Monitoring Highly Dynamic Land Surface Motion with Satellite Radar Interferometry

Conroy, Philip

DOI
10.4233/uuid:491585b8-860a-4e5b-8e3e-97b4e23ae4ec

Publication date
2025

Document Version
Final published version

Citation (APA)
Conroy, P. (2025). *Monitoring Highly Dynamic Land Surface Motion with Satellite Radar Interferometry*. [Dissertation (TU Delft), Delft University of Technology]. https://doi.org/10.4233/uuid:491585b8-860a-4e5b-8e3e-97b4e23ae4ec

Monitoring Highly Dynamic Land Surface Motion with Satellite Radar Interferometry

Monitoring Highly Dynamic Land Surface Motion with Satellite Radar Interferometry

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus Prof. dr. ir. T.H.J.J. van der Hagen,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op dag 28 maart 2025 om 12:30 uur

door

Philip Nicholas CONROY

Master of Science in Electrical Engineering
Kungliga Tekniska högskolan, Stockholm, Zweden
geboren te Cold Lake, Canada.

Dit proefschrift is goedgekeurd door de promotors.

Samenstelling promotiecommissie:

Rector Magnificus,	voorzitter
Prof. dr. ir. R.F. Hanssen	Technische Universiteit Delft, promotor
Dr. ir. F. Lopez-Dekker	Technische Universiteit Delft, copromotor
Onafhankelijke leden:	
Prof. dr. E. Stouthamer	Universiteit Utrecht
Prof. dr. A.V. Monti-Guarnieri	Politecnico di Milano
Dr. A. Kull	Tartu Ülikool
Prof. dr. ir. P.J.G. Teunissen	Technische Universiteit Delft
Overige leden:	
Dr. ir. F.J. van Leijen	Technische Universiteit Delft

Cover design: Frithjof Ehlers

NWO-NWA Living on Soft Soils: Subsidence and Society

ISBN 978-94-6384-752-0

Printed by Ridderprint

An electronic version of this dissertation is available at
http://repository.tudelft.nl/.

Preface

The first question you ask yourself when starting a PhD is "where do I even begin?" Somehow I'm still not sure how to answer that. You just start working on something, and somewhere, somehow, you find yourself suddenly in the thick of it. And before you know it, the time has passed and you are wrapping up! As I sit here now, thesis very nearly finished, I find myself asking the same question about this preface. So many people have been a part of this journey that it is difficult to even begin to acknowledge everyone.

So, first things first: I would like to thank my promoter and supervisor, Ramon Hanssen, for taking me on this journey and supporting my work for the past four years. Your constant energy kept me motivated, and your challenging questions kept me sharp and engaged throughout the project. One question above all was always a motivator: "Are we happy with this?", the answer always being "No", as there is always more to do, more problems to fix, and above all more to be learned. Yet, looking back over this entire project, seeing where we started and where we are now, I feel like I can finally answer "Yes". Also in the InSAR research group, I would like to thank Freek van Leijen. Without your constant support, this project would not have been possible. Thank you as well for such a careful reading of this thesis draft when you had so many other things to do. To the other PhDs in the group, Wietske Brouwer, Simon van Diepen, Yuqing Wang, Dita Lumban-Gaol, and Alex Lapadat, thank you for creating such a positive group to be a part of. I learned something from each one of you and I wish you all the best for your futures.

Next, I would like to thank Esther Stouthamer, the principle investigator of the Living on Soft Soils project which funded this work. Your enthusiasm and support for the development of an unproven technology like InSAR for peatland monitoring is amazing. The LOSS project gave me the opportunity to meet experts from so many different fields who I would otherwise have never encountered, and I grateful for having been able to learn so much about a topic which lies at the heart of Dutch society and culture. To all the other LOSS PhDs: best of luck with finishing up! Thank you as well to Sanneke van Asselen and Giles Erkens from Deltares for your support during this project, which would have been impossible without the high-quality extensometer data you shared freely with me, and thank you for your enthusiasm about my work!

At the TU Delft Department of Geoscience and Remote Sensing, I would like to thank my copromotor Paco Lopez-Dekker, who is more passionate about SAR than anyone else I have ever met. I feel like I learn something new every time we speak. Many thank-yous to Frithjof Ehlers, for serving as paranymph during my defence, for designing the cover art of this thesis, and of course for introducing me to the power of the squircle. Having a hungover breakfast together in Vienna during EGU, I wished you would just shut up

about this function to draw squares with rounded corners and let me enjoy my coffee, but who would have thought that it would actually help solve a nonlinear optimization problem I was having! To Yan Yuan, thank you for always keeping it real, being a reliable friend, and taking care of Abby even though she scratched the hell out of your apartment. To Ben Hernandez, thanks for all the memes. But more importantly, thanks for being a great friend and bouldering coach! I hope you liked proposition 10. Finally, I would like to thank everyone in the GRS department for making what has easily been the most positive and fun workplace I have ever been a part of. The energy in this department is something special, and it keeps me happy about showing up to work every day.

Thank you to my remaining committee members Andrea Monti-Guarnieri, Ain Kull and Peter Teunissen for reading my work and for your positive and constructive feedback.

To my parents, who supported me during my studies and always encouraged me. Thank you to my dad, Phil, for encouraging me to study engineering and to stick with it when the going got tough. I guess radar is becoming a family business. Thank you to my mom, Andrea, for always giving me a home to come back to and for always being so supportive.

Finally, thank you to Julia, for saying yes to going on this adventure and moving to Delft with me, and for always being there for me.

Philip Conroy
Delft, March 2025

Contents

List of Abbreviations

AOI	Area of interest
APS	Atmospheric phase screen
CCDS	Continuously coherent distributed scatterer
CCPS	Continuously coherent point scatterer
CRB	Cramer-Rao bound
DECADE	Delft Contextually-Aided Distributed Scatterer Environment
DEM	Digital elevation model
DePSI	Delft Persistent Scatterer Interferometry
DOGM	Direction of ground motion
DORIS	Delft Object-oriented Radar Interferometric Software
DOY	Day of year
DS	Distributed scatterer
EDC	Equivalent daisy-chain
ESM	Equivalent single master
ETRF	European Terrestrial Reference Frame
GHG	Greenhouse gas
GNSS	Global navigation satellite systems
GPS	Global positioning system
HMM	Hidden Markov model
IEEE	Institute of Electrical and Electronics Engineers
IGRS	Integrated geodetic reference station
ILS	Integer least-squares
InSAR	SAR interferometry
KNMI	Royal Netherlands Meteorological Institute
LEO	Low Earth orbit
LOS	Line of sight
LSTM	Long short-term memory
MCF	Minimum cost flow
MSE	Mean squared error
NOBV	Nationaal Onderzoeksprogramma Broeikasgassen Veenweiden
OMT	Overall model test
PDF	Probability density function
PS	Point scatterer
PSI	Point scatterer interferometry
RAR	Real aperture radar
RCS	Radar cross-section
RIVM	National Institute for Public Health and the Environment

RMSD	Root mean squared difference
RMSE	Root mean squared error
RNN	Recurrent neural network
SAR	Synthetic aperture radar
SBAS	Small baseline subset
SHP	Statistically homogeneous pixel
SLC	Single-look complex
SNR	Signal-to-noise ratio
SPAMS	Simple Parameterization for the Motion of Soils
SR	Success rate
SRTM	Shuttle Radar Topography Mission
SVD	Singular value decomposition
TCDS	Temporarily coherent distributed scatterer
TCPS	Temporarily coherent point scatterer

List of Symbols

Radar Terms

$\hat{a}$	Float ambiguity
$\check{a}$	Integer ambiguity
B	Interferometric baseline
$B_{\perp}$	Perpendicular baseline
$B_{\parallel}$	Parallel baseline
BW	Bandwidth
f_s	Sampling frequency
$f(z)$	Vertical scattering profile
k_z	Vertical wavenumber
L	Effective number of looks
r	Slant range
r_i	Slant range in image i
OSR	Oversampling ratio
δ_{az}	Azimuth resolution
δ_{r}	Range resolution
Γ	Coherence matrix
γ_{dc}	Daisy-chain coherence
$\gamma_{\text{i,j}}$	Coherence between epochs i and j
λ	Wavelength
ϕ	Interferometric phase
ϕ_{ij}^{mn}	Interferometric phase of epoch j w.r.t. i and scatterer n w.r.t. m
ϕ_{ij}	Interferometric phase of epoch j w.r.t. i
ϕ_0	Expectation value of interferometric phase
ϕ_{dc}	Daisy-chain interferometric phase
ϕ_{edc}	Equivalent daisy-chain interferometric phase
ϕ_{esm}	Equivalent single mother/master interferometric phase
ϕ_{sbas}	Subset of small baseline interferometric phases
ϕ_{ijk}	Closure phase of epochs i, j, k
ψ	SAR SLC pixel phase
θ	Look angle
θ_{inc}	Incidence angle
Ω	Set of selected pixels in interferometric stack

Soil-Related Terms

E	Daily mean evapotranspiration
f_o	Fraction of organic matter content in soil
f_c	Fraction of Carbon in soil organic matter
F	Fraction of subsidence caused by oxidation
I	Irreversible component of soil surface displacement
p	Psychrometric constant
R	Reversible component of soil surface displacement
s	Derivative of saturation water vapor pressure vs. temperature
S_{day}	Daily mean solar radiant exposure
T_{day}	Daily mean temperature
x_E	Evapotranspiration coefficient
x_I	Irreversible subsidence coefficient
x_P	Precipitation coefficient
z	Soil surface displacement
z_M	Modeled soil surface displacement
z'	Soil subsidence rate
$\bar{z'}$	Mean annual subsidence rate
ΔH_v	Latent heat of vaporization of water
ρ	Soil bulk density
ρ_w	Bulk density of water
τ	Integration time

Machine Learning Terms

h	Neural network prediction
J	Loss function
s_t	Input data sequence at epoch t
w	Model parameters (network weights)
w_{opt}	Optimal model parameters (network weights)
y_t	Training target value at epoch t

Observation Theory Terms

A	Design matrix
DOF	Degrees of freedom
$\hat{e}$	Model residuals
M	Number of observations
N	Number of parameters
Q	Variance-covariance matrix
T	Test statistic
T_{crit}	Critical value
x	Model parameters
$\hat{x}$	Estimated model parameters
y	Observations
$\hat{y}$	Adjusted observations
α	Significance level

General Mathematics

e	Euler's number
i	Imaginary unit
R^2	Coefficient of determination
t	Time or epoch
t_0	Initial time or epoch
x	Position
α	Significance level
λ	Eigenvalue
π	Pi
σ	Standard deviation
σ^2	Variance
ξ	Eigenvector
∞	Infinity
$\mathbb{C}$	Set of all complex numbers
$\mathbb{R}$	Set of all real numbers
$\mathbb{Z}$	Set of all integers

Mathematical Operators

$\arg\{\cdot\}$	Argument (complex phase)
$\arg\min\{\cdot\}$	Arguments of the minima
diag	Diagonal matrix
$E\{\cdot\}$	Expectation operator
$D\{\cdot\}$	Dispersion operator
$\in$	Element in the set
$\forall$	For all
$\exp\{\cdot\}$	Exponential function, equivalent to $e^{\{\cdot\}}$
$\ln\{\cdot\}$	Natural logarithm
$W\{\cdot\}$	Wrapping operator
$\Gamma(\cdot)$	Gamma function
$\Sigma\{\cdot\}$	Summation
$\{\cdot\}^*$	Complex conjugate
$\circ$	Hadamard product
$\{\cdot\}^{\mathrm{T}}$	Matrix transpose
$\{\cdot\}^{-1}$	Inverse
$\widehat{\{\cdot\}}$	Estimator/ estimated value
${}_pF_q(\cdot)$	Generalized hypergeometric function

Summary

Over the past three decades, synthetic aperture radar (SAR) interferometry (InSAR) has become one of the most important Earth observation technologies in the world, and its use has become common in applications such as topographic mapping, monitoring earthquakes and volcanoes, as well as the built environment. Despite these advances, many technical and scientific challenges remain unsolved in the field, which prevent its use across diverse regions and biomes. One such type of region are wetlands and peatlands, which are notoriously challenging to monitor remotely due to poor signal quality and rapidly changing conditions between SAR acquisitions. This problem is particularly relevant in the Netherlands, because a significant portion of the country is composed of drained peat and clay soils which lie below sea level. These "soft soils" exhibit highly dynamic temporal behaviour that is closely linked to the phreatic groundwater system. In addition, they also exhibit a slow, irreversible subsidence caused by compaction and oxidation, the latter of which is a greenhouse gas (GHG) emitting process. It is this slow, irreversible subsidence component which scientists, governments, farmers and other stakeholders are trying to better understand, and evaluate the risks it poses.

Previous efforts in monitoring the cultivated soft soil regions of the Netherlands by InSAR have been hampered by two main problems, which in this work are referred to as "cycle slips" and "loss-of-lock". The former refers to consistent errors made in ambiguity resolution due to signals which exhibit such highly dynamic behaviour that standard algorithms cannot correctly interpret the wrapped phase data. The latter term refers to a permanent and irreparable loss of coherence in an interferometric SAR data stack. It is common in peatland regions for coherence levels to rise and fall seasonally, and in general, no coherent interferometric combination exists between the coherent periods. This condition means that the interferometric time series is severed during these incoherent periods, and only intermittent, disconnected temporal subsets of data are useable.

An appropriate parameterization of the expected land surface motion is needed in order to overcome both of these challenges. By using the disconnected coherent subsets of wrapped phase data to estimate a set of model parameters that describe the expected land surface motion to within an acceptable level of accuracy, we can subsequently use such a model to both estimate the correct ambiguity levels of the unwrapped phase, and interpolate over the incoherent periods in order to align the coherent subsets of data. Therefore, the methodology developed in this research, called Delft Contextually-Aided Distributed Scatterer Environment (DECADE), focuses on how we can use the accumulation of sparse coherent data subsets over time to learn how a certain patch of the Earth behaves, rather than on simply supplying an InSAR time series of displacements without any context.

This work concludes by showing the first InSAR-derived estimates of land surface motion of the soft soil regions of the Dutch Green Heart (*Groene Hart*). These estimates are compared to in-situ measurements at three locations, and there is a remarkably good

level of agreement between the results of the two techniques. This indicates that the DECADE methodology is indeed capable of enabling the use of InSAR to monitor these highly dynamic and decorrelating peatland regions. Finally, by combining estimates of the average irreversible subsidence rate with soil data of the region, an approximate worst-case estimate of greenhouse gas emissions caused by soil oxidation is established.

Samenvatting

Radar interferometrie (InSAR) heeft zich gedurende de laatste dertig jaar ontwikkeld tot een van de belangrijkste hedendaagse ruimtegeodetische aardobservatietechnieken. De techniek wordt gebruikt voor het in kaart brengen van de topografie van de aarde in de vorm van hoogtemodellen, of voor het bestuderen van de deformatie van het aardoppervlak. Een zich snel ontwikkelende toepassing is het monitoren van de gebouwde leefomgeving. Alhoewel dit heeft geleid tot een veelheid van operationele producten, zijn er nog significante uitdagingen, zowel technisch als wetenschappelijk, die de directe toepassing in specifieke gebieden in de weg staan. Een pregnant voorbeeld is het veenweidegebied. De uitdagingen worden hier gevormd doordat de complexe radarreflectie van het grasland relatief snel en aanzienlijk kan veranderen tussen opeenvolgende radaropnames. Deze vegetatieëffecten zorgen voor ruis in de data, en in combinatie met de fasemeerduidigheden en het dynamisch gedrag van de veenbodems leidt dit tot significante problemen in de schatting van de bodembeweging. Dit probleem is zeer relevant voor Nederland, omdat een groot deel van het land bestaand uit veen- en kleigrond blootgesteld aan een kunstmatige variatie van de grondwaterstand. Deze 'zachte' bodem blijkt een zeer sterk dynamisch gedrag te hebben, sterk gecorreleerd met het freatisch grondwaterniveau. Tevens is er sprake van een langzame maar gestage onomkeerbare bodemdaling, ten gevolge van compactie en oxidatie, hetgeen tevens leidt tot de uitstoot van broeikasgassen. Het is met name deze langzame en onomkeerbare bodemdaling die van groot belang is om te begrijpen, om toekomstig gedrag te kunnen voorspellen en er adequate maatregelen tegen te nemen. Hiervoor moet het echter eerst mogelijk zijn om deze te kunnen meten.

Eerdere pogingen om de bodembeweging van het Nederlandse veenweidegebied te monitoren met behulp van InSAR legden twee beperkingen bloot: (i) het abusievelijk overslaan van een semi-golflengte in de schatting van de geheeltallige meerduidigheden (meerduidigheidssprong of 'cycle-slip') en (ii) het onomkeerbaar verbreken van het correlatieverband in de keten van radarwaarnemingen (correlatiebreuk of 'loss-of-lock'). Cycle-slips zijn het gevolg van de combinatie van een sterk dynamisch signaal en ruis, waardoor standaard algoritmes verkeerde resultaten geven. Loss-of-lock treedt op wanneer er een permanante (onomkeerbare) verandering in de reflectieëigenschappen van het terrein plaatsvindt, waardoor de interferometrische coherentie tussen opeenvolgende radaropnames tijdelijk onder een bepaald minimum komt. In het veenweidegebied is deze variatie in coherentiewaarden onvermijdelijk, bijvoorbeeld tussen de seizoenen als gevolg van gewasgroei en landgebruik. Als gevolg hiervan moet de serie van radaropnames als onderbroken of gefragmenteerd worden beschouwd.

Beide beperkingen zijn te ondervangen door een adequate parameterisatie van de verwachte bodembeweging. Door gebruik te maken van de gefragmenteerde coherente delen van de waarnemingstijdreeksen van de relatieve fase kunnen we de gekozen modelparameters schatten. Hierbij wordt bewust gekozen voor een niet-perfecte parameterisa-

tie zolang deze maar een acceptabele precisie en betrouwbaarheid heeft om de correcte meerduidigheden te kunnen schatten. Daarna worden de incoherente perioden in de waarnemingstijdsreeks overbrugd. De ontwikkelde methode, Delft Contextually-Aided Distributed Scatterer Environment (DECADE), optimaliseert de geaccumuleerde serie van temporeel schaars bemonsterde deelperioden, met als doel het afdoende beschrijven van het dynamisch gedrag van een bepaald gebied, zoals een perceel, in plaats van het op afroep produceren van een InSAR bewegingstijdserie zonder context.

Als eindresultaat wordt binnen dit onderzoek de eerste vlakdekkende schatting van bodembeweging van het veenweidegebied op basis van InSAR gepresenteerd, toegepast op het Groene Hart. De resultaten zijn vergeleken met onafhankelijke veldmetingen op drie locaties, waarbij er een overtuigende mate van overeenkomst wordt aangetoond. Dit betekent dat de DECADE methode in staat is om gebruikmakend van ruimtegeodesie (satellietwaarnemingen) het veenweide gebied adequaat te monitoren. Ook is aangetoond dat de combinatie van de geschatte onomkeerbare (irreversibele) bodemdaling met additionele bodeminformatie kan worden gebruikt om een bovengrens-scenario van broeikasgasemissies ten gevolge van oxidatie vast te stellen.

1 Introduction

God created the Earth,
but the Dutch created the Netherlands.

Proverb

1.1 Motivation

SAR Interferometry is the only geodetic technique that has the potential to monitor land surface motion over a wide spatial extent at frequent intervals with high precision, however, highly dynamic displacement signals and temporal decorrelation have prevented the technique from being successfully applied over cultivated peatlands. This research aims to overcome these challenges and enable land surface monitoring of these dynamic and noisy environments.

1.2 Background

Despite their critical importance for carbon emissions, water security and biodiversity, northern peatlands remain a poorly understood biome due to the difficulties in producing observations with sufficient quality and spatiotemporal resolution. Drained peatlands subside over time as the upper layers of soil dry out and oxidize, and due to consolidation of the deeper soil layers caused by loading conditions of the layers above [1]. This subsidence releases significant amounts of carbon dioxide and nitrous oxide into the atmosphere through soil oxidation, and creates risks for flooding due to land elevation loss [2], [3].

This is of particular importance in the Netherlands, where the topic of subsidence is part of a larger ongoing societal debate regarding land use. There are significant questions about balancing the cultural heritage and traditional agriculture of the region with long-term sustainability concerns such as nitrogen pollution and water security in the context of global climate change. Much of the coastal plain already lies below mean sea level and additional subsidence constitutes a significant threat to the security of the country. The

Dutch government has committed to a 1 Mt reduction in CO_2-equivalents per year originating from its peatlands before 2030 [4], corresponding to a reduction of approximately 25% according to a government report [5] (based on a methodology by Wageningen Environmental Research [6], [7]), a decision with significant impact on future land use, agriculture and the economy that requires adequate assessment of subsidence rates.

While it is clear that soil subsidence is occurring in the peatland regions of the Netherlands, producing accurate and reliable estimates of the phenomenon from geodetic measurements has proven challenging. This is due in part to a lack of understanding how these environments behave over short timescales. Until recently, most geodetic techniques used to study the problem are either not representative, do not provide the required spatial sampling density, or provide data too infrequently to capture the rapid seasonal and sub-seasonal movements of the soil. Therefore, geologists and soil scientists have had to make do with long- (1000–100 years [2]) and medium- (100–10 years [8], [9]) term estimates of ground subsidence over time, or at best, semi-annual surveys of a few specific locations [10]. Moreover, researchers require high spatial coverage in order to accurately assess and quantify the various land subsidence phenomena across the country. Soil heterogeneity in the Holocene as well as the discontinuous nature of the Dutch landscape which is split apart by drainage ditches and canals into separate polders with varying land usage and vegetation can result in a high spatial variability in the ground displacement signal.

The only potentially viable technique available to provide accurate land surface motion time series estimates with wide spatial coverage and the required temporal frequency is SAR interferometry (InSAR). However, several technical challenges have prevented the use of InSAR in this context. In particular, the soil dynamics under study exhibit both rapid temporal behaviour on a daily time scale, as well as a large dynamic range, a problem which is made worse by the very poor signal quality which is obtained from radar observations of these regions in combination with the wrapped nature of the SAR observations. The goal of this research is to identify, quantify and overcome these challenges in order to enable the use of InSAR as a tool for studying the land surface motion of Dutch peatlands.

1.3 Prior Knowledge

1.3.1 Land Motion and Subsidence in the Netherlands

Large regions of the Netherlands have been reclaimed by pumping water out of wetlands, exposing low-lying, fertile peat soils which are used for intensive agriculture. As much of the western half of the country lies below mean sea level, water must be continuously monitored, pumped and drained out of the ground to prevent flooding and damage. Once the ground water level is below the surface level, oxidation begins to take effect as the organic content in the soils dry out. Over time, this irreversible loss of soil becomes significant, and water levels must be lowered again once they become too high relative to the ground to prevent flooding and enable agriculture. Repetition of this procedure over several centuries has slowly caused the mean level of the ground to sink drastically, an effect which is compounded by sea level rise due to climate change [2], [11] (Fig. 1.1). Furthermore, the water management and flood defence infrastructure in the country have

Figure 1.1: Visualization of the cycle of land subsidence caused by active water management in the Netherlands. Figure taken from the Netherlands Council for the Environment and Infrastructure (RLI), 2020.

prevented river sediment from being naturally transported to these low-lying regions, creating additional susceptibility to flooding [12].

The first indications that these cultivated peatlands exhibited highly dynamic behaviour came in 1977, when a levelling study by Schothorst [10] showed that a test polder in Zegveld exhibited large seasonal fluctuations in elevation. However, it was not until over 40 years later that a study by van Asselen et al. [13] based on extensometers located throughout the country revealed that the motion of the land surface is much more dynamic that previously thought, and is largely composed of seasonal and sub-seasonal signals originating in the shallow subsurface which are approximately an order of magnitude greater than the slow, irreversible subsidence processes. Despite these new measurements, the link between precipitation and evapotranspiration and the surface elevation had not been made. Furthermore, the fact that precipitation and evapotranspiration are responsible for such rapid and relatively large shrinkage and swell has significant ramifications for the representativity of surveys by levelling or airborne laser scanning, which could potentially produce significantly biased estimations depending on when the campaigns are carried out, as motions on the order of 5–10 cm are possible within the time span of a few weeks.

1.3.2 Greenhouse Gas Emissions

Drained peat soils release greenhouse gases (GHGs) as they decompose in the presence of oxygen, and it is estimated that peat oxidation is responsible for approximately 2.4% of the total GHG emissions of the Netherlands (4.25 Mt-CO_2-equivalents/year, corresponding to 19 tons of CO_2-equivalents per hectare per year) according to the National Institute for Public Health and the Environment (RIVM) [14], [15]. Re-wetting these drained peatlands could significantly reduce these GHG emission rates despite the fact that wetlands produce CH_4, a more potent GHG, due to the fact that wetlands act as a carbon sink as part of the peat creation process. A recent study reviewing global in-situ peatland GHG emission measurements has found that the sum of total emissions from the three major GHGs emitted by peatlands and wetlands (CO_2, CH_4 and N_2O) is minimized when the water table is near the surface (-30 to -5 cm depth), resulting in near-zero net emissions [16]. Finding the optimal water table level for the Dutch peatlands will require a balancing between the interests of various stakeholders to allow for productive agriculture, meeting environmental targets and ensuring flood protection and water security. This again underscores the need to be able to accurately quantify the motion of the shallow soils on a regional scale.

Apart from the actual estimation of an accurate long-term average subsidence rate, a major source of uncertainty in deriving GHG emission estimates from subsidence is the fraction of subsidence which can be attributed to oxidation versus other mechanisms which do not create emissions [7]. A generally accepted range for this fraction is between 0.33–0.67 [17], however, other ranges have been reported, such as 0.25–0.71 in [18], and a study based on in-situ measurements in the Zegveld area of the Green Heart region found a fraction as large as 0.85 [10]. As GHG emissions scale linearly with this fraction of subsidence, the assumption of a certain fraction will have a significant impact on the resulting GHG emission estimate.

1.3.3 Previous InSAR Studies of Peatland Surface Motion

It has so far been impossible to monitor the motion of peatlands using InSAR due to the combination of very rapid displacement signals, poor signal quality, and the inherent InSAR phase ambiguity problem. This is not only the case in the managed and cultivated peatlands of the Netherlands [19], [20], but also in wild peatlands in other northern regions such as Scotland and Estonia [21]–[23].

In 2008, Caro Cuenca and Hanssen [24] indirectly observed effects of peat oxidation in the motion of shallow-founded buildings using point scatterer (PS) interferometry. Work on directly monitoring the land surface motion of the Dutch peatlands with distributed scatterer (DS) InSAR began with Morishita and Hanssen in 2015 [19], who applied a multi-satellite approach combining ALOS, Envisat, Radarsat-2 and TerraSAR-X data together in a joint small baseline subset (SBAS) inversion [25] (Sec. 2.4.8) to monitor subsidence in the Delfland region. The resulting time series strongly overestimated the true average subsidence rate due to challenges with phase unwrapping (Sec. 2.4.9) and rapid temporal decorrelation (Sec. 2.4.5). Nevertheless, the work was significant in demonstrating the combination of multiple disconnected datasets into a unified estimation framework. In the same year, the authors also documented the effects of temporal decorrelation of SAR observations of grasslands in the region at different frequency

Table 1.1: Overview of Soil Subsidence Monitoring Techniques

Technique	Precision	Spatial Coverage	Temporal Coverage	Representativity
Extensometer	High	Low	High	High
Leveling	High	Medium	Low	Low
Airborne Laser Scanning	Low	High	Low	Low
GNSS	Medium	Low	High	Low
InSAR	High	High	Medium	High

bands [26]. In 2018, Samiei-Esfahany and Hanssen conducted a feasibility study of applying DS InSAR techniques to monitor peatland surface motion in the Veendam region of the Netherlands [27], and concluded that there was not enough useful information content in either C- or X-band data stacks to successfully invert a complete time series of motion. However, the study does note that individual coherent periods occurring during wintertime do contain useful information that could be utilized, an insight key to developing the processing methodology developed in this work. In 2020, Heuff and Hanssen [28] developed a DS processing method employing the state of the art "EMI" estimator [29] (Sec. 2.4.8), however they encountered the same decorrelation and phase unwrapping challenges as in previous studies.

1.3.4 Other Geodetic Techniques

Several geodetic techniques apart from InSAR have been used to estimate land surface elevation change in the Netherlands, with their applicability for monitoring shallow soil subsidence summarized in Table 1.1. Monitoring the changes in measured elevation over time by levelling and airborne laser scanning has been used to produce subsidence estimates for cultivated peatlands [8], however, the accuracy of these techniques for this application are poor due to the sensitivity of laser scanning to changes in vegetation and the lack of reference benchmarks representative of the soil motion. Generally, for both techniques, only one survey is made approximately every 5–10 years due to the cost and effort involved. While permanent global navigation satellite systems (GNSS) stations provide a continuous stream of absolute position measurements over time, GNSS uncertainties are generally much higher in the local up component, and the data is limited to the locations where a permanent GNSS station is present. More importantly, these stations are often set on foundations for stability, which will make the stations insensitive to the top layer of shallow soil movement. Recently, new in-situ measurements by extensometers have become available, which provide accurate (sub-mm level) soil displacements at an hourly rate [13]. While this data is extremely valuable, and indeed used as validation in this work, it is limited to where and when the devices have been installed. InSAR is the only technique which can potentially provide both high spatial coverage with frequent enough temporal sampling in order to accurately track rapid ground displacement, making it a good candidate to monitor the soft soil dynamics over the entire country if the open technical issues preventing its implementation can be overcome.

1

1.4 Research Gaps Addressed by this Work

1.4.1 Overview

This work seeks to address several problems which have been identified as preventing the use of InSAR for monitoring cultivated peatland regions: the lack of a functional parameterization of the expected motion of the target, in this case the shallow soils, as well as the problems related to phase unwrapping termed *cycle slips*, and those related to temporal decorrelation termed *loss-of-lock*, which are explained in the following sections.

1.4.2 Unknown Displacement Model

While it is known that shallow soil motion is mainly driven by changes in phreatic groundwater levels [9], [30], [31], its short-term dynamics have not been accurately modeled, due in large part to a lack of observational data on the required time scales. The first sub-daily in-situ measurements by extensometer of soil motion in the Netherlands became available around 2020 [13], [32]. As a consequence, it is not known how to optimally parameterize the rapid motion of shallow soft soils, which is necessary to invert the wrapped SAR phase observations to useful displacement estimates. Previous InSAR studies [19], [24] have attempted to parameterize the functional model as a combination of sinusoids and linear functions, but this formulation was not sufficient to accurately capture the complex motion exhibited by the shallow soil motion. Ch. 3 details the development of two shallow soil motion models based on the extensometer measurements for this purpose.

1.4.3 Cycle Slips

The term cycle slip is used in this text to refer to the consistent phase unwrapping errors (Sec. 2.4.9) caused by algorithms that cannot correctly interpret displacements which exceed a quarter-wavelength (half of a 2π ambiguity) along the line of sight (LOS), which at C-band, corresponds to 1.4 cm LOS. It is very challenging to correctly estimate the ambiguity level for signals exceeding this level of displacement between acquisitions, as such large motions in certain direction are generally interpreted as small motions in the opposite direction. Conversely, small observed phase changes may in fact correspond to large vertical motions with an additional missed ambiguity level. This situation is made worse by the strong noise levels that are encountered, which makes it difficult to separate real displacement from unwanted noise. We found that these moments of rapid displacement are often due to soil uplift caused by rain-induced swell during the autumn and early spring, and incorrectly unwrapping these signals can lead to strong overestimation of the average subsidence rate in the region [19]. Ch. 4 describes the development of a phase unwrapping methodology that takes soil and water conditions into account in order to anticipate the correct ambiguity level.

1.4.4 Loss-of-Lock

This work also introduces the term *loss-of-lock*, a condition that arises from strong temporal decorrelation (Sec. 2.3.5) such that the chain of connected interferometric pairs in the time series is cut. This corresponds to a partitioned coherence matrix (Sec. 2.4.1) in which there are isolated blocks of coherent data, but with no coherent connection between them. In an agricultural context, this generally occurs seasonally during the summer due

to vegetation growth, or during periods of farming activity such as plowing and harvesting. While temporal decorrelation is a well-known effect, the implications of inverting a complete time series with loss-of-lock are rarely considered or addressed in the community of InSAR users, which can lead to severe misinterpretation of the observational data. This problem is addressed by Ch. 5.

1.5 Scope Limitations

1.5.1 L-Band Data

It is worth noting that the issues described above are expected to be mitigated by using L-band SAR data instead of C-band. In particular, at L-band frequencies, the quarter-wavelength phase wrapping threshold is approximately 6 cm along the LoS, meaning that the maximum expected signal gradient between subsequent acquisitions (approx. 2 cm) fits comfortably within this threshold. The temporal decorrelaton rate of L-band observations of grasslands is also significantly slower [26] than C-band. However, there are several open issues with using L-band data as well as project constraints which have prevented its use in this research. Apart from a small number of acquisitions by the Japanese ALOS 1 satellite made between 2007–2011, the only satellite mission currently making regular (sub-annual) L-band SAR acquisitions over the Netherlands is the Argentinian SAOCOM mission, which began acquiring imagery of the Netherlands in mid-July of 2023, and is providing a small but growing stack of about 20 images times two tracks (at time of writing). It should be noted that the integration of L-band data into the overall methodology is not trivial, nor does its inclusion guarantee that the problems discussed in Sec. 1.4 will be immediately solved, as shown by the challenges faced in previous studies using L-band data [19].

1.5.2 Deeply-Based Land Motion

The objective of this work is to enable the monitoring of land surface motion originated by shallow-based processes in the Holocene. Land surface motion originating from deeply-based processes such as mining, gas extraction and water extraction from deep aquifers are not in the scope of this research. This also limits the spatial extent of the analysis to peatland regions in the Netherlands in which no significant deep signal is expected; precluding analysis of the regions in the North of the country (ex. Groningen) which are also peatlands, but where gas extraction constitutes a significant component of the overall displacement signal.

1.5.3 Transformation from Relative to Absolute Displacements

The estimations of motion originating from InSAR are inherently relative, due to the double-differenced nature of the phase differences used (Sec. 2.3.3). Converting these estimates from a relative to an absolute reference frame (i.e. a frame not dependant on the motion of a reference pixel, such as the European Terrestrial Reference Frame (ETRF)) requires analysis of the motion of the reference point used with another geodetic technique and is out of the scope of this work.

2

A Review of Synthetic Aperture Radar and Interferometry

InSAR doesn't work.

Ramon Hanssen

This chapter will introduce the fundamental concepts of SAR and InSAR, with a focus on the techniques which are relevant to this research. Sec. 2.2 provides a basic introduction to SAR and describes the imaging scenario and geometry, as well as a description of the basic observables. Sec. 2.3 introduces the concept of InSAR, and how information can be extracted by interferometrically combining images. Finally, Sec. 2.4 provides an in-depth look at the techniques and challenges encountered in the state of the art of distributed scatterer InSAR analysis.

2.1 Introduction

There is a long chain of steps to go from a raw set of recorded radar pulses to a focused, complex-valued synthetic aperture radar (SAR) image, and to go from a stack of focused images to a ground displacement estimate with SAR interferometry (InSAR). This chapter will give a brief overview of these steps, starting with the basics of SAR in Sec. 2.2, in particular the imaging geometry, the radar phase observable and scatterer classifications. The basic concepts of InSAR are then introduced in Sec. 2.3, again starting with the geometry of multiple SAR acquisitions, and then moving to a breakdown of the different signal components which are present in an interferometric observation. This is to give the reader a basis on which to contextualize the more in-depth sections of the chapter and is by no means an exhaustive introduction to the topic. For a more detailed description of these concepts, the reader is directed to the following references from which the content of this section is adapted: [33]–[36]. Finally, Sec. 2.4 will introduce the key concepts of distributed scatterer (DS) InSAR, as well as provide the theoretical basis for

the work contained in the subsequent chapters. Interested readers can find more detail on this topic in the following references: [20], [37].

2.2 Synthetic Aperture Radar

This section introduces the basic concepts behind SAR imaging, including the geometry, the radar observables, and scatterer classification, which are the most relevant concepts for understanding InSAR.

2.2.1 Basic Principles of Radar Sensors

Radar imagers such as SARs produce complex-valued images whose intensities appear similar to passive optical images, but the principles behind generating them are fundamentally different. Passive optical sensors (like cameras) use a lens and mirror system to project radiation from a separate light source (like the sun) scattered by objects in the scene onto a two-dimensional array of detectors. This conserves the angular relationships between multiple targets and their images. Imaging radars produce the radiation used to illuminate the scene, and sample the backscattered signal in time to retrieve the range dimension, and their relative Doppler shifts to separate them in the azimuth direction.

2.2.2 SAR Imaging Geometry

A simplified diagram of the SAR imaging geometry is shown in Figure 2.1. The local coordinate system of a side-looking radar sensor (aka. the radar coordinates) is defined by three dimensions: the *slant range*, (also simply called the *range*), the *azimuth*, which is the distance from the radar antenna to the target along the flight path of the radar, and the *cross range*, which is the direction orthogonal to the plane spanned by the range and azimuth directions [38], [39].

Ultimately, the radar measures the delay time and Doppler shift of received echoes, which are used to estimate the distance from itself to targets on the ground along the line of sight (LOS) (i.e. the slant range) and the azimuth direction, respectively. The *ground range* is the projection of the slant range vector onto the horizontal plane. From the perspective of the satellite, the angle between the slant range vector and the nadir direction is called the *look angle* and is denoted by θ. This is closely related to the *incidence angle*, θ_{inc}, which is the angle between the normal vector of the surface at the target location and the zenith. If the surface of the Earth were flat, these would be equal as they would form alternate interior angles. However, due to the curvature of the Earth, they deviate (not depicted).

The radar illuminates a region on the ground called the *footprint*. The extent of the footprint in ground range is called the *swath*, and is an important parameter when designing SAR missions as it determines how large of a region will be observable at any given moment. The edge of the swath closest to the radar is called the *near range*, and the edge furthest from the radar is called *far range*. Note that the incidence angles will be different for different pixels within the footprint depending on their range.

The extent of the footprint in azimuth is called the *azimuth beamwidth*. The beamwidth

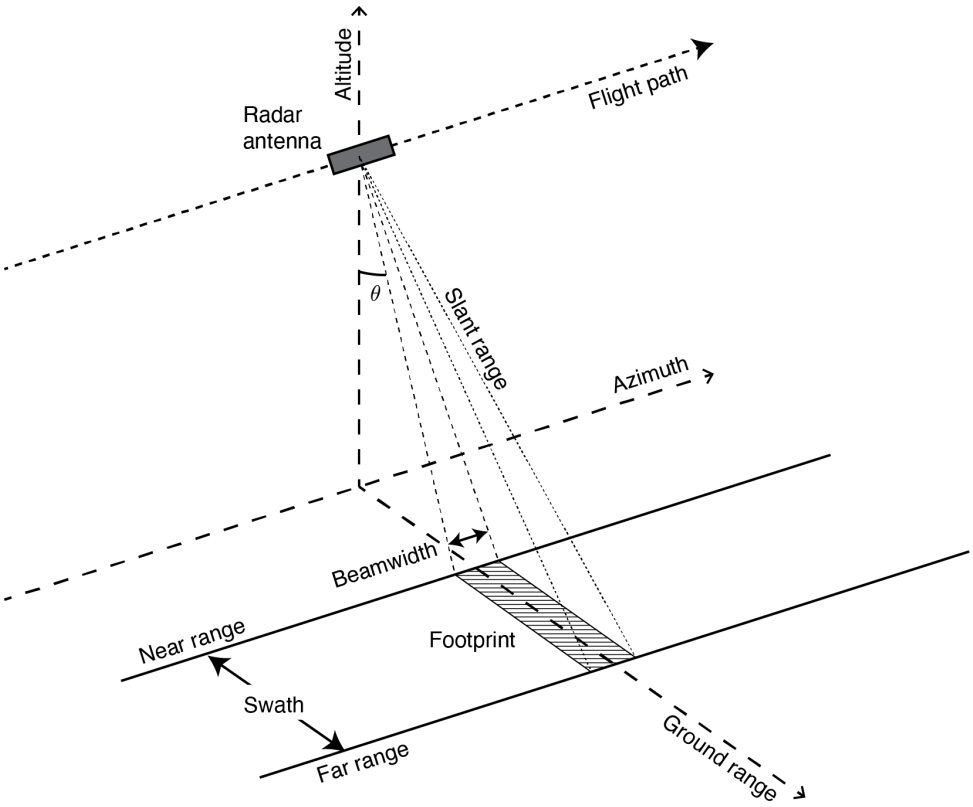

Figure 2.1: Simplified diagram showing a right-looking SAR imaging scenario. Adapted from [33] and [34].

of the system, W_{az} is

$$W_{\text{az}} = \frac{r\lambda}{L_a}, \tag{2.1}$$

where λ is the wavelength of the sensor, r is the slant range to the target, and L_a is the length of the antenna. If not for the innovation of the synthetic aperture, the beamwidth would be the best achievable resolution in the azimuth direction, which is explained in Sec. 2.2.4.

2.2.3 Range Resolution

The resolution of a SAR system is defined by the minimum separation distance needed to differentiate two targets on the ground. If two targets are separated in range by a distance Δr, then their respective echoes will be separated by a time

$$\Delta t = \frac{2\Delta r}{c}, \tag{2.2}$$

where c is the speed of light and the factor of 2 is due to the two-way travel time. The inverse of this time is the range bandwidth, BW_{r}. Thus the best achievable range resolution

of the system is

$$\delta_{\mathrm{r}} = \frac{c}{2BW_r}. \tag{2.3}$$

Note that this value is independent of distance to the target.

2.2.4 Azimuth Resolution and the Synthetic Aperture

The *synthetic aperture* is a powerful technique in which the motion of the radar sensor itself is used to synthetically enhance the effective antenna size, hence its name. Without this innovation, it would be impossible to create useable radar images of the ground from orbit, as the azimuthal resolution degrades with distance between the radar and the target. The azimuth resolution of a real aperture radar (RAR) system, $\delta_{\mathrm{az,rar}}$, is limited by the distance between the radar antenna and the target, such that

$$\delta_{\mathrm{az,rar}} = W_{\mathrm{az}} = \frac{r\lambda}{L_a}, \tag{2.4}$$

where r is the range of the sensor to the target, λ is the wavelength and L_a is the length of the antenna. For a typical C-band sensor in low Earth orbit (LEO) ($\sim$700 km altitude), this would result in a very poor resolution of around four kilometers.

In 1951, Carl Wiley discovered that a radar mounted perpendicular to the vehicle flight path would receive echos that have been Doppler-shifted as a result of its motion [40]. Combining the recorded echos while taking into account their Doppler-shift could allow one to separate the combined responses of targets within the footprint and improve the resolution of a spaceborne system one-thousandfold [41]. This effect can be derived in several ways; one simple way is to consider the geometry of the scene. Successive echoes are recorded as the sensor moves along the flight path, extending the effective length of the antenna, L_{sar}, to

$$L_{\mathrm{sar}} = \frac{r\lambda}{2L_a}, \tag{2.5}$$

where the subscript "sar" denotes synthetic aperture radar. Substituting the effective antenna length into the original equation for azimuth resolution, we obtain

$$\delta_{\mathrm{az,sar}} = \frac{L_a}{2}. \tag{2.6}$$

Now the azimuth resolution no longer depends on the range to the target, as any additional distance to the target is exactly compensated for by an increase in the size of the synthetic aperture.

