The application of machine learning to understand the dynamics of soil structure

By Caroline Roy, BSc. (Hons)

Thesis submitted to the University of Nottingham for the degree of Doctor of Philosophy

2025

Abstract

Machine learning and artificial intelligence are quickly changing the world. Just a few years ago, these concepts usually provoked visions of theoretical fiction. However, now they are becoming commonplace across a vast array of industries and even in our daily lives.

Soil structure is a very important factor in soil health and has strong applications in agricultural studies and improvement. It is also a dynamic, complex, and very hard-to-measure component. Scientific advancements which improve the accessibility of soil structure quantifications are of great importance and benefit to the soil science community.

X-ray Computed Tomography (X-ray CT) is a non-destructive and non-invasive technique for observing and analysing soil structure. It has been used in soil science for several years for various tasks, leading to a large volume of data to be analysed. A lot of this data is often discounted, due to the extreme volumes produced.

Machine learning allows for high-throughput data processing, unrestricted by human processing and labour limitations. Data processing and analysis via machine learning algorithms might also notice new trends or patterns in the data that would not be noticed with human analysis. This means that machine learning algorithms could be the perfect solution to take a deeper dive into the data collected from X-ray CT, get more information out of the data, and potentially find new patterns.

Soil structure can vary greatly due to the soil composition, weather conditions (temperature, rain, or lack thereof), crops, compaction, tillage management and the biodiversity in the soil. In recent years there has been a lot of interest in zero tillage, as part of a series of measures usually deployed under the banner of conservation or regenerative agriculture and how it may be beneficial over conventional tillage. This research looks at how the tillage practices impact the soil structure and in particular the pore network, and how these properties may vary over a growing season. Different machine learning techniques are looked at and applied to demonstrate how machine learning can be used alongside X-ray CT images and soil data to improve processes and observe new information.

Ultimately, the large volume of data produced by X-ray CT make excellent training datasets for machine learning algorithms. The neural networks demonstrated were mostly successful at identifying and predicting soil type, water content and tillage treatment based on soil structure. Shortcomings could be overcome by increasing the volume and variation of training data, as well as making use of more advanced machine learning techniques such as multi-resolution networks.

Keywords

Tillage, Zero Tillage, X-ray Computed Tomography, Soil Structure, Pore Network, Machine Learning

Acknowledgements

I would like to express my deepest appreciation to Prof. Sacha Mooney and Prof. Tony Pridmore for their invaluable and constant patience, feedback, and support. Sacha for all your guidance and support throughout extension requests and the help you provided during the difficult times I faced completing this thesis. I'm extremely grateful to Dr. Mohammadreza Soltaninejad, without you I reckon I'd still be having issues with my code.

Additionally, this project would not have been possible without the generous support from BBSRC, who financed my research.

Special thanks to Dr. Atkinson, for all the support with the X-ray CT equipment and specialist software, as well as some great mental health walks around campus. Prof. Andrew French was incredible with helping me through final corrections and making me feel like my work mattered. Thanks also to Dr Cooper, for being a friendly face when I first started, and always being happy to help.

Thank you to my family, for always supporting me no matter which direction I decided to go in life, or how long it took me to get there. Also, to Gwyneth, you'll never know how much you inspired me to keep going.

Lastly, I'd like to mention and thank Alex, for always believing in me even when I couldn't believe in myself, and for supporting me through the long and often difficult journey that this PhD has taken. His encouragement and patience have meant more than I can put into words.

Contents

Abbreviations and acronyms

µA	Microampere
µCT	X-ray Microtomography
µm	Micrometre
2-D	Two-Dimensional
3-D	Three-Dimensional
4-D	Four-Dimensional
AI	Artificial Intelligence
ANN	Artificial Neural Network
CNN	Convolutional Neural Network
DNN	Deep Neural Network
GAN	Generative Adversarial Networks
GPS	Global Positioning System
kPa	Kilopascal
kV	Kilovolts
LUCAS	Land Use/Land Cover Area Frame Survey
NLP	Natural Language Processing
RGB	Red, green, and blue
RL	Reinforcement Learning
RLHF	Reinforcement Learning with Human Feedback
RSA	Root System Architecture
VESS	Visual Evaluation of Soil Structure
X-ray CT	X-ray Computed Tomography
YOLO	You Only Look Once
ZT	Zero-Till/Zero-tilled

List of Tables

List of Figures

1. Introduction

1.1 General Introduction

Soil structure describes the arrangement and sizes of soil aggregates and pore space between and within aggregate structures. Structure can also be used to determine or define the quality of a soil (Dexter 1988) as well as being an indicator of water flow through a soil (Wilson *et al.* 1988, Dexter *et al.* 2001). As such, developing a deeper understanding of soil structure and new ways to determine the structure of a soil are important to future crop growth in a world with a rapidly changing climate and growing populations.

X-ray CT is a non-invasive and non-destructive imaging technique; originally it was designed for medical functions (Hounsfield 1980) but has also been applied to soils since the 1980s (Petrovic *et al.* 1982). Traditional methods of observing soil structure and plant growth are often destructive and disturb the sample. The non-destructive and non-invasive elements of X-ray CT mean that plant roots can be observed at several time points during the growth of the plant.

While X-ray CT is a highly useful tool for observing soil structure as well as plant and root growth in a non-destructive manner, it also creates a high volume of data, especially when you consider experimental designs involving replicates or samples at multiple time points. Although the exact number of images created from X-ray CT scans depends on the sample size and the scan settings, one scan could easily result in over 1000 images. When this is multiplied by number of samples and potential for repeat scans over time periods, the number of images for one experiment could quickly get into 5-figures, if not more. Oftentimes, a lot of this data would be disregarded to save on manual processing time. This large data volume issue could be solved by using machine learning to automate data analysis. Human eyes may not be able to differentiate between a few pixels difference when observing two adjacent images from a sample, but machine learning algorithms can identify these variations as highly significant.

At a basic level automated image analysis involves a computer being programmed to identify specific features within a set of images, usually relying upon human input for the initial feature labelling. This human input can restrict the effectiveness of the final algorithm and limits the algorithm to a specific task instead of a more general use. Artificial neural networks (ANNs) are networks inspired by the human brain and the way humans learn (Salu 1985, Anderson 1988, Heskes *et al.* 1991). They have significantly advanced machine learning by enabling systems to automatically learn to analyse complex patterns and representations from data, improving performance in tasks such as image and speech recognition. ANNs are often applied to identify data patterns which are too complicated or extensive for manual extraction by a human. The ability to automatically learn complex patterns from large datasets reduces the need for manual feature programming. Machine learning can go one step further than ANNs with the use of convolutional neural networks (CNNs). CNNs have many more layers of neural networks than ANNs and so are considered a form of deep learning. Deep learning is extensively applied in computer vision, particularly in object detection and image generation, as well as in natural language processing, speech recognition and investment decision making (Zhong *et al.* 2019).

1.2 Research Justification

In the general field of X-ray CT there is a lot of data generated with every sample scanned. One soil sample in a 300mm height tube with a diameter of 50mm scanned a resolution of 40μm could result in just under 2000 images from the top down, totalling around 3GB. The scan time for the above example was 30 minutes, followed by an additional 10 minutes for the automatic reconstruction of the 3-D image. The reconstruction, performed by the scanner's integrated software, involved stitching sequential 2D images into a coherent 3-D volume, aligning them based on their spatial coordinates to accurately represent the internal structure of the sample. It could then take several hours to process the data manually, depending on the aim of the study. Manually identifying and marking root systems architecture is even more time consuming and could easily take several hours or even days. Multiply this by the number of replicates and/or time points an experiment may have, and you easily end up with a research project that takes months or even years to analyse and totalling several hundred GB of data.

If a computer program could analyse X-ray CT data of soil quickly then that would support the ability to increase the amount of data collected, especially where replication is usually limited by analytical time. Machine learning also has the benefit of being able to identify patterns that may be overlooked by humans (Heung *et al.* 2016), so this approach could potentially uncover features previously unnoticed in soil types, soil treatments, root growth etc.

1.3 Methodological Approach

This thesis focuses on soil structure dynamics, specifically the physical changes in soil architecture that occur in response to external factors such as wetting and drying cycles and different tillage treatments. Soil structure dynamics are critical to understanding soil behaviour, as changes in pore connectivity and the soil pore network can influence water retention, root penetration and microbial activity.

The two main experimental approaches were undertaken to investigate soil structure changes and assess the potential of machine learning to classify different soil conditions. In both cases, CNNs were developed to explore whether machine learning could uncover meaningful patterns within the soil structural data and potentially improve the efficiency of current soil analysis workflows.

The first approach focused on soil structural dynamics during wetting and drying cycles. Image data for this study were obtained from the experiment conducted by Tracy *et al.* (2015). Soil cores were collected from two locations, one representing a predominantly sandy soil and the other a clay-based soil. Cores measuring 10mm in height and 10mm in diameter were scanned at a resolution of 10μm. Loose soil was also collected for the analysis of basic soil properties. During the experiment, soils were submerged in water and later scanned as they dried under controlled pressures within a vacuum chamber, allowing for precise management of water content. A CNN was trained and applied to determine whether machine learning could accurately differentiate between the changing structures of sand and clay soils at various water contents, and whether it could reliably identify the same soil type across different moisture states.

The second approach examined structural changes resulting from different tillage treatments. Soil cores were collected from an experimental field site that had been managed under zero-tillage and conventional tillage practices since 2014. Cores were 300mm in height and 50mm in diameter, scanned at a resolution of 40μm. Structural comparisons were initially conducted manually through visual and quantitative assessment. Subsequently, a CNN was employed to investigate whether machine learning could distinguish the soil structure resulting from the two different tillage practices at similar time points, and the analyse how the structure of a given soil evolved over time under a specific treatment.

Together, these studies provide insight into the capability of machine learning techniques, particularly CNNs, to automate and develop the characterisation and study of soil structural dynamics under varying environmental and management conditions.

1.4 Aims and Objectives

The overarching aim of this research was to see if machine learning was well suited to predicting structural changes in soil. This aim was addressed in the following manner:

- How is machine learning being used in soil science, and how might it be applied to X-ray CT studies of soil structure?
    1. To learn how machine learning is currently being used in soil science and similar areas.
    2. To observe how machine learning can be used alongside X-ray CT.
- How is artificial intelligence being used, and how might it develop in the future?
    1. To understand how artificial intelligence is currently applied.
    2. To study how artificial intelligence is constantly being improved and used in new ways.
    3. To summarise what the potential future of artificial intelligence will be in the context of soil science.
- Can machine learning predict the soil structure independent of soil texture? Aims:
    1. To develop a convolutional neural network (CNN) for the classification of sand and clay soils.
    2. To evaluate whether the CNN exhibits higher classification accuracy and superior performance for sand or clay soils.
    3. To accurately predict soil structure using a CNN for soils at different water contents.
- Do soil structural dynamics vary over a growing season under conventional and zero tillage? Aims:
    1. To take samples directly after cultivation, and throughout the growing season.
    2. To compare samples from the cultivated soil to the zero-till soil.
    3. To quantify soil structure and pore dynamics for each sample.
    4. To compare how structure changes over a growing season in tilled and zero-tilled soil.
- Can machine learning identify structural changes to tilled and zero-tilled soil over a growing season?

1. To develop a CNN to observe structural changes in tilled and zero-tilled soil over time.
2. To test if the network works better for certain times or cultivation techniques over others.
3. To accurately predict soil structure using a CNN for soils at different points of the year.
4. To accurately predict soil structure using a CNN for soils with different cultivations.

1.5 Thesis Outline.

The rest of the thesis is structured as follows.

Chapter 2 gives an overview of current literature, covering fundamental concepts in soil science, soil structure, and imaging techniques, with a particular focus on the use of X-ray CT. It also introduces the current landscape of image analysis and machine learning applications in environmental sciences.

Chapter 3 provides a focused discussion on artificial intelligence, particularly the rapid evolution and expanding influence of machine learning and deep learning. Given the pace of change and its growing impact across research disciplines and industries, this chapter provides a deeper insight into AI developments, applications and implications.

Chapter 4 addresses the quantification of soil structural changes under different agricultural management practices. This chapter evaluates how tillage systems impact soil structure over time using conventional image analysis techniques, whilst laying the groundwork for comparison with machine learning-based approaches in the following chapters.

Chapter 5 explores the use of convolutional neural networks to analyse changes in soil structure during wetting and drying cycles. This chapter investigates the ability of machine learning models to distinguish between soil types and detect subtle structural transitions, contributing to our understanding of soil response under changing environmental conditions.

Chapter 6 builds upon the findings of Chapters 4 and 5 by applying machine learning techniques to the seasonal tillage dataset. It examines whether automated models can detect differences in soil structure between tillage systems over time, and assesses how these findings compare with manual observations, offering insight into the practical application of AI in soil monitoring.

Chapter 7 concludes the thesis by detailing key findings, discussing their broader implications, and outlining directions for future research. The chapter also considers the environmental and computational impacts of machine learning in scientific research.

2. Developments in machine learning in soil science and its applications alongside X-ray Computed Tomography

2.1 X-ray Computed Tomography and its applications in soil science

2.1.1. What is X-ray Computed Tomography?

X-ray Computed Tomography (X-ray CT), also known as X-ray micro-computed tomography or X-ray microtomography (μCT), is a non-invasive and non-destructive imaging technique. The method of CT was, in part, developed by Godfrey Hounsfield for medical diagnoses (Hounsfield 1980). A conventional X-ray scan produces a two-dimensional image, cannot distinguish between soft tissues and does not allow a quantitative measure of the densities of substances through which the X-rays have passed (Hounsfield 1980). CT, however, takes a series of continuous cross sections to obtain 3-D information about the sample – including the valuable information of soft tissues and X-ray absorption of different substances (Hounsfield 1980).

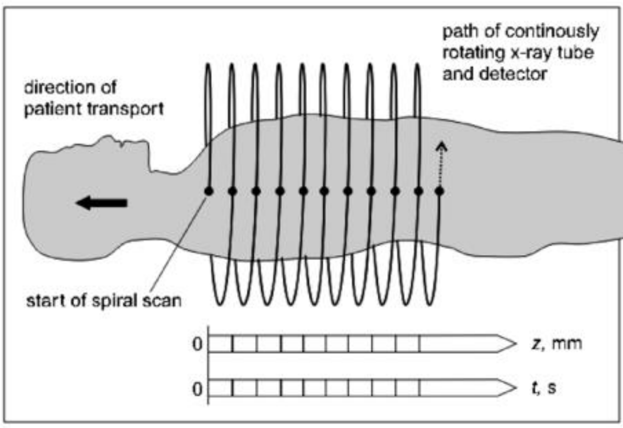

Fig. 2.1: Transportation of the sample (patient) through a continuously rotating X-ray tube and detector.
Reproduced from Kalender 2006, "X-ray computed tomography", *Physics in Medicine and Biology,* Vol 51, Issue 13, pp. 29-43, DOI: 10.1088/0031-9155/51/13/R03.

Originally developed and used for medical applications, including the first clinical CT images being used to detect a cystic frontal lobe tumour (Kalender 2006), this technique has been applied in a range of other fields, including materials science for analysing composite materials (Garcea *et al.* 2018), industrial inspection for detecting flaws in manufactured components (Szabo *et al.* 2017), and archaeological research for non-destructive examination of artefacts (Applbaum *et al.* 2005). Micro-CT scanning, a high-resolution variant of CT, has even been used to study internal organs within beetles(Li *et al.* 2018).

The 3-D data produced by CT can be extended into four dimensions (4-D) by incorporating changes over time or motion into the analysis. This enables the observation of processes such as fluid movement, structural changes, and biological growth within a sample. Favoured for its non-invasive and non-destructive capabilities, X-ray CT has been applied to the study of soil since the early 1980s (Petrovic *et al.* 1982), offering a novel insight into the complex, heterogeneous structures of soil. By providing a detailed, 3-D characterisation of pore spaces, aggregation and compaction, CT scanning has become a crucial tool in advancing soil science, particularly for studying soil-water interactions, root growth, and the impacts of agricultural management practices.

2.1.2 Soil Structure and the use of X-ray CT

Soil structure is used to describe the arrangement and sizes of soil aggregates and the pore space between and within them. (Dexter 1988) gives the following definitions of soil quality based on structure. A 'good' soil structure helps to promote soil hydraulic properties, allows free movement of air and roots, and provides an appropriate environment to support microbes, fungi and soil fauna. A 'poor' soil structure usually implies reduced movement of air and water through the soil, leading to a greatly reduced or less beneficial population of microorganisms. However, there is no universally accepted method to measure and quantify soil structure, with many previous efforts limited to qualitative observations.

Soil structure can be assessed through both indirect and direct methods, depending on the aspect of structure being examined and the level of precision required. Indirect qualification methods include determining water flow through a soil sample (Wilson *et al.* 1988, Dexter *et al.* 2001), where hydraulic conductivity or infiltration rates are measured to infer the size and connectivity of pore networks. Another indirect approach is the application of mechanical stress to a soil sample to observe the fragmentation pattern (Sandri *et al.* 1998), providing insights into aggregate stability and soil cohesion under external forces.

Direct observations involve preparing soil samples through impregnation with resin or wax to preserve the structure, allowing thin sections to be produced for imaging and analysis. Image analysis techniques can then be applied to quantify the size, shape, and spatial arrangement of aggregates and pores (Dexter 1988). While these methods offer detailed structural information, they are labour-intensive and require specific laboratory facilities.

For rapid in-field assessments, visual scoring systems such as the Visual Evaluation of Soil Structure (VESS) have been developed. Based on the Peerklamp test (Peerlkamp 1959), the VESS system provides a simple, easily taught and referenced framework for evaluating soil structure quality directly in the field (Ball *et al.* 2007). Assessments are made by examining the size, shape and stability of aggregates when a soil block is broken open, allowing farmers and researchers to monitor soil health without the need for specialised equipment.

The main downside to these techniques is that they in some way alter or disrupt the sample and could therefore influence the observed results. X-ray CT allows access to the internal 3-D structure of a soil column without damaging or disrupting the sample. This means that information on the structural parameters of the soil pore network can be obtained (Katuwal *et al.* 2015) including the pore network connectivity – a crucial property controlling many soil functions (Vogel *et al.* 2010, Schlüter *et al.* 2011, Dal Ferro *et al.* 2013). Knowledge from these X-ray CT observations of soils have highly useful applications in tillage and non-tillage experiments and management practices (Schlüter *et al.* 2011, Ferreira *et al.* 2018, Soto-Gomez *et al.* 2018).

2.1.3. Soil Hydrology

Being able to trace particle movement allows for prediction of particle movement through soils. This is potentially very useful when it comes to forecasting impacts of environmental contamination or calculating amount of fertiliser to apply to agricultural soil. Movement of water, and subsequently dissolved substances, is strongly reliant upon the pore network of a soil (Soto-Gomez *et al.* 2018). Soil pore networks are developed over time through earthworm burrows, root channels and wet-drying cycles. These macropores have been shown to allow movement of organic matter and microorganisms for several metres through soil (Cannell 1985, McDowellboyer *et al.* 1986).

X-ray CT images were used to visualise water movement within soil, highlighting the preferential flow paths pre and post watering (Mooney 2002). They have also been used to calculate pore diameter and tortuosity to determine soil air permeability (Naveed *et al.* 2013). The non-invasive and non-destructive nature of X-ray CT was used to great benefit by continuously monitoring a sample to quantify soil structure and preferential flow (Luo *et al.* 2008).

Traditional methods of observing and measuring transport in soil often make use of tracer materials (Gitis *et al.* 2002, Paradelo *et al.* 2013, Soto-Gomez *et al.* 2016), or measuring soil features such as soil structure, hydraulic conductivity and air permeability (Kjaergaard *et al.* 2004). These can be time consuming and depending on the tracer used, be expensive. They can also be destructive to the soil sample, eliminating chances of repeatedly making observations from a single sample. However, combining tracer materials with X-ray CT imaging proved to be successful in medical and veterinary research (Gunda *et al.* 2009). This concept was then applied to tracing particle movement in soil; by tracking movement of a tracer material used to represent to common fungicide compounds (Grayling *et al.* 2018). X-ray CT was used to map the 3-D pore geometry of a soil sample. When combined with a tracer material, this information was used to successfully track the movement of particles through the sample as shown in Figure 2.1 (Grayling *et al.* 2018). Using X-ray CT allowed the movement of the tracer material and the proxy agrochemical to be tracked over several hours. Visualisation of the progress and processes by which agrochemicals are moving through the soil in real-time can help when making agricultural management decisions, as well as improving the development of new agrochemicals. The knowledge of how particles, such as agrochemicals, move through soils means that optimum quantity of these chemicals could be calculated before application, potentially leading to a reduction in run-off of excess chemicals.

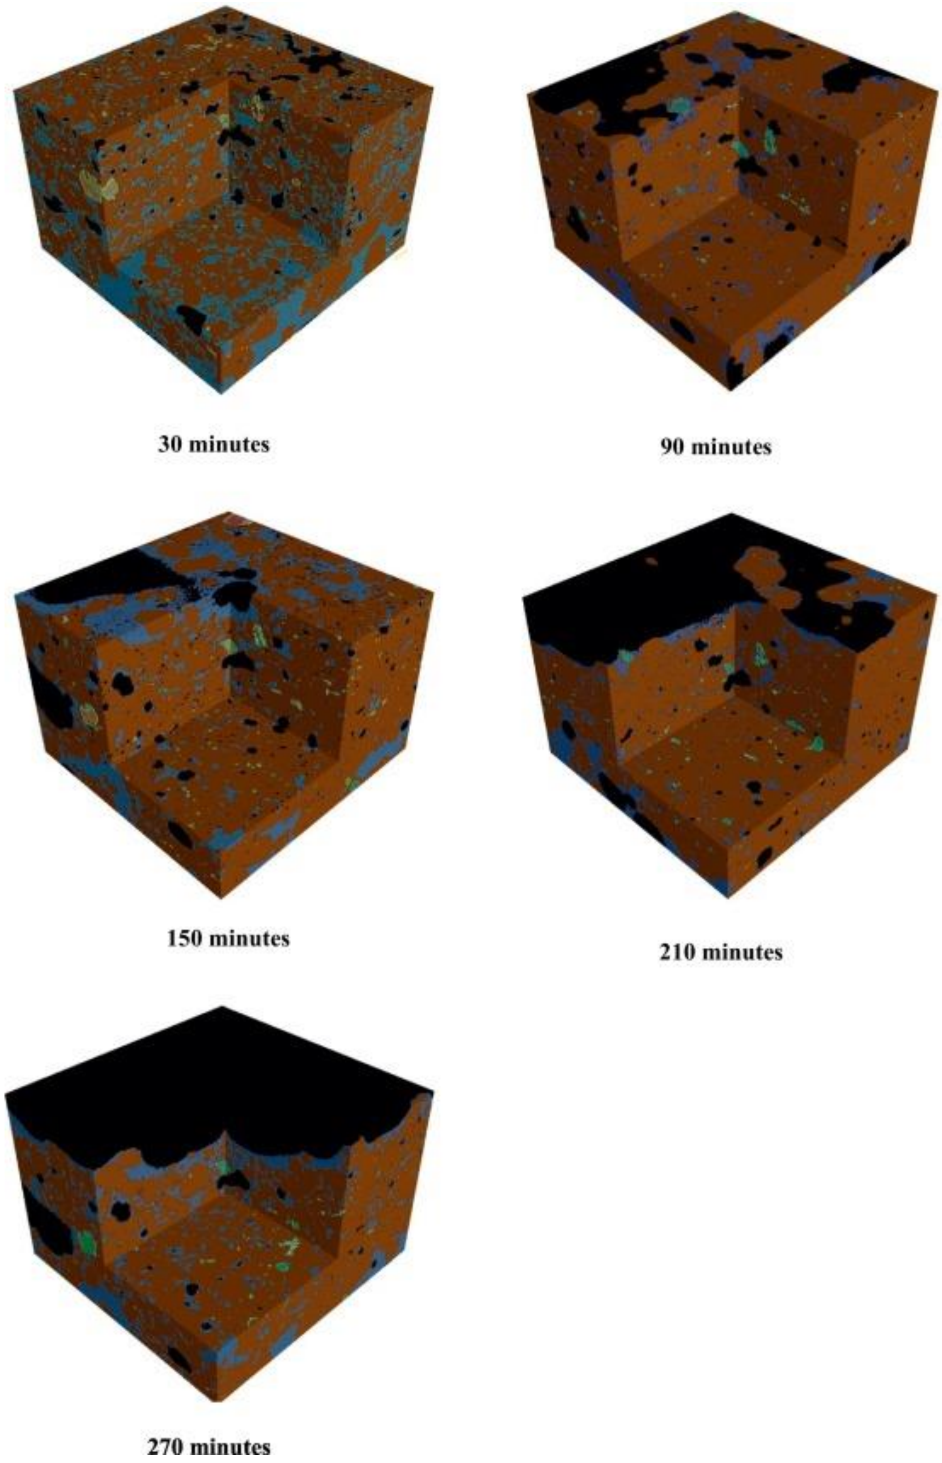

Fig. 2.2: 3-D visualisations of the distribution of tracer filled pores over time, where black is air-filled pore space, blue is tracer filled pore space, pale green is a proxy for agrochemicals and brown is soil. Reprinted with permission from Geoderma, Vol. 321, Grayling et al., *The application of X-ray micro Computed Tomography imaging for tracing particle movement in soil*, pp. 8-14, Copyright (2018), with permission from Elsevier.

One of the more well-known and easily observable parts of the wetting-drying cycle on soils is crack dynamics. Soil cracks occur due to soil deformation and pore network shrinking, resulting from evaporative water loss (Tang *et al.*

2011). Crack dynamics are mainly due to and depend on the clay content of a soil, as well as soil mineralogy. Water flow and solute transport will bypass the soil matrix and transition through the cracks instead, preventing water and nutrients from being absorbed completely into the soil matrix, and therefore impacting agricultural applications and their effectiveness (Allaire *et al.* 2009).

Diel *et al.* (2019) examined three German soils and the impact of wetting and drying on the soil structure. X-ray CT was used to observe the crack dynamics at different points in the wetting-drying cycle without destroying or interrupting the samples. Samples were compacted to two bulk densities and wetted (W1) and left to dry out (D1). The cycle was repeated with W2 and D2 respectively. These samples were scanned at each of these 4 stages (see Figure 2.3). Garnet particles were added to the samples to measure soil structure dynamics.

As well as providing visual data, as shown in Figure 2.3, the images could also be used to calculate pore size distribution, macroporosity, soil matrix density and the distance between pores.

Some of the results observed through X-ray CT would not have been possible to do with traditional methods. They concluded that the extent of crack dynamics from the wetting-drying cycle was dependent on a range of soil properties, including soil organic matter content, clay content and bulk density.

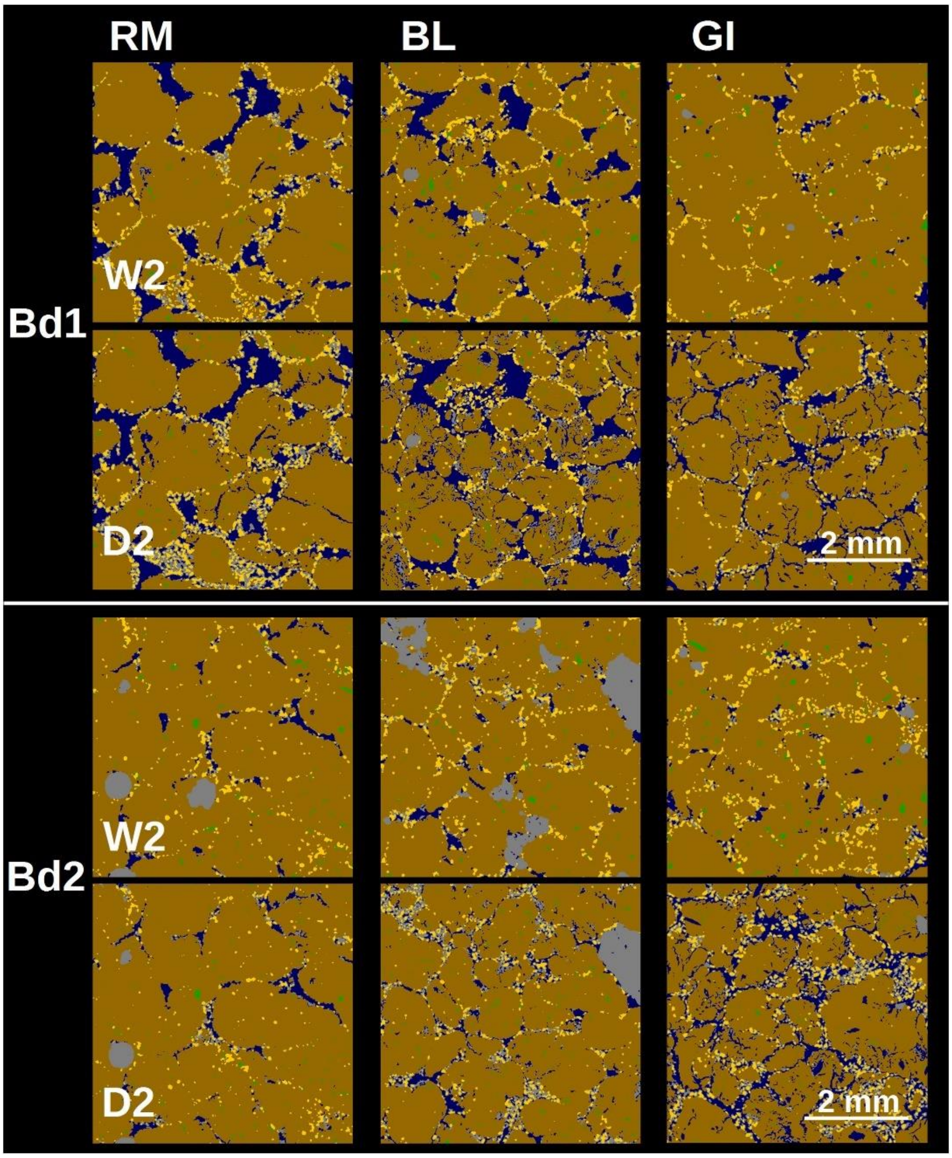

Fig. 2.3: Segmented X-ray CT images for two bulk densities (Bd1 and Bd2) of soil in a wet stage (W2) and dry stage (D2) of three different soil samples from three locations (RM, Bl, GI). For soil descriptions see Table 2.1. Brown: Soil, Blue: Air, Green: Occluded pores, Yellow: Garnet particles, Grey: Unclassified/masked. (Diel *et al.* 2019)

2.1.4. Roots

Roots play a crucial role in plant development, providing structural support, water and nutrient uptake, and interactions with soil microorganisms. Understanding root growth and behaviour under different environmental conditions is therefore important for advancing agricultural practices, soil management and ecological research. However, studying roots presents challenges due to their hidden position within opaque soil systems.

Traditional methods of examining roots involved either the destructive removal of the roots from their environment (Trachsel *et al.* 2011) or growth in artificial media (Downie *et al.* 2012). Removal may damage the sample and also prevents the roots from being observed over time, whilst artificial media may alter natural growth (Mairhofer *et al.* 2016). The use of X-ray CT offers a powerful alternative, allowing root systems to be non-invasively observed at multiple stages of development, thereby enabling the detailed study of root responses to varying environmental factors.

One of the simplest approaches to segmenting roots or pore space in CT images is the application of threshold-based segmentation. In this method, voxel intensities are classified into different phrases (e.g. soil, roots, pores) based on their greyscale values (Tuller *et al.* 2013). Voxels, or volumetric pixels, represent small cubic elements of 3-D image data, each with an associated greyscale intensity corresponding to the material density at that location. Voxels with intensity values above or below a certain threshold are assigned to the target structure (e.g., roots or pores), while others are assigned to background. This approach is considered to be straightforward and computationally efficient.

A slightly more advanced form of this is known as region growing. Starting from a user-defined start point within the target structure, the algorithm gradually adds neighbouring voxels that meet specific similarity criteria, typically based on greyscale intensity (Schlüter *et al.* 2010). This method is particularly useful for tracking continuous structures such as roots, as it can follow their shape and branching patterns through the 3-D volume. Region growing can be sensitive to noise and may require careful tuning of growth parameters to prevent expansion into adjacent non-root or non-pore regions.

However, both of these approaches face significant limitations, particularly when greyscale ranges overlap between phases. The presence of water within roots or soil pores can cause greyscale values to shift, complicating the separation of root and soil structures. As a result, simple and region-growing thresholding often struggles to accurately distinguish between phases in variable or moisture-variable soils, making this approach to automatic root tracking challenging.

Automatic root tracking refers to the use of image analysis algorithms to follow the shape and position of roots through a series of 3-D images without requiring manual tracing by the researcher. Software such as *RooTrak* achieve this by using greyscale values to segment and trace root systems throughout the 3-D volume (Mairhofer *et al.* 2012). Whilst *RooTrak* is able to recognise emerging lateral roots as being associated with the original root system, it has issues identifying root systems of different plants when they come into contact with each other. This was addressed by Mairhofer *et al.* (2015). The authors refined the original *RooTrak* algorithm by using shape information and multi-level set tracking. The shape of each root cross-section was recorded and aligned using the Iterative Closest Point algorithm during collisions, ensuring each root system remains distinct even when roots are in contact.

Being able to visualise the root system architecture (RSA) without destroying or damaging the sample allows for studies into the response of RSA to factors such as soil texture (Rogers *et al.* 2016) and soil compaction (Tracy *et al.* 2012), and these responses can be tracked through samples over time.

More recently, machine learning approaches – in particular CNNs, have been applied to root and pore segmentation tasks. Architectures such as U-Net are well-suited to soil imaging tasks due to their ability to learn complex features from relatively small datasets. CNN-based methods can capture subtle textural and structural patterns that traditional thresholding or region-growing techniques may miss, making them highly suitable in cases where root architectures are complex or CT images are noisy. Additionally, machine learning models can generalise across varying soil types and scanning conditions if adequately trained, offering a promising path towards more robust automated segmentation in soil science research.

2.2. Machine Learning

2.2.1 What is machine learning?

Machine learning allows automated measurement, counting and analysis of data, allowing high-throughput processing of data. Automated image analysis technologies make use of non-destructive imaging methods to capture information over time, enabling the monitoring of living organisms throughout their development. In plant science, image-based analysis has become particularly popular, with sensors covering a wide range of wavelengths, including visible to far-infrared (Lowe *et al.* 2017), thermal and fluorescence (Chaerle *et al.* 2004, Chaerle *et al.* 2007), and conventional trichromatic (RGB) imaging (Hernandez-Rabadan *et al.* 2014).

Traditional machine learning approaches applied to image analysis rely heavily on feature engineering, the manual design and selection of specific measurable properties from the raw data to help the model distinguish between different classes of objects (Pound *et al.* 2017). Examples of such features include edges, textures, colours or shapes extracted from regions of interest. This process is necessary because raw images contain vast amounts of pixel data, making direct analysis computationally impracticable and inefficient (Singh *et al.* 2016). After feature extraction, the reduced dataset is input into machine learning classifiers such as support vector machines (SVMs), decision trees, or random forests, which learn to separate classes based on the selected features (Pound *et al.* 2017).

While this approach is effective in certain tasks, traditional machine learning has limitations. Manual feature engineering is labour-intensive and often requires expert knowledge, which can introduce biases and restrict the model's applicability to specific tasks. Additionally, hand-crafted features may fail to capture complex patterns that are essential for accurate classification, particularly in varied datasets.

To overcome these limitations and advance the capabilities of machine learning, deep learning approaches have been developed. Artificial neural networks (ANNs), inspired by the structure and function of the human brain, offer a way to automatically learn relevant features from raw data (Salu 1985, Anderson 1988, Heskes *et al.* 1991). Rather than relying on manually engineered inputs, ANNs can learn classified representations of data, progressively identifying simple features at early layers (e.g., edges or textures) and more complex patterns at deeper levels.

Convolutional neural networks (CNNs) take this concept and extend if further and are well suited to image analysis tasks. CNNs automatically extract spatial features from images using convolutional layers, pooling layers, and non-linear activation functions, allowing for the efficient processing of high-dimensional image data. By eliminating the need for manual feature engineering, CNNs enable the direct use of raw image pixels as input, improving model generality and performance across a wide range of tasks such as object recognition, segmentation and classification (Wang *et al.* 2021).

2.2.2 Deep learning

Deep learning is a branch of machine learning that involves the use of neural networks with many layers, allowing models to learn increasingly complex and abstract representations of data. The concept of deep learning was first developed around 2006 and has rapidly advanced with improvements in computational power, algorithms and the availability of large datasets (Schmidhuber 2015).

Within the context of image analysis and processing, deep learning models automatically learn to detect important features within images, such as edges, textures, shapes or objects, without the need for manual feature engineering. Image analysis tasks typically include object classification (assigning images to categories e.g., distinguishing between sandy and clay soils), object detection (identifying and locating objects within an image e.g., detecting root tips withing a soil volume), segmentation (dividing an image into meaningful regions e.g., separating root systems from surrounding soil in a CT scan), and image generation (creating new images based on learned patterns e.g., generating synthetic images of soil pore networks or root structures for use in model training).

Different types of neural networks are suited to different tasks. Fully connected neural networks (FCNs) are typically used for structured or tabular data, whereas CNNs are designed specifically for image-based data. CNNs apply convolutional filters across an image to detect spatial order of features, from low-level patterns such as edges and boundaries to high-level structures like entire objects. This makes CNNs highly effective in computer vision tasks, as they preserve the spatial relationship between pixels, unlike standard ANNs where each input pixel is treated independently (Habibi *et al.* 2017).

Deep neural networks (DNNs), including CNNs, are particularly powerful in the field of computer vision due to their ability to handle the complexity and high dimensionality found within image data. Their depth allows them to model highly intricate relationships, making them suitable for applications such as object detection, image segmentation, and image generation (Zhong *et al.* 2019). Beyond image analysis, deep learning has been successfully applied in natural language processing, speech recognition, investment decision-making, and autonomous, driving systems, where the ability to extract and integrate complex patterns from large amounts of raw data whilst retaining context is crucial (Zhong *et al.* 2019).

2.3 Image analysis

2.3.1 Observing changes over time

X-ray CT images can be taken of a sample at various time points, allowing for observation of any changes to the sample through time. Automated image analysis becomes an extremely powerful tool at this point, with the high throughput nature allowing for many images of many plants to be observed. The ability to observe seed imbibition, a process that usually occurs unobserved in soil, is of great interest (Belin *et al.* 2018). By observing plants regularly over time, it can also be possible to detect early signs of stress or plant diseases (Lowe *et al.* 2017, Czedik-Eysenberg *et al.* 2018). Early detection of these biotic and abiotic stresses can allow treatment before any major damage to the plant or yield occurs.

X-ray imaging of soil over time allows for the detailed observation of dynamic processes within the soil matrix. For example, changes in soil saturation can be monitored by tracking the movement and distribution of water through pore spaces. Beyond moisture dynamics, structural changes such as aggregate swelling, shrinking, cracking and the formation of collapse of pore networks can also be quantified. Observeing these processes enables researchers to assess how different soil types respond to environmental changes, providing insights into soil stability, hydraulic behaviour, and mechanical properties under varying conditions.

2.3.2 Image analysis of X-ray CT images

X-ray CT provides a powerful, non-invasive method for obtaining detailed 3-D representations of internal structures. However, the sheer volume of image data generated presents a significant challenge for manual analysis, which can be time-consuming, labour-intensive and prone to variability between observers.

Quantitative image analysis techniques have therefore become increasingly important, enabling automated measurement, detection, and interpretation of features within CT scans. Automated methods reduce the workload for researchers and clinicians, increase reproducibility, and allow for high-throughput analysis of large datasets. Survey studies have indicated a growing preference among medical imaging professionals for the adoption of automated and computer-assisted analysis tools (Gajawelli *et al.* 2019).

Within clinical contexts, automated CT image analysis has been successfully applied to a wide range of tasks, including tumour detection, fracture identification, and disease classification. These applications have benefited from the development of both traditional machine learning methods, which rely on feature engineering, and more recently, deep learning approaches such as CNNs, which are capable of learning directly from raw image data.

In comparison to the significant advances that have been made in the field of medical imaging, the application of automated image analysis to soil science remains limited. In the context of soil CT imaging, challenges such as variable material properties, overlapping greyscale values based on moisture content, and the complex structure of roots and pore networks make automated analysis difficult. Notable examples of progress in this area include the development of *RooTrak's* automated tool for tracking root systems through soil volumes (Mairhofer *et al.* 2012).

More recently, deep learning has been used to show promising results for the structural analysis of soil CT images, with reduced manual annotation. In this approach, models were trained to classify soil structural features, and a gradient-based class activation map was used to confirm that the networks had learned meaningful representations of soil samples (Wieland *et al.* 2021). However, this study also highlighted that model performance varied between different soil types, with some samples fitting well and others requiring additional models or training, indicating that soil heterogeneity remains a challenge for generalisation within machine learning algorithms. These developments still highlight the growing potential for machine learning techniques to be used in soil image analysis.

2.4 Machine learning and environmental applications

2.4.1 Applications of machine learning in environmental sciences

The main benefit of using machine learning in any situation is that it removes or greatly reduces the manual work required. With this human error is reduced, and as long as the computer is running efficiently, a higher amount of work can be carried out. A high-throughput capability of a system means that a large number of plants i.e., in a commercial field or greenhouse can be screened regularly for early signs of stresses or disease.

Patterns in data that would typically be missed by human observers can be reliably identified by machine learning algorithms. Topics such as climate change have a vast amount of data collected over many years, with some data going back a few hundred years or further, in a wide range of variables. Machine learning can be applied to great benefit here by looking at all this data and spotting potentially previously unnoticed trends.

Liu *et al.* (2018) investigated using transfer learning to estimate soil clay content using hyperspectral data in combination with large-scale soil spectral libraries. In this study, a one-dimensional CNN was trained first on Land Use/Land Cover Area Frame Survey (LUCAS) mineral soil spectra, then fine-tuned using LUCAS organic soil and field samples derived from hyperspectral imagery. This use of transfer learning, where a model is trained on one dataset and then adapted to a related but distinct dataset, enabled the authors to improve model performance across varying soil types and sensor sources. Their results were very promising, with a reasonable accuracy (R^2) of 0.601. The authors proposed further work involving different regions and the use of expanded datasets to enhance the model's generalisability.

2.4.2 Current uses of machine learning in soil science

Machine learning is being used in a number of places to predict unknown soil properties based on known data. Sarmadian *et al.* (2009) used easily measurable properties such as clay, sand, silt, and organic carbon to predict bulk density and water saturation percentage, among other properties. Li *et al.* (2014) used known physical properties of the soil, such as content of organic matter and total nitrogen, to predict soil nutrient levels. This was done using support vector machines and a general regression neural network, a type of ANN specifically designed for function approximation and regression tasks.. Moreira de Melo *et al.* (2015) used a neural network to predict the water retention curve of soils. Their results were successful, and sensitivity analysis showed that the parameters required for predicting water contents varied based on the moisture conditions of the soils in question. This additional data means that inputs can be adjusted to suit the situation, in order to get better results out of the networks at the end.

As mentioned previously, the issue of automatically separating roots from soil in X-ray CT images is of high importance in soil science. Manually doing this can take many hours per scan, and some lateral roots may be missed. Douarre *et al.* (2018) demonstrated the use of transfer learning to teach the network how to identify root from soil and follow this along to efficiently segment root from soil. The training was done using synthetic data, but the final network was capable of performing segmentation on real data.

Transfer learning was successfully used to estimate soil clay content of hyperspectral soil data (Liu *et al.* 2018). Soil spectroscopy allows for rapid and non-destructive analysis of soil properties. Combining this with transfer learning allows for high-throughput processing of data, meaning whole fields can be analysed in a short amount of time. Although this experiment was only tested on a limited area and the experimental design was limited, further investigation into this could yield great results.

SoilJ is a plugin for the imaging software ImageJ, which aims at automatically processing cylindrical soil columns for X-ray CT data sets (Koestel 2018). This plugin is able to recognise the column outline, segment the image, extract any roots and detect the soil surface topography, all automatically. SoilJ reduces the time and expert knowledge required to evaluate X-ray CT images.

With technological advances meaning that smartphones now have high quality digital cameras built-in, Aitkenhead *et al.* (2016) worked on using smartphone imagery to automatically assess soil parameters. Images were taken using an iPhone 2 with a colour correction card placed in each shot to allow normalisation and corrections between photos later. The phone's GPS was also used to capture the site location. An artificial neural network was used to estimate the soil structure and textures. Whilst the approach was not able to obtain an accuracy good enough to replace laboratory or in-field expert analysis, this could be improved with further work.

Recent years have seen a huge increase in the development of and use of machine learning. This has resulted in huge technological developments in a large range of industries, including soil science. This includes the development of machine learning models to classify and predict a suitable crop based on soil nutrition levels (Swathi *et al.* 2023) as well as integrating machine learning in soil nutrition analysis to improve crop productivity and yield, as well as sustainable soil practices (Venkateswara Reddy *et al.* 2024).

2.4.3 Current applications of neural networks and soil structure

There is currently not much research in the area of how neural networks could be used to quantify or analyse soil structure. This is especially true when considering the use of neural networks with X-ray CT images of soils to look at soil structure.

Farfani *et al.* (2015) used Wilson (1998)'s data to design a network to analyse soil structure and determine suitability of a soil for building on. Artificial intelligence has been increasingly applied across geotechnical engineering due to their ability to model complex nonlinear relationships and deal with uncertain material behaviour (Baghbani *et al.* 2022). Neural networks have been used to estimate the shear strength of soil for road construction in Vietnam (Tien Bui *et al.* 2019), predict tunnel deformation for an underground excavation project (Lai *et al.* 2016) and assess settlement caused by shield tunnelling (Ye *et al.* 2022). A recent review of over 1,200 studies found ANNs to be the most widely applied AI technique in geotechnical engineering, with prominent applications in slope stability, tunnelling, foundation design, and unsaturated soil modelling (Baghbani *et al.* 2022).

