efficient Image Evidence Analysis of CNN Classification Results”. In: *arXiv preprint arXiv:1801.01693*.

## Chapter 2

# Blind Signal Separation Methods for InSAR: The Potential to Automatically Detect and Monitor Signals of Volcanic Deformation

M. Gaddes<sup>1</sup>, A. Hooper<sup>1</sup>, M. Bagnardi<sup>2,3</sup>, H. Inman<sup>4,5</sup>, and F. Albino<sup>6</sup>

<sup>1</sup> *COMET, School of Earth and Environment, University of Leeds, United Kingdom*

<sup>2</sup> *Formerly, COMET, School of Earth and Environment, University of Leeds, United Kingdom*

<sup>3</sup> *Now at Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA*

<sup>4</sup> *Formerly at School of Earth and Environment, University of Leeds, Leeds, UK*

<sup>5</sup> *Now at Shearwater Geoservices, CA, USA*

<sup>6</sup> *School of Earth Sciences, University of Bristol, Bristol, UK*

Keypoints:

- Spatial independent component analysis (sICA) outperforms principal component analysis (PCA) and non-negative matrix factorisation (NMF) as the best blind source separation method for identifying sources in InSAR time series, but difficulties remain in applying it to stratovolcanoes.
- sICA is able to automatically isolate three previously documented signals that were present at the 2015 eruption of Wolf Volcano (Galapagos Archipelago, Ecuador).
- We present a prototype automatic detection algorithm, which identifies the initiation of the 2015 eruption of Wolf Volcano.

## Abstract

There are some 1500 volcanoes with the potential to erupt, but most are not instrumentally monitored. However, routine acquisition by the Sentinel-1 satellites now fulfils the requirements needed for InSAR to progress from a retrospective analysis tool to one used for near real time monitoring globally. However, global monitoring produces vast quantities of data and consequently, an automatic detection algorithm is therefore required that is able to identify signs of new deformation, or changes in rate, in a time series of interferograms.

On the basis that much of the signal contained in a time series of interferograms can be considered as a linear mixture of several latent sources, we explore the use of blind source separation methods to address this issue. We consider principal component analysis (PCA) and independent component analysis (ICA) which have previously been applied to InSAR data, and non-negative matrix factorisation (NMF) which has not. Our systematic analysis of the three methods shows independent component analysis (ICA) to be best suited for most applications with InSAR data. However, care must be taken in the dimension reduction step of ICA not to remove important smaller magnitude signals. We apply ICA to the 2015 Wolf Volcano eruption (Galapagos Archipelago, Ecuador) and automatically isolate three signals, which are broadly similar to those manually identified in other studies. Finally, we develop a prototype detection algorithm based on ICA to identify the onset of the eruption.

## 2.1 Introduction

The world's  $\sim 1500$  subaerial active volcanoes (Siebert and Simkin, 2013) pose a diverse set of geophysical hazards, which range from those such as pyroclastic flows, which impact people living near the volcano, to those such as ash clouds, which can impact people across the globe. Data gathered from methods or instruments such as seismol-

ogy, tiltmeters, strain gauges, levelling lines, satellite derived thermal imagery, satellite derived gas emissions, satellite derived deformation measurements, and geochemical analysis of fumarolic gases can be used to identify the signs of volcanic unrest that are indicative of a certain hazard (Sparks et al. (2012), and references therein). However, the majority of the world’s volcanoes are not routinely monitored (Sparks et al., 2012), and the vast majority of the methods previously mentioned would be prohibitively expensive to extend to all volcanoes, due to the requirement for more instruments to be acquired and deployed.

The ability of interferometric synthetic aperture radar (InSAR) to measure ground displacements allows deformation of both a volcano’s flanks and the area surrounding it to be constrained (e.g. Pinel et al. (2014) and references therein), and these measurements of displacement can be a valuable indicator of volcanic unrest that may lead to an eruption (Biggs et al., 2014; Ebmeier et al., 2018). Therefore, a satellite with a suitable acquisition strategy could provide the measurements that would allow monitoring of all of the world’s subaerial volcanoes.

Until recently, InSAR was not suited to real-time global monitoring as SAR satellites in operation before 2015 rarely routinely acquired images over areas of geophysical interest, and data were generally not available until days after it was acquired. However, the European Space Agency’s most recent SAR satellites (the Sentinel-1 constellation) have addressed the limitations previously listed and have allowed InSAR to evolve from a retrospective analysis tool into one that can be used for near real time monitoring. Other advances that make this possible include the construction of new processing facilities to automatically and rapidly create interferograms for large parts of the globe (e.g. González et al. (2016) and Meyer et al. (2016)), and new methods for rapidly updating a time series when new interferograms are created (e.g. Spaans and Hooper (2016)).

However, using data derived from the Sentinel-1 satellites to monitor the world’s subaerial active volcanoes produces too many interferograms for them to be manually inspected for signs of unrest. Consequently, we present results of work to construct an algorithm to automatically detect signs of unrest in a time series of interferograms over a given volcano. To avoid the time consuming nature of labelling data (e.g. manually inspecting interferograms and assigning information such as “contains deformation”), our algorithm must be an unsupervised one (i.e. one that is trained on unlabelled data). One class of unsupervised learning algorithms are termed blind signal separation (BSS) methods, and one of these (ICA) has been utilised to isolate signals of geophysical interest in an InSAR time series by Ebmeier (2016). Therefore, we aim to construct an algorithm based on the BSS premise that latent sources of interest can be extracted from a time series of interferograms.

However, as ICA is one of a suite of BSS methods, we endeavour to ascertain which of the methods is best suited to InSAR data. We do this by first introducing the fundamentals of several BSS techniques (Section 2.2), explaining our novel method for applying a new BSS method to InSAR data (Section 2.3), and comparing the results of applying several BSS methods to a synthetic dataset (Section 2.4). In Section 2.5 we validate that the chosen method works with real data through the use of two time series of Sentinel-1 interferograms, before describing and demonstrating our prototype automatic detection algorithm in Section 2.6.

## 2.2 Blind Signal Separation

### 2.2.1 Interferograms as mixtures of signals

Interferograms consist of measurements of amplitude and phase at pixel locations. After correcting for geometric terms, the phase consists of contributions from several sources:

$$\varphi = W \{ \varphi_{def} + \varphi_{orb} + \varphi_{atm} + \Delta\varphi_{\theta} + \varphi_N \} \quad (2.1)$$

where  $\varphi_{def}$  is the phase change due to deformation of the ground surface,  $\varphi_{orb}$  is the phase due to errors in the location of the satellite at each acquisition,  $\varphi_{atm}$  is the phase change due to changes in the atmospheric delay,  $\Delta\varphi_{\theta}$  is the phase due to misestimation of the look angle,  $\varphi_N$  is the phase noise, and  $W$  is a wrapping operator that results in the phase lying between  $-\pi$  and  $\pi$  (Hooper et al., 2012). In geophysical applications  $\varphi_{def}$  is usually the signal of interest and a suite of methods exist to reduce the contributions from other terms. Signals that are considered to be dominated by deformation have been attributed to a variety of volcanic processes, including pre-eruptive inflation of a magma chamber, subsidence due to flank loading by new material, subsidence due to cooling of a magma body below a volcano, and subsidence due to changes in a volcano's geothermal system (Ebmeier et al., 2018).

If multiple latent signals combine in unknown quantities to form an interferogram, recovering the original signals can be viewed as a blind signal/source separation problem (BSS, Jutten and Herault (1991)). In this class of problems, the observed mixed signals are assumed to be generated using the following mixing model:

$$\mathbf{X} = \mathbf{A}\mathbf{S} \quad (2.2)$$

Where, using the standard nomenclature of BSS literature,  $\mathbf{X}$  contains the mixtures as row vectors,  $\mathbf{S}$  contains the unknown sources as row vectors, and  $\mathbf{A}$  is the unknown mixing matrix that combines varying amounts of the sources to create each mixture.

In this work, we refer to scalars as lower case letters (e.g.  $a$ ), row vectors as bold case lower letters (e.g.  $\mathbf{a}$ ), and matrices as bold upper case letters (e.g.  $\mathbf{A}$ ). The sources,  $\mathbf{S}$ , can be recovered if we could calculate the unmixing matrix,  $\mathbf{W}$ :

$$\mathbf{S} = \mathbf{W}\mathbf{X} \quad (2.3)$$

where:

$$\mathbf{A} = \mathbf{W}^{-1} \quad (2.4)$$

The preceding description of linear mixing can also be expressed in terms of Euclidean geometry and is demonstrated in Figure 2.1.  $m$  variables measured at  $n$  time points can be considered as points in an  $m$  dimensional ( $mD$ ) space (e.g. two  $10^5$  pixel interferograms would be  $10^5$  points in a 2D space). If the sources are stored in a data matrix ( $\mathbf{S}$ ) with each row containing a new variable, each column of this matrix is a  $mD$  position vector determining that observations point in the space. Matrix multiplication of these sources ( $\mathbf{S}$ ) by a mixing matrix ( $\mathbf{A}$ ) is equivalent to the inner product of the row vectors of  $\mathbf{A}$ , and the column vectors of  $\mathbf{S}$ . Consequently, the row vectors of  $\mathbf{A}$  can be considered as the axes defining a new subspace that the data ( $\mathbf{S}$ ) are projected into. The unmixing process ( $\mathbf{S} = \mathbf{W}\mathbf{X}$ ) can be considered in a similar way, with the rows of  $\mathbf{W}$  containing the basis vectors required to recover the sources. The goal of BSS is to find the basis vectors required to recover each source from the mixtures (i.e. the rows of  $\mathbf{W}$ ). However, it must be noted that the sign of the sources that are recovered by BSS methods (the rows of  $\mathbf{S}$ ) remains ambiguous, as the opposing sign may be present in the column of the mixing matrix that controls the strength of a given source in each mixture.

Linear mixing can be complicated by differences in the number of mixtures relative to the number of latent sources. In the simplest case, the number of mixtures is equal to the number of latent sources and the mixing and unmixing matrices ( $\mathbf{A}$  and  $\mathbf{W}$ ) are square. However, in what is termed the over-complete case (under-determined in inverse theory terminology) there are more latent sources than mixtures, and in the under-complete case there are more mixtures than latent sources (Amari, 1999). A time series of InSAR data at a subaerial volcano is likely to consist of tens to hundreds of interferograms but of substantially fewer latent sources (such as deformation and a topographically correlated atmospheric phase screen), and we therefore expect it to be under-complete.

In this line of reasoning we do not expect the unique turbulent atmospheric phase screens (APS) present in each interferogram to be recoverable as sources. An attempt

to recover the turbulent APSs would increase the number of sources to more than the number of mixtures (a significant change), and shift the problem from the under-complete to over-complete case. The sparse nature of the time course for each turbulent APS and the huge increase in the number of sources that would have to be sought leads us to discount the turbulent APSs as sources, and instead treat them as noise terms in the under-complete case. However, other atmospheric signals such as those due to atmospheric pressure gradients associated with weather systems or the rain shadow effect are not unique, and we may therefore expect to recover them as components. Methods to correct for atmospheric phase screens are also routinely applied to InSAR data, but as these can introduce new and erroneous signals, we apply BSS to uncorrected time series.

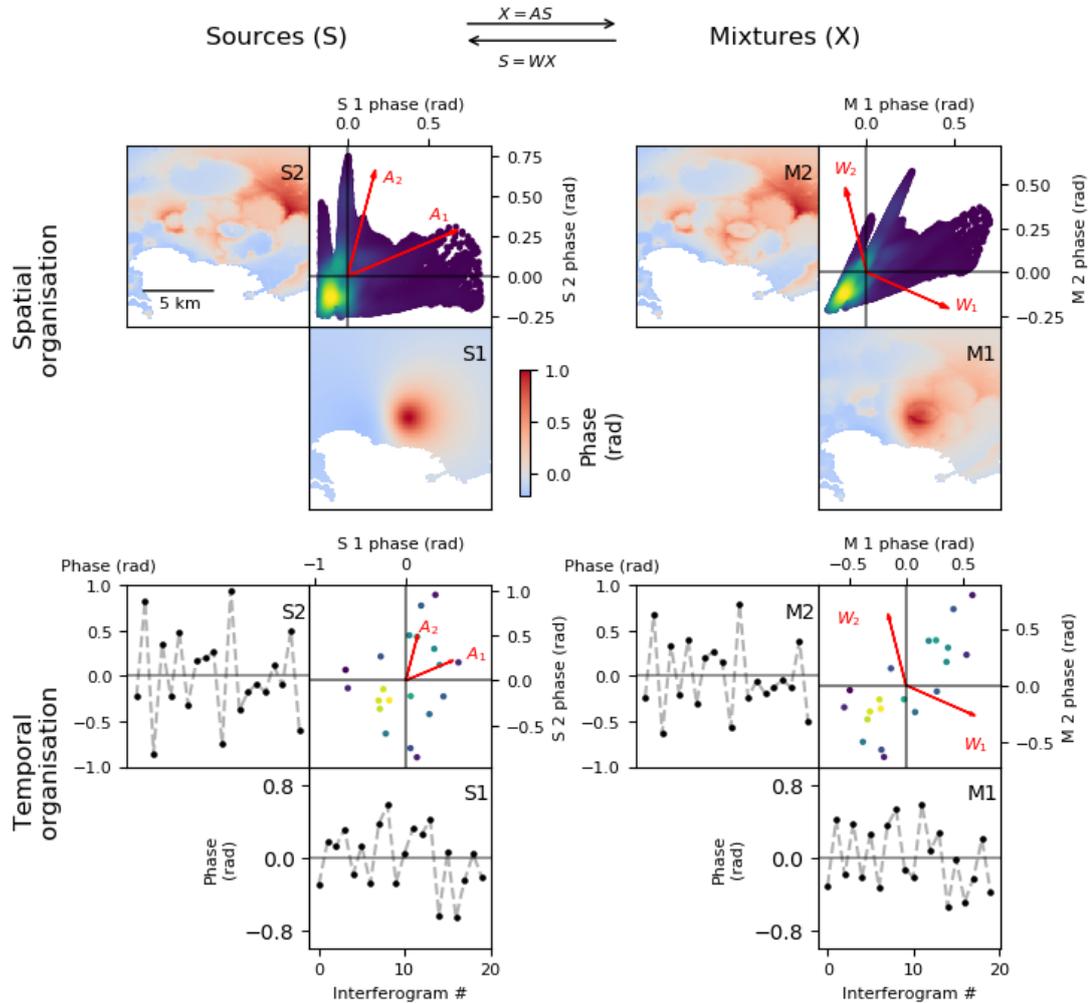
BSS for the under-complete case requires extra processing when certain algorithms are used and these caveats are discussed for each method. To implement BSS, we have investigated non-negative matrix factorization (NMF), principal component analysis (PCA), and independent component analysis (ICA). These methods are introduced in the following subsections, before our novel application of NMF to InSAR data is demonstrated in Section 2.3.

### 2.2.2 Organising and synthesising data

Before considering how to apply NMF, PCA, and ICA, we must first transform a time series of interferograms into a form that these methods can be applied to. The three methods consider the statistics of multiple observations of several variables and consequently do not require the spatial (or temporal) relationships between pixels (or interferograms) to be conserved. Therefore, the information contained within the time series can be converted to row (or column) vectors, providing that this reorganisation is performed consistently. We also refer to the time history of a spatial map (which can be as small as one pixel, or more commonly as large as a latent source) as a time course, in the style of BSS literature.

How these row vectors are formed has important implications, and consequently the two different approaches are termed as architecture I and architecture II in BSS literature (Bartlett et al., 2002). However, this nomenclature is opaque when applied to InSAR data and we instead refer to architecture I as spatial organisation, as when ICA is applied to architecture I data, the recovered latent sources are spatially independent. When ICA is applied to architecture II data the discovered latent sources are temporally independent, so we instead refer to architecture II as temporal organisation.

In spatial organisation, each image contains multiple realisations of a single random variable, with the number of random variables being equal to the number of images, whilst in temporal organisation, each pixel is a random variable, with as many observa-



**Figure 2.1:** Linear mixing for spatially and temporally organised data. Top row: Spatial organisation, in which a deformation signal (S1) is mixed with a topographically correlated signal (S2, generated from the topography of Campi Flegrei, Italy) through re-projecting the data in the directions  $A_1$  and  $A_2$ , to produce two mixtures, M1 and M2. The sources can be recovered from the mixtures by reprojecting the data in the directions  $W_1$  and  $W_2$ . In this architecture, the mixture space has as many dimensions as there are interferograms, and as many points in the space as there are pixels in the interferograms. Bottom row: Temporal organisation, in which a pixel with a phase change due to deformation over 20 epochs (S1) is mixed with a pixel with a phase change due to a topographically correlated atmospheric signal (S2) in a similar manner to the spatial case to produce two pixels (M1 and M2) that exhibit both deformation and atmospheric signals. In this architecture, the mixture space has as many dimensions as the interferograms have pixels, and as many points in this space as there are interferograms. Progression from blue to yellow is used to indicate areas of high point density in these (and any following) scatter plots.

tions as there are images. In this work, we adhere to the conventions of BSS literature and place variables as rows in our data matrix, and each observation of these variables occupies a new column. Therefore, for a time series of  $t$  interferograms each of  $p$  pixels, our data matrix would be  $t \times p$  for spatial organisation, and  $p \times t$  for temporal organisation. In Euclidean space, a time series organised spatially is likely to consist of  $10^1 - 10^2$  interferograms and require a space of equal dimensions, whilst for temporal organisation, an interferogram is likely to consist of  $10^3 - 10^7$  pixels and require a space of very high dimensionality. Figure 2.1 demonstrates the difference of spatial and temporal organisation, and how the (un)mixing vectors described in the previous section can be interpreted. As the number of interferograms controls the number of data points when the data are organised temporally, a hypothetical time series of only three interferograms would provide a very sparsely populated space in which to perform the analysis, and would be equivalent to the scatter plots in Figure 2.1 having only three data points.

To both introduce and compare PCA, ICA and NMF, we generate a synthetic time series using equation 2.2, with one matrix (either  $\mathbf{A}$  or  $\mathbf{S}$ ) containing the spatial patterns of the two synthetic InSAR signals, whilst the other matrix contains the strengths (termed time courses) with which each spatial pattern contributes to a each interferogram. To generate spatially-organised data, we postulate that the spatial maps are statistically independent and so place these in  $\mathbf{S}$ , whilst for temporally organised data we postulate that the time courses are statistically independent and so place these in  $\mathbf{S}$ . We choose our two synthetic InSAR signals to be ground deformation due to a volcanic process, and apparent ground movement due to a topographically correlated atmospheric phase signal (APS) as separation of signals of this type has been shown to be important (e.g. Delacourt et al. (1998)).

The spatial pattern of the two synthetic signals is shown in Figure 2.2 and, as these are sources when the data are organised spatially, provides insights into the challenge of recovering them. The area corresponds to Campi Flegrei (Italy), with the subaerial caldera walls apparent in the atmospheric signal. The surface deformation is created by modelling the inflation of a point source in an elastic half space (Mogi, 1958) as this has been used to successfully model observations of deformation at Campi Flegrei (Lundgren et al., 2001), and is strongly non-Gaussian by all the measures utilised. The spatial pattern for the topographically correlated delay is synthesized by assuming a linear relationship between phase and altitude for each pixel (calculated using the SRTM 30m DEM (Farr et al., 2007)).

The temporal nature of the two synthetic signals is also shown in Figure 2.2 and, as these are sources when the data are temporally organised, provides insights into the challenge of recovering them. The time course for the deformation is generated to approximate a period of inflation at a volcano by varying in strength around a

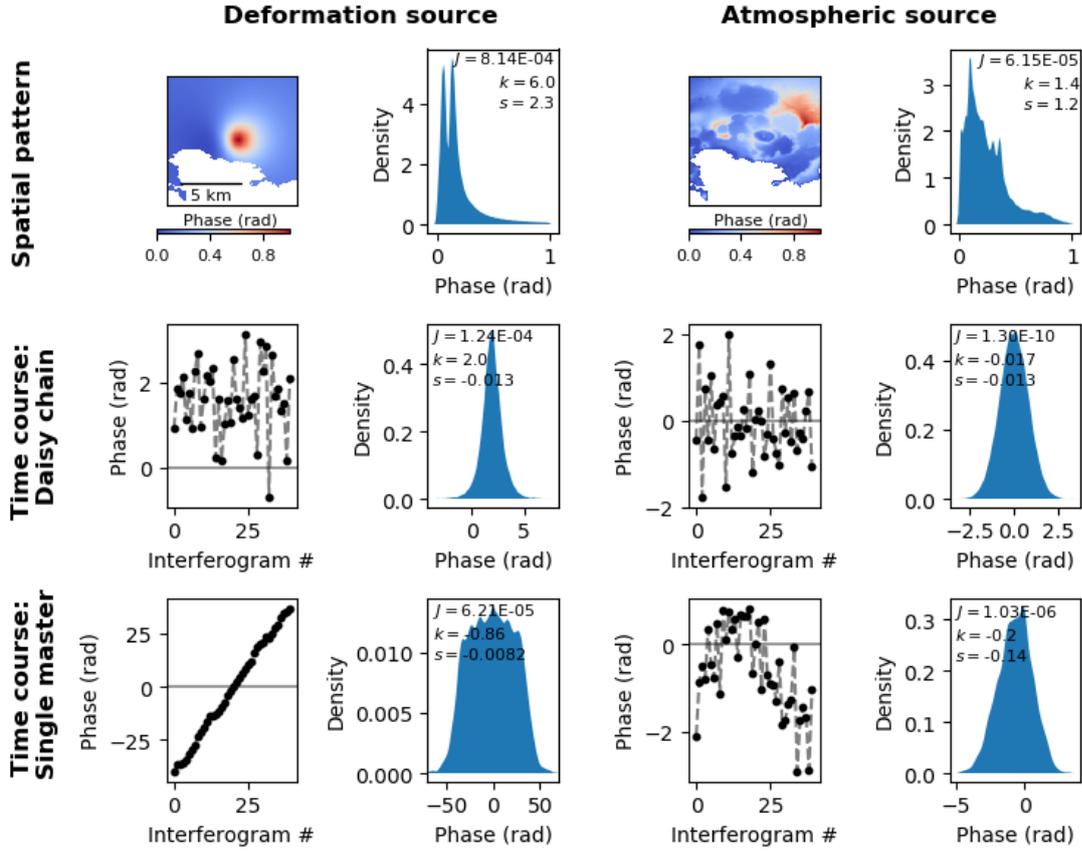
mean value that is above zero. The time course is generated by drawing values from a hyperbolic secant distribution with mean of 0.1 and variance of 1. This distribution is chosen as its excess kurtosis ( $k = 2$ ) is similar to that found during analysis of the size of changes in displacement over 6 day intervals at a variety of GPS stations ( $k = 2.4$ ) that experienced deformation due to volcanic unrest, using a method similar to that described in Liu et al. (2018). Whilst some of the deformation that occurred within these 6 day intervals was generally smaller than could be detected in a single interferogram, the use of GPS instead of InSAR as the source for this data allowed for these smaller signals to be more easily detected. Additionally, our choice of a time series approach is motivated by the ability to detect signals that may not be visible in a single 6 or 12 day interferogram, and therefore we believe that synthesising time courses that match these smaller magnitude signals creates more realistic synthetic data (see supplementary material for further details).

The time course for the topographically correlated tropospheric phase delay is synthesised as a sinusoidal wave (wavelength: 1 yr, amplitude: 4 rad/km) combined with Gaussian noise as this closely resembles the temporal evolution of the delay/elevation ratio at Colima Volcano measured by Pinel et al. (2011). Whilst this function may not be applicable to all volcanoes, it provides a challenging example for our synthetic tests as when a daisy chain of interferograms is formed, the long term sinusoidal trend is removed and only the synthetic Gaussian noise remains.

In Section 2.2.1 we also addressed why the signal introduced by the turbulent atmospheric phase screen (APS) could only be treated as a noise term (in contrast to recovering the turbulent APS for each interferogram). For synthetic tests, we generate turbulent APSs as spatially correlated noise, and then difference pairs of these to make signals that would be expected in either a single master or daisy chain time series. As the turbulent APS acts as noise and complicates the recovery of latent sources, we do not include it in the examples presented in this section. However, in the more complex synthetic tests performed in Section 2.4, a turbulent APS is included. Figures 2.3 and 2.4 show examples of spatially organised and temporally organised synthetic data, and the results of applying ICA and PCA to them.

### 2.2.3 Principal Component Analysis

Principal component analysis (PCA, also termed the Karhunen-Loève expansion (Karhunen, 1947), the Hotelling transform (Hotelling, 1933), and empirical orthogonal functions (Lorenz, 1956)) has been applied to InSAR data by several authors, such as Kositsky and Avouac (2010), Rudolph et al. (2013), Chaussard et al. (2014), and Kositsky and Avouac (2010) to isolate signals of geophysical interest. The first principal component is the direction in which the maximum amount of variance in the mean centered mixtures can be explained, the second seeks the same result given that it is orthogonal to



**Figure 2.2:** Sources and estimated PDFs for synthetic deformation source (left) and atmospheric source (right). Top row: The spatial pattern of the two sources, corresponding estimated PDF (Gaussian kernel density estimate), and measures of non-Gaussianity (Introduced in Section 2.2.4) of sources, where  $J$  is approximate negentropy (0 for Gaussian),  $k$  the excess kurtosis (0 for Gaussian), and  $s$  the skewness (0 for Gaussian). Campi Flegrei was used for this synthetic example, with the Gulf of Pozzuoli visible as the masked (white) pixels in the lower part of the image. The spatial pattern of the deformation source is highly non-Gaussian for all measures, whilst the atmospheric source is closer to Gaussian. Middle row: Temporal nature of the synthetic sources when a daisy chain of interferograms are created. The estimate of the PDF is generated from 100 synthetic time courses each of 40 interferograms, but for clarity only the first time course is shown in the left hand plots. Bottom row: Temporal nature of the synthetic sources when a single master is used to create the interferograms. The deformation time course appears non-Gaussian as it approximates a uniform distribution, whilst the atmospheric source is again closer to Gaussian. Note that the non-Gaussian nature of the temporal signals is a result of how they were synthesised, and that the temporal trend in the atmospheric source is due to the sinusoidal term described in the text, which is intended to approximate seasonal changes in the signal's strength.

the first, and so on. The data can then be projected in the basis defined by the principal components, with the result that the data are now uncorrelated. The principal components can be found in a variety of ways (e.g. singular value decomposition), but are routinely found by calculating the eigenvectors and eigenvalues of the mean centred data's covariance matrix. Once the eigenvectors have been calculated, the change of basis can be achieved using:

$$\mathbf{S}_{reco} = \mathbf{E}^T \mathbf{X} \quad (2.5)$$

where  $\mathbf{S}_{reco}$  are the reconstructions of the sources,  $\mathbf{E}$  is a matrix of eigenvectors as columns (termed  $\mathbf{W}$  in equation 2.3), and  $\mathbf{X}$  is the matrix of mixtures. A consequence of PCA requiring the data to be first mean centred, is that a reference pixel need not be chosen for the time series as all the interferograms are adjusted so that when the data are spatially organised each interferogram has a mean of zero, and when temporally organised the time history for each pixel has a mean of zero. The strength of PCA lies in the nature that the first principal components (or eigenvectors) contain the majority of the variance of the data, and so by discarding the later components, a large proportion of the variance of the data can be expressed in relatively few dimensions. However, due to the orthogonality of the principal components, it is apparent that should the rows of  $\mathbf{A}$  not be orthogonal, PCA cannot separate the two sources fully. Consequently, in some fields independent component analysis (ICA, Comon (1994) and Jutten and Herault (1991)) is preferred to PCA.

PCA can be applied to both spatially-organised data (to find uncorrelated images), and temporally organised data (to find uncorrelated time courses). However, when performing PCA on images (such as interferograms) organised temporally (that is, where each pixel is a variable), many datasets will have more dimensions than points populating the space (e.g. for a time series of 19 interferograms, each of 10,000 pixels, this would result in 19 points in 10,000 dimensions). Consequently, both the calculation of the covariance matrix (of size  $10,000 \times 10,000$  in this example) and its eigendecomposition becomes very computationally expensive. Previous applications of PCA to temporally organised interferograms by Ebmeier (2016) have circumvented this by spatially down-sampling the data, and so reducing the number of pixels and therefore dimensions. However, we instead calculate the principal components of these sparsely populated high dimensional spaces using the PCA “compact trick” (Solem, 2012), which is computationally efficient as it uses the constraint that when the number of dimensions is greater than the number of samples,  $s$ , there will only be  $s - 1$  eigenvectors. Considering PCA performed on two data points in  $3D$  (i.e. the case in which there are more dimensions than data), we can see that the direction of maximum variance lies along the line joining the two points. However, the second principal component cannot

be defined, as there are an infinite number of vectors which are perpendicular to the first in a  $3D$  space. Therefore, the number of principal components is limited to be one less than the number of data points, and when considering temporal InSAR data, as there are significantly more dimensions than data, much computational cost can be avoided through finding only the reduced number of principal components.

Minimal considerations of whether the data are complete or under-complete are required when applying PCA. In the complete case, all the eigenvectors are retained and  $\mathbf{E}^T$  is square. In the noiseless under-complete case with  $m$  mixtures and  $s$  latent sources, the observations of the mixtures will lie on an  $sD$  hyperplane (demonstrated in Figure 2.3) and the last  $m - s$  eigenvalues will be zero, indicating the corresponding eigenvalues can be discarded (reducing  $\mathbf{E}^T$  to an  $s \times m$  matrix). In the noisy under-complete case, the data create a volume of equal dimension to the space and all eigenvalues are non-zero, but further interpretation depends on the signal-to-noise ratio (SNR). In the case that the SNR is high enough, the first  $s$  eigenvalues are significantly larger than the remaining  $m - s$  and the number of sources can be identified, but as the SNR decreases, the difference between the two sets of eigenvalues decreases until separation and constraint of the number of sources becomes difficult. The difficulty in choosing the number of components to retain is demonstrated in McKeown et al. (1998), where a small but interesting signal was identified in a very low eigenvector. Spatially small signals contained within large (250 km wide) Sentinel-1 interferograms may contribute little to the overall variance of the time series and be at risk of omission due to occurring in a low ranked eigenvector, yet may be of geophysical interest. We expect the strength of a signal, the proportion of an image that it covers, and the number of interferograms that it features in, to determine how highly the signal is ranked within the eigenvectors extracted by PCA. Consequently, we postulate that in order to maximise the chance of a signal of a given strength and spatial size to be retained in the highest eigenvectors, the interferograms should be cropped to the smallest practical area of interest around a volcano. Given that deformation sources are commonly offset from a volcanic centre (Ebmeier et al., 2018) by tens of kilometres, the smallest practical area is likely to be around 20 – 50 km in size, but to remain a parameter that is tuned for each application of a BSS method to a volcanic centre.

## 2.2.4 Independent Component Analysis

ICA provides a new basis for the mixtures such that they are no longer statistically dependent (i.e. the independent components, or sources, have been recovered), and is the dominant method to apply BSS (Stone, 2002). In a similar fashion to PCA, it has been applied to InSAR data by Ebmeier (2016) in order to isolate signals of volcanic interest. ICA can be used to identify latent sources as these usually have probability density functions (PDFs) that are less Gaussian than mixtures. This is due

to the central limit theorem, which stipulates that by summing several independent sources, the resulting mixture has a more Gaussian PDF than any of the constituent sources (Hyvärinen and Oja, 2000). ICA algorithms measure the Gaussianity of linear combinations of the mixtures, before adjusting this mixture in order to create a signal that has maximum non-Gaussianity, and is therefore likely to be one of the original sources. This point also highlights the limitation that ICA algorithms cannot be applied to Gaussian signals (or random variables), as the latent sources would not be more non-Gaussian than the mixtures. The wealth of successful applications of ICA to BSS problems can be used to justify the expectation that it would outperform PCA as two physical processes that are unrelated (such as deformation at a volcano and atmospheric delay) are likely to be statistically independent, as opposed to merely uncorrelated. However, signals such as deformation and atmospheric delay at a stratovolcano may be physically unrelated, but are likely to be spatially similar. We explore the search for spatially uncorrelated and independent sources at a stratovolcano in more detail in Section 2.4.5.

Several measures of the non-Gaussianity of random variables exist and have been used by various ICA algorithms. Kurtosis provides one of simplest measures to implement and measures how “peaked” a distribution is. As Gaussian random variables have a kurtosis of three, it is common for kurtosis to instead refer to excess kurtosis, which is a measure in which all values are reduced by three, shifting the kurtosis of a Gaussian distribution to zero. A random variable with a high kurtosis (i.e.  $k > 0$ ) has a “peaked” or “spikey” PDF with many values closely grouped together and long tails, whilst a random variable with a low kurtosis (i.e.  $k < 0$ ) has a very broad PDF (such as a uniform distribution, with a kurtosis of  $-1.2$ ). Gradient descent can be used to maximise the non-Gaussianity of the signals to be recovered, but the more complex fixed-point iteration presented as FastICA in Hyvärinen and Oja (1997) has two advantages in that it has been shown to converge more quickly, and also does not require a learning rate to be chosen.

Whilst kurtosis is an intuitive and computationally efficient way of measuring the non-Gaussian nature of a random variable, it is not robust and can be heavily influenced by outliers (Hyvärinen, 1999). Therefore, differential entropy is used as a measure of non-Gaussianity in more recent versions of the FastICA algorithm (Hyvärinen, 1999). Differential entropy is a measure of the information that a variable conveys and is largest for more unpredictable (or random) variables, and lowest for more predictable ones (in the case of a coin that always landed on heads, the entropy would be 0). For a collection of random variables with equal variance, the maximum differential entropy is achieved by that with a Gaussian distribution (Hyvärinen et al., 2001). Therefore, it can be used as a measure of how Gaussian a random variable is. This can be simplified by defining a new quantity, negentropy, which is 0 for a Gaussian

random variable and always non-negative. However, as differential entropy (and so negentropy) requires the PDF of a variable to be known, it remains difficult to measure efficiently. Therefore, approximations for differential entropy have been developed for the FastICA algorithm, but a full discussion of these is outside the scope of this paper (see e.g. Hyvärinen (1998) and Hyvärinen et al. (2001)). As per using kurtosis, gradient descent can be used to maximise the approximated negentropy, but a faster and more robust fixed-point algorithm is presented in a newer version of FastICA in (Hyvärinen, 1999).

As per PCA, ICA can be applied to both spatially and temporally-organised data. When ICA is applied to spatially-organised data, spatially independent sources and unconstrained time courses are recovered, and the method is referred to as spatial ICA, or sICA. When ICA is applied to temporally-organised data, independent time courses and unconstrained spatial sources are recovered, and the method is referred to as temporal ICA, or tICA. The independent sources that ICA seeks to recover are similar to the uncorrelated sources that PCA seeks to recover, but as independence is a more robust measure than the uncorrelatedness which PCA seeks (discussed further in Section S3), we may expect the sources recovered by ICA to be more faithful reconstructions of the latent sources that generated the data. ICA has been applied across a variety of fields, ranging from tICA on speech data (Bell and Sejnowski, 1995), tICA on electroencephalographic (EEG) data (Makeig et al., 1996), sICA on functional magnetic resonance imaging (fMRI) data (McKeown and Makeig S, 1998), and sICA on facial images (Bartlett et al., 2002). More recently, sICA and tICA have also been performed on geophysical data using the FastICA algorithm (Frappart et al., 2010; Ebmeier, 2016; Chaussard et al., 2017; Cohen-Weber et al., 2018).

Unlike PCA, applying ICA to under-complete data requires extra considerations. FastICA can only operate with square mixing and unmixing matrices and, whilst this makes them very suitable for the complete case, it makes their application to the under-complete case more challenging. An example of performing ICA with rectangular matrices is presented in Porrill and Stone (1998) and termed under-complete ICA (or uICA), whilst the error-gated Hebbian rule (EGHR) proposed by Isomura and Toyozumi (2016) retains square mixing and unmixing matrices, but is capable of recovering repeated versions of the original sources if the data are under-complete. However, we choose to implement the commonly used method of dimension reduction, as the uICA method presented in Porrill and Stone (1998) has not been applied to data similar to geophysical data, and the EGHR ICA algorithm of Isomura and Toyozumi (2016) requires tuning of a learning rate, which may prohibit the automating of the implementation of ICA that is required for ICA to be used in an automatic detection algorithm.

Dimension reduction seeks to compress the data by expressing it using a new and

smaller set of variables which, in the ideal case, are still able to convey the essential features of the data. Using this as a preprocessing step for ICA allows us to reduce the number of input variables to equal the number of sources that we wish to recover, and then allows ICA to be performed on the new lower dimension data (i.e. where  $\mathbf{A}$  and  $\mathbf{W}$  are now square). PCA is commonly used for this process, but includes the important caveat that some signals may be discarded in dimensions that were thought of as unimportant (McKeown et al., 1998). Most ICA algorithms require the data to be whitened (or sphered) prior to ingestion as this greatly simplifies the unmixing matrix as it changes from a full rank to orthogonal matrix (Hyvärinen and Oja, 1997), and as PCA can be used to perform whitening, it is usually incorporated into the dimension reduction step. As tICA requires temporally organised data to be whitened, we again use the “compact trick” (discussed in Section 2.2.3) to allow tPCA to be performed without downsampling the data.

### 2.2.5 PCA and ICA example

To demonstrate the functioning of PCA and ICA we present a low dimension example of under-complete data (three mixtures of two sources). Whilst these are of significantly lower dimension than would be encountered for real data, they remain useful examples as they allow for the data and (un)mixing vectors to be plotted clearly, and the role of PCA as a pre-processing step for ICA to be demonstrated. For the time series to be of such low dimension that it can be plotted in 3D, spatially-organised data must consist of three interferograms, whilst temporally organised data must consist of three pixels (in a similar fashion to the 2D case shown in Figure 2.1). Whilst three mixtures could be generated from more than two sources, this is the opposite of the under-complete case that we expect to encounter with InSAR data. Therefore, we use two sources and a  $3 \times 2$  mixing matrix ( $\mathbf{A}$ ) to generate under-complete data. To generate spatially organised data, the spatial pattern of the two sources (deformation and topographically correlated delay) are placed as rows in  $\mathbf{S}$ , and the time courses placed as columns in  $\mathbf{A}$  (see equation 2.2), whilst for temporal data, the opposite is performed.

