Shifting focus from quantity to quality: Mapping, understanding, and managing within-field variability in crop quality in cotton and grain systems

MIKAELA JANE TILSE

BScAgr (Hons I)

Supervisor: Professor Thomas F. A. Bishop
Associate Supervisor: Dr Patrick Filippi
Associate Supervisor: Associate Professor Brett M. Whelan

A thesis submitted in fulfilment of
the requirements for the degree of
Doctor of Philosophy

School of Life and Environmental Sciences
Faculty of Science
The University of Sydney
Australia

12 May 2025

Abstract

Australian cotton and grain growers are world-renowned for producing high-yielding, high-quality fibre and grains. However, there is still considerable variation in both yield and quality within and between fields, farms, and seasons. Grain quality, namely the grain protein content (GPC), and cotton fibre quality, including length and micronaire (a composite measurement of fibre fineness [diameter] and maturity), are key determinants of the prices that growers receive due to the introduction of a premium and discount system for Australian growers. Thus, there is an onus on growers to manage for both quality and quantity to attain premium prices.

Site-specific crop management (SSCM) is the practical application of precision agriculture (PA) principles, and involves the allocation of resources and agronomic practicse to match spatiotemporal variability in the crop growing environment. However, uncertainty regarding the amount of within-field variation necessary to justify investment in PA technologies, and a lack of understanding regarding the drivers of this variation to support improved decision-making, is a considerable limitation to the adoption of PA for growers and advisors. Today, more data is being collected on farms and by the industry than ever before (e.g. yield data, variable-rate inputs), and there is also an enormous amount of public data that is free to access (e.g. remote sensing imagery) which can be used to describe or represent variability in GPC and cotton fibre quality. By understanding how and why cotton fibre and grain quality varies within-fields, growers and advisors can be equipped with the necessary information and tools to make better management decisions for more profitable and environmentally sustainable production systems.

This thesis explores the application of on-farm and publicly-available spatial data layers for the description, characterisation, and quantification of within-field variability in cotton yield and

fibre quality (length and micronaire) and GPC, and to understand the drivers of this variability within fields. Chapter 1 provides an overview and background of the Australian cotton and grains industries and the current role of PA in understanding and managing for variability in cotton fibre and grain quality. Chapter 2 presents a generalised geostatistical approach using area-to-point kriging to map and downscale areal observations of crop production data, which is illustrated using cotton yield and fibre quality (length and micronaire) data which is measured as a module (areal/block) average. Chapter 3 demonstrates how a combination of readily-available yield, agronomic, and publicly-available data layers can be used to create a model to predict GPC within-fields to fill-in gaps in the absence of a protein sensor. Chapter 4 investigates the relationship between wheat grain yield and GPC and applies interpretive machine learning approaches using existing spatial data layers to understand the drivers of spatial variability in GPC within-fields. In Chapter 5, the opportunities for SSCM for wheat grain yield and GPC are compared by quantifying the magnitude and spatial structure of within-field variability using the Opportunity Index (OI).

While the interpretation and application of the growing plethora of spatial data layers for decision-making is a challenge for growers and advisors, this research demonstrates the how a PA approach can use these data layers to better understand the nature and drivers of within-field variability in cotton fibre and grain quality to make better management decisions for more profitable and environmentally sustainable production systems that optimise both yield and quality.

Statement of Originality

This thesis has not been submitted for any degree or other purposes. I certify that the intellectual content of this thesis is the product of my own work and that all the assistance received in preparing this thesis and the sources have been acknowledged.

Mikaela J. Tilse | 12 May, 2025

Author Attribution

This is to certify that to the best of my knowledge, the intellectual content of this thesis is a product of my own work, except when otherwise acknowledged.

Chapter 2 is published in *Precision Agriculture*. Mikaela Tilse (MT) co-designed the study with Patrick Filippi (PF), Brett Whelan (BW), and Thomas Bishop (TB). Data was provided from growers via Precision Cropping Technologies (PCT) and Anthony Rudd from I-AG. The manuscript and data analysis was completed by MT but all co-authors provided feedback to manuscript drafts.

Chapters 3, 4, and 5 were co-designed by MT, PF, and TB. Data was provided from growers via PCT. The manuscript and data analysis was completed by MT but all co-authors provided intellectual contributions to the manuscripts.

Chapter 3 is an extended version of peer-reviewed conference papers presented at the *16th International Conference on Precision Agriculture* with the title "Predicting, mapping, and understanding the drivers of grain protein content variability – utilizing harvester mounted grain protein sensors", the *Society of Precision Agriculture Australia (SPAA) 2024 Australasian Precision Ag Symposium* with the title "Decoding the protein puzzle: understanding variability in grain protein content", and the *Grains Research and Development Corporation (GRDC) 2024 Northern GRDC Grains Research Update* with the title "Predicting and mapping grain protein content to better understand variability – utilising John Deere's new Harvestlab™ 3000 grain sensing system".

Chapter 4 is an extended version of a peer-reviewed conference paper for the *21st Australian Agronomy Conference* under the title "Understanding the drivers of within-field grain protein content variability with spatial data layers".

Appendix A is a book chapter published in *Soil Constraints to Crop Production* titled "Soil Constraint Diagnosis and Mapping". The manuscript was completed by MT, but PF, TB, and BW provided intellectual contributions and feedback to the manuscript.

Appendix B is a book chapter published in *Encyclopedia of the Soils in the Environment* titled "Proximal soil sensing in the field". The manuscript was completed by MT, and PF, and Uta Stockmann provided intellectual contributions and feedback to the manuscript.

In addition to the statements above, in cases where I am not the corresponding author of a published item, permission to include the published material has been granted by the corresponding author.

Mikaela J. Tilse | 12 May, 2025

As supervisor for the candidature upon which this thesis is based, I can confirm that the authorship attribution statements above are correct.

Thomas F.A. Bishop | 17 January, 2025

Acknowledgements

Completing this PhD is a culmination of over 20 years of education. It has been challenging, but also incredibly rewarding.

Firstly, I would like to acknowledge the financial support of the organisations who aided my research, travel, and learning opportunities throughout this PhD. Funding for this research was provided by the Christian Rowe Thornett Stipend Scholarship and the Thomas Lawrance Pawlett Scholarship (Supplementary) within the Sydney Institute of Agriculture at The University of Sydney. Research activities conducted as part of this thesis were supported by the Cotton Research and Development Corporation (CRDC) for a PhD research project "Assessing yield and fibre quality variability in cotton systems through data science for improved management" (Grant ID: US2104). The CRDC also provided a PhD travel grant to attend the 2022 Australian Cotton Conference (Grant ID: US2302). The industry has welcomed me with open arms and enthusiasm, and I am very grateful for the opportunities that have been provided. Support to attend the 16th International Conference on Precision Agriculture and visit to Cotton Incorporated in the USA during 2024 was provided by the Farrer Memorial Trust International Travelling Scholarship.

I would like to thank my supervisors, Patrick and Tom, for your guidance, patience and wisdom throughout. You've pushed me to hone my research skills and encouraged me to pursue ideas down rabbit holes, but also pulled me in the right direction when I was straying too far. Thanks must also go to Brett for your supervision during the early stages of my PhD. Thankyou for helping me to build solid foundations as a researcher, and for setting me up on the right path.

None of this research would have been possible without the growers, agronomists, and professionals who grew the crops and provided the data. In particular, a big thanks goes to

Precision Cropping Technologies, John Deere, and Anthony Rudd from i-Ag for connecting me with growers, facilitating access to the data, and answering my many questions. Also, a huge thanks goes to Kieran in Narrabri for your assistance when implementing field trials.

To the original 'clump' members – Jono, James, and Si Yang – thankyou for your friendship, guidance, banter, and beers. It's been a pleasure to work together. A huge thanks also goes to Sally for answering countless questions, field trips, cheese and crackers, and cups of tea. An honourable mention goes to Ryan for the many mid-work chats, post-work beers, and continued friendship. To housemates and friends, thankyou for your support over the years. And to all the other PhD candidates, post-docs, staff and students at ATP, thankyou for your guidance, assistance, and comradery.

An immeasurable about of thanks goes to James. Thankyou for listening at all hours, for the passionate discussions, and the reminders to take breaks. It has been a wild and wonderful time navigating our PhDs together, and I thank you for all of your support over the last few years. You have seen all of the ups and downs throughout this journey, and I thank you for being there through all of it.

Finally, thankyou to my family – Alison, Steven, Ellie, grandparents, and beyond – for instilling a passion for agriculture, and for your continued support and encouragement. I could not have done this without you.

Contents

List of Figures

List of Tables

List of Acronyms

10FCV: 10-fold cross-validation.

2FCV: Two-fold cross-validation.

A2PK: Area-to-point kriging.

ADR: Australian Premium Durum.

AFIS: Advanced Fibre Information System.

AFV: Average total field variation.

AGP: General Purpose Grade.

AH: Australian Hard.

AIC: Akaike Information Criterion.

AIW: Australian Innovative Wheat.

ANN: Artificial neural network.

ANW: Australian Noodle Wheat.

APH: Australian Prime Hard.

APW: Australian Premium White.

APWN: Australian Premium Noodle Wheat.

ASFT: Australian Soft Wheat.

ASW: Australian Standard White.

ASWN: Australian Standard White Noodle.

AWW: Australian White Wheat.

CBAM: Carbon Border Adjustment Mechanism.

CCCI: Canopy Chlorophyll Content Index.

CEC: Cation exchange capacity.

COSMOS: COsmic-ray Soil Moisture Observing System.

CRDC: Cotton Research and Development Corporation.

CRNS: Cosmic ray neutron sensing.

CSIRO: Commonwealth Scientific and Industrial Research Organisation.

CV: Coefficient of variation.

CV_a: Areal coefficient of variation.

DAP: Days after planting.

DD: Day degrees.

DEM: Digital elevation model.

DRS: Diffuse reflectance infrared spectroscopy.

DSM: Digital soil mapping.

EC: Electrical conductivity.

EC_a: Apparent electrical conductivity.

EM: Electromagnetic.

EMI: Electromagnetic induction.

EPO: External parameter orthogonalization.

ER: Electrical resistivity.

EU: European Union.

EVI: Enhanced Vegetation Index.

FDR: Frequency-domain reflectometry.

FineRes: Fine-resolution.

GEE: Google Earth Engine.

GHG: Greenhouse gas.

GIS: Geographic Information System.

GPC: Grain protein content.

GPR: Ground penetrating radar.

GPS: Global Positioning System.

GRDC: Grains Research and Development Corporation.

GRS: Electrical resistivity.

GRVI: Green Red Vegetation Index (GRVI).

HID: Harvest Identification.

HVI: High-Volume Instrument.

IML: Interpretive machine learning.

IR: Infra-red.

LCCC: Lin's concordance correlation coefficient.

LOFYOCV: Leave one Field-Year out cross-validation.

MAP: Monoammonium phosphate.

Max: Maximum.

MgmtClass: Management class.

Min: Minimum.

MIR: Mid infra-red.

MLR: Multiple linear regression.

ModRes: Module-resolution.

MSI: MultiSpectral Instrument.

N: Nitrogen.

NDRE: Normalised Difference Red-Edge.

NDVI: Normalised Difference Vegetation Index.

NFF: National Farmers Federation.

NIR: Near-infrared.

NMR: Nuclear magnetic resonance.

NNSW: northern New South Wales.

NSW: New South Wales.

OI: Opportunity Index.

OSAVI: Optimised Soil Adjusted Vegetation Index.

P&D: Premium and discount.

PA: Precision agriculture.

PAWC: Plant Available Water Capacity.

PCT: Precision Cropping Technologies.

PSS: Proximal soil sensing.

pXRF: Portable X-Ray Fluorescence.

$q_{50}(CV_a)$: Median magnitude of autocorrelated variation.

R^2: R-squared.

RE/R: Red-Edge Ratio.

RF: Random Forest.

RFID: Radio Frequency Identification.

RGB: Red Green Blue.

RMSE: Root mean square error.

ROI: Region of interest.

RSI: Ratio Spectral Index.

S: Practical range.

s: Operational machinery constant.

S_v: Spatial structure of variation.

SCANS: Soil Condition Analysis System.

SD: Standard deviation.

SHAP: SHapley Additive exPlanations.

SNSW: southern New South Wales.

SNV: Standard Normal Variate.

SOC: Soil Organic Carbon.

SPAA: Society of Precision Agriculture Australia.

SRTM: Shuttle Radar Topography Mission.

SSCM: Site specific crop management.

SWIR: Short-wave infrared.

TC: Total count.

TDR: Time-domain reflectometry.

TNM: Thermalized neutron methods.

UAV: Unmanned aerial vehicle.

USA: United States of America.

VI: Vegetation index.

VRA: Variable rate application.

WA: Western Australia.

WHC: Water holding capacity.

XGBoost: eXtreme Gradient Boost.

XRF: X-ray fluorescence.

CHAPTER 1

Challenges and opportunities in managing crop quality in Australian cotton and grain systems

---

1.1 Precision agriculture and the Australian cotton and grains industries

The National Farmers Federation (NFF), the peak national advocacy organisation representing farmers and agriculture across Australia, has set an ambitious target for the Australian agricultural industry; to exceed AUD$100 billion in farm gate output by 2030 (NFF 2019). The gross value of Australian agricultural production has increased by 51% in the past 20 years from ~AUS$62.2 billion (2003/04) to AUD$94.3 billion (2022/23) (ABS 2024). However, growers are faced with considerable economic, agronomic, and environmental challenges, including price fluctuations and increasingly unstable climate conditions, that are putting pressure on Australia's agricultural industry to meet this $100 billion target.

Following the signing of the Paris Agreement (UN 2015), the Australian Government is developing a plan to guide the transition of the Australian economy to net-zero greenhouse gas emissions by 2050 (DCCEEW 2024). Agriculture can play a pivotal role in this through emissions reduction and sequestration strategies, including soil carbon farming and increases in Nitrogen (N) and water-use efficiencies (Roth *et al*. 2013). Across the globe, farmers are acutely aware of the impacts of climate change, including prolonged heatwaves and more erratic rainfall events. Practice changes in farm management will be driven by risk-mitigation strategies to address the impacts of climate change on-farms (Hayman *et al*. 2019), and by external market forces such as the European Unions (EUs) Carbon Border Adjustment

Mechanism (CBAM; Deloitte Access Economics Pty Ltd 2023), which is a carbon border tax on products imported into the EU.

Australia is the driest inhabited continent in the world, and only ~5% of Australia's total land area is used for broadacre crop production (DAFF 2024). The vast majority of broadacre cropping is conducted under dryland (also known as rainfed) conditions, where there is no supplementary irrigation (Anderson *et al*. 2016). Irrigated cropping accounts for 0.2% of Australia's land area (DAFF 2024), but this fluctuates considerably season-to-season depending on water allocations. The volatility of returns experienced by Australian farmers are among the highest in the world, and this volatility is largely driven by seasonal conditions, namely water availability, alongside fluctuations in temperature, pests, and disease (Keogh 2012). Production costs are also among the highest in the world, largely driven by high labour, agrichemicals (including fertiliser and pesticides), and machinery costs (Herbert 2012). However, these challenges have also spurred innovation and investment in new technologies and techniques to enhance the value, efficiency, and sustainability of Australia's agricultural sector.

Wheat is the largest broadacre crop grown in Australia in terms of value, accounting for 60% of all winter crops sold (ABS 2024), and wheat exports from Australia account for 10 – 15% of the global wheat trade (AEGIC 2022). The vast majority of grains in Australia (including wheat) are grown under dryland conditions. Cotton is predominately grown for its fibre and is Australia's largest summer crop by area, with 766 thousand hectares planted in 2022/23 (ABS 2024) and has a gross value of AUS$3.1 billion (2023/24) (DAFF 2024). While Australia is a relatively minor producer on a global scale, over 99% of Australian-grown cotton is exported (DAFF 2023). Around 83% of Australia's cotton is irrigated, but this fluctuates season-to-season depending on water availability (Roth *et al*. 2013). At least 80% of Australia's cotton growing area is irrigated using gravity surface-irrigation systems, namely furrow irrigation or siphonless/bankless channels, but there is an increasing use of centre-pivot and lateral-move irrigation machine systems, and only a small proportion is irrigated using subsurface drip-irrigation systems (Roth *et al*. 2013).

In Australia, wheat is grown across three broad regions; the north (Queensland and New South Wales, NSW), the south (Victoria, Tasmania, South Australia), and the west (Western Australia, WA; GRDC 2024; Figure 1.1). Cotton is grown in Queensland, NSW, northern Victoria, and parts of the Northern Territory and WA (Figure 1.2). If seasonal conditions are favourable, cotton and wheat are often grown as part of a summer-winter cropping rotation in the northern grains region where they are used to diversify cropping rotations to help manage disease, weeds, herbicide options, and increase farming returns (Baird *et al*. 2022). While both cotton and wheat are grown in broadacre production systems, agronomic management differs between each crop, and they face unique challenges in terms of market specifications, spatiotemporal variability, and management.

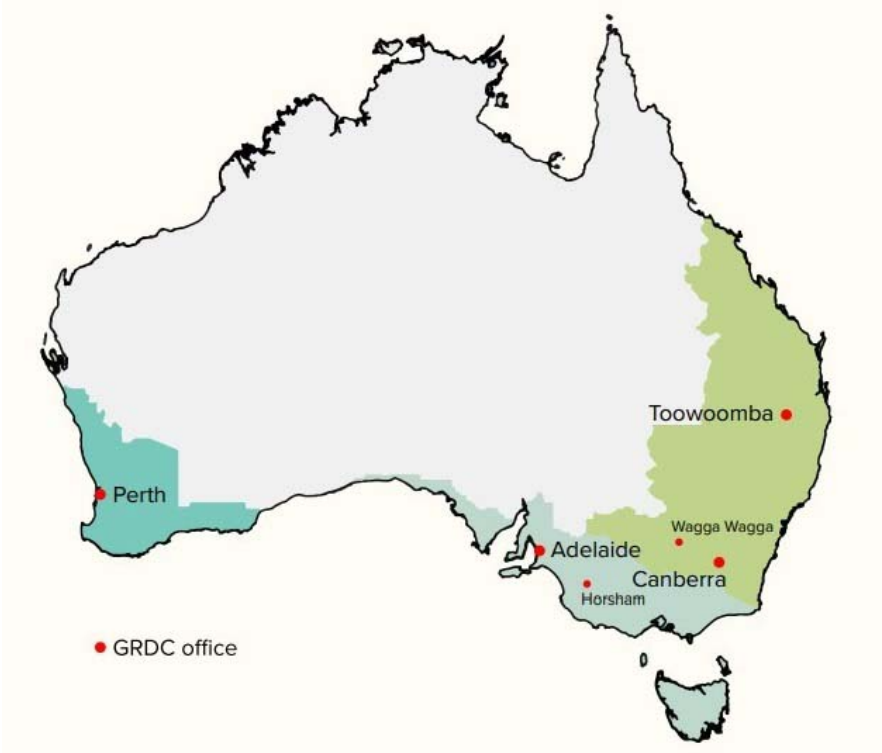

FIGURE 1.1. Grain growing regions of Australia (GRDC, 2024).

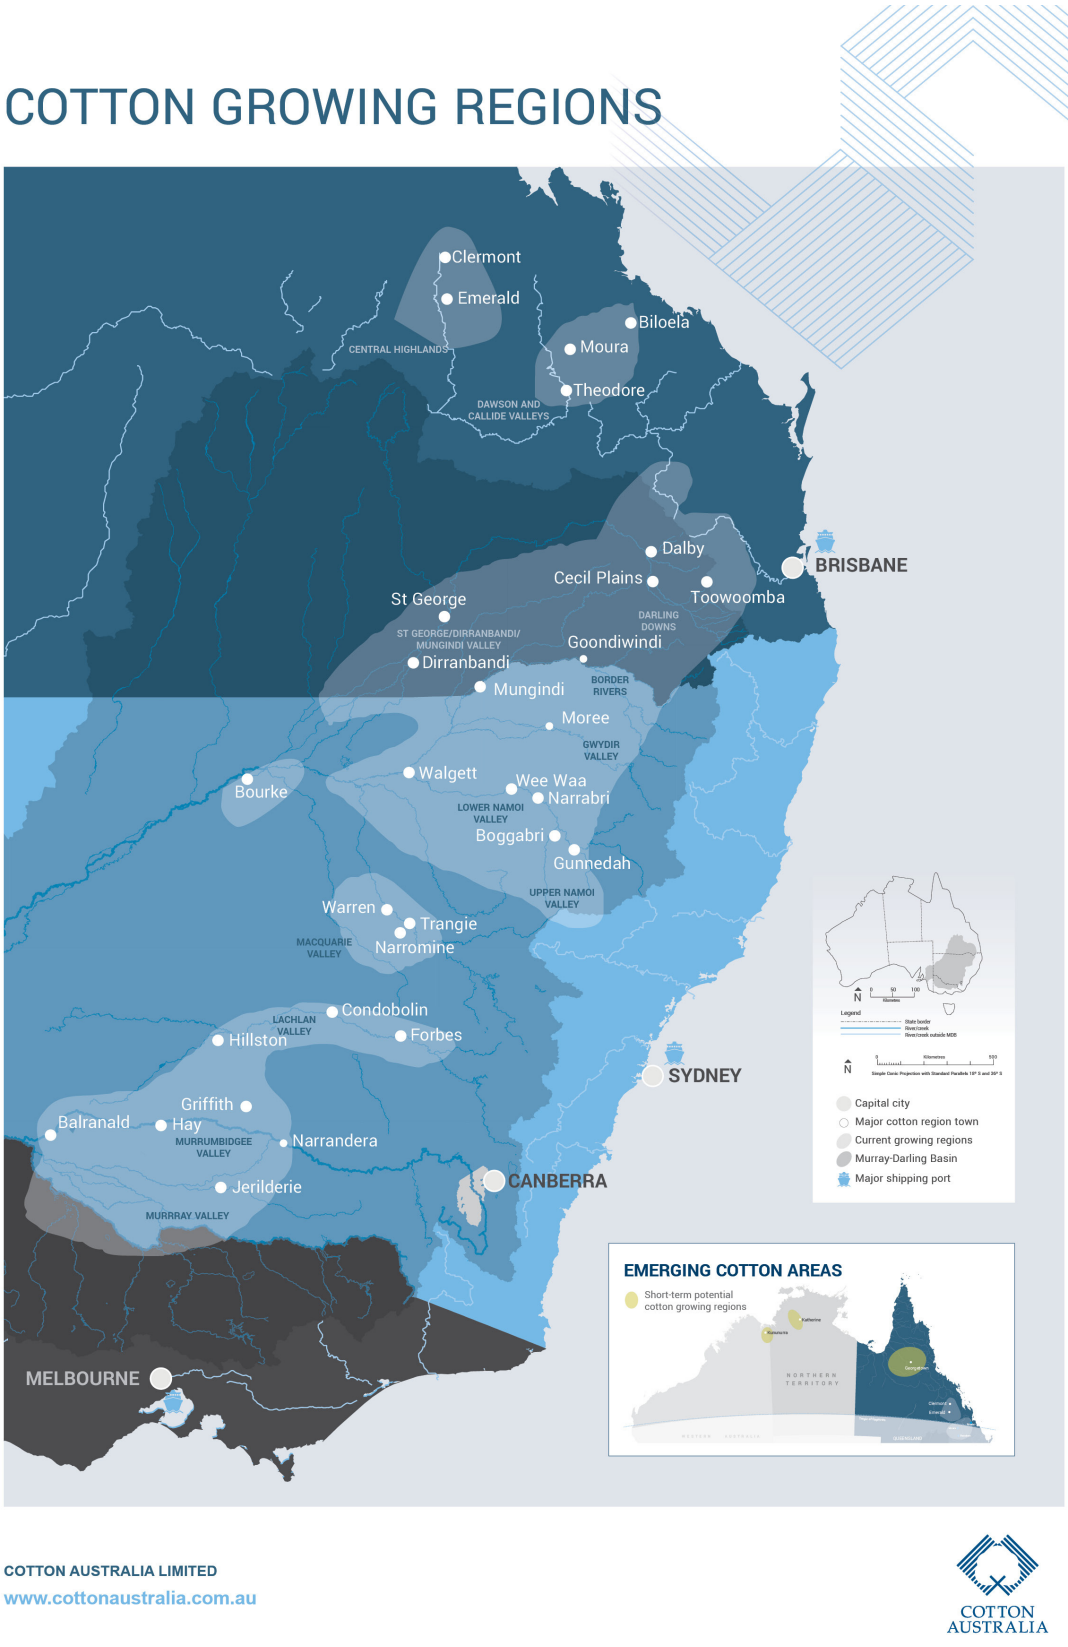

FIGURE 1.2. Australian cotton growing region (Cotton Australia 2024).

Australian growers have, and continue to be, at the forefront of innovation and the adoption of new technologies and management practices. This is largely attributed to economic necessity (Pannell *et al*. 2006), poor soils (Davies *et al*. 2022), and low rainfall conditions (Anderson *et al*. 2016; Keogh 2012). High labour and other input costs (e.g. fertiliser) have also driven increased mechanisation and led to a trend of increasing farm size. The adoption of precision agricultural practices, including the variable rate application of fertilisers, has been a large contributor to reduced costs and improvements in yields and quality (Anderson *et al*. 2016).

Precision agriculture (PA) is a discipline and management approach that involves the observation, assessment, and response to variability in agricultural systems. In broadacre cropping, PA can involve site-specific crop management (SSCM), in which differential treatments are applied to better match soil and crop requirements as they vary within a field. By managing fields on a within-field, site-specific basis, PA has the potential to increase input-use efficiency, reduce environmental impacts, and increase the return on investment to growers. The Australian grains industry was one of the earliest adopters of PA and SSCM, and adoption in other farming industries has expanded since the late 1990s (Whelan *et al*. 2013), including into the Australian cotton industry.

Many of the tools and technologies used to implement PA in the cotton and grains industries are similar, such as electronic controllers for variable-rate application (VRA) of nutrients and herbicides. However, unique challenges differentiate the adoption and practice of PA between the cotton and grain industries. Broadly speaking, the adoption of PA in a dryland wheat system is more straightforward compared to in gravity surface-irrigation (i.e. furrow or bankless channel) systems for cotton. This is due to the scale and method of application of inputs (e.g. flood irrigation is applied to an entire cotton field with limited opportunities for variation in application), and logistical constraints including vehicular access to fields at critical times following an irrigation event, for example. Even so, PA can play a critical role in both cotton and grain production to understand and respond to site-specific variability.

Most PA research and adoption in the industry has focused on improving and optimising crop yields because yield remains a primary factor in determining profitability. There has been considerable gains in the water-use efficiency of Australian cotton (Roth *et al*. 2013) and

some progress in closing the yield gap for Australian wheat (Monjardino *et al*. 2019), which is the difference between potential and achieved water-limited yield potential. However, actual cotton and grain yields have stagnated in the last 20 years (Muleke *et al*. 2022; Hochman *et al*. 2017). Increasing volatility and changing export market demands are pushing growers to do more with less and maximise the return on investment (Muleke *et al*. 2022). Enhancing and optimising cotton fibre and grain quality presents an opportunity for Australian growers to remain competitive. This introduction chapter will focus on cotton fibre and grain quality in primarily an Australian context, but will also consider international research, and will provide an overview of how quality is measured, the drivers of variability, and the role of PA in understanding and managing for cotton fibre and grain quality.

1.2 Shifting focus from quantity to quality

Producing high-yielding, high-quality grains and cotton fibre is a key factor for growers to attain premium prices (Stewart *et al*. 2002; Long *et al*. 2013b). Australian wheat has a reputation for high quality among export markets, where it is preferred for making noodles, steamed bread, and flat breads (Cato *et al*. 2020). Likewise, Australian cotton is highly sought-after by international merchants, spinners, and manufacturers due to its high-quality and white colour (McVeigh 2017). In Australia, cotton and grain growers are paid a premium (or receive a discount) depending on whether their grain or fibre meets a designated quality threshold or standard (Sections 1.2.1 and 1.2.2).

The premium & discount (P&D) system used for wheat and cotton has placed an onus on growers to manage for both quality and quantity. For wheat, premiums can vary anywhere from 3 to 20% above the base price (Anderson *et al*. 2016). As an example, the percentage of the WA wheat crop that received a price premium has risen substantially since the 1990s (Anderson *et al*. 2016; Anderson *et al*. 2005; Figure 1.3). This has largely been driven by optimisation of N fertiliser regimes and other improvements in agronomic research (Anderson *et al*. 2016). It is possible to blend or segregate wheat grain from different areas of a field based on the grain protein content (GPC) recorded during harvest to increase uniformity or

achieve higher returns. On the other hand, the blending of cotton during ginning is impractical (Sluijs *et al*. 2019) and for some fibre properties like colour, will not improve the quality grade. For cotton growers, there is a greater onus to manage quality to minimise discounts rather than achieve premiums, as only small premiums are awarded (~$12/bale), while discounts are large (up to $100/bale for a substantial downgrade on colour, McVeigh 2017). Further, as around 72% of Australian agricultural production is exported (as of 2019/20, ABS 2024), there is pressure on growers to maintain high-quality cotton and grain standards to uphold a reputation with international buyers and markets and remain competitive.

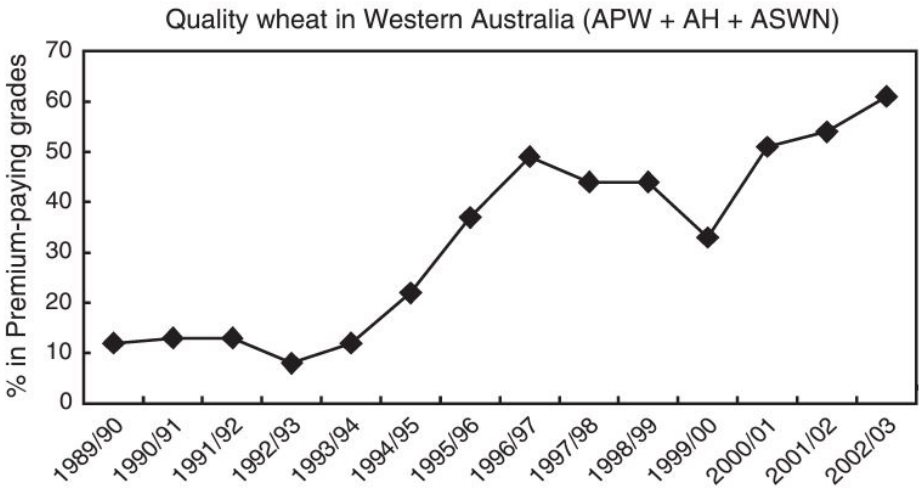

FIGURE 1.3. Percentage of Western Australian wheat crop that was delivered into premium-paying grades: Australian Premium White (APW), Australian Hard (AH) and Australian Standard White Noodle (ASWN) grades. From Anderson *et al*. (2005).

1.2.1 Grain protein content

Grain quality is a broad term that refers to the characteristics of a grain that determine its suitability for a specific end-use, encompassing both physical (e.g. moisture content, kernel size) and compositional (e.g. protein or oil content) characteristics. For each of the many end uses of cereal grains, there are specific quality requirements (Nuttall *et al*. 2017).

For wheat, the wheat GPC, which is a compositional trait representing the percentage of protein within a grain, is a key determinant of grain quality. Wheat grain quality is also

determined by other characteristics including moisture content, screenings (the % of matter passing through a 2 mm slotted screen), or insect damage, but the GPC is of high importance to millers and processers (Nuttall *et al*. 2017; GTA 2024).

In Australia, wheat is classified into 12 classes to meet distinct quality attributes for processors and end-uses (1.1; GTA 2024). Each variety is categorised into a class based on processing and end-product quality, and the class determines the highest grade that a variety can be accepted into at delivery. Within each class, there are multiple bin grades, known as the bin grade cascade, that a delivery can be classified into depending on the quality standards at receival. Each year, grain quality standards are provided by Grain Trade Australia (GTA 2024). Higher quality, premium wheat grains such as Australian Prime Hard APH wheat typically receive premiums, whereas lower grade feed wheats may face discounts. Premium and discount systems also apply for other grain crops in Australia including canola (discounts applied when oil < 42%), and barley (premium malt grades = 9 – 12% protein).

1.2.2 Cotton fibre quality

The quality of cotton can be expressed by a number of different measurements which are conducted post-harvest during classing. Depending on the trait, fibre quality is measured visually (e.g. colour and leaf), using a High-Volume Instrument (HVI) (e.g. micronaire, length, strength), or with an Advanced Fibre Information System (AFIS) (e.g. neps). In Australia, cotton fibre is classed using six key characteristics; colour, leaf, staple length, micronaire, strength, and extraneous matter (Table 1.2), which are based on the standards set by the United States Department of Agriculture. Pricing adjustments may also be made for other undesirable quality characteristics, including neps (fibre entanglements), contamination (e.g. plastic wrap), and stickiness (contamination of cotton from silverleaf whitefly, cotton aphid, or solenopsis mealybug, Table 1.2). Other fibre quality characteristics that are measured by HVI instruments and are important for spinners but do not have a direct impact on price include length uniformity index, elongation, short fibre index, and maturity. The price received for cotton is dependent on the quality of each bale (227 kg) of cotton lint (CRDC *et al*. 2023). Cotton prices are quoted for “base grade” 31-3-36, G5 (Figure 1.4), and premiums or discounts

TABLE 1.1. Australian wheat classes, including protein thresholds (minimum, min; maximum, max), attributes and guidelines, and end-uses (GTA 2024; Grains Australia 2024; GRDC 2016).

Class	Protein threshold (%)	Attributes and guidelines	End-uses
Australian Prime Hard (APH)	Min 13%	Hard white wheat, high protein, strong balanced dough Hard white wheat, high protein, exceptional milling quality	Yellow alkaline noodles; Japanese ramen noodles; high volume breads
Australian Hard (AH)	Min 11.5%	Hard white wheat, high protein, excellent dough quality	Pan, hearth and flat breads; steamed products; yellow alkaline noodles
Australian Premium White (APW)	Min 10%	Hard white wheat, high protein, strong balanced dough properties	Multipurpose; flat and steamed breads; instant and fresh noodles
Australian Standard White (ASW)	NA	Hard white wheat, low protein, minimum dough and end product testing requirements	Range of breads including European breads and rolls, flat and steamed breads
Australian White Wheat (AWW)	NA	Versatile, medium-to-low protein white wheat	General purpose requirements and blending by flour millers
General Purpose Grade (AGP)	No min	Wheat that has failed to meet minimum receival standards for milling wheat grades, but higher minimum standards at receival compared to Feed class.	Animal feed
Australian Premium Durum (ADR)	Min 13%	Very hard vitreous amber-coloured kernels, excellent colour	Wet and dry pasta; couscous; hearth and flat breads
Australian Soft Wheat (ASFT)	Max 9.5%	White soft wheat, weak and extensible dough properties, low water absorption and starch damage	Biscuits; cakes; pastry; steamed breads
Australian Innovative Wheat (AIW)	NA	Specific to variety for specified market function, unique characteristics according to variety. No grade specifications within class.	Distinctive use validated by market opportunity
Australian Noodle Wheat (ANW)	9.5 – 11.5%	Soft kernel hardness, white grain, exceptional noodle sheet colour stability, high flour paste attributes	Specific end product use for Udon-style noodles
Australian Premium Noodle Wheat (APWN)	10 – 11.5%	Hard wheat, good noodle colour and brightness stability	White salted and instant noodles
Feed	No min	Wheat that has failed to meet minimum receival standards for milling wheat grades.	Animal feed

apply for higher and lower grades, respectively (CRDC *et al*. 2023). While Australian cotton receives only small premiums for its high quality in the present market, there is an onus on growers to manage for quality to minimise discounts and maintain Australia's reputation as a producer of high-quality cotton (McVeigh 2017).

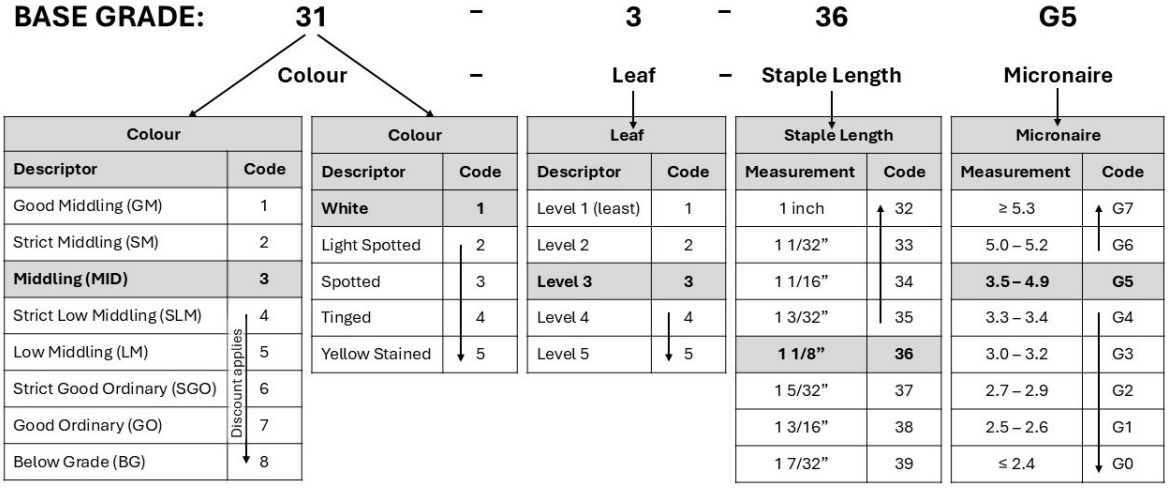

Colour	
Descriptor	**Code**
Good Middling (GM)	1
Strict Middling (SM)	2
Middling (MID)	**3**
Strict Low Middling (SLM)	4
Low Middling (LM)	5
Strict Good Ordinary (SGO)	6
Good Ordinary (GO)	7
Below Grade (BG)	8

Colour	
Descriptor	**Code**
White	**1**
Light Spotted	2
Spotted	3
Tinged	4
Yellow Stained	5

Leaf	
Descriptor	**Code**
Level 1 (least)	1
Level 2	2
Level 3	**3**
Level 4	4
Level 5	5

Staple Length	
Measurement	**Code**
1 inch	32
1 1/32”	33
1 1/16”	34
1 3/32”	35
1 1/8”	**36**
1 5/32”	37
1 3/16”	38
1 7/32”	39

Micronaire	
Measurement	**Code**
≥ 5.3	G7
5.0 – 5.2	G6
3.5 – 4.9	**G5**
3.3 – 3.4	G4
3.0 – 3.2	G3
2.7 – 2.9	G2
2.5 – 2.6	G1
≤ 2.4	G0

FIGURE 1.4. Interpretation of Australian base grade for upland cotton; 31 – 3 – 36, G5. Adapted from CRDC *et al*. (2023).

TABLE 1.2. Cotton fibre quality traits, adapted from CRDC *et al*. (2023).

Quality trait	Description	Ideal range
Colour	The degree of yellowness and reflectance or whiteness.	Ideal colour grade >31 (Middling White). Small premiums for good grades, and significant discounts for poor grades.
Leaf	A measure of the amount of leaf (trash) material remaining in the cotton lint sample after ginning.	$\leq$ leaf grade 3 is ideal. High levels of trash (e.g. grass or bark) incur large price discounts.
Staple Length	Length varies with variety, and is affected by stress during fibre development and stress during mechanical processes at and after harvest. Measured and reported in 100^{ths} and 32nds of an inch.	Upper half mean length >36/32nds (1.125 inch) is ideal. Premium range >40/32nds (1.250 inch). Significant price discounts <33/32nds.
Micronaire	A combination of fibre linear density and fibre maturity. Test measures the resistance offered by a weighted plug of fibres in a chamber of fixed volume to a metered airflow. Value is unitless.	3.5 – 4.8 is desirable. Premium range 3.8 – 4.2. Significant price discounts apply <3.5 and >5.0
Strength	Fibre strength is reported in gram per tex (g/tex), where one tex is equal to the weight in grams of 1000 m of fibre. Strength represents the force required to break a bundle of fibres one tex unit in size.	>29 g/tex is ideal. Small premiums for values >29 g/tex. Premium fibre >34 g/tex.

Continued on next page

Table 1.2 continued from previous page

Quality trait	Description	Ideal range
Uniformity	The ratio between the mean length of fibres, and the upper half mean length of fibres. The lower the variation in the length distribution, the higher the length uniformity index.	>80% is ideal. Small price discounts <78%, no premiums apply.
Neps	Small entanglements or knots of fibres.	<250 neps per gram (neps/g) is ideal. For premium fibre, <200 neps/g. Moderate price discounts.
Short fibre content	The amount of fibres (in percent by weight) less than 12.7 mm (0.5 inches). No premiums or discounts apply, but the metric is used by mills to determine laydowns (bale selection).	<8% is ideal
Moisture	Moisture content of seed cotton.	6 – 10% is ideal. >12% considered excessive
Contamination	Foreign matter such as plastic, rocks, sticks present in the sample.	Low/none is ideal. A reputation for contamination has a negative impact on price, sales, and future exports.
Stickiness	Excessive sugars due to insect secretions (e.g. silverleaf whitefly, cotton aphid, or Solenopsis mealybug), mould, or plant sugars.	Low/none is ideal. High levels of stickiness incur significant price discounts and can lead to rejection by the buyer.

1.3 The role of Precision Agriculture

Precision agriculture is both a management approach and a research discipline that aims to increase the long-term site-specific and whole-farm production efficiency and profitability by addressing and responding to variation in crop production (Whelan *et al*. 2013). Through the collection and utilisation of spatial data layers from a range of different sources – from farms, cotton gins, grain silos, industry, and freely available data – PA can play a vital role in mapping, understanding, and managing for variation in cotton fibre and grain quality.

1.3.1 Mapping cotton fibre and grain yield

While Australian cotton and grain growers are renowned for producing high-quality cotton and grains (Cato *et al*. 2020; McVeigh 2017), variability within and between fields, farms, and seasons remains a concern. Farmers and advisors often anecdotally know or understand spatial variability within their fields, where this "informal" knowledge has been collected through observation of the agronomic context over multiple seasons (Appendix A, Tilse *et al*. 2022). Grain yield monitors were first developed in the early 1990s, and the first cotton yield monitors became available in 1997 (Roades *et al*. 2000). The development of cotton yield monitors lagged behind grain yield monitors because of challenges with the accuracy, calibration, and maintenance of optical sensors that tended to become blocked by dust (Durrence *et al*. 2005; Roades *et al*. 2000; Watcharaanantapong *et al*. 2014), but more accurate and reliable cotton yield monitors were introduced in 2000 (Larson *et al*. 2005). Today, 87% of Australian grain growers have access to yield monitors (Bramley *et al*. 2019), and more than 90% of Australia's cotton crop is harvested using round-module harvesters which are fitted with yield monitors (Miao *et al*. 2020). This yield data has helped to provide an objective quantification of variability within fields. Yield maps derived from on-the-go yield monitors have been used alongside a combination of other on-farm (e.g. electromagnetic (EM) surveys) and publicly-available (e.g. remotely-sensed satellite imagery) spatial data layers in research and commercial applications. This data has been used to predict (Habibi *et al*. 2024; Nguyen *et al*. 2019), forecast (Filippi *et al*. 2019b; Filippi *et al*. 2020b), and

understand the drivers of variability in wheat grain and cotton fibre yields (Jones *et al*. 2022; Poole *et al*. 2024) at a range of different spatial and temporal scales.

1.3.2 Mapping cotton fibre and grain quality

Unlike for cotton fibre and grain yields, measuring and mapping variability in wheat GPC and cotton fibre quality is more challenging. This is due to complex interactions between genetic, environmental, and management factors that drive variability which may be difficult to measure, or there is a lack of commercially-available sensing technologies in the case of mapping cotton fibre quality. Maps describing variability in cotton fibre quality and grain GPC are important PA tool.

1.3.2.1 Grain protein content

Wheat GPC can be mapped on-the-go at harvest using combine harvester-mounted grain protein sensors. In the late 1990s and early 2000s, combine harvester-mounted protein sensors based on near-infrared (NIR) spectroscopy were being developed (e.g. the ProSpectra™ Grain Analyser, not commercialised; Rosenberg *et al*. 2000) and made available commercially (e.g. Zeltex AccuHarvest® Analyser, https://zeltex.com/). In 2013, the CropScan 3000H On Combine Analyser (formerly Next Instruments, now trading as CropScanAg, https://cropscanag.com/) was commercially released in Australia for the real-time, on-the-go measurement of protein, oil, and moisture in grains and oil seeds during harvest, and subsequent models have since followed (e.g. CropScan 3300H, 4000VT). More recently in 2023, John Deere commercially released the HarvestLab 3000™ Grain Sensing System in Australia, which is an NIR spectroscopy-based sensor mounted onboard a combine harvester for the measurement of continuous grain flow. Onboard the harvester, the auger extracts a threshed grain sample, passes it along and presents it to the NIR constituent sensor where the sample is then illuminated by a broadband light source. Specific wavelengths of this light are absorbed by the grain sample, depending on the molecular composition of the grain, and a spectrometer inside the sensor measures the wavelengths reflected by the sample which is then used to estimate properties such as protein, oil, or moisture content. This prediction

is performed using an inbuilt calibration function which was trained using the spectra from a large number of samples across a range of conditions, and the measurement accuracy of moisture and protein estimation with the sensor is ~$\pm$0.6% (90% confidence interval). Samples are taken continuously during harvest, with one measurement performed per second (Schade *et al*. 2023).

While harvester-mounted, on-the-go sensors are available to map wheat GPC, not all growers have access to these sensors and maps of GPC are not available for every field, farm, or season. The Australian grains industry is a world leader in adopting technologies, and a huge amount of data is now being collected on-farms and by the industry including yield maps, EM induction surveys, and as-applied fertiliser maps, for example. There is also an enormous amount of public data now available (e.g. Sentinel/Landsat satellites, Soil and Landscape Grid of Australia) that is free to access across every field and farm across Australia. The quality and quantity of this information being collected is constantly increasing, and we are now in the age of "Big Data" in agriculture (Himesh *et al*. 2018) which has led to an increased focus on using empirical models to model crop yields. This often involves building a spatiotemporal (i.e. space-time) dataframe/datacube, which collates data of varying spatial and temporal resolutions from different on-farm and off-farm sources into a single georeferenced database. Machine learning approaches are increasingly being used in crop yield modelling to capture the complex and nonlinear nature of the interactions between variables in these large space-time cubes (e.g. Fajardo *et al*. 2021), and the application of machine learning in agriculture is discussed by Liakos *et al*. (2018). There is an opportunity to use similar machine learning approaches with a large datacube of on-farm and/or publicly-available information to map wheat GPC within fields to fill-in gaps in the absence of grain protein sensors.

1.3.2.2 Cotton fibre quality

While cotton yields can be mapped easily with harvester-mounted yield monitors, there is no commercially available sensor to map cotton fibre quality on-the-go at harvest. Instead, cotton fibre quality is measured post-harvest using fibre samples collected from a ginned (227 kg) bale, and is typically presented as an average across many bales or for an entire field that is

not directly linked to in-field locations. Fine-resolution maps of cotton fibre quality have been produced for research purposes using point-based samples collected by-hand within fields, which were then interpolated to the field-scale (Gemtos *et al*. 2005), but this is not realistic in commercial settings. Prototype harvester-mounted fibre quality sensors has been development to measure micronaire (Schielack *et al*. 2016) and trash (Long *et al*. 2020) on-the-go, but these have not been commercialised yet as the presence of seed, foreign matter, and varying moisture content on the harvester means more work is needed to implement these sensors onboard harvesters for on-the-go measurement (Hardin *et al*. 2022).

Post-harvest module-resolution maps of cotton fibre quality can be produced by tracking cotton fibre quality data for each bale from the gin, to the module, and back to the in-field locations through the use of Radio Frequency Identification (RFID) technology (Fuhrer *et al*. 2024, Figure 1.5). More than 90% of Australia's cotton crop is harvested using round-module harvesters (Miao *et al*. 2020), which use RFID tags to label and track data from harvest through to the gin and fibre quality classing via John Deere's HID system. As the harvester moves throughout the field, the seed cotton is harvested and formed into a round-module that weighs around 2.5 tonnes (t). An RFID tag is incorporated into the cylindrical module wrap which covers each round-module, and includes time-stamped and georeferenced data unique to each module describing the module completion location and time, moisture content, and area harvested, among other information. During the ginning process, the cotton seeds and lint are separated, and the lint is pressed into rectangular bales, each weighing 227 kg. Each round-module holds ~4 cotton bales and individual bale numbers are assigned and associated with the module from which each bale came from. At classing, cotton fibre quality data for each bale can then be traced back to the module. This information has been used for logistics to track module progression and location during storage and transport to the gin, but can also be used to produce module-resolution fibre quality maps to visualise the distribution of fibre quality attributes across a field (Fuhrer *et al*. 2024). Module-resolution maps of cotton fibre quality have recently been made available from a commercial PA company "Precision Cropping Technologies" (PCT, https://pct.ag/), but the extent to which growers have adopted this technology remains unclear at this stage.

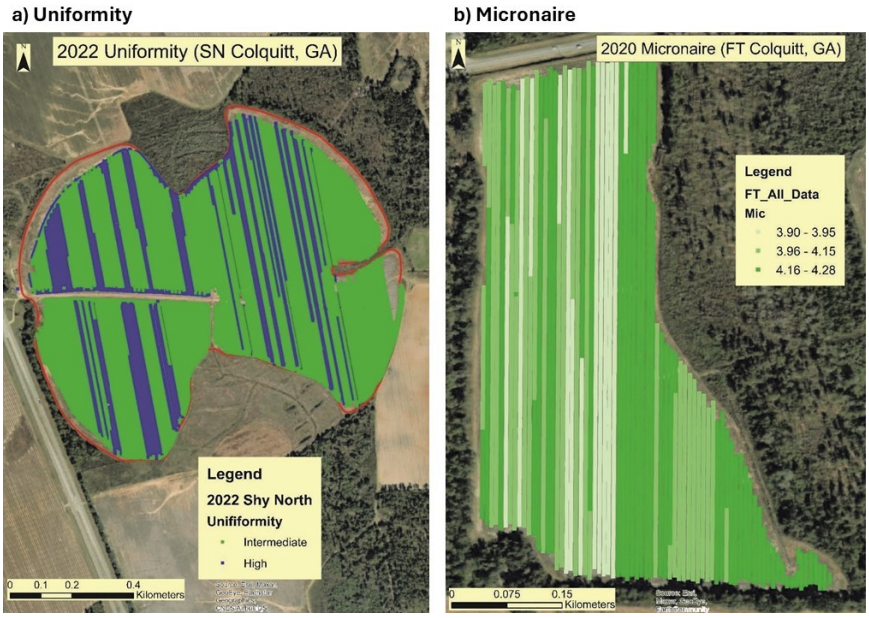

FIGURE 1.5. Cotton fibre quality maps for a) uniformity and b) micronaire from Fuhrer *et al*. (2024).

1.3.3 Understanding variability in cotton fibre and grain quality

1.3.3.1 Grain protein content

Variability in wheat grain yield and protein content is driven by complex interactions between genetic, environmental, and management factors, including variety, rainfall, temperature, nutrition, and soil and landscape features. Within-fields, spatial variability in yield and GPC is largely driven by N availability in the soil and applied as fertiliser, and soil moisture availability throughout the growing season (Whelan *et al*. 2013). The presence of soil physical and chemical constraints can reduce crop growth by restricting root growth and/or water or nutrient uptake (Dang *et al*. 2010). Generally, an inverse relationship between wheat grain yield and GPC is expected due to the grain protein "dilution effect" where, as N increases, yield and GPC increase up until a certain point, after which yield begins to plateau while GPC continues to increase. At very high N levels, a yield penalty may be incurred. When soil moisture is limited, GPC will increase while yields decrease (Holford *et al*. 1992; Simmonds 1995).

