

Department of Construction Sciences
Solid Mechanics

ISRN LUTFD2/TFHF-26/5272-SE(1-72)

Analysing Strain and Pore Evolution in Cement Mortar Using X-ray Tomography

Master's Dissertation by
Tanja Peric

Supervisors:
Stefanos Athanasopoulos, Division of Solid Mechanics
Thawatchart Chulapakorn, European Spallation Source ERIC
Jonas Engqvist, Division of Solid Mechanics

Examiner:
Stephen Hall, Division of Solid Mechanics

Copyright © 2026 by the Division of Solid Mechanics
and Tanja Peric

Printed by Media-Tryck AB, Lund, Sweden

For information, address:

Division of Solid Mechanics, Lund University, Box 118, SE-221 00 Lund, Sweden

Webpage: www.solid.lth.se

Acknowledgements

From the European Spallation Source (ESS), I would like to thank my supervisors Stefanos Athanasopoulos and Thawatchart (Toon) Chulapakorn for their support throughout the project, from invaluable scientific guidance to the smaller practical things that made the work possible. I would also like to thank Karin Rathsmann for responding to my initial email of interest, inviting me to ESS and connecting me with the Science group. Thank you Karin for your guidance, encouragement and for being a supportive shoulder when needed.

I would like to extend a special thank you to Johan Hektor at Malmö University (MAU), the brilliant mind behind **evoSegment**. Thank you for joining all my weekly presentations, asking thoughtful questions, providing valuable feedback and answering emails at lightning speed. I would also like to thank Katja Frid for joining meetings, sharing insight about the material and offering encouragement throughout the project.

From the Division of Solid Mechanics, I would like to thank Jonas Engqvist for his time, guidance and help during the writing of this thesis. I would also like to thank my examiner, Stephen Hall, for his support and valuable input.

Lastly, I would like to thank my family, friends, teachers, colleagues and bosses, who have contributed with knowledge, support, lessons and perspective along the way. You have all left marks, some gentle, and some not, but all part of the person I have become.

Lund, June 2026
Tanja Peric

Abstract

Freeze-thaw cycling is a major durability issue in cement-based materials, as water within its pores can freeze, expand and generate internal stresses that contribute to cracking and material degradation. Understanding how pores, water distribution and deformation evolve during repeated temperature cycles is therefore important for improving the durability of these materials. This thesis investigates freeze-thaw behaviour in a cement mortar specimen using X-ray tomography data acquired during repeated freezing and thawing. Although the experiment included both X-ray and neutron tomography, the scope of this work focuses only on the X-ray data.

A Python-based workflow was developed to analyse the tomographic datasets. The workflow involves image preprocessing, global rigid registration, Digital Volume Correlation (DVC), strain analysis and joint histogram-based pore segmentation. The DVC and strain analysis were performed using SPAM, while segmentation was performed using `evoSegment`. The pipelines were designed to be reusable for future analysis at the European Spallation Source (ESS).

The initial global registration study showed that using a representative image range, excluding regions affected by artefacts, improved both alignment quality and computational efficiency. The DVC parameter study identified a half-window size of 20 pixels, node spacing of 20 pixels and search range of 20 pixels as a good compromise between convergence, resolution and computational cost. The strain analysis showed that DVC can identify regions of deformation and cracking, especially near the specimen perimeter. However, the absolute strain magnitudes should be interpreted with caution due to edge effects caused by the strain algorithm, cracking, fragment movement and image noise.

The joint histogram-based pore segmentation showed differences between inner, intermediate and outer pores. Outer pores were generally more affected by freeze-thaw cycling, showing water loss and nearby crack formation. In contrast, inner and intermediate pores were mainly air-filled and showed no visible cracking, with water appearing to redistribute internally rather than leaving the pore.

Overall, the workflow provides a useful basis for analysing deformation, cracking and pore evolution during freeze-thaw cycling. Future work should focus on validating the results with neutron data, improving boundary and registration handling, and analysing a larger number of pores and specimens to quantify water redistribution and damage development more systematically.

Contents

1	Introduction	1
1.1	ILL NeXT	1
1.1.1	Dual-Modality Imaging: X-ray and Neutron Tomography	2
1.1.2	Tomography	2
1.2	Cement Mortar	2
2	Theory	4
2.1	Motion and Deformation Analysis	4
2.1.1	Continuum Description of Motion	4
2.1.2	From Displacement Fields to Deformation Gradients	5
2.1.3	Digital Image and Volume Correlation	6
2.1.4	Strain Analysis	7
2.1.5	Artefacts	8
2.2	The Joint Histogram Approach	9
3	Method	10
3.1	Dataset Description	10
3.2	The Tools	12
3.2.1	SPAM	12
3.2.2	evoSegment	18
4	Results	21
4.1	Parameter and Range Selection	21
4.1.1	Evaluation of Registration Ranges	21
4.1.2	Evaluation of DVC Parameters	23
4.1.3	Strain Value Ranges and Non-Converged Points	27
4.2	Strain Analysis	29
4.2.1	Global, Half-Specimen and Quarter-Specimen Analysis	29
4.2.2	Local Strain and Crack Development Analysis	31
4.3	Pore Segmentation	34
4.3.1	Analysis of Outer Pores	35
4.3.2	Comparison of Inner, Outer and Intermediate Pores	36
4.4	Pipelines for ESS	37

5	Discussion	38
5.1	Motion and Deformation Analysis	38
5.1.1	DVC Parameter Analysis	38
5.1.2	Strain Analysis	39
5.2	Pore Segmentation Analysis	40
6	Conclusions and Future Work	43
6.1	Conclusions	43
6.2	Future Work	44
	Bibliography	45
A	Overview of the Experimental Setup and Experiment at ILL NeXT	49
B	Example of All Three Strain Algorithm Outputs	52
C	Interpretation of Pore Segmentations	53
D	Registration Result Tables	55
E	Strain Analysis Tables and Figures	57
E.1	Changes in the Respective Ranges 1	57
E.2	Changes Examples 2	59
E.3	Changes in the Respective Ranges 2	60
F	Outer Pore Segmentations	62
G	Inner Pore Segmentations	66

Chapter 1

Introduction

Concrete forms the backbone of our built world, serving as the fundamental material upon which modern infrastructure depends. This dependency comes with an environmental cost, as the production of one of concrete’s most common binders, Portland cement, accounts for approximately 7 % of global carbon dioxide emissions. This creates a strong incentive to not only use concrete efficiently but also to extend its lifespan [1, 2].

Concrete is a multiphase composite of water, cement and aggregates, valued for its durability, yet in cold climates its lifetime is often shortened by persistent and recurring freeze-thaw cycles. Water trapped within the material’s pores freezes and expands by up to 9 % in volume, generating internal stresses that crack and weaken the material over time. Understanding how these processes propagate and how the material behaves during temperature cycling is crucial for understanding and improving concrete’s durability, resilience and longevity [2–5].

The thesis is part of a larger project investigating freeze-thaw processes in cement mortar using combined X-ray and neutron tomography, although this work is limited to the X-ray tomography data. The objective of this thesis is to analyse freeze-thaw processes in cement mortar by processing X-ray data from controlled temperature cycles, with a focus on changes in pore structure, water distribution and internal deformations. To achieve this, a Python-based workflow to perform registration, pore segmentation and strain analysis is developed, providing a useful and adaptable framework for evaluating future datasets at the European Spallation Source (ESS).

This chapter serves as an introduction for the study by introducing the NeXT instrument, principles of dual-modality X-ray and neutron tomography, and the material properties of the cement mortar system to be examined.

1.1 ILL NeXT

The experimental datasets used in this study were acquired from the Neutron and X-ray Tomography (NeXT) instrument at the Institut Laue-Langevin (ILL) in Grenoble, France. NeXT is a dual-modality imaging instrument that enables combined neutron and X-ray

computed tomography (CT). The combination allows the solid structure, air and water or ice within the material to be observed simultaneously, making the instrument well suited for studying freeze-thaw behaviour [6].

1.1.1 Dual-Modality Imaging: X-ray and Neutron Tomography

Dual-modality imaging combines two complementary imaging techniques that highlight different material components based on how they interact with matter. In this study, the provided dataset combines X-ray and neutron computed tomography, where X-ray CT produces three-dimensional tomographies with contrast based on material density. Dense components such as mineral grains and the cement matrix absorb X-rays strongly, producing clear images of pores, cracks and solid phases. In contrast, materials with low density, such as water, absorb fewer X-rays and are therefore poorly resolved. Neutron CT instead produces three-dimensional tomographies with contrast based on the interaction of neutrons with light atomic nuclei, such as hydrogen, making water and air visible. The information obtained from neutron imaging depends on neutron energies, where lower-energy neutrons (cold neutrons) are commonly used for imaging due to their high sensitivity to hydrogen-containing materials such as water. While higher energy neutrons (thermal neutrons) are of shorter wavelength and therefore useful for investigating atomic-scale structures and interactions. The provided datasets from the ILL NeXT instrument were acquired using a white neutron beam, covering a broad energy range and providing access to both types of information simultaneously. Together, X-ray and neutron CT provide complementary three-dimensional information about the specimen, where X-ray CT maps the solid matrix and pore structure, while neutron CT is expected to provide information on water movement and phase changes [7–9].

1.1.2 Tomography

In both X-ray and neutron imaging, tomography refers to the reconstruction of the internal structure of a specimen from a series of two-dimensional projections acquired at different angles. Each projection represents a line integral of the material’s attenuation along the beam path. The projections can be arranged into sinograms, which represent the measured signal as a function of rotation angle and detector position. The sinograms provide an intermediate representation of the data and form the basis for reconstruction. Reconstruction algorithms, such as filtered back projection (FBP), combine these projections into cross-sectional slices, which can then be combined into a complete three-dimensional volume. A schematic of the process is presented in Figure 1.1 [10, 11].

1.2 Cement Mortar

Cement is a mass produced binding material used to hold construction materials together and forms the basis of both concrete and mortar. When mixed with water, it hydrates and hardens into a porous cement paste that binds to surrounding particles. Mortar consists of cement, water and sand, while concrete additionally contains coarse aggregates. Air-entraining agents (AEAs) are often added to intentionally introduce small air voids within




Figure 1.1: Schematic overview of an X-ray tomographic reconstruction workflow. 1. Projection images are acquired at multiple rotation angles as the beam passes through the specimen. 2. For one cross-sectional slice, the same detector row is extracted from each projection image (red line). 3. These rows are stacked as a function of projection angle to form a sinogram. 4. The same procedure is repeated for all detector rows. 5. Each sinogram is reconstructed, here using filtered back-projection (FBP), to obtain cross-sectional slices of the specimen [10, 11].

the cement paste in order to improve the resistance to freeze-thaw damage by providing space for the water to expand upon freezing, thus reducing internal stresses. The hardened cement paste contains a complex pore system, through which water, air and dissolved ions move, spanning several orders of magnitude in size from nanometre-scale gel pores to larger capillary pores to micrometre-scale air voids. Pore size, distribution, connectivity and porosity affect strength and durability. The freezing point of water is also dependent on pore size, where smaller pores freeze at lower temperatures, a relationship governed by the Gibbs-Thomson effect [3, 12–14].

Water and ice movement within the pore system, and the associated pressure build-up, have been explained through different hypotheses. Rather than one mechanism fully explaining the build-up and subsequent damage, the process is likely controlled by a combination of interacting mechanisms. The hydraulic model describes how unfrozen water is forced through the capillary and gel pore network as ice forms and expands. The osmotic model relates the movement of water to concentration differences which develop when ions are left as pure water freezes to ice. The crystallisation pressure theory relates damage to the overpressure exerted by growing ice crystals on the surrounding pore walls. The micro-ice lens model describes how the temperature difference between frozen and unfrozen pores generates a negative pressure in the gel water during cooling, acting like a piston that drives unfrozen water out of the gel matrix and toward the ice-filled larger pore. These mechanisms show that freeze-thaw damage is controlled not only by the presence of water, but also by pore structure, degree of water saturation, ion concentration and the duration of water exposure. The longer the material remains in contact with water, the higher the degree of saturation becomes, progressively reducing the capacity of the air voids accommodating ice expansion [4, 15–18].

Chapter 2

Theory

This chapter introduces the theoretical basis for the employed image registration and analysis techniques, including Digital Image Correlation (DIC), Digital Volume Correlation (DVC), strain analysis and segmentation based on a joint histogram approach.

2.1 Motion and Deformation Analysis

Image registration is the procedure of aligning two or more images so that a given point in a specimen appears at the same location in each image. Registration is typically performed in two main situations. First, images acquired at different time points within the same modality can be aligned to quantify how the specimen has moved and/or deformed over time. Second, images from different modalities acquired at similar time points can be aligned to combine complementary information. In both cases, registration involves determining a transformation that maps a reference image to a deformed or target image. From this transformation, displacement and deformation can be extracted [19–21].

In this study, the transformation is estimated using correlation-based methods, specifically DVC. To provide the general framework for understanding and interpreting the output, a continuum description of motion is introduced [21].

2.1.1 Continuum Description of Motion

The continuum description assumes that the specimen of interest can be represented as a continuously distributed body even though it is composed of discrete microstructures. The advantage of this assumption is that motion and deformation can be expressed as smoothly varying fields over the whole specimen, rather than by having to track individual particles*. The continuum assumption allows for the use of differential calculus, enabling quantities such as strain and stress to be defined through derivatives. Different regions of the specimen may therefore deform differently, and the full deformation is reconstructed into a continuous transformation field [20–22].

* For computational purposes, however, the specimen will be sampled at a finite number of points or regions, where each region is allowed to undergo its own local motion.

2.1.2 From Displacement Fields to Deformation Gradients

Motion describes the change in position of a point between two configurations. The overall motion between a reference configuration, described by coordinate vector $\mathbf{X}$, and a deformed configuration, described by coordinate vector $\mathbf{x}$, can be written as

$$\mathbf{x} = \Phi(\mathbf{X}),$$

where Φ represents the transformation from the reference configuration to the deformed configuration. The displacement field, $\mathbf{u}(\mathbf{X})$, is defined as the difference between the deformed and reference positions,

$$\mathbf{u}(\mathbf{X}) = \mathbf{x} - \mathbf{X} = \Phi(\mathbf{X}) - \mathbf{X}.$$

The displacement field provides a measure of how much, and in which direction, each point has moved. In three-dimensional DVC, this displacement has one component in each coordinate direction. Using the coordinate order z, y, x , the displacement vector can be written as

$$\mathbf{u}(\mathbf{X}) = \begin{bmatrix} u_z(\mathbf{X}) \\ u_y(\mathbf{X}) \\ u_x(\mathbf{X}) \end{bmatrix} = \begin{bmatrix} z' - z \\ y' - y \\ x' - x \end{bmatrix},$$

where (z, y, x) defines the position of a point in the reference configuration, while (z', y', x') defines the position of the same point in the deformed configuration [20].