Whilst there is currently little on the use of this sort of work on soils, crack and fracture detection in bones and steel have been used in clinical and engineering environments. Lindsey *et al.* (2018) used a deep learning model to detect fractures on radiograph images. The results showed similar a diagnostic accuracy to trained orthopaedic professionals, and when used alongside medical staff, can improve the detection of wrist fractures. Networks have also been used to process images to detect cracks in concrete, allowing essential works to be completed on infrastructure before any serious damage or risk occurs (Cha *et al.* 2017). Wang *et al.* (2018) demonstrated crack detection and analysis on real-time CT scanning of soil samples under compression. This work shows how cracks initially formed and developed in response to the stress applied.

2.4.4 Current issues of machine learning and soil structure applications

The main issues found with machine learning approaches to any problem is that if the data being tested is different to the training data, the network can struggle to make accurate predictions. With soil, this means if the network had been trained on data from a certain location or soil type and then used on soil from a different location, results could have a very low accuracy. This is discussed in Aitkenhead *et al.* (2016), where smartphone photography was used to assess soil properties, and in Zhai *et al.* (2006), where remote sensing data was used to classify soil textures. Making a network capable of being flexible enough to analyse a variety of soil types or soils from different locations is a key obstacle to overcome.

3. Artificial Intelligence: Recent advancements and potential future applications

3.1 Introduction

3.1.1 How has artificial intelligence (AI) advanced in the last 5 years?

In recent years, AI has made significant progress in several areas, including natural language processing (NLP), computer vision, and reinforcement learning. NLP models, such as BERT, GPT-3 and GPT-4 have achieved high-level results on a range of language tasks, including the very popular ChatGPT (Zheng *et al.* 2021, Gubelmann 2023). These NLP models can be used to improve language translators, improve voice assistants such as Siri and Alexa, and data processing to extract important information from medical and legal documents (Sezgin *et al.* 2023). Computer vision, including tasks similar to those mentioned within the rest of this thesis, have used convolutional neural networks (CNNs) and other deep learning techniques to improve image and video recognition and generation (Jiao *et al.* 2019, Lai 2019, Vayadande *et al.* 2023). CNNs are a class of deep learning models specifically designed to process grid-like data such as images. They use convolutional layers to automatically detect patterns such as edges, textures, and shapes, allowing them to identify and extract spatial orders of features directly from raw pixel data. Reinforcement learning is a machine learning technique where the program isn't taught what to do, but over time discovers which actions return the best result in the situation (Sutton *et al.* 2018). Reinforcement learning's most recognisable application is in training computer programs to play games (such as chess) to a higher level than the world's top human players. It is also used in engineering, particularly in the development of autonomous vehicle systems, where reinforcement learning algorithms enable cars to make decisions about navigation, obstacle avoidance, and adaptive driving strategies based on real-time feedback from their environment. AI is used in a wide variety of industries and applications, including healthcare, finance, customer service, marketing, and scientific research.

In healthcare, AI is utilised in assisting medical professionals with detecting abnormalities in medical images such as X-rays, MRIs, CT scans and mammograms (Esmaeilzadeh 2020). By automating this process, test results can be produced much quicker, allowing for swifter action to be taken if required. Using AI can cut down tumour classification time to roughly 3 minutes (Tadiboina 2022), when a tumour specimen is analysed, and this can even be performed safely in the surgery room straight after the specimen is extracted from the patient.

NLPs can be used to extract information from documents, providing summarised reports as well as analysis of this information. This could be of great use in legal work (Mowbray *et al.* 2020) as well as fraud detection for transactions and banks (Cao 2022).

Artificial intelligence is now embedded in a wide range of everyday and industrial applications, many of which may go unnoticed in daily life. Virtual personal assistants such as Siri, Alexa, and Google Assistant rely on NLP to interpret commands, answer questions and control connected devices. Streaming platforms like Netflix, Spotify, and YouTube use machine learning algorithms to generate personalised recommendations based on user activity. Search engines use AI to interpret queries and prioritise relevant results as well as summarise search results. Email providers apply AI to filter spam by analysing message content, sender behaviour, and engagement history.

Outside of these familiar examples, AI is used extensively across sectors such as retail, healthcare, and finance. Online shops make use of AI for product recommendation, dynamic pricing, and customer support chatbots (Pillarisetty *et al.* 2022). In finance, AI supports fraud detection, credit scoring, and high-frequency trading strategies. Healthcare applications include diagnostic support, such as tumour identification, as well as predictive modelling for patient outcomes and drug development through large-scale data analysis.

In transportation and logistics, AI helps to optimise delivery routes, monitor vehicle performance, and support the development of autonomous vehicles with real-time decision making. Agriculture also benefits from AI, with tools for crop monitoring, automated irrigation, pest detection, and yield prediction (Sharma 2021). Manufacturing industries use AI for predictive maintenance, quality control through image recognition, and managing production lines (Kim *et al.* 2022).

On a broader scale, AI is being used in energy systems for demand forecasting and smart grid management, and in smart cities is used for analysing traffic flow, improving waste collection, and monitoring environmental conditions (Allam *et al.* 2019, Herath *et al.* 2022). Social media platforms rely on AI to recommend content, detect hate speech, and moderate user-generated material. Education platforms use AI for adaptive learning, automated marking, and providing students with targeted feedback (Zhai *et al.* 2021, Holmes *et al.* 2022). AI is also increasingly used in creative industries, generating text, music, and artwork, and even supporting virtual environment design in gaming and film production.

3.1.2 The advancements of computer vision

Computer vision, in particular, has advanced greatly in recent years, with several interesting developments. Much of this progress has been driven by the increasing availability of large datasets, as well as growing computational power, which has improved processing capabilities and made it possible to train more powerful networks. A lot of these breakthroughs have been made possible by CNNs, which are particularly effective at analysing visual data. CNNs work by passing images through multiple layers of filters that detect features like edges, shapes, and textures. These layers build on each other, allowing the network to learn increasingly complex visual patterns, making

CNNs well suited to tasks like image classification, object detection and segmentation.

Object detection models such as You Only Look Once (YOLO)(Redmon *et al.* 2016) and RetinaNet (Lin *et al.* 2017) models are built on CNNs and are used to identify and locate objects within images or video. This technology is widely used in the development of self-driving cars, where it helps to recognise traffic signs, detect hazards, and track other vehicles (Flores-Calero *et al.* 2024).In agriculture, CNN-based models have also been applied to crop monitoring. A YOLO model was used to detect sweet peppers growing on the plant, count the number of peppers and classify the growth stage (Paul *et al.* 2024). As 3-D computer vision and video recognition continue to develop, these models are increasingly moving beyond 2-D images and into real-world scenarios that involve depth, motion, and potentially changing environments.

3.1.3 Generative Adversarial Networks

Generative Adversarial Networks (GANs) are a class of deep learning models designed to generate new, synthetic data that mimics the patterns found in real datasets. First introduced by Goodfellow *et al.* (2014), GANs have become a widely recognised and one of the most applied generative models in artificial intelligence. They are particularly popular for image generation tasks but are also being explored across a wide range of creative, scientific, and industrial applications.

GAN's work by training two neural networks in competition with one another: a generator and a discriminator. The generator's task is to create synthetic data, typically starting from randomised noise, that resembles the training data as closely as possible. The discriminator is trained to distinguish between real data (from the original training dataset) and the fake data produced by the generator. The two networks are trained simultaneously, with the generator improving its output to fool the discriminator, and the discriminator becoming better at identifying which inputs are real. Over time, this antagonistic process results in a generator that can produce data that is highly realistic (Hong *et al.* 2019).

GANs have allowed for more powerful generative AI approaches, such as diffusion. Models such as DALL-E and programs like Midjourney have demonstrated the ability to use diffusion models to create high-resolution images from simple text prompts, transforming short descriptions into visual representations (Wakefield 2021). These tools are not limited to research environments; they are now widely accessible to the public through user-friendly platforms, making powerful generative capabilities accessible outside of specialist domains. While there are ongoing debates around whether these image outputs can be considered “art” (Coeckelbergh 2023), the fact that such models can take words and turn them into visual concepts is an incredible development. GANs are also used in a range of other areas, including image super-resolution (Bulat 2018, You *et al.* 2023), face generation (Kammoun *et al.* 2022), and the creation of synthetic data for training machine learning models where real data may be limited or sensitive (Vaz *et al.* 2024).

In agricultural and environmental research, GANs have been used to generate synthetic crop images, simulate plant growth under different conditions, and enhance datasets where high-quality labelled data is scarce. In medical imaging, GANs have been applied to create realistic scans for training diagnostic algorithms and to improve resolution in noisy or low-quality images (You *et al.* 2023). While applications in soil science are still emerging, the potential for generating artificial pore structures or root system images for use in training segmentation models is clear.

Despite their usefulness, GANs are not without challenges. Training can often be unstable and can lead to issues such as mode collapse, where the generator demonstrates limited variation in output. The models are also computationally intensive, requiring careful balancing of the generator and discriminator during training. In addition to technical challenges, ethical concerns have been raised regarding the misuse of GANs for creating realistic fake content, such as deepfakes or manipulated images.

Nevertheless, the ability of GANs to learn complex patterns and produce novel data makes them a powerful tool in modern AI. Their continued development is likely to pay an important role not only in creative industries, but also in scientific modelling, data simulation, and visual analysis.

3.1.4 Artificial intelligence and research

Artificial intelligence has become a popular topic for researchers as well. Used across a range of fields, from soil science to food science (Hassoun *et al.* 2024), to the research and development of cosmetics (Elder *et al.* 2024). As seen in Figure 3.1, in the past 5 years there is an increasing number of publications released with terms such as "artificial intelligence", "machine learning" and "deep learning" included in the title, abstract or keywords.

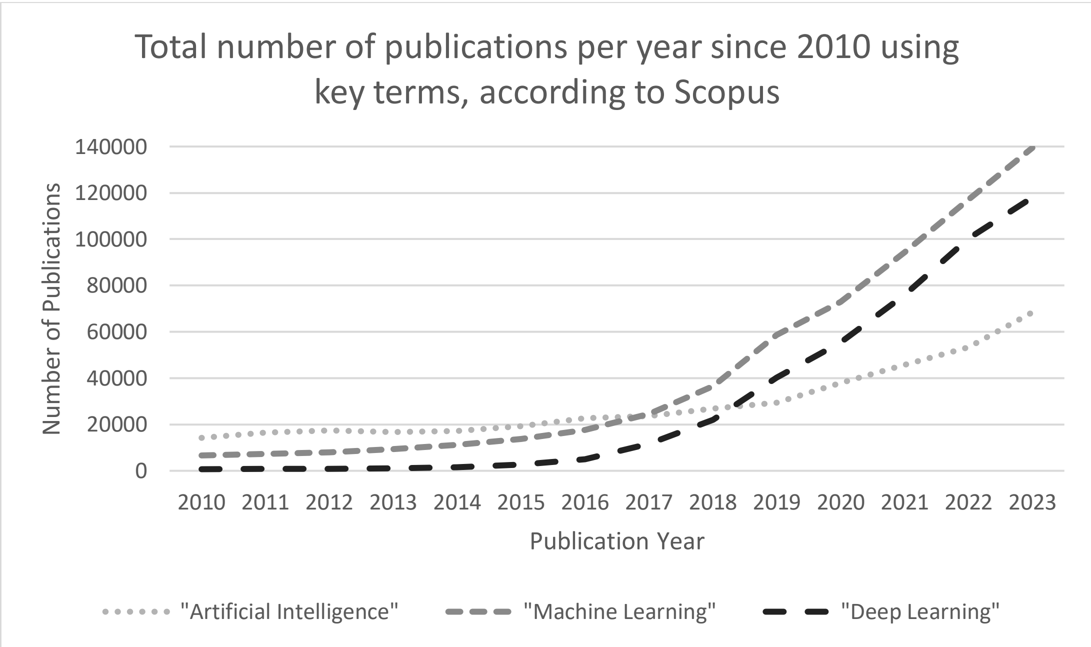

Figure 3.1: Trend showing the increase in the number of publications year on year found on Scopus when searching for several AI related terms.

AI has a large range of uses in research, from identifying trends and patterns that may otherwise go unnoticed by manual processing, to assisting with monotonous tasks, such as cell counting (Guo *et al.* 2019).

As AI develops and reaches new milestones, the impact upon scientific research could be huge if properly utilised. Automation of tasks will not only allow more work to be done in less time, saving many hours of work, and freeing up people to complete other tasks, but also eliminate human error.

3.2 Discussion

A major concern about AI is that it will replace humans in many job roles, from customer service and data-based roles to more creative roles such as story writing and creating artworks.

One very popular text-based AI is ChatGPT (OpenAI 2022), a chatbot that was launched by OpenAI in November 2022. The program uses a language model trained to produce text and was optimised by using Reinforcement Learning with Human Feedback (RLHF). RLHF is a training method that requires human demonstrations to guide the model towards the desired behaviour. The model was trained on large amounts of data written by humans, including conversations, so it can hold a "human-like" discussion.

AI has been popular and common in mainstream media lately (late 2022 into early 2023) due to the emergence of AI-generated art as well as ChatGPT. There are news articles and websites covering the topic of language bots such as ChatGPT being advanced enough to output articles and blog posts, as well as create scripts and captions for social media as well as copy write for products and businesses. Whilst other text-based and language programs do exist, ChatGPT is the most easily accessible and popular model currently available.

There are however several shortcomings to this otherwise impressive tool. Earlier versions of ChatGPT were not connected to the internet, so it was unable to lookup answers in the same way a search engine would. More recent versions now have browsing capabilities, which means they can retrieve up-to-date information and provide source links. This has significantly improved their ability to cite references and attribute information, although accuracy and reliability still depend on the quality of the sources accessed. Another limitation is the model's knowledge base, which remains restricted by the original training data. While internet access allows it to update some information, there are still constraints in how effectively it integrates new knowledge.

One significant concern surrounding the increasing development and uptake of AI in content generation is its potential to replace professional writers, including journalists and freelance authors. With sufficient and appropriate training data, large language models are capable of producing coherent text on a wide range of topics, raising questions about the future role of human authors in industries that rely heavily on written communication.

In academic and scientific contexts, the possibility that AI could be used to draft research papers or scholarly articles introduces similar concerns, particularly with plagiarism and the reliability of generated information. While recent improvements mean that ChatGPT can now provide references more effectively, it is not yet meeting the expected academic standards when it comes to traceability. The model may still return incomplete or incorrect citations, particularly when asked to reference highly specific or less widely accessible materials.

This latter phenomenon is an example of what is referred to as an AI hallucination, a situation in which the model produces a confident and detailed response that is not grounded in its training data. In essence, when the model lacks access to verifiable information, it may generate an answer by combining pieces of related data in a way that appears coherent and credible, despite being inaccurate or entirely fictitious.

It is likely that, in the future, other models or tools will be produced which are able to overcome these limitations.

When asked for Soil Science papers, the model returned the names of several academic journals, which it claimed were commonly associated with the field. This list included the journal name and a summary of what they publish. Below is the output it gave:

“There are many research papers and articles on soil science that have been published in a variety of scientific journals. Here are a few examples of journals that often publish articles on soil science:

- “Soil Science Society of America Journal" publishes articles on a wide range of topics in soil science, including soil physics, soil chemistry, soil biology, soil genesis, soil classification, and soil management.
- “Geoderma" publishes articles on the study of soils and the critical zone, including soil physics, soil chemistry, soil biology, soil genesis, soil classification, soil management, and biogeochemical cycling.
- “European Journal of Soil Science" publishes research on the chemical, physical and biological properties of soils, and their interactions with the atmosphere, hydrosphere, and biosphere.
- "Pedosphere" is focused on the study of soils and their interactions with the environment, including soil genesis, soil classification, soil management, soil-plant-atmosphere interactions, and soil-water-air-biota interactions.
- "Soil Research" is a scientific journal publishing papers on all aspects of soil science including soil biology, soil chemistry, soil physics, soil classification, land use and land management, and environmental and conservation issues.”

After conducting some preliminary research, the names of journals and their descriptions appear to be very accurate in this case. However, there are other scenarios where the AI Hallucination can be observed. When asked for some papers on soil texture in drought conditions, the model hallucinated some very convincing sounding papers.

1. "Drought-induced changes in soil texture and water retention properties in a Mediterranean agro-ecosystem" by M.R. Rueda, L. García-Orenes, A. Mataix-Solera, S. Gómez-Gómez, published in the journal "Catena" in 2010.

2. "Effect of soil texture on water retention and plant growth under drought conditions" by J. Wang, H. Zhang, H. Zhang, and Q. Zhang, published in the journal "Soil Science and Plant Nutrition" in 2019.
3. "Soil texture and water-holding capacity under different land use and drought conditions" by Y. Zhang, Y. Wang, X. Liu, and Y. Liu, published in the journal "Agriculture, Ecosystems & Environment" in 2018.
4. "Soil texture and water retention: implications for drought tolerance in Mediterranean shrubs" by E. García-Sánchez, J. P. Sánchez-Hernández, A. L. Sánchez-García, published in the journal "Environmental and Experimental Botany" in 2017.

As credible as those references sound, they do not exist. Many of the authors have also been meshed or combined and do not (seem to) exist. While being limited in being able to access or look up this information, this limits the current applications of this sort of software. As the software develops even further over time, it will be interesting to observe how these limitations are overcome.

The advancements of AI are of great concern to many people, who worry that it will impact their careers and livelihoods. With computers and AI now able to assist with writing computer code (Barke *et al.* 2022, Perry *et al.* 2022, Sandoval *et al.* 2022), as well as run shops such as Amazon Go stores. While the use of AI may be able to improve many areas of daily life, it could also take over many tasks, reducing the number of jobs available for people.

The introduction of AI art has caused a particular stir and divide in popularity. Whilst some people love the tools for quick generation of fantasy art, others are worried it could replace artists. Another issue is how the model generates art. Some models seem to take parts of art and merge them together into a final art piece, would could count as plagiarism or theft (Cetinic *et al.* 2021).

Whilst experts currently play a crucial role for initial training of AI, their part may become lesser as time goes on. If AI becomes so advanced that it can code and train further computer models, will there become a point when humans aren't needed to input anything? While this sounds like a Sci-Fi movie plot, with the way AI has advanced in recent years it could become a possibility.

As well as improving daily life by taking over monotonous tasks, AI could provide new opportunities and inclusivity for those with difficulties or disabilities. This could be through the use of self-driving wheelchairs with systems similar to self-driving cars, image recognition systems being used to give those with low-vision live descriptions of their surroundings or live captioning for those who are hard-of-hearing (Touzet 2023).

There are several areas where AI is expected to make significant advances soon. Some of the most interesting and impactful include:

1. Healthcare. AI is becoming heavily used to analyse medical images and assist diagnoses. It can also be used in administration tasks to improve patient care and improve timeframes for communication and follow-ups.
2. Drug discovery. The improved efficiency, accuracy and speed offered by AI is hugely popular within the pharmaceutical field, and could lead to many exciting breakthroughs (Mak *et al.* 2022).
3. Agriculture. From analysing soil nutrition state and outputting recommendations for treatments and crops for that year (Swathi *et al.* 2023), to creating systems capable of identifying crops and their development stage (Paul *et al.* 2024), AI has the potential to revolutionise agriculture (Wakchaure *et al.* 2023) in a time when efficient food production is of high importance.
4. Robotics: Combining AI with robotics will lead to the creation of robots capable of performing a wide range of tasks autonomously. This will have applications in a wide range of uses, from creating intelligent prosthetics for lost limbs to automation of tasks within manufacturing and beyond (Soori *et al.* 2023).
5. Live environmental monitoring. AI can model and predict events based on data. This could have incredible real-world applications to monitor

and predict natural disasters, ensuring evacuations and other precautions can happen well before disaster strikes (Velev *et al.* 2023).

It's important to note that AI is a major interdisciplinary field that is constantly evolving, so it's difficult to predict exactly what the future will hold. However, these are some of the areas that are expected to see significant progress soon.

As the use of artificial intelligence continues to expand across research, industry, and daily life, growing attention is being brought to the environmental impact of developing and running these models. While AI offers many potential efficiencies, particularly in automation and decision-making, these benefits must be weighed against the substantial resources required for model training, data storage, and ongoing operation.

One of the most significant concerns is the energy demand associated with training large-scale AI models. Training deep learning architectures, such as large language models, can involve running thousands of GPU units for extended periods. For example, the training of the language model GPT-3 is estimated to have required several hundred megawatt-hours of electricity (Gupta *et al.* 2025). This model in particular is projected to be 1,000 times larger than a typical language model at that time, and took 355 years in computing time to train (MIT News 2020). In other words, it would have taken a single high-end processor running continuously for three and a half centuries to complete the training task. However, in reality, this task was achieved by using thousands of processors in parallel to cut this timeframe down to a few weeks or months. These figures can be amplified further by considering the trend of repeatedly training and fine-tuning models to improve the performance of specific tasks.

In addition to the energy used during training, there are substantial ongoing costs linked to storage and inference. The large datasets used to train AI models, which could easily compromise of terabytes or even petabytes of data, must be maintained and frequently accessed. Storage and retrieval of this data across dispersed systems demands significant infrastructure, and inference (the process of running a trained model to generate outputs) also incurs energy costs, particularly for models deployed at such a large scale.

An often-overlooked aspect of this infrastructure is the water usage required for cooling data centres. High-performance computing facilities operating continuously generate considerable amounts of heat, and so cooling systems are essential to prevent hardware failure. Many large-scale data centres utilise water-intensive evaporative cooling systems. This has raised concerns about the local environmental impact of data centres, especially in water-scarce locations.

However, it can be tricky to get an accurate estimate of the energy usage and corresponding environmental impact of AI, due to models being incredibly variable, and most of the big companies running AI, such as Meta, Microsoft, and OpenAI, aren't making this information available.

These environmental impacts are not unique to AI, but the rapid uptake in AI applications, as well as the arms race to develop increasingly larger and more complicated models, has captured intense scrutiny. More sustainable approaches are being explored, including the optimisation of model architectures for efficiency, the use of renewable energy to power data centres, and improving hardware design to reduce cooling requirements. As it stands, the environmental footprint of AI remains substantial and should be considered when evaluating the adoption of machine learning for new applications.

3.3 Conclusions

The past five years have seen remarkable advancements in A, particularly in the fields of NLP, computer vision, and generative modelling. These technological advances have not only enhanced existing systems, but have also enabled entirely new forms of automation, data analysis, and human-computer interaction. Models such as ChatGPT exemplify how NLP systems are now capable of generating human-like responses and supporting real-time communication across numerous platforms. Similarly, developments in computer vision have improved the ability of machines to interpret and respond to visual data, supporting innovations in sectors ranging from autonomous transport to medical imaging.

Machine learning is increasingly embedded into everyday life and industrial operations. Virtual assistants, recommendation engines, real-time translation, and smart home devices have already demonstrated the potential consumer-facing applications of AI. In professional environments, AI enhances efficiency and accuracy through tasks such as fraud detection in finance, predictive maintenance in manufacturing, crop and yield analysis in agriculture, and early diagnostic support in healthcare. These applications highlight AI's ability to extract insights from vast and complex datasets at speeds and scales that are unachievable by human means alone.

Generative AI technologies, particularly GANs, have introduced novel capabilities in content creation, from generating artwork and music to synthesising training data in low-resource fields. These tools are also being applied in medical and environmental sciences, helping to simulate conditions, produce synthetic imagery for model training, and enhance low-resolution data. As these technologies continue to improve and evolve, their creative, industrial, and scientific applications are expected to broaden, creating new opportunities and challenges.

At the same time, the rise of AI technologies presents significant considerations for society and industry. One key area of concern is employment, as the automation of tasks previously carried out by humans may shift the structure of the workforce. While AI can improve productivity and remove barriers for individuals, such as those with disabilities through innovations such as live captioning or autonomous mobility, it may also contribute to job displacement in sectors reliant on repetitive or process-driven work. Balancing automation with job availability, and ensuring access to AI-enhanced tools, should be central to future policy and frameworks.

Also, the environmental impact of developing and deploying large-scale AI models requires careful consideration. Training and running these models requires substantial computational resources, often resulting in high energy consumption and water usage for data storage cooling. As AI adoption scales, sustainable development practices, including improved model efficiency, use of renewable energy, and better data management, will be crucial to mitigate long-term environmental effects.

Looking ahead, AI is expected to play a foundational role in shaping future advancements across a wide range of sectors. In healthcare, AI may revolutionise diagnostics and drug discovery, improving and saving lives. In agriculture, it has the potential to support sustainable food production through precision farming and predictive analytics. In environmental monitoring, AI can be used to detect, model, and mitigate climate events. The integration of AI with robots will expand possibilities in areas such as prosthetics and autonomous systems.

In conclusion, AI is no longer an emerging concept, but a present-day reality with a strong influence and a lot of potential. While the trajectory of AI brings new uncertainties and environmental concerns, it also offers huge potential for societal and technological progress. It is crucial that researchers, developers, industries and policy makers alike consider the adoption of this new technology whilst also considering responsibilities and sustainable advancements.

4. Quantification of soil structural dynamics over a growth season for soils under conventional and zero tillage

4.1 Introduction

Tillage practices, particularly those using ploughs were recorded as early as the 1800s (Lal *et al.* 2007), and are still widely used today. The aim of tillage is to: 1) to prepare a desirable soil structure for seed to germinate in, 2) to support the flow of water, air and heat into and through the soil, 3) to control weeds and plant diseases, 4) to remove foreign materials such as rocks and old roots, and 5) to establish surface structure (such as beds and furrows) for irrigation and drainage (Gebhardt *et al.* 1985). However, these conventional tillage methods have been shown to cause increased surface run-off and erosion by wind and rain (Lindstrom *et al.* 2001), as well as compaction by repeated use of large farm machinery. In more recent years, it has also been demonstrated that conventional tillage increases the release of organic carbon from soil contributing to global CO_2 emissions (Guo *et al.* 2020).

Zero-Tillage (ZT) is an agricultural practice where crop seeds are sown directly onto the surface of the soil; without ploughing beforehand. ZT has been found to produce better pore connectivity in soils (Galdos *et al.* 2019) and leading to enhanced water infiltration (Su *et al.* 2007), despite these both being goals of tillage. It has also been found that ZT leads to reduced surface run off and erosion (DeLaune *et al.* 2013) when compared to conventionally tilled soils. ZT has also been shown to reduce CO_2 emissions and support carbon sequestration, something that has previously been disputed in the literature (Cooper *et al.* 2021).

To understand the impacts of different tillage treatments on soil structure, X-ray Computed Tomography (CT) has emerged as a popular tool to non-destructively observe the structure of soil. The technique facilitates visualisation and quantification of aggregates and pores as they are structured in the field. This technique has been used to observe short-term ZT structural effects (Guo *et al.* 2020, Alskaf *et al.* 2021), as well as long-term (30+years) of ZT (Galdos *et al.* 2019). Prior to the popularity of X-ray CT, Atkinson *et al.* (2009) used a resin impregnation method to observe structural changes over a season under conventional and zero tillage management and found that the mean pore size of soil varied based on cultivation treatment and time since cultivation, with power harrowing producing the best pore structure for seed establishment.

It has already been recorded that the physical properties of soil including soil structure vary seasonally. Bryk *et al.* (2017) observed that seasonal weather impacts soil structure. Between November 2009 and March 2010, the aggregate structure changed from having well-developed sub-angular blocks to a structure with less well-developed sub-angular blocks. Between May and June, the soil had a non-aggregate structure as a result of aggregate welding due to compaction by rainfall (Bryk *et al.* 2017). Earthworms have been well documented to influence soil structure by supporting soil aggregates and the associated micropore development (Zangerlé *et al.* 2011) as well as creating macropores as they move through the soil (Lamandé *et al.* 2011). Schon *et al.* (2017) observed 70% more macropores in soil under the influence of high earthworm abundance and diversity. Soil moisture and temperature influence earthworm activity (Perreault *et al.* 2006), so during the colder and wetter winter months they exert less influence in terms of soil structure development.

The objectives of this study are:

- To quantify changes in soil structural properties over the course of a growing season using X-ray CT imaging
- To compare soil structural dynamics between conventional ploughing and long-term zero tillage management systems
- To assess changes in porosity, pore size distribution, and structural stability over time for each management practice.

- To evaluate how soil structure changes over time under conventional and zero tillage management in field conditions.

4.2 Methods and Materials

4.2.1 Sampling site

Soil samples were collected by the author from an established experimental trial located at the University of Nottingham farm. The trial includes plots that have been under zero-tillage and conventional tillage systems since 2014. The site is located in Kingston on Soar, Nottinghamshire, England (52.841677, -1.254445). The soil at the site is classified as a sandy loam belonging to the Dunnington Heath Series, corresponding to a Brown Earth according to the FAO classification system (Alskaf *et al.* 2021).

4.2.2 Treatments

The trial contains conventional, minimum till and zero till treatments along with plus or minus plant residues. However, in this study only the no-residue plots were selected for analysis in order to simplify sampling and isolate the structural effects of tillage alone, without the additional variability introduced by surface residue cover. This approach also allowed for consistent sampling to be conducted across both tillage treatments.

Each plot was 12m long by 9.6m wide, and each treatment had 4 replicates established in 4 separate blocks. For further details see Alskaf *et al.* (2021).

Conventional tillage involved ploughing to a 30cm depth whereas zero-till plots were not cultivated before sowing. Cultivation was conducted in September, a few days before the first sampling session. The crop grown during this experiment was *Skyfall* winter wheat *(Triticum aestivum)*, with sowing in mid-October and harvest in August.

4.2.3 Sampling

The first samples were taken in September, a few days after cultivation had occurred. One sample per treatment was taken from each block at each sampling session. This led to a total of eight samples (two treatments x four replicates) each sampling session. To minimise variability caused by spatial heterogeneity within the plots, sampling locations were selected near the centre of each plot while avoiding wheel tracks and visible residues. Where possible, cores were taken from undisturbed areas representative of the plot condition.

Samples were taken on 23rd September 2020 (shortly after cultivation), 16th October (shortly after drilling), 14th December, 25th February, 23rd April, and 21st June 2021, leading to a total of forty-eight samples.

Soil samples were collected using a soil corer with a 5 cm diameter with transparent plastic sleeves inside, which was inserted into the soil to a depth of approximately 30 cm. Extracted samples were then sealed within plastic bags to maintain moisture levels. These were stored at 4°C until X-ray scanning took place, typically within a few days.

Due to the high proportion of stones present in the field, some difficulty was encountered during core collection. In cases where the samples were incomplete or compromised by obstructions, resampling was carried out immediately to ensure that a complete, undisturbed core was successfully collected from each plot on the scheduled sampling day.

4.2.4 X-ray Computed Tomography (CT)

X-ray CT was used to visualise the internal soil structure. The soil cores were scanned using a Phoenix v|tome|x L Custom® μCT scanner (GE Sensing and Inspection Technologies, Wunstorf, Germany) located in the Hounsfield Facility at the University of Nottingham. Each core was scanned for 30 minutes at 150kV with a resolution of 40μm with 1800 images per scan. Phoenix datos x software (GE Sensing and Inspection Technologies) was used to reconstruct the scan data into 3-D greyscale volumes. These volume files were then exported into VG Studio Max software to view the 3-D image and to normalise the brightness and contrast. 2D images were then exported from VG Studio for processing.

4.2.5 Image Processing

The reconstructed 3-D CT volumes were first processed in VG Studio Max, where a central region of interest was selected from each core to minimise edge effects and any potential disruption from the sample process. This region of interest was then exported as a stack of 2-D cross-sectional images. These 2-D image sets were subsequently processed in ImageJ (Schindelin *et al.* 2012) for segmentation and analysis. Examples of the extracted 2-D slices are shown in Figure 4.1.

Before being analysed, the 2-D images were segmented to distinguish between solid soil material and pore space. Several thresholding algorithms available within ImageJ were tested on a subset of images, and segmentation outcomes were visually compared to the original greyscale images. The Huang algorithm (Huang *et al.* 1995) was selected as it consistently provided the most accurate separation of soil matrix and pore space for these samples. The Huang method is effective for images with relatively low contrast or overlapping greyscale values, making it highly suitable for soil CT data.

Once thresholded, the images were converted to binary format, where black pixels represented solid soil material and white pixels represented pore space (Figure 4.2). These binarised images were then analysed using the BoneJ plugin (Doube *et al.* 2010). This plugin was used to calculate the pore structural properties from the full 3-D reconstructions, including total pore volume, total porosity, and the volume of the largest continuous pore. Here, total porosity refers to the proportion of pore space detected within the imaged volume at the resolution used in this study.

During sampling and scanning, it became apparent that a number of soil cores contained large stones. In some cases, these stones occupied a significant proportion of the imaged volume or created artefacts due to their high X-ray attenuation, which resulted in artificial brightness and misclassification of pore space during thresholding. To avoid introducing bias, any images where stones occupied more than 40% of the 2-D slice area, or where segmentation was clearly compromised, were excluded from the analysis. This exclusion procedure was applied consistently across all samples. Approximately 15% of the images were removed from the dataset following this quality control stage.

4.2.6 Statistical Analysis

The statistical software package GenStat (International 2021) was used to analyse data using a repeated measures analysis of variance (ANOVA) to test for significant differences between and among treatments. RStudio (RStudio Team (2020) was used to calculate standard errors of differences and create graphical outputs.

4.3 Results

4.3.1 Soil Structure

Figures 4.1-4.6 display single 2-D images from X-ray CT scans of tilled and zero-tilled soil cores for each month of the sampling period: September, October, December, February, April, and June, respectively. Each figure displays both the greyscale CT image and binary images, enabling a visual example of soil structure and pore distribution between treatments.

In the greyscale images, notable structural differences were observed in the earlier months (September and October), with conventional tillage displaying a more fragmented structure and a grainy soil texture, often characterised by several clear macroaggregates. In contrast, zero-till soils lacked clear aggregation but exhibited more defined cracks and channels, suggesting better connected pore networks. This trend was similar from September to December, with zero-till consistently showing clearer pore pathways.

From February onwards, the structural differences between the two treatments becoming less distinct. In both April and June, soil porosity between conventional and zero-till becomes more visually comparable, with few obvious distinctions in the greyscale and binary images. This convergence may reflect seasonal influences such as weathering and biological activity gradually normalising structural differences over time.

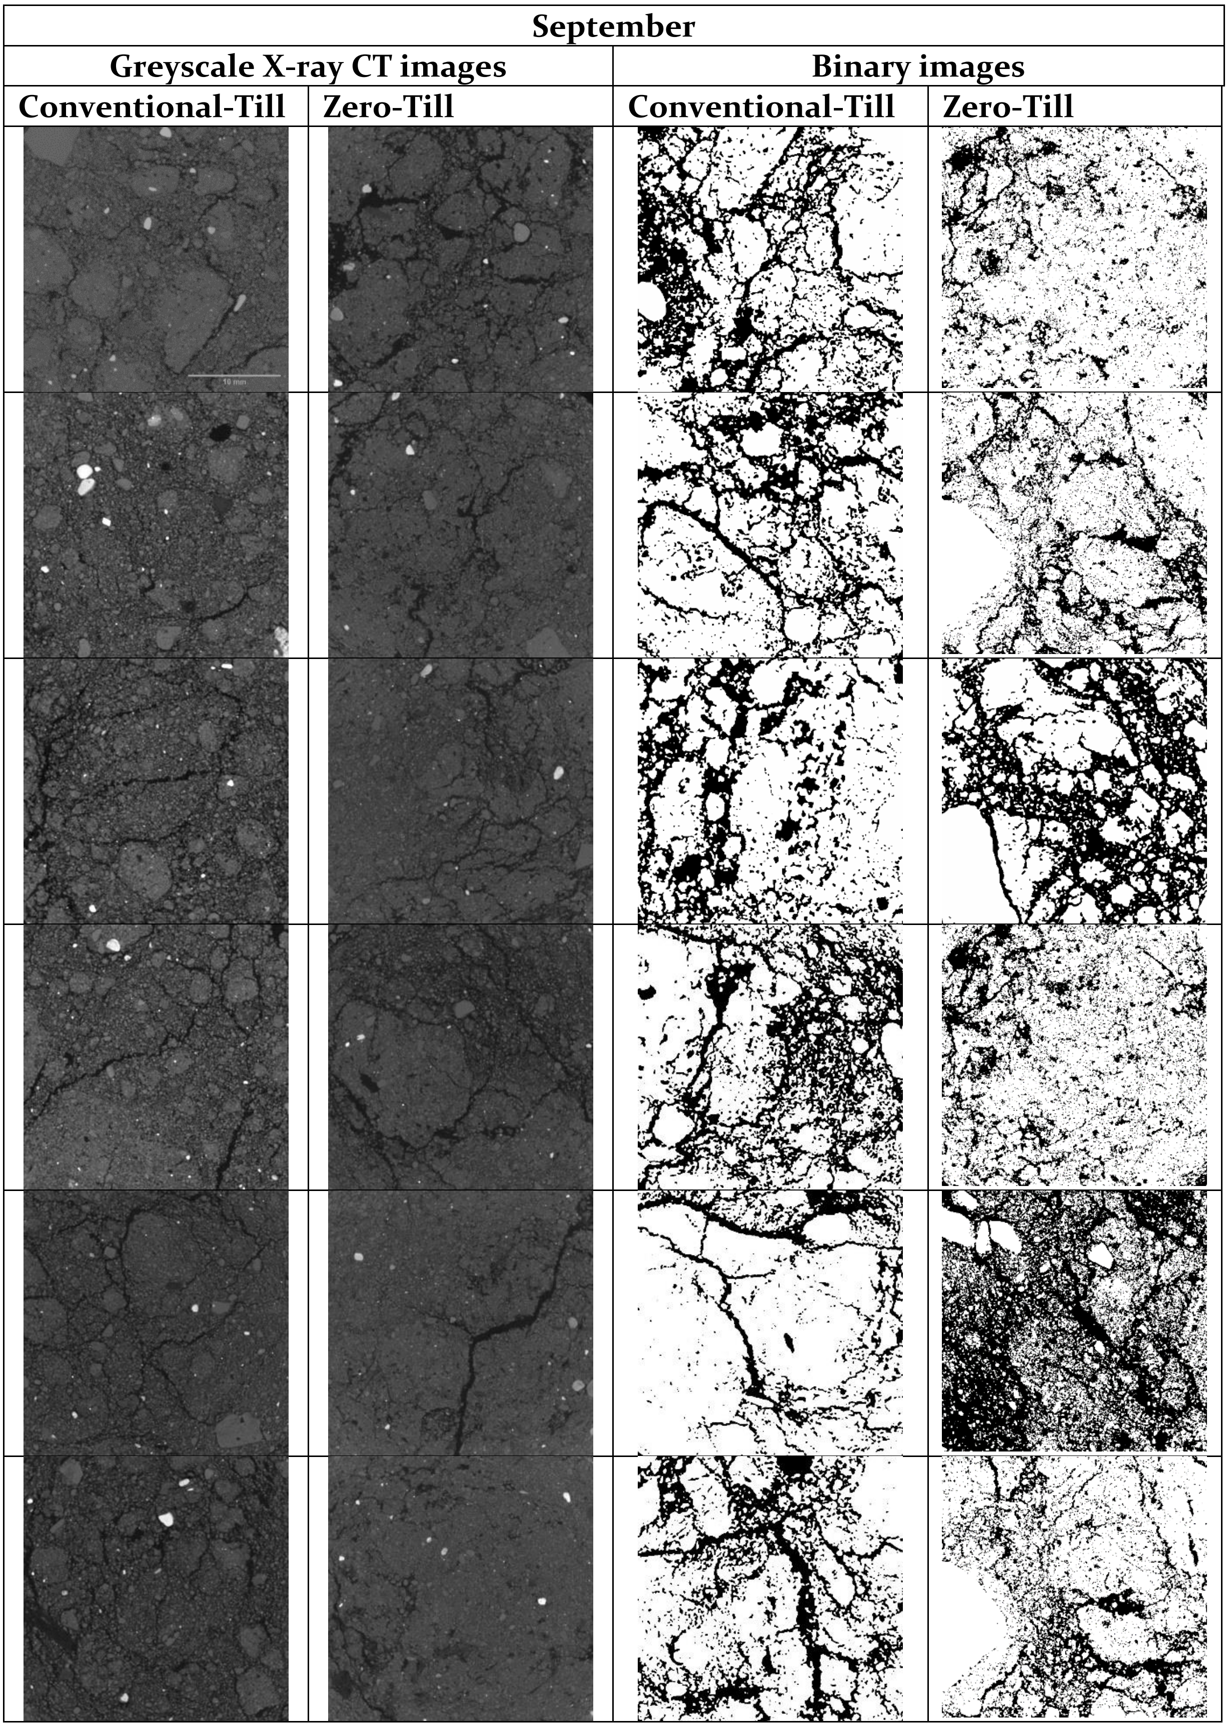

Fig. 4.1: Greyscale X-ray CT images and binary images for September, comparing Conventional Till and Zero-Till treatments. All images are taken from similar depths (approx. 10-20cm). In the greyscale CT images, lighter shades represent denser, less porous material, while darker areas indicate more porous regions (e.g. air). In the binary images, black areas represent pore space.

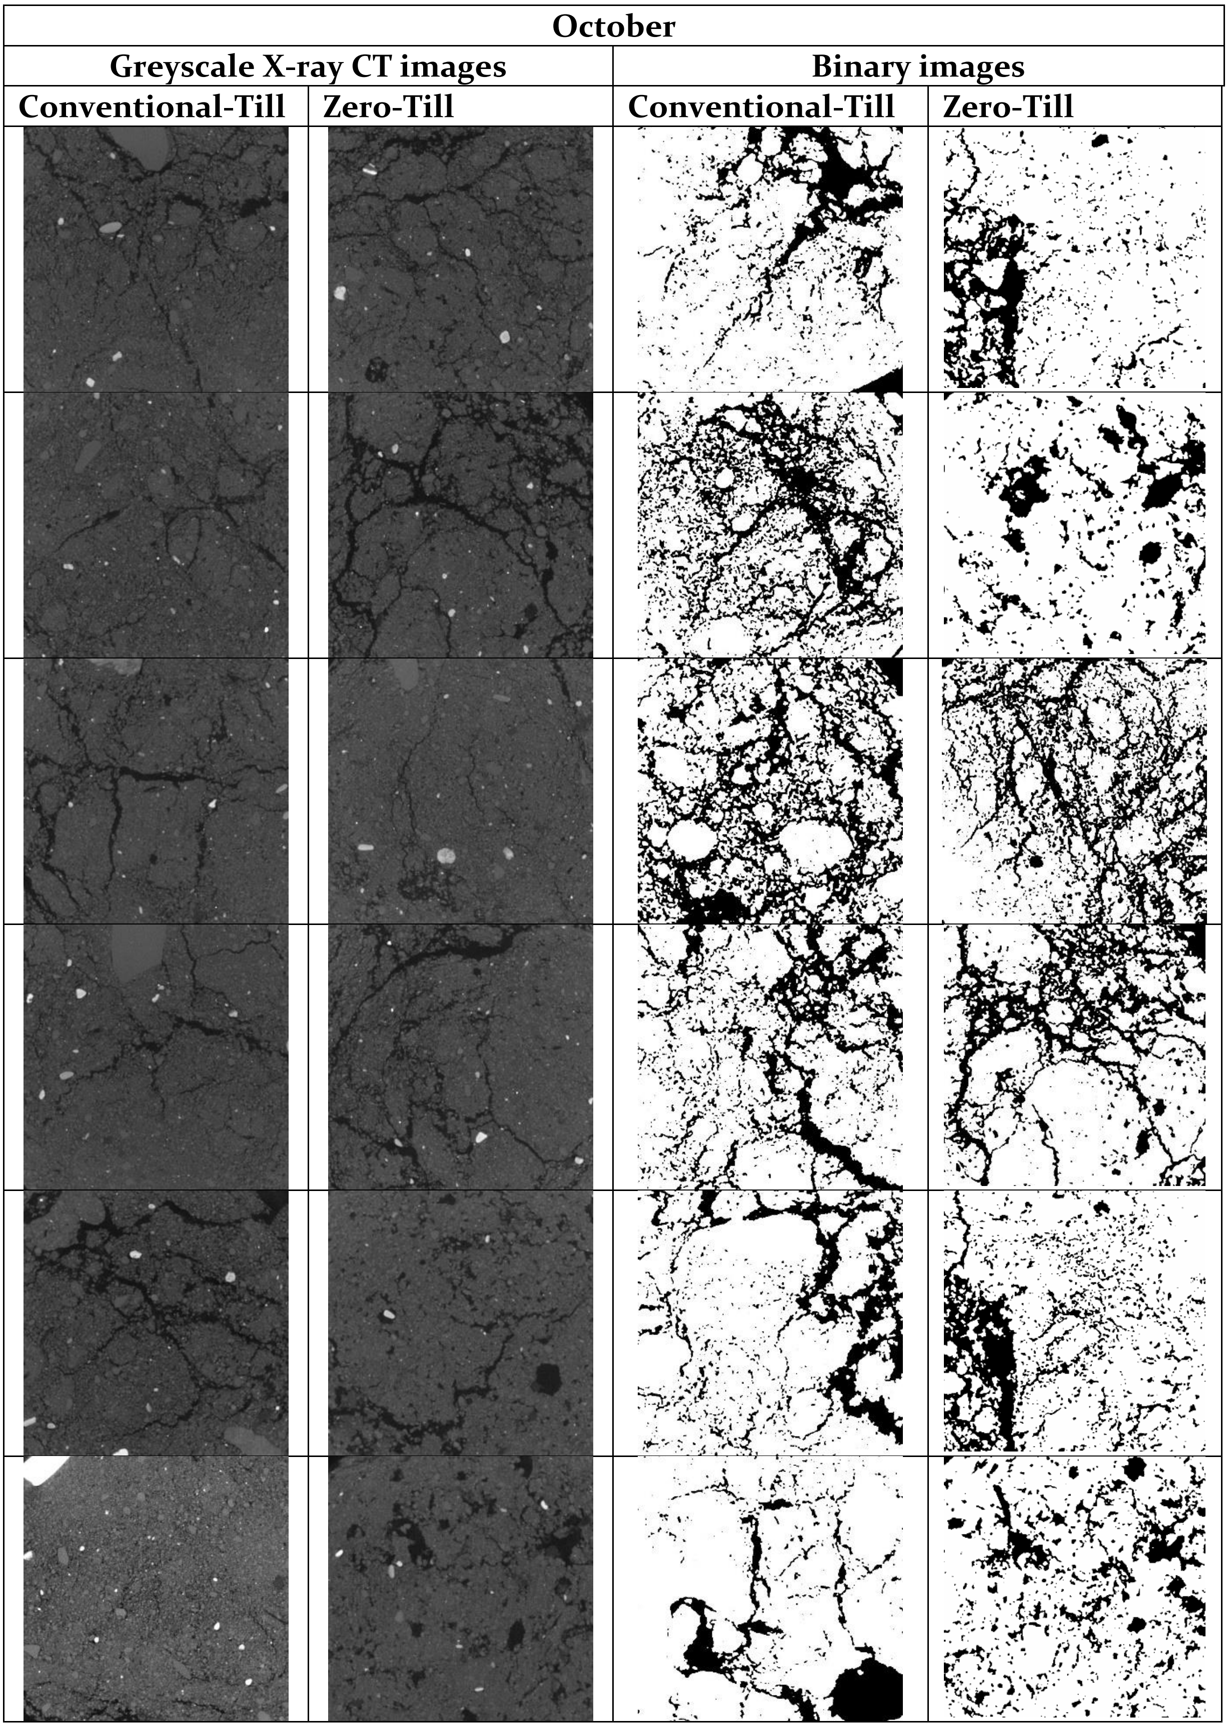

Fig. 4.2: Greyscale X-ray CT images and binary images for October, comparing Conventional Till and Zero-Till treatments. All images are taken from similar depths (approx. 10-20cm). In the greyscale CT images, lighter shades represent denser, less porous material, while darker areas indicate more porous regions (e.g. air). In the binary images, black areas represent pore space.