Figure 2.3 shows the Euclidean representation of mixing sources to create spatially organised data, and the results of using PCA and ICA to attempt to recover the sources. As only two sources and a negligible amount of noise (visible as PC3 in Figure 2.3) are used to generate the data, the data lie on a plane in the 3D space of the three mixtures. The first two principal components (PCs) lie in this plane, and due to the orthogonality constraints imposed on the un-mixing vectors of PCA, the third lies perpendicular to the plane. This results in the variance accounted for by the third PC being approximately zero, and we are therefore able to conclude that the three mixtures were generated by two sources. To perform ICA, we rescale the data projected in the direction of the first two PCs such that the variance in each direction is one (i.e. the data are whitened),

and then seek two un-mixing vectors in this 2D space. Should noise be added to the above case, the three mixtures would no longer lie in a plane and would instead create a volume within the 3D space. If the signal is of significantly larger magnitude than the noise, identifying the plane in which the sources lie remains possible, and ICA is able to recover the sources accurately. However, as more noise is added, the plane that the sources creates becomes harder to identify using PCA, and if the vectors found by PCA are not aligned correctly (i.e. the first 2 lie in this plane), some signal may be present in the 3rd PC. When this component is discarded the signal it contained is also lost, and ICA may therefore fail to recover the original sources faithfully.

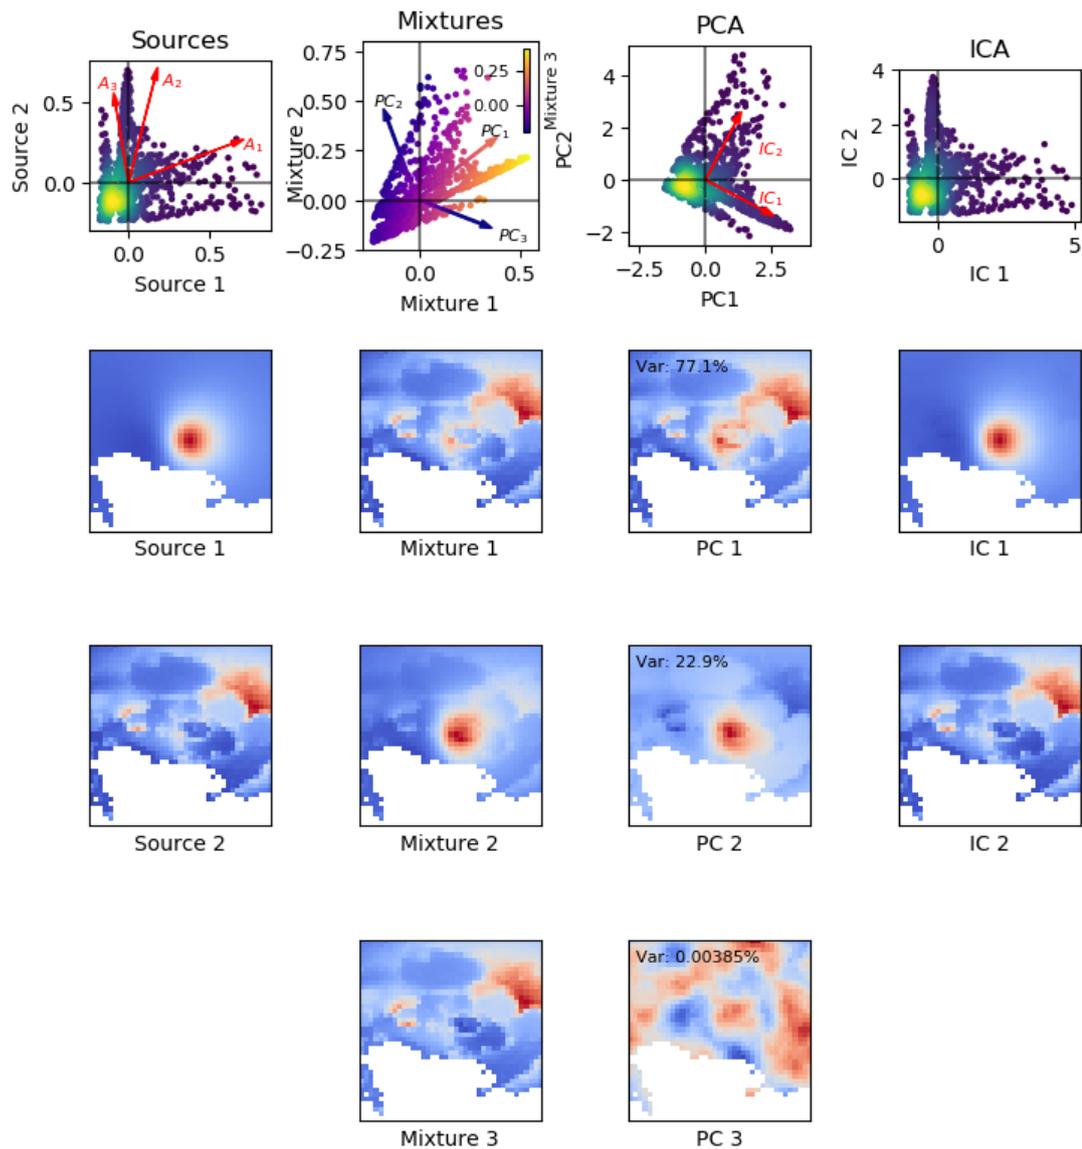
Figure 2.4 shows the Euclidean representation of mixing sources to create temporally organised data, and the results of using PCA and ICA to attempt to recover the sources. As this scenario is limited to 3 pixels, these do not have any spatial meaning, so the elements of  $\mathbf{A}$  are chosen from uniform distributions  $-0.5$  to  $+0.5$ . In a similar manner to the spatial example discussed in the previous paragraph, the very low variance of the third principal component ( $1.85 \times 10^{-15}\%$  of the total variance) indicates that it is numerical noise and should not be retained for use by ICA. Whilst comparison of the original and recovered sources is not as clear as for the spatial patterns, examination shows that the sources recovered by PCA are visibly different to the original sources, whilst those recovered by ICA are not.

### 2.2.6 Non-negative Matrix Factorisation

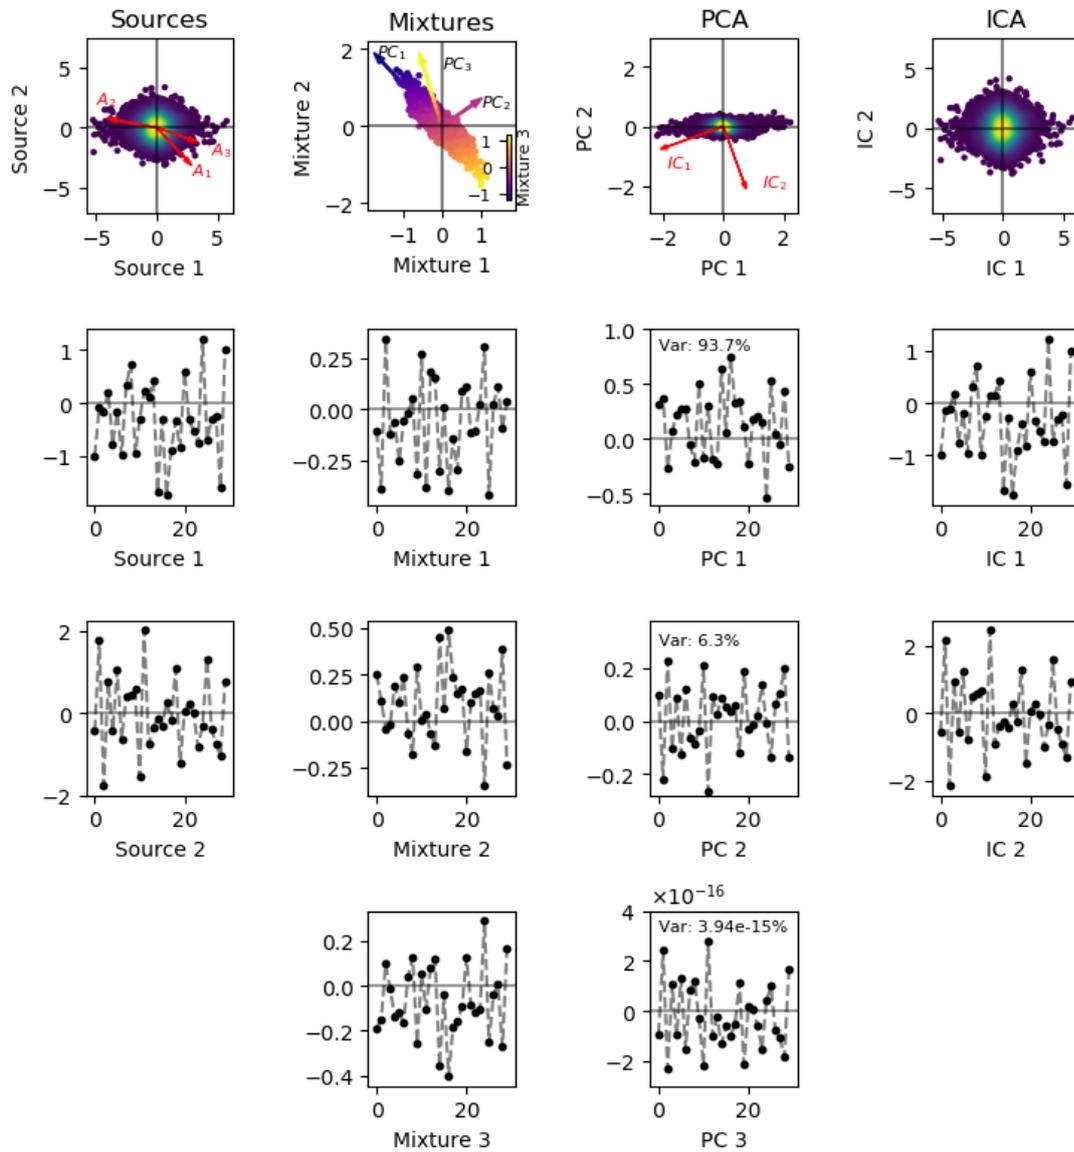
Non-negative matrix factorisation (NMF) developed from positive matrix factorisation (Paatero and Tapper, 1994) and factorises a non-negative data matrix of mixtures,  $\mathbf{X}$ , into two non-negative matrices,  $\mathbf{A}$  and  $\mathbf{S}$  (often termed  $\mathbf{W}$  and  $\mathbf{H}$  in NMF literature). Non-negative matrix factorisation became well known when Lee and Seung (1999) showed that a collection of 2500 facial images, each of 361 pixels, could be expressed as 49 sources that were easily interpretable to a human observer (i.e. the sources corresponded to parts of a face, such as eyes or a nose). In terms of Euclidean geometry, this corresponds to a 361 dimensional space populated by 2500 points being condensed to a 49 dimensional space, and corresponds with the under-determined case described in Section 2.2.1. However, to our knowledge, NMF has not been applied to InSAR data.

The multiplicative update rules of Lee and Seung (1999) find the local minimum of an objective function that measures the misfit between  $\mathbf{X}$  and  $\mathbf{A} \times \mathbf{S}$ . If the mixtures are linear combinations of a smaller number of latent sources, gradual identification of these sources (and their mixing matrix) occurs as the objective function reduces. In the case of noiseless data, the objective function may approach 0.

However, in the complete case (i.e. the number of sources and mixtures are equal)



**Figure 2.3:** Column 1: A scatter plot comparing the values for each pixel in the two sources (row one), and the two sources (rows two to three). These are mixed with the 3 mixing vectors of  $\mathbf{A}$  (plotted in red). As per the previous scatter plots, yellow indicates areas of high point density. Column 2: Scatter plot of the 3 mixtures (top), and the three mixtures. As the data lie on a plane in the 3D space, colour is used to for the third axis. The three 3D un-mixing vectors found by PCA are plotted with the colour indicating their vertical component.  $PC_1$ , and  $PC_2$  progress through the colour spectrum of the vertical axis at the same rate as the data points, and consequently lie in the plane. However,  $PC_3$  progresses through the colour spectrum more rapidly as it lies perpendicular to the plane. Column 3: The three components recovered by PCA, and the variance in each direction. As  $PC_3$  can be seen to correspond to noise, only components 1 and 2 are retained for use by ICA. The upper scatter plot shows  $PC_1$  and  $PC_2$  rescaled in such a way that the variance in each direction is 1 (i.e. the data are whitened). The un-mixing vectors found by ICA are shown in red. Column 4: The two sources recovered by ICA, which can be seen to be a good approximation of the original sources. Note that due to the ambiguity of the sign of the recovered sources found by BSS methods, some of the recovered sources are sign-flipped versions of the original. However, to aid in quick analysis of the recovered sources, we remove any obvious sign flips from figures within this work.



**Figure 2.4:** Column 1: The two temporal sources that are mixed with the 3 mixing vectors of  $\mathbf{A}$  (plotted in red). As per the previous scatter plots, yellow indicates areas of high point density. Column 2: Scatter plot of the 3 mixtures (top), the unmixing vectors found by PCA, and the three mixtures. These show the same features as those discussed in detail in Figure 2.3, but the salient point remains that the points lie on a plane, with  $PC_1$  and  $PC_2$  in the plane, and  $PC_3$  orthogonal to it. Column 3: The three components recovered by PCA, and the variance in each direction.  $PC_2$  can be seen to recover source 2 well, but  $PC_1$  is a poor approximation of source 1. As component 3 can be seen to correspond to noise, only components 1 and 2 are retained for use by ICA. The un-mixing vectors found by ICA are shown in red. Column 4: The two sources recovered by ICA. Visual inspection shows these to be a good approximation of the original sources.

any positive and non-linearly dependent vectors can be used to fit the data exactly, providing the values of  $\mathbf{S}$  are adjusted correctly by the algorithm. Consequently, the original sources are not recovered, and the data are instead reconstructed from near random permutations of the latent sources. This situation can be demonstrated by considering observations of 3 mixtures that are fit when the rows of  $\mathbf{A}$  contain vectors in the direction of the space's coordinates (i.e.  $[1, 0, 0]$  etc.), and the 3 recovered sources would be the 3 mixtures. This arrangement would be able to reconstruct the data exactly, but would not recover meaningful or accurate sources. Therefore, unlike PCA and ICA, NMF can only be applied to under-complete data.

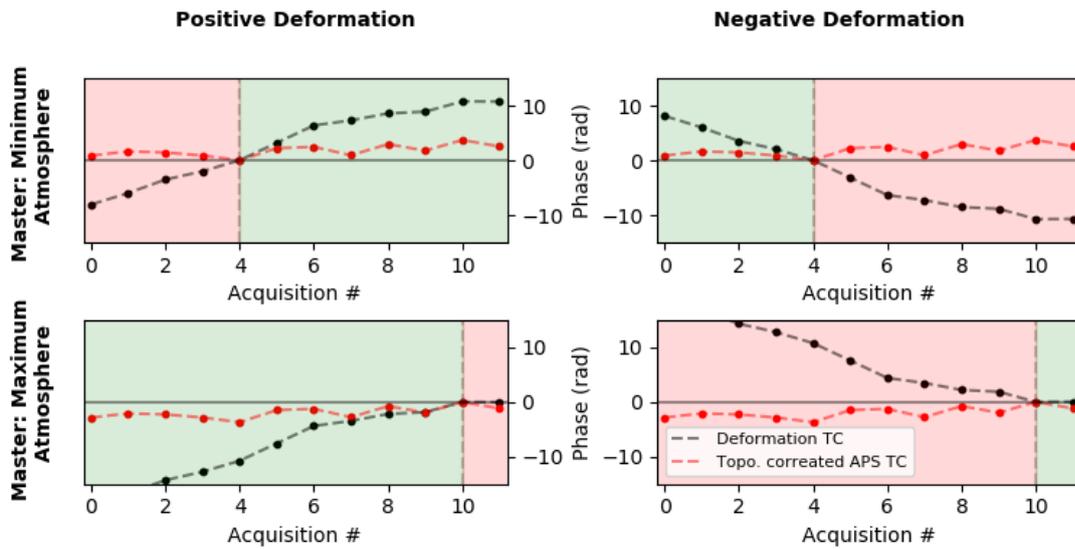
### 2.3 Application of NMF to InSAR time series

As NMF cannot be used on data that contains negative values, we instead construct a time series of interferograms relative to a single master. This master image is chosen to be the date on which either the strongest/weakest topographically correlated atmosphere occurs as this ensures that this signal is either entirely positive/entirely negative. We synthesise a volcanic signal that is steadily inflating/deflating at a varying rate, to ensure that all the data after/before the master is positive. NMF can then be applied to regions of the time series in which both signals are positive, or by applying a trivial sign flip, both signals are negative. Figure 2.5 demonstrates these possible scenarios which, whilst somewhat limited, we believe remain useful. Dates on which strong topographically correlated APS signals occur may be estimated from outside deforming regions, but more complex methods are required for cases in which the deformation is not of a constant sign before or after the master image, which are outside the remit of this initial study.

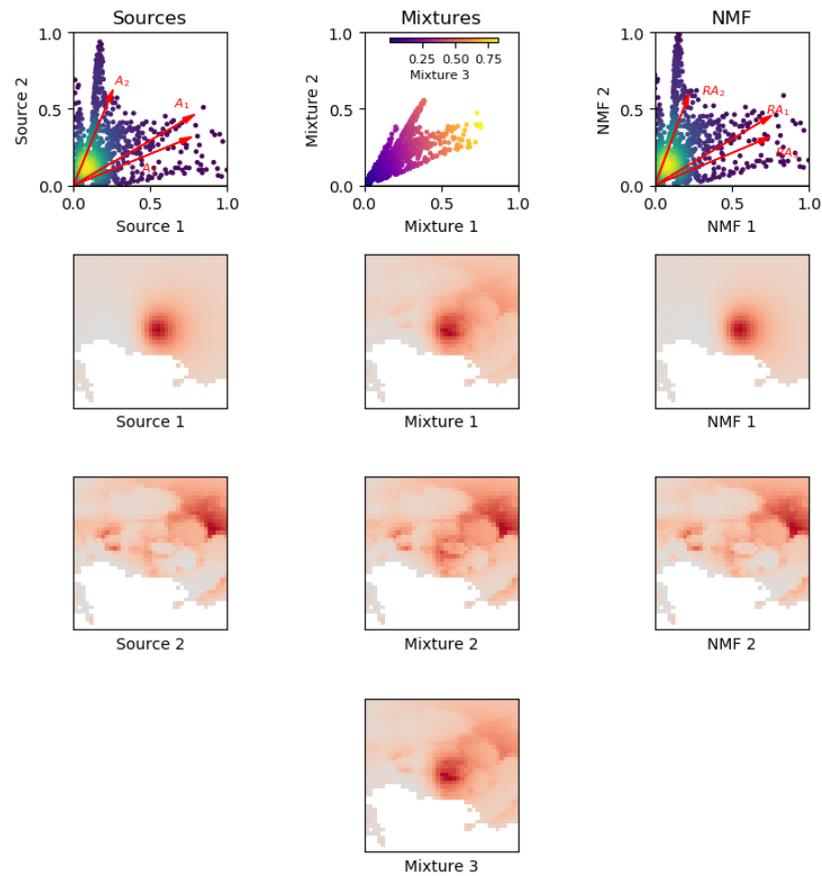
Figure 2.6 shows the results of applying NMF to spatially-organised under-complete 3D data that corresponds to any of the green regions of Figure 2.5. Very slight correlations remain between the two sources due to incomplete separation, but NMF can be seen to recover the sources well (mean residual per pixel: 0.0001). For brevity, we do not show the results of applying the same process to temporally organised data as it is very similar to Figure 2.4. The addition of noise in the form of a turbulent APS and any associated negative values is addressed in Section 2.4.1 through adding small values to all of the pixels within an interferograms (or time points within a time courses, if the data are temporally organised).

### 2.4 Comparison of PCA, ICA and NMF

To determine which of the BSS methods is most suited to isolating signals of geophysical interest, we construct a collection of differing time series in which we vary whether (1) the interferograms are constructed as a daisy chain or relative to a single master, (2) the



**Figure 2.5:** Possible scenarios for application of NMF to single master InSAR time series. Top left: Uplift (black points) and a topographically correlated APS (red points) with the master interferogram on the acquisition with the minimum atmosphere (number 4). Data after the master (green shading) is non-negative and can be used, whilst data before the master contains both positive and negative data and cannot be used (red shading). Top right: As above, but with subsidence. Data before the master data is non-negative. Bottom left: Uplift and a master chosen for the maximum atmosphere (number 10). The data before the master is non-positive, and can be utilised by NMF providing a trivial sign flip is performed. Bottom right: Subsidence and a master chosen for the maximum atmosphere. A sign flip of the data after the master facilitates application of NMF.



**Figure 2.6:** Results of application of NMF to spatially-organised data. Column 1: To create 3 mixtures from 2 sources, 3 2d vectors are required (rows of  $\mathbf{A}$ ). Column 2: Scatter plot of the 3 mixtures. The points form a plane in the 3d space which allow them to be visualised adequately using colour for the 3rd axis. Column 3: The 2 recovered sources and mixing vectors, showing the near exact recovery of the original sources.

number of sources recovered, (3) the length of the time series, and (4) the strength of the turbulent atmospheric phase screen (APS). These time series consist of three types of synthetic signals aimed to mimic those seen in real data, with one corresponding to a deformation source, a second to a topographically correlated APS, and a third to a turbulent APS. Whilst this is not a comprehensive list of types of APS seen in InSAR data, it contains enough complexity for illuminating experiments with BSS methods to be performed.

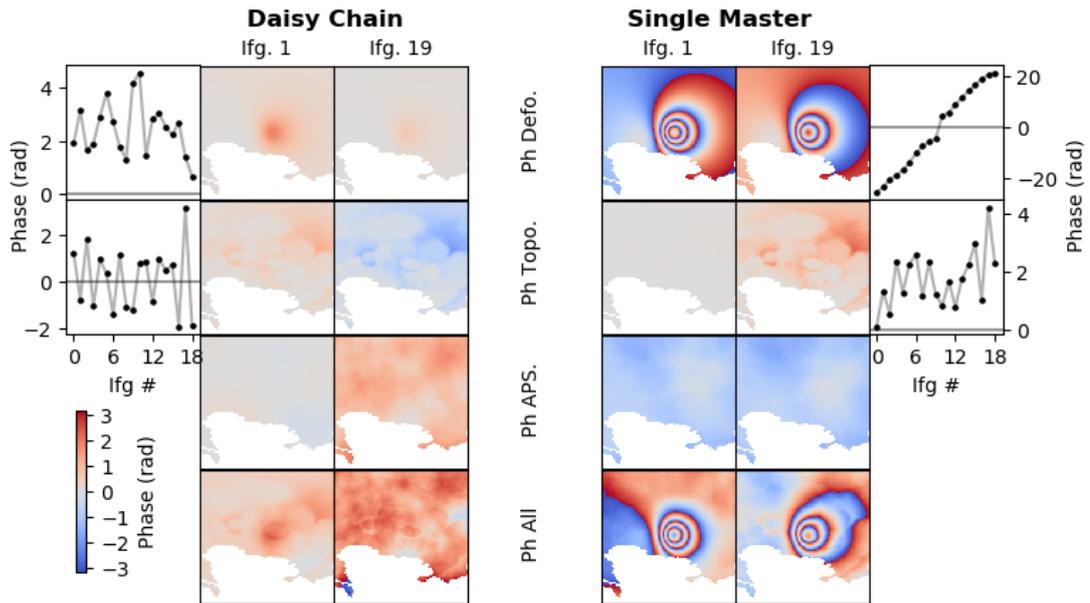
To evaluate the fidelity of the sources recovered by a BSS method, we first multiply each time course by its associated spatial map to construct a time series of only that signal. The mean residual per pixel is then calculated between each synthesised source and each recovered source, before finally calculating the total residual when different recovered sources are used to match each synthesised source (i.e. one recovered source cannot be used to fit both synthesised sources).

#### 2.4.1 Single master vs daisy chain time series

When creating interferograms from SAR acquisitions, the interferograms can be constructed to show the signal between two acquisitions with a short spatial or temporal baseline, or between each acquisition to a single master acquisition (see Hooper et al. (2012) for a more complete description). In the case that spatial baselines remain small, interferograms can be created between acquisitions with the shortest temporal baselines, making what is often termed a daisy chain of interferograms. Figure 2.7 shows how the same signals manifest themselves when constructed using either the daisy chain or single master approach, whilst Figure 2.8 shows the results of using the BSS methods described previously on each case.

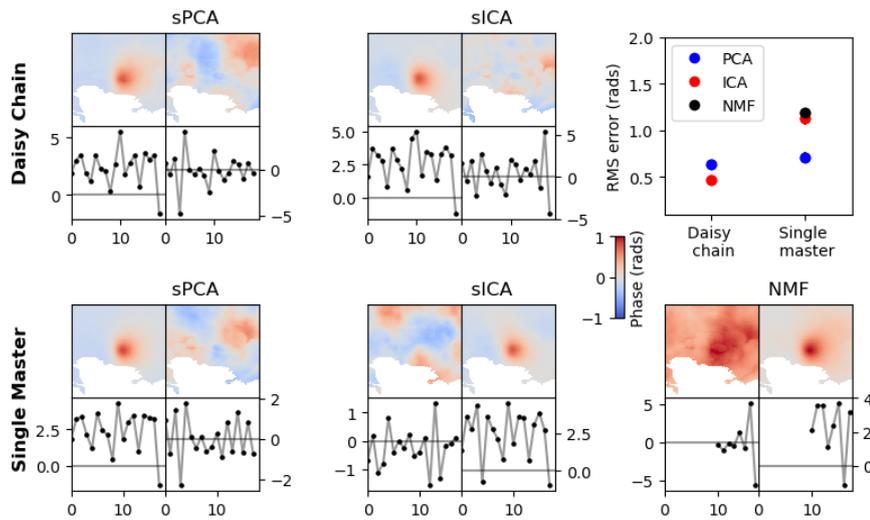
For the single master case, NMF can be seen to recover the sources poorly as even though the deformation source is successfully isolated in one recovered source, the second recovered source contains a mixture of the two synthetic sources. This incorrect separation creates a high residual (RMS error:  $\sim 1.2$  rad). We interpret this as a result of the addition of a turbulent APS in the more complex synthetic data that we use in this experiment (in contrast to that used in Section 2.3) introducing some pixels that are negative, which requires constants to be added to the entire interferogram to allow it to still be used with NMF. As the linear mixing model of NMF does not allow for this, our result of poor performance by NMF when a turbulent APS is added is unsurprising, and consequently we do not pursue its use through the remaining tests.

When applied to single master data, sPCA and sICA are also able to recover the deformation signal accurately and the topographically correlated APS poorly, producing a moderate to high overall error (RMS error:  $\sim 0.8$  rad and  $\sim 1.1$  rad, respectively). This result seems plausible as when the data are organised relative to a single master,



**Figure 2.7:** Comparison of a synthetic time series of interferograms arranged as either a daisy chain of interferograms (left), or relative to a single master (right) formed from 20 acquisitions. The top row shows the signal due to inflation of a Mogi source (and its associated time course), the second due to a topographically correlated atmospheric signal (and time course), the third due to a turbulent APS, and the fourth due to the combination of these. In the single master time series, the master date was chosen to be when the topographically correlated atmospheric signal was a minimum, in order to ensure it remains positive in all the interferograms and NMF can be used. The phase is wrapped to between  $-\pi$  and  $\pi$  as the scale varies significantly between the two cases. In the daisy chain case, the signal due to the Mogi source is of comparable magnitude to the other signals, but in the single master case dwarfs them. Figure 2.8 shows the results of applying the suite of BSS techniques to each case.

the deformation signal becomes around an order of magnitude larger than the topographically correlated APS signal (see the time courses on the upper right of Figure 2.7). In the daisy chain case, both synthetic signals have a comparable magnitude and are recovered well, producing a low overall error (RMS error:  $\sim 0.5$  rad). We therefore conclude that through organising the data in a daisy chain both atmospheric and deformation signals retain comparable variances, and are more accurately recovered than in the single master case. Consequently, we select this method as optimal and use it in the following scenarios. It should be noted that when comparing the RMS error, both the recovery of the deformation source and the topographically correlated APS are weighted equally, as we wish to judge the algorithms not only on their ability to isolate deformation, but to also constrain the strength of the topographically correlated APS in any given interferogram.



**Figure 2.8:** The results of applying PCA, ICA, and NMF to the spatially-organised version of the data (i.e. sPCA and sICA) shown in Figure 2.7. To ease interpretation and comparison of the recovered time courses, the single master cases are differentiated to produce the equivalent daisy chain time courses. Application of the BSS methods to temporally organised data (i.e. tPCA and tICA) produced very poor results and so are omitted for clarity, but further figures (e.g. 2.9, 2.10, and 2.11) show these methods. The time courses recovered by NMF (lower right) only show the signals after the master date as these signals are predominantly positive, but are generally very poor reconstructions of the synthetic sources. In the daisy chain case, both PCA and ICA are seen to recover both the sources well, whilst in the single master case PCA and ICA are seen to recover the deformation source well and the topographically correlated APS poorly.

### 2.4.2 Number of sources sought

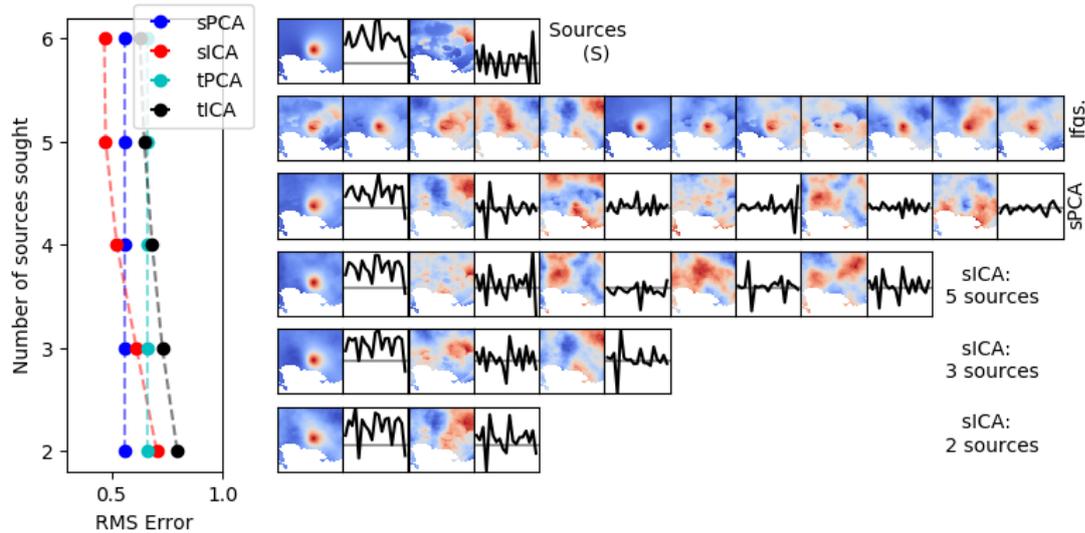
In Section 2.2.1 we introduced the under-complete nature of InSAR data, and discussed the importance of PCA as a pre-processing step for ICA. In the noiseless case introduced in Figure 2.3, using PCA to reduce the dimensionality of the data to equal the number of sources was trivial as the sudden change in variance between the 2nd and 3rd components suggested that the data formed an almost flat 2D feature in the 3D space. In the example shown in Figure 2.9, the data are very under-complete (25 interferograms from 2 sources) and contain significant noise (in the form of a turbulent APS), which presents a more realistic challenge for PCA and ICA.

As the number of sources that PCA recovers does not vary, it is only applied once to produce one set of recovered sources and time courses. In contrast to the data used in the previous section, we increase the variance of the turbulent atmospheric signal by 10% to visibly relegate the topographically correlated APS to the 3rd and 4th principal component (shown in Figure 2.9). Consequently, when sICA is performed to recover two sources, the signal contained in the 3rd and 4th PCA component is discarded and sICA can only recover the deformation source, producing a high RMS residual. However, as sICA is performed on more of the first sPCA sources, the signal contained in the higher sPCA sources is available to the FastICA algorithm and the fidelity of the recovered sources increases (producing a lower residual). A consequence of this is that the ICA algorithm then also recovers some sources that are just noise (a turbulent atmospheric signal), but these can potentially be separated from those of interest either by eye, or using clustering methods (Ebmeier, 2016). Figure 2.9 shows that the minimum residual is found when three extra sources are sought.

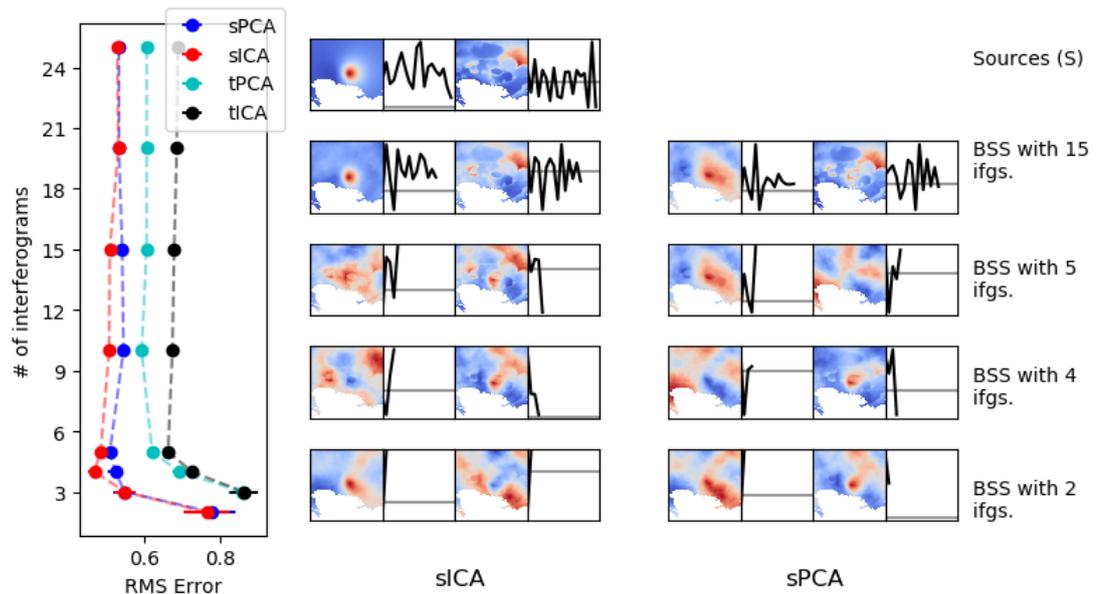
When applying ICA to real world examples, the problem becomes more complex as the number of sources is not necessarily known. However, this example shows that for ICA to perform optimally, the number of sources need be known only approximately so that the number to be recovered can be set correspondingly higher.

### 2.4.3 Length of time series

The length of a time series of interferograms over a volcanic centre can vary between a single interferogram spanning an event, to a time series spanning several years. Whilst the case of a single interferogram is not suitable for BSS methods to be applied, identification of the most suitable BSS method for time series of different lengths is important for an automatic detection algorithm. Consequently, in a similar fashion to the previous sections, we construct a suite of synthetic time series and crop these to different lengths to analyse the performance of different methods. In light of the findings of previous sections, we construct the time series as a daisy chain of interferograms and set the FastICA algorithm to retrieve 4 sources (i.e. 2 more than the 2 used to create



**Figure 2.9:** The results of sPCA, sICA, tPCA, and tICA applied to a time series of 25 daisy chain interferograms when the number of recovered sources is varied. The RMS residual between each synthesised and recovered case is shown on the left after averaging the results over 50 synthetic time series, whilst the right hand section shows the results from one particular time series. The two sources (top row) each have variances of 1 and are mixed with a turbulent atmospheric signal also of variance 1 to produce the time series of 25 interferograms in row 2. The third row shows the 6 largest sPCA components, with the first showing elements of the synthesised Mogi source, and the fourth showing a mixture of elements of the topographically correlated APS and noise created by the turbulent APS. The remaining 3 rows show the results of sICA when a decreasing number of sources are sought. As ICA does not place sources in a significant order (unlike PCA), the recovered sources most like the synthesised sources have been placed in the first 2 columns for clarity. sICA performs approximately equally well when 5 or 6 sources are recovered, which we attribute to being a consequence of part of the topographically correlated APS residing in PC4 and PC5 and only being accessible to the FastICA algorithm when the number of sources to be recovered is increased. In contrast to the preceding subsections, the length of the time series changes in this experiment, and consequently the number of interferograms over which the RMS residual is calculated increases. This causes the “well” seen at four interferograms, but if the RMS residual were to be standardised as “per interferogram” it would be seen to continue to reduce. However, this would prohibit comparison with the other figures presented within this section.



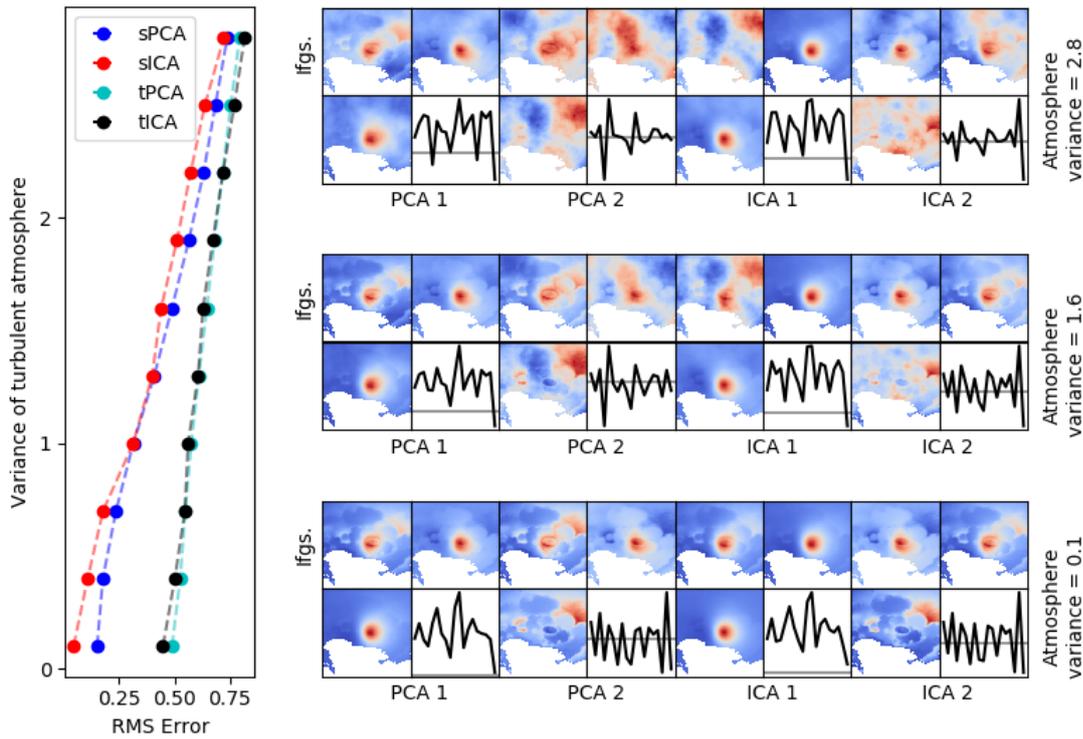
**Figure 2.10:** The results of sPCA, sICA, tPCA, and tICA applied to time series of interferograms of increasing length. The mean residual per pixel in the differing cases is shown on the left after averaging the results over 50 synthetic time series (with error bars showing the variance), whilst the right hand section shows the results from one particular time series. The top row shows the two synthesised sources and their associated time courses, whilst the second to fifth rows show the results of sPCA and sICA applied to time series of different lengths. Performance of all the BSS methods increases as the time series increases to around 5 interferograms in length, before then remaining constant. In the time series depicted on the right hand side, the results when using longer time series show some aspects of the turbulent APS signal relegating the topographically correlated APS to below the 2nd component.

the time series).