The motion described by the transformation Φ may include rigid body motion, shape-changing deformation, or a combination of both. Rigid body motion represents the case where the specimen undergoes translation, denoted $\mathbf{t}$, and/or rotation, denoted $\mathbf{R}$, without any change in shape, meaning distances and/or angles between points remain unchanged. In contrast, deformation occurs when distances and/or angles between points change, indicating a local internal mechanical response such as stretching, compression, or shearing [23], [24]. To describe how the motion changes locally within the specimen, the deformation gradient tensor, $\mathbf{F}$, is introduced as

$$\mathbf{F} = \frac{\partial \Phi(\mathbf{X})}{\partial \mathbf{X}}.$$

Since $\mathbf{x} = \Phi(\mathbf{X})$ and $\mathbf{u}(\mathbf{X}) = \mathbf{x} - \mathbf{X}$, the deformation gradient can also be expressed in terms of a displacement field,

$$\mathbf{F} = \mathbf{I} + \frac{\partial \mathbf{u}(\mathbf{X})}{\partial \mathbf{X}}, \quad (2.1)$$

where $\mathbf{I}$ is the identity tensor and $\frac{\partial \mathbf{u}(\mathbf{X})}{\partial \mathbf{X}}$ is the displacement gradient. For a three-dimensional DVC analysis, $\mathbf{F}$ is a 3×3 tensor,

$$\mathbf{F} = \begin{bmatrix} F_{zz} & F_{zy} & F_{zx} \\ F_{yz} & F_{yy} & F_{yx} \\ F_{xz} & F_{xy} & F_{xx} \end{bmatrix} = \begin{bmatrix} 1 + \frac{\partial u_z}{\partial z} & \frac{\partial u_z}{\partial y} & \frac{\partial u_z}{\partial x} \\ \frac{\partial u_y}{\partial z} & 1 + \frac{\partial u_y}{\partial y} & \frac{\partial u_y}{\partial x} \\ \frac{\partial u_x}{\partial z} & \frac{\partial u_x}{\partial y} & 1 + \frac{\partial u_x}{\partial x} \end{bmatrix},$$

where the diagonal terms describe changes along the same coordinate directions, like stretching and compression, while the off-diagonal terms describe coupling between different directions, such as shearing and rotational effects [20, 21].

In DVC, the local transformation may be written in matrix form using homogeneous coordinates. This form combines the deformation gradient $\mathbf{F}$ and the translation vector $\mathbf{t}$ in one transformation matrix,

$$\Phi = \begin{bmatrix} F_{zz} & F_{zy} & F_{zx} & t_z \\ F_{yz} & F_{yy} & F_{yx} & t_y \\ F_{xz} & F_{xy} & F_{xx} & t_x \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

Using this 4×4 matrix representation allows translation, rotation, stretching and shearing to be represented within a single matrix [20].

2.1.3 Digital Image and Volume Correlation

Motion and deformation are within the continuum framework described as continuous displacement fields, from which quantities such as strain can be derived. However, in studies such as this thesis, motion cannot be observed directly as a continuous field. Instead, it must be estimated from the image data, which are inherently discrete and consist of pixels or voxels with associated greyscale values. Therefore, image correlation methods are required to translate changes in grey-value patterns between different states into displacement fields, which can then be used to evaluate deformation and strain [20, 21, 25].

2.1.3.1 Digital Image Correlation (DIC)

DIC is an image correlation method used to estimate displacements by tracking changes between a reference image and a deformed image. This is done by extracting a small image subset from the reference image and shifting it over a search region in the deformed image to find the displacement that gives the best grey-value match. Repeating this process over many subsets provides a set of discrete displacement measurements, which can then be interpreted within the continuum framework to analyse deformation and strain. Conventional DIC is applied to surface images, either in two dimensions or in three dimensions using multiple cameras to reconstruct surface motion [21, 26].

2.1.3.2 Digital Volume Correlation (DVC)

Extending the DIC principles to volumetric image data is referred to as DVC. DVC enables the estimation of local transformations between three-dimensional datasets, thereby quantifying how different regions of the specimen move between two configurations. DVC can be implemented in several different ways, however, all approaches share a common order of steps [21, 26].

Local correlation approach DVC acts locally by dividing the reference image into a grid of points, referred to as nodes. Around each node, a small volumetric region, known as a subset or correlation window, is defined. Within each subset, the deformation is assumed to be homogeneous, meaning that a single transformation describes the motion of all voxels in that region. For each subset in the reference image, a corresponding region is sought in the deformed image within a predefined search window. This search window defines the region over which possible displacements are evaluated and is selected based on the expected magnitude of motion [21, 26].

Similarity measure and displacement estimation For each candidate displacement within the search window, a similarity measure is computed based on the greyscale values. Typically, this is done using a cross-correlation coefficient, which quantifies how well the greyscale value patterns match. The displacement that maximises this similarity measure is selected as the discrete estimate of motion [21, 26, 27].

Sub-voxel refinement Since real displacements are not restricted to integer voxel values, a sub-voxel refinement step is required. Otherwise, the displacement field may appear stepped, with large strains corresponding to the steps and missing strains between. This is typically achieved through interpolation and iterative optimisation which allows the transformation parameters to be determined with higher precision [21, 26].

From displacement to deformation measures Finally, the displacement field obtained from all nodes can then be used to compute local deformations [21].

In these steps, a few key parameters must be carefully considered. The subset size governs the trade-off between resilience and resolution, where a larger subset improves correlation stability by averaging over more voxels, while smaller subsets increase resolution but are more sensitive to noise and image artefacts. The search range governs the maximum displacement allowed during correlation, where larger ranges allow detection of greater motion at the expense of increased computational cost and risk of incorrect matches, while smaller ranges improve efficiency but may miss larger displacements. Node spacing governs sampling of the displacement field, where smaller overlapping spacings increase sampling density and robustness, but might introduce redundancy and higher computational cost. Finally, a choice has to be made regarding the similarity measure, where normalised measures are often preferred as they account for variations in lighting and image contrast, reducing the risk of incorrect matches [21, 28].

2.1.4 Strain Analysis

Within the continuum framework, deformation is described through the displacement field, which represents the motion of material points within the specimen. Strain can be defined in different ways depending on the assumptions made about the deformation. In this thesis, strain is interpreted using a finite-deformation description, where deformation is described through the deformation gradient tensor, as defined in Equation 2.1. This tensor

relates a small region in the reference configuration to its deformed state. This description can account for local stretching, compression, shear and rotation without assuming that these changes are infinitesimally small. The small-strain tensor can be obtained as a simplification by neglecting higher order terms in the displacement gradient, which is only appropriate when deformations and rotations are small. In this work, strain is mainly evaluated in terms of volumetric and deviatoric strain. Volumetric strain describes local volume changes, where positive values indicate expansion and negative values indicate contraction, which makes it the most relevant measure for this analysis. Deviatoric strain, in contrast, describes shape change or distortion without considering volume change. Deviatoric strain, in contrast, describes shape change or distortion without considering volume change. In this work, deviatoric strain is evaluated using an equivalent (von Mises-type) scalar measure, which is always non-negative, where a value of zero indicates no distortion and larger values indicate increasing shape change [21, 26].

2.1.5 Artefacts

Artefacts are a common issue in X-ray images and affect image quality and the overall analysis. Most artefacts relevant for this study are associated with high-density regions, such as metallic components. Metal artefacts degrade the X-ray image quality and typically appear as bright regions and streaking patterns that reduce structural details and distort the reconstructed image. Their severity depends on size, shape and composition of the metal, with high-density, high atomic number materials producing more pronounced effects [29].




Figure 2.1: Metal-induced artefacts in an X-ray CT slice from the provided dataset. The embedded thermocouple (1.red) generates streak artefacts (3.purple) due to photon starvation, resulting in dark regions of unreliable data. The streak artefacts can be seen throughout the slice. A smaller inclusion (2.green) appears bright but does not significantly affect image quality.

In Figure 2.1, the source of an artefact is circled in the red, that being a metal thermocouple embedded within the specimen, which strongly attenuates the X-ray beam and leads to local saturation of the signal. It gives rise to streak artefacts highlighted in purple, primarily caused by photon starvation. Photon starvation occurs because only a small number of photons passing through the thermocouple reach the detector, resulting in increased noise and dark streaks corresponding to missing or unreliable data. Regions 1 and 3 are thus unstable for analysis. In contrast, region 2, although bright, also indicates a high

attenuation feature, but it does not generate the same streak artefacts, and thus appears suitable for analysis [29].

2.2 The Joint Histogram Approach

The segmentation approach used in this study is based on a joint histogram method that enables the analysis of material evolution by comparing voxel intensities between different states. Instead of segmenting individual images independently, the method considers how each voxel changes over time [30].

A joint histogram is constructed by plotting the greyscale values of each voxel in the reference image against its corresponding greyscale value in an evolved image. Generally, the horizontal axis represents the reference intensity, while the vertical axis represents the evolved intensity. Each voxel is therefore mapped to a single point in this two-dimensional space. If a voxel does not change between the two states, the corresponding point lies along the diagonal of the joint histogram. Deviations from the diagonal indicate changes in intensity, where points above the diagonal correspond to an increase in greyscale value, while points below correspond to a decrease. These intensity changes can be associated with physical transitions in the material, such as fluid redistribution or phase changes. Segmentation is then performed by identifying regions within the joint histogram that correspond to different types of intensity changes. Once identified, the classification can be mapped back to the image domain [30].

In contrast to conventional segmentation methods such as thresholding or clustering in image space, the joint histogram approach preserves the temporal history of each voxel. As a result, voxels that exhibit different evolution pathways can be distinguished even if they have identical greyscale values in the final state [30].

The method is modality-agnostic and can be applied to both X-ray and neutron datasets. However, its accuracy depends on proper image registration, as misalignment between datasets may introduce artificial intensity changes that do not reflect real material change [30].

Chapter 3

Method

This chapter describes the datasets, preprocessing steps and workflows used to analyse the freeze-thaw deformation and pore evolution, including DVC and strain analysis in SPAM and joint histogram-based segmentation using *evoSegment* [30, 31].

3.1 Dataset Description

The study is based on an experiment involving repeated freeze-thaw cycles of a cylindrical cement mortar specimen. During the experiment, neutron and X-ray tomography scans were acquired at selected stages of the cycling process. The experiment began with an initial tomography scan of the specimen, followed by repeated freeze-thaw cycling, where in total, six acquisition stages spanning 47 cycles were recorded, as summarised in Table 3.1. A ramp refers to a single temperature event, either freezing or thawing, while a cycle consists of one complete freeze-thaw sequence. Tomographic acquisitions are grouped into sets, each corresponding to one cycle and including both neutron and X-ray measurements during freezing and thawing. For more information regarding the experiment providing the datasets see Appendix A.




Figure 3.1: Example of dataset preprocessing. a) The initial reference dataset (slice 1054/1150), b) the last thawed dataset (slice 1054/1150) showing differences in structural regions (1), thermocouple-induced artefacts (2), and background variations (3). c) After preprocessing (slice 486/977), the dataset is cropped to a circular region of interest and the background is removed, improving consistency for further analysis.

Table 3.1: Overview of the thirteen X-ray tomography datasets used in this study.

Dataset	Acquisition stage	Cycle number	State
0	Set 0	0	Reference
1	Set 1	1	Freezing
2	Set 1	1	Thawing
3	Set 2	12	Freezing
4	Set 2	12	Thawing
5	Set 3	22	Freezing
6	Set 3	22	Thawing
7	Set 4	32	Freezing
8	Set 4	32	Thawing
9	Set 5	42	Freezing
10	Set 5	42	Thawing
11	Set 6	47	Freezing
12	Set 6	47	Thawing

Table 3.2: Overview of dataset properties before and after preprocessing. Cropping and filtering steps were applied to remove regions with low structural information and imaging artefacts, improving the reliability of subsequent analysis.

Property	Initial dataset	Preprocessed dataset
Number of slices	1150	980 (after cropping)
In-plane image size	1130×1118 px	1130×1118 px (unchanged)
Bit depth	8-bit / 16-bit	16-bit used
Bottom slices removed	–	45 slices (low information)
Top slices removed	–	125 slices (artefact region)
Cross-section	Full field of view	Circular crop (background removed)
Background	Present	Set to 0 (black)
Thermocouple region	Included	Excluded (only slices 0 to 510 considered)
Dataset 12	Included	Excluded from analysis, used only in Section 3.2.2

The dataset initially consisted of 1150 slices with an in-plane size of 1130×1118 pixels, stored in both 8-bit and 16-bit formats. However, the 8-bit data had been truncated during reconstruction to use only 7 of its bits. The difference between using 16-bit and 8-bit data was later evaluated. To improve data quality and computational efficiency some preprocessing steps were applied, as summarised in Table 3.2. These included the removal of low-information slices at the beginning of the scans, as well as removal of the final slices due to inconsistency between datasets, where some contained a saturated component while others did not. As shown in Figure 3.1, differences between datasets can be observed in several regions. The central region (1) exhibits significant variation between scans, while the thermocouple (2) introduces beam hardening and photon starvation artefacts that differ between datasets and lead to unreliable grey values. Regions including the thermocouple were later excluded, as they would be found to affect the correlation and thus also the strain analysis. In addition, the background (3) shows noticeable variations in greyscale intensity, despite containing no meaningful information. To mitigate these effects the datasets were cropped to a circular region of interest to remove the background, which was set to pure black, ensuring a clear separation between background and specimen.

3.2 The Tools

Two main software tools were used in this study: **SPAM** and **evoSegment**. These tools were used for image registration, displacement and strain analysis as well as segmentation.

3.2.1 SPAM

SPAM is a Python-based library designed for applications such as DIC, DVC and strain analysis. In this work it was used for two purposes:

1. implementing DVC to quantify local deformations and compute strains, and
2. estimating initial global rigid transformations between datasets.

The global transformations calculated with **SPAM** were later used to align datasets, which is a requirement for using **evoSegment**, as explained in section 2.2.

3.2.1.1 Implementing DVC using SPAM

The **SPAM**-based workflow used in this thesis is summarised and visualised in Figure 3.2 as follows:

1. Perform a global registration between the reference and evolved volumes.
2. Define a regular grid of correlation nodes in the reference image and extract a local subset around each node.
3. Perform a discrete pixel search around the globally predicted positions.
4. Refine the results by iterative local affine registration to obtain sub-voxel displacement and deformation.
5. Filtering to remove non-converged points.
6. Use the resulting displacement or transformation field to compute deformation gradients and strain.