2.2.5 SAR Observables

The resulting SAR image after range and azimuth focusing is a high-resolution complex-valued image covering a large spatial extent. A given pixel in this image is a complex number in which the amplitude and phase both have a physically interpretable origin. The amplitude is essentially the brightness of the pixel, and corresponds to the amount of radiated energy that is reflected back to the sensor. The phase component of the pixel corresponds to (i) the travel time of the received radar pulse which is a function of the

distance of the target to the antenna and the refractive index of the medium along the propagation path, and (ii) the random summation of the reflections of elementary scatterers within the resolution cell.

The value of an arbitrary pixel, S, in a SAR image is a complex number and is generally written in polar form as

$$S = A \exp(i\psi), \tag{2.7}$$

where A is the amplitude, i is the imaginary unit, and ψ is the phase. Interferometry is based on combining the phase observations of multiple radar acquisitions, and the phase is therefore the observable we will focus on. The value of the phase is a summation of several effects which determine the optical path length the pulse must travel, and ultimately its total travel time. The phase components are represented as [33]

$$\underline{\psi} = \psi_{\text{geom}} + \psi_{\text{scat}} + \psi_{\text{atmo}} + \underline{\psi}_{\text{noise}}, \tag{2.8}$$

where

- $\underline{\psi}$ is the observed phase, where the underline indicates its stochastic nature;
- ψ_{geom} is the geometric phase, the range-dependent phase which is determined by the separation distance between the radar antenna and the target;
- ψ_{scat} is the scattering phase, which is a result of the coherent summation of all the backscattered echos of the elementary scatterers contained within the resolution cell, as well as the effects of volume scattering which determine the effective scattering phase centre of the target;
- ψ_{atmo} is the atmospheric phase, which affects the optical path length the pulse must travel through due to variations in the refractive index of the troposphere and ionosphere. Variations in the temperature, pressure, and humidity will create variations in the refractive index along the path through the troposphere, and variations in free electron density in the ionosphere;
- $\underline{\psi}_{\text{noise}}$ is phase noise caused primarily by thermal noise in the radar transmitter and receiver.

Although the phase has a well understood physical origin, the cumulative effects of these components, particularly the scattering phase, results in a uniformly distributed random observed phase between 0 and 2π radians (360 degrees). The 2π interval on which the phase is defined is arbitrary, and without loss of generality we use the notation of defining the phase on the interval between $-\pi$ (-180 degrees) and $+\pi$ ($+180$ degrees), which is denoted as $[-\pi, \pi)$. For simplicity and readability, the underline is omitted from further equations containing the term $\underline{\psi}$, but it should be understood that it contains both deterministic and stochastic components.

2.2.6 Scatterer Classification

2.2.6.1 Overview

The observed value of a resolution cell in a SAR image is the sum of all the backscattered energy of the objects within that cell, called *elementary scatterers* [33]. Certain objects will reflect a radar pulse back to the sensor more effectively due to their material properties, shape, size, and orientation. Different configurations of elementary scatterers will result in different observed behaviours. It is therefore helpful to categorize these

Table 2.1: Delft SAR Scatterer Taxonomy (After [42])

	Stable Over Time	**Variable Over Time**
Point Target	Continuously coherent point scatterer (CCPS)	Temporarily coherent point scatterer (TCPS)
Distributed Target	Continuously coherent distributed scatterer (CCDS)	Temporarily coherent distributed scatterer (TCDS)

scatterers into groups according to their spatial an temporal characteristics, summarized in Table 2.1 [42].

2.2.6.2 Point Scatterers

Resolution cells containing an object with a high radar cross-section (RCS) will be dominated by the reflections from that object, as shown in Fig. 2.2. As a result, the behaviour of the entire resolution cell can be localized to the point where that object is located, with all other elementary scatterers acting as clutter. Such a scatterer configuration is called a *point scatterer* (PS). Generally, point scatterers (PSs) provide stable, low-noise reflections which are very useful in time series analysis because the geometric components of the scattering phase are relatively easily extracted. This type of analysis is called point scatterer interferometry (PSI) [1] and has been developed into a mature technology over the previous two decades [43], [44]. PSs can be natural or man-made objects, however they are most often formed by reflections from buildings or infrastructure, which makes PSI particularly well-suited to monitoring the built environment, see [45]–[48].

Since PSs are often formed by the reflections of rigid structures, they tend to exhibit a consistent and comparable reflection mechanism over time. If the reflection of a PS is consistent throughout the entire image stack used for a given analysis, it is called a "continuously coherent point scatterer (CCPS)". The term coherent refers to the consistency of the scatterer's reflection, a concept that is expanded upon in Sec. 2.3 and 2.4. However, the reflection of PSs does not always remain consistent over time, for instance due to other objects obscuring the backscatter in certain images, or due to the construction date of the object in question in the case of the built environment. Such a PS is referred to as a "temporarily coherent point scatterer (TCPS)" [42].

2.2.6.3 Distributed Scatterers

Regions that do not contain an object which dominates the backscattered signal of its surroundings are known as *distributed scatterers* (DS), shown if Fig. 2.3. The diffuse reflected signals from these resolution cells tend to have a much lower signal-to-noise ratio (SNR) than a PS. This is apparent in the right-hand side of Fig. 2.3, where the clutter signal of the weak elementary scatterers in Fig. 2.2 have become the primary signal which is observed. Therefore, it is common to coherently average multiple DS pixels together in a process called *multilooking* in order to reduce the noise level of

[1] Note that PSI is also known as Permanent Scatterer Interferometry, or Persistent Scatterer Interferometry, implying a permanently unvarying scattering mechanism, and conflating the point-like nature of the scatterers with this temporal invariability, which are two unrelated concepts. In the Delft taxonomy [42], we prefer to disentangle these concepts.

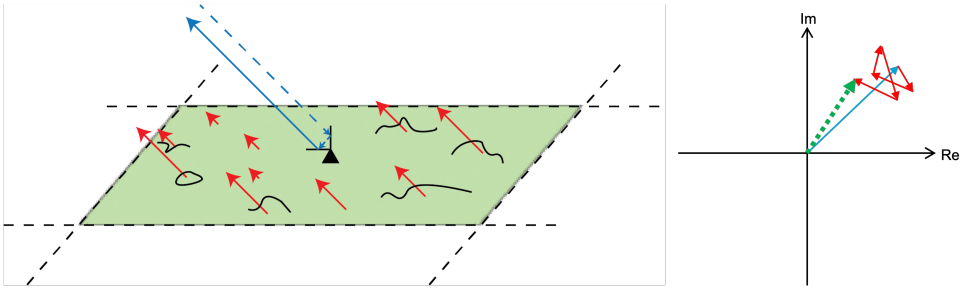

Figure 2.2: Diagram showing the basic configuration of a point scatterer. Left: simplified spatial view of a resolution cell containing elementary scatterers with a dominant scattering object. Right: diagram showing the magnitude and phase of the example resolution cell in the complex plane. Blue: reflected signal of dominant PS. Red: reflected signals of other elementary scatterers acting as clutter. Green: resultant observed signal. After [49].

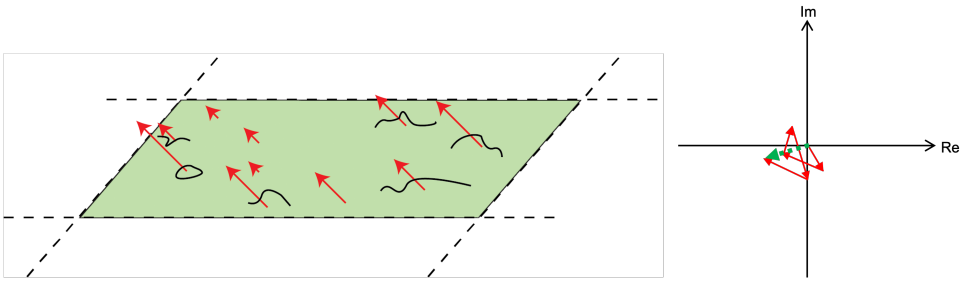

Figure 2.3: Diagram showing the basic configuration of a distributed scatterer. Left: simplified spatial view of a resolution cell containing elementary scatterers without a dominant scattering object. Right: diagram showing the magnitude and phase of the example resolution cell in the complex plane. Red: reflected signals of the elementary scatterers. Green: resultant observed signal. After [49].

the observed signal. In non-arid regions, DSs often do not remain stable over time due to changes in vegetation, human activity, etc. and generally tend to lose coherence over time, adding a significant challenge to using them for monitoring purposes [50]. These are denoted as “temporarily coherent distributed scatterer (TCDS)” [42]. In regions where the ground and land cover remain stable over time such as uninhabited deserts, one may encounter stable DSs, which are denoted as “continuously coherent distributed scatterer (CCDS)”. This work is focused on the use of DS techniques to monitor dynamic regions over time, and Sec. 2.4 goes into detail on the state of the art.

2.2.7 Common SAR Frequency Bands

There are several possible frequency bands in which SAR instruments generally operate. These bands are designated by a letter code and the definition is maintained by the Institute of Electrical and Electronics Engineers (IEEE) [51]. Different frequencies are sensitive to different objects, geometries and processes on the ground. The choice of band of a SAR instrument is dictated by its mission objectives and/or system requirements. In general, higher frequencies provide higher resolution, but will tend to decorrelate faster

Table 2.2: Common SAR Frequency Bands

Designation	L	C	X
Frequency (GHz)	1–2	4–8	8–12
Wavelength (cm)	15–30	3.8–7.5	2.4–3.8
Common uses	Biomass Ice Land subsidence Soil moisture	Compromise between L and X band	High resolution imagery Urban monitoring
Satellite missions	ALOS 1, 2 SAOCOM NISAR (upcoming) ROSE-L (upcoming)	Sentinel-1 Radarsat-2	TerraSAR-X Paz COSMO-SkyMed

(Sec. 2.3.5). Lower frequencies will generally penetrate deeper into solid materials. Table 2.2 provides a summary of the most commonly-used frequency bands and some well known satellite missions. This list is not exhaustive and is provided as background information for the reader.

2.3 SAR Interferometry

This section introduces the concept of SAR interferometry, paying special attention to the geometry of repeat-pass acquisitions and the various components of the interferometric phase observation.

2.3.1 Introduction

As seen in Eq. (2.8), the phase of a given pixel in a SAR image is a combination of several components, resulting in what is essentially a random number uniformly distributed on the interval $[-\pi, \pi)$. In a simplified view, the most important of these components are the geometric phase and the scattering phase. If two images of the same region were to be combined, and the scattering phase components were similar in both images, then combining them would make the scattering components cancel out and would yield a product with a phase value which would be very sensitive to the geometric components of the scene [33].

The concept of coherently combining multiple SAR acquisitions of the same scene in order to glean geometric information about targets located within that scene is known as SAR interferometry (InSAR). In this context, the term *coherent* refers to the fact that the phases of the SAR acquisitions are preserved; therefore a coherent combination of two SAR image pixels can be said to be the *interference* of the two sampled wave returns of the given target. The coherent combination of two SAR phase images is called an *interferogram*, and the phase difference between two SAR pixels is called the *interferometric phase*.

In order for the interferometric phase to contain any useful geometric information, there needs to be a difference between the two SAR images that are combined, either spatial or temporal. This difference is called the *interferometric baseline*, and the choice

of this baseline will affect the type of information that is contained within the interferometric phase.

The two most basic types of InSAR are essentially the result of this choice of baseline, and are called *single-pass* and *repeat-pass*. In single-pass interferometry, there is no temporal baseline, which means that there are two radars attached to the same vehicle (aircraft or spacecraft) acquiring imagery of the same location from different viewing geometries. This technique is commonly employed to produce accurate digital elevation models (DEMs), such as the well known Shuttle Radar Topography Mission (SRTM) DEM [52] which was acquired by the space shuttle Endeavour in the year 2000. In repeat-pass interferometry, there is a temporal baseline, which refers to the fact that some time has passed between SAR acquisitions of the same target. This technique is usually used to monitor changes in a scene or target over time. While a spatial baseline still exists as well, these are usually kept as small as possible, as is the case in the Sentinel-1 mission which as of the time of writing has produced the most commonly used dataset of SAR imagery in the world. Repeat-pass interferometry is therefore a very powerful tool for monitoring geometric changes in a scene over time, and is the focus of this work.

2.3.2 InSAR Geometry

A diagram showing the geometry of a typical repeat-pass interferometric configuration is shown in Fig. 2.4. The target in position P_H is imaged by the observing satellites(s) at positions x_i and x_j. The spatial interferometric baseline, B, is defined as the distance between these positions. The baseline can be decomposed into components $B_\perp$ and $B_\parallel$, which are the perpendicular and parallel baselines respectively. The ranges from the target to the satellite(s) in positions x_i and x_j are r_i and r_j, respectively. The range to the target is often expressed with respect to the reference (or "mother") image, (when the satellite is in position x_i), and is denoted as r without any subscript and simply referred to as "the range". The target is located at a certain height, H. However, without a DEM, the height of the topography is unknown, and the target will appear to be located at position P_0 on the reference ellipsoid, H_0. The look angle, θ, is defined as the angle between the the line-of-sight vector of the satellite in position x_i to the target at position P_0, and the normal vector of the reference ellipsoid. The baseline orientation, α, is the angle between the baseline vector and the direction of the tangent line of the reference ellipsoid [33].

2.3.3 The Interferometric Phase

In order to understand what information is contained within the interferometric phase, we start by defining two SAR acquisitions of a given target, with indices i and j. The value of the two pixels of the two acquisitions of the same target are

$$\begin{aligned} S_i &= A_i \exp(i\psi_i) \\ S_j &= A_j \exp(i\psi_j), \end{aligned} \tag{2.9}$$

where A is the amplitude of the pixel and ψ is the phase. When coherently combining the two SAR images to produce an interferogram, one image is multiplied with the complex conjugate of the other image. The value of a pixel in the resulting interferogram is

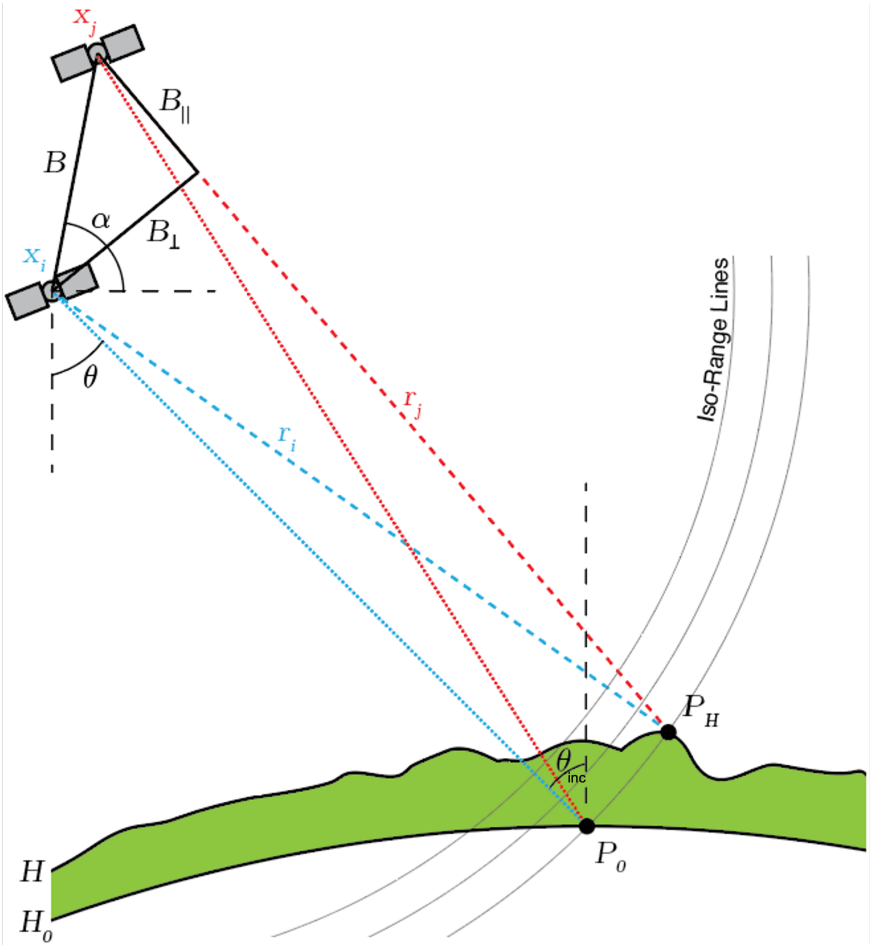

Figure 2.4: Simplified diagram showing the geometry of a repeat-pass interferometric configuration, with acquisitions occurring at locations x_i and x_j. B is the spatial baseline, with perpendicular component $B_{\perp}$ and parallel component $B_{\|}$. The target is located at position P_H with height H, corresponding to the point P_0 on the reference ellipsoid, H_0. θ is the look angle, θ_{inc} is the incidence angle, and α is the baseline orientation angle. The along-track (azimuth) coordinate is into the page. Note that this diagram does not depict any motion of the target. Adapted from [33].

therefore

$$\begin{aligned} I_{ij} &= S_i S_j^* \\ &= A_i A_j \exp\left(i(\psi_i - \psi_j)\right) \\ &= A_i A_j \exp\left(i\phi_{ij}\right), \end{aligned} \tag{2.10}$$

where $\{\cdot\}^*$ denotes conjugation and ϕ_{ij} is the interferometric phase of acquisition j with respect to i. The sign of the interferometric phase depends on which acquisition is conjugated and which is kept as-is. In our notation, this is noted in the order of the interferogram subscript. The unconjugated acquisition in the first position of the subscript is denoted in this text as the reference acquisition, which is also known as the "master" or "mother" image.

The observed interferometric phase $\underline{\phi}_{ij}$ is the summation of the various phase components from Eq. (2.8) of both images comprising the interferogram:

$$\underline{\phi}_{ij} = \phi_{\text{geom}} + \phi_{\text{scat}} + \phi_{\text{atmo}} + \underline{\phi}_{\text{noise}}, \tag{2.11}$$

where ϕ_{geom} is the phase contribution due to differences in geometry, ϕ_{scat} is the phase contribution due to differences in scattering, ϕ_{atmo} is the phase contribution due to differences in atmosphere, and $\underline{\phi}_{\text{noise}}$ is the sum of the noise in both image pixels. The underline indicates the stochastic nature of the variable. The following sections will examine the different components of the interferometric phase in detail.

2.3.4 Interferometric Geometric Phase

This component of the interferometric phase is based on the difference between the geometric components of the two respective acquisitions, and is caused by motion of the target and differences in the orbital position of the satellite(s) occurring between acquisitions. The geometric phase can be written as the sum of three sub-components:

$$\phi_{\text{geom}} = \phi_{\text{disp}} + \phi_{\text{ref}} + \phi_{\text{topo}}, \tag{2.12}$$

where ϕ_{disp} is the displacement phase, ϕ_{ref} is the so-called "reference" phase, and ϕ_{topo} is the topographic phase. This component is therefore very dependent on the interferometric baseline. If we are interested in observing the relative motion of a target over time, which is the focus of this work, then we want to have a non-zero temporal baseline and to minimize the spatial baseline as much as possible.

2.3.4.1 Displacement Phase

The phase contribution of the motion of the target between acquisitions is called the *displacement phase*, and can be written as

$$\begin{aligned} \phi_{\text{disp}} &= \frac{-4\pi}{\lambda}(r_i - r_j) \\ &= \frac{-4\pi}{\lambda}\Delta r_{ij}, \end{aligned} \tag{2.13}$$

where Δr_{ij} is the change in distance from the target to the radar along the line-of-sight between epochs i and j. The scaling factor $2\pi/\lambda$ converts the geometric distance in meters to an angular phase value in radians. There is also an additional factor of 2 which corresponds to the fact that the radar pulse must travel along the range distance twice. This is the signal we are ultimately trying to measure in repeat-pass time series InSAR.

2.3.4.2 Reference Phase

This phase contribution is caused by the differences in the positions of the satellite(s) at epochs i and j, when evaluated for positions on a reference surface, such as a global reference ellipsoid. This is essentially the difference in path length along the line of sight caused by a non-zero spatial baseline, as shown in Fig. 2.4. The reference phase typically appears as a smoothly varying trend in the interferogram. It can be computed with knowledge of the parallel baseline and the reference surface by using the far-field approximation, which assumes that the line-of-sight vectors of acquisitions i and j are parallel [33], [53], resulting in

$$\begin{aligned}\phi_{\text{ref}} &= \frac{4\pi}{\lambda} B \sin(\theta - \alpha) \\ &= \frac{4\pi}{\lambda} B_{\parallel},\end{aligned} \tag{2.14}$$

where $B_{\parallel}$ is the parallel baseline when evaluated for a point on the reference surface.

2.3.4.3 Topographic Phase

This component is due to the effects of topography, i.e., the difference between the actual elevation of the target and the reference surface. Fig. 2.4 shows how a target at point P_H will appear at point P_0 on the reference ellipsoid in the reference image (image i), following the iso-range lines. The overall path length contribution of the topography in both images is

$$\phi_{\text{topo}} = \frac{-4\pi}{\lambda}\Big[\Big(|\vec{x}_i - \vec{P}_H| - |\vec{x}_j - \vec{P}_H|\Big) - \Big(|\vec{x}_i - \vec{P}_0| - |\vec{x}_j - \vec{P}_0|\Big)\Big]. \tag{2.15}$$

Because points P_H and P_0 lie on the same iso-range lines with respect to image i, the two terms containing the ranges with respect to point x_i are in fact equal and cancel each other out, leaving the difference in range with respect to point x_j, giving the equation

$$\phi_{\text{topo}} = \frac{-4\pi}{\lambda}\Big(|\vec{x}_j - \vec{P}_0| - |\vec{x}_j - \vec{P}_H|\Big), \tag{2.16}$$

which is visualized as the difference in length between the red dashed and dotted lines in Fig. 2.4. This equation can be written in terms of the reference image and interferometric baseline using the far-field approximation [33], [53],

$$\begin{aligned}\phi_{\text{topo}} &= \frac{-4\pi}{\lambda} \cdot \frac{B\cos(\theta - \alpha)}{r \sin\theta_{\text{inc}}} \cdot H \\ &= \frac{-4\pi}{\lambda} \cdot \frac{B_{\perp}}{r \sin\theta_{\text{inc}}} \cdot H.\end{aligned} \tag{2.17}$$

In displacement monitoring applications, this term must be estimated and removed using a-priori knowledge of the topography, i.e. a DEM. This data is generally imperfect (ex. additional and unmodelled elevation of structures) and a residual topographic phase component will remain after DEM correction, referred to as the DEM error.

Multi-temporal InSAR can also be used to create DEMs. In this case, Eq. (2.17) becomes the signal of interest and the target is assumed to be stationary. The displacement phase in Eq. (2.13) becomes an error term, which is minimized by averaging subsequent acquisitions together.

2.3.5 Interferometric Scattering Phase

This component of the interferometric phase is based on the difference between the scattering components of a given pixel in the two respective acquisitions, i and j, and essentially acts as an additional noise term, and is written as

$$\phi_{\text{scat}} = \psi_{\text{scat},i} - \psi_{\text{scat},j}. \tag{2.18}$$

If the surface and volume scattering components $\psi_{\text{scat},i}$ and $\psi_{\text{scat},j}$ remain stable over time, then ϕ_{scat} will remain small, i.e., close to zero, hence with minimal effect in Eq. (2.12), even if they form a significant part of the backscattered echo of the individual acquisitions.

These scattering components will inevitably change over time [26], [50], causing the contribution of ϕ_{scat} to become more significant, i.e., increasingly non-zero, until it cannot be disentangled from the geometric component in Eq. (2.12) anymore. This effect is called *temporal decorrelation* [50] and is a significant consideration in any distributed scatterer InSAR analysis, which is covered in depth in Sec. 2.4 and Ch. 5.

2.3.6 Interferometric Atmospheric Phase

This component of the interferometric phase is caused by differences in tropospheric and ionospheric conditions in which the radar signal has travelled during the two acquisitions, i and j, which is written for a certain pixel as

$$\phi_{\text{atmo}} = \psi_{\text{atmo},i} - \psi_{\text{atmo},j}. \tag{2.19}$$

Changes in weather patterns and ionospheric conditions between acquisitions will lead to differences in the refractive index along the LOS between the radar and the target. Because of its relatively smooth spatial characteristics, the atmospheric signal forms a spatially variable atmospheric phase screen (APS) that can be difficult to distinguish from the displacement phase or orbital errors [33]. However, the atmospheric phase can be considered to be uncorrelated between interferograms, and so can be mitigated by filtering once initial estimates for the geometric interferometric phase are obtained for a network of high-quality point scatterers [54]. Because the tropospheric and ionospheric conditions primarily affect the time delay of the radar pulse, and these delays are near-identical over short distances, the atmospheric delays do not affect the interferometric coherence or its estimator, see Sec. 2.4.1.

2.3.7 Double-Difference Arc Phase

Secs. 2.3.3–2.3.6 have shown that there are several components which make up the observed interferometric phase. In order to extract useful information about the displace-

ment phase from the overall observation, the other components must be removed. While satellite orbits and atmospheric conditions can be modeled to some degree of accuracy, it is impossible to estimate the absolute effect of these components in the interferometric phase with the millimetric precision that is required.

Points nearby one another on the ground will share very similar orbital and atmospheric phases. By taking the difference between their respective interferometric phases, i.e., forming an arc between two points, the majority of these components will cancel out, leaving only differences in the displacement, topographic, and scattering phases. The double-difference arc phase of a point, n, with respect to a another point, m, is denoted as

$$\phi_{ij}^{mn} = \phi_{ij}^{m} - \phi_{ij}^{n}, \tag{2.20}$$

where the point m is called the reference point. Because of this, all InSAR-derived displacement products are relative with respect to a reference point. Therefore it is very important that care is taken when selecting a reference point. It is often advantageous to choose a point with a very high SNR, such that the scattering phase components of the point can be assumed to be very low. By using a DEM, the topographic phase can be mostly removed, leaving only the differential displacement phase, as well as the residual atmospheric and orbital phases. These residuals can be removed by using external information, i.e., models [55], or by converting assumptions to filtering operations [33], [54]. Note that in subsequent chapters, the superscript mn is excluded for readability, and it should be understood that all interferometric phases are taken with respect to a certain reference point.

2.4 Distributed Scatterer InSAR

This section will describe the state of the art of distributed scatterer (DS) InSAR techniques and provide the theoretical basis for the work in the subsequent chapters of this thesis, as well as highlight the open research gaps addressed by this thesis. A comprehensive review of the subject of DS processing and analysis is available in [20].

2.4.1 The Coherence Matrix

Coherence is a useful metric for the quality of an interferometric combination. The (complex) coherence of a given interferometric combination of a given pixel is defined as [33], [56]

$$\gamma_{ij} = \frac{E\{S_i S_j^*\}}{\sqrt{E\{|S_i|^2\}E\{|S_j|^2\}}} = |\gamma_{ij}| \exp(i\phi_{0_{ij}}), \tag{2.21}$$

where $E\{\cdot\}$ is the expectation operator. This results in a complex value with a magnitude, $|\gamma_{ij}|$, in the interval [0,1] and a phase value, $\phi_{0_{ij}} \in [-\pi, \pi)$, equal to the expectation of the interferometric phase. The magnitude of the coherence is the normalized covariance between corresponding pixels of two SAR images combined interferometrically. Thus it can be viewed as a measure of the useful information available in a specific interferometric combination. This is because the differences in the images caused by a change in scattering phase of the elementary scatterers also affects the amplitude values. If the scattering phase components are similar in both images, their effects will cancel out in

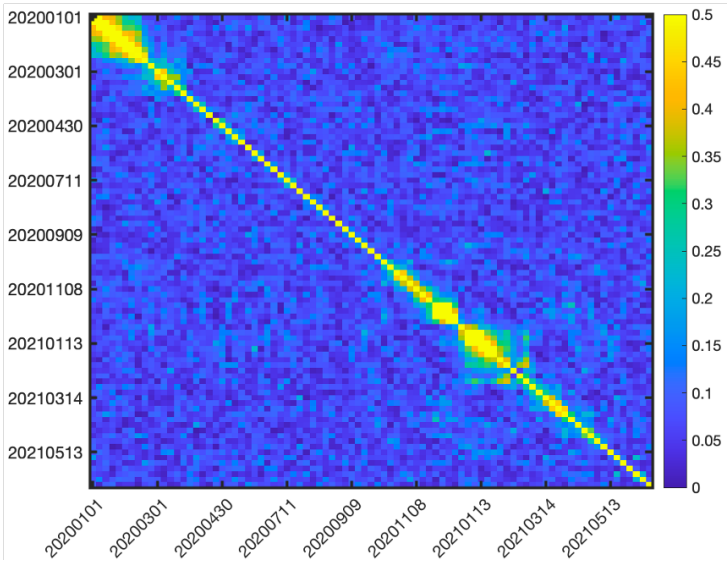

Figure 2.5: Example of an observed (magnitude) coherence matrix, using data acquired by Sentinel-1 ascending track 088 imagery of a grassland DS near Zegveld, the Netherlands. The colour axis is limited to the range [0, 0.5] in order to emphasize the low coherence values.

the interferogram, leaving only the geometric (and atmospheric) phase terms, and the coherence will be high. On the other hand, if the scattering phase components change significantly between images, the effect of this term will dominate and the coherence will be low. Such an interferometric combination is said to be *decorrelated* (Sec. 2.3.4 and 2.3.5).

The coherences of each possible interferometric combination in a stack of N SAR images can be expressed as elements in an $N \times N$ matrix called the *coherence matrix*, Γ, which is defined as

$$\Gamma = \begin{bmatrix} 1 & \gamma_{12} & \cdots & \gamma_{1N} \\ \gamma_{21} & 1 & \cdots & \gamma_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_{N1} & \gamma_{N2} & \cdots & 1 \end{bmatrix}. \tag{2.22}$$

The main diagonal of Γ is always unity, as it corresponds to each SAR image combined with itself.

The definition of coherence in Eq. (2.21) depends on the expectation value, which means that in order to estimate this value with real data a statistical approach is required. The expectation operator in Eq. (2.21) is replaced with the statistical average over a selected group of pixels, Ω, resulting in the (complex) coherence estimator

$$\hat{\gamma}_{ij} = \frac{\sum_{n\in\Omega} S_i^n S_j^{n*}}{\sqrt{\left(\sum_{n\in\Omega} |S_i^n|^2\right)\left(\sum_{n\in\Omega} |S_j^n|^2\right)}}, \tag{2.23}$$

where the $\hat{\cdot}$ notation indicates an estimator and $S^n_{i,j}$ is the value of the nth SAR SLC pixel at epoch i or j. This has multiple implications: 1) the coherence of a single pixel cannot be estimated using this definition (which is why it is not used in PS analysis), and 2) that the statistical average is used to estimate the expected value, making Eq. (2.23) a biased estimator for low numbers of pixels [57]. The process of coherently averaging multiple pixels in an interferometric SAR stack is called *multilooking*, and means that the phase of $\hat{\gamma}_{ij}$ is equal to the multilooked interferometric phase of the pixels in Ω.

Evaluating Eq. (2.23) for every combination of images in the interferometric stack results in the sample coherence matrix, $\hat{\Gamma}$, which is very useful for graphically representing the coherence behaviour of a distributed scatterer over time, an example of which is shown in Fig. 2.5. The matrix is triangular symmetric in magnitude, and the interferometric phases $\phi_{ij} = -\phi_{ji}$. The main diagonal is unity. The matrix shows at a glance how the coherence evolves over time and which interferometric combinations contain useful information. Sec. 2.4.8 discusses how the information contained within the complex coherence matrix can be extracted to estimate the interferometric phase of a DS.

2.4.2 Statistically Homogeneous Pixels

The choice of the set of pixels Ω to be included in the multilooking window is critical, as it determines which observations are to be averaged together and be treated as a single combined observation. Thus, care must be taken to ensure that the pixels chosen for inclusion are representative of the same process occurring on the ground, else the risk of phase inconsistencies becomes significant, as shown in Fig. 2.6. To that end, several procedures for the a-priori testing of the homogeneity of the pixels have been developed [58], [59]. In general, the idea is to test if the amplitudes of the selected pixels are likely to be from the same distribution, and therefore representative of the same process. Thus these pixels are called statistically homogeneous pixels (SHPs).

This idea was introduced by [58], who used a Kolmogorov-Smirnov test to estimate whether or not a given pixel belongs to the same distribution as one another, and is the most commonly used SHP test. This method does however have several shortcomings: the test is computationally expensive and can drastically increase the time needed to process a large dataset. Secondly, just because the amplitudes are estimated to belong to the same distribution, there is no guarantee that the phases do in fact belong together. Finally, the test does not filter out extreme values from the statistically homogeneous distribution (if they do not occur often enough to reject the distribution). For example, in an agricultural context, if a tractor were to be imaged in a given pixel that is otherwise identical to a neighbour, the Kolmogorov-Smirnov test will not filter out the affected pixel containing the tractor, despite the fact that that pixel phase clearly has a completely uncorrelated geometric component in the affected epoch. Conversely, if a pixel is excluded because enough epochs are affected by some inhomogeneity, the pixel is removed for all epochs, even those which may contain useful information. Other methods, such as the Anderson-Darling test [60] can improve on the sensitivity to the tails of the distribution, yet the other drawbacks of the methodology remain. Thus the interpretation of the results of such a test, and indeed the entire analysis, depends on these underlying assumptions.

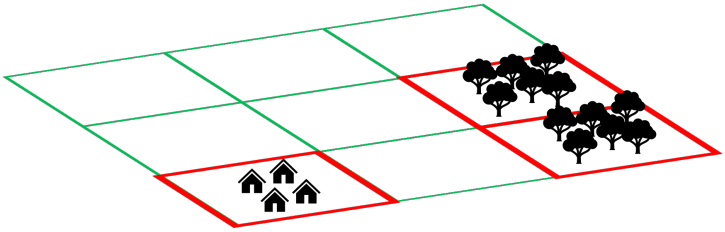

Figure 2.6: Simplified diagram of statistically homogeneous pixels (green). Pixels containing reflections from unwanted/ unrepresentative objects on the ground (red) can introduce different components into the geometric and scattering phases, leading to erroneous interpretations. Adapted from [20].

2.4.3 Pixels versus Looks

Multilooking is defined as the combination of multiple resolution cells together, however, functionally speaking we have access to pixels, not the resolution cells. The size of a resolution cell is a parameter which depends on the SAR system, see Sec. 2.2, whereas the pixel size is determined by the digital sampling (or posting) of the obtained radar data in azimuth and range. It is common practice to oversample the radar data by a small degree, which means that the pixel size is smaller than the resolution cell sizes, and that the pixels are not completely independent samples when considering the statistical quantities of the DS. Thus we use the term *effective number of looks* to refer to the number of independent samples which is derived from the number of pixels [33]. The effective number of looks, L, is

$$L = \frac{N}{\text{OSR}}, \tag{2.24}$$

where N is the number of pixels and OSR is the oversampling ratio given by

$$\text{OSR} = \frac{\text{PRF}}{\text{BW}_{\text{az}}} \cdot \frac{f_s}{\text{BW}_{\text{r}}}, \tag{2.25}$$

where PRF is the pulse repetition frequency, f_s is the range sampling frequency, and BW_{az} and BW_{r} are the azimuth and range bandwidths, respectively.

2.4.4 Coherence Estimation Bias

The coherence estimator in Eq. (2.23) is biased for low values of coherence and for a low effective number of looks, as reported by [57], which gives an expected value of the coherence estimator as

$$E\{|\hat{\gamma}|\} = \frac{\Gamma(L)\Gamma(3/2)}{\Gamma(L+1/2)} \cdot {}_3F_2(3/2, L, L; L+1/2, 1; |\gamma|^2)(1-|\gamma|^2)^L, \tag{2.26}$$

where $\Gamma(\cdot)$ is the gamma function and ${}_pF_q(\cdot)$ is the generalized hypergeometric function. Evaluating this expression for varying levels of coherence and number of looks one obtains the plot shown in Fig. 2.7. This shows the risk of interpreting weakly coherent combinations with a low degree of multilooking, as it is likely that the coherence is significantly overestimated.

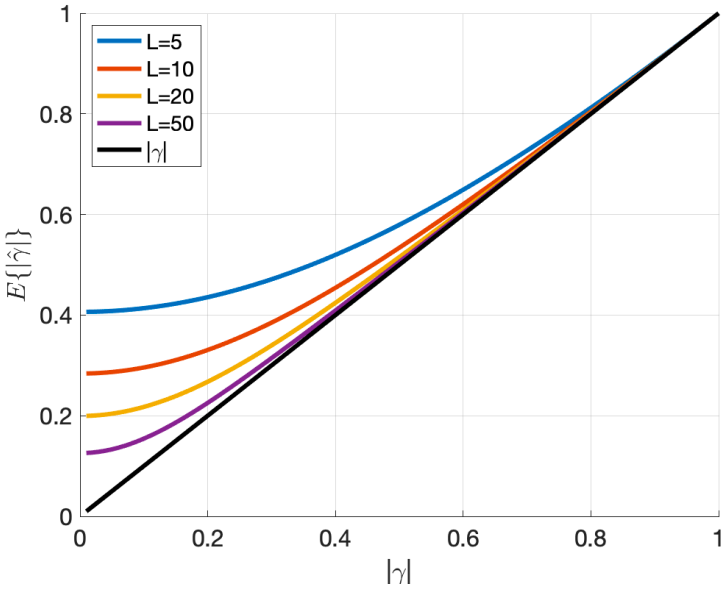

Figure 2.7: True versus estimated coherence evaluated for several degrees of multilooking according to Eq. (2.26). After [61].

2.4.5 Effect of Coherence on Phase Distribution

The stochastic components of a SAR acquisition are typically modelled as a complex Gaussian process. Thus, for a given SAR acquisition, the amplitude component is Rayleigh distributed, and the phase component is uniformly distributed on the interval $[-\pi, \pi)$. This is no longer the case for an interferogram, or interferometric combination of a subset of pixels in the case of correlation between the SAR acquisitions. The phase distribution of a given interferometric combination is described by a probability density function (PDF) as a function of the coherence and the effective number of looks, L, see [33, Eq. (4.2.24)], and [62], [63]

$$f(\phi|\gamma, L, \phi_0) = \frac{\Gamma(L+1/2)(1-\gamma^2)^L\beta}{2\sqrt{\pi}\Gamma(L)(1-\beta^2)^{L+1/2}} + \frac{(1-\gamma^2)^L}{2\pi} \cdot {}_2F_1(L, 1; \frac{1}{2}; \beta^2), \qquad (2.27)$$

where ϕ_0 is the mean or expected interferometric phase and $\beta = |\gamma| \cos(\phi - \phi_0)$. The distribution is plotted for varying levels of coherence and effective number of looks in Fig. 2.8. As the coherence and effective number of looks increase, the distribution of the phase becomes more localized. At zero coherence, the distribution becomes uniform, indicating complete information loss, and at unity coherence the distribution becomes a delta function, indicating that the two combined images are identical, and the interferometric phase is zero (not shown). It is important to note that Eq. (2.27) assumes that each look is an independent sample of the same process. In reality, each resolution cell on the ground will be subject to different effects and processes which affect the geometric and scattering phases differently.

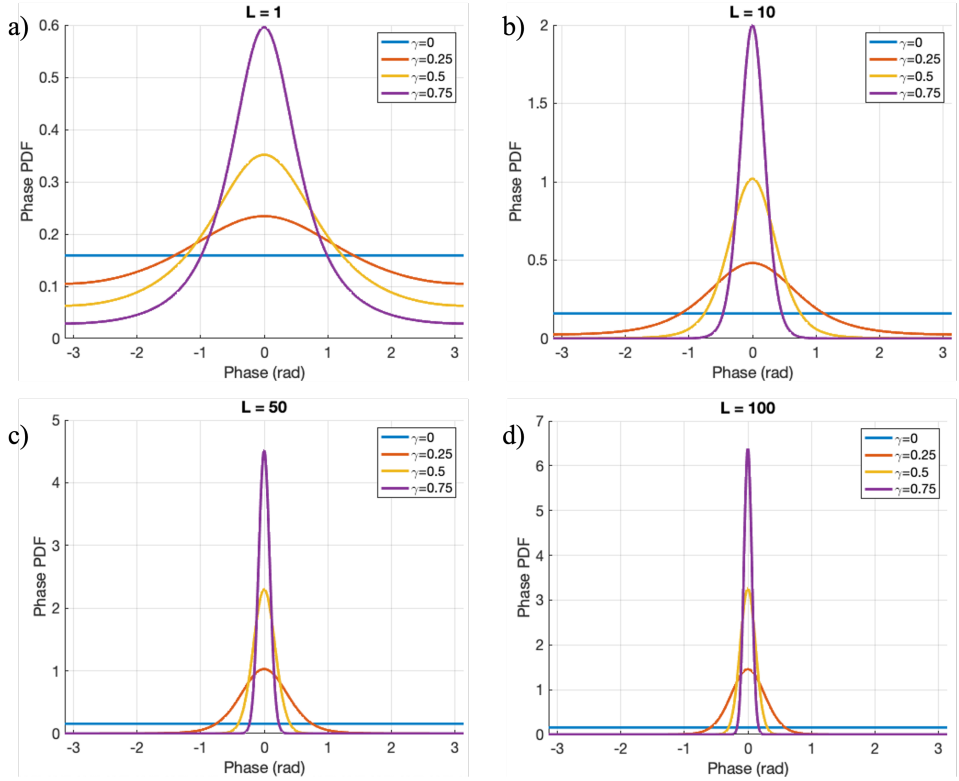

Figure 2.8: Phase probability density functions obtained by evaluating Eq. (2.27) assuming $\phi_0 = 0$ with varying levels of coherence and effective number of looks. a): L=1, b): L=10, c): L=50, d): L=100. Adapted from [33].

2.4.6 Stochastic Model for Distributed Scatterers

Eq. (2.27) can be evaluated to determine the expected phase standard deviation, σ_ϕ, for a given level of coherence and effective number of looks, see [33, Eq. 4.2.27]. Thus one can obtain an estimate for the level of precision of an interferometric combination of a given distributed scatterer, shown in Fig. 2.9. This allows for an uncertainty estimate of the interferometric observation and any derived quantities using standard error propagation methods (APS and phase unwrapping errors notwithstanding) [20], [33].

Fig. 2.9 clearly shows the effect of multilooking on the obtainable precision of an interferometric observation. Under the assumption of representativity, increasing the effective number of looks can drastically improve the quality of the observed DS signal. However, the diminishing returns of increasing multilooking are also visible: moving from $L = 50$ to $L = 100$ provides only a marginal uncertainty reduction despite the very large increase in the number samples used. It should also be noted that at low coherence, no amount of multilooking can improve the signal quality, as σ_ϕ converges to $2\pi/\sqrt{12} \approx 1.8138$ rad at zero coherence, which is the standard deviation of a uniform distribution on the interval $[-\pi, \pi)$. Further limitations of multilooking are discussed in Sec. 2.4.7.