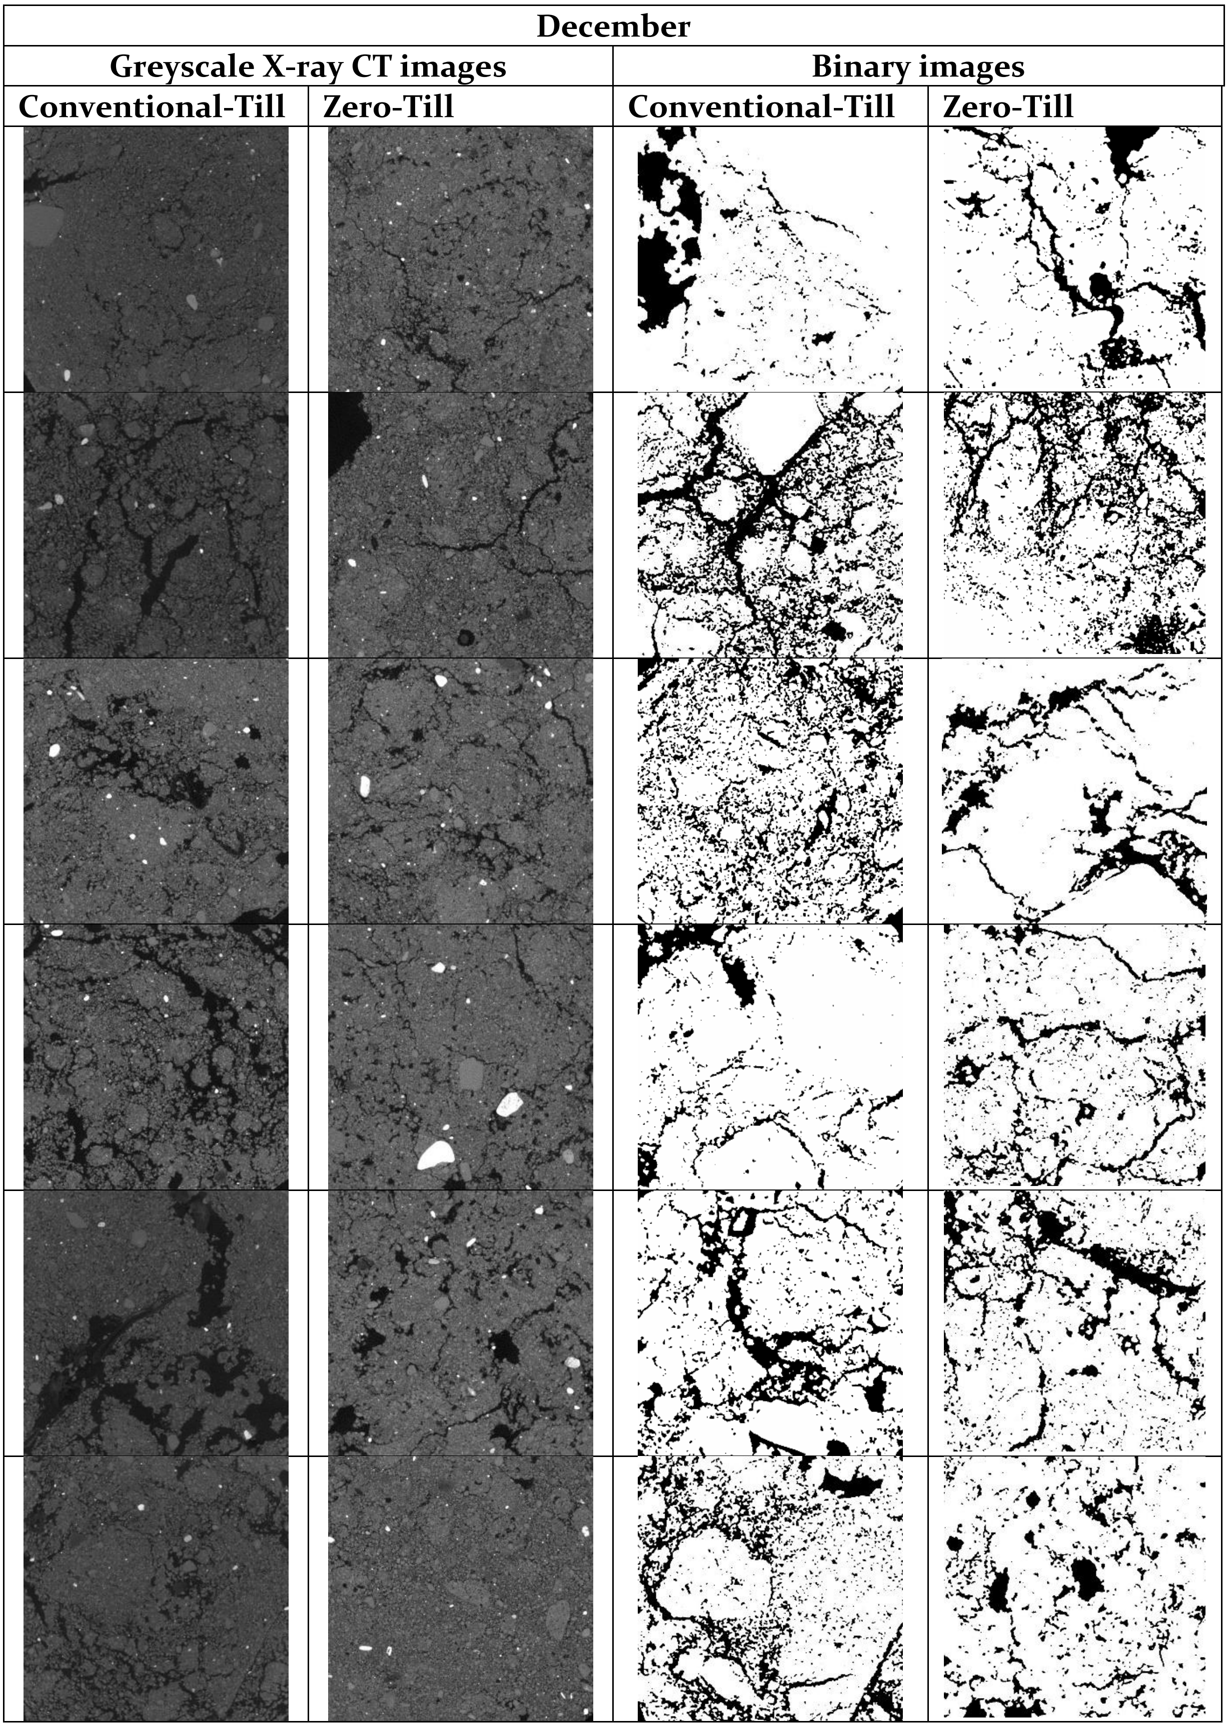

Fig. 4.3: Greyscale X-ray CT images and binary images for December, comparing Conventional Till and Zero-Till treatments. All images are taken from similar depths (approx. 10-20cm). In the greyscale CT images, lighter shades represent denser, less porous material, while darker areas indicate more porous regions (e.g. air). In the binary images, black areas represent pore space.

February			
Greyscale X-ray CT images		**Binary images**	
Conventional-Till	**Zero-Till**	**Conventional-Till**	**Zero-Till**

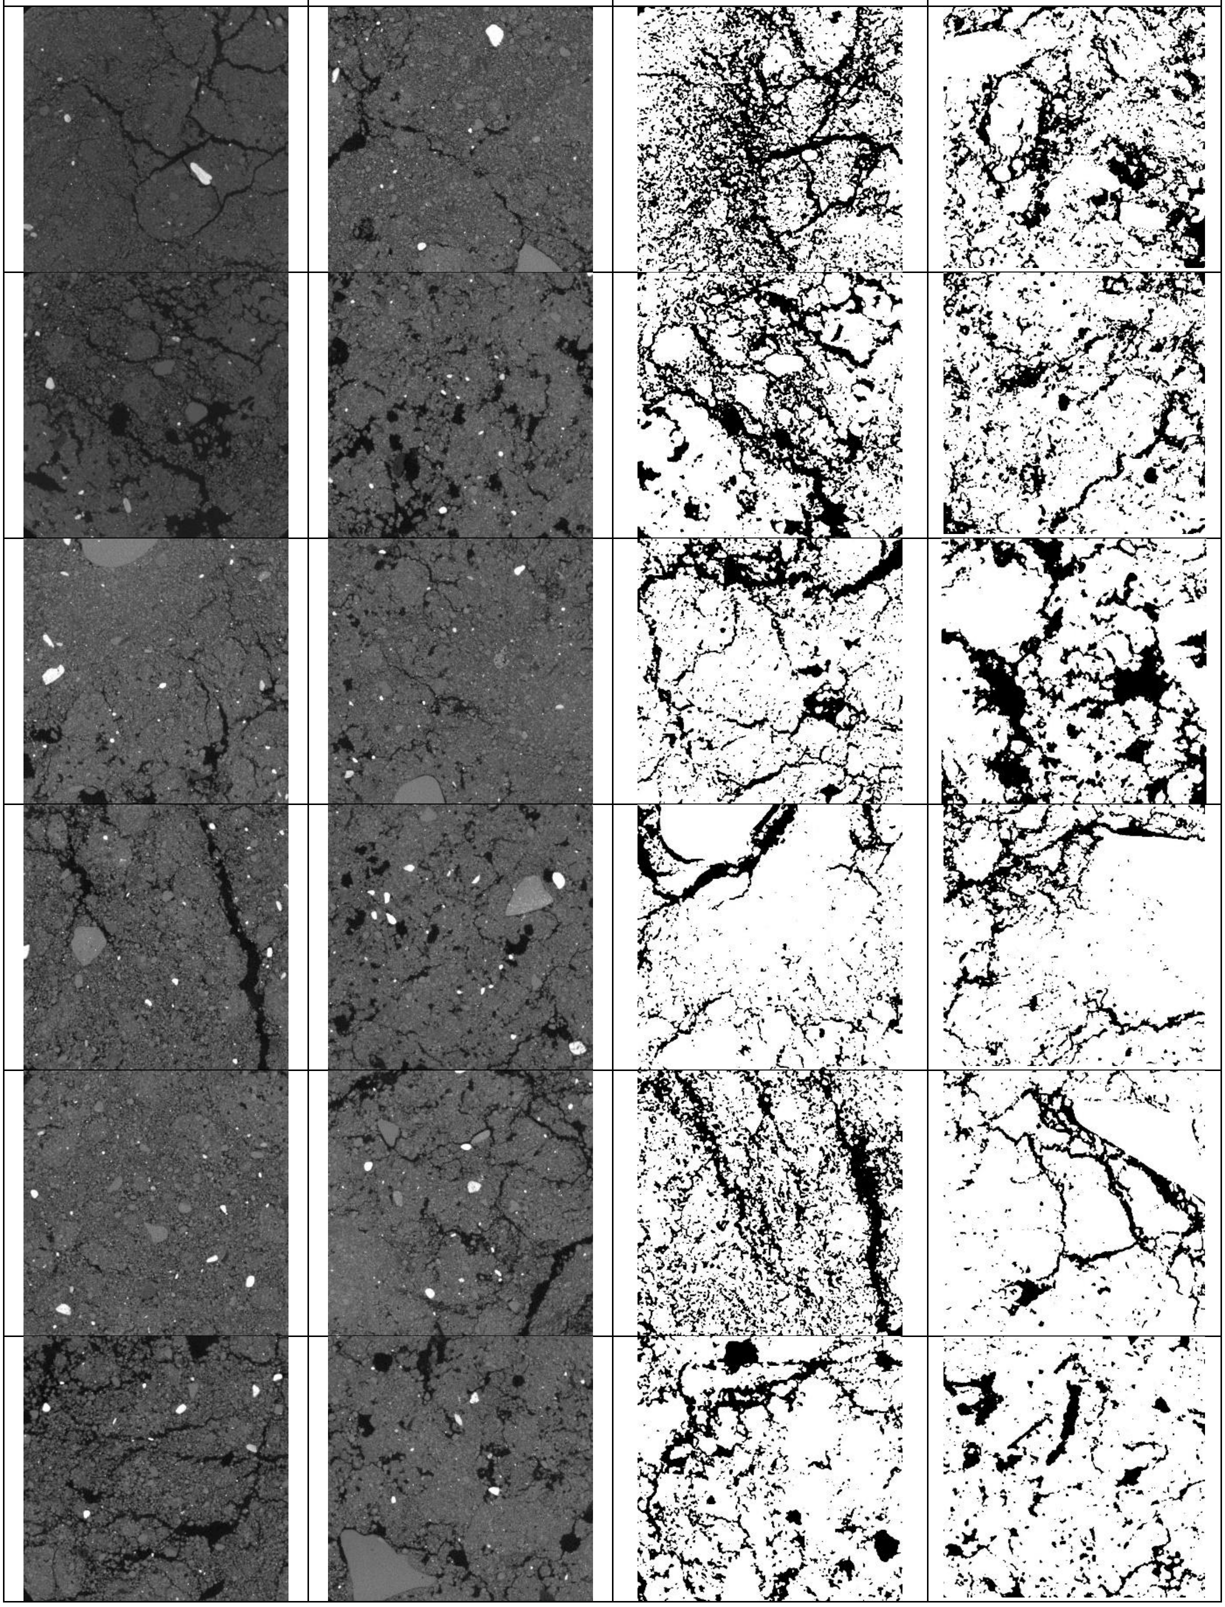

Fig. 4.4: Greyscale X-ray CT images and binary images for February, comparing Conventional Till and Zero-Till treatments. All images are taken from similar depths (approx. 10-20cm). In the greyscale CT images, lighter shades represent denser, less porous material, while darker areas indicate more porous regions (e.g. air). In the binary images, black areas represent pore space.

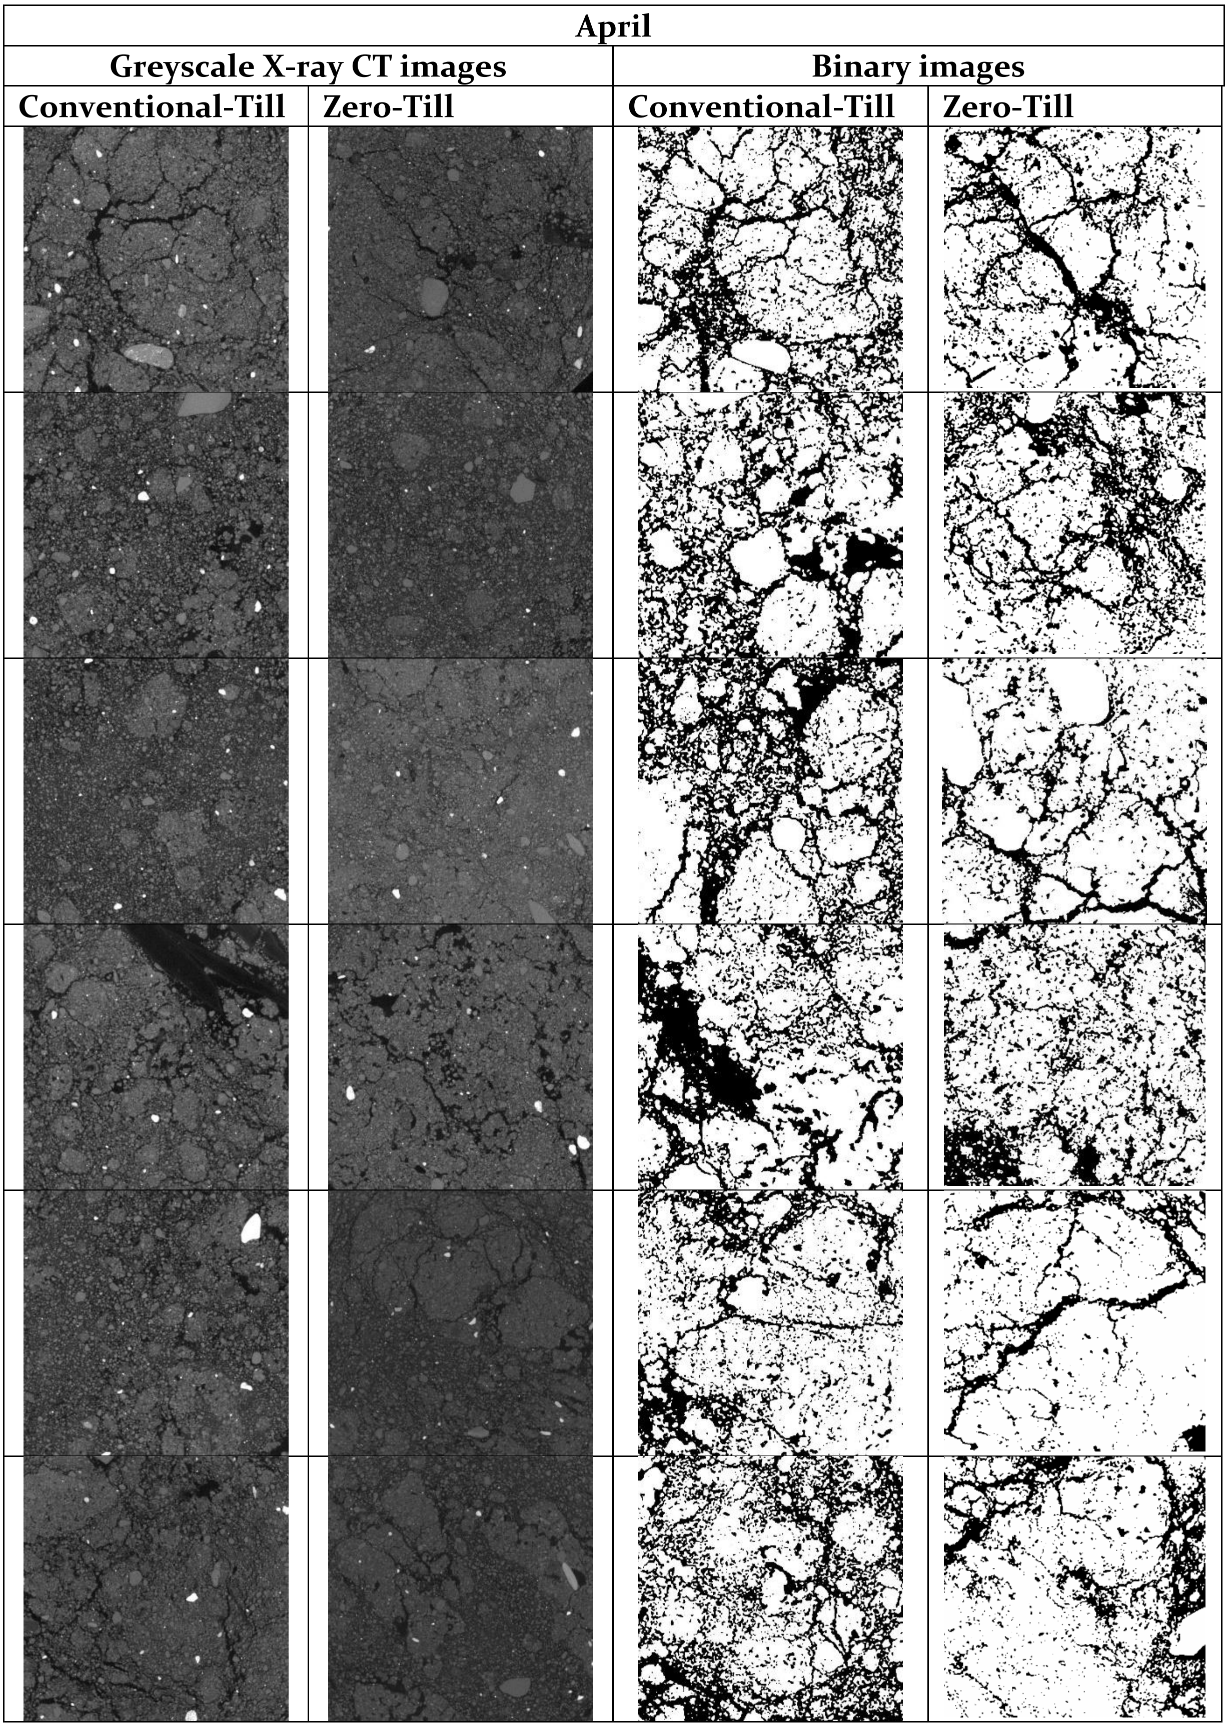

Fig. 4.5: Greyscale X-ray CT images and binary images for April, comparing Conventional Till and Zero-Till treatments. All images are taken from similar depths (approx. 10-20cm). In the greyscale CT images, lighter shades represent denser, less porous material, while darker areas indicate more porous regions (e.g. air). In the binary images, black areas represent pore space.

June

Greyscale X-ray CT images		Binary images	
Conventional-Till	**Zero-Till**	**Conventional-Till**	**Zero-Till**

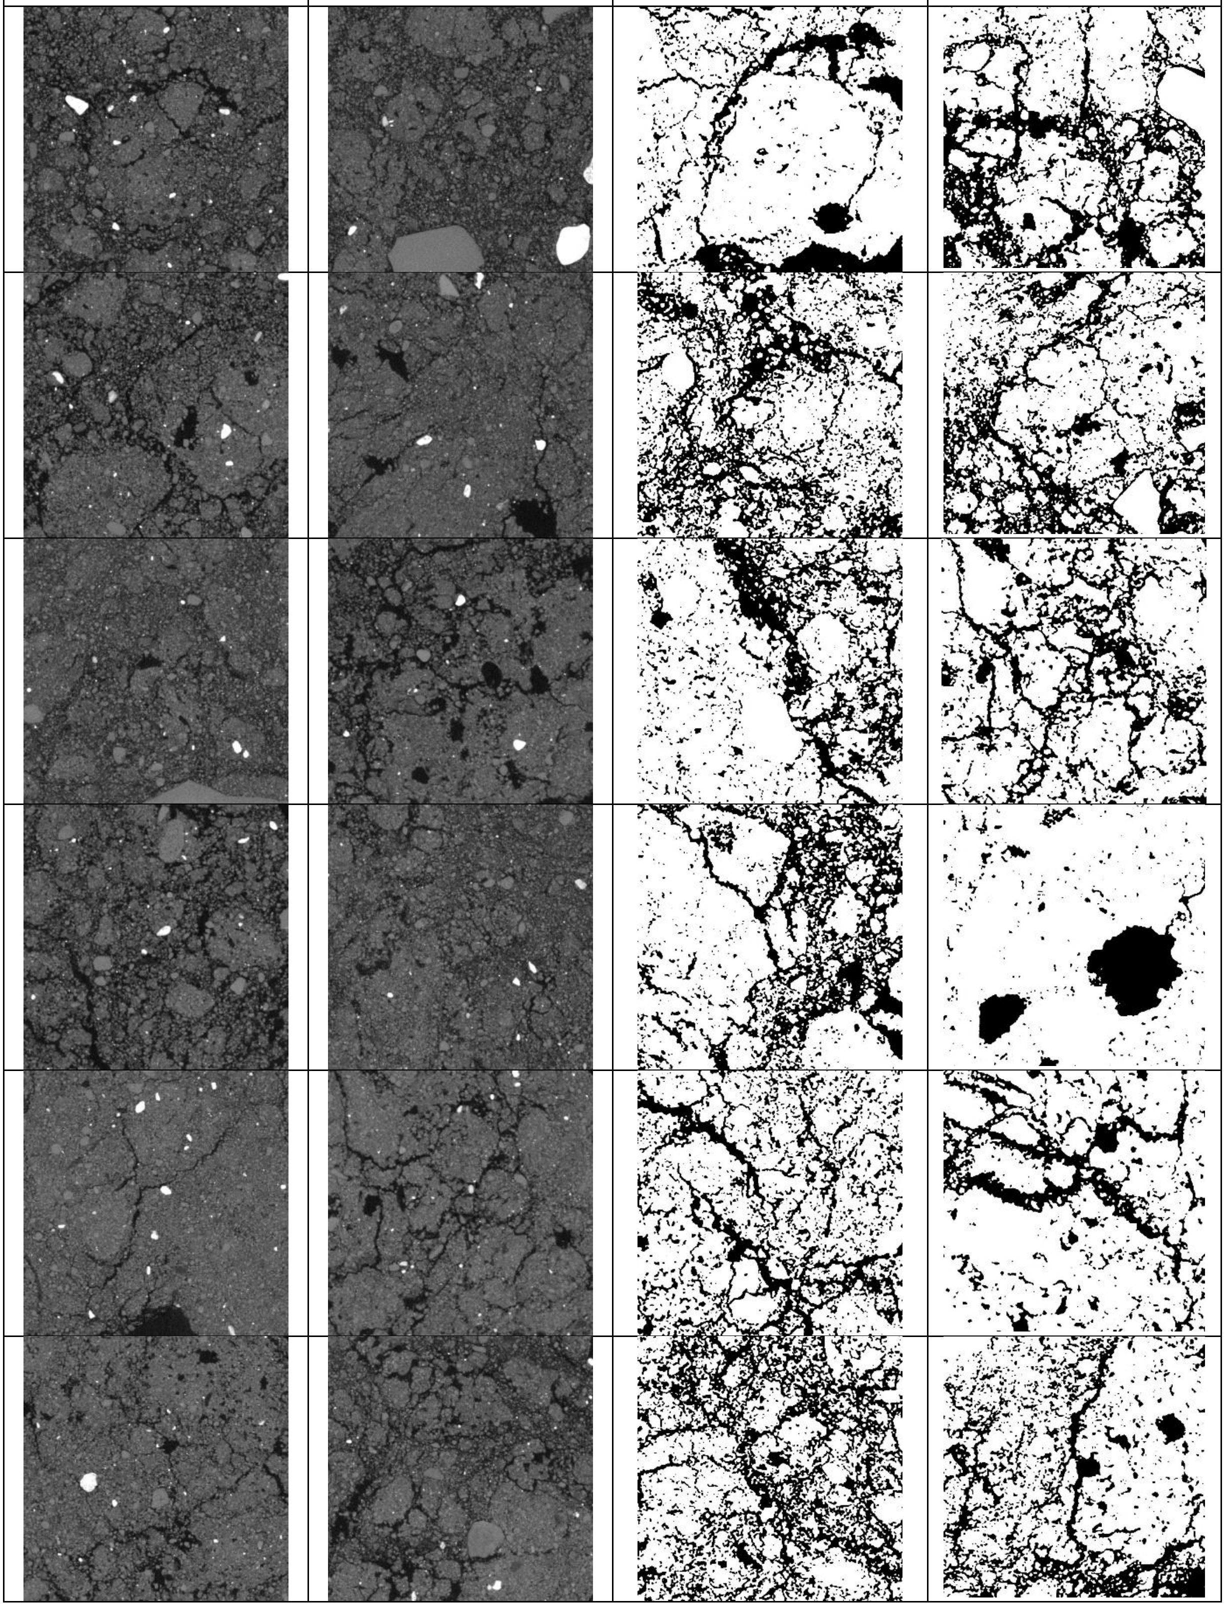

Fig. 4.6: Greyscale X-ray CT images and binary images for June, comparing Conventional Till and Zero-Till treatments. All images are taken from similar depths (approx. 10-20cm). In the greyscale CT images, lighter shades represent denser, less porous material, while darker areas indicate more porous regions (e.g. air). In the binary images, black areas represent pore space.

4.3.2 Total Porosity

Total porosity indicates the percentage of pore space within a sample. These values were calculated from the binarised image stacks using Bone J (Doube *et al.* 2010) by dividing the number of pixels classified as pore space by the total number of pixels within the region of interest, and expressing this as a percentage.

Figure 4.7 shows that tilled soil has a higher average porosity than ZT soil throughout the season. In September, directly after cultivation, this difference in porosity is most noticeable, with tilled soil being about 20% more porous than zero-tilled. The difference between tilled and ZT decreases as the season goes on. Similar to Figure 4.7, porosity for both tillage treatments decreasing by 21% for tilled and 12% for ZT from September to December (Autumn to Winter) and then increasing again from February to April (Winter to Spring), can be observed. In June, there is a decrease of about 6% for tilled and a very small increase of 0.5% for ZT. Over the season, the porosity of tilled soil decreased by 12%, whilst the porosity of ZT soil increased by 3%.

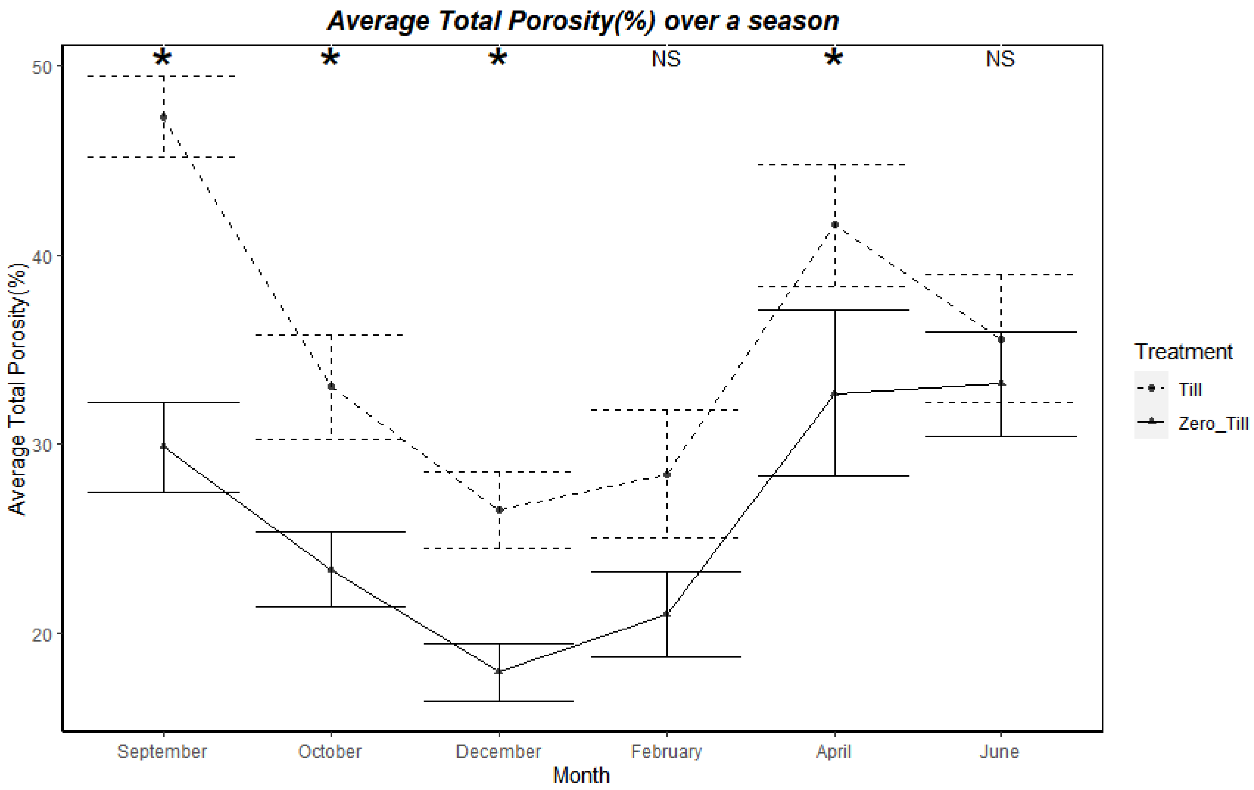

Fig. 4.7: Comparison of how the average porosity varies between conventional tillage (till) and zero-tillage (Zero Till) over a season. Error bars represent standard error of differences (S.E.D.)

4.3.4 Largest Individual Pore

One of the most recorded soil structural differences observed under different tillage management is the impact on the pore network connectivity. By measuring the volume of the largest pore, an indication of how connected the pore network is provided. Figure 4.8 shows that the largest pore volume does change over a season, with the pore volume decreasing in both tilled and zero-tilled from September until February and then increasing again in April. Tilled has a larger largest pore volume than zero-tilled until June, and the difference is significant in all months except for February and June. This indicates that tilled soils are more connected than the zero-tilled soils.

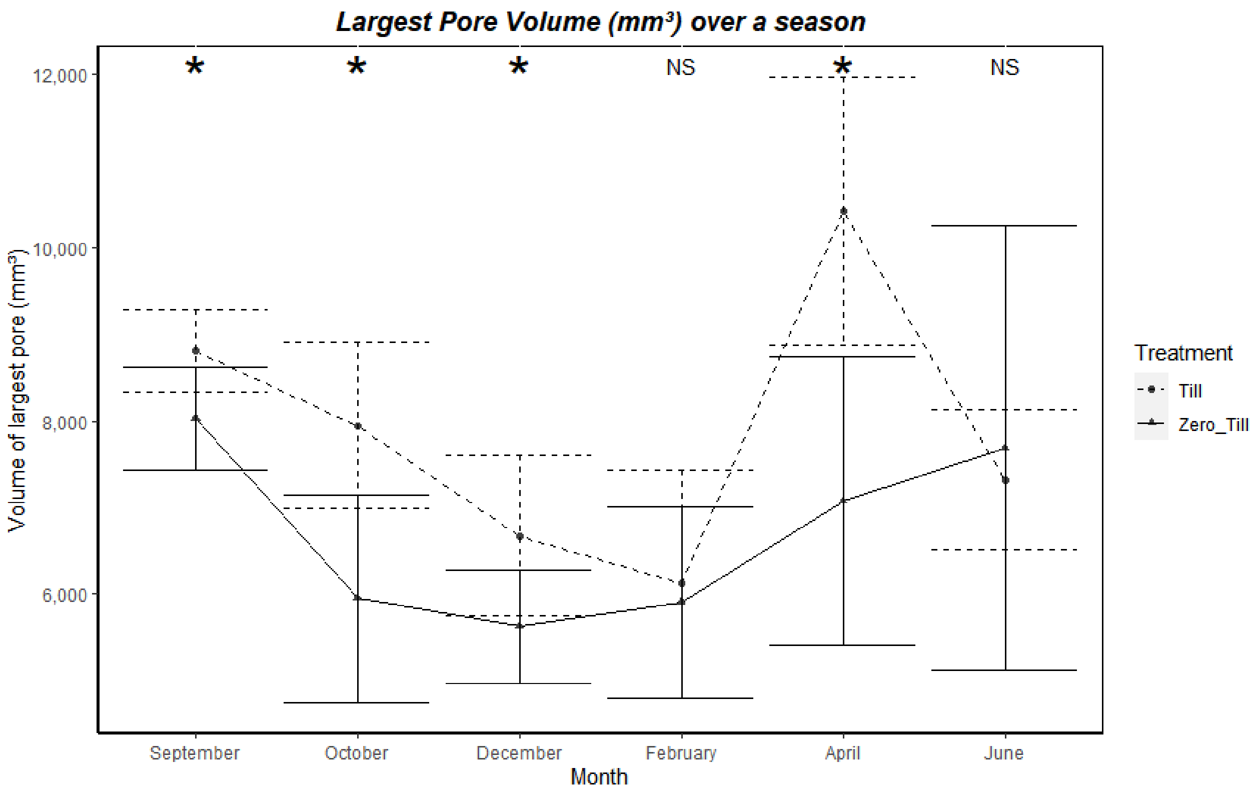

Fig. 4.8: Comparison of how volume of the largest pore in each sample varies between conventional tillage (till) and zero-tillage (Zero Till) over a season. Error bars represent standard error of differences (S.E.D.)

4.3.5 Average Pore Size

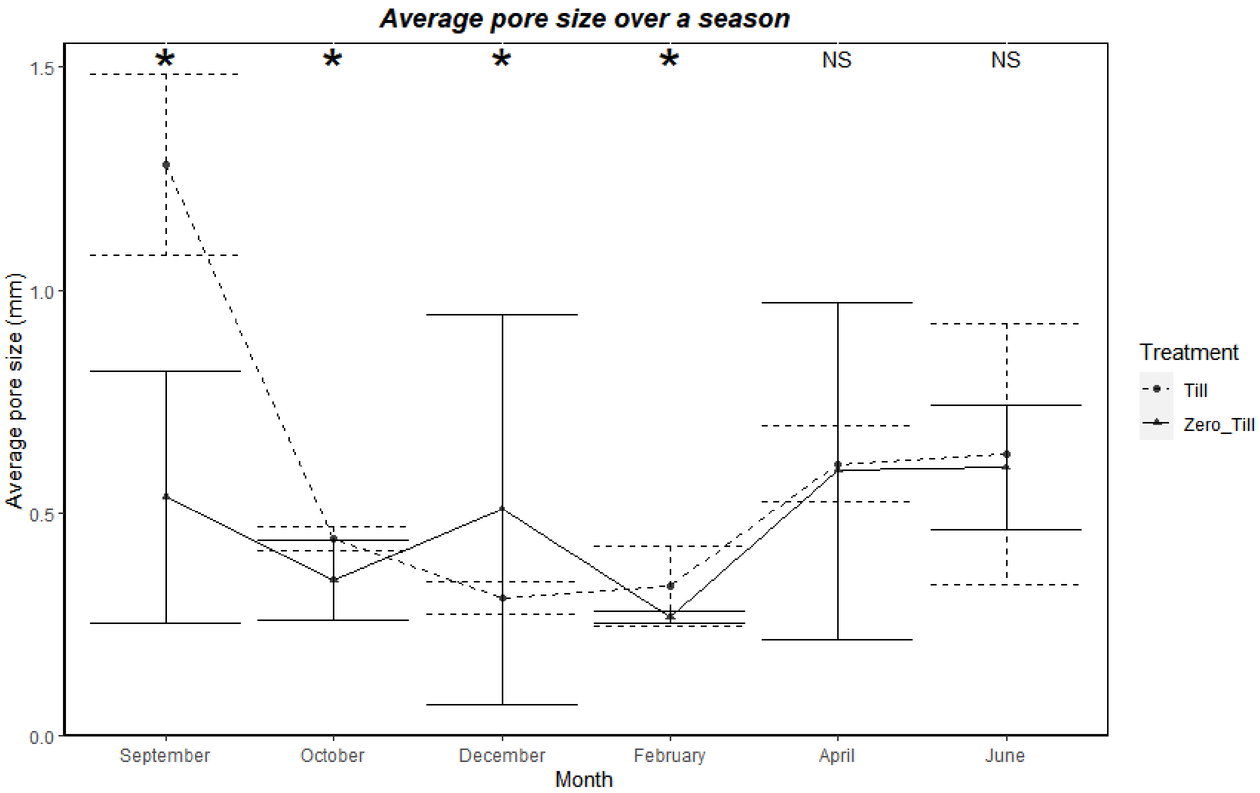

Fig. 4.9: Comparison of how pore size varies between conventional tillage (till) and zero tillage (Zero Till) over a season. Error bars represent standard error of differences (SED).

Figure 4.9 represents the average pore size and how it changes over a season between conventional and ZT. In September directly after tilling, the average pore size in tilled soil is over twice as large as those in the ZT soil. However, by October the two treatments appear to be more similar in average pore size. Apart from September, the two tillage treatments had similar average pore sizes and tilled soil has a slightly larger average size than ZT soil. As seen with the Figures 4.7-4.9, there was a decrease observed during the winter months, with increases beginning between February and April.

4.4 Discussion

Seasonal change caused variations in measurements between tillage treatments. The decreases observed in Figures 4.3 and 4.5 from September to December and increase from February to June could indicate both soils reaching a similar dormant state with few differences in soil structure. Winter months in the UK are known to be wet and cold, leading to reduced earthworm borrowing (Perreault *et al.* 2006) and microbial activity (Van Elsas *et al.* 2019), both of which are known contributors to the soil pore network. This suggests that as soils become dormant, there is an impact upon the pore network structure.

June in Figures 4.7, 4.8 and 4.9 also display non-significant results in June, which could be due to equilibrium rebalancing 10 months after cultivation. In the spring and summer months, crop roots could also be impacting in the samples, although best efforts were made to avoid including full plants in the samples, it is highly likely the soils sampled were influenced by roots.

The initial impact of conventional tillage compared to ZT on soil structure is fairly well understood, and this explains the large differences between the September results for Figures 4.1-4.6. Then over 3 months of Autumn/Winter they become more similar. Weathering of soil can include surface erosion and compaction (Connolly 1998). The samples examined here are from 20-30cm deep in the soil, which shows the potential depth of weathering impact. When you combine the potential compaction impact of weather with the impact of conventional tillage methods on soil compaction, you seem to get a compounding effect on compaction to at least 30 cm depth (Liu *et al.* 2022). Liu *et al.* (2022) showed the effect of soil compaction on grain yield under higher and lower tillage compaction over 2 years, with one drier and one wetter year. However, none have shown this weathering effect on soil structure over a season. Considering the impact compaction has on hydraulic conductivity and water uptake by crops (Gupta *et al.* 1989, Jayawardane *et al.* 1994, Connolly 1998), many papers may need to be mindful of when sampling occurs, due to the impact seasonal time may have on results. As seen in the figures in this paper, the soil structure results collected may be interpreted differently based on when in the season these samples were collected.

Figure 4.7 shows that the changes in soil porosity over a season change very similarly in both cultivations. The conventionally tilled soil displays the initially much higher porosity and remains higher than zero-till throughout the experiment. The two soil treatments follow a very similar pattern whilst conventional till remains higher, although the results in June seem to be moving towards a similar level. None of the other measurements taken in this study show this same pattern. This suggests that conventional tillage has a large and lasting impact on this aspect of soil structure, at least for one growing season. In Figures 4.8 and 4.9, which show the largest pore volume and the average pore size however, after the initial higher values shown by conventionally tilled soils, the following months seem to show these measurements becoming very similar between the two treatments: leading to no significant difference between treatments in June. This could be the soil returning to its own state of equilibrium after disruption caused by cultivation in September.

4.5 Conclusions

This study demonstrated that soil structure varies significantly throughout the growing season under both conventional tillage and zero tillage systems. The clearest differences between treatments were, as expected, observed immediately after cultivation, where conventional tillage resulted in substantially higher total porosity, larger pore sizes, and larger individual pores compared to zero tillage. However, these structural differences lessened over time, with both treatments showing a general decrease in porosity and pore size through the autumn and winter months, followed by partial recovery in spring. By June, the final sampling point, differences between tillage systems in pore size and largest pore volume had largely converged, while porosity remained consistently higher in conventionally tilled soils.

These seasonal patterns suggest that soil structure following tillage is dynamic and highly responsive to both management practices and environmental factors such as rainfall, temperature, biological activity, and root growth. Compaction from natural weathering processes may account for the observed decline in porosity over winter, even at depths of 20-30cm, while root development may contribute to partial structural recovery in spring.

Importantly, the results also indicate that the timing of sampling can strongly influence the significance of observed differences between tillage treatments. Measurements taken immediately after cultivation may overstate the longer-term implications of tillage on soil structure, while measurements taken later in the season may underrepresent these effects. This highlights the need for careful consideration of sampling timing when assessing the structural effects of tillage practices.

While this study was able to quantify key structural changes using established image analysis techniques, the complexity and variability of soil structure, particularly in response to wetting-drying cycles and biological activity, suggest that additional structural indicators may exist which are not easily captured using conventional methods. The large datasets generated by X-ray CT scanning provide a potential wealth of unused structural information beyond basic porosity and pore size measurements. In this context, machine learning offers significant potential to extract new structural features, recognise subtle patterns, and automate the analysis of these complex image datasets. The application of machine learning techniques in the following chapter will explore whether such methods can provide additional insights into the structural dynamics of soils under different tillage treatments and potentially identify previously unrecognised patterns of structural development throughout the growing season.

5. The use of machine learning to identify structural changes during the wetting-drying cycle of soils.

5.1 Introduction

5.1.1 The wetting and drying cycle of soils

The wetting and drying cycle of a soil is of great importance to the physical structure of a soil. These cycles influence the stability, formation, and degradation of soil aggregates (Telfair *et al.* 1957), which are critical to maintaining structure, porosity, and water movement within soil systems. There have been many studies which have demonstrated that repeated wetting and drying can both enhance and degrade aggregate stability depending on the frequency, duration, and intensity of cycles, as well as the soil's composition and organic matter content (Denef *et al.* 2001, Le Bissonnais 2016). Earlier work by Telfair *et al.* (1957) and Dexter *et al.* (1984) established that wetting-drying cycles directly impact the strength and integrity of aggregates, particularly in clay-rich soils, where the shrink-swell behaviour controls structural change (Wang *et al.* 2018).