1. Multiscale Global Registration and Binning To ensure computational efficiency, global registration was performed using a multiscale approach, via the **SPAM** function `regis-terMultiscale()`. This involved temporarily downsampling the images through binning, where neighbouring voxels are averaged to create a lower resolution representation. Registration was first carried out on coarsely binned data to estimate large-scale motion and then progressively refined at higher resolutions, up to full resolution images. This initial step aligned the deformed image with the reference image, capturing bulk motion such as translation and rotation before solving the local correlation problem. The global transformation is represented by the operator Φ and defined as

$$\Phi = \begin{bmatrix} \mathbf{F} & \mathbf{t} \\ \mathbf{0}^T & 1 \end{bmatrix},$$




Figure 3.2: Schematic of the SPAM-based DVC workflow. A global registration is followed by local subset correlation on a regular grid, discrete search, sub-voxel affine refinement, and finally filtering and computation of deformation gradients and strain from the resulting displacement field.

where $\mathbf{t}$ is the translation vector and $\mathbf{F}$ is, in this case, the linear component describing the rotation and deformation. Here, linear means that the transformation does not depend on position and acts the same everywhere. The resulting transformation was used as an initial guess for the local search, ensuring that correlation starts from a well-aligned configuration.

2. Definition of the Correlation Grid Following global registration, using the SPAM function `makeGrid()`, a regular grid of correlation nodes were defined within the reference image, where motion is evaluated. The spacing between nodes, controlled by the parameter `nodeSpacing`, determined the sampling density, an example is later shown in Figure 4.5. Around each node, a local subset (or correlation window) was extracted, with its size defined by the parameter half-window size, (`hws`). For a half-window size of

$$(h_z, h_y, h_x),$$

the corresponding subset size is

$$(2h_z + 1) \times (2h_y + 1) \times (2h_x + 1)$$

[31].




Figure 3.3: Example of correlation grid for Dataset 0 (reference) with `hws` = 20 pixels, `searchRange` = 20 pixels and `nodeSpacing` = 20 pixels.

3. Discrete Pixel Search For each node, SPAM performs a discrete pixel search, using the function `pixelSearch()`, comparing a reference subset with candidate subsets in the evolved image over a predefined search region, defined by the parameter `searchRange`. The search was centred around the position predicted by the initial global registration.

For each candidate displacement $\mathbf{p}$, a similarity measure was computed using the normalised-cross-correlation measure (NCC), and the optimal discrete displacement was obtained as

$$\mathbf{p} = \arg \max_{\mathbf{p} \in \mathcal{S}} \text{NCC}(\mathbf{p}),$$

where $\mathcal{S}$ denotes the search region [31].

Prior to the optimisation, a mask was applied to exclude background regions by thresholding voxels with zero grey values, thereby reducing computational cost.

4. Local Iterative Refinement Since true displacements are not restricted to integer voxel values, the discrete pixel-search result is further refined to sub-voxel precision using the SPAM function `ldic()`. The integer-voxel displacements obtained from the pixel search serve here as an initial guess for the local optimisation. The previously described mask is also applied here. This step performed an iterative affine registration for each node, where the local transformation, Φ_i , is expressed as

$$\Phi_i = \begin{bmatrix} \mathbf{F}_i^{(i)} & \mathbf{t}_i^{(i)} \\ \mathbf{0}^T & 1 \end{bmatrix},$$

where $\mathbf{t}_i$ is the local translation of node i and $\mathbf{F}_i$ is the linear component approximating the deformation within the subset. The transformation Φ_i was estimated by minimising the difference between the reference subset, I_r and the transformed subset I_d ,

$$I_r(\mathbf{x}) \approx I_d(\Phi_i \cdot \mathbf{x}).$$

This was solved iteratively using a gradient-based quasi-Newton scheme. At each iteration the transformation was updated by a small increment $\|\Delta\Phi\|$, which represents the adjustment of the current state. The magnitude of this update is quantified by its norm

$\|\Delta\Phi\|$, which provides a measure of how much the solution changes between iterations. Convergence was assumed when the update became smaller than a threshold $\delta\Phi_{min}$

$$\|\Delta\Phi\| < \delta\Phi_{min}.$$

If the updates remained large or the error increased, the solution was considered unstable and may diverge. The process was terminated if the maximum number of iterations was reached.

Grid Point Statuses After correlation each grid point was assigned a status value that describes the outcome and reliability of the local registration. The possible statuses were defined as follows:

- **Status 2: Successful convergence.** The iterative procedure reached the convergence criterion, meaning that the transformation updates became sufficiently small.
- **Status 1: Non-successful convergence.** Maximum number of iterations reached. The solution did not fully converge, but no clear divergence was detected.
- **Status -1: Non-successful convergence.** Divergence due to increasing error. The similarity between the reference and deformed subsets worsened during optimisation.
- **Status -3: Non-successful convergence.** Divergence due to unrealistic deformation. This typically corresponds to excessive volumetric changes, indicating that the local affine model is no longer valid.
- **Status -7: Invalid solution.** The computed transformation contains non-physical values such as NaN or infinity.

Masked regions were excluded from the correlation, and nodes located within these regions were assigned a status of -7 .

DVC assumes locally continuous deformation within each subset. When this assumption breaks, such as near cracks or discontinuities, the correlation may fail or produce unreliable results. The status values were therefore used to identify and exclude such regions from further analysis.

5. Filtering the Displacement Field The displacement field obtained from the correlation may contain unreliable points, identified through their status value. Since strain is computed from spatial derivatives, it is very sensitive to noise, outliers and missing values. Therefore, a filtering step is applied prior to strain computation, excluding background regions and nodes without a successful convergence status (status 2), using the script `SPAM-filterPhiField`.

Invalid points were corrected by interpolation based on neighbouring valid points. For each such point, a set of neighbouring nodes was used, and displacements were estimated as a distance-weighted average. The deformation gradient was not interpolated but instead

replaced by the identity matrix to avoid introducing artificial strain. Although this approach assumes no local deformation at corrected points, it does not guarantee zero strain in the final strain field, as this depends on the strain computation method [31].

Since non-converged points after filtering were essentially treated as missing values, minimising their number is an important criterion when selecting suitable DVC parameters. A large number of non-converged points indicates a greater loss of information, as these points are estimated using information from neighbouring points. Following the recommendation in the SPAM documentation, suitable choices of DVC parameters therefore helped reduce the number of non-converged points, ensuring that as much true information as possible is retained.

6. Strain Computation Methods in SPAM In SPAM, strain can be obtained in three ways: (i) directly from the local transformation during correlation, (ii) from the displacement field using the Q8 method, or (iii) from the displacement field using the Geers method. These approaches differ in how the deformation gradient was estimated, affecting their interpretation, sensitivity to noise and robustness.

Local (Subset-Based) Approach Each correlation window is assigned a local affine transformation, Φ , where an estimate of the deformation gradient is provided. The strain can therefore be computed independently for each subset, providing high resolution. However, this approach is sensitive to noise and local correlation errors, often resulting in noisy strain fields.

Displacement-Based Methods Both the Q8 and Geers methods estimate the deformation gradient from the displacement field, as defined in Equation 2.1. The difference between the two methods lies in how neighbouring information is used to estimate spatial derivatives.

Q8 Approach In the Q8 method, the deformation gradient is computed within cells defined by 8 neighbouring nodes. These nodes form a hexahedral element, within which the displacement field is assumed to vary linearly. Under this assumption, the displacement gradient can be estimated using trilinear interpolation of the nodal displacements, equivalent to an 8-node finite element formulation.

The deformation gradient is then obtained from the displacement gradient as defined in Equation 2.1. Since the gradient is evaluated per cell, the resulting deformation gradient is defined at cell centres rather than at the nodes. Consequently, the deformation field has reduced dimensions: for a displacement field of size (n_z, n_y, n_x) , the Q8 field has size $(n_z - 1, n_y - 1, n_x - 1)$. While this approach improves resilience by averaging information over a finite volume, it also reduces spatial resolution compared with nodal methods.

Geers Approach In the Geers method, the deformation gradient is computed at each node using a set of neighbouring nodes. For a central node, c , neighbouring nodes $k = 1, \dots, N$

are considered. The relative positions in the reference and evolved configurations are defined as

$$\Delta \mathbf{X}_k = \mathbf{X}_k - \mathbf{X}_c, \quad \Delta \mathbf{x}_k = (\mathbf{X}_k + \mathbf{u}_k) - (\mathbf{X}_c + \mathbf{u}_c).$$

Assuming that the deformation is locally affine, the relation

$$\Delta \mathbf{x}_k \approx \mathbf{F} \Delta \mathbf{X}_k$$

is enforced. Since this relation is affected by noise, the deformation gradient is obtained by solving the least-squares problem

$$\sum_i \|\Delta \mathbf{x}_i - \mathbf{F} \Delta \mathbf{X}_i\|^2.$$

This corresponds to minimising the difference between the observed evolved neighbour vector and those predicted by the linear transformation $\mathbf{F}$. By incorporating multiple neighbours the method reduces sensitivity to noise and produces smoother results [32].

For the scope of this thesis, the strain algorithm was limited to Q8 and Geers, since the Local approach produces results that are too noisy. In this thesis, the Geers approach was used because it is the most intuitive and requires the least processing. It is also considered the most resilient option, which makes it suitable for the potentially noisy dataset. The pipelines were, however, designed to handle all three choices. An example image of all three strain outputs can be found in Appendix B.

Volumetric and Deviatoric Strain Measures Once the deformation gradient has been estimated, it can be used to compute strain measures that emphasise different aspects of the local deformation. Volumetric strain describes local volume change and is computed from the determinant of the deformation gradient,

$$\varepsilon_{\text{vol}} = \det(\mathbf{F}) - 1.$$

Since the determinant of the deformation gradient represents the ratio between the deformed and reference volume, taking minus one gives an estimate of the relative volume change [31].

Deviatoric strain describes the shape-changing part of the deformation after the volume changing part has been removed. It is computed from the stretch part of the deformation gradient, where the volumetric contribution from the stretch tensor, $\mathbf{U}$, is removed via polar decomposition, giving

$$\mathbf{U}_{\text{dev}} = \mathbf{U} \det(\mathbf{F})^{-1/3}.$$

The deviatoric strain is then computed as

$$\varepsilon_{\text{dev}} = \|\mathbf{U}_{\text{dev}} - \mathbf{I}\|,$$

where a value of zero indicates no distortional shape change, while larger values indicate increasing distortion [31].

Displacement Arrows The displacement vectors were overlaid on selected image slices to visualise the local motion of the specimen. Each arrow represents the displacement at one node. The direction of the arrows shows the direction in which the material has moved within the slice, while the arrow length indicates the relative size of the displacement. The colour of the arrow represents the total displacement magnitude, calculated from all three displacement components

$$|\mathbf{u}| = \sqrt{u_x^2 + u_y^2 + u_z^2}.$$

3.2.2 evoSegment

evoSegment is a Python-based library used for the segmentation and analysis of evolving structures. In this work, it was used to segment and thus analyse phase changes within the pore structure during the repeated freeze-thaw cycles.

3.2.2.1 Preprocessing: Global Rigid Registrations

Before using **evoSegment**, the datasets must be spatially aligned to ensure voxel-wise correspondence between states. The analysis was performed on X-ray tomography datasets where changes in greyscale values were expected to reflect variations in fluid occupancy within pores. A global rigid registration was performed relative to the reference Dataset 0, using the **SPAM** library, employing the **registerMultiscale()** function with **PhiRigid = True**, restricting the transformation to translations and rotations only.

Performing registration on the full-resolution dataset is computationally expensive and can be affected by reconstruction artefacts. To test efficiency, the transformation was estimated using three different subvolumes. Those included:

- **Representative range:** 505 slices containing clear structural features and minimal artefacts (excluding the thermocouple).
- **Thermocouple range** 505 slices, where approximately one third of the volume is affected by the thermocouple artefacts.
- **Full range:** the entire dataset of 997 slices.

A subset size of 505 slices was chosen as the smallest range that ensured stable convergence, with a margin for robustness. The same size was used for the representative and thermocouple ranges to enable fair comparison. The resulting transformations were compared using the final optimisation error, number of iterations, computational time and resulting translation vector. The final optimisation error describes how well the transformed image matches the reference image, calculated as the normalised summed squared difference between their grey values, where a lower error indicates a better match.

3.2.2.2 Joint Histogram-Based Segmentation

Cropped regions containing pores were extracted and used as input into the following **evoSegment** pipeline.

Preprocessing Prior to histogram construction, all datasets were median filtered to reduce noise and improve cluster stability. In addition, global normalisation was applied, where all volumes were scaled using a shared intensity range (excluding zero-valued background voxels). This ensured consistent intensity mapping across datasets, which is important for later analysis.

Joint Histogram Construction A reference volume was selected, and all remaining volumes were treated as evolved states, concatenated along the vertical axis. This ensured that all greyscale values present in all datasets were captured, and segmentation only had to be performed once. The histogram was then created using the `make2DHistogram()` function. The histogram was visualised on a logarithmic scale to capture both dense and sparse regions.




Figure 3.4: Overview of the joint histogram-based segmentation flow using `evoSegment`. 1. Manual detection of regions of interest (ROIs) in the image domain (water, air, background). 2. Projection of ROI voxel intensities into the joint histogram, relating reference and evolved greyscale values. 3. Identification and segmentation of clusters in histogram space. 4. Assignment of labels to histogram regions. 5. Mapping of labels back to the image domain and post-processing to improve spatial consistency. 6. Final segmentation results across multiple states, illustrating the evolution of pore structures over freeze-thaw cycles.

ROI-Guided Labelling To help navigate cluster identification, regions of interest (ROIs) were manually defined in the image domain (Figure 3.4 1.), targeting relevant features such as water, air and background. The corresponding voxel intensities were projected into a joint histogram (Figure 3.4 2.), linking histogram regions to physical structures. Based on the guidance of the points in the joint histogram, clusters were manually segmented in histogram space by drawing polygons around them (Figure 3.4 3-4.).

Mapping to Image Domain Once labelled in histogram space, the segmentation was mapped back to the image domain, assigning each voxel a label based on its histogram position. This resulted in a representation of current phases and their evolution across all

datasets (Figure 3.4 6.). An example of how the final segmentations were interpreted is provided in Appendix C.

Post-processing To improve spatial consistency, a constrained label expansion was implemented. Unlabelled voxels were assigned to the nearest label within a limited neighbourhood, restricted to non-zero intensity regions, avoiding propagation into the background. This reduced the gaps caused by discretisation while preserving cluster boundaries (Figure 3.4 5.).

For this part of the analysis, two groups of pores were defined: inner pores and outer pores. The distinction between the two groups was based on the radial distance of each pore from the centre of the specimen, with pores located within approximately 250 pixels (≈ 5 mm) from the centre being classified as inner pores, and pores located beyond this distance being classified as outer pores.

Chapter 4

Results

This chapter presents the results from the registration, DVC parameter analysis, strain analysis and pore segmentation. The first part evaluates suitable registration ranges and DVC parameters. The second part analyses displacement and strain fields to identify local deformation and cracking during freeze-thaw cycling. The final part presents the pore segmentation results and summarises the developed pipelines intended for future ESS datasets.

4.1 Parameter and Range Selection

4.1.1 Evaluation of Registration Ranges

The performance of the global rigid registrations was evaluated using three different ranges for estimating the transformation: a thermocouple range, a full dataset range and a representative range excluding slices affected by the thermocouple artefacts. The ranges were compared using optimisation error, number of iterations, computational time and resulting translation and rotation vectors, for both binned and non-binned datasets. In the following tables, green indicates the best-performing result, yellow intermediate performance and red the poorest performance, while grey-highlighted rotation values could not be visually validated.