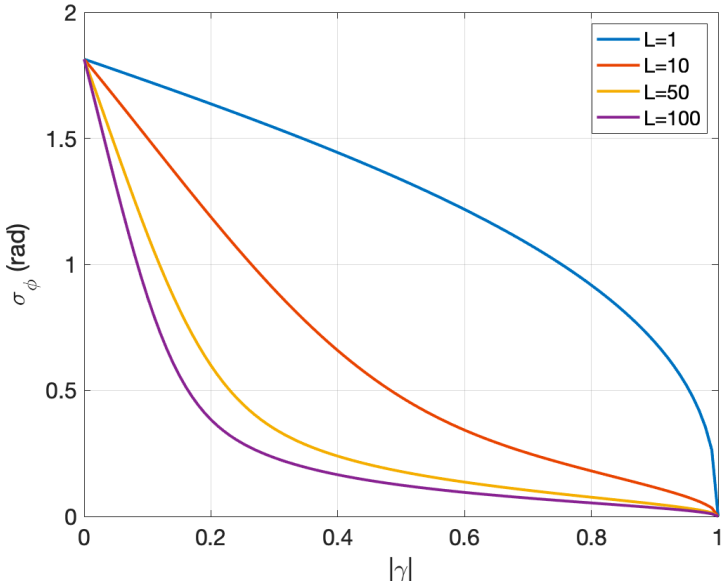

Figure 2.9: Phase standard deviation versus coherence by evaluating the standard deviation of Eq. (2.27) with varying levels of coherence and effective number of looks. After [64].

2.4.7 Phase Triangularity

The concept of phase triangularity (sometimes called phase consistency) states the fact that the interferometric phase combinations of a given pixel must be self-consistent, that is, for the epochs i, j and k,

$$\phi_{ij} + \phi_{jk} - \phi_{ik} = 0. \tag{2.28}$$

While this relation holds for a single pixel, it may no longer be valid once multilooking is applied, and the right-hand side of Eq. (2.28) becomes non-zero, which is known as the closure phase [20], [65], [66]. Closure phase can be caused by multilooking random noise, or by systematic components in the scattering phase term. A basic statistical description of closure phase arising from skewed volume scattering profiles was developed in [65] as

$$\phi_{ijk} \approx \frac{-1}{2} E\Big\{(f(z) - \mu_f)^3\Big\} k_{z_{ij}} k_{z_{jk}} k_{z_{ki}}, \tag{2.29}$$

where ϕ_{ijk} is the non-zero closure phase term that arises from Eq. (2.28) in the presence of multilooking, $f(z)$ is the vertical scattering profile, μ_f is the mean of $f(z)$, and k_z is the vertical wavenumber. Despite this description, it is unclear how this term affects a given observed interferometric phase, ϕ_{ij}, see Eq. (2.11), however, it has been shown that closure phases can bias interferometric time series estimates [67].

This brief summary reflects the field's current state of understanding (or lack thereof) of this phenomenon. Current research is mainly focused on extracting useful information from observed closure phases and relating it to processes occurring within the scene, such as crop growth [68] or soil moisture changes [69]. Despite this lack of clarity, the risk that such a term can bias interferometric time series estimates should inform the choice of

how redundant interferometric combinations within the coherence matrix are interpreted, which is elaborated in Sec. 2.4.8.

2.4.8 Consistent Phase Estimation

2.4.8.1 Basic Approaches

The coherence matrix contains the phases of every possible combination of SAR images in the stack, however, much of this information is redundant and we would like to have only one phase per epoch in order to interpret the motion of the object under observation. There are multiple approaches to accomplish this [25], [29], [44], [58], depicted in Fig. 2.10.

Methods for PSs [44], which remain coherent over time, simply take the difference in phase between each epoch and a single designated reference epoch, which is called the "single mother stack" or "single master stack". This is equivalent to looking at a single row/column in the coherence matrix.

This is not possible for DS applications due to the effects of temporal decorrelation [25], [26]. A few epochs before or after the reference image, the obtained interferometric combinations may be completely decorrelated. This effect is visible in Fig. 2.5. The simplest method for DS phase estimation is to take the "daisy-chain" of phases, that is, to take the difference in phase between each consecutive acquisition, which corresponds to looking at the first off-diagonal of the coherence matrix. This differential phase can then be integrated with respect to the reference epoch to obtain the so-called "Equivalent single master (ESM)", or "Equivalent Single Mother".

2.4.8.2 Small Baseline Subset Approaches

Since DS-based observations are sensitive to temporal decorrelation, one of the first methods to build upon the daisy-chain concept was to use all of, but only those, interferometric combinations which contain coherent information. This is known as the small baseline subset (SBAS) technique [25]. At this time (in the early 2000s), SAR missions tended to have larger perpendicular baselines, so baseline decorrelation was also a concern. Now, in the Sentinel-1 era of consistently small perpendicular baselines, SBAS generally refers only to picking coherent temporal interferometric combinations.

SBAS requires the user to first spatially unwrap each selected small-baseline interferogram before inverting the ESM phase time series. The ESM phases are found by first writing the vector of the subset of used small baseline phases as different combinations of the ESM phases. In a stack of N SAR acquisitions, there are $N(N-1)/2$ possible interferometric combinations. The vector of coherent, small baseline subset of these combinations is denoted as ϕ_{sbas}, and we can write

$$A\phi_{\text{esm}} = \phi_{\text{sbas}}, \tag{2.30}$$

where A is a design matrix that indicates which combination of ESM phases is used to obtain the subset of observed phases. A consistent set of ESM phases which satisfies Eq. (2.30) can then be estimated in the least-squares sense by

$$\hat{\phi}_{\text{esm}} = (A^T A)^{-1} A^T \phi_{\text{sbas}}, \tag{2.31}$$

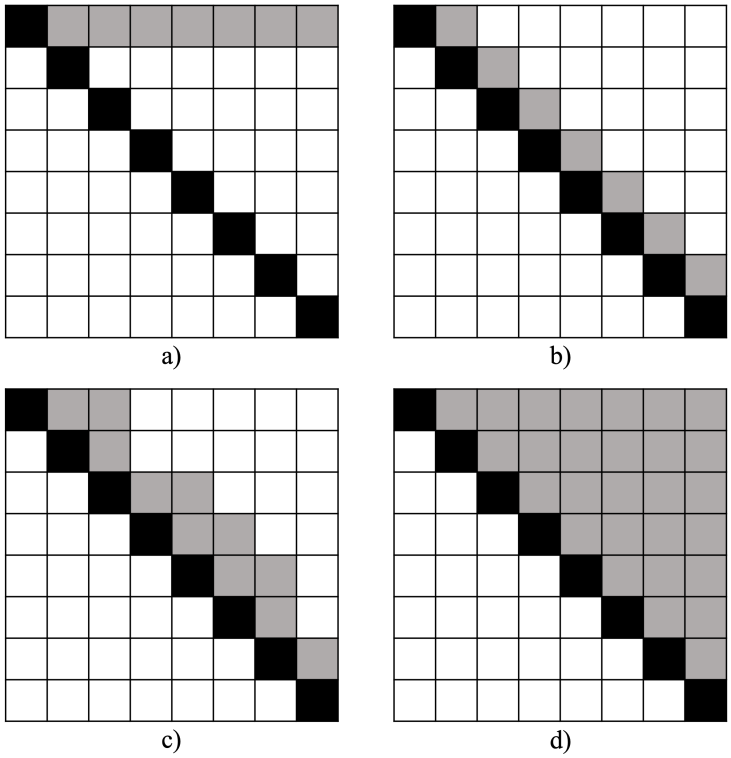

Figure 2.10: Canonical coherence matrices depicting various consistent phase estimation approaches: a) single-reference (PS) b) daisy-chain c) SBAS d) full-rank. Black squares: unity, grey squares: used interferometric combinations, white squares: unused interferometric combinations. Adapted from [37].

where $\{\cdot\}^{\mathrm{T}}$ indicates the matrix transpose. Depending on the choice of interferometric subsets, the term $A^T A$ may be rank deficient and not invertible. To solve this problem, the authors of [25] use the singular value decomposition (SVD) to find the pseudoinverse of $A^T A$, corresponding to the L_2-norm solution. There are several versions of SBAS-based processing algorithms, however, all of them share the common methodology of first spatially unwrapping the selected interferograms and subsequently computing the pseudoinverse of Eq. (2.31).

SBAS is a computationally robust method which will always find a solution given a choice of interferogram subsets, and is the most commonly used DS processing methodology. However, due to the phase triangularity problem (Sec. 2.4.7), the method is sensitive to the choice of interferometric subsets used. Once multilooking is applied, the different interferometric combinations are no longer self-consistent, and it has been shown that SBAS will provide biased solutions in cases where there are significant phase closures [67].

2.4.8.3 Full-Rank Approaches

An alternative school of thought is use all $N(N-1)/2$ possible interferometric combinations. This guarantees an overdetermined system of equations to be solved, and mitigates the risk of biasing the solution by excluding certain interferometric combinations in the event of significant closure phases, as is the case in SBAS processing [67]. There are several ways in which the full coherence matrix is reduced, using either maximum-likelihood estimation [58], eigenvalue decomposition, or a combination of both [29]. Unfortunately these full-rank methods have been given a number of unclear names such as "phase triangulation" or "phase linking". The state-of-the art full-rank estimator is the "Eigendecomposition based maximum likelihood estimator of interferometric phase (EMI)" method proposed in [29], which is used later in this work. This text will simply use the term ESM phase estimation when referring to this step.

The EMI procedure reduces the full set of all interferometric combinations to a single set of consistent ESM phases, $\hat{\phi}_{\text{esm}}$, as estimated by the phase of the minimum eigenvector of the Hadamard product of the inverse of the sample coherence matrix with the complex sample coherence matrix, as given by

$$(|\hat{\Gamma}|^{-1} \circ \hat{\Gamma})\xi = \lambda\xi, \tag{2.32}$$

where $\circ$ denotes the Hadamard product, λ is the eigenvalue, and ξ is the eigenvector. The estimated ESM phases are given by

$$\hat{\phi}_{\text{esm}} = \arg\{\xi\}. \tag{2.33}$$

One of the major advantages of the EMI method is that it works on the wrapped phase observations, allowing for an ESM phase estimation before any phase unwrapping, which is generally the most error-prone step in InSAR processing. Strong decorrelation can hinder the effectiveness of the EMI estimation, but it is unclear to what extent this has an effect, and whether or not the full stack of interferograms should always be used even in the case of complete decorrelation.

2.4.9 Phase Unwrapping

2.4.9.1 Overview

Interferometric phase is an angular measurement bounded by the interval $[-\pi, \pi)$, where the linear LOS displacements are mapped to the phase following Eq. (2.13). Displacement histories exceeding these bounds are wrapped back onto the interval. Stated plainly, an angular measurement cannot measure displacements larger than a full circle (2π, 360°), which corresponds to the linear distance $\lambda/2$. Therefore, there exists an unknown integer number of full phase cycles in each phase observation, called the *ambiguity* [33]. In order to go from a set of angular phase measurements to linear displacements, these ambiguities must first be estimated, which is called *phase unwrapping* or *ambiguity resolution* [70].

There are several different approaches to phase unwrapping, which will not be covered exhaustively by this summary. Two of the most common approaches are to either spatially unwrap interferograms, and then estimate the overall time series change over all sets of multilooked pixels, which is the most common approach in DS processing; or to assume a deformation model or constraint in time, and unwrap the phases according this

model, and subsequently apply spatial checks with neighbouring pixels to ensure a spatially continuous solution, which is the most commonly used technique in PS processing. In both cases, phase unwrapping must make assumptions about the motion of the scatterer in order to produce a solution, because it is an "ill-posed" problem, i.e., there are an infinite number of solutions which satisfy the mathematical constraints. This means that the unwrapped phase is no longer a measurement, but an estimated value based on certain underlying assumptions. If these assumptions do not accurately reflect the true motion phenomena, then the correct displacement time series cannot be inverted from the wrapped phase observations. Therefore, some knowledge (in the form of smoothness constraints or parametric model) of the displacement phenomenon is needed a-priori in order to make the correct assumptions.

2.4.9.2 Spatial Phase Unwrapping

In many DS applications, interferograms have smoothly-varying phases. In such cases, the overall unwrapped phase can be estimated by spatially integrating the wrapped phases in the interferogram. The most commonly used method for this is SNAPHU [71]. This method is based on the "minimum cost flow (MCF)" algorithm [72], which examines the phase residues of loops of pixels/multilooked cells in order to find inconsistencies in the unwrapped result.

Despite their widespread applicability, MCF-based algorithms have some significant drawbacks. First, by their spatial nature, these algorithms assume a spatial smoothness constraint between adjacent pixels, however, in highly discontinuous environments this assumption may not always be true. In the canalized agricultural regions of the Netherlands, there exist numerous sharp discontinuities in the landscape and water level heights. Secondly, while capable of correctly unwrapping spatial phase gradients larger than π, previous studies have shown that the approach fails to catch very rapid motions in highly dynamic and noisy environments such as peatlands when integrating the cumulative unwrapped phases over time, which results in extreme overestimation of the true displacement rate [19].

2.4.9.3 Temporal Phase Unwrapping

Unwrapping the temporal evolution of the phase of each individual scatterer is most commonly done in PS processing methods, however, it can also be applied in the DS case if a set of ESM phases is available. In this case, a parametric model for the expected motion of the scatterer is required, called the *functional model*. The fit of the phase is tested against the model using a temporal coherence estimator a-priori, and/or against other unwrapped scatterers using smoothness constraints a-posteriori [54]. The temporal coherence (also known as the ensemble coherence) estimator of a scatterer is defined as

$$\hat{\gamma}_{\text{temp}} = \frac{1}{N}\left|\sum_{n=1}^{N} \exp\left(i[\phi(t_n) - \phi_{\text{M}}(\hat{x}, t_n)]\right)\right|, \tag{2.34}$$

where N is the number of interferometric phase observations and ϕ_{M} is the modeled phase. This is a useful metric for model fitting, because it can be applied on the wrapped phases. Several methods exist for estimating the most likely integer ambiguity for each epoch, such as integer least-squares (ILS) or integer bootstrapping [70], [73], [74].

This method is primarily used for PS analysis, where the observed points usually lie on rather stable objects, such as in the built environment. Consequently, the functional model usually has a simple form, such as a model with linear and seasonal components. For more complex phenomena, such a simple description may not be adequate to accurately capture the dynamics contained in the wrapped phase observations. On the other hand, a functional model should not contain too many unknown parameters, or else fitting can become challenging as it may be too difficult to constrain the unknowns, or a model may be overfit to the data.

2.4.9.4 Implicit vs. Explicit Unwrapping

It is also useful to distinguish between so-called "implicit" and "explicit" phase unwrapping methods. The term explicit refers to a method which attempts to directly estimate the most likely ambiguity level, which allows one to obtain an unwrapped phase time series, from which the ground displacement can be inferred. SNAPHU is an example of an explicit method. Alternatively, implicit methods try to find the most likely displacement model (the functional model), which is fit to the wrapped phase observations. The estimated displacement model then implies a certain ambiguity as a consequence of that model. There are advantages and disadvantages to either approach, and the target/signal of interest as well as the SAR sensor should be taken into account when deciding which method to use. For example, in the case of peatlands which are of interest in this work, the SNAPHU method will be more reliable and robust when used on L-band data than C-band, due to the fact that the signal dynamic range (approx 5–10 cm) is much larger than the C-band wavelength (5.5 cm), but smaller than the L-band wavelength (24 cm).

2.5 Summary

This chapter has provided a brief overview of the relevant background of InSAR as it pertains to the monitoring of shallow-based ground motion. A description of SAR imaging geometry is provided along with an outline of basic scatterer designations in the Delft taxonomy. The concept of interferometry is introduced, and the various terms in the interferometric phase are described. Finally, topics surrounding the quality of DS InSAR are discussed, with focus on the effect of coherence on the expected phase distribution, how to combine phase observations into one time series of consistent phases, and finally phase unwrapping.

The following chapters will address some of the open issues in the field, in particular those pertaining to the monitoring of dynamic land surface motion of rapidly decorrelating areas. This includes the development of models for the expected motion of the land surface in Ch. 3, temporal phase unwrapping of dynamic and noisy signals in Ch. 4, the reconstruction of interferometric time series estimates which have been split by temporal decorrelation in Ch. 5, and finally, scaling the analysis to a large region over the entire archive of Sentinel-1 imagery in Ch. 6.

3

Modeling the Motion of Soils

Sometimes science is more art than science.
A lot of people don't get that.

Rick and Morty

This chapter describes the development of two models for the motion of shallow soft soils, in particular the motion of deformable soils such as peat and clay in the Vadose zone. One model is based on machine learning techniques, and the other is based on a simple parameterization. Both models use meteorological data as input for making predictions about the relative position of the surface height over time, described in Sec. 3.2. Sections 3.3 and 3.4 describe the machine learning and parametric model implementation and parameterization respectively. The performance of both models is shown and discussed in Sec. 3.5.

Highlights:

1. *The feasibility of employing simple models for the kinematic behaviour of shallow soil motion is demonstrated.*
2. *Precipitation and evapotranspiration are shown to be the dominant drivers of shallow soil motion.*
3. *The effects of climate stresses are visible in the irreversible subsidence component.*

Parts of this chapter have been published in *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 60, No. 5234611 [75]; and *Geoderma* Vol. 440, No. 116699 [76].

3.1 Introduction

3.1.1 Motivation

Although InSAR is capable of producing measurements of ground displacement with precision on the order of millimetres, as we have seen in Ch. 2, the technique provides an underdetermined set of observations with ambiguities that must first be resolved in order to produce a useable time series. This inevitably leads to the need for making assumptions about the most likely relative displacement rate, either spatial or temporal. A major pillar of this thesis is that the quality of InSAR as a technique can be significantly improved when these assumptions are made explicit and when they are made "smarter"; i.e. suited to the environment of study, as opposed to making blanket assumptions without consideration of the processes being monitored Accordingly, a significant part of this research was put towards creating models for the expected displacement of shallow soft soils. This chapter details the development of two models for shallow soft soil displacement, a machine learning model and a parametric model, which are later used in the InSAR processing steps detailed in Ch. 4 and 5.

3.1.2 Model Development Philosophy

While it is clear that shallow soil movements are caused primarily by changes in phreatic groundwater level [30], [31], modelling the expected motion of soft soils such as peat and clay is an ongoing effort and generally involves accurately parameterizing the material and hydrological properties of every layer within the modeled soil strata. Most studies relating soil subsidence to groundwater focus on the effects of extraction from deep confined aquifers, and/or on the effects of settlement in an urban context, for example: [77]–[79]. While modelling unconfined (phreatic) subsurface groundwater is an ongoing effort in The Netherlands [80]–[83], the relationship between the shallow groundwater system and the corresponding soil displacement is not well understood, as so far the focus has been on studying the effect of phreatic groundwater levels on greenhouse gas emissions, and because in-situ measurements of the phenomenon with adequate temporal sampling were not available until recently [13].

Despite these past efforts, there is also a need to be able to describe and predict the motion of these soils simply, with as few model parameters or assumptions as possible. This is particularly important in inverse problems such as those encountered in InSAR, in which one has only one ambiguous observable per location and measurement epoch, rendering any highly multivariate model too complex for inversion, as the problems are too unconstrained to be solved without making many assumptions. This was the motivation to develop a soil motion model with the following requirements:

1. The model should have as few parameters as possible (in the case of a parametric model).
2. All input data should be readily available.
3. The model should be accurate, allowing for minor variations caused by higher-order effects which are not captured.
4. The model should be validated at all available test locations.

The goal is to create an empirical model for observed shallow soil displacement, in first approximation, rather than a complete description of all processes occurring in the shallow

subsurface. Obviously, potential anthropogenic interventions in the water management would not be covered. This resulted in the development of two models: a machine learning model, and a parametric model which was named Simple Parameterization for the Motion of Soils (SPAMS). Both models use the same input data and produce estimates of the same quantity; they differ only in their parameterizations. We will first look at what quantities are used for the input of the models in Sec. 3.2, then each model will be described in detail in Sections 3.3 and 3.4.

3.2 Model Inputs

To simplify the problem, only the two most dominant drivers of soil movement are considered: precipitation and evapotranspiration [30]. These values respectively constitute the primary source and sink of groundwater in the shallow Holocene. In the the first iteration of the machine learning approach [75], temperature and day of year (DOY) inputs were used instead of evapotranspiration. The use of evapotranspiration is a refinement which captures the effects of temperature and seasonality.

The value for evapotranspiration reported by the Royal Netherlands Meteorological Institute (KNMI) is the so-called "De Bruin-Makkink" reference evapotranspiration [84], [85]. This model is applicable to grasslands in the Netherlands, and can be rescaled to model the effects of other vegetation types [86]. A major advantage of the model is that only two easily obtainable input values are required: average daily temperature and daily solar radiant exposure. The De Bruin-Makkink model for reference evapotranspiration, E, is given by

$$E = C \cdot \frac{s}{s+p} \cdot \frac{S_{\text{day}}}{\Delta H_v \rho_w}, \tag{3.1}$$

where C is a constant equal to 0.65 [84], s is the first derivative of the saturation water vapour pressure with respect to temperature (kPa/°C) as given by [87], p is the psychrometric constant (kPa/°C), S_{day} is the daily solar radiant exposure (J/m^2), ΔH_v is the latent heat of vaporization of water (J/kg), and ρ_w is the bulk density of water (i.e. 1000 kg/m^3). s is a function of the mean daily temperature, T_{day} (°C), and is given by

$$s = \frac{7.5 \cdot 237.3}{(237.3 + T_{\text{day}})^2} \cdot \ln 10 \cdot 0.6107 \cdot 10^{\frac{7.5 \cdot T_{\text{day}}}{(237.3 + T_{\text{day}})}}, \tag{3.2}$$

the psychometric constant, p, is also a function of mean daily temperature and is given by

$$p = 0.0646 + 6 \cdot 10^{-5} \cdot T_{\text{day}}, \tag{3.3}$$

as well as the latent heat of vaporization of water:

$$\Delta H_v = (2501 - 2.375 \cdot T_{\text{day}}) \cdot 1000. \tag{3.4}$$

For more information about these quantities, the reader is referred to [84].

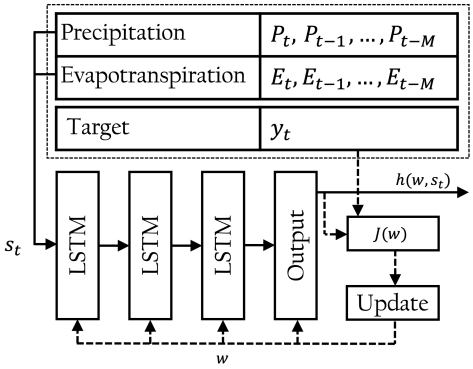

Figure 3.1: RNN flow diagram with an instance of a training data sequence s_t at epoch t. The feedforward prediction path is shown with solid lines, and the feedback path used for training the network is shown with dashed lines.

3.3 Machine Learning Model

3.3.1 Basic Architecture

A recurrent neural network (RNN) is optimized for the task of processing sequences of input data in time and predicting an output time series, and is therefore a suitable choice for the goal of modeling soil motion based on input sequences of meteorological data [88].

The RNN is shown in Fig. 3.1 and is comprised of an input layer, three long short-term memory (LSTM) hidden layers, and an output layer. LSTM layers are chosen because of their ability to use both current and past values of a time series in computation. Multiple layers are used in order to capture the nonlinear relationships between the input sequences and the target value. The outputs of each layer are renormalized, and a 10% dropout is added to help prevent overfitting [89]. The RNN is created using the TensorFlow Python library. An instance of the training set (also pictured in Fig. 3.1) is a sequence of daily precipitation and evapotranspiration values starting from the day in question (the current epoch, t) going back to the M^{th} previous day.

There are two options for the configuration of the output node:

1. A 1xn size node, providing a categorical prediction about the state of the height at time t relative a previous epoch t_{-m}.
2. A 1x1 size node providing a scalar, real-valued output corresponding to the estimated height at time t.

These options are elaborated in Secs. 3.3.2 and 3.3.3.

3.3.2 Categorical Prediction

In this setup, the objective of the model is not to reproduce a time series of surface elevation, but rather to predict whether or not the surface level has moved upwards or downwards with respect to a previous epoch. This concept is helpful for SAR interferometry (InSAR) phase unwrapping, and is covered in depth in Ch. 4.

The input data, is a sequence of daily precipitation and evapotranspiration (alternatively DOY and temperature) values starting from the day in question (the current epoch,

m) going back to the N^{th} previous day. However, now these sequences are coupled with one integer target value, corresponding to the relative state of the elevation with respect to a previous epoch, six days prior to match the Sentinel-1 revisit time. This means we have an estimation of whether or not the surface level shifted upwards, downwards, or did not move significantly between SAR image acquisitions. An integer label of 0, 1, or 2 is assigned to the target value y_t corresponding to either insignificant, upward, or downward motion respectively.

A set of M training sequences is used to train the network, which takes the precipitation and temperature inputs and attempts to predict the target value at the output. Error in the predictions is quantified through the categorical cross-entropy loss function J, given in [90] as

$$J(w) = -\frac{1}{M}\sum_{t=1}^{M}\sum_{j=1}^{C} y_{t,j} \cdot \log\left[h_j\left(w, s_t\right)\right], \tag{3.5}$$

where w is the vector of weights in the neural network model, y_m is the one-hot vector representation of the true target at epoch t and $h_j(w, s_t)$ is the output value. C is the length of the one-hot vectors, i.e. the number of categories. A one-hot vector is a mapping of an integer category to a vector with a length equal to the number of categories, and in which all elements are zero apart from the element corresponding to the given category. For example, in the set of categories {0,1,2}, the category "2" maps to the vector [0,0,1]. The output of the network, $h_j(w, s_t)$, is therefore a real-valued estimate of the probability of the surface level being in a certain state.

3.3.3 Real-Valued Prediction

Another possible configuration of the output node allows for a real-valued prediction. A set of M training sequences is used to train the network, which takes the precipitation and temperature inputs and attempts to predict the target value at the output. These are coupled with one real-valued target, y_t, the soil height as measured in-situ. Error in the predictions is quantified through by the mean squared error (MSE) loss function, J, given by

$$J(w) = -\frac{1}{M}\sum_{t=1}^{M}\left[h\left(w, s_t\right) - y_t\right]^2, \tag{3.6}$$

where w is the vector of weights in the neural network model (also known as the model parameters) and $h(w, s_t)$ is the model output.

3.3.4 Network Optimization

In both configurations, the network weights are updated in each training iteration so as to minimize J. The problem statement can then be formulated as finding the optimal set of weights, w_{opt}, that minimize J as

$$w_{\text{opt}} = \arg\min_{w} J(w). \tag{3.7}$$

The first part of the measurement time series spanning dates from June 2020 – October 2022 is used as the training and validation set which is used to train the network

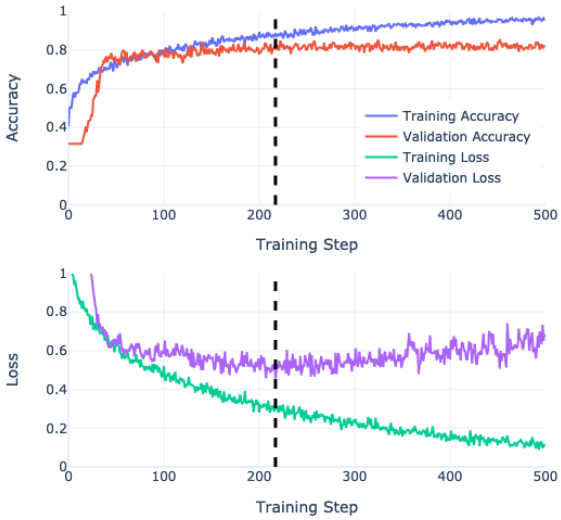

Figure 3.2: RNN learning curves. The black dashed line indicates the training iteration in which the optimum model weights are achieved. Accuracy is defined as the ratio of the number of correct predictions to the total number of predictions, and loss is defined by Eq. (3.5).

(i.e. optimize the network weight parameters). The final year of the measurement time series, from October 2022 - October 2023 is used as a testing set in order to assess the performance of the model.

The training dataset is split between 80% training and 20% validation. The network is trained over many iterations as the algorithm attempts to reduce the training loss. The final model which is saved is the iteration in which the validation loss is minimized. The graph of loss versus training iteration is known as the network's "learning curve". An example set of these learning curves for the categorical case is shown in Fig. 3.2. The model iteration used is given by the vertical black dashed line.

3.4 Parametric Model (SPAMS)

3.4.1 Model Description

The second model to be developed was a parametric model with the objective of describing the shallow soil behaviour as simply as possible. By keeping the model simple, we sacrifice some potential accuracy for better generality, as it becomes less likely to overfit models with simple parameterizations. In SPAMS, the overall relative soil surface height, z, in a given reference system and at a given time t relative to a start time t_0 in days, is modeled as a combination of reversible (shrinkage and swell) and irreversible processes (oxidation):

$$z(x,t) = R(x,t) + I(x,t), \tag{3.8}$$

where R is the reversible component and I is the irreversible component. x is the set of lithology dependent parameters which will vary with location. We do not consider irreversible subsidence due to compaction or creep, only model the behaviour of unloaded

soils with respect to changes in phreatic groundwater level and soil moisture. The reversible component is obtained by considering the balance between the dominant source and sink of ground water, i.e., precipitation and evapotranspiration, respectively. This balance is sometimes referred to as the rainfall or water surplus and its integral over time as the cumulative rainfall/water surplus [30], [91]. We modify this concept by introducing a scaling factor between the precipitation and evapotranspiration terms in order to model the material and hydrological properties of a given region. Different locations will have different soil stratigraphies, land use, and land parcel geometries, thus resulting in different responses to meteorological conditions. This also rescales the reference evapotranspiration value to one better suited to the vegetation cover of the area. Thus, the reversible component is modeled as the scaled difference between precipitation, $P(t)$, and evapotranspiration, $E(t)$, integrated over a period of time τ (in days):

$$R(x,t) = \sum_{t'=t-\tau}^{t} [x_P P(t') - x_E E(t')], \tag{3.9}$$

where P and E are daily mean precipitation and evapotranspiration in millimetres, respectively, as reported by KNMI, and x_P and x_E are unknown relative scaling factors which will differ per location. These factors reflect the relative effect on soil height each respective process has, i.e., their relative strengths based on seepage and infiltration, as well as the scaling from cumulative groundwater balance to soil surface height. The unknown integration time τ is also different for different soils, as the memory/hysteresis of the system will differ based on material properties and geometry.

The irreversible component of soil subsidence is often modeled and reported as a constant linear rate [8]. However, we note that this ignores the effect of water in the system, and the fact that oxidation primarily occurs while there is a net loss of water in the system. We can make a simple improvement to this approximation by taking into account when the soil is wetting or drying. We retain a constant linear rate, but modulate its activity based on the scaled water surplus of Eq. (3.9). Thus the irreversible component is estimated by

$$I(x,t) = \sum_{t'=t_0}^{t} x_I \cdot f(x,t), \tag{3.10}$$

where x_I is an unknown constant rate of irreversible subsidence, and

$$f(x,t) = \begin{cases} 0, & \text{for } R(x,t) > 0 \\ 1, & \text{for } R(x,t) \leq 0. \end{cases} \tag{3.11}$$

When $R > 0$ (see Eq. (3.9)), the precipitation term dominates and the soil is considered to be wetting. When $R \leq 0$, the evapotranspiration term dominates and the soil is considered to be drying, and undergoing oxidation. As-is, the model ignores the effects of compaction below the Vadose zone; that is, compaction in the saturated zone caused by the mass of the water above. This could be included by setting the zero-term in Eq. (3.11) to an unknown constant.

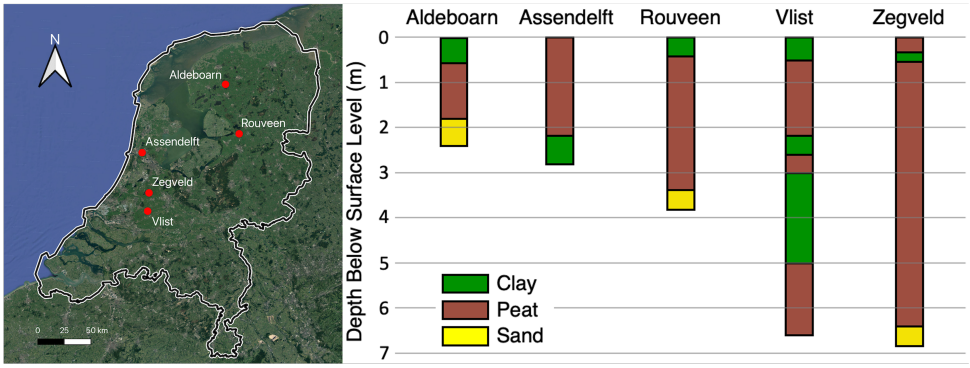

Figure 3.3: Map of The Netherlands annotated with extensometer locations in red, along with corresponding simplified Holocene borehole lithographies.

3.4.2 Parameter Estimation

There are four unknown parameters to estimate for a given location: the scaling factors x_P, x_E, and x_I, and the integration time τ. While these parameters are clearly linked with the physical makeup of the area of study, at this point we simply use them as empirical factors; additional study would be necessary to link the parameter values with soil and hydrological properties, as well as other factors such as parcel size and shape, land use, or ground water management factors such as freeboard or ditch water levels.

The parameters are estimated by minimizing the MSE between the model output and a set of training data. The MSE is quantified in a similar manner to the machine learning model "real-valued prediction" loss in Eq. 3.6:

$$J(\hat{x}) = \frac{1}{M}\sum_{t=1}^{M}\left[z(\hat{x},t) - y(t)\right]^2, \tag{3.12}$$

where $\hat{x}$ is the set of estimated parameters, M is the length of the time series, and y is the ground truth data, described in Sec. 3.5.1. The optimal set of parameters, $\hat{x}_{\text{opt}}$, is found by minimizing J such that

$$\hat{x}_{\text{opt}} = \arg\min_{\hat{x}} J(\hat{x}). \tag{3.13}$$

The first part of the measurement time series spanning dates from June 2020 – October 2022 is used as the training set which is used to fit the model parameters. The final year of the measurement time series, from October 2022 - October 2023 is used as a testing set in order to assess the performance of the model.

3.5 Results and Discussion

3.5.1 Test Locations

The two models are tested at five locations, shown in Fig. 3.3 along with corresponding publicly available borehole log data from the immediate vicinity [92]. The test locations are distributed across various different parts of the Dutch coastal plain and have different

Table 3.1: Categorical Test Confusion Matrix

	True STAY	True UP	True DOWN
Predicted STAY	0.61	0.12	0.22
Predicted UP	0.14	0.88	0.02
Predicted DOWN	0.24	0	0.76

combinations of clay, peat and sand in the Holocene sequence, thus providing a representative set of conditions for the region. All locations are managed grasslands used for agriculture.

The test locations were chosen by the project consortium *National research program on greenhouse gas emissions in peatlands* (Dutch: Nationaal Onderzoeksprogramma Broeikasgassen Veenweiden (NOBV)). Each site is equipped with an extensometer which are used as the source of the training and testing datasets at each location (available: [32]). The extensometers are permanent devices based in the Pleistocene layer, and consist of several measurement anchors at different vertical levels in the above Holocene. Thus they provide continuous measurements of the movement of the Holocene layer above the base of the device. We use the topmost anchor located at 5 cm depth, which is the shallowest depth at which an anchor can be reliably fixed, which almost completely captures the full motion of the Holocene layer(s). A full description of the system is provided in [13].

3.5.2 Machine Learning Model Results

3.5.2.1 Categorical Prediction

The categorical machine learning model is evaluated using 977 predictions on data gathered from October 15, 2018 to July 4, 2021 at the Zegveld site. The performance of the model is determined by creating a confusion matrix of the real vs. predicted movement classes, shown in Table 3.1. It can be seen that the network can very accurately differentiate between significant upward and downward motion (UP vs. DOWN) with little to no error (bottom-right quadrant of Table 3.1. It has greater difficulty in distinguishing between UP/DOWN vs. STAY. This is due to the fact that the choice of threshold which defines what amount of movement is considered large enough to be an UP/DOWN state vs. what is not large enough and is considered a STAY is a somewhat arbitrary distinction. This is not necessarily a major problem given the application of the categorical prediction in ambiguity resolution, because the important distinction to make is between significant upward and downward motion, i.e. correctly predicting UP vs. DOWN.

This was the first model to be developed over the course of this research and demonstrated the feasibility of using a machine learning model to make predictions about the state of ground motion. This model is used in tandem with the phase observations in a temporal phase unwrapping algorithm which is described in detail in Ch. 4.

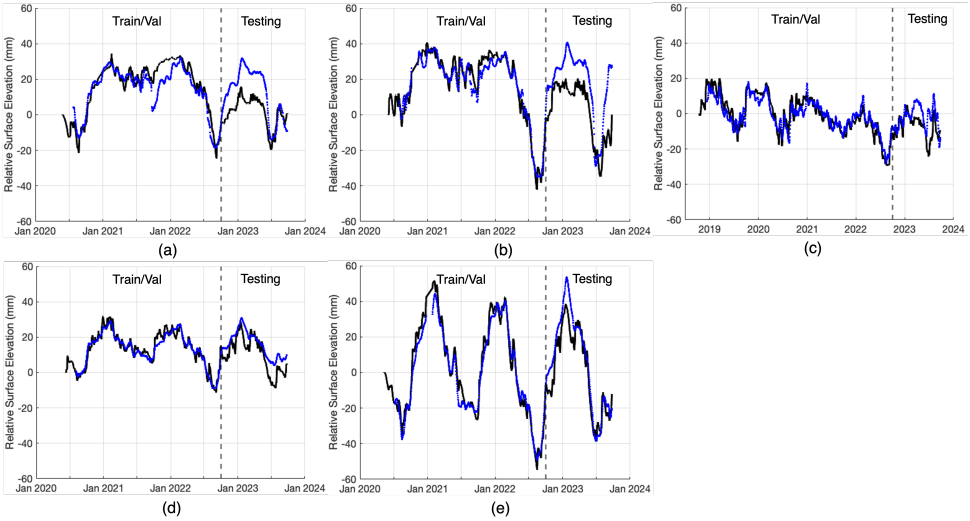

Figure 3.4: Machine learning model training/validation and testing results (blue) plotted against ground truth data obtained from extensometer readings at each location (black). Locations: (a) Aldeboarn, (b) Assendelft, (c) Rouveen, (d) Vlist, (e) Zegveld.

3.5.2.2 Real-Valued Prediction

Time series results of the real-valued version of the machine learning model are shown in Fig. 3.4 and Table 3.2. The model is able to make accurate predictions of the ground motion time series in three of the five test locations: (c), (d) and (e). The model struggled to correctly predict the correct annual seasonal uplift for locations (a) and (b), largely overestimating the fall/winter rebound by quite a large margin. In fact this effect is visible to some degree in all the predicted time series, which raises concerns about this model's ability to reliably predict an accurate value for the irreversible component of subsidence. While this model can generate a time series that "looks" realistic, the black-box nature of the model's predictions and the fact that it is easy to overfit such models during training means that the reliability of this technique is in question. However, it may be possible to improve the performance with a different RNN architecture, for instance by removing a hidden layer or by performing fewer training epochs to prevent overfitting. These measures would likely come at the cost of reduced accuracy.

3.5.3 Parametric Model Results

The SPAMS model is validated by comparing the output to in-situ measurement data, i.e. testing data, taken by extensometer readings from the five test locations, shown in Fig. 3.5 and Table 3.3. This demonstrates that the model is able to reliably approximate the relative soil displacement at every test location, each with different Holocene soil stratigraphies and depths.

While it is clear that the model is too simple to perfectly capture all the high-frequency components of the surface motion, the mean seasonal and sub-seasonal effects are accurately modelled at every location, which is sufficient for our objectives. The best perfor-

Table 3.2: Real-Valued Machine Learning Model Performance

Location Designation	Aldeboarn (a)	Assendelft (b)	Rouveen (c)	Vlist (d)	Zegveld (e)
Distance site-station [km]	17.4	9.80	34.8	7.20	19.4
Training RMSE [mm]	6.05	4.42	4.01	2.53	4.87
Testing RMSE [mm]	13.1	17.9	8.64	6.82	8.66
Training RMSE/σ [mm/mm]	0.41	0.23	0.39	0.26	0.20
Testing RMSE/σ [mm/mm]	0.89	0.92	0.83	0.70	0.35
Overall R^2 [ ]	0.65	0.71	0.75	0.81	0.94

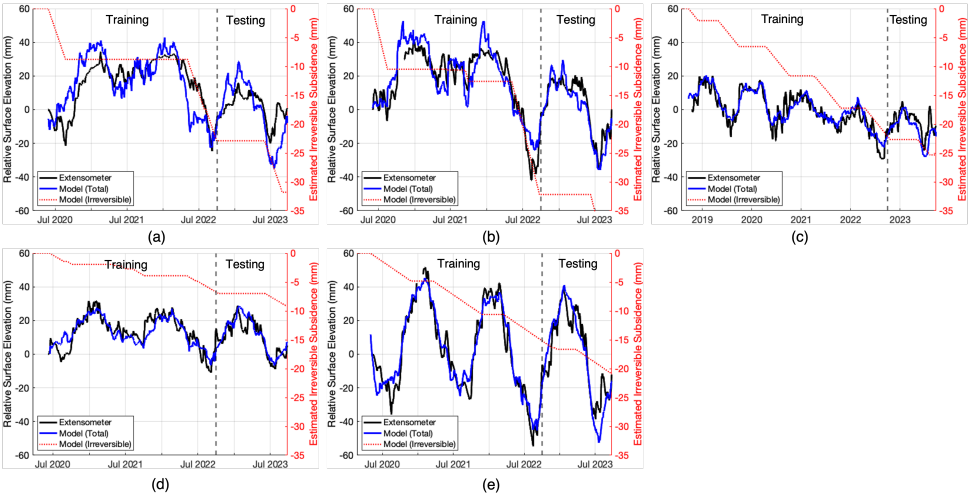

Figure 3.5: SPAMS model training and testing results (blue) plotted against ground truth data obtained from extensometer readings at each location (black). Red: estimated irreversible component. Locations: (a) Aldeboarn, (b) Assendelft, (c) Rouveen, (d) Vlist, (e) Zegveld.

mance is found at sites b) and d), which display the lowest overall root mean squared error (RMSE), and the lowest RMSE normalized to the standard deviation of the in-situ data (RMSE/σ) respectively. These sites also have the shortest distance to their corresponding weather station, so it is likely that they have the most accurate input data. The worst performance is encountered at site a), which exhibits the most complex displacement history, with very large differences between subsequent years. Nevertheless, the model is still able to capture the large variations between seasons.

A major benefit to parameterizing the irreversible component separately from the overall surface displacement is that it allows for an estimation of when the most significant soil volume loss occurs. For example, the irreversible component was significantly greater at sites a), b) and d) in the hot and dry summer of 2022 compared to the previous wetter and cooler summer of 2021. Thus we can compare the irreversible rates between years to gain understanding of how climate stresses on the soil and water system can affect subsidence.