Research using a range of different statistical and geostatistical approaches has revealed considerable variability in wheat grain yield and GPC within fields and between seasons, and that wheat grain yield and GPC do not always vary in the same way, to the same magnitude, or with the same spatial structure within fields (Whelan *et al*. 2009). However, the slow, but increasing uptake of grain protein sensors means that much of this work has been limited to a small number of fields or seasons. In the absence of on-the-go grain protein sensors, satellite data has been used to predict GPC within fields (Hernandez *et al*. 2023; Stoy *et al*. 2022; Tan *et al*. 2020), but this work still requires ground-truthing. This has often been limited to either hand-samples in the absence of protein sensors (e.g. Gozdowski *et al*. 2017; Hernandez *et al*. 2023; Tan *et al*. 2020; Zhou *et al*. 2021), or only a small number of fields or seasons (e.g. Wang *et al*. 2019; Stoy *et al*. 2022).

On their own, maps of GPC can describe spatial variation, but additional interpretation and analysis is needed to understand the drivers of this variability, and to translate this knowledge into tangible management decisions. Improvements in data processing capabilities means that there is potential to use data layers from a wide variety of sources – from farms, grain silos, industry, and freely available data – in a large datacube to understand variability across the whole farming system. Machine learning approaches have been used to understand the drivers of variability in cotton and soybean yields at the within-field and field level (Jones *et al*. 2022; Nehbandani *et al*. 2023). As the adoption of grain protein sensors increases, there is an opportunity to use similar machine learning approaches to understand variability in wheat GPC within fields. Rather than using data layers in isolation to understand variation in a single field at a time, there is also an opportunity to use the growing plethora of on-farm and publicly-available data to identify, characterise, and understand variability in GPC across an entire farm, collection of farms, or region.

1.3.3.2 Cotton fibre quality

Over the past 15 – 20 years, Australian cotton fibre quality has greatly improved. Australian cotton fibre is now regarded as one of the top-quality fibres in upland cotton (*Gossypium hirsutum*) (McVeigh 2017), which is a medium-to-short stapled cotton and is the most common

type of cotton grown in Australia and the United States of America (USA). However, there is still the issue of significant variation in fibre quality, namely micronaire (a composite measurement of fibre fineness [diameter] and maturity), length, and strength, which exists both within and between fields. Some fibre quality concerns previously outlined by mills and spinners, including substandard neps (fibre entanglements which do not absorb dyes well), micronaire (fibre that is too thick or thin affects the spinning process), and short fibres (which negatively affect evenness and yarn strength, Gordon *et al*. 2004), are determined by stress during fibre development (e.g. moisture) or the presence of immature bolls at harvest (e.g. premature defoliation, CRDC *et al*. 2023). As cotton fibres are largely comprised of cellulose, factors that negatively impact photosynthesis, such as low temperatures or water stress, may limit fibre thickening and fibre maturation. The prevalence of immature fibres and nep formation would be particularly high if the cotton plant experienced stress during critical periods, especially towards the end of the growing season (Bel-Berger *et al*. 1998). Likewise, water stress from early-to-mid boll filling may significantly increase micronaire and detrimentally affect length, which can prevent growers from meeting quality standards (Ballester *et al*. 2019; Bange *et al*. 2018). As an indeterminant plant, cotton has some capacity to compensate for early stress events (e.g. water stress or high insect pressure), but the late-maturation of new bolls may have increased susceptibility to cool temperatures which further delays maturation, resulting in immature bolls at harvest and quality discounts (CRDC *et al*. 2023). In most cases, management decisions that promote higher yields will also contribute to higher quality cotton (CRDC *et al*. 2023).

While there has been considerable research conducted to understand cotton fibre quality on individual plants or at discrete locations (e.g. Wang *et al*. 2014), and assessment of the spatial variability of cotton yield across a field (Ge *et al*. 2008; Ping *et al*. 2007), mapping and assessing spatial variability in cotton fibre quality attributes within a field has been limited. Fine-resolution maps of cotton fibre quality have been produced for research purposes using hand-samples collected within a field (Gemtos *et al*. 2005; Johnson *et al*. 2002), and there has been some research conducted using unmanned areal vehicles (UAVs) to predict cotton fibre quality parameters within fields (Xu *et al*. 2021). However, ground-truthing is limited to hand-sampling which restricts the scale of this research. Overall, there has been limited work

on PA for cotton fibre quality in an Australian or international context, and much of the work that has been conducted has not shifted past the research phase.

1.3.4 Using data for decision-making

While the commercial uptake of on-the-go grain protein sensors by growers has been slow, their adoption has been steadily increasing as more and more growers are becoming interested in ways to utilise this data for improved decision-making (Bastos *et al*. 2021). Knowing the magnitude and extent of spatial variability in GPC can be used by growers to make management and logistical decisions, including to segregate grain harvest for premium markets (Long *et al*. 2013a; Stewart *et al*. 2002), or as a decision-making tool for nutrient application (Scott 2022; Long *et al*. 2011). For example, maps of yield and GPC can be used in combination to calculate total N removal by the crop, which can then be used to inform an N replacement strategy for the following season which has been shown to increase grain quality, grain yield, and profitability (e.g. Bonfil *et al*. 2008). In cotton, if areas consistently producing high or low yields or fibre quality can be identified and segregated into management zones, opportunities for selective harvest (Boydell *et al*. 2002) or module staging (Sluijs *et al*. 2019) can be explored to minimise discounts. If yield or quality data can be collected over multiple seasons, long-term trends can be identified and more temporally-stable drivers of variation can be differentiated from seasonal fluctuations.

Uncertainty regarding the amount of within-field variation necessary to justify investment in PA technologies has long been a concern for growers and advisors (McBratney *et al*. 1997). A large magnitude of variation should allow for a greater differentiation of input applications, which will likely result in greater economic and environmental benefits in comparison with uniform management (Pringle *et al*. 2003). However, the spatial structure – whether variability forms coherent patterns or appears as random noise – is also a key consideration. The presence of a trend in a yield or protein map may indicate an increased opportunity for SSCM because farm machinery such as variable-rate fertiliser sprayers are best directed when proposed application patterns are smooth and broad, and can operate within the spatial patterns and structures of variation. Therefore, a field with a large magnitude of variation with a strong

spatial structure will be more amenable to SSCM than a field with large variation but poor spatial structure (Pringle *et al*. 2003).

1.3.5 Key challenges and research gaps

While there is considerable interest in utilising PA to better understand cotton fibre and grain quality within fields, and to improve management outcomes, several challenges and research gaps remain. One of the primary challenges is the absence of a commercially-available on-the-go sensor for mapping cotton fibre quality. Before further research can be conducted into the use of PA for better understanding and managing variability in cotton fibre quality, a map of fibre quality at a fine-resolution is needed. Similarly, while on-the-go sensors for measuring wheat GPC are becoming more widely adopted, they are not universally deployed. This means that maps of GPC are not available for every field, farm, or season. The increasing availability and diversity of on-farm and publicly-available data, alongside advancements in computational tools and modelling approaches (e.g. machine learning), presents an opportunity to address these gaps. By leveraging these resources, it may be possible to predict and map cotton fibre quality and GPC at fine spatial scales within fields, even in the absence of direct measures.

Even with access to grain protein sensors and maps of GPC, understanding the drivers of variability within fields and across seasons remains a challenge. Variability in GPC is influenced by complex interactions between numerous factors, including management practices, rainfall and temperature, and soil properties. While management strategies such as variable-rate N fertilisation can be used to manage variability in GPC (Wang *et al*. 2023), identifying and isolating the root causes of this variability remains a challenge. For instance, certain patterns of variability may be due to inherent soil or climatic limitations that are beyond the growers ability to control. Therefore, further research is needed to untangle the complex drivers of variability in GPC within fields and determine the extent to which this variability can be effectively mitigated through management interventions.

Finally, while growers now have access to more data than ever before, interpreting and using this information to make informed decisions remains a significant challenge. Publicly-available datasets, such as remote sensing imagery, and on-farm data layers, including soil surveys, and yield and GPC maps, provide valuable insights into how and why variability is occurring. However, more work is needed to translate this data into actionable management strategies that enhance farm profitability, production efficiency, and environmental sustainability. Tools and frameworks need to be developed to simplify the integration of these diverse data sources into simple decision-making tools that assist growers towards the stepwise adoption of PA.

1.4 Research questions and themes

This thesis explores the application of on-farm and publicly-available spatial data layers, and advanced modelling approaches (e.g. machine learning, geostatistics), for the description, characterisation, and quantification of within-field variability in cotton fibre and grain quality, and to understand the drivers of this variability for improved management. For cotton, research is focused on developing a map of cotton fibre quality at a fine-resolution, as further research into the use of PA for better understanding and managing variability in cotton fibre quality cannot be conducted without this. The remainder of this thesis explores the use of PA to better understand and manage for variability in wheat GPC. This research leverages off new and emerging technologies (e.g. John Deere's round-module cotton harvesters and on-the-go protein sensors) and collaborations with the commercial precision agriculture company PCT to facilitate access to on-farm data.

This thesis seeks to address five specific research questions based on the gaps identified:

(1) Can high-resolution spatial maps of variability in cotton fibre quality be produced?
(2) Can the wealth of available spatiotemporal data be used to predict GPC at a fine-resolution:
    (a) within parts-of a field where two headers with yield monitors but only one protein sensor are operating; or

(b) across whole fields in the complete absence of grain protein sensor data but there is protein sensor data available for surrounding fields or in past seasons?

(3) Can the drivers of variability in wheat GPC be better understood?

(4) What is the opportunity for managing variability in GPC?

(5) If there is an opportunity for managing variability in GPC, what can be done?

The five research questions will be discussed across three broad themes:

HOW? How does quality vary?
WHY? Why does quality vary?
SO, WHAT? What can be done to optimise quality?

These research questions and themes will be addressed within the four research chapters of this thesis. Specifically, "How?" will address research questions 1 and 2 by developing methods to map cotton fibre and grain quality, as well as research question 5 by quantifying the opportunity for SSCM. "Why?" will address research question 3 by using spatial data layers to understand the drivers of variability in grain quality; and "So what?" will address research question 5 by identifying management strategies to optimise quality.

1.5 Thesis structure

This thesis is divided into six chapters, incorporating one published article in a peer reviewed journal, and numerous peer-reviewed conference papers that have been extended upon within relevant chapters. Two published book chapters have been included as appendices.

Chapter 1 (this one) provides an overview and background of the Australian cotton and grains industries, and the current role of PA in understanding and managing for variability in cotton fibre and grain quality.

Chapter 2 presents a generalised geostatistical approach using area-to-point kriging to map and downscale areal observations of crop production data, which is illustrated using cotton yield and fibre quality (length and micronaire) data which is measured as a module (areal/block)

average. This chapter has been published in the peer-reviewed journal *Precision Agriculture* (Tilse *et al*. 2024e).

Chapter 3 demonstrates how a combination of readily-available yield, agronomic, and publicly-available data layers can be used to create a model to predict GPC within-fields and fill-in gaps in the absence of a protein sensor. This chapter presents an extended version of conference papers presented at the *16th International Conference on Precision Agriculture* (Tilse *et al*. 2024c), the *Society of Precision Agriculture Australia (SPAA) 2024 Australasian Precision Ag Symposium* (Tilse *et al*. 2024a), and the *Grains Research and Development Corporation (GRDC) 2024 Northern GRDC Grains Research Update* (Tilse *et al*. 2024b).

Chapter 4 investigates the relationship between wheat grain yield and GPC and applies interpretive machine learning approaches using existing spatial data layers to understand the drivers of spatial variability in GPC within-fields. The utilisation of this knowledge for optimal N management is also explored. This chapter presents an extended version of a peer-reviewed conference paper presented at the *21st Australian Agronomy Conference* (Tilse *et al*. 2024d).

Chapter 5 compares the opportunities for SSCM for wheat grain yield and GPC by quantifying the magnitude and spatial structure of within-field variability using the Opportunity Index (OI) metric.

Finally, Chapter 6 presents a general discussion and overview of the thesis, including how each of the research questions and themes were addressed, alongside limitations and opportunities for future research.

Two book chapters are also included as Appendices. The presence of soil physical or chemical constraints can reduce crop growth, yield and quality by restricting root growth and/or water or nutrient uptake, and are a big driver of quality variation. Appendix A was published in the book *Soil Constraints to Crop Production* (Dang *et al*. 2022) and provides an overview of the range of soil sampling strategies and techniques available to understand soil physical, chemical, and biological constraints in both the surface and/or the subsoil (Tilse *et al*. 2022). Proximal sensing technologies, including EM surveys, are useful for describing variability

within fields, and spatial data layers obtained via the use of proximal sensors were used throughout this thesis. Appendix B was published in the book *Encyclopedia of Soils in the Environment, Second Edition* (Goss *et al*. 2023) and provides an overview of available proximal soil sensors available and their use in the field (Tilse *et al*. 2023).

CHAPTER 2

Downscaling crop production data to fine scale estimates with geostatistics and remote sensing: a case study in mapping cotton fibre quality

---

Abstract

A generalised approach to downscale areal observations of crop production data is illustrated using cotton yield and fibre quality (length and micronaire) data which is measured as a module (areal/block) average. Two features of the downscaling algorithm are; (i) to estimate spatial trends in yield and quality using regression with fine resolution predictors such as remote sensing imagery, and (ii) use area-to-point kriging (A2PK) to downscale either the observations in the absence of a useful spatial trend model or the residuals from the trend model (if useful) from areal averages. Correlations with remote sensing covariates were stronger for cotton fibre yield than for cotton fibre micronaire, and much stronger compared to those for cotton fibre length. Spatial trends in cotton fibre yield and micronaire could be estimated with good model quality using regression with remote sensing covariates with or without A2PK in almost all fields. Conversely, model quality was poorer for cotton fibre length and there was only a small difference in model performance between the null and trend models. When the downscaling approach was tested using fine-resolution yield observations, model performance was poorer at a fine-resolution compared to the module-resolution, which was to be expected. This approach enables the creation of high-resolution raster maps of variables of interest with a much finer spatial resolution compared to the areal observations, and can be applied for any areal averaged crop production data in a range of broadacre and horticultural industries (e.g. sugarcane, apples, citrus). The finer spatial resolution may allow

growers or agronomists to better understand the drivers of variability within fields, assess management implications, and create management plans at a higher resolution.

2.1 Introduction

Downscaling, which is the disaggregation of coarse spatial-resolution data to finer-resolution predictions, areal data is a challenge in agricultural and geographical applications, including for crop yield prediction (Brus *et al*. 2018), and when working with remote sensing imagery at different spatial resolutions (Wang *et al*. 2015). Harvester mounted, on-the-go yield monitors are not available for all crops, or in all areas of the globe. Instead, areal observations at the block, sub-, or whole-field scale is often the only available yield data for many commodities, including sugarcane and manually harvested horticultural crops. While cotton yields can be mapped easily with harvester-mounted yield monitors, cotton fibre quality cannot be mapped on-the-go like quality variables of other crops (e.g. grain protein content). Instead, cotton fibre quality data is collected post-harvest and is typically returned to growers as averaged values across whole or part of a field that are not directly linked to in-field locations.

Cotton fibre quality is expressed by a number of different measurements, including fibre length, strength, micronaire (a composite measurement of fibre fineness (diameter) and maturity), colour, and trash (any material other than cotton fibre in a sample, e.g. leaf and plant matter) content, and are described in a number of grades which affect the final price that is paid for a (227 kg) bale of cotton (CRDC *et al*. 2023). In the context of the Australian and United States of America (USA) cotton industries, a bale of cotton refers to 227 kg (500 pounds) of pressed and bound cotton. Fibre quality data for each bale is used by merchants and spinners, and growers are paid based on this quality.

Producing high-yielding, and high-quality cotton fibre is a key factor for growers to attain premium prices (Long *et al*. 2013b). While Australian cotton growers are renowned for producing high yields, the costs of production are continuing to increase (CRDC *et al*. 2023). Also, there is still the issue of significant variation in fibre quality, namely length and micronaire, both within and between fields. While soil water availability (Ballester *et*

al. 2019), variety selection (Eaton 1947; Meredith 1986; Meredith 1994), and management factors including nitrogen (N) applications (Sluijs 2022) are understood to be among the drivers of variation in cotton fibre length and micronaire, it is still often difficult to predict or understand the causes of this variation, particularly spatially within fields.

Together with yield, fibre quality is a key determinant of the prices that growers receive. Cotton fibre with quality characteristics below the Australian 'base grade' (CRDC *et al*. 2023) are subject to considerable discounts which negatively impact grower returns. For example, the average yield for the 2021-22 cotton growing season in Australia was 11.31 bales/hectare (ha) under irrigated conditions (CRDC 2022). Given this, for an Australian grower with a 100 ha field, the difference between a Middling (31) and Strict Low Middling (41) colour grade may result in loss of AUD$75/bale (~US$50/bale, based on AUD rate of 0.66 at time of writing) or AUD$84,825 (~US$56,250) per field, while a micronaire discount from G5 (3.5 – 4.9 micronaire) to G4 (3.3 – 3.4 micronaire) may result in a loss of AUD$50/bale (~US$33/bale) or AUD$56,550 (~US$37,500) per field (McVeigh 2017). Given that Australian cotton receives only a small premium for high quality in the present market (McVeigh 2017), there is an opportunity for growers better understand and manage for cotton fibre quality to minimise discounts and maximise their return on investment.

In commercial broadacre cotton in countries like Australia, the USA, and Brazil, it is possible to create post-harvest areal (module-aggregated) maps of cotton fibre quality using module completion locations, yield, and quality data from John Deere's Harvest Identification (HID) system onboard round-module harvesters (Fuhrer *et al*. 2020; Fuhrer 2022). While these maps are rarely produced, they do give some indication to growers of the broad spatial trends. However, these coarse module-aggregated fibre quality maps are at insufficient spatial resolutions to understand how and why fibre quality varies within fields, and for management decisions.

Yield maps are a valuable component of precision agriculture (PA). Broadly, the different uses of yield maps can be broken down into three main categories: 1) understanding spatial variability, 2) auditing production, and 3) guiding decisions for site-specific management (Longchamps *et al*. 2022). Collecting and accessing high-quality data at harvest that is

presented in a useful and informative way for any, or all, of these applications is a key consideration for researchers, growers, and advisors. The adoption of PA practices is limited for many manually harvested horticultural crops and for managing cotton fibre quality due to a lack of available sensing technologies and available data, or coarse spatial resolutions of the data when available. Downscaling yield and quality data can help to overcome these limitations for implementing site-specific management strategies, and for better understanding the drivers of variability in crop yield and quality within fields.

Downscaling areal yield data is a particular issue for manually harvested fruit and vegetable crops, and broadacre commodities like sugarcane. For many hand-harvested fruit crops, yield data is aggregated into bins across several trees, to management zones, or to a block, where a block may refer to a field or stand of trees that are managed together. Likewise, while yield monitors for sugarcane are commercially available, they are rarely used, and yield data is typically aggregated to a block.

Downscaled yield maps have been produced for manually harvested fruit crops (e.g. citrus) by georeferencing boxes, bins, or bags across an orchard and using ordinary kriging to interpolate estimates (Colaço *et al*. 2020). However, ordinary kriging assumes a point support, meaning it operates under the assumption that the data or measurements are collected at specific, discrete locations without any spatial averaging, and the long, narrow, unusual shape of cotton round-modules presents a challenge. A round-module of cotton can weigh between 2000 – 2600 kg, producing f̃our bales of ginned lint. Given that the average yield for the 2021-22 cotton growing season in Australia was 11.31 bales/hectare (bales/ha) under irrigated conditions (CRDC 2022), a single round module of cotton has a harvest area of ~3500 m^2. For a standard cotton picker front of 6 m and an irrigated cotton field with a row length of 500 m, this equates to almost two passes of the field, typically from alternating rows, for a single round-module of cotton.

Area-to-point kriging (A2PK) is a geostatistical framework used to disaggregate areal observations to the point support and is useful for creating high-resolution raster maps (Brus *et al*. 2018). Area-to-point kriging presents an alternative approach for downscaling areal harvest

data as it incorporates the variable sizes and shapes of areal data into the variogram deconvolution and kriging process (Kerry *et al*. 2012). This is particularly useful for downscaling irregular blocks of data, or cotton modules in this case. A block in geostatistics is distinct from the management area, field, or block referred to previously, where the block support in geostatistics refers to the area at which a measurement is taken or a prediction made.

Auxiliary information such as remote sensing imagery may also be used to capture spatial-variability within blocks or areal units via a regression-kriging approach. The incorporation of fine-scale covariate data can capture within-block variability and represent the factors influencing crop growth and development, such as soil physical and chemical constraints (Pallegedara Dewage *et al*. 2020). Unlike bespoke, on-farm data such as electromagnetic (EM) surveys, management information, or soil data, satellite-based remote sensing imagery is publicly-available across all fields and farms across the globe. Remote sensing platforms such as Sentinel-2, which has a spatial resolution of up to 10 m and a revisit time of ~five days, can capture within-field crop reflectance variations throughout the growing season and have been used to model cotton yield (e.g. He *et al*. 2019; Leo *et al*. 2021).

Here, an approach to downscale and unlock existing module-aggregated cotton quality data for more precise management of cotton fibre quality is presented.

Two features of the downscaling approach are:

(1) To estimate spatial trends (using regression) in quality variables with publicly available, fine-resolution remote sensing imagery; and
(2) Use A2PK to downscale either:
    (a) The observations in the absence of a useful spatial trend model, or
    (b) The residuals from the trend model (if useful) from areal averages.

The downscaling approach is illustrated across 11 fields in New South Wales (NSW), Australia for two cotton fibre quality attributes (length and micronaire), and is further validated using fine-resolution yield observations.

The aims of this study were to:

(1) Assess the ability of remote sensing imagery to model variability in cotton yield and fibre quality variables. This is the foundation of the downscaling approach.
(2) Illustrate and test the downscaling approach using fine-resolution yield data.

The goal of this study is not to forecast cotton fibre quality prior to harvest.

This approach enables the creation of high-resolution raster maps of variables of interest with a much finer spatial resolution compared to the areal observations, and can be applied for any areal averaged data (e.g. sugarcane, manually harvested horticultural crops). The finer spatial resolution may allow growers or agronomists to better understand the drivers of variability within fields, assess management implications, and create management plans at a higher resolution.

2.2 Methods

2.2.1 Study area and available data

Cotton fibre quality and yield data was collected from 11 fields across two cotton growing regions of NSW, Australia for the 2021/22 growing season. The dataset included three flood-irrigated cotton fields from southern NSW (SNSW) and eight flood-irrigated cotton fields from northern NSW (NNSW) (Figure 2.1, Table 2.1). The cotton was sown in October/November 2021, and was harvested as late as July 2022 due to a wet finish to the season.

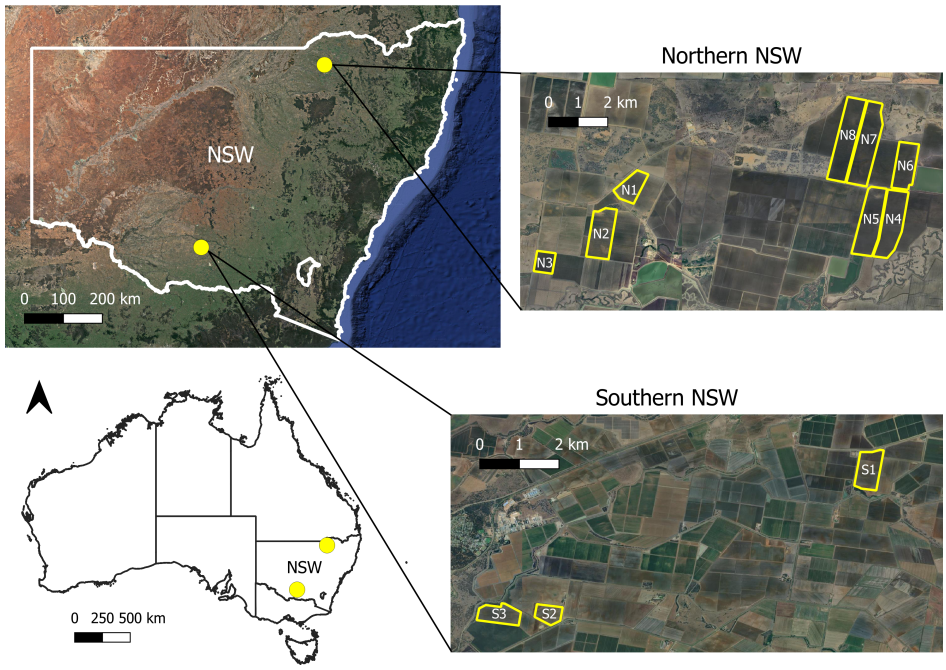

FIGURE 2.1. Location of the study fields relative to New South Wales (NSW), Australia, and proximity of study field locations to each other, in northern NSW and northern NSW, respectively.

TABLE 2.1. Information describing fields with available cotton fibre quality and yield data across southern New South Wales (SNSW) and northern New South Wales (NNSW), including field area (in hectares, ha), the number of round modules, and the number of individual (227 kg) bales within each field.

Region	Field	Area (ha)	Number of modules	Number of (227 kg) bales
SNSW	S1	66	124	598
	S2	33	72	364
	S3	53	116	574
NNSW	N1	79	290	1301
	N2	145	419	2495
	N3	51	152	646
	N4	177	314	2198
	N5	185	384	2464
	N6	120	256	1556
	N7	216	484	2945
	N8	210	438	2714

Cotton fibre quality variables were measured post-harvest at a commercial classing facility using a High-Volume Instrument (HVI) and cotton quality data was collected for length (inches) and micronaire (unitless), among other fibre quality variables not examined in this study. Cotton fibre quality data is measured and reported for each bale, and Radio Frequency Identification (RFID) tags enable each bale to be linked back to the respective round module. There are approximately four bales per one round module and fibre quality data for each bale was aggregated (averaged) for each round module, resulting in one value for each fibre quality variable for each module. For each field, there were between 72 and 484 individual cotton modules of data available across fields ranging between 33 and 216 hectares. In total, there was between 364 and 2945 bales worth of cotton fibre quality data for each field (Table 2.1). Yield data was recorded on-the-go at harvest by harvester-mounted yield monitors.

Time- and geo-referenced information collected from John Deere's HID system (including module wrap location, time, and yield data) was used to map module-aggregated areal cotton fibre quality data for length and micronaire across each field using the methodology adapted from Fuhrer *et al*. (2020) and Fuhrer (2022). This was done by using the round module and harvest data time-stamps, whereby yield data points recorded after the time-stamp of the previous module through to the time-stamp of the current module were allocated to the current module. This effectively gives the area harvested by which the module came from and maps each module's travel path. Module-aggregated fibre quality data was then linked back to each module area, and fibre quality variables were mapped across each field to produce a "floorboard map" of module-resolution fibre quality data (Figure 2.2). Yield data was also aggregated to the module-resolution by averaging yield data point values within each module, based on its spatial footprint. All data processing and analysis was performed in R, version 4.3.1, unless specified otherwise.

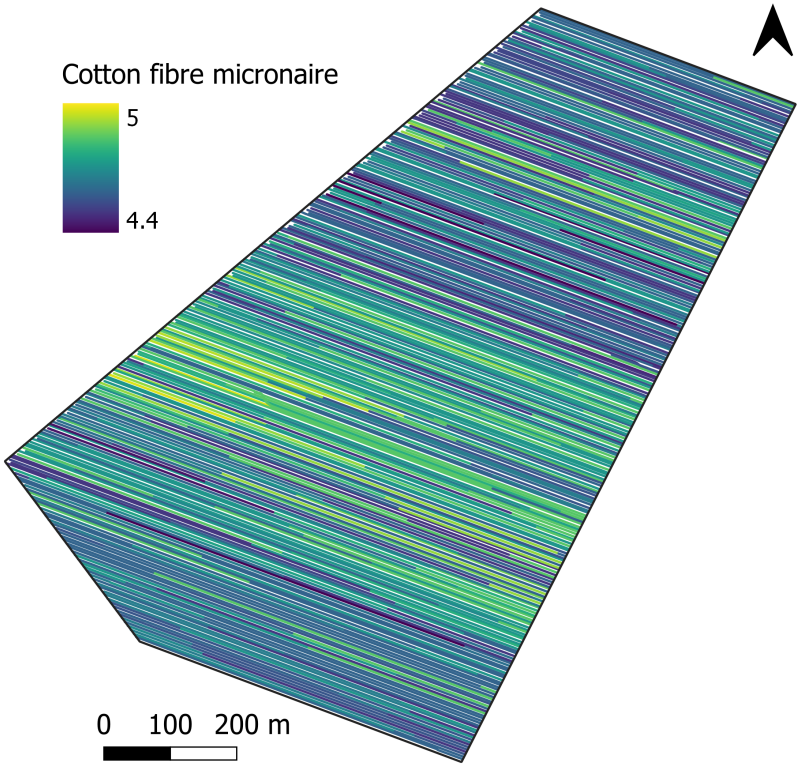

FIGURE 2.2. A module-aggregated (averaged) "floorboard map" for cotton fibre micronaire for field N1. Cotton fibre micronaire (unitless) is a composite measurement of fibre fineness (diameter) and maturity.

2.2.2 Remote sensing covariates

All remote sensing images throughout the growing season (~October 2021 – July 2022) were obtained from Sentinel-2 at a 10 m spatial resolution with a five-day revisit time. Sentinel-2 surface reflectance imagery was accessed through the Google Earth Engine (GEE) (Gorelick *et al*. 2017) via the R package 'rgee', version 1.1.6. (Aybar *et al*. 2020). Sentinel-2 was chosen as it provides data at a high spatial (up to 10 m) and temporal (up to five-day revisit time) resolution. A cloud-masking filter was applied to all images to remove pixels affected by cloud cover, and images affected by cloud cover were removed from the analysis. In total, there was ~25 images available for the NNSW fields and ~70 images for SNSW, after

cloud-cover removal. The Sentinel-2 imagery was then used to derive a suite of vegetation indices, described below (Table 2.2), including the Normalised Difference Vegetation Index (NDVI), Normalised Difference Red Edge (NDRE), Enhanced Vegetation Index (EVI), Canopy Chlorophyll Content Index (CCCI), Green Red Vegetation Index(GRVI), Red-Edge Ratio (RE/R), Ratio Spectral Index (RSI), and Optimised Soil Adjusted Vegetation Index (OSAVI).

All vegetation indices used in this study have been reported for cotton growth, nutrition, yield or fibre quality monitoring or prediction within existing literature. Cotton fibre yield and quality variables including length and micronaire may be impaired when the crop is water stressed at particular points in the growing season (Bange *et al*. 2018; Brodrick *et al*. 2012; Wiggins *et al*. 2014), and crop water stress can be observed by changes in photosynthesis and chlorophyll content (Ballester *et al*. 2019). The NDVI has been widely reported for water stress detection (Ihuoma *et al*. 2017), plant biomass and crop N status monitoring (Ballester *et al*. 2017), and has been shown to be a good predictor of cotton lint yield from early-to-mid season (Ballester *et al*. 2019). However, the NDVI may saturate at full canopy cover later in the season. Alternatively, the EVI has been used for cotton water status monitoring (Lin *et al*. 2020) and cotton canopy N estimation (Chew *et al*. 2023). The GRVI and RE/R have been shown to be well correlated with and are better predictors of cotton fibre micronaire compared to the NDVI (Ballester *et al*. 2019). The GRVI is sensitive to changes in green colour and canopy structure, particularly at full canopy cover (Ballester *et al*. 2019). The red-edge, or the transition region of reflectance between the red and near infrared regions, is sensitive to changes in leaf chlorophyll content (Ju *et al*. 2010) and the RE/R and NDRE both include information from the red edge and red bands (Ballester *et al*. 2019; Thompson *et al*. 2019). During early growth stages, the OSAVI incorporates soil-adjustment factors to reduce the influence of soil background on vegetation signals and has improved sensitivity to high biomass conditions compared to the NDVI (Rondeaux *et al*. 1996; Zarco-Tejada *et al*. 2005). The CCCI uses the NDVI, as an estimate of canopy cover/biomass, and the NDRE, as a measure of chlorophyll content, and has been used to estimate cotton N status and yield prediction (Oosterhuis *et al*. 2015; Ballester *et al*. 2017). The RSI, which represents the ratio of infrared and red light intensities, has been used for cotton yield estimation (Xu *et al*. 2021).

Given the varying performance of these eight different vegetation indices for monitoring, modelling, and predicting cotton growth, nutrition, yield and/or fibre quality, this study aimed to assess the ability of different vegetation indices to capture and describe variability in cotton fibre yield and quality spatially within fields, with a particular application for downscaling cotton fibre quality data to a fine-resolution. Also, if one vegetation index is more strongly correlated with a particular fibre quality variable or shows clear temporal trends, for example, this index may be useful for future research outside the scope of this present study such as for fibre quality forecasting.

TABLE 2.2. Vegetation indices calculated in the study for cotton yield and fibre quality modelling, obtained from Sentinel-2 satellite imagery.

Vegetation Index	Abbreviation	Formula	Reference	Application
Normalised Difference Vegetation Index	NDVI	$\frac{NIR - Red}{NIR + Red}$	Rouse *et al*. (1974)	Vegetation health and biomass
Normalised Difference Red Edge	NDRE	$\frac{NIR - RedEdge}{NIR + RedEdge}$	Gitelson *et al*. (1994)	Vegetation health and biomass, sensitivity to changes in chlorophyll content
Enhanced Vegetation Index	EVI	$2.5\frac{NIR - Red}{NIR + 6 \times Red - 7.5 \times Blue + 1}$	Liu *et al*. (1995)	Vegetation health and biomass, sensitivity in high biomass/full canopy cover
Canopy Chlorophyll Content Index	CCCI	$\frac{\frac{NIR - RedEdge}{NIR + RedEdge}}{\frac{NIR - Red}{NIR + Red}}$	Barnes *et al*. (2000)	Vegetation health and biomass, yield estimation, N status
Green Red Vegetation Index	GRVI	$\frac{Green - Red}{Green + Red}$	Tucker (1979)	Vegetation cover, sensitivity to changes in canopy structure and green colour
Red-Edge Ratio	RE/R	$\frac{RedEdge}{Red}$	Ballester *et al*. (2013)	Sensitivity to changes in chlorophyll content

All remote sensing covariates were resampled to two resolutions: 1) The first was to a 5 m resolution for each field to build the covariate grid; and 2) the second was to module-aggregated (mean) values which were extracted within each module polygon (block). The module-aggregated covariates and fine-resolution covariate grid were used for trend model building and prediction, respectively. A 5 m grid was used to best represent cotton module widths, which are ~5 m wide.

Spearman-rank correlation coefficient values were calculated between each module-aggregated covariate, and yield and fibre quality variables at the block support within each field. The Spearman-rank correlation coefficient is a non-parametric test, meaning that it does not rely on the assumption of normality, that works with rank-ordered variables and is more robust to outliers compared to Pearson's correlation coefficient. Given that more images were available in SNSW compared to NNSW as these fields were less impacted by cloud cover during the 2021/22 growing season, only the top-four most-correlated images (based on absolute values) for each vegetation index were used in the modelling and downscaling approach, described below, to reduce computation intensity and for direct comparison between both growing regions. This totalled a maximum of 32 covariates that could be used within each regression model for each field.

2.2.3 Modelling approach and the decision-making framework

The dataset of module-aggregated areal yield, fibre quality, and covariate data was split with a 70% calibration, 30% validation split for variable selection, model calibration and validation for all modelling and downscaling approaches. Both the trend (multiple linear regression, MLR or Random Forest, RF) and A2PK modelling was performed at the spatial support of the modules (block support), and all models were built using the calibration dataset. Trend (MLR and RF) models were then applied to the fine-resolution covariate grid to produce downscaled predictions at the point support.

The general form of the modelling approach consists of two components;

- a trend, which involves fitting either a MLR or RF model to the data. The modelling is performed at the spatial support of the modules; and
- modelling the spatial auto-correlation of the residuals from each trend model based on deconvoluting a point-scale variogram from the areal average data, which can be used for prediction using A2PK (Goovaerts 2008).

The decision-making framework can be summarised in Figure 2.3 and is described in further detail below. For each cotton fibre quality variable, regression (MLR or RF) models were used to estimate spatial trends in the data. If spatial trends could not be estimated, A2PK was used to downscale the module-resolution fibre quality variable observations. If spatial trends could be estimated, the model quality of each regression model was compared, and the best model was chosen (either MLR or RF). For the chosen regression modelling approach, trend model residuals were downscaled using A2PK and model quality was compared with and without the addition of the A2PK residuals. A final downscaled map of each cotton fibre quality variable was then produced using the chosen approach from the decision-making framework (Figure 2.3).

2.2.4 Null model

The null model is to use the the field mean value of the response variable for prediction. This was compared with the MLR and/or RF trend models to assess which was superior. This was assessed by calculating the root mean square error (RMSE) using the regression model predictions and comparing this to the RMSE of the null model (i.e. if the field mean of the response was used as the predictor). Fine-resolution fibre quality data for length and micronaire is not available to further validate downscaling predictions, however the fine-resolution (5 m) observed yield data was used as a surrogate throughout to assess how well downscaled predictions of areal data reflect fine-resolution observations. For fine-resolution cotton fibre yield, two null models were compared; the module aggregated (mean) yield value was used for prediction (FineRes A), and the field mean value of fine-resolution yield data

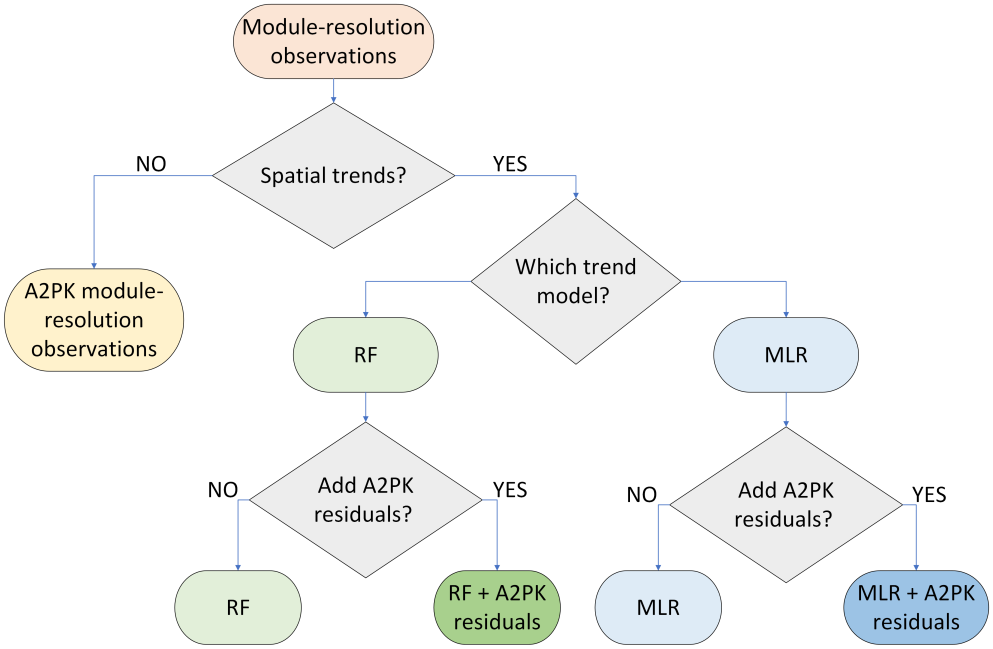

FIGURE 2.3. Decision tree framework describing the modelling and downscaling process, including area-to-point kriging (A2PK), Random Forest (RF), and multiple linear regression (MLR).

was used for prediction (FineRes B) which is equivalent to the null model used for length and micronaire.

2.2.5 Multiple linear regression

Multiple linear regression with stepwise backward elimination based on the Akaike Information Criterion (AIC) (Akaike 1998) was used to model each cotton yield and fibre quality variable for each field individually via the 'lm' function from the R package 'stats', version 4.2.3, and the 'stepAIC' function from the R package 'MASS', version 7.3.58.2. Backwards selection begins with the full model and sequentially removes the least statistically significant variables until the model with the lowest AIC is determined, indicating a more parsimonious model. Multiple linear regression was then performed using the final selection of covariates to make predictions within each field using the covariate grid.

2.2.6 Random forest models

Random Forest models (Breiman 2001), which are a suite of decision trees, were used as the alternate regression method to model the relationship between each cotton yield and fibre quality variable, and their respective top-four most-correlated images for each vegetation index (Section 2.2.2) for each field. Random Forest models are efficient at identifying and describing complex interactions and non-linear, non-parametric relationships between response and predictor variables (Breiman 2001). The R package 'ranger', version 0.14.1 (Wright *et al*. 2017), was used as it enables the fast implementation of RF models and is well-suited to high dimensional data (Wright *et al*. 2017). The 'permutation' variable importance measure from 'ranger' (Strobl *et al*. 2008; Wright *et al*. 2016) was used as part of an iterative process for variable selection. Within this process, an initial RF model was built using all 32 available covariates for each yield or fibre quality variable in each field. The permutation variable importance was then calculated, and the lowest ranked (i.e. least important) variable was removed before the process was repeated and a new RF model was built. At each iteration, out-of-bag model quality statistics (the RMSE and Lin's concordance correlation coefficient (LCCC); described in Section 2.2.9) were calculated. Out-of-bag evaluation is used to estimate the performance of the RF model by using samples not included in the bootstrap sample for training. After all iterations were complete, the change in model quality with an increasing number of covariates removed was plotted and the point where model quality substantially decreased was deemed to be the optimum number of covariates to retain within each model. While much of the covariate acquisition and modelling processes was automated, user discretion was still necessary to determine the optimum number of covariates.

2.2.7 Area-to-point kriging

Based on the RMSE of the models, if the MLR or RF model was better than the null model (i.e. a lower RMSE), it was used as the trend and the residuals were kriged using A2PK, described below. In the case when the null model was better, A2PK was used directly on the observations, which is analogous to ordinary kriging where the trend is the local mean of the observations.

In preparation for A2PK, module-resolution areal values (either observations or trend model residuals) were discretized into small regular point-samples at a 5 m resolution using the 'discretizePolygon' function in the 'atakrig' package in R, version 0.9.8. The 'atakrig' package enables the deconvolution of point-scale variograms from irregular/regular spatial areal data that does not require boundary alignment of the input data (Hu *et al*. 2020). By default, equal weight is assigned to all discrete points of each polygon area (Hu *et al*. 2020). Additionally, a discretized grid to be predicted onto was produced at a 5 m resolution, which is the same grid used throughout. A point-scale variogram deconvolution algorithm was then implemented onto the discretized point coordinates and their weights according to Goovaerts (2008) using the 'deconvPointVgm' function in 'atakrig'. The iterative procedure is described by Hu *et al*. (2020). Area-to-point kriging was then implemented using the 'ataKriging' function in 'atakrig'.

2.2.8 Predictions

If chosen (based on the RMSE), the MLR or RF trend model was then used to make predictions across each field at a 5 m spatial resolution using the covariate grid. Area-to-point kriging was then used to downscale the residuals of each chosen MLR or RF model onto the same 5 m grid. The final map output was a downscaled map of model predictions with A2PK of the residuals (here-on referred to as MLR + A2PK or RF + A2PK, depending on the chosen trend model approach) for each yield and fibre quality variable on a 5 m grid. If the fit of the null model was better than the MLR or RF trend models, A2PK was used on its own to downscale module-aggregated observations from the block to point (fine-resolution, 5 m, grid) support.

2.2.9 Model validation

To assess model quality, downscaled predictions were aggregated (mean) to the module-resolution and extracted for each module polygon area (i.e. to the block support). These aggregated predictions were then compared to module-aggregated observed yield and fibre

quality values at the module locations from the validation dataset. In the absence of fine-resolution cotton fibre quality data for length and micronaire, cotton yield data was used as a surrogate to assess how well downscaled predictions of areal data perform at a fine-resolution and reflect within-module variability. The LCCC was used as an assessment of model quality. The LCCC is a measure of both precision (how close the predictions are to each other) and accuracy (how close the predictions are to the observed values). The LCCC explains the fit of the observed and predicted values to the 1:1 line, where values of 0 are a poor fit and 1 for a perfect fit. The LCCC is unit-less, and is useful for comparing the quality of predictions between variables with different magnitudes (Lin 1989).

2.2.10 Day degrees calculation

Crop development is strongly influenced by temperature, and the rate of crop development is commonly estimated using thermal time as a function of growing day degree's (DD). In this study, all variable selection was conducted based on correlations for individual dates and images and did not factor in crop growth stages. Further reporting and discussion of variable selection and model predictions (Sections 2.3 and 2.4) were made with reference to crop development stages, but were not used in the models.

Day degrees is the accumulation of heat units related to the daily maximum and minimum temperatures. The 1532 DD method (Bange *et al*. 2022) has a base temperature of 15.6°C (average) and an optimum temperature of 32°C (maximum) (Equations 2.1 and 2.2):

$$DD = \frac{(T_{max} - T_{min})}{2} - 15.6 \tag{2.1}$$

Where maximum temperature (T_{max}) is < 32°, or

$$DD = \frac{(32 - T_{min})}{2} - 15.6 \tag{2.2}$$

Where T_{max} is > 32°. T_{min} is the minimum temperature.

The accumulated DD targets for cotton development stages were used as follows (Table 2.3).

TABLE 2.3. Crop development targets for day degrees (DD) based on the 1532 DD method (Bange *et al*. 2022; CRDC *et al*. 2023) and corresponding growth stage referred to in this study.

Development stage	Accumulated DD after planting	Growth stage
Emergence	50	Planting to Emergence
First square	339	Emergence to First Square
First Flower	584	Square to First Flower
First Open Boll	1093	Flower to Open Boll
	>1093	Open Boll to Harvest

2.3 Results

2.3.1 Exploratory data analysis

Observed cotton fibre length values were within or above the Australian Base Grade (i.e. 1.11 – 1.13 inches) for all fields (Figure 2.4). Over 90% of cotton fibre micronaire values in all NNSW fields, and SNSW fields S2 and S3, were within the Australian Base Grade (i.e. 3.5 – 4.9; Figure 2.5). Conversely, over 90% of modules in SNSW field S1 has observed cotton fibre micronaire values below the Australian Base Grade (i.e. < 3.5; Figure 2.5). Observed yield values ranged between ~2 and 16 bales/ha. Observed yield values in fields N1 and N2 were significantly higher than for all other fields, and yield values in SNSW (fields S1, S2, S3) were significantly lower compared to fields in NNSW (Figure 2.6).

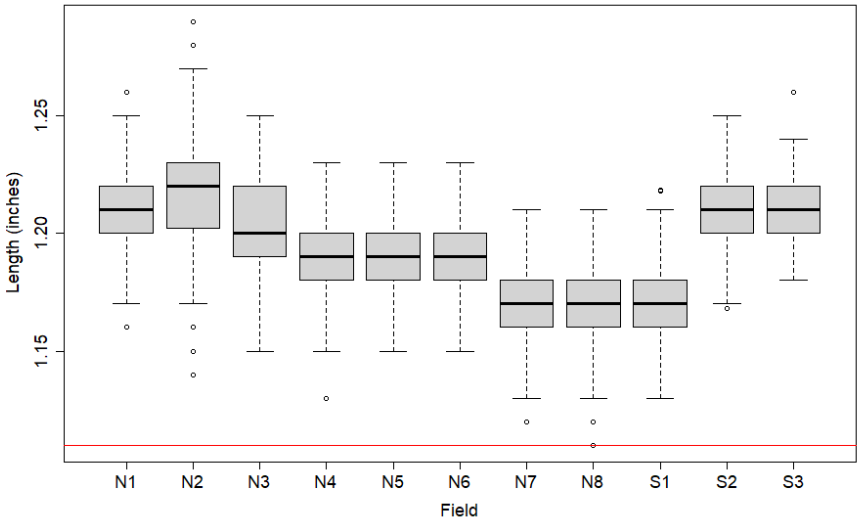

FIGURE 2.4. Distribution of observed cotton fibre length (inches) values aggregated to the module resolution for each field. Red line indicates the lower threshold of the Australian Base Grade for cotton fibre length (1.11 – 1.13; CRDC *et al*. 2023), where values below this range are penalised with discounts.

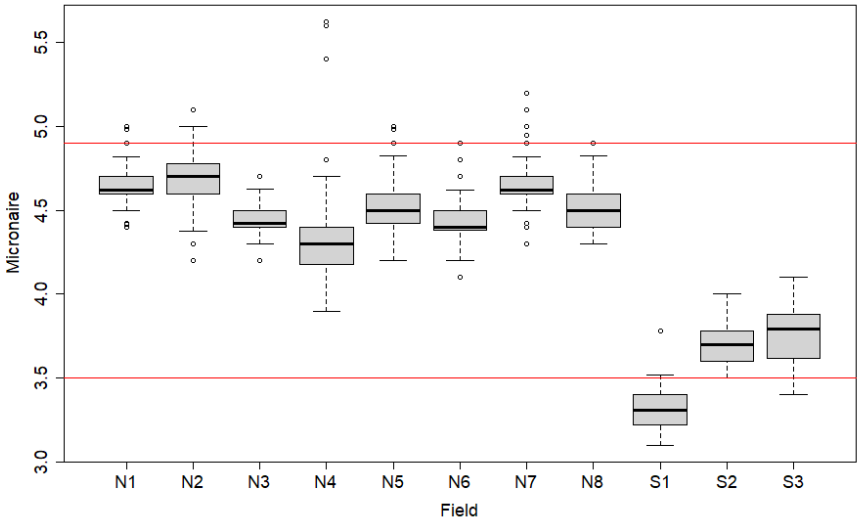

FIGURE 2.5. Distribution of observed cotton fibre micronaire (unitless) values aggregated to the module resolution for each field. Red lines indicate the boundaries of the Australian Base Grade for micronaire (3.5 – 4.9; CRDC *et al*. 2023), where values outside of this range are penalised with discounts and values within this range may be awarded premiums.

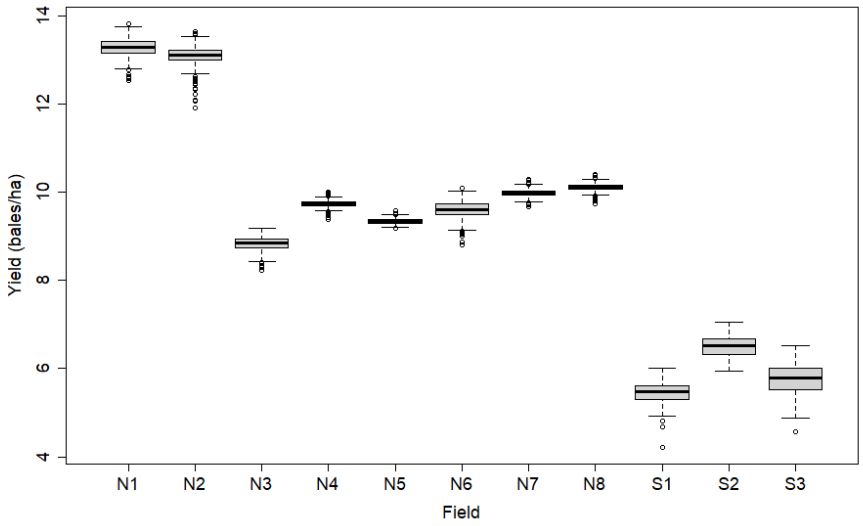

FIGURE 2.6. Distribution of observed cotton fibre yield (bales/hectare, ha) values aggregated to the module resolution for each field.

Crop development stages were calculated based on growing DD using the DD function developed by Bange *et al*. (2022). The rate of growth and development in SNSW was approximately 20 – 40 days slower than for NNSW (relative to days after planting, DAP) (Figure 2.7). Notably, the duration of Square to First Flower was considerably shorter than other crop development stages and NNSW progressed between the different growth stages earlier compared to SNSW (Figure 2.7). This had implications for the relationship between fibre quality attributes and variable selection for covariates across the growing season.