Table 4.1: Registration results for Dataset 0 (reference) and Dataset 12 (thawed) (binning = 2). Expected axial shift from visual alignment: ~ 15 slices (t_z).

Range	Error	Iter.	Time	Translation $t = [t_z, t_y, t_x]$	Rotation $R = [r_z, r_y, r_x]$
Thermocouple	279.64	64	23m 4s	[10.109, 0.183, 0.402]	[2.893, -0.085, -0.020]
Full	249.15	508	32m 55s	[10.884, -0.050, 0.589]	[2.028, -0.160, -0.095]
Representative	219.40	39	7m 7s	[12.734, 0.025, 0.212]	[1.029, -0.026, -0.024]

Overall, the representative range provided the most robust results, with the lowest optimisation error, fewest iterations and shortest computational times. This was also true for the most challenging case, Dataset 12 (Table 4.1) that was deliberately excluded from analysis, because it differed too much with regard to alignment and how it was affected by the different artefacts. In this case, the representative range reduced the computational

time to approximately 7 minutes, compared to over 23 minutes for the thermocouple range and more than 30 minutes for the full dataset. The resulting axial translation ($t_z = 12.7$) was also closest to the visually estimated shift of approximately 15 slices.

Table 4.2: Registration results for Dataset 0 (reference) and Dataset 8 (frozen) (binning = 2). Expected axial shift from visual alignment: ~ 1 slice (t_z).

Range	Error	Iter.	Time	Translation $t = [t_z, t_y, t_x]$	Rotation $R = [r_z, r_y, r_x]$
Thermocouple	178.66	147	16m 32s	[-0.318, 0.425, 0.322]	[-0.284, -0.051, 0.059]
Full	160.62	149	32m 3s	[-0.388, 1.444, 1.103]	[-0.434, -0.037, 0.013]
Representative	164.02	68	10m 40s	[-0.576, 0.272, 0.139]	[-0.515, -0.028, 0.040]

For datasets with smaller axial shifts, the difference between the representative and full ranges became less pronounced. In some cases, the full range produced a similar or slightly lower optimisation error, but visual inspection indicated that the representative range still gave a more reliable alignment (Table 4.2), which was of highest importance. This suggested that when translations were small, differences in cross-sectional translation and rotation had a stronger influence on the error value, making visual validation more difficult.

Table 4.3: Registration results for Dataset 0 (reference) and Dataset 1 (frozen) (binning = 2). Expected axial shift from visual alignment: ~ 4 slices (t_z).

Range	Error	Iter.	Time	Translation $t = [t_z, t_y, t_x]$	Rotation $R = [r_z, r_y, r_x]$
Thermocouple	130.01	30	6m 3s	[-4.418, 0.122, 0.091]	[0.184, -0.036, -0.050]
Full	121.19	36	12m 56s	[-4.330, 0.214, 0.085]	[0.189, -0.055, -0.074]
Representative	115.85	28	5m 31s	[-4.205, 0.205, 0.011]	[0.194, -0.041, -0.048]

Table 4.4: Registration results for Dataset 0 (reference) and Dataset 4 (thawed) (binning = 2). Expected axial shift from visual alignment: ~ 2 slices (t_z).

Range	Error	Iter.	Time	Translation $t = [t_z, t_y, t_x]$	Rotation $R = [r_z, r_y, r_x]$
Thermocouple	151.71	70	7m 21s	[-1.591, 0.309, 0.218]	[-0.259, -0.007, 0.010]
Full	144.93	26	5m 36s	[-1.507, 0.746, 0.529]	[-0.308, 0.002, -0.009]
Representative	143.28	51	5m 27s	[-1.519, 0.260, 0.141]	[-0.348, 0.001, 0.010]

The thermocouple range performed worst across all evaluated cases (Tables 4.1- 4.4). It consistently required the most iterations and produced the highest errors, indicating that reconstruction artefacts from the thermocouple negatively affected registration. The full range performed better than the thermocouple range because it included a larger dataset with more reliable regions. However, due to its size, the computational times did in some cases become much longer. These results then showed that using a smaller representative range, excluding artefact-affected regions, was important both for alignment quality and computational time, especially for larger translations. When translations were closer to one pixel the difference between the representative range and full range became smaller, but the representative range remained preferable because it avoided unreliable artefact affected regions.

Table 4.5: Registration results for Dataset 0 (reference) and Dataset 12 (thawed) (without binning). Expected axial shift from visual alignment: ~ 30 slices (t_z). The full range registration could not be executed due to memory limitations.

Range	Error	Iter.	Time	Translation $t = [t_z, t_y, t_x]$	Rotation $R = [r_z, r_y, r_x]$
Thermocouple	727.28	453	2h 53m	[20.447, 0.320, 0.720]	[3.283, -0.072, -0.011]
Full	-	-	-	-	-
Representative	429.12	221	59m 38s	[25.349, 0.082, 0.409]	[1.518, -0.019, -0.018]

Table 4.6: Registration results for Dataset 0 (reference) and Dataset 1 (frozen) (without binning). Expected axial shift from visual alignment: ~ 8 slices (t_z). The full range registration could not be executed due to memory limitations.

Range	Error	Iter.	Time	Translation $t = [t_z, t_y, t_x]$	Rotation $R = [r_z, r_y, r_x]$
Thermocouple	332.34	113	36m 48s	[-8.797, 0.236, 0.175]	[0.211, -0.034, -0.046]
Full	-	-	-	-	-
Representative	299.46	100	27m 50s	[-8.398, 0.394, 0.021]	[0.212, -0.037, -0.045]

The non-binned results further supported the choice of using a representative region (Table 4.5 and Table 4.6), where the full range registration could not be performed due to memory limitations. In this case the thermocouple range required much longer computational times and, again, produced very poor alignment results. Based on these findings, the choice of excluding the thermocouple regions was confirmed to be beneficial for the future analysis, as the presence of the thermocouple clearly affected the first step of DVC, and consequently the alignment needed for the joint histogram approach. Excluding the thermocouple region from further analysis ensured that the analysis was not affected by its presence and artefacts.

4.1.2 Evaluation of DVC Parameters

4.1.2.1 Half-window Size

An initial visual assessment of the specimen was used to estimate relevant feature sizes, as shown in Figure 4.1. The pores ranged from approximately 7 to 53 pixels in diameter, while visible cracks extended approximately 50 pixels. Based on these scales, **hws** values of 15 pixels, 20 pixels and 25 pixels were considered reasonable. A smaller value, **hws** = 10 pixels was also tested to verify the expected loss of resilience for too small subsets.




Figure 4.1: a) Dataset 0 (reference), slice close to the top of the specimen. b) Dataset 7 (frozen), same slice as reference. c) Examples of relevant feature sizes in the cement specimen, including pores, cracks, the thermocouple region and dense inclusions, same slice as reference. 1 mm $\approx$ 50 pixels.

Table 4.7: Summary of tested half-window sizes for 16-bit data. All runs used `nodeSpacing` = $2 \cdot \text{hws}$ and `searchRange` = $\pm 2 \cdot \text{hws}$. The convergence rate was defined as the proportion of converged points, status 2, relative to all valid points, excluding invalid masked points, status -7: Convergence rate = $\frac{N_{\text{status}=2}}{N_{\text{total}} - N_{\text{status}=-7}}$. Based on Dataset 0 (reference) and Dataset 7 (frozen).

hws	Nodes	Convergence rate	Observation
10	18144	0.708	High resolution, but many non-converged points
15	5184	0.791	Improved convergence, still relatively noisy
20	2184	0.823	Good balance between convergence and resolution
25	1089	0.846	Best convergence, but too few nodes and strong smoothing




Figure 4.2: Comparison of filtered strain results for different half-window sizes. Because different `hws` values generate grids with different node positions the displayed slices/indices are not identical but selected to be as spatially close as possible, the respective strain indices are circled in black. (a) `hws` = 15 shows the highest amount of noise and several non-converged regions inside the specimen, in addition to points along the perimeter. (b) `hws` = 20 reduces internal noise while preserving local deformation features, particularly in the upper-left and lower-right regions. (c) `hws` = 25 produces the smoothest results, but also suppresses smaller-scale features visible for `hws` = 20. All strain images are shown after filtering.

The results in Table 4.7 showed that increasing `hws` generally improved the convergence rate, as expected. However, this improvement came at the cost of a reduced number of grid points and lower spatial resolution. For `hws` = 10 the number of nodes was high, but a large proportion of points did not converge, making the resulting displacement and strain field unreliable. In contrast, `hws` = 25 gave the highest convergence rate, but produced a sparse displacement field and increased smoothing. When investigating the increased number of non-converged points (most of which had a status of -3) for smaller `hws`, one could note that they were mainly located along the specimen perimeter. Another pattern was that the first and last slices contained the most non-converged points, shown in Figure 4.2. Smaller `hws` also introduced more internal noise, while larger values increasingly smoothed local features. Based on the convergence behaviour, `hws` = 20 was selected as the best compromise between resilience, spatial resolution and the preservation of deformation features.

4.1.2.2 Comparing 8-bit and 16-bit Data

The comparison between the provided 8-bit and 16-bit data was performed as a part of the half-window size evaluation, using the same **hws** values and DVC settings for both bit depths. As shown in Table 4.8, both bit depths showed similar trends across the tested half-window sizes, where increasing **hws** improved convergence rate, while smaller subsets produced more non-converged points. The convergence rates were also comparable between the 8-bit and 16-bit data, where the 8-bit data performed slightly better in 2/3 of the cases, an effect that may be related to the increased detail and noise sensitivity in the 16-bit data. However, these results did indicate that the overall convergence behaviour was mainly controlled by the DVC parameters and image structure, rather than bit depth alone.

Table 4.8: Summary comparison of 8-bit and 16-bit DVC results for selected half-window sizes. Pixel search and LDIC times are reported as computation times for each step. Based on dataset 0 (reference) and dataset 7 (frozen).

hws	Bit depth	Convergence rate	Pixel search time	LDIC time
15	16-bit	0.791	16 min 25 s	2 min 14 s
15	8-bit	0.754	16 min 10 s	2 min 27 s
20	16-bit	0.823	36 min 46 s	2 min 22 s
20	8-bit	0.834	37 min 1 s	2 min 12 s
25	16-bit	0.846	2 h 11 min 42 s	3 min 16 s
25	8-bit	0.850	2 h 3 min 24 s	3 min 35 s




Figure 4.3: Comparison between 16-bit and 8-bit representations of the same regions. Images a) Dataset 10 (thawed) and b) Dataset 12 (thawed) show the 16-bit data, where local greyscale detail and artefacts are clearly visible. Images c) Dataset 10 (thawed) and d) Dataset 12 (thawed), show the corresponding 8-bit data, where these variations are compressed and the artefacts are less distinct. This difference is important for interpretation as reduced artefact visibility in the 8-bit data could lead to apparent strain being incorrectly interpreted as real deformation.

The computational times were also similar. In some cases, the 16-bit data was even slightly faster, suggesting that bit depth did not have a systematic effect on the computational times. However, the 16-bit data preserved a larger dynamic range and more subtle greyscale variations than the 8-bit data (both inherently and due to the fact that the 8-bit data was truncated), which is important for correlation-based methods like DVC. Since DVC relies on local greyscale values to identify corresponding regions between volumes, a loss of intensity in the 8-bit data may reduce the reliability of the visual interpretation and correlation results. An example of this difference is shown in Figure 4.3, where in the 16-bit images local details and artefacts are more clearly visible, whereas the corresponding 8-bit images compress these variations. This is important because reduced visibility of artefacts may lead to incorrect interpretation of apparent strain and deformation.

4.1.2.3 Search Range

The results in Table 4.9 show that the optimal search range depended on the magnitude of the expected motion. For the moderate deformation case (Dataset 5 and 6), the smaller search range performed best, with ± 10 providing the highest convergence rate. However, for the comparison with the largest difference, between Dataset 0 and Dataset 12, ± 20 provided the best results, both with regards to convergence rate and the mean number of non-converged points per slice. This indicates that ± 10 might be sufficient for smaller changes but could become too restrictive when large displacement or structural changes were present. A search range of ± 20 was therefore selected as the more resilient compromise.

Table 4.9: Summary of tested search ranges for `hws` = 20 and `nodeSpacing` = 40 using 16-bit data. Convergence rate was calculated as $\frac{N_{\text{status}=2}}{N_{\text{total}} - N_{\text{status}=-7}}$. The mean number of non-converged points per slice was computed after excluding the first and last dataset, since these were shown in Section 4.1.2.1 to be particularly noisy.

Datasets	Range	Convergence rate	Mean non-conv./slice	Observation
5-6	± 40	0.823	22.18	Large range; more non-converged points
5-6	± 20	0.943	4.55	Good convergence and robust search
5-6	± 10	0.949	4.09	Best for moderate changes
0-12	± 40	0.640	39.30	Large range did not improve convergence
0-12	± 20	0.731	29.45	Best for large changes
0-12	± 10	0.637	44.45	Too restrictive for large changes

4.1.2.4 Node Spacing

As it can be noted in Table 4.10, reducing the node spacing did not improve the convergence rate but did increase the sampling density of the displacement field. For `hws` = 20 and `searchRange` = ± 20 , reducing `nodeSpacing` from 40 to 20 increased the number of nodes from 2184 to 18144 while the convergence rate remained nearly unchanged (decreasing only from 0.943 to 0.941). This resulted in a denser displacement field with a similar proportion

of non-converged points. As expected, the computational times increased substantially. Therefore, `nodeSpacing` = 20 was selected to obtain a denser displacement and strain field while maintaining comparable convergence behaviour. The increased computational time was considered acceptable for the improved spatial detail.

Based on the parameter study, the final DVC parameters were set to `hws` = 20, `nodeSpacing` = 20 and `searchRange` = ± 20 . This combination provided the best balance between convergence, spatial resolution, robustness to larger displacements and computational cost.

Table 4.10: Summary of tested node spacings for 16-bit data with `hws` = 20 and `searchRange` = 20. Overlapping grids use `nodeSpacing` = 20, while non-overlapping grids use `nodeSpacing` = 40. Based on Dataset 0 (reference) and Dataset 7 (frozen).

Node spacing	Nodes	Convergence rate	Mean non-conv./slice	Pixel search time	LDIC time
40	2184	0.943	4.55	9 min 39 s	5 min 49 s
20	18144	0.941	20.05	38 min 2 s	58 min 15s

4.1.3 Strain Value Ranges and Non-Converged Points

Strain values could not be interpreted reliably without filtering, as non-converged points generate unrealistically large values. As shown in Figure 4.4, these extreme values coincide with regions where the DVC correlation failed. Even after filtering large strain values remained. Percentile clipping could therefore be applied to improve visual interpretation of strain fields. As shown in Table 4.11 the full strain range was dominated by extreme values. These extreme values did coincide with the non-converged points, found mainly along the perimeter of the specimen, as well as in the first and last slice of the strains. This indicates that the large strain values could be related to the Geers strain computation algorithm.