During a wetting event, water infiltrates and displaces air within soil pores and networks, leading to slaking and the potential disintegration of weak aggregates due to rapid hydration and trapped air. Slaking refers to the rapid breakdown of soil aggregates when dry soil is suddenly exposed to water (Le Bissonnais 2016). When drying occurs, capillary forces draw particles closer together, potentially reforming aggregates, but also increasing the risk of compaction and cracking, especially in soils with high clay content (Peng *et al.* 2016, Sayem *et al.* 2016, Qi *et al.* 2022). Repeated cycles of swelling (wetting) and shrinkage (drying) in clay soils often leads to the formation of visible surface cracks, which can alter infiltration pathways, impact root growth, and contribute to preferential flow of water, and therefore solutes (Chertkov 2002, Braudeau *et al.* 2014). This preferential flow can result in uneven wetting, reduced water retention, and the rapid transport of nutrients, agrochemicals, or contaminants through the soil profile, potentially impacting both crop uptake and surrounding groundwater quality (Jarvis 2007).

In agricultural settings, these processes play a crucial role in determining soil workability, compaction risk, and the timing of tillage. Rajaram *et al.* (1999) highlighted that soil moisture history, specifically the stage of the wetting-drying cycle at the time of tillage, impacts the mechanical resistance of the soil, influencing seedbed preparation and subsequent root development. Research has also shown that these cycles influence the soil microbial community, with rewetting events often triggering pulses of microbial respiration and nutrient mineralisation, referred to as the "Birch effect" (Fierer *et al.* 2002, Borken *et al.* 2009). This biological activity can further contribute to aggregate stability or breakdown, depending on how organic matter is being cycled and redistributed during wetting-drying stages.

The implications of these processes are becoming increasingly important under changing climatic conditions. Climate models predict greater variability in precipitation, with more intense rainfall as well as prolonged dry spells (IPCC 2023). These changes are expected to intensify the frequency and severity of wetting-drying cycles in many agricultural regions. Understanding how a soil responds structurally and biologically to such changes is essential for informing sustainable land management practices, including the selection of crop varieties, irrigation strategies, and tillage timing and management (Vogel *et al.* 2018). Further knowledge in this area can help mitigate risks and improve long-term soil resilience and health in the face of meteorological extremes.

5.1.2 The use of X-ray CT to study soil structure

X-ray CT provides a non-destructive and non-invasive method to observe soil structure in 3-D, allowing detailed investigation of the internal structure of aggregates and pores without disturbing the sample. This is a distinct advantage over many traditional structural assessment techniques, which often require destructive or damaging preparation steps such as resin or wax impregnation to stabilise soil samples for imaging (Dexter 1988). Other methods, such as VESS (Ball *et al.* 2007) provide useful information but alter the sample during preparation, which limits repeated measurements and the ability to monitor structural changes over time.

The ability of X-ray CT to generate high-resolution 3-D visualisations makes it especially valuable for studying the structural changes of soil under different environmental conditions and management practices. Structural properties such as aggregate size and distribution, pore connections, tortuosity and macropore geometry can be calculated and quantified directly from CT data (Schlüter *et al.* 2010, Vogel *et al.* 2010, Dal Ferro *et al.* 2013, Katuwal *et al.* 2015). These parameters are crucial for understanding key soil functions, such as hydraulic conductivity, gas exchange and root penetration (Ferreira *et al.* 2018, Soto-Gomez *et al.* 2018).

One important application of X-ray CT has been to examine structural changes in soil subjected to repeated wetting and drying cycles, which are common and natural in agricultural environments. As soils are exposed to these cycles, aggregates may weaken or change, and pore networks may be altered through shrinkage, swelling, or collapse. Studies have successfully used CT images to visualise pore network distribution, crack formation, and aggregate breakdown during wetting-drying processes, especially in clay-rich soils prone to shrink-swell behaviour (Ma *et al.* 2015, Pires *et al.* 2020, Zaidi *et al.* 2021). These structural changes change water infiltration, storage and transport within the soil, often leading to the development of preferential flow paths (Mooney 2002). Preferential flow allows water and solutes to bypass finer soil matrix pores, resulting in altered transport of nutrients, agrochemicals, or contaminants.

As well as observing structural changes, X-ray CT has also been applied to investigate how soil type can influence structural behaviour. For example, (Luo *et al.* 2010)compared pore structures across two different soil types, revealing differences in pore size distribution, macropore formation, and structural resilience under varying moisture conditions. Understanding the how different soil types are affected is important for soil management practices.

Beyond static imagery, X-ray CT also allows dynamic tracking of fluid within the soil, offering further insight into hydrological processes. Mooney (2002) demonstrated how X-ray CT can visualise and quantify soil macroporosity and potential flow pathways. Whilst most early studies focussed on static imaging, Grayling *et al.* (2018) combined CT imaging with tracer materials to directly observe solute transport, preferential flow pathways, and agrochemical movement through the pore network.

In summary, X-ray CT has become a very useful tool for investigating both the static and dynamic aspects of soil structure. The ability to capture high-resolution, 3-D information across different soil types and treatment practices provides critical insight into how soils respond to management and environmental stresses. This knowledge is central for developing effective strategies to maintain soil health, optimise agricultural productivity, and alleviate environmental risks under changing climate conditions.

5.1.3 Machine learning

The complicated and dynamic nature of soil structure makes it challenging to fully capture and analyse using conventional soil science methods alone. Recent advances in machine learning offer new opportunities to better quantify and interpret structural changes that occur during wetting-drying cycles, particularly when working with large and complex datasets, such as those generated from X-ray CT.

Machine learning encompasses a broad set of computational techniques that allow algorithms to learn patterns and relationships within data, without being explicitly programmed to predefined rules or categories. The algorithms can be broadly categorised into supervised learning, where models are trained using labelled datasets, and unsupervised learning, where models identify inherent patterns within unlabelled data. When applied to soil structural analysis, machine learning can identify features such as pore size distribution, aggregate formations, and crack propagation based on patterns learned directly from CT image data. In previous studies, algorithms have successfully predicted physical and chemical properties of soils, as well as estimating soil water retention properties. (Sarmadian *et al.* 2009, Li *et al.* 2014, Moreira de Melo *et al.* 2015).

In the context of structural changes during wetting-drying cycles, deep learning approaches show considerable promise for extracting new structural forms from CT image data. CNNs are specifically designed to recognise spatial features in image datasets and can be trained to detect patterns such as those associated with pore collapse, crack formation, and aggregate breakdown. Theses structural responses to wetting-drying are not always easily capture by basic pore volume measurements alone, but may be reflected in the geometry, distribution, and connectivity of the pore network that CNNs can quantify automatically (Wang *et al.* 2018).

Transfer learning has already been applied successfully in related soil imaging contexts, such as root-soil segmentation (Douarre *et al.* 2018) and clay content estimation from hyperspectral data (Liu *et al.* 2018), suggesting that similar approaches could be adapted to analyse structural changes observed from CT images of soils under wetting-drying cycles. Automated segmentation tools such as SoilJ (Koestel 2018) have demonstrated the feasibility of semi-automated image processing for soil columns, but neural networks may enable a higher level of automation and feature recognition.

The practicality of machine learning for this application comes from its ability to detect complex structural changes that may not be capture by simple volume-based metrics or human observation. Wetting-drying cycles induce changes in soil saturation, which directly impacts aggregate stability, pore connectivity, and crack development. By training machine learning models to associate structural changes with variations in saturation, it may become possible to develop predictive models that describe how different soils respond to repeated wetting and drying events under field conditions.

Although research applying machine learning directly to structural changes in wetting-drying soil cycles is limited, related work in geotechnical engineering has demonstrated the suitability of these methods for modelling soil deformation, strength and fracture processes (Farfani *et al.* 2015, Baghbani *et al.* 2022). Deep learning methods have also been used to detect cracks in medical imaging (Lindsey *et al.* 2018) and infrastructure monitoring (Cha *et al.* 2017), providing a strong basis for applying similar techniques to the analysis of structural changes in soils under wetting-drying cycles.

In summary, the integration of machine learning with X-ray CT presents a highly promising approach for the non-destructive detection and quantification of previously overlooked features of soil structural changes during wetting and drying cycles. By enabling more detailed and automated analysis of complex 3-D pore networks and structural changes, these tools may improve the ability to understand soil structure impacts under increasingly variable climate conditions.

5.1.4 Aims and objectives

The aim of this chapter is to explore whether machine learning can be used to automatically identify structural changes that occur in soils as they undergo wetting and drying cycles. Using X-ray CT imaging data, the potential for machine learning models to detect features such as changes in pore structure, aggregate formation and breakdown, and crack development will be tested. In addition to assessing the structural changes, the study will also investigate whether machine learning can classify soil types (such as sand and clay) based on differences in structure. Once these aspects have been examined, the application of machine learning will be developed to estimate soil water saturation based on structural characteristics visible in the CT data.

Hypotheses:

- Machine learning algorithms can be trained to identify and classify key components of soil structure, such as aggregate stability, pore network characteristics, and crack formation, independent of soil texture.
- Machine learning can be used to identify the stage of the wetting-drying cycle based on quantifiable structural features observed in the soil (e.g. pore volume, aggregate formation and breakdown, crack development).
- The application of machine learning to analyse soil structure and wetting-drying stages offers greater efficiency and objectivity compared to manual image analysis techniques.
- The application of machine learning may enable the discovery of previously unrecognised structural indicators or patterns within wetting-drying cycles, contributing novel quantitative data to the understanding of soil structural dynamics.

5.2. Methodology

5.2.1 Sample collection and imaging

This study uses both pre-existing and newly collected soil sample data. The first dataset is from (Tracy *et al.* 2015), who conducted a detailed study at the University of Nottingham farm at Bunny, Nottinghamshire, UK (52.52°N, 1.07°W). The site displays a gradient in soil texture, ranging from a sand-heavy loam to a clay-heavy loam. Tracy *et al.* (2015) selected two representative points along this gradient, which they referred to as “sand” and “clay” samples, respectively. These designations are continued in this study.

To expand upon this data and validate the machine learning model, a new set of samples (six sand and six clay) were obtained from the same locations. Unlike the Tracy *et al.* (2015) samples used in the model, these newly collected soils were of unknown water content, intended for use in blind testing to evaluate how well the trained model could classify unseen data. However, due to the model's limited performance in classifying known moisture states from the Tracy *et al.* (2015) dataset, these blind predictions would not be accurately interpreted.

Tracy *et al.* (2015) classified their soil textures via particle size analysis, resulting in the profiles shown in Table 5.1. The new samples were similarly analysed and are shown in Table 5.2.

	Sand (%)	**Clay (%)**	**Silt (%)**	**Texture**
Tracy *et al.* "Sand"	83	13	4	Sandy Loam
Tracy *et al.* "Clay"	36	33	31	Clay Loam

Table 5. 1: Particle size analysis of soil samples from Tracey et al. (2015's).

	Sand (%)	**Clay (%)**	**Silt (%)**	**Texture**
New "Sand" Samples	57	32	11	Sandy Clay Loam
New "Clay" Samples	7	56	37	Clay

Table 5.2: Particle size analysis of newly collected soil samples.

Tracy *et al.* (2015) subjected their samples to a controlled water retention process during scanning. The soil cores were saturated with distilled water and then held at specific matric potentials (measured in kilopascals, kPa) using a small vacuum chamber. Matric potential refers to the tension with which water is being held within the soil sample, with 0 kPa representing full saturation, and increasingly negative values corresponding to gradually drier conditions. The lowest matric potential used was -75 kPa, representing very dry conditions, particularly in the sandy soils.

In contrast, the newly collected samples used for testing in this study were scanned in their field state without additional wetting or drying treatments. These samples were intended to function as independent, blind tests of the model's classification capabilities on unprocessed data and so were left "as is" following extraction from the field.

Samples were scanned using a Phoenix Nanotom 180NF (GE Sensing & Inspection Technologies GmbH, Wunstorf, Germany). Voxel resolution was set to 50 μm, with a maximum energy of 160 kV and a current of 180 μA.

Once the X-ray CT images had been reconstructed, VG StudioMAX®2.2 (Volume Graphics GmbH, Heidelberg, Germany) was used to save each scan as an image stack, with each stack having approximately 1,800 images. Each stack was individually loaded into Image J (Schindelin *et al.* 2012), starting with image 1 and in increments of 2 to reduce final number of images. A large amount of each image contained the blank space around the sample in the scanner, as well as the plastic tube around the soil core (Figure 5.1). A rectangular selection was made and cropped to inside the core so that the final images would only show the soil. Images were then cropped to a square format and resized to 256 x 256 pixels for input into the neural network. This ensured that all images had consistent dimensions, regardless of their original aspect ratio. The same process was repeated for each image stack. With Tracy et al 2015's data as well as the new soil samples scanned for this paper, this was a total of 88 image stacks cropped and squared ready for use in the neural network.

The training image set for each model was randomly split into 70% training and 30% validation.

Image	Description
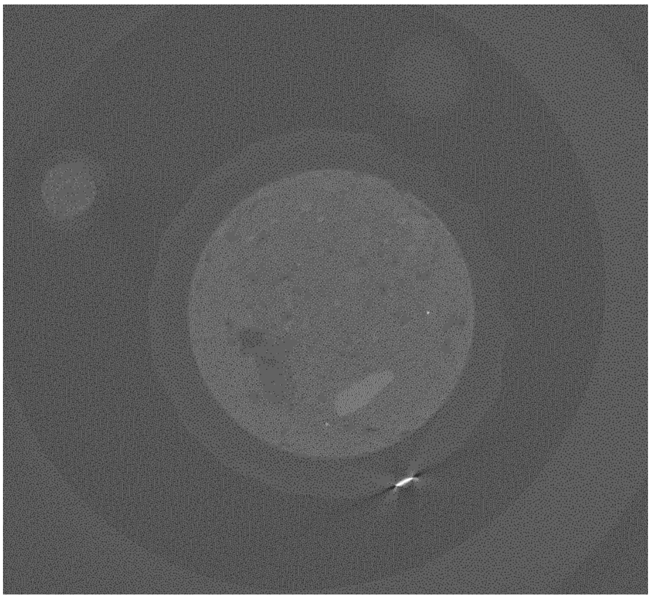	Image of Clay at -0 kPa, straight from the scanner. The light-grey circle in the middle is the soil. The slightly dark circular outline around this is the plastic tube in which the sample is held. Surrounding the rest of the image is blank space.
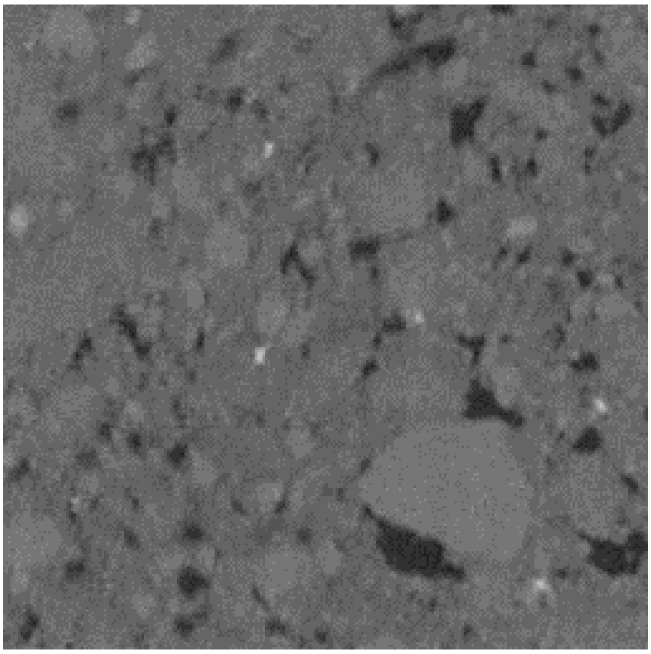	Each image was cropped and squared to just inside the circle containing the soil. Images were also resized to 256px by 256px for use in the network.

Fig. 5.1: Image processing for scanned data before being put through the CNN.

5.2.2 Machine Learning Models

Two different neural networks were used and evaluated in this study:

a. **Custom CNN**

   A personalised CNN was constructed using TensorFlow and Keras, designed to balance computational efficiency with classification performance. The architecture includes three convolutional blocks with increasing filter depths and batch normalisation, ending with a dense classification layer modified to match the number of soil classes (2 or 14).

The model was compiled with the Adam optimiser and categorical cross-entropy loss. Class weights were calculated and applied for each experiment to account for class imbalance.

Compared to the EfficientNetV2 model described below, which contains over 300 layers, the custom CNN developed in this study is substantially shallower, consisting of only 9 trainable layers. This simpler architecture reduces computational demands and training time, though may limit the network's ability to capture highly abstract features.

Layer (type)	**Output Shape**
Conv2D (32)	(224, 224, 32)
MaxPooling2D	(112, 112, 32)
Conv2D (64)	(110, 110, 64)
MaxPooling2D	(55, 55, 64)
Conv2D (128) x 2	(53, 53, 128)
MaxPooling2D	(26, 26, 128)
Flatten	-
Dense (128)	
Output (Dense)	[Number of classes]

Table 5.3: Custom CNN structure.

b. **EfficientNetV2 B0 (Transfer Learning)**

To benchmark the performance of the custom CNN, a pre-trained EfficientNetV2 B0 model (Tan *et al.* 2021) was also used. This model was chosen due to its higher performance on image classification tasks whilst also exhibiting relatively efficient training times and computational usage.

The final classification layer was replaced with a dense layer matching the number of soil classes. Input images were resized to 224x224 to match the EfficientNetV2 architecture.

The use of EfficientNetV2 also allows evaluation of the trade-offs between model complexity, and therefore environmental cost and classification accuracy. While the EfficientNetV2 models are much more efficient with computational resources than other available models, they still require significantly more computational resources than the custom-built model. However, the EfficientNetV2 models generally outperform simpler architectures in cases with subtle features.

5.2.3 Saliency Maps

Saliency maps were used as a method to visualise the network and identify what patterns or features the network was using to classify images. These images are gradient-based heatmaps that highlight areas of each image the model relied on during prediction.

For each image, the gradient of the predicted class score with respect to the input pixels was calculated. TensorFlow's GradientTape function was used to compute the gradients. These gradients were then processed to create a single 2-D saliency map image, which was normalised and overlaid using a colour scale. Warmer colours indicate regions (patterns, edges, textures etc.) that were more influential in classification.

The final image (e.g. Figure 5.2) displays the original soil sample image alongside its saliency map, with the predicted class label shown above. These maps help to interpret the model's focus and verify whether the predictions (correct or otherwise) are being made on meaningful structural features in the soil.

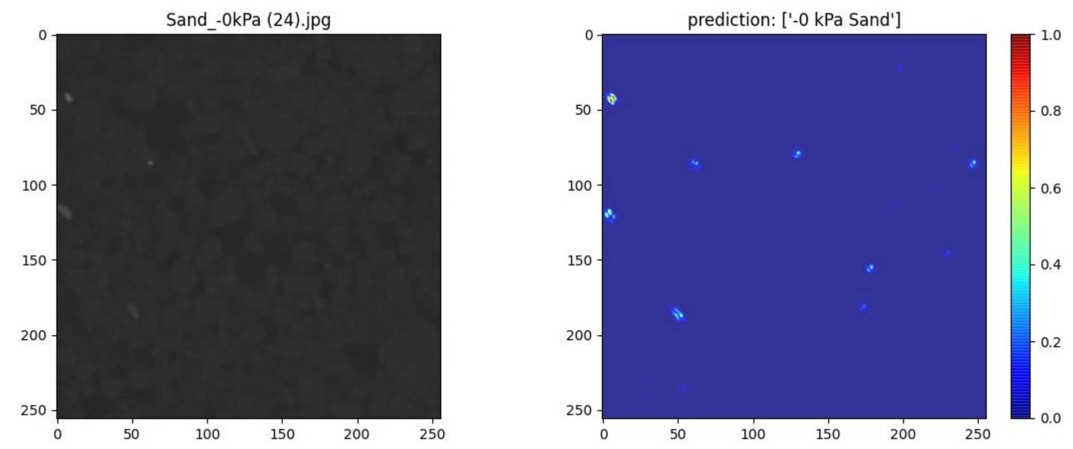

Fig. 5.2: A saliency map (right) showing what features were used to predict the soil type and water content of the image on the left. Blue pixels show they were not used in identification, whilst the brighter pixels show how useful those pixels were with the prediction.

5.2.4 Approach to the classification problem

The dataset used in this study compromised two soil types (sand and clay) and seven levels of water content, represented by matric potential values (kPa): -0kPa, -5kPa, -10kPa, -20kPa, -40kPa, -60 kPa, -75kPa). These conditions resulted in 14 distinct classes (2 soil types x 7 water contents).

To simplify the classification task and evaluate the baseline capabilities of each network, a series of binary classification problems were first conducted. These included:

- Sand (wet) vs Sand (dry)
- Clay (wet) vs Clay (dry)
- Sand (wet) vs Clay (wet)
- Sand (dry) vs Clay (dry)

Where "wet" is used to describe the -0 kPa samples, while "dry" is used to describe the -75 kPa samples.

These two-class problems allowed for focused testing of how well each model could distinguish between soil type or water content when the other variable was held constant. This step also helped to identify the strengths and limitations of each network before attempting the more complicated 14-class classification task, where the network was required to distinguish between both soil texture and moisture level simultaneously.

5.3 Results

5.3.1 Binary Classification: Sand vs Clay (Custom CNN)

Two binary classification tasks were conducted to assess the ability of the custom CNN to distinguish between soil textures (sand and clay) under controlled moisture conditions. These experiments only used the 0 kPa (wet) and -75 kPa (dry) samples to isolate the influence of soil type from water content.

Wet comparison (Sand vs Clay at 0 kPa):

The model achieved an overall accuracy of 63%, with a precision of 0.67 for clay and 0.61 for sand. While the network showed some ability to extinguish between the two textures, performance was limited, particularly in correctly identifying clay samples, with a recall of only 0.51. This suggests that, under saturated conditions, the structural features differentiating sand and clay may appear less distinct or may be less consistently captured by this model.

Dry comparison (Sand vs Clay at -75 kPa):

Performance improved greatly when classifying dry samples, with an overall accuracy of 87%. The model achieved high precision and recall across both classes, with clay (precision 0.75, recall 0.91) and sand (precision 0.95, recall 0.84). These results indicate that drying accentuates the textural and structural differences between sand and clay (which we know to be true from traditional analysis of X-ray CT images), making them more separable by the CNN.

Overall, the results suggest that the model is more capable of distinguishing between soil textures when moisture content is low, likely due to enhanced visual contrast in pore structure, or additional structural features such as cracking.

5.3.2. *Binary Classification: Wet vs Dry (Custom CNN)*

To investigate the ability of the model to distinguish between wet and dry soils within the same texture class, two further binary classification tasks were carried out, using only wet or dry samples. These were tested on unseen data from Tracy *et al.* (2015) to assess model generalisability.

Sand (wet vs dry):

The model achieved an overall accuracy of 65%, with balanced performance across both classes (precision and recall both around 0.65). While not highly accurate, these results suggest that the network was moderately able to differentiate between saturated and dry sand soil samples. The improved recalled for the dry class (0.76) may indicate that the structural changes you'd expect to see in drying sand soil, such as increased visible porosity due to the poor water retention found in sandy soils, were detectable by the network.

Clay (wet vs dry):

Performance on the clay samples was less successful. While the network achieved a high recall (1.00) and reasonable precision (0.62) for the wet class, it failed entirely to recognise dry clay samples, achieving 0% recall and precision. This suggests that the model overfitted to features associated with saturated clay and did not learn to generalise well to the drier clay images. It is possible that the drier clay samples exhibited more subtle structural patterns that were underrepresented in the training data, or that the visual textural changes at this level of drying were not sufficiently distinct enough for the model to effectively learn.

These results suggest that water content influences model performance differently depending on soil texture. The model was successful at identifying water-related structural changes in sand, but struggled with clay, potentially due to less visible structural variations at these saturation levels, or a class imbalance within the dataset.

A summary of performance for all four binary classification experiments using the custom CNN model is shown in Table 5.4.

5.3.3 Binary Classification: Sand vs Clay (EfficientNetV2 B0)

Efficient Net was used to evaluate how well a high-capacity transfer learning model could distinguish between soil textures at different saturations.

Wet comparison (Sand vs Clay at 0 kPa):

EfficientNetV2 B0 achieved an overall accuracy of 52% on saturated samples. While the recall for clay was high (0.87), showing a strong sensitivity to saturated clay features, the model performed poorly on sand (0.17 recall), frequently misclassifying it as clay. This pattern was consistent across the predictions of unseen data. These results suggest that under wet conditions, the visual features of clay dominated the network's predictions, possibly due to differences in density or texture that the model overly prioritised during training.

Wet comparison (Sand vs Clay at -75 kPa):

Model performance on the dry samples was worse than on wet. The model achieved 34% accuracy, with 0.00 precision and recall for the sand class. All sand samples were misclassified as clay, leading to a severely skewed confusion matrix. While clay recall was perfect (1.00), this came at the expense of any ability to identify the sand class. This suggests significant overfitting to clay-like features, even with class weighting used, or an inability of the pre-trained network to extract discriminative features from dry sand images.

5.3.4 Binary Classification: Wet vs Dry (EfficientNetV2 B0)

EfficientNet V2 B0 was also evaluated on the binary classification task of distinguishing wet and dry soil conditions within each soil type.

Sand (wet vs dry)

For the sand samples, the model achieved an overall accuracy of 65%, comparable to that of the custom CNN. However, performance was unbalanced. While the model achieved 96% recall for dry sand, it showed a poor recall of 29% for wet sand. This skew indicates a bias toward predicting dry conditions even when class weighting was used, likely due to more prominent structural features. The low recall for the wet class may indicated subtle structural features that were not strongly captured by the model.

Clay (wet vs dry)

EfficientNet achieved an overall accuracy of 62%. However, this result was based entirely on the performance for the wet samples, achieving a precision of 0.62 and recall of 1.00, respectively. Even with class weighting, the model failed to correctly identify any of the dry clay samples, with a precision and recall of 0.00. This suggests that the model may have overfitted to features associated with saturated clay or was unable to identify consistent structural features in the dry clay.

These results show that while EfficientNet was able to classify dry sand well, its performance on wet clay and wet sand was limited, highlighting challenges in generalising across subtle structural variations in moisture conditions, especially in clay.

Experiment	Accuracy	Precision	Recall
Sand vs Clay (Wet)	63%	Clay: 0.67 Sand: 0.61	Clay: 0.51 Sand:0.75
Sand vs Clay (Dry)	87%	Clay: 0.75 Sand: 0.95	Clay: 0.91 Sand: 0.84
Sand (Wet vs Dry)	65%	Wet: 0.65 Dry: 0.66	Wet: 0.52 Dry: 0.76
Clay (Wet vs Dry)	62%	Wet: 0.62 Dry: 0.00	Wet: 1.00 Dry: 0.00

Table 5.4: Results of the four binary classification experiments using the custom CNN model, tested on unseen data.

Experiment	Accuracy	Precision	Recall
Sand vs Clay (Wet)	52%	Clay: 0.51 Sand: 0.58	Clay: 0.87 Sand:0.17
Sand vs Clay (Dry)	34%	Clay: 0.34 Sand: 0.00	Clay: 1.00 Sand: 0.00
Sand (Wet vs Dry)	65%	Wet: 0.87 Dry: 0.62	Wet: 0.29 Dry: 0.96
Clay (Wet vs Dry)	62%	Wet: 0.62 Dry: 0.00	Wet: 1.00 Dry: 0.00

Table 5.5: Results of the four binary classification experiments using the custom EfficentNetV2 B0, tested on unseen data.

5.3.5 Multi-Class Classification: Soil Type and Water Content (Custom CNN)

After successfully training the network on the binary classification tasks, a more complex situation was trialled – distinguishing between 14 classes representing combinations of soil texture (sand and clay) and water content (classified by matric potential). Each soil type was associated with seven matric potentials (0, -5, -10, -20, -40, -60, and -75 kPa), resulting in 14 unique classes.

During training, the model demonstrated strong performance, with high accuracy, precision and F1 scores for most classes (Table 5.4). The average training accuracy reached 97.8%, with an average F1-score of 0.978, indicating the model was able to learn distinguishing features across the full range of classes. Several classes achieved perfect or near-perfect precision and recall during training, suggesting strong discriminative learning.

Class		**Precision**	**Recall**	**F1-score**
Clay	0 kPa	1.000	0.899	0.947
	-5 kPa	0.972	0.979	0.975

	-10 kPa	1.000	0.968	0.983
	-20 kPa	1.000	0.961	0.980
	-40 kPa	0.948	0.993	0.970
	-60 kPa	0.906	0.973	0.939
	-75 kPa	0.987	0.987	0.987
Sand	0 kPa	0.919	1.000	0.958
	-5 kPa	0.994	1.000	0.997
	-10 kPa	1.000	0.949	0.974
	-20 kPa	1.000	1.000	1.000
	-40 kPa	0.993	1.000	0.997
	-60 kPa	0.987	0.994	0.991
	-75 kPa	1.000	0.993	0.997
Average:		0.979	0.978	0.978
Accuracy:		0.978		

Table 5.6: Network training performance on 14 classes.

To fully examine model behaviour and usability, the network was evaluated on a set of previously seen images, to validate performance before testing on unseen data. Figure 5.3 displays the percentage of correct predictions across each of the 14 classes, based on 25 images per class. While clay classes performed consistently well across all water contents, sand classes showed a more variable performance. For instance, prediction accuracy for fully saturated sand (-0 kPa) reached 100%, but declined sharply for intermediate water contents, dropping to just 17% accuracy by -10 kPa.

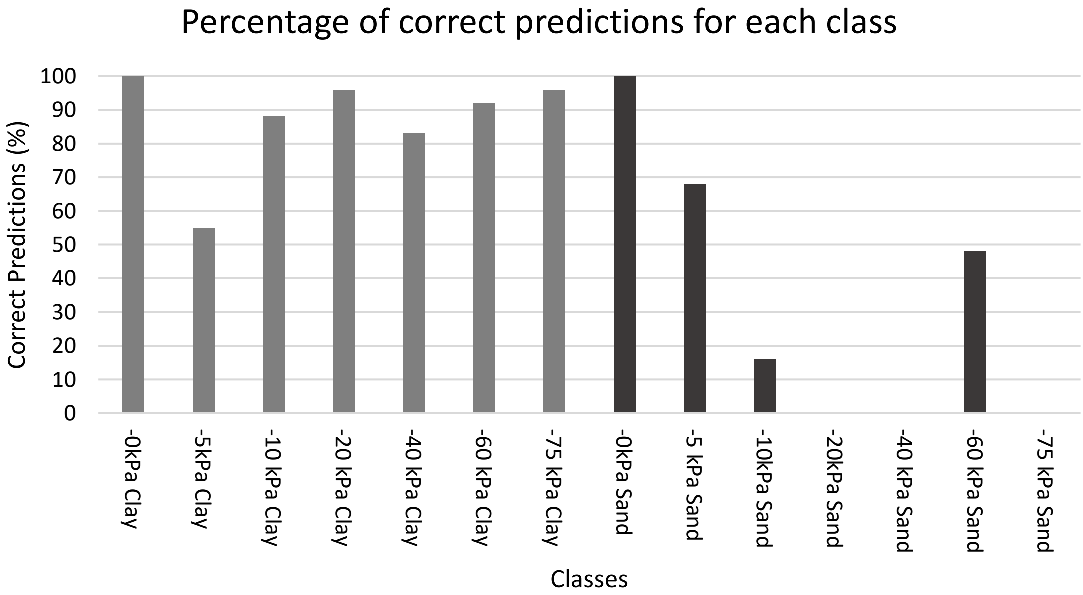

Fig. 5.3: Percentage of correct predictions for each class on previously seen data

This decline likely reflects the low water retention of sandy soils (Jackisch *et al.* 2020), which result in subtler structural differences between adjacent classes as the water content approaches very low levels. These findings align with the water retention patterns reported by Tracy *et al.* (2015), who observed a steep decline in water content for sandy soils between 0 kPa and -10 kPa, compared to a more gradual release in clays (Figure 5.4). As soil dries, especially in sandy soils, structure may appear more visually similar across multiple matric potentials, making classification more difficult.

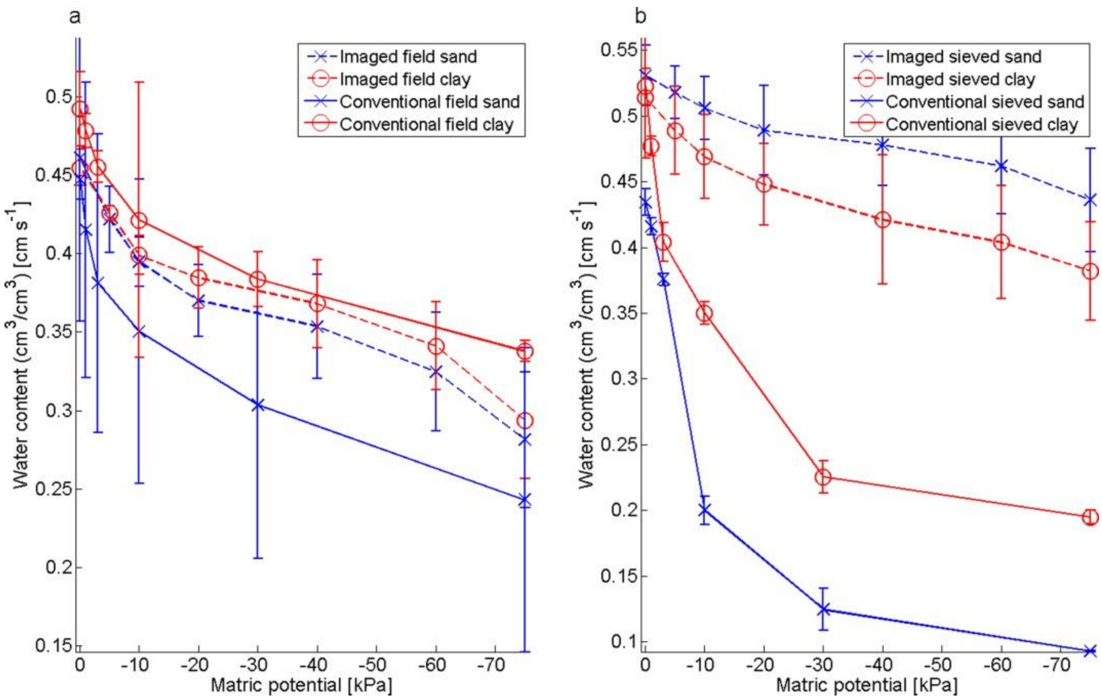

Fig. 5.4: Water release curves for (a) sieved sand and clay and (b) field-structured soils. Note the different scaling on the y-axis. (Tracy et al. 2015).

Given the difficulty in distinguishing between -60 kPa and -75 kPa, and the poor predictive performance for -75 kPa sand, this class was removed from the dataset to trial how model performance would be impacted. While -60 kPa sand still posed challenges, with an accuracy of around 50%, removing -75 kPa helped reduce class confusion and improved model clarity, particularly for mid-range sand moisture levels.

The network was retrained on the now 13-class problem and tested using unseen images, with approximately 250 images per class. This updated class set-up displayed improved predictive performance for both clay and sand across the remaining classes (Figure 5.7). Prediction accuracy for clay remained close to 100% across all water contents, while accuracy for sand classes showed improvement, particularly in the mid-range matric potentials. Interestingly, approximately 80% of misclassified sand samples were predicted as clay, suggesting a potential overlap in pore structure under drier conditions.

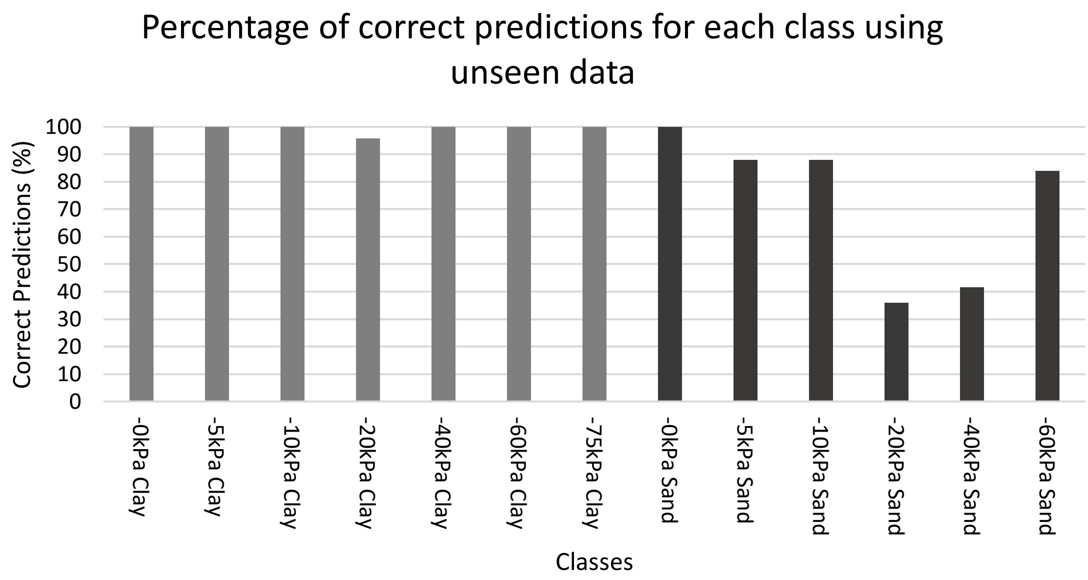

Fig. 5.5: Percentage of correct predictions for each class on unseen data after removing -75 kPa sand.

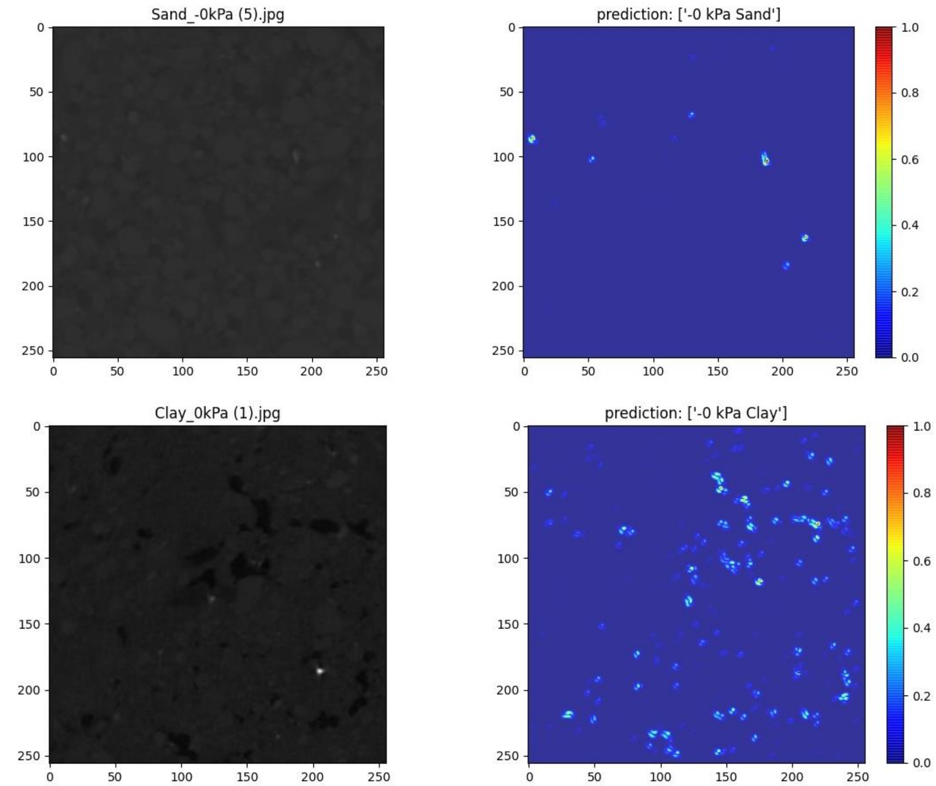

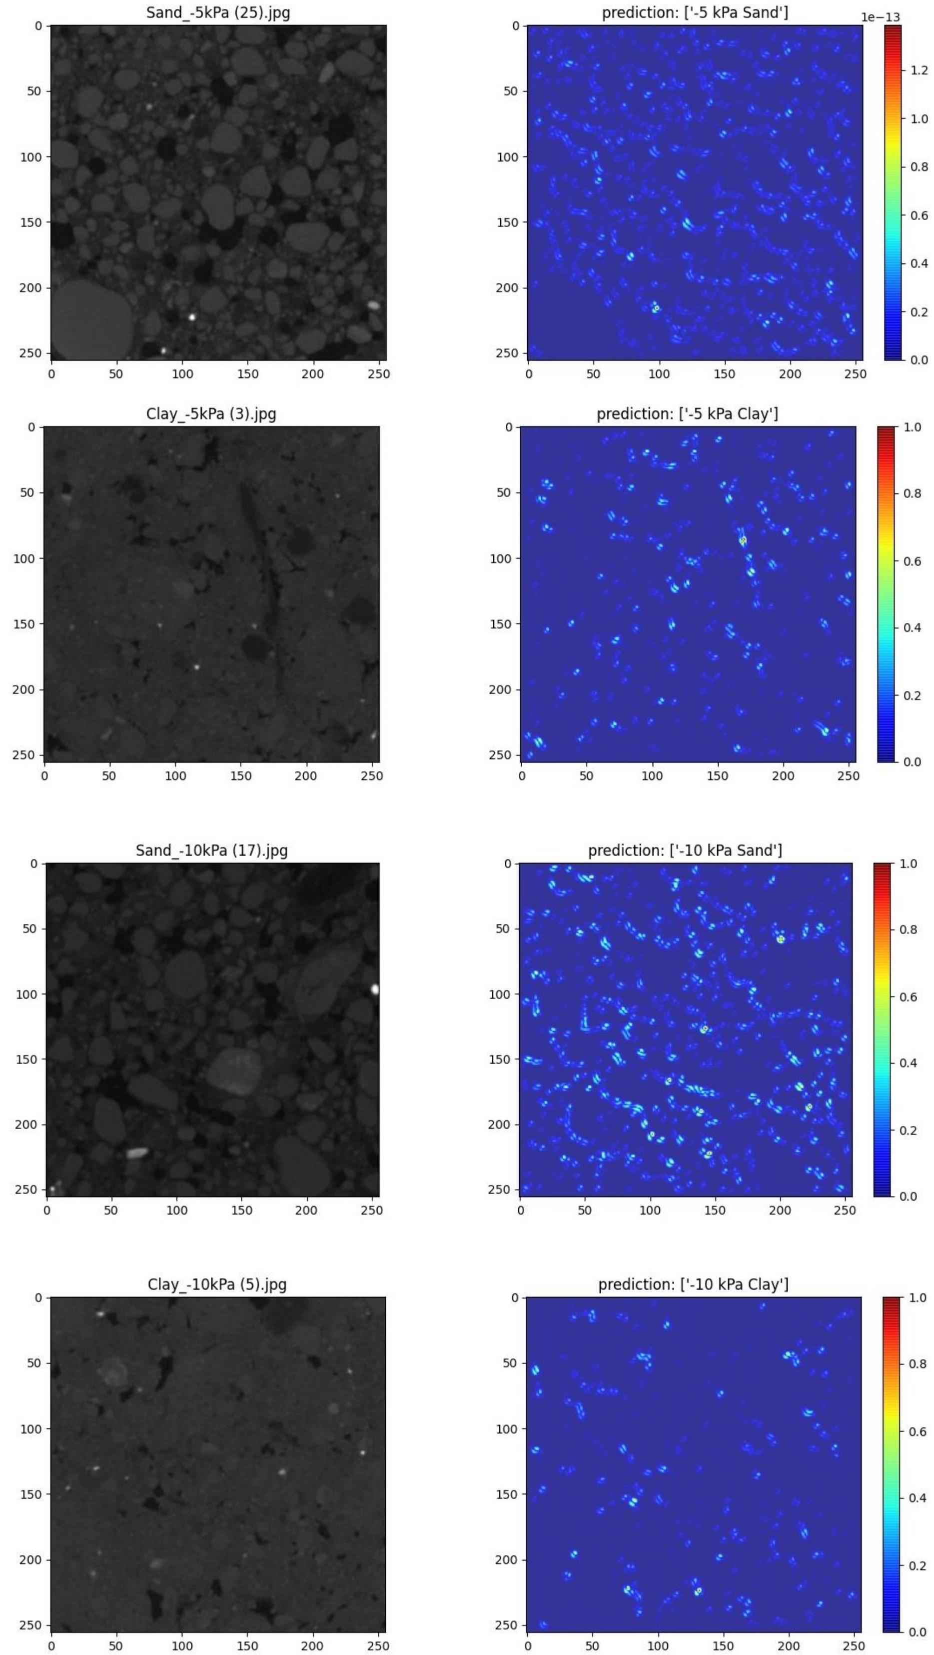
Sand_-5kPa (25).jpg
prediction: ['-5 kPa Sand']
1e−13
Clay_-5kPa (3).jpg
prediction: ['-5 kPa Clay']
Sand_-10kPa (17).jpg
prediction: ['-10 kPa Sand']
Clay_-10kPa (5).jpg
prediction: ['-10 kPa Clay']