Figure 2.10 shows the results of applying sPCA, sICA, tPCA, and tICA to these time series. As per the previous experiments, application of PCA and ICA to temporally organised data produces poor results (mean residuals of 0.6 – 0.2, respectively), whilst to spatially-organised data produces significantly improved results (mean residuals of 0.1 – 0.03, respectively). sICA is seen to outperform sPCA at all time series lengths, though the difference decreases in magnitude as the time series increases in length.

#### 2.4.4 Strength of turbulent atmosphere contribution

In contrast to a topographically correlated atmospheric signal, a turbulent atmospheric signal cannot be retrieved by PCA, ICA, or NMF and consequently appears as noise in the time series. To determine whether PCA or ICA are best suited to dealing with time series with a strong turbulent atmospheric signal, we synthesise a suite of time series with differing strengths of turbulent atmospheric signal. In light of the findings of the previous section, we again arrange the time series as daisy chain of interferograms and seek 4 sources with the FastICA algorithm.



**Figure 2.11:** The results of sPCA, sICA, tPCA, and tICA applied to a suite of time series with different strengths of turbulent APS. The strength of the turbulent APS signal is quantified in terms of its variance which was set to 1 for the previous examples. The mean residual per pixel in the differing cases is shown on the left after averaging the results over 50 synthetic time series, whilst the right hand side shows 3 cases for differing strengths of turbulent APS. For each case, the top row shows the time series of interferograms (showing the differing contribution of the turbulent APS), whilst the second row shows the sources and associated time courses recovered by sPCA and sICA.

Figure 2.11 shows the results of applying sPCA, sICA, tPCA, and tICA to these time series. As per the previous experiments, the methods that are applied to spatially-organised data produce results with significantly lower residuals. Of these, sICA is seen to outperform sPCA in the majority of cases, although there are certain strengths of atmospheric noise at which the results are comparable.

#### 2.4.5 Spatial independence of sources

When using PCA and ICA to recover spatial maps (i.e. images) as sources, we are seeking sources that are either uncorrelated or statistically independent. Consequently, a key assumption of ICA is that the latent sources are statistically independent (or uncorrelated for PCA), which may not be the case for signals that may be encountered at certain volcanic centres. Figure 2.12 shows the results of performing PCA and ICA in a similar manner as described in the preceding sections, but instead of at Campi Flegrei, on Mt. Vesuvius (a  $\sim 1000\text{m}$  high stratovolcano  $\sim 15\text{km}$  east of Naples). In

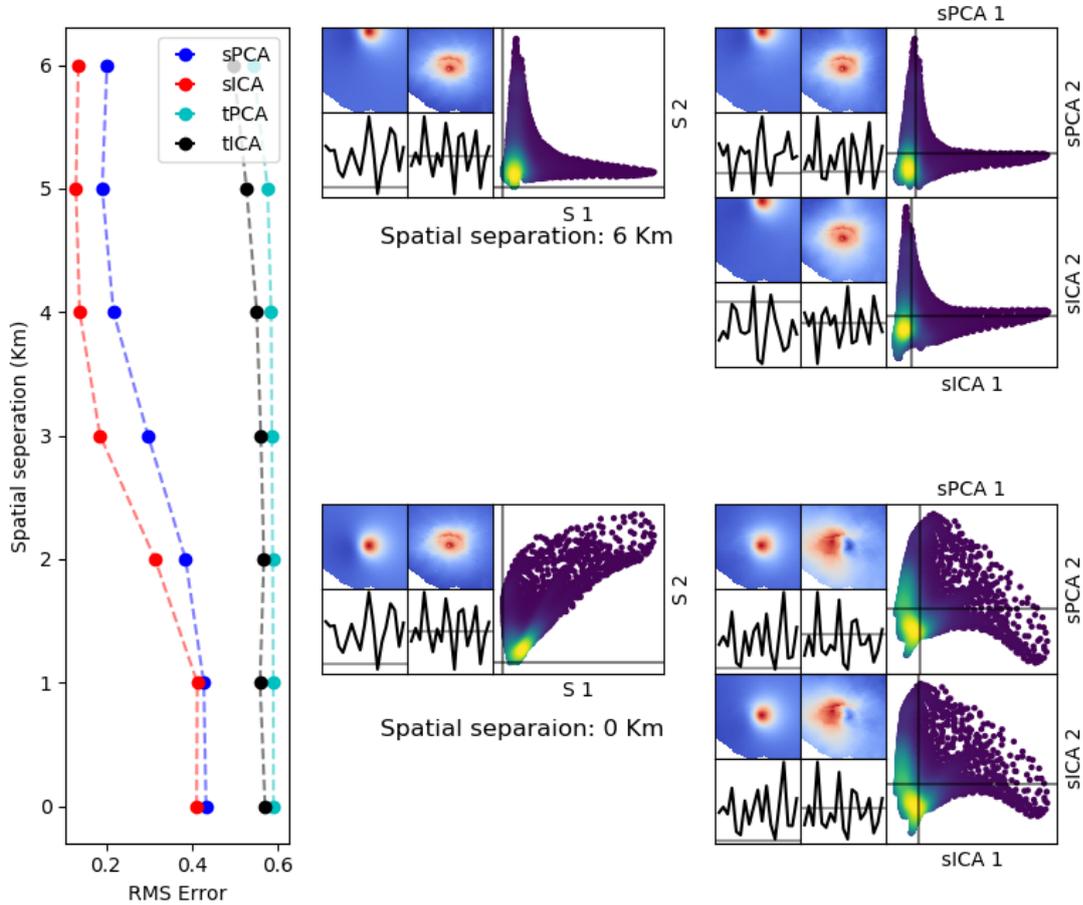
such a case, the topographically correlated APS is dominated by the signal produced by the cone of the stratovolcano and, if the synthetic deformation source is located under this cone, the two signals lie in the same location. When considering the pixels of the interferograms, those at the top of the cone are now likely have high values in both the deformation and topographically correlated APS, whilst those at the edges are likely to have low values in the two sources. Consequently, the two sources are now no longer statistically independent, and we would not expect PCA or ICA to be able to separate them. This result is seen in Figure 2.12, which shows that as the two synthetic sources are brought closer together, they cease to be statistically independent and the results of sPCA and sICA decrease in quality. This issue is discussed and explored in more detail in Section 2.5.2, in which sICA is applied to Mt. Etna.

Figure 2.12 also shows the results of performing tPCA and tICA on the data with non-statistically independent spatial maps. The results of these methods remain generally constant at all spatial separations, which we interpret as being due to PCA and the FastICA algorithm finding time courses that remain statistically independent regardless of changes to the spatial sources. However, in accordance with the preceding sections, the results of tPCA/tICA remain poor, and are worse than those found by sICA when the sources are not statistically independent.

#### 2.4.6 Comparison conclusions and method limitations

From the experiments carried out in the previous four subsections, we conclude from the suite of BSS methods studied that sICA is the most suited to use with InSAR time series. When sICA is used, performance is optimal when the number of sources sought is set to be slightly larger than the number expected to exist (such as recovering five sources when two are postulated to exist), and when interferograms are constructed to minimise the temporal baselines, such as through creating a daisy chain of interferograms.

However, limitations in the application of sICA remain. The performance of the algorithm reduces significantly when the spatial statistical independence of the sources is reduced, such as may happen at a stratovolcano where a topographically correlated APS and a broad deformation signal may be spatially similar. This issue is explored further through application of sICA to real data at a stratovolcano in Section 2.5.2. The performance of sICA is also limited in cases where noise from the turbulent atmosphere may dwarf signals of geophysical interest, but we do not discover any thresholds at which sICA fails, and instead see a gradual degradation of the accuracy of the recovered sources as the signal to noise ratio decreases.



**Figure 2.12:** The results of sPCA, sICA, tPCA, and tICA applied to a suite of time series in which the overlap of the two sources is varied. The left hand section shows the mean residual per pixel decreasing for sPCA and sICA as the spatial separation is increased, and changing little for tPCA and tICA. The upper right hand section shows the results with 6km of separation. The two sources can be seen not to overlap and no correlations are seen in the scatter plot comparing the two sources. The sources recovered by sPCA and sICA can be seen to be good reconstructions, and the associated scatter plots shows the same lack of correlation as the original sources. The lower right hand section shows the results with 0km of separation. The two sources overlap and cause many of the pixels to be correlated (as demonstrated in the scatter plot). The sources recovered by sICA are not good reconstructions as the scatter plot shows that the FastICA algorithm has sought sources that are statistically independent (and therefore uncorrelated, as seen in the scatter plot with the majority of the points now lying along the coordinate axis). The results from sPCA are similar, except with one source a negative version of the source found by sICA (causing the data in the scatter plot to be mirrored around the y axis).

## 2.5 Application to real data

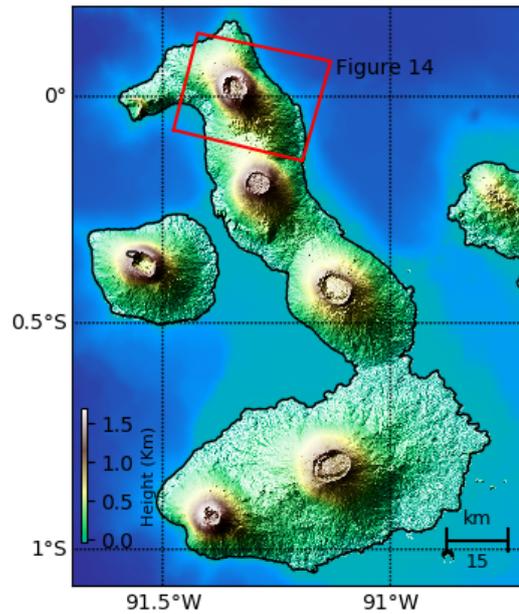
To further explore the ability of sICA to recover latent signals from a time series of interferograms, we present results from its application to two time series. The first spans the 2015 eruption at Wolf volcano (Galapagos archipelago, Ecuador) and was chosen as we further develop the use of sICA in an automatic detection algorithm that is able to detect the onset of this eruption. The second is centred on Mount Etna, and was chosen as an example of a stratovolcano at which application of sICA is likely to be problematic, due to the considerations of spatial independence discussed in Section 2.4.5.

### 2.5.1 Wolf volcano

Several existing studies present detailed results of modelling the observed surface deformations (e.g. Novellis et al. (2017) and Xu et al. (2016)), but we instead focus on the ability of sICA to automatically isolate the signals discussed in these papers. A detailed schematic of the timing of Sentinel-1 acquisitions and the two phases of the eruption is presented in Novellis et al. (2017), but we include the salient features and an overview map (Figure 2.13) here. The first phase involved the opening of a circumferential fissure on May 25th (2015) on the south-eastern caldera rim. This fissure produced two lava flows down the south eastern flanks of the volcano (Venzke, 2015), but by June 2nd activity ceased (Bernard et al., 2015). The second phase initiated around October 11th and involved an intra-caldera fissure with lava flows that covered the caldera floor (Bernard et al., 2015). The surface deformation associated with these events was attributed to two dykes (one circumferential and one intra-caldera) and two magma chambers ( $\sim 1\text{km}$  and  $\sim 5\text{km}$  below sea level) by Xu et al. (2016), and to one dyke and one shallow magma chamber ( $\sim 1.5\text{km}$  below the caldera floor) by Novellis et al. (2017).

We formed a time series of 20 daisy-chain descending Sentinel-1 interferograms covering the Isabela and Fernandina islands in the western Galapagos archipelago from December 13th, 2014, to October 21st, 2015. The unwrapped interferograms were formed using LiCSAR (González et al., 2016) and include filtering with a Goldstein filter (Goldstein and Werner, 1998). Figure 2.14 shows a subset of this time series focussed on the eruptive period, with pixels with a post-filtering coherence of  $< 0.8$  masked. The dominant features of this time series are range increase for the caldera floor (interpreted as predominantly subsidence) and range decrease for the area around the circumferential fissure (likely to be eastward motion). However, interferogram eight also shows a broader signal (more visible in the wrapped interferogram), that corresponds to a signal attributed to the deeper magma chamber by Xu et al. (2016).

Figure 2.15 shows the results of applying sICA to the time series. As around 3 – 4

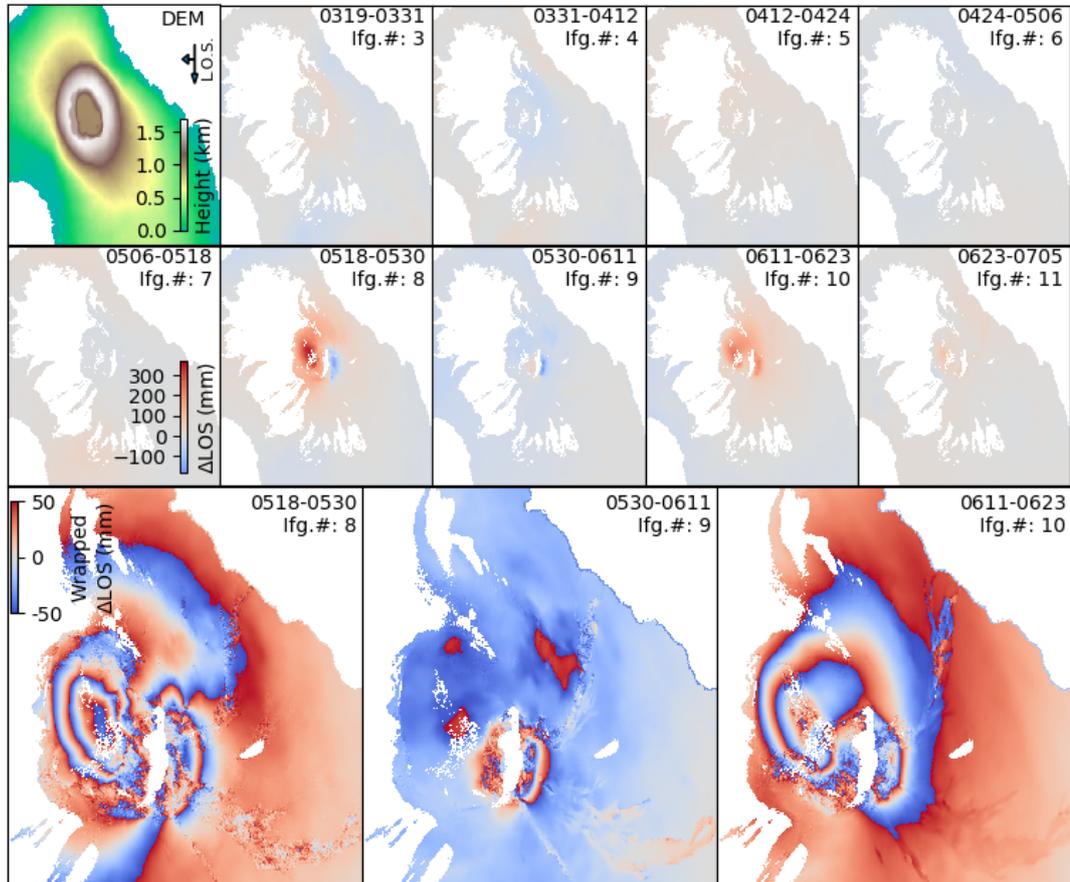


**Figure 2.13:** A shaded relief DEM of the western isles of the Galapagos Archipelago, with the calderas of 6 shield volcanoes visible as “upturned soup bowls”. Figure 2.14 focuses on Wolf Volcano, and the extent of the figure is depicted by the red region. Topography is taken from the SRTM DEM (Farr et al., 2007), and bathymetry from the GEBCO’2014 Grid (GEBCO, 2015).

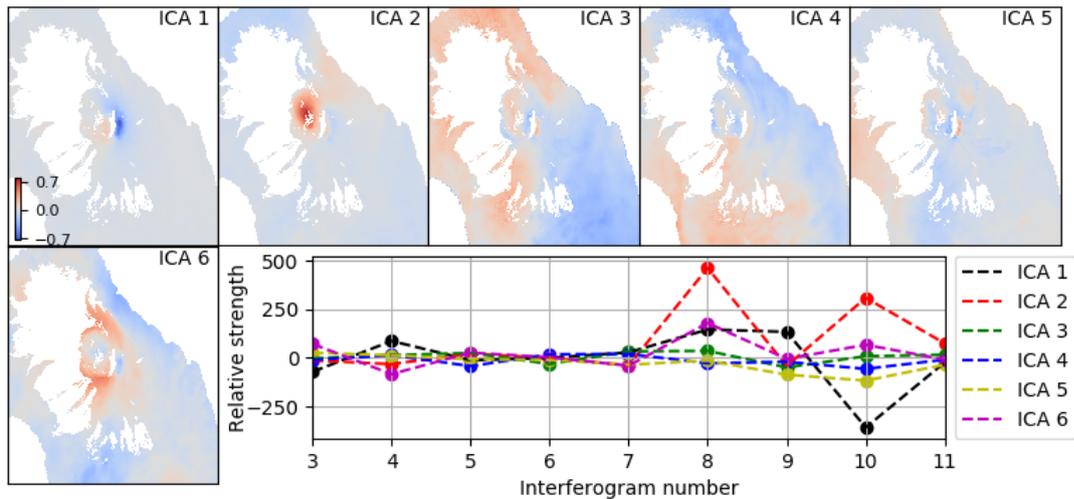
geophysical signals are expected, we set the FastICA algorithm to recover 6 components (in light of the results of Section 2.4.2). Visual inspection of the higher order principal components (6-20) suggests that these contain only turbulent atmospheric signals which are not at risk of containing important geophysical signals. sICA is able to recover the spatial pattern and time courses of the caldera floor subsidence, movement of the region surrounding the circumferential dike, and broad subsidence around the volcano. Recombination of the time courses and spatial patterns allows the time series to be reconstructed with minimal atmospheric signals.

### 2.5.2 Etna volcano

Mount Etna is a stratovolcano located on the eastern edge of the island of Sicily. It comprises of over  $\sim 3300\text{m}$  of elevation between its summit and eastern flanks, and is considered to be one of the world’s two most monitored volcanoes (González and Palano, 2014). InSAR has been used to measure deformation centred under the cone of Etna during the later portion of the 1991 – 1993 eruption (Massonnet et al., 1995), but a portion of the signals attributed to deformation by this study were later revised to be due to a topographically correlated atmospheric phase screen (Delacourt et al., 1998). Both the deformation and atmospheric phase screen (APS) described in these studies were centred under the topographic expression of this volcano and, as they



**Figure 2.14:** Top left: DEM showing the northern section of the island of Isabella (Galapagos archipelago), and the cone of Wolf Volcano with the satellite line-of-sight (L.O.S.) vector in the top right corner. Interferograms 3-11: A subset of the Sentinel-1 time series spanning the period of unrest that occurred in May and June of 2015. Numbers in the top left of each interferogram depict the dates of the two acquisitions that the image spans (mmdd-mmdd), pixels with an average coherence  $< 0.8$  are masked (predominantly removing the pixels on the western vegetated slopes of Wolf), and line-of-sight (LOS) change is measured in mm, with positive values indicating an increase in range (corresponding to subsidence of the ground). Lower section: Interferograms 8-10 re-wrapped to the interval  $-50\text{mm}$  to  $50\text{mm}$  in order to highlight more subtle features of the deformation pattern (such as the broad deflation signal in interferogram 8).



**Figure 2.15:** Results of sICA applied to the time series shown in Figure 2.14, showing the 6 components recovered, and the strength of each one throughout a subset of the time series (lower right). We interpret components 1,2, and 6 as representing deformation, and the remainder as representing atmospheric signals. Component 1 appears to capture the signal near to the circumferential fissure, component 2 the subsidence of the caldera floor, and component 6 the broad subsidence associated with the deeper chamber. The remaining signals (3-5) contain traces of the other signals (such as the circumferential dike signal), but we interpret them as containing predominantly atmospheric signals.

are therefore unlikely to be spatially independent, we expect application of sICA at Etna to be challenging. Subsequent geodetic studies at Etna have also measured other deformation processes that may be recoverable by sICA, such as eastward movement of sections of the faulted eastern flank (Solaro et al., 2010), and westward movement of the western flank (Aloisi et al., 2007; Lundgren and Rosen, 2003)

We utilised a time series of Sentinel-1 interferograms that were formed as a preliminary result of work to use the LiCSAR processor (González et al., 2016) to automatically form interferograms at all of the world’s volcanoes that have been active during the Holocene. Data storage and processing was performed at the Climate, Environment and Monitoring from Space (CEMS) facility and, as some images are not yet available at this facility, our time series is split into two distinct sections (2016/09/03-2017/03/08 and 2018/01/14-2018/05/02). However, as we expect these two time series to contain the same signals, we are able to perform sICA on them as a single time series. In contrast to the Galapagos time series used in the preceding section, the Etna time series contains interferograms created between each acquisition and around three following it, which created a network of overlapping interferograms. Pixels with a coherence below 0.7 in any of the interferograms were masked throughout the entire time series of 76 interferograms (which are shown in Figure S2). The time courses recovered by sICA describe the strength that each component was used in each of the 76 interferograms, and a simple least-squares inversion was performed to invert for the strength that each

component was used in a daisy chain of 28 interferograms linking the acquisitions dates in the manner described by Lundgren et al. (2001).

Figure 2.16 shows the independent components (ICs) recovered by sICA, their strengths through the daisy chain of 28 interferograms, and a comparison between each IC and the DEM. We interpret the spatial pattern of IC0 as capturing eastward movement of the densely faulted eastern flank of the volcano. The cumulative nature of this component's time course is also indicative of it capturing deformation, but the final value attained by the cumulative time course remains low when compared to the changes at each time step, which is exemplified by the large change seen in the last data point. Inspection of the interferograms used in this analysis (Figure S2) shows the penultimate interferograms contain a broad negative signal, which is likely to have caused this IC to have been used to attempt to fit it. Given the short nature of this time series ( $\sim 12$  months), the confidence in this measurement has the potential to be improved upon through the use of longer time series, but this remains beyond the scope of this paper. We also interpret IC2 as capturing broader east-west spreading of the volcano due to both the spatial pattern seen, and the cumulative nature of the time course. In contrast to IC0-2, the time courses of IC3 and IC4 do not show any cumulative motion, and we conclude that these capture purely atmospheric signals.

When the correlations between phase and elevation are considered for each IC, only IC1 is seen to exhibit a strong linear relationship, which we interpret as suggesting that it may be capturing a topographically correlated APS. However, inspection of its cumulative time course shows that, with the exception of the last data point, an approximately linear increase occurs, which we interpret as suggesting that the IC is capturing broad, volcano wide deformation of the type first measured by Massonnet et al. (1995). Through use of a longer time series, we envisage that trends in the cumulative time course would become clearer and, should the cumulative signal continue to return towards zero as it does in the last data point, we would be more confident that the signal is solely capturing a topographically correlated APS. However, in the case that cumulative motion continues, we would be more confident that the IC contained a signal due to deformation, but could not rule out a contribution from a topographically correlated APS.

From our initial analysis, we conclude that when applying sICA at stratovolcanoes, some signals of geophysical interest may be isolated, but separation of a broad inflation/deflation signal that is centred under the cone from a topographically correlated APS may not be possible. However, as we seek to use sICA as the foundation for an automatic detection algorithm, these results do not dissuade us. Through isolating the combination of broad deformation and a topographic APS to one IC, we are still able to characterise the baseline behaviour of this combined signal through analysis of the IC's time course. Taking the example of a stratovolcano in which both a topographically

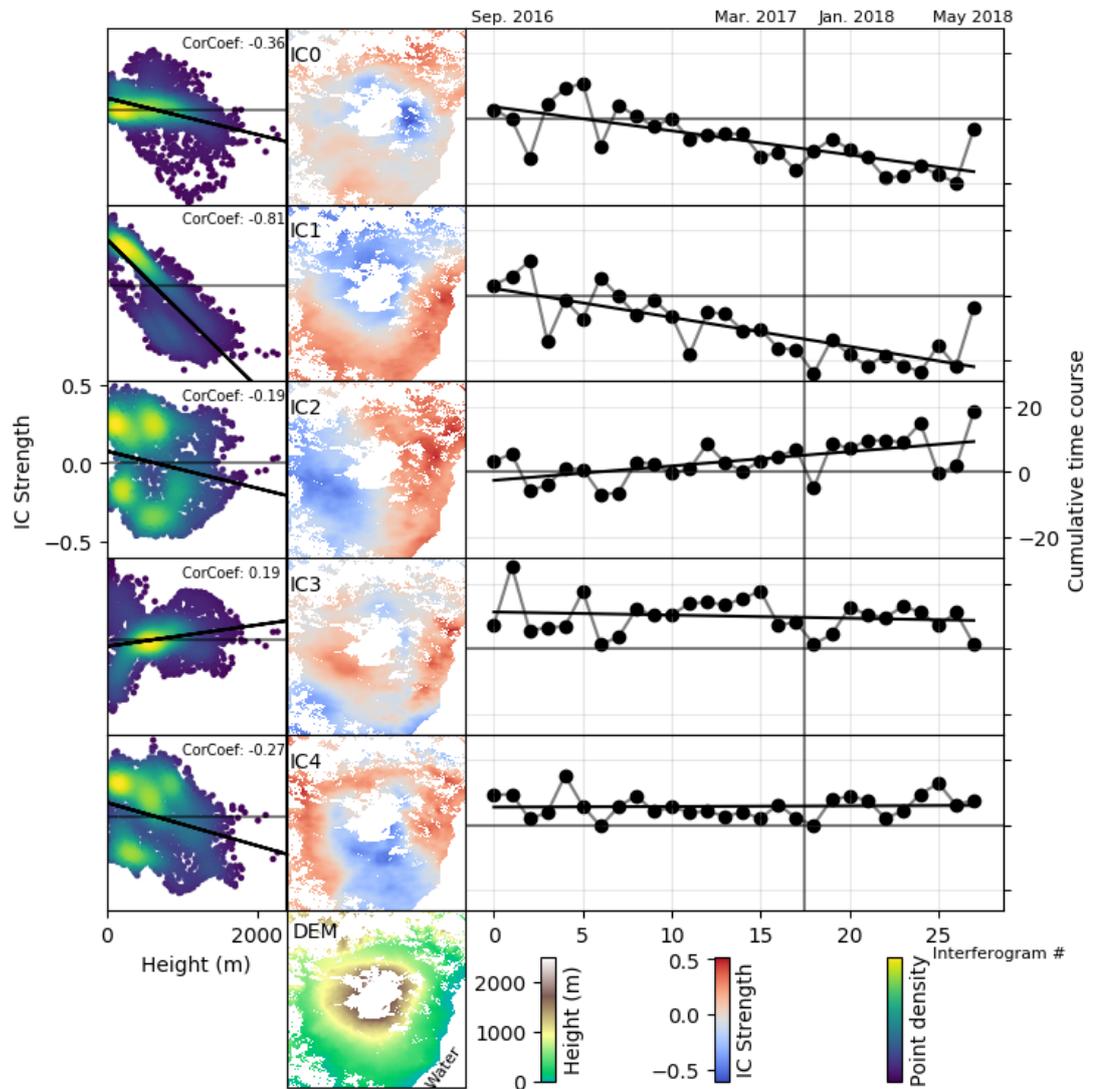
correlated APS and a gradual inflation signal are isolated into one IC, any change in the rate of inflation would lead to a change in the time course of the IC, which we could then seek to flag as an indicator of the volcano entering a period of unrest. However, further developments, such as the use of weather models to estimate the strength of the topographically correlated APS independently, remain beyond the remit of this paper.

## 2.6 Use of sICA in a monitoring algorithm

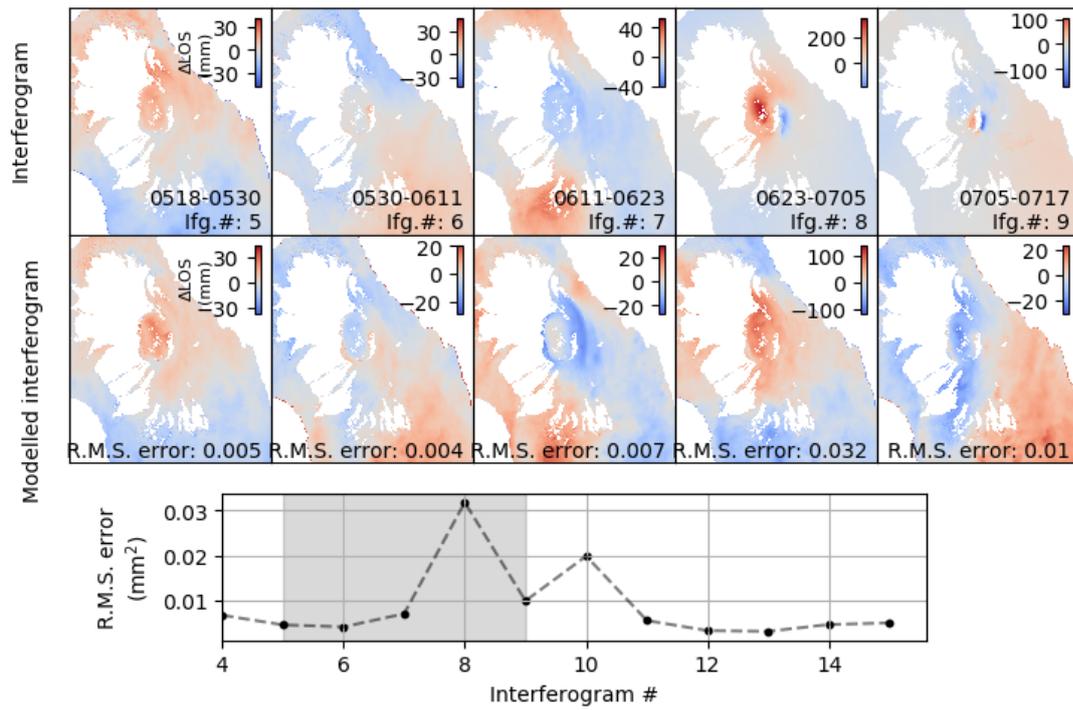
Our initial hypothesis was that the signals present in an InSAR time series at a volcanic centre could be expressed as a linear combination of a small number of latent signals, and that through isolation of these signals we could implement a way to monitor signs of unrest at a volcano. The results of the preceding section and those of Ebmeier (2016) have demonstrated the validity of this approach, and we now present results of a prototype monitoring algorithm that incorporates sICA.

Tools such as LiCSAR (González et al., 2016) are now producing time series of Sentinel-1 data at almost all of the world's active volcanoes. With routine acquisitions, a daisy-chain of interferograms is quick to build to the minimum six images required for sICA to be applied in the manner described in the preceding section. sICA can be applied to these six images to determine six latent components that characterise the atmospheric and geophysical signals for a period of steady state activity at the volcano, which may or may not include background deformation. When the next daisy-chain interferogram is added to the time series, we perform a simple least squares inversion to fit this image using a combination of the learned components, before calculating the mean absolute residual for each pixel between the actual and recreated data. We postulate that if no new deformation source is present, the mean absolute residual (henceforth referred to as residual) will be low, as no new signals are present. However, should the volcano enter a period of unrest leading to a new deformation source, the residual is likely to increase, which can be used to flag potential activity. Alternatively, if a background deformation source changes rate, this will lead to an uncharacteristic contribution from the component in which it lies, which can also serve as a flag.

Figure 2.17 shows the results of applying this algorithm to the Wolf time series presented in the previous section. Interferograms from before the period of unrest can be fit using the sICA components with a residual of  $\sim 1 \text{ mm}^2$ , yet when new signals are encountered during the period of unrest (such as subsidence of the caldera floor), the residual increases markedly and provides a clear flag that the volcano has entered a period of unrest. After this abates, the residual returns to pre-unrest levels.



**Figure 2.16:** Column two shows the five independent components (ICs) recovered by sICA at Etna and the DEM covering the area processed. Pixels in the area marked “water” do not have a useable radar return, whilst the remaining white areas of the DEM are masked due to low coherence. Column one shows the results of plotting each IC against the DEM, with IC1’s graph showing the clearest linear relationship which we interpret as being a result of IC1 capturing a topographically correlated APS. Column three shows the cumulative time courses, with IC1-3 showing cumulative motion throughout the two time periods that the data span.



**Figure 2.17:** Results of our prototype monitoring algorithm. The algorithm attempts to reconstruct a given interferogram (top row) using a linear mixture (second row) of six independent components determined from the first six interferograms. The mean of the absolute value of the residual is then calculated, and plotted for each new acquisition (bottom row). The five interferograms in the top two rows correspond to the shaded region. Before the period of unrest, the residual is seen to hover around  $\sim 0.005 \text{ mm}^2$  as the reconstructions approximate the originals well, but with the introduction of a new signal the residual increases approximately 6 fold.

## 2.7 Discussion

From our comparison of NMF, (t/s)PCA, and (t/s)ICA we have found that in almost all of the synthetic tests performed, sICA is the method most capable of recovering latent signals from a time series of interferograms. In application to real data (the Sentinel-1 Galapagos time series), the results obtained using sICA were highly plausible as the spatial patterns of the recovered sources agreed with signals identified as being due to geophysical processes in other studies.

This result provides justification for the construction of an automatic monitoring algorithm based on sICA, as it facilitates processes such as isolation of a signal of interest, or separation of geophysical signals from atmospheric ones.

The recovery of temporally independent time courses initially appears the most attractive approach as we are confident that the geophysical processes of interest at a volcanic site are temporally statistically independent from atmospheric processes, yet is hindered by two constraints: Firstly, the results presented in Pinel et al. (2011) suggest that a topographically correlated APS signal will be Gaussian in time (rendering it unrecoverable by ICA in most cases), and secondly (and more importantly) the large number of pixels observed at relatively few times produces a data matrix which is unsuited to ICA.

In contrast, spatially-organised data utilises the transpose of the temporal data and is well suited for the FastICA algorithm. However, statistical independence of the spatial nature of the sources is not to be expected in some cases (and in Section 2.4.5 we show that violation of this assumption does affect the fidelity of the sources recovered) and, whilst many sources of geophysical interest are likely to be non-Gaussian, this is not necessarily true for atmospheric signals. Despite these trade-offs, we show that in the majority of the synthetic cases considered, sICA outperforms all the other methods considered, and in all but the synthetic stratovolcano case, is able to recover useful latent signals. The vulnerability of sICA to signals that are not independent remains problematic for some applications, but as we seek only to characterise baseline behaviour for our automatic detection algorithm, we believe sICA can still be used for our goals. Through application of the method to other datasets in the future, we expect that more information on the importance of these two potential limitations will come to light.

Scope remains for the refinement of sICA with InSAR data. A key part of most ICA algorithms is how the non-Gaussian nature of a signal is measured and various approaches for this exist. In this study, we use a measurement similar to kurtosis, but inspection of estimates of the PDFs of signals (e.g. Figure 2.2) suggests that for many signals (such as a Mogi source), they may be more clearly identified through using a measure such as skewness. This has been implemented for medical imaging data by

Stone et al. (2002), and may be applicable for InSAR data.

ICA also requires the same pixels to be used throughout the time series. However, as the number of coherent pixels changes between interferograms, our method of masking all pixels with an average coherence  $< 0.8$  does not make use of all the information available. A more complex strategy to incorporate the information in pixels that are only coherent in some interferograms remains beyond the scope of this work, but may allow for more subtle signals to be recovered with sICA.

## 2.8 Conclusion

Our study suggests that sICA is the most suited BSS method for use with an InSAR time series at a volcanic centre. This is shown through synthetic tests, and application to a time series of Sentinel-1 data that spans the 2015 Wolf eruption, in which 3 signals of geophysical interest were isolated. However, aspects of the FastICA algorithm appear suitable for fine tuning to further increase its suitability for use with InSAR data. We introduce a simple algorithm that incorporates sICA to detect when a volcano enters a period of unrest, and demonstrate that it would have automatically identified the May 2015 eruption at Wolf volcano. Building from this point, future work on the automatic detection algorithm could allow for identification of different types of unrest (e.g. caused by the acceleration of a previously steady state process), and form an integral part of a system to automatically monitor all of the world's sub-aerial volcanoes.

## 2.9 Acknowledgments

The two sets of Copernicus Sentinel data (covering Wolf and Etna volcanoes) 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/>). MEG is supported by the LiCS grant. Figures were prepared in Matplotlib (Hunter, 2007). The Python implementation of FastICA used to perform ICA was modified from Varoquaux et al. (2012).

## References

- Aloisi, Marco, Mauro Amore, Mario Mattia, and Domenico Patane (2007). "Faulting on the western flank of Mt Etna and magma intrusions in the shallow crust". In: *Terra Nova* 19.1, pp. 89–94. DOI: 10.1111/j.1365-3121.2006.00724.x.
- Amari, S. (1999). "Natural Gradient Learning for Over- and Under-Complete Bases in ICA". In: *Neural Computation* 11.8, pp. 1875–1883. ISSN: 0899-7667. DOI: 10.1162/089976699300015990.