Table 3.3: SPAMS Estimated Model Parameters and Performance

Location Designation	Aldeboarn (a)	Assendelft (b)	Rouveen (c)	Vlist (d)	Zegveld (e)
Distance site-station [km]	17.4	9.8	34.8	7.2	19.4
x_P [m/mm]	$1.7{\cdot}10^{-4}$	$1.5{\cdot}10^{-4}$	$6.3{\cdot}10^{-5}$	$8.0{\cdot}10^{-5}$	$9.7{\cdot}10^{-5}$
x_E [m/mm]	$1.3{\cdot}10^{-4}$	$9.2{\cdot}10^{-5}$	$8.2{\cdot}10^{-5}$	$6.4{\cdot}10^{-5}$	$2.7{\cdot}10^{-4}$
x_I [m/d]	$-1.0{\cdot}10^{-4}$	$-1.4{\cdot}10^{-4}$	$-2.9{\cdot}10^{-5}$	$-2.0{\cdot}10^{-5}$	$-2.3{\cdot}10^{-5}$
τ [d]	80	80	54	86	69
Training RMSE [mm]	10.6	8.39	4.18	4.80	6.64
Testing RMSE [mm]	12.2	6.38	6.03	4.51	10.4
Training RMSE/σ [mm/mm]	0.72	0.43	0.40	0.49	0.27
Testing RMSE/σ [mm/mm]	0.83	0.33	0.58	0.46	0.42
Overall R^2 [ ]	0.43	0.84	0.80	0.77	0.90

3.6 Conclusion

In this chapter we have demonstrated that the motion of shallow soft soils is predictable and how it can be modeled using local weather data. The feasibility in employing simple models for the kinematic behaviour of shallow soil motion without invoking complex hydrological simulations is demonstrated. Depending on their application, these models can take different forms and output either categorical or time series predictions of soil motion. Ch. 4 and 5 will describe how these models are employed within the InSAR processing chain to resolve ambiguities and help overcome data loss caused by temporal decorrelation.

In the future, the models can also be further developed; for example, the SPAMS model can be improved by considering an additional compaction term which acts on the saturated soils below the Vadose zone. This would require an additional parameter to be estimated, but could provide more long term accuracy in the estimation of irreversible subsidence occurring over decadal time scales.

4

Temporal Phase Unwrapping and Cycle Slips

A leap of faith – yes, but only after reflection.

Søren Kierkegaard

This chapter describes a methodology which allows for additional environmental information to be integrated into temporal phase unwrapping procedures, mitigating the effects of rapid displacements which exceed the standard unwrapping threshold of a quarter wavelength. Problematic epochs that cause errors in the temporal phase unwrapping process can be anticipated by machine learning algorithms which create categorical predictions about the relative ambiguity level based on readily-available meteorological data. These predictions significantly assist in the interpretation of large changes in the wrapped interferometric phase and enable the monitoring of environments not previously possible using standard minimum-gradient phase unwrapping techniques.

Highlights:

1. *Standard (explicit) phase unwrapping techniques are not able to cope with signals exceeding the critical unwrapping threshold.*
2. *Previous studies of Dutch peatlands showed extremely large subsidence rates, often because the rapid upward motion of the peat following rainy periods was missed by standard phase unwrapping.*
3. *Augmenting the phase unwrapping process with additional information about the expected behaviour of the environment can prevent phase unwrapping errors.*

Parts of this chapter have been published in *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 60, No. 5234611 [75].

4.1 Introduction

Phase unwrapping, also known as ambiguity resolution, is an underdetermined problem in which assumptions must be made in order to obtain a result in InSAR time series analysis. This problem is particularly acute for distributed scatterer InSAR, in which noise levels can be so large that they are comparable in magnitude to the signal of investigation. Additionally, deformation rates can be highly nonlinear and orders of magnitude larger than neighboring point scatterers, which may be part of a more stable object.

The combination of these factors has often proven too challenging for conventional InSAR processing methods to successfully monitor the low-lying peatland regions of the Netherlands. The soft soils which comprise the majority of the country's agricultural lands are prone to rapid deformation and are very difficult, or even impossible to directly monitor using standard distributed scatterer (DS) InSAR techniques [19], [28] which use smoothness constraints to perform spatial phase unwrapping [71], [93], [94]. Recently, a significant amount of research has been conducted to improve two-dimensional phase unwrapping [95], and there is great interest in the applicability of machine learning techniques in solving the problem [96]; however, this research has largely been limited to the spatial case and does not consider the temporal effects of a rapidly deforming region. Rapid soil motion and non-stationary coherence cause large fluctuations affecting both the functional and stochastic models, which rules out the use of time-domain phase unwrapping techniques such as integer integer least-squares (ILS) [73], as the stochastic model will become too large to reliably evaluate an assumed deformation model.

Using newly available ground truth measurements, we find that in the case of Sentinel-1 observations, rapid soil uplift due to increases in the ground water level between satellite overpasses can cause phase displacements larger than half a cycle ($\lambda/4$) at C-band. This renders time series analysis of the region using Sentinel-1 imagery impossible using standard techniques, because phase unwrapping algorithms will typically assume the smaller phase displacement in the opposite direction to be the correct solution. This introduces a systematic error in all subsequent points of the time series corresponding to one 2π ambiguity level. While this problem may potentially be avoided by using L-band SAR data, there is currently no operational L-band radar mission observing the Netherlands at the temporal frequencies required for monitoring this environment.

It is therefore necessary to augment the standard unwrapping procedure with additional information in order to prevent errors in the unwrapping direction. This is done by considering the Direction of ground motion (DOGM), and therefore the correct phase unwrapping direction, as states in a hidden Markov model (HMM). By using a modification of the widely-used Viterbi algorithm [97], we can integrate additional contextual information about the system into a generalized probabilistic framework which can be used to guide the unwrapping of the interferometric phase in the time domain. In our case, this additional information comes from a Recurrent Neural Network (RNN) which takes meteorological data as its input to predict the DOGM. This was chosen because meteorological data is easily accessible and interpretable for our areas of study, but in general, any model which predicts the state of ground motion can be used. We show that this framework is able to reliably anticipate rapid soil deformation events and correctly unwrap the interferometric phases.

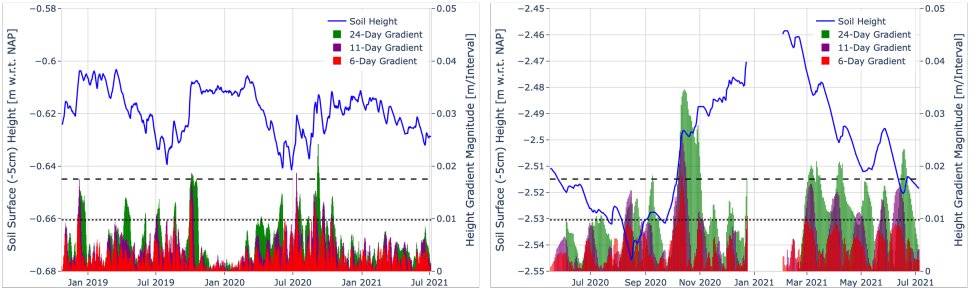

Figure 4.1: (a) Rouveen and (b) Zegveld daily extensometer soil height measurements with respect to Netherlands vertical datum (Dutch: Normaal Amsterdams Peil) (blue trace, left-hand y-axis) and corresponding height gradient magnitudes computed over 6, 11 and 24-day intervals (red, purple, green bars, right-hand y-axis). Representative C- and X-band unwrapping thresholds projected from slant range to the vertical are indicated by the horizontal dashed and dotted lines respectively (right-hand y-axis). Not visible: L-band unwrapping threshold projected to the vertical (0.0726 m). The gap in data around Jan. 2021 in Zegveld is caused by uplift which exceeded the range of the extensometer scale and could not be recorded.

4.2 Extensometer Data and Signal Simulation

4.2.1 Extensometer Data

Extensometers placed in various locations across the Netherlands have been continuously monitoring the vertical movement of the surface and shallow subsurface of peaty soft soils for several years [13]. The longest running of these stations is in Rouveen, which has been collecting data since October 2018, and exhibits relatively modest levels of movement, within a range of 4 cm, see Fig. 4.1 (a). On the other hand, some of the largest measured soil variations are observed in Zegveld, within a range of 9 cm, see Fig. 4.1 (b).

These two locations therefore show a representative range of ground movements which can be reasonably expected to be encountered in Dutch agricultural areas. Inspection of the extensometer measurements (blue trace) shows a very high degree of reversible deformation following a strong seasonal trend. The most stable periods are during the winter months, when cold, rainy conditions in the Netherlands keep the soils saturated with moisture. As temperatures rise and precipitation levels decrease, the soils dry out and become more unstable. A significant amount of movement over short time scales can also be observed. This creates an additional complication for InSAR observations, as significant movement can occur even between the relatively frequent 6-day Sentinel-1 overpass cycle, which is shown by the 6-day height gradient magnitude (red bars). In the case of Zegveld, shifts in soil height can cause strong gradients in the measured signal which exceed the unwrapping threshold (projected from slant range to the vertical using a typical Sentinel-1 incidence angle of 37°, and shown by the dashed horizontal line), even in an ideal, noiseless case. We can also infer that we do not need to expect the deformation to exceed multiple ambiguity levels at C-band within one overpass.

For other missions, the outlook is significantly worse. In the case of Radarsat2 (corresponding to the 24-day gradient shown by green bars), which also operates at C-band, the

gradients routinely exceed the unwrapping threshold in both locations. With TerraSAR-X (corresponding to the 11-day gradient shown by purple bars, and the dotted line representing the X-band unwrapping threshold height projected from slant range to vertical), the signal shifts are consistently larger than the unwrapping threshold, rendering any time-series interpretation extremely challenging. On a more hopeful note, the future ROSE-L mission operating at L-band, with a 6-day revisit time should be able to comfortably follow the observed signal without significant risk of phase unwrapping errors (The L-band unwrapping threshold is significantly larger than the dynamic range of the signals plotted in Fig. 4.1 and is therefore not shown). While this mission will undoubtedly be a very useful asset for subsidence monitoring in the Netherlands, it is expected to be launched in the year 2027, and will then require several additional years to build up a dataset of observations to allow for time-series analysis. This means scientists and policymakers will need to wait about another 10 years before the benefits of the ROSE-L mission become truly available to them, and many decisions until that time will need to be based on available Sentinel-1 data. While other L-band missions have been flown, notably ALOS-1, ALOS-2 and SAOCOM, there have not been enough acquisitions to provide the necessary coverage.

Note that the gap in data from end Dec. 2020 to mid Jan. 2021 in Fig. 4.1 in Zegveld is caused by uplift which exceeded the range of the extensometer scale and could not be recorded. We conservatively do not use the data from this time in our subsequent analysis; however it could be interpreted that an additional spike in the observed deformation gradients would be present here as well.

4.2.2 Simulated InSAR Signal

In order to assess the ability of an algorithm to correctly unwrap the observed phase, we simulate an InSAR signal based on the ground truth provided by the extensometers. This allows us to control the level of noise in the data and reference the obtained solutions to a known true value. We produce this simulated signal, ϕ_{sim}, by downsampling the extensometer data to one observation per six days, converting vertical displacement into phase, projecting from the vertical axis onto the slant range direction, adding noise, and wrapping the resultant phase on the interval $[-\pi, \pi)$. The simulated signal is given by

$$\phi_{\text{sim},ij} = W\left\{\Delta z_{\text{ext}}(t_i, t_j) \cdot \frac{4\pi \cos\theta}{\lambda} + n(t_i, t_j)\right\}, \tag{4.1}$$

where $W\{\cdot\}$ is the phase wrapping operator [33], $\Delta z_{\text{ext}}(t_i, t_j)$ is the change in the extensometer signal between acquisition epochs t_i and t_j, λ is the radar wavelength of 0.0556 m, corresponding to the wavelength of the Sentinel-1 radar, θ is the radar incidence angle and $n(t_i, t_j)$ is additive heteroscedastic noise. The variance of $n(t_i, t_j)$ differs for each interferogram and is governed by the coherence γ_{ij}, and equivalent number of looks, L [33]. The coherence of a multilooked set of pixels Ω in an interferogram is estimated

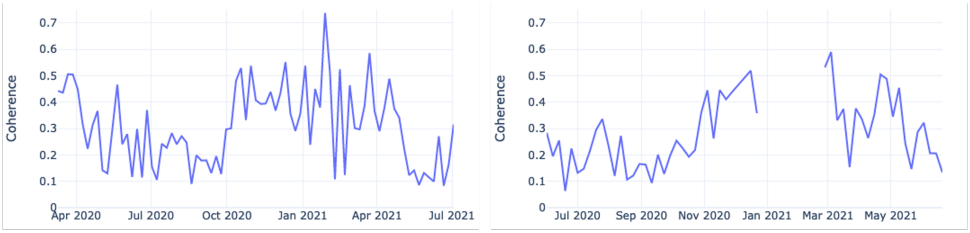

Figure 4.2: Daisy-chain coherence values used to generate additive noise in (a) Rouveen and (b) Zegveld. The term "daisy-chain" refers to the off-diagonal elements of the coherence matrix showing the relative coherence between consecutive epochs.

from the magnitude of the complex sample coherence given in [33] by

$$|\gamma_{ij}| = \frac{\left|\sum_{n\in\Omega} S_i^n S_j^{n*}\right|}{\sqrt{\left(\sum_{n\in\Omega} |S_i^n|^2\right)\left(\sum_{n\in\Omega} |S_j^n|^2\right)}}, \tag{4.2}$$

where $S_{i,j}^n$ is the nth SAR single-look complex (SLC) pixel in acquisitions i and j. We focus on the daisy-chain coherence, which refers to the coherence observed in interferograms of consecutive acquisitions ($j = i-1$), i.e. the magnitudes of the first off-diagonal of the full coherence matrix, as shown in Fig. 4.2. For simplicity, we denote the daisy-chain phase $\phi(t_i, t_{i-1})$ as ϕ_{dc}. The additive phase noise is modelled as a random process which follows a circular Gaussian distribution. The probability density function (PDF) of the interferometric phase is [33], [62], [63]

$$f(\phi|\gamma, L, \phi_0) = \frac{\Gamma(L+1/2)(1-|\gamma|^2)^L\beta}{2\sqrt{\pi}\Gamma(L)(1-\beta^2)^{L+1/2}} + \frac{(1-|\gamma|^2)^L}{2\pi} \cdot {}_2F_1(L, 1; \frac{1}{2}; \beta^2), \tag{4.3}$$

where $\beta = |\gamma|\cos(\phi - \phi_0)$, $\Gamma(\cdot)$ is the gamma function and $F(\cdot)$ is the hypergeometric function. The PDF is evaluated over the interval $\phi = [-\pi, \pi)$ and the mean, ϕ_0, is taken as 0 to center the additive noise around the signal. An equivalent alternative formulation is possible in which one sets $\phi_0 = W\{\Delta z_{\text{ext}} \cdot 4\pi\cos\theta/\lambda\}$, due to the fact that $f(\phi|\gamma, L, \phi_0)$ produces random samples of wrapped phases. In this case, the distribution would describe the entire signal, ϕ_{sim}, and not only the noise component, and the phase wrapping would be stated implicitly. We choose to use the $\phi_0 = 0$ formulation because it allows the reader to more clearly see all the steps which are taken to produce the simulated phases.

Thus, every epoch, i, has a different noise variance σ_n^2, which is based on the daisy-chain coherence value between epochs i and $i-1$, and determines the shape of the distribution $f(\phi_{\text{dc}}|\gamma_{\text{dc}}, L)$, which is randomly sampled. The simulated coherence values used to generate these noise values are shown in Fig. 4.2, which are representative values for Sentinel-1 observations over Rouveen and Zegveld during the periods in which the extensometers were active. Epochs with coherences of less than 0.05 can occasionally occur,

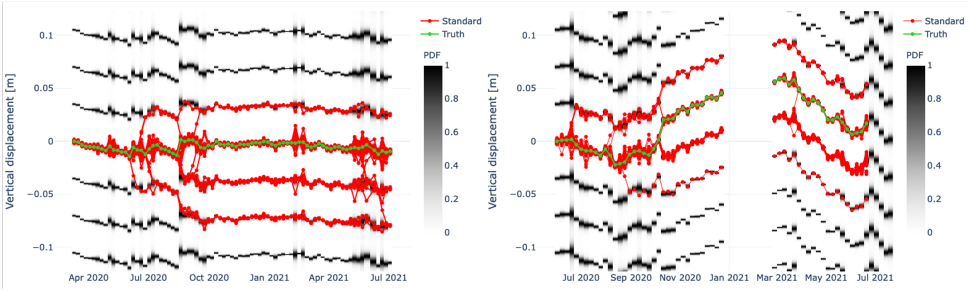

Figure 4.3: Normalized multi-modal PDF (greyscale) based on Eqs. (4.1) and (4.3) from downsampled extensometer soil height measurements (green trace) with additive noise, corresponding to the Sentinel-1 revisit time for (a) Rouveen and (b) Zegveld. During times of high coherence, the PDF is clearly defined and a clear peak is visible at each ambiguity level. During low coherence epochs, the PDF becomes "smeared-out" and distinguishing ambiguity levels becomes challenging. 50 realizations of random noise are generated and a solution obtained using the standard minimum gradient approach (red traces).

4

but are omitted from this study because this results in almost complete information loss in the interferogram, and phase unwrapping becomes redundant because only noise remains in the observation. In a fully-realized InSAR processor, these interferograms must be omitted, or bridged using some form of assumption or pseudo-observation. We use a multilooking factor $L = 100$, as strong multilooking is an important part of a strategy for noise-reduction to enable InSAR over grasslands on peaty soil. Without such high levels of multilooking, noise levels in the summer months of low coherence can become so strong that it is impossible to interpret the interferometric phase.

To visualize the effect of noise, Fig. 4.3 shows 50 unwrapped time series (red traces) obtained by 50 realizations of random noise generated by the PDF (greyscale). The PDF is multimodal, which reflects that the observed phase is wrapped and the true ambiguity level is unknown. Epochs of low coherence widen the PDF such that it is no longer a set of clearly defined peaks, which can cause the observed phase change between epochs to exceed the unwrapping threshold. The divergence of the red traces illustrates how noise can have a significant effect on unwrapping decisions, and therefore on the inferred vertical displacement. Strong noise can either introduce new unwrapping errors, as in the case of Rouveen, or it can exacerbate existing problems already present, as in the Zegveld case, cf. Fig. 4.1 (a) and (b), respectively.

4.2.3 Implications of Rapid Ground Motion and Noise

Figs. 4.1, 4.2, and 4.3 demonstrate the difficulty of directly producing InSAR time series estimates of the Dutch polderlands. First, from a signal perspective, Fig. 4.1 shows that the ground motion which we attempt to measure is highly rapid and nonlinear. Large shifts can occur even within the 6-day Sentinel-1 revisit interval, creating problems for phase unwrapping even in a noise-free scenario. The second additional challenge is that of noise, which is related to that of temporal decorrelation, shown in Figs. 4.2 and 4.3. Additive noise can occasionally become very large during periods of low coherence and

Table 4.1: Categorical RNN Confusion Matrix

	True STAY	True UP	True DOWN
Predicted STAY	0.61	0.12	0.22
Predicted UP	0.14	0.88	0.02
Predicted DOWN	0.24	0	0.76

can cause phase wrapping to occur when there would be none in the hypothetical noise-free scenario. This is of particular importance during the late summer and autumn months, when the soil height gradients are at their largest. The observer, and by extension, the algorithm, needs more information in order to correctly unwrap the observed phase in this environment. The following section will describe how this may be accomplished by using readily available meteorological data.

4.3 RNN-Aided Phase Unwrapping

4.3.1 Introduction

The categorical RNN model described in Sec. 3.3.2 is able to estimate the DOGM, however, these predictions are not 100% accurate (as shown in Table 4.1), and all this information must still be integrated into a phase unwrapping routine. Additionally, the wrapped phases and coherences are also important information which should be used. To integrate all this information together, the problem can be posed as an HMM, which is a general formulation describing how a system may transition between various unobservable "hidden" states (the unwrapped phase transitions) given a set of noisy observations (the recurrent neural network (RNN) estimation).

4.3.2 Hidden Markov Model

The system of wrapped interferometric phase observations and estimated RNN classes are considered to act together as an HMM, pictured in Fig. 4.4. For simplicity, we consider three hidden states corresponding to the average ground movement: uplift (UP), subsidence (DOWN), or no significant movement (STAY), which correspond closely to (but are distinct from) the RNN output classes. These states are hidden, because the nature of wrapped phases is such that one cannot discern the DOGM from them directly. The additional observations in this HMM therefore come from the RNN output classes.

To complete the model, a set of *transition* and *emission* probabilities are required. The transition probabilities represent the probability of moving to a given state from the current state. In our implementation, these are estimated using the wrapped phases, ϕ_{dc}. The probabilities are estimated in two steps: (i) estimating the probability of downward vs. upward motion, and (ii) estimating the significance of the motion, by comparing the magnitude of the phase change with the distribution given in Eq. (4.3). The up/down probabilities are determined by considering the two nearest ambiguity levels as branches:

$$\begin{aligned} b_1 &= \phi_{\mathrm{dc}} \\ b_2 &= \phi_{\mathrm{dc}} - \mathrm{sign}(\phi_{\mathrm{dc}}) \cdot 2\pi, \end{aligned} \tag{4.4}$$

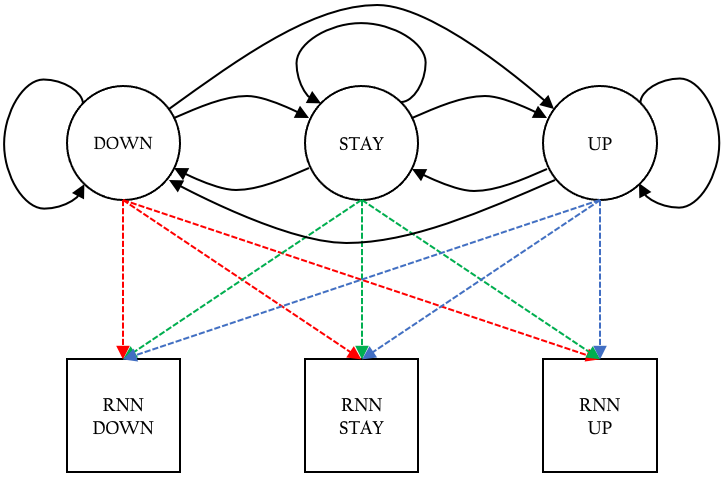

Figure 4.4: Hidden Markov Model used to integrate wrapped phase information with RNN output. Solid lines represent transition probabilities, dashed lines represent emission probabilities.

where ϕ_{dc} is the observed change in the wrapped phases between epochs, which is equal to the equivalent daisy-chain phase. These branches correspond to the UP or DOWN state, depending on the sign of the observed phase change. The initial probabilities of the two complementary branches are estimated based on the magnitude of ϕ_{dc} and are given by

$$\begin{aligned} p_{b_1} &= 1 - \frac{1}{2}\left[\text{erf}\left(|\phi_{\text{dc}}| - \pi\right) + 1\right] \\ p_{b_2} &= 1 - p_{b_1}, \end{aligned} \tag{4.5}$$

where $\text{erf}(\cdot)$ denotes the error function. As $|\phi_{\text{dc}}|$ increases, the probability of remaining on branch b_1 smoothly decreases from 1. When $|\phi_{\text{dc}}| = \pi$, both branches are equally likely. In this initial estimation, the branch corresponding to the minimum phase gradient solution is favoured. In the absence of any additional information, this branch would always be selected by the algorithm, corresponding to the standard unwrapping solution. Which branch corresponds to UP and which to DOWN depends on the sign of ϕ_{dc}.

While the probabilities estimated by Eq. (4.5) provide the relative likelihoods of transitioning to the UP vs DOWN state, to complete the set of transition probabilities, the likelihood of the STAY state must also be estimated. This is done by using the temporal coherence to compare the magnitude of the phase to three times the estimated phase noise standard deviation, σ_{ϕ_n}, of that epoch, which is determined using the circular Gaussian distribution given in Eq. (4.3). This provides an estimate as to whether or not an observed phase change contains significant movement with respect to the estimated level of noise in the measurement. This probability of significance is estimated by comparing

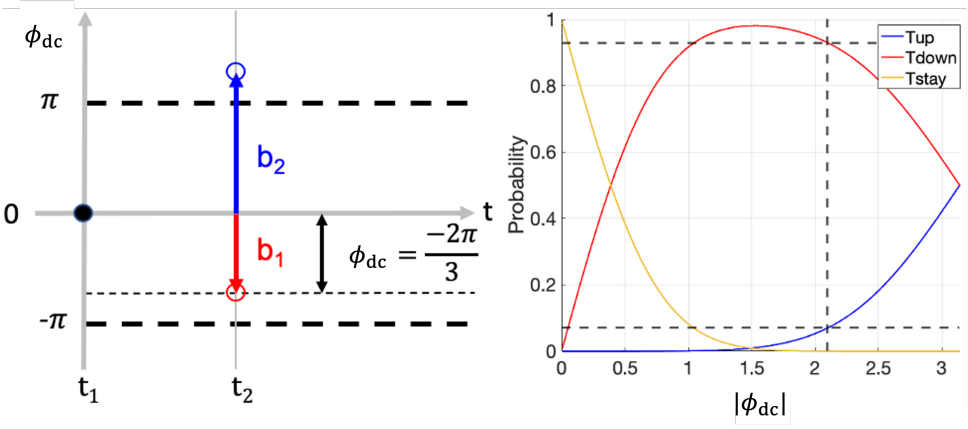

Figure 4.5: Example estimation of the transition probabilities based on an observed phase change $\phi_{dc} = -2\pi/3$. Left: blue indicates the upward branch and red indicates the downward branch. Right: blue indicates the T(UP), red T(DOWN), and yellow T(STAY). The black dashed lines correspond to the UP and DOWN probabilities estimated for ϕ_{dc} (not shown: T(STAY)=0).

the magnitude of the phase change with the *n*-sigma noise level:

$$p_{sig} = \text{erf}\left(\frac{1}{\sqrt{2}} \cdot \frac{|\phi_{dc}|}{n\sigma_{\phi_n}}\right). \tag{4.6}$$

The best value of n is empirically found to be 1.5. Finally, to obtain the set of transition probabilities, the probabilities of Eq. (4.5) are conditioned on p_{sig}:

$$\begin{aligned} T(\text{UP}) &= p_{b_{1,2}} \cdot p_{sig} \\ T(\text{DOWN}) &= p_{b_{2,1}} \cdot p_{sig} \\ T(\text{STAY}) &= 1 - p_{sig}, \end{aligned} \tag{4.7}$$

where the subscript $b_{1,2}$ refers to the branch corresponding to upward motion, and $b_{2,1}$ is the opposite branch corresponding to downward motion depending on the signs of the terms in Eq. (4.4). A graphical example of how an observed phase change relates to the transition probabilities is shown in Fig. 4.5. The *emission* probabilities describe the likelihood of an external observation (in our case the RNN output) while being in a given state. These probabilities are estimated by taking the values of the appropriate column of the confusion matrix obtained from testing the RNN (Table 3.1). There are two key differences which differentiate our implementation and a typical HMM example: (i) the transition probabilities are usually assumed to be static, but in our case they change with every epoch, and (ii) in our case, the transition probabilities are independent of the current state, while usually they are not. The effect this has on the calculation is described in Sec. 4.3.3.

4.3.3 Modified Viterbi Algorithm

The Viterbi algorithm [97] describes how to predict the most likely sequence of states in an HMM. It does this by predicting the most likely current state and subsequently using that state's transition probabilities in combination with the external observations to predict the next state. Thus, the sequence of most-likely states, called the Viterbi path, is determined. The output of this calculation is a "*trellis*" of probabilities P which describe the likelihood of being in a given state s at epoch t (corresponding to t_i in Eq. (4.1)), which can be written as

$$P(t, s) = P(t-1, k) \cdot T(k, s) \cdot E(t, s), \tag{4.8}$$

where k is the most likely state at time $t-1$, T is the matrix of state transition probabilities and E is the matrix of emission probabilities. This calculation is simplified in our implementation by the fact that the next state does not depend on the current state; that is, while the set of state transition probabilities change from epoch to epoch, they are the same for each state. Since $P(t-1, k)$ is the same for each choice of k, it becomes a scalar which no longer affects the outcome and is dropped. We therefore calculate the trellis of probabilities as

$$P(t, s) = T(t, s) \cdot E(t, s), \tag{4.9}$$

which shows how the RNN estimation is combined with the wrapped phases to inform the algorithm how the phase should be unwrapped. T represents the information present in the wrapped phases, and E the information from the RNN. If one were to imagine the case in which E were not present, that is, $P(t, s) = T(t, s)$, the phase unwrapping would be completely informed by the wrapped phases and the algorithm would behave identically to the standard minimum gradient implementations. Inversely, if one were to ignore the wrapped phases and take $P(t, s) = E(t, s)$, the phase unwrapping would be driven completely by the RNN. Therefore Eq. (4.9) shows how the final probabilities are assessed based on the balance of the relative levels of confidence in the two sources of information at that epoch.

4.3.4 Phase Unwrapping

The phase unwrapping works by stepping through the trellis to find the most likely state at each epoch and integrating the corresponding phase. The two branches in Eq. (4.4) are considered as the two possible *paths* the algorithm can take to unwrap the phase. The algorithm steps through each epoch and selects the appropriate ambiguity level, i.e., -2π, 0, or $+2\pi$, to add to the phase at that epoch according to the value of the most likely state. In the case of an UP state, the upper branch is chosen, and in the DOWN state, the lower branch is chosen. In the STAY case, the branch corresponding to the smallest phase change is selected. This way, the algorithm attempts to find the correct ambiguity level without modifying the observed phases. When the observed phase change is low, or the estimated level of noise is high, the algorithm becomes more conservative and the STAY state becomes more likely, causing the algorithm to default back to the standard minimum gradient solution.

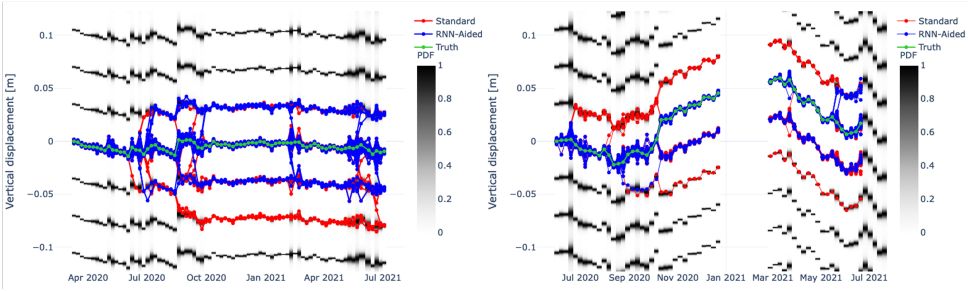

(a) (b)

Figure 4.6: Normalized multi-modal PDF (greyscale) based on Eqs. (4.1) and (4.3) from downsampled extensometer soil height measurements (green trace) with additive noise, corresponding to Sentinel-1 revisit time for (a) Rouveen and (b) Zegveld. 50 realizations of random noise are generated and solutions obtained using the standard minimum gradient (red traces) and the RNN-aided (blue traces) approaches.

4

4.4 Results

4.4.1 Optimal Path Through Noisy Data

The RNN-aided algorithm is used to unwrap fifty instances of noisy phases originally shown in Fig. 4.3, shown again with the RNN-aided solution in Fig. 4.6. While the RNN-aided algorithm solutions contain significantly less error than their standard counterparts, it can be observed that noise still has a significant effect on the solution. We therefore include the option to retain multiple solutions which branch out in the problematic epochs where the algorithm cannot confidently choose a certain path.

The effect of introducing the RNN-aided algorithm is clearly shown in Fig. 4.7,which shows histograms of the error in the final epoch of the unwrapped time series with respect to the true value in 1000 realizations of noise. Here, a downward skew can be seen in the standard algorithm's results at both locations, but which is particularly pronounced in the Zegveld case. This downward skew is strongly mitigated by the RNN-aided algorithm, and the correct solution is more consistently obtained.

4.4.2 Unwrapping Success Rates Versus Coherence

The RNN-aided phase unwrapping algorithm is tested against a standard minimum-gradient unwrapping algorithm to test its ability to correctly interpret the simulated deformation phase as given in Eq. (4.1). Several values of temporal coherence ranging from 0.05 to 0.95 are swept through to generate the test cases, see Fig. 4.8 for Rouveen and Zegveld. For each coherence level, 1000 independent noise simulations are generated according to the circular Gaussian distribution in Eq. (4.3). Note that for this particular test, while the noise is random from epoch to epoch, the coherence is held constant for the entire run through the time-series. This simplification is made in order to generate the test statistics and accurately compare the two algorithms in a standardized manner. For each time-series run, the integer number of unwrapping errors is tallied and the unwrapping error rate is determined by dividing the total number of errors by the number of unwrapping operations performed for each coherence level ($N_{\text{errors}}/N_{\text{epochs}} \cdot N_{\text{iterations}}$). The position

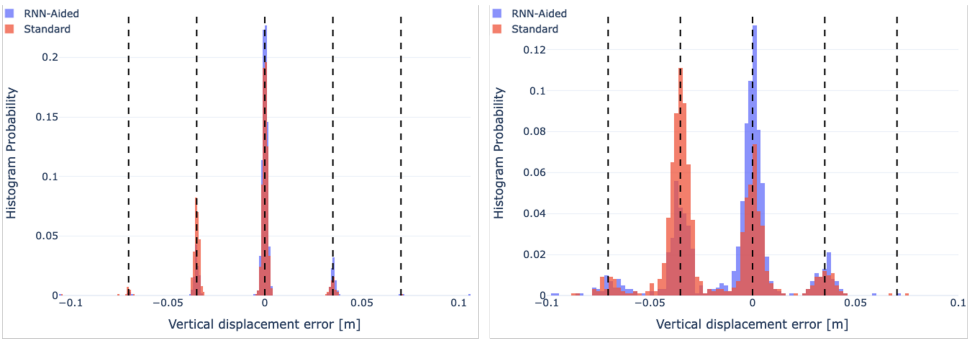

Figure 4.7: Histograms of the error in the final epoch of the obtained time-series in (a) Rouveen and (b) Zegveld of 1000 realizations of simulated noise. Red: standard algorithm. Blue: RNN-aided algorithm. The nearest ambiguity levels are indicated by the black dashed lines. The standard algorithm produces results which are downwardly skewed. This skew is corrected by the RNN-aided algorithm.

of the unwrapping error in the time series is not considered in this tally.

The RNN-aided unwrapping algorithm consistently outperforms the standard algorithm in both cases. In the case of Rouveen, both are able to correctly unwrap the time-series and retrieve the correct deformation phase in high-coherence scenarios, because the deformation phase gradient never exceeds the unwrapping threshold. But as coherence drops and noise levels increase, errors begin to appear in the solution. The RNN-aided algorithm is more resistant to this noise, as it does not solely rely upon the wrapped phases to estimate the correct ambiguity level. In the Zegveld case, the deformation gradients are stronger and do exceed the unwrapping threshold even in a noiseless scenario. Thus we conclude that the standard algorithm will intrinsically never be able to estimate the correct ambiguity level. However, the RNN-aided algorithm is able to consistently unwrap the time series correctly. As coherence decreases, the likelihood of making an error increases, as spikes of noise can affect the algorithm's ability to interpret the phase. Nevertheless, the RNN-aided algorithm is still able to reject some of the effects of noise and perform better, or as well as, the standard algorithm.

4.5 Discussion

Extensometer measurements gathered throughout the Dutch polderlands have shown that there is a very close relationship between the phreatic groundwater level and soil height [13]. While ground water measurements are not directly available, we are able to anticipate large shifts in the soil height which disrupt standard phase unwrapping by means of a motion-classifier RNN. This model can very reliably anticipate the large changes in the ground motion of nearby independent locations, as shown by the lower-right quadrant of the confusion matrix in Table 3.1. The model has greater difficulty in differentiating between significant and insignificant levels of motion in both directions, which is a somewhat arbitrary distinction. This can be avoided by removing the STAY class altogether, however in that case the model predictions can become unstable, as even tiny or zero

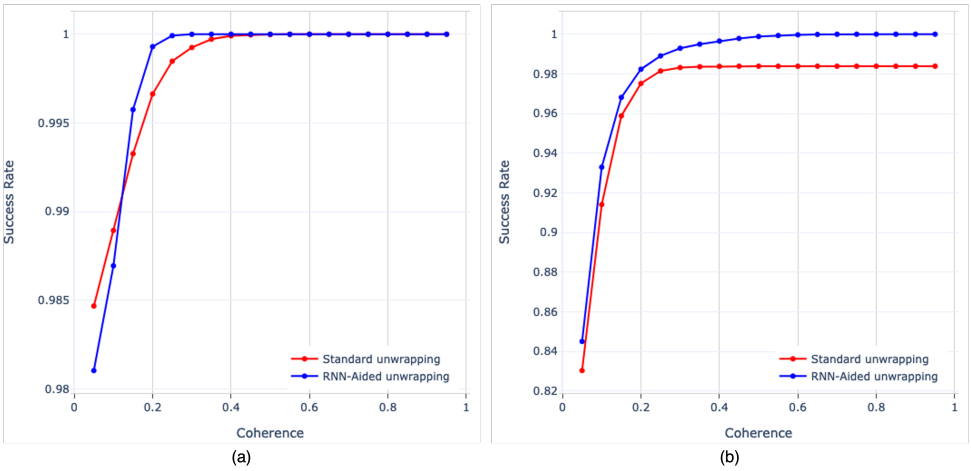

Figure 4.8: Algorithm performance curves of unwrapping success rate vs. coherence in (a) Rouveen and (b) Zegveld. Red: standard algorithm. Blue: RNN-aided algorithm. In the Zegveld case, the standard algorithm cannot exceed a success rate of approx. 98%. This means that at least one phase unwrapping error can be expected for a time series containing 50 epochs (just under one year of Sentinel-1 data).

deformations must be given an UP or DOWN label by the model. We find that including a STAY class corresponding to small deformations (< 3 mm) provides more accuracy in making the distinction between large UP/DOWN events, which we choose to prioritize.

Previous InSAR studies of the Dutch peatlands which have attempted to directly monitor the ground via DS processing [19] have failed to match the observations of other geodetic techniques such as levelling, lidar surveys, and extensometer measurements, and often strongly overestimate the true linear subsidence rate. The reason for this (as discussed in Sec. 4.1) is due in large part to conventional phase unwrapping techniques being unable to cope with the rapid movements which occur within the shortest temporal baselines, as shown in Fig. 4.1, which results in estimates skewed downward, as shown in Fig. 4.7. By removing this downward skew, we hope to enable accurate monitoring of this region with DS InSAR for the first time.

The results of Fig. 4.8 show that the RNN can provide a significant performance gain by supplying the unwrapping procedure with more information. The deformation test scenario chosen based on the Zegveld extensometer measurements as shown in Fig. 4.1 shows a period of strong uplift in the autumn of 2020. When downsampling these measurements to 6-day intervals corresponding to the Sentinel-1 overpass cycle, there is at least one epoch in which the deformation phase exceeds half a cycle ($\lambda/4$) and the standard unwrapping procedures fail, even in the presence of little or no noise. On the other hand, the RNN-aided algorithm is able to consistently retrieve the correct time series. In the Rouveen case, both algorithms are able to achieve a 100% success rate (SR) given a high enough coherence level. However, the RNN-aided algorithm is able to achieve 100% SR at lower levels of coherence, which shows how integrating more information into the processing workflow will enable greater performance. Coherence itself can be viewed as a measure of the information content in an interferogram (or lack of coherence

as a measure of entropy). Thus, introducing new information to the algorithm creates an increase in the "*effective coherence*" of the interferogram. For example, in the Rouveen case, Fig. 4.8 shows an effective coherence gain of $\Delta\gamma = 0.175$ at SR=100% when moving from the standard unwrapping algorithm ($\gamma = 0.4$) to the RNN-aided algorithm ($\gamma = 0.225$).

One can also observe that the RNN-aided algorithm performs differently between locations, with better performance in Rouveen than in Zegveld. This is likely due to the fact that the available dataset in Rouveen is much longer than in Zegveld. The Rouveen dataset begins 1.5 years before the Zegveld one, which means that the RNN has more than twice as much training data available to it. Therefore, one may also infer that the potential for improvement as available datasets expand is very large, and that the performance of these models will continue to improve as more data becomes available.

4.6 Conclusion

Phase unwrapping is an underdetermined problem which is conventionally sidestepped by assuming that the correct solution is the one which corresponds to the smallest phase change between epochs. While this is often a reasonable approach, ground-based measurements have shown that this is not a good enough solution in certain regions, such as the Dutch peatlands. If InSAR is to be used to monitor ground deformation in these regions, a different approach is required. This necessitates the integration of additional information to help solve the phase unwrapping problem.

This chapter demonstrated how ground-based measurements can be used to create a model that anticipates large shifts in the ground level based on readily available environmental information such as precipitation and temperature. Predictions from this model are integrated into the phase unwrapping process by considering the system of wrapped phases and model predictions as an HMM. The relative probabilities of which direction the observed phases should be integrated into the time series can then be estimated using the Viterbi algorithm. While this scenario focused on a model to anticipate ground movement caused by changes in precipitation and temperature, this HMM framework can also be used in other cases in which a researcher would wish to integrate additional information into a temporal phase unwrapping application.

5

Loss-of-Lock

Why don't you just use L-band?

Reviewer 2

This chapter introduces the term loss-of-lock to describe a specific form of coherence loss which results in the breakage of an InSAR time series. Loss-of-lock creates a specific pattern in the coherence matrix of a multilooked distributed scatterer by which it may be detected. Along with identification, a new distributed scatterer processing methodology is introduced which is designed to mitigate the effects of loss-of-lock by introducing contextual data to assist in the time series processing. This methodology is of particular relevance to regions which suffer from severe temporal decorrelation, such as northern peatlands. The new method is applied to two subsiding cultivated peatland regions in The Netherlands which previously proved impossible to monitor using distributed scatterer InSAR techniques. Results show a very good agreement with in-situ validation data as well as spatial correlation between regions and the natural terrain.

Highlights:

1. *Loss-of-lock is defined and shown to be a common occurrence in distributed scatterer InSAR problems.*
2. *A novel distributed scatterer InSAR processing methodology is introduced with a strategy to mitigate the effects on loss-of-lock.*
3. *Initial results on test locations demonstrate the feasibility of the new methodology.*

Parts of this chapter have been published in *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 61, No. 5220911 [98].

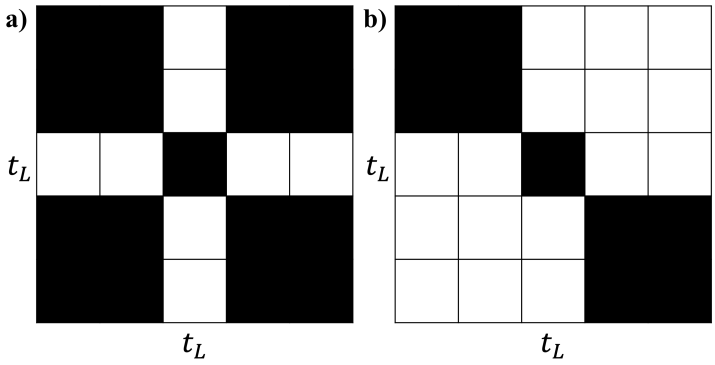

Figure 5.1: Canonical coherence matrices showing different types of coherence losses for a set of five subsequent SAR acquisitions. Shaded cells: interferometric combinations that are sufficiently coherent to produce a useful phase estimation. Empty cells: insufficiently coherent combinations. a) An intermittent loss of coherence at epoch t_L that does not produce a loss-of-lock, because coherent interferometric combinations exist which connect epochs preceding and following t_L. Archetype: intermittent snow cover. b) A loss of coherence resulting in a loss-of-lock. There are no sufficiently coherent interferometric combinations connecting the epochs preceding and following t_L. Archetype: plowing, harvesting.