Table 4.11: Strain ranges depending on the chosen percentile range.

Percentile range	Volumetric strain range	Deviatoric strain range
0-100	-0.497, 0.361	0.004, 13.712
2-98	-0.062, 0.088	0.006, 0.186
10-90	-0.013, 0.042	0.009, 0.060




Figure 4.4: Overview of non-converged points and strain outliers for the Geers strain calculation. Dataset 0 (reference) versus Dataset 7 (frozen). (a) Number of non-converged points per strain slice index. (b) Absolute maximum volumetric and deviatoric strain values before filtering, showing extreme outliers that dominate the scale. (c) Corresponding absolute maximum values after filtering, where the strongest outliers are reduced but boundary related values still remain.




Figure 4.5: Effect of percentile clipping on the visualisation of volumetric and deviatoric strain for strain slice indices 100-140, for Dataset 0 versus Dataset 7. (a) The full 0-100 % range is dominated by extreme values, making internal strain variations difficult to interpret. (b) The 2-98 % range reduces the influence of outliers while preserving most visible strain features. (c) The 10-90 % range further enhances contrast but saturates more of the strain effects. (d) Non-converged points shown for comparison.

The results from the percentile clipping presented in Figure 4.5 show that the chosen range had a large impact on the interpretation of the strain fields. Without clipping, local strain variations inside the specimen remained hidden. A 2-98 % percentile range provided a more balanced representation, reducing the influence of extreme values by saturating them. However, the optimal percentile range did differ between slices and datasets, and must be assessed for each case.

These results showed some limitations of the strain analysis, where large boundary strains and extreme values could dominate the visualisation thereby masking smaller changes. Dividing the specimen into smaller regions and analysing local strain variations may therefore provide more reliable results.

4.2 Strain Analysis

4.2.1 Global, Half-Specimen and Quarter-Specimen Analysis

By first analysing the two datasets that were compared, one could determine where the expected changes were, and based on that investigate what was actually seen from strain and displacement analysis. Figure 4.6 exemplifies three positions that visually differed between the two datasets: two voids that seemed to become emphasised and even slightly extended (circled red and blue) and a crack that became more apparent (yellow).




Figure 4.6: a) Dataset 6 (slice 360). b) Dataset 8 (slice 360/977). Both have three regions (red, yellow and blue) circled, where visual changes are seen between the evolution from dataset 6 to dataset 8.)




Figure 4.7: Binning = 2. a) Strain index 9 overlaid on slice 180 (Dataset 6), included in the subvolume are slices 180-220. b) Dataset 6, slice 220, c) Dataset 6, slice 200. d) Dataset 6, slice 180. Emphasises the importance of investigating multiple slices within the node subvolume.

In Figure 4.7 the strains were overlaid on top of slice a) 220, and corresponding displacement vectors overlaid on slices b) 220, c) 200, d) 180. This figure emphasises the importance of investigating all three displacement vector slices included in the computation of a single

strain index, otherwise, the strains could be incorrectly interpreted*. In Figure 4.7 a) the displacement vectors responsible for the large strain output in the upper left part of the specimen were presented. Visually, no clear structural change could be identified in this region to explain the large displacement and strain. This emphasised the need for further filtering the displacements even after using SPAM's own filtering script. This could be done using percentile clipping, where displacement arrows over or under a certain percentile were disregarded. Interpreting the strains and displacements from a global perspective we risk




Figure 4.8: Binning = 2. a) Left half of slice index 9, overlaid on slice 220 (Dataset 6) with its corresponding displacement field overlaid as vectors on top of the same slice. b) Quarter left half of slice index 9, overlaid on slice 220 (Dataset 6) with its corresponding displacement field overlaid as vectors on top of the same slice.

only seeing the absolute largest changes, thus, by switching the view to one half at the time, or even one quarter at the time, it becomes easier to interpret the behaviours, as can be seen in Figure 4.8 a) and b). However, the characteristic ways of the Geers approach also removed additional information along the perimeter of the chosen views, which has to be taken into account when choosing view. This became especially apparent for Figure 4.8 b). Without additional clipping, both the a) and b) cases showed displacement vectors indicating cracking in the red region of Figure 4.6. In the blue region of Figure 4.6, larger displacement vectors and strains for a) and the other half of the quarter are shown (presented in Figure 4.9). The yellow region in Figure 4.6 did not seem to change as much as the other two regions, at least when evaluating the displacements and strains.




Figure 4.9: a) No clipping. b) 95 % percentile clipping. c) 90 % percentile clipping.

* In reality, because of how the strains are computed with the Geers approach, neighbouring strain values, and thus, their slice ranges, would also need to be considered to get the full story. But for this analysis, only the slices in the subvolume of the chosen node will be considered.

For the yellow region in Figure 4.6, using percentile clipping, as shown in Figure 4.9, reduced edge effects and isolated local strain features. However, the presence of noise and edge effects varied depending on region and view. No single clipping value could therefore be applied universally, and the appropriate level had to be tested for each case. The clipping helped reduce the extreme values of the strain, but most of the observed effects remained close to the edge, limiting the detail visible in the strain field. The approach was therefore most useful in cases where the strains were separated further from the perimeter of the specimen.

4.2.2 Local Strain and Crack Development Analysis

The strain analysis was performed using the half-specimen view, based on the evaluation in Subsection 4.2.1. The dataset comparisons were divided into frozen-frozen (F-F), frozen-thawed (F-T) and thawed-thawed (T-T) groups. The evaluated regions are shown in Figure 4.10, and the observations between the F-T groups are presented in Figure E.2, while the rest are summarised in Appendix E, Tables E.1 and E.3, with examples of the strain results presented in Figure 4.11.




Figure 4.10: a) Dataset 0 (reference), slice 180 b) Dataset 11 (thawed), slice 180. In both images three regions where we expect to see a lot of changes are indicated. These regions then correspond to the same regions that were marked out in section 4.2.1 Figure 4.6, but showing a representative initial slice and a later slice, further emphasizing how large these differences will be.




Figure 4.11: Generally much lower strains than in the next example. This is a slice that is actually closer to the top, but still, less strains. This shows that we have most changes along the perimeter, this makes it difficult to actually evaluate how well the strains correspond to what we can confirm visually. In 7-9 we see larger changes, but it also affects the whole specimen.

Table 4.12: Summary of strain localisation in the frozen-thawed (F-T) comparisons. The colours refer to the regions marked in Figure E.1. Green indicates clear changes in strain and/or displacement vectors in expected regions, yellow indicates weaker or uncertain changes that may be real change or noise, and red indicates no visible change.

Datasets/State	Region	Observation
1-2 / F-T	Red	Indication of cracking, mixed strain.
	Yellow	Indication of change, negative strain.
3-4 / F-T	Red	Indication of cracking, negative strain.
	Yellow	Indication of change, positive strain.
5-6 / F-T	Red	Indication of change, no strain; edge effect.
	Yellow	Indication of change, positive strain.
7-8 / F-T	Red	Indication of cracking, no strain; edge effect.
	Yellow	Indication of change, positive strain.
9-10 / F-T	Red	No change.
	Yellow	No change.
11-12 / F-T	Red	Inconclusive displacement; large negative strain.
	Yellow	Inconclusive displacement; large positive strain.

The clearest changes were generally observed in comparisons involving Datasets 6, 7 and 8. Some comparisons showed mainly positive strain, some mainly negative, and several regions associated with cracking showed mixed positive and negative strain.

A second set of regions was also evaluated, as shown in Figure E.1, with the results summarised in Tables E.4, 4.12 and E.5. This second set showed a similar pattern, with stronger changes observed in the comparisons involving Datasets 5, 7 and 9. In several T-T comparisons, regions with a lot of cracking mainly showed negative strain, while regions with less cracking showed mostly positive strain. In both evaluations, smaller changes were generally observed in the later comparisons involving datasets 9, 10 and 11. This may be related to the experimental procedure, since fewer freeze-thaw cycles were performed between these acquisitions, resulting in smaller accumulated changes between datasets.

An additional example of local damage during these cycles is shown in Figure 4.12. In the previous analyses, the specimen was mainly examined by varying the slice position in the z-direction, in this example, a fixed position in the y-direction was instead selected, close to the perimeter, $y = 40$. This allowed damage to be viewed along the height of the specimen, where cracks and void-like regions became more pronounced throughout the freeze-thaw cycling. In later datasets, small material fragments appeared to detach or shift within the cracked region. This indicates that the specimen did not deform smoothly in these regions, but instead developed local discontinuities and possible fragment movement. The corresponding displacement vectors in Figure 4.13 show that a large amount of percentile clipping was required before the expected local changes became easier to identify.




Figure 4.12: Slice at fixed position $y = 40$, viewed along the specimen height across freeze-thaw cycles. The green circle in a) marks an air region present at the start, which appears to move downward during cycling. From Dataset 9 two detaching fragments, circled in red, become visible and shift position as cycling continues, indicating local material chipping.




Figure 4.13: Displacement vectors (Dataset 6 to Dataset 8), overlaid on Dataset 6 (slice $y = 40$). Different levels of percentile clipping: a) 100 % (no clipping) b) 95 %, c) 90 %, d) 85 %, e) 80 %, f) 75 %.

Analysis of Dataset 12 stood out from the other datasets, with very large and inconclusive displacement vectors, as shown in Appendix E Figure E.3. This was likely related to the alignment difficulties observed in the initial registration step, Subsection 3.2.2. However, the strain results did confirm deformation in the regions where changes were expected.

A limitation of these observations was that many of the visually identified changes were located close to the specimen perimeter. This is also the region where the strain calculation is considered least reliable, as discussed in Section 4.1.3. These observations are further analysed in the discussion, Section 5.1.2, with focus on the meaning of positive, negative and mixed strain, and on the reliability of strain values close to the specimen perimeter.

4.3 Pore Segmentation

Figure 4.14 shows three different approaches to pore segmentation, providing different levels of detail in how the grey-value clusters in the joint histograms were interpreted. In Figure 4.14 a), a simple segmentation approach was used. The main clusters along the diagonal were identified, corresponding to background, water and air phases, together with transition regions beyond the diagonal. In this approach, both the darker and brighter air voids were grouped together into one pore phase, providing less detailed information, but the results were clearer and easier to interpret.




Figure 4.14: Example of three pore segmentation approaches. a) Simple segmentation, where “darker” and “brighter” voids are segmented as one, histogram in d). b) Segmentation where only the “darker” voids are segmented, histogram in e). c) Complex segmentation where “darker” and “brighter” voids are segmented separately, histogram in f). Due to the previously noted misalignment issues during the initial registration step, the complex segmentation becomes more difficult to read and interpret.

In Figure 4.14 b), only the darker void regions were classified as air voids, while the brighter air void regions were assigned to belong to the water phase. To determine whether both the




Figure 4.15: Dataset 0 (reference). Cross-section of the pore visualised in Figure 4.14, supporting the interpretation that both the darker and brighter void regions may represent air. Red curve emphasising the meniscus shape of the water inside the pore.

darker and brighter void regions represented air, the pore was examined in cross-section, as shown in Figure 4.15. The meniscus shape separated the brighter and darker void regions from the clearly brighter water phase. The interpretation should however be confirmed using the neutron data. Figure 4.14 c) showed a more detailed segmentation, where the darker and brighter void regions, as well as their transition regions, were segmented separately. However, this segmentation was also more sensitive to the quality of the initial alignment. Slight misalignments, especially in the final thawed dataset and subsequent segmentations, could be incorrectly interpreted as pore movement. As a result, the more detailed segmentation became harder to interpret. Regardless of the choice of segmentation style, these results showed that the joint histogram approach could be used to isolate specific parts of the pore region, making it useful for future analysis. For the scope of this thesis, the main interest was the movement and redistribution of water and air within the pores, rather than the detailed separation between darker and brighter voids regions, therefore, a simpler segmentation approach was used.

4.3.1 Analysis of Outer Pores

Outer pores were analysed using selected examples shown in Appendix F, Figure F.1. In several of these pores, cracks extending from the pore became more pronounced throughout the freeze-thaw cycles. Some pores consisted mainly of air and showed little visible change, while other pores were initially fully saturated and appeared to empty during the cycles. In these cases, the water appeared to move into nearby cracks, leaving a void behind. The number of cycles required for an initially fully saturated pore to become empty varied, and therefore could not be defined as a single fixed value. In this analysis, it became apparent that pores located slightly closer to the centre of the specimen showed behaviours more similar to those of the inner pores. This indicated that the previously assumed distance separating outer and inner pores might not represent the true boundary between pore groups. Instead, either an intermediate group of pores should be considered between the outer and inner regions, or the boundary separating the two groups should be redefined altogether. It is also possible that this boundary was not fixed in time but shifted as cycling continued, with the outer region growing further inwards the more freeze-thaw cycles the specimen underwent. In that case, no single distance could describe the separation, as pores initially behaving like inner pores would gradually take on outer-pore behaviour over the course of the experiment.

4.3.2 Comparison of Inner, Outer and Intermediate Pores

The analysed inner pores mainly contained air, as shown in Appendix G, Figure G.1. When water was present, it generally appeared in smaller amounts and did not seem to escape during the freeze-thaw cycling. No clear cracking was segmented close to these pores. The pore behaviour was strongly dependent on the height within the pore at which the analysis was performed. Figure 4.16 showed the same pore at three different heights. At the top of the pore, a), water was mainly observed during the thawed states, whereas air was mostly present during the frozen states. Further down, closer to the middle of the pore, b), the pore was mainly air-filled and showed very little variation between the freeze-thaw states. At the bottom of the pore, c), the opposite behaviour was observed: water was mainly present during the frozen states, whereas the thawed states mostly contained air. This behaviour became increasingly apparent after multiple cycles.




Figure 4.16: Segmented outer pores at three positions: a) top slice 342/977, b) middle slice 331/977 and c) bottom slice 306/977. The corresponding histograms are shown in d), e) and f), respectively.

A similar analysis was performed on an intermediate pore located closer to the inner region of the specimen, shown in Appendix F, Figure F.3. This pore showed similar behaviour to the inner pores. Water appeared closer to the top of the pore during the thawed states and closer to the bottom during the frozen states, while the middle part remained mainly unchanged and air-filled.

Outer pores located closer to the specimen edge showed a different behaviour, as shown in Appendix F, Figure F.4. In these pores, the same internal redistribution of water was not observed, probably because the pore was fully saturated. After 12 cycles (corresponding to Dataset 3), the pore appeared to crack, as an air void became visible. During the following cycles, the pore seemed to lose most of its water content and was eventually fully emptied.

4.4 Pipelines for ESS

Python-based pipelines were developed to make the workflow reusable for future ESS datasets. The DVC pipeline uses **SPAM** for global registration, local DVC, filtering, displacement analysis and strain computation. These can be applied to individual dataset pairs or complete folders, with the option to analyse frozen-frozen, thawed-thawed or frozen-thawed comparisons. Individual processing steps can also be run separately, allowing intermediate outputs such as the correlation grid and correlation results.