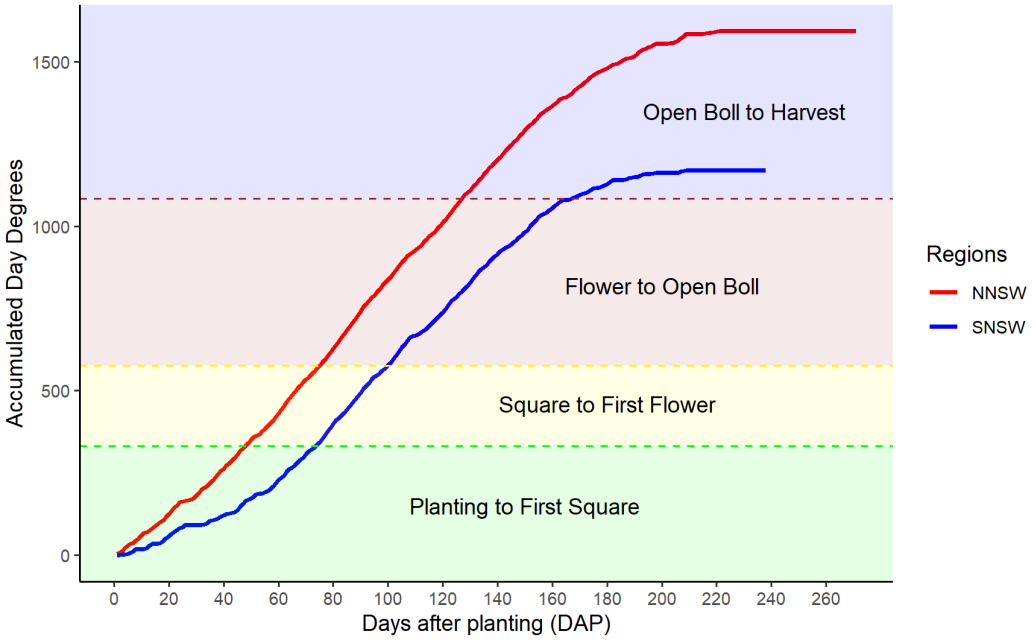

FIGURE 2.7. Accumulated Day Degrees against days after planting (DAP) for NNSW compared to SNSW. Accumulated Day Degrees calculated based on the 1532 DD method (Bange *et al*. 2022; CRDC *et al*. 2023).

2.3.2 Correlations over the length of the growing season

To explore the relationship between cotton fibre yield and quality variables, and remotely-sensed vegetation indices, over the length of the growing season, spearman-rank correlation coefficient values were calculated between all available module-aggregated covariates, and yield and fibre quality variables at the block support within each field. Overall, correlations between available remotely sensed vegetation indices and cotton fibre yield were stronger compared to those with cotton fibre length and micronaire (Table 2.4). More images were available over the growing season for SNSW compared to NNSW as NNSW has summer-dominant rainfall, resulting in more clouds during the cotton growing season and thus more images were filtered out due to cloud cover (Section 2.2.2). For cotton fibre length and yield, correlations were stronger in SNSW compared to NNSW, based on the largest absolute value (most negative/most positive). However, for cotton fibre micronaire, correlations were marginally stronger in NNSW compared to SNSW, based on largest absolute values (Table 2.4).

TABLE 2.4. Summary statistics of Spearman-Rank Correlation Coefficient values between cotton fibre length, micronaire, and yield, and all remotely-sensed vegetation indices available across the entire growing season for fields in NNSW and SNSW. Mean values calculated from and presented as absolute values.

			RSI	RE/R	OSAVI	NDVI	NDRE	GRVI	EVI	CCCI
Length	NNSW	Min	-0.24	-0.24	-0.27	-0.25	-0.31	-0.27	-0.27	-0.34
		Max	0.35	0.35	0.32	0.32	0.36	0.36	0.24	0.32
		Mean	0.09	0.10	0.09	0.09	0.09	0.10	0.08	0.08
	SNSW	Min	-0.44	-0.48	-0.43	-0.43	-0.44	-0.47	-0.47	-0.43
		Max	0.41	0.57	0.41	0.42	0.46	0.50	0.44	0.44
		Mean	0.14	0.15	0.14	0.15	0.15	0.13	0.14	0.15
Micronaire	NNSW	Min	-0.69	-0.72	-0.67	-0.68	-0.67	-0.73	-0.60	-0.59
		Max	0.68	0.76	0.70	0.68	0.67	0.60	0.71	0.59
		Mean	0.21	0.21	0.22	0.21	0.20	0.21	0.21	0.19
	SNSW	Min	-0.54	-0.51	-0.55	-0.55	-0.56	-0.55	-0.54	-0.54
		Max	0.53	0.52	0.52	0.52	0.51	0.31	0.43	0.41
		Mean	0.17	0.17	0.18	0.17	0.18	0.15	0.17	0.17
Yield	NNSW	Min	-0.39	-0.62	-0.34	-0.37	-0.46	-0.71	-0.17	-0.40
		Max	0.78	0.78	0.81	0.81	0.79	0.76	0.84	0.78
		Mean	0.18	0.14	0.19	0.19	0.20	0.14	0.17	0.19
	SNSW	Min	-0.67	-0.66	-0.66	-0.64	-0.68	-0.60	-0.68	-0.64
		Max	0.72	0.80	0.67	0.73	0.73	0.76	0.65	0.53
		Mean	0.33	0.26	0.31	0.31	0.28	0.23	0.29	0.20

Out of all available images over the growing season, the strongest negative relationship with cotton fibre length was -0.47 (GRVI) and the strongest positive relationship was 0.57 (RE/R; Table 2.4). Based on largest absolute values, correlations with cotton fibre length were stronger for the RE/R, GRVI, EVI and NDRE compared to other vegetation indices (Table 2.4). When averaged across all fields, there was a weak temporal trend shifting from negatively to positively correlated as the season progressed for most vegetation indices (Figure

2.8), and correlations with cotton fibre length were typically stronger for SNSW fields (Figure 2.8b).

Out of all available images over the growing season, the strongest negative relationship with cotton fibre micronaire was -0.73, and the strongest positive relationship was 0.76 (Table 2.4). Correlations generally shifted from positive to negative as the season progressed, and were strongest towards the end of the growing season during Flower to Open Boll, and Open Boll to Harvest (negatively correlated, Figure 2.9). Compared to all other vegetation indices, largest absolute correlations with cotton fibre micronaire were for the RE/R, GRVI, and EVI in NNSW, and for the NDRE, GRVI, NDVI and OSAVI in SNSW (Table 2.4).

Correlations were strongest between cotton fibre yield and remotely sensed vegetation indices, compared to those with cotton fibre length and micronaire. The strongest negative relationship was -0.71 and the strongest positive relationship was 0.84 (Table 2.4). Correlations with cotton fibre yield showed a clear temporal trend, with the strength and direction of mean correlations across all fields shifting from positive to neutral (NNSW; Figure 2.10a) or positive to negative (SNSW; Figure 2.10b) as the season progressed. This temporal trend was clearer for SNSW fields compared to those in NNSW, and correlations between vegetation indices and cotton fibre yield were stronger for SNSW overall. Correlations were strongest during Flower to Open Boll for NNSW (positively correlated; Figure 2.10a), and from Planting to First Square (positively correlated) and Open Boll to Harvest (negatively correlated) for SNSW (Figure 2.10b). Overall, largest absolute correlations with cotton fibre yield were for the EVI, OSAVI, NDVI and RE/R (Table 2.4).

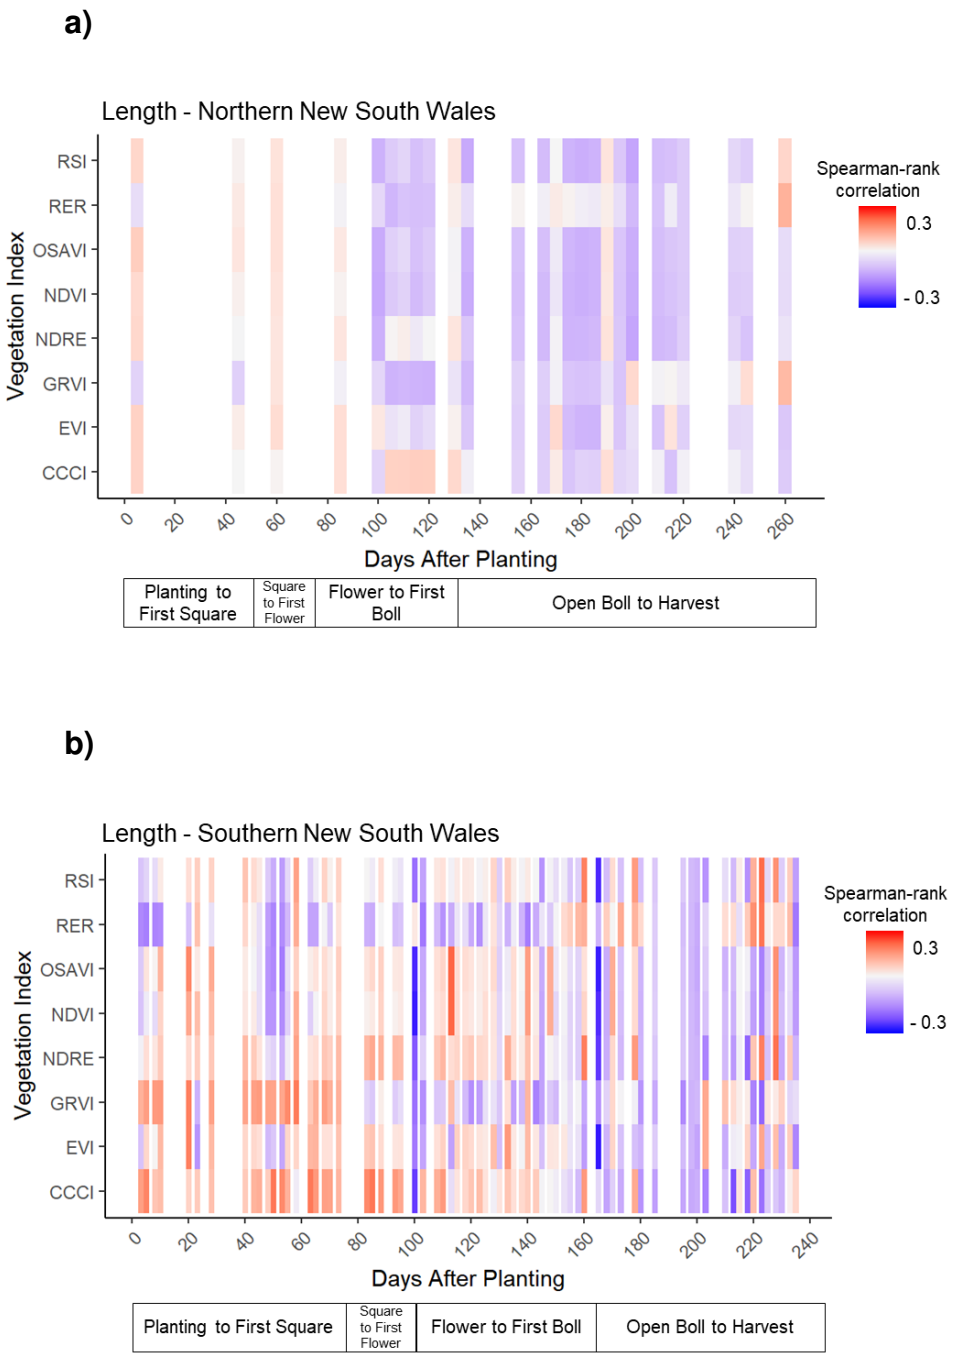

FIGURE 2.8. Average Spearman-Rank Correlation Coefficient values between cotton fibre length and remotely sensed vegetation indices for each date across the growing season for (a) NNSW and (b) SNSW. Days after planting and growth stage targets based on the 1532 day degree (DD) method (Bange *et al.* 2022; CRDC *et al.* 2023) are presented on the x-axis.

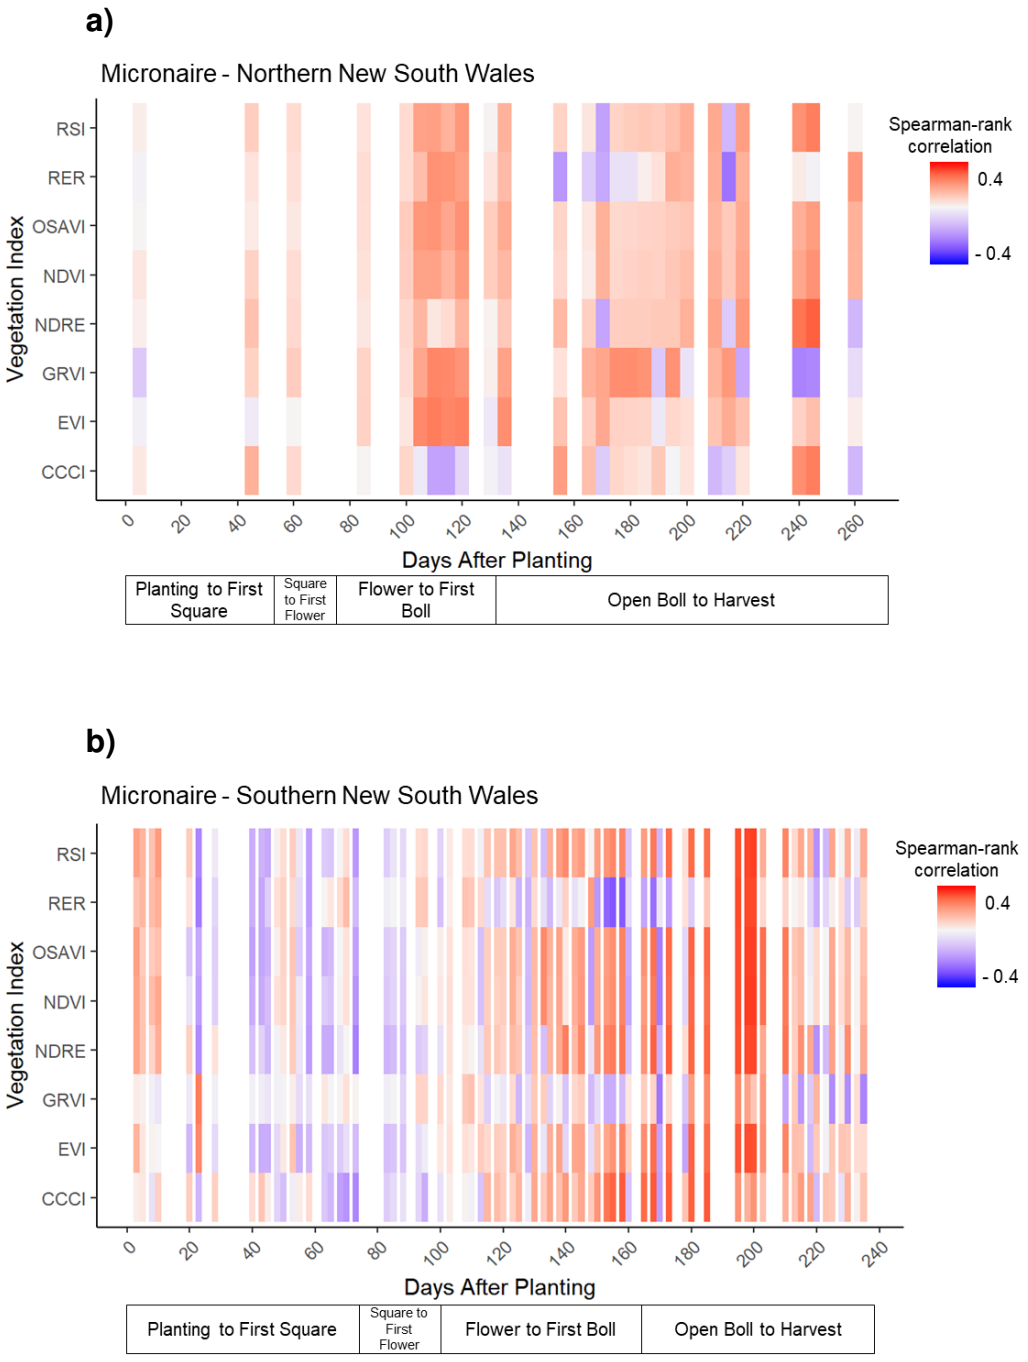

FIGURE 2.9. Average Spearman-Rank Correlation Coefficient values between cotton fibre micronaire and remotely sensed vegetation indices for each date across the growing season for (a) NNSW and (b) SNSW. Days after planting and growth stage targets based on the 1532 day degree (DD) method (Bange *et al*. 2022; CRDC *et al*. 2023) are presented on the x-axis.

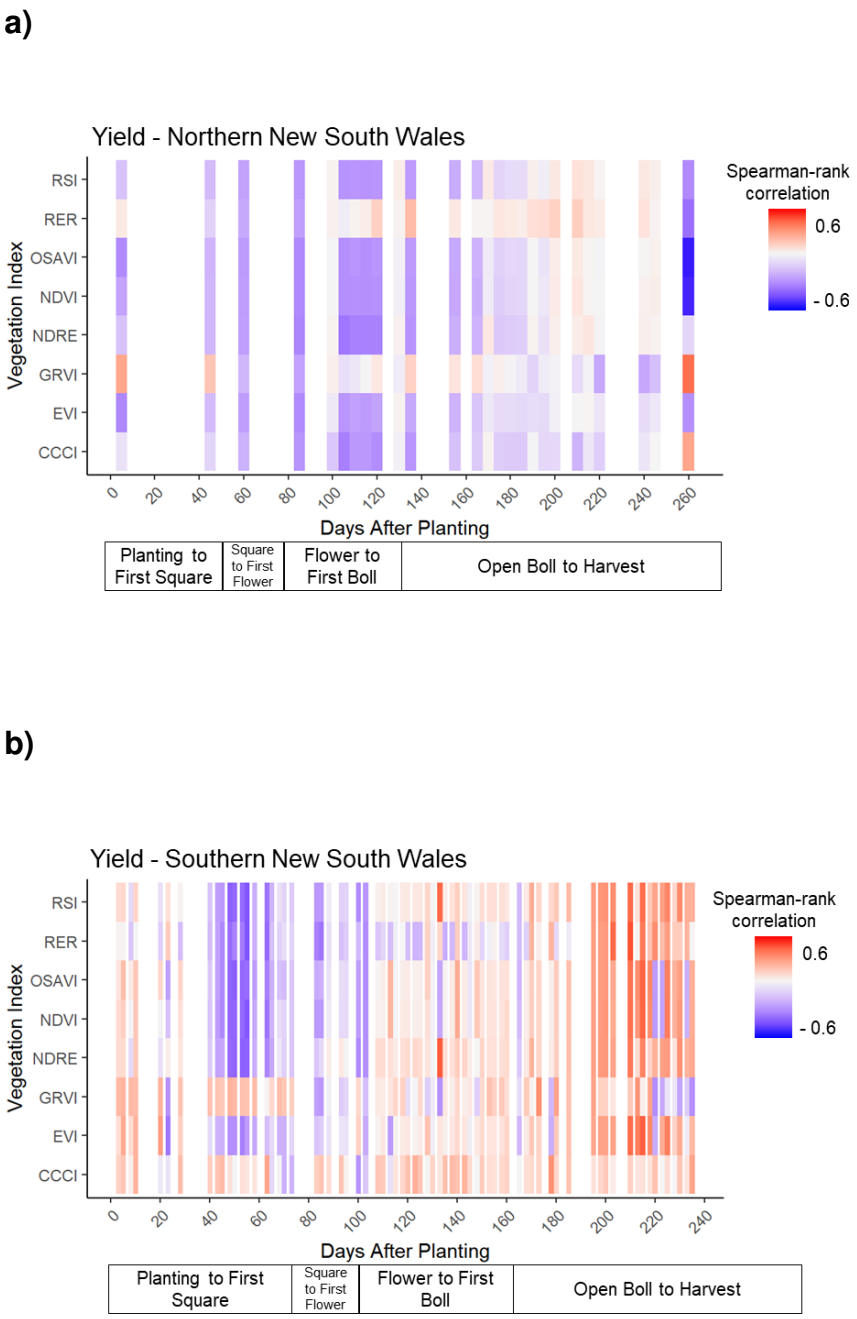

FIGURE 2.10. Average Spearman-Rank Correlation Coefficient values between cotton yield and remotely sensed vegetation indices for each available date across the -growing season for (a) NNSW and (b) SNSW. Days after planting and growth stage targets based on the 1532 day degree (DD) method (Bange *et al*. 2022; CRDC *et al*. 2023) are presented on the x-axis.

2.3.3 Variable selection

Based on the spearman-rank correlation coefficient values calculated between each module-aggregated remote sensing covariate, and cotton fibre length, micronaire, and yield variables (Section 2.3.2), the top-four most-correlated images for each vegetation index for each response variable were then selected to reduce computation intensity and for direct comparison between both growing regions. This totalled 32 possible covariates that could be used within each model prior to further variable selection during model calibration. Of those 32 images that were selected for each variable, the majority were from later in the growing season during Flower to Open Boll, and Open Boll to Harvest. Consequently, after the iterative variable selection process during model calibration, covariates used in the final regression models for each variable were also predominately from later in the growing season as well.

Of covariates chosen in the final MLR models for cotton fibre length, the strongest negative relationship (as judged by the correlation) was -0.48 and the strongest positive relationship was 0.57 (both the RE/R; Table 2.A1). For the RF models, the strongest negative relationship was –0.47 (EVI) and the strongest positive relationship was 0.80 (GRVI; Table 2.A2). Correlations were strongest during Flower to Open Boll, and most covariates were selected from the second half of the growing season. Overall, the CCCI, NDVI and NDRE were selected more frequently than other vegetation indices in the MLR models, and the NDVI was the only covariate chosen in every MLR model. For the RF models, the OSAVI, NDVI, RE/R and NDRE were selected more frequently.

Of the covariates chosen in the final MLR models for cotton fibre micronaire, the strongest negative relationship was -0.70 (GRVI) and the strongest positive relationship was 0.66 (OSAVI; Table 2.A1). For the RF models, the strongest negative relationship was –0.70 and the strongest positive relationship was 0.76 (both the RE/R; Table 2.A2). As for length, most covariates were selected from the second half of the growing season and correlations were strongest during Flower to Open Boll, and Open Boll to Harvest. The RSI, OSAVI, NDVI, and NDRE were chosen more frequently than other vegetation indices in the cotton fibre

micronaire MLR models, and the OSAVI, NDVI and EVI were chosen more frequently in the RF models.

The strongest negative relationship between cotton fibre yield and covariates chosen in the final MLR and RF models was -0.68 (NDRE) and the strongest positive relationship was 0.84 (EVI; Tables 2.A1 and 2.A2). The NDRE, EVI and NDVI were selected more frequently in the final RF models, and the NDVI, CCCI, and EVI were selected more frequently in the MLR models. Overall, correlations were strongest during Flower to Open Boll.

2.3.4 Model quality and the decision-making framework

In following the decision-making framework (Figure 2.3), the decision to use a regression model, either on its own or with A2PK of the residuals, compared to downscaling module observations using A2PK alone, was based on the RMSE. If the RMSE of the null model (using the field mean) was lower than either the MLR or RF regression model, A2PK or the null model was the chosen downscaling approach. If the RMSE of either regression model (MLR or RF) was lower, the regression model with or without A2PK of the residuals was the chosen approach.

A trend could be modelled for cotton fibre length in seven (N1, N2, N3, N6, N7, N8, S1) out of the total 11 fields, although the difference in RMSE between the trend models and the null model was very small (Table 2.5, Table 2.A3). On average, the MLR or RF trend model alone (RMSE 0.0159 – 0.0187) performed better than the trend model + A2PK (RMSE 0.0165 – 0.0350), but again there was little difference in RMSEs, particularly for RF models with or without A2PK of the residuals (Table 2.6, Table 2.A3). There was a negligible difference in the strength of correlations between cotton fibre length and available covariates across the difference downscaling approaches (Tables 2.A1 and 2.A2).

TABLE 2.5. Mean model quality statistics for northern New South Wales (NNSW) and southern New South Wales (SNSW) fields when downscaling cotton fibre length data for each downscaling approach. All models were validated against module-resolution cotton fibre length observations (ModRes). * denotes if null or trend model is better. In the cases where either the MLR or RF model was better than the null model, **bold** indicates if the chosen model alone or model + A2PK is better. In cases when the null model was better than either regression model, **bold** indicates if the null model or A2PK is better

Model	Validation support	NNSW		SNSW		Overall	
		RMSE	LCCC	RMSE	LCCC	RMSE	LCCC
Null	ModRes	0.0169	–	0.0156	–	0.0165	–
A2PK	ModRes	9.3074	0.0000	4.7479	0.0001	8.0639	0.000
MLR	ModRes	0.0172	0.2346	0.0230	0.1791	0.0187	0.2195
MLR + A2PK	ModRes	0.0392	0.3088	0.0240	0.1773	0.0350	0.2729
RF	ModRes	**0.0160***	0.2482	0.0156*	0.2664	**0.0159***	0.2532
RF + A2PK	ModRes	0.0169	0.2299	**0.0154**	0.2720	0.0165	0.2414

TABLE 2.6. Mean model quality statistics for northern New South Wales (NNSW) and southern New South Wales (SNSW) fields when downscaling cotton fibre micronaire data for each downscaling approach. All models were validated against module-resolution cotton fibre micronaire observations (ModRes). **Bold** indicates the best model. * denotes if null or trend model is better. In the cases where either the MLR or RF model was better than the null model, **bold** indicates if the chosen model alone or model + A2PK is better. In cases when the null model was better than either regression model, **bold** indicates if the null model or A2PK is better.

Model	Validation support	NNSW		SNSW		Overall	
		RMSE	LCCC	RMSE	LCCC	RMSE	LCCC
Null	ModRes	0.1208	–	0.1043	–	0.1163	–
A2PK	ModRes	5.9811	0.0000	2.3741	0.0000	4.9973	0.0000
MLR	ModRes	0.1271	0.3519	0.4206	0.0438	0.2072	0.2679
MLR + A2PK	ModRes	0.1399	0.3265	0.4401	0.0635	0.2218	0.2548
RF	ModRes	**0.1035***	0.4338	**0.1030***	0.2751	**0.1034***	0.3905
RF + A2PK	ModRes	0.1091	0.4089	0.1426	0.1086	0.1182	0.3270

The quality of the final downscaling approach chosen in each field was fair-to-moderate when validated against module-resolution cotton fibre length observations (LCCC 0.2070 – 0.5368, Table 2.A3). In all cases where the null model was chosen over the trend model (N4, N5, S2, S3), the null model performed better than the A2PK approach (Table 2.5).

A trend model for cotton fibre micronaire could be built for all but one field (S3). In most cases, either the RF or MLR trend model alone performed better than the trend model + A2PK, based on the RMSE (Table 2.6, Table 2.A3). This indicates that all the spatial structure in the cotton fibre micronaire data was captured by the trend model in most cases. In all cases where the trend model alone performed better than the trend model + A2PK, correlations between available covariates and cotton fibre micronaire were higher. The quality of the final trend model chosen in each field was fair-to-good when validated against module-resolution cotton fibre micronaire observations (LCCC 0.2283 – 0.6879, Table 2.A3). For field S3, a trend model to downscale cotton fibre micronaire could not be built using remotely sensed vegetation indices, and A2PK performed worse (RMSE 2.1597) than the null model (RMSE 0.1128; Table 2.A3). Correlations with available covariates were lower in this field compared to others in SNSW where a RF model was chosen.

When predicting cotton fibre yield, on average the RF + A2PK downscaling approach performed best overall when validated at the module-resolution for both SNSW (LCCC = 0.7880) and NNSW (LCCC = 0.7332) (Table 2.7, Table 2.A3). The chosen downscaling approach for individual fields across SNSW and NSNW was varied, but always involved some level of regression; either the RF (fields N1, N2, N5, N8), RF + A2PK (fields N7, S1, S2, S3), or MLR + A2PK (fields N3, N4, N6) approaches (Table 2.A3). Correlations between available covariates and cotton fibre yield were generally stronger in fields where the regression model alone was the chosen downscaling approach. However, there was no difference in the strength of correlations between the RF model on its own and the MLR + A2PK for NNSW.

TABLE 2.7. Mean model quality statistics for northern New South Wales (NNSW) and southern New South Wales (SNSW) fields when downscaling cotton yield data for each downscaling approach. All models were validated against module-resolution yield observations (ModRes) and against fineresolution yield observations (FineRes). The Null model for FineRes yield data was validated by comparing FineRes yield observations against the mean yield value for each module (A) and the mean value for each field (B). * denotes if null or trend model is better. In the cases where either the MLR or RF model was better than the null model, bold (for the ModRes validation support) or *italics* (for the FineRes validation support) indicates if the chosen model alone or model + A2PK is better. In cases when the null model was better than either regression model, **bold** (for the ModRes validation support) or *italics* (for the FineRes validation support) indicates if the null model or A2PK is better.

Model	Validation support	NNSW		SNSW		Overall	
		RMSE	LCCC	RMSE	LCCC	RMSE	LCCC
Null	ModRes	0.1403	–	0.3245	–	0.1905	–
	FineRes A	*0.1630**	–	0.4013*	–	0.2280*	–
	FineRes B	0.49998	–	0.49997	–	0.49998	–
A2PK	ModRes	0.0773	0.7956	0.1394	0.8946	0.0942	0.8226
	FineRes	0.2350	0.4450	*0.3667*	0.6033	0.2709	0.4882
MLR	ModRes	0.1327	0.5421	0.2409	0.6386	0.1622	0.5684
	FineRes	0.6223	0.2587	0.9411	0.1980	0.7092	0.2421
MLR + A2PK	ModRes	0.1055	0.7121	0.2236	0.6856	0.1377	0.7048
	FineRes	0.6187	0.3148	0.9364	0.2099	0.7053	0.2862
RF	ModRes	0.0894*	0.6640	0.2000*	0.7150	0.1196*	0.6779
	FineRes	0.3326*	0.2274	0.4056*	0.3201	0.3525*	0.2527
RF + A2PK	ModRes	**0.0835**	0.7126	**0.1654**	0.7880	**0.1059**	0.7332
	FineRes	0.3047	0.3801	*0.3689*	0.4907	*0.3222*	0.4103

In the absence of fine-resolution cotton fibre quality data for length and micronaire, cotton yield data was used as a surrogate to assess how well downscaled predictions of areal data perform at a fine-resolution and reflect within-module variability. For fine-resolution cotton fibre yield, two null models were compared; the module aggregated (mean) yield value was used for prediction (FineRes A), and the field mean value of fine-resolution yield data was

used for prediction (FineRes B). For FineRes A, no trend model could be built for eight (N1, N2, N3, N4, N5, N7, N8, S2) out of the 11 total fields. Of these, A2PK performed better than the null for two NNSW fields (N7, N8; Table 2.A3). The trend model performed better than the null model in one NNSW field (N6) and two (S1, S3) out of the three SNSW fields (Table 2.A3), and the RF + A2PK model performed best in these cases (LCCC = 0.3520 – 0.7449; Table 2.A3). For FineRes B, the trend model always performed better than the null, and on average the RF + A2PK downscaling approach performed best overall when validated at a fine-resolution (LCCC = 0.3222; Table 2.7).

An example of observed, predicted, and residual values for NNSW field N1 at a fine-resolution, and at the module-resolution, are presented for cotton fibre yield (Figure 2.11).

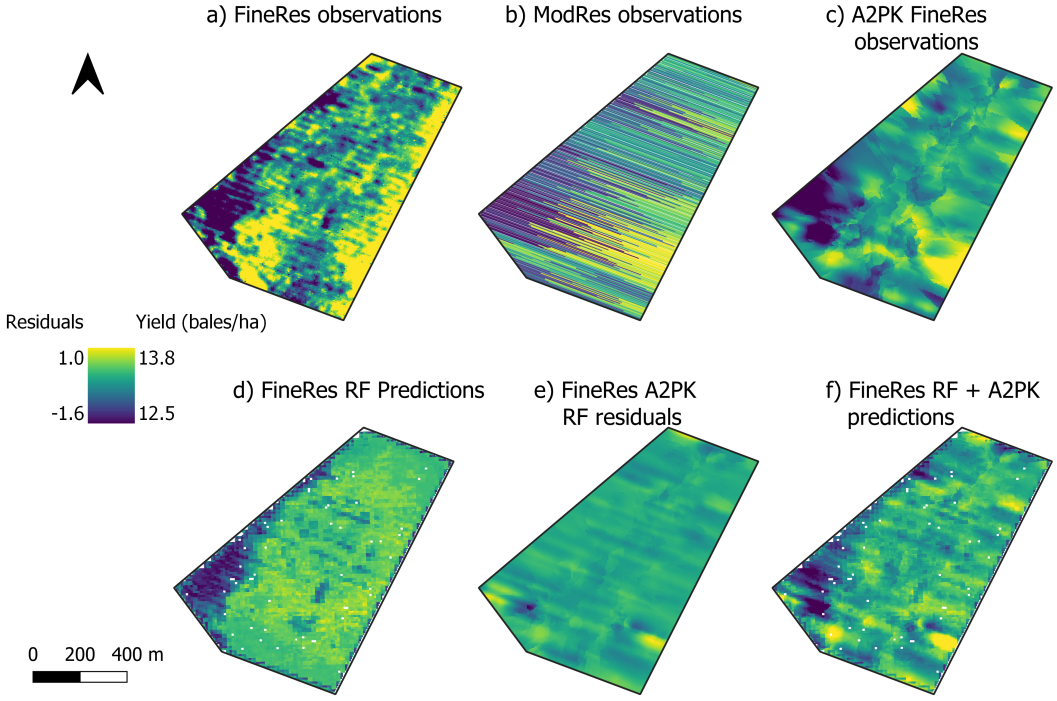

FIGURE 2.11. Maps of cotton fibre yield values are presented for NNSW field N1 for a) observed values at a fine-resolution (FineRes), b) observed values aggregated (mean) within each module (ModRes), c) FineRes area-to-point kriging (A2PK) downscaled predictions, d) FineRes Random Forest (RF) model predictions, e) FineRes A2PK of the RF model residuals, and f) FineRes RF model predictions with A2PK of the residuals (RF + A2PK).

2.4 Discussion

The downscaling approach presented in this study is flexible and can be applied to many different manually harvested horticultural (e.g. citrus) or broadacre (e.g. sugarcane) crops without yield monitors and/or with areal averaged yield or quality data, including cotton fibre quality as shown in this study. The approach presented in this study considers publicly available, remotely sensed covariates to predict spatial variability in regression models, and accounts for additional variation by kriging the MLR or RF residuals with A2PK to produce a final map of downscaled predictions. If these spatial trends cannot be estimated or variability is high, it may be more appropriate to downscale the areal data using A2PK alone, or not try to downscale the data at all.

2.4.1 Downscaling cotton yield and fibre quality data

Firstly, this study aimed to assess the ability of remote sensing imagery to describe variability in cotton fibre yield and quality variables. Correlations with remote sensing covariates were stronger for cotton fibre yield than for cotton fibre micronaire, and much stronger compared to those for cotton fibre length. An understanding of how well remote sensing imagery represents variability in cotton fibre yield and quality is an important first step in implementing this downscaling approach.

To test the downscaling approach, fine-resolution yield data was used as a surrogate in the absence of fine-resolution cotton fibre quality data for length and micronaire. There was a dramatic improvement in the downscaling performance when fine-resolution predictions were validated against module-averages (FineRes A) compared to the field resolution (FineRes B). While the module-averaged (FineRes A) null model performed better than any trend or A2PK downscaling approach for six out of the 11 total fields, this highlights that there is still value in utilising module-resolution yield (and fibre quality) data, even without downscaling. However, it is not clear if variability in cotton fibre yield is the same as for cotton fibre length or micronaire, and we cannot truly test this downscaling approach on cotton fibre quality data. on cotton fibre quality data.

When validated at the fine-resolution, a trend in cotton fibre yield could be modelled in three out of the total 11 fields, indicating that, in certain situations, RF models built using remote sensing imagery can be used to downscale module-resolution cotton fibre yield data. For these fields, correlations with remote sensing images were relatively high in most cases, particularly for SNSW. From this, we can speculate that this downscaling approach *may* work for cotton fibre micronaire data, and likely not work at all for cotton fibre length, given the weaker correlations with remote sensing imagery compared to those for cotton fibre yield. In future, defining a threshold for the correlation strength may be necessary to determine if and when the downscaling approach should be implemented.

The unique long and narrow size and shape of cotton modules presents a challenge for kriging and interpolation. While A2PK was presented here as an alternative approach for downscaling areal harvest data, the quality of downscaled predictions with A2PK was very poor for cotton fibre length and micronaire and the null model outperformed A2PK when predictions were validated against module-aggregated observations. It is expected that A2PK would perform better on more regularly-shaped areal data, such as fruit bin yields in tree crops (Colaço *et al*. 2020).

While spatial trends in length could be modelled using regression with remote sensing imagery in some fields, this was not the case for all and the difference in RMSE values between the trend and null models was minimal. There was no clear difference in the strength of correlations between cotton fibre length and available covariates in fields where a trend model could or could-not be built. These results are not surprising given the unique narrow and long size and shape of the cotton modules. Further, given that 75% of the variation in fibre length is due to variety selection, with the remaining 25% of the variation attributed to weather, environment, and management factors including N applications (Eaton 1947; Meredith 1986; Meredith 1994), it is not surprising that trend models performed poorest when downscaling cotton fibre length compared to those for cotton fibre micronaire and yield. While remotely sensed vegetation indices were used as a proxy to represent variability, it appears that much of the variation in cotton fibre length could not be captured by remote sensing imagery alone. Additional drivers of variation, such as the availability of N in the

soil (Sluijs 2022), may not have been captured by single crop reflectance images, and future work should consider all available production data to improve predictions and help identify potential drivers of variability within and between fields.

Quality issues with micronaire, namely high micronaire, have been raised by spinners (Gordon *et al*. 2004), and low micronaire is a concern for growers in SNSW. Approximately 51% of the variation in cotton fibre micronaire is attributed to weather (e.g. temperature, cloud cover) and management (e.g. nutrients, growth regulators, irrigation) (Eaton 1947; Meredith 1986; Cathey *et al*. 1988; Bradow *et al*. 2000). Understanding how micronaire varies within fields, and what may be driving this variability, can enable growers to make better management decisions pre- or post-harvest. For example, if multiple seasons of yield and quality data can be collected for the same field, long-term trends can be identified and areas consistently producing high or low yields or fibre quality can be delineated into management zones for selective harvest (Boydell *et al*. 2002), module staging (Sluijs *et al*. 2019), or site-specific agronomic management (Velandia *et al*. 2008).

2.4.2 The downscaling approach and variable selection

The spatial dependence of some cotton fibre quality attributes is understood, and micronaire has been suggested as the most predictable fibre quality variable in the field (Wang *et al*. 2017). Water stress early in the season and during boll filling and maturation can have detrimental effects on lint quality variables such as micronaire (Ballester *et al*. 2019). Micronaire has been shown to correlate with spatial variation in soil properties that influence soil water holding capacity (Ballester *et al*. 2019), including clay content and apparent electrical conductivity (EC_a) (based on an EM induction survey; Ge *et al*. 2008; Wang *et al*. 2017. However, bespoke proximally sensed, fine-resolution soil data like EC_a maps are not available for all fields and farms, and publicly-accessible remote sensing imagery presents a possible alternative to represent within-field variation.

Incorporating fine-scale covariate data into the downscaling approach is considered superior to A2PK alone as it allows for a representation of within-module variability, by including

covariates representing factors affecting crop growth and development such as soil physical and chemical constraints (Pallegedara Dewage *et al*. 2020). In fields where correlations between cotton fibre micronaire, yield, and available covariates were high, trend models alone were the chosen downscaling approach, whereas the addition of A2PK of the residuals was necessary when correlations were weaker.

While high-resolution soil and plant physiology data is not available for many (or all) fields and farms, publicly available satellite imagery can act as a surrogate for this information by capturing a range of soil, plant, and environment interactions within and between fields. While high-resolution unmanned areal vehicle (UAV)-based multispectral imagery and associated vegetation indices have been used to model crop growth and biomass accumulation over a growing season for cotton breeding research applications (Dodge 2023), UAV technology is logistically and computationally prohibitive for data collection across whole fields or farms on a commercial scale, particularly in an Australian context. Satellite-based remote sensing imagery is publicly-available for all fields and farms across the globe, and new and emerging high-resolution (e.g. Planet; Inc. 2019) and hyperspectral (e.g. Prisma; Cogliati *et al*. 2021) satellite-based remote sensing platforms may replace UAV technologies in the future.

There is a trade-off between model performance, complexity, and transferability. Multiple linear regression assumes a linear relationship between the response and predictor variable. If assumptions are not met, MLR models may struggle to capture complex nonlinear patterns, leading to poor predictive performance. Conversely, RF models are non-parametric and can capture non-linear relationships, are more robust to outliers, and naturally account for interactions between predictors. Random Forest models were the most-frequently chosen regression approach overall, predominately on their own without A2PK of the residuals.

While no single vegetation index was found to be used more frequently than any other when predicting cotton yield and fibre quality, images from the second half of the season were generally more strongly correlated and were subsequently used more often in the regression models. Cotton has an indeterminate growth habit, and there is the potential for the plant to compensate for stress-induced boll shedding early in the season. As the season progresses into flowering and boll fill, there is less time available for the plant to compensate for stress

and boll losses, and crop spectral data towards the end of the season may more accurately represent cotton yield and fibre quality at harvest.

It should be emphasised that the goal of this study was not to forecast cotton fibre quality prior to harvest. Rather, this study aimed to assess how well remote sensing imagery can capture variability in cotton fibre length, micronaire, and yield, and then illustrate and test the downscaling approach using fine-resolution cotton fibre yield data. While the timing of particular dates and covariates used in the regression models can provide an insight into how cotton fibre yield and quality can be represented both spatially and temporally, including if particular growth stages are more correlated and better reflect fibre quantity and quality at harvest, we did not use this information to make end-of-season forecasts. Interpretive machine learning approaches have been used to identify the causes of crop yield variability in cotton within fields (Jones *et al*. 2022). Future work should consider additional spatial covariates, including N applied maps and soil surveys, within an interpretive machine learning approach to better understand the factors that may be driving variability in cotton fibre quality, both spatially and temporally.

Building a universal model applicable to every field across both northern and southern NSW was not the objective for this study. Rather, a near-fully automated data collection, modelling and prediction method was used here to build a bespoke model for every field. Publicly-available data platforms such as GEE or Data Harvester (Haan *et al*. 2023; Harianto *et al*. 2023) can provide access to every single remote sensing image taken since the release of Landsat 1 in 1972. In this study, ~25 images for NNSW, and ~70 images for SNSW, were available across the 2021/22 cotton growing season from Sentinel-2. The development of R packages like 'rgee' now mean that remote sensing image extraction and vegetation index calculation can be fully automated within a script and user input is limited to specifying the date range, extraction boundary, and necessary Sentinel-2 bands to calculate each vegetation index, for example. Thus, there is minimal difference in computation time and effort required to extract one single vegetation index compared to eight or more vegetation indices in this case. With the ongoing development of data platforms, coding languages, programs, and

plugins, the potential to build a bespoke model unique to each field using a near-automated process is now being realised and has been demonstrated here.

2.4.3 Model validation

Predictions obtained with A2PK are coherent or mass-preserving, meaning that the average of the point predictions within a given area is equal to that of the known areal mean. While this is desirable when the areal data is assumed to be an errorless observation of the spatial mean (e.g. pixel values of remote sensing imagery; Brus *et al*. 2018), this may not be the case for cotton fibre quality measurements. A round module of cotton is broken down into approximately 4 bales during ginning, and 2 sub-samples weighing a minimum of 200 grams each are taken from each side of the bale. These samples and bale averages for fibre quality variables may not capture the potential variability that can occur within the bales (Bradow *et al*. 2000; Gourlot *et al*. 2005), and there is greater variability in fibre quality variables between bales from round modules compared to larger conventional modules as there is less blending during module building and ginning (Sluijs *et al*. 2015). At this stage, it is not possible to accurately validate these downscaling approaches for cotton fibre quality data, given that no "true" fine-resolution fibre quality data is available. Fine-resolution yield observations were used as a surrogate to assess the downscaling algorithm performance, but further investigation is needed to assess the fine-resolution variability of cotton fibre quality variables within fields.

2.4.4 Future work

2.4.4.1 Interpretive machine learning and forecasting cotton fibre quality

In future, a larger suite of available covariates, including on-farm soil data, nutrient prescription/application maps, and remotely sensed vegetation index growth curve derivatives (e.g. Dodge 2023), for example, could be considered to improve downscaling estimates and model quality. Interpretive machine learning approaches (e.g. Jones *et al*. 2022) should also be utilised to better understand the drivers of variability in cotton fibre quality variables within

fields. If sufficient data is available, there is the potential to apply interpretive machine learning approaches across multiple fields, farms, regions, and seasons, and incorporate additional climatic and management information including the timing of planting and defoliation, crop variety, or temperature, to better understand the drivers of cotton fibre quality variability both spatially and temporally. Further, there may be the potential to utilise this information to forecast cotton fibre quality variables during the season or prior to harvest. Forecasts of cotton fibre quality variables can be used to inform management decisions throughout the season, such as nutrient applications, defoliation, selective harvest, or module staging at the gin for blending to minimise discounts and reduce variability (Sluijs *et al*. 2019).

2.4.4.2 Applications beyond cotton

The introduction of RFID technology in cotton has enabled cotton modules to be labelled and tracked from cotton harvest to the gin, and means that cotton fibre quality can now be mapped within fields. Longchamps *et al*. (2022) highlight the potential for RFID technologies to be utilised for yield mapping in perennial horticultural crops. Given the strong spatial organisation of cultivated plants in rows, pots, panels or bays, there is the ability to georeference these spatial units and associated harvesting units (bins or bags) with the harvested location using RFIDs (e.g. Ampatzidis *et al*. 2009). There is the potential to apply the downscaling approach demonstrated in this study to downscale areal-averaged horticultural yield or quality data to produce a higher resolution map.

A key feature of perennial horticultural crops is that plants or rows of plants can be individualised, resulting in remote sensing-based yield mapping methods in horticulture primarily utilising UAVs and high-resolution imagery to map individual plants or segment individual fruits, for example (Longchamps *et al*. 2022). However, the required accuracy of the yield map will depend on its application. Areal yield maps for several trees or over a block may be sufficient to audit production (e.g. to audit harvested tonnage against delivered tonnage) or for understanding general spatial patterns in yield response. In cases where yield data is use to quantify and validate the effect of management on production, high(er)-resolution, accurate, and precise yield maps are required (Longchamps *et al*. 2022). There is the potential

to utilise this A2PK and/or regression approach for yield (or quality) mapping in horticultural crops, provided that the resolution of the final yield map can be captured by remote-sensing platforms (e.g. Sentinel-2).

2.5 Conclusion

Fine-resolution yield and quality data is important for understanding variability and for informing management decisions including PA practices. However, areal observations at the block, sub-, or whole-field scale is often the only available yield or quality data for many commodities including for manually-harvested horticultural crops, sugarcane, and cotton. In this study, a downscaling approach was demonstrated where spatial trends in cotton fibre yield and quality were estimated using regression with fine-resolution remotely sensed vegetation indices. If spatial trends could be estimated, A2PK was used to downscale the residuals from the trend model. However, in the absence of a spatial trend model, A2PK was compared with the null model (field mean) to determine the best method for downscaling areal observations from the block to the point support. Correlations with remote sensing covariates were stronger for cotton fibre yield than for cotton fibre micronaire, and much stronger compared to those for cotton fibre length. Broadly, covariates from the second half of the growing season were more strongly correlated and were used more frequently within the regression models. Spatial trends in cotton fibre yield and micronaire could be estimated with good model quality using regression with remote sensing covariates in almost all fields. Conversely, model quality was poorer for cotton fibre length and there was only a small difference in model performance between the null and trend models. The addition of A2PK of the residuals improved model performance for downscaling areal cotton fibre yield observations, but not for cotton fibre micronaire. Fine-resolution cotton fibre yield data was used as a surrogate to test the downscaling approach in the absence of fine-resolution length or micronaire data. It was not surprising that model performance was poorer at the fine-resolution compared to the module-resolution, but it is expected that the downscaling approach will perform better on more regularly-shaped areal data, such as fruit bin yields in tree crops. This approach enables the creation of high-resolution raster maps of variables of

interest with a much finer spatial resolution compared to the areal observations, and can be applied for any areal averaged crop production data in a range of agricultural and horticultural industries (e.g. sugarcane, apples, citrus). In future, there are opportunities to incorporate a greater diversity of covariates in interpretive machine learning approaches to investigate the drivers of variability in cotton fibre yield and quality across larger spatial and temporal scales, and for in-season forecasts. Further, there may be the opportunity to apply a similar approach for yield or quality mapping in manually harvested fruit and vegetable crops.

2.6 Appendices

TABLE 2.A1. Summary statistics of Spearman-Rank Correlation Coefficient values between cotton fibre length, micronaire, and yield, and remotely-sensed vegetation indices chosen in the final Multiple linear regression models for field fields in NNSW, SNSW, and for all fields.

			RSI	RE/R	OSAVI	NDVI	NDRE	GRVI	EVI	CCCI
Yield	NNSW	Min	-0.21	-0.62	-0.16	-0.16	-0.19	-0.34	-0.17	-0.24
		Max	0.75	0.75	0.81	0.80	0.79	0.34	0.84	0.78
		Mean	0.32	0.29	0.31	0.35	0.41	0.20	0.34	0.37
	SNSW	Min	-0.67	-0.61	-0.59	-0.64	-0.68	-0.59	-0.65	-0.64
		Max	0.70	0.77	0.67	0.70	0.73	0.76	0.65	0.53
		Mean	0.59	0.60	0.48	0.63	0.64	0.62	0.60	0.49
	All	Min	-0.67	-0.62	-0.59	-0.64	-0.68	-0.59	-0.64	-0.64
		Max	0.75	0.77	0.81	0.80	0.79	0.76	0.78	0.78
		Mean	0.41	0.40	0.36	0.40	0.48	0.34	0.40	0.49
Length	NNSW	Min	-0.24	-0.23	-0.27	-0.25	-0.10	-0.27	-0.27	-0.33
		Max	0.35	0.33	0.25	0.26	0.36	0.32	0.24	0.21
		Mean	0.22	0.19	0.18	0.20	0.19	0.19	0.20	0.19
	SNSW	Min	-0.42	-0.48	-0.43	-0.43	-0.44	0.26	-0.23	-0.42
		Max	0.41	0.57	0.41	0.42	0.46	0.26	0.44	0.44
		Mean	0.31	0.38	0.35	0.34	0.34	0.26	0.27	0.37
	All	Min	-0.42	-0.48	-0.43	-0.43	-0.44	-0.27	-0.27	-0.42
		Max	0.41	0.57	0.41	0.42	0.46	0.32	0.44	0.44
		Mean	0.25	0.27	0.25	0.26	0.22	0.20	0.23	0.26
Micronaire	NNSW	Min	-0.55	-0.58	-0.67	-0.55	-0.62	-0.70	-0.60	-0.58
		Max	0.26	0.49	0.66	-0.18	0.49	0.29	0.64	0.54
		Mean	0.35	0.38	0.37	0.36	0.37	0.30	0.41	0.28
	SNSW	Min	-0.54	-0.33	-0.53	-0.55	-0.55	-0.47	-0.51	-0.54
		Max	-0.23	0.52	-0.22	-0.24	-0.24	0.27	0.27	-0.25
		Mean	0.41	0.43	0.32	0.41	0.41	0.34	0.33	0.38
	All	Min	-0.55	-0.58	-0.67	-0.55	-0.62	-0.70	-0.60	-0.58
		Max	0.26	0.52	0.66	-0.18	0.49	0.29	0.64	0.54
		Mean	0.37	0.38	0.36	0.38	0.38	0.31	0.39	0.31

TABLE 2.A2. Summary statistics of Spearman-Rank Correlation Coefficient values between cotton fibre length, micronaire, and yield, and remotely-sensed vegetation indices chosen in the final Random Forest models for field fields in NNSW, SNSW, and for All fields.