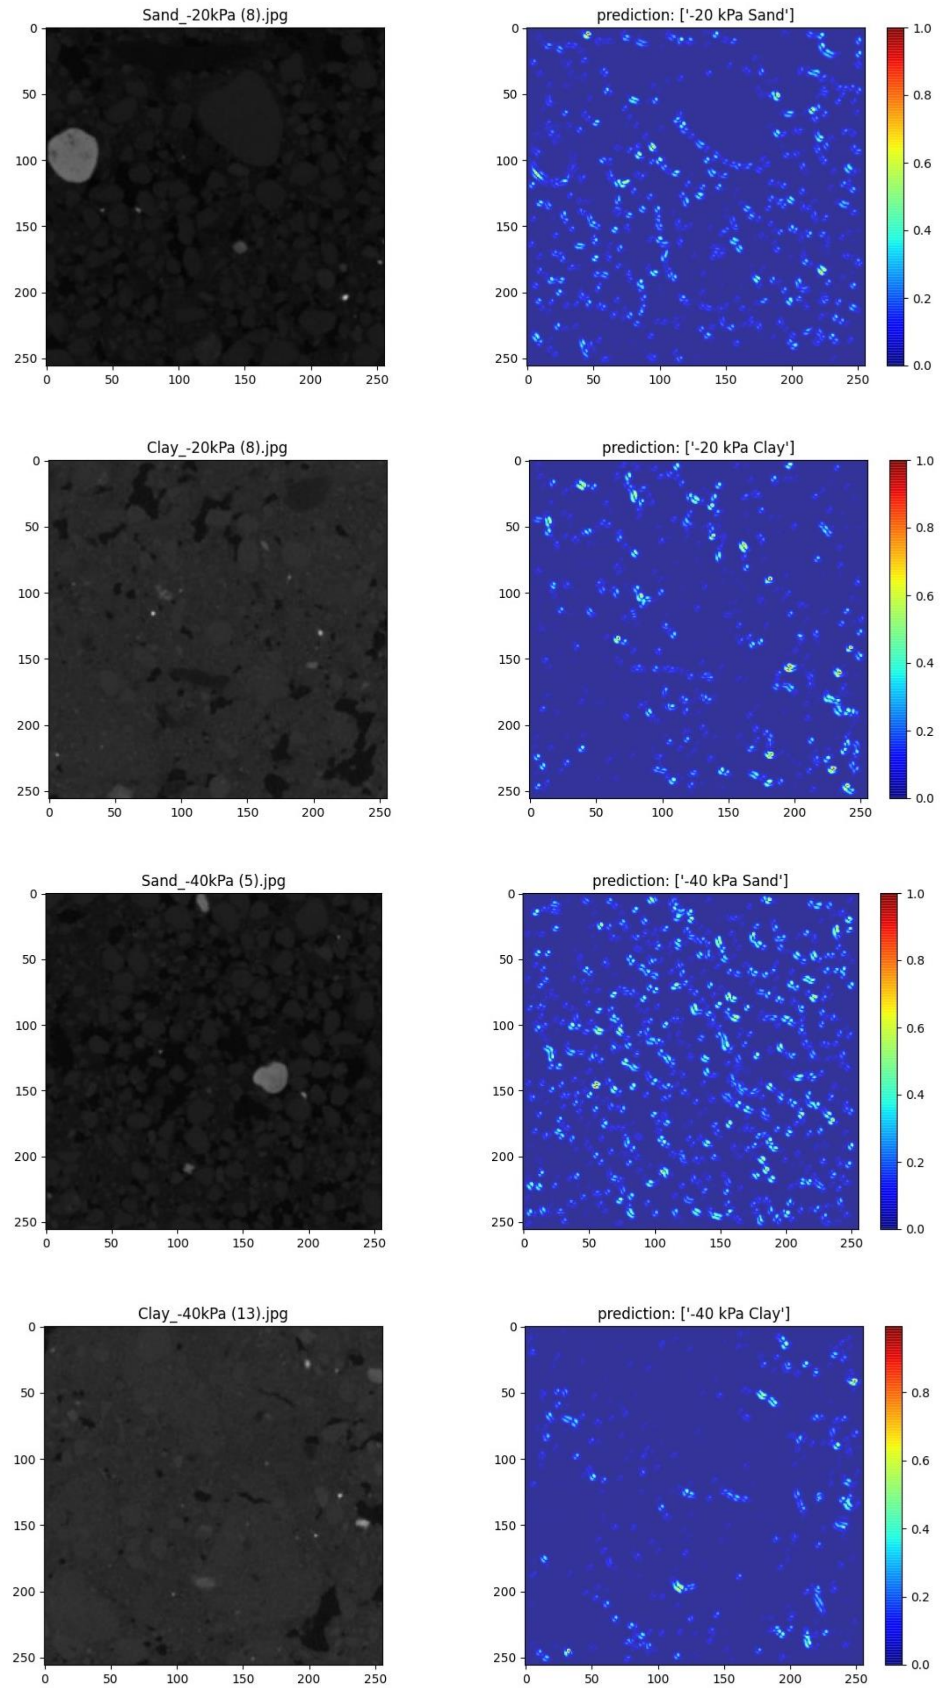
Sand_-20kPa (8).jpg
prediction: ['-20 kPa Sand']
Clay_-20kPa (8).jpg
prediction: ['-20 kPa Clay']
Sand_-40kPa (5).jpg
prediction: ['-40 kPa Sand']
Clay_-40kPa (13).jpg
prediction: ['-40 kPa Clay']

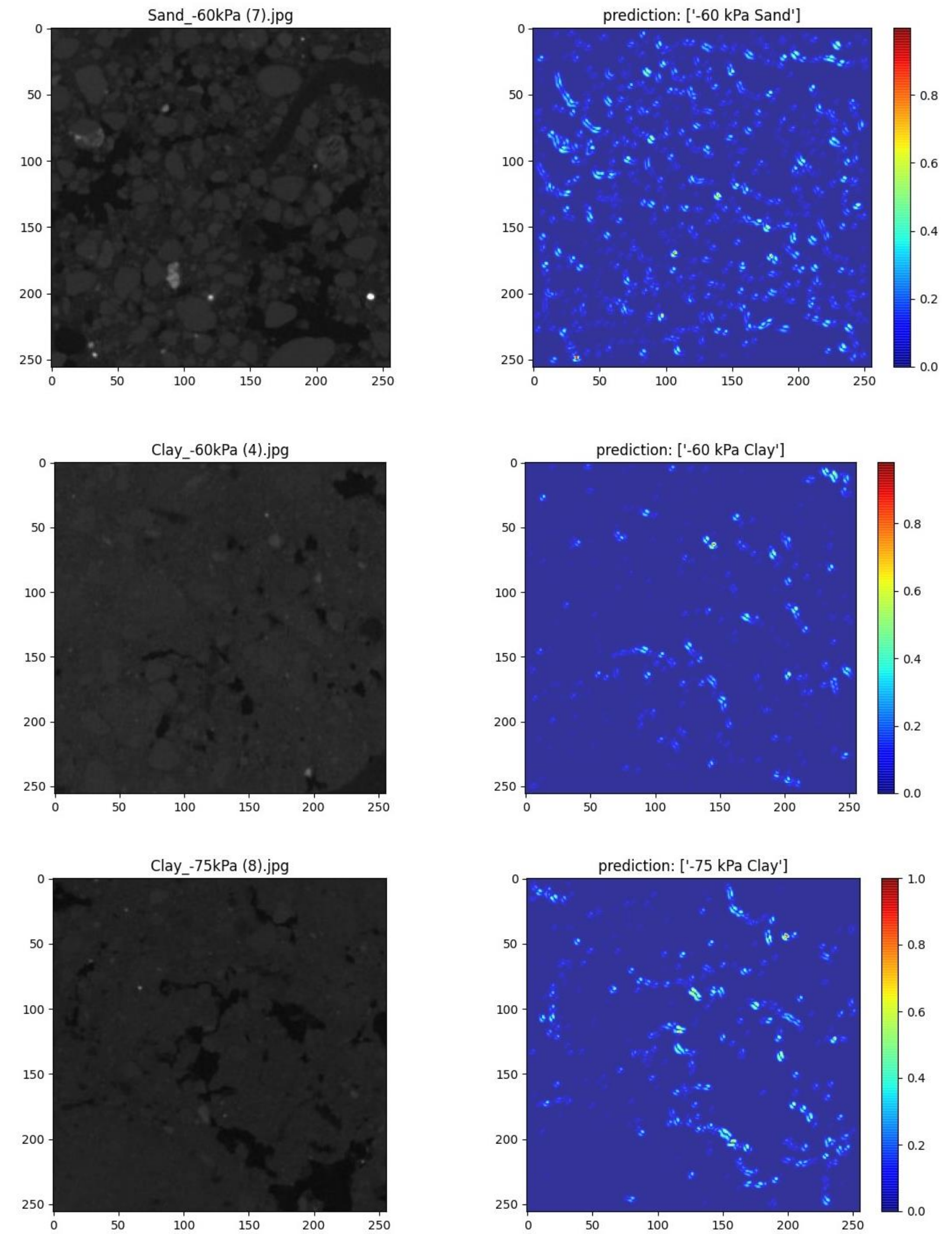

Fig. 5.6: Saliency maps showing one correct prediction for each of the 13 classes presented here.

Figure 5.5 displays one correct prediction for each of the 13 classes used in the neural network. Overall, the saliency maps for the sand images seem to be a lot busier than those for clay, displaying a lot more brightly lit areas on the map than shown for clay.

5.3.6 Multi-Class Classification: Soil Type and Water Content (EfficientNet V2 B0)

The same 14-class soil classification task was repeated using EfficientNetV2 B0. The model was trained on the same dataset to classify the same combinations of soil type and water content.

Training performance appeared promising initially, with training accuracy steadily increasing to over 93% by epoch 10 (Table 5.6). However, validation accuracy remained extremely low throughout, never exceeding 8.7%, and the final text accuracy dropped to just 8%, indicated the model failed to generalise beyond the training data. This discrepancy between training and validation performance indicates overfitting, where the model memorises training features but failed to learn meaningful patterns applicable to new samples.

Epoch	**Training Accuracy (%)**	**Validation Accuracy (%)**	**Validation Loss**
1	35.3	8.7	5.91
5	84.6	8.5	153.68
10	93.9	8.1	17.89

Table 5.7: EfficientNetV2 B0 training summary for the 14-class problem.

Despite its deeper and more advanced architecture, the model struggled to correctly classify the test images. Only two classes, Sand -60 kPa and Sand -75 kPa, received any non-zero recall, with sand -60kPa dominated predictions across nearly all inputs. This is most clearly shown in the confusion matrix (Figure 5.7), where almost every sample, regardless of actual class, is misclassified as Sand -60 kPa. While precision for this one class reached 0.08 and recall 0.88, the model entirely failed to recognise any of the remaining 13 classes.

This suggests a failure in training generalisation rather than an issue with data quantity, as class distributions were balanced and test data came from the overall dataset. It is likely that EfficientNetV2 B0 was ill-suited to this particular problem, even with data augmentation and class weighting.

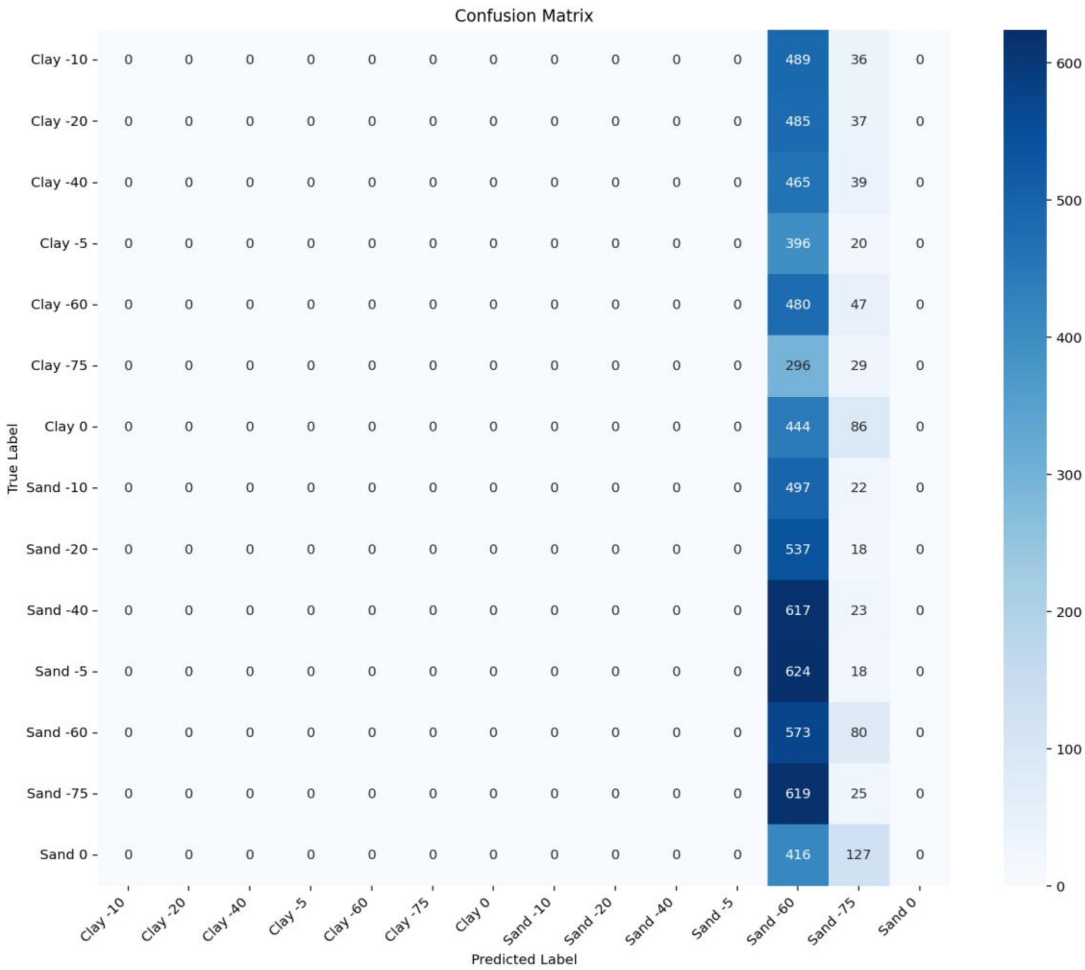

Figure 5.7: Confusion matrix showing prediction distribution of the EfficentNetV2 B0 model.

These results highlight the importance of model complexity matching the required task. In contrast to the custom CNN used, EfficientNet showed very poor performance on this multi-class problem without a lot of additional customisation and optimisation.

5.4 Discussion

This chapter set out to evaluate the potential of machine learning to classify soil type and moisture content using X-ray CT images. A custom-built CNN was developed to perform both binary and multi-class classification tasks. Its performance was compared against EfficientNetV2B0, a highly popular, pre-trained model designed for large-scale image classification. Across six experiments, both models were tested on unseen data, providing a robust comparison of generalisability and reliability.

The custom CNN demonstrated a moderate to strong performance across all binary classification tasks. It was especially strong at distinguishing between wet and dry sand, with 65% accuracy and balanced precision and recall. In contrast, performance was slightly weaker when classifying wet and dry clay (62%), as the model failed to identify dry clay samples at all. Based on soil physics, this could likely be because clay soils retain moisture more effectively than sand, and often exhibit less visually obvious structural differences during gradual drying, making classification more difficult (Jackisch *et al.* 2020). The models recall for saturated clay (1.00) and complete failure for dry clay (0.00) suggests that overfitting to the wet samples may have occurred.

The custom CNN also showed promising results in the texture classification tasks. When comparing saturated samples (sand and clay at 0 kPa), accuracy reached 60%, although performance was skewed towards correct clay predictions, with the model struggling to consistently identify saturated sand. In contrast, performance in dry texture classification (sand vs clay at -75 kPa) reached 87% accuracy, showing that the structural differences between soil types become more distinct under drier conditions. Overall, the results suggest that while water content can obscure textural features, it also introduces additional structural variance that the network can identify.

The 14-class problem posed the most difficult task, requiring the model to correctly predict both soil type and specific water content levels. Surprisingly, the custom CNN achieved extremely high training performance (98% accuracy), with strong precision and recall across nearly all classes. However, initial testing on seen data revealed a clear imbalance, with clay predictions remaining high across all water contents, but sand performance dropping significantly beyond 0 kPa. This was especially noticeable between -10kPa and -60 kPa, aligning with known water release curves for sand, where water loss occurs more easily than in clay. Upon manual inspection of the images by experts in X-ray CT imaging of soil images, it was decided that the images in these drier classes (-60 kPa and -75 kPa) were very visually similar, likely contributing to prediction errors. Based on this, the -75 kPa sand class was removed. When the network was retrained, this resulted in improved accuracy across the remaining sand classes, especially in the mid-range. While removing problematic classes may seem debateable, the significant overlap in visual similarity and structure between -60 and -75 kPa sand provided a valid justification for this refinement, especially with the goal being to improve generalisability.

In contrast, EfficientNetV2 B0 underperformed across nearly all tasks. Despite high training accuracy (94% by the final epoch), test accuracy for the 14-class problem was 8%, and the classification report and confusion matrix revealed that nearly all classes were misclassified as sand at -60 kPa. This overfitting to a single dominant class highlights the limitations of applying a large, general-purpose model to a small, specific dataset without significant tuning and customisation. Similar issues were observed in the binary tasks. Although EfficientNet achieved a comparable 65% accuracy for wet and dry sand, this was skewed by a very high recall for the dry class, suggesting a poor ability to balance predictions across both categories. Its performance in classifying wet and dry clay was similar, with complete failure to detect the dry class. Textural classifications were also poor, with just 52% accuracy in the wet texture comparison and 34% for dry.

These results show a major limitation of using large, general-purpose transfer learning models on small, domain-specific datasets. While architectures like EfficientNetV2 offer strong theoretical performance, achieving reliable results with them can require substantial modifications and adjustments. These may include fine-tuning multiple layers of the network, adjusting learning rates, applying targeted data augmentation, and potentially providing a greatly expanded dataset. However, introducing these adjustments undermines the primary motivation for utilising a pre-trained model in this study: ease of use and reduced development time. In this study, the simper custom CNN, specifically tailored for this sort of CT data, not only performed better across most tasks, but also required fewer resources and demonstrated less overfitting.

The saliency maps offered great insight into the features used for classification. For clay samples, especially the wet samples, the model highlighted larger aggregates and pore structures as distinguishing features. In sand, particularly under drier conditions, the model focused more on pore networks. This suggests the network was learning relevant physical features, rather than overfitting to just noise, helping to support the credibility of the predictions it was making.

While this study presents a promising application for machine learning in soil structural analysis, it still has limits. Currently, this algorithm is mainly limited by its training data. As data was only sourced from one field, despite collections being several years apart, the network has a limited knowledge of soil to work with. If shown images from a soil sample from elsewhere, it might struggle to correctly identify soil type and water content. It would be beneficial to expand the data set with samples from a range of locations and ensure a large variety of soil textures are included. The clay shown here may be visually very different to a similar soil texture collected from a few miles away.

With improvements, this sort of technology being able to accurately predict soil structure for soils at different water contents could be highly useful. Being able to evaluate the soil type of a sample, despite moisture levels, could help farmers make decisions about which crops to grow with which treatments. By picking the best crops for the soil type and applying optimal treatments such as cultivation and fertiliser based on the soil data, could greatly increase crop yield as well as reduce run-off of excess products, thereby reducing costs and improving environmental impact.

5.5 Conclusion

This chapter has demonstrated that machine learning, specifically CNNs, can be effectively applied to X-ray CT images of soil to classify both soil texture and water content. The custom-built CNN showed strong performance across a range of binary and multi-class classification tasks, especially when predicting soil texture under dry conditions and distinguishing between moisture levels in sandy soil. While limitations were observed, especially in correctly identifying drier clay or similar moisture levels in sandy soils, overall, the model was able to detect meaningful structure features, demonstrating the feasibility of this approach.

In contrast, the "off-the-shelf" EfficientNetV2 B0 model struggled greatly, showing signs of overfitting and failing to generalise, particularly in the 14-class identification task. These results reinforce the value of developing specific models tailored to the data, especially when working with small and potentially structurally subtle datasets, such as soil CT images.

This technology has the potential for a range of applications in precision agriculture, environmental monitoring, and even construction. Machine learning models that can rapidly assess soil structure could support decision-making in irrigation scheduling, fertiliser application, and land suitability assessments. However, for these tools to become widely and commercially usable, significant further development is required.

A key limitation of this study is the restricted scope of the training dataset, which was sourced entirely from a single location. This limits the generalisability of the trained models to other soils with different compositions or structural characteristics. Future research should prioritise the collection of a more geographically diverse dataset, including soils from multiple regions and even climates. This would allow the development of a robust, general-purpose model capable of analysing a broader range of conditions within soil.

With the expected continued advancements in machine learning, this area of research has the potential to contribute meaningfully to sustainable land management. Combining high-resolution imaging with intelligent, adaptable algorithms may offer new ways to monitor, classify and manage soil health.

6. The use of machine learning to identify structural changes to tilled and zero-tilled soil over a growing season

6.1 Introduction

The results gathered in Chapter 4 demonstrated that soil structure is highly dynamic over the course of a growing season, with both tillage treatments exhibiting seasonal fluctuations in porosity and pore size. While conventional image analysis methods were effective for quantifying these broad structural trends, they rely on a limited set of pre-defined metrics, potentially overlooking more subtle or complex structural changes within the pore network. The large volumes of high-resolution X-ray CT data generated from this experiment may contain considerably more detailed information than can be fully captured through traditional analysis alone.

In this chapter, machine learning is applied to the same CT dataset to explore whether additional patterns or structural changes can be identified that were not apparent when using standard image processing techniques. Machine learning methods offer the advantage of being able to analyse complex, high-dimensional data without requiring prior assumptions about which features may be most important. This allows the algorithms to potentially detect structural indicators that may not have been identified using conventional porosity or pore size analysis.

Applying machine learning to soil structure analysis has potential practical information for agricultural management. A more detailed understanding of how soil structure evolves under different tillage regimes could support more informed decisions about crop management, soil health, water retention, and overall field productivity and output. By identifying seasonal trends across the growing season, machine learning models may help optimise future tillage timing or intensity, improving both agronomic and environmental outcomes.

In the longer term, the development of automated, machine-learning based image analysis techniques may also allow rapid processing of large CT datasets, reducing the time and manual input currently required for soil structure studies, and potentially making this sort of soil analysis more commercially available.

The objectives for this chapter are:

- Apply machine learning methods to X-ray CT data to identify structural changes in soil under conventional tillage and zero-tillage systems over a growth season.
- Evaluate whether machine learning can detect structural differences or trends that may not have been identified through conventional image processing methods.
- Assess the potential usefulness of machine learning in automating and enhancing soil structure analysis for future research and agricultural management approaches.

6.2 Methodology

6.2.1 Sample collection and treatment

The dataset used for this chapter is the same as that described in Chapter 4. Soil samples were collected, scanned using X-ray CT, and pre-processed following the procedures outlined in Sections 4.2.1 to 4.2.5.

In brief, soil cores were collected from conventionally tilled and zero-tilled plots at multiple times across a growing season and imaged using an X-ray CT scanner at a resolution of 40µm. Detailed descriptions of the sampling site, treatments, scanning procedures, and initial image processing are provided in Chapter 4.

In this chapter, the previously processed image datasets were used as inputs for machine learning analysis. The machine learning methods applied are detailed below.

6.2.2 Convolutional Neural Network (CNN)

A CNN was developed using TensorFlow to classify and analyse the soil structural images generated from X-ray CT scanning. The architecture of the network is shown in Table 6.1 and Figure 6.1. The design and structure of the CNN were selected following initial testing of several different model configurations in order to balance model performance, computational efficiency, and suitability for the available dataset size.

The network consists of 3 convolutional layers, each responsible for extracting increasingly abstract spatial features from the input images. The first convolutional layer applied 64 layers to detect low-level image features such as edges and textures. The second and third convolutional layers used 32 and 16 filters respectively to capture more complex and higher-level patterns within the soil structure images. Following the convolutional layers, a flattening layer was used to convert the now multidimensional feature maps into a one-dimensional vector suitable for classification. This vector was then passed to a dense, fully connected layer containing 12 output nodes, each one corresponding to the classes defined for the analysis task (6 months x 2 tillage treatments).

Layer (type)	**Output Shape**	**Param #**
conv2d (Conv2D)	(None, 254, 254, 64)	640
conv2d_1 (Conv2D)	(None, 252, 252, 32)	18464
conv2d_2 (Conv2D)	(None, 250, 250, 16)	4624
flatten (Flatten)	(None, 2032128)	0
dense (Dense)	(None, 12)	28449806

Table 6.1: TensorFlow CNN architecture.

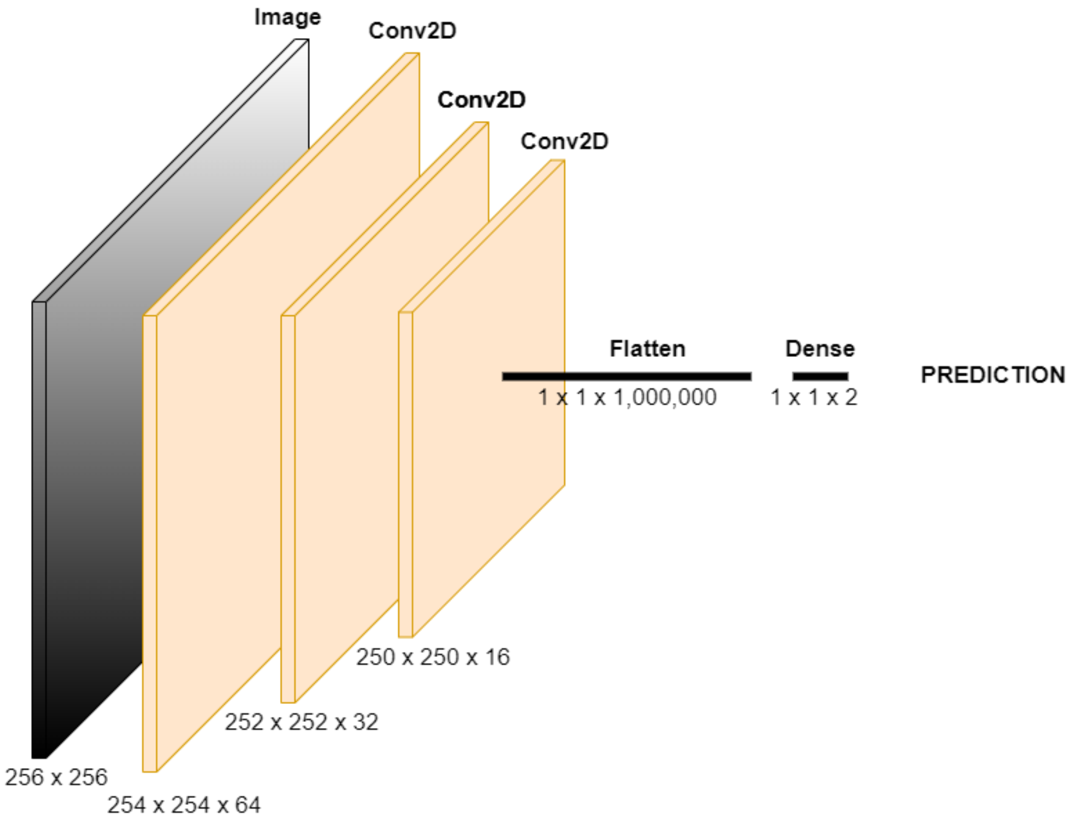

Fig. 6.1: Convolutional Neural Network (CNN) architecture visualisation.

Initial experimentation involved testing CNN architectures with both fewer and additional convolutional layers, along with different filter sizes and network depths. The final architecture was chosen as it produced stable classification accuracy. Adding additional convolutional layers was found to result in diminishing returns or increased overfitting.

The network was trained using a binary cross-entropy loss function, which was appropriate for the classification task being performed. Class weighting was applied to ensure equal weighting across classes, preventing bias towards the more frequently represented classes in the training data. Model optimisation was performed using the Adam optimiser with default TensorFlow parameters.

To increase the size and variability of the training dataset and improve generalisation, data augmentation techniques were applied during training. These included random horizontal and vertical flipping of images, as well as random rotations. This all helped introduce additional variability in spatial orientation without altering the structural features of the soil samples.

While several state-of-the-art CNN architectures exist for image classification tasks, such as ResNet, VGG, Inception, they were not utilised in this study. These networks are designed for very large, highly diverse image datasets, and require intensive training data, substantial computational resources, and careful adjustment and parameter optimisation to avoid overfitting. Given the relatively small size and specialised nature of the X-ray CT soil structure dataset, a custom-built, smaller-scale CNN was deemed more appropriate. In additional to the practical benefits of using a simpler model architecture, smaller networks also offer a lower computational and energy cost during training and inference. As machine learning models continue to grow in size and energy demand, there is an increasing recognition of their potential contribution to environmental impacts. Given that climate change, partly driving by rising global energy usage, it itself placing increasing pressure on agricultural systems, the development of more resource-efficient modelling approaches is both scientifically and environmentally desirable.

6.2.3 Dataset Preparation and Classification Design

The dataset used for this study was derived from the X-ray CT image data described in Chapter 4. Soil samples from both conventional tillage and zero tillage treatments were collected at six time points across the growing season (September, October, December, February, April, and June), producing a total of 48 samples (6 months x (2 treatments x 4 replicates)). After image processing and segmentation, as described in Chapter 4, these samples were converted into 2-D image slices which formed the input data for the training and testing of the neural network.

Two main classification approaches were implemented:

1. **Temporal classification (by sampling month):**
   In the first approach, all available samples from tilled samples, and the CNN was trained to classify each image based on the sample month. This model included six output classes, corresponding to the six months of the growing season. After evaluating the performance of this

approach, it was decided not to repeat this experiment using zero-till only data, as preliminary testing indicated limited additional benefit in separating the treatments for this task.

2. **Treatment classification (by tillage type within each month):** In the second approach, independent models were trained for each sampling month to classify tillage treatment. For each month, the network was trained to distinguish between conventional tillage and zero tillage based on soil structure present at that time point. This resulted in six separate binary classification models.

In both approaches, the dataset was divided into training, validation, and test subsets for model evaluation. Data augmentation, including random rotations and horizontal and vertical flips, was applied during training to increase data variability and reduce overfitting.

6.3 Results

6.3.1 Temporal classification (tillage images)

The CNN was initially trained on the full dataset containing just tillage samples, with the aim of classifying samples by sampling month based on soil structure alone. Table 6.2 shows the precision, recall and F-1 scores for each month, alongside the overall accuracy after 15 training epochs.

	Precision	**Recall**	**F1-score**
September	1.000	0.012	0.023
October	0.000	0.000	0.000
December	0.309	0.496	0.381
February	0.000	0.000	0.000
April	0.083	0.030	0.044
June	0.364	1.000	0.534
Average	0.270	0.328	0.534
Accuracy	0.33		

Table 6.2: Network performance on month classification task.

The overall performance on this task was poor, with an accuracy of 33%. The results indicate that the network struggled to reliably differentiate between sampling months based solely on the soil structure visible in the CT images.

For some months, such as October and February, the network was unable to correctly classify any images, resulting in zero precision, recall and F-1 score. For other months, performance varied. September, which represented the point immediately after tillage, achieved perfect precision (1.000), meaning that whenever the network predicted a sample as September, it was always correct. However, recall for September was very low (0.012), indicating that it correctly identified only a small fraction of the true September samples, while most were misclassified into other months.

December and April achieved modest recall but low precision, suggesting that the network was often predicting these months incorrectly, generating both false positives and false negatives. June achieved the highest recall (1.000), correctly identifying all true June samples. However, its precision (0.364) was still relatively low, suggesting that a high number of false positives where samples from other months were incorrectly classified as June.

The overall pattern suggests that certain months were more easily identifiable than others, likely reflecting the varying degree of structural change occurring in the soil over the season. The particularly poor performance for October and February could be due to periods where structural differences between months were less distinct or more variable within the dataset, making classification more difficult for the network.

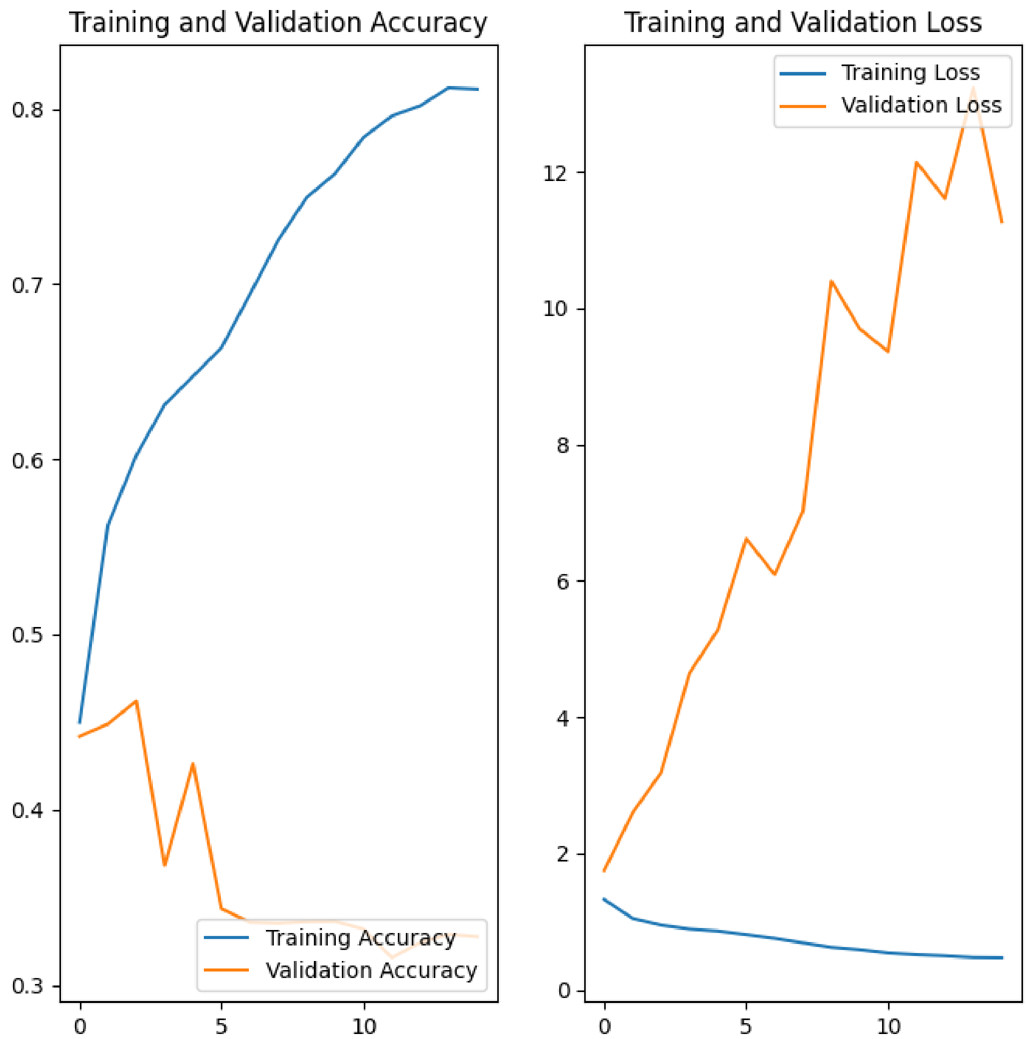

Fig. 6.2: Training and Validation graphs over 15 epochs on the tilled image data.

6.3.2 Treatment classification (Month-by-Month models)

Following the poor performance observed when training the network to classify sampling month across the full dataset, a second experimental approach was implemented. In this approach, separate models were trained for each month to classify tillage treatment (zero-till vs till) based on structural features observed within sample images. Each dataset contains approximately 3000-4,500 images per treatment group, depending on the month.

Training and validation accuracy and loss curves for each model are shown in Figure 6.3. In contrast to the previous classification task, these models displayed more consistent training behaviour, with accuracy and loss curves generally following the expected patterns. However, some variability between months was observed, both in model stability and final accuracy.

Table 6.3 summarises the final performance of each monthly model.

Month	Images per class (approx.)	Accuracy (%)	Till Precision (%)	Zero-Till Precision (%)
September	3000	94	94	94
October	3800	85	88	82
December	3000	92	95	89
February	4500	85	97	75
April	3000	87	83	97
June	3000	84	81	87

Table 6.3: CNN performance for tillage classification across months.

Across all months, the CNN was able to distinguish between zero-till and conventional till with moderate to high accuracy. The best overall performance was observed in September, immediately following tillage, where the network achieved 94% accuracy and balanced precision scores of 94% for both tillage treatments. This suggests that structural differences between tillage and zero-till soils were most distinct directly after cultivation – as expected.

In the following months, classification performance remained generally high, though some variability between months was observed. October, December, and February all achieved accuracy scores between 85% and 92%, showing that the network could still reliably identify structural differences as the season progressed. December produced the second highest accuracy at 92%, with strong precision scores of 95% for tilled soils, and 89% for zero-till soils.

Some imbalance in precision between classes was observed in certain months. In February, for example, the network achieved very high precision for the tilled class (97%) but had much lower precision for zero-till (75%). A similar, but reversed, pattern as seen in April, where zero-till precision reached 97%

and tilled precision was lower at 83%. These variations suggest that, from the machine learning perspective, structural differences between treatments may become more or less distinct at different stages of the growing season, potentially influenced by seasonal biological and environmental processes such as root development, or wetting-drying cycles.

By June, the final sampling point, the network achieved 84% overall accuracy, with precision scores of 81% for tilled soils and 87% for zero-till. This small reduction in classification performance towards the end of the season may reflect structures becoming more similar between treatments over time, as previously suggested by the traditional image analysis presented in Chapter 5.

Overall, these results demonstrate that while differences between tillage treatments remain detectable throughout the season, the magnitude of and consistency of these differences vary as the soil structure evolves over time.

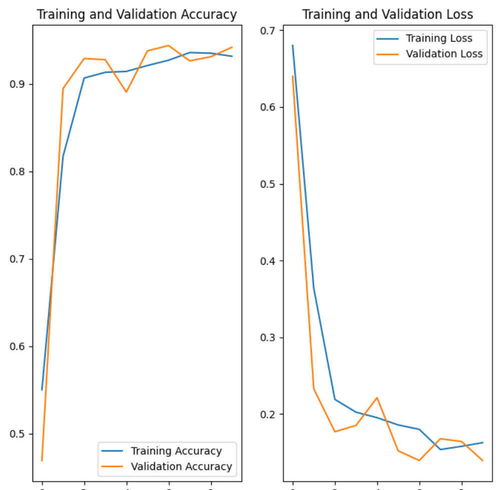
Training and Validation Accuracy
Training Accuracy
Validation Accuracy
Training and Validation Loss
Training Loss
Validation Loss

a.

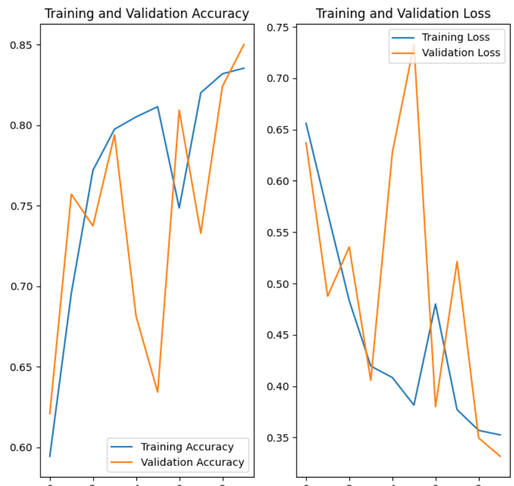
Training and Validation Accuracy
Training Accuracy
Validation Accuracy
Training and Validation Loss
Training Loss
Validation Loss

b.


Training and Validation Accuracy
Training Accuracy
Validation Accuracy
Training and Validation Loss
Training Loss
Validation Loss

c.


Training and Validation Accuracy
Training Accuracy
Validation Accuracy
Training and Validation Loss
Training Loss
Validation Loss

d.


Training and Validation Accuracy
Training Accuracy
Validation Accuracy
Training and Validation Loss
Training Loss
Validation Loss

e.

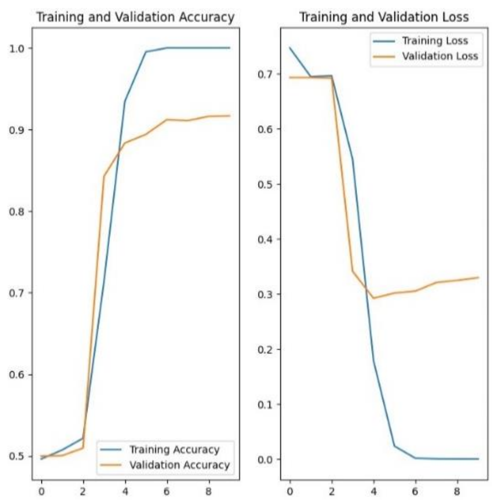
Training and Validation Accuracy
Training Accuracy
Validation Accuracy
Training and Validation Loss
Training Loss
Validation Loss

f.

Fig. 6.3: Training and validation accuracy and loss for each monthly tillage classification model. September (a), October (b), December (c), February (d), April (e), and June (f). The graphs display the network's performance against training epochs for each month-specific model.

6.3.3 Testing and Saliency Maps

Once training for each month-specific model had been completed (using the optimal number of training epochs, as identified in Section 6.3.2), the models were evaluated on independent test datasets comprising of previously unseen image data. The performance of each model on these unseen test sets is summarised in Table 6.4.

Month	**Class**	**Images Tested**	**Correct Predictions**	**% Correct**
September	Till	182	178	97.8
September	Zero-Till	210	203	96.7
October	Till	232	7	3.0
October	Zero-Till	181	179	98.9
December	Till	238	86	36.1
December	Zero-Till	186	152	81.7
February	Till	236	195	82.6
February	Zero-Till	278	36	12.9
April	Till	215	213	99.1
April	Zero-Till	203	0	0
June	Till	198	194	98.0
June	Zero-Till	204	0	0

Table 6.4: CNN prediction performance on independent, previously unseen test images for each month.

Figure 6.5 gives a visual summary of these results.

The test set results reveal some considerable variation in model performance between months and between classes. For some months, such as September, April, and June, the models performed well, correctly classifying almost all test images for the tilled samples. In contrast, classification of zero-till images was highly variable, with complete failure to correctly classify any zero-till test images in April and June, despite good performance on training and validation data.

The October model displayed particularly inconsistent behaviour between classes. While it achieved a high test accuracy for zero-till samples (98.9%), the model performed very poorly on tilled samples, correctly classifying only 3% of test images. This suggests that the network struggled to consistently identify structural features distinguishing tilled soil from zero-till at this point in the growing season, despite apparently good learning behaviour during training. The underlying cause of this variability may reflect either complex structural overlap between classes at this time point, or insufficient robustness in the model due to limited training data variability.

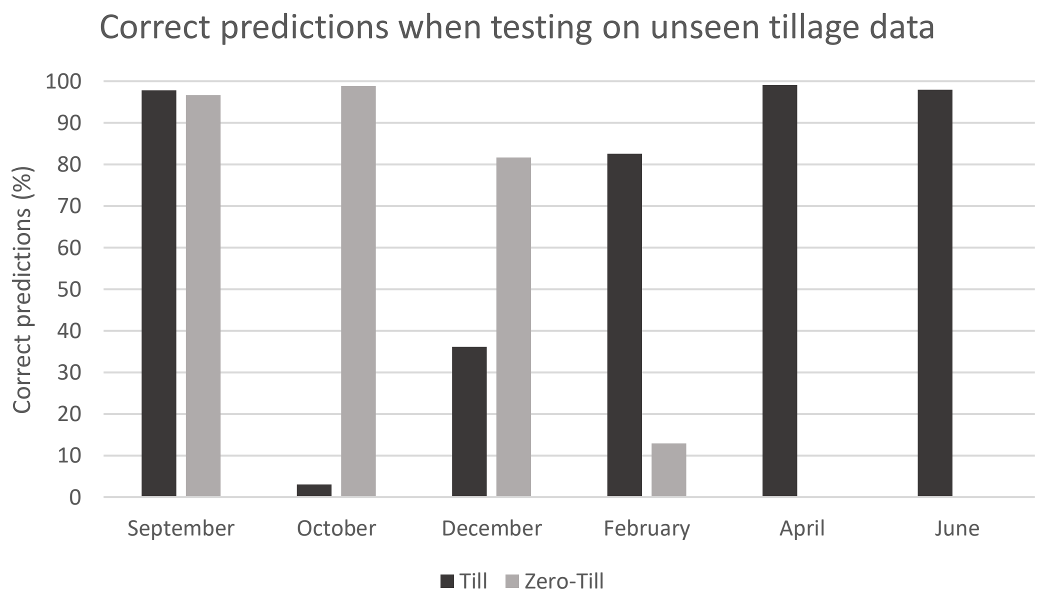

Fig. 6.4: Chart showing the percentage of correct predictions made by the CNN after being shown previously unseen tillage data for each month.

To further investigate the features used by the CNN during classification, saliency maps were generated for both correctly and incorrectly classified images (Figures 6.5-6.20). Saliency maps highlight regions within the input image that contributed most strongly to the network's classification decision, providing insight into which structural features the model focussed on.

In the majority of correctly classified samples across months, saliency maps identified clear structural boundaries, pore spaces, and aggregate edges as important features. However, saliency maps for October presented a different pattern. In both correctly and incorrectly tilled images, the saliency maps frequently displayed a scattered 'star map' appearance, as shown in Figures 6.8-6.10. These maps appeared noisier than saliency maps from other months, and lacked a consistent focus on identifiable structural features, suggesting the network may have struggled to extract stable features for tilled samples at this time point.

In contrast, saliency maps for correctly classified zero-till images from October (Figure 6.10) were more consistent and displayed less focus on stone materials within the sample. While some blank regions corresponding to stones were observed, these maps showed clearer attention to the pore structures and aggregates, aligning more closely with expectations based on soil structure.

The saliency maps therefore provide an additional insight into and evidence that, particularly in October, the model struggled to consistently extract meaningful features from the tilled data, potentially contributing to the poor classification performance observed in the test set for this class.

In summary, the CNN showed limited ability to classify samples by sampling month but performed more successfully when classifying tillage treatments on a month-by-month basis. Treatment classification accuracy was generally high across most months, particularly immediately after cultivation, though performance varied between months and between classes. Additional testing on unseen data and analysis of saliency maps highlighted both stable structural features the network could identify, and months where model performance was less consistent, particularly in October and later in the season.