A strain visualisation pipeline was also implemented for displaying the computed strain field and displacement vectors. The strain results can be shown for all three available strain computation methods or for one selected method, and overlaid on the corresponding image slice. Interactive sliders allow the user to move through the affected slices and compare the results in different directions.

Finally, a complete **evoSegment** pipeline was developed for joint histogram-based segmentation. Additional implemented functions included ROI selection, projection of selected regions onto the joint histogram and label filling.

Chapter 5

Discussion

This chapter discusses the results from the deformation and pore segmentation analyses. The DVC parameter choices were first evaluated in relation to convergence, resolution and strain reliability. The strain results were then interpreted qualitatively, with a focus on the deformation localisation, strain sign and crack development. Finally, the pore segmentation results were discussed in relation to pore position, water redistribution, cracking and segmentation limitations.

5.1 Motion and Deformation Analysis

5.1.1 DVC Parameter Analysis

The DVC parameter study showed that no single parameter choice optimised all aspects of the analysis. Instead, the final parameters had to balance resolution, convergence behaviour and computational cost. The selected combination of `hws` = 20 pixels, `nodeSpacing` = 20 pixels and `searchRange` = 20 pixels was therefore not chosen because it provided the highest convergence rate in every case, but because it provided the most useful compromise for this study. The use of overlapping nodes increased the spatial density of the displacement field, which improved the visualisation of deformation around pores and cracks. However, when the node spacing was smaller than the subset size, neighbouring subsets overlapped and therefore produced measurements that were not fully independent. The denser field should therefore be interpreted as improved spatial sampling rather than completely independent measurements at every node.

An important outcome of the parameter evaluation was that convergence should not be interpreted only from the convergence rate. A larger number of non-converged points did not necessarily mean that the selected parameters were unsuitable. Across the analysed datasets, many of the non-converged points were located near the specimen boundary, where the correlation window may have included both specimen and background, or contained too little reliable image texture. Non-converged points may also appear in regions where the assumptions of DVC were no longer valid. DVC assumes that the same local pattern can be tracked between the reference and evolved datasets. When a crack forms, opens, or when fragments move independently, the local pattern may change too much for

the algorithm to describe it as a continuous deformation. A pre-existing crack forms part of the trackable pattern and can often still be correlated, whereas a newly formed crack has no counterpart in the reference image and therefore tends to produce non-converged points. In such regions, non-convergence should be interpreted as an indication of discontinuities or damage, rather than only numerical failure.

A separate limitation appeared during the strain computations. Even after filtering, unrealistically large strain values remained close to the specimen perimeter. This was likely related to the Geers strain algorithm, which estimated strain from neighbouring nodes. Near the specimen boundary, the neighbouring node configuration was incomplete or asymmetric. As a result, the strain calculation became less reliable and may have amplified boundary effects, producing a frame-like pattern around the edge. Therefore, the large strain values near the perimeter should not be interpreted as exact quantitative strain magnitudes.

The reliability issue therefore occurred in two stages. First, during DVC correlation, some nodes failed to converge, most of these carried a status of -3, meaning they diverged due to unrealistic deformation, which also accounts for the large strain values observed before filtering, Figure 4.4 b). However, even after filtering, large strain values persisted near the specimen perimeter (Figure 4.5), which can be explained by the second stage, the Geers strain computation. Here, boundary nodes produced unreliable strain values because the available neighbouring node information was incomplete, resulting in the frame-like pattern of larger strains along the specimen perimeter. For this reason, the strain magnitudes close to the specimen perimeter were interpreted qualitatively. They were used to identify where changes occur and to compare whether these regions correspond to displacement vectors and visually observed cracks.

The bit-depth comparison showed that convergence statistics alone were not sufficient to determine the most suitable representation. More greyscale detail is beneficial for DVC analysis as it increases sensitivity, though too much detail can introduce noise and therefore negatively affect correlation. The 16-bit data was considered to provide the better balance between these two factors and was therefore selected for the main analysis.

5.1.2 Strain Analysis

The strain analysis showed that displacement vectors and strain fields could be used together to qualitatively identify where local changes developed during freeze-thaw cycling. Damage was not distributed evenly throughout the specimen. Instead, the strongest strain localisations were found in regions where visual changes were also observed, especially close to the specimen perimeter. Several possible explanations for this pattern were considered. The outer region likely had the highest degree of water saturation and was probably more exposed to the water source. Even though the water was located at the top of the specimen, the experimental setup may have allowed water to be in contact with the sides as well, and prolonged contact would have contributed to higher saturation and greater pressure build-up during freezing, as shown in Section 4.3.1. Investigating the upper part of the specimen, which was in direct contact with water would be of particular interest in

this regard. The outer region was also more directly exposed to the imposed temperature changes, meaning it experienced the cooling and heating effects earlier. The sign of the volumetric strains provided additional information about the type of local deformation. Positive strain may indicate local expansion, opening or dilation, while negative strain may indicate compression, contraction or local closing. Mixed positive and negative strain was frequently observed in regions where cracking was also identified. This is plausible, since crack formation may cause local opening on one side of the crack while the surrounding material is compressed or displaced in another direction. The first strain analysis, presented in Subsection 4.2.2 and Appendix E, Tables E.1 and E.3 showed that mixed positive and negative strain often were associated with regions of stronger deformation, especially in the blue region. Positive strain alone appeared more often in regions with early or weaker indications of change, while negative strain was observed in some regions where cracking or local closing may occur. However, the strain sign should not be interpreted alone, as some regions were affected by the previously mentioned edge effects or inconclusive displacement vectors.

The second strain analysis, Appendix E, Tables E.4-E.5 supported the previously seen behaviour, but also showed a clearer contrast between regions. In several T-T comparisons, the red region mainly showed negative strain, while the yellow region mainly showed positive strain. This suggests that neighbouring regions may deform differently during the same interval, with one region undergoing local compression or closing while another undergoes expansion or opening.

The example with the detached material fragments presented in Figure 4.13, further emphasises the limitation of applying DVC to these datasets, where cracks form or small pieces of material move independently, breaking the assumption of DVC. The resulting displacement and strain fields therefore cause uncertainty in the interpretation, making these tools less suitable in these cases.

Based on visual comparison with image data, regions containing both positive and negative strain appeared to be the clearest indicators of local cracking or stronger deformation. Positive and negative strain may provide important information, but they should be interpreted in relation to the displacement vectors and the development of the visual crack. The results should therefore be used to identify where deformation occurs, rather than quantifying the exact strain magnitude.

5.2 Pore Segmentation Analysis

The brighter and darker void regions described in Section 4.3 (Figure 4.14) may arise from how each voxel integrates attenuation over a discrete volume. A voxel containing both air and water would then appear brighter than one filled with only air, but darker than the water-filled voxels. The brighter voids might therefore correspond to partially filled voxels rather than a separate phase. Figure 4.14 also showed that the pore behaviour varied with height, meaning that the slice at which a pore is examined determines what is observed.

Comparisons between pores should therefore sample each pore consistently and across several heights rather than relying on a single slice, which was also done in this analysis.

The pore segmentation results showed a clear difference between pores located near the outer region of the specimen and pores located closer to the centre. The outer pores were in general more saturated with water, while the inner pores were mainly air-filled, with water present only locally, often near the top or bottom of the pore. These results showed that the degree of saturation may be related to the pore position, with pores located closer to the perimeter appearing more saturated, as indicated in the analysis in Section 5.1.2. Examples of outer pores in Appendix F, Figure F.1 showed how fully saturated pores, which lacked air, were less able to accommodate the expansion associated with freezing, increasing local stresses and contributing to crack formation. The inner and intermediate pores were less saturated and more air-filled, and showed no visible cracking or water loss. This indicated that in these cases, the air acted as intended, reducing the stresses of the pore enough to prevent it from cracking. The number of cycles required for an outer pore to crack or become fully emptied could not be determined from the current analysis. This behaviour, and the number of cycles required, likely depends on several factors, including distance from the specimen edge, presence of other smaller pores nearby, presence of nearby cracks, saturation level and the size of the pore. For example, one of the clearest cases of a pore cracking and being emptied was shown in Appendix F, Figure F.4, which is a rather larger pore located close to the edge of the specimen. From these examples, it was not possible to determine whether the behaviour was mainly caused by the pore being located close to the specimen edge, by its relatively large size and water content, or by a combination of these factors. A larger number of pores would need to be analysed to separate these effects and provide a more statistically reliable interpretation.

The separation of water at the top of the pore during thawed states and at the bottom during frozen states indicates that gravity was not the dominant driver of this redistribution. Something actively moved water in opposite directions depending on thermal state, pointing instead toward mechanisms governed by the surrounding micropore network. As discussed in Section 1.2, the cement paste surrounding the larger air-entrained pores contains a network of smaller pores, such as capillary and gel pores. Due to surface tension effects, the smaller pores will freeze later than the larger pores. This means that during freezing, as ice forms and grows within the larger pore, water in the surrounding smaller pores can remain liquid at the same temperature. This could produce a driving force toward the growing ice and draw the unfrozen water from the pores into the larger pore. Under gravity, it could be assumed that the water would then accumulate at the bottom. Then, during thawing, as temperatures would rise, the ice would melt. The smaller surrounding pores that lost water during freezing would now draw the water back through capillary suction. However, this does not explain the presence of water at the top. One possible explanation could be that the micropore connections are predominantly located at the top of the pore. However, this remains a hypothesis and neutron data, and would need to be confirmed through analysis of the neutron data and further investigation of the surrounding pore network. The fact that the middle part of the pore remains air-filled further supports the hypothesis that water is not moving through the interior of the pore, but rather entering and exiting through connections in the pore wall. This would, again,

need to be confirmed through analysis of the neutron data.

A direct limitation of the pore segmentation was its sensitivity to initial image alignment. This dependency was especially visible for Dataset 12, where the misalignment made it impossible to gather any useful information from the segmentation. Despite these limitations, the overall pipeline was successful in processing the datasets and provided a way of segmenting and comparing pore behaviour during freeze-thaw cycling. The method made it possible to distinguish between pores that remained mostly stable, pores where water redistributed internally, and pores where water appeared to move into nearby cracks.

Both the strain analysis and pore segmentations indicate that most changes were observed along the perimeter of the specimen, supporting the hypothesis that the outer part of the specimen is more strongly affected by the freeze-thaw cycling.

Chapter 6

Conclusions and Future Work

This chapter presents the main conclusions from the developed workflow and the analysis of freeze-thaw behaviour in cement mortar. It summarises the findings from the registration, DVC, strain analysis and pore segmentation, and identifies future work needed to improve reliability and validate the results.

6.1 Conclusions

The thesis developed and evaluated a workflow for analysing freeze-thaw behaviour in cement mortar using X-ray tomography data. The workflow combined global registration, DVC-based displacement and strain analysis using **SPAM**, and joint histogram based pore segmentation using **evoSegment**. The developed pipelines were also made reusable for future ESS datasets.

The initial registration study showed that using a representative range, excluding an artefact-affected region, provided the most reliable and computationally efficient alignment. This was important as it showed that alignment of larger datasets could be made without using the full range of the dataset, which otherwise involved memory limitations.

The DVC parameter study showed that **hws** = 20 pixels, **nodeSpacing** = 20 pixels and **searchRange** = 20 pixels provided the best compromise between convergence, resolution and computational cost for this dataset. The 16-bit data was selected for the main analysis, because it preserved more greyscale detail and made artefacts easier to identify, even though the convergence behaviour was similar to the 8-bit data.

The strain analysis showed that DVC could be used to identify regions where deformation and cracking occur during freeze-thaw cycling. The clearest changes were mainly observed near the specimen perimeter, where most cracking was also visually identified. However, the absolute strain magnitudes are to be interpreted with caution, since boundary effects from the chosen strain algorithm, cracking and fragment movement and alignment uncertainty affect the reliability of the strain values. Therefore, the strain fields were used qualitatively to locate regions of change rather than to quantify exact deformation.

The pore segmentation showed a clear difference between outer and inner pores. Outer pores were generally more affected by the freeze-thaw cycles, often showing water loss and crack formation. In contrast, inner and intermediate pores were more air-filled, showed no cracking and showed clear water redistribution during the cycling. These results suggest that pore behaviour depends on the position within the specimen and water saturation.

Overall, the results show that the combined DVC and pore segmentation workflow can provide useful information about deformation, cracking and phase evolution within pores during repeated freeze-thaw cycles. However, the analysis also shows that the reliability of the results strongly depends on image alignment, artefacts, segmentation choice and the limitation of DVC in cracked regions.

6.2 Future Work

Future work should first focus on improving the experimental set-up. A stable reference feature outside the specimen could be included to make visual alignment between datasets easier, provided that it can be excluded from the final analysis. The thermocouple should also, if possible, be replaced or repositioned, since it introduced artefacts and prevented reliable analysis in the upper part of the specimen. In addition, all experimental changes should be documented to improve reproducibility and avoid unexplained differences between datasets.

For the DVC and strain analysis, future work should investigate methods for improving the reliability of results close to the specimen boundary, where the largest changes were often observed but where the strain calculation was least reliable. Padding or alternative boundary treatments could help recover more information near the specimen edges and at the first and last slice. The temperature data should also be analysed to help explain the observed effects.

For pore segmentation, a larger number of pores should be analysed to confirm whether the observed differences between inner, intermediate and outer pores are representative. As the present study was limited to 47 freeze-thaw cycles, extending beyond this number would help establish whether the outer region continues to grow inwards over time, and whether a fixed boundary, an intermediate pore group, or a cycle-dependent boundary best describes the separation. Future analysis should include pore size, estimate of water saturation, radial position, height within the specimen and distance to nearby cracks. Quantifying the segmented air and water volumes between freeze-thaw states would make it possible to estimate how much water remains, redistributes or escapes from the pores. Smaller pores should be investigated, such as capillary pores and gel pores, which are expected to behave differently than the air-filled voids, and might help explain when and why the pores crack. Here, improved initial registration and higher-resolution imaging would be needed. Finally, the segmentation results should be compared with neutron data to confirm whether the observed grey-value changes correspond to actual water redistribution.