			RSI	RE/R	OSAVI	NDVI	NDRE	GRVI	EVI	CCCI
Yield	NNSW	Min	-0.21	-0.62	-0.16	-0.16	-0.19	-0.24	-0.17	-0.23
		Max	0.61	0.32	0.81	0.55	0.62	0.34	0.84	0.78
		Mean	0.32	0.29	0.42	0.27	0.34	0.20	0.43	0.35
	SNSW	Min	-0.65	-0.61	-0.59	-0.63	-0.68	-0.59	-0.65	-0.56
		Max	0.72	0.80	0.66	0.73	0.73	0.76	0.60	0.42
		Mean	0.68	0.72	0.61	0.64	0.66	0.63	0.38	0.49
	All	Min	-0.65	-0.62	-0.59	-0.63	-0.68	-0.59	-0.65	-0.56
		Max	0.72	0.80	0.81	0.73	0.73	0.76	0.84	0.78
		Mean	0.38	0.39	0.45	0.36	0.41	0.37	0.59	0.36
Length	NNSW	Min	-0.24	-0.24	-0.27	-0.24	-0.24	-0.27	-0.22	-0.34
		Max	0.35	0.35	0.32	0.32	0.36	0.80	0.17	0.32
		Mean	0.22	0.24	0.19	0.20	0.21	0.28	0.17	0.25
	SNSW	Min	-0.27	-0.34	-0.22	-0.21	0.37	0.46	-0.47	0.29
		Max	0.41	0.57	0.41	0.42	0.46	0.46	0.39	0.43
		Mean	0.37	0.36	0.34	0.37	0.43	0.49	0.32	0.36
	All	Min	-0.27	-0.34	-0.27	-0.20	-0.24	-0.27	-0.47	-0.34
		Max	0.41	0.57	0.41	0.42	0.46	0.80	0.39	0.43
		Mean	0.26	0.27	0.22	0.24	0.25	0.30	0.21	0.27
Micronaire	NNSW	Min	-0.69	-0.7	-0.67	-0.68	-0.67	-0.57	-0.56	-0.46
		Max	0.52	0.76	0.27	0.50	0.49	0.29	0.22	0.54
		Mean	0.40	0.47	0.42	0.43	0.42	0.40	0.42	0.34
	SNSW	Min	-0.54	0.32	-0.55	-0.54	-0.56	-0.55	-0.54	-0.54
		Max	-0.54	0.32	-0.22	-0.22	0.23	0.28	-0.21	-0.25
		Mean	0.54	0.32	0.44	0.38	0.40	0.40	0.40	0.34
	All	Min	-0.69	-0.72	-0.67	-0.68	-0.67	-0.57	-0.56	-0.54
		Max	0.52	0.76	0.27	0.50	0.49	0.29	0.22	0.54
		Mean	0.41	0.46	0.42	0.41	0.41	0.40	0.42	0.34

TABLE 2.A3. Model quality statistics for downscaling cotton fibre length, micronaire, and yield data for each downscaling approach. All models were validated against module-resolution observations (ModRes), and yield predictions were also validated against fine-resolution (5 m) yield observations (FineRes). The Null model for FineRes yield data was validated by comparing FineRes yield observations against the mean yield value for each module (A) and the mean value for each field (B). In the cases where either the MLR or RF model was better than the null model, **bold** (for the ModRes validation support) or *italics* (for the FineRes validation support) indicates if the chosen model alone or model + A2PK is better. In cases when the null model was better than either regression model, **bold** (for the ModRes validation support) or *italics* (for the FineRes validation support) indicates if the null model or A2PK is better.

				NNSW								WA		
	Model	Validation support		N1	N2	N3	N4	N5	N6	N7	N8	S1	S2	S3
Length	Null model	ModRes	RMSE	0.0170	0.0199	0.0199	**0.0158***	**0.0151***	0.0163	0.0150	0.0161	0.0162	**0.0157***	**0.0149***
	A2PK	ModRes	RMSE	12.0602	11.9083	7.5981	8.5591	8.1523	8.4204	8.8090	8.9517	4.2961	5.2904	4.6571
			LCCC	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0001	0.0001
	MLR	ModRes	RMSE	0.0168	0.0176	**0.0167***	0.0170	0.0223	0.0178	**0.0143***	0.0148	0.0313	0.0199	0.0177
			LCCC	0.2768	0.3925	0.4835	0.0983	0.0395	0.0536	0.2005	0.3321	0.2635	0.2330	0.0407
	MLR + A2PK	ModRes	RMSE	0.0177	0.1880	**0.0161**	0.0168	0.0235	0.0174	0.0150	0.0187	0.0313	0.0232	0.0175
			LCCC	0.3681	0.4909	0.5368	0.2602	0.0809	0.1678	0.2864	0.2791	0.2695	0.1955	0.0669
	RF	ModRes	RMSE	0.0164*	0.0167*	0.0179	0.0166	0.0154	**0.0155***	0.0148	**0.0147***	**0.0139***	0.0173	0.0157
			LCCC	0.2933	0.4313	0.3102	0.1053	0.1172	0.2070	0.1660	0.3555	0.4650	0.2343	0.0999
	RF + A2PK	ModRes	RMSE	**0.0163**	0.0195	0.0189	0.0162	0.0149	0.0155	0.0157	0.0181	0.0143	0.0164	0.0156
			LCCC	0.3235	0.4038	0.1446	0.1972	0.1553	0.1989	0.1925	0.2232	0.5777	0.1297	0.1087
Micronaire	Null model	ModRes	RMSE	0.1074	0.1346	0.0979	0.1715	0.1158	0.1005	0.1218	0.1168	0.1079	0.0923	**0.1128***
	A2PK	ModRes	RMSE	8.6487	8.4468	4.3554	5.4623	4.8321	5.1542	5.3327	5.6164	2.1608	2.8017	2.1597
			LCCC	0.0000	0.0000	0.0000	0.0001	0.0000	0.0000	0.0000	0.0000	0.0000	0.0001	0.0000
	MLR	ModRes	RMSE	0.1149	**0.1262***	0.1114	0.1464	0.2069	**0.0967***	**0.0949***	0.1195	0.9930	0.1353	0.1336
			LCCC	0.2411	0.3326	0.1071	0.5920	0.0505	0.5215	0.5784	0.3923	0.0011	-0.0115	0.1417
	MLR + A2PK	ModRes	RMSE	0.1280	0.1365	0.1196	0.1611	0.2224	0.1222	0.1082	0.1214	1.0410	0.1402	0.1391
			LCCC	0.2406	0.3381	0.0380	0.5907	0.0506	0.3719	0.5463	0.4360	-0.0022	-0.0131	0.2059

Table 2.A3 continued from previous page

	RF	ModRes	RMSE	0.1046*	0.1264	0.0929*	**0.1213***	0.0899*	**0.0989**	0.0999	0.0942*	**0.1069***	**0.0740***	0.1281
			LCCC	0.2436	0.2523	0.2429	0.6879	0.5902	0.3948	0.5176	0.5408	0.3182	0.5147	-0.0077
	RF + A2PK	ModRes	RMSE	**0.1021**	0.1248	**0.0887**	0.1439	0.0966	0.0955	0.1123	0.1085	0.2244	0.0890	0.1145
			LCCC	0.2283	0.3197	0.2803	0.5736	0.5836	0.4385	0.4792	0.3679	-0.0469	0.1284	0.2442
Yield	Null model	ModRes	RMSE	0.1930	0.1869	0.1890	0.0953	0.0601	0.2113	0.0901	0.0969	0.3084	0.2337	0.4313
		FineRes A	RMSE	*0.1473**	*0.2099**	*0.2465**	*0.1004**	*0.0663**	0.3533	0.0779*	0.1023*	0.3854	*0.2093**	0.6092
		FineRes B	RMSE	0.49998	0.49999	0.49997	0.49999	0.49999	0.49998	0.49999	0.49999	0.49998	0.49996	0.49997
	A2PK	ModRes	RMSE	0.1291	0.0734	0.0848	0.0577	0.0432	0.1337	0.0464	0.0500	0.1460	0.0736	0.1986
			LCCC	0.7727	0.9125	0.8874	0.7516	0.6427	0.7342	0.8306	0.8327	0.8572	0.9543	0.8724
		FineRes	RMSE	0.5000	0.3900	0.3000	0.2500	0.2100	0.1300	*0.0500*	*0.0500*	0.3800	0.4200	0.3000
			LCCC	0.4100	0.5600	0.6000	0.4400	0.2400	0.5400	0.3200	0.4500	0.5900	0.4500	0.7700
	MLR	ModRes	RMSE	0.3903	0.1327	0.0804*	0.0670*	0.0850	0.1476*	0.0886	0.0698	0.3108	0.1485	0.2634
			LCCC	0.2585	0.6957	0.8984	0.6880	0.1776	0.6775	0.2839	0.6570	0.3306	0.7838	0.8013
		FineRes	RMSE	1.1748	0.5035	0.4233	1.3232	0.2512	0.6545	0.3725	0.2752	1.1326	0.9601	0.7305
			LCCC	0.1516	0.4573	0.4682	0.0081	0.2484	0.1970	0.1436	0.3951	0.0524	0.1952	0.3464
	MLR + A2PK	ModRes	RMSE	0.3320	0.0797	**0.0648**	**0.0515**	0.0799	**0.1221**	0.0599	0.0543	0.3079	0.1111	0.2517
			LCCC	0.3525	0.9051	0.9379	0.8300	0.3200	0.8028	0.7378	0.8104	0.3597	0.8745	0.8226
		FineRes	RMSE	1.1732	0.4822	0.4167	1.3709	0.2433	0.6436	0.3580	0.2615	1.1304	0.9563	0.7224
			LCCC	0.1613	0.5346	0.4913	0.0072	0.3178	0.2484	0.2636	0.4940	0.0584	0.2020	0.3692
	RF	ModRes	RMSE	**0.1115***	**0.0994***	0.0814	0.0812	**0.0404***	0.1656	0.0713*	**0.0645***	0.2205*	0.1256*	0.2538*
			LCCC	0.8081	0.8303	0.8938	0.4467	0.6627	0.5166	0.4988	0.6553	0.5802	0.8076	0.7573
		FineRes	RMSE	0.4736*	0.4312*	0.2867*	0.3852*	0.2176*	0.3184*	0.2831*	0.2650*	0.3850*	0.4468*	0.3849*
			LCCC	0.3120	0.3675	0.5509	-0.0010	0.0693	0.1944	0.1420	0.1841	0.2371	0.2566	0.4666
	RF + A2PK	ModRes	RMSE	0.1328	0.1140	0.0635	0.0744	0.0430	0.1145	0.0582	0.0678	0.2102	0.1080	0.1781
			LCCC	0.6365	0.7726	0.9326	0.5501	0.6272	0.7945	0.7067	0.6804	0.6077	0.8670	0.8894
		FineRes	RMSE	0.4688	0.3901	0.2674	0.2614	0.2099	0.2831	0.2642	0.2930	0.3658	0.4320	0.3088
			LCCC	0.3837	0.5461	0.6563	0.0059	0.1962	0.5051	0.3169	0.4307	0.3520	0.3752	0.7449

CHAPTER 3

Predicting within-field grain protein content at scale using agronomic and remote sensing variables, and machine learning

---

Abstract

Grain protein content (GPC) is a key determinant of the prices that grain growers receive, but there is considerable variability within and between fields, farms, and seasons. Despite growing interest in measuring and mapping within-field GPC variability, the uptake of grain protein sensors has been slow, resulting in considerable knowledge gaps. Building a predictive model to map GPC in areas of a farm without a GPC sensor can provide growers with valuable insights for better management decisions. This chapter presents a data-driven, machine learning (random forest) approach to predict GPC and yield within agricultural fields using 63 paired yield and protein maps collected over four seasons (2020 – 2023) in Western Australia and northern New South Wales, Australia. Model performance for yield and GPC predictions using different combinations of yield, on-farm agronomic (e.g. sowing and harvest dates, cropping history, variety) and publicly-available (e.g. digital elevation model, radiometric surveys, remotely-sensed satellite imagery) spatial data layers were tested using two validation approaches: leave one Field-Year out cross validation (LOFYOCV) and two-fold cross validation (2FCV) at either a fine-resolution (30 m) or across management classes. The 2FCV method, which simulates interpolating GPC within fields to fill-in unsampled areas, outperformed LOFYOCV, which tested extrapolation across unsampled fields. Combining yield, agronomic, and publicly-available data layers produced the best quality predictions of GPC. Providing growers with GPC maps can inform management decisions to optimise

both yield and quality, leading to more profitable and environmentally sustainable production systems.

3.1 Introduction

Grain protein content (GPC) is an important measure of wheat quality that directly influences both the utilisation and commercial value of wheat (Xu *et al*. 2020). In Australia, wheat varieties are classed into nationally standardised grades primarily based on protein content and the grains intended end-use. For example, Australian Hard (AH) wheat is a high-protein wheat with a base protein content of 11.5 – 13% that is typically used for bread making. Each grade has a baseline price, with premiums or discounts applied if the GPC at receival is above or below the threshold. A similar premium/discount system applies in the United States of America (USA), although premiums and discounts vary regionally and with market conditions. In Europe, while GPC is important, other quality factors, including gluten strength and test weight, also influence grading. National grading systems differ between countries but in general, premiums are awarded to higher protein wheat and discounts are applied to lower protein wheat. For Australian grain growers, optimising GPC is key to attaining premium prices and improving variability.

Wheat GPC is determined by complex interactions between genetic, environmental, and management factors, including crop variety, rainfall, and soil physical, chemical, and nutritional characteristics. Within fields, nitrogen (N) in the soil and applied as fertiliser, and soil moisture availability throughout the growing season, are the two key drivers of GPC variability (Whelan *et al*. 2013). Knowing the magnitude and extent of variability in GPC within fields, across a farm, and over multiple seasons, can be useful for managing the quality of marketed grain, improving N nutrition decisions, and evaluating management outcomes (Whelan *et al*. 2013; Bastos *et al*. 2021).

Laboratory (or benchtop) near-infrared (NIR) spectroscopy sensors have been used for more than 20 years by grain delivery points and silos across Australia to evaluate the protein content of grain being delivered (Taylor *et al*. 2007b). However, only one measurement is taken per

received load of grain and spatial variability within fields is not captured using benchtop sensors. Since the late 1990s and early 2000s, a range of combine harvester-mounted NIR spectroscopy sensors for GPC have been developed and commercialised to measure wheat GPC on-the-go at harvest (e.g. Zeltex AccuHarvest® Analyser, https://zeltex.com/; CropScan 3300H, https://cropscanag.com/).

Most recently, John Deere commercially released the HarvestLab 3000™ Grain Sensing System in Australia, which is an NIR spectroscopy-based sensor mounted onboard a combine harvester for the measurement of continuous grain flow. Onboard the harvester, the auger extracts a threshed grain sample, passes it along and presents it to the NIR constituent sensor where the sample is then illuminated by a broadband light source. Specific wavelengths of this light are absorbed by the grain sample, depending on the molecular composition of the grain, and a spectrometer inside the sensor measures the wavelengths reflected by the sample which is then used to estimate properties such as protein, oil, or moisture content. This prediction is performed using an inbuilt calibration function which was trained using the spectra from a large number of samples across a range of conditions, and the measurement accuracy of moisture and protein estimation with the sensor is ~$\pm$ 0.6% (90% confidence interval). Approximately one measurement is performed per second (Schade *et al*. 2023).

Despite growing interest in measuring and mapping GPC, the uptake of grain protein sensors has been slow and comprehensive GPC datasets are not available for every field, farm, or season. This is resulting in considerable knowledge gaps. For example, GPC maps may not be available for a field or for an entire farm if the grower lacks a harvester fitted with a grain protein sensor, or if the sensor was purchased recently and there is no historical GPC records. Additionally, data may only be collected for part of a field, particularly in large fields where multiple combine harvesters are operating (e.g. in Australia or the USA). If one harvester is equipped with both a grain protein sensor and a yield monitor, while others only have yield monitors, GPC data will only cover part of the field. This lack of data can hinder understanding GPC variability, but it may be possible to fill these knowledge gaps by extrapolating or interpolating data from other layers to unsampled areas.

Farmers and the agricultural industry are now collecting an unprecedented volume of information, including yield data, soil surveys, and management details such as the timing of sowing and harvest. There is also an enormous amount of public data that is free to access, such as remote-sensing imagery and terrain attribute data (e.g. digital elevation models, DEMs). These data layers can help explain or proxy the factors influencing GPC variability. For example, vegetation indices like the normalised difference red edge (NDRE) are useful in estimating chlorophyll content and may describe variability in N sufficiency (Wang *et al*. 2019). While collecting on-farm data generally yields more accurate results, collecting bespoke surveys (e.g. electromagnetic [EM] induction) or soil sampling is often an expensive and unfeasible task for most farmers (Filippi *et al*. 2020a). However, existing on-farm data describing agronomic decisions, such as the timing of sowing or variety selection, considerably impacts GPC and can be used for predictions (Fettell *et al*. 2012). Hybrid approaches combining on-farm and publicly-available data layers have been explored for digital soil mapping (e.g. Filippi *et al*. 2020a) and crop yield forecasting (Filippi *et al*. 2019b). This growing abundance of information offers an opportunity to model and map GPC in the absence of on-the-go protein sensors, and to compare the applications of different on-farm and publicly-available data layers, to fill knowledge-gaps and map GPC in unsampled areas of a field or farm.

On-the-go yield monitors have been used extensively in research and commercial applications to map yield variability within fields. Yield monitor data has been used alongside a combination of other on-farm and/or publicly-available data layers, including proximal and remote sensing technologies, to build models to predict (Habibi *et al*. 2024), forecast (Filippi *et al*. 2019b), and understand the drivers of cereal grain yields (Jones *et al*. 2022) at a range of different spatial and temporal scales. Comparatively, there has been less research utilising combine harvester-mounted, on-the-go grain protein sensors (Bastos *et al*. 2021).

Much of the research that has utilised high-resolution grain protein data derived from combine-harvester mounted protein sensors has been limited to a small number of fields, often for a single season (e.g. Morari *et al*. 2018; Wang *et al*. 2019). There has been some work on modelling and mapping GPC at the regional or county level for logistics and planning (e.g.

Wang *et al*. 2014), but few regional studied utilise GPC data derived from on-the-go sensors and most large-scale modelling and mapping exercises use limited grain protein data sampled by hand from experimental plot trials or sites (e.g. Song *et al*. 2023). The slow commercial uptake of grain protein sensors is one of the key limiting factors to much of the existing research. In the absence of on-the-go grain protein sensors, remote-sensing imagery has been explored to (directly or indirectly) predict GPC within fields (e.g. Stoy *et al*. 2022; Hernandez *et al*. 2023), for disease detection and the impacts on protein variability (Nutter *et al*. 2002), or for mapping N uptake (e.g. Wang *et al*. 2019). Again however, these analyses are often limited to data from a field for one season which does not capture sufficient spatio-temporal variation that occurs between different fields, farms, regions, and seasons, or are limited to point-based sampling approaches which have a limited ability to assess within-field variability at a high-spatial resolution. Whelan *et al*. (2009) conducted one of the largest studies using high-resolution GPC maps collected within-fields using a combine harvester-mounted sensing system. In that study, the authors measured, mapped, and compared the site-specific variation in wheat GPC and yield within-fields using 27 paired yield and protein maps collected over 3 seasons (2003 – 2005) in Australia.

The growing abundance and diversity of on-farm and publicly-available data layers presents an opportunity to fill knowledge-gaps and map wheat GPC in areas of a field or farm without a grain protein sensor. These data layers can be used to build a spatial model to interpolate GPC and fill-in gaps within-fields left by headers without a grain protein sensor, extrapolate GPC across entire unsampled fields, or hindcast to map GPC for an entire farm prior to the purchase of a protein sensor to build a historical record of GPC.

In this chapter, a data-driven, machine learning approach is presented which utilises a combination of yield, agronomic, and publicly-available spatial data layers to model and predict GPC within fields and fill data gaps across farms. Agronomic and publicly-available spatial data layers are also used to model and predict yield within fields as a benchmark to compare GPC predictions against. This will be useful to better understand the predictability of GPC against yield which has been the subject of much more research. The aim of this research was

to create models to predict GPC and yield in areas of a farm without a grain protein or yield sensor, by:

(1) Using readily available yield, agronomic, and/or publicly-available data;
(2) Validating model performance using two different validation approaches; and
(3) Predicting at different spatial resolutions (30 m grid and management classes).

By building a predictive model to predict GPC in areas of a farm without a grain protein sensor, growers and advisors can be equipped with the necessary information and tools to make better management decisions for more profitable and environmentally sustainable production systems. In the face of increasing input costs and variability, maps of GPC can be used to better understand the costs of variable GPC, understand N dynamics and agronomy, and assess the implications of fertiliser decisions prior to or during the growing season. Improving this understanding can have positive outcomes for on-farm economics, production efficiencies, and environmental sustainability.

3.2 Methods

This paper presents a data-driven, machine learning (random forest) approach to predict GPC and yield within agricultural fields using 63 paired yield and protein maps collected over four seasons (2020 – 2023) in Western Australia (WA) and northern New South Wales (NNSW), Australia. The methods are summarised below and are explained in further detail in Sections 3.2.1 – 3.2.6.

Briefly, five experiments (Section 3.2.3) using different combinations of yield, agronomic (variety, sowing and harvest dates, cropping history) and/or publicly-available (DEM, radiometric surveys, remotely-sensed vegetation indices, barest earth satellite imagery) data layers were used to build models to predict wheat GPC within fields. Yield was also modelled within fields as a benchmark to compare GPC predictions against. Each experiment was validated using two different validation approaches (Section 3.2.5) at two validation spatial resolutions (Section 3.2.6) to test different protein sensor data availability scenarios

and applications for precision agriculture. In summary, for each experiment for each farm (NNSW or WA), random forest models for GPC and yield were separately calibrated and validated using either a leave one field-year out cross-validation (LOFYOCV) or a two-fold cross-validation (2FCV) approach, and were also validated at a fine-resolution (30 m) or within management classes.

3.2.1 Study area

The study was conducted on two large broadacre, dryland farms in eastern and western Australia (Figure 3.1). The first farm was in NNSW (Figure 3.1) and was comprised of 16 broadacre dryland fields ranging from ~44 – 1287 hectares (ha) in size. Wheat, barley, canola, chickpeas, faba beans, and sorghum are the dominant crops grown in rotation. The NNSW farm has a humid subtropical climate (Beck *et al*. 2018), with warm-to-hot summers and cool winters. The annual average rainfall is 579 mm and is summer dominant (BOM 2024d). The second farm was in WA (Figure 3.1), comprising of 30 broadacre dryland fields from ~38 – 504 ha in size. Wheat, barley, canola, and lupins are the dominant crops grown. The WA farm has a Mediterranean climate (Beck *et al*. 2018) with cool, wet winters and hot, dry summers. The annual average rainfall for the WA farm is between 493 mm with winter dominant rainfall (BOM 2024c).

3.2.2 Available data

3.2.2.1 Yield and protein data

Wheat grain yield and protein sensor data was collected at harvest for the 2020 – 2023 growing seasons on harvesters equipped with the John Deere HarvestLab 3000^{TM} NIR spectroscopy sensor. During harvest, samples were taken every ~1 – 3 seconds, and all yield and protein measurements were collected at the same location. The yield data was cleaned to remove values < 0 or > 10 tonnes per hectare (t/ha; Taylor *et al*. 2007a), and the protein data was cleaned to remove values < 6 and > 22% to exclude extreme outliers but still represent protein responses to different seasonal conditions (Blakeney *et al*. 2009; Žilić *et al*. 2011).

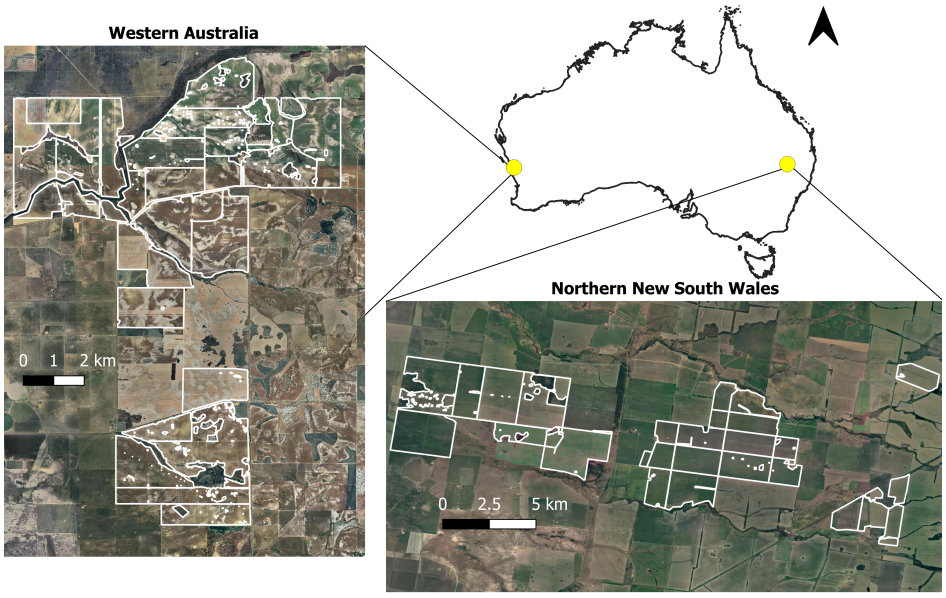

FIGURE 3.1. Map of farms in northern New South Wales and Western Australia.

Measurements more than 2.5 standard deviations (SDs) above and below the field mean were removed (Taylor *et al*. 2007a), and concurrent duplicate protein values were also removed based on their time and location to account for potential data errors at harvest. All data processing and analysis was performed in R, version 4.3.1 (R Core Team 2021).

Yield and GPC data was collected for 16 fields across the NNSW farm, and 30 fields across the WA farm, for four seasons (2020 – 2023). All yield and protein data for each season was allocated a unique field-year combination representing the corresponding field and season. There were 22 unique field-years worth of data available from the NNSW farm, and 41 from the WA farm. The yield and protein data points were then rasterised onto a continuous 30 m grid using block kriging via the gstat R package (v2.1.1; Pebesma 2004). The same 30 m grid was used throughout to build the covariate grid, and for model calibration, validation, and prediction.

3.2.2.2 On-farm agronomic data

On-farm agronomic data layers, including growing season, sowing and harvest dates, wheat variety, and cropping history for the previous three seasons (Table 3.1), was accessed for each field-year from farm machinery records and historical yield maps.

3.2.2.3 Publicly-available data

All publicly-available data layers, including satellite remote sensing imagery, geophysical and soil attribute data, were accessed via the dataharvester R package (v0.1.2; Haan *et al*. 2023; Harianto *et al*. 2023). The dataharvester R package is an R interface to the Geodata-Harvester, an open-source software that facilitates automated geodata harvesting from a wide range of geospatial and environmental data sources and provides users with spatial-temporal aligned raster maps and data frames.

Satellite remote-sensing imagery

Sentinel-2 MultiSpectral Instrument (MSI) Level-2A satellite imagery from 1 August to 31 October for each year was accessed at a 30 m spatial resolution from the Google Earth Engine (GEE; Gorelick *et al*. 2017). Sentinel-2 has a high spatial resolution of up to 10 m and a high temporal resolution with a 5-day global revisit frequency, meaning that it is well-suited to the within-field analysis of spatial variability in crop-reflectance throughout the growing season (Segarra *et al*. 2020). A cloud-masking filter was applied to all images to remove pixels affected by cloud cover. The Normalised Difference Vegetation Index (NDVI; Rouse *et al*. 1974; Equation 3.1), NDRE (Gitelson *et al*. 1994; Equation 3.2), and Enhanced Vegetation Index (EVI; Liu *et al*. 1995; Equation 3.3) were then calculated for each image and the peak (maximum) value for each vegetation index between August 1 – October 31 was extracted for each growing season.

$$NDVI = \frac{NIR - Red}{NIR + Red} \tag{3.1}$$

$$NDRE = \frac{NIR - RedEdge}{NIR + RedEdge} \tag{3.2}$$

$$EVI = 2.5(\frac{NIR - Red}{NIR + 6 \times Red - 7.5 \times Blue + 1}) \tag{3.3}$$

Remotely-sensed vegetation indices are a useful proxy for crop biomass, which in turn is a good proxy for crop growth, yield, quality, and the factors driving variability (Stoy *et al*. 2022). The NDVI was used as a benchmark as it is the most widely used vegetation index that correlates with yield (Ulfa *et al*. 2022) and has been used to estimate GPC (Tan *et al*. 2020; Stoy *et al*. 2022). The red-edge band, which is used to calculate the NDRE (Equation 3.2), is sensitive to changes in chlorophyll content and plant N concentration (Cao *et al*. 2018; Meloni *et al*. 2024) and is less sensitive to saturation at high biomass loads compared to the NDVI. The NDRE has been used to map grain yield (Wang *et al*. 2019) and represent variability in GPC (Szigeti *et al*. 2024). The EVI is an index sensitive to biomass (Huete *et al*. 2002) and also does not saturate at high biomass loads, and has been used in the estimation of GPC (Tan *et al*. 2020; Xu *et al*. 2020).

Imagery from August – October for each growing season was used as it (approximately) captures the reproductive and grain development phases (i.e. stem elongation to flowering, and flowering to maturity), which have been found to be the best periods for estimating grain yield and protein content (Zhao *et al*. 2005; Wang *et al*. 2014). During this time, rapid growth is taking place and there is a high water and nutrient demand from the crop during the development, storage, remobilisation and translocation of resources including carbohydrates and nitrogen (Edwards *et al*. 2008; Manning *et al*. 2008). Current-season imagery was used to capture seasonal-interactions between genetic, management, and environmental factors, while long-term imagery was largely used to represent more temporally-stable soil attributes and management history.

All Landsat 7 Tier 1 Surface Reflectance satellite imagery from 1 January 2000 to 31 December 2020 was accessed from the GEE at a 30 m spatial resolution. No seasonal filtering was applied, meaning that images from both cropping and non-cropping months were included

to capture long-term trends in surface reflectance. A cloud-masking filter was applied to all images to remove pixels affected by cloud cover. For each pixel, the 5th, 50th (median) and 95th percentile values were calculated across the entire 2000 – 2020 time-series, rather than per year, to reflect the full range of variability. Given the many images obtained over the available time period, the banding error issue generated by the failure of the Scale Line Corrector in Landsat 7 was negated. The Landsat 7 imagery was used to assess the underlying productivity, independent of seasonal effects, and was used to derive Red Band and NDVI imagery. The Red band imagery represented soil variability (i.e. topsoil colour; Filippi *et al*. 2019a; Jones *et al*. 2022), and the NDVI imagery represented biomass production.

Terrain, geophysical and soil attributes

All terrain, geophysical and soil attribute data was accessed via the dataharvester R package. The Geoscience Australia Shuttle Radar Topography Mission (SRTM) derived DEM 1 arc second grid (Gallant *et al*. 2009), and Sentinel-2 Barest Earth imagery (Wilford *et al*. 2021) were accessed at a 30 m spatial resolution. Airborne gamma radiometric dose rate, uranium, thorium, and potassium data (geophysical attribute) was sourced from the Geoscience Australia's Radiometric Map of Australia (Minty *et al*. 2009) at a 100 m spatial resolution. While vegetation is temporally variable and reflects seasonal responses to the environment and management factors, the Sentinel-2 Barest Earth imagery is sensitive to surface minerology and soil attributes (Wilford *et al*. 2021) and can be used to represent variability in soil type. Topographic variation, described through terrain attributes like elevation, and its influence on the movement and accumulation of water and nutrients has been shown to have a significant influence on crop yield, and this relationship is stronger for drier seasons (Kumhálová *et al*. 2011). Gamma radiometric data provides information about soil parent material and relates to soil forming processes (Pracilio *et al*. 2003).

All yield, agronomic, and publicly-available covariates were then extracted onto a 30 m grid for each field via nearest neighbour extraction. For each field, the same 30 m standard grid was used to rasterise yield and protein data, build the covariate grid, and for model calibration, validation, and prediction. An iterative variable selection process, which is described in Section 3.2.4, was used to select the final number of covariates to use within each model.

Table 3.1 shows the final covariates that were used within each model, but does not include those covariates that were available but were not selected.

3.2.3 Experiments

Five experiments involving different combinations of yield, on-farm agronomic (variety, sowing and harvest dates, cropping history for the three past seasons) and/or publicly-available data layers (DEM, radiometric surveys, remotely sensed vegetation indices, barest earth satellite imagery) were used to model wheat GPC and yield.

(1) **Experiment 1**: Yield + Agronomic + Publicly-available
(2) **Experiment 2**: Agronomic + Publicly-available
(3) **Experiment 3**: Yield + Publicly-available
(4) **Experiment 4**: Publicly-available
(5) **Experiment 5**: Yield

For each experiment, one model was built for each farm. Experiments 1 – 5 were assessed for GPC, and experiments 2 and 4 were assessed for yield.

[H]

TABLE 3.1. Final spatial covariates used in the random forest models for wheat grain protein content (GPC) and yield across the northern New South Wales (NNSW) and Western Australia (WA) farms for each of the five experiments. X denotes if the covariate was chosen in the final model for each experiment. This table does not include those covariates that were available but were not selected for any experiment during the iterative model building process. DEM, Digital Elevation Model; NDVI, Normalised Difference Vegetation Index; NDRE, Normalised Difference Red Edge; EVI, Enhanced Vegetation Index; SWIR, short-wave infra-red. The following experiments are described in detail in Section 3.2.3: E1, Experiment 1 – Yield + Agronomic + Publicly-available; E2, Experiment 2 – Agronomic + Publicly-available; E3, Experiment 3 – Yield + Publicly-available; E4, Experiment 4 – Publicly-available; E5, Experiment 5 – Yield.

			NNSW							WA						
			GPC					Yield		GPC					Yield	
Source	Data type	Covariate	E1	E2	E3	E4	E5	E2	E4	E1	E2	E3	E4	E5	E2	E4
On-farm	Harvest	Yield	X		X		X			X		X		X		
	Agronomic	Year	X	X				X		X	X				X	
		Sowing Date	X	X				X		X	X				X	
		Harvest Date	X	X						X	X					
		Variety	X	X				X		X	X					
		Crop history – 1 year prior (Crop -1)	X	X				X		X	X					
		Crop history – 2 years prior (Crop -2)	X	X				X		X	X				X	
		Crop history – 3 years prior (Crop -3)								X	X					
Publicly-available	Remote sensing	Current season peak NDVI (peak NDVI)	X	X	X	X		X	X	X	X	X	X		X	X
		Current season max NDRE (max NDRE)	X	X	X	X		X	X	X	X	X	X		X	X
		Current season max EVI (max EVI)	X	X	X	X		X	X	X		X	X		X	X
		Mean EVI, 1 year prior (EVI 1)	X	X	X	X		X	X		X	X	X		X	X

Table 3.1 continued from previous page

	Mean EVI, 5 years prior (EVI 5)	X	X	X	X	X	X			X	X		X	X
	Mean EVI, 10 years prior (EVI 10)	X	X	X	X	X	X		X	X				
	NDVI 50th percentile (NDVI 50)	X	X	X	X	X								
	NDVI 95th percentile (NDVI 95)		X	X	X								X	X
	RED 50th percentile (RED 50)										X			
Radiometrics	Dose Rate			X	X	X		X	X	X	X		X	X
	Potassium (Rad K) (%)							X	X	X				
	Thorium (Rad Th) (%)							X	X	X	X		X	
	Uranium (Rad U) (%)							X	X				X	X
Terrain attribute	DEM	X	X	X	X	X	X	X	X	X	X		X	X
Bare Earth	Blue					X	X							
	SWIR 1	X	X	X	X	X	X							
	SWIR 2					X								

3.2.4 Modelling

Random forest models (Breiman 2001), which are a suite of decision trees, were used to model the relationship between GPC or yield and the predictor variables in each farm for each experiment. Random forest models have been used to forecast grain crop yields (e.g. Filippi *et al*. 2019b) and Zhou *et al*. (2021) demonstrated that random forest models outperformed a traditional linear regression approach and artificial neural networks (ANNs) when predicting GPC. Compared to more complex machine learning approaches (e.g. ANNs), random forest models have a relatively low training time and computational cost (Roßbach 2018). Random forest models can also handle interactions between predictor variables, capture non-linear relationships between response and predictor variables, and are less sensitive to noise and outliers in the data (Breiman 2001).

The ranger R package (v0.14.1; Wright *et al*. 2017), was used to implement the random forest models as it enables the fast implementation of random forest models and is well-suited to high dimensional data (Wright *et al*. 2017). For each random forest model, the number of trees was set to 500 and the number of variables to possibly split at each node (mtry) was set to the default option, which was the rounded-down square root of the number of variables within each model (Wright *et al*. 2017).

An iterative process was used for variable selection and modelling, akin to recursive feature elimination but was guided by permutation-based variable importance rather than coefficient weights or statistical criteria. For each farm, an initial model for GPC or yield was built for each experiment using all available covariates for that experiment. The 'permutation' variable importance measure from the ranger R package (Strobl *et al*. 2008) was then used to rank covariates. In each iteration, the least important covariate was removed based on the permutation variable importance ranking, and a new model was built without it. This process was repeated iteratively, recalculating the variable importance at each step. Out-of-bag model quality statistics were recorded at each iteration, and after all iterations were complete, the Lin's Concordance Correlation Coefficient, LCCC) was compared against the number of covariates removed. The optimal number of covariates to exclude was determined

by identifying the point where the rate of the change in LCCC stabilised and/or reached the highest value. There was no predefined cutoff values for the maximum LCCC, and a parsimonious model was sought to balance model complexity and quality.

3.2.5 Validation approaches

Grain protein sensors may not be available across all fields or farms. In some cases, entire fields may not have maps of GPC. In other cases across large fields where multiple harvesters are operating (e.g. in Australia or the USA), only one harvester may be equipped with a grain protein sensor. This leaves information gaps across parts-of or for an entire field (Filippi *et al.* 2025). These two scenarios were tested using two different validation approaches (Figure 3.2):

(1) A LOFYOCV method, which was used to simulate cases where grain protein sensor data was not available for an entire field; or
(2) A 2FCV method, which was used to simulate cases where only one harvester is equipped with a grain protein sensor and GPC data is only available for part of a field.

The LOFYOCV approach was used as a between-field cross-validation approach. For each iteration, all data from a single field-year was removed and set aside as the validation dataset. The remaining data from all other field-years was used as the calibration dataset to train the random forest model, incorporating the selected features from the iterative variable selection process for each experiment (Section 3.2.4). A new model was built for each calibration data subset and then applied to the corresponding validation field-year. This process was repeated across all field-years (i.e. a total of 22 times for the NNSW farm, and 41 times for the WA farm).

The objective of LOFYOCV was to demonstrate how existing data could be used to fill in gaps, both spatially and temporally. Spatially, LOFYOCV maintains spatial structure in the data by leaving out entire field-years during cross validation, thus testing the models ability to extrapolate onto unseen areas. While the LOFYOCV prevents the model from

using direct yield or protein observations from the validation year, some temporally-stable covariates (e.g. elevation, barest earth imagery, radiometrics) persist in both the calibration and validation datasets if there are multiple years of data for the same field. However, this was not a concern as the model must still infer protein content without direct access to field-year specific observations (e.g. peak vegetation index values or agronomic details for that season). This assesses the models ability to use data from other seasons and fields to map GPC across unseen areas and in new seasons, and addresses the temporal gap-filling objective. For example, if a grower has grain protein data for certain seasons but lacks data for other seasons, information from the same field in other seasons can be used to infer historical protein levels.

The 2FCV method was used as a within-field cross-validation approach to evaluate model interpolation to unseen areas within a field data from part of that same field, from other fields, and other seasons. Each field was divided into 10 equal longitudinal sections to simulate different combine header passes, which were alternately labelled as "even" or "odd". For each field-year, the even-labelled sections were removed and set aside as the validated dataset, while the remaining field-years and the odd-labelled sections of the current field-year were used as the calibration dataset to train the random forest model. The process was then repeated, this time using the odd-labelled sections as the validation dataset and the even-labelled sections for calibration. This allowed the model to interpolate missing data within a field while still utilising full-field datasets from all other field-years to strengthen predictions. While field-year specific data was shared between the "odd" and "even" portions of each field-year (e.g. sowing and harvest dates), potential data leakage was not a concern as the objective of this cross-validation approach was to interpolate GPC data within fields to fill-in incomplete grain protein maps, not predict GPC across unseen areas (as for LOFYOCV). The spatial structure is still maintained by splitting the data into spatial blocks based on longitude, simulating the challenge of multiple harvesters but only one protein sensor, but agronomically-important information is still retained (Filippi *et al*. 2025). This 2FCV process was conducted for all field-year combinations (i.e. a total of 44 times in NNSW, and 82 times in WA).

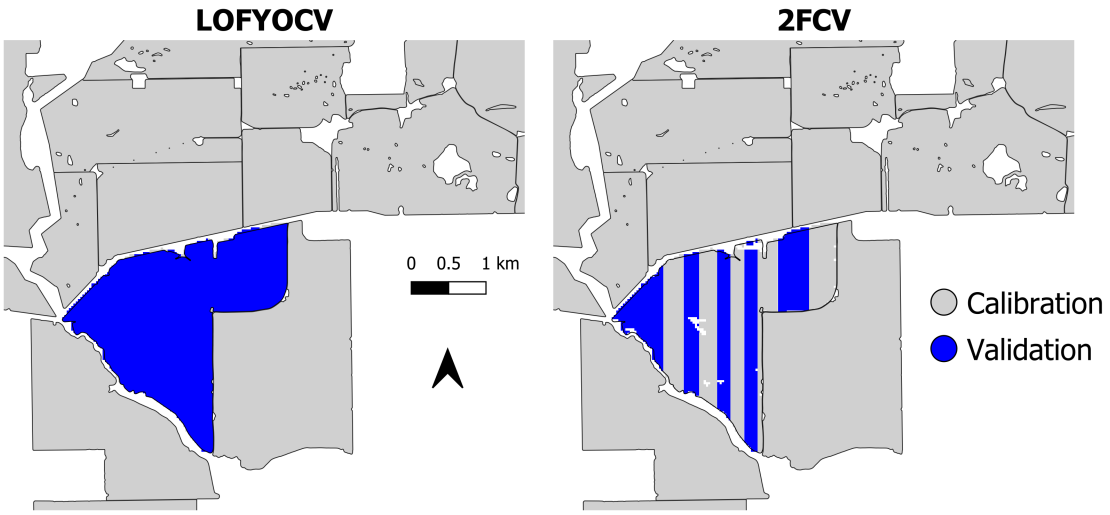

FIGURE 3.2. The leave-one-Field-Year-out cross-validation (LOFYOCV, a) and two-fold cross-validation (2FCV, b) approaches used for an example field. The LOFYOCV and 2FCV approaches were repeated for all Field-Year combinations (i.e. a total of 22 times for the northern New South Wales farm, and 41 times for the Western Australia farm).

Model quality was assessed using the LCCC (Lin 1989) and the root mean squared error (RMSE) to enable a comprehensive evaluation of model performance.

The LCCC is a measure of both the precision (how close the predictions are to each other) and the accuracy (how close the predictions are to the true values). It quantifies the fit of the observed and predicted values to a 1:1 line, where values of 0 are a poor fit (poor agreement between observed and predicted values) and 1 for a perfect fit (perfect agreement between observed and predicted values). The LCCC is unitless, making it particularly useful for comparing the predictive performance of models across variables of different magnitudes (e.g. between yield in t/ha and GPC %; Lin 1989). Thus, the LCCC was used to compare the overall quality of model predictions, considering both precision and accuracy, between yield and GPC.

The RMSE provides a direct measure of the accuracy of the predictions in the variable's units (e.g. % for GPC, or t/ha for yield). It quantifies the magnitude of prediction errors, with lower RMSE values indicating better model accuracy. The RMSE was used to compare the

performance of different experiments for predicting either GPC or yield, with a focus on minimising the error between predicted and observed values.

3.2.6 Validation spatial resolution

Model calibration, validation, and predictions were made at a fine (30 m) resolution. Model quality statistics (RMSE and LCCC) were then assessed at both a fine-resolution and across management classes.

Aggregating fine-resolution data to management classes can reduce noise and provide maps that are more informative for management decisions such as for N removal and prescription maps. The delineation of fields into management classes is common when implementing precision agriculture and site-specific crop management (Taylor *et al*. 2007a). The terms "management zone" and "management class" are frequently used interchangeably. A management class refers the area which a particular treatment may be applied, whereas a management zone is a contiguous area to which a specific management class treatment is applied. Thus, a management class may consistent of numerous management zones (Taylor *et al*. 2007a; Whelan *et al*. 2013). Here, management classes are referred to as being those areas that share a similar yield response, based on yield data for the current season.

Yield data from the current season was used to create five management classes per field using k-means clustering. While multiple seasons of data are typically used to define management classes, this study focused on aggregating the data for model evaluation only, rather than for creating operational classes. Using only the current season's yield data ensured that the management classes reflected the actual growing conditions influencing both yield and GPC during that specific year, even in atypical seasons. Five management classes were used to maintain validity when calculating the model quality statistics (LCCC and RMSE). Observed and predicted GPC and yield values were then aggregated (mean) into each class, and model quality was assessed using the RMSE and LCCC.

3.2.7 Prediction accuracy

Grain protein grades are a key determinant of the prices that growers receive. In Australia, GPC is classed into four grades; < 10.5%, 10.5 – 11.5%, 11.5 – 13%, and 13%. Depending on the nature of variation in GPC grades within and between fields, different protein grades may determine how grain is harvested, stored, and transported, or how pre- and in-season management decisions are made such as N fertiliser applications, for example. Understanding how well the model performs in terms of accurately predicting the correct (i.e. observed) protein grade is of importance to growers and decision makers.

Cohens kappa (κ, Cohen 1960; Equation 3.4) was used to assess how well the model predicts the correct protein grade (i.e. if observed and predicted values fall within the same protein grade) while also accounting for random chance. Cohens kappa calculates the agreement between observed and predicted classes (the proportion of times the observed and predicted classes are the same), and takes into account the possibility of random agreement, which is the agreement you would expect if the model were randomly assigning categories, given the distribution of the classes in the observed and predicted data. For example, if a class is very common, the model might predict that class correctly simply because it is guessing that class more frequently. Unlike simple accuracy metrics, which may be inflated by the majority class, Cohens Kappa can handle class imbalance. Cohens kappa is calculated as follows;

$$\kappa = \frac{p_0 - p_e}{1 - p_e} \tag{3.4}$$

Where:

$p_0 =$ the observed agreement (the percentage of times the observed and predicted data agree), and

$p_e =$ the expected agreement by chance (the probability that the observed and predicted values agree randomly, based on their individual distributions).

Similar to correlation coefficients, Cohens kappa is a unitless value ranging from -1 to 1, where 0 represents the agreement expected from random chance, and 1 represents perfect agreement. Cohens kappa can be interpreted as follows (Cohen 1960);

- κ 0: no agreement (the agreement expected by random chance or worse),
- κ 0.01 – 0.20: none-to-slight agreement,
- κ 0.21 – 0.40: fair agreement,
- κ 0.41 – 0.60: moderate agreement,
- κ 0.61 – 0.80: substantial agreement, and
- κ 0.81 – 1.00: (almost) perfect agreement.

3.3 Results and discussion

3.3.1 Experiments and available data

3.3.1.1 Exploratory data analysis of observed yield and protein data

Overall, observed yields were higher in NNSW (mean = 3.69 t/ha) compared to WA (mean = 2.20 t/ha), and GPC was higher in WA (mean = 12.38%) compared to NNSW (mean = 11.72%, Table 3.2). Both yield and GPC showed a normal distribution with mean and median values being close to each other. The number of GPC observations within each of the Australian protein grades was more evenly distributed in NNSW than in WA. In NNSW, the GPC grade < 10.5% had the most observations, whereas in WA, the GPC grade 11.5 – 13% had the most GPC observations (Figure 3.3).

TABLE 3.2. Summary statistics for grain yield (tonnes per hectare, t/ha) and grain protein content (GPC, %) for the northern New South Wales (NNSW) and Western Australia (WA) farms. Summary statistics are provided for minimum (min) and maximum (max) values, 1st and 3rd quartiles (Q1 and Q3, respectively), mean, median, and coefficient of variation (CV).

	NNSW		WA	
Statistic	GPC (%)	Yield (t/ha)	GPC (%)	Yield (t/ha)
Min	6.11	0.00	6.04	0.04
Q1	10.29	2.84	11.10	1.31
Median	11.50	3.90	12.20	2.06
Mean	11.72	3.69	12.38	2.20
Q3	12.96	4.70	13.35	3.00
Max	21.76	9.91	21.29	8.12
CV (%)	17.27	36.04	14.08	49.41

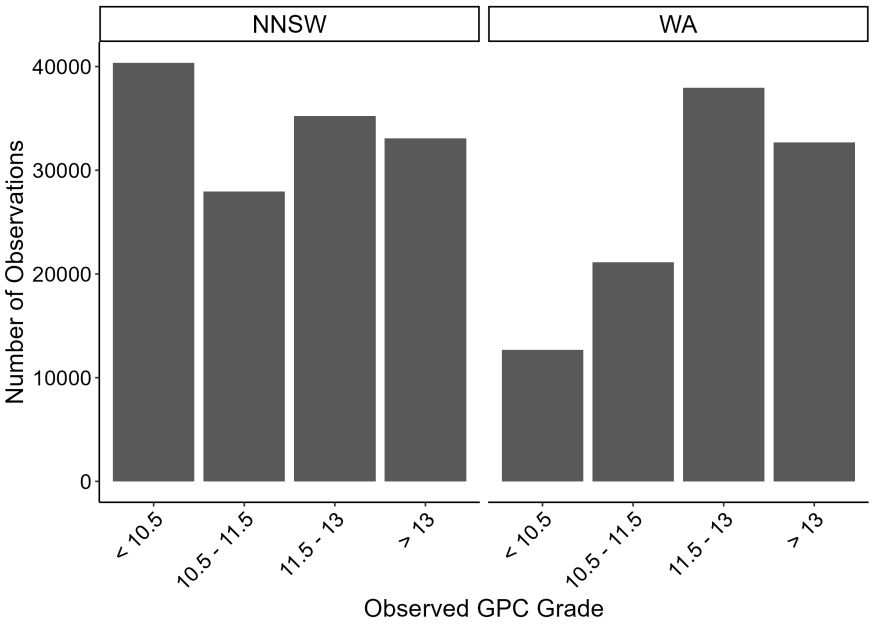

FIGURE 3.3. Number of observations in Australian wheat grain protein content (GPC) grades for northern New South Wales (NNSW) and Western Australia (WA).

Yield was more variable than GPC, based on the coefficient of variation (Table 3.2), and there was considerable variability in both GPC and yield within and between fields, and season-to-season (Appendix 3.A1). This variability is not unexpected in annual dryland (rainfed) farming systems where fluctuations in seasonal rainfall year-to-year is the primary

driver of temporal variation (French *et al*. 1984). It is important to note that yield, being a continuous variable, has a broader potential range and is therefore more variable based on the CV compared to GPC, which is a proportional variable that cannot exceed 0% or 100%.