- Bartlett, Marian Stewart, J Movellan, and Terrence J. Sejnowski (2002). “Face recognition by Independent Component Analysis”. In: *IEEE Transactions on Neural Networks* 13.6, pp. 486–504.
- Bell, Anthony J. and Terrence J. Sejnowski (1995). “Information-Maximization Approach to Blind Separation and Blind Deconvolution”. In: *Neural Computation* 1159.1994, pp. 1129–1159.
- Bernard, B., P. Ramon, H. Wright, A. Guevara, S. Hidalgo, D. Pacheco, D. Narvaez, and F. Vasconez (2015). “Preliminary results on the 2015 eruption of Wolf volcano, Isabela island, Galápagos: Chronology, dispersion of the volcanic products, and insight into the eruptive dynamics”. In: *Abstract V31B-3022 presented at 2015 Fall Meeting, AGU, San Francisco, Calif.*
- Biggs, J, S K Ebmeier, W P Aspinall, Z Lu, M E Pritchard, R S J Sparks, and T a Mather (Jan. 2014). “Global link between deformation and volcanic eruption quantified by satellite imagery.” In: *Nature communications* 5, p. 3471. ISSN: 2041-1723. DOI: 10.1038/ncomms4471.
- Chaussard, E., P. Mililo, Roland Burgmann, D. Perissin, Eric J. Fielding, and B Baker (2017). “Remote Sensing of Ground Deformation for Monitoring Groundwater Management Practices: Application to the Santa Clara Valley During the 2012–2015 California Drought”. In: *Journal of Geophysical Research : Solid Earth*, pp. 8566–8582. DOI: 10.1002/2017JB014676.
- Chaussard, E, R Bürgmann, M Shirzaei, E J Fielding, and B Baker (2014). “Predictability of hydraulic head changes and characterization of aquifer-system and fault properties from InSAR-derived ground deformation”. In: *Journal of Geophysical Research : Solid Earth* 119, pp. 6572–6590. DOI: 10.1002/2014JB011266.
- Cohen-Weber, J., Roland Burgmann, E. Chaussard, C. Giannico, and A. Ferretti (2018). “Spatiotemporal Patterns of Precipitation-Modulated Landslide Deformation From Independent Component Analysis of InSAR Time Series”. In: *Geophysical Journal International*, pp. 1878–1887. DOI: 10.1002/2017GL075950.
- Comon, Pierre (1994). “Independent component analysis, A new concept?” In: *Signal Processing* 36.3, pp. 287–314. ISSN: 01651684. DOI: 10.1016/0165-1684(94)90029-9.
- Delacourt, C., P. Briole, and J. a. Achache (1998). “Tropospheric corrections of SAR interferograms with strong topography. Application to Etna”. In: *Geophysical Research Letters* 25.15, p. 2849. ISSN: 0094-8276. DOI: 10.1029/98GL02112.
- Ebmeier, S K, B J Andrews, M C Araya, D W D Arnold, J Biggs, C Cooper, E Cottrell, M Furtney, J Hickey, J Jay, R Lloyd, A L Parker, M E Pritchard, E Robertson, E Venzke, and J L Williamson (2018). “Synthesis of global satellite observations of magmatic and volcanic deformation : implications for volcano monitoring & the lateral extent of magmatic domains”. In: *Journal of Applied Volcanology*, pp. 1–26. DOI: 10.1186/s13617-018-0071-3.

- Ebmeier, S. K. (2016). “Application of Independent Component Analysis to multi-temporal InSAR data with volcanic case studies”. In: *Journal of Geophysical Research: Solid Earth*, pp. 8970–8986. ISSN: 21699313. DOI: 10.1002/2016JB013765. URL: <http://doi.wiley.com/10.1002/2016JB013765>.
- Farr, T.G., P.a. Rosen, Edward Caro, and Robert Crippen (2007). “The Shuttle Radar Topography Mission”. In: *Reviews of Geophysics* 45.2005, pp. 1–33. DOI: 10.1029/2005RG000183.1.
- Frappart, F., G. Ramillien, P. Maisongrande, and M. P. Bonnet (2010). “Denoising satellite gravity signals by independent component analysis”. In: *IEEE Geoscience and Remote Sensing Letters* 7.3, pp. 421–425. ISSN: 1545598X. DOI: 10.1109/LGRS.2009.2037837.
- GEBCO (2015). *The GEBCO 2014 Grid, version 20150318*. URL: [www.gebco.net](http://www.gebco.net) (visited on 02/16/2018).
- Goldstein, Richard M. and Charles L. Werner (1998). “Radar interferogram filtering for geophysical applications”. In: *Geophysical Research Letters* 25.21, p. 4035. ISSN: 0094-8276. DOI: 10.1029/1998GL900033.
- González, Pablo J. and Mimmo Palano (2014). “Mt. Etna 2001 eruption: New insights into the magmatic feeding system and the mechanical response of the western flank from a detailed geodetic dataset”. In: *Journal of Volcanology and Geothermal Research* 274, pp. 108–121. ISSN: 03770273. DOI: 10.1016/j.jvolgeores.2014.02.001.
- González, Pablo J., Richard J. Walters, E. L. Hatton, Karsten Spaans, A. McDougall, Andrew Hooper, and Tim J. Wright (2016). “LiCSAR: Tools for automated generation of Sentinel-1 frame interferograms”. In: *AGU Fall meeting*.
- Hooper, Andrew, David Bekaert, Karsten Spaans, and Mahmut Arıkan (Jan. 2012). “Recent advances in SAR interferometry time series analysis for measuring crustal deformation”. In: *Tectonophysics* 514-517, pp. 1–13. ISSN: 00401951. DOI: 10.1016/j.tecto.2011.10.013.
- Hotelling, H (1933). “Analysis of a complex of statistical variables into principal components”. In: *Journal of Educational Psychology* 24(7), pp. 498–520.
- Hunter, J. D. (2007). “Matplotlib: A 2D graphics environment”. In: *Computing In Science & Engineering* 9.3, pp. 90–95.
- Hyvärinen, A. and E. Oja (2000). “Independent component analysis: Algorithms and applications”. In: *Neural Networks* 13.4-5, pp. 411–430. ISSN: 08936080. DOI: 10.1016/S0893-6080(00)00026-5.
- Hyvarinen, Aapo (1998). “New Approximations of Differential Entropy for Independent Component Analysis and Projection Pursuit”. In: 1.
- Hyvärinen, Aapo (1999). “Fast and robust fixed-point algorithms for independent component analysis”. In: *IEEE Transactions on Neural Networks* 10.3, pp. 626–634. ISSN: 10459227. DOI: 10.1109/72.761722.

- Hyvärinen, Aapo, Juha Karhunen, and Erkki Oja (2001). *Independent Component Analysis*. Vol. 21. 1. John Wiley and Sons, INC., Hoboken, New Jersey, pp. 135–144. ISBN: 047140540X. DOI: 10.1002/0471221317.
- Hyvärinen, Aapo and Erkki Oja (1997). “A Fast Fixed-Point Algorithm for Independent Component Analysis”. In: *Neural Computation* 9.7, pp. 1483–1492. ISSN: 0899-7667. DOI: 10.1162/neco.1997.9.7.1483.
- Isomura, Takuya and Taro Toyoizumi (2016). “A Local Learning Rule for Independent Component Analysis”. In: *Nature Publishing Group*, pp. 1–17. ISSN: 20452322. DOI: 10.1038/srep28073.
- Jutten, C and J Herault (1991). “Blind separation of sources, part I: an adaptive algorithm based on neuromimetic architecture”. In: *Signal Processing* 24, pp. 1–10. ISSN: 01651684. DOI: 10.1016/0165-1684(91)90079-X.
- Karhunen, Kari (1947). “Über lineare Methoden in der Wahrscheinlichkeitsrechnung”. In: *Universitat Helsinki* 37.
- Kositsky, a. P. and J. P. Avouac (2010). “Inverting geodetic time series with a principal component analysis-based inversion method”. In: *Journal of Geophysical Research: Solid Earth* 115.3, pp. 1–19. ISSN: 21699356. DOI: 10.1029/2009JB006535.
- Lee, Daniel D and H Sebastian Seung (1999). “Learning the parts of objects by non-negative matrix factorization”. In: *Nature* 401.October 1999, pp. 788–791. ISSN: 0028-0836. DOI: 10.1038/44565. arXiv: arXiv:1408.1149. URL: <http://www.ncbi.nlm.nih.gov/pubmed/10548103>.
- Liu, Bin, Matt King, and Wujiao Dai (2018). “Common mode error in Antarctic GPS coordinate time-series on its effect on bedrock-uplift estimates”. In: *Geophysical Journal International* 214, pp. 1652–1664. DOI: 10.1093/gji/ggy217.
- Lorenz, E N (1956). *Empirical Orthogonal Functions and Statistical Weather Prediction*.
- Lundgren, P., S. Usai, E. Sansosti, R. Lanari, M. Tesauro, G. Fornaro, and P. Berardino (2001). “Modeling surface deformation observed with synthetic aperture radar interferometry at Campi Flegrei caldera”. In: *Journal of Geophysical Research* 106.June, pp. 19355–19366. ISSN: 0148-0227. DOI: 10.1029/2001JB000194.
- Lundgren, Paul and Paul A Rosen (2003). “Source model for the 2001 flank eruption of Mt . Etna volcano”. In: *Geophysical Research Letters* 30.7, pp. 8–11. DOI: 10.1029/2002GL016774.
- Makeig, Scott, Anthony J. Bell., Tzyy-Ping Jung, and Terrence J. Sejnowski (1996). “Independent Component Analysis of Electroencephalographic Data”. In: *Advances in Neural Information Processing Systems* 8, pp. 145–151. ISSN: 10495258. DOI: 10.1109/ICOSP.2002.1180091.
- Massonnet, Didier, Pierre Briole, and Alain Arnaud (1995). *Deflation of Mount Etna monitored by spaceborne radar interferometry*. DOI: 10.1038/375567a0.
- McKeown, M J, T P Jung, S Makeig, G Brown, S S Kindermann, T W Lee, and T J Sejnowski (1998). “Spatially independent activity patterns in functional MRI data

- during the stroop color-naming task.” In: *Proceedings of the National Academy of Sciences of the United States of America* 95.3, pp. 803–810. ISSN: 00278424. DOI: 10.1073/pnas.95.3.803. URL: <http://www.pnas.org/content/95/3/803%7B%5C%7D5Cnhttp://www.ncbi.nlm.nih.gov/pubmed/9448244%7B%5C%7D5Cnhttp://www.pnas.org/content/95/3/803.full.pdf%7B%5C%7D5Cnhttp://www.pnas.org/content/95/3/803.long>.
- McKeown, M and Makeig S (1998). “Analysis of fMRI data by blind separation into independent spatial components”. In: *Human Brain Mapping* 6.June 1997, pp. 160–188.
- Meyer, F. J., P. W. Webley, J. Dehn, S. A. Arko, D. B. McAlpin, and W. Gong (2016). “The SARVIEWS Project: Automated SAR Processing in Support of Operational Near Real-time Volcano Monitoring”. In: *AGU Fall meeting*.
- Mogi, Kiyoo (1958). “Relations between the eruptions of various volcanoes and the deformations of the ground surfaces around them”. In: *Bulletin of the Earthquake Research Institute* 36, pp. 99–134. ISSN: 0040-8972. DOI: 10.1016/j.epsl.2004.04.016.
- Novellis, V De, R Castaldo, C De Luca, S Pepe, I Zinno, F Casu, R Lanari, and G Solaro (2017). “Source modelling of the 2015 Wolf volcano ( Galápagos ) eruption inferred from Sentinel 1-A DInSAR deformation maps and pre-eruptive ENVISAT time series”. In: *Journal of Volcanology and Geothermal Research* 344, pp. 246–256. ISSN: 0377-0273. DOI: 10.1016/j.jvolgeores.2017.05.013.
- Paatero, Pentti and Unto Tapper (1994). “Positive Matrix Factorization - A Nonnegative Factor Model with Optimal Utilization of Error Estimates of Data Values”. In: *Environmetrics* 5.April 1993, pp. 111–126. ISSN: 11804009. DOI: 10.1002/env.3170050203.
- Pinel, V., A. Hooper, S. De la Cruz-Reyna, G. Reyes-Davila, M. P. Doin, and P. Bascou (2011). “The challenging retrieval of the displacement field from InSAR data for andesitic stratovolcanoes: Case study of Popocatepetl and Colima Volcano, Mexico”. In: *Journal of Volcanology and Geothermal Research* 200.1-2, pp. 49–61. ISSN: 03770273. DOI: 10.1016/j.jvolgeores.2010.12.002.
- Pinel, V., M.P. Poland, and Andrew Hooper (Nov. 2014). “Volcanology: Lessons learned from Synthetic Aperture Radar imagery”. In: *Journal of Volcanology and Geothermal Research* 289, pp. 81–113. ISSN: 03770273. DOI: 10.1016/j.jvolgeores.2014.10.010.
- Porrill, John and Jv James V Stone (1998). “Undercomplete independent component analysis for signal separation and dimension reduction”. In: pp. 1–10.
- Rudolph, M L, M Shirzaei, M Manga, and Y Fukushima (2013). “Evolution and future of the Lusi mud eruption inferred from ground deformation”. In: *Geophysical Research Letters* 40.January, pp. 1089–1092. DOI: 10.1002/grl.50189.

- Siebert, L. and T. Simkin (2013). *Volcanoes of the World: an Illustrated Catalog of Holocene Volcanoes and their Eruptions*. URL: [www.volcano.si.edu](http://www.volcano.si.edu).
- Solaro, Giuseppe, Valerio Acocella, Susi Pepe, Joel Ruch, Marco Neri, and Eugenio Sansosti (2010). “Anatomy of an unstable volcano from InSAR: Multiple processes affecting flank instability at Mt. Etna, 1994-2008”. In: *Journal of Geophysical Research: Solid Earth* 115.10, pp. 1994–2008. ISSN: 21699356. DOI: 10.1029/2009JB000820.
- Solem, Jan Erik (2012). *Programming Computer Vision with Python*.
- Spaans, Karsten and Andy Hooper (2016). “InSAR processing for volcano monitoring and other near-real time applications”. In: *Journal of Geophysical Research: Solid Earth*, n/a–n/a. ISSN: 21699313. DOI: 10.1002/2015JB012752.
- Sparks, R. S. J., J. Biggs, and J. W. Neuberg (2012). “Monitoring Volcanoes”. In: *Science* 335.6074, pp. 1310–1311. ISSN: 0036-8075. DOI: 10.1126/science.1219485.
- Stone, J.V, J Porrill, N.R Porter, and I.D Wilkinson (2002). “Spatiotemporal Independent Component Analysis of Event-Related fMRI Data Using Skewed Probability Density Functions”. In: *NeuroImage* 15.2, pp. 407–421. ISSN: 10538119. DOI: 10.1006/nimg.2001.0986.
- Stone, James (2002). “Independent component analysis: an introduction.” In: *Trends in cognitive sciences* 6.2, pp. 59–64. ISSN: 1364-6613. DOI: 10.1016/S1364-6613(00)01813-1.
- Varoquaux, G., P. Lafaye de Micheaux, and S. van der Walt (2012). *Python implementation of the fast ICA algorithms*. <https://github.com/GaelVaroquaux/canica/blob/master/canica/algorithms/fastica.py>.
- Venzke, E (ed.) (2015). “Global Volcanism Program, 2016. Report on Wolf (Ecuador)”. In: *Bulletin of the Global Volcanism Network* 41:10. Smithsonian Institution.
- Xu, Wenbin, Sigurjón Jónsson, Joël Ruch, and Yosuke Aoki (2016). “The 2015 Wolf volcano (Galápagos) eruption studied using Sentinel-1 and ALOS-2 data”. In: *Geophysical Research Letters*, pp. 9573–9580. DOI: 10.1002/2016GL069820..



## Chapter 3

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

M. Gaddes<sup>1</sup>, A. Hooper<sup>1</sup>, and M. Bagnardi<sup>2,3</sup>

<sup>1</sup> *COMET, School of Earth and Environment, University of Leeds, United Kingdom*

<sup>2</sup> *Formerly, COMET, School of Earth and Environment, University of Leeds, United Kingdom*

<sup>3</sup> *Now at Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA*

Keypoints:

- We have developed an algorithm that incorporates spatial independent component analysis (sICA) to detect signs of deformation generating volcanic unrest in a time series of interferograms at a volcanic centre.
- When our algorithm is applied to a time series of Sentinel-1 data covering the Galapagos Archipelago, we are able to detect the increase in inflation that preceded the 2018 eruption of Sierra Negra.
- One component of our algorithm performs sICA robustly in order to separate geophysical and atmospheric signals in a time series of interferograms, and we release this freely to the scientific community under the name ICASAR.

## Abstract

The latest generation of SAR satellites produce measurements of ground deformation at the majority of the world’s sub-aerial active volcanoes, and can be used to detect signs of deformation generating volcanic unrest. We present a detection algorithm that uses these data to automatically warn when deformation at a volcano departs from the background. We demonstrate our approach on synthetic datasets, and the unrest leading to the 2018 eruption of Sierra Negra (Galapagos). Our algorithm encompasses spatial independent component analysis (sICA), and uses a significantly improved version of the ICASO algorithm, which we term ICASAR, to robustly perform sICA. 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.

## 3.1 Introduction

There are  $\sim 1400$  volcanoes globally with the potential to erupt, but with only  $\sim 100$  volcano observatories, many volcanoes remain unmonitored (Loughlin et al., 2015). However, the routine global acquisition of the latest generation of 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 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 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 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 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 datasets. When utilising ICA in an automatic detection algorithm, this approach cannot be relied upon, and we must instead implement other 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), whilst 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 unmixing matrix,  $\mathbf{W}$ , randomly at each run of the FastICA algorithm, and seeks to ensure that a variety of local minima are sampled. 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 analysed using 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 randomisation 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 realisations of the recovered sources can then be analysed through the same clustering approach that was discussed in the preceding paragraph. In practise, 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 and Ramkumar, 2013), which tests the reliability of the recovered components on separate datasets 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 analysing the similarities of the mixing matrices (Hyvärinen, 2011), but a subsequent version performs this analysis on the recovered sources instead (Hyvärinen and Ramkumar, 2013). The algorithm was originally used with inter-session or inter-subject 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 one did not share any acquisition dates with those in group two) before performing the analysis.

We choose to implement the ICASO algorithm rather than the ISCTEST algorithm as our goal of creating a detection algorithm centres on characterising the baseline behaviour 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 characterise a volcano's background behaviour. The former approach would be more suited to the ISCTEST algorithm as we could assign p-values to recovered components, but through use of the ICASO algorithm we can recover latent sources confidently, and avoid the need to subdivide our data into two independent datasets. In applications in which the signals of interest are of low magnitude, the ability of the ICASO algorithm to retain all of the input images is also likely to be useful as this approach increases the signal-to-noise ratio.

## 3.2 Methods: Detection algorithm

Our 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 one to characterise the baseline data, whilst 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, whilst 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, whilst interferograms created after this point we consider as testing data.

### 3.2.1 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 modernised Python version that is specialised for use with InSAR data. We term our Python based algorithm ICASAR, and make it freely available via GitHub (Gaddes, 2017).

Figure 3.1 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 6 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, whilst 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), whilst using the absolute value of the correlation between each source as the distance metric. This is used as when clustering images, there are as many data points as there are images, residing within a space with as many dimensions as there are pixels. Whilst the Euclidean distance could be calculated between points, the ambiguity of the sign of sources recovered by ICA would result in sign flipped versions of the same source being treated as having a large distance between them, despite featuring essentially the

same source. However, by using the absolute value of the correlation between images, identical sources with a correlation of (1) and sign flipped versions with a correlation of (-1) are treated in the same manner, whilst dissimilar sources have low correlations. A trivial step can then be performed to transfer this measure of similarity into one of distance (i.e. a high correlation translates to a low distance), which 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. The cutting of the cluster tree is calculated through considering how clusters decay in size as the algorithm moves from the case in which all points are considered as one cluster, to the case in which each point is an individual cluster. As the algorithm moves down the hierarchy, points leave a cluster and, providing the number of points to leave a cluster in a step is lower than the minimum cluster size hyperparameter, they are considered noise. However, in the case that the number of points leaving per step is larger than the minimum cluster size, the cluster can be said to have divided, with the exiting points having now created a new cluster, instead of being labelled as noise. The tree can then be cut to maximise the longevity of the clusters that it contains, as a compact and isolated cluster will remain stable throughout much of the process, whilst a cluster comprised of two dense regions is likely to divide into two clusters midway through the process, and therefore be relatively short lived. 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.

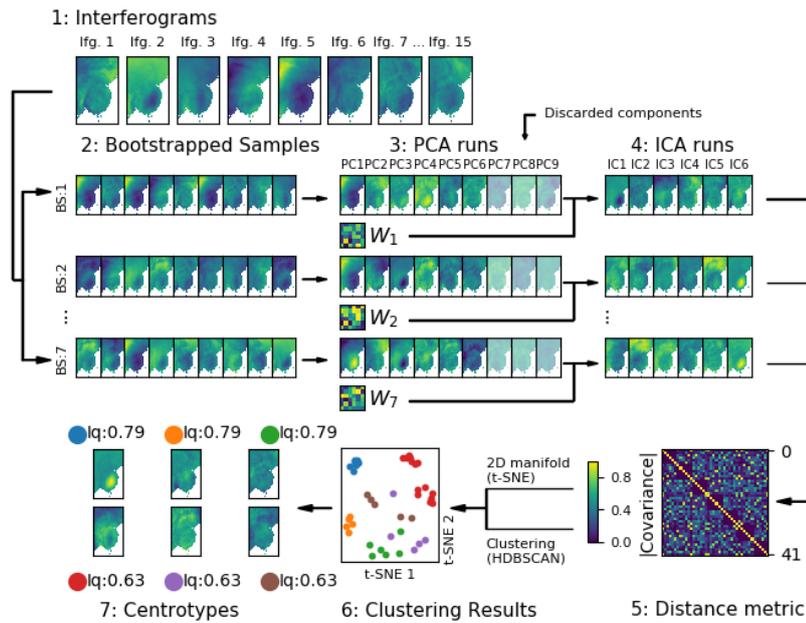
The second method seeks to provide a representation of the similarities between each recovered source in a manner that is easily interpretable to a human. This is achieved through considering each recovered source as a sample within a space with as many dimensions as the sources have pixels, and then fitting a  $2D$  manifold through this space which preserves the distances between pairs of points in the high dimensional space and on the  $2D$  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 (the absolute value of the correlation between sources) when finding the manifold. The original ICASO algorithm utilised curvilinear component

analysis (CCA, Demartines and Herault (1997)), but in ICASAR we substitute this algorithm for the newer t-distributed stochastic neighbour embedding (t-SNE, Maaten and Hinton (2008)), as this has been shown to produce 2D maps that reveal structures within the data more robustly than a variety of methods, including CCA (Maaten and Hinton, 2008).

These two methods can be used in a complementary fashion to create the 2D plot shown in part six of Figure 3.1, in which the points that represent each recovered source are coloured depending on which cluster they are a member of, or if they are noise. Whilst inspection of this plot is not required, it can be used by a human observer to qualitatively ascertain the robustness of recovered sources. An advantage of this approach over the ISCTEST algorithm is that it provides a simple and powerful visualisation that allows a human interpreter to easily understand the relationships between the latent sources recovered in each run of the FastICA algorithm. Himberg et al. (2004) suggest that compact and isolated clusters are considered to contain robust sources, whilst those that form more indistinct clusters are considered to contain more spurious sources. In addition, through colouring 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 intra cluster 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 centrotpe. This is calculated as the point within a cluster that minimises the distance between it and the remaining points within that cluster. Combining the centrotypes 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 detection algorithm.

In the first part of our detection algorithm, we utilise the ICASAR algorithm with both bootstrapping and random initiation of the unmixing matrix, and seek 200 runs that converge. 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



**Figure 3.1:** Depiction of the ICASAR algorithm when applied to a time series of InSAR data that covers Sierra Negra (Galapagos Archipelago, Ecuador). This time series features uplift of the caldera floor which is clearest in interferogram two, and poor coherence requiring the masking of pixels on the south-eastern flank of the volcano (i.e. the white area). Interferograms are chosen from the mean centred baseline data (1) randomly to create multiple bootstrapped samples (2). PCA is performed on each of these samples (3), and lower order components are discarded (shown with reduced opacity) to reduce the dimensionality to the required level. sICA is then performed on each whitened sample (4), before the similarity/distance between each recovered source is estimated (5). The distance matrix is then used in parallel by the HDBSCAN clustering algorithm, and the t-SNE manifold learning algorithm to produce a 2D visualisation of the recovered sources (6). The centrotypes of each cluster are then recovered (7), and ordered by their cluster quality index,  $I_q$ .

principal components are not discarded (Gaddes et al., 2018). Consequently, we set the FastICA algorithm to recover around six sources, as we postulate that the majority of the world’s volcanic centres are likely to contain several consistent atmospheric signals (e.g. a topographically correlated atmospheric phase screen), 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 case that independent information or the inspection of interferograms suggests that several processes may be occurring at a volcano, this value may need to be increased.

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 are chosen with replacement to create a sample containing  $m$  interferograms. The data now lie within a  $n$ D hyperplane in the  $m$ D space, and only  $n$  principal components can be found. If we 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 deformation, 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 centre, 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, 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.

### 3.2.2 Characterisation of the baseline data

The second stage of our algorithm seeks to characterise the temporal nature of the sources recovered in the previous stage. Figures 3.2 and 3.3 depict the application of our algorithm to synthetic data similar to that described in Chapter 2 (except also featuring the addition of east-west phase ramps), 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 baseline 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 characterise 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 3.2 and 3.3 as the colours 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. In the work presented in Chapter 2, only six interferograms were required for the baseline stage, but given that sICA is able to recover sources more accurately with longer time series, longer baseline stages are likely to be advantageous.

Our ability to fit the baseline data using our learned sources can be also characterised 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”. We believe that this second method will avoid the false positives that may be caused by strong atmospheric signals in a specific SAR image, as the interferograms before and after this image will show the signal with the opposite sign. Consequently, the residuals produced when we are unable to fit this new signal in the two interferograms will also be of opposite sign, and therefore sum to approximately zero when considering each pixel. However, an added complication of this is that after the RMS residual indicates a new signal has entered the time series, an additional interferogram is required before the algorithm can ascertain if it was caused by an unusual atmosphere on the date shared by the two interferograms. 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.

### 3.2.3 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 behaviour 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 baseline phase, but shifted appropriately vertically, the data points after the period of transient deformation again lie close to the line of best fit and are no longer flagged as expressing a significant deviation from the baseline stage, providing that the physical process operating during the baseline stage are again active in the same manner. This is demonstrated in the time course of IC2 in Figure 3.2, and the RMS cumulative residual in Figure 3.3.

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 (modelled as the inflation of a point source in an elastic half space (Mogi, 1958)), a turbulent atmospheric phase screen (APS), an East-West phase gradient, and a topographically correlated APS (discussed in Gaddes et al. (2018)).

#### 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 dataset, we know that only three signals were used to generate the time series, whilst 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 3.2 shows the results of applying first the ICASAR algorithm and then our detection algorithm to these data, with the  $2D$  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  $2D$  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), whilst the fourth appears to capture an aspect of the turbulent atmosphere.

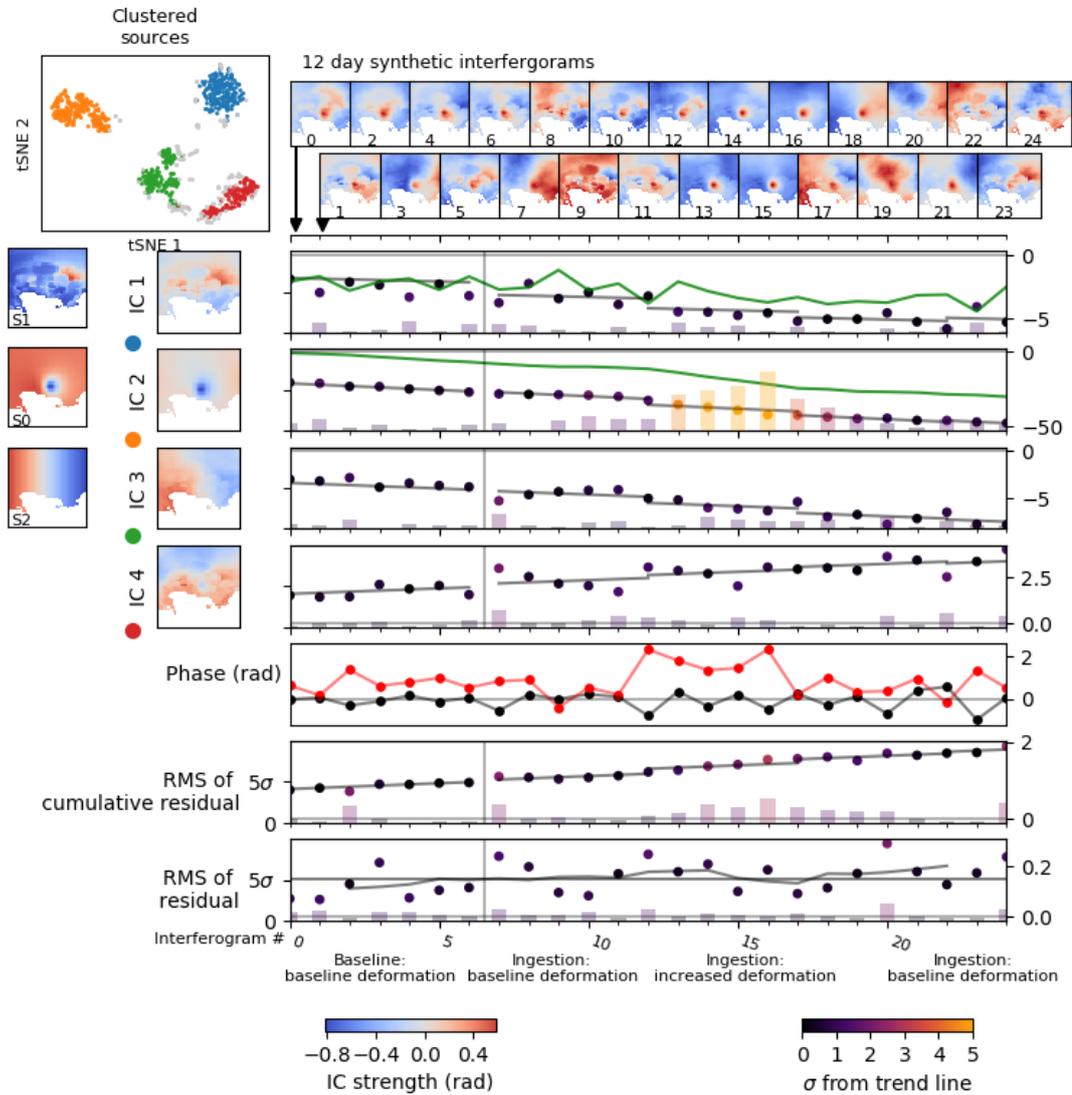
During the “Ingestion: baseline deformation” phase, the time series continues in a manner similar to that of the baseline data and no significant deviations from the extrapolated trend lines are seen for either the cumulative time courses and cumulative residual. However, in interferograms 12 to 16, 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 caused by a process such as increased flux of magma from depth to an area of shallow storage. 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 colours). 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 behaviour seen during baseline stage, and the algorithm returns to assigning insignificant deviations to each data point.

### **Emergence of a new signal**

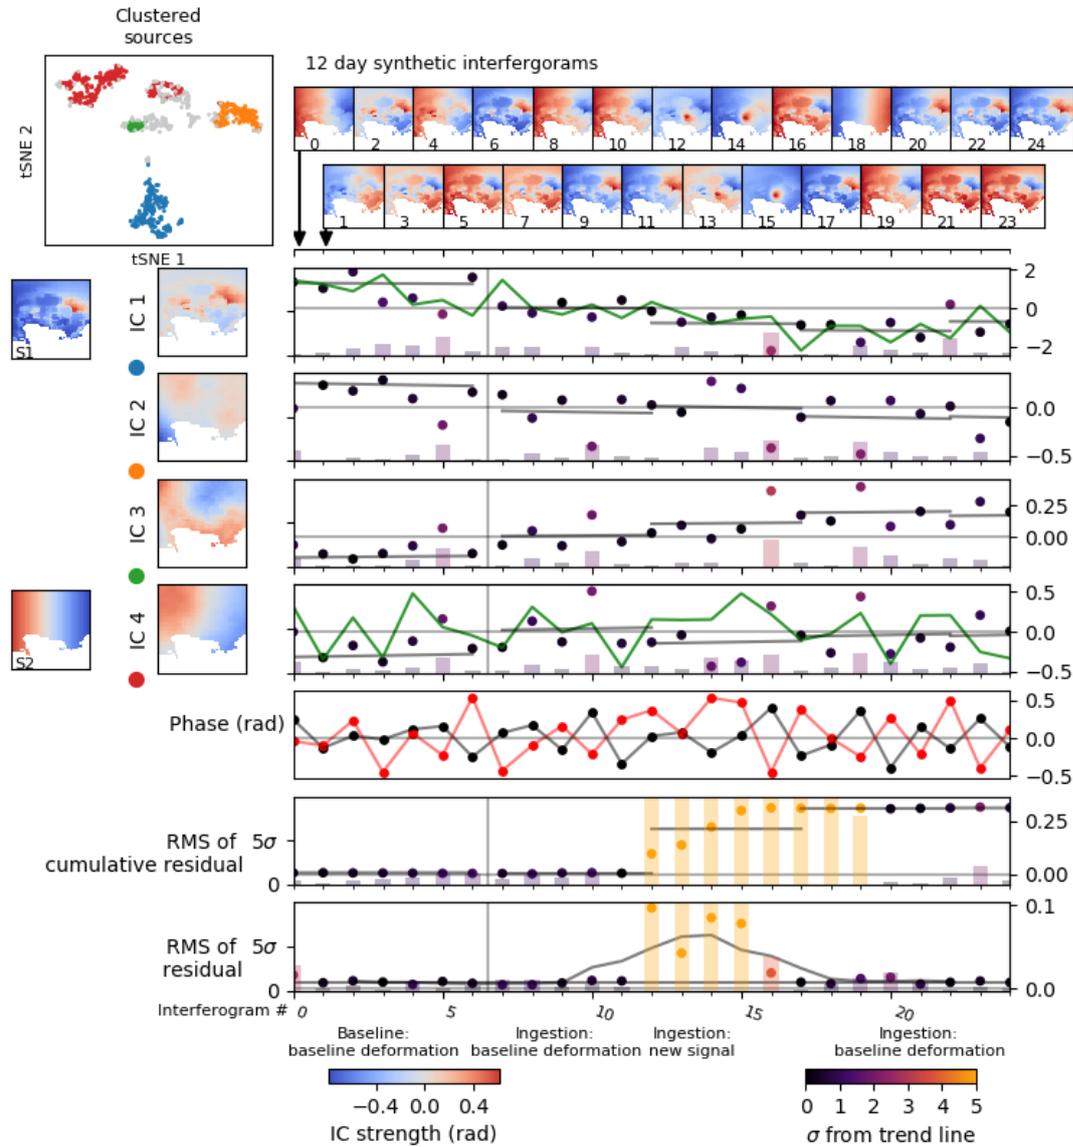
This scenario aims to demonstrate our algorithm’s ability to detect a period of volcanic unrest created when a new deformation signal enters the time series. Figure 3.3 shows the results of applying first the ICASAR algorithm and then our detection algorithm to this data. We observe that the two most significant sources recovered (IC1 and IC2) correspond to the two synthetic sources, whilst 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.

## **3.3 Application to Sierra Negra**

To demonstrate our 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.



**Figure 3.2:** The results of applying the 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 centrotypes of the four clusters are shown as IC1-IC4, of which three can be seen to correspond to synthetic sources (S0-S2). 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 baseline stage. For comparison, the cumulative time courses used to synthesis the data are shown in green. The deviations between each point and the line of best fit are shown as both the colour of the data point, and as the height and colour the low opacity bars (black to orange). 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 labelling of points). The fifth line graph shows the values for a pixel outside the deforming region in black, and for a pixel within the deforming region in red.



**Figure 3.3:** 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 3.2, 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.

Sierra Negra is a basaltic shield volcano located on the southern half of Isabela island in the Galapagos Archipelago (Reynolds et al., 1995). Figure 3.4 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 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 et al., 2005; Jónsson, 2009) and GPS (Chadwick et al., 2006), with total uplift between 1992 and 2005 reaching just below five metres prior to the eruption. During this period, uplift was in part accommodated by trapdoor faulting, which may have acted to reduce the build up of strain and delay the eventual eruption (Chadwick et al., 2006). Modelling of the observed deformation by different studies suggested that it was caused by the inflation of a sill at a depth of  $\sim 2$  km below the caldera floor, or  $\sim 1$  km below sea level (Amelung et al., 2000; Yun et al., 2006; Chadwick et al., 2006; Jónsson, 2009).

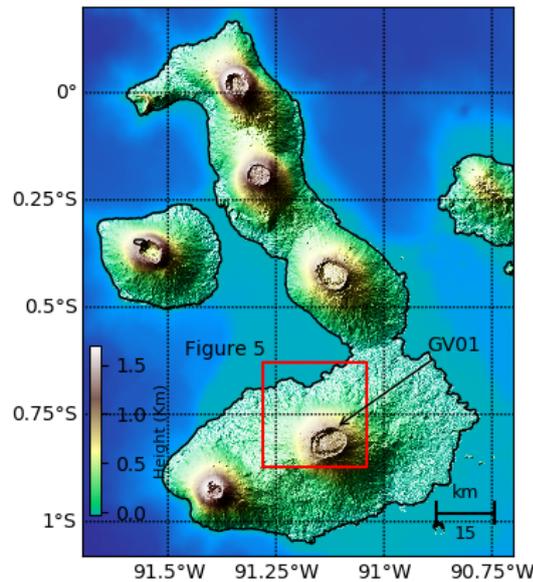
Between the 2005 and 2018 eruptions, the deployment of a temporary broadband seismic network termed the SIGNET array (Tepp et al., 2014) has provided additional insights into the structure of the crust beneath Sierra Negra. Body wave tomography has been used to infer the existence of large low velocity 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), whilst attenuation tomography has been used to infer the existence of separate low velocity zones from 0.5 – 3 km 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 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 the 26th of 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 the 23rd of August activity had ceased (Sennert, 2018b). During the eruption,  $\text{SO}_2$  emissions were visible to the Deep Space Climate Observatory satellite (Carn et al., 2018).

To test our detection algorithm on Sierra Negra, we first create a “daisy chain” of the shortest possible temporal baseline interferograms from 98 Sentinel-1 synthetic aperture RADAR images (shown in Figure S1) using the LiCSAR processor (González et al., 2016), which includes filtering with a Goldstein filter (Goldstein and Werner, 1998) and unwrapping using SNAPHU (Chen and Zebker, 2001). The data span the thirteenth of December 2014 to the first of July 2018, with the last 12 day interferogram capturing the co-eruptive 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 (Dec 2014 - April 2015) and recover the suite of sources shown in Figure 3.5. 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 (Iq) is displayed in the legend and can be used as a metric to rank the sources based on their robustness, as discussed in Section 3.2.1. Whilst further human analysis of these components is not required for our automatic detection algorithm, we present a brief discussion here, as standalone use of the ICASAR algorithm to isolate signals of geophysical interest may be useful for motives other than volcano monitoring. Figure 3.6 shows the centrotypes of each cluster ordered by their cluster quality index, and compares them against the Shuttle Radar Topography Mission Digital Elevation Model (SRTM DEM, Farr et al. (2007)). 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, and it is likely that this component is capturing a topographically correlated APS. Visual inspection of IC1 suggests that this signal contains the uplift signal centred at the caldera, whilst the remaining components show no easily interpretable patterns and are likely to correspond to atmospheric signals.