5

5.1 Introduction

5.1.1 Definition of Loss-of-Lock

Interferometric coherence is a measure of how much information is available in an interferogram. While losses of coherence are a common phenomenon, we differentiate between intermittent losses of coherence, where an event results in the loss of one or more epochs, but the overall time series is unaffected (Fig. 5.1a); and *loss-of-lock*, a more serious loss of coherence which results in an irreparable discontinuity of the time series (Fig. 5.1b). A loss-of-lock is defined at a given epoch t_L such that a sustained loss of coherence is observed, and no coherent interferometric combination exists which connects observations across the epoch in question. At low sample coherences ($<$0.1), the distribution of the interferometric phase approaches a uniform distribution [64]. This means that all useful interferometric information (i.e. the displacement component of the interferometric phase) at that epoch is lost. Thus, the time-series is effectively cut at t_L, and the information content in the SAR image stack alone is not sufficient to estimate a connected set of interferometric phases spanning the entire observation period without additional information or assumptions.

Loss-of-lock is a diagnostic term in that it is defined based on conditions in an observed coherence matrix, rather than the occurrence of a particular scattering phenomenon, although the coherence losses are ultimately related to physical changes in the scattering object(s). For instance, a short-lived snowfall on an otherwise undisturbed and stationary grassland will result in an intermittent loss of coherence, while a loss-of-lock may be caused by agricultural activities such as plowing, or changes in vegetation such as harvesting of crops, resulting in a drastic reconfiguration of the scattering geometry of the ground, without implying any subsidence.

It is important to note that the presence of a loss-of-lock event may not be readily

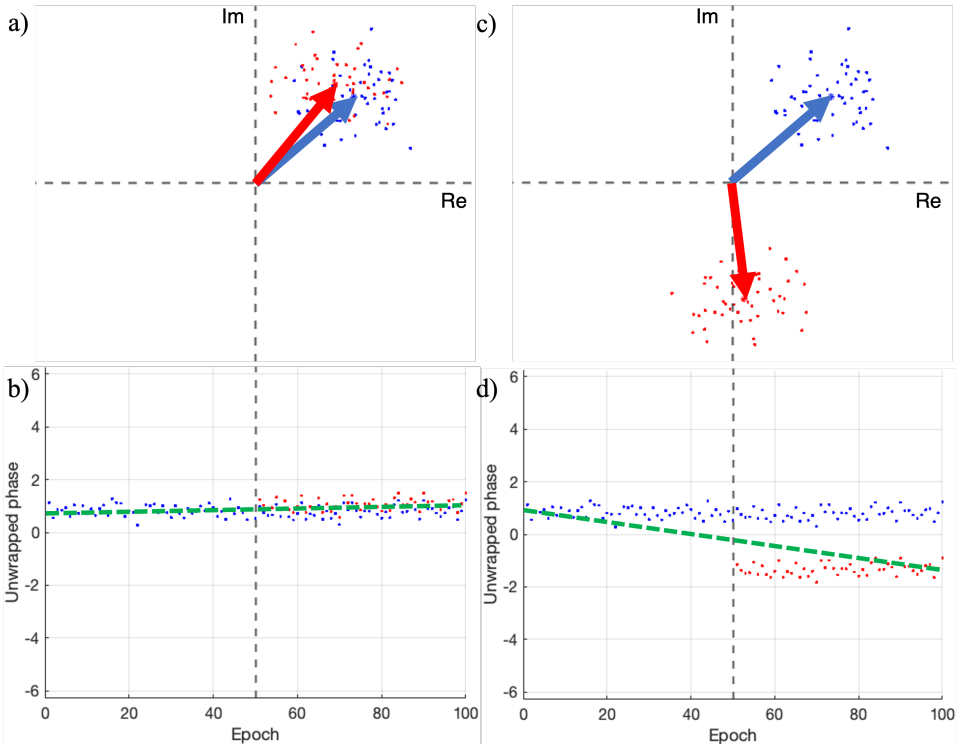

Figure 5.2: Two examples of a loss-of-lock which result in different interpretations, resulting in a perceived change in the displacement phase at epoch 50. Top: interferometric complex phasors, with arrows indicating the mean value pre- and post- loss-of-lock (blue and red, respectively). Bottom: resulting unwrapped phases. Blue dots: true (noisy) displacement phase of the ground level. Red dots: observed (noisy) displacement phase in the presence of loss-of-lock. Green dashed line: estimated linear velocity.

apparent from inspecting a displacement (or phase) time series. If one considers an event in which the mean surface level of the region under observation remains constant, but the scattering geometry changes drastically (for instance by ploughing), due to the wrapped nature of phase observations, the wrapped phase observation following the event may quite likely fall close to that of the previous epoch, and both phase unwrapping algorithms and manual inspection will overlook the change (as shown in Fig. 5.2 a) and b)). Alternatively, it is also possible for large phase differences to be observed, due to changes in the scattering surface which are then misattributed as displacements (as shown in Fig. 5.2 c) and d)).

Different sensors will be sensitive to different phenomena occurring on the ground and in the atmosphere, i.e. a loss-of-lock observed at C-band may not be observed at L-band. A more practical description of the phenomena affecting our study area is provided in Sec. 5.1.2.

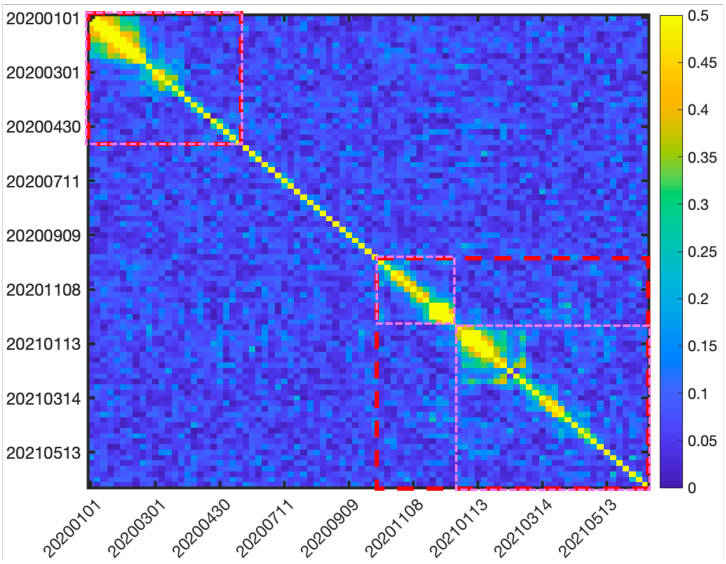

Figure 5.3: Observed coherence matrix showing loss-of-lock in Sentinel-1 ascending track 88 of a multilooked region near Zegveld, the Netherlands. The red and magenta dashed lines are added to indicate the disconnected coherent periods. When or if a loss-of-lock occurs depends on the minimum allowable coherence threshold (discussed in 5.1.2).

5.1.2 Observed Loss-of-Lock Events

A practical example of loss-of-lock is shown in Fig. 5.3, showing sample coherences for the period between January 2020 to May 2021. Two coherent periods (identified with the red dashed lines) are separated by a substantial amount of time, but more importantly, it can be seen that there is no significantly coherent interferometric combination linking them. This means that the two periods are effectively disconnected, and additional information will be required to estimate a time series spanning the entire observation period. Depending on the threshold used, an additional loss-of-lock event can be observed, resulting in three coherent periods, as indicated by the magenta lines. This shows that: i) detecting a loss-of-lock is dependent on the choice of allowed level of noise vs. amount of data used, and ii) a loss-of-lock can be both a sudden and/or a sustained condition. This choice of threshold is discussed further in Sec. 5.2.5.

5.2 Distributed Scatterer Processing Methodology

5.2.1 Overview

A high-level end-to-end process flow diagram is given in Fig. 5.4. The system makes use of two well-established InSAR software packages, Delft Object-oriented Radar Interferometric Software (DORIS) [99] and Delft Persistent Scatterer Interferometry (DePSI) [54] (blue and red sections of Fig 5.4, respectively). DORIS is used to align, resample, and geolocate the level-1 single-look complex (SLC) SAR image stack. DePSI is used to create a network of point scatterers (PS) and estimate an atmospheric phase screen (APS), which

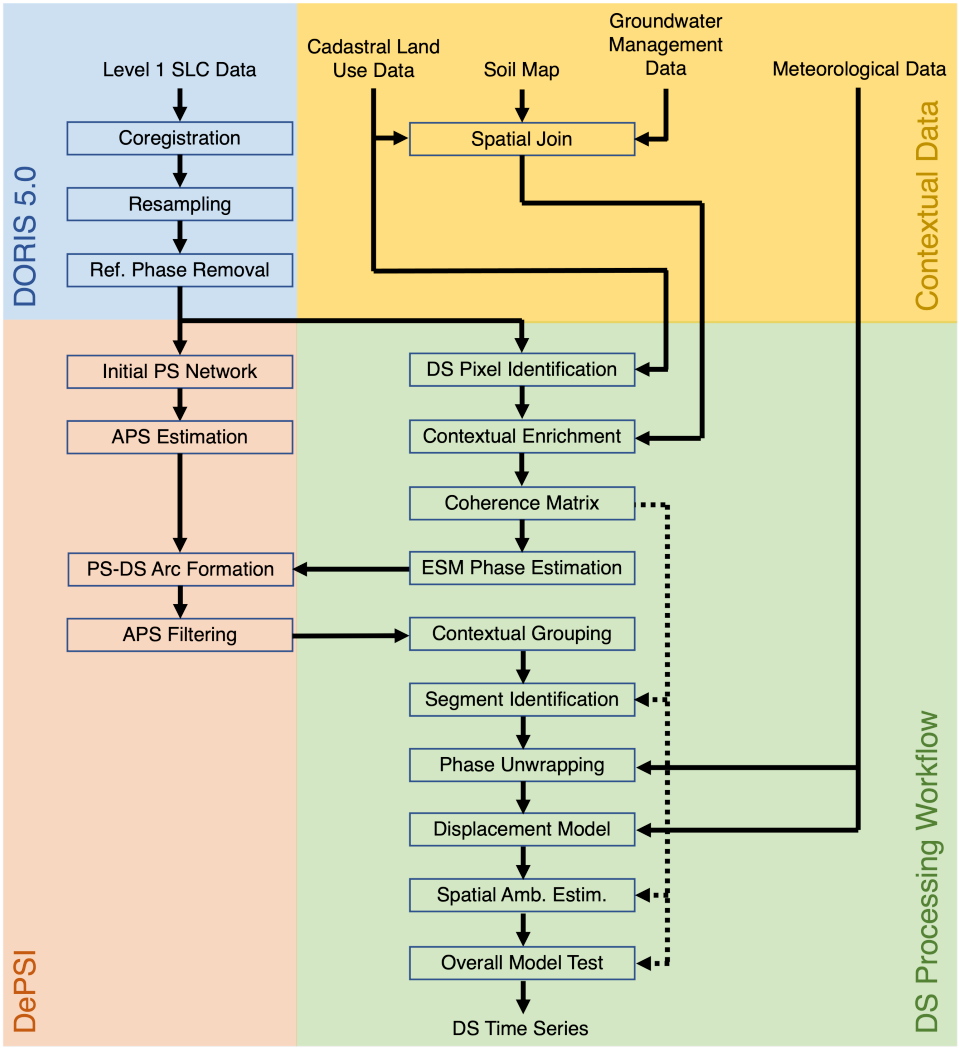

Figure 5.4: Simplified process flow diagram showing the major steps taken to create DS time series estimates from Sentinel-1 Level 1 SLC SAR data with the aid of spatial and temporal contextual data. DECADE v0 software implementation.

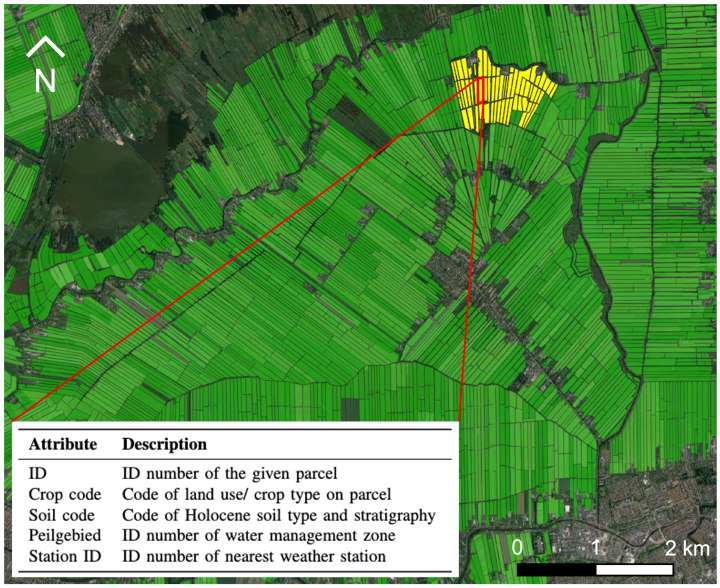

Figure 5.5: Graphical visualization of spatial contextual data in QGIS, based on land parcel polygons of the region surrounding Zegveld, The Netherlands. The attributed polygons are shown in green over a background optical satellite image of the region. A parcel of interest is highlighted with a red border, and the corresponding contextual group is highlighted in yellow.

can be applied to both the point scatterer (PS) and distributed scatterer (DS) phase observations. The remaining parts of this section are dedicated to describing the "Contextual Data" and "DS Processing Workflow" sections of the diagram (yellow and blue sections of Fig 5.4, respectively), which are the novel aspects of this methodology. This method has been implemented in a software package called Delft Contextually-Aided Distributed Scatterer Environment (DECADE).

5.2.2 Spatial Contextual Data and DS Pixel Identification

Three publicly available spatial datasets are combined via a spatial join operation: cadastral land use (parcel) polygons, soil maps, and groundwater management zones (Dutch: *peilgebied*) [100]. The datasets are provided as vector geometries in geopackage format, allowing for a straightforward combination of the data. This is accomplished by taking the land parcel delineations in the cadastral dataset as the base layer for the spatial join and performing a one-to-one attribution with the features with the largest overlap. While the cadastral and groundwater management zones follow similar geographic boundaries, the soil map has a different spatial structure and this one-to-one attribution results in some information loss. For instance, a parcel mostly composed of peat with a smaller vein of sand running through it will simply be labelled with the peat soilcode. This is done to constrain the problem variables and have only one value per region. This method could be extended for regions with larger or more heterogeneous parcels by subdividing them either geographically or based on other relevant contextual information. Finally, the near-

est weather station is found by Voronoi polygonization, and its corresponding ID is added in order to enable attribution of (temporal) meteorological data. An example of such a combined dataset is shown in Fig. 5.5.

DORIS provides a coregistered SLC stack along with the geolocation of each pixel in a grid. Each pixel in the stack can therefore be assigned an ID corresponding to the polygon it belongs to (Fig. 5.5). Each polygon is assigned a coordinate according to its centroid.

5.2.3 Coherence Matrix and ESM Phase Estimation

Multilooking is performed on a per-polygon basis. As can be seen in Fig. 5.5, the Dutch peatlands are divided into large rectangular parcels surrounded by drainage ditches, which provides us with a natural set of multilooking boundaries. While parcel sizes vary in shape and size, in general, the groundwater level and land cover within a parcel is consistent, atmospheric delay variability will be negligible (at the sub-mm level) [33], [55], and a parcel will typically contain 100 pixels. A minimum number of 50 pixels per polygon is enforced for noise suppression. Thus we are able to ensure ergodicity and representativity while maximizing the number of equivalent looks. Following this parcel selection, we also optionally apply a statistically homogeneous pixel (SHP) test as outlined in [58]. This can filter out misattributed pixels due to geolocation errors in the radar and contextual data, as well as the effects of unwanted scatterers within the region, such as electrical masts, light posts, trees, etc.

The complex sample coherence matrix of a multilooked region, $\hat{\Gamma}$, consisting of elements $\hat{\gamma}_{ij}$ is given by

$$|\gamma_{ij}| = \frac{\left|\sum_{n\in\Omega} S_i^n S_j^{n*}\right|}{\sqrt{\left(\sum_{n\in\Omega} |S_i^n|^2\right)\left(\sum_{n\in\Omega} |S_j^n|^2\right)}}, \tag{5.1}$$

where $S_{i,j}^n$ is the nth single-look complex (SLC) SAR pixel nth pixel acquired at epochs i and j, the asterisk denotes the complex conjugate, and Ω is the set of all selected pixels within the multilooked region. We differentiate between the complex coherence matrix Γ, and the coherence matrix, $|\Gamma|$, which is the matrix of the magnitudes of the elements of Γ. An example of $|\hat{\Gamma}|$ of a multilooked parcel is shown in Fig. 5.3.

An Equivalent single master (ESM) [101] set of phases is estimated using the "EMI" method, as described in [29]. This procedure reduces the full set of all interferometric combinations to a single set of consistent phases, $\hat{\phi}_{\text{esm}}$, as estimated by the phase of the minimum eigenvector of the Hadamard product of the inverse of the sample coherence matrix with the complex sample coherence matrix, as given by

$$(|\hat{\Gamma}|^{-1} \circ \hat{\Gamma})\xi = \lambda\xi, \tag{5.2}$$

where $\circ$ denotes the Hadamard product, λ is the eigenvalue, and ξ is the eigenvector. The estimated interferometric phases are given by

$$\hat{\phi}_{\text{esm}} = \arg\{\xi\}. \tag{5.3}$$

Strong decorrelation can hinder the effectiveness of the ESM phase estimation. In cases in which coherence is completely lost during a loss-of-lock event, it may be advantageous to only perform the estimate within the identified coherent blocks. This can reduce the amount of noise at the input to the estimator, however, one risks losing useful long-term coherent information (see Sec. 6.2.1). The decision to perform block-wise estimation could also be driven by contextual data, i.e. a-priori information about the land use/cover which indicates that a loss-of-lock has occurred, such as knowledge of plowing or harvesting events.

The estimated ESM phases $\hat{\phi}_{\text{esm}}$ per polygon are then imported into DePSI as virtual points into the secondary network of scatterers in order to apply atmospheric phase screen (APS) filtering. The locations of the virtual representative points are given by the centroids of the given parcel polygon. The APS estimation is based on an initial primary network of PS's. The filtered phases $\hat{\phi}_{\text{aps}}$ are read back out of DePSI following the APS filtering stage.

5.2.4 Contextual Enrichment and Grouping

We have now obtained a set of wrapped, multilooked and filtered DS phases which are each characterized by the set of attributes shown in Fig. 5.5, along with a point coordinate given by the polygon centroid. The parcels and their estimated phases are grouped together according to their shared attributes, establishing *contextual groups*. We contend that parcels which share the same land use, soil classification, and belong to the same groundwater management regime should be expected to behave in a similar fashion. That is, although we expect to see variations in phase according to differing noise and clutter conditions, local variations in soil stratigraphy, and variations in the optical depth due to land cover, we expect that the parcels in a contextual group can be expected to move according to the same displacement model in the mean sense. This grouping becomes critical in the context of bridging loss-of-lock in the parcel time series, which is described in Secs. 5.2.5, 5.2.6, and 5.2.7. The contextual group corresponding to the red highlighted parcel of Fig. 5.5 is indicated in yellow. The identified contextual groups are then filtered by their number of members: we have found a minimum of 30 members is needed to ensure sufficient coverage throughout the year, however, this value will change depending on the coherence behaviour of the area under investigation.

5.2.5 Segment Identification

Due to the loss-of-lock phenomenon, attempting to interpret the entire ESM time series of phases at once is not possible and will result in several types of error, such as interpreting a noise-dominated signal as real deformation, or phase unwrapping errors when transitioning from incoherent to coherent interferograms [19]–[23]. Thus, a different approach is required.

We begin by identifying which parts of a time series are of sufficient quality that they contain physically interpretable information which can be unwrapped with an acceptable degree of error. Despite using a full-rank method to estimate the ESM phases (Sec. 5.2.3), we find that the best quality indicator we have available is the so-called daisy-chain coherence, $|\hat{\gamma}_{\text{dc}}|$, which is the magnitude of the first off-diagonal of the coherence matrix (corresponding to the indices $j = i - 1$ in Eq. (5.1)). These are the coherence magni-

tudes of the interferograms with the shortest temporal baseline in the dataset, which for Sentinel-1 data is six days. In general, we expect these to be the most coherent interferograms in the dataset, as less time has passed for decorrelation effects to occur [26], while orbital baselines for Sentinel-1 are always small, resulting in negligible baseline decorrelation [33]. We threshold the daisy-chain coherence to identify sufficiently coherent subsections of the full time-series, which we term (temporal) *segments*. A segment is a contiguous subset of a time-series in which the coherence is sufficiently high to estimate a consistent set of interferometric phases. Thus a segment is defined by two thresholds: the minimum coherence, and the minimum number of consecutive coherent epochs, which can be determined experimentally. In our case, a minimum of five consecutive epochs with $||hat\gamma_{\text{dc}}| > 0.12$ is used as a threshold.

Each contextual group therefore contains many coherent segments: one for each contiguously coherent period of a parcel, times the number of parcels, times the number of satellite tracks covering the area of interest (AOI), which can be combined. By considering such a large number of segments, we are able to span loss-of-lock events in one parcel time series with coherent observations from a neighbouring one from the same contextual group.

5.2.6 Temporal Ambiguity Resolution

The identified segments are initially treated as independent time-series. Temporal phase unwrapping (or ambiguity resolution) is performed independently on each segment using a method aided by a machine learning model, as described by [75] (Ch. 3). The ground surface level of peatlands is extremely unstable and prone to rapid fluctuations depending on temperature and precipitation levels, so we use a recurrent neural network (RNN) to aid in making predictions about which ambiguity level is correct. This RNN model uses temperature, precipitation, and day of year as inputs, which is publicly available daily weather data. Detailed information about the implementation and testing of the methodology is provided in [75] (Ch. 4).

5.2.7 Displacement Model

We have now obtained a collection of temporally unwrapped segments, which are internally consistent but disconnected from one another by an unknown vertical shift, Δz, which represents the unknown displacement history of the DS during the loss-of-lock period. Thus, in order to recombine the coherent segments, this unknown shift must be estimated. This can be accomplished with the aid of a displacement model, which can be used to align all the segments of a contextual group.

A parametric model which relates precipitation and evapotranspiration to soil surface displacement at a particular location has been developed in [76], and is detailed in Ch. 3. The modeled soil surface height, z_M, is a function of precipitation, evapotranspiration, and the Holocene stratigraphy at the modeled location. It is combination of reversible processes, such as shrinkage and swell, and irreversible processes, such as soil oxidation:

$$z_M(x,t) = R(x,t) + I(x,t), \tag{5.4}$$

where the model is parameterized by the lithology dependent unknowns in x and t is time in days. R represents the reversible component and I is the irreversible component

of the relative soil surface position. The reversible component is estimated by the scaled cumulative difference between precipitation and evapotranspiration:

$$R(x,t) = \sum_{t'=t-\tau}^{t} [x_P \cdot P(t') - x_E \cdot E(t')], \tag{5.5}$$

where x_P and x_E are empirical scaling factors, and τ is the integration time. P is daily mean precipitation [mm], and E is daily mean reference evapotranspiration [mm] [84]. Daily values for P and E are provided at every weather station in the Netherlands [102].

The irreversible component is approximated as a linear rate, which is only considered active when R is negative, indicating drying soil conditions:

$$I(x,t) = \sum_{t'=t_0}^{t} x_I \cdot f(x,t), \tag{5.6}$$

where x_I is a constant, and

$$f(x,t) = \begin{cases} 0, & \text{for } R(x,t) > 0 \\ 1, & \text{for } R(x,t) \leq 0. \end{cases} \tag{5.7}$$

Thus, the model is parameterized by the four unknowns

$$x = [x_P, x_E, x_I, \tau]. \tag{5.8}$$

These parameters depend on the depth and stratigraphy of the Holocene sequence at a given location, i.e. the lithology of that location. Details on the validation of the model is provided in [76]. For the test locations shown in Sec. 3.5.3, the RMSE of the model with validation data is 6.9 mm in Zegveld, and 4.1 mm in Rouveen.

Now we will show how to accurately estimate these model parameters, given the sparse unwrapped measurements we have available. This result can then be used to align the unwrapped segments of a given contextual group and estimate a continuous displacement time series. The relationship between the unwrapped ESM phases of the nth segment of a given DS polygon, $\phi_{\text{esm},n}$, and the group displacement model, z_M, is given by

$$\phi_{\text{esm},n}(t) = \frac{-4\pi \cos\theta_{\text{inc}}}{\lambda} \cdot [z_M(x, P(t), E(t)) + \Delta z_n] + \epsilon; \quad \forall\ t \in T_n, \tag{5.9}$$

where t is time, θ is the incidence angle, λ is the wavelength, Δz_n is the unknown vertical shift (constant for a given segment), ϵ is a combination of noise, phase unwrapping errors and model residuals, and T_n is the set of all epochs in the nth coherent segment. Eq. (5.9) cannot be solved in its current form, as the model parameters x must be known *a-priori* in order to evaluate the correct Δz. While they can theoretically be estimated simultaneously, the high degree of correlation between these unknowns can result in a very poor estimation. Instead, we note that Δz is common for all phases within a given

segment. Thus by taking the difference in time between phases (i.e. the equivalent daisy-chain (EDC) phase), the Δz term drops out and the model parameters x can be estimated directly by solving

$$\phi_{\text{edc},n}(t) = \frac{-4\pi \cos\theta_{\text{inc}}}{\lambda} \cdot \Delta z_M(x, P(t), E(t)) + \epsilon_\Delta; \quad \forall\ t \in T_n. \tag{5.10}$$

Now that the model parameters have been estimated, the Δz for each coherent segment can subsequently be estimated by taking the average difference between the model and the ESM phase time series over the coherent period T_n:

$$\Delta\hat{z}_n = \left\langle \frac{-\lambda}{4\pi \cos\theta_{\text{inc}}} \cdot \phi_{\text{esm},n}(t) - z_M(\hat{x}, P(t), E(t)) \right\rangle; \quad \forall\ t \in T_n, \tag{5.11}$$

where $\hat{x}$ are the estimated model parameters, and $\langle\cdot\rangle$ denotes averaging. This process is repeated for each contextual group described in Sec. 5.2.4, so there is one model for every identified contextual group.

This method can also be used to align the phase observations of multiple satellite tracks together, provided there is no significant horizontal motion, or else that the vertical component of the displacement phase can be accurately estimated, and that care is taken to ensure that the same object is used as a reference point across all tracks.

5.2.8 Spatial Ambiguity Resolution

The typical approach to DS InSAR processing involves applying a minimum cost flow spatial unwrapping algorithm to the data, such as the well-known SNAPHU algorithm [71]. However, this approach is not well-suited to peatland observations due to rapid soil movements and the high degree of multilooking required [21], [23], [75]. Heterogeneity in both the type and depth of the soft soil layer of the Holocene will result in different responses to the seasonal weather conditions which the ground is exposed to, leading to spatial differences in the seasonal amplitude of the reversible displacement. When combined with high degrees of multilooking, this can create sharp discontinuities in the downsampled interferogram, which will essentially lead to aliasing if strong spatial continuity constraints, such as those in SNAPHU, are applied [103].

For these reasons, direct spatial comparison of phase changes between adjacent parcels is an error-fraught process, and could result in introducing additional phase unwrapping errors instead of improving the result. We therefore take advantage of the mean displacement of the contextual group, which is the best estimate of how the contextual group of parcels should behave on average. Since a time series for the expected mean behaviour of the contextual group has already been estimated, it is now a straightforward process to apply integer least-squares (ILS) or an integer bootstrapping estimation[73], [104] to refine the estimated ambiguities of each DS polygon belonging to the contextual group. The ambiguities are estimated by first obtaining a *float* solution, given by

$$\hat{a} = (A^T Q_y^{-1} A)^{-1} A^T Q_y^{-1} (\hat{\phi}_p - \hat{\phi}_{\text{group}}), \tag{5.12}$$

where $\hat{a}$ are the real-valued float ambiguity corrections, A is 2π times the $n \times n$ identity matrix, $\hat{\phi}_p$ is the vector of unwrapped interferometric phases of the pth DS polygon,

and $\hat{\phi}_{\text{group}}$ is the vector of mean unwrapped phase of the entire contextual group. Q_y is the variance-covariance matrix of the phase observations, and can be approximated by the Cramer-Rao bound (CRB) [20]. Thus, the covariance between two interferometric phases, ϕ_{ij} and ϕ_{kl} is approximated by

$$\text{Cov}\{\phi_{ij}, \phi_{kl}\} \approx \frac{|\gamma_{i,k}||\gamma_{j,l}| - |\gamma_{i,l}||\gamma_{j,k}|}{2L \cdot |\gamma_{i,j}||\gamma_{k,l}|}, \tag{5.13}$$

where L is the effective number of looks [33], and $|\gamma|$ is the magnitude of the sample coherence (as determined by Eq. (5.1)) of the given interferometric combination, as indicated by the epoch subscripts i, j, k and l. Next, integer bootstrapping [105] is be applied, which provides the most likely integer ambiguities as

$$\breve{a} = \begin{bmatrix} \breve{a}_1 \\ \breve{a}_2 \\ \vdots \\ \breve{a}_n \end{bmatrix} = \begin{bmatrix} [\hat{a}_1] \\ [\hat{a}_2 - l_{21}(\hat{a}_1 - \breve{a}_1)] \\ \vdots \\ [\hat{a}_2 - \sum_{i=1}^{n-1} l_{ni}(\hat{a}_i - \breve{a}_i)] \end{bmatrix}, \tag{5.14}$$

where $[\cdot]$ is the rounding operator, $\breve{a} \in \mathbb{Z}^n$ is the vector of estimated integer ambiguities, and l are the entries of a lower triangular matrix L obtained by decomposing the matrix Q_y^{-1} into L and a diagonal matrix D, such that $LDL^T = Q_y^{-1}$.

5.2.9 Overall Model Test

Finally, a quality check is performed on the estimated contextual group results to ensure reliability. Overall statistics of the estimated contextual group parameters are generated for the entire AOI, and groups are flagged whose parameters deviate significantly (i.e. greater than 2σ). The unwrapped parcel phases of the member parcels are compared to the estimated group model in flagged groups in which it is suspected that the contextual group model has been poorly estimated by means of an overall model test (OMT) to assess the level of agreement between the overall group model and the constituent parcels.

The OMT is performed by comparing the model residuals $\hat{e}$ to the estimated precision of the observations Q_y to generate the test statistic T for each DS polygon:

$$T = \hat{e}^T Q_y^{-1} \hat{e}, \tag{5.15}$$

for a time series with N total epochs, the nth element of $\hat{e}$ is given by

$$\hat{e}_n = \Delta\phi(t = t_n) - \frac{-4\pi \cos\theta}{\lambda} \cdot \Delta M(\hat{x}, t = t_n). \tag{5.16}$$

The operator Δ refers to the fact that we use the differential daisy-chain phase as defined in Eq. (5.10) in order to remove the estimated vertical displacement shifts (the displacement occurring during the loss-of-lock periods) from the equation.

The test statistic T follows a central chi-squared distribution with four degrees of freedom, corresponding to the four unknown model parameters (Eq. (5.8)), and is compared to a critical value which follows from a chosen significance level α. If T exceeds the critical value, then the model does not follow the observations to within the estimated

precision of the observations at that significance level. In our case, the precision estimation comes from the CRB, which is the theoretical lower bound on the best achievable uncertainty. Thus while it is correlated with the true uncertainty, estimating the CRB based on the sample coherence (Eq. (5.13)) will systematically overestimate the uncertainty of the phase observations. Therefore the significance level is chosen more strictly to compensate for this.

The OMT is performed recursively on flagged groups by choosing an initial α and removing points which are rejected by the test. The model parameters of the contextual group are then re-estimated with the rejected points removed. If the new model parameters fall below the acceptable threshold then the group is sustained. If the parameters still deviate, α is slowly decreased and the procedure performed again. If after several iterations (ex. 5) the estimated model parameters still fall outside the accepted bounds, it is concluded that the model is not suitable for the terrain in question, and the group result is discarded. In a multiple hypothesis testing context, this procedure could be reiterated with an alternative model.

5.3 Results

5.3.1 Description of Satellite Data Used

Sentinel-1 imagery of two 10×10 km regions around Zegveld and Rouveen, The Netherlands, are used as test sites for the time period spanning Jan. 2017 - Dec. 2022. In Zegveld, four tracks are used: ascending 088 and 161, and descending 037 and 110. In Rouveen, three tracks are used: ascending 015 and 088, and descending 037. The unwrapped segments of all available tracks are combined (as discussed in Sec. 5.2.7) by projecting them onto the vertical after ensuring that all phases are referenced to the same object. The common reference point is found by identifying common PSs, i.e. points that are visible in all tracks and located on the same object. The PS with the lowest total normalized amplitude dispersion across all tracks is chosen as a common reference.

5.3.2 Multilooking Based on Contextual Data vs SHP Test Only

A comparison between a standard multilooking approach which employs 300×300 m regions and the parcel-based multilooking approach is shown in Figs. 5.6 a) and b), respectively. As can be seen in the standard approach, despite the use of an SHP test, pixels from a number of objects which we do not expect to behave the same way are still averaged together. This is particularly apparent in the NE and NW corners of the image, where agricultural fields, residential yards, and greenery along roadways are all grouped together.

By including parcel cadastral information, we can help ensure that we are indeed averaging pixels which belong to the same objects or regions. An SHP test can also still be applied to remove unwanted pixels from within the parcel boundaries.

5.3.3 Coherent Segment Identification and Commonalities

Fig. 5.7 illustrates the advantage of grouping similar parcels together into contextual groups. While almost all regions provide sufficiently coherent (i.e. $|\gamma_{\text{dc}}| > 0.12$) data over the winter period, from approximately October to April, the coherence of most re-

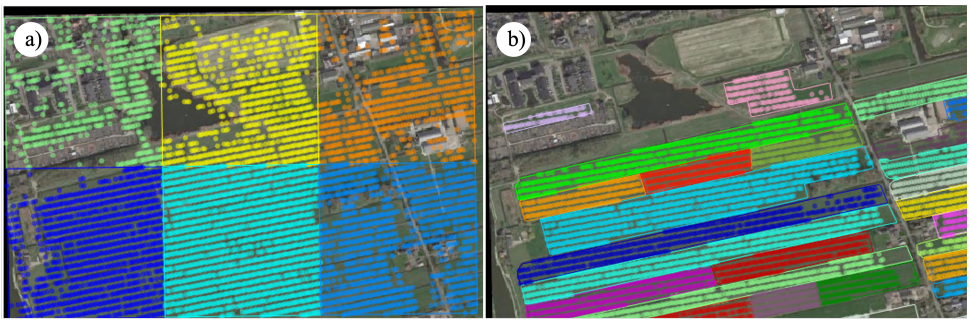

Figure 5.6: Comparison between two multilooking strategies of the same area, colour-coded by region. Solid lines: boundaries of a given multilooking region. Circles: pixels included in the multilooking. a): Standard square multilooking procedure using 300×300 m areas. b): Parcel-based multilooking.

gions drops significantly in the spring and is only intermittently present throughout the summer period until the following October. However, by combining the observations of enough similar parcels, we are able to have a year-round set of data with which to estimate the parameters of the displacement model as per Eq. (5.10). The coherence threshold of $|\gamma_{dc}| > 0.12$ was experimentally found to be the highest value which still ensured sufficient data coverage year-round.

It is interesting to note that there is both a systematic and a random aspect to the coherence behaviour of these regions. A systematic loss of coherence from April to October is clearly visible in the majority of parcels, however, the exact timing of this loss, as well as the intermittent recovery of coherence during the summer, seems to be a random event. This distribution is visualized by the shading of the background of Fig. 5.7. While it is clear that losses of coherence in these regions are caused by agricultural activities such as mowing and grazing, as well as changes in the scattering properties of the medium [26] caused by the drying of the soil and vegetation over the summer periods, it is unclear why some parcels seem to show higher coherence levels than others from the same contextual group at the same moments in time. This may be caused by some fields being used more intensively for agriculture than others, for instance differences in the level of grazing between various fields.

5.3.4 Time Series Estimation

An example group time series result is given in Fig. 5.8. This result demonstrates how the displacement estimates of several temporarily coherent regions can be combined together to produce an unbroken time series of the overall region. The result matches very well with the available in-situ validation data. Note that the validation data is not available for the entire span of the time series due to their installation dates. The difference between the contextual group median result and the validation data is quantified by the root mean square difference (RMSD) in Table 5.1. However, it should also be noted that we do not expect an exact match between the InSAR and ground-based results, because the InSAR result shows the average behaviour over a large spatial extent, whereas the ground-based measurement is of a single point. Moreover, the ground-based results do not capture

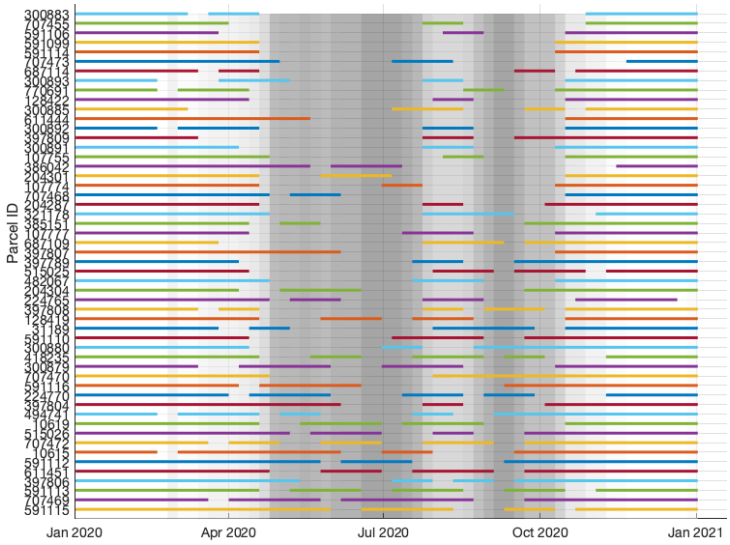

Figure 5.7: Chart showing the availability of coherent data over time for a period of one year for a selection of parcels belonging to the same contextual group. The y-axis indicates the ID number of a given parcel, and the presence of a solid line indicates the presence of sufficiently coherent data. The background is shaded to indicate the relative degree of availability (i.e. the number of coherent parcels divided by total number of parcels) such that a white background indicates complete availability with darker shading as availability decreases.

Table 5.1: Difference between InSAR and extensometer estimates

Location	Zegveld	Rouveen
Parcel RMSD (mm)	7.9	6.6
Group RMSD (mm)	6.9	5.3

the influence of the top five centimeters of soil, due to the position of the extensometer anchors. Nevertheless, as the major factors driving the motion are the same for both cases, we see that the agreement between them is very close, particularly in the observed short-term dynamics.

5.3.5 Effective Number of Looks Over Time

An important factor governing the quality of the result is the effective number of looks [33], shown in Fig. 5.9. This number fluctuates throughout the year due to the availability of coherent segments in the contextual group, as discussed in Sec. 5.3.3. It is important to ensure that there remain enough coherent observations during the periods in which most regions are incoherent. If too few coherent observations are present, then the overall contextual group result can become biased by the behaviour (and noise) of only a few pixels. The effective number of looks L used at a given time is given by

$$L = (\text{No.segments}) \times (\text{No.pixels/segment}) \times \text{OSR}, \tag{5.17}$$

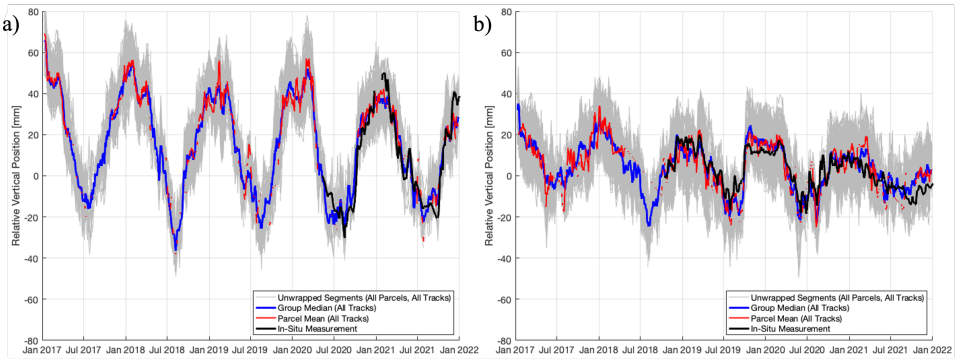

Figure 5.8: Surface level time series results plotted against in-situ extensometer measurements for the period Jan. 1, 2017 – Dec. 31, 2022 at the a) Zegveld and b) Rouveen regions. Grey lines: all segments of all parcels belonging to the contextual group. Blue line: contextual group median time series. Red line: Mean of time series segments of a selected parcel in the contextual group. Black line: in-situ measurement by extensometer of the same parcel. All the individual coherent segments belonging to the contextual group are shown in grey for readability.

5

where OSR is the oversampling rate given by

$$\mathrm{OSR} = \frac{\mathrm{PRF}}{\mathrm{BW_{az}}} \cdot \frac{f_s}{\mathrm{BW_r}} \tag{5.18}$$

where PRF is the pulse repetition frequency, f_s is the range sampling frequency, and BW_{az} and BW_{r} are the azimuth and range bandwidths, respectively.

5.3.6 Estimated Linear Rates

Approximate linear subsidence rates are shown in Fig. 5.10. These rates are estimated by linear regression of the contextual group mean time series results shown in Sec. 5.3.4, however, it should be noted that the total length of the observation period (five years) is too short to establish a robust estimate of the rate. Thus, these results provide an order of magnitude estimate and can be used to assess the spatial distribution of subsidence in the area.

5.4 Discussion

5.4.1 On the Absence of Contextual Data

It is quite likely that in some cases, additional contextual data may not be available for the region under investigation, for instance in peatland regions in remote locations. In such a case, additional remote sensing data may be integrated into the processing workflow in order to identify and group common pixels together, such as the SAR backscatter data, as is done in the established squeeSAR [58] method, or through the use of semantic segmentation techniques on co-located optical or hyperspectral imagery. A-posteriori techniques such as t-SNE [106], [107] can potentially be used to group similarly behaving scatterers together into contextual groups.

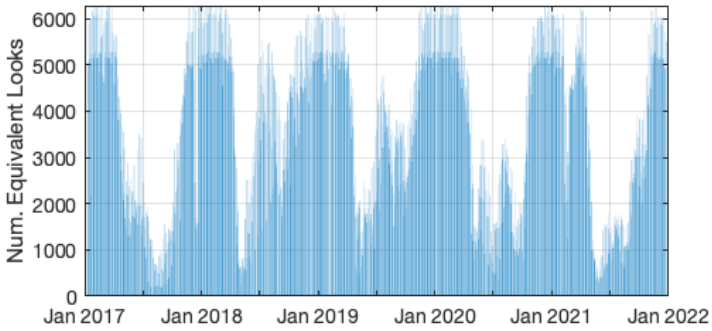

Figure 5.9: Chart showing the effective number of looks over time for the period Jan. 1, 2017 – Dec. 31, 2022 for the same contextual group shown in Fig. 5.8a.