Bibliography

- [1] International Energy Agency, “Breakthrough agenda report 2025,” International Energy Agency, Paris, 2025. [Online]. Available: <https://www.iea.org/reports/breakthrough-agenda-report-2025>.
- [2] S. O. Mitchell, “Concrete science,” *Concrete Construction - World of Concrete*, vol. 55, pp. 1–62, 4 Jan. 2001, ISSN: 15337316. DOI: [10.1016/B978-081551437-4.50004-7](https://doi.org/10.1016/B978-081551437-4.50004-7). [Online]. Available: <https://www.sciencedirect.com/science/chapter/edited-volume/pii/B9780815514374500047>.
- [3] S. Pilehvar et al., “Effect of freeze-thaw cycles on the mechanical behavior of geopolymer concrete and portland cement concrete containing micro-encapsulated phase change materials,” *Construction and Building Materials*, vol. 200, pp. 94–103, Mar. 2019, ISSN: 0950-0618. DOI: [10.1016/J.CONBUILDMAT.2018.12.057](https://doi.org/10.1016/J.CONBUILDMAT.2018.12.057). [Online]. Available: <https://doi.org/10.1016/J.CONBUILDMAT.2018.12.057>.
- [4] S. Guler and Z. F. Akbulut, “A comprehensive review of concrete durability in freeze-thaw conditions: Mechanisms, prevention, and mitigation strategies,” *Structures*, vol. 75, p. 108 804, May 2025, ISSN: 2352-0124. DOI: [10.1016/J.ISTRUC.2025.108804](https://doi.org/10.1016/J.ISTRUC.2025.108804). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2352012425006186>.
- [5] L. Xu, Y. Wang, Y. Wang, and T. Cheng, “A new concrete freeze–thaw damage model based on hydraulic pressure mechanism and its application,” *Materials*, vol. 18, p. 3708, 15 Aug. 2025, ISSN: 19961944. DOI: [10.3390/MA18153708](https://doi.org/10.3390/MA18153708). [Online]. Available: <https://pmc.ncbi.nlm.nih.gov/articles/PMC12348689/>.
- [6] A. Tengattini et al., “Next-grenoble, the neutron and x-ray tomograph in grenoble,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 968, p. 163 939, Jul. 2020, ISSN: 0168-9002. DOI: [10.1016/J.NIMA.2020.163939](https://doi.org/10.1016/J.NIMA.2020.163939). [Online]. Available: <https://doi.org/10.1016/J.NIMA.2020.163939>.
- [7] A. Kalra, “Developing fe human models from medical images,” *Basic Finite Element Method as Applied to Injury Biomechanics*, pp. 389–415, Jan. 2018. DOI: [10.1016/B978-0-12-809831-8.00009-X](https://doi.org/10.1016/B978-0-12-809831-8.00009-X). [Online]. Available: <https://doi.org/10.1016/B978-0-12-809831-8.00009-X>.
- [8] I. A. E. Agency, “Cold neutron sources: Practical considerations and modern research,” *IAEA Tecdoc Series*, 2023. [Online]. Available: https://www-pub.iaea.org/MTCD/Publications/PDF/TE-2025_web.pdf.

- [9] N. Kardjilov, I. Manke, R. Woracek, A. Hilger, and J. Banhart, “Advances in neutron imaging,” *Materials Today*, vol. 21, pp. 652–672, 6 Jul. 2018, ISSN: 1369-7021. DOI: [10.1016/J.MATTOD.2018.03.001](https://doi.org/10.1016/J.MATTOD.2018.03.001). [Online]. Available: <https://doi.org/10.1016/J.MATTOD.2018.03.001>.
- [10] R. Schofield et al., “Image reconstruction: Part 1 – understanding filtered back projection, noise and image acquisition,” *Journal of Cardiovascular Computed Tomography*, vol. 14, pp. 219–225, 3 May 2020, ISSN: 1934-5925. DOI: [10.1016/J.JCCT.2019.04.008](https://doi.org/10.1016/J.JCCT.2019.04.008). [Online]. Available: <https://doi.org/10.1016/J.JCCT.2019.04.008>.
- [11] P. J. Withers et al., “X-ray computed tomography,” *Nature Reviews Methods Primers*, vol. 1, p. 18, 1 2021, ISSN: 2662-8449. DOI: [10.1038/s43586-021-00015-4](https://doi.org/10.1038/s43586-021-00015-4). [Online]. Available: <https://www.nature.com/articles/s43586-021-00015-4>.
- [12] J. A. Elliott, “Gibbsian surface thermodynamics,” *The Journal of Physical Chemistry B*, vol. 124, pp. 10 859–10 878, 48 Dec. 2020, ISSN: 15205207. DOI: [10.1021/ACS.JPCB.0C05946](https://doi.org/10.1021/ACS.JPCB.0C05946). [Online]. Available: [/doi/pdf/10.1021/acs.jpcb.0c05946?ref=article_openPDF](https://doi.org/10.1021/acs.jpcb.0c05946?ref=article_openPDF).
- [13] M. G. Matar, A. N. Aday, B. C. Acarturk, and W. V. Srubar, “A review of air-entraining agents for concrete: Chemistry, air-void parameters, and freeze-thaw resistance,” *Journal of Building Engineering*, vol. 114, p. 114 262, Nov. 2025, ISSN: 2352-7102. DOI: [10.1016/J.JOBE.2025.114262](https://doi.org/10.1016/J.JOBE.2025.114262). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2352710225024994>.
- [14] S. Diamond, “Physical and chemical characteristics of cement composites,” *Durability of Concrete and Cement Composites*, pp. 10–44, Jan. 2007. DOI: [10.1533/9781845693398.10](https://doi.org/10.1533/9781845693398.10). [Online]. Available: <https://doi.org/10.1533/9781845693398.10>.
- [15] P. Zhang, F. H. Wittmann, M. Vogel, H. S. Müller, and T. Zhao, “Influence of freeze-thaw cycles on capillary absorption and chloride penetration into concrete,” *Cement and Concrete Research*, vol. 100, pp. 60–67, Oct. 2017, ISSN: 0008-8846. DOI: [10.1016/J.CEMCONRES.2017.05.018](https://doi.org/10.1016/J.CEMCONRES.2017.05.018). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0008884617300418>.
- [16] H. Wang et al., “Comprehensive assessment and freeze–thaw durability prediction of wet-sprayed concrete for cold-region tunnels,” *Materials 2025, Vol. 18, Page 2955*, vol. 18, p. 2955, 13 Jun. 2025, ISSN: 1996-1944. DOI: [10.3390/MA18132955](https://doi.org/10.3390/MA18132955). [Online]. Available: <https://www.mdpi.com/1996-1944/18/13/2955/htm%20https://www.mdpi.com/1996-1944/18/13/2955>.
- [17] T. C. Powers, “The air requirement of frost-resistant concrete,” *Proceedings of the Highway Research Board*, vol. 29, pp. 184–211, 1949.
- [18] T. C. Powers and R. A. Helmuth, “Theory of volume changes in hardened port-land-cement paste during freezing,” *Proceedings of the Highway Research Board*, vol. 32, pp. 285–297, 1953.

- [19] M. Chen, N. J. Tustison, R. Jena, and J. C. Gee, “Image registration: Fundamentals and recent advances based on deep learning,” *Neuromethods*, vol. 197, pp. 435–458, Jul. 2023, ISSN: 19406045. DOI: [10.1007/978-1-0716-3195-9_14](https://doi.org/10.1007/978-1-0716-3195-9_14). [Online]. Available: <https://www.ncbi.nlm.nih.gov/books/NBK597490/>.
- [20] SPAM core dev. team, *Tutorial: Refreshments in continuum mechanics spam documentation*, 2023. [Online]. Available: <https://www.spam-project.dev/docs/tutorials/tutorial-00-refreshments.html>.
- [21] G. Viggiani, S. Hall, and E. Romero, Eds., *ALERT Doctoral School 2012: Advanced experimental techniques in geomechanics*, English. ALERT Geomaterials, 2012, ISBN: 978-3-00-039683-0.
- [22] G. G. Gray, J. K. Morgan, and P. F. Sanz, “Overview of continuum and particle dynamics methods for mechanical modeling of contractional geologic structures,” *Journal of Structural Geology*, vol. 59, pp. 19–36, Feb. 2014, ISSN: 0191-8141. DOI: [10.1016/J.JSG.2013.11.009](https://doi.org/10.1016/J.JSG.2013.11.009). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0191814113002125>.
- [23] D. Rees, *Basic Engineering Plasticity: An Introduction with Engineering and Manufacturing Applications*. Butterworth-Heinemann, 2006, ISBN: 0-7506-8025-3.
- [24] J. A. Zimmerman, D. J. Bammann, and H. Gao, “Deformation gradients for continuum mechanical analysis of atomistic simulations,” *International Journal of Solids and Structures*, vol. 46, pp. 238–253, 2 Jan. 2009, ISSN: 0020-7683. DOI: [10.1016/J.IJSOLSTR.2008.08.036](https://doi.org/10.1016/J.IJSOLSTR.2008.08.036). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S002076830800348X>.
- [25] SPAM core dev. team, *Tutorial: Image correlation, theory, spam documentation*, 2023. [Online]. Available: <https://www.spam-project.dev/docs/tutorials/tutorial-02a-DIC-theory.html>.
- [26] G. M. Hassan, “Deformation measurement in the presence of discontinuities with digital image correlation: A review,” *Optics and Lasers in Engineering*, vol. 137, p. 106394, 3 Feb. 2021, ISSN: 01438166. DOI: [10.1016/j.optlaseng.2020.106394](https://doi.org/10.1016/j.optlaseng.2020.106394). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0143816620309830>.
- [27] B. Pan and B. Wang, “Some recent advances in digital volume correlation,” *Optics and Lasers in Engineering*, vol. 135, p. 106189, 1-2 Dec. 2020, ISSN: 01438166. DOI: [10.1016/j.optlaseng.2020.106189](https://doi.org/10.1016/j.optlaseng.2020.106189). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0143816619315982>.
- [28] H. Zhang et al., *On the upper and lower bounds of correlation window size in digital image correlation analysis*, Jul. 2017. [Online]. Available: <https://researchportal.bath.ac.uk/en/publications/on-the-upper-and-lower-bounds-of-correlation-window-size-in-digit/>.

- [29] S. P. Osipov, S. V. Chakhlov, V. Y. Zhvyrbliia, D. A. Sednev, O. S. Osipov, and E. Y. Usachev, “The nature of metal artifacts in x-ray computed tomography and their reduction by optimization of tomography systems parameters,” *Applied Sciences* 2023, Vol. 13, Page 2666, vol. 13, p. 2666, 4 Feb. 2023, ISSN: 2076-3417. DOI: [10.3390/AP13042666](https://doi.org/10.3390/AP13042666). [Online]. Available: <https://www.mdpi.com/2076-3417/13/4/2666/htm%20https://www.mdpi.com/2076-3417/13/4/2666>.
- [30] J. Hektor, J. Engqvist, and S. A. Hall, “Evosegment: 4d image segmentation of microstructural evolution using joint histograms,” *Tomography of Materials and Structures*, vol. 4, p. 100 023, Mar. 2024, ISSN: 2949673X. DOI: [10.1016/j.tmater.2023.100023](https://doi.org/10.1016/j.tmater.2023.100023).
- [31] O. Stamati et al., “Spam: Software for practical analysis of materials,” *Journal of Open Source Software*, vol. 5, p. 2286, 51 Jul. 2020. DOI: [10.21105/joss.02286](https://doi.org/10.21105/joss.02286).
- [32] M. G. Geers, R. D. Borst, and W. A. Brekelmans, “Computing strain fields from discrete displacement fields in 2d-solids,” *International Journal of Solids and Structures*, vol. 33, pp. 4293–4307, 29 Dec. 1996, ISSN: 0020-7683. DOI: [10.1016/0020-7683\(95\)00240-5](https://doi.org/10.1016/0020-7683(95)00240-5). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0020768395002405?via%3Dihub>.

Appendix A


Overview of the Experimental Setup and Experiment at ILL NeXT

The relevant experiment for this study was performed at the NeXT instrument at ILL in Grenoble, France, using an in situ freeze-thaw system, allowing simultaneous neutron and X-ray imaging during controlled temperature cycling. A cylindrical cement specimen (10 mm height, 10 mm diameter), fully saturated with water, was fixed inside a custom-built temperature rig positioned in a multimodal neutron and X-ray beamline. Approximately 0.5 ml of water was added on top of the sample.


The temperature rig was based on a Peltier thermoelectric element coupled to a circulating cooling-fluid loop, allowing the specimen to be cooled below 0 °C and subsequently thawed. The experimental setup was designed to allow uninterrupted rotation during tomography. The temperature of the system was monitored at three positions: in the water above the sample, in the aluminium holder close to the Peltier element, and inside the sample. See example images of the setup in Figure A.2 and schematic in Figure A.1.

After an initial reference scan (20 °C), the sample underwent one stronger freezing step, using a lower target temperature than the later cycles (-25 °C), followed by a stronger thawing step (stepwise heating of first 10 °C and then 20 °C). Repeated freeze-thaw cycling was then performed using ordinary cycles with a cooling target of -20 °C and a heating target of 10 °C, where cooling initially took 30 minutes and heating 10 minutes, but the cooling time was then changed to 35 minutes. Tomographies were acquired at selected frozen and thawed states throughout the cycling sequence.

X-ray tomographic acquisitions were performed with a voxel size of approximately 10 μm . Neutron tomographic acquisitions were performed using a 23 mm pinhole and a binning of 2, with a voxel size of approximately 25 μm .



a)



b)

Figure A.2: Photographs of the experimental set-up at the NeXT instrument at ILL. a) Overview of the rig mounted in the neutron/X-ray imaging beamline before complete mounting. b) Close up of the mounted specimen and temperature cooling system.

- Temperature sensor position




Figure A.1: Schematic illustration of the in situ freeze-thaw system used at the NeXT instrument at ILL. The saturated cement specimen mounted in a temperature controlled rig with approximately 0.5 ml water above the sample. Freezing and thawing were controlled by the Peltier element coupled to a circulating cooling-fluid loop through a slip-ring connection. This set-up allowed for tomographic rotation during the experiment. The temperature was monitored at three positions: in the water above the sample (T1), in the aluminium holder above the Peltier element (T2), and inside the sample (T3).

Appendix B

Example of All Three Strain Algorithm Outputs




Figure B.1: Both volumetric and deviatoric strain calculated for the same index as Figure 4.7 using the three different ways of computing the strain. Comparing the three we immediately notice that the two right strains are much more similar than the furthest left one, as expected, due to it using the noisy displacement fields for computing the strains. As we know, there is some noise in the upper left part of the specimen, causing this type of strain artefact. It can be argued that it is much more emphasised in the Q8 method and that we overall have more variations in the Q8 variant than in the Geers variant. In the Geers variant we also get this characteristic cross- which is good, as it is easy to then identify noise when we have it present, that in contrast to the Q8 approach that mostly makes the effect of the noise affect more neighbours and stand out more. However, one can argue that the edge effects are more predominant in the Geers approach.

Appendix C

Interpretation of Pore Segmentations




Figure C.1: a) Segmented pores in image domain. b) Segmented joint histogram.

The resulting segmentation seen in Figure C.1 can be interpreted based on the position of clusters in the joint histogram. Voxels located along the diagonal correspond to regions that do not change between the reference and deformed states. In Figure C.1 b), three such (unchanged) clusters can be seen:

- **1. Green:** background (unchanged),
- **2. Blue:** water-filled pore regions (unchanged),
- **3. Velvet:** air-filled regions (unchanged).