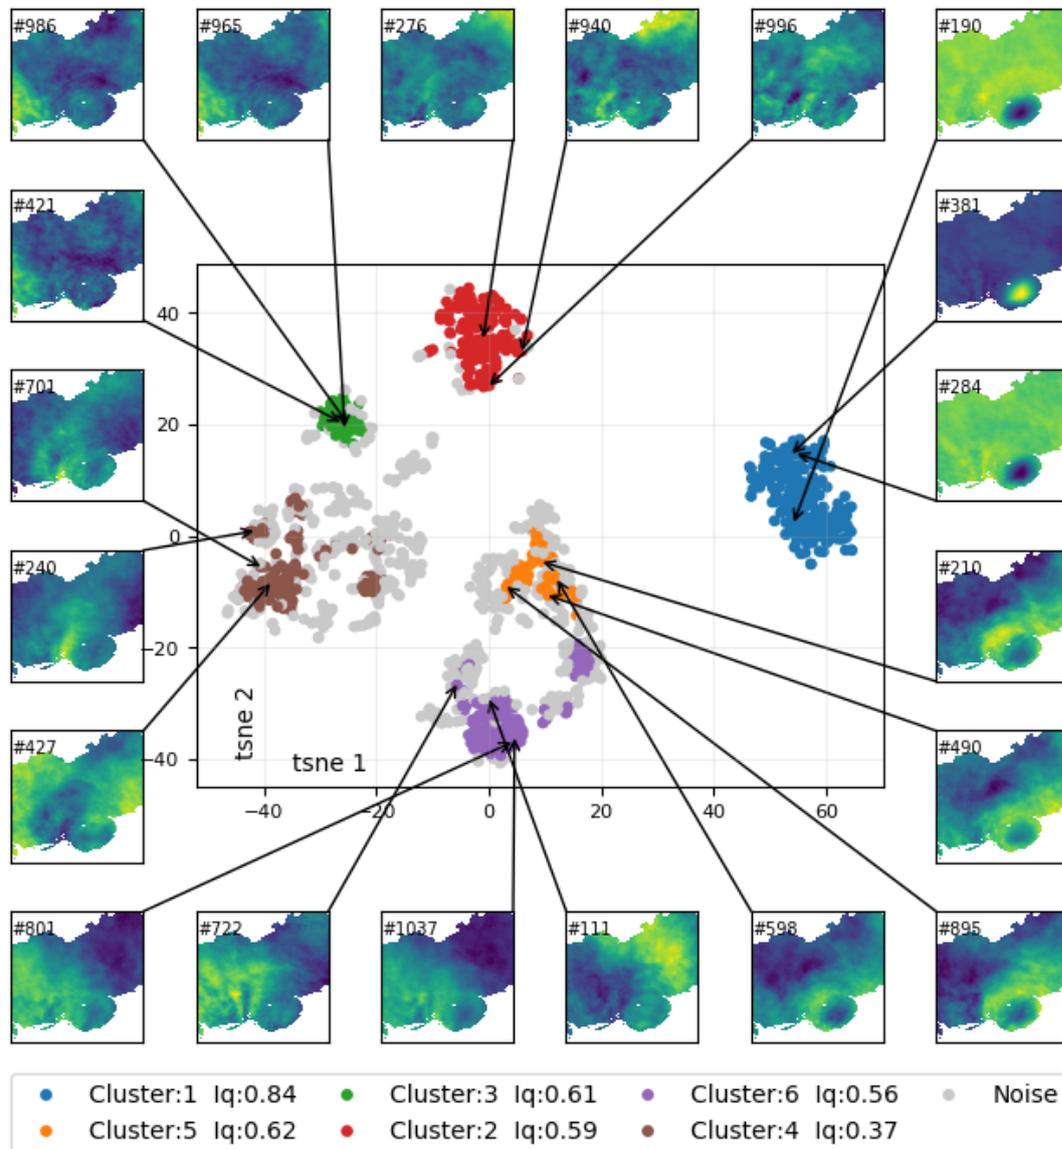
Figure 3.7 shows the results of applying our detection algorithm to the time series. The most striking feature is the flagging of the time course of IC1, as indicated by the orange colouring of the points, due to the rate of inflation increasing. Closer to the eruption, other time courses also exhibit unusual behaviour which is flagged as a sign of unrest (e.g. the time course of IC3 from interferogram 88 onwards), and may be due to processes such as slip on the intra-caldera 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



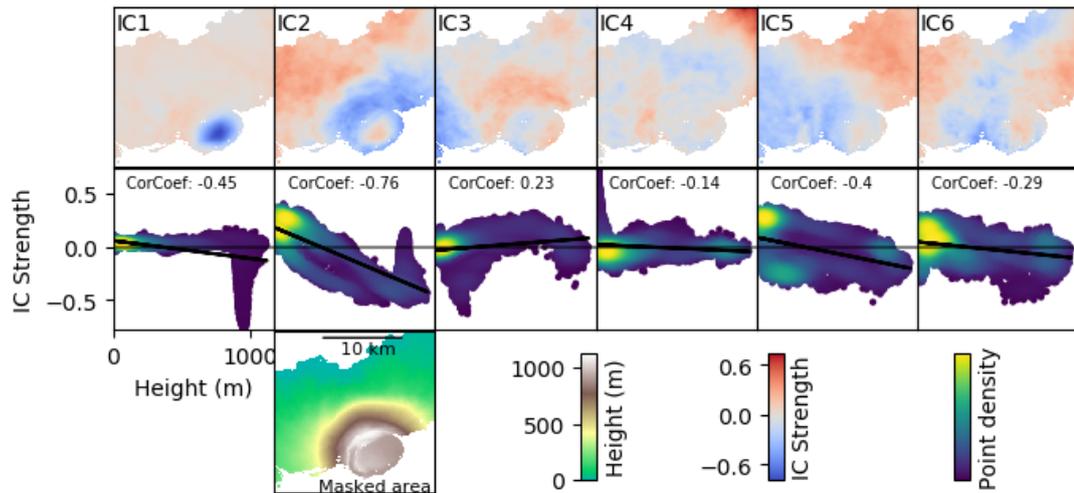
**Figure 3.4:** Overview of the Western Galapagos Islands of Fernandina and Isabella (Ecuador). Sierra Negra is visible as the area of high topography delineated by the red box, which shows the extent of the interferograms presented in the remaining figures. An arcuate area of high topography is visible 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 3.4 is labelled as GV01.

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 3.7), 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. Whilst 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. Additionally, increasing the length of the baseline data used at each volcano is likely to allow the ICs selected to characterise 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 3.5:** 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, whilst a random subset of these are shown in full around the edge of the plot. The position of each point within the 2D scatter plot is found by the manifold learning method t-SNE, whilst the labels and associated colours 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 labelled 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.



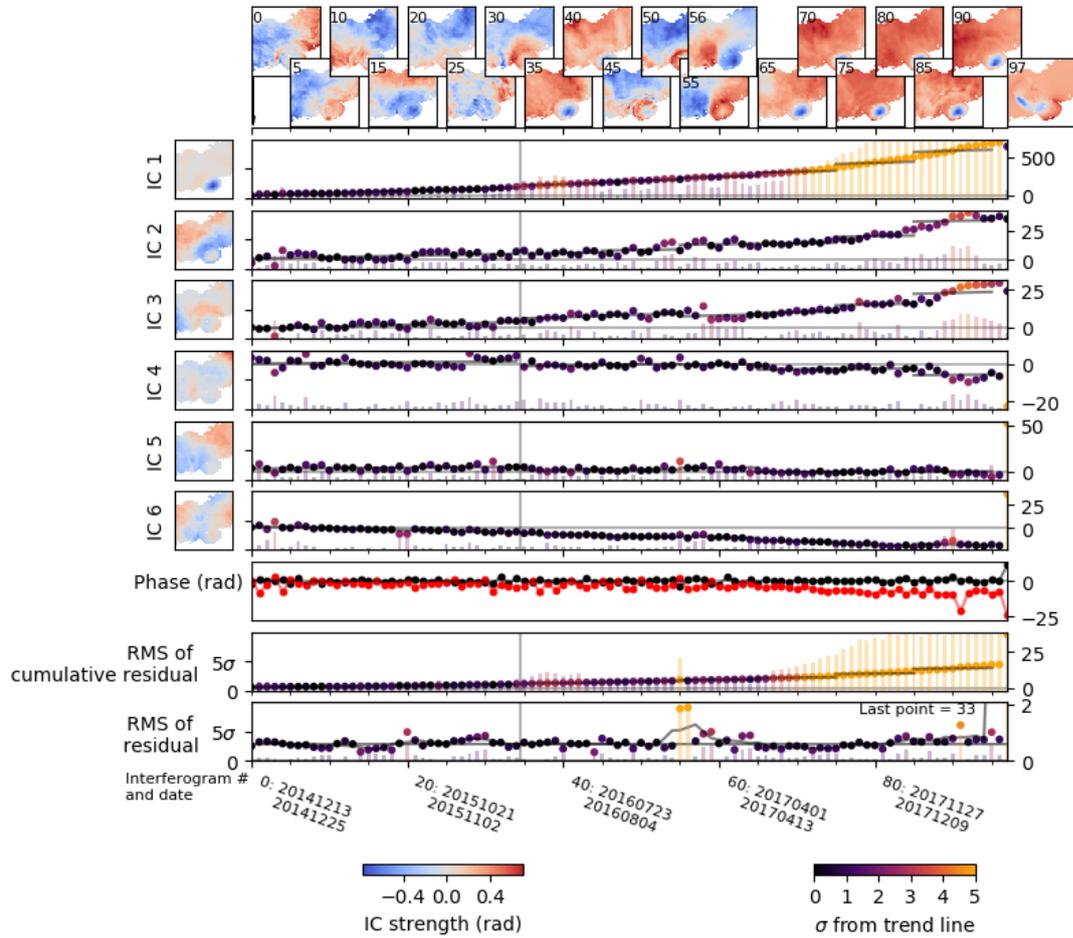
**Figure 3.6:** The centropypes of the six clusters from Figure 3.5 (top row), the topography as described by the SRTM-1 DEM (Farr et al. (2007), bottom row), and  $2D$  Gaussian kernel density estimates between the pixels of the DEM and each source (middle row). Visual comparison of the scatter plots of each IC versus DEM height indicates that IC2 most closely resembles the DEM, whilst the IC2 to DEM scatter plot also shows a clear trend and has the highest correlation coefficient ( $-0.76$ ).

### 3.4 Modelling 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 detection algorithm will identify volcanoes deforming in a similar manner, and modelling of the overpressure could then allow for forecasts of the timing of an eruption (Bato et al., 2017).

Through modelling IC1 as the signal of the caldera pre-eruptive uplift, we can infer information about the source of this deformation, such as the pressure change. We assume that the inflating sill that was modelled as the source of the 2005 eruption remained active, and we seek to constrain the pressure change in this sill between the 2005 and 2018 eruptions.

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 and Brodsky, 2006). This wide range is due to variations of viscosity of magmas of different com-



**Figure 3.7:** The results of applying our automatic detection algorithm to a time series of Sentinel-1 interferograms which cover the final  $\sim 3.5$  years of inflation before the 2018 eruption of Sierra Negra. 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 co-eruptive interferogram is fitted is an order of magnitude larger than seen previously, and is omitted from the RMS residual plot for clarity. In the remaining plots, the point can be seen in the extreme top right. 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.

positions, which impacts the ability of a dyke to continue to propagate without being arrested due to freezing of its walls (Jellinek and 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 ( $8.6 \pm 1.4$  MPa for basalt, to  $13.8 \pm 2.1$  MPa for Granite (Touloukain, 1989; Albino et al., 2018; Pinel et al., 2010)). 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 and Geist, 1995), we expect an overpressure required for eruption at the lower end of the reported values. Through calculating the pre-eruptive 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 originally modelled the deformation using a kinematic approach, with a horizontal rectangular dislocation embedded within an elastic half space with uniform opening (Okada, 1985), and perform a Bayesian inversion using the GBIS software (Bagnardi and Hooper, 2018). Whilst 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 rectangular dislocation (Okada, 1992) implemented in GBIS. Using this approach, we solve for the location (x,y,z), length, width, and pressure change. 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 centred interferograms for the time spanned by the Sentinel-1 time series can be reconstructed using:

$$\mathbf{X}_{ic1} = \mathbf{A}_{ic1} \times \mathbf{S}_{ic1} \quad (3.1)$$

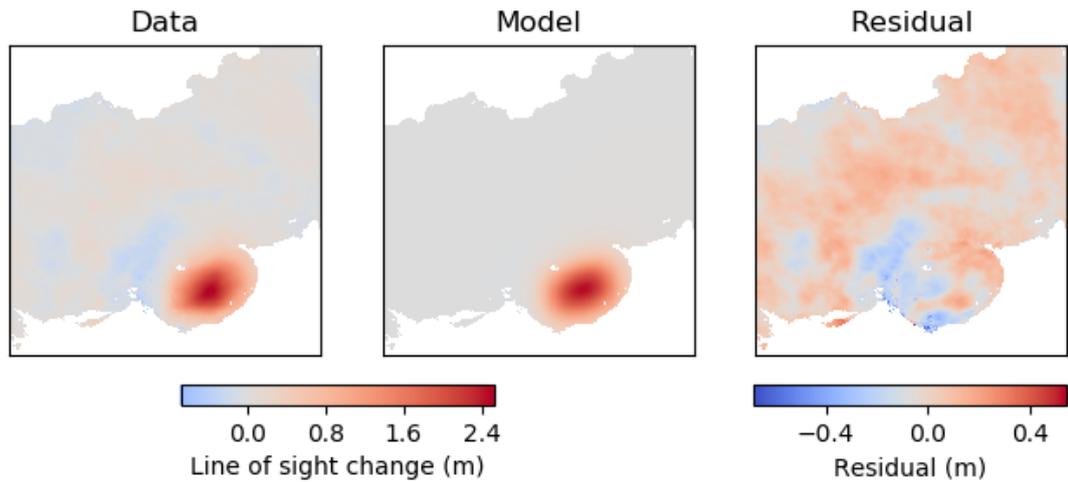
Where  $\mathbf{X}_{ic0}$  are the mean centered interferograms (i.e. the mixtures, in ICA terminology),  $\mathbf{A}_{ic1}$  is the first column of the matrix of time courses, and  $\mathbf{S}_{ic1}$  is the first row of the matrix of recovered sources. As the interferograms were originally created relative to a reference pixel outside the deforming region (located at 91.2 west, 0.7 south) but mean centred 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 centring processes. We then sum the phase change for each pixel in our reconstructed time series to calculate the cumulative motion of IC0 through the time series (shown in Figure 3.8), which we use as the input for our modelling.

Figure 3.8 shows the modelled deformation, and the residual between it and the input data. Our best-fitting model places a  $6.2 \times 3.7 \text{ km}^2$  rectangular dislocation within the caldera at a depth of  $\sim 2.0 \text{ km}$  which, when the shear strength of the crust is set as  $10 \text{ GPa}$  (Jónsson, 2009), has undergone a pressure change of  $10.4 \text{ MPa}$  (probability density functions are provided for model parameters in the supporting information). However, this modelling finds only the pressure change within the modelled sill between the 2018 eruption and the first Sentinel-1 acquisition (2014/12/13), and not since the end of the 2005 eruption. Whilst we could extrapolate the linear inflation seen in the early part of the Sentinel-1 time series (shown as IC0's cumulative time course in Figure 3.7, and in Figure 3.9) back to the end of the 2005 eruption, we instead utilise 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 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 inter-eruptive 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 ten continuous GPS station 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 data with 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 before and after the rate change if they feature the same proportional increase, which we observe them to. We then apply this operation to the entire GPS eastward cumulative displacement time series, and set the displacement to zero after the 2005 eruption. Figure 3.9 shows the results of this process, and visual inspection shows that GPS data is in broad agreement with IC1's time course, with features such as the change in rate in early 2017 occurring in both time series. For comparative purposes, we also fit a linear trend to the initial part of IC1's cumulative time course, and whilst 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 pre-eruptive pressure change, we rescale the pressure change calculated during the Sentinel-1 time series by the ratio of the total GPS derived dis-



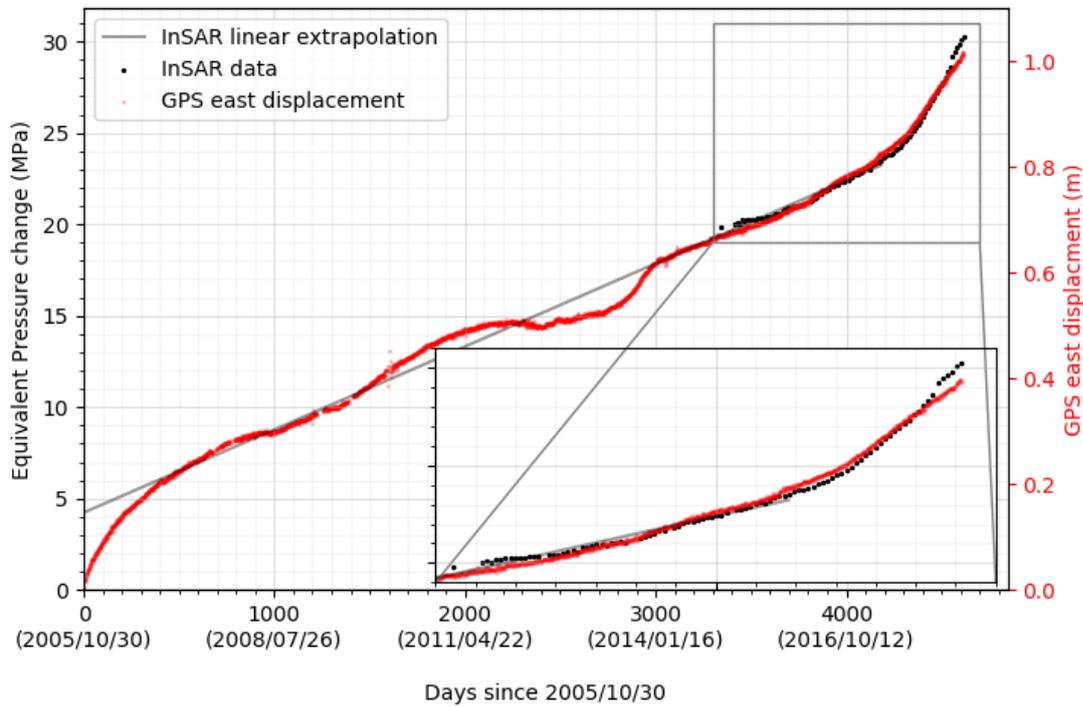
**Figure 3.8:** Data: The signal contained in IC0 throughout the Sentinel-1 time series, showing  $\sim 2.4\text{m}$  of motion towards the satellite during the Sentinel-1 time series. Model: The result of our optimal forward model, which treats the magma chamber as a  $6.2 \times 3.7 \text{ km}^2$  rectangular dislocation at a depth of  $\sim 2.0 \text{ km}$ . Residual: The misfit between our model and the data, which is dominated by a mottled pattern across the majority of the scene which ICA is unable to remove from IC0, and our model is unable to fit.

placement, which we show as the y-axis of Figure 3.9. We find that between the 2005 and 2018 eruptions, the observed surface deformation would translate to a change in magma chamber pressure of  $\sim 30 \text{ MPa}$ .

## 3.5 Discussion

### 3.5.1 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), a newer clustering algorithm (HDBSCAN), and is optimised 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  $I_q$ ) provide three independent measures of the significance and robustness of the recovered sources. Those deemed of geophysical interest can then be utilised for further investigation, in a manner similar to our use of the deformation source in Section 3.4. We utilise ICASAR as an integral part of our 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.



**Figure 3.9:** Temporal evolution of the modelled source. The cumulative time course for component IC1 is shown as black points and is rescaled in terms of the pressure change in the modelled rectangular dislocation, which can be seen to attain a maximum value of  $\sim 30$  MPa before the 2018 eruption. The pre-acceleration section of this time course (prior to day  $\sim 4200$ ) can be fit with a linear trend (grey line), and extrapolated back to the 2005 eruption. However, using GPS station GV01 to constrain the temporal evolution of the deformation shows that whilst 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  $\sim 4$  MPa.

The application of our detection algorithm to Sierra Negra’s pre-eruptive time series has demonstrated its ability to detect both changes in established signals (i.e. the acceleration of the caldera floor uplift), and to detect the emergence of new signals (i.e. those associated with the movement of magma to the surface during the 2018 eruption). Our algorithm is also computationally inexpensive, and can therefore be applied to Sentinel-1 data as they are acquired in order to begin the monitoring of sub-aerial 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 may therefore be able to detect subtle signals that are not clear in individual images (discussed further in the following paragraph and supporting information [Appendix B Text S2 and Figure 5.6]). Additionally, our algorithm’s central tenet of seeking deviations from the baseline behaviour 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. In order to avoid excess false positives, our algorithm is also able to detect when a new signal reverses in a subsequent interferogram, and to interpret this as being due to a strong APS and not require flagging as indicative of unrest.

It is inevitable that during routine monitoring, our algorithm will be faced with both more subtle signals, and noisier interferograms. 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 value of every 10 interferograms. As discussed in Section 3.2.3, the redrawing of the lines of best fit ensure that any small mis-estimations 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 3.2 but featuring a more subtle unrest signal that is not visible in a single 12 day interferogram is provided in the supporting information. During development of the detection algorithm, initial tests suggested that a more sensitive algorithm (i.e. redrawing lines infrequently) reduced the occurrence of false negatives, but increased the occurrence of false positives. Therefore, tuning of the algorithm for use by a specific monitoring agency may involve the calibration of this parameter to achieve the desired operation.

### 3.5.2 Overpressure before the 2018 Sierra Negra Eruption

Our estimation of a pre-eruptive overpressure of  $\sim 30$  MPa in our modelled sill below the Sierra Negra caldera is significantly larger than the overpressure values of  $\sim 1$  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 3.4). 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 pre-eruptive overpressure reduces to  $\sim 13.5$  MPa. Whilst 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 pre-eruptive 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 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 3.7), as both deviate from their baseline rates in the latter stages, and, whilst they are unlikely to represent faulting, may be being used by the algorithm in a new way as it attempts to fit any new signals that may be present in later interferograms. However, this 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.

In addition to faulting reducing the overpressure within the chamber, viscoelastic processes may occur in the country rock and also act to reduce the overpressure. Viscoelastic behaviour of the rocks surrounding the magma chamber below Sierra Negra is likely, as the relatively long-lived nature of the chamber will have raised the temperature of the surrounding country rocks, and this is thought to occur to such an extent that they no longer behave purely elastically (Segall, 2016). Bonafede et al. (1986) showed that for a magma chamber within a purely viscoelastic full space, lower overpressures are required for a given deformation when compared to a model using a purely elastic rheology. This is in agreement with the results of Jellinek and DePaolo (2003), who find that a viscoelastic country rock can inhibit the formation of dykes around large magma chambers through limiting the overpressure within a chamber. However, in this study the authors consider longer time scales, such as  $10^6$  years, which are required for catastrophic caldera forming (CCF) eruptions. The expectation that the omission

of viscoelastic processes leads to overestimations of the overpressure within a chamber is in broad agreement with the results presented in this study (i.e. we estimate an unrealistically high overpressure), and our result provides motivation for further study into viscoelastic processes at volcanoes with long-lived magma chambers.

Application of our 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 3.7). 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 caused 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.

### 3.6 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, is able to automatically determine the number of sources present in the recovered data.

We use our ICASAR algorithm as a constituent of a detection algorithm, which we demonstrate is able to detect signs of volcanic unrest due to both a change in rate of a pre-existing 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 flag this volcano as entering a period of increased activity when the rate of inflation increased approximately one year before the eventual eruption.

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 inter-eruptive inflation to an increase in pressure of 30 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 the spatial pattern of the measured deformation may have changed due to processes such as slip on the intra-caldera faults acting to reduce the overpressure within the sill.

### 3.7 Acknowledgments

MEG is supported by the LiCS grant. M. B. was supported by the NERC Centre for the Observation and Modelling of Earthquakes, Volcanoes and Tectonics (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. 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).

### References

- Albino, F, Falk Amelung, and P. Gregg (2018). “The Role of Pore Fluid Pressure on the Failure of Magma Reservoirs : Insights From Indonesian”. In: *Journal of Geophysical Research : Solid Earth* 123, pp. 1328–1349. DOI: 10.1002/2017JB014523.
- Amelung, Falk, S Jónsson, Howard Zebker, and Paul Segall (2000). “Widespread uplift and ‘trapdoor’ faulting on Galápagos volcanoes observed with radar interferometry.” In: *Nature* 407.6807, pp. 993–996. ISSN: 0028-0836. DOI: 10.1038/35039604. URL: <http://www.ncbi.nlm.nih.gov/pubmed/11069176>.
- Anantrasirichai, N, Juliet Biggs, F. Albino, P. Hill, and D. Bull (2018). “Application of Machine Learning to Classification of Volcanic Deformation in Routinely Generated InSAR Data”. In: *Journal of Geophysical Research : Solid Earth*, pp. 1–15. DOI: 10.1029/2018JB015911.
- Bagnardi, Marco and Falk Amelung (Oct. 2012). “Space-geodetic evidence for multiple magma reservoirs and subvolcanic lateral intrusions at Fernandina Volcano, Galápagos Islands”. In: *Journal of Geophysical Research: Solid Earth* 117.B10, n/a–n/a. ISSN: 01480227. DOI: 10.1029/2012JB009465. URL: <http://doi.wiley.com/10.1029/2012JB009465>.
- Bagnardi, Marco and Andy Hooper (2018). “Inversion of Surface Deformation Data for Rapid Estimates of Source Parameters and Uncertainties : A Bayesian Approach”. In: *Geochemistry, Geophysics, Geosystems* 19, pp. 2194–2211. DOI: 10.1029/2018GC007585.
- Bato, M Grace, Virginie Pinel, and Yajing Yan (2017). “Assimilation of Deformation data for eruption forecasting: potentiality assessment based on synthetic cases”. In: *Frontiers in Earth Science* 5, p. 48.
- Biggs, J., E.Y. Y. Anthony, and C.J. J. Ebinger (Oct. 2009). “Multiple inflation and deflation events at Kenyan volcanoes, East African Rift”. In: *Geology* 37.11, pp. 979–982. ISSN: 0091-7613. DOI: 10.1130/G30133A.1. URL: <http://geology.gsapubs.org/cgi/doi/10.1130/G30133A.1>.

- Blewitt, G, W Hammond, and C Kreemer (2018). “Harnessing the GPS Data Explosion for Interdisciplinary Science”. In: *EOS*. URL: [https://eos.org/project-updates/harnessing-the-gps-data-explosion-for-interdisciplinary-science?utm\\_source=rss&utm\\_medium=rss&utm\\_content=harnessing-the-gps-data-explosion-for-interdisciplinary-science](https://eos.org/project-updates/harnessing-the-gps-data-explosion-for-interdisciplinary-science?utm_source=rss&utm_medium=rss&utm_content=harnessing-the-gps-data-explosion-for-interdisciplinary-science).
- Bonafede, M, M Dragoni, and F Quarenì (1986). “Displacement and stress fields produced by a centre of dilation and by a pressure source in a viscoelastic half-space: application to the study of ground deformation and seismic activity at Campi Flegrei, Italy”. In: *Geophysical Journal International* 87.2, pp. 455–485.
- Campello, R, D Moulavi, A Zimek, and J Sander (2015). “Hierarchical density estimates for data clustering, visualization, and outlier detection and Outlier Detection”. In: *ACM Transactions on Knowledge Discovery from Data, New York*, pp. 1–5.
- Carn, S A, N A Krotkov, B L Fisher, C Li, and A J Prata (2018). “First observations of volcanic eruption clouds from the L1 Earth-Sun Lagrange point by DSCOVR / EPIC”. In: *Geophysical Research Letters*. DOI: 10.1029/2018GL079808.
- Chadwick, William W, D Geist, Sigurjon Jonsson, M.P. Poland, Daniel Johnson, and C. Meertens (2006). “A volcano bursting at the seams: Inflation, faulting, and eruption at Sierra Negra volcano, Galapagos”. In: *Geology* 12, pp. 1025–1028. DOI: 10.1130/G22826A.1.
- Chen, Curtis W and Howard A Zebker (2001). “Two-dimensional phase unwrapping with use of statistical models for cost functions in nonlinear optimization”. In: *Journal of the Optical Society of America* 18.2, pp. 338–351.
- Demartines, Pierre and J Hérault (1997). “Curvilinear Component Analysis: A Self-Organizing Neural Network for Nonlinear Mapping of Data Sets”. In: *IEEE Transactions on Neural Networks* 8.1, pp. 148–154.
- Ebmeier, S. K. (2016). “Application of Independent Component Analysis to multi-temporal InSAR data with volcanic case studies”. In: *Journal of Geophysical Research: Solid Earth*, pp. 8970–8986. ISSN: 21699313. DOI: 10.1002/2016JB013765. URL: <http://doi.wiley.com/10.1002/2016JB013765>.
- Farr, T.G., P.a. Rosen, Edward Caro, and Robert Crippen (2007). “The Shuttle Radar Topography Mission”. In: *Reviews of Geophysics* 45.2005, pp. 1–33. DOI: 10.1029/2005RG000183.1.
- Fialko, Yuri, Yakov Khazan, and Mark Simons (2001). “Deformation due to a pressurized horizontal circular crack in an elastic half-space, with application to volcano geodesy”. In: *Geophys. J. Int.* 146, pp. 181–190.
- Gaddes, M. E. (2017). *ICASAR*. <https://github.com/mattcyp88/ICASAR>.
- Gaddes, ME, A Hooper, M Bagnardi, H Inman, and F Albino (2018). “Blind signal separation methods for InSAR: The potential to automatically detect and monitor signals of volcanic deformation”. In: *Journal of Geophysical Research: Solid Earth*.

- Goldstein, Richard M. and Charles L. Werner (1998). “Radar interferogram filtering for geophysical applications”. In: *Geophysical Research Letters* 25.21, p. 4035. ISSN: 0094-8276. DOI: 10.1029/1998GL900033.
- González, Pablo J., Richard J. Walters, E. L. Hatton, Karsten Spaans, A. McDougall, Andrew Hooper, and Tim J. Wright (2016). “LiCSAR: Tools for automated generation of Sentinel-1 frame interferograms”. In: *AGU Fall meeting*.
- Himberg, Johan, Aapo Hyva, and Fabrizio Esposito (2004). “Validating the independent components of neuroimaging time series via clustering and visualization”. In: *NeuroImage* 22, pp. 1214–1222. DOI: 10.1016/j.neuroimage.2004.03.027.
- Hooper, Andrew, Benedikt Ófeigsson, Freysteinn Sigmundsson, Björn Lund, Páll Einarsson, Halldór Geirsson, and Erik Sturkell (2011). “Increased capture of magma in the crust promoted by ice-cap retreat in Iceland”. In: *Nature Geoscience* 4.11, p. 783.
- Hooper, Andrew, Paul Segall, Kaj Johnson, and Justin Rubinstein (2002). “Reconciling seismic and geodetic models of the 1989 Kilauea south flank earthquake”. In: *Geophysical research letters* 29.22.
- Hunter, J. D. (2007). “Matplotlib: A 2D graphics environment”. In: *Computing In Science & Engineering* 9.3, pp. 90–95.
- Hyvärinen, Aapo (2011). “Testing the ICA mixing matrix based on inter-subject or inter-session consistency”. In: *NeuroImage* 58.1, pp. 122–136. ISSN: 1053-8119. DOI: 10.1016/j.neuroimage.2011.05.086. URL: <http://dx.doi.org/10.1016/j.neuroimage.2011.05.086>.
- (2012). “Independent component analysis : recent advances”. In: *Philosophical Transactions of The Royal Society A*.
- Hyvärinen, Aapo and Pavan Ramkumar (2013). “Testing independent component patterns by inter-subject or inter-session consistency”. In: *Frontiers in Human Neuroscience* 7.March, pp. 1–15. DOI: 10.3389/fnhum.2013.00094.
- Jellinek, A Mark and Donald J DePaolo (2003). “A model for the origin of large silicic magma chambers: precursors of caldera-forming eruptions”. In: *Bulletin of Volcanology* 65.5, pp. 363–381.
- Jónsson, Sigurjón (June 2009). “Stress interaction between magma accumulation and trapdoor faulting on Sierra Negra volcano, Galápagos”. In: *Tectonophysics* 471.1-2, pp. 36–44. ISSN: 00401951. DOI: 10.1016/j.tecto.2008.08.005. URL: <http://linkinghub.elsevier.com/retrieve/pii/S0040195108003880>.
- Jónsson, Sigurjón, Howard Zebker, and Falk Amelung (2005). “On trapdoor faulting at Sierra Negra volcano , Galapagos”. In: *Journal of Volcanology and Geothermal Research* 144, pp. 59–71. DOI: 10.1016/j.jvolgeores.2004.11.029.
- Loughlin, Susan C, Robert Stephen John Sparks, Sarah K Brown, Susanna F Jenkins, and Charlotte Vye-Brown (2015). *Global volcanic hazards and risk*. Cambridge University Press.

- Maaten, Laurens Van Der and Geoffrey Hinton (2008). “Visualizing Data using t-SNE”. In: *Journal of Machine Learning Research* 9, pp. 2579–2605.
- Manga, Michael and Emily Brodsky (2006). “Seismic Triggering of Eruptions in the Far Field : Volcanoes and Geysers”. In: *Annual Review of Earth and Planetary Science*. DOI: 10.1146/annurev.earth.34.031405.125125.
- McInnes, Leland, John Healy, and Steve Astels (2017). “hdbscan : Hierarchical density based clustering”. In: *The Journal of Open Source Software* 2, pp. 11–12. DOI: 10.21105/joss.00205.
- Meinecke, Frank, Andreas Ziehe, and Motoaki Kawanabe (2002). “A Resampling Approach to Estimate the Stability of One-Dimensional or Multidimensional Independent Components”. In: *IEEE Transactions on Biomedical Engineering* 49.12, pp. 1514–1525.
- Mogi, Kiyoo (1958). “Relations between the eruptions of various volcanoes and the deformations of the ground surfaces around them”. In: *Bulletin of the Earthquake Research Institute* 36, pp. 99–134. ISSN: 0040-8972. DOI: 10.1016/j.epsl.2004.04.016.
- Okada (1985). “Surface deformation due to shear and tensile faults in a half-space”. In: *International Journal of Rock Mechanics and Mining Sciences Geomechanics Abstracts* 75.4, pp. 1135–1154. ISSN: 00371106. DOI: 10.1016/0148-9062(86)90674-1. URL: <http://linkinghub.elsevier.com/retrieve/pii/0148906286906741>.
- Okada, Yoshimitsu (1992). “Internal deformation due to shear and tensile faults in a half-space”. In: *Bulletin of the seismological society of America* 82.2, pp. 1018–1040.
- Pagli, Carolina, Freysteinn Sigmundsson, Thóra Arnadóttir, Páll Einarsson, and Erik Sturkell (2006). “Deflation of the Askja volcanic system: constraints on the deformation source from combined inversion of satellite radar interferograms and GPS measurements”. In: *Journal of Volcanology and Geothermal Research* 152.1-2, pp. 97–108.
- Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay (2011). “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12, pp. 2825–2830.
- Pinel, V, C Jaupart, and F Albino (2010). “On the relationship between cycles of eruptive activity and growth of a volcanic edifice”. In: *Journal of Volcanology and Geothermal Research* 194.4, pp. 150–164. ISSN: 0377-0273. DOI: 10.1016/j.jvolgeores.2010.05.006. URL: <http://dx.doi.org/10.1016/j.jvolgeores.2010.05.006>.
- Reynolds, Robert W, Dennis Geist, and M. Kurz (1995). “Physical volcanology and structural development of Sierra Negra volcano , Isabela Island , Galapagos archipelago”. In: *GSA Bulletin* 107.12, pp. 1398–1410.

- Reynolds, Robert W and J Geist (1995). “Petrology of lavas from Sierra Negra volcano, Isabela Island, Galapagos archipelago”. In: *Journal of Geophysical Research* 100.
- Rodd, Rebecca L, Jonathan M Lees, and Gabrielle Tepp (2016). “Three-dimensional attenuation model of Sierra Negra Volcano , Galápagos Archipelago”. In: *Geophysical Research Letters* 43, pp. 6259–6266. DOI: 10.1002/2016GL069554. Received.
- Schultz, RA (1995). “Limits on strength and deformation properties of jointed basaltic rock masses”. In: *Rock Mechanics and Rock Engineering* 28.1, pp. 1–15.
- Segall, Paul (2016). “Repressurization following eruption from a magma chamber with a viscoelastic aureole”. In: *Journal of Geophysical Research: Solid Earth* 121.12, pp. 8501–8522.
- Sennert, S K (ed.) (2018a). “Global Volcanism Program, 2018. Report on Sierra Negra (Ecuador)”. In: *Smithsonian Institution and US Geological Survey Weekly Volcanic Activity Report*, 18 July-24 July 2018.
- (2018b). “No Title Global Volcanism Program, 2018. Report on Sierra Negra (Ecuador)”. In: *Smithsonian Institution and US Geological Survey Weekly Volcanic Activity Report*, 29 August-4 September 2018.
- Stock, Michael J, Marco Bagnardi, David A Neave, John MacLennan, Benjamin Bernard, Iris Buisman, Matthew LM Gleeson, and Dennis Geist (2018). “Integrated petrological and geophysical constraints on magma system architecture in the western Galápagos Archipelago: insights from Wolf volcano”. In: *Geochemistry, Geophysics, Geosystems* 19.12, pp. 4722–4743.
- Tepp, Gabrielle, Cynthia J. Ebinger, Mario Ruiz, and Manahloh Belachew (2014). “Imaging rapidly deforming ocean island volcanoes in the western Galápagos archipelago, Ecuador”. In: *Journal of Geophysical Research : Solid Earth* 119, pp. 442–463. DOI: 10.1002/2013JB010227. Received.
- Touloukain, Y. S. (1989). “Physical properties of rocks and minerals”. In: *CINDAS Data Series on Materials Properties , 63- 76*, pp. 63–76.
- Vasconez, By Francisco, Patricio Ramón, Stephen Hernandez, Silvana Hidalgo, and Benjamin Bernard (2018). “The different character of the recent eruptions of Galápagos shield volcanoes ( Ecuador ): La Cumbre vs . Sierra Negra”. In: *EarthArXiv Preprints*, pp. 1–11. DOI: <https://doi.org/10.31223/osf.io/t7f6k>.
- Yun, S, P Segall, and H Zebker (2006). “Constraints on magma chamber geometry at Sierra Negra Volcano , Galapagos Islands , based on InSAR observations”. In: *Journal of Volcanology and Geothermal Research* 150, pp. 232–243. DOI: 10.1016/j.jvolgeores.2005.07.009.