5.4.2 Model Reliability and Goodness of Fit

Steps are taken to ensure goodness of fit and reliability of the estimated model with the overall model test (Sec. 5.2.9). The T-score of a given parcel (eq. (5.15)) is shown in Fig. 5.11, however whether or not that parcel is used in the final rate estimation depends on the procedure outlined in Sec. 5.2.9. When comparing Fig. 5.10 to Fig. 5.11, it can be seen that some parcels with a high T-score, and therefore a poor agreement with the contextual group model, are discarded from the final result. These are the ones detected by the iterative testing procedure. Other parcels with high T-scores are flagged for further re-evaluation but are not discarded immediately because their corresponding groups fall within expected bounds. There are several main causes of error which make a parcel deviate from the estimated contextual group model:

- Misattribution within the contextual dataset: for example, errors in the soil map, or incorrect land use classifications
- Phase unwrapping errors
- Noise and decorrelation
- Model parameter estimation errors

Often these causes are correlated; a misattributed parcel may be grouped with a set of other parcels in which it should not belong, and introduce error into the contextual group model estimation. One region where this is evident is in a group of central-northern parcels (approx. coordinates: 52.62°N, 6.13°E) in the Rouveen area. Although classified as grassland, it is in fact a large rewilded “Natura-2000” region. Some of the parcels in this region are more densely covered with vegetation as opposed to being simple grasslands. This means that the phase behaviour in these parcels is possibly different from the surrounding areas, and in some cases the estimated displacement model may not be valid there. The OMT procedure is able to identify this and re-estimate a valid model with the remaining parcels not discarded by the test. A similar situation is visible in the SW corner of the Zegveld region (approx. coordinates: 52.095°N, 4.75°E).

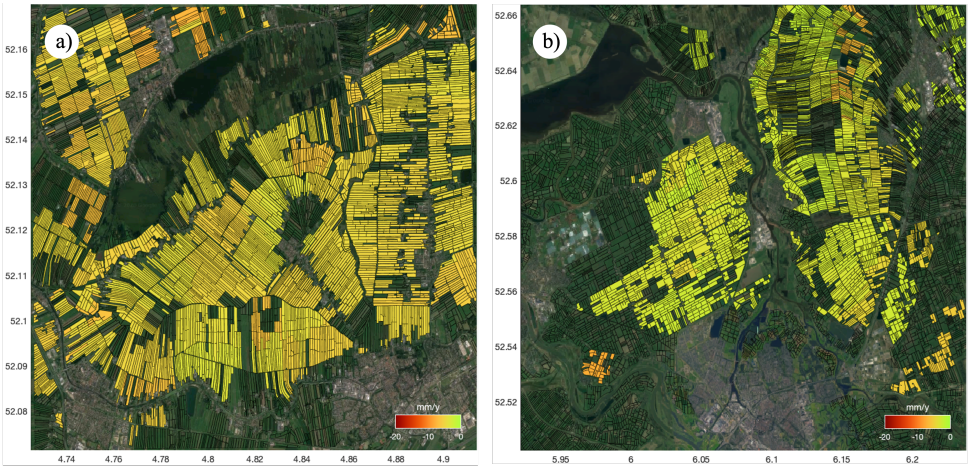

Figure 5.10: Colour-coded map of estimated linear subsidence rates for the period Jan. 1, 2017 – Dec. 31, 2022 at the a) Zegveld and b) Rouveen regions. Parcels with no fill indicate that no estimation has been made at that location.

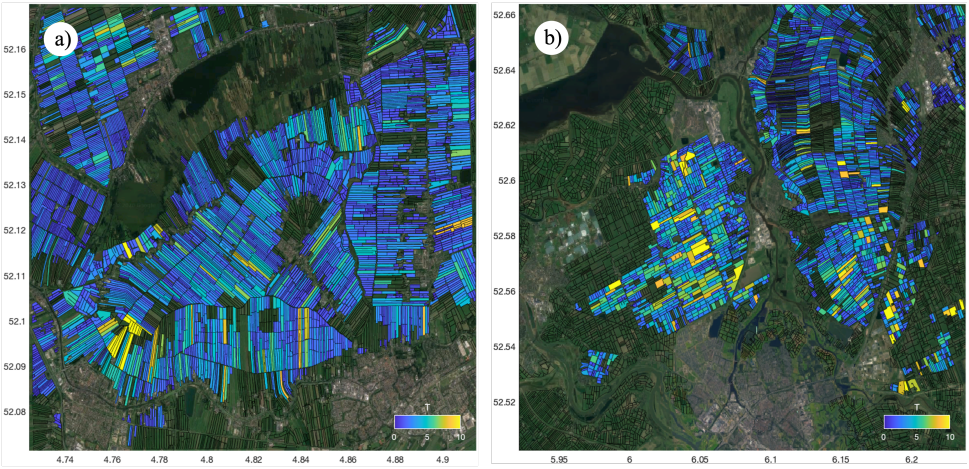

Figure 5.11: Colour-coded map of estimated goodness of fit according the T-score value at the a) Zegveld and b) Rouveen regions. Parcels with no fill indicate that no estimation has been made at that location.

5.4.3 Mean Displacement Model versus Mean Phase Change

The use of a mean displacement model is a choice which in theory could be omitted. One could for example simply take the mean of the daisy-chain differential phase $\phi_{\text{dc}}(t)$ of the entire contextual group and integrate it to obtain a relative position time series. However in that case, one becomes very dependent upon a select few sets of measurements during the low-coherence times, as shown in Figs. 5.7 and 5.9. Any biases, noise or phase unwrapping errors in these observations would then be directly propagated into the mean contextual group time series. Therefore, using the set of all $\phi_{\text{dc}}(t)$ observations to estimate a set of global model parameters is a safer option, provided the model is valid for the contextual group.

5.5 Conclusion

Loss-of-lock is a permanent loss of coherence between two or more parts of a time series which is impossible to repair using the data in the SAR image stack alone. While decorrelation is a topic that has been discussed at length in the past, the specific implications of a loss-of-lock event are not well understood nor has a name been given to the phenomenon despite its very common occurrence in certain regions around the world, such as northern peatlands.

This chapter introduced a new DS processing methodology which makes use of contextual data in order to reconnect coherent observations separated by loss-of-lock. In this methodology we perform multilooking based on polygons which mark physically existing divisions in the terrain, and assign a set of attributes and multilooked phases to each polygon. As is observable from their coherence matrices, most of these phase histories suffer from loss-of-lock. The observations of different polygons which we expect to behave in a similar manner are combined in order to parameterize a common functional model. This model is used to align the disparate observations to estimate a single unbroken time series for the contextual group.

Using this methodology, we have successfully been able to estimate accurate InSAR displacement time series in several subsiding peatland regions in The Netherlands which was previously not possible with InSAR. To our knowledge, this is the first time that an accurate and validated time series has been estimated based on direct observation of the peatland pixels using DS techniques.

6

Land Surface Monitoring at Regional Scale

60 percent of the time, it works every time.

Anchorman

This chapter documents the development of the DECADE methodology for regional scale use over the entire Sentinel-1 archive up to January 2023. The results of applying DECADE for estimating the land surface motion of the rapidly moving and decorrelating peatlands of the Green Heart region of the Netherlands for eight years from January 2015–January 2023 are provided. Lastly, greenhouse gas emissions are estimated by combining the per-parcel subsidence estimates with soil type information.

Highlights:

1. *The challenges of moving from a short time- and local spatial scale analysis to long time- and regional spatial scale are described.*
2. *The DECADE distributed scatterer processing methodology is documented.*
3. *Per-parcel subsidence rates are estimated for the Green Heart region of the Netherlands. An average irreversible subsidence rate of 6.3 mm/year is obtained over the region.*
4. *Greenhouse gas emission estimates due to soil oxidation are provided for the region. An upper bound average GHG emission intensity of 21.5 tonnes of CO_2-equivalents per hectare per year is obtained, or 2.3 Mt CO2-eq/year for the entire Green Heart, 54% of the total national estimated GHG emissions caused by peatland oxidation in the Netherlands, or 1.3% of the entire GHG emissions of the Netherlands in 2019 [14].*

Parts of this chapter have been submitted to *Geophysical Research Letters*, [108].

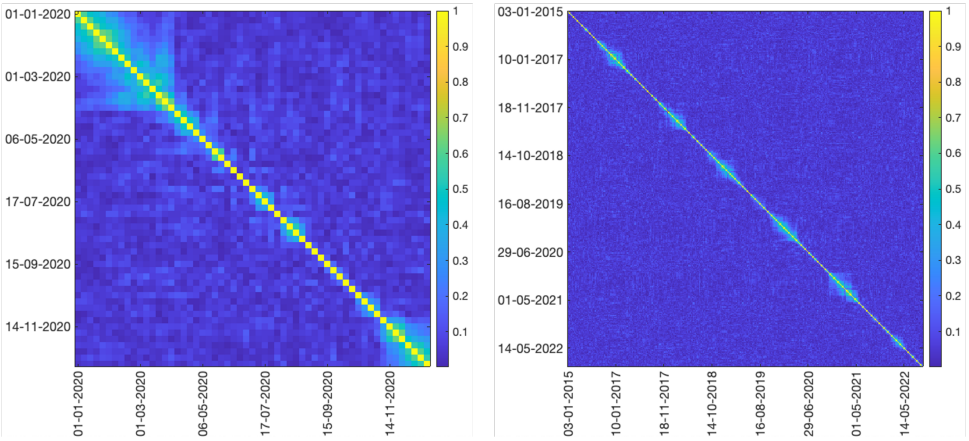

Figure 6.1: Comparison of coherence matrices of a one-year stack (left) vs. an eight-year stack (right) of the same location. The respective proportions of interferometric combinations with $|\hat{\gamma}| > 0.15$ is 13% and 2.4%.

6.1 Introduction

The previous chapters describe methodological developments aimed at enabling the use of C-band radar interferometry for monitoring shallow-based land surface motion. While they all provide results of test cases to demonstrate the feasibility of the proposed methods, they are limited in both their spatial and temporal extents. One of the key advantages of InSAR is the ability to provide estimates over large spatial extents, which has not yet been taken advantage of, nor has the entire archive of Sentinel-1 been used in those initial results. The goal of this chapter is to apply the newly developed methods to a large region of peatlands in the Netherlands to provide land surface displacement estimates at a regional scale using the entire available Sentinel-1 archive. Establishing a multi-year average subsidence rate estimate also enables a preliminary analysis of greenhouse gas (GHG) emissions from the region, which is presented following the estimation of the region-wide subsidence rates.

6.2 Challenges of Scaling Up to Regional Monitoring

6.2.1 ESM Phase Estimation in Large Stacks

The Sentinel-1 archive of imagery contains four tracks covering the Green Heart region, of approximately 400 SAR images each, spanning eight years (at the time of assessment). This large stack size introduces an additional level of complexity when performing an Equivalent single master (ESM) phase estimation [20], in particular for regions in which coherence is rapidly lost. The performance of full-rank phase estimators such as EMI [29] is unknown for such large stacks, due to the fact that they had not yet grown to such sizes when the estimators were developed. For example, the EMI method was published in 2018, only two years into the era of 6-day Sentinel-1 revisit times.

One major problem with applying the EMI method on rapidly decorrelating C-band data is the simple fact that most of the possible $N(N-1)/2$ combinations of single-

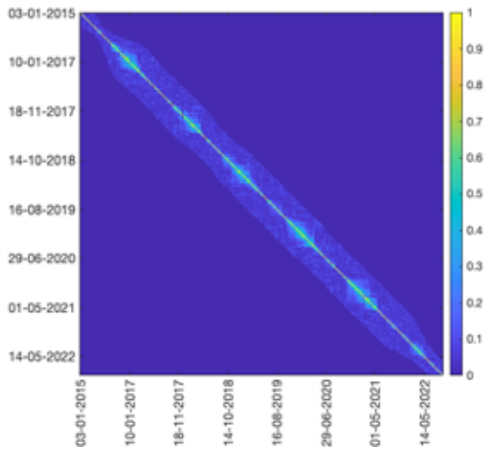

Figure 6.2: Example coherence matrix of an eight-year stack where all combinations with a temporal baseline greater than six months are masked. The proportion of (used) interferometric combinations with $|\hat{\gamma}| > 0.15$ is 18.4%.

look complex (SLC) images will be decorrelated, and that the proportion of coherent to incoherent data in the coherence matrix of a stack decreases geometrically (quadratically) with the length of the stack, as shown in Fig. 6.1. Due to noise and the biased nature of the coherence estimator, see Eq. (2.26) [57], combinations with long temporal baselines still have a non-zero coherence, even when it is expected (or known) [26] to be zero, which means they are still given some weight in the estimation process despite the fact that they are completely noise dominated-data.

One simple solution to this problem is to mask the coherence matrix by forcing all combinations with a temporal baseline beyond a certain threshold to be zero. This can drastically decrease the amount of noise going into the estimator even with a very conservative threshold. For example, the masked coherence matrix in Fig. 6.2 is set such that all combinations longer than six months are masked and removed from the estimation. This approach can also decrease the computational burden of calculating the full complex coherence matrix as the long-baseline interferometric combinations can be omitted.

6.2.2 Point Scatterer Networks

A network of high-quality point scatterers (PSs) [44] is used for atmospheric phase screen (APS) estimation and removal, and for the selection of an appropriate reference point, and is implemented using the Delft Persistent Scatterer Interferometry (DePSI) methodology [54]. DePSI works by first creating a "primary" network of high quality points through which the APS is estimated and subsequently interpolated by Kriging before being removed for all primary and secondary (all remaining identified PS candidates) scatterers in the stack. In DECADE, the multilooked DSs are input to DePSI to be treated as additional virtual secondary points positioned at the DS centroid. For large datasets, the primary network formation step is very sensitive to parameters such as the point selection criteria and the network's spatial density. Additionally, as the time series gets longer, the

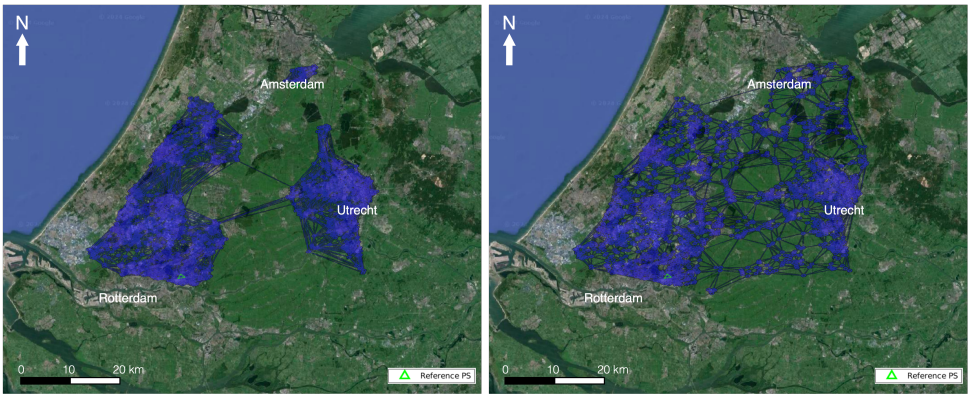

a) Spider b) Delaunay

Figure 6.3: Comparison of primary PS networks over the approximately 85×65 km Green Heart AOI obtained by different networking strategies in DePSI (Sentinel-1 ascending track 088). Left: network obtained by using "spider" algorithm [109]. Right: network obtained by using Delaunay triangulation.

probability of a given PS remaining coherent over the entire time span decreases.

At regional scales, hundreds of thousands of PSs will initially be identified. Arcs are then formed and initial parameters are estimated, including phase ambiguity resolution. These arcs are subsequently checked for spatial consistency by forming closed loops and checking if all estimated ambiguities in the loop sum to zero, over the entire primary network. Forming a consistent primary network is a challenge which requires balancing spatial density and regularity against computational overhead. If too many points are selected, checking spatial consistency between all the arcs becomes an unmanageably large computation. If too few points are selected, the network will be too sparse and the resulting APS will be poorly estimated. A lack of high-quality arcs between points can in fact also lead to split networks, which will significantly compromise the resulting APS estimate.

Fig. 6.3 visualizes the differences in the primary PS networks obtained by two different network formation strategies. The area of interest (AOI) approximately covers the Randstad region of the Netherlands, which is a highly urbanized area which encircles the Green Heart, which consists mostly of rural grassland pastures. This creates a lack of primary network points in the centre of the AOI. In Fig. 6.3 a), a spatially inconsistent set of points (and split networks) is obtained using the "spider" algorithm [109]. Split networks result in the need for multiple reference points, and no geodetic connection exists between the two networks, which will introduce error into the APS estimate. In Fig 6.3 b), a more spatially homogeneous network is obtained by using Delaunay triangulation. The more regular spatial coverage in the network results in one network, and the APS is interpolated over shorter distances from the primary network points. While the spider method creates a more dense and redundant network at local scales, it also tends to form longer arcs between distant high-quality points in lieu of shorter arcs which are less affected by atmospheric effects. For this reason, the Delaunay triangulation was found to be a more suitable method for this context.

6.3 DECADE InSAR Methodology

6.3.1 Overview

The Delft Contextually-Aided Distributed Scatterer Environment (DECADE) methodology developed in this research is pictured in the flowchart in Fig. 6.4 and is an adaptation of Fig. 5.4, with changes made as a result of the challenges described in Sec. 6.2. The simultaneous development of the SPAMS model (Ch. 3) also contributed to changes in the initial phase unwrapping and model estimation. Having only four model parameters to be estimated means that the model can be tailored for each individual parcel, as opposed to relying on a neural network trained at only a few locations with ground truth. Having fewer model parameters also helps ensure that overfitting is less likely, as discussed in Ch. 3. In addition, the fact that SPAMS provides daily displacement estimates which could be subsampled to match any revisit time made it a more effective choice than the RNN-based model for processing large interferometric stacks [1]. This section outlines the main processing steps taken in the DECADE v1.0 distributed scatterer (DS) processing block of Fig. 6.4, following APS filtering, as the preceding steps are not significantly changed.

6.3.2 Reference Point Selection and Motion Analysis

The selection of a common reference point is a crucial step for combining the observations of multiple viewing geometries (i.e., satellite tracks) together. Because the basic InSAR observable is a spatiotemporal double-difference phase, the displacement signal will always be relative to the motion of the reference point, which will be different for every track because they are processed independently. Therefore, in order to combine the observations of multiple tracks together, all observations should first be re-referenced to the same object on the ground. Generally, the only way to ensure this is to use ascending- and descending-oriented corner reflectors attached to the same structure. This is one of the purposes of the integrated geodetic reference stations (IGRSs) installed throughout the Netherlands. However, such a station cannot be used directly as the reference point for the entire time series because it was constructed in the year 2020, and we require the time series to go back to 2015.

As an alternative, we try to find a structure which appears in each track. First, a list of all PSs with the same (within a given threshold) estimated geolocation across all tracks is made. The PS with the lowest overall amplitude dispersion across all tracks is chosen from this list as the new reference point for all tracks. Thus we can be sure that a high quality (low noise) scatterer is chosen which approximately represents the same physical object in every track.

As mentioned above, these phase observations are all relative to the motion of the reference point. To ensure that the average annual soil subsidence rates are not affected by the motion of the reference point, we estimate its motion in the European Terrestrial Reference Frame (ETRF). This is accomplished by making use of the IGRS stations [110], as they consist of a radar corner reflector attached to a permanent global positioning system (GPS) station on the same rigid structure. Thus the relative motion between the chosen

[1] The loss of the Sentinel-1b satellite presented an additional challenge for the initial RNN-based phase unwrapping routine (Ch. 4), which assumed a constant 6-day revisit time to make predictions.

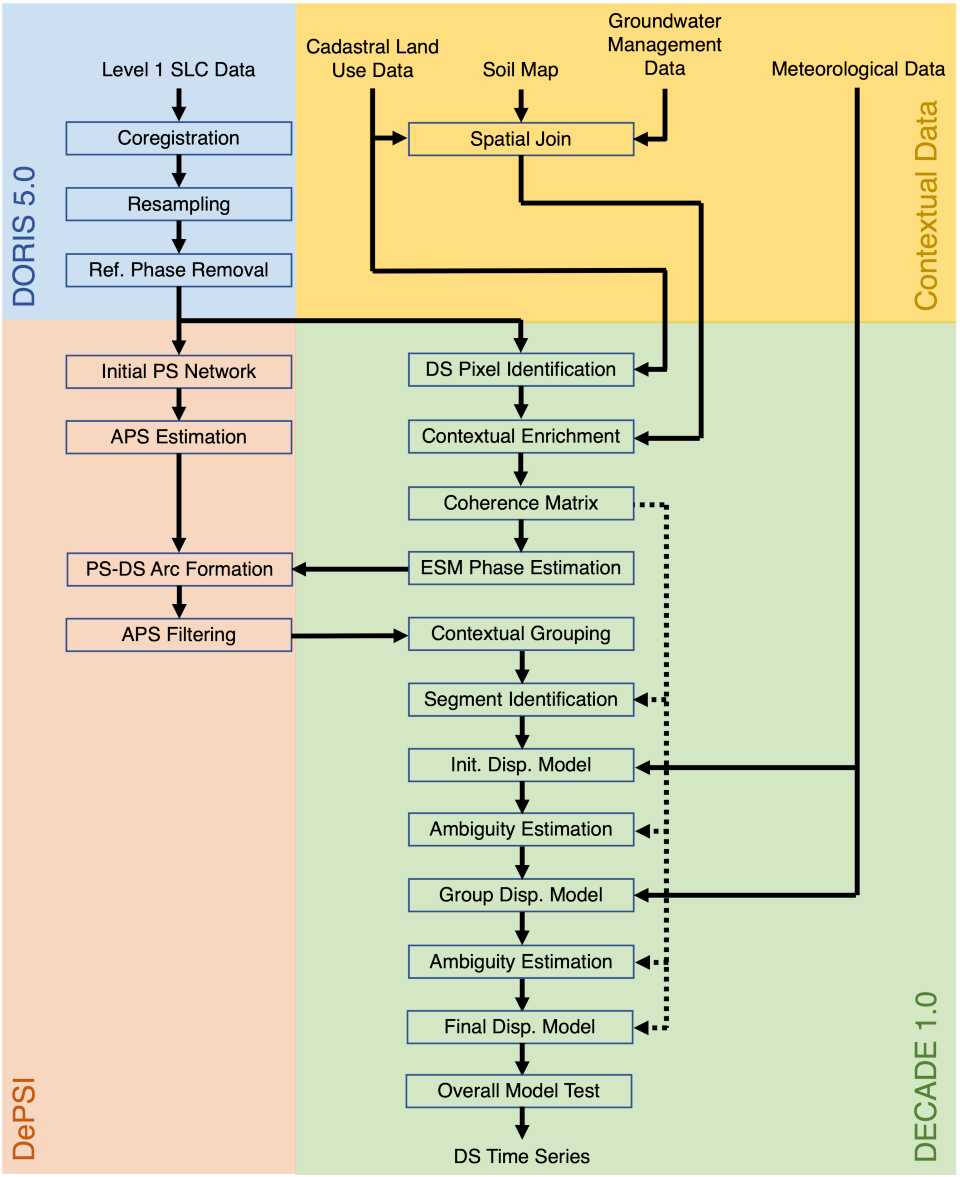

Figure 6.4: Simplified process flow diagram showing the major steps taken to create DS time series estimates from Sentinel-1 Level 1 SLC SAR data with the aid of spatial and temporal contextual data. DECADE v1.0 software implementation.

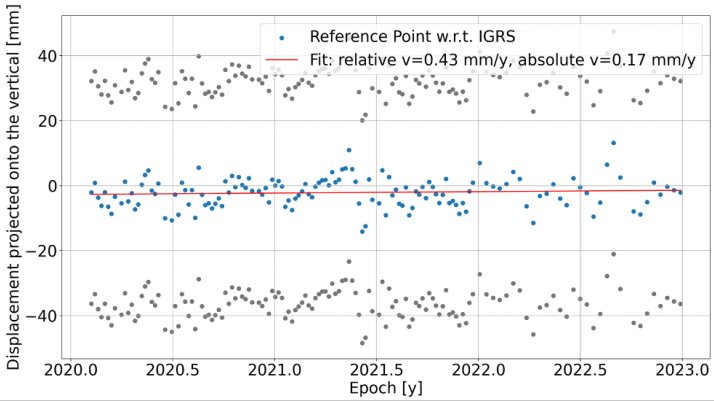

Figure 6.5: Plot showing estimated reference point motion relative to the IGRS station for the period 2020–2023 (blue dots, 2π ambiguities shown in grey). An overall relative linear trend (red line) 0.43 mm/year is estimated, corresponding to a linear trend of 0.17 mm/year in the ETRF (reference: S.A.N van Diepen, in prep.).

reference point and the IGRS station is analyzed for the period following 2020. Fig. 6.5 shows the phase observations and a linear trend of the motion of the reference point relative to the IGRS station. Using the GPS observations over the same period allows us to estimate the averaged motion of the reference point in the ETRF. A trend of 0.47 mm/year is obtained relative to the IGRS over the temporal subset, corresponding to a vertical motion of 0.17 mm/year in ETRF following removal of the GPS estimated trend. The value of 0.17 mm/year is not considered significant when compared to the expected average subsidence rates of the land surfaces under investigation (approx. 5–10 mm/year), and we can therefore assume that all the phase observations can be attributed to the DS. Residual reference point motion, noise and clutter can be mitigated by the use of a reference point noise estimation algorithm [111].

6.3.3 Initial Parcel-Wise Model Estimation and Phase Unwrapping

As elaborated in Ch. 5, the presence of loss-of-lock results in a sparse time series of wrapped interferometric phases. The time series is sparse because it consists of disconnected segments (subsets) of coherent data, interspersed by incoherent periods. On their own, these sparse, wrapped phases cannot be interpreted due to the unknown displacement during the incoherent period and the highly nonlinear nature of the surface motion signal [27], [75], [98]. By assuming that the true vertical surface displacement can be described by a model, z_M, with unknown parameters, x, the nth coherent segment of the signal can be written as

$$W\Big\{\phi_{\text{esm},n}(t)\Big\} = W\Big\{\frac{-4\pi\cos\theta_{\text{inc}}}{\lambda}\cdot\Big[z_M(x,t)+\Delta z_n\Big]+\epsilon\Big\};\quad \forall\ t\in T_n, \tag{6.1}$$

where $\phi_{\text{esm},n}$ is the ESM interferometric phase of the nth coherent segment, t is time, T_n is the set of all epochs in the nth coherent segment, θ_{inc} is the incidence angle, λ is the wavelength, Δz_n is the nth segment's constant offset caused by the unknown vertical displacement which occurred during the incoherent period, ϵ is a combination of noise and

model residuals, and $W\{\cdot\}$ is the wrapping operator. By taking the first-order difference between subsequent phases within each coherent segment, we can omit any unknown constant in the time series caused by displacement during an incoherent period, resulting in the equation

$$W\Big\{\phi_{\text{edc},n}(t)\Big\} = W\Big\{\frac{-4\pi\cos\theta_{\text{inc}}}{\lambda}\cdot\Delta z_M(x,t)+\epsilon_\Delta\Big\};\ \ \forall\ t\in T_n, \tag{6.2}$$

where $\phi_{\text{edc},n}$ is the (differential) equivalent daisy-chain (EDC) interferometric phase of the nth coherent segment and ϵ_Δ is the differenced noise and model residuals.

Now the time series can be implicitly unwrapped by estimating the optimal model parameters $\hat{x}_{\text{opt}}$. This is done by maximizing the temporal coherence, $\hat{\gamma}_{\text{temp}}$, between the modeled and observed phases in all coherent segments of all tracks by iteratively tuning the model parameters $\hat{x}$:

$$\hat{x}_{\text{opt}} = \arg\max_{\hat{x}}\hat{\gamma}_{\text{temp}}\Big(\phi_{\text{edc}}(t),\phi_M(\hat{x},t)\Big), \tag{6.3}$$

where the subscript M indicates the forward-modelled phase, and

$$\hat{\gamma}_{\text{temp}} = \frac{1}{N}\left|\sum_t \exp\Big(i[\phi_{\text{edc}}(t)-\phi_M(\hat{x},t)]\Big)\right|, \tag{6.4}$$

where N is the total number of coherent EDC phases in all segments and i is the imaginary unit. Eq. (6.4) can be used to combine the observations of multiple viewing geometries (or indeed sensors) together in order to derive a single common displacement model, provided that all observations are referenced to the same object. Once the optimal set of model parameters is found, the time series segments are unwrapped by estimating the closest integer ambiguity that matches the observed and modelled EDC phases under the assumption that the optimized model is representative of the vertical motion of the scatterer. The unwrapped ESM phases are obtained by taking the cumulative sum of the unwrapped EDC phases.

Following ambiguity estimation, the coherent segments are reconnected by estimating the constant offset Δz_n by

$$\Delta\hat{z}_n = \left\langle\frac{-\lambda}{4\pi\cos\theta_{\text{inc}}}\cdot\phi_{\text{esm},n}(t)-z_M(\hat{x}_{\text{opt}},t)\right\rangle;\ \ \forall\ t\in T_n, \tag{6.5}$$

where $\langle\cdot\rangle$ is the averaging operator. Now we have initial estimates for the unwrapped phase time series within each coherent segment, as well as an estimate for the relative positions of the segments with respect to one another.

6.3.4 Group Model Estimation

Following the initial parcel-wise estimation, the unwrapped and reconnected phase time series data of a given contextual group are combined and a joint group model is estimated. Contextual groups are identified based on their land cover/land use, soil type, elevation and water management zone classifications (Sec. 5.2.4). Despite the fact that

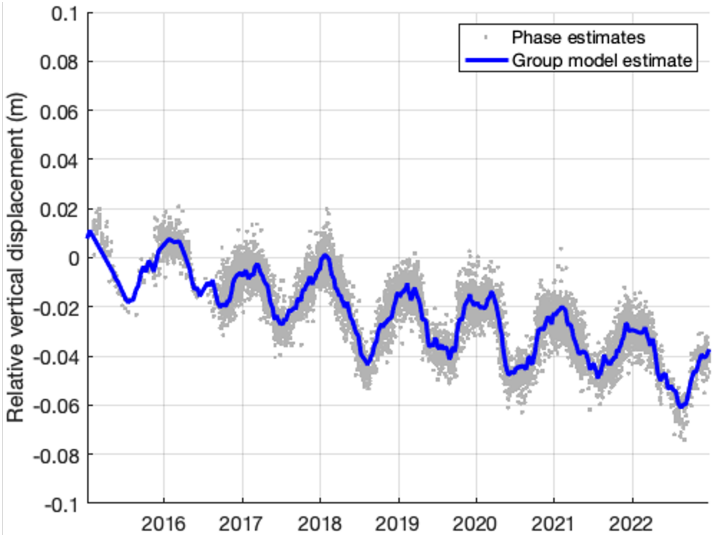

Figure 6.6: Plot of example contextual group model estimation. Grey dots: combined initial unwrapped phases of contextual group (projected to vertical displacement). Blue line: estimated group displacement model.

each parcel in a given contextual group shares these same attributes, there can be differences between the parcel seasonal shrinkage and swell amplitudes due to their differing geometries or other unconsidered factors, such as the installation of subsidence mitigation systems. Nevertheless, parcels within a contextual group are expected to show comparable irreversible subsidence coefficients (Eq. (3.10)), i.e. the x_I parameter, which we make use of in order refine the initial per-parcel models. Initially, the entire estimated group model was used to re-unwrap the phases (Sec. 5.2.8), however, this was found to be too restrictive and resulted in poor estimations in the parcel-wise x_P and x_E parameters.

The x_I parameter is challenging to estimate because it corresponds to a process that is approximately an order of magnitude smaller than the reversible shrinkage/swell (corresponding to the x_P and x_E parameters), and will be more affected by noise when moving to the differential EDC phase observations (Eq. (6.2)). Therefore, the group x_I parameter is used to re-estimate the phase ambiguities and segment shifts following the group model estimation, shown in Fig. 6.6. The group model is estimated by minimizing the root mean squared error (RMSE) between the model and the combined contextual group phases. Parcels with low temporal coherence ($\hat{\gamma}_{\text{temp}} < 0.15$, Eq. (6.4)) are excluded from the estimation to reduce noise and poorly-fit initial models from affecting the group estimate. The ambiguities and segment shifts are then re-estimated in the same way as described in Sec. 6.3.3.

As the combined observations of the contextual group represent independent samples of a similar process, the dispersion of the phase observations per epoch, $\sigma^2_{\text{group},i}$, of the entire contextual group is used as an approximation of the uncertainty of the phase observations (corresponding to the dispersion of the grey points in Fig 6.6). This has the benefit of modeling uncertainty due to coherence levels, but also captures the effects

of the imperfect ESM phase estimation (Sec. 6.2.1), APS estimation error (Sec. 6.2.2), unmodelled reference point motion and phase unwrapping errors without the need for storing the entire coherence matrix of every parcel, which would require a prohibitively large amount of data storage. It should be noted that this is a first-order approximation of the stochastic model and is not independent from the functional model, and therefore should be regarded as an upper bound. The covariance matrix of the phase observations of a given parcel is thus approximated by

$$Q_\phi \approx \mathrm{diag}(\hat{\sigma}^2_{\text{group}}), \tag{6.6}$$

where $\hat{\sigma}^2_{\text{group}}$ is the vector of sample variances of the contextual group phases at every epoch. In this ESM form, we expect the phase observations at each epoch to be independent [112].

6.3.5 Final Parcel-Wise Model Estimation and Uncertainty Propagation

Following the second implicit phase unwrapping, a final displacement model is fit to the unwrapped phases for each parcel. The final model is fit by minimizing the mean squared error (MSE) between the unwrapped phases of each track (projected to vertical displacement) and the model, such that

$$\hat{x}_{\text{opt}} = \arg\min_{\hat{x}} \left\langle \left(\frac{-\lambda}{4\pi \cos\theta_{\text{inc}}} \cdot \phi_{\text{unw}}(t) - z_M(\hat{x}, t) \right)^2 \right\rangle, \tag{6.7}$$

where ϕ_{unw} are the unwrapped interferometric phases. The fitting is performed using the Nelder-Mead Simplex method, which is a derivative-free method [113].

The per-parcel vectors of unwrapped phases of the N independent satellite tracks, ϕ_n, are combined in a single vector, y_{all}, by projecting from the LOS phase to vertical displacement (no horizontal motion is expected due to no deep subsidence sources in the region [114]) as

$$y_{\text{all}} = \begin{bmatrix} \phi_1 \cdot -\lambda/4\pi \cos\theta_{\text{inc},1} \\ \vdots \\ \phi_N \cdot -\lambda/4\pi \cos\theta_{\text{inc},N} \end{bmatrix}, \tag{6.8}$$

where λ is the wavelength, and $\theta_{\text{inc},n}$ is the incidence angle of the location corresponding to the nth track's viewing geometry. The covariance matrix of the final model parameters is estimated per parcel by propagating the covariance matrix of the unwrapped phases, Q_ϕ, which is in turn approximated as described in Sec. 6.3.4. The combined covariance matrix of all tracks is obtained by diagonally combining the individual covariance matrices of each track. Each viewing geometry is assumed to be independent, giving

$$Q_{y,\text{all}} = \begin{bmatrix} Q_{\phi,1} \cdot (\lambda/4\pi \cos\theta_{\text{inc},1})^2 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & Q_{\phi,N} \cdot (\lambda/4\pi \cos\theta_{\text{inc},N})^2 \end{bmatrix}. \tag{6.9}$$

Following the formulation of [19], the SPAMS model is converted to a linearized matrix form, such that

$$\hat{y} = AM\hat{x}, \tag{6.10}$$

where

$$\hat{x} = \begin{bmatrix} \hat{x_P} \\ \hat{x_E} \\ \hat{x_I} \end{bmatrix}, \tag{6.11}$$

and

$$M = \begin{bmatrix} \sum_{t_1-\tau}^{t_1} P(t) & -\sum_{t_1-\tau}^{t_1} E(t) & z_1 \\ & \vdots & \\ \sum_{t_K-\tau}^{t_K} P(t) & -\sum_{t_K-\tau}^{t_K} E(t) & z_K \end{bmatrix}, \tag{6.12}$$

and where z_k is 0 if $\sum_{t_k-\tau}^{t_k} P(t) > -\sum_{t_k-\tau}^{t_k} E(t)$, and 1 otherwise (see Eq. (3.10)). The daily model values are mapped to the satellite acquisition epochs by the design matrix

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \ldots \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \ldots \\ & & & \vdots & & & \end{bmatrix}, \tag{6.13}$$

where an element of A is unity when the model and observation dates coincide. The rows of A correspond to satellite observation epochs and the columns correspond to the model epochs. Thus the size of A is the total number of observation epochs times the number of model epochs, and the product AM is the overall functional model which maps the satellite observations to model estimations. In this linear form, the integration time τ is part of the functional model and cannot be written as a parameter in $\hat{x}$. Therefore, no uncertainty is estimated for this parameter. The covariance matrix of the model parameters is then propagated as

$$Q_{\hat{x}} = \left((AM)^T Q_{y,\text{all}}^{-1} (AM) \right)^{-1}, \tag{6.14}$$

and the covariance matrix of the (daily) modeled displacement is obtained by

$$Q_{\hat{y}} = M^T Q_{\hat{x}} M. \tag{6.15}$$

6.3.6 Overall Model Test

An overall model test (OMT) [115] is applied to assess the suitability of the assumed displacement model based on the model residuals and uncertainties. The model residuals, $\hat{e}$, are

$$\hat{e} = y - \hat{y}. \tag{6.16}$$

The test statistic, T, is a measure of the relative magnitude of the model residuals with respect to the covariance matrix, and is given by

$$T = \hat{e}^T Q_{\hat{y}}^{-1} \hat{e}. \tag{6.17}$$

The OMT compares the test statistic to a critical value, T_{crit}, which corresponds to a χ^2 distribution with a chosen significance level and number of degrees of freedom,

$$T_{\text{crit}} = \chi^2(\text{DOF}, \alpha), \tag{6.18}$$

where DOF is the number of degrees of freedom (i.e. the redundancy), and is given by

$$\text{DOF} = P - N, \tag{6.19}$$

where P is the number of observations and N is the number of unknown model parameters. In our case, P is the total number of coherent interferometric phase observations across all tracks, and N is three, see Eq. (6.11). α is the significance level, which is the probability of falsely rejecting the model. A 5% significance level is used, which is a conservative threshold to prevent false acceptance of the model. The model is rejected if

$$T > T_{\text{crit}}, \tag{6.20}$$

which indicates that the model residuals do not agree with the model uncertainties. This means that either the model does not adequately describe the observed behaviour, or else, the stochastic model is too optimistic.

6.4 Results and Discussion

6.4.1 Description of Satellite Data and Study Area

The area of study is known as the "Green Heart", an agricultural area comprised mainly of grasslands which is situated on mainly peat and marine/river clay soils (see Appendix A). While the region is largely rural, it is encircled by the "Randstad"; a metropolitan area made up of the cities of Amsterdam, Rotterdam, the Hague, Utrecht and others. Final estimations were made for 34908 grassland parcels in the region, covering a total area of approximately 600 km^2. A total of 1515 Sentinel-1 images from observation tracks (ascending 088 and 161, and descending 037 and 110) between January 2015 – January 2023 were used in the analysis.

Table 6.1: Summary of satellite data and study area

Parameter	Value
SAR sensor	Sentinel-1
Period	Jan. 2015 – Jan. 2023
Num. SAR images	1515
Num. SAR tracks	4
Num. parcels	34908
Total parcel area	601 km^2

6.4.2 Displacement Time Series Estimates

The estimated time series displacements of the three parcels with available in-situ measurements by extensometer [13], [32] are shown in Fig. 6.7. The combined unwrapped InSAR displacement estimates from the four overpassing Sentinel-1 tracks are shown as red points along with the corresponding 1σ estimated error bars. The estimated model is shown by the solid blue line, and the corresponding 3-sigma confidence region obtained by Eq. (6.14) is shown by the shaded blue area. The average subsidence rate is given

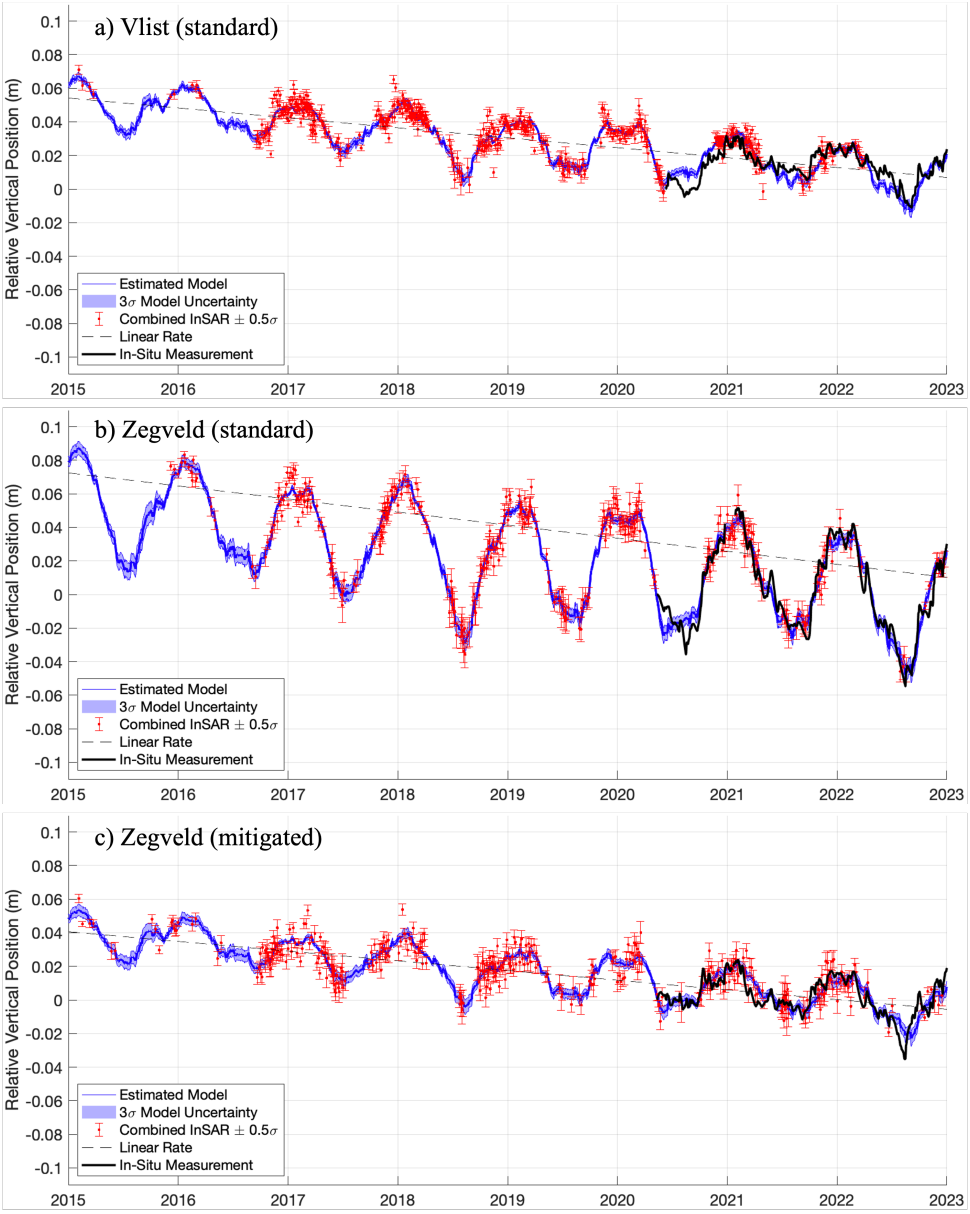

Figure 6.7: Estimated InSAR displacements (projected from LOS to vertical) and associated model against in-situ measurements at test sites a) Vlist (standard), b) Zegveld (standard), and c) Zegveld (mitigated) (Table 6.2) for 2015–2023. Red dots: combined (see Sec. 6.3) InSAR displacement estimates of Sentinel-1 ascending tracks 088 and 161, and descending tracks 037 and 110, with estimated 1σ error bars. Blue line: estimated InSAR displacement model. Shaded blue region: estimated 3-sigma confidence region. Dashed black line: estimated average subsidence rate. Bold black line: in-situ measurement [13], [32] (error bars too small to plot).