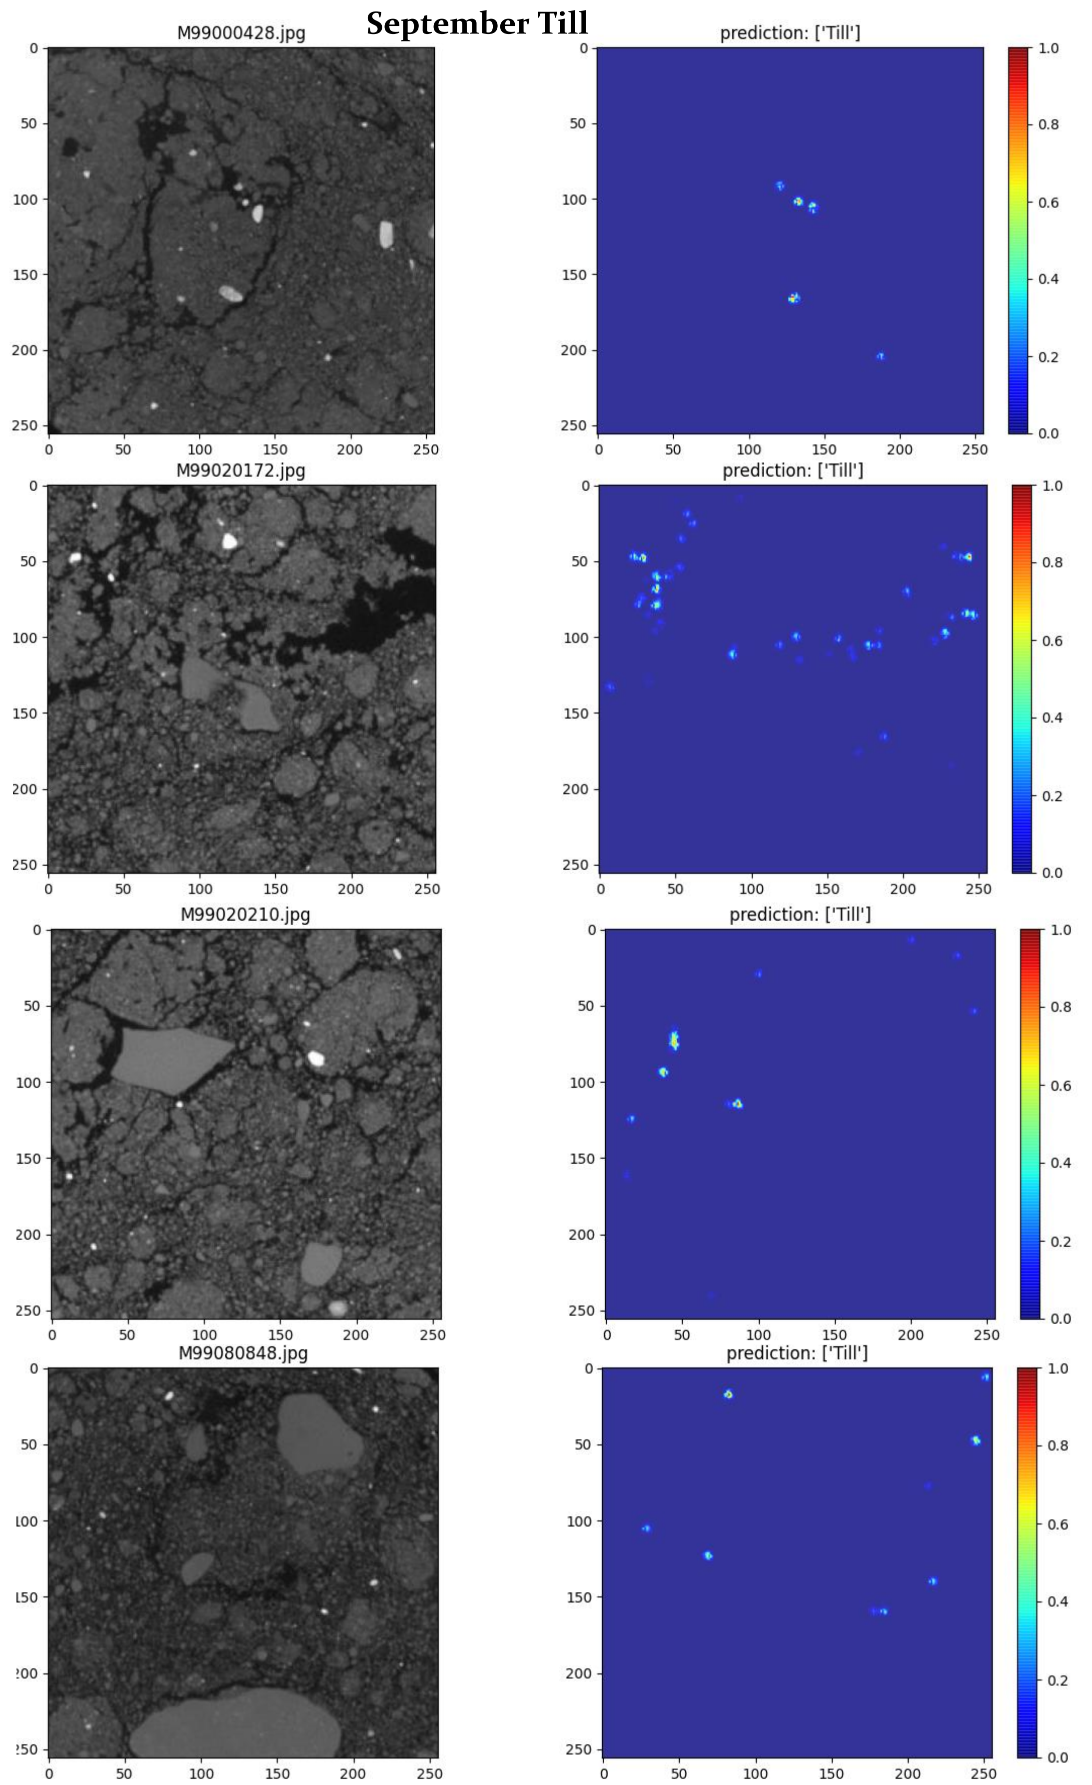

Fig. 6.5: Saliency Maps for correct till predictions on unseen September tillage data.

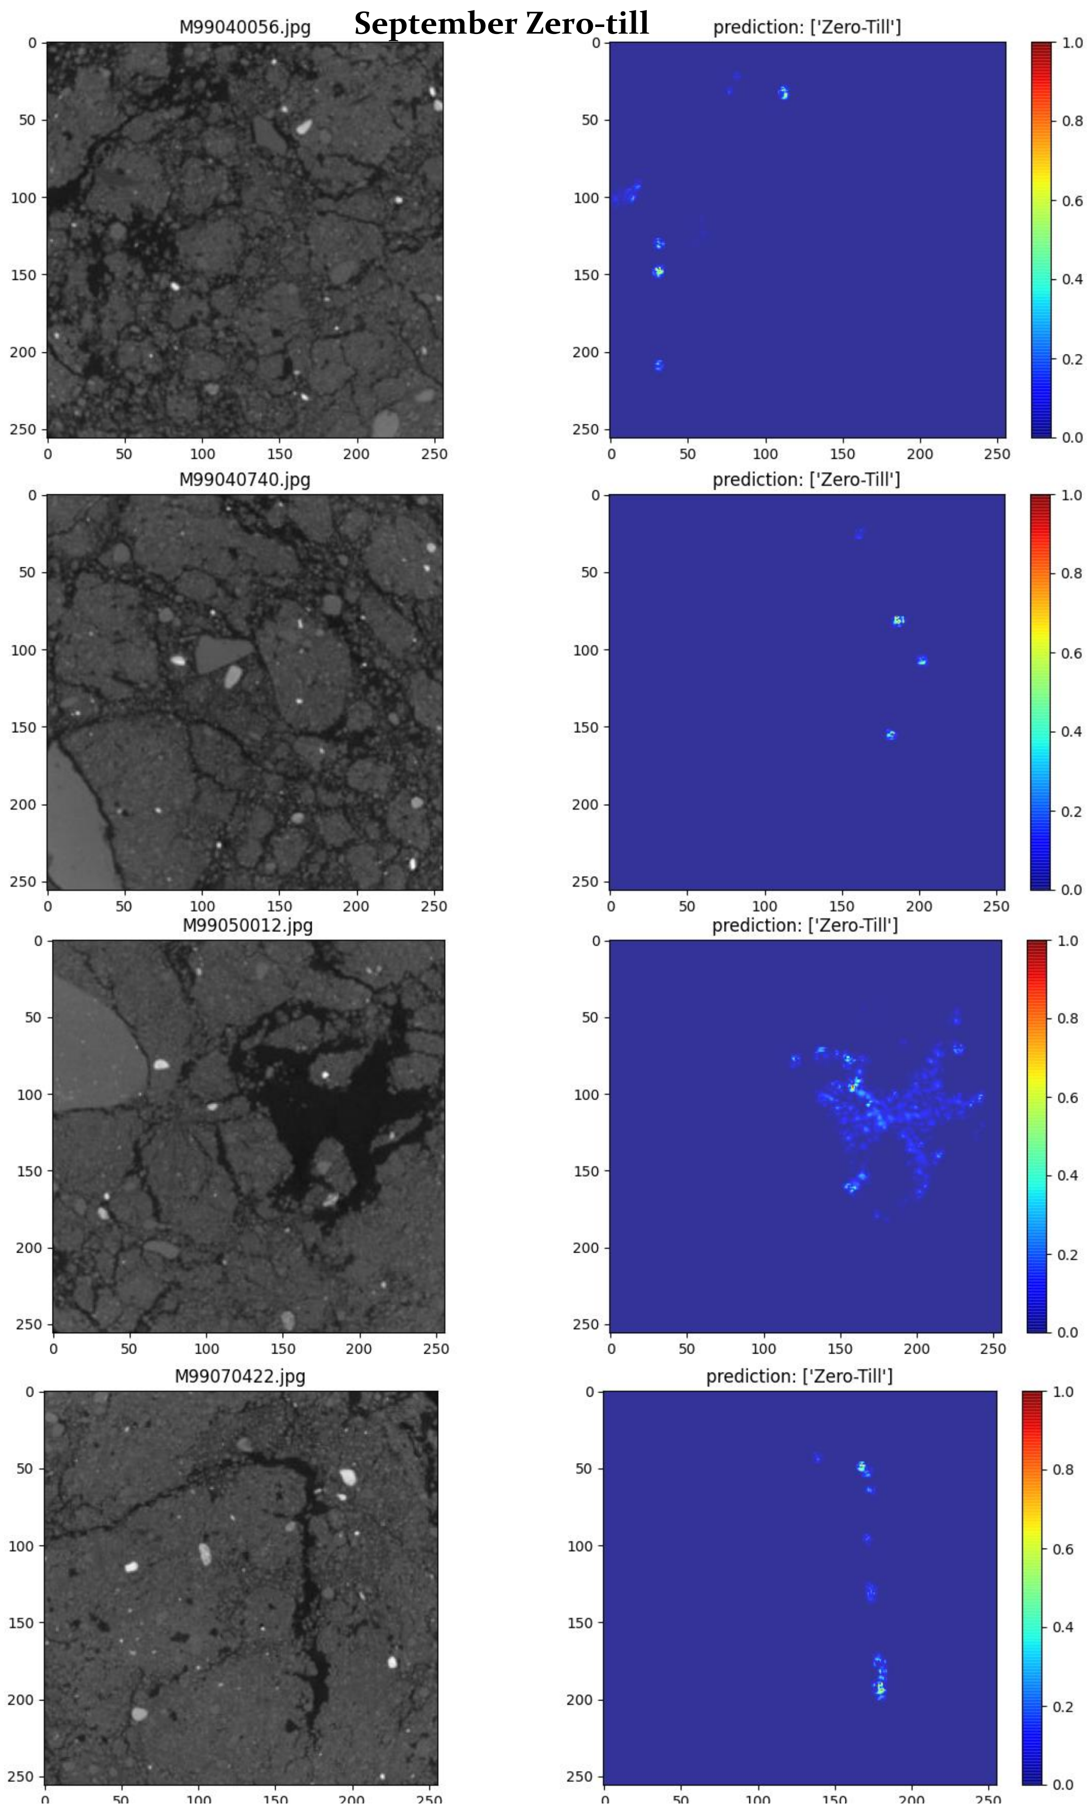

Fig. 6.6: Saliency Maps for correct zero-till predictions on unseen September tillage data.

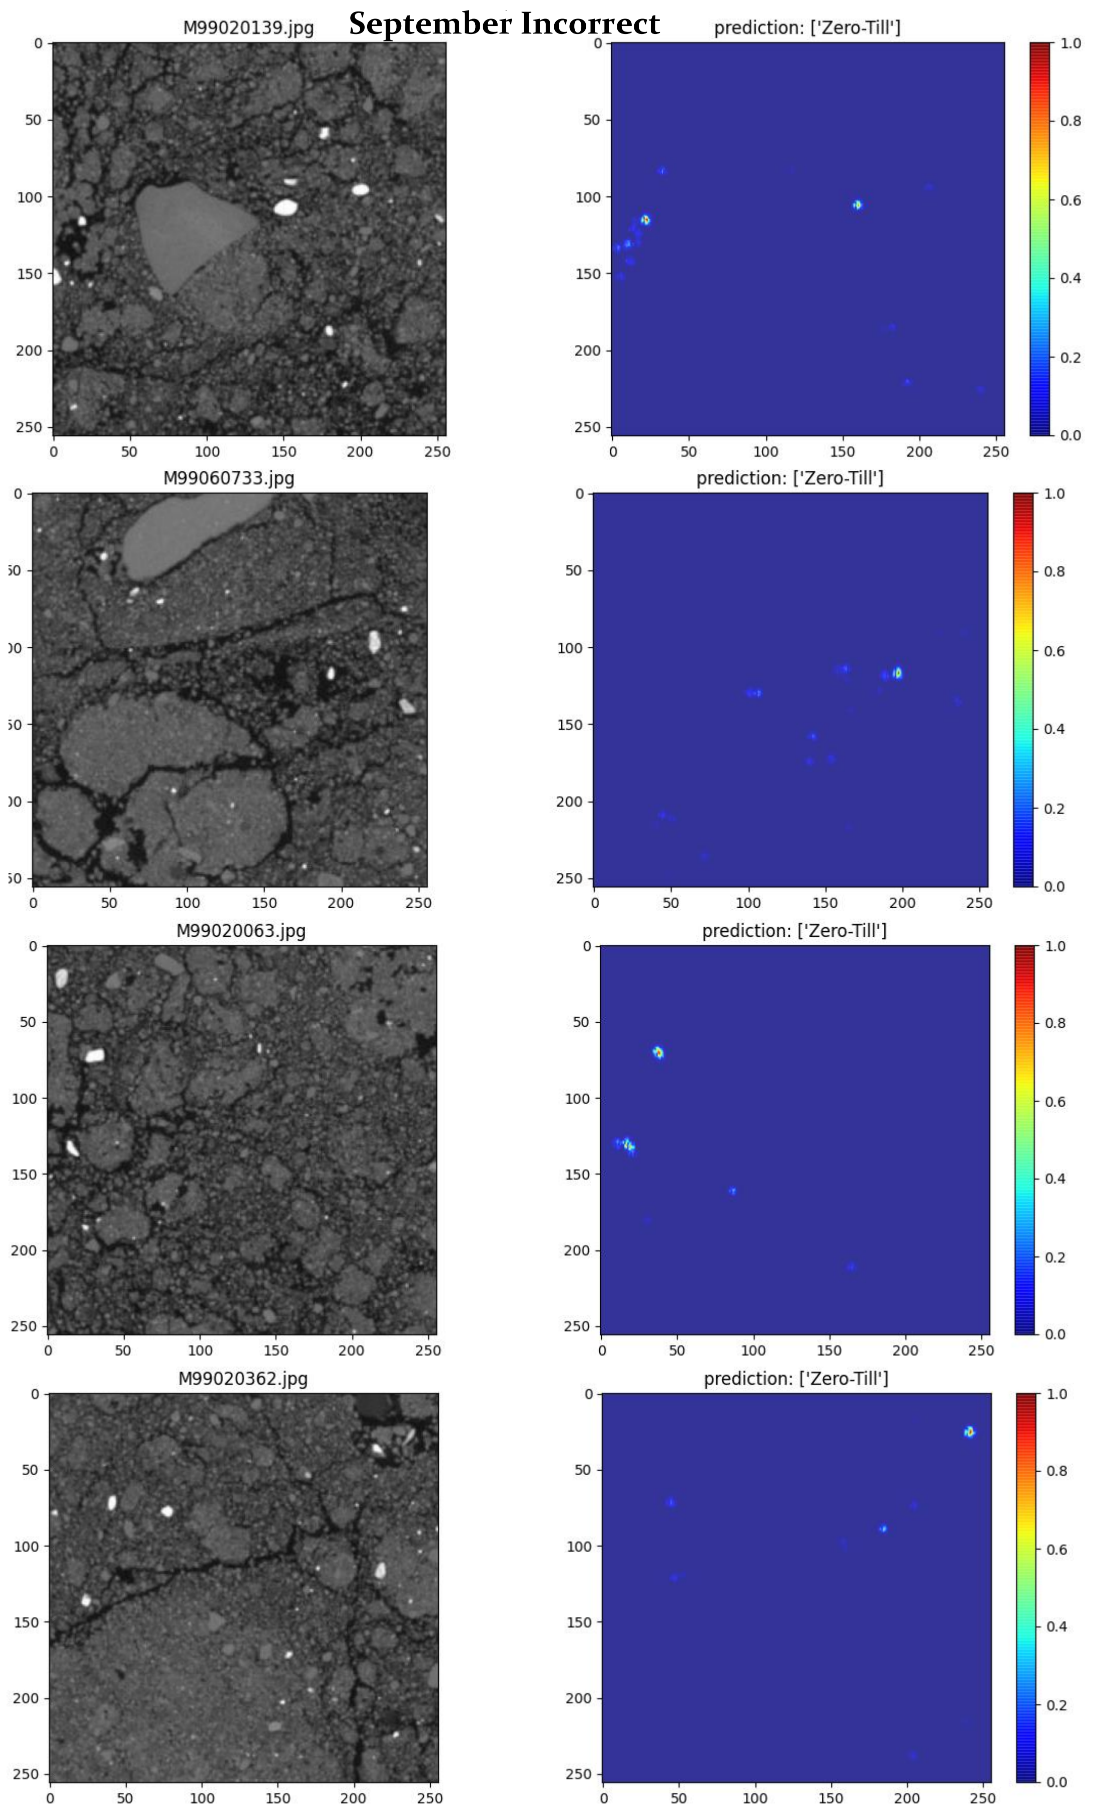

Fig. 6.7: Saliency Maps for incorrect predictions on unseen September tillage data.

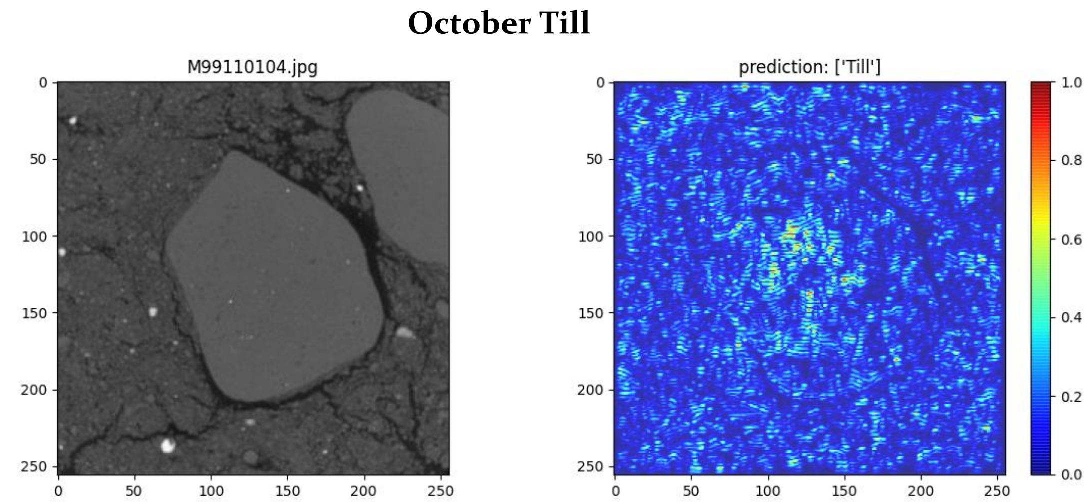

Fig. 6.8: Saliency Maps for correct till predictions on unseen October tillage data.

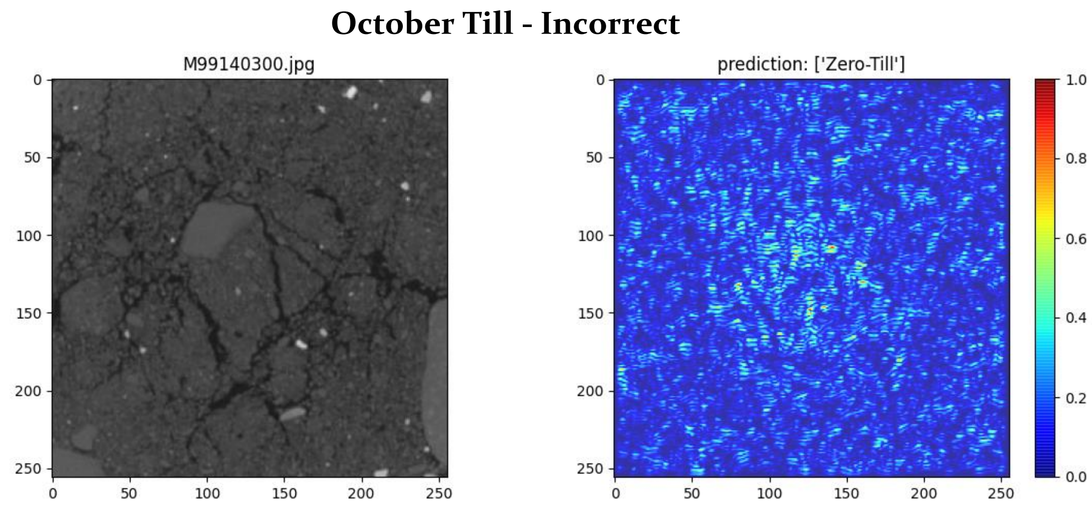

Fig. 6.9: Saliency maps for incorrect predictions on unseen October tillage data.

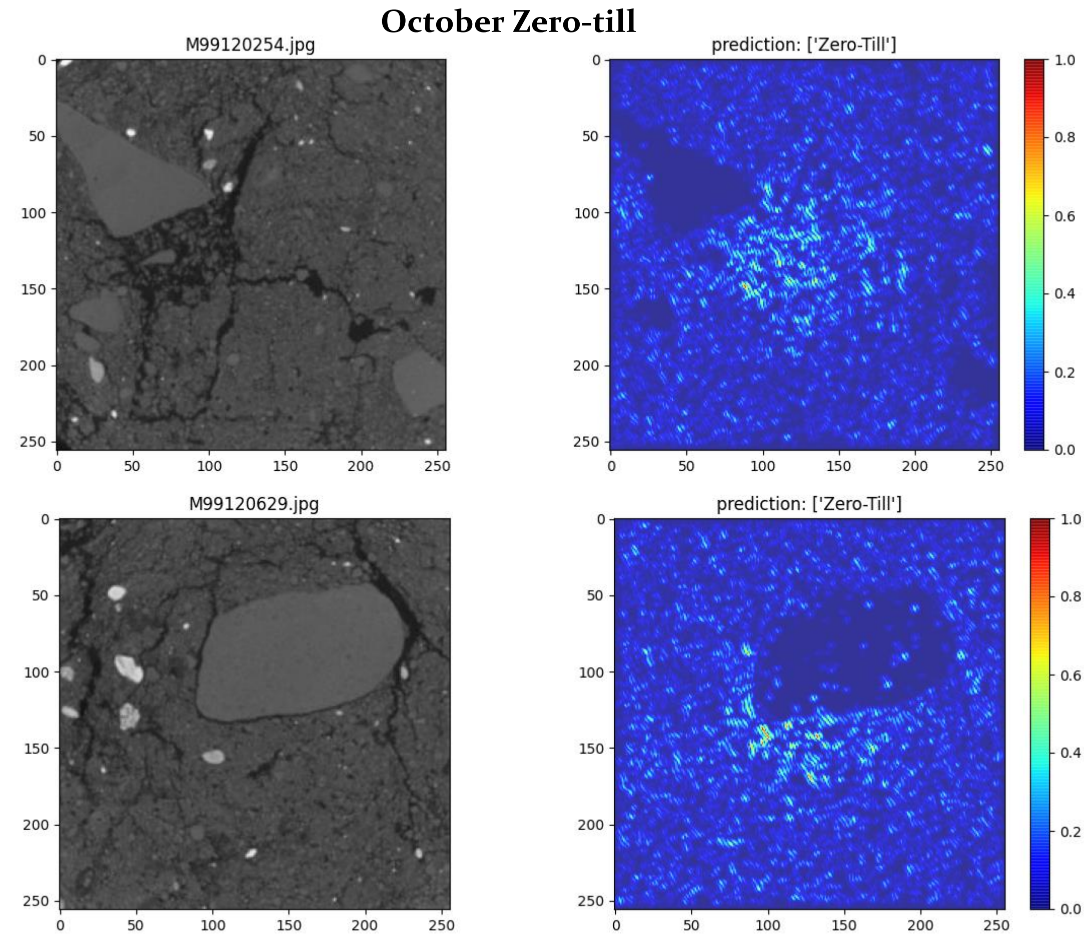

Fig. 6.10: Saliency maps for correct zero-till predictions on unseen October tillage data.

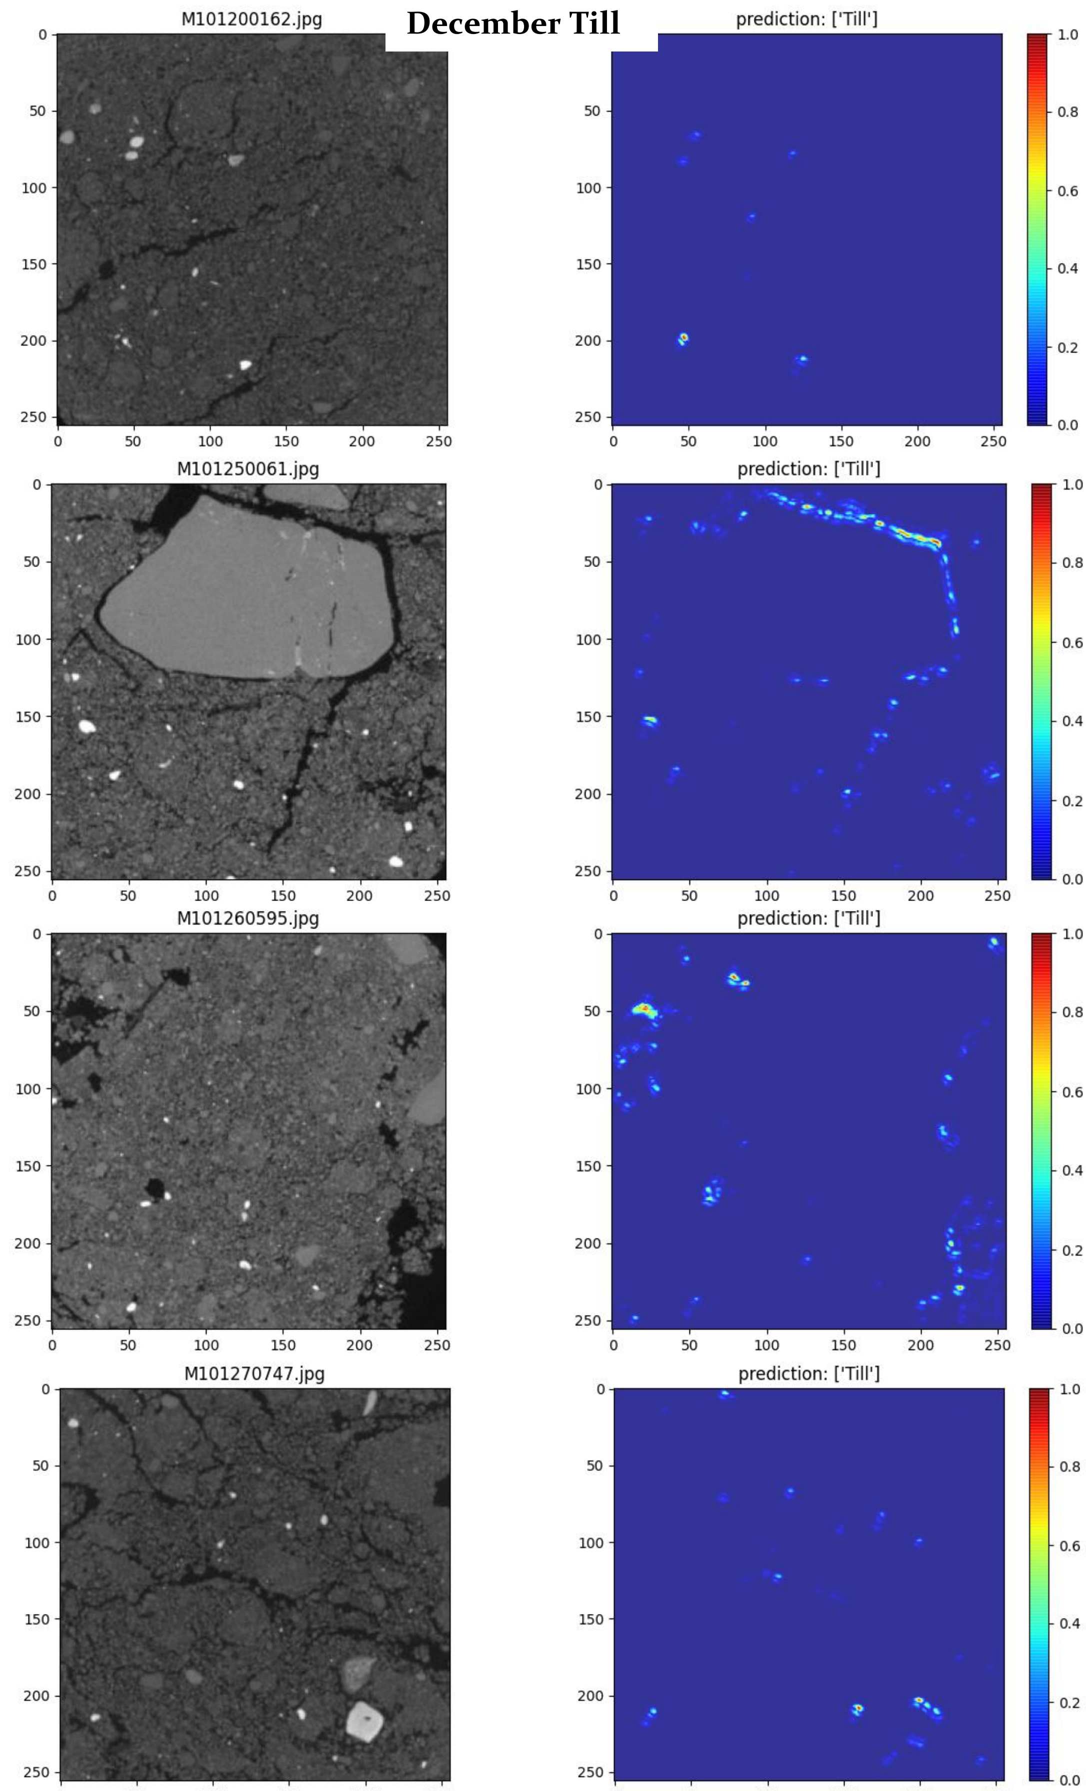

Fig. 6.11: Saliency maps for correct till predictions on unseen December tillage data.

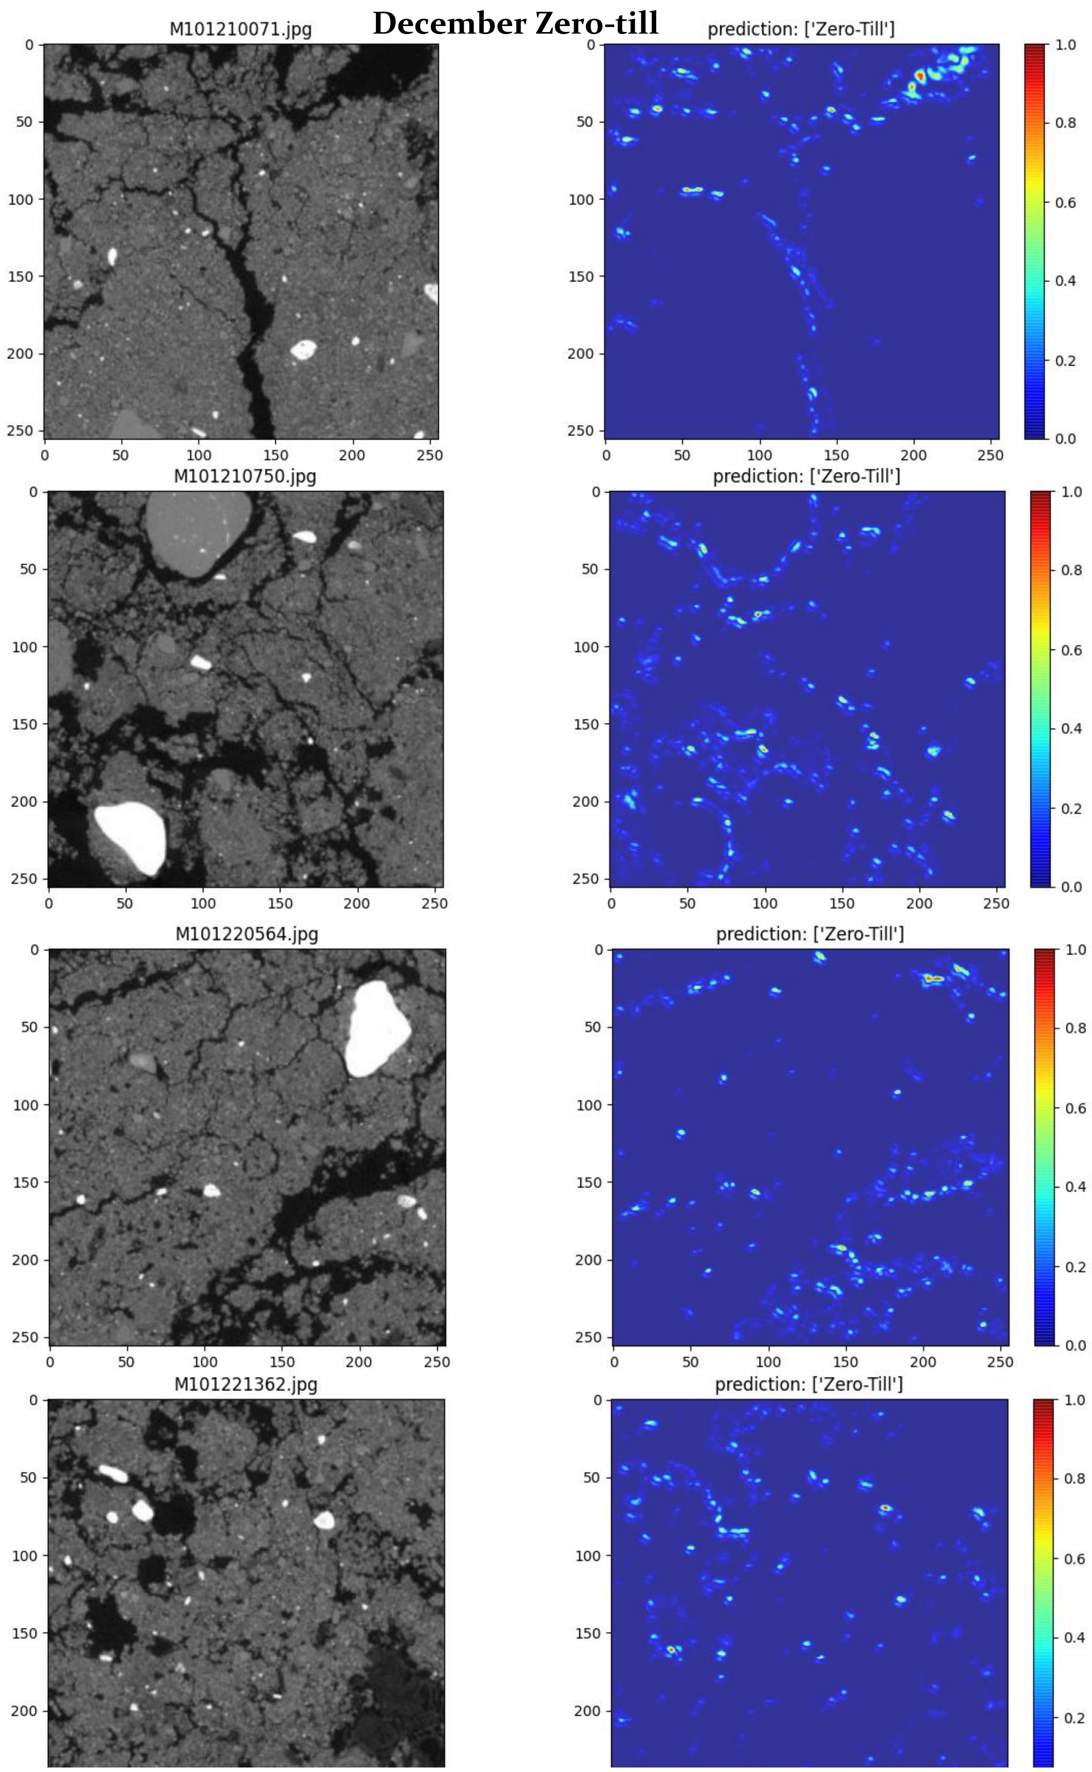

Fig. 6.12: Saliency maps for correct zero-till predictions on unseen December tillage data.

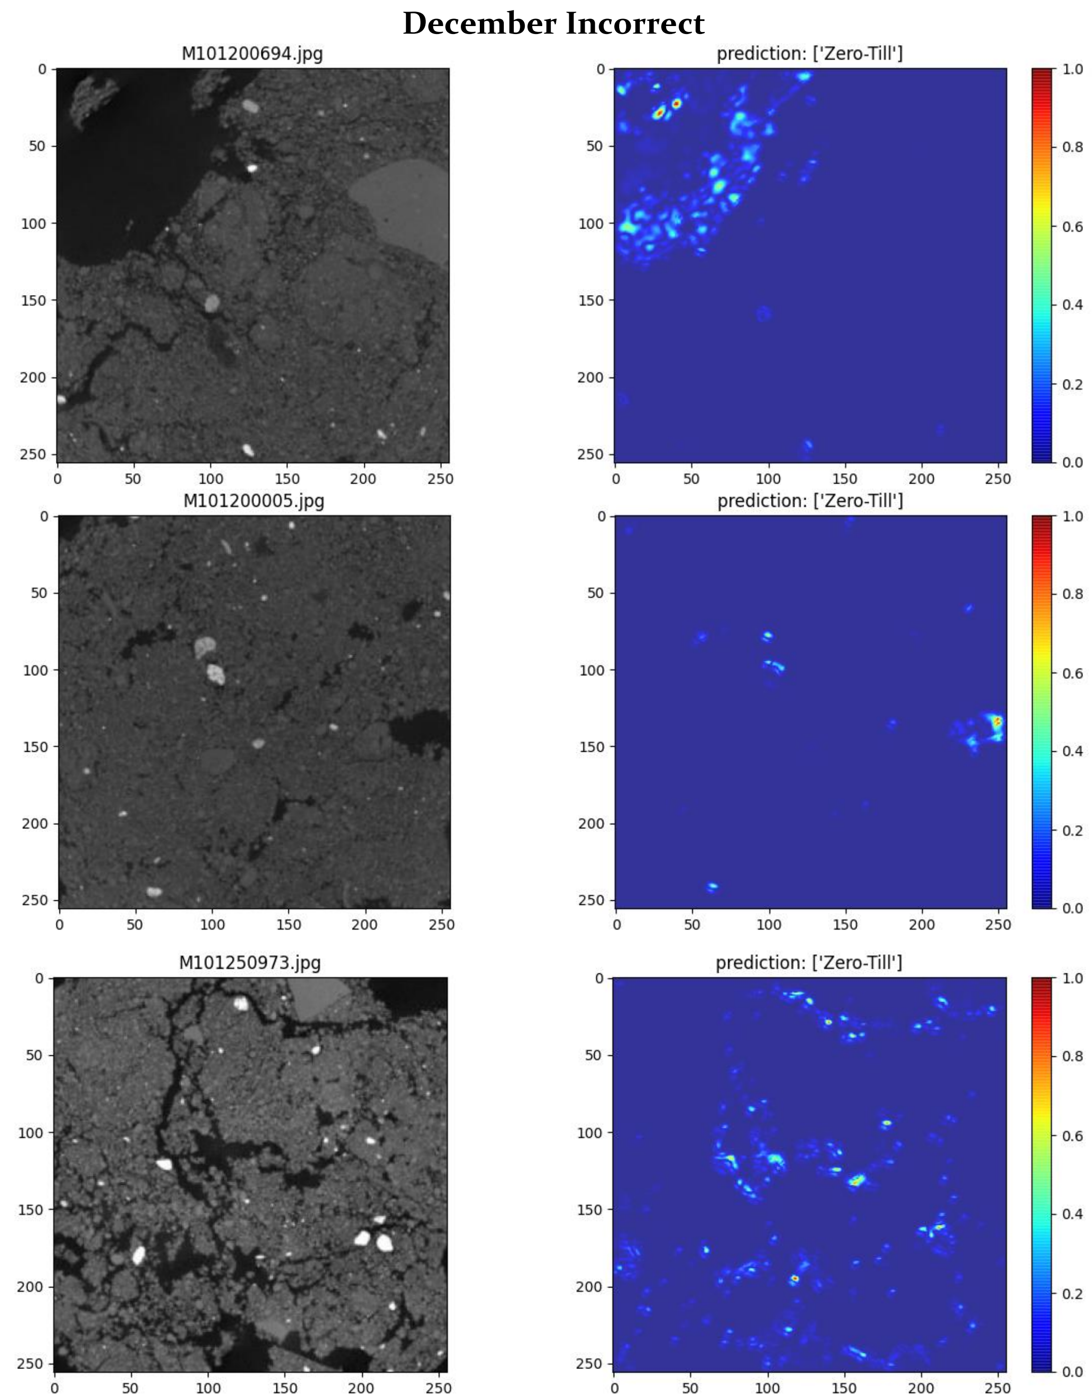

Fig. 6.3: Saliency maps for incorrect till predictions on unseen December data.

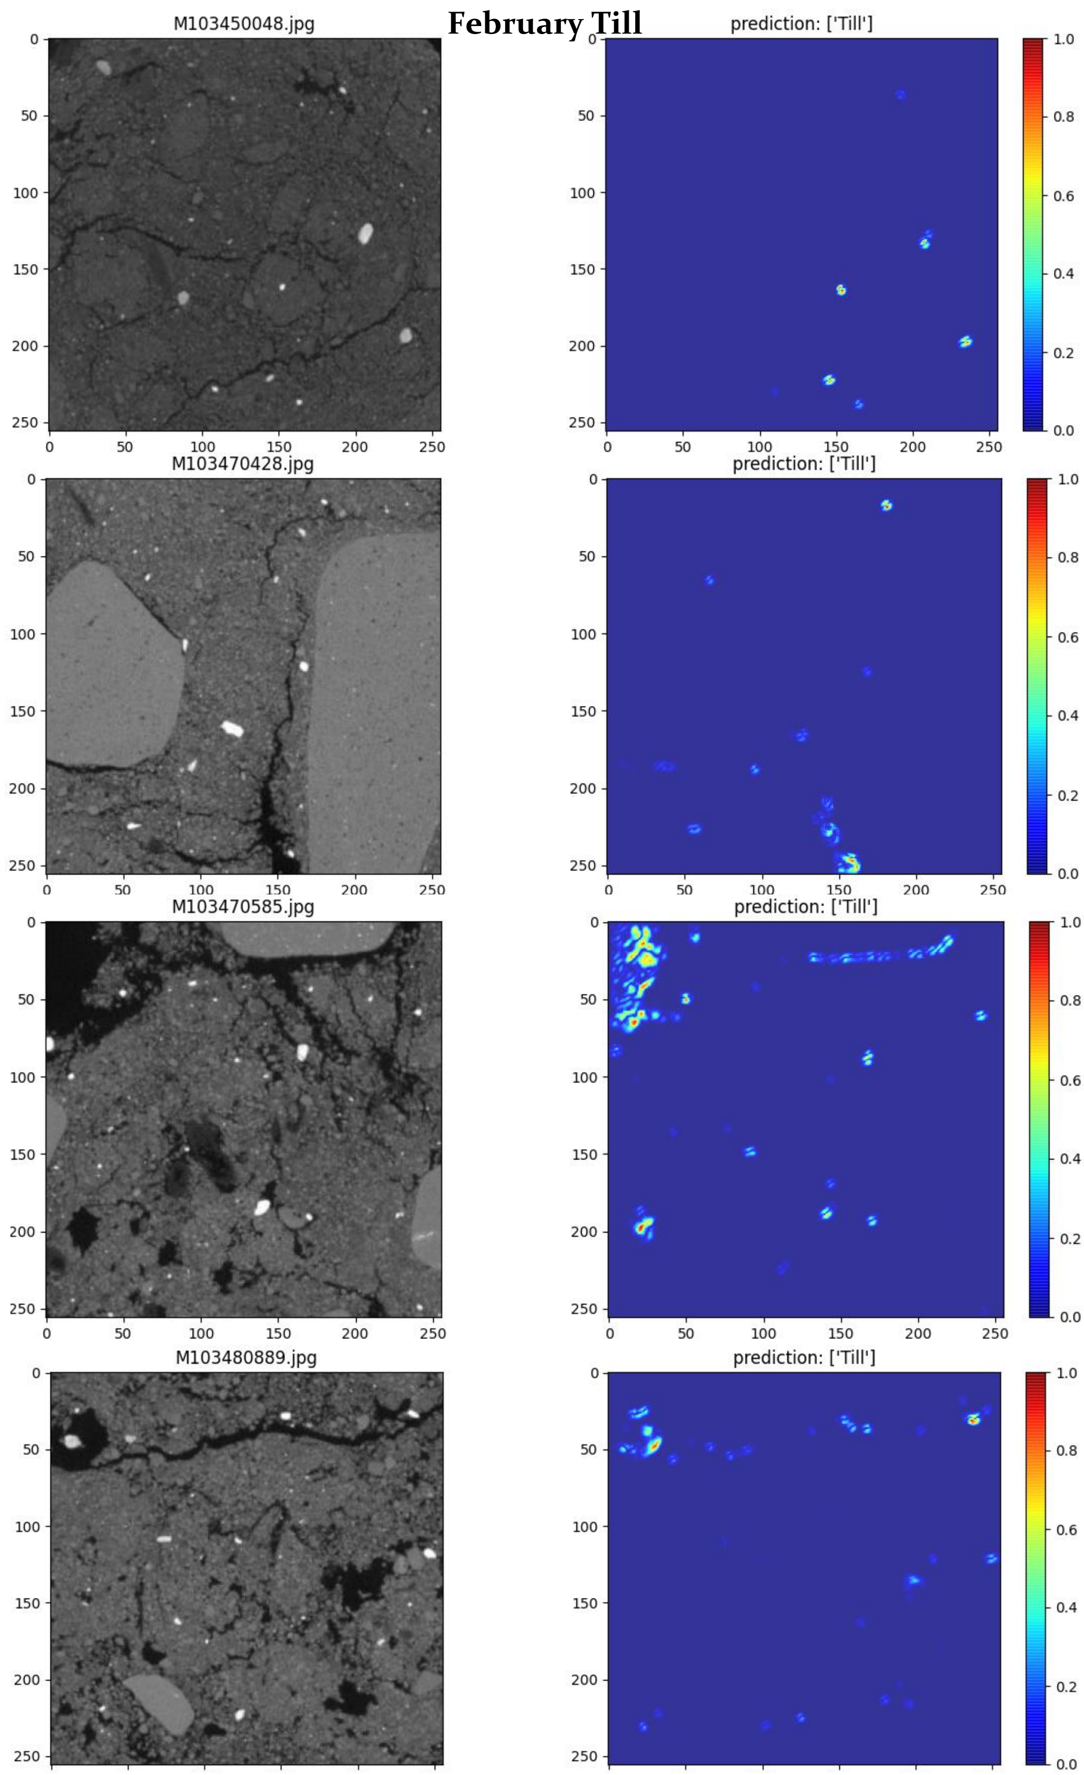

Fig. 6.4: Saliency map for correct till predictions on unseen February tillage data.