Voxels located away from the diagonal represent voxels that undergo changes in greyscale intensity, and are therefore interpreted as transitions between states:

- **4. Yellow:** voxels that became brighter and are classified as background, indicating that pore regions have shrunk or shifted,

- **5. Purple:** voxels that became darker and are classified as water, indicating that background regions have transitioned into pore space,
- **6. Light brown:** voxels that became brighter and are classified as water, corresponding to a transition from air to water,
- **7. Teal:** voxels that became darker and are classified as air, indicating that background or water regions have transitioned into air-filled regions.

In regards to phase evolution, the following analysis can be made. During the initial freeze-thaw cycles, limited redistribution between water and air is observed, with the pore predominantly remaining air-filled (velvet). In later cycles, subtle changes emerge, at "Thawed 2" we note how most of the air-filled regions are replaced by water (light brown), indicating fluid redistribution. At Dataset 5, the thawed state shows an increased amount of water content in the given slice, compared to the frozen state.

Here, we can also clearly see how the misalignment effects are segmented. Particularly in "Frozen 5", where the yellow regions at the top and the purple regions at the bottom indicate an apparent downwards movement of the pore, where voxels are incorrectly classified as background above and as pore space below. A more severe misalignment is shown in the final dataset, "Thawed 6", where these misalignments dominate the segmentation, and where the least teal segmentation is seen representing the transition from pore to air. The severe misalignment makes it difficult to interpret the real physical transition with respect to the reference state.

Appendix D

Registration Result Tables

Table D.1: Registration results for Dataset 0 compared with Dataset 1 to Dataset 6. Binning of 2.

Range	Error	Iter.	Time	Translation $t = [t_z, t_y, t_x]$	Rotation $R = [r_z, r_y, r_x]$
Dataset 0 and Dataset 1. Expected axial shift: ~ 4 slices (t_z).					
Thermocouple	130.01	30	6m 3s	[-4.418, 0.122, 0.091]	[0.184, -0.036, -0.050]
Full	121.19	36	12m 56s	[-4.330, 0.214, 0.085]	[0.189, -0.055, -0.074]
Representative	115.85	28	5m 31s	[-4.205, 0.205, 0.011]	[0.194, -0.041, -0.048]
Dataset 0 and Dataset 2. Expected axial shift: ~ 2 slices (t_z).					
Thermocouple	132.12	46	4m 49s	[-1.655, 0.110, 0.058]	[0.145, -0.002, -0.073]
Full	123.24	40	8m 27s	[-1.611, 0.225, 0.075]	[0.139, -0.002, -0.107]
Representative	121.27	39	4m 15s	[-1.782, 0.222, 0.033]	[0.097, -0.002, -0.067]
Dataset 0 and Dataset 3. Expected axial shift: ~ 4 slices (t_z).					
Thermocouple	150.78	63	6m 29s	[-3.680, 0.296, 0.253]	[0.240, -0.029, -0.018]
Full	142.13	91	18m 32s	[-3.627, 0.689, 0.513]	[0.206, -0.036, -0.060]
Representative	136.83	105	10m 53s	[-3.696, 0.237, 0.100]	[0.052, -0.023, -0.013]
Dataset 0 and Dataset 4. Expected axial shift: ~ 2 slices (t_z).					
Thermocouple	151.71	70	7m 21s	[-1.591, 0.309, 0.218]	[-0.259, -0.007, 0.010]
Full	144.93	26	5m 36s	[-1.507, 0.746, 0.529]	[-0.308, 0.002, -0.009]
Representative	143.28	51	5m 27s	[-1.519, 0.260, 0.141]	[-0.348, 0.001, 0.010]
Dataset 0 and Dataset 5. Expected axial shift: ~ 3 slices (t_z).					
Thermocouple	164.25	114	12m 31s	[-2.722, 0.362, 0.256]	[0.121, -0.035, 0.008]
Full	154.52	47	9m 41s	[-2.708, 1.031, 0.710]	[0.078, -0.045, -0.044]
Representative	149.69	59	6m 9s	[-3.025, 0.264, 0.112]	[-0.004, -0.030, 0.013]
Dataset 0 and Dataset 6. Expected axial shift: ~ 1 slice (t_z).					
Thermocouple	169.25	96	10m 6s	[-0.515, 0.401, 0.312]	[0.149, -0.032, 0.030]
Full	155.12	114	23m 20s	[-0.481, 1.119, 0.863]	[0.082, -0.035, 0.007]
Representative	139.48	53	5m 48s	[-0.734, 0.272, 0.156]	[0.010, -0.022, 0.023]

Table D.2: Registration results for Dataset 0 compared with Dataset 7 to Dataset 12. Binning of 2.

Range	Error	Iter.	Time	Translation $t = [t_z, t_y, t_x]$	Rotation $R = [r_z, r_y, r_x]$
Dataset 0 and Dataset 7. Expected axial shift: ~ 3 slices (t_z).					
Thermocouple	175.48	106	11m 20s	[-2.834, 0.339, 0.241]	[-0.506, -0.044, 0.016]
Full	165.49	54	11m 40s	[-2.870, 1.097, 0.746]	[-0.598, -0.051, -0.052]
Representative	164.53	58	6m 24s	[-3.279, 0.232, 0.072]	[-0.669, -0.034, 0.018]
Dataset 0 and Dataset 8. Expected axial shift: ~ 1 slice (t_z).					
Thermocouple	178.66	147	16m 32s	[-0.318, 0.425, 0.322]	[-0.284, -0.051, 0.059]
Full	160.62	149	32m 3s	[-0.388, 1.444, 1.103]	[-0.434, -0.037, 0.013]
Representative	164.02	68	10m 40s	[-0.576, 0.272, 0.139]	[-0.515, -0.028, 0.040]
Dataset 0 and Dataset 9. Expected axial shift: ~ 3 slices (t_z).					
Thermocouple	187.77	143	15m 31s	[-2.419, 0.378, 0.251]	[-0.378, -0.060, 0.060]
Full	174.21	61	13m 13s	[-2.440, 1.468, 1.038]	[-0.573, -0.084, 0.023]
Representative	174.10	56	6m 40s	[-3.074, 0.212, 0.067]	[-0.688, -0.033, 0.035]
Dataset 0 and Dataset 10. Expected axial shift: ~ 1 slice (t_z).					
Thermocouple	187.87	166	18m 29s	[-0.207, 0.432, 0.286]	[-0.381, -0.058, 0.082]
Full	170.52	124	28m 9s	[-0.227, 1.655, 1.170]	[-0.544, -0.059, 0.058]
Representative	168.28	66	7m 28s	[-0.746, 0.253, 0.107]	[-0.608, -0.027, 0.050]
Dataset 0 and Dataset 11. Expected axial shift: ~ 3 slices (t_z).					
Thermocouple	190.20	147	16m 53s	[-1.800, 0.351, 0.243]	[-0.398, -0.043, 0.043]
Full	177.65	53	12m 12s	[-1.806, 1.387, 0.996]	[-0.556, -0.082, 0.026]
Representative	176.11	55	6m 4s	[-2.538, 0.210, 0.071]	[-0.616, -0.030, 0.032]
Dataset 0 and Dataset 12. Expected axial shift: ~ 15 slices (t_z).					
Thermocouple	279.64	64	23m 4s	[10.109, 0.183, 0.402]	[2.893, -0.085, -0.020]
Full	249.15	508	32m 55s	[10.884, -0.050, 0.589]	[2.028, -0.160, -0.095]
Representative	219.40	39	7m 7s	[12.734, 0.025, 0.212]	[1.029, -0.026, -0.024]

Appendix E

Strain Analysis Tables and Figures

E.1 Changes in the Respective Ranges 1

Table E.1: Summary of strain localisation in the frozen-frozen (F-F) comparisons. The colours refer to the regions marked in Figure 4.6. Green indicates clear changes in strain and/or displacement vectors in expected regions, yellow indicates weaker or uncertain changes that may be real change or noise, and red indicates no visible change.

Datasets/State	Region	Observation
1-3 / F-F	Red	Change, mainly positive strain.
	Yellow	Change, cannot be distinguished from bulk movement.
	Blue	Clear cracking, mixed positive/negative strain.
3-5 / F-F	Red	Small change.
	Yellow	No visible change.
	Blue	Small change, mixed positive/negative strain.
5-7 / F-F	Red	Cracking, cannot be distinguished from edge strain.
	Yellow	Change, cannot be distinguished from edge strain.
	Blue	Clear cracking, mixed positive/negative strain.
7-9 / F-F	Red	Change, mixed positive/negative strain.
	Yellow	Change, mixed positive/negative strain.
	Blue	Large change, mixed positive/negative strain.
9-11 / F-F	Red	No visible change.
	Yellow	No visible change.
	Blue	No visible change.

Table E.2: Summary of strain localisation in the frozen-thawed (F-T) comparisons. The colours refer to the regions marked in Figure 4.10. Green indicates clear changes in strain and/or displacement vectors in expected regions, yellow indicates weaker or uncertain changes that may be real change or noise, and red indicates no visible change.

Datasets/State	Region	Observation
1-2 / F-T	Red	Indication of change, positive strain.
	Yellow	No change.
	Blue	No change.
3-4 / F-T	Red	Clear cracking, mixed positive/negative strain.
	Yellow	Change, cannot be distinguished from bulk movement.
	Blue	Indication of change, negative strain.
5-6 / F-T	Red	Indication of cracking, mixed positive/negative strain.
	Yellow	No change.
	Blue	Indication of change, mixed positive/negative strain.
7-8 / F-T	Red	Indication of change, some negative strain.
	Yellow	No change.
	Blue	Indication of change, negative strain.
9-10 / F-T	Red	Change, some mixed positive/negative strain.
	Yellow	Change, cannot be distinguished from bulk movement.
	Blue	Change, mixed positive/negative strain.
11-12 / F-T	Red	Inconclusive displacement; large positive strain.
	Yellow	Inconclusive displacement; large negative strain.
	Blue	Inconclusive displacement; large positive strain.

Table E.3: Summary of strain localisation in the thawed-thawed (T-T) comparisons. The colours refer to the regions marked in Figure 4.10. Green indicates clear changes in strain and/or displacement vectors in expected regions, yellow indicates weaker or uncertain changes that may be real change or noise, and red indicates no visible change.

Datasets/State	Region	Observation
2-4 / T-T	Red	Indication of change, mixed strain.
	Yellow	Indication of change, positive strain.
	Blue	No change.
4-6 / T-T	Red	Indication of change, mixed strain.
	Yellow	Bulk movement/change; strain not comparable.
	Blue	Indication of change, positive strain.
6-8 / T-T	Red	Clear cracking, negative strain.
	Yellow	Bulk movement/change; strain not distinguishable.
	Blue	Indication of change, mixed strain.
8-10 / T-T	Red	Indication of change, negative strain.
	Yellow	No change.
	Blue	Indication of change, mixed strain.
10-12 / T-T	Red	Inconclusive displacement; large positive strain.
	Green	Inconclusive displacement; large negative strain.
	Blue	Inconclusive displacement; large positive strain.

E.2 Changes Examples 2




Figure E.1: a) Dataset 0 (reference), slice 120. b) Dataset 11 (thawed), slice 120. Five regions are circled in each image, representing regions where we visually can confirm changes between the two datasets.

E.3 Changes in the Respective Ranges 2

Table E.4: Summary of strain localisation in the frozen-frozen (F-F) comparisons. The colours refer to the regions marked in Figure E.1. Green indicates clear changes in strain and/or displacement vectors in expected regions, yellow indicates weaker or uncertain changes that may be real change or noise, and red indicates no visible change.

Datasets/State	Region	Observation
1-3 / F-F	Red	Indication of change, positive strain.
	Yellow	Indication of change, negative strain.
3-5 / F-F	Red	Small indication of change, mainly bulk movement; positive strain.
	Yellow	Small indication of change, mainly bulk movement; positive strain.
5-7 / F-F	Red	Indication of change, mixed positive/negative strain.
	Yellow	Cracking, mixed positive/negative strain.
7-9 / F-F	Red	Bulk movement/change, positive strain.
	Yellow	Bulk movement/change, mainly negative strain.
9-11 / F-F	Red	No change.
	Yellow	No change.

Table E.5: Summary of strain localisation in the thawed-thawed (T-T) comparisons. The colours refer to the regions marked in Figure E.1. Green indicates clear changes in strain and/or displacement vectors in expected regions, yellow indicates weaker or uncertain changes that may be real change or noise, and red indicates no visible change.

Datasets/State	Region	Observation
2-4 / T-T	Red	Clear change, negative strain.
	Yellow	Clear change, positive strain.
4-6 / T-T	Red	Indication of change, negative strain.
	Yellow	Indication of change, positive strain.
6-8 / T-T	Red	Indication of cracking, negative strain.
	Yellow	Clear cracking, positive strain.
8-10 / T-T	Red	Clear cracking, mixed positive/negative strain.
	Yellow	Indication of change, positive strain.
10-12 / T-T	Red	Inconclusive displacement; large negative strain.
	Yellow	Inconclusive displacement; large positive strain.




Figure E.2: a) Visualises a change that is also visible in scenario b). c) shows the displacement vectors corresponding to b). d) shows a case with little change, at least in the two regions where change was expected. In all of these cases a persistent strain pattern appears at the bottom right, contributing to a large positive strain. The large magnitude would suggest noise but its consistency across many cycles makes it more likely that it reflects a real change that cannot be confirmed visually. Unlike the previous examples, the strains here are more spread out, though still connected to the specimen perimeter.




Figure E.3: Dataset 12. a) Strains overlaid on slice 220. b) Displacement vectors overlaid on slice 220.

Appendix F

Outer Pore Segmentations




Figure F.1: Four segmented outer pores, c) and d) located along the outer perimeter of the specimen, b) located closer to the outer perimeter of the specimen and a) located the furthest towards the centre of the specimen (considered an intermediate pore). a) Slice 435/977. b) Slice 32/977. c) Slice 239/977. d) Slice 315/977.




Figure F.2: Segmented histograms belonging to Figure F.1.






Figure F.3: Segmented intermediate pore, visualising the pore segmentation at the a) top (slice 454/977), b) middle (slice 435/977) and c) bottom (slice 418/977).

a)



b)



c)




Figure F.4: Segmented outer pore, visualising the pore segmentation at the a) top (slice 262/977), b) middle (slice 239/977) and c) bottom (slice 222/977). The subtle crack that starts to appear around Frozen Dataset 3 might be the culprit behind the loss of water.




Figure F.5: Segmented histograms a) belonging to Figure F.3 a). b) belonging to Figure F.3 c). c) belonging to Figure F.4 a). d) belonging to Figure F.4 c).

Appendix G

Inner Pore Segmentations




Figure G.1: Four segmented inner pores, d) being closest to centre, a) being the second closest, then b) and c) being equally far away from the centre. a) Slice 32/977. b) Slice 289/977. c) Slice 330/977. d) Slice 222/977.




Figure G.2: Segmented histograms a) belonging to Figure G.1 b). b) belonging to Figure G.1 c). c) belonging to Figure G.1 d). Histogram a) is already printed in Section 4.3 Figure 4.14.