There are key environmental and agronomic differences between the two growing regions. The western growing region (WA) is characterised by having low foil fertility and lower yields relative to Australia's southern (Victoria, Tasmania, South Australia) and northern (Queensland and NSW) grain growing regions. Comparatively, the northern grain growing region has relatively high seasonal rainfall and greater production variability compared to the southern and western growing regions (Davies *et al*. 2022). Given these distinct soil and climatic differences, a global model was not built, and instead, models were built at the farm-level, rather than bespoke modelling for individual fields. This focused on the use of existing, readily-available data to ensure model generalisability within a region without requiring additional data collection by growers.

An inverse relationship between grain yield and GPC is expected at the field-scale (Simmonds 1995), and this was observed overall in both NNSW (r = -0.06) and WA (r = -0.37; Figure 3.4). This inverse relationship is due to the grain protein "dilution effect", where under non-limiting soil moisture conditions, as the amount of carbohydrates in the grain (i.e. yield) increases, protein concentration decreases (Simmonds 1995). However, this is not always the case within fields (Whelan *et al*. 2009) and there was considerable variability in the relationship between yield and GPC within and between fields, and season-to-season (Appendices 3.A2 and 3.A3).

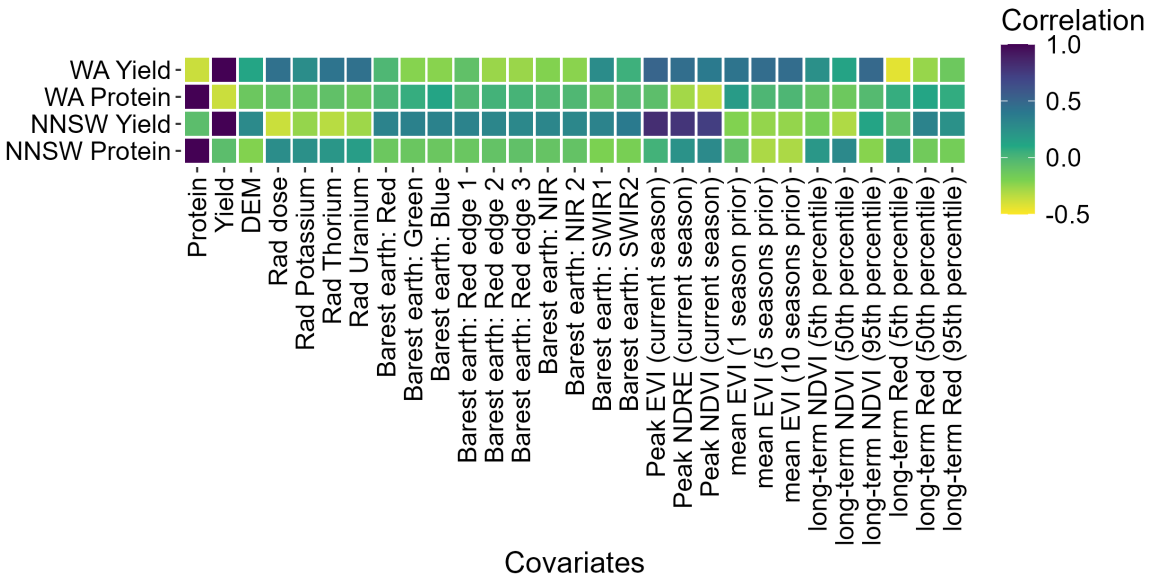

FIGURE 3.4. Spearman-rank correlations between wheat grain protein content and yield and all available covariates in northern New South Wales (NNSW) and Western Australia (WA). DEM, digital elevation model; Rad, radiometric; NIR, near-infrared; SWIR, short-wave infrared; EVI, enhanced vegetation index; NDRE, normalised difference red-edge; NDVI, normalised difference vegetation index.

3.3.1.2 Experiments and variable selection

The LCCC and RMSE were used to assess model performance. The LCCC is a unitless measure of both the precision and accuracy, enabling benchmarking across different variables (yield and GPC), whereas the RMSE quantifies the magnitude of prediction errors in the same units as the variable.

For both grain yield and GPC, experiments which utilised all possible data available performed best for both NNSW and WA. This means that for grain yield, Experiment 2 (Agronomic + Publicly-available) performed the best, and for GPC, Experiment 1 (Yield + Agronomic + Publicly-available) performed the best (Figures 3.5 and 3.5). In comparison between yield and GPC (using the LCCC), the performance for yield was consistently better than GPC (Figures 3.5 and 3.6). By far, Experiment 5 (Yield) performed worst when predicting GPC. Variable importance plots (Figure 3.7) and scatterplots of observed and predicted values (Figures 3.8

and 3.9) are only shown for the best final models from Experiment 1 for GPC and Experiment 2 for grain yield.

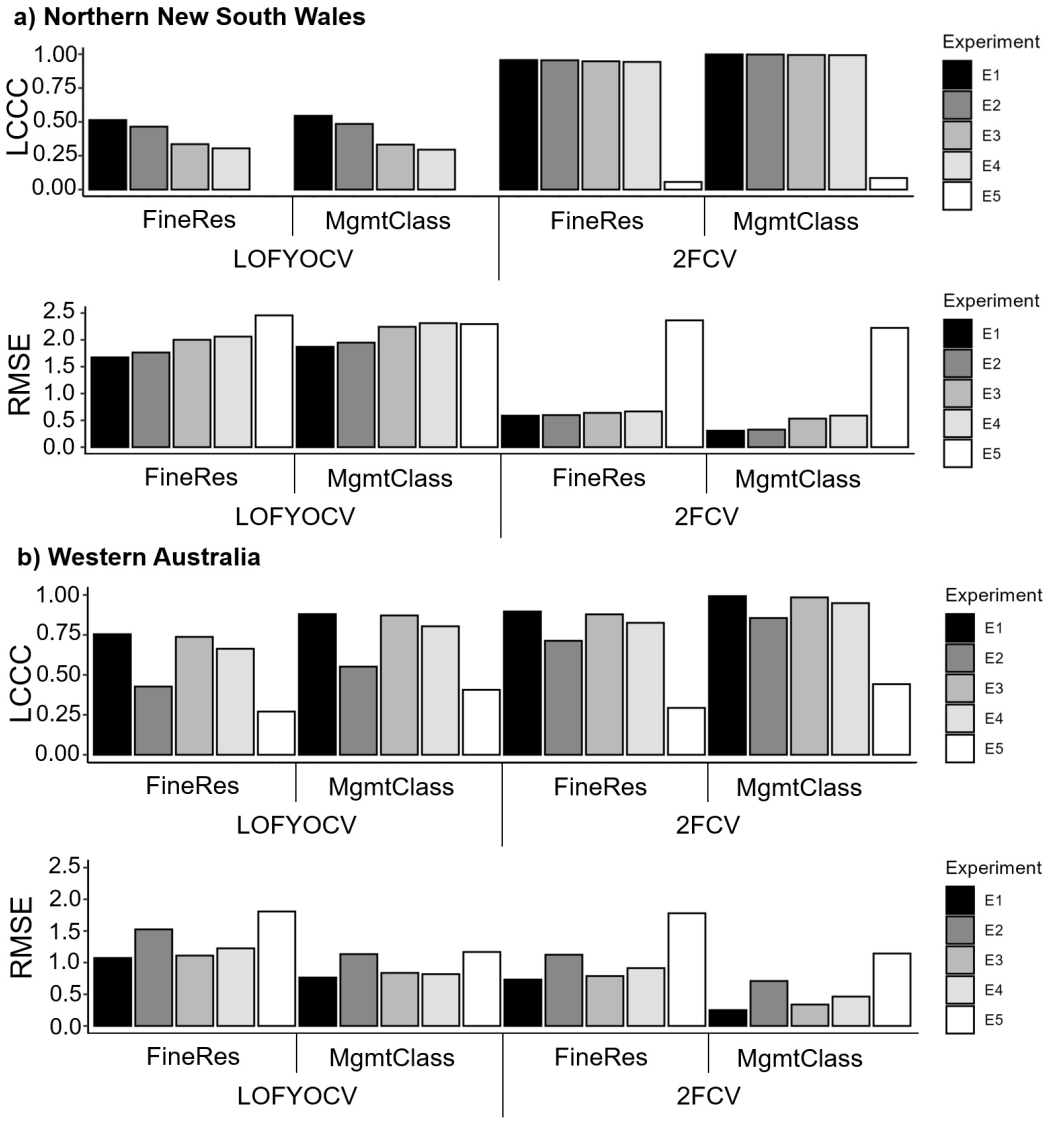

FIGURE 3.5. Wheat grain protein content model quality statistics for Experiments 1 – 5. Lin's Concordance Correlation Coefficient (LCCC) and root mean square error (RMSE) values are presented for a) Northern New South Wales and b) Western Australia farms for: Experiment 1 (E1: Yield + Agronomic + Publicly-available), Experiment 2 (E2, Agronomic + Publicly-available), Experiment 3 (E3, Yield + Publicly-available), Experiment 4 (E4, Publicly-available), and Experiment 5 (E5, Yield). Models were validated at a fine-resolution (FineRes) and aggregated to management classes (MgmtClass) using either a leave one field-year out cross-validation (LOFYOCV) and two-fold cross-validation (2FCV) approach.

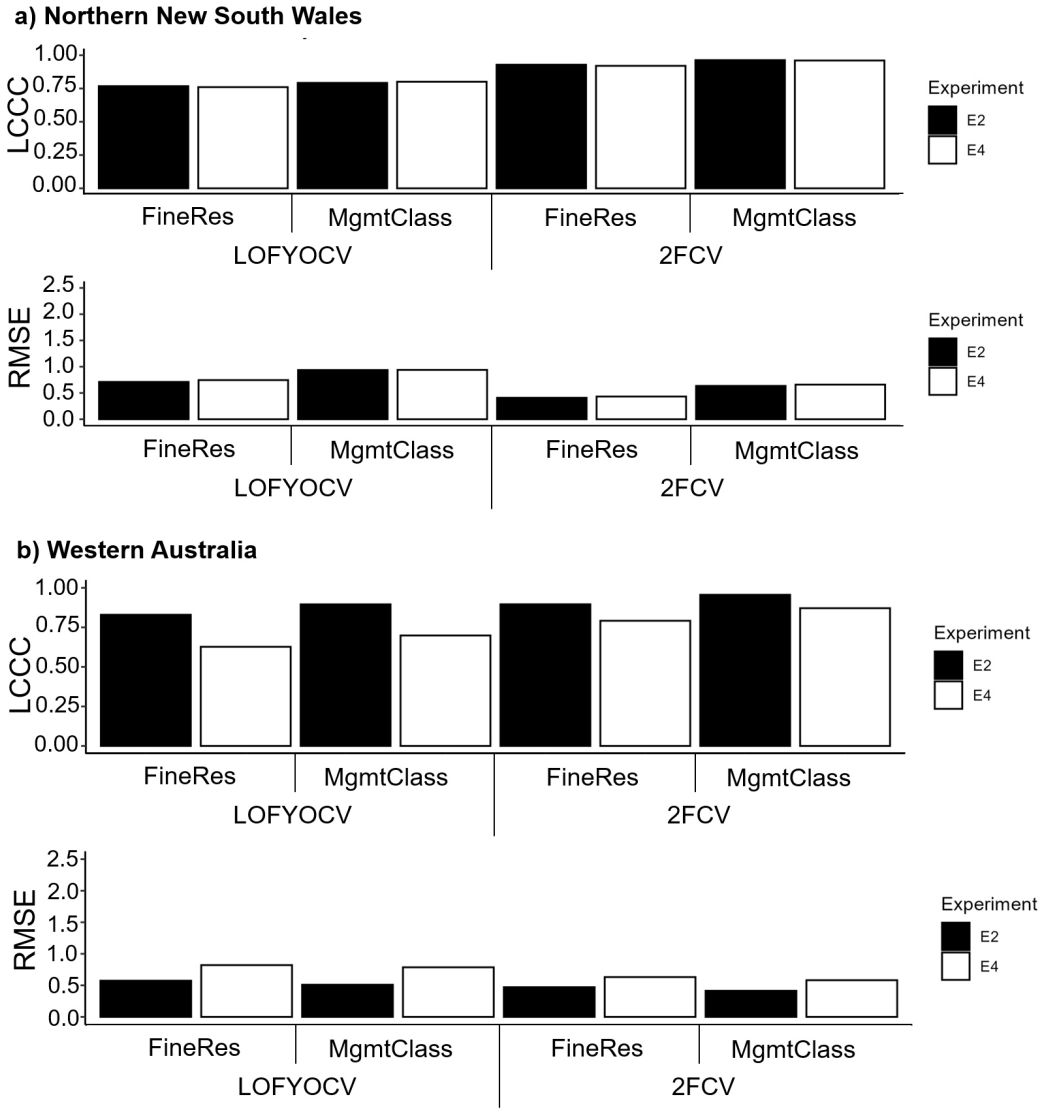

FIGURE 3.6. Wheat grain yield model quality statistics for Experiments 2 and 4. Lin's Concordance Correlation Coefficient (LCCC) and root mean square error (RMSE) values are presented for a) Northern New South Wales and b) Western Australia farms for: Experiment 2 (E2, Agronomic + Publicly-available), and Experiment 4 (E4, Publicly-available). Models were validated at a fine-resolution (FineRes) and aggregated to management classes (MgmtClass) using either a leave one field-year out cross-validation (LOFYOCV) and two-fold cross-validation (2FCV) approach.

Important features

While the uptake of grain protein sensors is increasing, it is unlikely that we will see a map of GPC for every field, farm, or season in the near future. This study highlights the potential to use existing yield, agronomic and publicly-available data layers to model and map GPC to fill-in previously unmapped areas of a farm.

Publicly-available data layers, including DEMs, remote sensing imagery, and radiometric surveys, are freely accessible across most fields and farms, making them valuable for large-scale applications. While models using only publicly-available data layers (Experiment 4) performed poorly, their inclusion alongside on-farm agronomic data layers considerably improved model performance. This reinforces the need to account for complex interactions between vegetation, soils, management, genetics, and seasonality, among other factors.

The best performing models incorporated all available data sources: Experiment 1 (Yield + Agronomic + Publicly-available) for GPC and Experiment 2 (Agronomic + Publicly-available) for yield. Publicly-available data layers served as proxies for key drivers of variability, such as soil attributes and nutrient availability. In NNSW, indices (e.g. mean NDVI and NDRE) and barest earth imagery (e.g. short-wave infra-red [SWIR] bands 1 and 2) were frequently selected in the best models, whereas their influence was weaker in WA (Table 3.1, Figure 3.7). Notably, agronomic factors like year, variety, and cropping history were among the most important predictors for GPC and yield, highlighting the key role of genetics and seasonal conditions (Matthews *et al*. 2024; Shackley *et al*. 2024; French *et al*. 1984).

In WA, radiometric data (dose rate, uranium, thorium, potassium) played a more prominent role in the best models for both GPC and yield. These layers are sensitive to changes in parent material and soil texture (Pracilio *et al*. 2003), and the soils across the NNSW farm are comparatively more homogenous than those in WA (Davies *et al*. 2022). Gamma radiometrics have been demonstrated to be useful predictors of clay content in weathered soils with variable soil texture in southern WA (Taylor *et al*. 2002; Pracilio *et al*. 2003).

Interestingly, the relationship between remote sensing indices and crop performance varied by region. Yield correlations with peak seasonal vegetation indices were positive in both

regions, but correlations were stronger in NNSW where higher rainfall and better soil moisture retention typically supports greater biomass conversion into yield (Davies *et al*. 2022). For GPC, peak seasonal vegetation indices were positively correlated in NNSW but negatively correlated in WA. In WA, this likely reflects the higher GPC driven by lower yields and water-limited conditions later in the season. Higher vegetation index values in WA may indicate more vigorous early growth, but the limited water availability later leads to lower yields and higher GPC.

The importance of previous cropping history was greater for GPC than for yield. Soil moisture and N are the two key drivers of variability in GPC (Whelan *et al*. 2013), and management decisions made during the previous season(s), including nutrient inputs, crop rotations, fallows, or row spacings, may have influenced soil moisture or residual N availability. However, data describing soil moisture or nutrient levels during previous seasons was not available to quantify or confirm this. Model performance for yield consistently exceeded that of GPC, reflecting stronger correlations observed between yield and current-season vegetation indices (r = 0.73 to 0.82, Appendices 3.A2 and 3.A4) compared to GPC (r = 0.03 to 0.28, Appendices 3.A3 and 3.A4). These results align with Wang *et al*. (2019), who found that while the NDRE and NDVI were useful predictors of grain yield and N uptake, their relationship with GPC was weak. Remote sensing indices provide a broad overview of vegetation vigour and biomass, but they primarily reflect cumulative growth patterns rather than the timing of key stress events, such as water or N limitations, particularly towards the end of the growing season, which have a more direct impact on GPC.

For some fields in NNSW, there was considerable variability in GPC that appeared to reflect prior field boundaries where smaller fields had been aggregated into one large field, despite these aggregations occurring ~10 years previously. While it is not clear why these patterns of variability persist, long-term remotely-sensed vegetation indices and barest earth imagery appears to capture this management history and field boundaries in the absence of seasonal fluctuations in vegetation. As the size and structure of Australia's farming enterprises change and smaller fields are aggregated into larger management areas, it is expected that similar patterns of variability will be revealed as more protein sensors are adopted commercially.

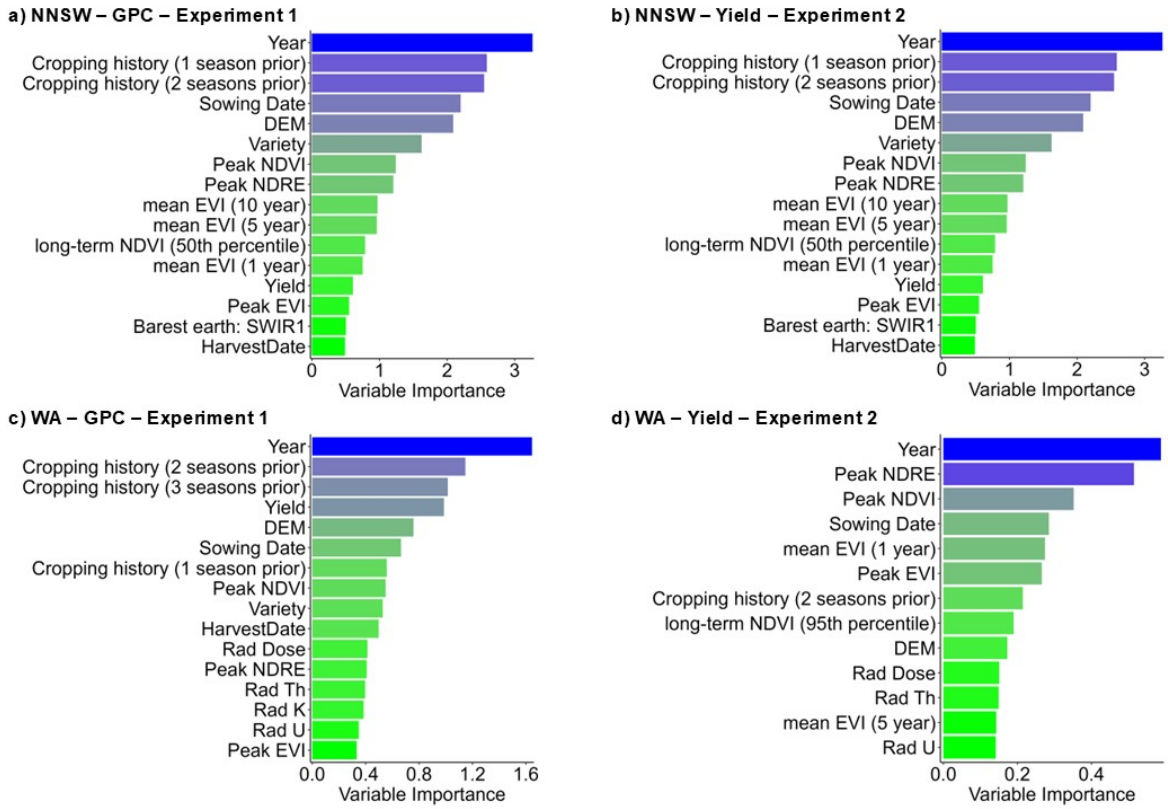

FIGURE 3.7. Variable importance plots for final wheat grain yield (B and D) and grain protein content (GPC; A and C) models in northern New South Wales (NNSW; A – B) and Western Australia (WA; C – D). Experiment 1: Yield + Agronomic + Publicly-available data layers. Experiment 2: Agronomic + Publicly-available data layers. DEM, Digital Elevation Model; NDVI, Normalised Difference Vegetation Index; NDRE, Normalised Difference Red Edge; EVI, Enhanced Vegetation Index; SWIR, short-wave infra-red.

3.3.2 Validation approaches

Both the LOFYOCV and 2FCV approaches are "spatially-aware" (Habibi *et al*. 2024) validation strategies, in that they divide the data into calibration and validation subsets based on their spatial location. The LOFYOCV approach involved removing an entire field-year to use as the validation subset, whereas the 2FCV approach involved splitting each field-year into two halves based on longitude (Figure 3.2) and using each half as the validation subset during each iteration. Compared to random k-fold cross-validation strategies, which overlook spatial structure or dependencies in the data, thereby potentially underestimating model

prediction errors and reducing performance (Ruß *et al*. 2010; Ferraciolli *et al*. 2019; Habibi *et al*. 2024), spatial cross-validation methods like LOFYOCV and 2FCV can provide more realistic accuracy measures for predicting unknown locations (Christy 2008; Stevens *et al*. 2012).

Overall, the 2FCV approach consistently outperformed LOFYOCV. This was unsurprising given that the 2FCV method was not entirely spatially independent due to shared field-year-level management information (e.g. variety, sowing, and harvest dates) across the "odd" and "even" halves of each field-year. However, this was desirable given the study objectives, and the 2FCV approach demonstrated that existing data layers can be used to fill gaps within fields with incomplete grain protein maps. In 2FCV, spatial structure is still maintained by splitting the data into spatial blocks based on longitude, simulating the challenge of multiple harvesters but only one protein sensor, but there is some shared data for the entire field-year. Therefore, 2FCV is more an interpolative evaluation rather than extrapolative as the calibration and validation subsets are close and the spatial domain of the model is similar to where predictions are being made (Smith *et al*. 2023). Future research should explore within-field interpolation using only data from the field in question, particularly when mapping grain protein in incomplete harvest data scenarios.

Fields with multiple seasons of data had better validation statistics than those with only one season, which was expected. As already described in Section 3.2.5, temporally-stable covariates (e.g. elevation, barest earth imagery), persist in both the calibration and validation datasets if there are multiple years of data for the same field, even for LOFYOCV. Given that the objective of this study was to assess how existing data could be used to fill gaps, both spatially and temporally, using existing data from other seasons to predict for unsampled seasons in the same field will be beneficial for improving predictions and provide growers with a comprehensive GPC dataset to assess long-term trends. Incorporating data from more fields and seasons should improve model performance by capturing a greater range of growing conditions. and improve the models ability to predict yield or GPC during seasons with no data. It should be noted that the predictions made in this study are not forecasts, as data for the

same season (e.g. peak vegetation indices) were used as model input features, but forecasting GPC should be explored in future work (Section 3.3.5).

3.3.3 Validation spatial resolution

While fine-resolution maps of grain protein provide a high degree of detail describing the spatial variability of GPC, these may be difficult to use to make operational decisions. Fields can be divided into management classes, which represent areas of a field that have similar constraints or production potential and can be differentially managed according to their requirements (Whelan *et al*. 2000; Taylor *et al*. 2007a). For example, management classes can be used for nutrient budgeting and replacement through variable rate applications to ensure correct applications of inputs that address production potential within each class (Whelan *et al*. 2013).

Validating across management classes performed better than at a fine-resolution for all models in WA, and for the GPC model validated using the 2FCV approach in NNSW (Figures 3.4 and 3.5). This is not surprising, as it is expected that aggregating predictions to more homogenous classes within fields will reduce variation and small-scale noise is smoothed from the analysis. However, validation across management classes did not always produce better predictions. In NNSW, yield models always performed better when validated at a fine-resolution, as was the case for the GPC models validated using the LOFYOCV approach (Figures 3.4 and 3.5). The model quality statistics from these predictions were the poorest overall at a fine-resolution, and there were clear outlying fields when scatterplots of observed and predicted values were examined (Figures 3.8 and 3.9). These outlying fields generally had very weak correlations with covariates, or had more extreme yield or GPC values that were not captured well during the LOFYOCV approach in particular (Figures 3.8 and 3.9, Appendices 3.A2 and 3.A3). These outlying fields also had fewer observations at the point support, whereas each field had an equal number of management classes. This means that when cross-validations were performed across management classes, these outlying fields likely had a greater impact on the validation statistics than at a fine-resolution.

[h]

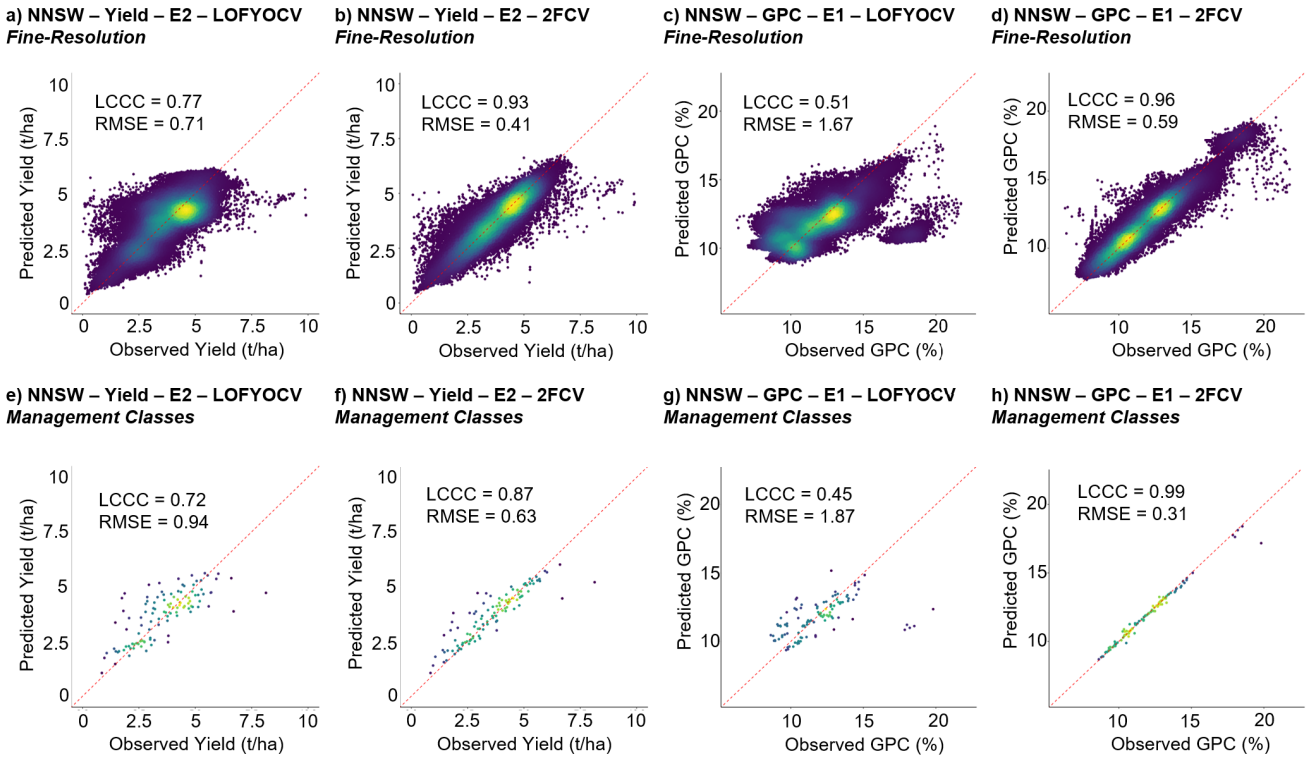

FIGURE 3.8. Observed and predicted values for wheat grain yield using Experiment 2 (E2) and grain protein content (GPC) using Experiment 1 (E1) for northern New South Wales (NNSW). Models were validated at a fine (30 m) resolution (FineRes; A, C, E, G) and across management Classes (MgmtClass; B, D, F, H) using either a leave-one-Field-Year-out cross-validation (LOFYOCV; A – B, E – F) or two-fold cross-validation (2FCV; C – D, G – H) approach.

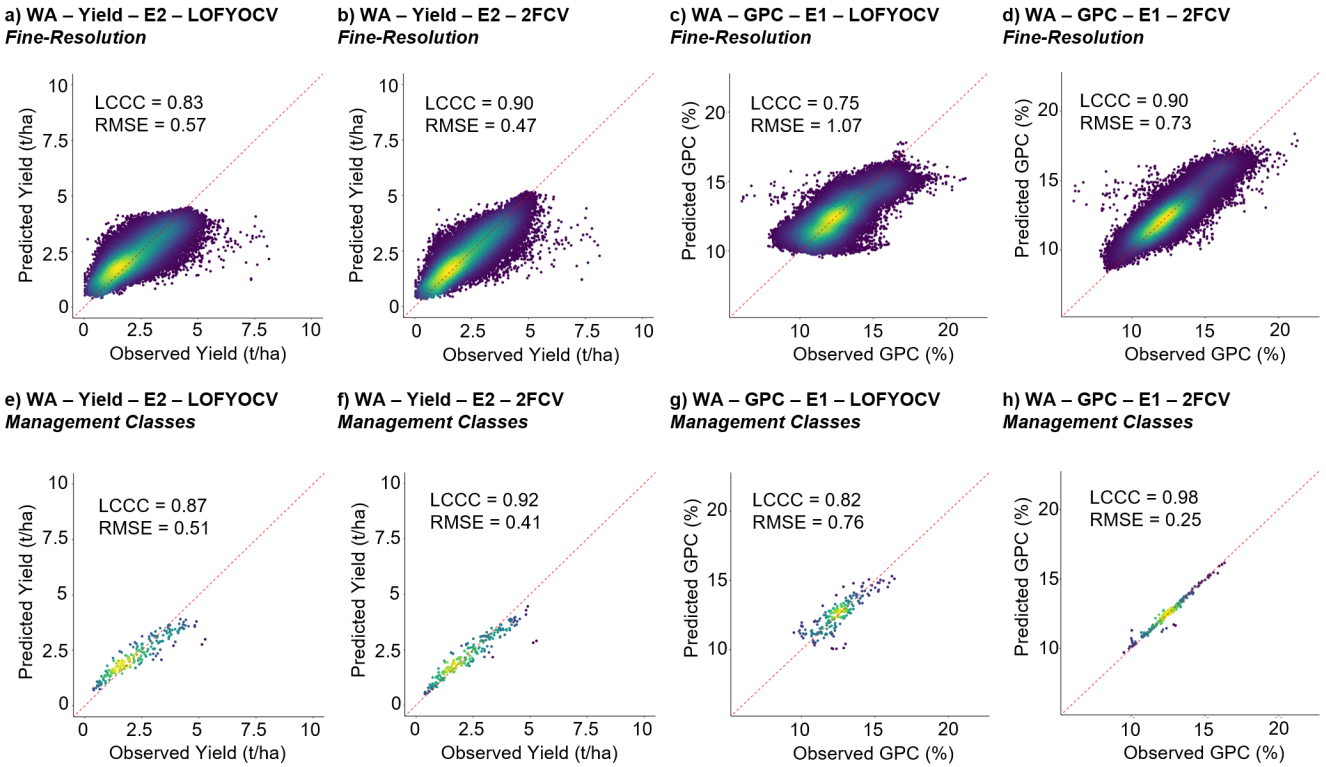

FIGURE 3.9. Observed and predicted values for wheat grain yield using Experiment 2 (E2) and grain protein content (GPC) using Experiment 1 (E1) for Western Australia (WA). Models were validated at a fine (30 m) resolution (FineRes; A, C, E, G) and across management classes (MgmtClass; B, D, F, H) using either a leave-one-Field-Year-out cross-validation (LOFYOCV; A – B, E – F) or two-fold cross-validation (2FCV; C – D, G – H) approach.

Management classes were delineated using a simple approach which used yield data for the current season. Management classes were used to test the impact of spatial resolution on the models performance, and this research did not aim to provide a recommendation to a grower or advisor as to how fields should be delineated into management classes. Thus, no optimisation of inputs or number of classes was conducted besides ensuring that there was a minimum of 100 data points in total to calculate the LCCC (i.e. 5 classes within each field). Given that the degree and drivers of variability in wheat GPC can be different between fields, the actual delineation of management classes should be conducted on a field-specific basis. The practical delineation of management classes should utilise an appropriate suite of data layers to represent variability, and a suitable number of management classes should be used that can practically be implemented and capture sufficient variability (Taylor *et al*. 2007a; Whelan *et al*. 2013).

3.3.4 Prediction accuracy

Cohens kappa is a measure of model accuracy, relative to random chance. Values range from -1 to 1, where 0 represents the agreement expected from random chance, and 1 represents perfect agreement. For both NNSW and WA, predictions made using Experiment 1 most accurately predicted the correct protein grade. Predictions aggregated to management-classes were more accurate than those at a fine-resolution, and the 2FCV approach was more accurate than the LOFYOCV approach (Figure 3.10). For Experiment 1, validating across management classes using the 2FCV approach had an almost perfect agreement between observed and predicted protein grades (i.e. $\kappa > 0.80$, Cohen 1960), and there was very little difference in the accuracy of predictions within each protein grade (results not shown). This indicates that the model captures and explains sufficient variability in the data and can handle an uneven distribution of observations between protein grades.

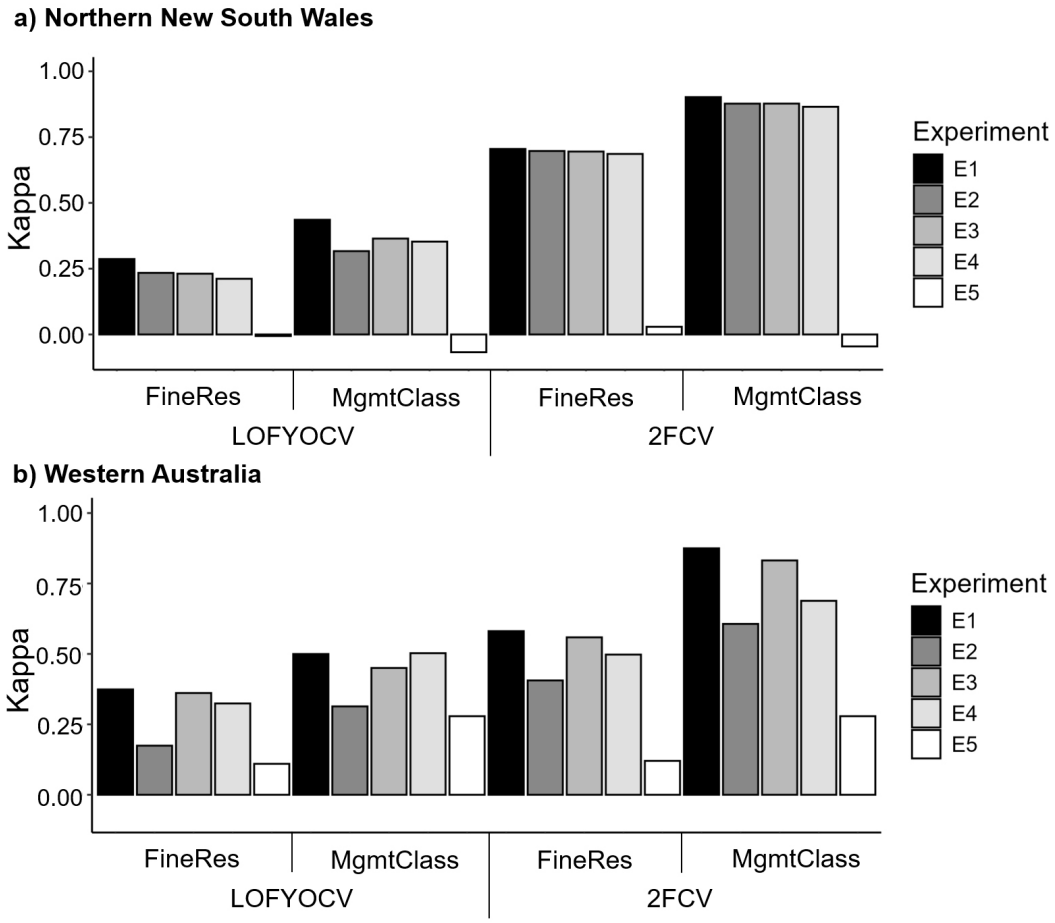

FIGURE 3.10. Prediction accuracy of model predictions of grain protein content (GPC) for five different Experiments in A) northern New South Wales (NNSW) and B) Western Australia (WA). Validated at a fine-resolution (FineRes) and across management classes (MgmtClass) using a leave one Field-Year out cross validation (LOFYOCV) or two-fold cross validation () approach.

Given the price premiums paid to Australian growers, accurately predicting the correct (observed) protein grade is important. In future, the ability to forecast grain protein grades may be of great benefit to growers to inform management operations such as the selective harvest of grain to attain premium prices (Bastos *et al*. 2021). Selective harvest, or the harvesting of grain according to different yield/quality criteria to exploit observed variation (Bramley *et al*. 2005), has been a major focus for high-value commodities including viticulture and horticulture. Bramley *et al*. (2012) explored the potential of selective harvest for cereals

to take advantage of on-the-go protein sensors and price premiums paid to Australian growers, but were limited by poor reliability of sensing technology at the time. Given considerable improvements made in the quality of grain protein sensors in recent years, despite their (slow but) gradual uptake, there is the potential to revisit the logistical and economic feasibility of selective harvest for GPC in the future.

3.3.5 Future work

3.3.5.1 Forecasting grain protein content

In the future, there is an opportunity to utilise the abundance of publicly-available data layers to build predictive models to forecast GPC prior to the end of the current season, or for future seasons (Bastos *et al*. 2021). It is possible to forecast grain yields within fields using publicly-available data layers, including satellite remote sensing imagery or climate data (e.g. Filippi *et al*. 2019b; Amin *et al*. 2024). While there has been some work on forecasting wheat GPC (Magney *et al*. 2016; Zhao *et al*. 2019) and soybean seed protein (Hernandez *et al*. 2023) prior to harvest, these studies are limited to a small number of fields or seasons with point-based samples which limit the ability to capture within-field variability. The addition of weather data, soil properties, or fertiliser input data may improve predictions by capturing the environmental and management factors that drive variability in GPC, but field-specific, bespoke data is time-consuming and expensive to collect and process (Chapter A; Tilse *et al*. 2022). This may restrict the ability to forecast protein to data-rich regions or farms only. Future work should examine the use of the growing catalogue of on-the-go grain protein sensor maps from multiple seasons, alongside publicly-available data layers to build models to forecast GPC prior to the end of the season.

3.3.5.2 Understanding the drivers of variability

Overall, while a combination of yield, agronomic, and publicly-available data layers could be used to predict grain yield and GPC, it is still unclear what is driving this variability. The drivers of variability in wheat grain yield, protein content, and their relationship within

and between fields, farms, and seasons should be explored in future research. This may include the utilisation of additional/different data layers and new approaches including interpretive machine learning. Typically, machine learning models like random forest models are considered a "black box", where it can be difficult to understand what factors are driving predictions within the model. Interpretive machine learning can be used to overcome this limitation. Interpretive machine learning refers to a collection of techniques developed to identify the importance of individual predictors in a model and determine what was used to make a prediction (Jones *et al*. 2022). Interpretive machine learning has been used to identify the causes of crop yield variability in cotton (Jones *et al*. 2022), where digital soil maps and terrain information was used to map cotton lint yield and interpretive machine learning was then used to identify the contribution of each predictor variable to the modelled yield prediction. Interpretive machine learning can be used to understand what may be driving variations in GPC and grain yield, and what may explain the changing relationship between yield and protein within and between fields, farms, and seasons. By identifying the contribution of each variable to modelled grain protein predictions, the major drivers of grain protein content within a field and across farms could then be mapped. Applying this over multiple seasons may also help to identify any seasonal fluctuations or changes in the drivers of grain protein over time.

3.4 Conclusion

Despite growing interest in measuring and mapping variability of GPC within fields and across farms, the uptake of grain protein sensors has been slow, resulting in considerable knowledge gaps. This chapter aimed to address these gaps by using different combinations of existing, readily-available yield, agronomic, and publicly-available data layers in a machine learning (random forest) approach to predict wheat grain yield and GPC within fields in the absence of a grain protein or yield sensor. Overall, predictions were better for grain yield than for GPC, and a combination of all available data sources (i.e. yield, agronomic, and/or publicly-available) produced the best predictions of grain yield and GPC. Estimates of grain yield and GPC were poorer when predicting onto an entirely new field (LOFYOCV), rather

than using some existing information from part-of a field (2FCV), and predictions were better when they were aggregated to management classes compared to when they were made at a fine-resolution. Grain protein grades were also predicted with a high degree of accuracy, which is important to growers as grain protein grades determine the price that growers receive. By building a predictive model to map GPC in areas of a farm without a grain protein sensor, growers and advisors can be equipped with the necessary information and tools to make better management decisions for more profitable and environmentally sustainable production systems that optimise both yield and quality. In the future, there is the opportunity to utilise the growing abundance of on-farm and publicly-available data layers to forecast GPC prior-to the end of the season, but more work is needed to better understand the drivers of variability within and between fields, farms, and seasons.

3.5 Appendices

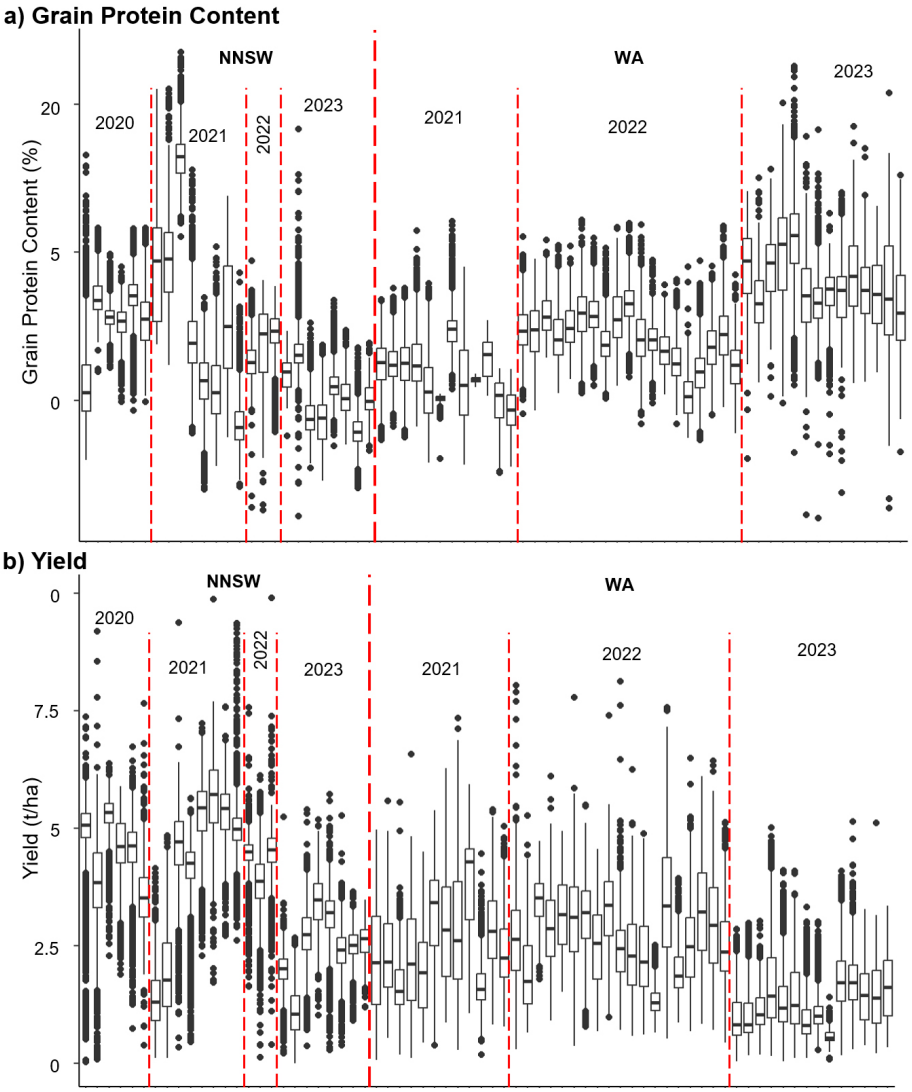

FIGURE 3.A1. Distribution of observed a) grain protein content (GPC) (%) and b) yield (tonnes per hectare, t/ha) data for all field-years in northern New South Wales (NNSW) and Western Australia (WA).

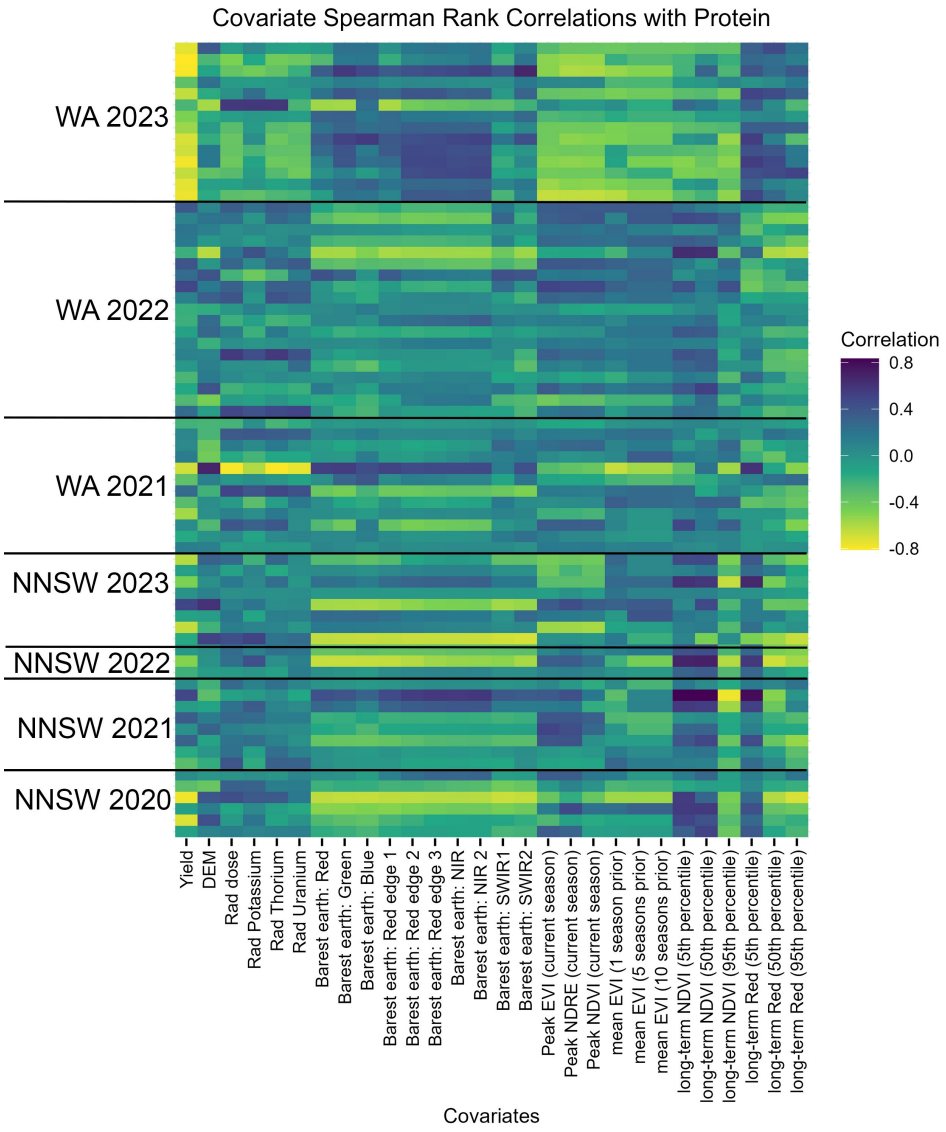

FIGURE 3.A2. Spearman-rank correlations between wheat grain protein content (GPC) and available covariates for each Field-Year in northern New South Wales (NNSW) and Western Australia (WA) for 2020 – 2023 seasons.

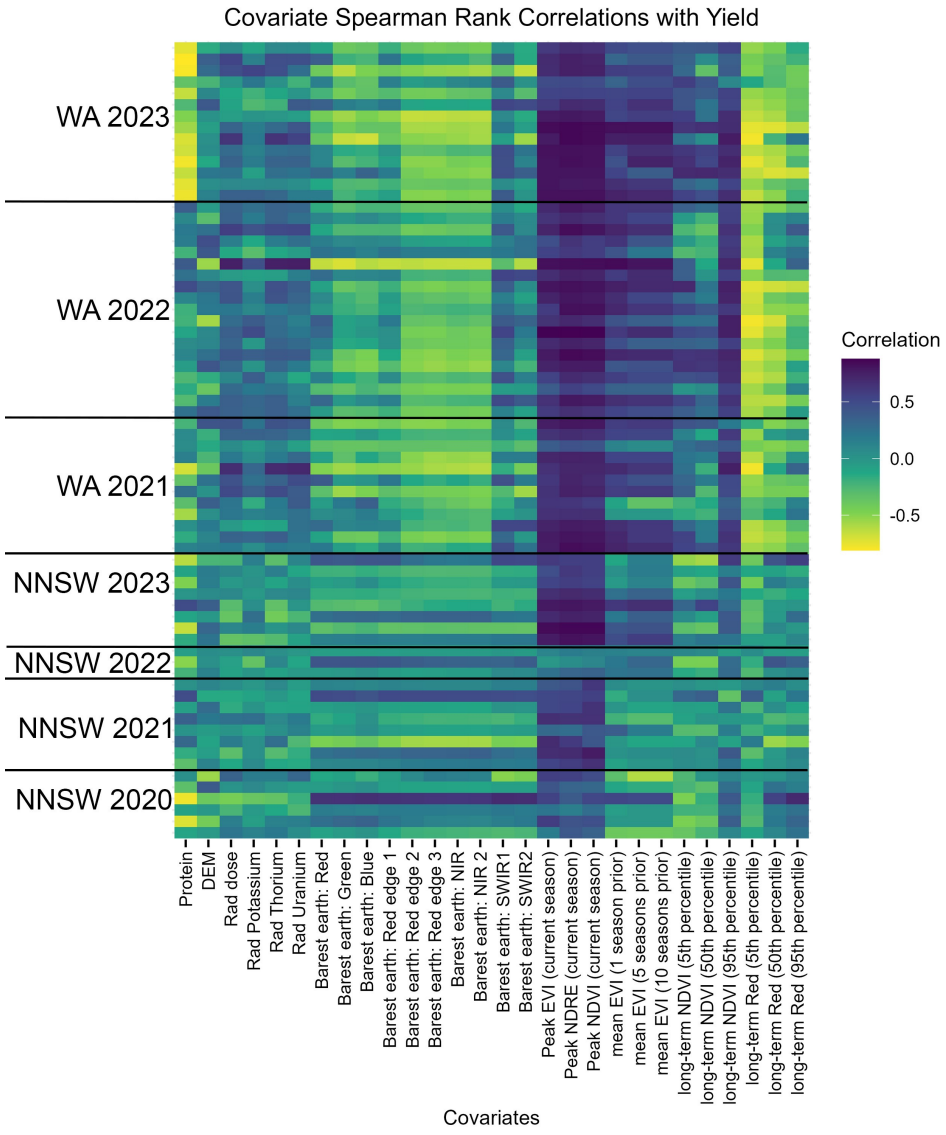

FIGURE 3.A3. Spearman-rank correlations between wheat grain yield and available covariates for each Field-Year in northern New South Wales (NNSW) and Western Australia (WA) for 2020 – 2023 seasons.