## Chapter 4

# Classifying and Locating Deformation in SAR Interferograms Using Deep Learning

M. Gaddes<sup>1</sup>, A. Hooper<sup>1</sup>, and F. Albino<sup>2</sup>

<sup>1</sup> *COMET, School of Earth and Environment, University of Leeds, United Kingdom*

<sup>2</sup> *University of Bristol, United Kingdom*

## Abstract

With the evolution of InSAR into a tool for active hazard monitoring, new methods are sought to quickly and automatically interpret the large number of interferograms that are created. In this work, we present a convolutional neural network (CNN) that is able to classify types of deformation seen in interferograms, and to locate deformation signals within an interferogram. We achieve this through creating a large dataset of synthetic interferograms which feature labels of both the type and location of any deformation, and can be used to train our model. We also find that our model's performance is improved through the inclusion of a small amount of real data. When building models of this type, it is common for some of the weights within the model to be transferred from other models designed for different problems. Consequently, we also investigate how to best organise interferograms such that the filters learned in models such as VGG16 are sensitive to the signals of interest in interferograms.

## 4.1 Introduction

In recent years, work to extend volcano monitoring to all of the world's  $\sim 1400$  subaerial volcanoes has resulted in the application of a diverse set of machine learning methods to ground deformation maps produced by interferometric RADAR satellites (InSAR). Work by Anantrasirichai et al. (2018) has used convolutional neural networks (CNNs) to determine if individual interferograms contain deformation, whilst work by Gaddes et al. (2018) has used blind signal separation methods to determine if a time series of interferograms show signs of unrest. However, in both of the examples detailed above, each algorithm demonstrates very limited knowledge of the diverse types of deformation that may be measured at volcanoes. The algorithm presented in Anantrasirichai et al. (2018) assigns all data containing deformation to one label, whilst the algorithm presented in Gaddes et al. (2018) is only able to alert a user to changes in the signals present. Consequently, we seek to improve upon these approaches by developing a CNN that is able to differentiate between different types of deformation, and to detect the spatial extent of it without using a sliding window approach.

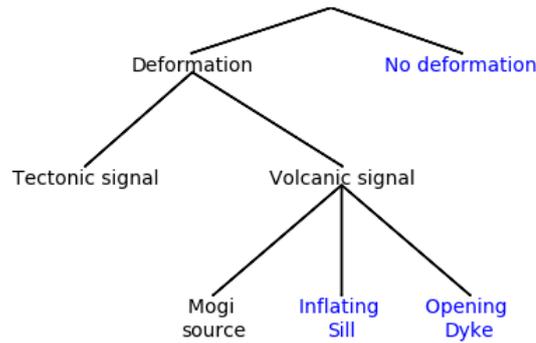
Detecting the spatial extent of an object is referred to as localisation in machine learning parlance, and a variety of methods exist to perform it. For the simple case in which only one classification driving object features in an image, this is commonly approached using one of two methods. In the first, the CNN is trained on relatively small images of the objects of interest (e.g.  $224 \times 224$ ), before the trained model is then used on larger images (e.g.  $1000 \times 500$ ) that are subdivided into smaller patches of equal resolution to the original training data. This approach is utilised in Anantrasirichai et al. (2018), who avoid the potentially large computation cost of the repeated forward passes by using the AlexNet CNN (Krizhevsky et al., 2012), which requires relatively

few operations to complete a forward pass through the model (Canziani et al., 2016). Additionally, this approach has the limitation that the CNN does not need to learn how to determine the location of the object of interest, and at a more fundamental level remains a classification model which has not learned to evaluate the spatial extent of different signals.

However, in the field of computer vision, CNNs have been developed that are able to both classify an image as containing an object, and describe the object's location. The location of an object is either indicated through encompassing it in a rectangle (e.g. Simonyan and Zisserman (2014) and Redmon et al. (2016) ) or, in more complex algorithms, indicating the exact outline of an object by identifying which pixels comprise it (e.g. He et al. (2017)). It would be expected that both types of models would provide more detailed information on the spatial extent of a signal of interest than a classification model that is repeatedly used on different areas of an image. Consequently, we endeavour to develop an algorithm that is able to both classify types of deformation, and localise it within an interferogram in one step. Figure 4.1 shows our initial division of deformation patterns into different classes, and can be considered similar to the WordNet hierarchy (Fellbaum, 1998) that underpins ImageNet (Deng et al., 2009).

When seeking to build a CNN to perform a classification or localisation problem, common approaches can be divided into one of three broad categories depending on the utilisation of pre-existing models. In the most fundamental case, a new model is designed before all the parameters within it are trained (e.g. Rauter and Winkler (2018)), but this approach has the risk of failing to utilise the successful applications of CNNs to other problems. Consequently, it is possible for the majority of the architecture of a model that is (or was) state of the art for a certain problem to be re-trained to solve the new problem. As many CNNs feature a fully connected network after the convolutional layers, it is common to retain the convolutional layers and design a new fully connected network that outputs the classes of interest. However, this approach still requires the training of a CNN that is likely to contain tens of millions of parameters, which will be both computationally expensive, and require a large volume of training data. AlexNet, a previously state-of-the-art image classification CNN, has 60 million parameters, was trained on 1.2 million images, and even when implemented on GPUs took around one week to train (Krizhevsky et al., 2012). Therefore, a common approach termed transfer learning is to retain both the structure and weights of the initial convolutional layers, and to train only the last fully connected part of the network. This approach was successfully used by Anantrasirichai et al. (2018), who evaluated several CNNs, before using the structure and weights of AlexNet with their own fully connected classifier to output whether an interferogram contained deformation or not.

The weights learned in the convolutional filters of a CNN are of great importance



**Figure 4.1:** Proposed hierarchy for signals of interest in interferograms at volcanic centres. We propose a model that is able to classify interferograms as either containing only atmospheric signals, or as containing deformation due to inflating sills or opening dykes. We choose these classes for our initial study as they are likely to present significantly different signals, in contrast to the harder problem of differentiating between signals that are best modelled as spherical chambers (i.e. Mogi sources) and those as sills. As our proposed model will work with only data from one look angle, we envisage that deformation due to processes that could be modelled as a Mogi source are likely to be incorporated in the inflating sill label. We do not present this hierarchy as complete, and envisage that future studies may add further subtrees, such as signals due to the cooling and contraction of emplaced lava flows.

to a network’s ability to detect features, as the filters must be sensitive to the patterns that these features present in an image. As networks such as AlexNet (Krizhevsky et al., 2012) and VGG16 (Simonyan and Zisserman, 2014) were originally developed to compete in the ImageNet competitions (Deng et al., 2009), the filters have been trained to detect the type of features present in natural images (e.g. photographs of a person, or car). When performing transfer learning, it is these filters that must be sensitive to the patterns presented in a deformation signal if the network is to correctly classify and locate it. However, as interferograms can be expressed in differing formats we also seek to explore which of these formats allows for the filters in models trained on natural images to excel.

## 4.2 Classification with different data formats

As the most common CNNs for computer vision are trained on images comprising of a channel for each of the red, green, and blue values for each pixel, other data that are to be used with the network must also be the channel. However, when considering an image of interferometric phase, these images contain only a single value for each pixel, and so consist of only one channel, and are analogous to a greyscale image. This difference in the number of channels can be circumvented through duplicating the one channel interferogram in each of the three input channels of a CNN, but in this section of our study we wish to determine if this approach can be improved upon.

When two SAR images are combined to form a single interferogram, the resulting

image is a  $2D$  array of complex numbers. Whilst the magnitude of each of these complex numbers relates to the brightness of a given pixel, it is common for only the argument to be displayed, as these phase values can be used to infer ground movement. However, the phase values of an interferogram are wrapped in the range  $[-\pi \pi]$  as only the fractional part of the phase value can be measured, but this ambiguity can be solved for to produce an unwrapped interferogram (Chen and Zebker, 2001). We postulate that in addition to the use of either wrapped or unwrapped data duplicated to fill three channels, the original complex numbers of an interferogram could be used in two channels in order for our network to also access information about the brightness of each pixel.

However, this approach can be expanded to feed more data into the CNN. When a human observer interprets an interferogram, they are likely to use data such as a digital elevation model (DEM) as this can be used to help determine if a signal is due to deformation, or due to a topographically correlated atmospheric phase screen (Bekaert et al., 2015). Consequently, we postulate that the inclusion of a DEM to our CNN will improve its performance, and seek to investigate this whilst varying the inputs across different channels.

To perform this analysis, we first synthesise a dataset of 40,000 labelled interferograms. The collection of enough labelled data to train a CNN is commonly time consuming or expensive, and we find that the addition of localisation labels to our data makes it more time consuming than in previous studies, due to the need to manually draw rectangles that outline each deforming region. Additionally, due to the large number of data that CNNs require to train and our expansion to classification of different types of deformation, procuring enough real data to do this may be not possible. Consequently, we perform this analysis using only synthetic data. Following the hierarchy proposed in Figure 4.1, we create interferograms that contain either no deformation, deformation due to an opening dyke, or deformation due to an inflating sill. We model the dykes and sills as approximately vertical and horizontal dislocations, respectively, with uniform opening in an elastic half space (Okada, 1985). For the set of sills, we randomly selects strikes in the range  $0 - 359^\circ$ , dips in the range  $0 - 5^\circ$ , openings in the range  $0.2 - 1$  m, depths in the range  $1.5 - 3.5$  km, and widths and lengths in the range  $2 - 6$  km. For the set of dykes, we randomly select strikes in the range  $0 - 359^\circ$ , dips in the range  $75 - 90^\circ$ , openings in the range  $0.1 - 0.7$  m, top depths in the range  $0 - 2$  km, bottom depths in the range  $0 - 8$  km, and lengths in the range  $0 - 10$  km. These deformation patterns are then combined with a topographically correlated atmospheric phase screen (APS), and a turbulent APS, which we discuss generating in more detail in Gaddes et al. (2018). We calculate the topographically correlated APS using a random selection of DEMs that cover areas of volcanic interest from the the SRTM 90m DEM (Farr et al., 2007), and use the coastline information contained within the

product to mask areas of water. We also synthesise areas of incoherence within our interferograms as spatially correlated noise with different length scales, which we mask in order for our synthetic interferograms to be as similar as possible to the Sentinel-1 interferograms automatically created by the LiCSAR processor (González et al., 2016). Figure 4.2 shows the results of mixing these different elements to create our synthetic interferograms.

This process creates unwrapped data which can be converted to wrapped data through finding modulo  $2\pi$  of the unwrapped phase. However, to synthesise both the real and imaginary part of a complex interferogram requires knowledge of both the brightness of a pixel and its phase. To achieve this, we again use the SRTM DEM, and calculate the intensity of reflected electromagnetic radiation at the angles of incidence used by the Sentinel-1 satellites ( $29.1 - 46.0^\circ$ ), before adding speckle noise. As inputs to CNNs that are to be trained using transfer learning must be rescaled to the inputs used in the original training, we use only relative values in the range  $(-1) - 1$  for the synthetic intensities. With knowledge of the modulus (relative intensity) and argument (wrapped phase) of each pixel of our synthetic interferogram, the real/imaginary components are simply the products of the modulus and cosine/sine of the argument, respectively. Figure 4.3 shows five different ways we can represent an interferogram using the three channels available.

The CNN we build to classify the synthetic interferograms uses the five convolutional blocks of VGG16 (Simonyan and Zisserman, 2014), with our own fully connected network after this. When an interferogram of shape  $(224 \times 224 \times 3)$  is passed through the convolutional layers of VGG16, it is transformed into a tensor of shape  $(7 \times 7 \times 512)$ . This is then flattened to make a vector of size 25088, before being passed through fully connected layers of size 256, 128, and an output layer of size three (i.e., dyke, sill, or no deformation). To produce a set of outputs that can be used as probabilities, we use a softmax activation for the last layer (Bridle, 1990), as for a given input this produces a set of outputs that sum to one, and represent the confidence that the model places on its prediction. For the remaining layers we use rectified linear units (ReLus) to reduce computation time (Agostinelli et al., 2014). As our model seeks to solve a classification problem, we use categorical cross entropy for the loss function, which we seek to reduce using the Nadam optimizer as this does not require the choice of a learning rate (Dozat, 2016).

A common problem of CNNs that are used for classification can be overfitting of the training data, which results in a model that generalises to new data poorly. We endeavour to limit this through the use of dropout (Srivastava et al., 2014) before both the 256 and 128 neuron layers, as through randomly removing some connections during each pass of the data through our model, we hope that our model is forced to learn more robust representations of the training data. As we use synthetic data, we are not

limited by the usual cost of collecting labelled data, and therefore are able to generate 20000 unique interferograms without the use of data augmentation.

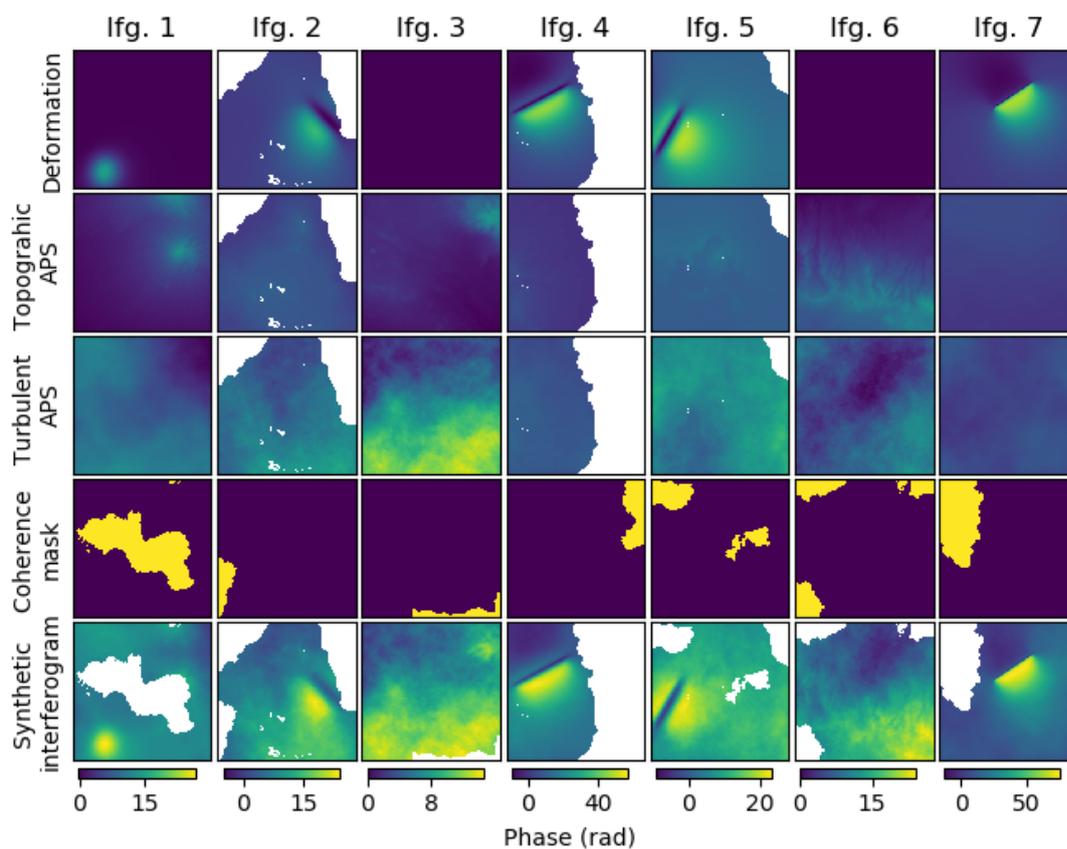
Figure 4.4 shows the results of training five models with each of the data formats previously discussed. The highest classification accuracy achieved is  $\sim 0.95$ , which is achieved when the models are trained with either wrapped or unwrapped data repeated across the three input channels. Inclusion of the DEM as the third channel appears to reduce classification accuracy, whilst very low accuracies are achieved in the real and imaginary channel case. We discuss these results in more detail in Section 4.4, but for the remainder of the paper we choose to work with data that is unwrapped and repeated across the three input channels. We choose this approach as no significant differences are seen between the classification accuracy ultimately achieved with either wrapped or unwrapped data, but the use of unwrapped data may allow for a model to be used with unwrapped time series, and so detect subtle signals produced by low strain rate processes. Additionally, a model that works with unwrapped data may also provide the opportunity to be expanded to locate and classify unwrapping errors automatically, but it should be noted that unwrapping is a computationally expensive process.

## 4.3 Classification and localisation

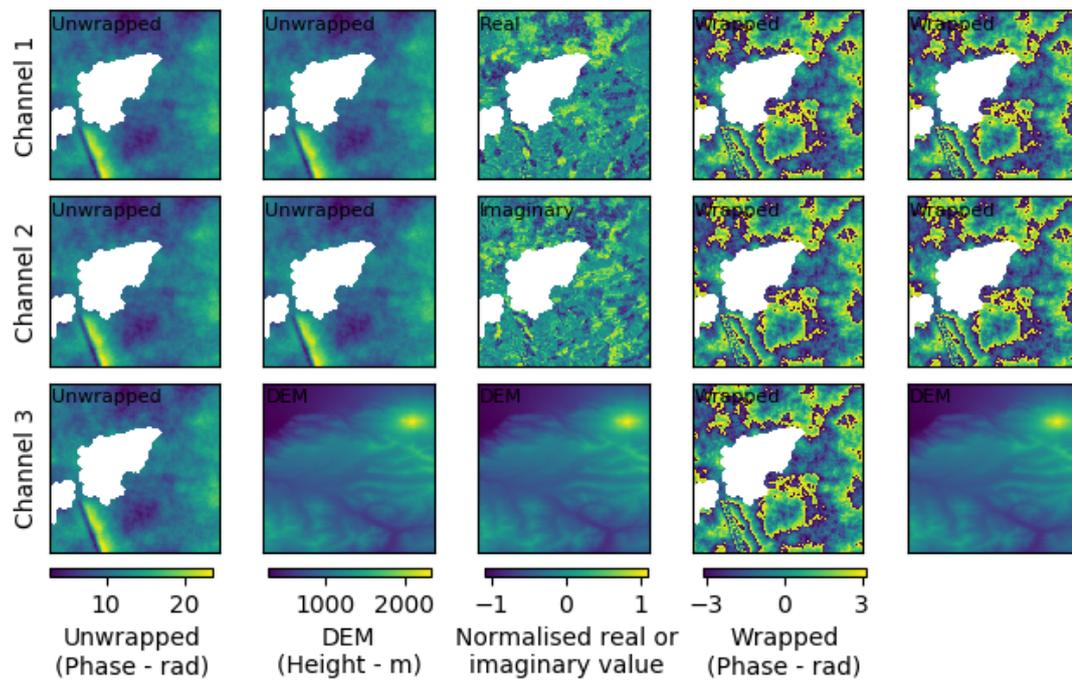
### 4.3.1 Using synthetic data

In the previous section, we demonstrated that when using VGG16 with convolutional weights learned on ImageNet data, optimal performance for classifying synthetic interferograms is achieved when the unwrapped phase is repeated across the three input channels. In this section, we build on the model used to perform classification by adding localisation output. We also endeavour to ascertain if the expense of collecting labelled data can be avoided entirely through the continued use of synthetic data when training our model.

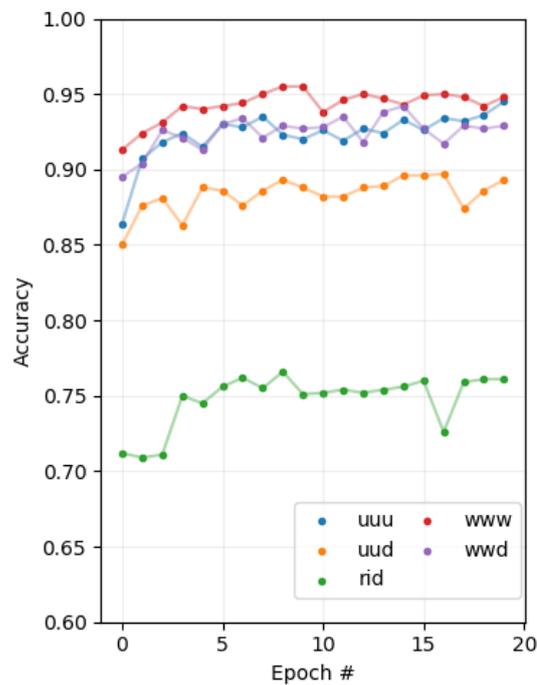
We achieve both classification and localisation through dividing the fully connected section of our model to produce two distinct outputs. One output returns the class of the input data in the manner described in Section 4.2, whilst the second returns the location of any deformation within the scene. In machine learning parlance, models of this type are termed double headed, and we subsequently refer to either of the outputs and their corresponding preceding layers as either the classification head or localisation head. Figure 4.5 shows the structure of the two heads, and how they diverge after the output of the fifth block of VGG16 has been flattened. The localisation head is structured in a similar manner to the model described in Simonyan and Zisserman (2014), in which the model conveys the location of any deformation through outputting a column vector containing four values. Two of these values determine the centre of the deformation pattern and two display its horizontal and vertical extent, and together



**Figure 4.2:** An example of the constituent parts of seven synthetic interferograms. Two of the interferograms do not feature deformation (e.g. interferogram 3), a third feature deformation due to an inflating sill (e.g. 1), and a third feature deformation due to an opening dyke (e.g. 2). These signals are geocoded and areas of water masked, before being combined with a topographically correlated APS, and a turbulent APS. Areas of incoherence are also synthesised, and these are used to mask the combination of the three signals to create the final synthetic interferograms.



**Figure 4.3:** Organisation of an interferogram into three channel form. Columns one and two feature unwrapped data that is repeated, and in column two the DEM is included as the third channel. In column three the real and imaginary elements of the complex values of each pixel of an interferogram occupy channels one and two, whilst the DEM is included in the third. Columns three and four feature wrapped data that is repeated, and in column five the DEM is included as the third channel.



**Figure 4.4:** Accuracy of classifying validation data during training using three channel data arranged in different formats. “u”: unwrapped data, “w”: wrapped data, “d”: DEM, “r” real component of interferogram, “i”: imaginary component of interferogram. Low accuracy is seen for the “rid” data, and in both the wrapped and unwrapped cases inclusion of the DEM in the third channel is seen to degrade classification accuracy. At the end of the 20 epochs of training, only a small difference is seen in accuracy between wrapped and unwrapped data, with both classifying ~95% of the validation data correctly.

can be used to construct a box encompassing a deformation pattern.

However, we find the localisation problem significantly harder to solve than the classification problem, which results in a more complex classification head. To reduce the time taken to develop and test possible localisation heads, we perform what is termed bottleneck learning in machine learning literature. This involves first computing the results of passing our entire dataset through the first five blocks of VGG16, before then training only the fully connected parts of our network. As the pass of the data through the convolutional layers of VGG16 is computationally expensive yet we do not seek to update the weights within these layers, this method avoids these costly and unnecessary repeat passes. Experimentation finds that the simplest model capable of good performance is achieved with five layers consisting of 2048, 1024, 512, 128, and 4 neurons.

When training our model, we use the mean squared error between the predicted location vector and the labelled location vector as our localisation loss function, which we seek to minimise. However, when using a double headed network, training is complicated by the fact that the model's overall loss is now a combination of the classification and localisation loss, which must be balanced using a hyperparameter commonly termed loss weighting. We experiment with this hyperparameter, and find that a value of 0.95 for the classification loss and 0.05 for the localisation loss provides a good balance between the two outputs. We believe that this value proves optimal as the localisation loss is significantly larger than the classification loss, but by weighting them unequally they then contribute to the overall loss approximately equally. In a similar manner to the design of a localisation head, the time required for the repeated model runs required to fine tune this hyperparameter is greatly reduced by first computing bottleneck features.

Figure 4.6 shows the results of training our classification and localisation model, which due to the computation of bottleneck features took under two days without the use of GPUs (a machine equipped with dual core Intel Xeon E5-2640s was used). During the training of our model, inspection of both the training and validation loss does not show the characteristic minima in validation loss being passed, despite continued decrease in the training loss that is indicative of a model that is overfitting. To improve the performance of our network, we also seek to improve the filters learned within the convolutional blocks in order for them to be better suited to our task. We perform this by changing the style of learning after the 10th epoch, and switch from updating only the fully connected layers, to also including the 5th convolutional block in our updates. However, if we were to resume training the network with an optimiser such as Nadam, which features an adaptive learning rate, only a small number of initial steps at too high a learning rate would quickly destroy the finely tuned values in both the convolutional blocks of VGG16, and our fully connected classification and localisation heads. We circumvent this through switching the optimizer to stochastic

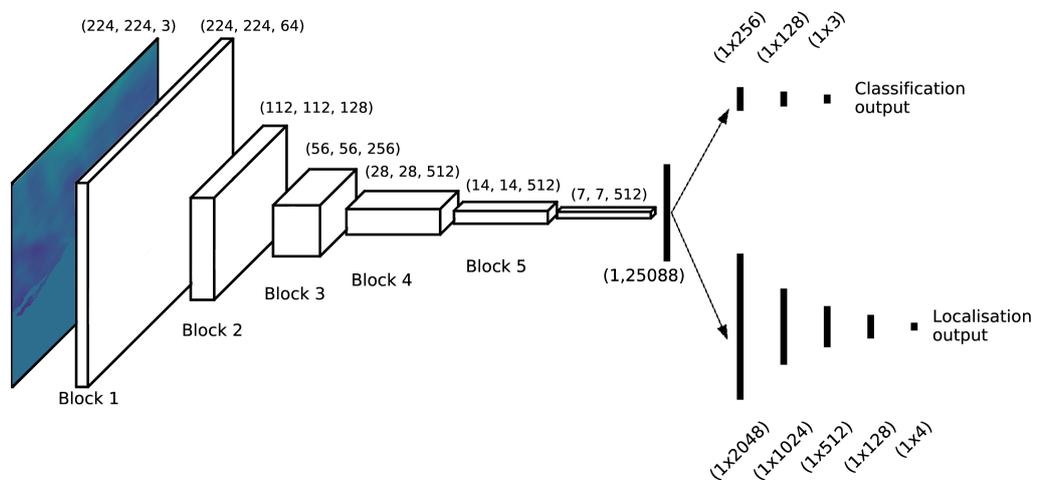
gradient descent (SGD) and setting the learning rate manually. However, as we are now updating the convolutional blocks of VGG16, we cannot simply use the bottleneck features we previously computed, and must instead perform the time consuming pass of the data through VGG16 at each step. This complicates the search for an optimal learning rate, but we find that a value of  $1.5 \times 10^{-8}$  does not degrade the performance already gained during training, but still allows for an increase in the performance of the localisation head.

Figure 4.7 shows the results of applying our trained classification and localisation model to a random selection of the testing data (i.e., data that the model was not exposed to during training). In each case, the classification can be seen to be accurate, and the localisation approximately correct. When considering entire test set of data, the classification accuracy is 0.89, whilst the localisation loss is 169. It should be noted that we could also report the classification loss (0.31), but we believe this is less useful than the accuracy. However, in the localisation case, accuracy is not a meaningful measure of the fidelity of the output, as it is instead a regression problem in which we aim to approximate the correct values which are continuous in nature. As the localisation loss is the mean squared error of the difference between the localisation coordinates and the predicted coordinates, a value of 169 indicates that the predictions are typically misplaced by around 13 pixels, which translates to 1.2 km when using 3 arcsecond pixels.

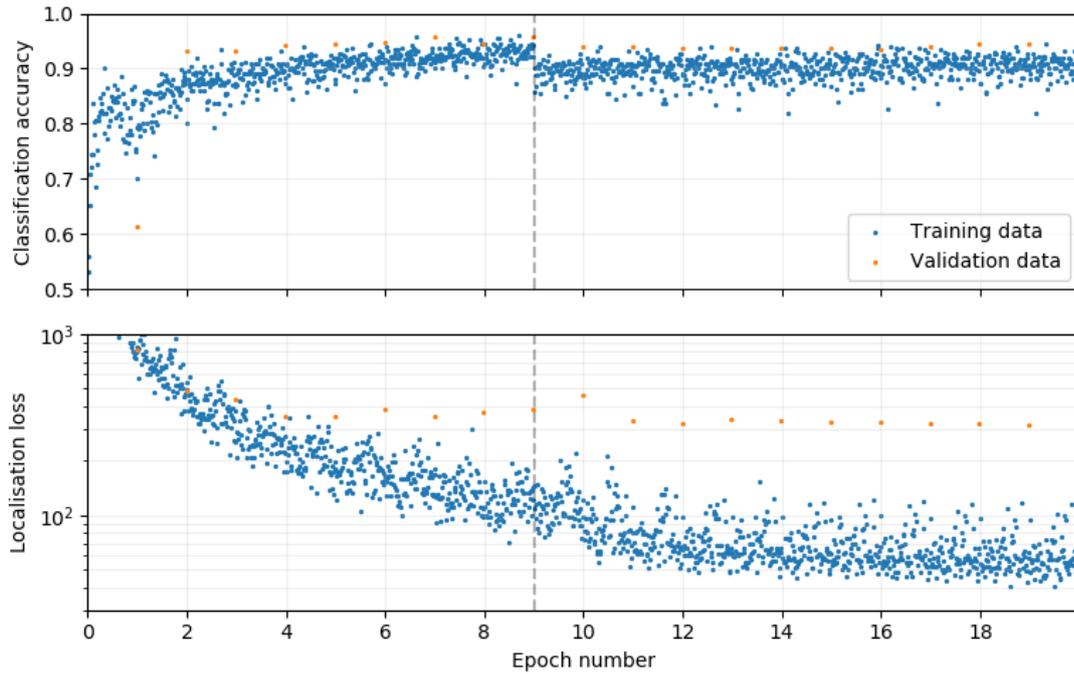
### 4.3.2 Application to real data

Whilst the model described in the previous section achieved good performance when classifying and locating deformation in synthetic interferograms, for use in automatic detection algorithms we wish for our CNN to be able to work with Sentinel-1 data. To test this, we apply our CNN to a collection of 52 Sentinel-1 interferograms on which, to allow for easier evaluation, we have performed the time consuming task of labelling both the class and location of deformation within them. However, in some examples assigning a single class to a complex deformation pattern is difficult, and we instead assign what we deem the dominant class to be. This problem is most evident in interferograms seven, nine and ten of Figure 4.7 that span the 2015 eruption of Wolf Volcano (Galapagos, Ecuador), in which signals were attributed to both the deflation of a sill and the opening of a dyke (Novellis et al., 2017; Xu et al., 2016).

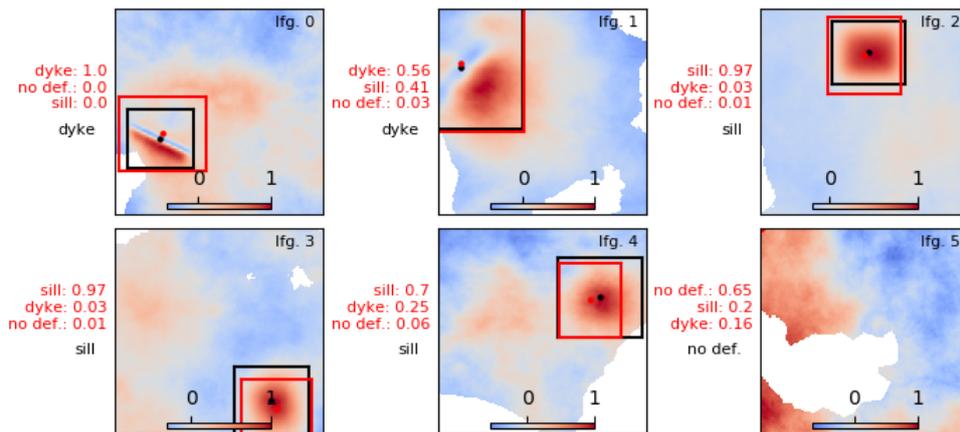
The interferograms used come from either a collection of time series that were either created by the authors of this study, or by the LiCSAR automatic interferogram processor, and feature the volcanoes Campi Flegrei, Agung, Wolf, Sierra Negra, and Alcedo, and contain interferograms that feature both inflating sills, opening dykes, and topographically correlated atmospheric signals. For the Galapagos volcanoes (Wolf, Sierra Negra, and Cerro Azul), deformation is visible in some of the 12 day inter-



**Figure 4.5:** Structure of our classification and localisation CNN. Input interferograms are first passed through the first five convolutional blocks of VGG16 to transform them from size  $(224 \times 224 \times 3)$  to size  $(7 \times 512)$ . These are flattened to create a large fully connected layer featuring 25088 neurons, which is connected to both the upper branch/head which performs classification, and the lower branch/head which performs localisation. We find the localisation problem more complex than classification, and consequentially our localisation branch/head features more layers each of more neurons. The output of the localisation head is a vector of four values determining the position and size of the deformation, whilst the output of the classification head is a vector of three values which indicate the probability for each class and sum to one.



**Figure 4.6:** Summary of training the two-headed model with synthetic data. The upper plot shows the accuracy of the classification head, whilst the lower plot shows the loss function for the localisation head. After the ninth epoch (marked by the vertical dashed line) the optimizer is switched from Nesterov Adam (NADAM) to stochastic gradient descent (SGD) with a manually chosen learning rate, and the weights in the fifth convolutional block of VGG16 are unfrozen. This extra learning stage allows the localisation loss for the validation data to decrease from  $\sim 400$  to  $\sim 300$ . This step can also be seen to initially damage the classification head, but this gradually recovers to an approximately equal accuracy ( $\sim 0.95$ ).



**Figure 4.7:** Results of our classification and localisation CNN on the testing data. Black class labels and location boxes were generated with the synthetic data, whilst red depicts those predicted by the CNN. As the model outputs a probability for each label, these are included as decimals with the predicted classes. Inspection of the results shows that in all of the randomly chosen cases, the classification is correct, and the localisation appears broadly correct.

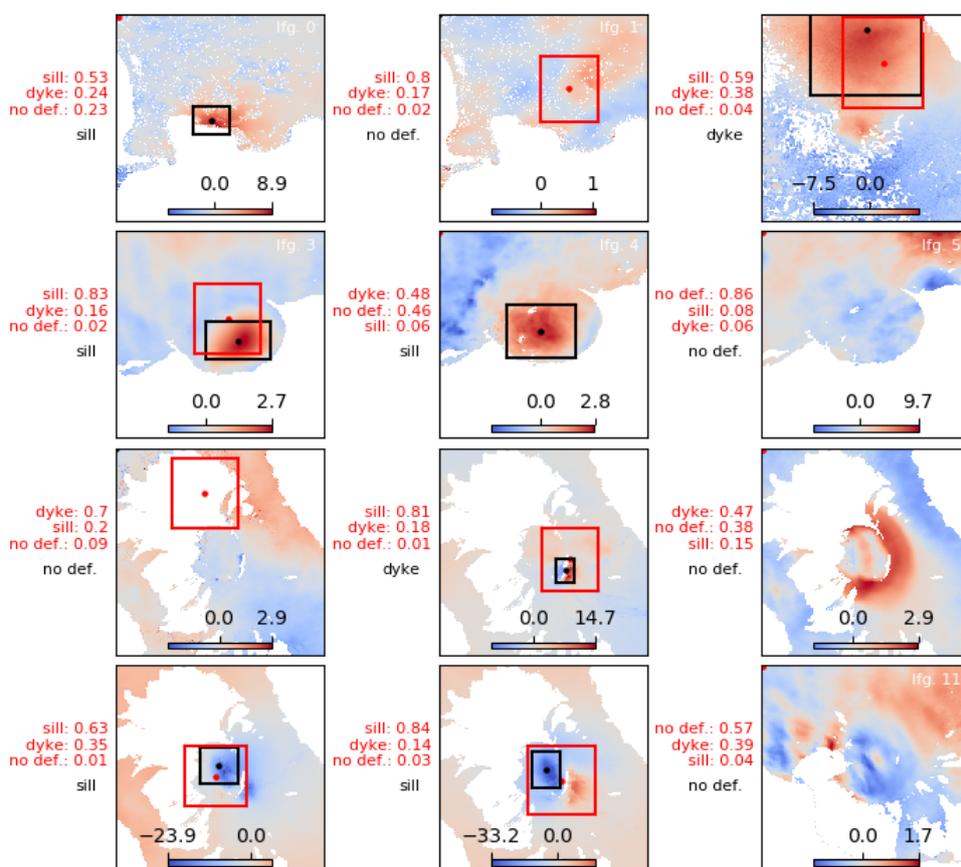
ferograms, and only filtering with a Goldstein filter (Goldstein and Werner, 1998), unwrapping using SNAPHU (Chen and Zebker, 2001), and masking of pixels with an average coherence below 0.7 is required. However, the deformation signal at Campi Flegrei is more subtle, and we are required to manually create interferograms with temporal baselines of 24/36/48/60 days in order for the deformation to be visible in a single interferogram. The deformation signal at Agung was attributed to the opening of a dyke (Albino et al., 2019), but due to the short lived nature of this event, is only visible in a relatively small number of the “daisy chain” of short temporal baseline interferograms. To increase the number of interferograms available, we again produce a selection of 24/36/48/60 day interferograms that span the event. In a manner similar to the Galapagos interferograms, we mask pixels with an average coherence below 0.7.

Figure 4.7 shows the results of applying our trained classification and localisation model to a quasi-random selection of Sentinel-1 interferograms. Interferograms such as 3 show a very clear inflation signal at Sierra Negra, and are correctly classified by the CNN, whilst the localisation is broadly correct. Other promising results include the labelling of the three Wolf coeruptive interferograms (seven, nine and ten) as containing a sill, which is also localised well. However, some interferograms are poorly classified, such as the signal in interferogram eight which shows what we interpret to be a strong topographically correlated APS. The divergent nature of our CNN’s two heads also leads to outputs that show disagreement between them. Interferogram zero demonstrates this, in which it is correctly classified as containing a sill, but features no localisation output.