Table 6.2: Estimated displacement model parameters and performance against extensometer (ext.)

Location Designation	Vlist (Standard) a)	Zegveld (Standard) b)	Zegveld (Mitigated) c)
Coordinates (lat., lon.)	(51.9827°, 4.8269°)	(52.1358°, 4.8399°)	(52.1372°, 4.8398°)
$\hat{x}_P \pm 1\sigma$ (mm/mm)	$7.4{\cdot}10^{-2} \pm 3.6{\cdot}10^{-3}$	$9.4{\cdot}10^{-2} \pm 6.4{\cdot}10^{-3}$	$7.4{\cdot}10^{-2} \pm 5.7{\cdot}10^{-3}$
$\hat{x}_E \pm 1\sigma$ (mm/mm)	$1.2{\cdot}10^{-1} \pm 4.0{\cdot}10^{-3}$	$2.7{\cdot}10^{-1} \pm 7.1{\cdot}10^{-3}$	$1.0{\cdot}10^{-1} \pm 5.6{\cdot}10^{-3}$
$\hat{x}_I \pm 1\sigma$ (mm/day)	$-2.7{\cdot}10^{-2} \pm 5.8{\cdot}10^{-4}$	$-2.7{\cdot}10^{-2} \pm 9.5{\cdot}10^{-4}$	$-2.8{\cdot}10^{-2} \pm 1.1{\cdot}10^{-3}$
$\hat{\tau}$ (days)	65	70	70
Av. subsidence (mm/year)	-5.9	-7.8	-5.6
RMSD Obs.–Ext. (mm)	6.5	7.4	7.5
RMSD Model–Ext. (mm)	5.5	6.7	4.1
RMSE Obs.–Model (mm)	4.7	5.9	6.0
Temporal Coherence	0.69	0.22	0.17
Area (m^2)	47428	11704	9053
Effective Num. Looks	409	99	78

by the dashed black line. Finally, the in-situ measurements by extensometer are shown by the solid black line (the expected precision of the extensometer is sub-millimetre, and corresponding error bars are too small to plot).

A comparison of estimated InSAR and available in-situ validation data is provided in Table 6.2. Sub-centimetre agreement (i.e. the root mean squared difference (RMSD)) is achieved across all three validation sites for both the InSAR displacement estimates, as well as the derived displacement model, which is approximately 5–10% of the overall signal variation. The Zegveld test sites b) and c) show a larger RMSE between observations and model due to the lower effective number of looks (i.e. independent samples) as a consequence of moving down to a sub-parcel scale resolution. The effective number of looks L is given by

$$L = n \times \mathrm{OSR}, \tag{6.21}$$

where n is the number of pixels and OSR is the oversampling rate, given by

$$\mathrm{OSR} = \frac{\mathrm{PRF}}{\mathrm{BW_{az}}} \cdot \frac{f_{s,\mathrm{r}}}{\mathrm{BW_r}} \tag{6.22}$$

where PRF is the pulse repetition frequency, $f_{s,\mathrm{r}}$ is the range sampling rate, and BW_{az} and BW_{r} are the azimuth and range bandwidths, respectively.

Differences are expected to be found between the InSAR estimates and the in-situ measurements, as the former are based on the average of many pixels covering a certain spatial extent (approximately 10^4 m^2) of the soil surface, and the latter are point-based measurements anchored at -5 cm in the soil subsurface. Nevertheless, close agreement between the results of the two techniques suggests that the InSAR estimations are reliable.

An experimental subsidence mitigation system [116] is installed at site c) (see Appendix B). This installation is a pressurized drainage system that is intended to reduce the overall fluctuations in groundwater level in the centre of the parcel, thereby reducing the amount of soil oxidation incurred during the dry summer periods. This site is in fact the northern part of the same test parcel as site b), which displays the most strongly

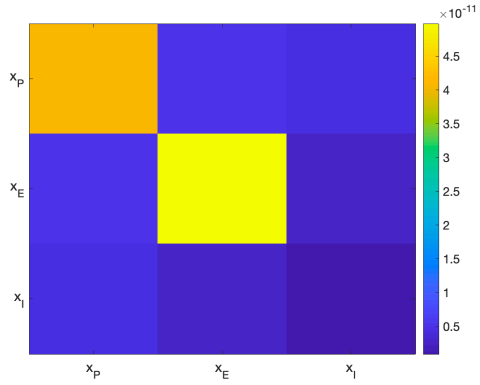

Figure 6.8: Covariance matrix $Q_{\hat{x}}$ obtained for test site b) (Zegveld, standard drainage). It can be seen that the x_P and x_E parameters are approximately uncorrelated. The x_I parameter is correlated with x_P and x_E to approximately the same order of magnitude as the variance in x_I. Note that the units of the covariance matrix vary depending on the parameters in $\hat{x}$, see Table 6.2.

dynamic behaviour in the region. This experiment provides a natural test case for our method to test whether or not it can resolve differences in the sub-parcel scale and be of use as a tool for assessing the efficacy of such subsidence/climate mitigation measures. A significant reduction in the seasonal amplitude is visible between sites b) and c), as plotted in Figs. 6.7 b) and c), from approximately 6 cm peak-peak to 2 cm peak-peak, respectively. A reduction of approximately 30% in the average subsidence rate was detected, from 7.8 mm/year at site b) to 5.6 mm/year at site c). The biggest losses in elevation during the summers of 2018 and 2023 are mitigated as well, by approximately 2 cm.

The effect of drainage measures on the displacement model parameters is shown in Table 6.2. In terms of these parameters, the effect of the measures is mainly characterized by a drastic reduction in the x_E parameter, which determines the magnitude of the response of the soil to evapotranspiration and the overall water retention ability of the parcel. Thus it is clear from our observations that despite the intrinsic susceptibility of the soil to oxidation (as characterized by the x_I parameter) remaining unchanged, the retention of water in the system leads to a reduction in the overall subsidence rate. This is what we expect to see, as the biochemical properties of the soil are the same in both cases; only the extrinsic groundwater conditions have changed.

A colour-coded visualization of a covariance matrix $Q_{\hat{x}}$ is shown in Fig. 6.8. It can be seen that the parameters governing reversible shrinkage and swell, x_P and x_E, are uncorrelated. The irreversible parameter, x_I, is correlated with x_P and x_E to approximately the same order of magnitude as the variance in x_I. This is expected because of how irreversible subsidence is defined in SPAMS, which depends on the cumulative balances of water gains and losses, see Eq. (3.10).

6.4.3 Regional Scale Soil Subsidence Rates

The average irreversible subsidence rate is estimated at parcel scale and visualized in Fig. 6.9. The most actively subsiding areas occur around the area between the cities of

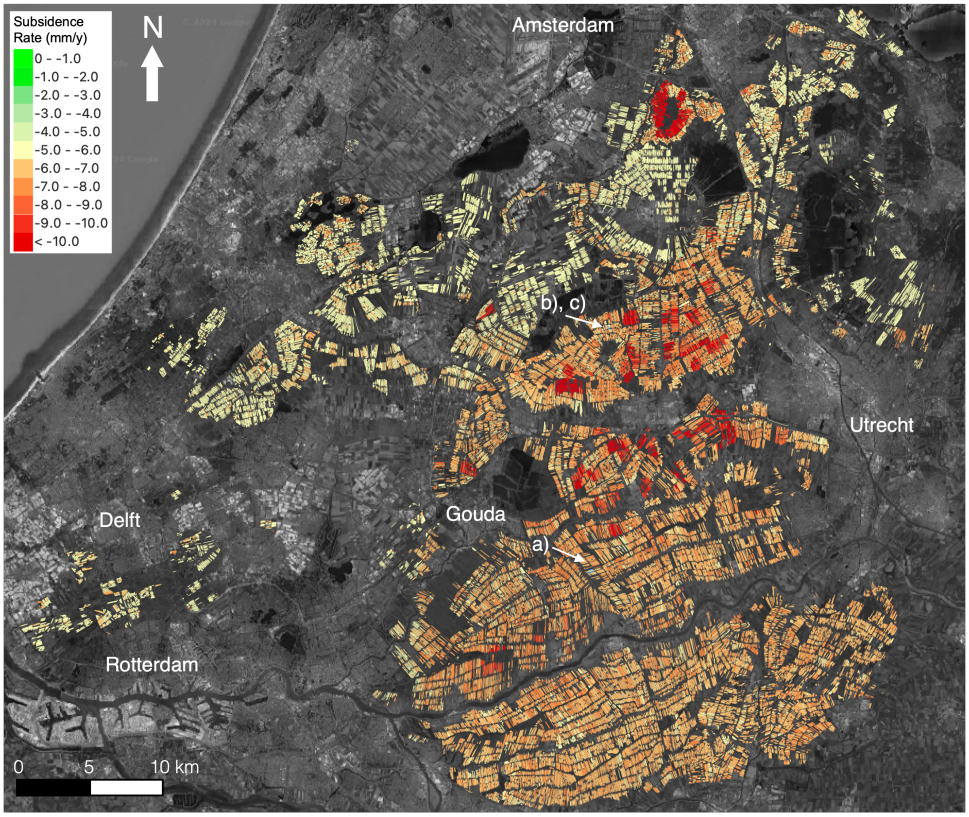

Figure 6.9: Map of parcel-scale average subsidence rates in the Green Heart region estimated over the period 2015–2023. The locations of test sites a) Vlist (standard), b) Zegveld (standard) and c) Zegveld (mitigated) are marked with white arrows.

Gouda and Utrecht, in the centre of the Green Heart region, which is a peat dominated region. A general slight decrease in subsidence rates is visible in the Northwest areas, where more marine clay is found in the Holocene, versus the Southeast, where peat is more prevalent (see Appendix A). An overall average subsidence rate for the region is found by weighted average of parcel surface area:

$$\bar{z'}_{\text{all}} = \frac{1}{A_{\text{tot}}} \sum_i A_i \cdot \bar{z'}_i, \tag{6.23}$$

where A is a given parcel's area and $\bar{z'}$ is the ith parcel's average subsidence rate. An overall average subsidence rate of 6.5 mm/year is found for the region.

Comparing this result to previous studies, there is significant uncertainty in the results obtained from repeated levelling campaigns and airborne laser scanning. A 1996 report by Beuving and van den Akker [117] (Table 11) based on repeated levelling surveys found subsidence rates in the Zegveld region of the Netherlands (around and including site b)) ranging from 4.7–13.5 mm/year depending on location and date of measurement. In a

2009 report by Jansen et al. [118] (Table 2.2), a decadal-scale average subsidence rate was estimated by combining levelling data from the period 1946–1972 and airborne laser scanning from the period 1997–2004, with rates ranging from 1.0–4.6 mm/year for peat soils, and 1.8–3.0 mm/year for marine and river clay soils (see soil map in Appendix A). The uncertainties in these estimates are significant: standard deviations for the former range from 4.6–17.2 mm/year, and 3.4–16.0 mm/year for the latter, approximately 1–4 times the magnitude of the actual estimated value in both cases. It is apparent from Fig. 6.7 that the sub-seasonal temporal dynamics of the land surface motion are approximately one order of magnitude greater than the long-term (multi-year) subsidence rates. This demonstrates the need for high temporal sampling of the surface displacement signal in order to obtain reliable average irreversible subsidence rate estimates, and how these previous studies based on observations at very localized points in space and time are at risk of significant biases due to the rapidly changing shrinkage and swell (the effects of vegetation on the airborne laser observation notwithstanding).

It should be noted that these average subsidence values are considered to be "steady-state" subsidence rates which do not consider any significant changes in anthropogenic intervention in the phreatic groundwater system. That is, with no significant changes to the groundwater management, the current rates of oxidation and compaction/creep are expected to remain more or less constant. However, if the current subsidence rates are allowed to continue unabated, and the region's current primary land use of grazing cattle is continued, additional lowering of the groundwater level will be required in the coming decades to enable the continued usage of the land[2]. Following such a lowering, a significant increase in the subsidence rate will occur due to settlement before returning back to a steady-state. As such, when projecting land surface heights over large time scales, the effective subsidence rate will necessarily be higher. Conversely, allowing the phreatic groundwater level to rise with respect to the surface level should decrease the overall subsidence rate.

6.4.4 Greenhouse Gas Emission Estimates

The estimation of average subsidence rates enables an estimate of GHG emissions caused by the oxidation-driven subsidence component. The average GHG emission intensity in CO_2-equivalents per hectare per year of the ith parcel, C_i, as a function of the average irreversible subsidence rate is estimated by the equation [6], [7]

$$C_i = F \cdot f_o \cdot f_{c,i} \cdot \rho \cdot \bar{z'}_i \cdot 10^4 \cdot 3.67, \tag{6.24}$$

where F is the fraction of subsidence caused by peat oxidation, f_o is the fraction of soil organic matter, $f_{c,i}$ is the soil-dependant fraction of carbon contained within the soil organic matter of the ith parcel, ρ is the bulk density of the soil, and $\bar{z'}_i$ is the average irreversible subsidence rate of the ith parcel. The factor 10^4 is a conversion of square metres to hectares, and the factor 3.67 is a conversion from tonnes-C to tonnes-CO_2 equivalents.

There is considerable uncertainty in assessing accurate values for the coefficient F, which varies quite strongly by location and time scale. The fraction F is reported in litera-

[2] See, in this context, the EU's Regulation on Nature Restoration (Nature Restoration Law) that came into effect 18 August 2024 [119].

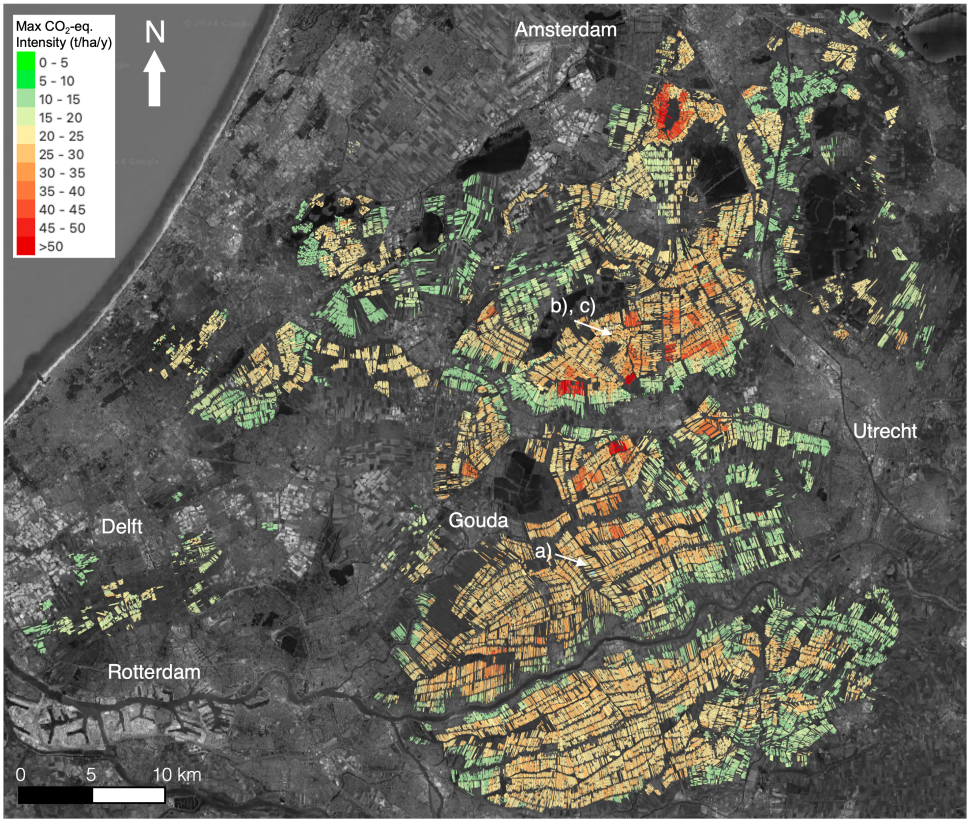

Figure 6.10: Map of parcel-scale maximum GHG emission intensities estimated over the period 2015–2023. The locations of test sites a) Vlist (standard), b) Zegveld (standard) and c) Zegveld (mitigated) are marked with white arrows.

ture to vary from 0.33–0.66 in [17], 0.25–0.71 in [18] and 0.85 in [10]. In [10], Schothorst determined a fraction of F of 0.85 at the Zegveld test site (site b)) by comparing the bulk densities of mineral elements in soils above and below the groundwater level, reasoning that the overall quantity of minerals should remain constant, but their bulk density should increase as the organic matter oxidizes. As 0.85 is an in-situ estimation from the region under study and also represents the upper bound of the reported values for F, we use this value as an estimate for the maximum expected level of subsidence caused by oxidation. We assume an average value of 220 kg/m^3 for the bulk density of organic matter (i.e. the product $f_o \cdot \rho$) in peat in the upper 30 cm of soil, as reported in [10]. The fraction of carbon contained within the soil organic matter, f_c, is generally assumed to be 0.55 for peat soils [15], [121]. The factor f_c is rescaled for other soil types by comparing the relative carbon stocks per soil type for grasslands, as inventoried in the 2023 report by [15] (Figure 11.2) on behalf of the Dutch government.

Applying these values to Eq. (6.24) for all estimated parcels in the region will provide a rough estimate of the maximum expected level of CO_2-equivalents we can expect to be

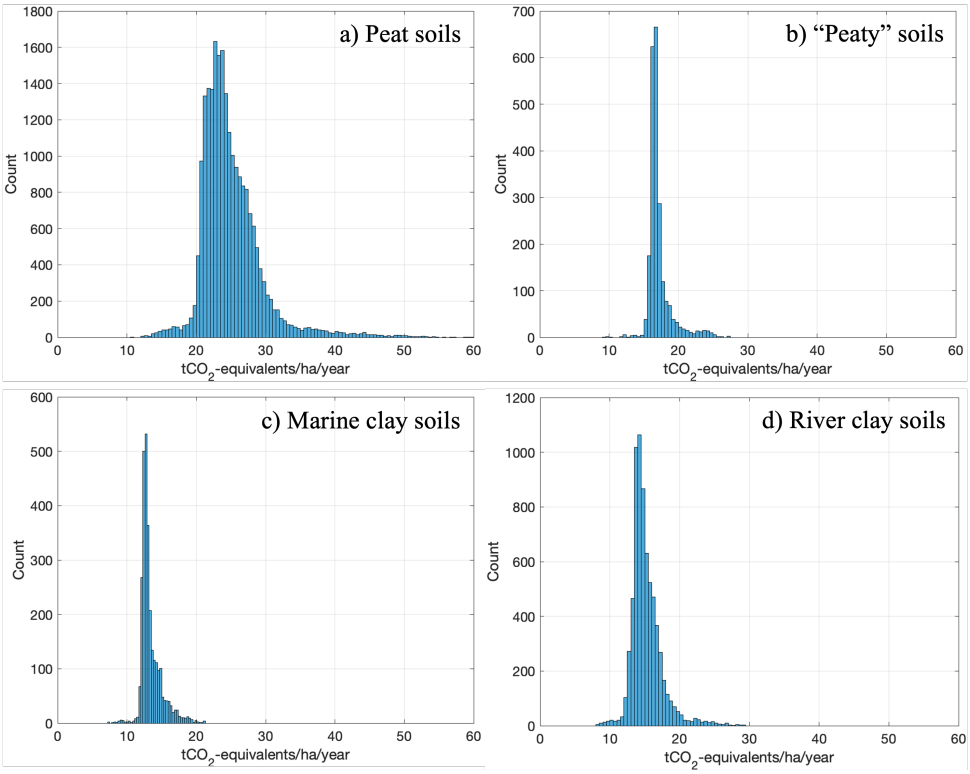

Figure 6.11: Histograms of GHG emission intensity obtained by Eq. (6.24) by main soil type classification [120]: a) Peat soils, b) "Peaty" soils, c) Marine clay soils, d) River clay soils.

released by peatland oxidation in the region, shown in Fig. 6.10. Certainly, the value of 0.85 for F will be an overestimation in many areas. Nevertheless, in doing so we establish an upper bound on the expected GHG emissions for the region based on satellite observations of ground motion. The average GHG emission intensity of the region can be established by taking the spatial average:

$$C_{\text{av}} = \frac{1}{A_{\text{tot}}} \sum_i A_i \cdot C_i. \tag{6.25}$$

We obtain an average maximum GHG emission intensity of 21.5 tCO_2-eq/ha/year, corresponding to 3.31 tCO_2-eq/ha/mm of subsidence. This is in good agreement with the National GHG Emission Inventory 1990–2019 [14] (based on the report [15]), which estimates an average 19 tCO_2-eq/ha/year for all peatlands in the entire Netherlands, considering we apply a blanket worst-case assumption across the entire region. Multiplying this rate by the combined total area of all land parcels within the region, a total maximum annual GHG emission rate of 2.3 $MtCO_2$-eq/year is obtained for the entire Green Heart. This figure corresponds to 54% of the total national estimated GHG emissions caused by peatland oxidation in the Netherlands (2.3 MtCO2-eq/year), or 1.3% of the entire GHG

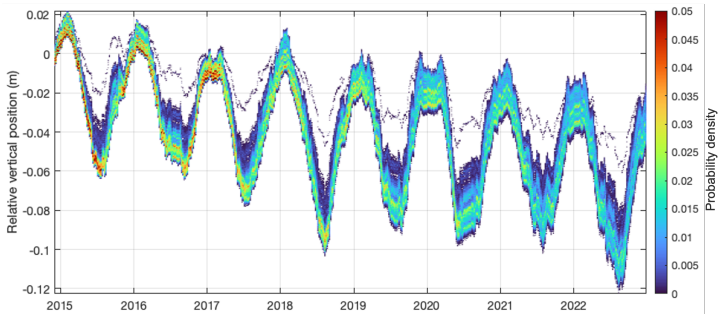

Figure 6.12: Heatmap showing the density of the initial model estimations over time at Zegveld site b) as a result of 500 runs in which 20% of the coherent data are randomly removed.

emissions of the Netherlands in 2019 [14]. The effect of soil type on GHG emissions is visible in Fig. 6.11, where peat soils show significantly higher levels of GHG emission intensity than all other soil types, due to a combination of higher subsidence rates and soil carbon content.

The effect of the pressurized drain subsidence mitigation measure installed at test site c) on the GHG emissions can be estimated by applying the difference between the estimated average subsidence rates between sites b) and c) (Table 6.2) to Eq. (6.24). We find that the subsidence mitigation measures are able to reduce the worst-case GHG emission intensity by approximately 7.6 tCO_2-eq/ha/year, approximately a 28% decrease.

6.4.5 Estimation Quality

6.4.5.1 Model Estimation Robustness

To evaluate the robustness of the initial model estimation to missing data, the unwrapped phases of parcel at site b) are taken and the model is re-estimated 500 times in which a random 20% of the coherent segments (Sec. 5.2.5) are removed. The resulting model estimates are plotted in a 2D histogram, shown in Fig. 6.12. In the majority of runs, comparable seasonal amplitude and irreversible subsidence rates are obtained, indicating that the method is generally robust and can handle losses of data. This also suggests that a the model estimation quality will continue to increase as the time series grows in length, provided that there are no significant changes to the parcel's biological and physical properties (the estimated parameters $\hat{x}$). The histogram is also positively skewed, which is particularly visible in the summer months. This indicates there is a tendency to underestimate the seasonal amplitude in the event of data loss. While the histogram diverges over time due to the differences in the irreversible component, it can be seen that the extreme values in the tails are relatively unlikely compared to the region containing the majority of the estimations.

6.4.5.2 Overall Model Test

The overall model test described in Sec. 6.3.6 is applied to the parcels with a model estimation, and is shown in Fig. 6.13. 27700 (approximately 80%) of the parcel models are sustained by the test at a significance level (α) of 5%, and 7208 parcels are rejected.

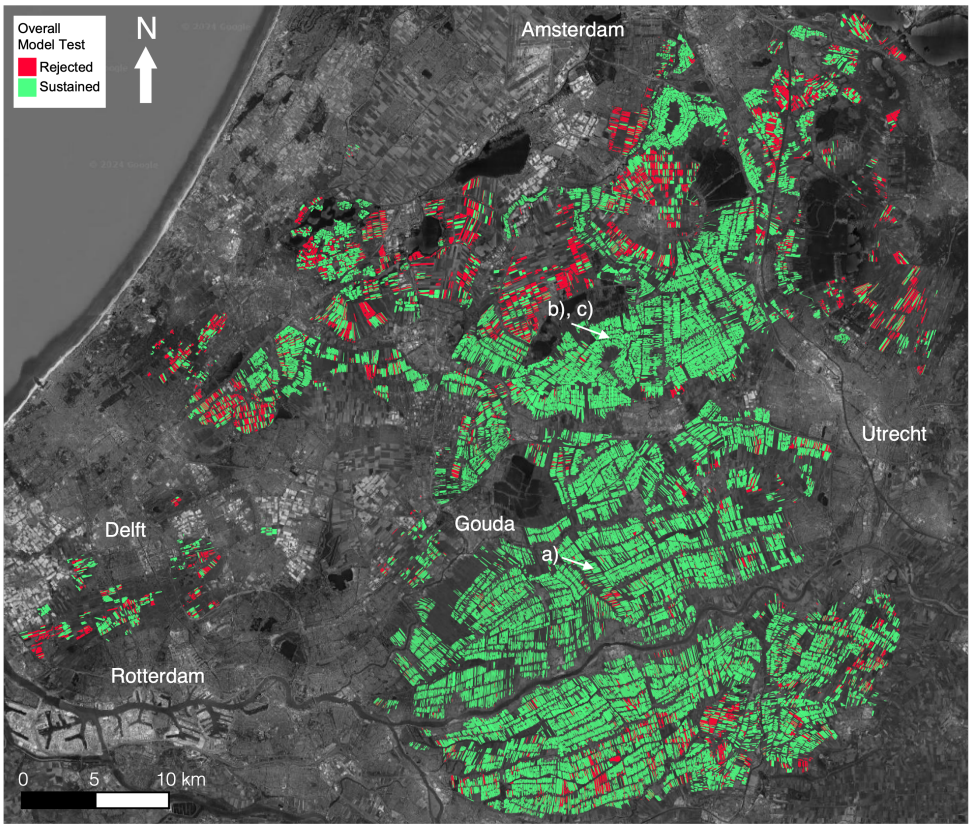

Figure 6.13: Map of parcel-scale overall model test results at 5% significance level in the Green Heart region estimated over the period 2015–2023. The locations of test sites a) Vlist (standard), b) Zegveld (standard) and c) Zegveld (mitigated) are marked with white arrows.

Interestingly, most of the rejected parcels are predominantly clay-dominated soils, suggesting that either the SPAMS model does not describe these soils as accurately as the peat-dominated soils, or else the estimated uncertainties are too optimistic for the clay-dominated parcels.

6.4.5.3 Temporal Coherence Between Observations and Model

The degree of fit between the model and wrapped phase observations as per the temporal coherence ($\hat{\gamma}_{\text{temp}}$) is visualized in Fig. 6.14. Generally speaking, a high degree of fit is obtained, ranging from approximately 0.4–0.8, values which are comparable to high-quality point scatterers. Most of the low-quality parcels ($\hat{\gamma}_{\text{temp}} < 0.4$) are attributed to parcels with small areas or irregular shapes. Indeed, as can be seen in Fig. 6.15, the temporal coherence is strongly dependent on the area of the parcel, which affects the number of independent looks used in the estimation procedure. The relationship between the temporal coherence and area appears to follow an error function curve, shown in red in Fig. 6.15. That is, it is possible to define a function relating the performance in temporal

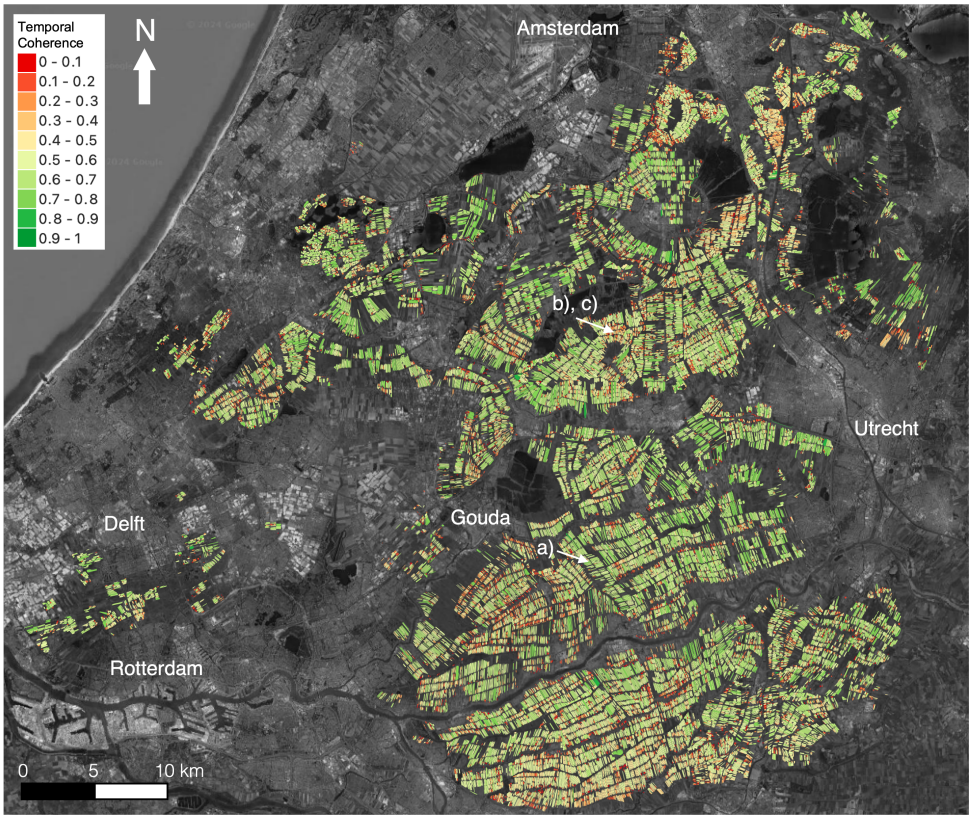

Figure 6.14: Map of temporal coherence between parcel-wise multilooked interferometric phase and estimated displacement model. The locations of test sites a) Vlist (standard), b) Zegveld (standard) and c) Zegveld (mitigated) are marked with white arrows.

coherence to the area of the multilooked region. This indicates that noise is the limiting factor in model estimation quality. Parcels with a larger area will be covered by more pixels, meaning a larger effective number of looks, allowing for a more accurate coherence estimation and better noise reduction by multilooking.

6.5 Conclusion

This chapter presented the results of applying a novel InSAR methodology to estimate land surface motion in the grassland areas of the Dutch Green Heart (*Groene Hart*) region. This is the first accurate and validated InSAR analysis performed at regional scale of these rapidly decorrelating peatlands. The methodology works by finding the set of parameters which optimizes the fit of a displacement model to the intermittently coherent wrapped phase observations. Thus the focus of this method is to use imperfect observational data y to learn how a given parcel reacts to hydrological inputs, which are described by the model parameters, x. The quality of the model fit to the data is generally very good,

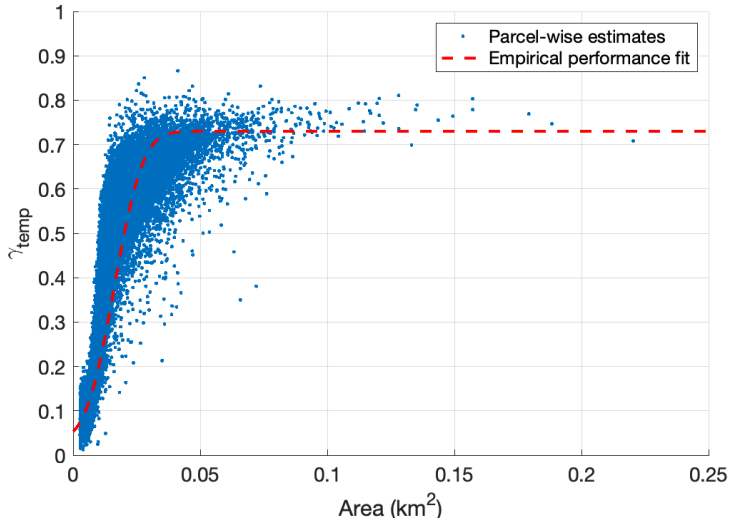

Figure 6.15: Temporal coherence vs corresponding parcel area. Blue dots: temporal coherence obtained for each parcel estimate. Red dashed curve: empirical fit to observed performance.

with most parcels having a temporal coherence ranging from 0.4–0.8, which is comparable to mature PS analysis methods which use only the lowest noise pixels in the image stack. Noise is shown to be the main limiting factor in the model fitting procedure, and that larger parcels with more pixel coverage show a correspondingly better fit to the data.

An average irreversible subsidence rate of 6.3 mm/year is obtained for the region, with the majority of parcels within the range of 5–8 mm/year. While the obtained estimates agree well with the available ground truth, the dynamics of the region show that any technique for monitoring the ground surface motion over time may not be representative if these dynamics are not taken into account. Currently, InSAR is the only technique which can provide both high spatial coverage with frequent enough temporal sampling in order to accurately monitor land surface displacement of these regions.

An upper bound on the GHG emissions due to the oxidation of organic matter in the soil is estimated from the per-parcel subsidence rates by assuming a worst-case fraction of subsidence caused by oxidation and the density of carbon contained within the soil. By spatial averaging, we obtain an average GHG emission intensity of 21.5 tonnes of CO_2-equivalents per hectare per year, and multiplying this value by the total sum of the areas of each parcel, we find a total upper bound emission rate of 2.3 Mt/year of CO_2-equivalents for the entire Green Heart region. This result demonstrates the feasibility of combining InSAR-derived motion estimates with soil data as a tool for assessing the GHG impact of land and water management in the region. The precision of the GHG estimates may be improved by combining the average irreversible subsidence rates with precise soil data such as that available from GeoTOP [122].

7

Conclusions and Recommendations

If you spend too much time thinking about a thing,
you will never get it done.

Bruce Lee

7.1 Conclusions

The overall conclusions of this work address the research gaps identified in Sec. 1.4. The research outcomes summarized below detail efforts made in order to enable the use of InSAR for the estimation of land surface motion of rapidly moving and decorrelating peatland environments, in particular the drained and cultivated peatlands of the Netherlands.

Shallow soft soil motion is driven by precipitation and evapotranspiration, and can be modelled to an accuracy within approximately 10% of the overall signal magnitude.

Two models for soft soil motion based on synoptic meteorological data were developed in Chapter 3. It was demonstrated that these are able to describe the observed motion of these soils to a remarkably accurate level, given their simplicity. Both models displayed testing RMSE values of approximately 10% of the overall signal magnitude or lower. A model based on machine learning techniques, a recurrent neural network (RNN), was able to capture sub-seasonal effects due to changing groundwater conditions within approximately 5 mm accuracy, but could not accurately model the multi-year irreversible subsidence signal. This is apparent during the winter periods, during which the model generally overestimated the surface level by 10–20 mm (Sec. 3.5.2), a significant deviation when considering annual subsidence rates in the range of 5–10 mm/year. This difficulty in generalizing to the multi-year time scale indicates that the model may be

overfit to the rapid sub-seasonal signal. On the other hand, a parametric model called Simple Parameterization for the Motion of Soils (SPAMS) was able to model both the highly dynamic sub-seasonal motions as well as the multi-year irreversible subsidence (Sec. 3.5.3) with an RMSE of approximately 1 cm or better (corresponding to about 10% of the signal magnitude) when validated against ground-based measurements at five test locations throughout the Netherlands. While the overall RMSE of this model was occasionally higher than that obtained by the RNN model, it was preferred for its overall simplicity and ability to generalize over both the sub-seasonal and multi-year time scales. SPAMS does not explicitly consider human interventions in the system such as pumping of excess water through the drainage ditches, and does not consider compaction, or oxidation occurring during wetting of the soil. Additional terms may be added to the model to account for these effects, at the cost of added complexity.

The ambiguities of dynamic signals exceeding a quarter-wavelength of LOS motion between acquisitions can be estimated by using contextual data to anticipate moments of strong subsidence and uplift.

Cycle slips are consistent phase unwrapping errors made by algorithms which cannot cope with displacements exceeding $\lambda/4$ between acquisitions. This condition is made worse when considering the strong noise levels encountered in cultivated peatland regions. A hidden Markov model based implementation for (explicit) temporal phase unwrapping is developed in Chapter 4 in order to overcome this problem. The method integrates categorical predictions about the state of relative motion from the machine learning model discussed in Sec. 3.3.2. Based on precipitation, temperature, and time inputs, the method compares an RNN prediction to the observed (wrapped) interferometric phase value in order to estimate whether or not an ambiguity level should be added or subtracted from the phase time series at the epoch in question. It is shown that standard minimum-gradient techniques fail to correctly handle rapid soil motions, which can create a strong bias (the magnitude of which corresponding to the number of cycle slips incurred) in the average annual subsidence rates estimated by C-band InSAR studies of peatland regions that rely on such phase unwrapping techniques (Sec. 4.4). This bias can be mitigated when correctly anticipating these large soil motions.

A full displacement time series can be inverted from interferometric data affected by loss-of-lock.

Temporal decorrelation strongly hinders the application of InSAR for monitoring land surface motion in highly dynamic areas. Previous studies focused on applying InSAR methods to peatland regions were constrained to monitoring motion within seasons or, worse, ignored the problem of decorrelation and produced physically unrealistic results. The concept of loss-of-lock was introduced in Chapter 5 to specifically describe coherence losses which result in a breakage of an interferometric time series, and describes how to identify and overcome them. The prototype of the Delft Contextually-Aided Distributed Scatterer Environment (DECADE) methodology was developed, which makes use of contextual data in order to reconnect coherent observations separated by loss-of-lock. The combined observations of groups of polygons which are expected to behave in a similar manner are used to estimate the parameters of a common displacement model.

This model is used to align the observations and estimate a single unbroken time series for the contextual group. The SPAMS model described in Sec. 3.4 was used for this task. While this method enabled the application of InSAR over decorrelating regions, the spatial resolution became limited by the contextual groups, and differently behaving parcels within a contextual group would not be detected.

It is possible to observe the shallow-based land surface motion of soft soils using InSAR. An average regional subsidence rate of 6.3 mm/year is obtained for the Green Heart for the period 2015–2023.

The DECADE methodology was further developed to scale up to large area (600 km^2) coverage over the full Sentinel-1 archive in order to estimate parcel-scale land surface motion in the Green Heart region of the Netherlands, and to supply estimates of multi-year average irreversible subsidence rates of the region. By identifying and mitigating the main problems preventing the use of (C-band) InSAR for this application, namely cycle slips and loss-of-lock, it becomes possible to use InSAR for land surface monitoring applications. By making use of the SPAMS model (Sec. 3.4), the DECADE InSAR methodology was further developed to enable parcel-scale time series estimation despite the loss-of-lock problem described in Chapter 5. This development represents a different paradigm in InSAR methodology, wherein the focus lies on estimating the parcel's intrinsic model parameters ($\hat{x}$) in order to learn how a given piece of land behaves, as opposed to focusing on obtaining as accurate of a time series as possible using a generic functional model ($\hat{y}$). Per-parcel land surface displacement time series estimates are provided for the Green Heart region, along with estimates of average irreversible subsidence rates and estimated displacement (SPAMS) model parameters. Sub-centimetre agreement is found for all test locations with in-situ measurements, corresponding to approximately 10% of the overall signal magnitude. The change in land surface motion caused by experimental subsidence mitigation measures is observable, which is also reflected in the estimated displacement model parameters.

An upper bound on greenhouse gas emissions from soil oxidation can be estimated from InSAR displacement data. An approximate upper bound of 21.5 t/ha/year of CO_2-equivalents are released by the oxidation of drying peat soils from the Green Heart, corresponding to a total annual rate of 2.3 Mt/year over the entire region.

Greenhouse gas (GHG) emissions caused by the oxidation of organic matter in the Green Heart's soft soils were estimated in Chapter 6, based on their known soil types, area, and estimated annual subsidence rates. There is considerable uncertainty in the fraction of subsidence caused by oxidation, versus subsidence caused by other processes which do not produce GHG emissions, which in turn creates uncertainty in GHG emission estimates. As such, the upper bound of observed values of this fraction was applied everywhere, and an upper bound on the expected level of GHG emissions from the oxidizing soils of the region was estimated. A total (maximum) value of 2.3 Mt/year of CO_2-equivalents for the entire Green Heart was estimated, which is approximately half of the national total estimated GHG emissions caused by oxidizing peatlands in the Netherlands, or 1.3% of the entire GHG emissions of the Netherlands in 2019 [14].

7.2 Contributions

This research has provided the following contributions to the field:

- Development of the SPAMS model for shallow soft soil motion
- Development of the DECADE methodology and software implementation for distributed scatterer InSAR analysis of decorrelating land surface motion
- Introduction of the term "loss-of-lock" into the InSAR literature to describe decorrelation events which result in breakages of an interferometric time series.
- First accurate and validated shallow soil surface motion time series estimates produced by InSAR in the Netherlands
- Estimated CO_2 impact of the period 2015–2023 from soil oxidation in the Dutch Green Heart region

7.3 Recommendations

Several areas for further study have been identified as a result of this research. As such, the following recommendations are identified.

The DECADE methodology has so far only been applied to grassland regions, which are not affected by agricultural activity such as plowing and harvesting (although the effects of vegetation growth and mowing are still present). While this was not a significant limitation in the Green Heart region where the vast majority of parcels were grasslands used for cattle grazing, this is not the case in other regions in the country where crop farming is more prevalent. Preliminary coherence studies of crop-covered parcels [123] indicate that high coherence levels are often encountered during the winter periods, corresponding to bare soils. During these periods, the coherence is generally higher than that of grasslands, suggesting that despite these farming activities, an estimation may still be possible. Furthermore, the DECADE methodology also does not consider any deep-based subsidence processes such as hydrocarbon extraction or mining, which are also not present in the Green Heart. Such deeply-based processes will need to be taken into account in order for a national-scale analysis to be performed.