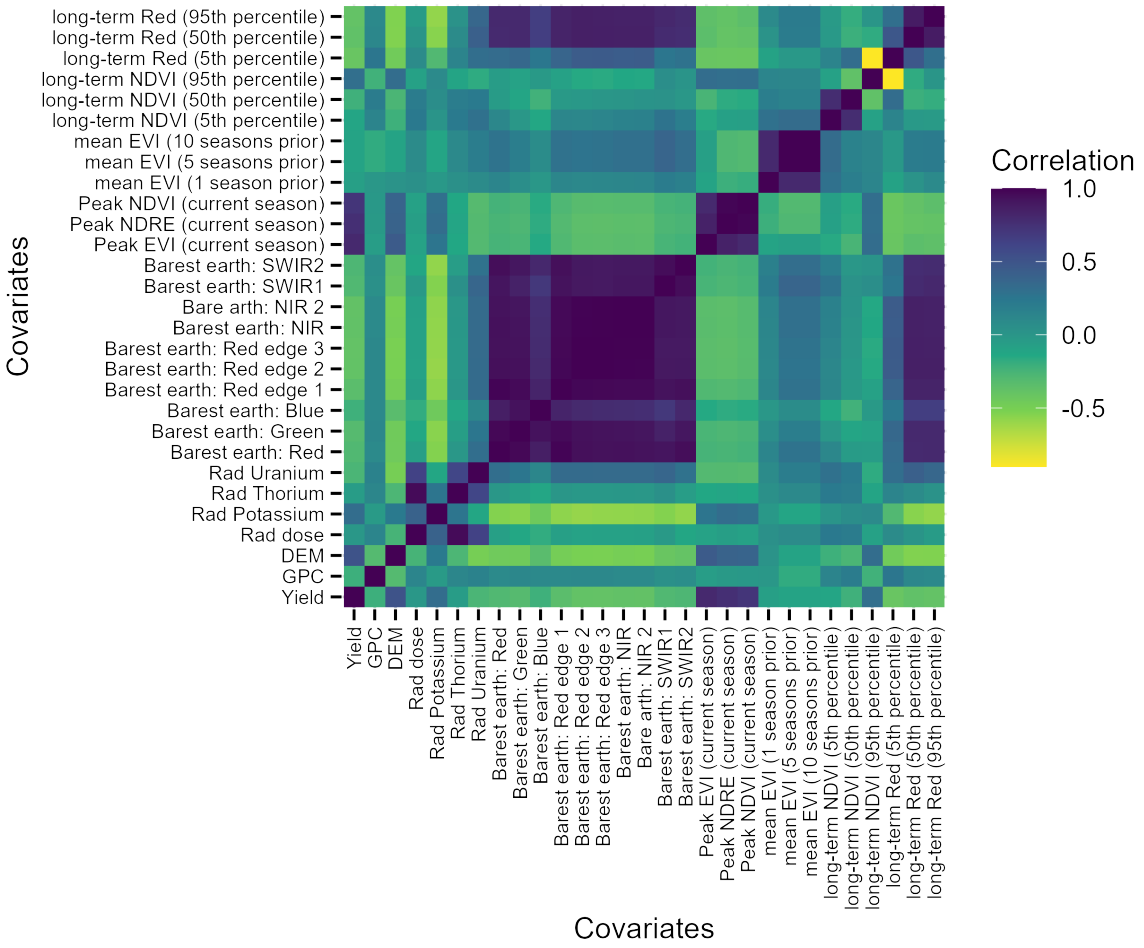

FIGURE 3.A4. Cross-correlations between covariates across both northern New South Wales (NNSW) and Western Australia (WA).

CHAPTER 4

Understanding the spatial drivers of grain protein content for optimal nitrogen management

---

Abstract

Variability in wheat grain protein content (GPC) is driven by complex interactions between genetic, environmental, and management factors, but it can be difficult to differentiate between factors within (e.g. nutrient management) or outside of a grower's control (e.g. seasonal conditions). This chapter applies an interpretive machine learning (IML) approach to understand the drivers of site-specific variability in wheat GPC across a dataset of 25 fields of dryland winter wheat from across the western grain growing region of Australia over the 2022 and 2023 growing seasons, equating to 38 field-years worth of data. First, an exploratory data analysis using a 90 m moving window was conducted to investigate local relationships between wheat grain yield and GPC within fields. Results showed that the relationship between yield and GPC was both positive and negatively correlated within fields, which differs from the expected inverse relationship. Interpretive machine learning was then used to better understand these interactions and the impacts of different variables. For all fields, IML was used to decompose GPC predictions into components driven-by features within (e.g. nitrogen N application) or outside the grower's control [i.e. inherent features including elevation, soil water holding capacity (WHC), and seasonal conditions]. Overall, a higher soil WHC and higher-rates of N increased GPC predictions, whereas high yield values decreased GPC. The impact of elevation on GPC was mixed. For one case-study field where an on-farm N-response experimental trial was conducted, a retrospective evaluation of the amount of N that should have been applied in the just completed season, given the seasonal conditions

experienced (i.e. given the inherent variability outside of the grower's control), revealed that N was overapplied by the grower. An analysis of the site-specific drivers of variability in wheat GPC can help grower's and advisors better understand variability in production systems and assess strategies to manage for this variability to improve profitability and sustainability.

4.1 Introduction

In Australia, grain grower's are paid for both the volume (yield) and quality of grain produced. The grain protein content (GPC) is an important determinant of quality and ultimately the economic value of wheat via the awarding of premiums or discounts. This ideally requires grower's to manage both GPC and yield to achieve premium prices (Stewart *et al*. 2002). However, many grower's don't measure protein on a site-specific basis, despite there being considerable variation in GPC within and between fields, farms, and seasons.

Nitrogen (N) is an important driver of wheat GPC and yield, and constitutes one of the single largest variable input costs for grain production in Australia (Agriculture Victoria 2017). Yet estimating the optimal fertiliser rate remains a challenging task for growers and advisors. There can be considerable spatial and temporal variability in crop nutrient demand, and in supply by the soil. Likewise, there are complex interactions between environmental and management factors both within and outside a grower's control, including the soils water holding capacity (WHC), rainfall, and temperature (Whelan *et al*. 2013). This variability and interactions mean that it is difficult to isolate and differentiate between the impact of N management, or if variability is due to other inherent factors outside of a grower's control.

Under uniform N fertiliser rates and non-limiting soil moisture conditions, an inverse relationship between wheat grain yield and GPC is expected due to the grain protein "dilution effect", where, as the amount of carbohydrates in the grain (yield) increases, the GPC decreases (Simmonds 1995; Terman *et al*. 1969). However, this inverse relationship does not always exist within fields (Whelan *et al*. 2009), and localised variations in the relationship between wheat grain yield and GPC within fields reflects the complex interactions between environmental and management factors throughout the season. The relationship between wheat grain yield

and GPC with variable N rates is non-linear and well-documented (Holford *et al*. 1992; Pan *et al*. 2020). At low N rates, increasing N supply typically increases yield but decreases GPC due to grain protein dilution as carbohydrate accumulation outpaces N uptake. As N supply increases, both yield and GPC may increase together, particularly if N remains available throughout grain fill. However, at high N rates, GPC will continue to increase as excess N is directed towards protein synthesis, whereas yields may plateau or decrease (Holford *et al*. 1992). Water availability and the timing of moisture stress also affects this relationship, which are largely driven by environmental and landscape factors such as rainfall and soil texture. Water stress during grain filling negatively affects grain yield but increases GPC (Campbell *et al*. 1977), whereas relatively high soil moisture can dilute GPC (Simmonds 1989).

Maps of GPC can be used to guide N management. The critical GPC, or the point at which N supply has been optimised for maximum grain yield, is ~12% in most Australian hard white wheats (Scott *et al*. 2012). However, this will vary between varieties and across seasonal conditions, and does not factor-in input costs or economic returns (Colaço *et al*. 2024). Using this simple 'rule-of-thumb' under favourable (non-drought) conditions, a GPC of less than 11.5% suggests that N supply was insufficient for the crop to meet its water-limited yield potential, a GPC of 11.5 – 12.5% suggests that N supply was optimal to achieve the water-limited yield potential, and a GPC of greater than 12.5% suggests that N supply was surplus to crop requirement. Based on this rule-of-thumb, maps of GPC can be used to guide N management for the coming season by identifying areas that had insufficient, optimal, or surplus N for the previous season. More N can be applied to low protein areas, and less N to high protein areas, to achieve a uniform critical GPC of 12% (Scott *et al*. 2012; Wang *et al*. 2023).

The adoption of on-the-go, harvester-mounted grain protein sensors in research and commercial settings is providing grower's, advisors, and researchers with high-resolution maps describing variability in GPC within fields. There is also a growing diversity and abundance of spatial data layers now being collected on-farms [e.g. soil electromagnetic (EM) induction surveys, variable-rate N fertiliser applications], and publicly-available data layers that are accessible across almost every field and farm in the world (e.g. digital elevation models,

DEMs). These spatial data layers can be used to describe variability in factors that influence GPC either directly, including environmental conditions like topography or soil characteristics, or indirectly through proxies that reflect differences in soil properties, nutrient availability, or crop growth conditions.

This growing diversity and complexity of spatial data layers are increasingly being leveraged in machine learning approaches to create crop models that estimate and map variability and account for the complex, non-linear relationships between variables (Liakos *et al*. 2018). Machine learning models have been used for a range of different applications including for the extrapolation of historical crop yields (Donohue *et al*. 2018), or forecasting crop yields (Filippi *et al*. 2020b). However, while these machine learning models excel at handling large, complex datasets, they are often criticised as being "black boxes" (Rudin 2019) because it can be difficult to discern how individual predictor variables contributed to the final model prediction, or the interactions between predictors.

Interpretive machine learning (IML) offers an approach to unpack these black boxes and decompose model predictions of the response (in this case, GPC) into the different components attributed to each feature. SHapley Additive exPlanations (SHAP) values are an IML technique which offer a local explanation of the model by quantifying the contribution of each feature to the prediction. The values are calculated for each prediction point, and the sum of SHAP values represents the deviation of the models predicted output at each prediction point from the field average prediction. Techniques such as SHAP values have been used by Jones *et al*. (2022) and Poole *et al*. (2024) to identify site-specific constraints to yield for a range of broadacre crops, including wheat, cotton, and chickpeas using spatial data layers such as digital-soil maps, proximally-sensed soil surveys, and topographic variables derived from DEMs. There is an opportunity to use IML to understand the drivers of site-specific variability in wheat GPC and decompose the contributions of each predictor variable to the final prediction of GPC. Further, if we can quantify how much these different variables contributed to an increase or decrease in GPC values, we can isolate the impact of N management alone and suggest how N could have been better managed to optimise GPC. Unlike conventional post-harvest assessments of N management using maps of GPC (e.g. Scott *et al*. 2012), IML

enables the impact of N to be isolated from the impacts of other variables outside of a growers control. This means that suggestions for how N could have been applied in the just completed season take into account inherent features in the environment.

This chapter seeks to investigate the drivers of within-field variability in winter wheat GPC using a range of spatial data layers which were collected on-farm or are publicly-available. This research was conducted across a dataset of 25 fields of dryland winter wheat from across the western grain growing region of Australia over the 2022 and 2023 growing seasons, equating to 38 field-years worth of data. This research aimed to:

(1) Explore the spatial variability of local relationships between wheat grain yield and protein content within fields by calculating correlations in a moving window;
(2) Apply IML to understand the site-specific drivers of variability in wheat GPC within fields; and
(3) Use the outputs from IML (i.e. SHAP values) to improve decision making by providing an estimate of the amount of N that should have been applied in the just completed season, given the seasonal conditions experienced.

A retrospective analysis of the site-specific drivers of variability in wheat GPC can help grower's and advisors better understand variability in their production systems. Optimising N management is important for reaching yield potentials and GPC targets. This enables growers to access price premiums and minimise adverse environmental and economic consequences associated with poor N management, thereby enhancing the sustainability of farming systems.

4.2 Methods

4.2.1 Study area

This study was conducted across 25 fields from three broadacre, dryland farms located in Western Australia (WA, Figure 4.1). Farms A and B were located approximately 250 – 300 km north-west of Perth, and Farm C was located approximately 170 km north-east of Perth.

All farms have a Mediterranean climate (Beck *et al*. 2018) with cool, wet winters, and hot, dry summers. Farm A has an annual average rainfall of 422 mm (BOM 2024b), Farm B has an annual average rainfall of 493 mm (BOM 2024c), and Farm C has an annual average rainfall of 339 mm (BOM 2024a).

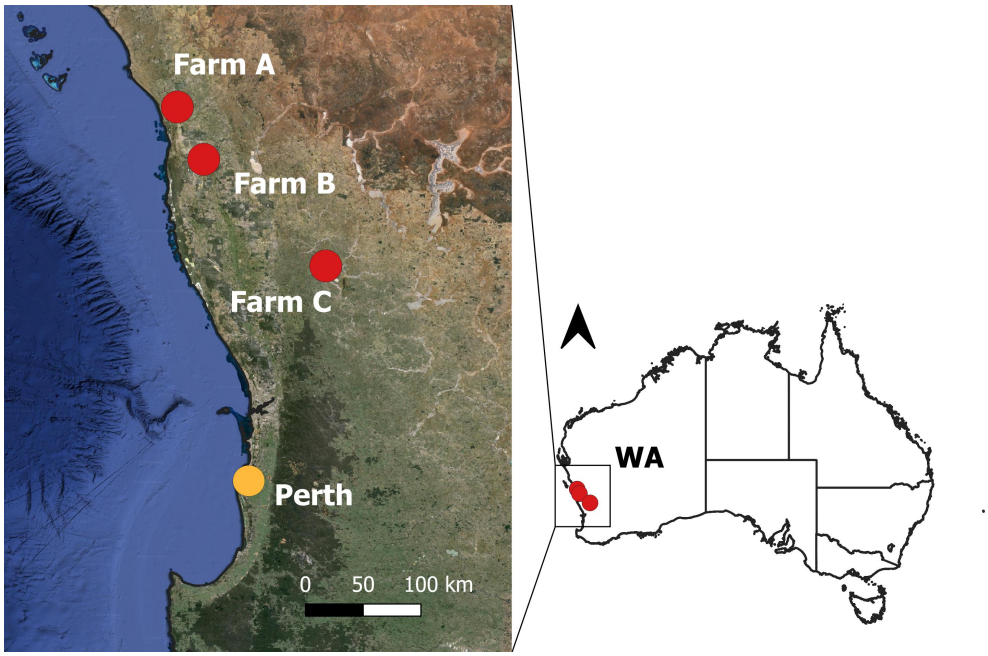

FIGURE 4.1. The study was conducted across three broadacre, dryland farms north of Perth in Western Australia.

4.2.2 Yield and protein data

Wheat grain yield and protein data was collected for each field at harvest for the 2022 and 2023 winter growing seasons using John Deere S790 combine harvesters equipped with the ActiveYield^{TM} automatic yield calibration system, and the HarvestLab 3000^{TM} near-infrared (NIR) spectroscopy sensor for (Tilse *et al*. 2024a). Approximately one measurement was performed per second (Schade *et al*. 2023), with a harvester header width of 12 m.

All yield and protein data for each field and season was cleaned as per the methods from Taylor *et al*. (2007a) and to values recommended by Taylor *et al*. (2007a), Blakeney *et al*. (2009), and Žilić *et al*. (2011). Yield values < 0 and > 10 tonnes per hectare (t/ha), and protein values < 8 and $> 22\%$, were removed. Measurements more than 2.5 standard deviations (SDs)

above and below the field mean were also removed, and concurrent duplicate protein values were excluded based on their time and location to account for data errors during harvest. Each yield and protein map was allocated a unique field-year combination describing the corresponding field and season of data. In total, there were 38 unique field-years worth of data available. For each field-year, the yield and protein data was rasterised onto a 10 m grid using block kriging via the 'gstat' R package (v2.1.1; Pebesma 2004). All data processing and analysis was performed using R Statistical Software (v4.3.1; R Core Team 2021).

4.2.3 Spatial covariates

As described in Section 4.1, the GPC is influenced by a range of genetic, environmental, and management factors including temperature, soil moisture and N availability, variety selection, and the timing of management decisions such as sowing (Whelan *et al*. 2013). Machine learning models can handle complex datasets with many predictor variables, enabling us to capture the impacts of these different factors (Liakos *et al*. 2018; Rudin 2019). However, simple models using fewer covariates are easier to interpret and gain a meaningful understanding of what drives GPC. For this study, four features were used to model wheat GPC (described in Section 4.2.6): total applied N, the soils apparent electrical conductivity (EC_a), elevation derived from a DEM, and yield. These features were chosen as they are easy to interpret and make agronomic sense as they can (in)directly describe the factors influencing variability in GPC, including water flow and accumulation (via elevation), soil water holding capacity (EC_a), nutrient supply (via total applied N), and the cumulative effects of all other seasonal conditions experienced by the crop, such as management decisions, weather, and inherent soil properties (i.e. naturally occurring characteristics that are not easily altered, such as soil depth), as represented by yield.

Every field-year used in this study had variable-rate N applied. This is common across the western growing region, but not widely practiced across the rest of Australia (Bramley *et al*. 2019). As-applied variable-rate urea maps (i.e. maps of the measured weight of urea applied during the season) for each field were used to calculate total applied N for each season (as

urea = 46% N). The variable-rate applications were determined by the grower/advisor for each farm.

The soils EC_a, which was measured and mapped across each field with an EM induction survey, was used as a proxy to describe relative variability in soil texture and WHC, based on prior knowledge and soil surveys of the farm (Wang 2024). In these fields, there is an absence of soil salinity and it is known that higher soil EC_a values represent higher clay content and thus, a higher soil WHC.

The Geoscience Australia Shuttle Radar Topography Mission (SRTM) derived 1 arc second DEM (Gallant *et al*. 2009) was accessed at a 30 m spatial resolution via the 'dataharvester' R package (v0.1.2; Harianto *et al*. 2023). The DEM was used to describe elevation change within each field and represents the patterns of water flow and accumulation. While high-resolution elevation data can be obtained from EM induction surveys, this high-resolution elevation data was not available for the fields used in this study.

While both yield and GPC are influenced by N applications, yield was used as a proxy for the cumulative effects of all seasonal conditions the crop experienced, including management decisions, weather, and inherent soil properties. In each field and for each season, the crop was managed by the grower/advisor to optimise yield. While weather conditions like temperature or rainfall are important drivers of GPC, the spatial resolution of these data layers is too coarse to use for a within-field analysis of variability. Likewise, management data like variety selection is uniform within fields, or was not available. The impact of weather, management decisions, and inherent soil features should be represented in the yield data. As the objective of this study was to differentiate sources of variability in GPC between N availability and all other factors, further interpretation as to how different factors interact and contribute to the observed yield was not required.

Elevation, total applied N and EC_a values were then extracted onto the same 10 m grid used for the yield and protein data to form a single dataframe/datacube.

4.2.4 On-farm N-response experimental trial

To generate a GPC response to a range of N fertiliser rates, an experimental N-response trial was conducted in one case-study field for the 2023 growing season. Adjacent strips of high [225 kg urea/hectare (ha)] and zero (0 kg urea/ha) N rate treatments were applied within each field as a mid-season top-up of urea, based on the methodology used by Colaço *et al*. (2024). The strips were positioned to run across management zones to capture variability in crop response, which were determined via k-means clustering using available spatial data layers including previous yield maps, gamma radiometric soil survey data, and the EM induction survey and DEM described in Section 4.2.3. The strips were 36 m wide, equivalent to three combine-harvester widths, and the field was managed consistently according to normal farmer practice in all aspects except for the mid-season urea top-up: one strip received a high rate of 225 kg urea/ha, while the other received no urea (0 kg urea/ha). All other management practices remained consistent across the field.

4.2.5 Moving window correlations and exploratory data analysis

Local correlations between wheat grain yield and GPC in a 90 m moving window were mapped across a 10 m grid within each field to investigate their relationship, based on the approach by Whelan *et al*. (2009). A 90 m window was chosen to smooth fine-scale noise while capturing meaningful spatial patterns at a scale that balanced computational efficiency with interpretability. This analysis was performed for every pixel in the dataset, resulting in a unique correlation value at each 10 m location. Identifying these spatial patterns is important for Site-Specific Crop Management (SSCM), as it helps reveal consistent zones where yield and protein are closely linked and can be used to guide targeted management interventions.

The moving window correlations were aggregated into broader categories (i.e. -1 to -0.75, -0.75 to -0.50, -0.50 to -0.25, -0.25 to 0, 0 to 0.25, 0.25 to 0.50, 0.50 to 0.75, and 0.75 to 1) to simplify visual interpretation of spatial patterns in the maps and to enable consistent comparison of correlation strength across fields using summary statistics (e.g. the proportions of each field-year covered by each bin). A positive yield-protein relationship represents a

general trend where yield and protein increase (or decrease) together, whereas a negative relationship represents a general trend where, as yield increases, protein tends to decrease, and vice versa.

4.2.6 Grain protein content modelling

An eXtreme Gradient Boost (XGBoost) model was used to create a predictive model of GPC for each field using the 'xgboost' R package (v1.7.7.1; Chen *et al*. 2020). An XGBoost model was used as they have proved valuable for modelling grain yields (Jones *et al*. 2022; Poole *et al*. 2024; Filippi *et al*. 2024b). XGBoost models also enable a simple calculation of SHAP values via the 'SHAPforxgboost' R package (v0.1.3; Liu *et al*. 2020). Total applied N, soil EC_a, elevation derived from a DEM, and yield were used as features within the model. An individual model was created for each field-year, so that the spatial drivers of variability could be assessed on a field-specific basis due to differences in management history and weather between fields and seasons. A 10-fold cross-validation (10FCV) was then used to assess the ability of each XGBoost model to predict GPC for each field. Model quality was assessed using the Lins concordance correlation coefficient (LCCC; Lin 1989), which is a measure of the agreement between observed and predicted values in relation to the 1:1 line where 1 is a perfect fit. The root mean square error (RMSE) was also used, which is a measure of the difference between observed and predicted value where the smaller the RMSE, the better.

4.2.7 Interpretive machine learning

For each model, SHAP values, which are an IML technique, were calculated using the 'SHAPforxgboost' R package (v0.1.3; Liu *et al*. 2020). SHAP values provide a local explanation of the model's prediction by quantifying the contribution of each input feature to the predicted output at each location. In this framework, the predicted GPC at each location is expressed as the sum of the models baseline prediction (i.e. the average predicted GPC across the field) and the SHAP values of all input features at that location (Equation 4.1).

$$Predicted\,GPC = \mu + \sum_{i=1}^{n} SHAP_{i,x} \tag{4.1}$$

Where:

μ = mean prediction of GPC across the field,

$SHAP_{i,x}$ = SHAP value for feature i at location x,

n = the total number of input features in the model.

Each SHAP value reflects the direction and magnitude of the features influence on the prediction, indicating whether it contributed to a higher or lower GPC compared to the field average. Together, these values allow of the interpretation of how different factors contribute to GPC variability within a field.

4.2.8 N recommendation using SHAP values

In this chapter, GPC values were used to estimate the site-specific N rates that would have been required to optimise GPC (i.e. reach a uniform GPC target of 12%) during the just completed season. This is a retrospective evaluation to better understand how variability could have been better managed, and this approach did not seek to make recommendations for future seasons. By calculating SHAP values for each input feature in the prediction model, the GPC prediction was decomposed into two components:

(1) The contribution from total applied N; and
(2) The combined contribution from all other features, defined as the inherent GPC (Equation 4.2).

The inherent GPC refers to the portion of GPC explained by non-N features, including spatial variability in soil properties, topography, management history, and environmental effects. For the purpose of this study, these factors are considered to be outside the grower's control and represent the natural or non-N-related management influences on GPC. The inherent GPC

was calculated as the sum of the field mean GPC prediction and the SHAP values for the non-N features (yield, elevation, EC_a; Equation 4.2).

$$Inherent\,GPC = \mu + SHAP_{Yield,x} + SHAP_{Elevation,x} + SHAP_{ECa,x} \tag{4.2}$$

Where:

$SHAP_{Yield,x}$ = SHAP value for yield at location x,

$SHAP_{Elevation,x}$ = SHAP value for elevation at location x,

$SHAP_{ECa,x}$ = SHAP value for ECa at location x.

To estimate the amount of N could have been applied to achieve a uniform GPC of 12% for the just completed season, the inherent GPC deviation was first calculated for each prediction location. The inherent GPC deviation was calculated as the difference between the inherent GPC and the target value of 12% (Equation 4.3). This represents how far the GPC was from a target of 12% in the absence of an N effect. The deviation indicates whether the GPC was above or below the target and the amount that would have needed to be corrected through N fertilisation.

$$Inherent\,GPC\,deviation = target\,GPC - inherent\,GPC \tag{4.3}$$

Where:

$target\,GPC$ = 12%

A spline function with four degrees of freedom was fitted between the actual applied N rates (which was used as an input feature in the original model) and their corresponding N SHAP values to characterise how incremental changes in applied N influenced predicted GPC across the field, according to the model. The spline was implemented using the 'splines' R package (R Core Team 2021), and four degrees of freedom were used to capture the non-linear relationship between changes in applied N and predicted GPC. This relationship was then

applied to the spatial map of inherent GPC deviation to determine the total N rate that could have been applied at each location at achieve the target GPC of 12%, given the environmental and management conditions of the previous season. The result is a retrospective, site-specific map of optimal N requirements.

It is important to note that this N recommendation is retrospective as it is based on previous seasons data. This provides an estimate of how much N should have been applied to optimise GPC, rather than a recommendation for current or future seasons.

4.3 Results and discussion

4.3.1 Exploratory data analysis

Correlations between wheat grain yield and GPC were calculated at different spatial extents, namely dataset-wide, on a field-basis, and using 90 m moving windows. Yield and GPC were negatively correlated (-0.54) across the entire dataset. When the correlation was calculated for each field, ~80% of field-years had a negative correlation between yield and GPC, and over 35% of field-years had a negative correlation between -0.75 and -0.50. Overall, correlations between yield and GPC at the field-year-level ranged between -0.90 and 0.55. For the eight field-years that had positive correlations between yield and GPC, both yield and GPC values tended to fall closer to the middle of their respective ranges. In contrast, negatively correlated field-years tended to have either relatively high or relatively low GPC. It is suspected that in fields where GPC was very high, soil moisture was likely a limiting factor as yields were very low, whereas in fields where GPC was very low, N was likely limiting as yields were high (Simmonds 1995; Holford *et al*. 1992).

When the relationship between yield and GPC was calculated using a 90 m moving window correlation within each field, the relationship varied in strength, direction, and spatial structure. An example field is provided in Figure 4.2. Within fields, negative correlations dominated and more than three-quarters of the field-years had negative correlations between yield and GPC that covered at least 50% of the within-field area (Figure 4.3). While this largely supports

the expected inverse yield-protein relationship (Simmonds 1995), these findings also align with those from Whelan *et al*. (2009) who found that the relationship between yield and GPC varied considerably within fields. Overall, every field-year had at least some area that displayed a positive correlation between yield and GPC.

These inverse relationships between yield and GPC are likely the result of a soil moisture limitation, as these visually corresponded to areas of low yields and high GPC (Figure 4.2). However on their own, these moving window correlation maps do not provide insight into the nature of the yield-protein relationship. This inverse relation may reflect either a low-yield/high-protein scenario, due to increased N supply in water-limited conditions, or a high-yield/low-protein scenario, due to increasing N-supply when soil moisture is not limiting. Adequate soil moisture and optimal N supply in areas likely resulted in positive yield-protein relationships, but both yield and GPC may have been very low in these areas as well (Terman *et al*. 1969; Whelan *et al*. 2009).

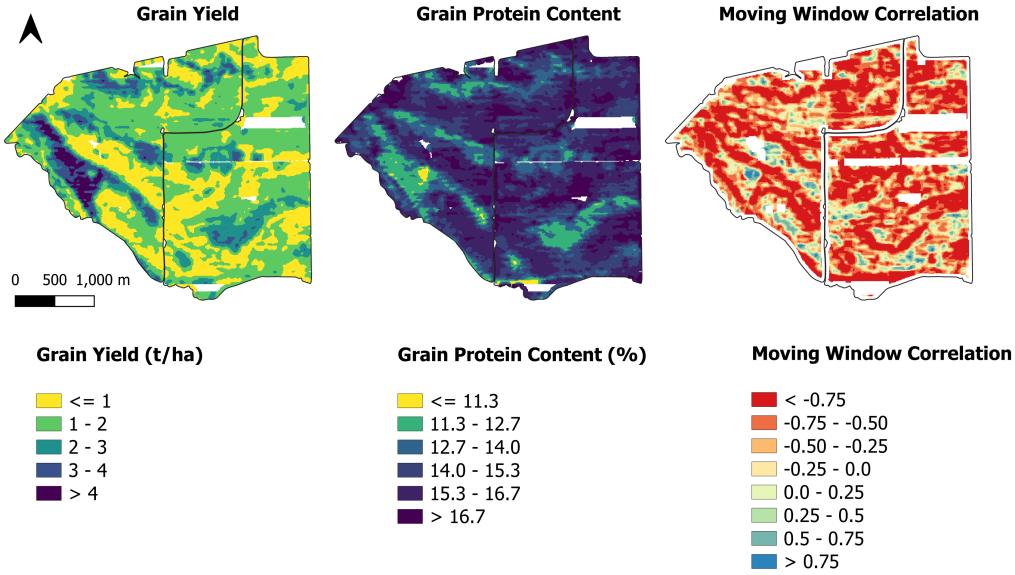

FIGURE 4.2. Observed wheat grain yield (tonnes per hectare, t/ha), grain protein content (%), and 90 m moving window correlations for an example field.

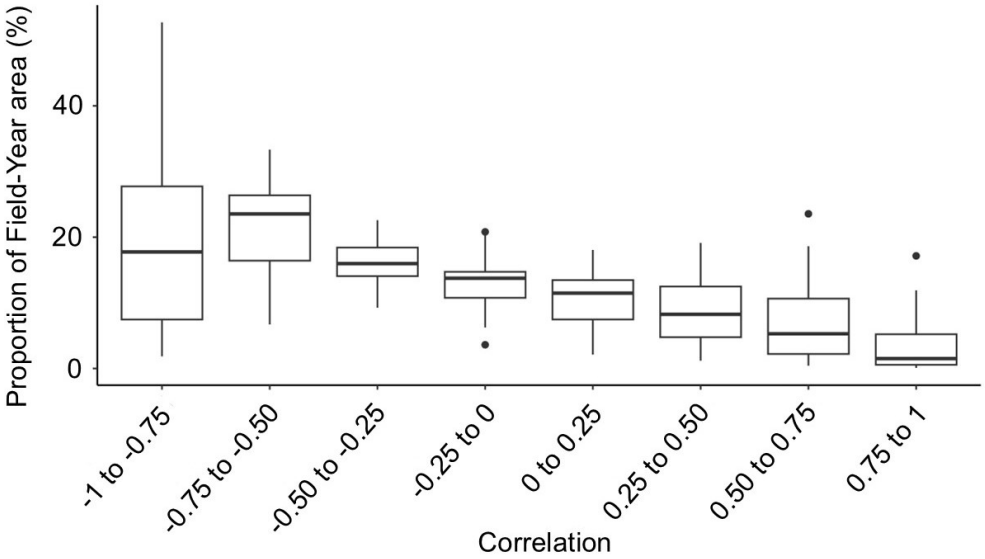

FIGURE 4.3. Proportion (%) of each Field-Year area covered by yield-protein moving window correlation categories for all Field-Years.

These moving window correlation analyses have revealed considerable variability in the relationship between wheat grain yield and GPC within fields, but this simple approach fails to capture how different variables interact and contribute to the relationship between yield and GPC at discrete locations. This is explored in Sections 3.2 and 3.3.

4.3.2 Grain protein predictions and model quality

An XGBoost model was used to model wheat GPC within fields. Overall, results from the 10FCV of the XGBoost model for GPC prediction produced an RMSE of 0.30%, and an LCCC of 0.98 when all field-years were compiled. When considering individual field-years, RMSE values ranged between 0.14 and 0.38%, and LCCC values ranged between 0.86 and 0.95 (Figure 4.4). It should be noted that the aim of this modelling approach was to describe the spatial patterns of GPC within fields, not to develop predictive model for forecasting. As such, the use of 10FCV was appropriate for evaluating model performance in capturing spatial structure, but the reported prediction accuracy should not be interpreted as indicative of forecasting utility (Filippi *et al*. 2025).

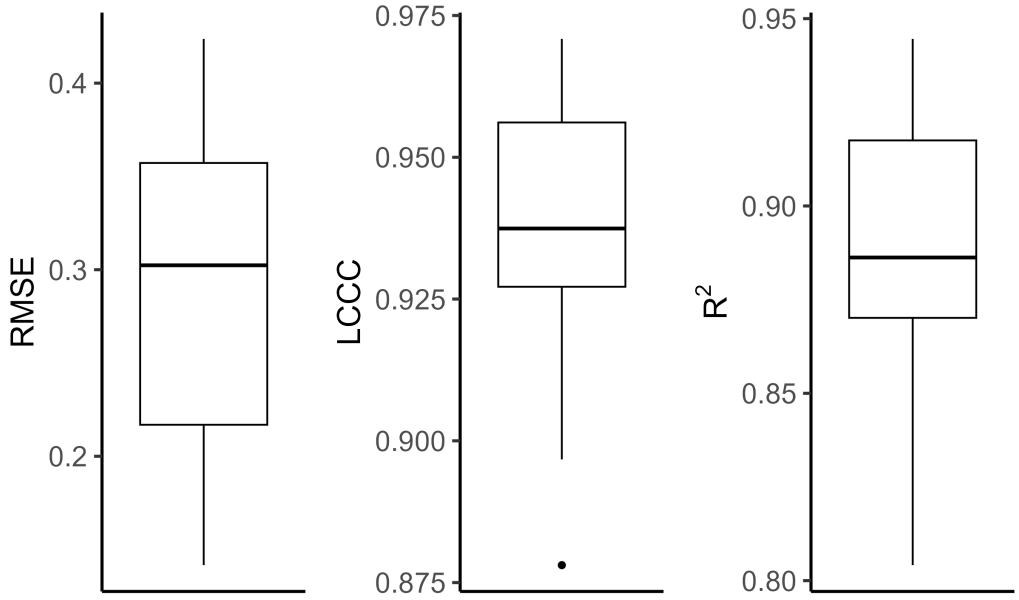

FIGURE 4.4. Boxplots showing distribution of 10-fold cross-validation (10FCV) for each field-year, including the root mean square error (RMSE), Lin's concordance correlation coefficient (LCCC), and r-squared (R^2).

4.3.3 Interpretive machine learning

4.3.3.1 SHAP summary plots

The SHAP summary plots describe the ranked feature importance and feature value (based on mean absolute SHAP value), and feature effects (deviation from field average prediction). The units are the same as for GPC (%). An individual model was developed and interpreted using IML for each field-year, and the results were then aggregated to provide an overall analysis. Overall, elevation was the highest ranked feature based on the mean SHAP value, followed by EC_a, yield, and total applied N (Figure 4.5). Between field-years, different features had different rankings. When the highest-ranked features were assessed within each field, yield consistently emerged as the top feature (i.e. with the highest mean SHAP value), followed by elevation, EC_a and total applied N (Figure 4.5). The relationship between yield and GPC varied considerably within and between fields (Section 4.3.1), but the SHAP plots for the

whole dataset (Figure 4.5a) and for an individual case study field (Figure 4.5b) showed that higher yield values reduced GPC.

While the mean feature importance value was low for total applied N, low rates of total applied N had the strongest decrease in GPC in some areas of the case-study field (Figure 4.5b). This localised importance of total applied N was observed across the whole study area (Figure 4.5a). This highlights the advantage of local feature importance analyses such as SHAP, as overall mean importance value misses this. Soil moisture and the capacity of the soil to hold moisture are key drivers of variability in GPC, and elevation influences how water moves and ponds within a field. Likewise, EC_a was used as a proxy for soils WHC, as higher WHC values are associated with increased clay content in these fields (Wang 2024) and is known to correlate with higher WHC. While the influence of elevation and EC_a at discrete locations was less clear than for yield or total applied N, it is clear that both features are important drivers of GPC overall.

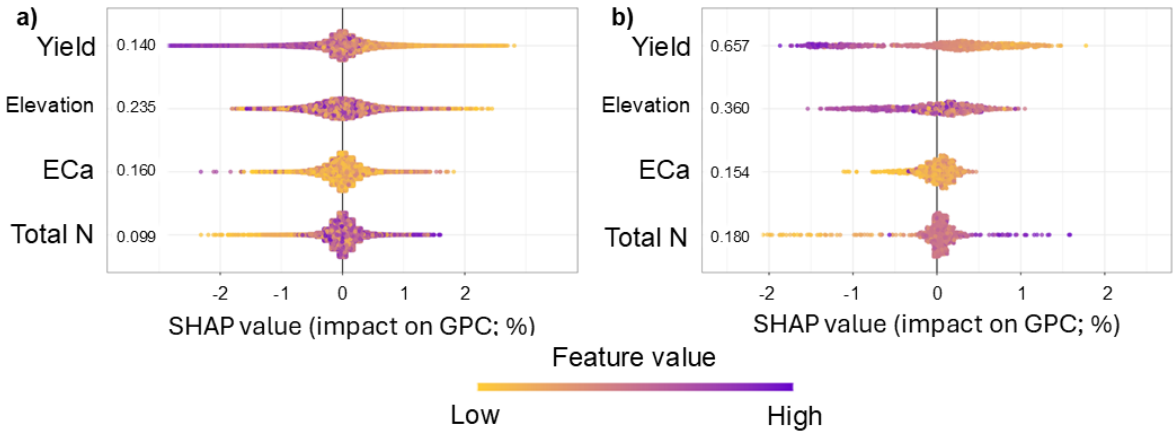

FIGURE 4.5. SHapley Additive exPlanations (SHAP) summary plots from the eXtreme Gradient Boosting (XGBoost) GPC model for a) the whole study area; and b) one case-study field. For each feature, the position of each discrete location on the x-axis reflects the SHAP value, which represents the feature effect on wheat grain protein content (GPC) at that particular location. The SHAP values are in the same units as GPC (%). The colour indicates the feature value from low to high on a relative scale. The value on the y-axis reflects the mean absolute SHAP value for each feature.

4.3.3.2 Feature SHAP maps

SHAP maps, which visualise the contribution of each feature to the GPC prediction for each point within the field, are presented for the case-study field where the on-farm N experimentation trial was conducted (Figure 4.6). Positive SHAP values indicate that a feature caused an increase in GPC at a location relative to the field mean, whereas negative SHAP values indicate a decrease in GPC. Overall, higher yields, lower total applied N, and lower soil EC_a, which is a proxy for the soils WHC, all had negative SHAP values and corresponded to lower GPC predictions, whereas areas with lower yields, higher total applied N, and higher soil EC_a had positive SHAP values and corresponded to higher GPC predictions. These feature contributions make sense and further reflect non-linear grain protein and yield interactions described in Sections 4.1 and 4.3.1 (Pan *et al*. 2020; Holford *et al*. 1992).

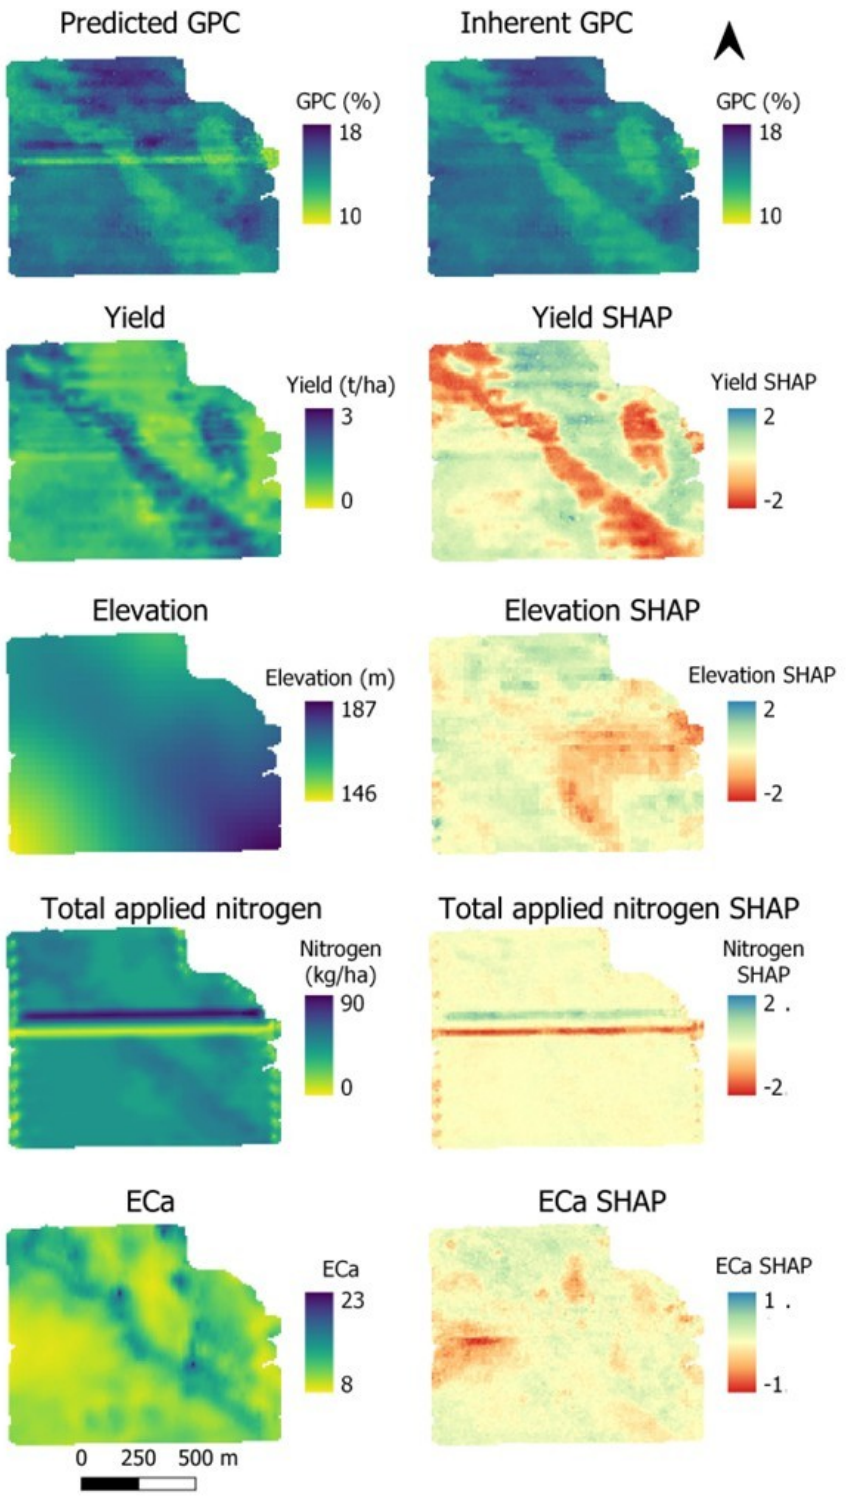

FIGURE 4.6. Observed feature values and their corresponding SHapley Additive exPlanations (SHAP) value maps within a case-study field, predicted grain protein content (GPC), and the inherent GPC variation (i.e. variability in GPC not including nitrogen, N; due to features outside of the growers control).

Given the variability in both yield and GPC observed along the zero-N strip (Figure 4.6), and that GPC values were close to the optimum of 12%, residual soil N may have already been sufficient in some areas of the field to achieve the optimum protein target. However, without soil N data, the extent of spatial variability in N supply remains unclear. Future studies should include pre- and post-season soil N tests to better understand the total available N.

While the interactions between yield, GPC, soil moisture, and available N are well-documented within the literature (e.g. Pan *et al*. 2020; Holford *et al*. 1992), this study is novel in that it demonstrates how existing spatial data layers that are already being collected on farms can be used as part of an IML approach to reveal and better understand GPC variability within fields. Unlike traditional N-response curves, which often assess N application effects under controlled conditions, this approach accounts for complex interrelationships between variables in real-world farming systems. Spatial data layers describing the soils WHC, elevation and yield were used to describe variability derived from inherent factors beyond the grower's control. The SHAP values were then used to decompose GPC variability between these inherent factors, and N which can be managed.

In terms of what is within the grower's control, the total applied N SHAP map showed values ranging from –2.5% and 2.5%. This indicates that this is the range of values over which GPC can be changed by varying applied N rates. In contrast, the inherent GPC, which was calculated as the sum of the field mean GPC prediction and the SHAP values for the non-N features (yield, elevation, EC_a; Equation 4.2), represents how factors outside of the grower's control, including soil and weather conditions experienced during the season, caused GPC to vary.

4.3.4 Applications for management decisions

Nitrogen is a key driver of variability in wheat GPC and is an important management input that can be adjusted by the grower to achieve an optimum GPC target (Wang *et al*. 2023). However, it can be difficult to differentiate between variability that can be addressed via N management, or if variability is due to inherent factors outside the grower's control. Given

that SHAP values decompose model predictions into different components attributable to each input feature, variability attributable to inherent (i.e. yield, elevation, soil WHC) and manageable (i.e. applied N) features could be differentiated to provide a retrospective analysis of the amount of N that should have been applied to achieve a uniform protein target of 12% for the season just finished, given the experienced seasonal conditions. This is shown in Figure 4.7 which shows how far each location is from 12%. It should be noted that any target can be used, but 12% was used here because a critical GPC of ~12% represents the point where N supply is optimised for maximum grain yield, with values below 11.5% indicating insufficient N supply and values above 12.5% suggesting surplus N relative to crop requirements (Scott 2022). For this field, due to inherent factors outside of the grower's control, GPC predictions were almost all greater than the optimum of 12% and deviated above the optimum by up to 5.2%. This is likely due to low yields as a result of limited moisture.

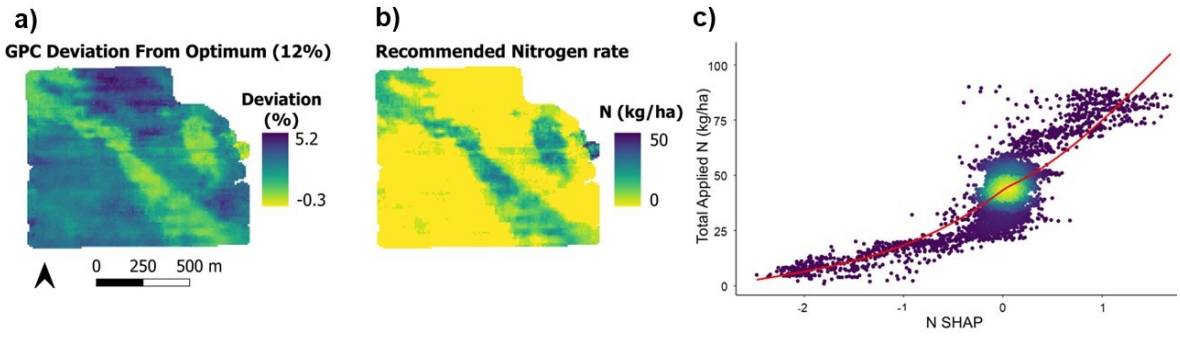

FIGURE 4.7. For the case-study field, maps describing a) the inherent grain protein content (GPC) deviation from a target GPC of 12%; b) a map of the retrospective recommended nitrogen (N) rate for the previous season to achieve a uniform GPC target of 12%, and c) the relationship between total applied N and N SHAP values, modelled using a spline.

To establish how much N should have been applied to achieve a target GPC of 12%, applied N was modelled against the N SHAP value with a spline which is shown in Figure 4.7c. The recommended N rate is presented in Figure 4.7b. Here, the lowest SHAP value is ~-2.5% at 0 N rate, which represents the most that N management could have lowered the GPC. This shows that 60% of the field should have received a 0 N rate, and implies that there was sufficient residual N in the soil for that season. However, soil tests were not conducted

and the actual available N in the soil is unknown. Recommended N rates from the IML approach ranged between 0 and 50 kg N/ha and averaged ~12 kg N/ha for the case-study field. Excluding areas where the experimental N-strip trials were placed, the N rates applied by the grower were on average ~30 kg N/ha higher than the IML retrospective recommendation.

Overall, throughout large parts of the case-study field, N was overapplied by the grower compared to the retrospective N recommendation made using this IML approach. Given relatively low yields and high GPC values throughout much of the field, it was assumed that N was not the limiting factor and that high GPC values were likely due to a lack of available soil moisture, and this was confirmed by the grower. Therefore, the higher rates of N that were applied by the grower compared to what are provided in this retrospective recommendation likely would not have increased yields, and rather, increased input costs without a return on investment. Reducing and reallocating N inputs to achieve a uniform protein target across a field can improve economic returns by reducing unnecessary expenditure on N fertiliser, and improve on-farm efficiencies by reducing the need for blending high and low protein grain to reduce discounts associated with poorer quality. However, an economic analysis was not conducted here, but should be incorporated into future research (e.g. Colaço *et al*. 2024).

While a map of GPC can be used to understand N adequacy for the season just completed (Scott 2022), variability in GPC may be attributable to many factors. This IML approach differentiates between variability attributable to features within and outside of the grower's control, provides an understanding of the site-specific drivers of variability for the previous season, and can be used to provide a retrospective recommendation of N that should have been applied to achieve a uniform protein target. While multi-year data is not currently available for this case-study field, if more data is collected for future seasons it may be possible to apply this IML approach to differentiate temporally-stable drivers of variability (e.g. soil texture) from those that vary season-to-season (e.g. rainfall and available soil moisture). In the absence of water-limitations, it is likely that GPC would more directly reflect the effects of applied N. In the future, multiple seasons worth of data for the same field could be used to better understand temporal variations in N fertiliser management strategies in response to differences in growing season precipitation patterns and different water-availability scenarios.

It may be possible to apply this methodology to forecasts of GPC prior-to or during the coming/current season (e.g. Chapter 3, Filippi *et al*. 2019b) to identify potential drivers of variability and develop N recommendation strategies based on seasonal conditions expected. This IML approach could also be generalised to understand the site-specific impact of other agronomic management practices such as defoliation strategies in cotton (e.g. Chen *et al*. 2022; Bryce 2025).

The implementation of on-farm crop response trials is necessary to generate variation in the GPC response to N. Another way to generate retrospective N recommendations is to model N-response curve using a local regression analysis with adjacent N-strips (e.g. Colaço *et al*. 2024) but an additional machine learning step is still necessary to extrapolate the response function for the rest of the field. In this approach we only use a machine learning model and do not have to model the N-response. The use of IML, specifically SHAP values, to derive a retrospective N recommendation eliminates the need to place N-strips adjacent to one another as the relationship between the N feature SHAP values and the total applied N rate is modelled using a spline, which is non-spatial. Here, it is assumed that the relationship between applied N and the N SHAP value is global for the field, but in the future, local response functions could be generated using moving windows to model the relationship between applied N and N SHAP values.

In the future, it may be possible to conduct a similar data-driven IML approach to understand the drivers of local yield-protein relationships, including why yield and GPC are positively correlated in some areas of a field and negatively correlated in others. Whelan *et al*. (2009) found that these local changes in the relationship are not random, and that the spatial pattern in the local relationship between yield and protein has an autocorrelation range of the same magnitude as yield and GPC individually. The exploratory results from this study (Section 4.3.1) suggest that IML could be used to better understand the drivers of variability in both yield and protein together within fields. However, a negative correlation can reflect a low-yield/high-protein scenario, or a high-yield/low-protein scenario. More work is needed to understand how IML could be implemented to better unpack the interactions between yield, protein, and different model features.