When considering the entire test set of real data, the classification accuracy is 0.46, whilst the localisation loss is  $\sim 2500$ . However, Table 4.1 shows that the distribution of accuracy between classes is strongly heterogeneous, with a high accuracy for the classification of sills (0.88), but a low accuracy for dykes and no deformation cases (0.33 and 0.25, respectively). We discuss the results of this model more fully in Section 4.4, but in the following section we seek to improve the performance of our model through the inclusion of real data during the training stage.

### 4.3.3 Using Sentinel-1 data

To increase the performance of our model further, we seek to incorporate real data into the training. We do this through revisiting the time series mentioned in the previous section, and labelling a further 173 interferograms which we use for training, whilst retaining the original set for further testing. It should be noted that the majority of these feature only atmospheric signals, and so are significantly less time consuming to label than those that feature deformation and require four localisation coordinates. However, 20000 synthetic interferograms were used to train the previous model, and the inclusion of 173 new interferograms is unlikely to impact the model significantly



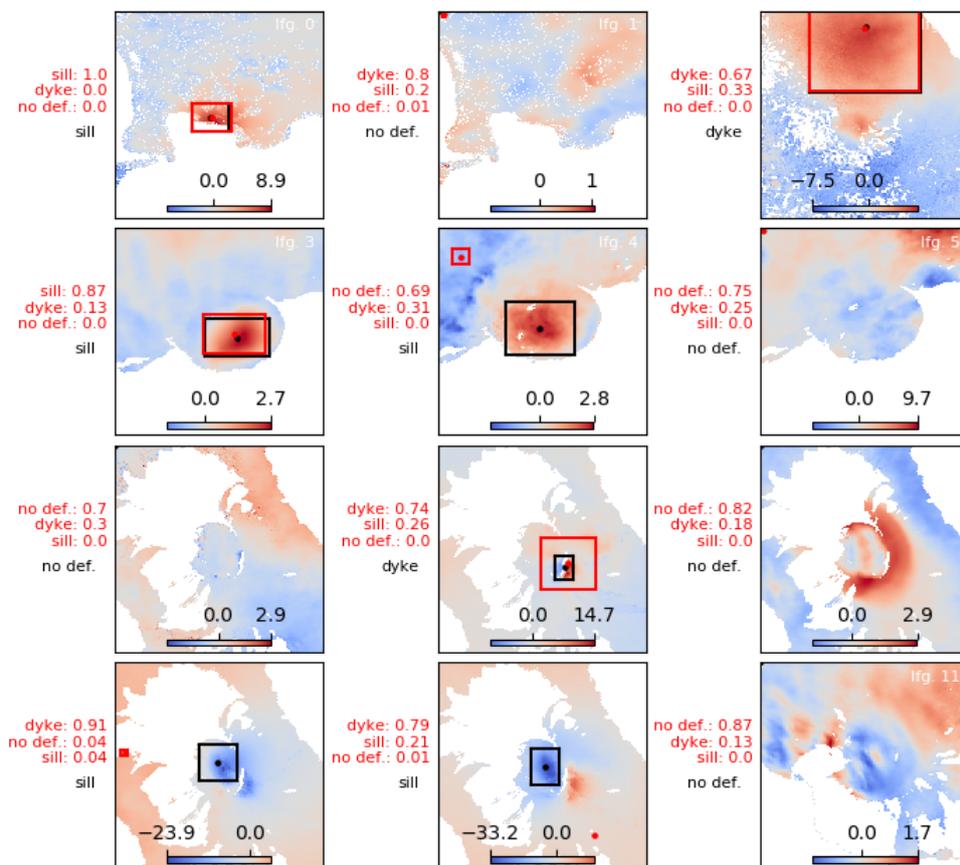
**Figure 4.8:** Results of our classification and localisation CNN on our testing set of Sentinel-1 interferograms. The labelling convention is as per the previous figure, but labels in black were manually created. Inspection of these results show that they vary between both the label and localisation being correct (e.g. 3, 10, 11), the localisation correct but the label incorrect (e.g. 2), the label correct but the localisation incorrect (e.g. 0), and both the label and localisation incorrect (e.g. 4). Interferograms 0 – 1 feature Campi Flegrei, 2 features Agung, 3 – 5 feature Sierra Negra, 6 – 10 feature Wolf, and 11 features Cerro Azul.

as these could still be classified poorly with minimal increase in the loss function. We therefore apply data augmentation, which involves creating random flips, rotations, and translations of the interferograms to extend our set of real training data to feature 20000 unique, though often highly correlated, Sentinel-1 interferograms.

Figure 4.9 shows the results of applying our CNN to the same set of test interferograms used in Section 4.3.2. Inspection shows greatly improved localisation, with very small errors for interferograms zero, two and three. False positives have also now been reduced, with the strong topographically correlated APS of interferogram eight now correctly classified as “no deformation”. However, some false positives remain, such as interferogram one, and the subtle signal seen in interferogram four is classified incorrectly. The complex deformation signals of interferograms seven, nine and ten also prove challenging as the deformation is not localised well by the model, but the CNN’s label of “dyke” is reasonable given that signals spanning the 2015 eruption of Wolf were attributed to both changes in the volume of a sill, and propagation of magma to the surface (Xu et al., 2016). Considering the entire real training dataset, performance has now increased, and the classification accuracy has risen to 0.62, whilst the localisation loss has decreased to 472. Table 4.1 shows that the inclusion of interferograms that image the opening of a dyke greatly increase the classification accuracy for this class (0.33 to 1.00), and that the relatively low overall classification accuracy remains a product of the model’s poor classification of interferograms that contain only atmospheric signals (0.62). Results of this type may be termed “false positives”, as they would cause a detection algorithm to erroneously flag interferograms as containing deformation, and methods to generate more complex synthetic atmospheres for training is discussed further in the following section.

## 4.4 Discussion

From the analysis performed in Section 4.2 we conclude that the incorporation of a DEM into our CNN cannot be achieved through the relatively easy step of using it as one channel in multichannel data. We believe this is because the weights in the first five convolutional blocks our model were transferred from VGG16 and, as this model was trained using natural images which are broadly similar across all three channels, they are not well suited for the case in which each channel is very different. However, if in future work the weights within the convolutional blocks of a classification and localisation model are trained from scratch, then these may easily allow for the incorporation of extra data in the different input channels. Should this approach not be feasible, information such as the DEM may be best incorporated through the use of a two input model, in which one set of convolutional filters are applied to the phase information, whilst a second is applied to the DEM. These two networks could then be merged at the fully connected stage, in much the same way as our fully connected



**Figure 4.9:** Results of our second classification and localisation CNN on our testing set of Sentinel-1 interferograms. Interferograms are shown with LOS displacement in cm, and with incoherent areas masked. The labelling convention and interferograms are as per Figure 4.8. This model can be seen to outperform the first, with improved classification and localisation. However, several errors remain; Interferogram four features a comparatively subtle uplift signal in comparison to others that preceded the 2018 eruption of Sierra Negra and is classified as “no deformation” by the model, whilst the complex co-eruptive signal of interferogram nine is not located accurately.

<b>Classification Accuracy</b>	Synthetic	Synthetic and Real
Dyke (3)	0.33	1.00
Sill (17)	0.88	0.82
No deformation (32)	0.25	0.47
Combined (52)	0.46	0.62

<b>Localisation Loss</b>	Synthetic	Synthetic and Real
Dyke (3)	533	68
Sill (17)	1589	514
No deformation (32)	3164	528
Combined (52)	2498	497

**Table 4.1:** Summary statistics for CNNs trained either with synthetic data, or with synthetic and real data. For both cases, the models can be seen to achieve good accuracy when classifying interferograms that contain deformation, but to misclassify interferograms that contain only atmospheric signals (accuracies of 0.46 and 0.62). The most significant reduction in localisation loss is also seen for interferograms that do not contain deformation (3164 to 528), suggesting that the inclusion of real data improves the model’s ability to correctly refrain from interpreting atmospheric signals as the location of deformation.

model diverges into two outputs. Should this be successful, it may also provide a method to add further inputs to a model, such as those outputted by a weather model, which may reduce false positives due to occurrences such as a strong topographically correlated APS. However, training the weights of a model from scratch and exploring more complex multi-input model architectures remains beyond the remit of this study.

The results presented in Figure 4.8 show that a model trained only with synthetic data is able to classify and locate deformation signals in Sentinel-1 data. However, it is only successful in cases with particularly clear deformation patterns, and is poor at classifying interferograms that contain only strong atmospheric signals. It is possible that both of these limitations may be overcome through the use of more realistic synthetic data, which our preliminary study may warrant further research into. The generation of more realistic deformation patterns may be achieved through steps such as more intelligent sampling of the parameters used in the forward models used to generate the deformation patterns, the use of different types of deformation models such as penny-shaped cracks (Fialko et al., 2001), and the superposition of multiple deformation patterns in a single interferogram such as was observed prior to the 2005 eruption of Sierra Negra (Jónsson, 2009). The generation of more realistic atmospheric signals could be achieved through increasing the complexity of synthetic data, such as through the use of phase-elevation ratios that are non-linear or spatially variable, or through using data from different sources. Interferograms that image regions with little deformation could be used to increase the complexity of the set of “no deformation” data, or combined with synthetic deformation patterns to produce more complex semi-synthetic data.

The results presented in Figure 4.9 show that a convolutional neural network model can be used to both classify different types of deformation, and to localise this within an interferogram. However, much scope for improvement remains, with several classification and localisation errors visible in this figure. The majority of the localisation errors are either in cases in which the deformation signal is slight (e.g. interferogram four of Figure 4.9), or in interferograms that span the 2015 eruption of Wolf volcano. In the former case, it is natural for a threshold in the signal to noise ratio to exist below which a method is not able to identify the signal of interest, and these interferograms appear to represent that. In the latter case, the interferograms in question contain complex deformation patterns due to both the opening of a dyke and the removal of magma from a sill below the caldera (Novellis et al., 2017), and the inclusion of either real or synthetic training data that contains multiple deformation patterns may alleviate this shortcoming.

The divergent nature of the two heads (classification and localisation) of our network also allows for discrepancies between their outputs. This is seen in interferogram 10 of Figure 4.9, in which a plausible label of “dyke” is attributed to the co-eruptive deformation signal, but the localisation head produces an erroneous location. However, we postulate that it may be possible to avoid errors of this type by using more complex model architectures. Models such as YOLO (Redmon et al., 2016) produce bounding boxes and classifications in one step, and have the added bonus of being able to work with images that contain multiple signals. If successfully applied to interferograms, a model of this complexity may avoid the discrepancy errors we encounter, and be able to handle interferograms that contain multiple deformation patterns.

Our approach to localisation avoids the need for repeated classification using a sliding window approach, and allows for our network to reason using the entire image. Whilst we believe this approach is beneficial, one caveat remains in that building a network that is able to utilise large interferograms can be complex. In our model, we use pixels of three arc second size and, with an input size of  $224 \times 224$ , the resulting model is able to “see” an approximately 20km square around a volcano. If we wish to proceed at this resolution, our model’s visual field could be increased through changing the input size to around  $400 \times 400$  which would not impact our ability to use VGG16’s filters (or convolutional blocks), but would increase the size of the first layer of the fully connected part of our network.

At present, an input with side length 224 is reduced to a feature map with side length 7 (shown in Figure 4.5) which, combined with a depth of 512, produces a flattened layer of size  $7 \times 7 \times 512 = 25088$ . However, doubling the input side length would double the feature map side length, increasing the flattened layer to a size of  $14 \times 14 \times 512 = 100352$ . Whilst our model contains millions of free parameters, connecting this layer to a subsequent layer would produce a significant increase in the total, and is likely to

require either more ingenuity or more data to be trained successfully. Analysis of the offsets of deformation patterns at volcanic centres by Ebmeier et al. (2018) finds that 8% of signals are located more than 10km from a volcanic edifice, and would therefore be missed by our current model. Future models that wish to perform localisation using a global approach may therefore require slight increases in size in order to capture all signals of interest, but the work required to train these larger models remains outside the remit of this study.

## 4.5 Conclusion

Our study finds that either wrapped or unwrapped data are approximately equally suited for use with the weights of VGG16’s filters that were trained on ImageNet data, whilst more complex use of the three channel format that these models support degrades performance. However, this may not be the case if the weights within VGG16’s filters are trained from scratch. We combine the five convolutional blocks of VGG16 with two fully connected networks to perform both classification and localisation, which allows our network to reason using the whole interferogram, and does not require a sliding window approach. Additionally, our network is able to differentiate between several different forms of deformation.

To minimise the costly nature of labelling data, we initially train our model using only synthetic data. We find that our model generalises well to some cases of Sentinel-1 data, but produces several false positive results when using interferograms that contain strong atmospheric signals. We alleviate this problem through the inclusion of a small amount of real data during the training phase, and present a model that is able to both classify and locate deformation within an interferogram of  $\sim 20$ km side length.

## 4.6 Acknowledgments

MEG is supported by the LiCS grant. Figures were prepared in Matplotlib (Hunter, 2007), and all CNN work was carried out in KERAS using the TensorFlow backend.

## References

- Agostinelli, Forest, Matthew Hoffman, Peter Sadowski, and Pierre Baldi (2014). “Learning activation functions to improve deep neural networks”. In: *arXiv preprint arXiv:1412.6830*.
- Albino, Fabien, Juliet Biggs, and Devy Kamil Syahbana (2019). “Dyke intrusion between neighbouring arc volcanoes responsible for 2017 pre-eruptive seismic swarm at Agung”. In: *Nature communications* 10.1, p. 748.

- Anantrasirichai, N, Juliet Biggs, F. Albino, P. Hill, and D. Bull (2018). “Application of Machine Learning to Classification of Volcanic Deformation in Routinely Generated InSAR Data”. In: *Journal of Geophysical Research : Solid Earth*, pp. 1–15. DOI: 10.1029/2018JB015911.
- Bekaert, D P S, A Hooper, and T J Wright (2015). “A spatially variable power law tropospheric correction technique for InSAR data”. In: *Journal of Geophysical Research: Solid Earth* 2, pp. 1–12. ISSN: 2169-9356. DOI: 10.1002/2014JB011557.A.
- Bridle, John S (1990). “Probabilistic interpretation of feedforward classification network outputs, with relationships to statistical pattern recognition”. In: *Neurocomputing*. Springer, pp. 227–236.
- Canziani, Alfredo, Adam Paszke, and Eugenio Culurciello (2016). “An analysis of deep neural network models for practical applications”. In: *arXiv preprint arXiv:1605.07678*.
- Chen, Curtis W and Howard A Zebker (2001). “Two-dimensional phase unwrapping with use of statistical models for cost functions in nonlinear optimization”. In: *Journal of the Optical Society of America* 18.2, pp. 338–351.
- Deng, Jia, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei (2009). “Imagenet: A large-scale hierarchical image database”. In: *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. Ieee, pp. 248–255.
- Dozat, Timothy (2016). “Incorporating nesterov momentum into adam”. In:
- Ebmeier, S K, B J Andrews, M C Araya, D W D Arnold, J Biggs, C Cooper, E Cottrell, M Furtney, J Hickey, J Jay, R Lloyd, A L Parker, M E Pritchard, E Robertson, E Venzke, and J L Williamson (2018). “Synthesis of global satellite observations of magmatic and volcanic deformation : implications for volcano monitoring & the lateral extent of magmatic domains”. In: *Journal of Applied Volcanology*, pp. 1–26. DOI: 10.1186/s13617-018-0071-3.
- Farr, T.G., P.a. Rosen, Edward Caro, and Robert Crippen (2007). “The Shuttle Radar Topography Mission”. In: *Reviews of Geophysics* 45.2005, pp. 1–33. DOI: 10.1029/2005RG000183.1.
- Fellbaum, C (1998). *WordNet: An Electronic Lexical Database: Bradford Book*.
- Fialko, Yuri, Yakov Khazan, and Mark Simons (2001). “Deformation due to a pressurized horizontal circular crack in an elastic half-space, with application to volcano geodesy”. In: *Geophys. J. Int.* 146, pp. 181–190.
- Gaddes, ME, A Hooper, M Bagnardi, H Inman, and F Albino (2018). “Blind signal separation methods for InSAR: The potential to automatically detect and monitor signals of volcanic deformation”. In: *Journal of Geophysical Research: Solid Earth*.
- Goldstein, Richard M. and Charles L. Werner (1998). “Radar interferogram filtering for geophysical applications”. In: *Geophysical Research Letters* 25.21, p. 4035. ISSN: 0094-8276. DOI: 10.1029/1998GL900033.

- González, Pablo J., Richard J. Walters, E. L. Hatton, Karsten Spaans, A. McDougall, Andrew Hooper, and Tim J. Wright (2016). “LiCSAR: Tools for automated generation of Sentinel-1 frame interferograms”. In: *AGU Fall meeting*.
- He, Kaiming, Georgia Gkioxari, Piotr Dollár, and Ross Girshick (2017). “Mask r-cnn”. In: *Proceedings of the IEEE international conference on computer vision*, pp. 2961–2969.
- Hunter, J. D. (2007). “Matplotlib: A 2D graphics environment”. In: *Computing In Science & Engineering* 9.3, pp. 90–95.
- Jónsson, Sigurjón (June 2009). “Stress interaction between magma accumulation and trapdoor faulting on Sierra Negra volcano, Galápagos”. In: *Tectonophysics* 471.1-2, pp. 36–44. ISSN: 00401951. DOI: 10.1016/j.tecto.2008.08.005. URL: <http://linkinghub.elsevier.com/retrieve/pii/S0040195108003880>.
- Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E Hinton (2012). “Imagenet classification with deep convolutional neural networks”. In: *Advances in neural information processing systems*, pp. 1097–1105.
- Novellis, V De, R Castaldo, C De Luca, S Pepe, I Zinno, F Casu, R Lanari, and G Solaro (2017). “Source modelling of the 2015 Wolf volcano ( Galápagos ) eruption inferred from Sentinel 1-A DInSAR deformation maps and pre-eruptive ENVISAT time series”. In: *Journal of Volcanology and Geothermal Research* 344, pp. 246–256. ISSN: 0377-0273. DOI: 10.1016/j.jvolgeores.2017.05.013.
- Okada (1985). “Surface deformation due to shear and tensile faults in a half-space”. In: *International Journal of Rock Mechanics and Mining Sciences Geomechanics Abstracts* 75.4, pp. 1135–1154. ISSN: 00371106. DOI: 10.1016/0148-9062(86)90674-1. URL: <http://linkinghub.elsevier.com/retrieve/pii/0148906286906741>.
- Rauter, Matthias and Daniel Winkler (2018). “Predicting Natural Hazards with Neuronal Networks”. In: *arXiv Electrical*, pp. 1–14. arXiv: [arXiv:1802.07257v1](https://arxiv.org/abs/1802.07257v1).
- Redmon, Joseph, Santosh Divvala, Ross Girshick, and Ali Farhadi (2016). “You only look once: Unified, real-time object detection”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 779–788.
- Simonyan, Karen and Andrew Zisserman (2014). “Very deep convolutional networks for large-scale image recognition”. In: *arXiv preprint arXiv:1409.1556*.
- Srivastava, Nitish, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov (2014). “Dropout: a simple way to prevent neural networks from overfitting”. In: *The Journal of Machine Learning Research* 15.1, pp. 1929–1958.
- Xu, Wenbin, Sigurjón Jónsson, Joël Ruch, and Yosuke Aoki (2016). “The 2015 Wolf volcano (Galápagos) eruption studied using Sentinel-1 and ALOS-2 data”. In: *Geophysical Research Letters*, pp. 9573–9580. DOI: 10.1002/2016GL069820..



## Chapter 5

# Discussion and Conclusions

In this thesis, my objective was to develop an algorithm to detect signs of deformation-generating volcanic unrest in a time series of interferograms. In Chapter 1 I divided this objective into four smaller aims, which I revisit in Section 5.1. In Section 5.2 I discuss the opportunities for further work on this topic, and in Section 5.3 I present my concluding remarks.

### 5.1 Project aims and key findings

#### 5.1.1 Characterising baseline signals

In Chapter 2, I explored the potential to characterise the baseline signals of a volcano imaged using InSAR by treating the individual interferograms as the combinations of several latent sources. This produced a set of spatial maps for each signal, and a set of vectors in which each value determined how each spatial map contributed to a given interferogram. If the interferograms are ordered temporally, the contribution vectors can be termed “time courses”, and viewed as a line graph. However, it should be noted that these line graphs determine the contribution per the temporal baseline of a given interferogram, and therefore it can be more intuitive to integrate a time course to produce cumulative time courses, such as are shown in Figure 2.16.

There exist a multitude of methods to decompose a time series of interferograms into spatial signals and (cumulative) time courses. The analysis that I performed in Chapter 2 found that independent component analysis set to recover spatially independent sources (sICA) using the FastICA algorithm (Hyvarinen, 1999) performed best in a selection of tests with synthetic data. However, this analysis detected two issues of potential concern. The first was concerned with the choice of the hyperparameter which determines the number of sources to be recovered. When considering  $n$  mixtures made from  $m$  sources, the mixtures will create  $m - \text{hypervolume}$  in the  $nD$  space, and

ICA requires that the dimensionality is first reduced to  $m$  using PCA. However, in the case with noise, the mixtures instead fill a  $n - \text{hypervolume}$ , and in reducing this to  $m$  dimensions using PCA, some aspects of the sources may be lost, depending on the signal to noise ratio. This problem is complicated as we do not generally know how many sources contributed to the mixtures, and in Chapter 2 I find that the best results are produced when the number of sources to be recovered (i.e.  $m$ ) is set to be  $\sim 2$  more than the estimated number of sources. This ensures that components in which signals of interest reside are not discarded, but also that the FastICA algorithm is not faced with too much noise. Whilst this approach provides satisfactory results, and is broadly similar to the approach utilised by Ebmeier (2016), scope remains for its improvement in further work. The new ICA algorithm presented in Isomura and Toyozumi (2016) does not require any preprocessing in the case in which the number of sources and mixtures is not equal, and may therefore greatly simplify the application of ICA to data in which the number of sources is not known.

The second issue of concern that my analysis detected was the application of sICA to situations in which a topographically correlated atmospheric phase screen (APS) is spatially similar to a deformation signal. This case is likely to occur at a stratovolcano undergoing broad, volcano wide, inflation or deflation, as this signal is likely to lie under the topographic expression of the volcano, which will also create a topographic APS in the same location. This issue is introduced in Section 2.7, in which I conclude that further work may be required to use sICA for in-depth studies at a stratovolcano, but for volcano monitoring in which I seek to only characterise baseline behaviour, slight inaccuracies in the sources recovered are not likely to be an issue.

After decomposing the baseline signals into a set of spatial patterns and associated time courses, a method was required to quantify baseline behaviour in terms of these. This is addressed in Sections 2.6 and 3.2, which detail how the algorithm I developed quantifies the use of spatial patterns through fitting lines of best fit through their cumulative time courses, and through the cumulative residual between each interferogram and its reconstruction made using the recovered sources. In this work, I fitted polynomials of order 1 (i.e.  $y = ax + b$ ), which provided adequate performance for all of the cases encountered. However, as the length of time series available from satellites such as Sentinel-1 increases, it is likely that the time courses of signals that capture atmospheric processes may be revealed to be non-linear. An example of such a case is provided by Pinel et al. (2011), who find that the strength of a topographically correlated APS at Colima volcano is best modelled as a sinusoid, due to seasonal changes in the atmosphere. Other processes, such as the cooling and contraction of a magma body below a volcano, are likely to produce a signal that is detected by my detection algorithm, and also feature a time course that, on long enough time scales, cannot be treated as linear. Therefore, further testing of my detection algorithm is likely to

indicate that it may also need to solve for the most appropriate type of line of best fit to characterise the temporal behaviour of the baseline data.

When utilising sICA, a common problem encountered is how to assess the statistical significances of the sources recovered (Hyvärinen, 2012). In Chapter 3 I addressed this through creating the ICASAR algorithm, which showed promising results when it was applied to both synthetic data, and a time series of Sentinel-1 data that imaged the pre-eruptive inflation of Sierra Negra (Section 3.3). Through using two independent methods in parallel, the ICASAR algorithm creates powerful visualisations of the similarities in the sources recovered by multiple runs of the FastICA algorithm, and so allows a user to easily build an understanding of which of the recovered sources are the most significant. Additionally, for use in an automatic detection algorithm, the ICASAR algorithm is able to automatically determine the number of sources that have been recovered, and to rank these recovered sources based on their robustness.

### 5.1.2 Detecting and displaying deviations from baseline behaviour

The prototype detection algorithm that I presented in Chapter 2 was able to detect the onset of the 2015 eruption of Wolf volcano by extrapolating the line of best fit for the cumulative residual, and determining how significant the discrepancy between a new value and the extrapolated line was (shown in Figure 2.17). In Chapter 3, I presented a more complete detection algorithm that applied the same approach to the cumulative time courses of each recovered source, and showed that this algorithm was able to detect the increase in uplift seen prior to the 2018 eruption of Sierra Negra (shown in Figure 3.7). As mentioned in the previous section, more complex lines of best fit (e.g. sinusoids) may be required to accurately describe the shapes of cumulative time courses/residuals, but this should not impact my algorithm's ability to meet this aim, providing that the function can be extrapolated to future dates. Figure 3.7 shows the output displayed when the monitoring algorithm is applied to a time series of Sentinel-1 interferograms that image Sierra Negra before and during the 2018 eruption. This figure provides a tool for a user who is familiar with deformation processes at a volcano to easily interpret this signals contained within the time series, and determine that a change in the style of deformation observed at the volcano had occurred around the middle of 2017 (interferogram  $\sim 67$  in Figure 3.7).

However, one limitation of this method is that the lines of best fit that are learned in the baseline stage must be redrawn periodically, as otherwise any small misestimation in the gradient of the cumulative time course/residual gradually leads to the flagging of all points as a large distance from the line, and so the algorithm producing false positive results. However, this approach introduces a hyperparameter to control how often the lines of best fit are redrawn. If the value is too low (e.g. the lines are redrawn every five interferograms), large changes in the time courses may not be detected producing false

negatives, whilst in the case that it is too high (e.g. the lines are redrawn every 100 interferograms), the false positives case previously mentioned occurs. I envisage that further testing of the algorithm at more volcanoes will provide insights into whether this parameter needs to be changed, and speculate that a separate module may be required to adaptively change this value for each volcano.

### 5.1.3 Complementary monitoring algorithm

I have refined my detection algorithm by producing a complementary model that works in parallel to it, which I detail in Chapter 4. In contrast to my detection algorithm that uses a time series of interferograms, this model is similar to that proposed by Anantrasirichai et al. (2018), and instead classifies new interferograms as they are produced by an automatic interferogram processing tool. My automatic detection model is a convolutional neural network that is able to classify interferograms as either not containing deformation, containing deformation due an inflating sill, or containing deformation due to an opening dyke. Additionally, the model is able to reason globally about an interferogram of size  $224 \times 224$  pixels, where each pixel is around 3 arcseconds in size, to locate a deformation signal within the image.

The outputs of this model are also optimised for use in hazard monitoring. The classification output uses a softmax function, which ensures that instead of producing one class label, the model outputs the probabilities for each type of class label. Therefore, in cases in which the model is relatively uncertain about a new interferogram, this information is also encoded into the classification output, and can be used to determine which interferograms need to be manually inspected as a priority. It is likely that if only one interferogram in a time series is labelled as containing deformation, and that label is associated with a high uncertainty, then the time series is of low priority for manual inspection, or a warning to be issued. However, if multiple interferograms contain the same label with low uncertainty, then the time series is of high priority for manual inspection, or a warning to be issued. Additionally, the model's ability to determine the spatial extent of a new signal works in a similar fashion, with signals with a large spatial extent likely to indicate a process that may pose a significant hazard.

However, whilst the model advances the state of the art through demonstrating that it is able to differentiate between different types of deformation, the classification output remains limited. It is likely that roughly symmetric processes such as the inflation of a magma chamber at depth would be classified as "sills" by the model, but other more complex processes would not be. In Chapter 2, I used sICA to isolate eastward movement of the eastern flank of Mt Etna, and signals such as these are likely to prove problematic for the model. In Section 5.2 I discuss how this model could be improved, including through providing more classification outputs.

#### 5.1.4 Study of a volcanic process

During development of my automatic detection algorithm, I tested parts of the algorithm on Wolf and Sierra Negra volcanoes. When sICA was applied to a time series of Sentinel-1 interferograms that spanned the 2015 eruption of Wolf, three signals of geophysical interest were automatically isolated. These are shown in Figure 2.15, and include a subsidence signal located in the caldera, an eastward movement signal located on the caldera floor, and a broader subsidence signal located under the edifice. Modelling by other studies (Novellis et al., 2017; Xu et al., 2016) attributed these signals to a deflating shallow sill ( $\sim 2$  km below the caldera floor), an opening dyke on the caldera rim, and a deflating sill at greater depth ( $\sim 6$  km below the caldera floor), respectively. However, given the multitude of co-eruptive studies, and that the time series of seven pre-eruptive interferograms was relatively short, I did not feel that this presented the best opportunity to study a volcanic process.

However, in contrast to the 2015 eruption of Wolf, the 2018 eruption of Sierra Negra provided a long time series of  $\sim 100$  pre-eruptive interferograms. This allowed for the study into the pre-eruptive overpressure that I describe in Chapter 3, in which I find a value for the inter-eruptive pressure change of  $\sim 13.5$  MPa. This value is significantly larger than the values thought to be required to cause the overlying elastic crust to fail in tension and mafic magma propagate to the surface, which are typically  $\sim 1$  MPa (Manga and Brodsky, 2006; Jellinek and DePaolo, 2003), even though I use a relatively low shear modulus of 4.5 GPa (Hooper et al., 2002) during the inversion to reduce the overpressure value I calculate. Previous studies of Sierra Negra have shown that prior to the 2005 eruption, movement on the faults at the edges of the caldera floor reduced the overpressure within the sill below by 3 MPa (Jónsson, 2009), which may provide a process to reduce the overpressure within the sill, and so postpone the eventual eruption. Evidence of this process occurring prior to the 2018 eruption is provided by the “RMS of cumulative residual” of Figure 3.7, in which the later interferograms ( $\sim 65$  onwards) cannot be fit as well by the six components learned during the baseline stage. This suggests that the signals present in the later interferograms differentiate from those earlier in the time series that can be fit well with the learned components, and may be due to a new signal entering the time series due to movement on the caldera floor faults.

The potential change in deformation pattern that is indicated by the decrease in my algorithm’s ability to fit the new interferograms using the components learned during the baseline stage also coincides with an increase in the rate of uplift that is seen as IC1’s time course in Figure 3.7, and in Figure 3.9. With the exception of a short period of time after the 2005 eruption, the rate of inflation remains approximately linear for 10 years (days  $\sim 500$  to  $\sim 4200$  in Figure 3.9), but within approximately one year of the change in rate, the volcano erupts. In light of this, I postulate that when the sill is

inflating at the lower rate, a process is occurring which limits the growth in overpressure within the sill. When the rate of inflation increases, this process cannot keep pace, and pressure within the sill grows, leading to events such as faulting within the caldera, and eventual eruption. However, further analysis of the processes occurring prior to the eruption of Sierra Negra remain beyond the scope of Chapter 3.

### 5.1.5 Further Remarks

In the examples considered within this thesis, both the detection algorithm that uses a time series approach (Chapter 3) and the detection algorithm that uses single interferograms (Chapter 4) are applied to similar data with similar results. However, tests with synthetic data show that the time series based detection algorithm may be able to detect changes due to signals that are not visible in a single interferogram, and may therefore be more applicable to the vast majority of deformation causing unrest events, in which significantly lower rates of deformation would be expected than those observed prior to the 2018 eruption of Sierra Negra. For a deep learning approach to also be applicable to such events, it is likely that further work will be required to advance from CNNs that use single interferograms, to those that can use a time series of data.

For each of the detection algorithms presented, further work is required for them to become fully automated. For the time series based approach, I envisage that a separate algorithm would be required to fulfil the role of a human interpreter analysing that data displayed by the algorithm (e.g. Figure 3.7). In the simplest case, this could merely collate the outputs from each time course and the cumulative residual at each time step, and determine if any of these show significant deviations. However, a more complex algorithm may allow for features such as a warning threshold to be set depending on the temporal stability of any deviations (i.e. if they persist through successive interferograms), or to determine if short-lived deviations are likely to be due to atmospheric signals. For the CNN approach, a relatively simple output that displays the probability of an interferogram containing deformation, and the spatial extent of any deformation it contains may allow for it be used for automatic detection. However, a more complex algorithm that is able to reason based on information such as if the location of the deformation has shifted, or if interferograms are consistently being labelled as containing deformation, may prove to be more useful for a non-expert user.

For application of the time series based detection algorithm to other volcanoes, tuning of the frequency with which lines of best fit are redrawn may be required to detect lower magnitude signals successfully. Additionally, as the application of sICA at stratovolcanoes was found to be challenging due to the lack of spatial independence between deformation and atmospheric sources (Chapter 2), the use of weather models to first remove a proportion of the topographically correlated APS may vastly improve

the functioning of the detection algorithm. The size of the interferogram that the algorithm is also applied to may also require tuning, as some deformation patterns may be larger than the sizes currently considered.

### 5.1.6 Key Findings

My work to construct an automatic detection algorithm resulted in several key findings:

1. sICA outperforms NMF, s/tPCA, and tICA to isolate signals of geophysical interest in a time series of interferograms.
2. The relative significance of the sources recovered when sICA is applied to InSAR can be automatically calculated using the ICASAR algorithm.
3. Convolutional neural networks can be trained to differentiate between different types of common volcanic deformation, and used to determine the spatial size and location of a deformation signal within an interferogram.
4. The overpressure within the sill that inflated before the 2018 eruption of Sierra Negra cannot be thought of in simple terms of a gradual increase in pressure until the tensile strength of the overlying crust is overcome.

## 5.2 Future work

During the research I have performed to construct an automatic detection algorithm, three distinct avenues for further work have become apparent. The first deals with refinements to my automatic detection algorithm, the second deals with further study of the 2018 eruption of Sierra Negra, and the third deals with the construction of more complex convolutional neural networks.

Refinements to my automatic detection algorithm are likely to be required to address issues that will be revealed during testing on a wider range of volcanoes. At present, the algorithm has only been tested on two Galapagos volcanoes, and these both provide large areas that remain coherent on long time scales. However, the ICA algorithm is central to my automatic detection algorithm, and has been tested at Etna (Chapter 2), and Calbuco and Parícutin Lava Fields (Ebmeier, 2016), which suggests that the detection algorithm is likely to generalise well to other volcanoes. However, tropical volcanoes of the type discussed in Ebmeier et al. (2013), which feature strong atmospheric signals and densely vegetated flanks, may prove challenging for sICA. Time series methods such as StaMPS (Hooper et al., 2012) may be required, or the use of L-Band data from planned SAR missions, such as NISAR (Rosen et al., 2015). Time series methods such as StaMPS require for a consistent set of pixels to be considered throughout the analysis, and therefore information may be discarded if pixels are coherent in only a few interferograms. However, a consistent set of pixels throughout

the time series is required when using FastICA/ICASAR, and therefore it is likely that removing noisy pixels using a method such as StaMPS may increase the performance of the ICASAR algorithm. Further work to combine these approaches in order to study signals with either smaller magnitudes than the 2015 eruption of Wolf such as the 2017 dyke intrusion at Agung, or with lower rates of pre-eruptive inflation than Sierra Negra, may illuminate previously overlooked signals.

The extension of the algorithm to cover volcanoes that feature periodic snow cover, such as those described in Spaans et al. (2015), may also prove challenging. The ICA algorithm requires a consistent choice of pixels throughout the time series under consideration, which results in the need to include pixels that are incoherent in snow covered images, if a user wishes to use the pixels in the remainder of the time series. The large areas of incoherent signal that significant snow fall would produce in a newly formed interferogram may therefore cause the automatic detection algorithm to try to fit it using unusual contributions of the baseline sources, which would then be highlighted as unrest. It remains to be seen, however, if through using a baseline of several years, the signals associated with snow falls may be incorporated into the baseline behaviour of a given volcano, allowing the algorithm to avoid false positive results.

Through separating signals of geophysical interest from other nuisance signals, such as those produced by changes in the atmosphere between SAR acquisitions, ICA can be considered as one of many methods that can reduce the impact of atmospheric phase delays on measurements made using InSAR. Other methods, such as the incorporation of data from weather models, may allow for some atmospheric signals to be removed prior to the application of ICA, and so allow it to recover signals of geophysical interest more accurately. However, a more complete analysis of how to use ICA with other tools designed to reduce the impact of atmospheric signals lies beyond the scope of this work.

As discussed in Section 5.1.1, tuning the frequency at which the lines fit to the cumulative time courses/residual are redrawn, and the use of more complex functions for the lines may allow for more accurate characterisation of baseline behaviour, and so improve the accuracy of the automatic detection algorithm. New ICA algorithms, such as the EGHR algorithm proposed in Isomura and Toyozumi (2016), may improve the components recovered by ICA, and so the performance of the automatic detection algorithm.

The convolutional neural network described in Chapter 4 provides many avenues for future work due to the large number of deep learning models produced each year by the computer vision community. Whilst the model I describe advances the state-of-the-art, errors are still produced when it is faced with some interferograms. Additionally, the design of a model that can only assign a single class label to each image is applicable

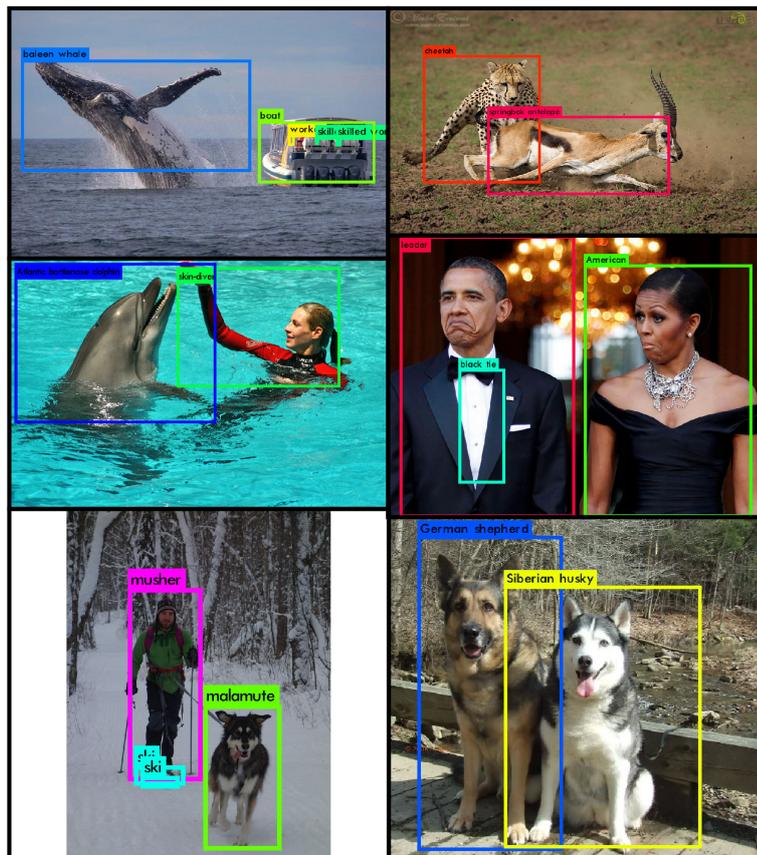
to many interferograms, but it is not uncommon for an interferogram to include deformation due to multiple processes. This is seen in Figure 4.9, in which the co-eruptive interferogram of Wolf volcano features areas deforming due to the deflation of a sill, and areas deforming due to the opening of a dyke.