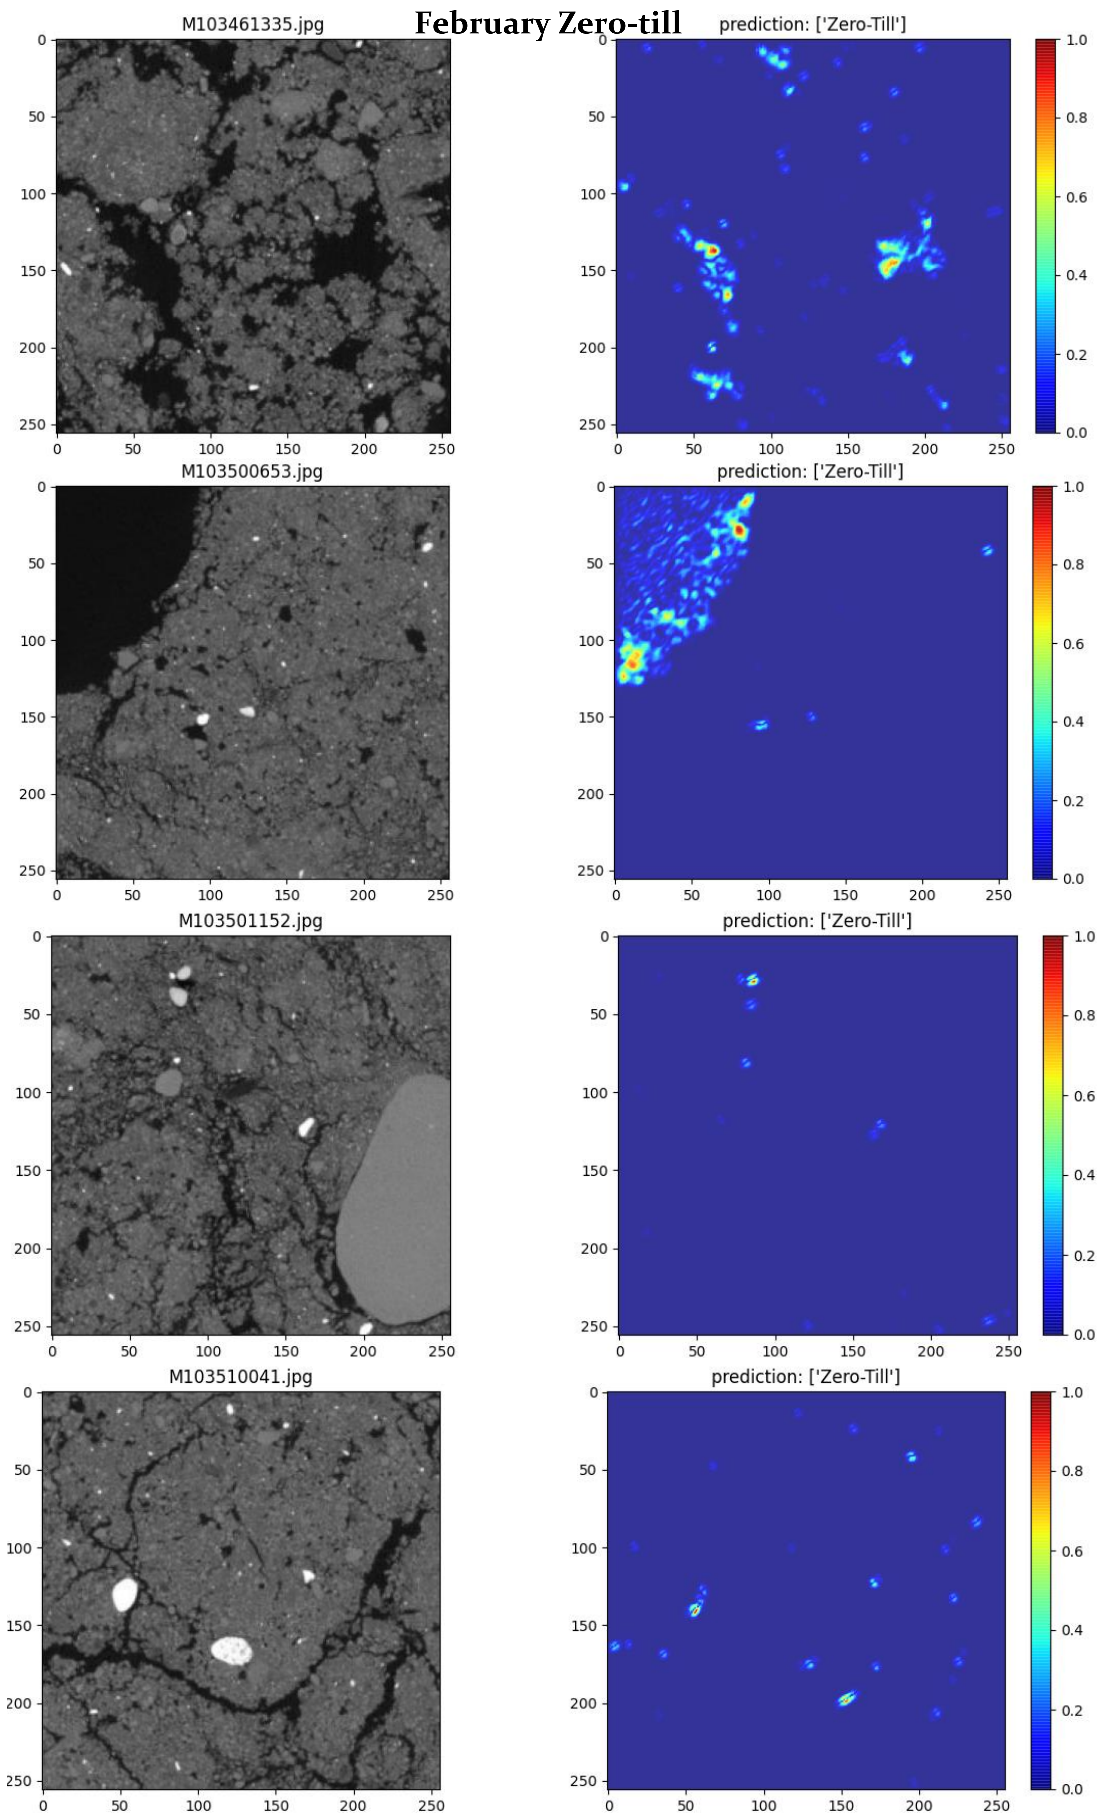

Fig. 6.15: Saliency maps for correct zero-till predictions on unseen February data.

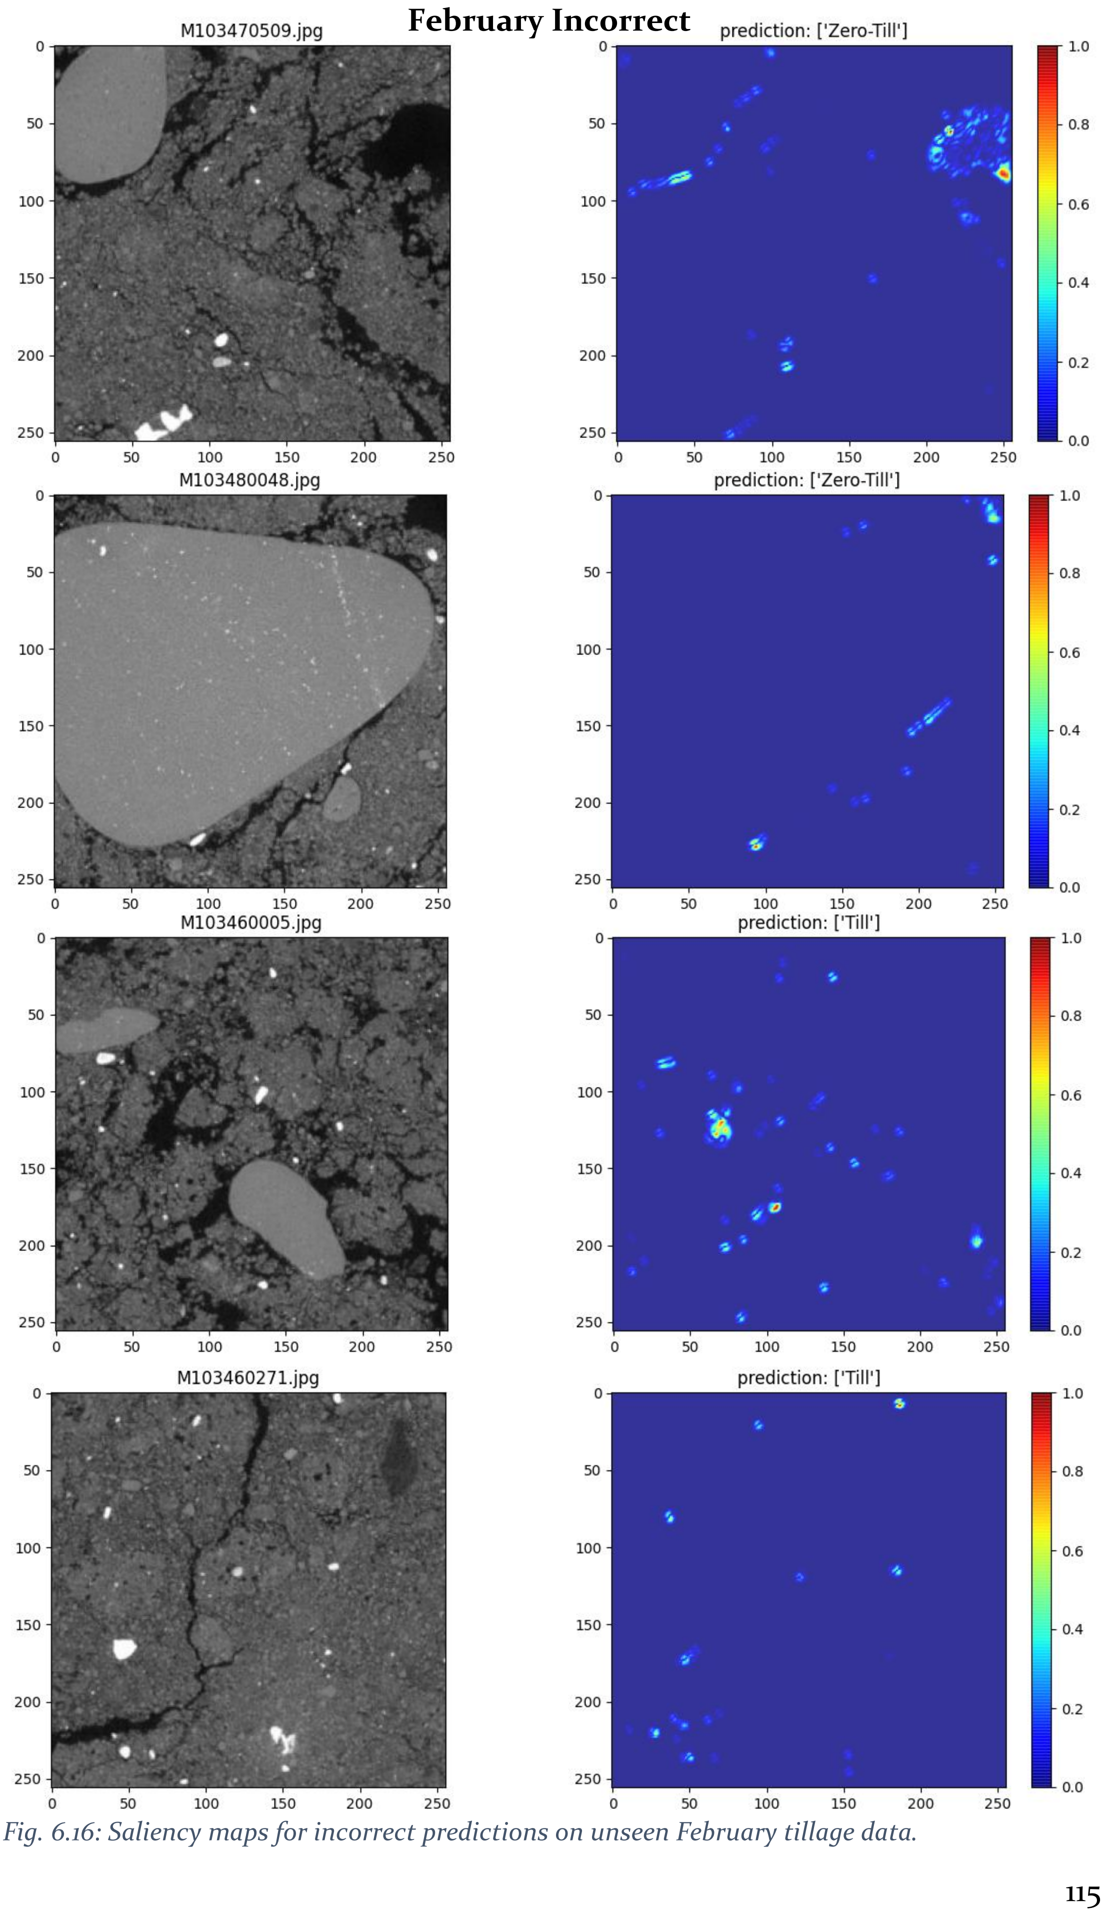

Fig. 6.16: Saliency maps for incorrect predictions on unseen February tillage data.

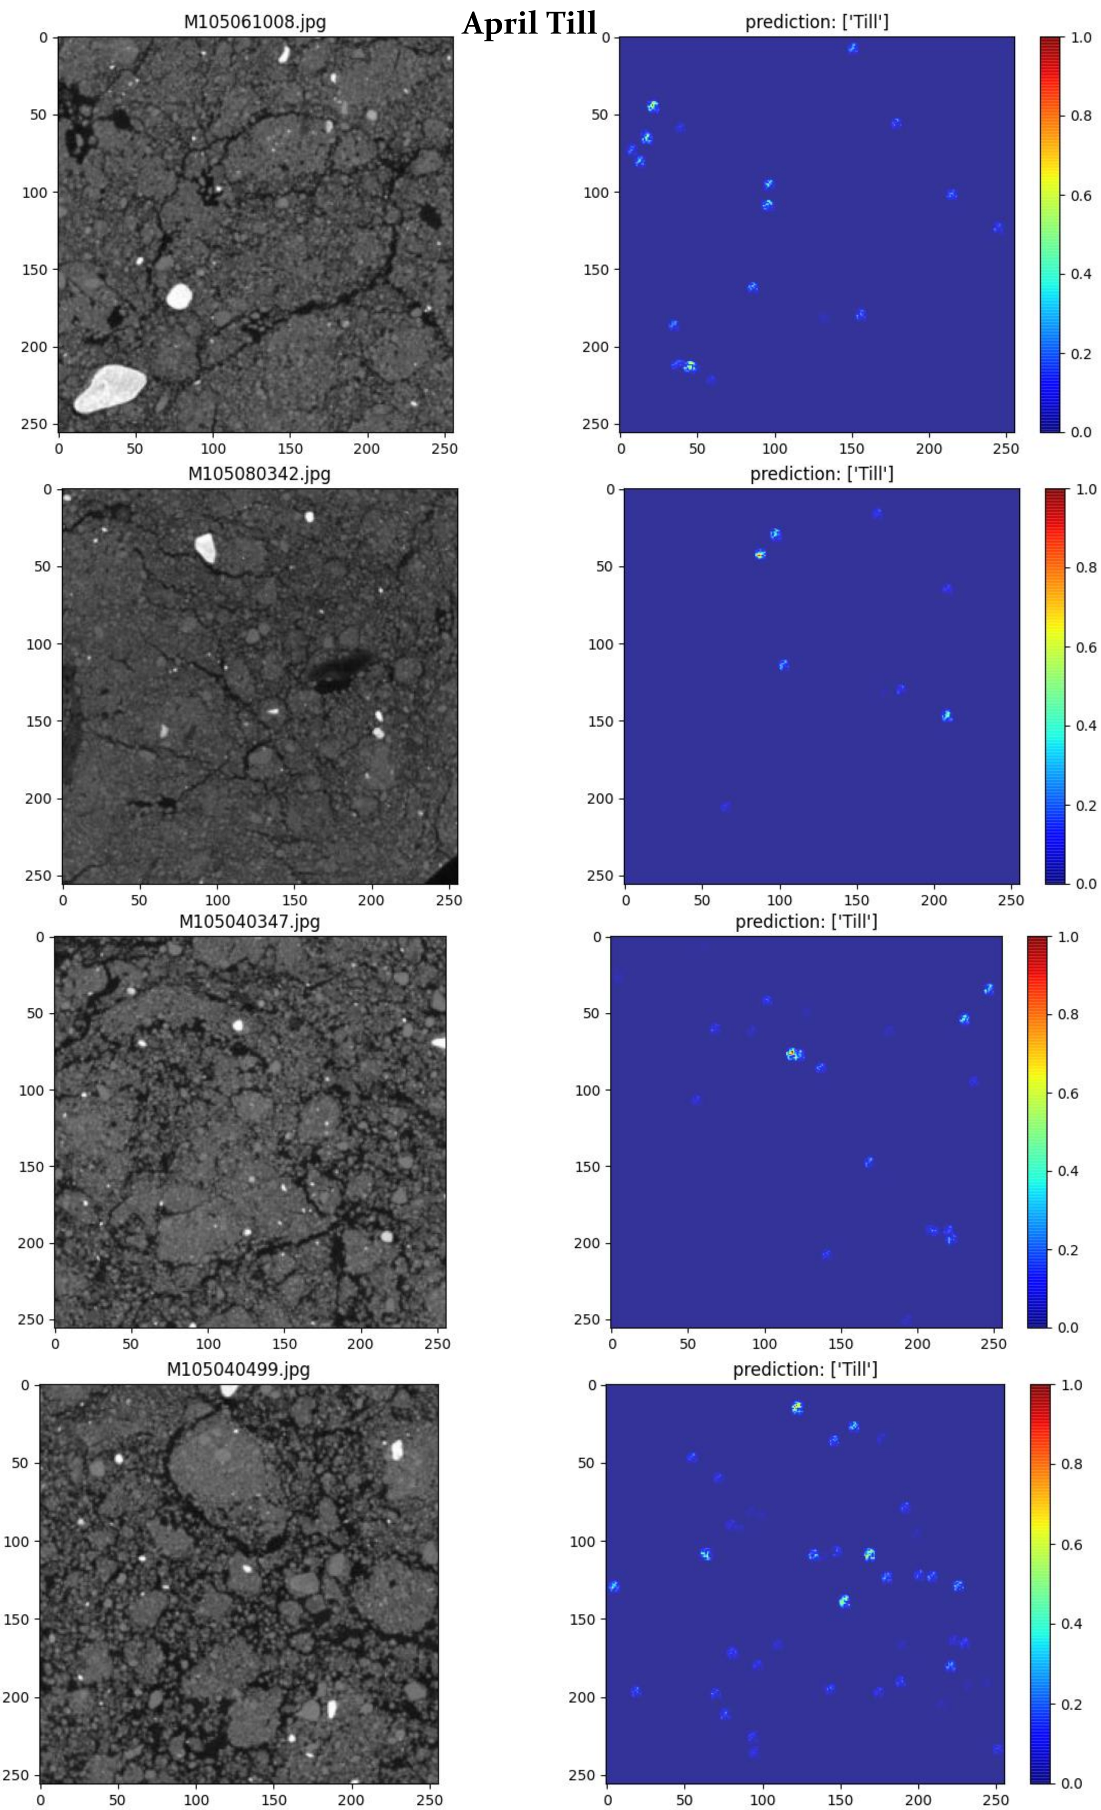

Fig. 6.17: Saliency maps for correct till predictions on unseen April tillage data.

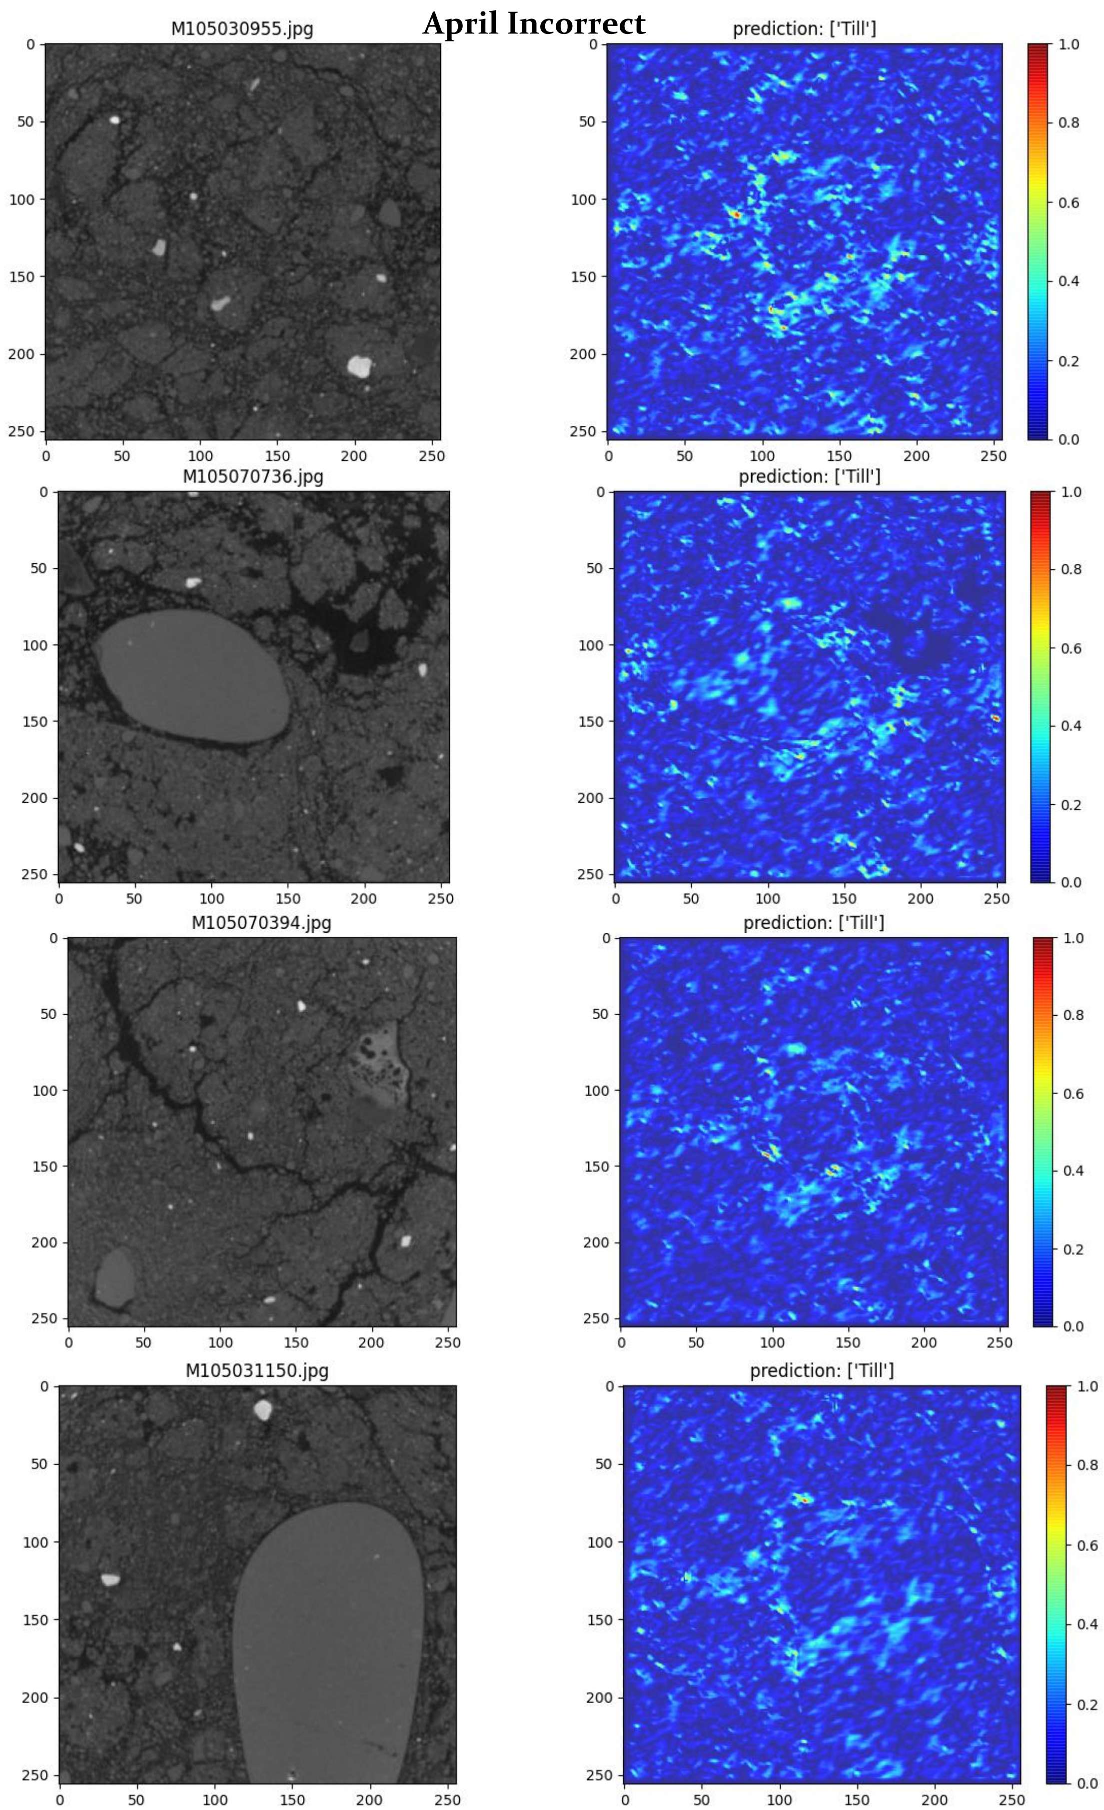

Fig. 6.18: Saliency maps for incorrect predictions on April tillage data.

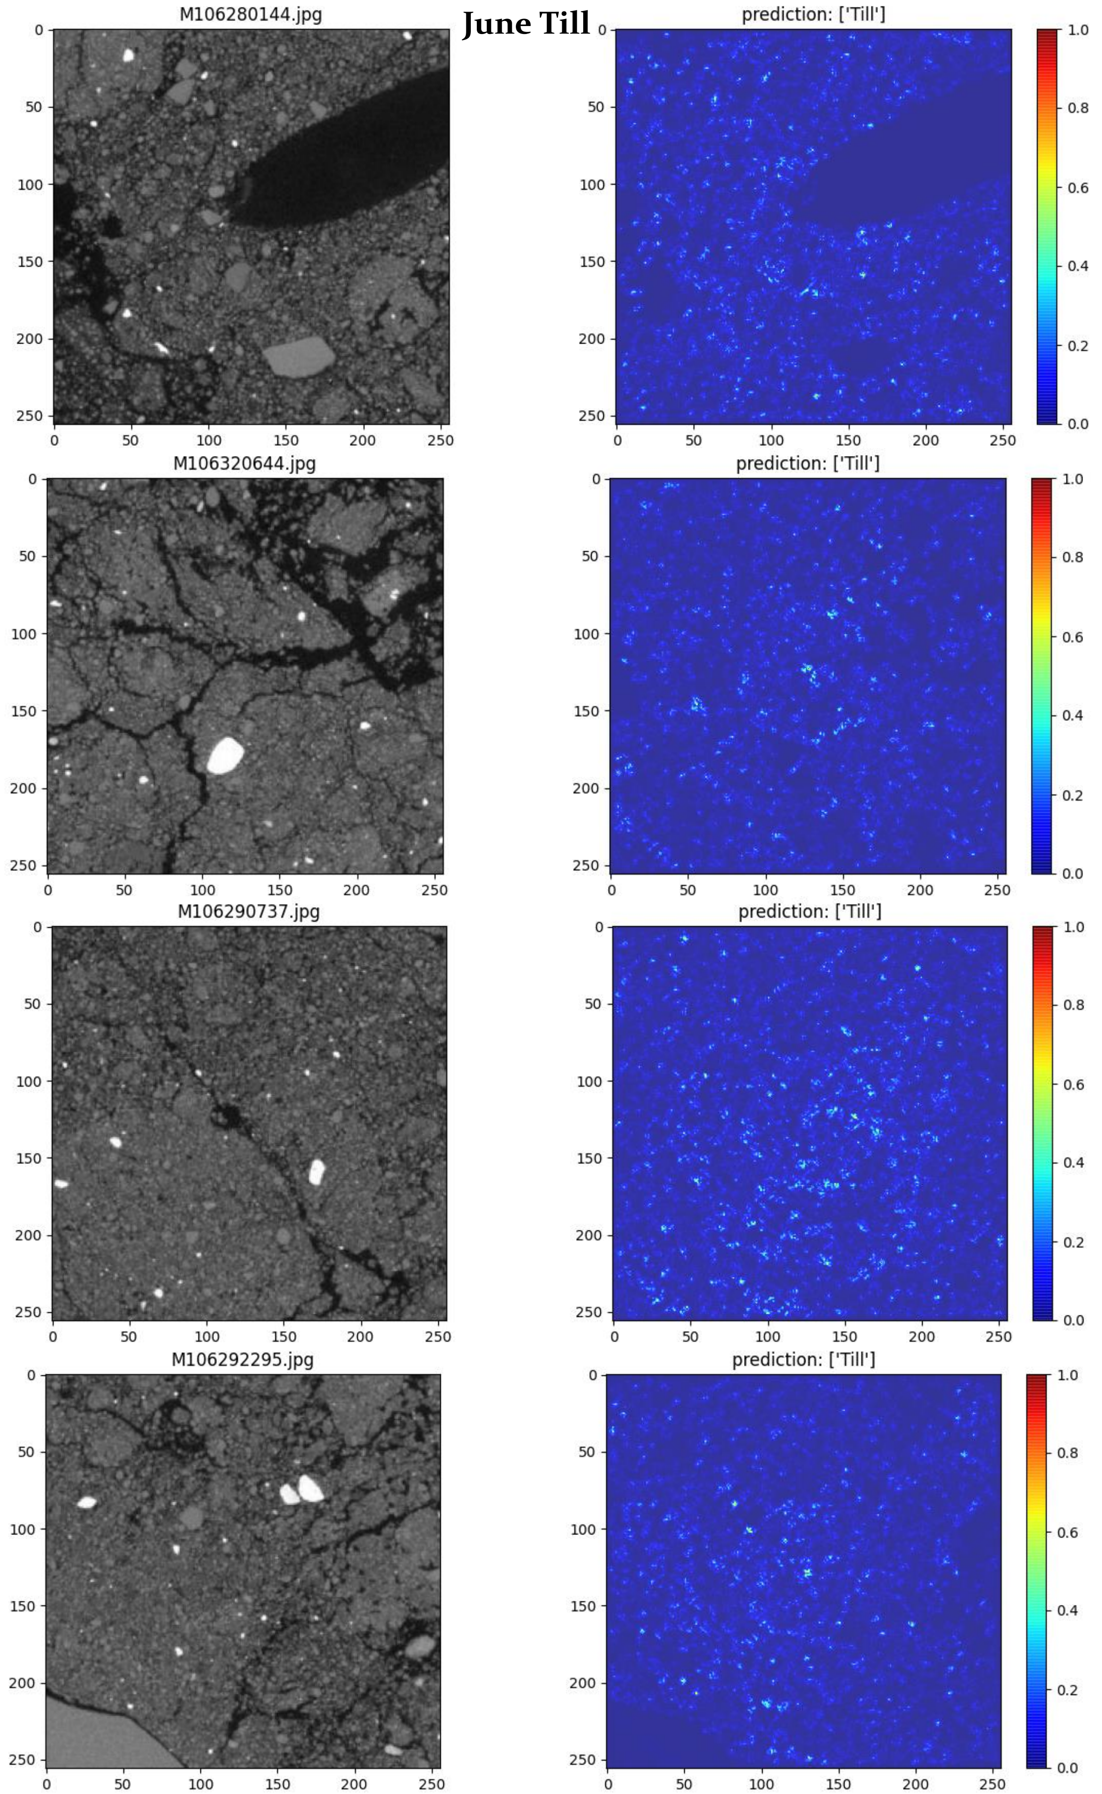

Fig. 6.19: Saliency maps for correct till predictions on June tillage data.

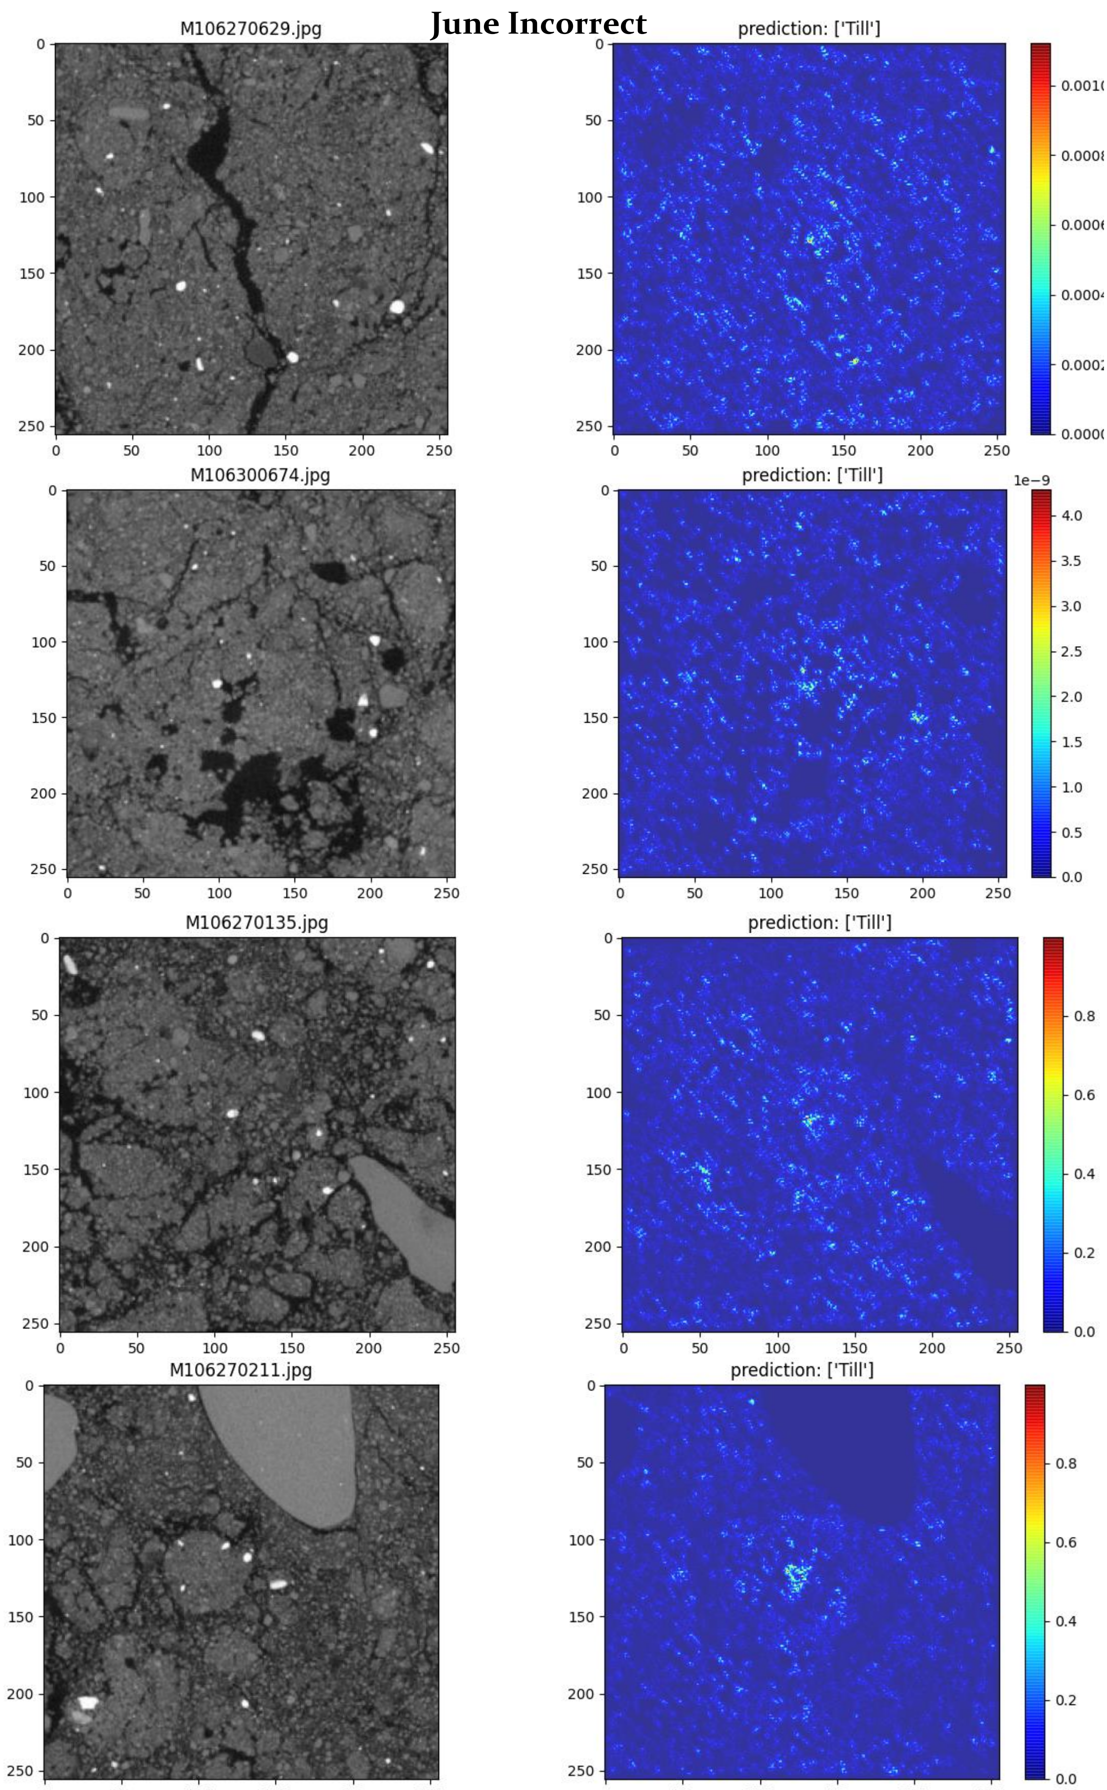

Fig. 6.20: Saliency maps for incorrect predictions in unseen June tillage data.

6.4 Discussion

This chapter explored the use of machine learning to analyse the same soil structure dataset previously studied through conventional image analysis in Chapter 4. The aim was to investigate whether a CNN could detect structural differences between tilled and zero-tilled soils across the growing season and identify patterns potentially missed by traditional porosity and pore size metrics.

The results demonstrate a mixed level of success. When trained to classify samples by month using all data combined, the network struggled to extract meaningful temporal patterns, achieving low classification accuracy overall. Certain months, such as October and February, were partially difficult for the network to differentiate from other months, likely reflecting subtler structural changes at these stages in the season. By contrast, September, immediately after cultivation, was classified most easily, with high precision, suggestion that this point in the season presented the clearest structural distinction.

When the problem was simplified to classifying tillage treatment month-by-month, the network performed substantially better. For most months, the CNN achieved high accuracy, particularly in September, where 94% accuracy was achieved on validation data and 97-98% on unseen test data. However, classification performance varied between months and between classes, with several months displaying strong performance on one treatment while struggling with the other. For example, in April and June, the network correctly classified nearly all tilled samples but failed entirely to correctly classify any zero-till samples in the test set. This suggests that while some structural differences between tillage treatments remain detectable later in the season, the magnitude of these differences varies, and at times becomes more challenging for the network to distinguish.

The saliency map analysis provides further insight into how the network made its classifications. In months with strong performance, such as September, saliency maps highlighted clear structural features such as pores, aggregates and cracks. In contrast, for problematic months such as October, saliency maps often displayed scattered or unfocused activation patterns, suggesting the network was unable to consistently identify robust structural features in these cases. Also, the fact that the network still misclassified some images even when one class was performing very poorly suggests that the network remained aware of both classes but lacked sufficient information about features in some months to separate them reliably.

When comparing these machine learning results to the conventional image analysis presented in Chapter 4, some parallels emerge. In both analyses, September consistently showed the clearest structural differences between tillage treatments, while differences became less pronounced as the season progressed. The conventional analysis showed a convergence of pore size and pore volume measurements later in the season, and this structural similarity likely contributed to the machine learning model's declining classification accuracy for certain months. However, while the traditional analysis focused on specific pre-defined parameters such as porosity and pore size, the CNN had access to the full complexity of the 2-D pore structure, including potentially subtle patterns not captured by conventional metrics. The variable performance across months suggests that some structural differences exist that are still not fully described or understood by either approach.

While the CNN architecture applied here was able to achieve promising results for treatment classification, the relatively simple design may limit its ability to capture more complex or subtle structural variations. A more advanced architecture, such as a DNN or a multi-resolution convolutional network could offer improvements. DNNs incorporate additional layers which allow the network to extract increasingly complex features. Initial layers typically identify simple features such as large shapes and edges between shapes (e.g. stones and crack/material boundaries in soil), while deeper levels may identify higher-order patterns such as aggregate arrangements, pore connectivity, or subtle textural variations that may be important for soil structure differentiation (Hou *et al.* 2022). The more subtle aspects of the soil structural differences observed in this dataset suggests that a deeper architecture may improve model performance by allowing the network to extract more detailed and informative features across multiple spatial scales.

Additionally, multi-resolution convolutional networks may be particularly well suited for this type of soil imaging data. These architectures combine both fine-scale and coarse-scale information simultaneously by processing high-resolution sub-regions alongside down sampled full images. This dual approach has shown strong performance in similar 3-D soil-root segmentation tasks (Soltaninejad *et al.* 2020), where complex biological structures are present at multiple spatial scales. Such models could allow future analysis to better capture both micro-scale pore structures and broader aggregate-level features, both of which are relevant for understanding tillage-induced structural changes.

While the present study focused on testing a relatively simple CNN architecture as a first step, future work could explore these more advanced models to improve classification accuracy, better handle the subtle seasonal variations observed, and potentially extract new structural insights not easily identified by conventional approaches.

Overall, this work demonstrates both the potential and current limitations of applying machine learning to soil structural analysis. While conventional image processing remains effective for capturing major structural changes, particularly immediately after tillage, machine learning offers the possibility of capturing more subtle or complex patterns within the large image datasets produced by X-ray CT. Combining both approaches may ultimately provide the most comprehensive understanding of how soil structure evolves over time in response to tillage and seasonal dynamics.

6.5 Conclusion

In this chapter, a CNN was applied to analyse soil structural data collected across a growing season from tilled and zero-till plots. The network successfully identified structural differences between tillage treatments when clear distinctions were present, particularly in September following cultivation. However, the network began to struggle, and classification performance became more variable as soil structure became more similar later in the season, highlighting the challenge of distinguishing between more subtle structures using this approach.

While the relatively simple CCN provided a useful initial framework for testing machine learning on soil structure, as in Chapter 5, its limitations became apparent as the structural differences between till and zero-till soils decreased in later months. The network also displayed some sensitivity due to overfitting during month-specific training, suggesting that the balance between model complexity and dataset size is critical when applying machine learning to soil structural analysis.

A smaller, custom network was intentionally selected in this study to limit computational complexity and reduce resource consumption. This decision reflects growing concerns about the environmental impact of increasingly large machine learning models, which consume significant energy and contribute indirectly to the same climate pressure that places stress on agricultural systems. However, for data of this complexity, where subtle and multiscale structural differences exist, more sophisticated models may ultimately be required to fully capture the relevant features and variances.

Future work may benefit from the use of more advanced architectures such as DNN or multi-resolution networks. Approaches such as those demonstrated by (Soltaninejad *et al.* 2020) have shown strong potential for extraction of both fine and coarse structural information in similar X-ray CT situations. The ability to capture multi-scale pore features and aggregate level organisation may allow improved classification performance in situations where the differences between treatments are less visually distinct.

Overall, while conventional image analysis remains highly effective for detecting major structural changes, machine learning offers the potential to extract additional structural insights from complex soil image datasets. With further refinement of model design and training approaches, machine learning may provide a valuable complementary tool for improving the analysis and understanding of soil structural dynamics, particularly in response to tillage and seasonal changes.