4.4 Conclusion

This study demonstrated that IML can be used to understand the site-specific drivers of variability in wheat GPC within-fields through an analysis of 38 field-years worth of data from the western grain growing region of Australia. The IML approach could be used to differentiate between variability attributable to features within and outside of a grower's control, and this study also goes further and extends the use of IML by providing a retrospective estimate of how much N should have been applied to achieve a uniform GPC target for the season just completed. This is a novel use of IML as it is applied to GPC to understand variability, uses the SHAP value composition to derive a management decision. Within-fields, local relationships between yield and GPC were not always negative, which is both contrary to their expected relationship but was also not surprising when considering other literature. Overall, from the IML analysis, higher-rates of N positively impacted GPC predictions, whereas high yield values had a negative impact, and the IML approach recommended a lower N-rate than what was applied by the grower in the case study field. In the future, there is an opportunity to utilise this IML approach over multiple seasons to understand spatial and temporal drivers of GPC, improve N recommendations by building more localised N-response models, and generalise this approach to other management options, such as variable-rate defoliation in cotton. From this understanding of the site-specific drivers of variability in GPC, management decisions can be tailored to optimise both yield and quality.

CHAPTER 5

Quality vs. quantity - does wheat grain protein content or yield have a greater opportunity for precision management?

Abstract

Fields characterised by large within-field variation stand to potentially gain from the differential application of inputs via Site-Specific Crop Management (SSCM). Maps generated from wheat grain yield and protein sensors are revealing considerable spatiotemporal variability within and between fields, farms, and seasons. Yet one of the biggest barriers to the successful adoption of SSCM is an assessment of the magnitude and spatial distribution of variability, and whether this is sufficient to warrant a change from uniform management. The Opportunity Index (OI) is a measure of the magnitude and structure (i.e. spatial cohesion) of within-field variability. The OI ranks fields as to their suitability for SSCM, so that fields with higher OI values are more amenable to SSCM approaches. This research used a database of 86 paired wheat grain yield and protein maps from 70 fields, collected over four seasons (2020 – 2023) across two farms in northern New South Wales (NSW) and four farms in Western Australia (WA), resulting in 86 field-years of data. This study aimed to compare opportunities for SSCM SSCM wheat grain yield and grain protein content (GPC) using the OI. The OI was calculated for every yield and protein map, and were then classed as having relatively low, medium, or high opportunities for SSCM. To assess which had a greater opportunity within each field, pairwise differences between the OI for yield and GPC were also calculated. Overall, results showed that there was a greater opportunity for SSCM for wheat grain yield than for GPC. While both yield and GPC had a similar spatial structure, yield had a greater magnitude of variation within fields. While the OI was greater in NSW than in WA, this is

likely due to variable-rate fertiliser applications already being employed in WA, but not in NSW. In calculating the OI across a farm or region, the OI ranks fields as to their suitability for SSCM and provides growers with a simple decision-making tool to determine where SSCM practices may be best directed for the stepwise adoption of precision agriculture. However, it does not provide recommendations for management options, and further investigations by the grower/advisor are needed to make SSCM decisions. Future work should include an economic analysis of the opportunities for SSCM with respect to the premium and discount system applied for grain quality in Australia, and more work is needed to better understand the drivers of the differences in the OI between yield and GPC within and between fields, farms, and seasons.

5.1 Introduction

Site-Specific Crop Management (SSCM) is the practical application of precision agriculture (PA) principles and involves the allocation of resources and agronomic practice to match spatiotemporal variability (Whelan *et al*. 2013). The application of SSCM is more advantageous in fields that have greater heterogeneity as the differential application of inputs has a greater potential to improve crop yields, optimise quality, and reduce environmental and economic costs (Ali *et al*. 2024). The Australian grains industry was among the earliest adopters of SSCM, and grain growers remain at the forefront of implementing and adopting new technologies and decision-making strategies (Anderson *et al*. 2016). As more and more data are being collected on-farm, we are continuing to capture and better understand more about the spatiotemporal variability being observed within fields. However, it can be challenging for growers and advisors to know what to do with this data, what data layers are most suitable for guiding management decisions, and where to best direct management strategies to achieve the greatest impact (either economically, environmentally, or logistically). Overall, one of the biggest barriers to the successful adoption of SSCM is an assessment of the magnitude and spatial distribution of variability, and whether this is sufficient to warrant a change from uniform management (Pringle *et al*. 2003).

Yield monitors have been commercially available for more than 20 years (Cook *et al*. 2001) and are widely used across the Australian grains industry (Bramley *et al*. 2019). Yield monitor data has revealed considerable variability within fields (Robertson *et al*. 2008; Whelan *et al*. 2009), and PA research has largely focused on understanding spatiotemporal variability in yields due to the large number of available yield maps. On the other hand, the uptake of grain protein sensors has historically been slow and much of the early research to understand variability in grain protein content (GPC) within-fields was limited to hand samples and a small number of protein maps, often for a single season (e.g. Skerritt *et al*. 2002; Stewart *et al*. 2002).

In Australia, the value of wheat is closely tied to the GPC, which is classified into grades and directly affects the prices that growers receive. Rising input costs (e.g. fertilisers; Hogan 2024) and volatile climatic conditions and export markets (Muleke *et al*. 2022) means that growers are increasingly interested in understanding variability in quality (i.e. GPC) and opportunities to leverage this for improved management. As a result, the adoption of grain protein sensors is slowly increasing and the commercial release of new on-the-go sensing platforms in recent years (e.g. HarvestLab 3000^{TM} from John Deere; Schade *et al*. 2023) are adding to the growing database of available protein maps. An analysis of 63 paired yield and protein maps collected from farms across Western Australia (WA) and northern New South Wales (NSW), Australia, has revealed considerable variability in wheat grain yield and GPC within fields and between seasons, and that yield and GPC generally do not vary in the same way, to the same magnitude, or with the same spatial structure within-fields (Chapter 4; Tilse *et al*. 2024a). This variability presents a challenge for growers and advisors when managing for both quality and quantity, but there is a need for further exploratory work to describe, compare, and understand variability in both yield and GPC for improved management.

Simple metrics (e.g. the coefficient of variation, CV, or standard deviation, SD) have been used to quantify variability in wheat grain yield and protein content within and between fields (e.g. Kelly *et al*. 2003), but these are non-spatial and fail to characterise the nature and source of variability (Pringle *et al*. 2003). This limitation is illustrated in Figure 5.1 (adapted from Leroux *et al*. 2019), in which three fields share the same global statistics (i.e. mean and

CV), but have distinct spatial structures ranging from weakly to strongly spatially structured. Geostatistics can provide the tools needed to quantify the structure of this spatial variability (Buttafuoco *et al*. 2016), such as an estimation of the range over which observations are related (i.e. the autocorrelation), or characterisation of variability arising from this spatial dependence or from measurement error.

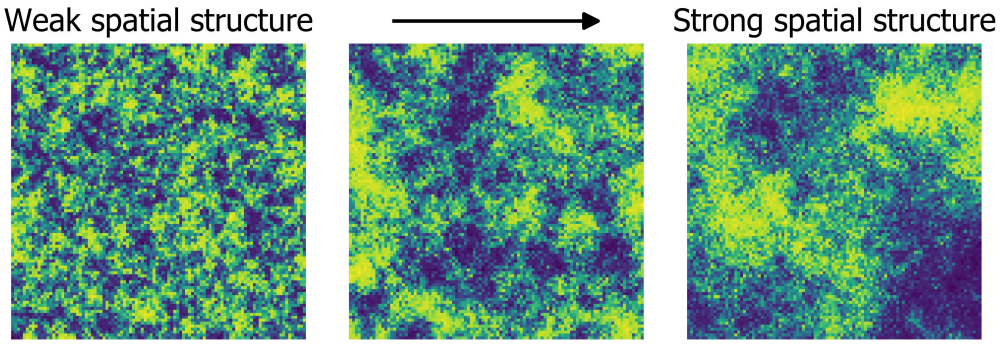

FIGURE 5.1. Examples of varying spatial structures within fields, ranging from weak to strong. Each field shares the same global statistics (i.e. mean and coefficient of variation). Adapted from Leroux *et al*. (2019).

Geostatistical methods have been used to assess and compare the variation in wheat grain yield and GPC within fields, but much of this work had been limited to georeferenced hand samples for a small number of fields (Skerritt *et al*. 2002; Stewart *et al*. 2002; Song *et al*. 2017). While the spatial variability in yield and GPC is often visually compared, this variability is often not quantified, or the range of spatial dependence is not directly compared between yield and protein (e.g. Skerritt *et al*. 2002; Miao *et al*. 2006). Both Kelly *et al*. (2004) and Whelan *et al*. (2009) directly compared variability in grain yield and protein using paired yield and protein maps within fields. Kelly *et al*. (2004) found that, for sorghum, the spatial dependence (which was quantified using a global variogram model) was higher for yield than for protein; for wheat, the spatial dependence was higher for protein than for yield, and, for barley, the spatial dependence was similar between yield and protein. A higher spatial dependence indicates that values exhibited stronger spatial continuity or structure across the field. However, the analysis by Kelly *et al*. (2004) was limited to only three fields for one season and did not utilise a protein sensor. Whelan *et al*. (2009) also found that wheat GPC was more strongly autocorrelated (i.e. more spatially dependent) than yield. The 27 paired

yield and protein maps used by Whelan *et al*. (2009) is among one of the largest studies to quantify and compare the spatial variability in wheat grain yields and GPC within fields using an on-the-go protein sensor.

Here, the Opportunity Index (OI) is presented as a measure of within-field variability and ranks fields as to their suitability for SSCM based on the structure and magnitude of spatial variation. Fields with a higher index value are more amenable to PA and SSCM approaches due to their high variability but strong spatial structure. The OI improves upon existing measures of variation (e.g. CV) by taking into account both the spatial structure of variability, and the ability to respond to this in terms of the operable distance of available machinery. The OI is based on the metric developed by Pringle *et al*. (2003) and improved upon by Oliveira *et al*. (2007), but Pringle *et al*. (2003) and Oliveira *et al*. (2007) only applied this to yield data. Here, the equation improved by Oliveira *et al*. (2007) is used but will be referred to as the OI as both grain yield and GPC will be compared.

The OI has been used to compare and assess the within-field yield variability and suitability of a range of broadacre crops for SSCM, including wheat, barley, canola, chickpeas, and lupins, cereal and pulse crops (Pringle *et al*. 2003; Oliveira 2009; Filippi *et al*. 2024a). However, published studies have only focussed on yield, with none that have assessed quality (e.g. protein or oil content). Given the increasing commercial adoption of grain protein sensors, there is now an opportunity to use the growing abundance of data being collected on-farm to quantify, characterise, and compare variability in wheat GPC and yield.

Depending on the variability in GPC, there are opportunities to manage grain harvest by segregating or blending grain from different areas of a field, to take advantage of the price premiums or reduce discounts (Bramley *et al*. 2012). Alternatively, inputs can be manipulated to optimise outputs and the return on investment (e.g. Colaço *et al*. 2024). Variability in GPC is largely managed by N fertiliser inputs (Wang *et al*. 2023), and reallocating and reducing fertiliser inputs to optimise quality has benefits beyond increasing grower returns, by minimising the environmental impacts of the overapplication of nutrients (e.g. runoff, emissions; Whelan *et al*. 2013). However, both the quality and quantity of grain determine the price received by growers. Depending on the magnitude and extent of variability in GPC

or yield, it may be more appropriate to focus inputs on achieving a maximum yield potential while maintaining a uniform GPC, for example. If variability in yield is greater than in protein, prioritising management of yield would likely result in higher returns. Conversely, if protein variability exceeds that of yield, it would be more effective to focus on optimising protein content through targeted management practices to reduce protein variability or reallocate N inputs. Nutrient applications and availabilities impact both yield and GPC, where, for example, low protein areas within a field can indicate an opportunity for higher N rates and likely a lost yield potential, meaning that both quality and quantity is lost (Scott 2022).

Uncertainty regarding the amount of within-field variation necessary to justify investment in PA technologies has long been a concern for growers and advisors (McBratney *et al*. 1997). Likewise, uncertainty over which is more variable (either GPC or yield) makes it difficult for growers to assess how SSCM should be implemented to achieve the best outcomes (both economically and environmentally). Comparing the variability in yield and quality, and ranking fields based on their variability and suitability for SSCM using the OI, presents an opportunity to streamline decision-making processes and encourage incremental steps in the adoption of PA to improve on-farm economics, production efficiencies, and environmental sustainability.

This chapter presents the use of the OI across six broadacre, dryland farms; four in Western Australia (WA), and two in northern New South Wales (NSW), Australia; to quantify variability in wheat grain yield and GPC, and compare opportunities for SSCM between grain quantity and quality. The aims of this study were to compare the OI between:

(1) Wheat grain yield and GPC, to assess which has a greater opportunity for SSCM within a field;
(2) Farming systems (NSW vs WA), to present case studies assessing the implementation of the OI under different environmental and management conditions; and
(3) Seasons (wet vs dry vs average), to consider the effect of seasonality on opportunities for SSCM.

5.2 Methods

5.2.1 Study area

This study was conducted across Australia's western and northern grain growing regions, with yield and protein data collected from 70 fields across six broadacre, dryland farms over four seasons (2020 – 2023), equating to 86 field-years worth of data. There were two farms in NSW and four farms in WA (Figure 5.2, Table 5.1). Farms A – D were located in WA, and farms E and F were located in NSW. Farms A – C were located approximately 250 – 300 km north-west of Perth, and Farm D was located approximately 170 km north-east of Perth. Farms A – D all have a Mediterranean climate (Beck *et al*. 2018) with cool, wet winters, and hot, dry summers. Farms E and F, which are located approximately 30 km north-east of Moree, have a humid subtropical climate (Beck *et al*. 2018), with warm-to-hot summers, cool winters and summer dominant rainfall.

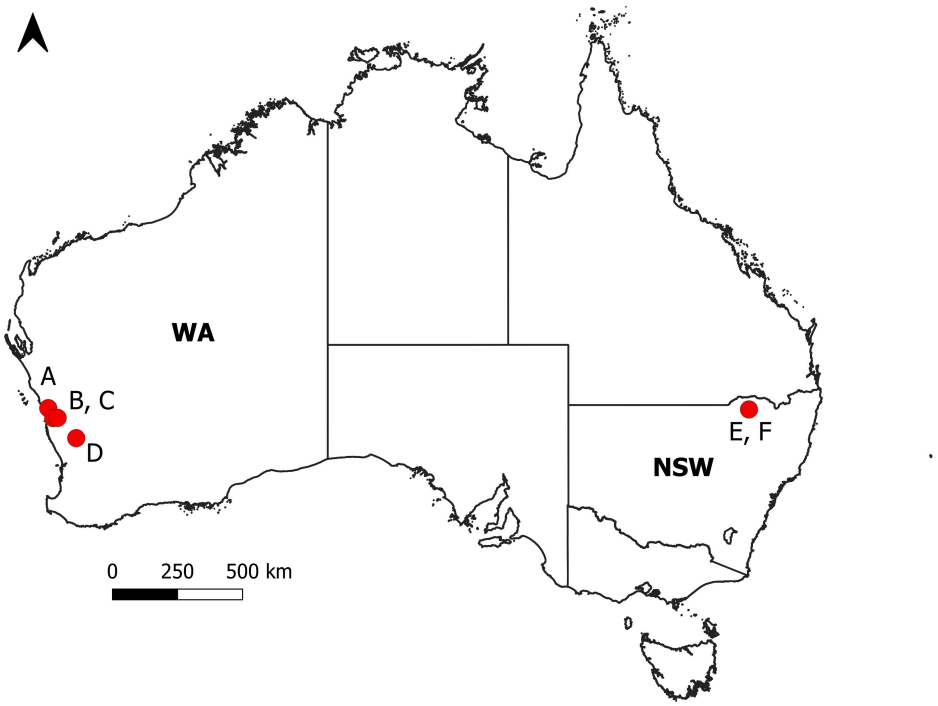

FIGURE 5.2. The study was conducted across six broadacre, dryland farms. Farms A – D were located in Western Australia (WA), and farms E and F were located in New South Wales (NSW).

TABLE 5.1. Available yield and protein data for each farm for each season. Mean field area (hectares), yield (tonnes per ha, t/ha), and grain protein content (GPC, %) and standard deviation (SD) are described. [a]BOM 2024b; [b]BOM 2024c; [c]BOM 2024a; [d]BOM 2024d.

Region	Farm	Number of fields per year					Mean field area (ha)	Annual average rainfall (mm)	Mean yield (t/ha) (± SD)	Mean GPC (%) (± SD)
		Total	2020	2021	2022	2023				
WA	Farm A	14	0	0	11	3	589	422[a]	3.63 (± 1.56)	12.11 (± 1.11)
	Farm B	39	0	5	20	14	234	493b	2.50 (± 1.12)	12.22 (± 1.54)
	Farm C	5	0	0	0	5	303	493[b]	1.32 (± 0.63)	11.90 (± 1.43)
	Farm D	6	3	0	0	3	149	339[c]	3.09 (± 0.92)	11.75 (± 1.95)
NSW	Farm E	11	5	0	3	3	587	579[d]	4.19 (± 1.25)	11.76 (± 1.44)
	Farm F	11	1	5	0	5	431	579[d]	3.92 (± 0.97)	11.32 (± 2.24)

5.2.2 Yield and protein data

Wheat grain yield and GPC data was collected at harvest for the 2020 – 2023 winter growing seasons using harvesters equipped with John Deere's HarvestLab 3000^{TM} near-infrared (NIR) spectroscopy sensor, which can be used to measure grain constituents like moisture, protein, starch, and oil values on-the-go at harvest in wheat, barley, and canola. The yield and protein data was cleaned as per the methods from Taylor *et al*. (2007a). Yield values < 0 or > 10 tonnes per hectare (t/ha) were removed (Taylor *et al*. 2007a), and protein values < 8 and > 22% were removed to exclude extreme outliers based on wheat protein content ranges from Blakeney *et al*. (2009) and Žilić *et al*. (2011), and reasonable thresholds from the observed data. Points more than 2.5 standard deviations (SDs) above and below the field mean were also excluded, and concurrent duplicate protein values were removed based on their time and location to account for potential data errors at harvest. All data processing and analysis was performed in R, version 4.3.1 (R Core Team 2021).

All yield and protein data for each field and season were allocated a unique field-year combination describing the corresponding field and season of data. In total, there were 86 unique field-years worth of data available. Each field-year had a yield and protein map available. Therefore, there were 172 maps available (86 yield, 86 GPC). All available field-years for each farm are described in Table 5.1, including the number of field-years, and mean field area (ha), yield (t/ha), and GPC (%).

5.2.3 Calculation of the OI

The OI (Equation 5.1) is a product of the magnitude and spatial structure of variation. The magnitude of variation (M_v) measures how much observations (either yield or protein) vary across a field, after correcting for random noise (i.e. the nugget effect), relative to a certain threshold. A higher M_v indicates greater differences between areas and suggests that a field may benefit from differential management. The spatial structure of variation (S_v) describes the size of the zones where observations are similar and assesses whether these zones are large enough for variable-rate technologies to operate effectively. A higher S_v indicates a strong

spatial structure in the data, in which observations vary over areas larger than the machinery's operational limits. Fields with a stronger spatial structure are considered more amenable to SSCM (Figure 5.1; Oliveira *et al*. 2007; Leroux *et al*. 2019).

Based on the methods described by Oliveira (2009), the OI was calculated as follows:

$$O_i = \sqrt{M_v \times S_v} = \sqrt{\frac{CV_a}{q_{50}(CV_a)} \times \frac{S}{s}} \tag{5.1}$$

Where;

$M_v =$ magnitude of variation,

$S_v =$ spatial structure of variation,

$CV_a =$ areal coefficient of variation (Equation 5.2),

$q_{50}(CV_a) =$ median areal coefficient of variation of all fields examined,

$S =$ maximum correlated distance within the field (Equations 5.5 and 5.5),

$s =$ minimum operatable distance of available machinery (Equation 5.5),

For each yield and protein map, an experimental variogram was computed using the 'vgm' function in the gstat R package (v2.1.1.; Pebesma 2004) and five variogram models (Exponential, Spherical, Gaussian, Matérn, and Linear) representing bounded and unbounded variogram models were fitted to the experimental variogram. The Akaike Information Criterion (AIC; Akaike 1998) was used to select the best variogram model, and variogram components including the nugget and range were extracted to calculate components of the M_v and S_v described below.

5.2.4 Magnitude of variation (M_v)

The average total field variation (AFV, Equation 5.3) and areal coefficient of variation (CV_a, Equation 5.2) were then calculated, and the practical range (S) was determined, with the

equation used depending on if the model type was unbounded (Equation 5.5) or bounded (Equation 5.4). The operational machinery constant (s, Equation 5.6) was calculated using the inputs suggested by Pringle *et al*. (2003) and used by Filippi *et al*. (2024a).

$$CV_a = \left(\frac{\sqrt{AFV}}{\bar{Y}} \right) \times 100 \tag{5.2}$$

Where;

$AFV =$ average of the total field variation,

$\bar{Y} =$ field mean yield or protein.

$$AFV = \frac{1}{n^2} \left(\left[\sum_{i=1}^{n} \sum_{j=1}^{n} \frac{(x_i - x_j)^2}{2} \right] - C_0 \right) \tag{5.3}$$

Where;

$n =$ the number of yield or protein observations,

$x_i =$ observation i,

$x_j =$ observation j,

$C_0 =$ nugget variation of the best fitting variogram.

5.2.5 Spatial structure of the variation (S_v)

The spatial structure of variation (S_v) was calculated by determining the practical range (S; Equations 5.5 and 5.4), defined as the maximum distance within a field where observations show meaningful spatial correlation. This value was then normalised against the minimum operational distance of available machinery (s; Equation 5.6). The method for estimating S depended on the whether the best-fit variogram model was bounded or unbounded (Equations 5.5 and 5.4). Specifically, distance (h) was determined as either:

(1) the distance where the semivariance reached 95% of the sill (C_1) plus the nugget effect (C_0), for bounded exponential models (Equation 5.5);
(2) the point where the semivariance equalled the sill (C_1) plus the nugget effect ($_0$), for bounded spherical models (Equation 5.5); or
(3) the point where the semivariance is equalled 95% of the variance at maximum lag minus the nugget effect (C_0), for unbounded models (Equation 5.4).

S was then calculated by dividing h by the square root of 10,000 to convert the units from meters to hectares.

$$S = \frac{h}{\sqrt{10000}}\ when : \begin{Bmatrix} \gamma(h) = C_0 + 0.95C_1(exp.) \\ \gamma(h) = C_0 + C_1(sph.) \end{Bmatrix} \tag{5.4}$$

Where;

$C_1 =$ partial sill,

$\gamma(h) =$ practical range,

$C_0 =$ nugget variation.

$$S = \frac{h}{\sqrt{10000}}\ when : \gamma(h) = C_0 + 0.95[\gamma_{MaxLag} - C_0] \tag{5.5}$$

Where;

$\gamma_{MaxLag} =$ semivariance at maximum lag of the experimental variogram,

Finally, the minimum operational distance of available machinery (s; Equation 5.6) was calculated using the constants for machinery width (β), vehicle speed (v), and time (τ) suggested by Pringle *et al*. (2003) and used by Filippi *et al*. (2024a). s indicates the distance over which machinery such as boom sprayers for variable-rate fertiliser applications can operate and change input applications to respond to different management areas within a field. A field with a larger practical range (S) indicates a strong spatial structure of variation. If the

practical range is greater than the minimum operational distance of the available machinery (s), then inputs can be differentially applied and the zones of variation within the field are large enough to justify SSCM.

$$s = \frac{\beta v \tau}{10000} \tag{5.6}$$

Where;

$\beta =$ width of machinery swath (m), set at 20,

$v =$ speed of vehicle (meters per second), set at 6.

$\tau =$ time to alter application rate (seconds, set at 3).

5.2.6 Classification of OI results

The OI is used to rank fields as to their suitability for SSCM, so that fields with a higher OI appear more amenable to variable management (Filippi *et al*. 2024a). Based on the approach by Filippi *et al*. (2024a), the OI for yield and protein for all fields was grouped into three classes using the 1st and 3rd quartiles for the full dataset of OI results. Yield and GPC were grouped together for the classification because they were being directly compared. Maps with an OI value of less than the 1st quartile were classed as having a relatively low opportunity for SSCM, maps with an OI between the 1st and 3rd quartiles were classed as having a relatively medium opportunity for SSCM, and maps with an OI above the 3rd quartile were classed as having a relatively high opportunity for SSCM. The OI values at these quartiles were rounded to the nearest 0.5 to provide intuitive threshold values (Filippi *et al*. 2024a).

5.2.7 Comparison of the OI between grain yield and protein content

The OI was calculated for every yield and protein map for all field-years. The CV was also calculated and compared with the OI as a baseline against standard measures of variation. The OI and CV was compared between wheat grain yield and GPC for different scenarios:

(1) Overall, to assess whether yield or GPC has a greater opportunity for SSCM;
(2) Between farming systems (NSW vs WA), to present case studies assessing the utilisation of the OI under different environmental and management conditions; and
(3) Between seasons, to consider the effect of seasonality on opportunities for SSCM.

Seasons were classed as wet, dry, or average based on the 1st and 3rd quartiles from the last 50 calendar years of rainfall data specific to each farm, which was accessed from SILO (Government 2024) via the dataharvester R package (V0.1.2; Harianto *et al*. 2023). Farms with total yearly rainfall below the 1st quartile were classed as "dry", farms between the 1st and 3rd quartiles were classed as "average", and farms above the 3rd quartile were classed as "wet". For all farms, 2023 was a dry season, 2020 was an average season, and 2021 and 2022 were both wet seasons.

Pairwise differences between the OI for wheat grain yield and GPC for each field-year were calculated and compared within and between regions and seasons to assess if there was a greater opportunity to manage for yield or protein (Equation 5.7). A positive pairwise OI result indicates that there was a greater opportunity to manage for yield in a particular field, whereas a negative pairwise OI result indicates that there was a greater opportunity to manage for GPC.

$$pairwise\ OI = Yield\ OI - GPC\ OI \tag{5.7}$$

The behaviour of OI over time was examined to determine whether yield or GPC changed more from season-to-season. For 15 fields with two seasons of available data, pairwise differences in the OI between seasons were calculated for both wheat grain yield and GPC (Equation 5.8). The absolute change in OI between seasons was then used to assess whether yield or GPC exhibited more change over time.

$$temporal\ change\ in\ OI = OI_{\text{x, t}} - OI_{\text{x, t-1}} \tag{5.8}$$

Where;

$x =$ Either yield or GPC,

$t =$ season of yield or protein data,

$t - 1 =$ earlier season of yield or protein data.

5.3 Results

Overall, yield OI values ranged between 3.13 and 24.61 and had a median OI of 9.12. The GPC OI values ranged between 1.64 and 14.81 and had a median value of 4.52. The CV was higher for yield (median CV = 0.30) than GPC (median CV = 0.07) (Figure 5.3, Table 5.2).

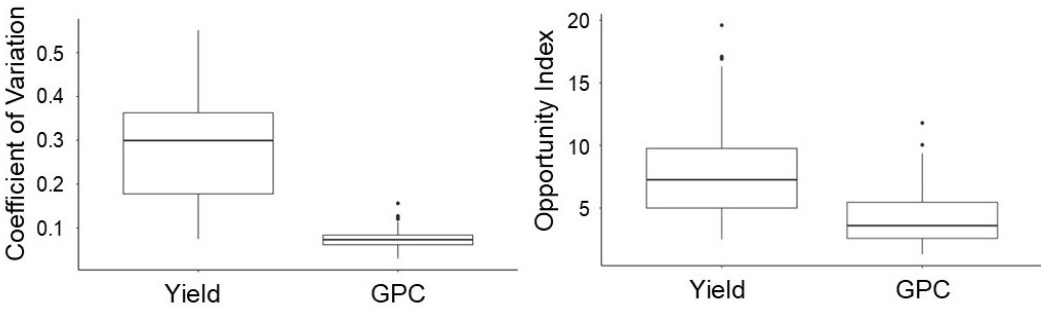

FIGURE 5.3. Distribution of Coefficient of Variation and Opportunity Index values for yield and grain protein content (GPC) maps.

Between regions, both the median CV and median OI for yield were higher in WA than in NSW, but the difference in the yield OI between NSW and WA was small (Figure 5.4, Table 5.2). Given that the OI is a product of both the magnitude (M_v) and structure of variation (S_v), the similarities in the OI between NSW and WA are likely due to a greater spatial structure in yield variation in NSW, but a greater magnitude of variation in WA (Table 5.3). For GPC, the CV was higher in WA than in NSW, whereas the GPC OI was higher in NSW compared to

WA. Therefore, although WA exhibits greater variability in both yield and GPC, the variability in WA may be less manageable due to a poorer spatial structure compared to NSW, meaning that the variation in WA may occur over smaller, less coherent areas.

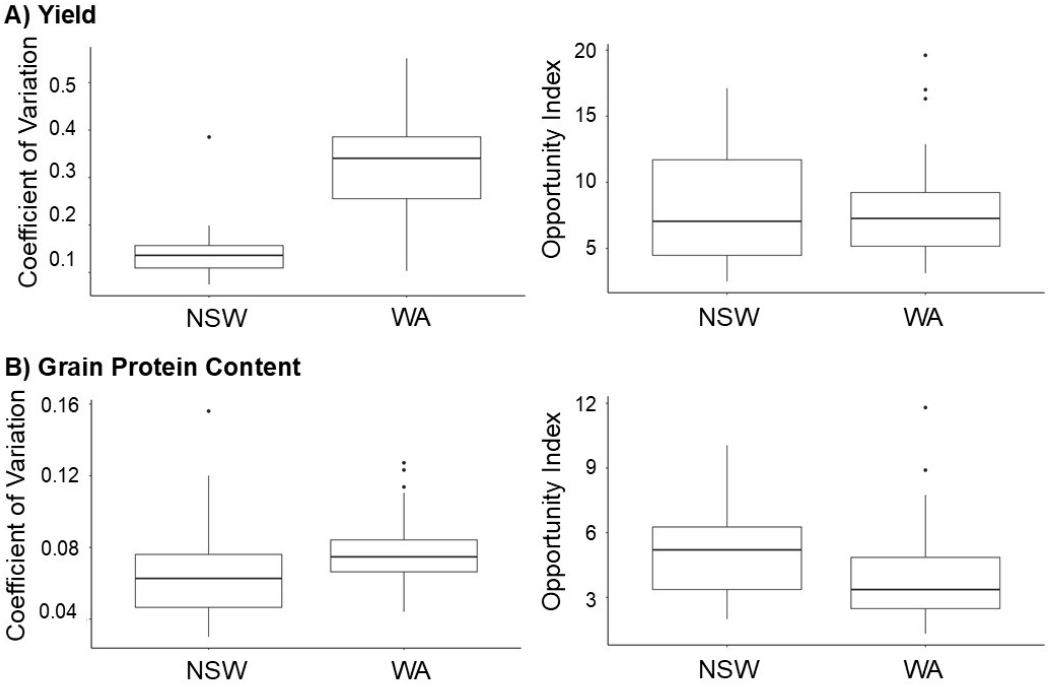

FIGURE 5.4. Distribution of Coefficient of Variation and Opportunity Index values for A) yield and B) grain protein content (GPC) maps in New South Wales (NSW) and Western Australia (WA).

Between seasons in WA, the yield CV was lowest during average seasons, while the yield OI was highest (Figure 5.5, Table 5.2). For GPC, both the CV and OI were highest during average seasons. These findings suggest that under "typical" growing conditions in WA, there is a greater opportunity for SSCM for both yield and GPC compared to more extreme seasons (i.e. wet or dry), likely due to the stronger spatial structure of the variation during average seasons compared to dry or wet seasons (Table 5.3). In NSW, yield CV and OI were highest during dry seasons and lowest during wet seasons, indicating that in wet years, yield variability is limited in both magnitude (M_v) and spatial structure (S_v). For GPC, CV was highest during wet seasons, whereas the OI was highest in average seasons, suggesting that while GPC variability may increase in wetter conditions, its spatial structure and, therefore,

the potential for SSCM is strongest under typical seasonal conditions. Overall, yield CV had greater variability season-to-season in WA compared to NSW, but the OI for both yield and GPC varied more season-to-season in NSW. This variability in the OI is largely attributed to seasonal variation in the spatial structure, as the magnitude of variation was comparatively more stable between seasons.

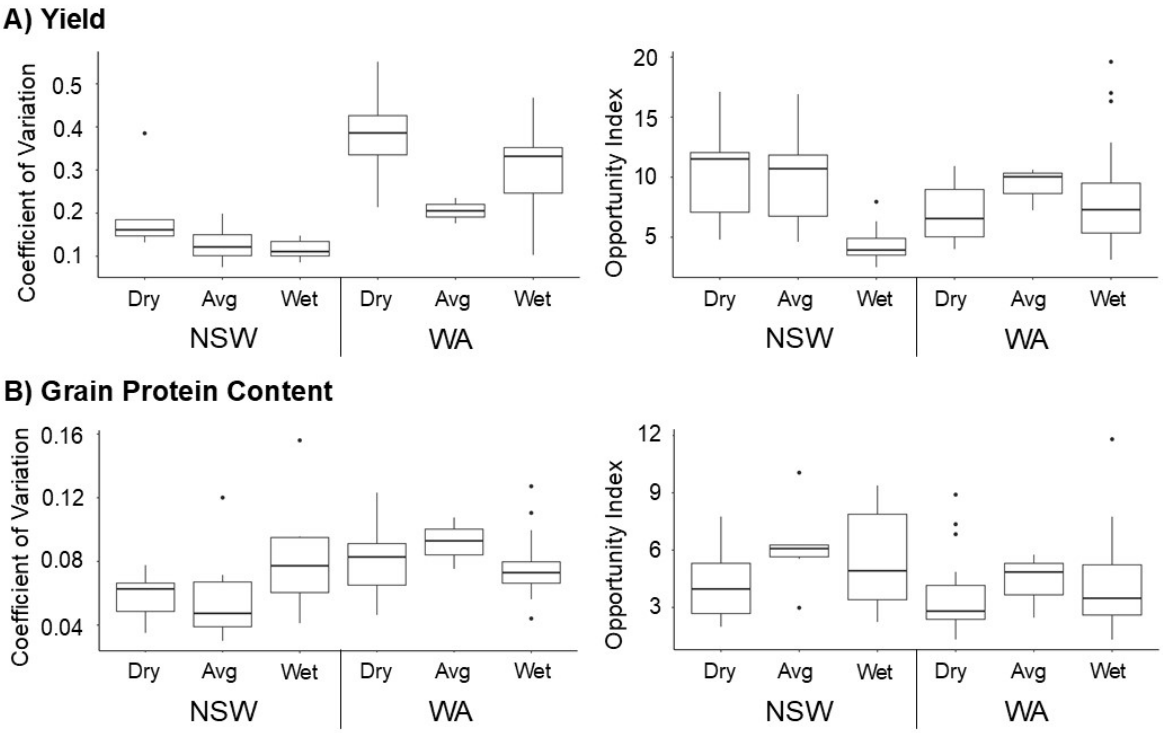

FIGURE 5.5. Distribution of Coefficient of Variation and Opportunity Index values for A) yield and B) grain protein content (GPC) maps in New South Wales (NSW) and Western Australia (WA) and between dry, average (avg), and wet seasons.

TABLE 5.2. Mean, median, and standard deviations (SD) of the magnitude of autocorrelated (spatially dependent) variation (M_v) and spatial structure of variation (M_v) for all wheat grain yield and grain protein content (GPC) data. Mean, median, and SD values are summarised overall, for each region (New South Wales, NSW; Western Australia, WA), and between wet, dry, and average seasons within each region.

			CV			OI		
Product	Region	Season	Mean	Median	SD	Mean	Median	SD
Yield	Overall		0.28	0.30	0.12	9.90	9.12	4.49
	NSW	All	0.15	0.14	0.06	10.20	8.85	5.57
		Dry	0.19	0.16	0.08	12.91	14.45	5.26
		Average	0.13	0.12	0.04	12.67	13.44	5.65
		Wet	0.12	0.11	0.02	5.63	4.93	2.24
	WA	All	0.33	0.34	0.09	9.79	9.12	4.10
		Dry	0.38	0.39	0.08	8.89	8.22	2.73
		Average	0.21	0.21	0.03	11.67	12.59	2.27
		Wet	0.31	0.33	0.08	10.16	9.15	4.75
GPC	Overall		0.07	0.07	0.02	5.32	4.52	2.70
	NSW	All	0.07	0.07	0.03	6.59	6.53	2.95
		Dry	0.06	0.06	0.01	5.35	4.98	2.48
		Average	0.06	0.05	0.03	7.75	7.63	2.85
		Wet	0.08	0.08	0.04	6.96	6.17	3.34
	WA	All	0.08	0.07	0.02	4.88	4.22	2.48
		Dry	0.08	0.08	0.02	4.50	3.52	2.45
		Average	0.09	0.09	0.02	5.47	6.09	2.14
		Wet	0.07	0.07	0.01	5.05	4.37	2.56

TABLE 5.3. Mean, median, and standard deviations (SD) of the magnitude of autocorrelated (spatially dependent) variation (M_v) and spatial structure of variation (S_v) for all wheat grain yield and grain protein content (GPC) data. Mean, median, and SD values are summarised overall, for each region (New South Wales, NSW; Western Australia, WA), and between wet, dry, and average seasons within each region.

			M_v			S_v		
Product	Region	Season	Mean	Median	SD	Mean	Median	SD
Yield	Overall	2.70		2.93	1.25	279.07	159.79	291.29
	NSW	All	1.25	1.12	0.67	513.28	420.64	409.95
		Dry	1.66	1.38	0.90	601.33	631.77	347.82
		Average	1.10	1.02	0.47	806.14	710.92	487.87
		Wet	0.96	0.95	0.31	205.59	140.04	150.34
	WA	All	3.20	3.30	0.98	198.56	130.91	180.88
		Dry	3.81	3.86	0.81	123.41	98.17	74.96
		Average	1.49	1.66	0.31	481.8	495.96	91.42
		Wet	3.00	3.22	0.86	219.16	131.26	201.62
GPC	Overall		0.70	0.66	0.23	266.11	190.40	239.33
	NSW	All	0.61	0.56	0.31	435.79	443.66	264.59
		Dry	0.48	0.50	0.16	361.49	255.72	273.48
		Average	0.58	0.46	0.35	587.92	606.34	206.45
		Wet	0.74	0.63	0.37	396.01	336.35	277.18
	WA	All	0.73	0.73	0.19	207.78	145.05	201.08)
		Dry	0.78	0.73	0.21	173.29	88.44	202.26
		Average	0.91	0.97	0.14	181.13	200.67	104.92
		Wet	0.69	0.69	0.17	229.29	186.71	206.46

When pairwise differences between the OI for yield and GPC were calculated to assess which had a greater opportunity for SSCM within each field-year, yield presented a greater opportunity (i.e. positive pairwise difference in OI) in over 87% of field-years overall (Figure 5.6). This trend was particularly prominent in WA, where more than 92% of field-years had a greater opportunity for SSCM with yield. In WA, pairwise differences in the OI between yield and GPC were relatively consistent across seasons, though there was greater variability in wet seasons. In NSW, a stronger seasonal effect was observed, with yield showing a greater opportunity for SSCM in dry and average years. However, during wet seasons, GPC had a greater opportunity for SSCM in NSW.

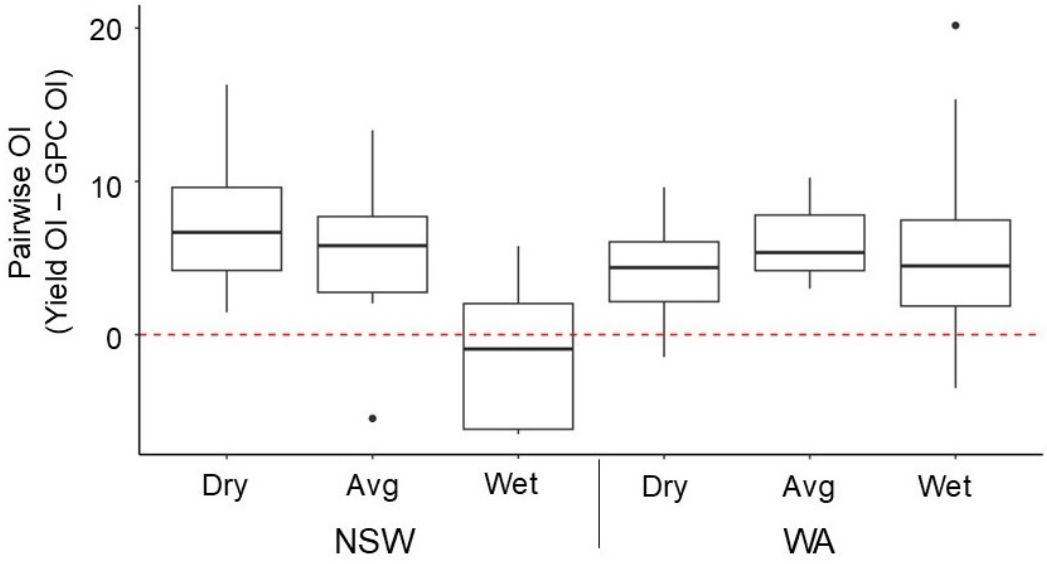

FIGURE 5.6. Pairwise differences between Opportunity Index (OI) for wheat grain yield and protein content for each Field-Year in northern New South Wales (NSW) and Western Australia (WA) in wet, dry, and average seasons. Positive pairwise OI values indicate that there was a greater opportunity for site-specific crop management (SSCM) for yield, whereas negative pairwise OI values indicate a greater opportunity for protein.

The OI results for both yield and GPC were grouped into low, medium, and high opportunities for SSCM based on the 1st and 3rd quartiles (Section 5.2.6), so that maps with an OI value less than 4.5 had a relatively low opportunity for SSCM, and maps with an OI greater than 10 had a relatively high opportunity (Figure 5.7). Yield and GPC were grouped together for the classification because they were being directly compared. For GPC, ~50% of field-years had a relatively low opportunity for SSCM, and less than 6% had a relatively high opportunity. For yield, almost 45% of field-years had a relatively high opportunity for SSCM, and less than 5% had a relatively low opportunity.

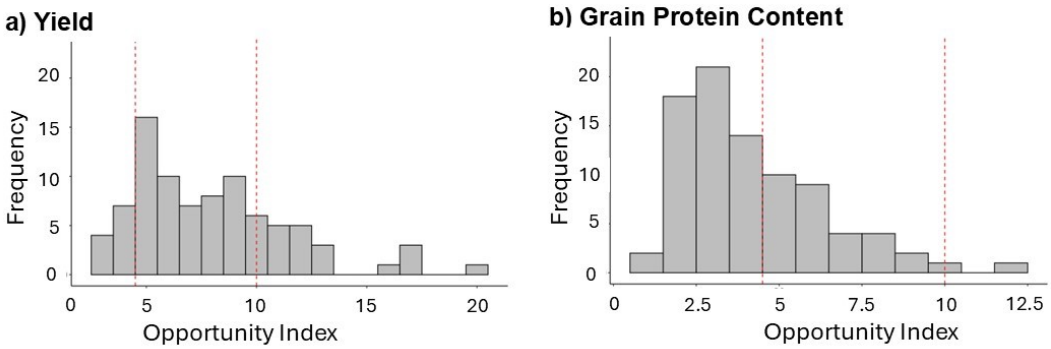

FIGURE 5.7. Distribution of Opportunity Index (OI) values for wheat grain yield (A) and protein content (B), with red lines marking the boundaries of the low (OI = 4.5) and high (OI = 10) classification of opportunity for site-specific crop management (SSCM).

Yield and protein maps are presented for three case study fields that are representative of relatively low, medium, and high opportunities for SSCM. For wheat grain yield, Field A had a low OI of 4.16, Field B had a medium OI of 8.28, and Field C had a high OI of 15.10 (Figure 5.8). For GPC, Field D had a low OI of 2.33, Field E had a medium OI of 6.90, and Field F had a high OI of 12.62 (Figure 5.9).

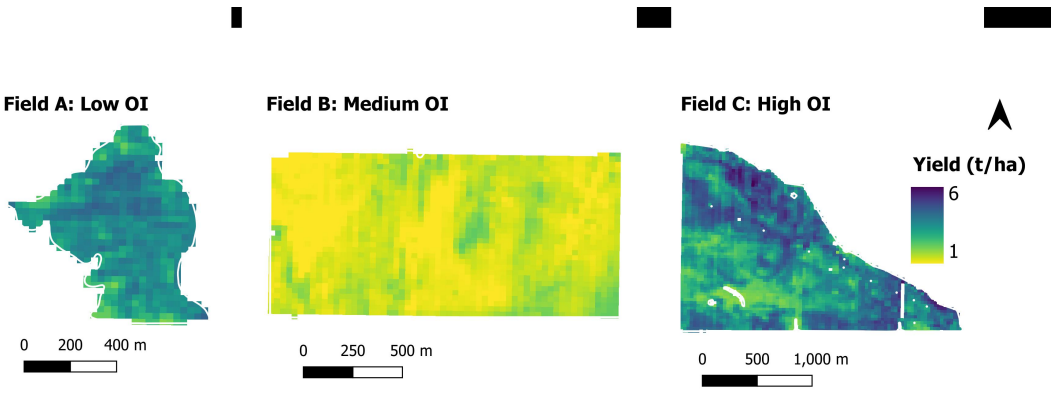

FIGURE 5.8. Three representative case study fields with low, medium, and high opportunity index (OI) values for wheat grain yield (tonnes per hectare, t/ha). Field A has a low OI (4.16), Field B has a medium OI (8.28), and Field C has a high OI (15.10).

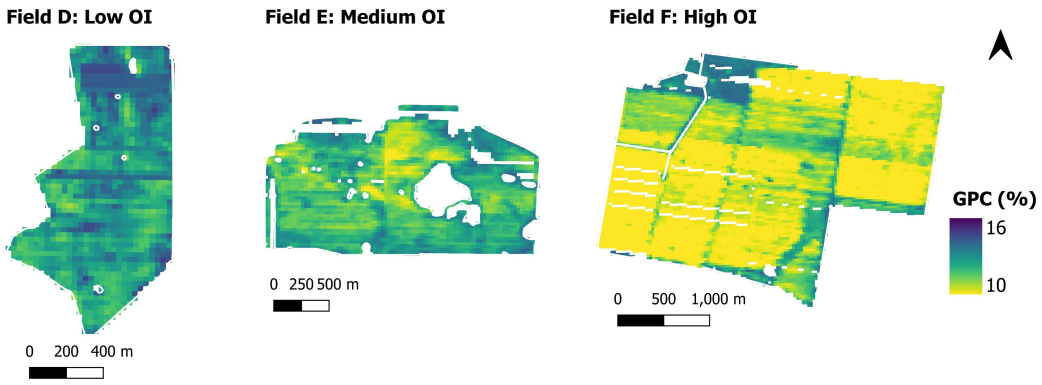

FIGURE 5.9. Three representative case study fields with low, medium, and high opportunity index (OI) values for wheat grain protein content (GPC, %). Field D has a OI (2.33), Field E has a medium OI (6.90), and Field F has a high OI (12.62).

The temporal change in the OI between seasons was calculated for both yield and GPC in the 15 fields with two seasons of available data (Equation 5.8). The results, presented in Table 5.4, indicate that overall, the temporal change in the OI for GPC was less than for yield. Specifically, the median temporal change in the OI for yield was 2.43 overall, with similar values observed in both NSW (median = 2.38) and WA (median = 2.48). In comparison, the GPC OI showed a median value of 2.14 overall, but the temporal change in the OI was higher in NSW (median = 3.71) than in WA (median = 1.18). Therefore, GPC exhibited more temporal variability in NSW than in WA as the OI changes more from season-to-season. In contrast, the yield OI showed a relatively similar level of temporal change between the regions.

TABLE 5.4. Mean, median, and standard deviation (SD) of the temporal change in the opportunity index (OI) OI seasons for yield and grain protein content (GPC) overall, and in New South Wales (NSW) and Western Australia (WA).

Product	Region	Mean	Median	SD
Yield	Overall	3.61	2.43	2.81
	NSW	4.42	2.38	3.50
	WA	2.80	2.48	1.68
GPC	Overall	3.16	2.14	3.14
	NSW	3.73	3.71	2.62
	WA	2.59	1.18	3.60

5.4 Discussion

While more data describing variation in GPC within-fields is needed to improve our understanding of spatiotemporal variability (e.g. on-the-go protein sensors), the value of this data and these technologies to growers has not been fully realised, as the translation of this information into tangible management decisions remains a challenge. The OI is an important tool for the implementation of PA on-farm. The OI quantifies both the structure and magnitude of variability, and assesses if there are opportunities for SSCM by considering if this variation can be addressed using available machinery. Given the potential costs involved with implementing SSCM, including in technologies, knowledge, or expertise, ranking fields by their suitability for SSCM can facilitate a stepwise adoption of PA and direct management interventions where they may be most-effective in addressing variability. For fields that have been identified as having a high opportunity for SSCM, further investigations should be conducted to understand the drivers of this variation and whether this variability is due to intrinsic factors (e.g. soil texture and water holding capacity), or can be controlled by specific management actions (e.g. variably applying nutrient inputs). Over time, the incremental identification, adoption, and evaluation of fields with SSCM may result in a change in whole-farm management.

5.4.1 Comparison of the OI between wheat grain yield and grain protein content

Across the entire dataset, there was a greater opportunity for SSCM for wheat grain yield than for protein content. This is not surprising when we consider the components of the OI equation (Equation 5.2), which is a product of both the magnitude (M_v) and spatial structure (S_v) of variation. While the spatial structure of variation for yield and GPC were similar, the magnitude of autocorrelated (spatially dependent) variation for GPC was much smaller than that for yield (Table 5.3), meaning that GPC is less variable within fields overall. This is likely a reflection of the nature of protein concentration as a proportional variable expressed as a percentage, which inherently limits its variability compared to yield which is a continuous variable in t/ha.

The OI provides a simple metric to quantify and rank fields as to their opportunity for SSCM, and can be used by growers and advisors to direct SSCM strategies to where they will have the greatest impact, both in terms of crop attributes and in which fields. Given that there was a greater opportunity for SSCM for wheat grain yield than for GPC, SSCM strategies should first focus on mitigating/managing for yield variability or to achieve yield potentials. This may include addressing potential constraints (e.g. soil chemical or physical properties) or the addition or reallocation of inputs (e.g. fertilisers). Increasing yields and/or reallocating inputs for a better return on investment will provide growers with better economic returns and improved sustainability outcomes. Then, after variability in yields have been addressed and yield potentials have been met, growers and advisors should then focus SSCM on optimising GPC to seek higher-value market premiums. However, management actions addressing yield variability, such as N fertiliser applications, will also impact GPC (Holford *et al*. 1992; Simmonds 1995) and these interactions should be considered when making management decisions.

5.4.2 Opportunities for SSCM between NSW and WA

The OI was compared between NSW and WA, but there are considerable agronomic, climatic, and landscape differences between the two farming systems as they operate within distinct agroecological zones (Davies *et al*. 2022). This means the OI should not be used to definitively state whether there are greater opportunities for SSCM in either region, but the OI can still reveal valuable insights into how SSCM may be implemented in each farming system for GPC and yield.

The western growing region of Australia (WA) is characterised by having low soil fertility and generally lower yields compared to the northern growing region (i.e. Queensland and NSW), whereas the northern growing region has relatively high seasonal rainfall and greater production variability (Davies *et al*. 2022). Much of the early PA research in Australia was conducted in WA (Cook *et al*. 1998; Cook *et al*. 1996; Adams *et al*. 1999), and the use of variable-rate technologies is significantly higher in the western growing region compared to the northern region (Bramley *et al*. 2019). In WA, spatial variability in wheat grain yield and protein content within-fields is predominately driven by variation in plant-available water capacity (Oliver *et al*. 2006), and the consistency in crop patterns in both space and time predispose WA to be advantageous for PA (Robertson *et al*. 2008).