To address these issues, future work could involve the use of significantly more complex CNNs, such as YOLO9000 (Redmon and Farhadi, 2017), or R-CNN (Girshick et al., 2013). These CNNs outperform older models such as VGG16 (used in Chapter 4), and AlexNet (used by Anantrasirichai et al. (2018)), by first proposing regions that objects lie in, before then classifying the object found in each region. Figure 5.1 shows an example of the results produced by YOLO9000, in which it can be seen to accurately determine the location of a diverse set of objects, and to classify them using a prodigious set of classes. Should a model of this performance be implemented with InSAR data, it may be able to detect regions containing deformation, before classifying these in a manner similar to that described in Chapter 4. Working in parallel to a time series approach, models of this complexity could provide a powerful tool for monitoring volcanoes.

The performance of deep learning models that use data such as the frames of a video, suggests that their application to a time series of interferograms may provide a novel way to monitor volcanoes. Models such as PredNet (Lotter et al., 2016) are able to predict subsequent frames of a video, and may also provide tools to forecast the likely evolution of a period of unrest. In addition to using more complex neural networks, it may also be possible design models that use domain specific inputs in a manner similar to a human interpreter of InSAR data. Inputs such as DEMs would allow a network to consider if a signal is likely to be due to deformation or a topographically correlated APS, whilst inputs such as both wrapped and unwrapped data may allow a network to identify unwrapping errors. The ability to train CNNs using only synthetic data (demonstrated in Chapter 4) may also provide an avenue for further exploration, and through generating training dataset of a similar size to ImageNet, it may become possible to train CNNs from scratch, and so develop filters specifically designed for identifying deformation signals.

### 5.3 Concluding remarks

In this thesis, I have presented two methodologies to monitor the world's subaerial volcanoes using InSAR. The first uses several different machine learning methods to detect changes in a time series of interferograms, of which sICA is the most fundamental. The nascent nature of the application of ICA to InSAR data coupled with the promising results generated, suggests that further application and refinement is likely, which may provide for ways to increase the performance of the automatic detection algorithm.



**Figure 5.1:** Results from YOLO9000, reproduced from Redmon and Farhadi (2017). The figure shows the results of the CNN when applied to ImageNet data, and its ability to both locate and classify multiple objects in a single image.

The second methodology I have presented uses deep learning to produce a model with a basic understanding of the signals that are present in an interferogram, and so could be used to detect signs of unrest.

The possibility to undertake this endeavour is a result of the new era of SAR, in which open data policies, vast acquisition strategies, short revisit times, and the automatic creation of interferograms are but several of the many recent advances that is allowing SAR to advance into a tool used for hazard monitoring. The recent advances in deep learning, such as the development of high level APIs such as Keras, and the availability of models with pre-trained weights to facilitate transfer learning, have also provided a vast array of tools that can be used to automate the interpretation of interferograms created during routine monitoring. At present, the future of applying powerful deep learning methods to the vast volumes of new SAR data produced looks full of promise.

## References

- Anantrasirichai, N, Juliet Biggs, F. Albino, P. Hill, and D. Bull (2018). “Application of Machine Learning to Classification of Volcanic Deformation in Routinely Generated InSAR Data”. In: *Journal of Geophysical Research : Solid Earth*, pp. 1–15. DOI: 10.1029/2018JB015911.
- Ebmeier, S. K. (2016). “Application of Independent Component Analysis to multi-temporal InSAR data with volcanic case studies”. In: *Journal of Geophysical Research: Solid Earth*, pp. 8970–8986. ISSN: 21699313. DOI: 10.1002/2016JB013765. URL: <http://doi.wiley.com/10.1002/2016JB013765>.
- Ebmeier, S.K., J. Biggs, T.a. Mather, and F. Amelung (2013). “Applicability of InSAR to tropical volcanoes: Insights from Central America”. In: *Geological Society, London, Special Publications* 380.1, pp. 15–37. ISSN: 0305-8719. DOI: 10.1144/SP380.2.
- Girshick, Ross B., Jeff Donahue, Trevor Darrell, and Jitendra Malik (2013). “Rich feature hierarchies for accurate object detection and semantic segmentation”. In: *CoRR* abs/1311.2524. arXiv: 1311.2524. URL: <http://arxiv.org/abs/1311.2524>.
- Hooper, Andrew, David Bekaert, Karsten Spaans, and Mahmut Arıkan (Jan. 2012). “Recent advances in SAR interferometry time series analysis for measuring crustal deformation”. In: *Tectonophysics* 514-517, pp. 1–13. ISSN: 00401951. DOI: 10.1016/j.tecto.2011.10.013.
- Hooper, Andrew, Paul Segall, Kaj Johnson, and Justin Rubinstein (2002). “Reconciling seismic and geodetic models of the 1989 Kilauea south flank earthquake”. In: *Geophysical research letters* 29.22.
- Hyvarinen, Aapo (1999). “Fast and robust fixed-point algorithm for independent component analysis”. In: *IEEE Transactions on Neural Networks and Learning Systems* 10, pp. 626–634. ISSN: 1045-9227. DOI: 10.1109/72.761722.

- Hyvärinen, Aapo (2012). “Independent component analysis : recent advances”. In: *Philosophical Transactions of The Royal Society A*.
- Isomura, Takuya and Taro Toyozumi (2016). “A Local Learning Rule for Independent Component Analysis”. In: *Nature Publishing Group*, pp. 1–17. ISSN: 20452322. DOI: 10.1038/srep28073.
- Jellinek, A Mark and Donald J DePaolo (2003). “A model for the origin of large silicic magma chambers: precursors of caldera-forming eruptions”. In: *Bulletin of Volcanology* 65.5, pp. 363–381.
- Jónsson, Sigurjón (June 2009). “Stress interaction between magma accumulation and trapdoor faulting on Sierra Negra volcano, Galápagos”. In: *Tectonophysics* 471.1-2, pp. 36–44. ISSN: 00401951. DOI: 10.1016/j.tecto.2008.08.005. URL: <http://linkinghub.elsevier.com/retrieve/pii/S0040195108003880>.
- Lotter, William, Gabriel Kreiman, and David Cox (2016). “Deep predictive coding networks for video prediction and unsupervised learning”. In: *arXiv preprint arXiv:1605.08104*.
- Manga, Michael and Emily Brodsky (2006). “Seismic Triggering of Eruptions in the Far Field : Volcanoes and Geysers”. In: *Annual Review of Earth and Planetary Science*. DOI: 10.1146/annurev.earth.34.031405.125125.
- Novellis, V De, R Castaldo, C De Luca, S Pepe, I Zinno, F Casu, R Lanari, and G Solaro (2017). “Source modelling of the 2015 Wolf volcano ( Galápagos ) eruption inferred from Sentinel 1-A DInSAR deformation maps and pre-eruptive ENVISAT time series”. In: *Journal of Volcanology and Geothermal Research* 344, pp. 246–256. ISSN: 0377-0273. DOI: 10.1016/j.jvolgeores.2017.05.013.
- Pinel, V., A. Hooper, S. De la Cruz-Reyna, G. Reyes-Davila, M. P. Doin, and P. Bascou (2011). “The challenging retrieval of the displacement field from InSAR data for andesitic stratovolcanoes: Case study of Popocatepetl and Colima Volcano, Mexico”. In: *Journal of Volcanology and Geothermal Research* 200.1-2, pp. 49–61. ISSN: 03770273. DOI: 10.1016/j.jvolgeores.2010.12.002.
- Redmon, Joseph and Ali Farhadi (2017). “YOLO9000: better, faster, stronger”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7263–7271.
- Rosen, Paul A, Scott Hensley, Scott Shaffer, Louise Veilleux, Manab Chakraborty, Tapan Misra, Rakesh Bhan, V Raju Sagi, and R Satish (2015). “The NASA-ISRO SAR mission-An international space partnership for science and societal benefit”. In: *2015 IEEE Radar Conference (RadarCon)*. IEEE, pp. 1610–1613.
- Spaans, Karsten, Sigrún Hreinsdóttir, Andrew Hooper, and Benedikt Gunnar Ófeigsson (2015). “Crustal movements due to Iceland’s shrinking ice caps mimic magma inflow signal at Katla volcano”. In: *Scientific Reports* 5.April, p. 10285. ISSN: 2045-2322. DOI: 10.1038/srep10285. URL: <http://www.nature.com/doifinder/10.1038/srep10285>.

- 
- Xu, Wenbin, Sigurjón Jónsson, Joël Ruch, and Yosuke Aoki (2016). “The 2015 Wolf volcano (Galápagos) eruption studied using Sentinel-1 and ALOS-2 data”. In: *Geophysical Research Letters*, pp. 9573–9580. DOI: 10.1002/2016GL069820..



# Appendix A: Supporting information for Chapter 2

## Contents of this file

1. Text S1 to S3
2. Figures S1 to S2

**Introduction** When synthesising temporal sources for testing with various blind signal separation methods (e.g. temporal ICA), an approximation of the temporal nature of a volcanic source during a period of unrest was required. The small study performed on GPS data to ascertain this is detailed in Text S1 and Figure S1. When performing ICA, the input data need to first be whitened. Whilst this is a routine step, we provide details of this in Text S2. ICA seeks statistically independent sources, whilst PCA seeks uncorrelated sources, and the differences between these two terms are explained in Text S3. Figure S2 contains the interferograms covering Mt Etna.

## Text S1

The temporal nature of the synthetic deformation source was chosen such that a sample of the change per 6 days (i.e. the deformation per 6 days) had a kurtosis of 2. The choice to deviate from a Gaussian distribution was the result of a small study conducted using GPS data.

Six Volcanoes (Soufrière Hills (Montserrat), Mount Ruapehu (New Zealand), Long Valley Caldera (California), Mt. St. Helens (S.W. Washington), Piton de la Fournaise (Reunion Island), and Santorini (Greece)) were used for the analysis, in which the vertical displacements at several GPS stations for each volcano were calculated for six day steps (to mimic the acquisition rate of the Sentinel-1 constellation). These time series were divided into either times of activity or quiescence using independent data, and the kurtosis for each type of period calculated. The average kurtosis for all stations at all volcanoes was then calculated, yielding a value of 2.46 for times of activity, and

1.12 for times of quiescence. Figure 5.4 shows part of this process for Piton de la Fournaise.

GPS data for Santorini was provided by Xanthos Papanikolaou (Dionysos Satellite Observatory of National Technical University of Athens), for Piton de la Fournaise by Dr Aline Peltier (Piton de la Fournaise volcano observatory), for Mount Ruapehu by the New Zealand GeoNet project, and for the remaining North American volcanoes by UNAVCO and the USGS.

## Text S2

Whitening the input data is required as an input for ICA and for completeness is summarised here. Consider a linear transformation of the mixtures,  $\mathbf{X}$ , such that they will then be whitened:

$$\mathbf{Z} = \mathbf{V}\mathbf{X} \quad (5.1)$$

Where  $\mathbf{Z}$  are the whitened mixtures, and  $\mathbf{V}$  is the whitening matrix. For  $\mathbf{Z}$  to be whitened, the covariance matrix,  $\mathbf{C}_Z$ , is equal to:

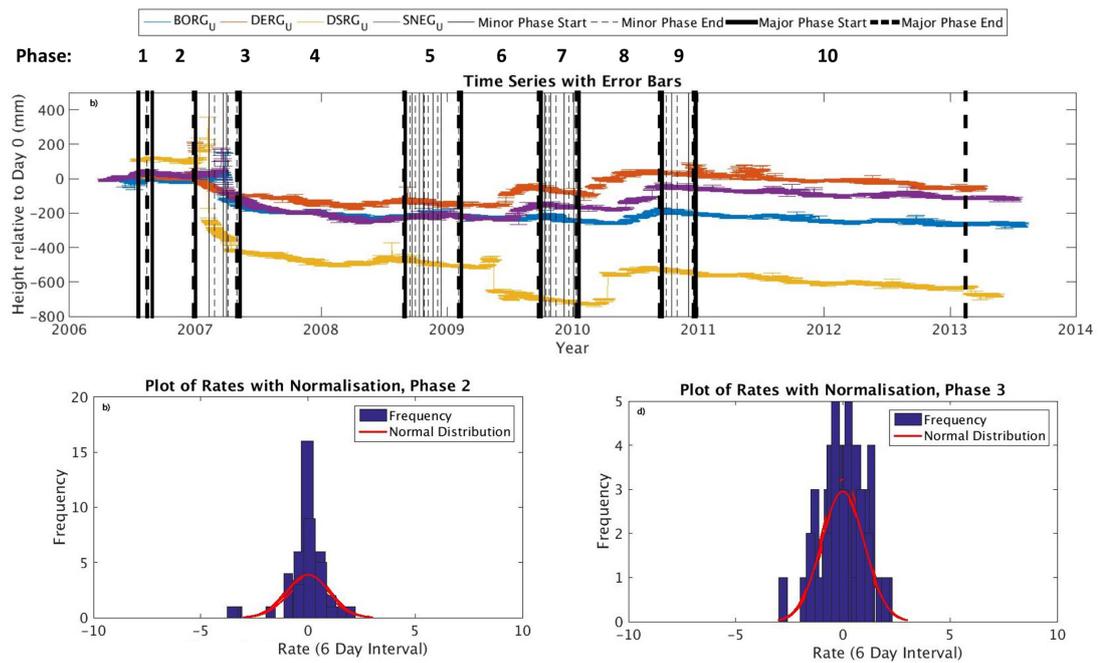
$$\mathbf{C}_Z = \mathbf{Z}\mathbf{Z}^T = \mathbf{I} \quad (5.2)$$

If we set  $\mathbf{E}$  to be a matrix with the normalised principal component axes as columns (i.e. the eigenvectors of  $\mathbf{X}\mathbf{X}^T$  for a  $t$  dimensional space), the least important column vectors can then be discarded such that only  $s$  remain, where  $s$  is the number of sources to be recovered. Scaling of the data to ensure unit variance can also be achieved through use of the corresponding eigenvalues  $\mathbf{D}$ , where  $\mathbf{D} = \text{diag}(d_1, d_2, \dots, d_s)$ :

$$\mathbf{V} = \mathbf{D}^{(-\frac{1}{2})}\mathbf{E}^T \quad (5.3)$$

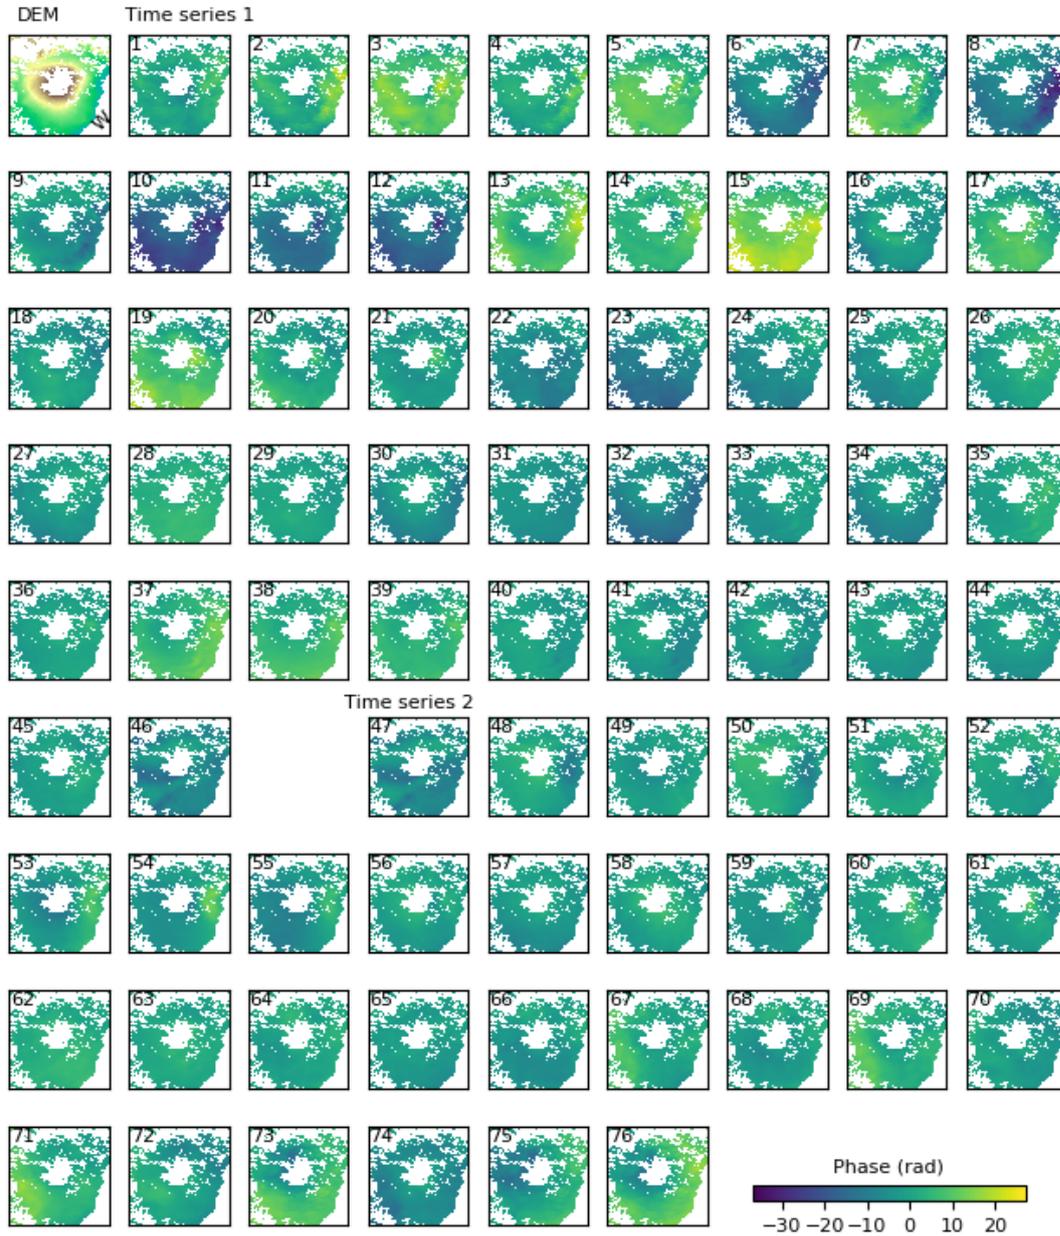
## Text S3

Implementing ICA to recover statistically independent sources is very closely related to applying PCA to recover uncorrelated sources. Informally, two signals are statistically independent if knowledge of the value of one signal does not convey any information of the value of the other, which contrasts with correlation that merely measures the linear relationship between variables. A similar measure applied to signals is that of correlation, and two signals are said to be uncorrelated if their covariance is 0. The difference between uncorrelatedness and the more important statistical independence can be demonstrated with two random variables that, when plotted in 2d space, form a



**Figure 5.2:** Top: Time series of the vertical component of displacement at four GPS stations at Piton de la Fournaise, with major periods of activity numbered and delimited by dark black lines. Bottom: histograms for the six day displacements during active phases two and three, with Gaussian distributions overlaid for comparison. Data from all of the three GPS stations were used to construct the histograms, and normalisation involved the mean centering of the data and scaling to unit variance in order to allow comparison between the six different volcanoes studied.

circle around the origin. Whilst the correlation of these values is 0, it is clear that they are not statistically independent as knowledge of one signal conveys information about the second signal (e.g. signal 1 attaining its maximum value conveys the information that signal 2 must also be at its median value). Therefore, two uncorrelated variables can still provide information about each other. The wealth of successful applications of ICA to BSS problems can be used to justify the expectation that it would outperform PCA as two physical processes that are unrelated (such as deformation at a volcano and atmospheric delay) are statistically independent, as opposed to merely uncorrelated.



**Figure 5.3:** The two time series of interferograms covering Mount Etna that sICA was applied to. The area to the east of the volcano which is marked "W" in the DEM (upper left panel) is masked as if does not include a usable radar return, whilst the remaining areas and pixels are masked due to low coherence.

# Appendix B: Supporting information for Chapter 3

## Contents of this file

1. Text S1 to S3
2. Figures S1 to S4

**Introduction** S1 contains a more complete discussion of clustering images, S2 and Figure S3 describe a synthetic example with a lower magnitude signal than used in the main text, S3 describes the attached animations, Figure S1 shows all the interferograms in our Sentinel-1 time series, Figure S2 shows the displacements recorded by several GPS stations surrounding Sierra Negra, and Figure S4 shows the probability density functions for the parameters estimated by our Bayesian inversion.

## Text S1

When clustering data it can be intuitive to think of each sample of our dataset as a vector with a length determined by the number of variables we have. Each sample can then be thought of as a point in a space with as many dimensions as we have variables and, our samples may form clusters that our algorithm of choice can identify. In the case of a dataset with 100 samples of three random variables, these clusters could be visualised in relative ease in a 3D space. However, when we wish to perform clustering with images (as are the output of sICA), our samples have as many variables as pixels, which in the case of interferograms tends to be of the order  $\sim 10^4$ . Consequently, we have relatively few samples in a very high dimensional space, and clustering becomes difficult. Therefore, clustering with images generally requires a specialised distance metric instead of using measures such as the Euclidean distance.

The ICASO algorithm uses the absolute value of the correlation between source pairs as a similarity measure. Through taking the absolute value, sign flipped versions of the same source attain a high similarity and do not form duplicate clusters. A trivial

step can be performed to convert similarities to distances (e.g.  $\mathbf{D} = 1 - \mathbf{S}$ , where  $\mathbf{D}$  is the distance matrix, and  $\mathbf{S}$  is our similarity matrix).

## Text S2

Figure S3 displays the results of applying our automatic detection algorithm to a time series similar to that used in Section 2 which features a more subtle signal and requires tuning of the frequency with which the lines of best fit are redrawn. The new signal which enters the time series from interferogram 90 onwards is of very small magnitude, and significant deviations from the lines of best fit are only observed for the RMS cumulative residual when the lines of best fit are redrawn every 60 interferograms, and not in the case when they are redrawn every 20 interferograms. However, in the case that the lines of best fit are redrawn infrequently, the algorithm flags several events erroneously (“false positive” results), such as the peaks of IC3’s time course. We believe this demonstrates the importance of the correct configuration of the parameter, as our algorithm’s sensitivity to small signals is achieved at the expense of an increased likelihood of false positive results. Further use of our algorithm during its future application to other volcanoes is likely to shed further light on this issue, but remains beyond the scope of this initial study.

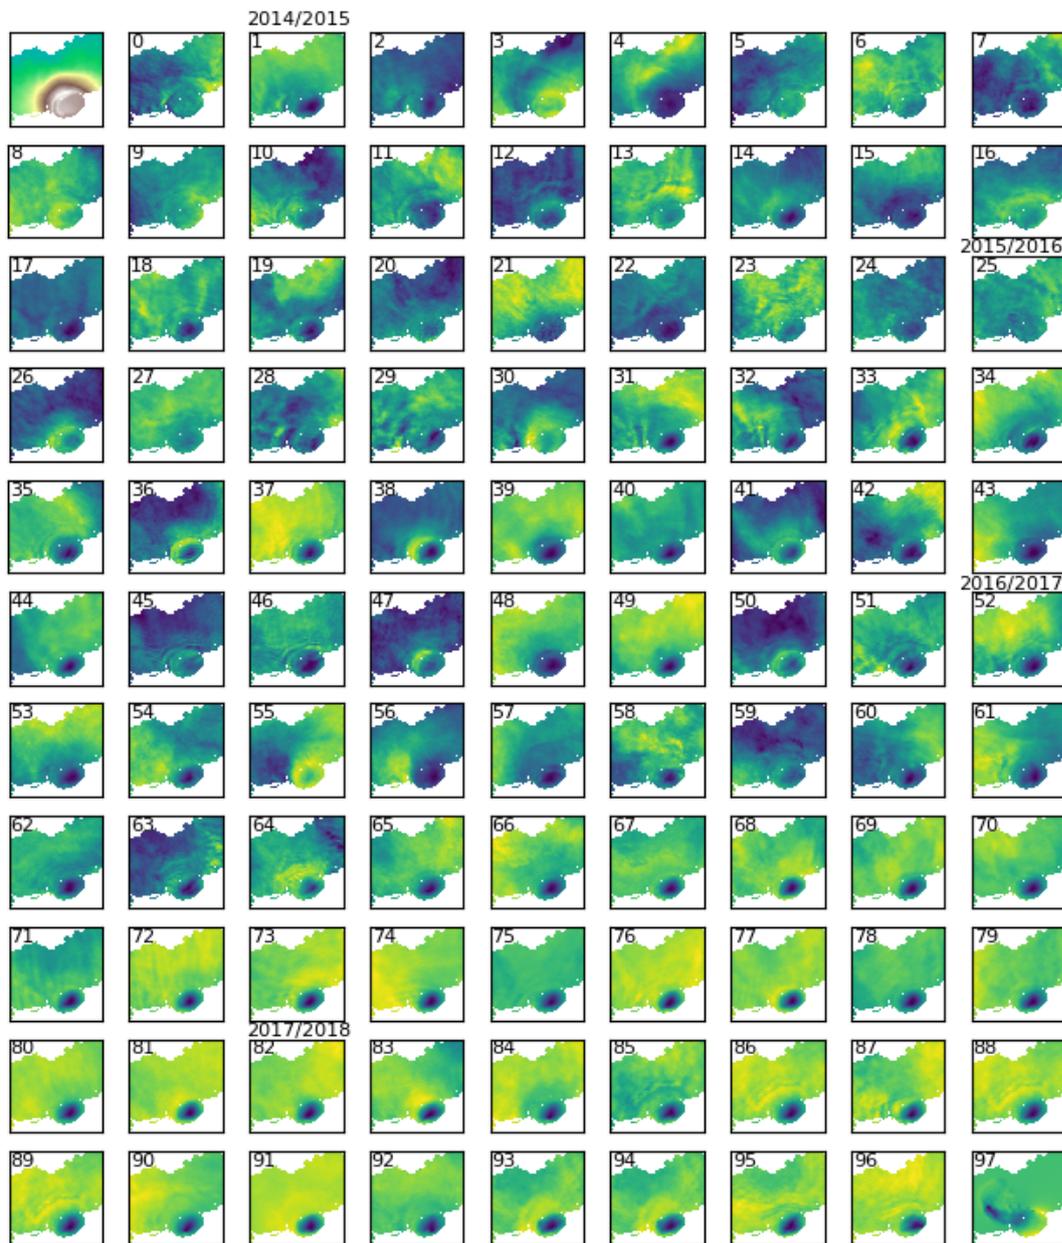
It should also be noted that the sinusoidal trend in IC3’s time course is due to the strength of our synthetic topographically correlated APS varying seasonally. The less sensitive case in which the lines of best fit are redrawn every 30 interferograms is more successful in avoiding the false positives that are seen when the lines are redrawn every 60 interferograms. This seasonal variation is an intrinsic part of our the nature of the baseline data, and through fitting a linear trend we fail to accurately characterise it. Therefore, we postulate that in future use, more complex functions may allow us to characterise the temporal nature of certain atmospheric signals, and so ultimately increase the sensitivity of the detection algorithm.

## Text S3

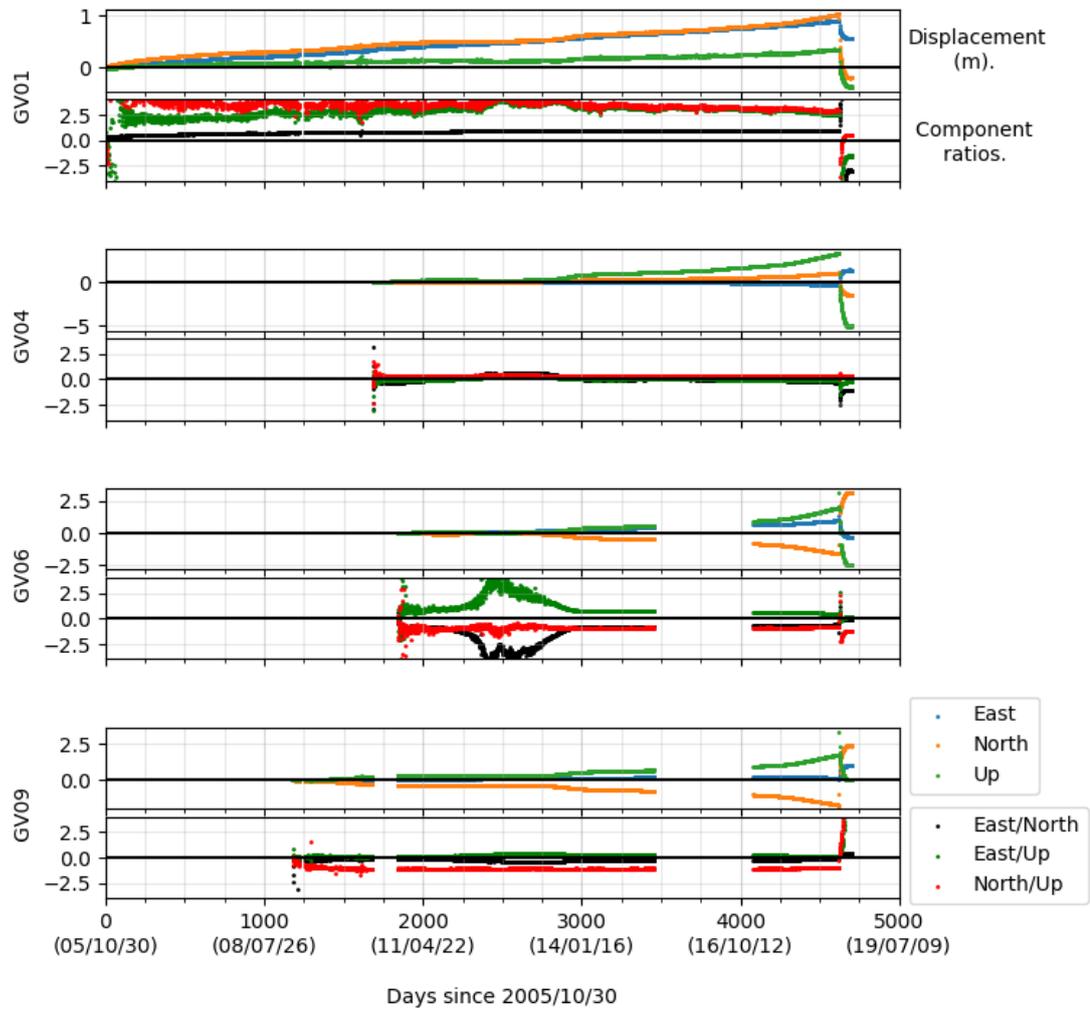
Implementing ICA to recover statistically independent sources is very closely related to applying PCA to recover uncorrelated sources. Informally, two signals are statistically independent if knowledge of the value of one signal does not convey any information of the value of the other, which contrasts with correlation that merely measures the linear relationship between variables. A similar measure applied to signals is that of correlation, and two signals are said to be uncorrelated if their covariance is 0. The difference between uncorrelatedness and the more important statistical independence can be demonstrated with two random variables that, when plotted in 2d space, form a circle around the origin. Whilst the correlation of these values is 0, it is clear that they

---

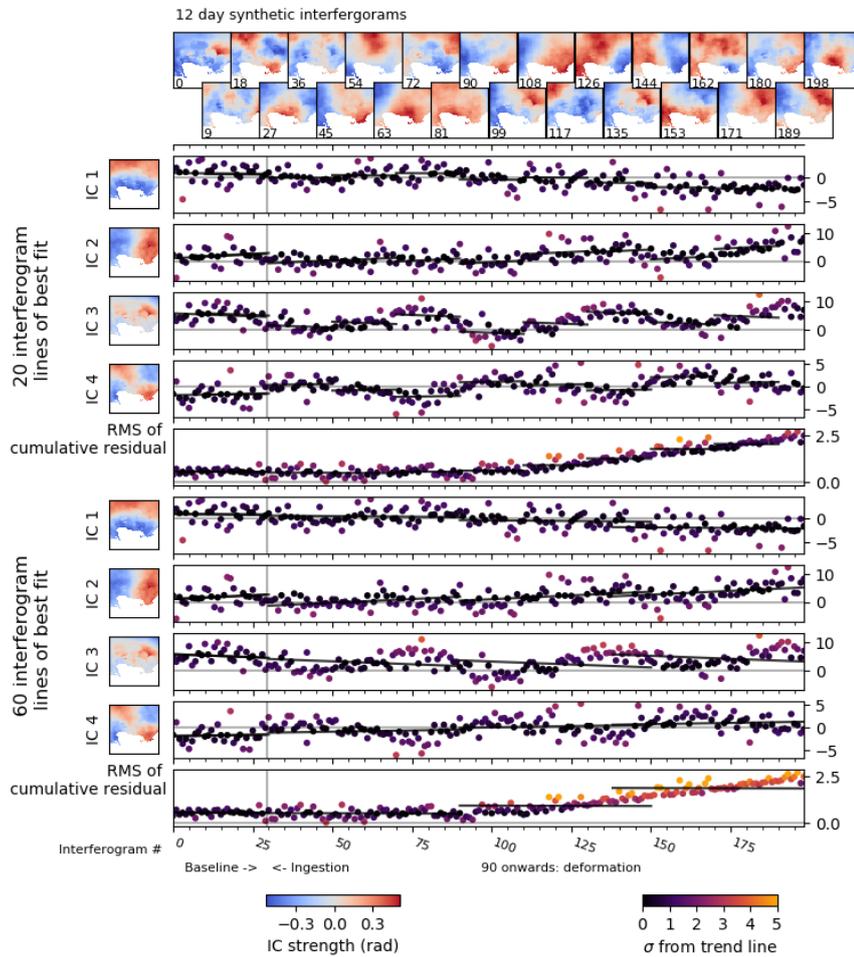
are not statistically independent as knowledge of one signal conveys information about the second signal (e.g. signal 1 attaining its maximum value conveys the information that signal 2 must also be at its median value). Therefore, two uncorrelated variables can still provide information about each other. The wealth of successful applications of ICA to BSS problems can be used to justify the expectation that it would outperform PCA as two physical processes that are unrelated (such as deformation at a volcano and atmospheric delay) are statistically independent, as opposed to merely uncorrelated.



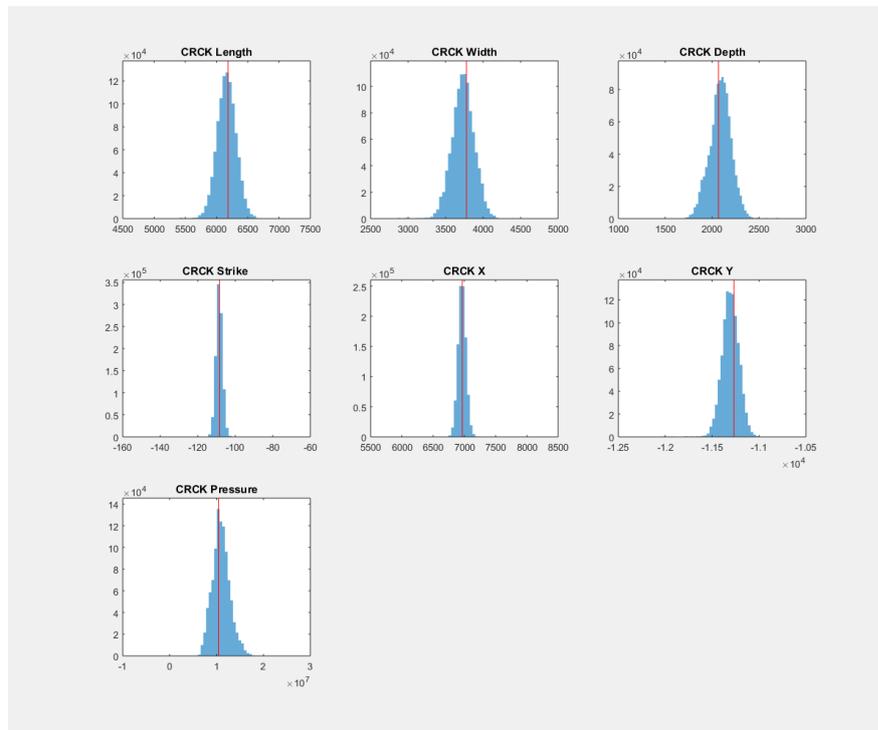
**Figure 5.4:** Time series of Sentinel-1 interferograms (0 – 97) and the corresponding DEM.



**Figure 5.5:** GPS data spanning the 2005 – 2018 inter-eruptive period. For each station (GV01, GV04, GV06, and GV09) the top plot shows the cumulative displacement in either East/North/Up directions, whilst the lower plot shows the ratios of each possible pair of directions (e.g. East vs. North). With the exception of a brief period around day 2500, the ratios remain constant, which we conclude shows that the style of deformation observed during the Sentinel-1 time series is likely to have remained similar for the entire inter-eruptive period.



**Figure 5.6:** Application of our automatic detection algorithm to a synthetic time series similar to that presented in Figure 3. The upper and lower halves of the figure show the effects of redrawing the lines of best fit every 20 and 60 interferograms, respectively. The first 90 interferograms contain signals from a topographically correlated APS (recovered as IC3), an east-west phase gradient (recovered as IC2), and a turbulent APS. The remaining interferograms (90–198) contain a small synthetic deformation signal in the centre of the frame which is difficult to identify by eye. However, as we are unable to fit this new signal with the learned components, the RMS cumulative residual increases in slope, but this is only detected (orange and yellow highlighting of points) when the lines of best fit are redrawn every 60 interferograms.



**Figure 5.7:** Results of our Bayesian inversion for our variable opening constant pressure horizontal dislocation. Red lines indicate the optimal values, which we report in the main text.