7. Conclusions and future work

7.1 Overview

The main aim of this thesis was to assess whether machine learning, particularly CNNs, could be used to identify and predict structural changes in soil from X-ray CT imagery. This investigation covered both temporal changes due to cultivation across a growing season and static differences in soil texture in moisture content. The study also explored the broader topic of the rapidly evolving applications of Artificial Intelligence, as well as how they can be applied in soil science. It also compared a custom-built CNN with a state-of-the-art (for the time) pretrained network (EfficientNetV2 B0) to evaluate the relative strengths and limitations of bespoke vs off-the-shelf solutions.

7.2. Summary of Findings

Across multiple chapters, the CNN models showed strong potential for detecting meaningful structural features in CT images of soil. The key findings include:

- In textural classification tasks, the custom CNN achieved high accuracy, especially when distinguishing between dry samples of sand and clay soils. This shows that visible structural differences between soil types, particularly under dry conditions, can be captured effectively by a lightweight neural network.
- Identifying structural differences between moisture levels proved to be more complicated. While sandy soil showed clear pore differentiation between wet and dry states, resulting in moderate classification accuracy, the network struggled with drier clay samples. The subtle structural differences between images, combined with a limited data set, likely reduced model performance.
- The 14-class classification experiment, involving both soil texture and water content, demonstrated the ability of the network to generalise

under complicated conditions. Despite the high number of classes, the custom CNN achieved nearly 98% training accuracy and showed improved feature identification through saliency maps. However, prediction accuracy on unseen data varied greatly, especially for visually similar classes, highlighting the need for further refinement.

- The comparison between the custom CNN and EfficientNet revealed significant issues when using a large pretrained model with extensive customisation and adjustments. Although EfficientNet performed well on training data, it suffered from overfitting, with a test accuracy of just 8% in the 14-class problem, and poor generalisation across all the binary classification tasks. The network consistently defaulted to a favoured class, demonstrating its inability to adapt to this specific dataset without extra tuning.

These findings align with the objectives set out at the beginning of this thesis, confirming that machine learning can indeed detect structural changes in soil, though success depends on model architecture, dataset balance, and the nature of the features being examined.

7.3. Broader Implications for Soil Science

Machine learning, when integrated with high-resolution imaging techniques such as X-ray CT, offers a promising approach to soil structural analysis. The datasets generated from such scans are often large, multidimensional, and rich in information that traditional metrics may fail to capture fully or require a large amount of training work to analyse. CNNs can help automate and accelerate the process of analysing these images, extracting meaningful patterns and features that may not be immediately accessible to human observers without extensive and specialised study.

This is especially valuable in precision agriculture, where real-time decisions about irrigation, fertiliser application, or crop choice may benefit from faster soil assessments. Similarly, in environmental and construction contexts, automated soil classification tools could be used to identify soil erosion risks or determine ground suitability.

However, for these tools to be applied reliably outside the laboratory, model generalisability must be improved. This will require datasets to be gathered across multiple soil types, locations and climates. This sort of expansion would allow for the development of robust, adaptable models capable of working across a wider variety of soil types and moisture levels.

The environmental costs of machine learning should also be considered. While smaller models such as those demonstrated here are relatively low-impact, training large models comes with significant energy and water suage. As AI adoption grows, sustainability in model design and deployment will become even more important. Choosing context-appropriate, efficient architectures can minimise environmental costs while still providing valuable insights.

7.4. Broader Technological Context

While the terms Artificial Intelligence and Machine Learning are often used interchangeably, there is a distinction between them: AI represents the overall goal of building intelligent systems, while ML describes on of the more popular methods used to achieve that goal. Most popular applications labelled as AI today, such as voice assistants, image recognition, and recommendation algorithms, are powered by machine learning techniques.

The rapid development of AI and ML are reshaping numerous fields, from scientific research and healthcare to finance and even the creative arts. Machine learning often offers substantial gains in efficiency, consistency and scalability. Repetitive and time-consuming tasks, such as the analysis of large and complicated datasets, can now be handled rapidly and reliably using trained models, freeing researchers to focus on interpretation and experimental design, and even designing larger scale experiments.

However, integrating AI into research and industries brings a number of challenges. Ethical concerns exist around job displacement, transparency, accountability, and the originality of AI-generated content. These issues are particularly of worry in research and creative fields, where the boundaries of ownership, creativity, and responsibility are all of concern.

An additional and increasingly pressing concern is the environmental impact of training and deploying large-scale machine learning models. Deep learning architectures, especially those with many layers and (potentially) millions of parameters, require significant computational power and energy. In this context, the decision to develop a custom, lightweight CNN for this thesis was not only practical given the dataset, but also an environmentally conscious choice. Although simpler in design, the custom model was more suited to the specific dataset challenges displayed in this research, and in future research could avoid the unnecessary environmental issues associated with larger, general-purpose networks.

7.4 Evaluating the use of Pretrained Models

This study also assessed the potential of using a pretrained model (EfficientNetV2 B0) for soil structure analysis. While EfficientNet is a powerful model trained on a vast dataset (ImageNet), its performance in this study was limited by:

- The small, specific nature of the dataset, which provided insufficient variation for the deeper layers of EfficientNet to generalise.
- Overfitting, which caused the model to pick a dominant class rather than learning meaningful distinctions.
- The lack of fine tuning and adapting the architecture, which are often needed to tailor a pretrained model to a specific task.

Although transfer learning is often praised for speeding up development, the amount of modification to make EfficientNet work effectively for the soil tasks here, such as additional data preprocessing or layer freezing, would have defeated the purpose of using a ready-made solution. In contrast, the custom CNN, although architecturally simpler, was well-adapted to the task and computationally lightweight.

7.5. Future Work

Building on the results of this thesis, several areas of future research are apparent:

- Larger and more diverse datasets. Acquiring soil samples from a range of sites with varying textures, structures, and biological activity would improve the model robustness and transferability.
- Multi-resolution neural networks. As demonstrated in other studies (Soltaninejad *et al.* 2020), networks capable of capturing both coarse and fine-scale features may be better suited for detecting subtle pore-level changes, such as seasonal dynamics.
- Longitudinal sampling. Regular sampling throughout the growing season proved to reveal important structural changes that would have been missed by only sampling at a single time point. Future work should continue this temporal approach or be aware of it when considering sample results.
- Integration with conventional analysis. Combining machine learning with traditional metrics and analysis could enhance the interpretability and model validation, further improving soil structural research.

7.6 Final Remarks

This thesis aimed to assess whether machine learning, particularly CNNs, could effectively identify structural differences in soils under a variety of conditions. The results strongly support this aim. Despite limitations in data diversity and subtle classification challenges, the networks successfully captured structural variation between soil textures, moisture levels, and tillage treatments.

The main barriers to further performance improvements lie in data availability, model generalisability, and model complexity. Future research should focus on building geographically and texturally diverse datasets, as well as exploring more advanced architectures, such as multi-resolution networks. Model complexity should also be balanced with environmental sustainability.

With these advancements, machine learning could offer transformative capabilities in soil science. This technology could support research, environmental monitoring, and more sustainable agriculture through powerful, interpretable and scalable tools.

Permissions

Parts of this thesis (Fig 2.1) are reproduced with permission from Kalendar, *X-ray Computed Tomography, Physics in Medicine & Biology,* Vol. 51, pp.29-43, © Institute of Physics and Engineering in Medicine. Reproduced by permission of IOP Publishing Ltd. All rights reserved.

Parts of this thesis (Fig. 2.2) are reproduced with permission from Grayling et al., *The application of X-ray micro Computed Tomography imaging for tracing particle movement in soil, Geoderma,* Vol. 321, pp. 8-14, © Elsevier (2018).

References

Aitkenhead, M., M. Coull, R. Gwatkin and D. Donnelly (2016). "Automated Soil Physical Parameter Assessment Using Smartphone and Digital Camera Imagery." Journal of Imaging **2**(4).

Allaire, S. E., S. Roulier and A. J. Cessna (2009). "Quantifying preferential flow in soils: A review of different techniques." Journal of Hydrology **378**(1): 179-204.

Allam, Z. and Z. A. Dhunny (2019). "On big data, artificial intelligence and smart cities." Cities **89**: 80-91.

Alskaf, K., S. J. Mooney, D. L. Sparkes, P. Wilson and S. Sjögersten (2021). "Short-term impacts of different tillage practices and plant residue retention on soil physical properties and greenhouse gas emissions." Soil and Tillage Research **206**: 104803.

Anderson, A. (1988). "Neural networks: learning from a computer cat." Nature **331**(6158): 657-659.

Applbaum, N. and Y. H. Applbaum (2005). The Use of Medical Computed Tomography (CT) Imaging in the Study of Ceramic and Clay Archaeological Artifacts from the Ancient Near East. X-rays for Archaeology. M. Uda, G. Demortier and I. Nakai. Dordrecht, Springer Netherlands**:** 231-245.

Atkinson, B. S., D. L. Sparkes and S. J. Mooney (2009). "The impact of soil structure on the establishment of winter wheat (Triticum aestivum)." European Journal of Agronomy **30**(4): 243-257.

Baghbani, A., T. Choudhury, S. Costa and J. Reiner (2022). "Application of artificial intelligence in geotechnical engineering: A state-of-the-art review." Earth-Science Reviews **228**: 103991.

Ball, B. C., T. Batey and L. J. Munkholm (2007). "Field assessment of soil structural quality – a development of the Peerlkamp test." Soil Use and Management **23**(4): 329-337.

Barke, S., M. B. James and N. Polikarpova (2022). "Grounded copilot: How programmers interact with code-generating models." arXiv preprint arXiv:2206.15000.

Belin, E., C. Douarre, N. Gillard, F. Franconi, J. Rojas-Varela, F. Chapeau-Blondeau, D. Demilly, J. Adrien, E. Maire and D. Rousseau (2018). "Evaluation of 3D/2D Imaging and Image Processing Techniques for the Monitoring of Seed Imbibition." Journal of Imaging **4**(7).

Borken, W. and E. Matzner (2009). "Reappraisal of drying and wetting effects on C and N mineralization and fluxes in soils." Global Change Biology **15**(4): 808-824.

Braudeau, E., A. T. Assi, H. Boukcim and R. Mohtar (2014). "Physics of the soil medium organization part 1: thermodynamic formulation of the pedostructure water retention and shrinkage curves." Frontiers in Environmental Science **Volume 2 - 2014**.

Bryk, M., B. Kołodziej, A. Słowińska-Jurkiewicz and M. Jaroszuk-Sierocińska (2017). "Evaluation of soil structure and physical properties influenced by weather conditions during autumn-winter-spring season." Soil and Tillage Research **170**: 66-76.

Bulat, A. Y., J.

Tzimiropoulos, G. (2018). To learn image super-resolution, use a GAN to learn how to do image degradation first. Proceedings of the European Conference on Computer Vision (ECCV), Munich, Germanmy.

Cannell, R. Q. (1985). "Reduced tillage in north-west Europe—A review." Soil and Tillage Research **5**(2): 129-177.

Cao, L. (2022). "AI in Finance: Challenges, Techniques, and Opportunities." ACM Comput. Surv. **55**(3): Article 64.

Cetinic, E. and J. She (2021). Understanding and Creating Art with AI: Review and Outlook.

Cha, Y.-J., W. Choi and O. Büyüköztürk (2017). "Deep Learning-Based Crack Damage Detection Using Convolutional Neural Networks." Computer-Aided Civil and Infrastructure Engineering **32**(5): 361-378.

Chaerle, L., D. Hagenbeek, E. De Bruyne, R. Valcke and D. Van Der Straeten (2004). "Thermal and chlorophyll-fluorescence imaging distinguish plant-pathogen interactions at an early stage." Plant Cell Physiol **45**(7): 887-896.

Chaerle, L., I. Leinonen, H. G. Jones and D. Van Der Straeten (2007). "Monitoring and screening plant populations with combined thermal and chlorophyll fluorescence imaging." J Exp Bot **58**(4): 773-784.

Chertkov, V. Y. (2002). "Modelling cracking stages of saturated soils as they dry and shrink." European Journal of Soil Science **53**(1): 105-118.

Coeckelbergh, M. (2023). "The Work of Art in the Age of AI Image Generation: Aesthetics and Human-Technology Relations as Process and Performance." Journal of Human-Technology Relations **1**.

Connolly, R. D. (1998). "Modelling effects of soil structure on the water balance of soil–crop systems: a review." Soil and Tillage Research **48**(1): 1-19.

Cooper, H. V., S. Sjögersten, R. M. Lark and S. J. Mooney (2021). "To till or not to till in a temperate ecosystem? Implications for climate change mitigation." Environmental Research Letters **16**(5): 054022.

Czedik-Eysenberg, A., S. Seitner, U. Güldener, S. Koemeda, J. Jez, M. Colombini and A. Djamei (2018). "The 'PhenoBox', a flexible, automated, open-source plant phenotyping solution." New Phytologist **219**(2): 808-823.

Dal Ferro, N., P. Charrier and F. Morari (2013). "Dual-scale micro-CT assessment of soil structure in a long-term fertilization experiment." Geoderma **204-205**: 84-93.

DeLaune, P. B., J. W. Sij and L. J. Krutz (2013). "Impact of soil aeration on runoff characteristics in dual-purpose no-till wheat systems." Journal of Soil and Water Conservation **68**(4): 315-324.

Denef, K., J. Six, H. Bossuyt, S. D. Frey, E. T. Elliott, R. Merckx and K. Paustian (2001). "Influence of dry–wet cycles on the interrelationship between aggregate, particulate organic matter, and microbial community dynamics." Soil Biology and Biochemistry **33**(12): 1599-1611.

Dexter, A. R. (1988). "Advances in characterization of soil structure." Soil and Tillage Research **11**(3): 199-238.

Dexter, A. R. and N. R. A. Bird (2001). "Methods for predicting the optimum and the range of soil water contents for tillage based on the water retention curve." Soil and Tillage Research **57**(4): 203-212.

Dexter, A. R., B. Kroesbergen and H. Kuipers (1984). "Some mechanical properties of aggregates of top soils from the IJsselmeer polders. 2. Remoulded soil aggregates and the effects of wetting and drying cycles." Netherlands Journal of Agricultural Science **32**(3): 215-227.

Diel, J., H.-J. Vogel and S. Schlüter (2019). "Impact of wetting and drying cycles on soil structure dynamics." Geoderma **345**: 63-71.

Douarre, C., R. Schielein, C. Frindel, S. Gerth and D. Rousseau (2018). "Transfer Learning from Synthetic Data Applied to Soil-Root Segmentation in X-Ray Tomography Images." Journal of Imaging **4**(5).

Doube, M., M. M. Kłosowski, I. Arganda-Carreras, F. P. Cordelières, R. P. Dougherty, J. S. Jackson, B. Schmid, J. R. Hutchinson and S. J. Shefelbine (2010). "BoneJ: Free and extensible bone image analysis in ImageJ." Bone **47**(6): 1076-1079.

Downie, H., N. Holden, W. Otten, A. J. Spiers, T. A. Valentine and L. X. Dupuy (2012). "Transparent Soil for Imaging the Rhizosphere." PLOS ONE **7**(9): e44276.

Elder, A., M. O. D. Cappelli, C. Ring and N. Saedi (2024). "Artificial intelligence in cosmetic dermatology: An update on current trends." Clinics in Dermatology.

Esmaeilzadeh, P. (2020). "Use of AI-based tools for healthcare purposes: a survey study from consumers’ perspectives." BMC Medical Informatics and Decision Making **20**(1): 170.

Farfani, H. A., F. Behnamfar and A. Fathollahi (2015). "Dynamic analysis of soil-structure interaction using the neural networks and the support vector machines." Expert Systems with Applications **42**(22): 8971-8981.

Ferreira, T., L. Pires, D. Wildenschild, R. Heck and A. Antonino (2018). "X-ray microtomography analysis of lime application effects on soil porous system." Geoderma **324**: 119-130.

Ferreira, T. R., L. F. Pires, D. Wildenschild, R. J. Heck and A. C. D. Antonino (2018). "X-ray microtomography analysis of lime application effects on soil porous system." Geoderma **324**: 119-130.

Fierer, N. and J. Schimel (2002). "Fierer N, Schmel JP.. Effects of drying-rewetting frequency on soil and nitrogen transformations. Soil Biol Biochem 34: 777-787." Soil Biology and Biochemistry **34**: 777-787.

Flores-Calero, M., C. A. Astudillo, D. Guevara, J. Maza, B. S. Lita, B. Defaz, J. S. Ante, D. Zabala-Blanco and J. M. Armingol Moreno (2024). "Traffic Sign Detection and Recognition Using YOLO Object Detection Algorithm: A Systematic Review." Mathematics **12**(2): 297.

Gajawelli, N., S. Tsao, M. Kromnick, M. Nelson and N. Leporé (2019). "Image Postprocessing Adoption Trends in Clinical Medical Imaging." Journal of the American College of Radiology.

Galdos, M. V., L. F. Pires, H. V. Cooper, J. C. Calonego, C. A. Rosolem and S. J. Mooney (2019). "Assessing the long-term effects of zero-tillage on the macroporosity of Brazilian soils using X-ray Computed Tomography." Geoderma **337**: 1126-1135.

Garcea, S. C., Y. Wang and P. J. Withers (2018). "X-ray computed tomography of polymer composites." Composites Science and Technology **156**: 305-319.

Gebhardt, M. R., T. C. Daniel, E. E. Schweizer and R. R. Allmaras (1985). "Conservation Tillage." Science **230**(4726): 625-630.

Gitis, V., A. Adin, A. Nasser, J. Gun and O. Lev (2002). "Fluorescent dye labeled bacteriophages—a new tracer for the investigation of viral transport in porous media: 1. Introduction and characterization." Water Research **36**(17): 4227-4234.

Goodfellow, I. J., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville and Y. Bengio (2014). "Generative Adversarial Networks." Communications of the ACM **63**(11): 139-144.

Grayling, K. M., S. D. Young, C. J. Roberts, M. I. de Heer, I. M. Shirley, C. J. Sturrock and S. J. Mooney (2018). "The application of X-ray micro Computed Tomography imaging for tracing particle movement in soil." Geoderma **321**: 8-14.

Gubelmann, R. (2023). "A Loosely Wittgensteinian Conception of the Linguistic Understanding of Large Language Models like BERT, GPT-3, and ChatGPT." Grazer Philosophische Studien **99**(4): 485-523.

Gunda, P., P. K. Thallapally, Y.-Y. Lin, M. Laird Forrest and C. J. Berkland (2009). "Nanoparticles for biomedical imaging AU - Nune, Satish K." Expert Opinion on Drug Delivery **6**(11): 1175-1194.

Guo, Y., R. Fan, X. Zhang, Y. Zhang, D. Wu, N. McLaughlin, S. Zhang, X. Chen, S. Jia and A. Liang (2020). "Tillage-induced effects on SOC through changes in aggregate stability and soil pore structure." Sci Total Environ **703**: 134617.

Guo, Y., J. Stein, G. Wu and A. Krishnamurthy (2019). *SAU-Net: A Universal Deep Network for Cell Counting*. Proceedings of the 10th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics. Niagara Falls, NY, USA, Association for Computing Machinery**:** 299–306.

Gupta, R., S. Tiwari and P. Chaudhary (2025). Computational Foundation of Generative AI Models. Generative AI: Techniques, Models and Applications. R. Gupta, S. Tiwari and P. Chaudhary. Cham, Springer Nature Switzerland**:** 23-44.

Gupta, S. C., P. P. Sharma and S. A. DeFranchi (1989). Compaction Effects on Soil Structure11Contribution from the Department of Soil Science and the Minnesota Agricultural Experiment Station, University of Minnesota, St. Paul, MN 55108. Paper No. 15636, Science Journal Series. Advances in Agronomy. N. C. Brady, Academic Press. **42:** 311-338.

Habibi, A. H. and H. E. Jahani (2017). Guide to Convolutional Neural Networks : A Practical Application to Traffic-Sign Detection and Classification. Cham, SWITZERLAND, Springer.

Hassoun, A., A. E.-D. Bekhit, A. R. Jambrak, J. M. Regenstein, F. Chemat, J. D. Morton, M. Gudjónsdóttir, M. Carpena, M. A. Prieto, P. Varela, R. N. Arshad, R. M. Aadil, Z. Bhat and Ø. Ueland (2024). "The fourth industrial revolution in the food industry—part II: Emerging food trends." Critical Reviews in Food Science and Nutrition **64**(2): 407-437.

Herath, H. M. K. K. M. B. and M. Mittal (2022). "Adoption of artificial intelligence in smart cities: A comprehensive review." International Journal of Information Management Data Insights **2**(1): 100076.

Hernandez-Rabadan, D. L., F. Ramos-Quintana and J. Guerrero Juk (2014). "Integrating SOMs and a Bayesian classifier for segmenting diseased plants in uncontrolled environments." ScientificWorldJournal **2014**: 214674.

Heskes, T. M. and B. Kappen (1991). "Learning processes in neural networks." Phys Rev A **44**(4): 2718-2726.

Heung, B., H. C. Ho, J. Zhang, A. Knudby, C. E. Bulmer and M. G. Schmidt (2016). "An overview and comparison of machine-learning techniques for classification purposes in digital soil mapping." Geoderma **265**: 62-77.

Holmes, W. and I. Tuomi (2022). "State of the art and practice in AI in education." European Journal of Education **57**(4): 542-570.

Hong, Y., U. Hwang, J. Yoo and S. Yoon (2019). "How Generative Adversarial Networks and Their Variants Work: An Overview." ACM Comput. Surv. **52**(1): Article 10.

Hou, L., W. Gao, F. der Bom, Z. Weng, C. L. Doolette, A. Maksimenko, D. Hausermann, Y. Zheng, C. Tang, E. Lombi and P. M. Kopittke (2022). "Use of X-ray tomography for examining root architecture in soils." Geoderma **405**: 115405.

Hounsfield, G. N. (1980). "Computed medical imaging." Science **210**(4465): 22.

Huang, L.-K. and M.-J. J. Wang (1995). "Image thresholding by minimizing the measures of fuzziness." Pattern Recognition **28**(1): 41-51.

International, V. (2021). *GenStat for Windows 20th Edition*. Hemel Hempstead, UK, VSN International.

IPCC (2023). *Climate Change 2023: Synthesis Report*. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, H. Lee and J. Romero (eds.)]. Geneva, Switzerland: 35–115.

Jackisch, C., K. Germer, T. Graeff, I. Andrä, K. Schulz, M. Schiedung, J. Haller-Jans, J. Schneider, J. Jaquemotte, P. Helmer, L. Lotz, A. Bauer, I. Hahn, M. Šanda, M. Kumpan, J. Dorner, G. de Rooij, S. Wessel-Bothe, L. Kottmann, S. Schittenhelm and W. Durner (2020). "Soil moisture and matric potential – an open field comparison of sensor systems." Earth Syst. Sci. Data **12**(1): 683-697.

Jarvis, N. J. (2007). "A review of non-equilibrium water flow and solute transport in soil macropores: principles, controlling factors and consequences for water quality." European Journal of Soil Science **58**(3): 523-546.

Jayawardane, N. and K. Chan (1994). "The management of soil physical properties limiting crop production in Australian sodic soils - a review." Soil Research **32**(1): 13-44.

Jiao, L. and J. Zhao (2019). "A Survey on the New Generation of Deep Learning in Image Processing." IEEE Access **7**: 172231-172263.

Kalender, W. A. (2006). "X-ray computed tomography." Physics in Medicine & Biology **51**(13): R29.

Kammoun, A., R. Slama, H. Tabia, T. Ouni and M. Abid (2022). "Generative Adversarial Networks for Face Generation: A Survey." ACM Comput. Surv. **55**(5): Article 94.

Katuwal, S., E. Arthur, M. Tuller, P. Moldrup and L. W. de Jonge (2015). "Quantification of Soil Pore Network Complexity with X-ray Computed Tomography and Gas Transport Measurements Soil Physics & Hydrology." Soil Science Society of America Journal **79**(6): 1577-1589.

Kim, S. W., J. H. Kong, S. W. Lee and S. Lee (2022). "Recent Advances of Artificial Intelligence in Manufacturing Industrial Sectors: A Review." International Journal of Precision Engineering and Manufacturing **23**(1): 111-129.

Kjaergaard, C., T. G. Poulsen, P. Moldrup and L. W. de Jonge (2004). "Colloid mobilization and transport in undisturbed soil columns. I. Pore structure characterization and tritium transport." Vadose Zone Journal **3**(2): 413-423.

Koestel, J. (2018). "SoilJ: An ImageJ Plugin for the Semiautomatic Processing of Three-Dimensional X-ray Images of Soils." Vadose Zone Journal **17**(1).

Lai, J., J. Qiu, Z. Feng, J. Chen and H. Fan (2016). "Prediction of Soil Deformation in Tunnelling Using Artificial Neural Networks." Comput Intell Neurosci **2016**: 6708183.

Lai, Y. (2019). "A Comparison of Traditional Machine Learning and Deep Learning in Image Recognition." Journal of Physics: Conference Series **1314**(1): 012148.

Lal, R., D. C. Reicosky and J. D. Hanson (2007). "Evolution of the plow over 10,000 years and the rationale for no-till farming." Soil and Tillage Research **93**(1): 1-12.

Lamandé, M., R. Labouriau, M. Holmstrup, S. B. Torp, M. H. Greve, G. Heckrath, B. V. Iversen, L. W. de Jonge, P. Moldrup and O. H. Jacobsen (2011). "Density of macropores as related to soil and earthworm community parameters in cultivated grasslands." Geoderma **162**(3): 319-326.

Le Bissonnais, Y. (2016). "Le Bissonnais, Y. (1996). Aggregate stability and assessment of crustability and erodibility: 1. Theory and methodology. European Journal of Soil Science, 47, 425–437." European Journal of Soil Science **67**(1): 2-4.

Li, H., W. Leng, Y. Zhou, F. Chen, Z. Xiu and D. Yang (2014). "Evaluation models for soil nutrient based on support vector machine and artificial neural networks." ScientificWorldJournal **2014**: 478569.

Li, Y., Y. Y. Ruan, M. T. Kasson, E. L. Stanley, C. P. D. T. Gillett, A. J. Johnson, M. N. Zhang and J. Hulcr (2018). "Structure of the Ambrosia Beetle (Coleoptera: Curculionidae) Mycangia Revealed Through Micro-Computed Tomography." Journal of Insect Science **18**(5).

Lin, T.-Y., P. Goyal, R. Girshick, K. He and P. Dollár (2017). Focal loss for dense object detection. Proceedings of the IEEE international conference on computer vision.

Lindsey, R., A. Daluiski, S. Chopra, A. Lachapelle, M. Mozer, S. Sicular, D. Hanel, M. Gardner, A. Gupta, R. Hotchkiss and H. Potter (2018). "Deep neural network improves fracture detection by clinicians." Proceedings of the National Academy of Sciences **115**(45): 11591.

Lindstrom, M. J., D. Lobb and T. Schumacher (2001). "Tillage Erosion: An Overview." Annals of Dryland Research **40**: 345-358.

Liu, H., T. Colombi, O. Jäck, T. Keller and M. Weih (2022). "Effects of soil compaction on grain yield of wheat depend on weather conditions." Science of The Total Environment **807**: 150763.

Liu, L., M. Ji and M. Buchroithner (2018). "Transfer Learning for Soil Spectroscopy Based on Convolutional Neural Networks and Its Application in Soil Clay Content Mapping Using Hyperspectral Imagery." Sensors (Basel) **18**(9).

Lowe, A., N. Harrison and A. P. French (2017). "Hyperspectral image analysis techniques for the detection and classification of the early onset of plant disease and stress." Plant Methods **13**: 80.

Luo, L., H. Lin and S. Li (2010). "Quantification of 3-D soil macropore networks in different soil types and land uses using computed tomography." Journal of Hydrology **393**(1): 53-64.

Luo, L. F., H. Lin and P. Halleck (2008). "Quantifying soil structure and preferential flow in intact soil using x-ray computed tomography." Soil Science Society of America Journal **72**(4): 1058-1069.

Ma, R., C. Cai, Z. Li, J. Wang, T. Xiao, G. Peng and W. Yang (2015). "Evaluation of soil aggregate microstructure and stability under wetting and drying cycles in two Ultisols using synchrotron-based X-ray micro-computed tomography." Soil and Tillage Research **149**: 1-11.

Mairhofer, S., J. Johnson, C. Sturrock, M. Bennett, S. Mooney and T. Pridmore (2016). "Visual tracking for the recovery of multiple interacting plant root systems from X-ray CT images." Machine Vision and Applications **27**(5): 721-734.

Mairhofer, S., C. J. Sturrock, M. J. Bennett, S. J. Mooney and T. P. Pridmore (2015). "Extracting multiple interacting root systems using X-ray microcomputed tomography." Plant J **84**(5): 1034-1043.

Mairhofer, S., S. Zappala, S. R. Tracy, C. Sturrock, M. Bennett, S. J. Mooney and T. Pridmore (2012). "RooTrak: Automated Recovery of Three-Dimensional Plant Root Architecture in Soil from X-Ray Microcomputed Tomography Images Using Visual Tracking." Plant Physiology **158**(2): 561-569.

Mak, K.-K., Y.-H. Wong and M. R. Pichika (2022). Artificial Intelligence in Drug Discovery and Development. Drug Discovery and Evaluation: Safety and Pharmacokinetic Assays. F. J. Hock and M. K. Pugsley. Cham, Springer International Publishing**:** 1-38.

McDowellboyer, L. M., J. R. Hunt and N. Sitar (1986). "Particle-Transport through Porous-Media." Water Resources Research **22**(13): 1901-1921.

MIT News (2020). "Shrinking deep learning's carbon footprint." https://climate.mit.edu/posts/shrinking-deep-learnings-carbon-footprint Accessed 22/11/2024 2024.

Mooney, S. J. (2002). "Three-dimensional visualization and quantification of soil macroporosity and water flow patterns using computed tomography." Soil Use and Management **18**(2): 142-151.

Moreira de Melo, T. and O. C. Pedrollo (2015). "Artificial Neural Networks for Estimating Soil Water Retention Curve Using Fitted and Measured Data." Applied and Environmental Soil Science **2015**: 16.

Mowbray, A., P. Chung and G. Greenleaf (2020). "Utilising AI in the legal assistance sector—Testing a role for legal information institutes." Computer Law & Security Review **38**: 105407.

Naveed, M., S. Hamamoto, K. Kawamoto, T. Sakaki, M. Takahashi, T. Komatsu, P. Moldrup, M. Lamande, D. Wildenschild, M. Prodanovic and L. W. de Jonge (2013). "Correlating Gas Transport Parameters and X-Ray Computed Tomography Measurements in Porous Media." Soil Science **178**(2): 60-68.

OpenAI. (2022). "ChatGPT (Generative Pre-trained Transformer)." from https://chat.openai.com/.

Paradelo, M., P. Moldrup, E. Arthur, M. Naveed, M. Holmstrup, J. E. Lopez-Periago and L. W. de Jonge (2013). "Effects of Past Copper

Contamination and Soil Structure on Copper Leaching from Soil." Journal of Environmental Quality **42**(6): 1852-1862.

Paul, A., R. Machavaram, Ambuj, D. Kumar and H. Nagar (2024). "Smart solutions for capsicum Harvesting: Unleashing the power of YOLO for Detection, Segmentation, growth stage Classification, Counting, and real-time mobile identification." Computers and Electronics in Agriculture **219**: 108832.

Peerlkamp, P. (1959). *A visual method of soil structure evaluation*: 216–221.

Peng, X., Z. B. Zhang, L. Gan and S. Yoshida (2016). "Linking Soil Shrinkage Behavior and Cracking in Two Paddy Soils as Affected by Wetting and Drying Cycles." Soil Science Society of America Journal **80**(5): 1145-1156.

Perreault, J. M. and J. K. Whalen (2006). "Earthworm burrowing in laboratory microcosms as influenced by soil temperature and moisture." Pedobiologia **50**(5): 397-403.

Perry, N., M. Srivastava, D. Kumar and D. Boneh (2022). "Do Users Write More Insecure Code with AI Assistants?" arXiv preprint arXiv:2211.03622.

Petrovic, A. M., J. E. Siebert and P. E. Rieke (1982). "Soil Bulk-Density Analysis in 3 Dimensions by Computed Tomographic Scanning." Soil Science Society of America Journal **46**(3): 445-450.

Pillarisetty, R. and P. Mishra (2022). "A Review of AI (Artificial Intelligence) Tools and Customer Experience in Online Fashion Retail." International Journal of E-Business Research (IJEBR) **18**(2): 1-12.

Pires, L. F., A. C. Auler, W. L. Roque and S. J. Mooney (2020). "X-ray microtomography analysis of soil pore structure dynamics under wetting and drying cycles." Geoderma **362**: 114103.

Pound, M. P., J. A. Atkinson, A. J. Townsend, M. H. Wilson, M. Griffiths, A. S. Jackson, A. Bulat, G. Tzimiropoulos, D. M. Wells, E. H. Murchie, T. P. Pridmore and A. P. French (2017). "Deep machine learning provides state-of-the-art performance in image-based plant phenotyping." Gigascience **6**(10): 1-10.

Qi, W., C. Wang, Z. Zhang, M. Huang and J. Xu (2022). "Experimental Investigation on the Impact of Drying–Wetting Cycles on the Shrink–Swell Behavior of Clay Loam in Farmland." Agriculture **12**(2): 245.

Rajaram, G. and D. C. Erbach (1999). "Effect of wetting and drying on soil physical properties1Joint contribution USDA–Agricultural Research Service and Iowa State University. Journal paper no. J-17482 of the Iowa Agricultural and Home Economics Experiment Station, Ames, IA, project no. 2462, and supported by Hatch Act and State of Iowa funds1." Journal of Terramechanics **36**(1): 39-49.

Redmon, J., S. Divvala, R. Girshick and A. Farhadi (2016). You only look once: Unified, real-time object detection. Proceedings of the IEEE conference on computer vision and pattern recognition.

Rogers, E. D., D. Monaenkova, M. Mijar, A. Nori, D. I. Goldman and P. N. Benfey (2016). "X-Ray Computed Tomography Reveals the Response of Root System Architecture to Soil Texture." Plant Physiol **171**(3): 2028-2040.

Salu, Y. (1985). "Learning and coding of concepts in neural networks." Biosystems **18**(1): 93-103.

Sandoval, G., H. Pearce, T. Nys, R. Karri, B. Dolan-Gavitt and S. Garg (2022). "Security Implications of Large Language Model Code Assistants: A User Study." arXiv preprint arXiv:2208.09727.

Sandri, R., T. Anken, T. Hilfiker, L. Sartori and H. Bollhalder (1998). "Comparison of methods for determining cloddiness in seedbed preparation." Soil and Tillage Research **45**(1): 75-90.

Sarmadian, F. and R. Taghizadeh Mehrjardi (2009). Modeling of Some Soil Properties Using Artificial Neural Network and Multivariate Regression in Gorgan Province, North Iran.

Sayem, H. and L.-W. Kong (2016). Effects of Drying-Wetting Cycles on Soil-Water Characteristic Curve.

Schindelin, J., I. Arganda-Carreras, E. Frise, V. Kaynig, M. Longair, T. Pietzsch, S. Preibisch, C. Rueden, S. Saalfeld, B. Schmid, J.-Y. Tinevez, D. J. White, V. Hartenstein, K. Eliceiri, P. Tomancak and A. Cardona (2012).

"Fiji: an open-source platform for biological-image analysis." Nature Methods **9**(7): 676-682.

Schlüter, S., U. Weller and H.-J. Vogel (2010). "Segmentation of X-ray microtomography images of soil using gradient masks." Computers & Geosciences **36**(10): 1246-1251.

Schlüter, S., U. Weller and H.-J. Vogel (2011). "Soil-structure development including seasonal dynamics in a long-term fertilization experiment." Journal of Plant Nutrition and Soil Science **174**(3): 395-403.

Schmidhuber, J. (2015). "Deep learning in neural networks: An overview." Neural Networks **61**: 85-117.

Schon, N. L., A. D. Mackay, R. A. Gray, C. van Koten and M. B. Dodd (2017). "Influence of earthworm abundance and diversity on soil structure and the implications for soil services throughout the season." Pedobiologia **62**: 41-47.

Sezgin, E., S.-A. Hussain, S. Rust and Y. Huang (2023). "Extracting Medical Information From Free-Text and Unstructured Patient-Generated Health Data Using Natural Language Processing Methods: Feasibility Study With Real-world Data." JMIR Form Res **7**: e43014.

Sharma, R. (2021). Artificial Intelligence in Agriculture: A Review. 2021 5th International Conference on Intelligent Computing and Control Systems (ICICCS).

Singh, A., B. Ganapathysubramanian, A. K. Singh and S. Sarkar (2016). "Machine Learning for High-Throughput Stress Phenotyping in Plants." Trends Plant Sci **21**(2): 110-124.

Soltaninejad, M., C. J. Sturrock, M. Griffiths, T. P. Pridmore and M. P. Pound (2020). "Three Dimensional Root CT Segmentation Using Multi-Resolution Encoder-Decoder Networks." IEEE Transactions on Image Processing **29**: 6667-6679.

Soori, M., B. Arezoo and R. Dastres (2023). "Artificial intelligence, machine learning and deep learning in advanced robotics, a review." Cognitive Robotics **3**: 54-70.

Soto-Gomez, D., P. Perez-Rodriguez, J. E. Lopez-Periago and M. Paradelo (2016). "Sepia ink as a surrogate for colloid transport tests in porous media." Journal of Contaminant Hydrology **191**: 88-98.

Soto-Gomez, D., P. Perez-Rodriguez, L. Vazquez-Juiz, J. E. Lopez-Periago and M. Paradeloa (2018). "Linking pore network characteristics extracted from CT images to the transport of solute and colloid tracers in soils under different tillage managements." Soil & Tillage Research **177**: 145-154.

Su, Z., J. Zhang, W. Wu, D. Cai, J. Lv, G. Jiang, J. Huang, J. Gao, R. Hartmann and D. Gabriels (2007). "Effects of conservation tillage practices on winter wheat water-use efficiency and crop yield on the Loess Plateau, China." Agricultural Water Management **87**(3): 307-314.

Sutton, R. S. and A. G. Barto (2018). Reinforcement Learning, second edition: An Introduction, MIT Press.

Swathi, T. and S. Sudha (2023). "Crop classification and prediction based on soil nutrition using machine learning methods." International Journal of Information Technology **15**(6): 2951-2960.

Szabo, I., J. Sun, G. Feng, J. Kanfoud, T.-H. Gan and C. Selcuk (2017). "Automated Defect Recognition as a Critical Element of a Three Dimensional X-ray Computed Tomography Imaging-Based Smart Non-Destructive Testing Technique in Additive Manufacturing of Near Net-Shape Parts." Applied Sciences **7**(11): 1156.

Tadiboina, S. N. (2022). "The Use of AI In Advanced Medical Imaging." Journal Of Positive School Psychology **6**(11): 1939-1946.

Tan, M. and Q. Le (2021). *EfficientNetV2: Smaller Models and Faster Training*. Proceedings of the 38th International Conference on Machine Learning. M. Marina and Z. Tong. Proceedings of Machine Learning Research, PMLR. **139:** 10096--10106.

Tang, C.-S., B. Shi, C. Liu, W.-B. Suo and L. Gao (2011). "Experimental characterization of shrinkage and desiccation cracking in thin clay layer." Applied Clay Science **52**(1): 69-77.

Team, R. (2020). *RStudio: Integrated Development for R.* PBC, Boston, MA.

Telfair, D., M. R. Garner and D. Miars (1957). "The Restoration of a Structurally Degenerated Soil." Soil Science Society of America Journal **21**(2): 131-134.

Tien Bui, D., N.-D. Hoang and V.-H. Nhu (2019). "A swarm intelligence-based machine learning approach for predicting soil shear strength for road construction: a case study at Trung Luong National Expressway Project (Vietnam)." Engineering with Computers **35**(3): 955-965.

Touzet, C. (2023). "Using AI to support people with disability in the labour market."

Trachsel, S., S. M. Kaeppler, K. M. Brown and J. P. Lynch (2011). "Shovelomics: high throughput phenotyping of maize (Zea mays L.) root architecture in the field." Plant and Soil **341**(1-2): 75-87.

Tracy, S. R., C. R. Black, J. A. Roberts, C. Sturrock, S. Mairhofer, J. Craigon and S. J. Mooney (2012). "Quantifying the impact of soil compaction on root system architecture in tomato (Solanum lycopersicum) by X-ray micro-computed tomography." Ann Bot **110**(2): 511-519.

Tracy, S. R., K. R. Daly, C. J. Sturrock, N. M. J. Crout, S. J. Mooney and T. Roose (2015). "Three-dimensional quantification of soil hydraulic properties using X-ray Computed Tomography and image-based modeling." Water Resources Research **51**(2): 1006-1022.

Tuller, M., R. Kulkarni and W. Fink (2013). Segmentation of X-Ray CT Data of Porous Materials: A Review of Global and Locally Adaptive Algorithms. Soil–Water–Root Processes: Advances in Tomography and Imaging**:** 157-182.

Van Elsas, J. D., J. T. Trevors, A. S. Rosado and P. Nannipieri (2019). Modern soil microbiology, CRC press.

Vayadande, K., S. Bhemde, V. Rajguru, P. Ugile, R. Lade and N. Raut (2023). AI-Based Image Generator Web Application using OpenAI's DALL-E System. 2023 International Conference on Recent Advances in Science and Engineering Technology (ICRASET).

Vaz, B. and Á. Figueira (2024). "GANs in the Panorama of Synthetic Data Generation Methods." ACM Trans. Multimedia Comput. Commun. Appl. **21**(1): Article 3.

Velev, D. and P. Zlateva (2023). "CHALLENGES OF ARTIFICIAL INTELLIGENCE APPLICATION FOR DISASTER RISK MANAGEMENT." Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci. **XLVIII-M-1-2023**: 387-394.

Venkateswara Reddy, L., D. Ganesh, M. Sunil Kumar, S. Gogula, M. Rekha and A. Sehgal (2024). "Applying machine learning to soil analysis for accurate farming." MATEC Web Conf. **392**: 01124.

Vogel, H. J., S. Bartke, K. Daedlow, K. Helming, I. Kögel-Knabner, B. Lang, E. Rabot, D. Russell, B. Stößel, U. Weller, M. Wiesmeier and U. Wollschläger (2018). "A systemic approach for modeling soil functions." SOIL **4**(1): 83-92.

Vogel, H. J., U. Weller and S. Schlüter (2010). "Quantification of soil structure based on Minkowski functions." Computers & Geosciences **36**(10): 1236-1245.

Wakchaure, M., B. K. Patle and A. K. Mahindrakar (2023). "Application of AI techniques and robotics in agriculture: A review." Artificial Intelligence in the Life Sciences **3**: 100057.

Wakefield, J. (2021). *AI draws a dog-walking baby radish in a tutu*. BBC News. https://www.bbc.co.uk/news/technology-55559463.

Wang, Y., C. Li and Y. Hu (2018). "X-ray computed tomography (CT) observations of crack damage evolution in soil-rock mixture during uniaxial deformation." Arabian Journal of Geosciences **11**(9).

Wang, Y., C. H. Li and Y. Z. Hu (2018). "X-ray computed tomography (CT) observations of crack damage evolution in soil-rock mixture during uniaxial deformation." Arabian Journal of Geosciences **11**(9).

Wang, Z. J., R. Turko, O. Shaikh, H. Park, N. Das, F. Hohman, M. Kahng and D. H. P. Chau (2021). "CNN Explainer: Learning Convolutional Neural Networks with Interactive Visualization." IEEE Transactions on Visualization and Computer Graphics **27**(2): 1396-1406.

Wieland, R., C. Ukawa, M. Joschko, A. Krolczyk, G. Fritsch, T. B. Hildebrandt, O. Schmidt, J. Filser and J. J. Jimenez (2021). "Use of deep learning for structural analysis of computer tomography images of soil samples." Royal Society Open Science **8**(3): 201275.

Wilson, D. W. (1998). Soil-Pile-Superstructure Interaction in Liquefying Sand Soft Clay Doctoral Dissertation, University of California at Davis.

Wilson, G. V. and R. J. Luxmoore (1988). "Infiltration, Macroporosity, and Mesoporosity Distributions on Two Forested Watersheds." Soil Science Society of America Journal **52**: 329-335.

Ye, X.-W., T. Jin and Y.-M. Chen (2022). "Machine learning-based forecasting of soil settlement induced by shield tunneling construction." Tunnelling and Underground Space Technology **124**: 104452.

You, S., B. Lei, S. Wang, C. K. Chui, A. C. Cheung, Y. Liu, M. Gan, G. Wu and Y. Shen (2023). "Fine Perceptive GANs for Brain MR Image Super-Resolution in Wavelet Domain." IEEE Transactions on Neural Networks and Learning Systems **34**(11): 8802-8814.

Zaidi, M., N.-D. Ahfir, A. Alem, S. Taibi, B. El Mansouri, Y. Zhang and H. Wang (2021). "Use of X-ray computed tomography for studying the desiccation cracking and self-healing of fine soil during drying–wetting paths." Engineering Geology **292**: 106255.

Zangerlé, A., A. Pando and P. Lavelle (2011). "Do earthworms and roots cooperate to build soil macroaggregates? A microcosm experiment." Geoderma **167-168**: 303-309.

Zhai, X., X. Chu, C. S. Chai, M. S. Y. Jong, A. Istenic, M. Spector, J.-B. Liu, J. Yuan and Y. Li (2021). "A Review of Artificial Intelligence (AI) in Education from 2010 to 2020." Complexity **2021**(1): 8812542.

Zhai, Y., J. A. Thomasson, J. E. Boggess and R. Sui (2006). "Soil texture classification with artificial neural networks operating on remote sensing data." Computers and Electronics in Agriculture **54**(2): 53-68.

Zheng, X., C. Zhang and P. C. Woodland (2021). Adapting GPT, GPT-2 and BERT Language Models for Speech Recognition. 2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU).

Zhong, G., X. Ling and L.-N. Wang (2019). "From shallow feature learning to deep learning: Benefits from the width and depth of deep architectures." Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery **9**(1): e1255.