The parameters of the SPAMS model are treated as static values which do not change with time. Variability in the displacement time series comes from the extrinsic precipitation and evapotranspiration values which are input to the model which has a certain set of (intrinsic) parameters. However, factors such as changing land use/land cover, anthropogenic interventions, and other aspects like changing parcel geometry and the introduction of subsidence mitigation systems could cause these parameters to change over time. Furthermore, as the time series of observations grows, it is desirable to avoid reprocessing of the entire SAR stack, but rather to simply recursively update the current estimates with the newest SAR image as it is acquired, as has been demonstrated with PS methods [124]. Such a method could potentially be used to account for changes in the intrinsic parcel properties over time.

The applicability of the SPAMS model may be extended to longer (multi-decadal) time scales by introducing a constant background subsidence rate related to long-term compaction and creep processes occurring in the Vadose zone. This could allow for the establishment of a more accurate linear subsidence rate, as well as allow for a more accurate separation of the component of subsidence caused by oxidation from the overall

irreversible subsidence rate, if the two rates can be appropriately distinguished during different times of the year (i.e. summer and winter). Furthermore, by comparing the relative magnitudes of such an additional term with the oxidation term, a more accurate fraction of subsidence caused by oxidation could potentially be retrieved, thereby reducing the uncertainty in GHG emissions based on the estimated subsidence rate.

The limitations of the EMI method need to be better understood in order to assure an accurately estimated consistent phase time series. It has been noted that the performance of EMI degrades as the number of images in an interferometric stack increases, as the ratio of coherent to incoherent data decreases quadratically with stack length (Sec. 6.2). The method can be further developed by applying the concept of the "sequential estimator" [125] for use in conjunction with the optimal stack size and the identification of loss-of-lock. Finally, the effects of closure phase on the ESM estimation should be better understood and quantified, as for example the implications of performing an ESM estimation over a region with changing land use/land cover over time have not been considered. Such changes may lead to additional closure phases or degraded estimation performance.

Identifying loss-of-lock is currently done by simply specifying a minimum threshold on the daisy-chain coherence level. While this is a practical solution which is easy to implement, it does not agree completely with how loss-of-lock is defined, and may result in sub-optimal separation of the time series into segments. An improved methodology that uses all the possible interferometric connections should be developed, and ideally combined with the EMI/sequential ESM estimation procedure (see recommendation above).

The stochastic model for distributed scatterers should be completed by formulating an estimate of the ESM and APS estimation quality to allow for error propagation from the observed coherence to the final estimated model parameters. Such a description would also require a model for the quality of the PS reference point to be included as well [126]. For large regions, the storage space requirements for saving the full covariance matrix for each multilooked distributed scatterer become very restrictive, however, space could be saved by only retaining the unmasked portion of the coherence matrix, as discussed in Sec. 6.2.1.

With the new availability of regular L-band SAR acquisitions of the Netherlands from SAOCOM as of mid-2023, a multi-sensor approach to monitoring is now possible in the region. The C-band Sentinel-1 and L-band SAOCOM missions are complementary in that Sentinel-1 provides frequent imagery that decorrelates quickly and is challenging to unwrap, whereas SAOCOM data is less prone to decorrelation and unwrapping errors but is provided more infrequently. These two sensors should be used in conjunction from 2023 onward in order to help constrain the model specification during C-band loss-of-lock events to ensure a more robust estimation of the irreversible subsidence rate.

Appendix A
Green Heart Soil Map

A simplified soil map of the Green Heart obtained from [100] based on the codes documented in [120] is shown in Figure A.1. Each of the four main classes shown contains numerous subclasses which have been combined in order to simplify interpretation of the figure.

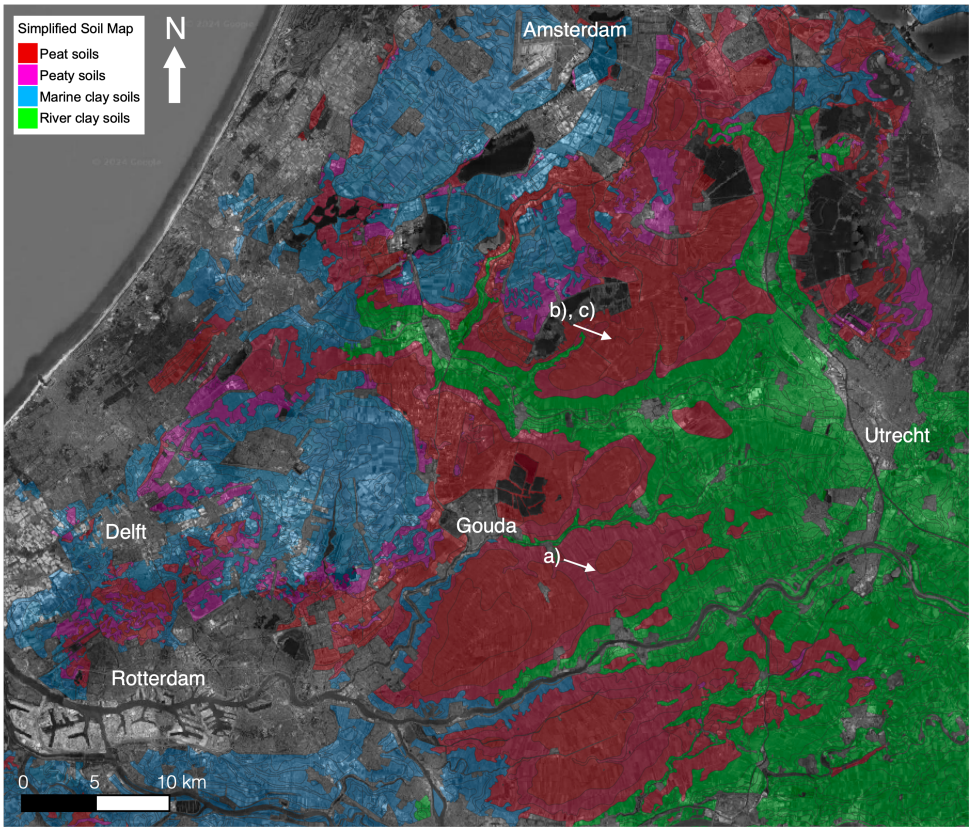

Figure A.1: Simplified soil map of the Green Heart region of the Netherlands showing peat soils, class “V” (red); “peaty” soils, class “W” (purple), marine clay soils, class “M” (blue); and river clay soils, class “R” (green). Classification from [120].

Appendix B
Subsidence Mitigation Measures

A simplified schematic diagram of the subsidence mitigation system [116] installed at the Zegveld location is shown in Fig. B.1. The objective of the system is to minimize groundwater fluctuations in the centre of the parcel by means of a pressurized drain, thereby reducing the effects of shrinkage and swell, and the drying of the upper layers of soils which leads to oxidation.

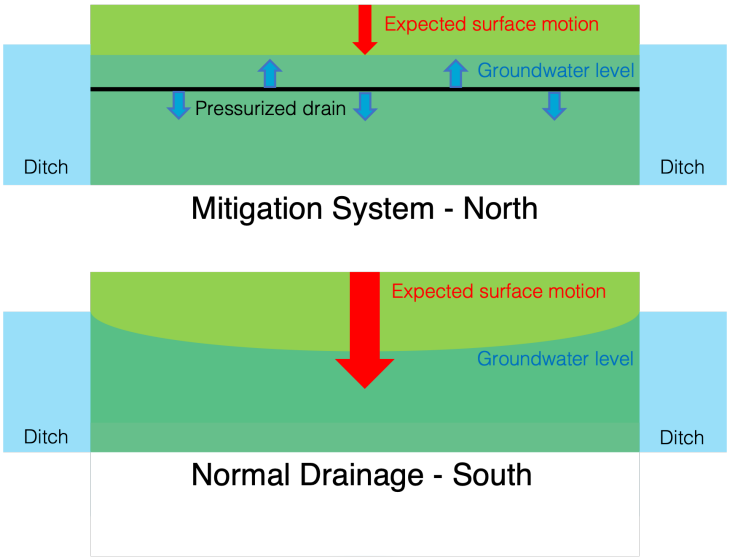

Figure B.1: Simplified schematic visualizing the difference between a normally managed parcel ('South' test plot) and the subsidence mitigation measure ('North' test plot).

An aerial view of the Zegveld test parcel is shown in Fig. B.2. The mitigation system is installed at the northern end of the parcel. In order to process the two halves separately, the parcel polygon is divided as shown by the blue lines.

Figure B.2: Aerial view of Zegveld test parcel, with polygon divisions visualized in blue. South sub-parcel: standard drainage system, site b). North sub-parcel: mitigated drainage system, site c).

Acknowledgements

This research is part of the Living on Soft Soils: Subsidence and Society (grant no.: NWA.1160.18.259) project, and is funded by the Dutch Research Council (NWO-NWA-ORC), Utrecht University, Wageningen University, Delft University of Technology, the Ministry of Infrastructure and Water Management, the Ministry of the Interior and Kingdom Relations, Deltares, Wageningen Environmental Research, the Geological Survey of the Netherlands (TNO), STOWA, Water Authority Hoogheemraadschap de Stichtse Rijnlanden, Water Authority Drents Overijsselse Delta, the Province of Utrecht, the Province of Zuid-Holland, the Municipality of Gouda, Platform Slappe Bodem, Sweco, Tauw BV, and NAM.

Bibliography

[1] S. van Asselen, E. Stouthamer, and T. van Asch, "Effects of peat compaction on delta evolution: A review on processes, responses, measuring and modeling", *Earth-Science Reviews*, vol. 92, no. 1, pp. 35–51, 2009. DOI: 10.1016/j.earscirev.2008.11.001.

[2] G. Erkens, M. J. van der Meulen, and H. Middelkoop, "Double trouble: Subsidence and CO2 respiration due to 1,000 years of Dutch coastal peatlands cultivation", *Hydrogeology Journal*, vol. 24, no. 3, pp. 551–568, 2016. DOI: 10.1007/s10040-016-1380-4.

[3] J. Pärn, J. Verhoeven, K. Butterbach-Bahl, *et al.*, "Nitrogen-rich organic soils under warm well-drained conditions are global nitrous oxide emission hotspots", *Nature Communications*, vol. 9, no. 1135, 2018. DOI: 10.1038/s41467-018-03540-1.

[4] C. Schouten, *Letter to Parliament by the Dutch Minister of Agriculture on the reduction of greenhouse gas emissions from peatlands (Dutch)*, 2020. [Online]. Available: https://www.rijksoverheid.nl/documenten/publicaties/2020/07/13/verplicht-format-bijlage-onderbouwing-en-evaluatie-veenplan-1e-fase.

[5] G. Van den Born, F. Kragt, D. Henkens, *et al.*, "Dalende bodems, stijgende kosten: Mogelijke maatregelen tegen veenbodemdaling in het landelijk en stedelijk gebied: Beleidsstudie", Tech. Rep., 2016. [Online]. Available: https://www.pbl.nl/downloads/pbl-2016-dalende-bodems-stijgende-kosten-1064pdf.

[6] P. Kuikman, J. van den Akker, and F. de Vries, "Emissie van N2O en CO2 uit organische landbouwbodems (in Dutch)", Dutch, Wageningen Environmental Research, Wageningen, the Netherlands, Tech. Rep. 1035-2, 2005. [Online]. Available: https://research.wur.nl/en/publications/emissie-van-n2o-en-co2-uit-organische-landbouwbodems.

[7] J. van den Akker, P. Kuikman, F. de Vries, *et al.*, "Emission of CO2 from agricultural peat soils in the Netherlands and ways to limit this emission", 2008, pp. 645–648, ISBN: 9780951489048.

[8] T. Hoogland, J. Van den Akker, and D. Brus, "Modeling the subsidence of peat soils in the Dutch coastal area", *Geoderma*, vol. 171, pp. 92–97, 2012.

[9] H. Bootsma, H. Kooi, and G. Erkens, "Atlantis, a tool for producing national predictive land subsidence maps of the Netherlands", *Proceedings of the International Association of Hydrological Sciences*, vol. 382, pp. 415–420, 2020. DOI: 10.5194/piahs-382-415-2020.

[10] C. Schothorst, "Subsidence of low moor peat soils in the western Netherlands", *Geoderma*, vol. 17, no. 4, pp. 265–291, 1977, ISSN: 0016-7061. DOI: 10.1016/0016-7061(77)90089-1.

[11] J. van Huissteden, R. van den Bos, and I. Marticorena Alvarez, "Modelling the effect of water-table management on CO2 and CH4 fluxes from peat soils", *Netherlands Journal of Geosciences - Geologie en Mijnbouw*, vol. 85, no. 1, pp. 3–18, 2006. DOI: 10.1017/s0016774600021399.

[12] M. J. van der Meulen, A. J. F. van der Spek, G. de Lange, *et al.*, "Regional sediment deficits in the dutch lowlands: Implications for long-term land-use options (8 pp)", *Journal of Soils and Sediments*, vol. 7, no. 1, pp. 9–16, 2007, ISSN: 1614-7480. DOI: 10.1065/jss2006.12.199.

[13] S. van Asselen, G. Erkens, and F. de Graaf, "Monitoring shallow subsidence in cultivated peatlands", *Proceedings of the International Association of Hydrological Sciences*, vol. 382, pp. 189–194, 2020. DOI: 10.5194/piahs-382-189-2020.

[14] P. Ruyssenaars, P. Coenen, J. Rienstra, P. Zijlema, E. Arets, *et al.*, "Greenhouse gas emissions in the Netherlands 1990–2019: National inventory report 2021", National Institute for Public Health and the Environment (RIVM), Tech. Rep., 2021. DOI: 10.21945/RIVM-2021-0007.

[15] E. Arets, S. van Baren, C. Hendriks, H. Kramer, J. Lesschen, and M. Schelhaa, "Greenhouse gas reporting of the LULUCF sector in the Netherlands. methodological background, update 2023", Statutory Research Tasks Unit for Nature and the Environment, Wageningen, Tech. Rep., 2023. [Online]. Available: https://www.emissieregistratie.nl/sites/default/files/2023-04/2023%20(WOt%20Natuur%20&%20Milieu)%20Metodology%20for%20report%20LULUCF%20ER%201990-2021.pdf.

[16] J. Zou, A. D. Ziegler, D. Chen, *et al.*, "Rewetting global wetlands effectively reduces major greenhouse gas emissions", *Nature Geoscience*, 2022. DOI: 10.1038/s41561-022-00989-0.

[17] T. V. Armentano and E. S. Menges, "Patterns of change in the carbon balance of organic soil-wetlands of the temperate zone", *Journal of Ecology*, vol. 74, no. 3, pp. 755–774, 1986.

[18] S. van Asselen, G. Erkens, E. Stouthamer, H. A. Woolderink, R. E. Geeraert, and M. M. Hefting, "The relative contribution of peat compaction and oxidation to subsidence in built-up areas in the Rhine-Meuse delta, the Netherlands", *Science of The Total Environment*, vol. 636, pp. 177–191, 2018, ISSN: 0048-9697. DOI: 10.1016/j.scitotenv.2018.04.141.

[19] Y. Morishita and R. F. Hanssen, "Deformation parameter estimation in low coherence areas using a multisatellite InSAR approach", *IEEE Trans. Geosci. and Remote Sens.*, vol. 53, no. 8, pp. 4275–4283, 2015. DOI: 10.1109/TGRS.2015.2394394.

[20] S. S. Esfahany, "Exploitation of distributed scatterers in synthetic aperture radar interferometry", Ph.D. dissertation, TU Delft, 2017. DOI: 10.4233/uuid:22d46f1e-9061-46b0-9726-760c41404b6f.

[21] L. Alshammari, D. J. Large, D. S. Boyd, *et al.*, "Long-term peatland condition assessment via surface motion monitoring using the ISBAS DInSAR technique over the Flow Country, Scotland", *Remote Sensing*, vol. 10, no. 7, 2018, ISSN: 2072-4292. DOI: 10.3390/rs10071103.

[22] T. Tampuu, J. Praks, R. Uiboupin, and A. Kull, "Long term interferometric temporal coherence and DInSAR phase in northern peatlands", *Remote Sensing*, vol. 12, no. 10, 2020, ISSN: 2072-4292. DOI: 10.3390/rs12101566.

[23] T. Tampuu, F. de Zan, R. Shau, J. Praks, M. Kohv, and A. Kull, "Reliability of Sentinel-1 InSAR distributed scatterer (DS) time series to estimate the temporal vertical movement of ombrotrophic bog surface", in *EGU General Assembly*, vol. EGU22-2387, Vienna, Austria, 2022. DOI: 10.5194/egusphere-egu22-2387.

[24] M. Caro Cuenca and R. F. Hanssen, "Subsidence due to peat decomposition in the Netherlands, kinematic observations from radar interferometry", in *Proceedings of the ESA Fringe Workshop*, Frascati, Italy, 2008, pp. 1–6.

[25] P. Berardino, G. Fornaro, R. Lanari, and E. Sansosti, "A new algorithm for surface deformation monitoring based on small baseline differential SAR interferograms", *IEEE Transactions on Geoscience and Remote Sensing*, vol. 40, no. 11, pp. 2375–2383, 2002.

[26] Y. Morishita and R. F. Hanssen, "Temporal decorrelation in L-, C-, and X-band satellite radar interferometry for pasture on drained peat soils", *IEEE Transactions on Geoscience and Remote Sensing*, vol. 53, no. 2, pp. 1096–1104, 2015. DOI: 10.1109/TGRS.2014.2333814.

[27] S. Samiei Esfahany and R. Hanssen, "Feasibility assessment of exploitation of distributed scatterers in InSAR stacks over pasture areas for different SAR satellite missions", *Earth Observation and Geomatics Engineering*, vol. 2, no. 2, pp. 64–73, 2018, ISSN: 2588-4352. DOI: 10.22059/eoge.2018.267528.1032.

[28] F. M. Heuff and R. F. Hanssen, "InSAR phase reduction using the remove-compute-restore method", in *2020 IEEE International Geoscience and Remote Sensing Symposium*, 2020. DOI: 10.1109/IGARSS39084.2020.9323720.

[29] H. Ansari, F. De Zan, and R. Bamler, "Efficient phase estimation for interferogram stacks", *IEEE Transactions on Geoscience and Remote Sensing*, vol. 56, no. 7, pp. 4109–4125, 2018. DOI: 10.1109/TGRS.2018.2826045.

[30] G. Kennedy and J. Price, "A conceptual model of volume-change controls on the hydrology of cutover peats", *Journal of Hydrology*, vol. 302, no. 1, pp. 13–27, 2005. DOI: 10.1016/j.jhydrol.2004.06.024.

[31] M. Camporese, S. Ferraris, M. Putti, P. Salandin, and P. Teatini, “Hydrological modeling in swelling/ shrinking peat soils”, *Water Resources Research*, vol. 42, no. 6, 2006. DOI: 10.1029/2005WR004495.

[32] NOBV, *Nationaal onderzoeksprogramma broeikasgassen veenweiden*, 2023. [Online]. Available: https://www.nobveenweiden.nl/en/what-we-measure/database/.

[33] R. F. Hanssen, *Radar Interferometry: Data Interpretation and Error Analysis*. Dordrecht: Kluwer Academic Publishers, 2001, ISBN: 978-0-7923-6945-5. DOI: 10.1007/0-306-47633-9.

[34] F. Ulaby and D. Long, *Microwave Radar and Radiometric Remote Sensing*. The University of Michigan Press, 2014, ISBN: 978-0-472-11935-6.

[35] K. Tomiyasu, “Tutorial review of synthetic-aperture radar (SAR) with applications to imaging of the ocean surface”, *Proceedings of the IEEE*, vol. 66, no. 5, pp. 563–583, 1978. DOI: 10.1109/PROC.1978.10961.

[36] A. Moreira, P. Prats-Iraola, M. Younis, G. Krieger, I. Hajnsek, and K. P. Papathanassiou, “A tutorial on synthetic aperture radar”, *IEEE Geoscience and Remote Sensing Magazine*, vol. 1, no. 1, pp. 6–43, 2013. DOI: 10.1109/MGRS.2013.2248301.

[37] H. Ansari, “Efficient high-precision time series analysis for synthetic aperture radar interferometry”, Ph.D. dissertation, Technical University of Munich, 2018.

[38] P. Dheenathayalan, D. Small, A. Schubert, and R. F. Hanssen, “High-precision positioning of radar scatterers”, *Journal of Geodesy*, vol. 90, no. 5, pp. 403–422, 2016. DOI: 10.1007/s00190-015-0883-4.

[39] M. Yang, P. Lopez-Dekker, P. Dheenathayalan, M. Liao, and R. F. Hanssen, “On the value of corner reflectors and surface models in InSAR precise point positioning”, *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 158, pp. 113–122, 2019. DOI: 10.1016/j.isprsjprs.2019.10.006.

[40] C. W. Sherwin, J. P. Ruina, and R. D. Rawcliffe, “Some early developments in synthetic aperture radar systems”, *IRE Transactions on Military Electronics*, vol. MIL-6, no. 2, pp. 111–115, 1962. DOI: 10.1109/IRET-MIL.1962.5008415.

[41] C. Wiley, “Pulsed Doppler radar methods and apparatus”, US3196436A, 1965.

[42] F. Hu, J. Wu, L. Chang, and R. F. Hanssen, “Incorporating temporary coherent scatterers in multi-temporal InSAR using adaptive temporal subsets”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 10, pp. 7658–7670, 2019. DOI: 10.1109/TGRS.2019.2915658.

[43] A. Ferretti, C. Prati, and F. Rocca, “Nonlinear subsidence rate estimation using permanent scatterers in differential SAR interferometry”, *IEEE Transactions on geoscience and remote sensing*, vol. 38, no. 5, pp. 2202–2212, 2000.

[44] A. Ferretti, C. Prati, and F. Rocca, “Permanent scatterers in SAR interferometry”, *IEEE Transactions on geoscience and remote sensing*, vol. 39, no. 1, pp. 8–20, 2001.

[45] M. Yang, P. Lopez-Dekker, P. Dheenathayalan, F. Biljecki, M. Liao, and R. F. Hanssen, “Linking persistent scatterers to the built environment using ray tracing on urban models”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 8, pp. 5764–5776, 2019. DOI: 10.1016/j.isprsjprs.2019.10.006.

[46] L. Chang and R. F. Hanssen, “Detection of cavity migration and sinkhole risk using radar interferometric time series”, *Remote sensing of environment*, vol. 147, pp. 56–64, 2014. DOI: 10.1016/j.rse.2014.03.002.

[47] L. Chang, R. P. Dollevoet, and R. F. Hanssen, “Nationwide railway monitoring using satellite SAR interferometry”, *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 10, no. 2, pp. 596–604, 2016. DOI: 10.1109/JSTARS.2016.2584783.

[48] L. Chang, R. P. Dollevoet, and R. F. Hanssen, “Monitoring line-infrastructure with multisensor SAR interferometry: Products and performance assessment metrics”, *IEEE journal of selected topics in applied earth observations and remote sensing*, vol. 11, no. 5, pp. 1593–1605, 2018. DOI: 10.1109/JSTARS.2018.2803074.

[49] A. J. Hooper, “Persistent scatterer radar interferometry for crustal deformation studies and modeling of volcanic deformation”, Ph.D. dissertation, Stanford, 2006.

[50] H. A. Zebker, J. Villasenor, *et al.*, “Decorrelation in interferometric radar echoes”, *IEEE Transactions on geoscience and remote sensing*, vol. 30, no. 5, pp. 950–959, 1992.

[51] “IEEE standard letter designations for radar-frequency bands”, *IEEE Std 521-2019 (Revision of IEEE Std 521-2002)*, 2020. DOI: 10.1109/IEEESTD.2020.8999849.

[52] J. J. van Zyl, “The shuttle radar topography mission (SRTM): A breakthrough in remote sensing of topography”, *Acta Astronautica*, vol. 48, no. 5, pp. 559–565, 2001, ISSN: 0094-5765. DOI: 10.1016/S0094-5765(01)00020-0.

[53] H. A. Zebker and R. M. Goldstein, “Topographic mapping from interferometric synthetic aperture radar observations”, *Journal of Geophysical Research: Solid Earth*, vol. 91, no. B5, pp. 4993–4999, 1986. DOI: 10.1029/JB091iB05p04993.

[54] F. van Leijen, “Persistent scatterer interferometry based on geodetic estimation theory”, Ph.D. dissertation, TU Delft, 2014, ISBN: 9789461862990.

[55] G. Mulder, F. J. van Leijen, and R. F. Hanssen, “A generic approach to parameterize the turbulent energy of single-epoch atmospheric delays from InSAR time series”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–13, 2023. DOI: 10.1109/TGRS.2023.3295898.

[56] M. Born, E. Wolf, and A. B. Bhatia, *Principles of optics: Electromagnetic theory of propagation, interference and diffraction of light*. New York: Pergamon Press, 1959.

[57] R. Touzi, A. Lopes, J. Bruniquel, and P. Vachon, “Coherence estimation for SAR imagery”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 37, no. 1, pp. 135–149, 1999. DOI: 10.1109/36.739146.

[58] A. Ferretti, A. Fumagalli, F. Novali, C. Prati, F. Rocca, and A. Rucci, “A new algorithm for processing interferometric data-stacks: SqueeSAR”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 9, pp. 3460–3470, 2011, ISSN: 1558-0644. DOI: 10.1109/TGRS.2011.2124465.

[59] M. Jiang, X. Ding, R. F. Hanssen, R. Malhotra, and L. Chang, “Fast statistically homogeneous pixel selection for covariance matrix estimation for multitemporal InSAR”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 53, no. 3, pp. 1213–1224, 2015. DOI: 10.1109/TGRS.2014.2336237.

[60] A. Parizzi and R. Brcic, “Adaptive InSAR stack multilooking exploiting amplitude statistics: A comparison between different techniques and practical results”, *IEEE Geoscience and Remote Sensing Letters*, vol. 8, no. 3, pp. 441–445, 2011. DOI: 10.1109/LGRS.2010.2083631.

[61] R. Touzi and A. Lopes, “Statistics of the Stokes parameters and of the complex coherence parameters in one-look and multilook speckle fields”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 34, no. 2, pp. 519–531, 1996.

[62] B. C. Barber, “The phase statistics of a multichannel radar interferometer”, *Waves in Random Media*, vol. 3, no. 4, pp. 257–266, 1993. DOI: 10.1088/0959-7174/3/4/002.

[63] J. S. Lee, K. Hoppel, S. Mango, and A. Miller, “Intensity and phase statistics of multilook polarimetric and interferometric SAR imagery”, *IEEE Trans. Geosci. and Remote Sens.*, vol. 32, no. 5, pp. 1017–1028, 1994. DOI: 10.1109/36.312890.

[64] D. Just and R. Bamler, “Phase statistics of interferograms with applications to synthetic aperture radar”, *Appl. Opt.*, vol. 33, no. 20, pp. 4361–4368, 1994. DOI: 10.1364/AO.33.004361.

[65] F. De Zan, A. Parizzi, P. Prats-Iraola, and P. López-Dekker, “A SAR interferometric model for soil moisture”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 52, no. 1, pp. 418–425, 2014, ISSN: 1558-0644. DOI: 10.1109/TGRS.2013.2241069.

[66] F. De Zan, M. Zonno, and P. Lopez-Dekker, “Phase inconsistencies and multiple scattering in SAR interferometry”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 53, no. 12, pp. 6608–6616, 2015. DOI: 10.1109/TGRS.2015.2444431.

[67] H. Ansari, F. De Zan, and A. Parizzi, "Study of systematic bias in measuring surface deformation with SAR interferometry", *IEEE Transactions on Geoscience and Remote Sensing*, vol. 59, no. 2, pp. 1285–1301, 2021. DOI: 10.1109/TGRS.2020.3003421.

[68] Y. Yuan, M. Kleinherenbrink, and P. Lopez-Dekker, "On crop growth and InSAR closure phases", *IEEE Transactions on Geoscience and Remote Sensing*, 2024. DOI: 10.1109/TGRS.2024.3432396.

[69] E. Wig, R. Michaelides, and H. Zebker, "Fine-resolution measurement of soil moisture from cumulative InSAR closure phase", *IEEE Transactions on Geoscience and Remote Sensing*, vol. 62, pp. 1–15, 2024. DOI: 10.1109/TGRS.2024.3399069.

[70] P. Teunissen, "Integer aperture GNSS ambiguity resolution", *Artificial Satellites*, vol. 38, no. 3, pp. 79–88, 2003.

[71] C. W. Chen and H. A. Zebker, "Two-dimensional phase unwrapping with use of statistical models for cost functions in nonlinear optimization", *J. Opt. Soc. Am. A*, vol. 18, no. 2, pp. 338–351, Feb. 2001. DOI: 10.1364/JOSAA.18.000338.

[72] M. Costantini, "A novel phase unwrapping method based on network programming", *IEEE Transactions on Geoscience and Remote Sensing*, vol. 36, no. 3, pp. 813–821, 1998. DOI: 10.1109/36.673674.

[73] B. Kampes and R. Hanssen, "Ambiguity resolution for permanent scatterer interferometry", *IEEE Transactions on Geoscience and Remote Sensing*, vol. 42, no. 11, pp. 2446–2453, 2004. DOI: 10.1109/TGRS.2004.835222.

[74] P. Teunissen, "On InSAR ambiguity resolution for deformation monitoring", *Artificial Satellites*, vol. 41, no. 1, pp. 19–22, 2006.

[75] P. Conroy, S. A. van Diepen, S. van Asselen, G. Erkens, F. J. van Leijen, and R. F. Hanssen, "Probabilistic estimation of InSAR displacement phase guided by contextual information and artificial intelligence", *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–11, 2022. DOI: 10.1109/TGRS.2022.3203872.

[76] P. Conroy, S. A. N. van Diepen, and R. F. Hanssen, "SPAMS: A new empirical model for soft soil surface displacement based on meteorological input data", *Geoderma*, 2023. DOI: 10.1016/j.geoderma.2023.116699.

[77] J. P. Hsi, J. P. Carter, and J. C. Small, "Surface subsidence and drawdown of the water table due to pumping", *Géotechnique*, vol. 44, no. 3, pp. 381–396, 1994. DOI: 10.1680/geot.1994.44.3.381.

[78] J. Mas-Pla, A. Rodríguez-Florit, M. Zamorano, C. Roqué, A. Menció, and D. Brusi, "Anticipating the effects of groundwater withdrawal on seawater intrusion and soil settlement in urban coastal areas", *Hydrological Processes*, vol. 27, no. 16, pp. 2352–2366, 2013. DOI: 10.1002/hyp.9377.

[79] D. Peduto, A. Prosperi, G. Nicodemo, and M. Korff, "District-scale numerical analysis of settlements related to groundwater lowering in variable soil conditions", *Canadian Geotechnical Journal*, vol. 59, no. 6, pp. 978–993, 2022. DOI: 10.1139/cgj-2021-0041.

[80] J. van Dam and R. Feddes, "Numerical simulation of infiltration, evaporation and shallow groundwater levels with the Richards equation", *Journal of Hydrology*, vol. 233, no. 1, pp. 72–85, 2000, ISSN: 0022-1694. DOI: 10.1016/S0022-1694(00)00227-4.

[81] J. van der Gaast, H. Vroon, and H. Massop, *Grondwaterregime op basis van karteerbare kenmerken (Dutch)*, Dutch. STOWA, 2010, ISBN: 9789057735011.

[82] W. J. De Lange, G. F. Prinsen, J. C. Hoogewoud, *et al.*, "An operational, multi-scale, multi-model system for consensus-based, integrated water management and policy analysis: The Netherlands hydrological instrument", *Env. Modelling and Softw.*, vol. 59, pp. 98–108, 2014. DOI: 10.1016/j.envsoft.2014.05.009.

[83] G. Erkens, "National research program on greenhouse gases in peatlands (NOBV) data analysis (Dutch)", NOBV Consortium, Tech. Rep., 2021.

[84] H. A. R. de Bruin, "From Penman to Makkink", in *Evaporation and Weather: Proceedings and Information*, TNO Committee on Hydrological Research, vol. 28, 1987.

[85] P. Hiemstra and R. Sluiter, “Interpolation of Makkink evaporation in the Netherlands”, Royal Netherlands Meteorological Institute, Tech. Rep., 2011.

[86] A. F. G. Jacobs and H. A. R. de Bruin, “Makkink’s equation for evapotranspiration applied to unstressed maize”, *Hydrological Processes*, vol. 12, no. 7, pp. 1063–1066, 1998.

[87] O. Tetens, “Über einige meteorologische Begriffe”, *Z. Geophys*, vol. 6, 1930.

[88] S. Hochreiter and J. Schmidhuber, “Long short-term memory”, *Neural Comput.*, vol. 9, no. 8, 1997. DOI: 10.1162/neco.1997.9.8.1735.

[89] G. E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Improving neural networks by preventing co-adaptation of feature detectors”, *arXiv e-prints*, arXiv:1207.0580, 2012.

[90] Z. Zhang and M. R. Sabuncu, “Generalized cross entropy loss for training deep neural networks with noisy labels”, in *32nd Conf. Neural Inf. Process. Syst.*, 2018. arXiv: 1805.07836.

[91] Y. Yihdego and J. Webb, “An empirical water budget model as a tool to identify the impact of land-use change in stream flow in Southeastern Australia”, *Water Resources Management*, no. 27, 2013. DOI: 10.1007/s11269-013-0449-2.

[92] DINOloket, *Data and information on the Dutch subsurface*, TNO Geological Survey of The Netherlands, 2023. [Online]. Available: www.dinoloket.nl/en/subsurface-data.

[93] R. M. Goldstein, H. A. Zebker, and C. L. Werner, “Satellite radar interferometry: Two-dimensional phase unwrapping”, *Radio Science*, vol. 23, no. 4, pp. 713–720, 1988. DOI: 10.1029/RS023i004p00713.

[94] R. Bamler and P. Hartl, “Synthetic aperture radar interferometry”, *Inverse Problems*, vol. 14, pp. 1–54, 1998.

[95] H. Yu, Y. Lan, Z. Yuan, J. Xu, and H. Lee, “Phase unwrapping in InSAR: A review”, *IEEE Geoscience and Remote Sensing Magazine*, vol. 7, no. 1, pp. 40–58, 2019. DOI: 10.1109/MGRS.2018.2873644.

[96] L. Zhou, H. Yu, Y. Lan, and M. Xing, “Artificial intelligence in interferometric synthetic aperture radar phase unwrapping: A review”, *IEEE Geoscience and Remote Sensing Magazine*, vol. 9, no. 2, pp. 10–28, 2021. DOI: 10.1109/MGRS.2021.3065811.

[97] A. Viterbi, “Error bounds for convolutional codes and an asymptotically optimum decoding algorithm”, *IEEE Transactions on Information Theory*, vol. 13, no. 2, pp. 260–269, 1967. DOI: 10.1109/TIT.1967.1054010.

[98] P. Conroy, S. A. van Diepen, F. J. van Leijen, and R. F. Hanssen, “Bridging loss-of-lock in InSAR time series of distributed scatterers”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, no. 5220911, 2023. DOI: 10.1109/TGRS.2023.3329967.

[99] B. Kampes, R. F. Hanssen, and Z. Perski, “Radar interferometry with public domain tools”, in *Proceedings of FRINGE 2003*, 2003.

[100] *Publieke Dienstverlening Op de Kaart (PDOK)*, www.pdok.nl.

[101] S. Samiei-Esfahany, J. E. Martins, F. van Leijen, and R. F. Hanssen, “Phase estimation for distributed scatterers in InSAR stacks using integer least squares estimation”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 10, pp. 5671–5687, 2016. DOI: 10.1109/TGRS.2016.2566604.

[102] KNMI, *Daggegevens van het weer in Nederland*, 2023. [Online]. Available: https://www.knmi.nl/nederland-nu/klimatologie/daggegevens.

[103] K. Pepin and H. Zebker, “Aliasing in InSAR and SBAS time series”, in *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, 2021, pp. 2663–2666. DOI: 10.1109/IGARSS47720.2021.9555161.

[104] P. Teunissen, “Least-squares estimation of the integer GPS ambiguities”, In Publications and annual report 1993, number 6 in LGR-Series, pages 59–74. Delft Geodetic Computing Centre., 1993.

[105] P. Teunissen, “Success probability of integer GPS ambiguity rounding and bootstrapping”, *Journal of Geodesy*, vol. 72, pp. 606–612, 1998. DOI: 10.1007/s001900050199.

[106] L. van der Maaten and G. Hinton, “Visualizing data using t-SNE”, *Journal of Machine Learning Research*, 2008, ISSN: 1532-4435.

[107] P. Conroy and R. F. Hanssen, “Towards automatic functional model specification for distributed scatterers using t-SNE”, in *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, 2021, pp. 2659–2662. DOI: 10.1109/IGARSS47720.2021.9553123.

[108] P. Conroy and R. F. Hanssen, “An upper bound on carbon emissions of drained peat soils from satellite radar interferometry (submitted for publication)”, *Geophysical Research Letters*, 2025.

[109] B. M. Kampes, *Radar Interferometry: Persistent Scatterer Technique*. Dordrecht, The Netherlands: Springer, 2006. DOI: 10.1007/978-1-4020-4723-7.

[110] R. Hanssen, “A radar retroreflector device and a method of preparing a radar retroreflector device (patent filed 2017)”, English, NL2019103B1, 2019. [Online]. Available: https://patents.google.com/patent/NL2019103B1/nl.

[111] L. Chang and R. F. Hanssen, “A probabilistic approach for InSAR time-series postprocessing”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 1, pp. 421–430, 2016. DOI: 10.1109/TGRS.2015.2459037.

[112] W. S. Brouwer, Y. Wang, F. J. van Leijen, and R. F. Hanssen, “On the stochastic model for InSAR single arc point scatterer time series”, in *IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium*, 2023, pp. 7902–7905. DOI: 10.1109/IGARSS52108.2023.10282629.

[113] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright, “Convergence properties of the Nelder-Mead simplex method in low dimensions”, *SIAM Journal on Optimization*, vol. 9, no. 1, pp. 112–147, 1998. DOI: 10.1137/S1052623496303470.

[114] W. S. Brouwer and R. F. Hanssen, “A treatise on InSAR geometry and 3-D displacement estimation”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–11, 2023. DOI: 10.1109/TGRS.2023.3322595.

[115] P. J. Teunissen, *Testing Theory: An Introduction*, 3rd ed. TU Delft OPEN Publishing, 2024. DOI: 10.59490/tb.96.

[116] I. Hoving, P. Vereijken, K. van Houwelingen, and M. Pleijter, “Hydrologische en landbouwkundige effectentoepassing onderwaterdrains bij dynamisch slootpeilbeheer op veengrond (Dutch)”, Wageningen UR Livestock Research, Tech. Rep., 2013. [Online]. Available: https://edepot.wur.nl/265293.

[117] J. Beuving and J. van den Akker, “Maaiveldsdaling van veengrasland bij twee slootpeilen in de polder Zegvelderbroek (in Dutch)”, Wageningen Environemental Research, Tech. Rep., 1996.

[118] P. Jansen, R. Hendriks, and C. Kwakernaak, “Behoud van veenbodems door ander peilbeheer : Maatregelen voor een robuuste inrichting van het westelijk veenweidegebied (in Dutch)”, Integrated Water Resources Management, Wageningen Environmental Research, Tech. Rep., 2009.

[119] European Commission, *Degraded ecosystems to be restored across Europe as Nature Restoration Law enters into force*, Press Release, Aug. 2024. [Online]. Available: https://environment.ec.europa.eu/news/nature-restoration-law-enters-force-2024-08-15_en.

[120] H. de Bakker and J. Schelling, *System van bodemsclassificattie voor Nederland (in Dutch)*, 2nd ed. Wageningen, the Netherlands: Winand Staring Centre, 1989, ISBN: 90-220-0997-1.

[121] R. Hendriks, “Nutrientenbelasting van oppervlaktewater in veenweidegebieden (in Dutch)”, Wageningen Environmental Research, Tech. Rep., 1993. [Online]. Available: https://edepot.wur.nl/303313.

[122] Basisregistratie Ondergrond, *Geotop*, Database, 2024. [Online]. Available: https://basisregistratieondergrond.nl/inhoud-bro/registratieobjecten/modellen/geotop-gtm/.

[123] P. Conroy, Y. Lumban-Gaol, S. van Diepen, F. van Leijen, and R. F. Hanssen, “First wide-area Dutch peatland subsidence estimates based on InSAR”, in *2024 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, Athens, Greece, 2024. DOI: 10.1109/IGARSS53475.2024.10642504.

[124] Y. Wang, W. Brouwer, F. van Leijen, and R. F. Hanssen, “Constrained recursive parameter estimation for InSAR arcs”, in *2024 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, Athens, Greece, 2024.

[125] H. Ansari, F. De Zan, and R. Bamler, “Sequential estimator: Toward efficient InSAR time series analysis”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 55, no. 10, pp. 5637–5652, 2017, ISSN: 1558-0644. DOI: 10.1109/TGRS.2017.2711037.

[126] W. Brouwer and R. F. Hanssen, “On the treatment of the reference image for InSAR parameter estimation for point scatterers”, in *2024 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, Athens, Greece, 2024.

List of Publications

Journal Papers

1. **P. Conroy**, S.A.N. van Diepen, S. Van Asselen, G. Erkens, F.J. van Leijen, R.F. Hanssen, "Probabilistic Estimation of InSAR Displacement Phase Guided by Contextual Information and Artificial Intelligence", *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 60, No. 5234611, 2022.
2. **P. Conroy**, S.A.N. van Diepen, F.J. van Leijen, R.F. Hanssen, "Bridging Loss-of-Lock in InSAR Time Series of Distributed Scatterers", *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 61, No. 5220911, 2023.
3. **P. Conroy**, S.A.N. van Diepen, R.F. Hanssen, "SPAMS: A New Empirical Model for Soft Soil Surface Displacement Based on Meteorological Input Data", *Geoderma*, Vol. 440, No. 116699, 2023.
4. **P. Conroy**, R.F. Hanssen, "An Upper Bound on Carbon Emissions of Drained Peat Soils from Satellite Radar Interferometry", *Geophysical Research Letters*, Submitted, 2025.

Conference Papers

1. **P. Conroy**, R.F. Hanssen, "Towards Automatic Functional Model Specification for Distributed Scatterers Using T-SNE", *Proceedings of the 2021 IEEE International Geoscience and Remote Sensing Symposium*, Brussels, Belgium, 2021.
2. **P. Conroy**, S.A.N. van Diepen, F.J. van Leijen, R.F. Hanssen, "Hybrid InSAR Processing for Rapidly Deforming Peatlands Aided by Contextual Information", *Proceedings of the 2022 IEEE International Geoscience and Remote Sensing Symposium*, Kuala Lumpur, Malaysia, 2022.
3. **P. Conroy**, S.A.N. van Diepen, F.J. van Leijen, R.F. Hanssen, "First Results of Dutch Peatland Subsidence Observations Using InSAR", *Proceedings of the Tenth International Symposium on Land Subsidence*, Delft, The Netherlands, 2023.
4. **P. Conroy**, S.A.N. van Diepen, F.J. van Leijen, R.F. Hanssen, "Bridging InSAR Coherence Losses Using Contextual Data Driven Processing", *Proceedings of the 2023 IEEE International Geoscience and Remote Sensing Symposium*, Pasadena, United States, 2023.
5. **P. Conroy**, Y.A. Lumban-Gaol, S.A.N. van Diepen, F.J. van Leijen, R.F. Hanssen, "First Wide-Area Dutch Peatland Subsidence Estimates Based on InSAR", *Proceedings of the 2024 IEEE International Geoscience and Remote Sensing Symposium*, Athens, Greece, 2024.