While at first glance the results suggest that there are lower opportunities for SSCM for wheat GPC and yield in WA compared to NSW (Figure 5.4), the greater CV in WA for both yield and GPC is likely due to the existing implementation of SSCM in WA. It is possible that differentially applying nutrients to match yield potential across different zones within a field has increased the magnitude of variation, but short-range variability, possibly due to within-zone heterogeneity, is resulting in a weaker spatial structure and, therefore, a lower OI.

Given that SSCM is already in practice across many WA farms, the spatial variability observed may actually reflect the impacts of ongoing management decisions. This presents a challenge when interpreting the OI as it becomes difficult to disentangle the effects of inherent soil or environmental variability from those introduced from variable-rate applications. As a result, the OI may not fully represent the potential opportunities for SSCM in WA, but rather than

current outcomes of SSCM. This may explain why in WA, GPC has a lower OI than yield overall. This is because GPC is typically managed to increase uniformity between zones, whereas yield is often managed to achieve yield potential within each zone which results in a larger magnitude of variation in yield. Despite this complexity, the OI still provides valuable insight into how SSCM is driving variation in yield and GPC in WA, particularly by revealing where spatial structure is weak and where a refinement in zone delineation or input application may be needed.

Interestingly, while the spatial structure of variation is overall stronger in NSW than in WA (Table 5.3), both regions show stronger spatial structure in average seasons compared to wet or dry years. This is likely because of extreme conditions such as drought or flooding which can suppress spatial patterns if crop performance is uniformly poor. Conversely in average seasons, underlying soil differences in texture and water holding capacity are more clear, resulting in a stronger spatial structure of variation. Of note, during wet seasons in NSW, there was greater opportunity for SSCM (i.e. a higher OI) for protein than for yield. This is likely because the soil N supply was limiting to GPC, whereas soil moisture was adequate.

In WA, the relative consistency in spatial structure in yield and GPC across years is largely due to persistent spatial patterns in soil texture that tend to define production zones and remain stable over time (Oliver *et al*. 2006). However, the spatial structure is still relatively weaker than in NSW. This may be due to within-zone heterogeneity, as discussed previously. On the other hand, the strong spatial structure in yield and GPC observed in NSW may be influenced by legacy field boundaries which continue to influence yield and protein patterns (Chapter 4, Tilse *et al*. 2024a). In addition, the average field size in NSW was approximately 200 hectares greater than in WA, meaning that fields in NSW may encompass a greater range of topographic and soil features which contribute to stronger, more coherent spatial patterns. Therefore, even though WA exhibits an overall greater magnitude of variation, NSW likely has a higher OI due to its stronger spatial structure.

We expect that fields that exhibit more temporally stable OI patterns should be more amenable to SSCM. There was greater season-to-season variability in the OI overall in NSW compared to WA, and temporal pairwise differences in the GPC OI within fields were also greater in

NSW. Therefore, we should expect that SSCM may be more suitable in WA as fields are more temporally stable (Whelan *et al*. 2000). However, only two seasons worth of grain yield and protein data were available across 15 fields, and only one field had three seasons of data available, meaning that long-term temporal trends could not be assessed. Likewise, depending on the SSCM strategy implemented (e.g. the amelioration of soil chemical constraints, or the variable-rate application of N-fertilisers), there may be carry-over effects beyond the current season which may alter the structure and magnitude of variation in future seasons. This necessitates a re-evaluation of the OI over time and may allow the assessment of the success of SSCM strategies implemented through benchmarking of variability. There is a growing appetite for the collection of grain protein sensor data by growers to better understand and manage for variability in GPC. As more and more seasons worth of paired yield and protein data are being collected by growers, there is the possibility in the future to conduct a more thorough temporal analysis of OI stability within fields over time. This would help to better understand the behaviour of variability within fields, and could also be used to better inform management decisions by directing SSCM strategies to those areas that are more temporally stable but spatially variable (Whelan *et al*. 2000).

5.4.3 Future work

While the OI captures and describes the magnitude and spatial structure of variation, and assesses if this variation can be addressed using available machinery, it does not consider the economics of potential management strategies (Oliveira 2009). For example, a field may exhibit a high degree of variability in yield with a strong spatial structure, indicating potential gains from SSCM to increase yields in poor-performing areas. However, if urea prices are high and yields cannot be increased beyond a magnitude that would result in a return on investment that is greater than not applying these inputs. Likewise, changes in N fertiliser inputs to address variability in yield will also impact GPC, and visa-versa (Holford *et al*. 1992; Simmonds 1995).

In the future, economics of these decision-making strategies should be considered. This could include the calculation of the OI based on a map of profit within a field, rather than calculating

the OI for OI or GPC separately. Given maps of GPC and yield, and fixed grain prices, profit could be calculated to assess which fields are economically more variable. From this, management recommendations such as N fertiliser inputs could be simulated (e.g. Chapter 4, Tilse *et al*. 2024a) to assess the cost of implementing SSCM given a yield and/or protein target for a particular field. The consideration of economics in decision-making strategies is imperative for growers, and is particularly important in Australia when dealing with quality due to the premium and discount system applied to Australian grain producers.

The results from this study indicate that there is less opportunity for SSCM for GPC than for yield. However, while grain yield and protein maps derived from on-the-go sensors are useful for characterising spatiotemporal variability, complex interactions between grain yield and GPC may be missed if variability is only assessed using yield or protein maps in isolation (Long *et al*. 2011). A map of GPC is considered to be analogous to an N adequacy map (Scott 2022), where a low GPC (< 12%) indicates N was insufficient for crop requirements, and a high GPC (> 12% indicates a surplus of N. At the same time, with the addition of N, both yield and GPC increase up until a certain point, which GPC continues to increase but yields are penalised (Holford *et al*. 1992; Simmonds 1995). Given the OI for SSCM for GPC in Field F in Figure 5.9, for example, we should expect that increasing N fertiliser inputs should raise the GPC closer to an optimum of 12% and will also result in a substantial yield increase. This should result in considerable economic benefits if profits associated with the yield increase and attainment of GPC premiums exceed input costs. To capture interactions between yield and GPC, future work could consider the OI from the joint spatial correlation of yield and GPC, calculated within moving windows of correlation (e.g. Chapter 4, Tilse *et al*. 2024a) or multivariate geostatistics (e.g. Casa *et al*. 2008; Buttafuoco *et al*. 2016).

There is a growing plethora of spatial data layers now available to describe the variability being observed within fields and across farms in Australia, including on-farm management records, publicly-available satellite imagery, weather and climate data. There is an opportunity to utilise these spatial data layers to better understand the differences in OI between yield and protein, including how and why the OI varies within and between fields, farms, and seasons.

Further work is needed to better understand the drivers of within-field variability in yield and protein content, and their interactions (e.g. Tilse *et al*. 2024a).

Overall, ranking fields by their suitability for SSCM can facilitate a stepwise adoption of PA on farms and guide management decisions to where they will be most effective at addressing or mitigating variation. The OI does not provide a recommendation for how to manage for this variation, and further investigations by the grower/advisor are needed to determine what management options may be suitable to address the variation observed within a field. These options will differ depending on their agronomic or economic goals. Even so, the OI is an important tool to help quantify and describe the variability occurring within fields and across farms, and for guiding management decisions. Over time, the incremental identification, adoption, and evaluation of fields using SSCM may result in a change in whole-farm management for improved sustainability, efficiency, and economic outcomes.

5.5 Conclusion

The OI provides a simple metric to quantify and compare the magnitude and spatial structure of variability in wheat grain yield and GPC within fields, and rank fields as to their suitability for SSCM In turn, this can support the stepwise adoption of PA on farm and guide management decisions to where they will be most effective at addressing or mitigating variation. Overall, there was a greater opportunity for SSCM for yield than for GPC, and there was a greater opportunity for SSCM in NSW compared to WA, but this is more likely a reflection of SSCM already being implemented in WA, rather than an actual indication of opportunity. While both yield and GPC had a similar spatial structure of variation, yield had a greater magnitude of variation overall and the spatial structure of variation was greater in NSW which is likely due to differences in soil texture between each region, as WA soils are more sandy whereas NSW is more clayey, alongside paddock history, and the larger average size of NSW fields. Given that the OI was higher for yield than for GPC, SSCM should focus on mitigating/managing for yield variability to achieve yield potentials before optimising GPC to seek higher-value market premiums. However, an economic analysis needs to be considered in future work.

Likewise, future work should focus on better understanding the drivers of the differences in the OI between yield and GPC. This could help to identify if variability can be controlled or ameliorated by growers (e.g. due to a soil pH constraint or nutrient deficiency), or if it is due to inherent features such as soil texture.

CHAPTER 6

General Discussion

Increasing volatility and changing export market demands are pushing Australian cotton and grain growers to do more with less (Muleke *et al*. 2022). Likewise, despite continued research and agronomic advancement, yield improvements appear to be approaching a plateau (Muleke *et al*. 2022; Hochman *et al*. 2017). Precision agriculture (PA) offers an opportunity to shift focus from yield to quality to maintain or enhance access to premium markets, and optimise cotton fibre and grain quality for improved productivity and sustainability. Increasing economic, agronomic, and environmental challenges faced by Australian cotton and grain growers have spurred innovation and investment in new technologies and management strategies, and more data is being collected on-farms and by the industry than ever before. Yet the application of this data and technologies for decision-making has been limited as technologies such as on-the-go fibre quality and grain protein sensors are not available at all (in the case of cotton), or are not widely adopted (in the case of grains). Further, translating this information into actionable management decisions has been a challenge.

This thesis highlights how the growing diversity and abundance of publicly-available and on-farm spatial data layers can be used to better understand and describe variability in cotton fibre and grain quality for improved management. Specifically, five research questions were addressed:

(1) Can a high-resolution spatial map of variability in cotton fibre quality be produced?
(2) Can the wealth of available spatiotemporal data be used to predict wheat grain protein content (GPC) at a fine-resolution:
    (a) within parts-of a field where two headers with yield monitors but only one protein sensor are operating; or
    (b) across whole fields in the complete absence of grain protein sensor data but there is protein sensor data available for surrounding fields or in past seasons?
(3) Can the drivers of variability in wheat GPC be better understood?
(4) What is the opportunity for managing variability in GPC?
(5) If there is an opportunity for managing variability in GPC, what can be done?

These five research questions were addressed within Chapters 2 – 5. Research question 1 was addressed in Chapter 2, which presented a generalised geostatistical approach using area-to-point kriging to map and downscale areal observations of cotton fibre yield and quality (length and micronaire) data. Research question 2 was addressed in Chapter 3, which demonstrated how a combination of readily-available yield, agronomic, and publicly-available data layers could be used to create a model to predict GPC within-fields and fill-in gaps in the absence of a grain protein sensor. Research questions 3 and 5 were addressed in Chapter 4, where an interpretive machine learning approach was used to understand the drivers of spatial variability in GPC within-fields, and the utilisation of this knowledge for retrospectively determining optimal nitrogen (N) management was also explored. Research questions 1 and 4 were addressed in Chapter 5 which quantified the magnitude and spatial structure of within-field variability in GPC to understand how it varies using the Opportunity Index (OI) metric, and also compared opportunities for site-specific crop management (SSCM) between wheat grain yield and GPC.

Each research question aligns with a specific chapter, while three overarching themes connect these questions:

HOW? How does quality vary?

WHY? Why does quality vary?

SO, WHAT? What can be done to optimise quality?

These themes can be considered a framework for the stepwise adoption of PA for understanding and optimising cotton fibre and grain quality (Whelan *et al*. 2013). First, a foundational understanding of **HOW** cotton fibre and grain quality varies is necessary, including identifying where this variability is occurring and quantifying how much variability there is. This can be understood by mapping quality within fields, which was conducted in Chapters 2 and 3, and also by quantifying the magnitude and spatial structure of this variability within fields which was described in Chapter 5. Once variability has been described, the drivers of this variability (**WHY?**) can be assessed. The drivers of variation were explored in Chapter 4. Then, if the drivers of this variation can be addressed or mitigated, management decisions to optimise cotton fibre and grain quality can be implemented (**SO WHAT?**). This includes identifying fields where SSCM may be more amenable (Chapter 5), and the delineation between features within and outside of a grower's control to provide a retrospective N recommendation (Chapter 4). In this final chapter, these themes will be discussed, and future directions will be explored.

6.1 HOW?

One of the first stages for the progressive adoption of PA is measuring where and how much crop production varies within a field, across a farm, and between seasons (Whelan *et al*. 2013). While growers often inherently know about the variation occurring within fields, including which areas consistently perform better than others or where quality issues are likely to arise, maps of cotton fibre and grain quality can provide growers and advisors with an objective description of this variability. These maps can then be used to quantify how much variability is occurring, better understand why this variability is occurring (6.2), and what can be done to manage for or mitigate this variability (6.3). However, not all growers have access to grain

protein sensors, and on-the-go sensors for cotton fibre quality are not commercially available, meaning that maps of cotton fibre or grain quality are often not available. This is one of the biggest factors limiting the consideration of quality in PA.

For cotton fibre quality, the lack of an on-the-go sensor to map fibre quality at harvest is one of the greatest challenges. However, in the absence of an on-the-go sensor, fibre quality measurements can still be tracked back to in-field locations through the use of Radio Frequency Identification (RFID) tags and John Deere's Harvest Identification (HID) system to produce a module-resolution map of cotton fibre quality measurements. Such a method would not be possible in grains due to differences in the way grain is harvested and a lack of RFID or tracking technology from the field to the combine-harvester, chaser-bin, and receival depot.

Chapter 2 presented a generalised geostatistical approach to downscale areal observations of cotton fibre quality data (length and micronaire) to a fine-resolution using publicly-available remote sensing imagery. Fine-resolution yield observations were also used as a surrogate to assess the downscaling algorithms performance, as no fine-resolution fibre quality data were available. The downscaling approach involved two key components; (i) the estimation of spatial trends in yield and quality using regression with fine resolution features such as remote sensing imagery, and (ii) the use of area-to-point kriging (A2PK) to downscale either the observations in the absence of a useful spatial trend model or the residuals from the trend model (if useful) from areal averages. A decision-making framework based on the root mean square error (RMSE) was used to determine whether random forest (RF) or multiple linear regression (MLR) models should be utilised for estimating spatial trends in the downscaling approach. Overall, RF models outperformed MLR models and thus were selected more frequently due to their lower RMSE values (RF RMSE = 0.0159 – 0.1059, MLR RMSE = 0.0187 – 0.2072). Correlations with remote sensing features were strongest for cotton fibre yield, followed by micronaire, and were weak for fibre length. As a result, spatial trends could be estimated with good model quality for cotton fibre yield and micronaire, but not for length. When fine-resolution yield data was used as a surrogate for cotton fibre quality, the downscaling approach performed better when validated at the module-resolution (RMSE

for the RF + A2PK approach = 0.1059) compared to the fine-resolution (RMSE for the RF + A2PK approach = 0.3222). This result was not surprising given the unique narrow and long size and shape of the cotton modules. In fields where strong correlations existed between cotton fibre yield, micronaire, and remotely-sensed features, trend models alone performed best. In fields with weaker correlations, combining trend models with A2PK of the residuals improved performance.

For wheat GPC, combine harvester-mounted, on-the-go grain protein sensor are commercially available. Yet, their slow uptake by growers means that maps of GPC are not available for every field, farm, or season. Chapter 3 presented a machine learning approach to predict GPC within fields to fill-in gaps where grain protein sensor data was not available using a combination of readily-available yield, agronomic, and publicly-available data layers. In some cases, grain protein sensors are not available for an entire field, but in other cases where multiple headers are operating (e.g. in Australia or the United States of America [USA]), only one header may be fitted with a protein sensor. This leaves gaps in data collection across parts-of or for an entire field, and may fail to capture or describe the variability occurring. The best quality predictions of GPC were achieve using a combination of yield, agronomic, and publicly-available data layers, with model performance ranging from Lin's concordance correlation coefficient (LCCC) 0.51 to 0.96, and RMSE 0.47 to 1.67. Model quality was generally higher when predictions were validated at the spatial resolution of management zones (LCCC = 0.45 – 0.99, RMSE = 0.25 – 1.87) compared to at a fine-resolution (LCCC = 0.51 – 0.96, RMSE = 0.59 – 1.67). Similarly, interpolation within fields (i.e. the two-fold cross validation [2FCV] approach) performed better (LCCC = 0.90 – 0.99, RMSE = 0.25 – 0.73) than extrapolation onto unsampled fields (i.e. leave one Field-Year out cross-validation [LOFYOCV] approach; LCCC = 0.75 – 0.82, RMSE = 0.76 – 1.87).

Both Chapters 2 and 3 utilised existing spatial data layers to map cotton fibre or grain quality, including information that was already being recorded during agronomic operations (i.e. sowing and harvest dates, variety) or was publicly-available (e.g. a digital elevation model [DEM] or remotely-sensed vegetation indices). Bespoke on-farm data such as point-based soil observations can be used to model and map cotton or grain crops within fields (e.g. Jones *et*

al. 2022), but the collection and processing of this on-farm data can also be time-consuming and expensive (Tilse *et al*. 2022). Chapters 2 and 3 demonstrated how existing data layers (either on-farm or publicly-available) can be used to model and map cotton fibre and GPC, and provide growers with access to maps describing variability in cotton fibre and grain quality in the absence of on-the-go sensors.

On their own, maps of yield or quality are useful to describe variation occurring within fields and can be used to identify consistently high-performing or high-quality areas, or regions with persistent issues over multiple seasons. However, the opportunity for managing this variation is based on the magnitude and spatial structure (i.e. organisation) of the variation. The magnitude of variation determines the potential to tailor input applications for greater economic and environmental gains compared to uniform practice. At the same time, the spatial structure – whether variability forms coherent patterns or appears as random noise – plays a critical role in management feasibility (Pringle *et al*. 2003). For example, cohesive and clear spatial patterns make it easier for farm machinery such as variable-rate sprayers to differentially apply inputs and match variability. In Chapter 5, the spatial structure and magnitude of variation of wheat grain yield and GPC was calculated using the OI metric to compare the opportunities for SSCM between yield and protein. While both yield and GPC had a similar spatial structure of variation, the magnitude of variation was greater for yield than for GPC. Unlike simple descriptive statistics of variability, such as the coefficient of variation (CV) which are non-spatial, the OI describes how variability is spatially structured in a field and uses this assessment of spatial structure to assess if SSCM can be employed based on the operating width and response time of available machinery. Importantly, the OI ranks fields as to their suitability for SSCM and provides both a quantification of how yield or quality varies within a field, but also provides a metric that can be used in the decision making process for the stepwise adoption of PA (see 6.3).

Data availability and resolution was one of the greatest limiting factors for validating the modelling and downscaling approach for cotton fibre quality. In the absence of high-resolution cotton fibre quality data, cotton yield maps were used as a surrogate to validate the down-scaling approach, thus it was assumed that yield and fibre quality vary in the same way. The

variability in cotton fibre yield and quality has been compared using simple, non-spatial metrics of variation including the CV. Studies by Elms *et al*. (2001), Ge *et al*. (2007), and Johnson *et al*. (2002) all concluded that cotton yields exhibited considerably greater CVs than cotton fibre quality attributes such as length and micronaire. However, as demonstrated in Chapter 4 and 6.1, non-spatial metrics do not sufficiently characterise or describe variation, nor the suitability of this variation for SSCM. Johnson *et al*. (2002) used variograms to summarise the spatial correlation of cotton fibre yield and quality attributes. The authors concluded that cotton fibre micronaire consistently displayed spatial correlation over multiple seasons, whereas fibre yield was spatially correlated in some seasons but not others (Johnson *et al*. 2002). However, this analysis was limited to one field and the spatial structure of variability is not well understood over multiple fields, regions, or seasons. Given that wheat grain yield and GPC do not always vary in the same way and are not always inversely related (Whelan *et al*. 2009, Chapter 4), it can be assumed that similar differences between cotton fibre yield and quality will also exist in different areas of a field or in different seasons (e.g. between wet and dry). In the future, more work is needed to improve the validation of the downscaling approach for cotton fibre quality, including the quantification and comparison of the spatial structure and magnitude of variation between cotton fibre quality and yield for multiple fields and seasons.

A lack of available yield or quality sensors is not a challenge that is unique to the cotton or grains industries. Manually-harvested horticultural crops such as citrus face a similar downscaling challenge, as yield or quality data is often captured for a bin which aggregates data from several trees, for example. Alternatively, yield or quality data may only be presented as a block or field average (e.g. sugarcane). This downscaling approach is generalised and can be transferrable across a range of crops with areal-averaged data if yield or quality information can be geolocated to a block or several trees. However for many horticultural crops, quality is as important as quantity in determining crop value and there can be a large difference between total production yield and the actual marketable yield as the minimum quality thresholds must be met for produce to be marketable (Longchamps *et al*. 2022). Likewise, collecting data on quality parameters at harvest is a considerable challenge for horticultural crops due to differences in quality parameters between specific markets, and differences in the way

quality parameters are measured (e.g. visually, destructively or non-destructively, in situ or postharvest; Longchamps *et al*. 2022). If total yield or quality data simply needs to be downscaled from a block or field average to a fine-resolution, this generalised approach may be applied. However, if particular quality thresholds and differences between total production and actual marketable yield need to be considered, additional work is needed to capture these nuances.

6.2 WHY?

Once variability has been described, the next step in the progressive adoption of PA is to determine the major causes of this variability. Yield or quality maps can describe variability in crop performance within a field, but they give little indication about the reasons for this variability (Whelan *et al*. 2013). As widely discussed throughout this thesis, variability in wheat GPC is determined by complex interactions between genetic, environmental, and management features, including soil moisture and N availability, variety, and temperature (Whelan *et al*. 2013). This variability can occur across a range of spatial and temporal scales (e.g. within and between fields, farms, and seasons), and can be influenced by features both inside (e.g. nutrient applications) or outside (e.g. rainfall, temperature) a grower's control. Grain yield and GPC are generally understood to be negatively correlated, which is the result of grain protein dilution that is largely driven by soil moisture and N availability (Simmonds 1995; Holford *et al*. 1992). Additional factors such as the presence of soil constraints also limits a crops access to soil moisture and N (Dang *et al*. 2006). However, this is not always the case, and both Whelan *et al*. (2009) and Chapter 4 in this thesis demonstrate that the relationship between yield and GPC can be highly variable within and between fields, and season-to-season . The site-specific nature of this variability means that the drivers of variation also need to be considered on a field-by-field basis.

Chapter 4 sought to address research question 3 and used a range of different approaches to better understand why GPC varies. These approaches ranged from simple moving window correlations between yield and GPC to investigate local relationships, through to more

complex interpretive machine learning (IML) approaches which decomposed predictions of GPC into different input features within a model. Over the entire dataset, yield and GPC were negatively correlated (r^2 = -0.54) which reflects the expected grain protein dilution effect by yield (Simmonds 1995). Yet, when a 90 m moving window was used to calculate local correlations between wheat grain yield and GPC within fields, yield and GPC were both positively and negatively correlated in different areas of a field. While this simple approach highlighted considerable variability in the yield-GPC relationship within fields, it failed to capture interactions between features or clarify the nature of the relationship. For example, it was unclear whether the inverse relationship reflected a low-yield/high-protein scenario, likely due to moisture stress or high N rates, or a high-yield/low-protein scenario, associated with increasing N supply when soil moisture is not limiting (Terman *et al*. 1969; Holford *et al*. 1992; Simmonds 1995; Whelan *et al*. 2009).

SHapley Additive exPlanations (SHAP) values are an IML technique and were used to discern how a prediction of GPC was derived, identify the importance of individual features in the model, and decompose GPC predictions into the contributions made by different input features. The IML analysis was run for each field individually, but overall low total applied N rates decreased GPC predictions, whereas low yield values increased GPC predictions. These results were not surprising and further reflected the non-linear interactions between GPC, yield, soil moisture, and N. At low N rates, increasing N supply usually increases yield but reduces GPC due to carbohydrate dilution. As N availability rises, both yield and GPC increase together, and at high N levels, GPC typically increases whereas yields decrease (Holford *et al*. 1992; Simmonds 1995; Terman *et al*. 1969).

While the results from the IML approach were not surprising, it did allow for the differentiation between variability attributable to different features. Nitrogen is commonly employed to enhance wheat production to maximise yields and optimise GPC (Wang *et al*. 2023), and maps of GPC are often used to understand N adequacy for the previous season (Scott 2022). However, N influences grain yield and GPC in different ways under different soil moisture conditions (Terman *et al*. 1969), and using a GPC alone to guide N application will likely fail to capture these complex interactions. While the implementation of an IML approach by

growers for understanding variability in GPC is a way off, this research highlights the value in untangling the complex interactions between environmental and management features to understand how individual features drive variability in GPC for improved management (see 6.3).

6.3 SO, WHAT?

Even with the vast amounts of data being collected on-farms and by the industry, it remains that many growers 'don't know what to do with their maps' (Bramley *et al*. 2019). The translation of this data into useful decision-making tools is the next step in the adoption of PA on-farm. Once variability in quality has been described (how), and the drivers of this variability identified (why), management decisions to optimise quality can be implemented. These decisions may include optimising inputs to reduce variability, maximise gross margins, minimise environmental impacts, or improve the quality of marketed grain through selective harvest or logistics management (Whelan *et al*. 2013).

6.3.1 Stepwise adoption of PA

The OI metric from Chapter 5, and the IML approach from Chapter 4, can be used together to guide decision-making on-farm. First, the OI can be used to identify fields that are most suitable for SSCM. The OI quantifies both the structure and magnitude of variability, and assesses if there are opportunities for SSCM by considering if this variation can be addressed using available machinery. The OI provides a clear ranking of fields where SSCM is more amenable, and provides growers and advisors with a direct assessment of where the stepwise adoption of SSCM and PA would be best directed on-farms. Then within fields that have been identified as having a high opportunity for SSCM, the SSCM of variability can be identified, and features can be differentiated between those within and outside of the growers control. For those features that can be controlled by growers, a retrospective recommendation of inputs can be calculated and compared to actual inputs applied by the grower in the previous season (see below). If these analyses can be conducted over multiple seasons, more temporally stable

variability can be differentiated from seasonal effects, and this knowledge base can be used to inform future management decisions based on past experience.

6.3.2 Optimising inputs

Suboptimal N management is a main contributor to the yield gap in Australia (Hochman *et al*. 2018), and N application is commonly used to optimise GPC and maximise yields (Wang *et al*. 2023). Yet estimating the optimal N fertiliser rate remains a challenging task for farmers and advisors. There can be considerable spatial and temporal variability in crop nutrient demand, and in supply by the soil (Colaço *et al*. 2024). Also, complex interactions between genetic, environmental, and management features throughout the growing season mean that it is difficult to isolate and differentiate between the impact of N management, or if variability is due to other inherent features outside of a growers control. Volatile input costs including the rising cost of urea (Hogan 2024) are also increasing pressure on growers to reduce or optimise inputs.

In Chapter 4 of this thesis, IML and SHAP values were used to decompose GPC predictions into different components attributable to each feature. This was then used to provide a retrospective analysis of the amount of N that should have been applied for the previous season. The goal was to assess how N should have been managed, given the inherent landscape and seasonal conditions experienced. A DEM was used for elevation which relates to water flows, frost, and soil type, and a soil apparent electrical conductivity (EC_a) survey was used as a proxy for the soils water holding capacity. Wheat grain yield maps can provide a comprehensive overview of crop production throughout the whole growing season and were used to represent all other inherent seasonal conditions that the crop experienced (e.g. temperature, rainfall, soil constraints). Overall, the IML approach recommended a lower N-rate than what was applied by the grower and that N was overapplied for the previous season. Varying the total applied N rate could adjust GPC values by a maximum of ~$\pm$2.5%, but an economic analysis was not conducted and should be incorporated in future research to evaluate the feasibility of management approaches (e.g. Colaço *et al*. 2024.

While the IML focused on how N inputs could be adjusted to optimise GPC, a similar approach could be used for other manageable features that influence variability in GPC or yield. For example, foliar fungal diseases such as yellow rust can result in considerable yield losses (Whetton *et al*. 2018). Conventionally, fungicides are applied at a uniform rate across fields, but variable rate fungicide applications have been shown to increase gross margins and improve the quantity and quality of grain yields (Whetton *et al*. 2018). There may be the potential to apply a similar IML approach to understand the impact of variable rate fungicide applications and how application rates can be adjusted to avoid disease with minimum fungicide use, or reduce variability in grain yield or quality due to the occurrence of disease.

6.4 Future directions

6.4.1 Applications beyond cotton fibre quality and wheat grain protein content

The adoption of PA practices is limited for many manually harvested fruit and vegetable crops, and broadacre commodities like sugarcane. This is because, like for cotton modules, data is aggregated into bins across several trees, to management zones, or to a block (i.e. a field or standard of trees managed together). The downscaling approach presented in Chapter 2 has the potential to be applied to any areal-averaged yield or quality data for a range of crops, as long as there is the ability to georeference the harvesting units (i.e. bins or bags) with the in-field location using systems such as RFID tags (Longchamps *et al*. 2022; Ampatzidis *et al*. 2009).

The within-field spatial variability in quality attributes in other grain crops, including protein content in barley or oil content in canola, is poorly understood. John Deere's HarvestLab 3000^{TM} can also be used to measure barley protein or canola oil content on-the-go at harvest. There is an opportunity to use this data to better understand how and why quality attributes

for other grain crops vary within fields, or what SSCM strategies are available to optimise inputs that maximise both yield and quality across the whole farming system.

6.4.2 Interpretive machine learning for cotton fibre quality

There is an opportunity to apply IML approaches to understand the drivers of spatial variability in cotton fibre quality. However, high-quality, high-resolution data is needed (Jones *et al*. 2022). Prototype harvester-mounted fibre quality sensors have been developed to measure micronaire (Schielack *et al*. 2016) and trash (Long *et al*. 2020) on-the-go using near-infrared (NIR) spectroscopy, but they have not been commercialised. This is because the presence of seed, foreign matter, and varying moisture content on the harvester means more work is needed to implement these sensors onboard harvesters for on-the-go measurement (Hardin *et al*. 2022). Until on-the-go cotton fibre quality sensors are developed and made commercially available, mapping module-resolution fibre quality data (e.g. from commercial PA company "Precision Cropping Technologies", (PCT, https://pct.ag/) and improving methods to downscale this data to a fine-resolution remains the approaches that can be used by growers, advisors, and researchers to better understand variability in cotton fibre quality. Fibre quality data for Chapter 2 was limited to one season, but if multiple seasons worth of cotton fibre quality data can be collected for the same field, both spatial and temporal trends in fibre quality can be assessed. Ongoing and future collaborations with commercial companies such as PCT and John Deere present an opportunity to build a large database of fibre quality data to understand these temporal trends, and improve the extension of research outcomes to growers and advisors in industry.

6.4.3 Using diverse spatiotemporal data for interpretive machine learning

The spatial datasets used in the IML approach (Chapter 4) were limited to existing information available for each farm, including yield maps, electromagnetic (EM) surveys, a publicly-available DEM, and total applied N maps. The DEM, yield and EM survey maps were used

as proxies to represent all inherent features influencing variability in GPC throughout the growing season, but they cannot be used to directly quantify the impact of soil, landscape, or climatic features on GPC within fields. Alternatively, spatial data layers including digital soil maps of constraints and maps of soil available water content are direct measures of soil properties. Variability in soil water and constraints has a strong impact on yield and protein (Orton *et al*. 2018; Dang *et al*. 2010). These spatial data layers could be used in future within a large space-time datacube (Filippi *et al*. 2019b) to better understand and quantify the impact of different soil, landscape, or management features on variability in GPC within-fields.

A range of different approaches are available to estimate N application rates, and Colaço *et al*. (2024) compared 13 different approaches to develop N recommendations. These ranged from simplified mass-balance approaches, through to data-rich digital methods based on empirical models. While empirical, multivariate, data-driven methods were more accurate and increased the profitability of fertiliser management, they also relied on upon extensive on-farm digital databases including historical yield data, soil surveys, and remotely-sensed vegetation indices (Colaço *et al*. 2024). There is a trade-off between accuracy, model complexity, and interpretation, and the addition of lots of features into the model may make interpretation more difficult. While SHAP values are useful for explaining machine learning models, they struggle to capture the interactions between features. For example, in the IML approach presented in Chapter 4, both total applied N and the soils water holding capacity (WHC) will interact to determine the final GPC. If the soils WHC is high but total applied N is low, then it is assumed that grain yields would be high and the GPC will be low. Conversely, if both the soils WHC the total applied N is high, then both yield and GPC will likely be high (Simmonds 1995; Holford *et al*. 1992). The SHAP values merge the individual and interaction effects together into a single number, which means we cannot tell what features interact, nor differentiate between the individual and interactive influence of a feature (Muschalik *et al*. 2024; Muschalik 2024). Shapley interactions offer an approach to address these limitations and decompose the effects of features into individual features and the interactions between features. Shapley interactions should be explored in future research to better understand how different features interact to drive variability in wheat GPC within fields.

6.4.4 Modelling interactions between wheat grain protein content and yield

Wheat grain yield and GPC are strongly interrelated, and an inverse relationship is generally expected. Under non-limiting soil moisture conditions, as the amount of carbohydrates in the grain (yield) increases, the protein concentration decreases (Simmonds 1995). Likewise, as N increases, yield and GPC will both increase up until a certain point, after which yields begin to plateau while GPC continues to increase. At very high N levels, a yield penalty may be incurred (Holford *et al*. 1992). While this relationship is relatively well-understood, yield and protein are generally modelled independently when predicting or interpreting model results. Yield may be included as a feature within a model for GPC (e.g. Chapter 3), but the interrelationship between yield and protein is often not considered. The moving window correlation analysis between yield and protein revealed that yield and GPC do not always vary in the same way, to the same magnitude, or with the same spatial structure within fields, and that yield and GPC can be positively and negatively correlated in different parts of a field (Chapter 4).

It is possible to model both yield and GPC together to capture their interactions. Yield and protein data points are co-located for the same season, meaning that they are influenced by the same conditions, but are impacted in different ways. This includes how different levels of soil moisture and applied N influence yield and GPC differently (Holford *et al*. 1992; Simmonds 1995). Modelling both yield and GPC can provide an understanding of how they interact, how different features drive these interactions, and could also be used to provide a retrospective N recommendation to maximise yields while also optimising GPC. Filippi *et al*. (2018) used bivariate linear mixed models to model co-located soil pH samples taken at two different time points (2002 and 2015), and Orton *et al*. (2014) used a multivariate geostatistical approach (linear model of coregionalization) to model the covariance between soil organic carbon and bulk density data which were sampled at the same locations. Both approaches can handle two (or more) correlated variables that exhibit spatial correlation, and treat each variable (either the different time points Filippi *et al*. 2018, or soil property Orton *et al*. 2014) as separate columns within a dataframe/datacube. This is conceptually similar to modelling

with co-located yield and GPC observations from the same season. Both approaches had high predictive power with low uncertainty (Filippi *et al*. 2018; Orton *et al*. 2014), but they are computationally intensive and difficult to implement for large datasets across a large area. Alternatively, co-located yield and protein data can be stacked in a datacube, with crop property (i.e. yield or protein) used as a feature in the model. A similar approach was used by (Filippi *et al*. 2019b) to forecast grain crop yield for wheat, barley, and canola, where crop type was included as a feature. A machine learning approach such as that used by Filippi *et al*. (2019b) is less computationally intensive compared to the bivariate linear mixed model geostatistical approaches used by Filippi *et al*. (2018) and Orton *et al*. (2014). This would mean we could model the GPC and yield jointly in the same model and get spatially coherent predictions.

6.4.5 Forecasting grain protein content

The modelling approaches described in this thesis focused on hindcasting yields, GPC, and cotton fibre quality for the previous season (Chapter 2 – 4). There has been growing interest in forecasting cotton and grain yields to guide adaptive crop management strategies at critical growth stages. For example, mid-season yield forecasts may be useful for the variable-rate application of fertilisers to match expected crop nutrient requirements (Filippi *et al*. 2020b; Han *et al*. 2022), or forecasts made early in the season can be useful for logistics or futures contracting (Filippi *et al*. 2019b). At this stage, more work is needed to better how and why cotton fibre varies spatially within fields before forecasting cotton fibre quality can be explored. In grains, it may be possible to apply data-driven modelling approaches similar to Filippi *et al*. (2019b) and Filippi *et al*. (2020b) to forecast GPC within fields. Early-to-mid season forecasts of GPC would be beneficial to guide mid-season top-up applications of N, or to plan selective harvest or blending strategies of high and low-quality grain.

6.4.6 Translating data to decisions

Fertiliser accounts for ~20 – 30% of total farm cash costs in Australian cropping systems (DAFF 2023), and prices fluctuate considerably due to global energy markets and trade. The world is also transitioning towards net-zero and there is increased scrutiny on emissions from agriculture by governments and consumers. A significant portion of this scrutiny is on synthetic N fertilisers which accounted for 2.1% of global greenhouse gas (GHG) emissions (Menegat *et al*. 2022) and 8.3% of farm-gate emissions in 2019 (FAO 2021). Suboptimal N management is also a main contributor to the yield gap in Australia (Hochman *et al*. 2018), and N management is a key strategy to optimise GPC (Wang *et al*. 2023). The optimisation of N fertiliser inputs offers growers the opportunity to maximise profits while minimising GHG emissions. This will be a key focus as we move into the future of farming, especially if reduced emissions are needed to access certain markets or achieve premiums as the world moves to net zero. Precision agriculture will play an integral role in addressing these future challenges for Australian agriculture, including for optimising input applications, or for the recording of inputs within fields for supply-chain and emissions traceability.

6.5 Concluding remarks

This thesis has demonstrated how the growing abundance and diversity of on-farm and publicly-available spatial data layers can be leveraged to better understand, characterise, and manage variability in cotton fibre and grain quality, and address key challenges faced by Australian growers. Maps of cotton fibre and GPC at a fine-resolution provide growers with the necessary tools to describe variability within fields. By integrating PA principles with a range of analytical techniques, this research builds upon these maps to provide insights into how and why grain quality varies within fields, and applications of this knowledge for improved management. These findings highlight the important role PA can play in enhancing the profitability, efficiency, and sustainability of Australian farming systems. This is particularly imperative as growers are faced with increasing pressures from economic, environmental, and market challenges. This thesis highlights the increasing need to shift

management focus from quantity to quality, and the knowledge that can be gained within farming systems through the utilisation and application of PA.

References

Adams, M.L., Cook, S.E., Caccetta, P.A. and Pringle, M.J. (1999). 'Machine Learning Methods in Site-Specific Management Research: An Australian Case Study'. In: *Proceedings of the fourth international conference on precision agriculture*. Wiley Online Library, pp. 1321–1333.

Agriculture Victoria (2017). *Summary report: Reducing on-farm nitrous oxide emissions through improved nitrogen use efficiency in grains*. Economic Development, Jobs, Transport and Resources.

Akaike, H. (1998). 'Information Theory and an Extension of the Maximum Likelihood Principle'. In: *Selected Papers of Hirotugu Akaike*. Ed. by Parzen, E., Tanabe, K. and Kitagawa, G. New York, NY: Springer New York, pp. 199–213. ISBN: 978-1-4612-1694-0. DOI: 10.1007/978-1-4612-1694-0_15. URL: https://doi.org/10.1007/978-1-4612-1694-0_15.

Ali, A., Hassan, M.U. and Kaul, H.P. (2024). 'Broad Scope of Site-Specific Crop Management and Specific Role of Remote Sensing Technologies Within It—A Review'. In: *Journal of Agronomy and Crop Science* 210.4, e12732. DOI: 10.1111/jac.12732.

Amin, E., Pipia, L., Belda, S., Perich, G., Graf, L.V., Aasen, H., Van Wittenberghe, S., Moreno, J. and Jochem, V. (2024). 'In-season forecasting of within-field grain yield from Sentinel-2 time series data'. In: *International Journal of Applied Earth Observation and Geoinformation* 126, p. 103636. ISSN: 1569-8432. DOI: 10.1016/j.jag.2023.103636. URL: https://www.sciencedirect.com/science/article/pii/S1569843223004600.

Ampatzidis, Y.G. and Vougioukas, S.G. (2009). 'Field experiments for evaluating the incorporation of RFID and barcode registration and digital weighing technologies in manual fruit harvesting'. In: *Computers and Electronics in Agriculture* 66.2, pp. 166–172. ISSN:

0168-1699. DOI: 10.1016/j.compag.2009.01.008. URL: https://www.sciencedirect.com/science/article/pii/S0168169909000210.

Anderson, W.K., Hamza, M.A., Sharma, D.L., D'Antuono, M.F., Hoyle, F.C., Hill, N., Shackley, B.J., Amjad, M. and Zaicou-Kunesch, C. (2005). 'The role of management in yield improvement of the wheat crop—a review with special emphasis on Western Australia'. In: *Australian Journal of Agricultural Research* 56.11, pp. 1137–1149. DOI: 10.1071/AR05077.

Anderson, W.K., Stephens, D. and Siddique, K.H.M. (2016). 'Dryland agriculture in Australia: experiences and innovations'. In: *Innovations in dryland agriculture*, pp. 299–319. DOI: 10.1007/978-3-319-47928-6_11.

Australia, Grain Trade (2024). 'Section 2 – Wheat Trading Standards 2024/25 Season'. In: *GTA Wheat Trading Standards 2024/25*.

Aybar, C., Wu, Q., Bautista, L., Yali, R. and Barja, A. (2020). 'rgee: An R package for interacting with Google Earth Engine'. In: *Journal of Open Source Software*. DOI: 10.21105/joss.02272. URL: github.com/r-spatial/rgee/.

Baird, J., Bell, L., Lawrence, D., Gentry, J., Erbacher, A., Hertel, K., Aisthorpe, D. and Lester, D. (2022). 'Comparing grain and cotton in northern NSW. Impacts on the cropping systems and the advantages of growing summer crops to improve $/mm and as a disease break from winter cereal dominated systems'. In: *GRDC 2022 Grains Research Update – North*. Grains Research and Development Corporation (GRDC). URL: https://grdc.com.au/resources-and-publications/grdc-update-papers/tab-content/grdc-update-papers/2022/03/comparing-grain-and-cotton-in-northern-nsw.-impacts-on-the-cropping-systems-and-the-advantages-of-growing-summer-crops-to-improve-$mm-and-as-a-disease-break-from-winter-cereal-dominated-systems.

Ballester, C., Brinkhoff, J., Quayle, W.C. and Hornbuckle, J. (2019). 'Monitoring the Effects of Water Stress in Cotton Using the Green Red Vegetation Index and Red Edge Ratio'. In: *Remote Sensing* 11.7. ISSN: 2072-4292. DOI: 10.3390/rs11070873. URL: https://www.mdpi.com/2072-4292/11/7/873.

Ballester, C., Hornbuckle, J., Brinkhoff, J., Smith, J. and Quayle, W. (2017). 'Assessment of In-Season Cotton Nitrogen Status and Lint Yield Prediction from Unmanned Aerial System Imagery'. In: *Remote Sensing* 9.11. ISSN: 2072-4292. DOI: 10.3390/rs9111149. URL: https://www.mdpi.com/2072-4292/9/11/1149.

Ballester, C., Jiménez-Bello, M.A., Castel, J.R. and Intrigliolo, D.S. (2013). 'Usefulness of thermography for plant water stress detection in citrus and persimmon trees'. In: *Agricultural and Forest Meteorology* 168, pp. 120–129. ISSN: 0168-1923. DOI: 10.1016/j.agrformet.2012.08.005. URL: https://www.sciencedirect.com/science/article/pii/S0168192312002572.

Bange, M., Nunn, C., Mahan, J., Payton, P., Milroy, S., Finger, N., Caton, J., Dodge, W. and Quinn, J. (2022). 'Improving temperature-based predictions of the timing of flowering in cotton'. In: *Agronomy Journal* 114.5, pp. 2728–2742. DOI: 10.1002/agj2.21086.

Bange, M.P., Sluijs, M.H.J. van der, Constable, G.A., Gordon, S.G., Long, R.L. and Naylor, G.R.S. (2018). *FIBREpak. From Seeds to Good Shirts: a Guide to Improving Australian Cotton Fibre Quality*. English. Narrabri, NSW: Cotton Research and Development Corporation (CRDC). ISBN: 9780977531714.

Barnes, E.M., Clarke, T.R., Richards, S.E., Colaizzi, P.D., Haberland, J., Kostrzewski, M., Waller, P., Choi, C., Riley, E., Thompson, T., Lascano, R., Li, H. and Moran, M. (2000). 'Coincident detection of crop water stress, nitrogen status and canopy density using ground based multispectral data'. In: *Proceedings of the fifth international conference on precision agriculture, Bloomington, MN, USA*. Ed. by Robert, P., Rust, R. and Larson, W.E. Vol. 1619. 6.

Bastos, L.M., Reis, A.F.B., Sharda, A., Wright, Y. and Ciampitti, I.A. (2021). 'Current status and future opportunities for grain protein prediction using on-and off-combine sensors: A synthesis-analysis of the literature'. In: *Remote Sensing* 13.24, p. 5027. DOI: 10.3390/rs13245027.

Beck, H., Zimmermann, N., McVicar, T., Vergopolan, N., Berg, A. and Wood, E.F. (2018). 'Present and future Köppen-Geiger climate classification maps at 1-km resolution'. In: *Scientific Data* 5.180214. DOI: 10.1038/sdata.2018.214.

Bel-Berger, P. and Roberts, G. (1998). 'Neps: How do they impact cotton quality?' In: *1998 Australian Cotton Conference*. Cotton Research and Development Corporation (CRDC). URL: http://27.111.91.222/xmlui/handle/1/10.

Blakeney, A.B., Cracknell, R.L., Crosbie, G.B., Jefferies, S.P., Miskelly, D.M., O'Brien, B.L., Panozzo, J.F., Suter, D.A.I., Solah, V., Watts, T., Westcott, T. and Williams, R.M. (2009). *Understanding Australian Wheat Quality*. Grains Research and Development Corporation (GRDC).

Bonfil, D.J., Mufradi, I., Asido, S. and Long, D.S. (2008). 'Precision nitrogen management based on nitrogen removal in rainfed wheat'. In: *Proceedings of the 9th International Conference on Precision Agriculture, Denver, CO, USA*, pp. 20–23.

Boydell, B and McBratney, A.B. (2002). 'Identifying potential within-field management zones from cotton-yield estimates'. In: *Precision Agriculture* 3, pp. 9–23. DOI: 10.1023/A:1013318002609.

Bradow, J.M. and Davidonis, G.H. (2000). 'Quantitation of fiber quality and the cotton production-processing interface: a physiologist's perspective'. In: *Journal of Cotton Science* 4.1, pp. 34–64. DOI: 10.1023/A:1010933404324.

Bramley, R.G.V., Mowat, D., Gobbett, D., Branson, M., Wakefield, A. and Wilksch, R. (2012). 'Mixing grapes and grain - Scoping the opportunity for selective harvesting in cereals'. In: *Capturing Opportunities and Overcoming Obstacles in Australian Agronomy: Proceedings of the 16th Australian Agronomy Conference*. Ed. by Yunusa, I. Australian Society of Agronomy.

Bramley, R.G.V. and Ouzman, J. (2019). 'Farmer attitudes to the use of sensors and automation in fertilizer decision-making: Nitrogen fertilization in the Australian grains sector'. In: *Precision agriculture* 20, pp. 157–175. DOI: 10.1007/s11119-018-9589-y.

Bramley, R.G.V., Proffitt, A.P.B., Hinze, C.J., Pearse, B. and Hamilton, R.P. (2005). 'Generating benefits from Precision Viticulture through selective harvesting'. In: *Precision Agriculture '05*. Wageningen Academic Publishers, pp. 891–898. DOI: 10.3920/978-90-8686-549-9_109.

Breiman, L. (2001). 'Random forests'. In: *Machine learning* 45, pp. 5–32. DOI: 10.1023/A:1010933404324.

Brodrick, R., Yeates, S., Roth, G., Gibb, D., Henggeler, S. and Wigginton, D. (2012). 'Managing irrigated cotton agronomy'. In: *WATERpak. A guide for irrigation management in cotton and grain farming systems*. Ed. by Dugdale, H., Harris, G., Neilsen, J., Richards, D., Wigginton, D. and Williams, D. 3rd ed. CottonInfo, pp. 248–263.

Brus, D.J., Boogaard, H., Ceccarelli, T., Orton, T.G., Traore, S. and Zhang, M. (2018). 'Geostatistical disaggregation of polygon maps of average crop yields by area-to-point kriging'. In: *European Journal of Agronomy* 97, pp. 48–59. ISSN: 1161-0301. DOI: 10.1016/j.eja.2018.05.003. URL: https://www.sciencedirect.com/science/article/pii/S1161030118301047.

Bryce, A. (2025). 'Vegetation indices maps for targeted product applications in cotton'. In: *Precision Ag News, Summer 2025*. Society of Precision Agriculture Australia. URL: https://www.spaa.com.au/membership-spaa/precision-ag-news-pan/latest-issue/.

Buttafuoco, G. and Lucà, F. (2016). 'The contribution of geostatistics to precision agriculture'. In: *Annals of Agricultural & Crop Science* 1.2, pp. 1–2.

Campbell, C.A., Davidson, H.R. and Warder, F.G. (1977). 'Effects of fertilizer N and soil moisture on yield, yield components, protein content and N accumulation in the aboveground parts of spring wheat'. In: *Canadian Journal of Soil Science* 57.3, pp. 311–327.

Cao, Q., Miao, Y., Shen, J., Yuan, F., Cheng, Sh. and Cui, Z. (2018). 'Evaluating Two Crop Circle Active Canopy Sensors for In-Season Diagnosis of Winter Wheat Nitrogen Status'. In: *Agronomy* 8.10. ISSN: 2073-4395. DOI: 10.3390/agronomy8100201. URL: https://www.mdpi.com/2073-4395/8/10/201.

Casa, R. and Castrignanò, A. (2008). 'Analysis of spatial relationships between soil and crop variables in a durum wheat field using a multivariate geostatistical approach'. In: *European Journal of Agronomy* 28.3, pp. 331–342.

Cathey, G.W. and Meredith, W.R. (1988). 'Cotton response to planting date and mepiquat chloride'. In: *Agronomy Journal* 80.3, pp. 463–466. DOI: 10.2134/agronj1988.00021962008000030014x.

Cato, L. and Mullan, D. (2020). 'Wheat quality: wheat breeding and quality testing in Australia'. In: *Breadmaking (Third Edition)*. Elsevier, pp. 221–259. DOI: 10.1016/B978-0-08-102519-2.00008-6.

Centre, Australian Export Grains Innovation (2022). 'Australian Wheat'. In: *AEGIC*.

Chen, P., Xu, W., Zhan, Y., Yang, W., Wang, J. and Lan, Y. (2022). 'Evaluation of cotton defoliation rate and establishment of spray prescription map using remote sensing imagery'. In: *Remote Sensing* 14.17, p. 4206.

Chen, T. and He, T. (2020). 'xgboost: Extreme Gradient Boosting'. In: *R package version 1.7.7.1*. DOI: 10.48550/arXiv.1603.02754.

Chew, B.J., Wiratama, W. and Goh, M.H. (2023). 'Canopy nitrogen estimation on cotton plant using satellite imagery'. In: *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* 48, pp. 73–79. DOI: 10.5194/isprs-archives-XLVIII-M-1-2023-73-2023. URL: https://isprs-archives.copernicus.org/articles/XLVIII-M-1-2023/73/2023/.

Christy, C.D. (2008). 'Real-time measurement of soil attributes using on-the-go near infrared reflectance spectroscopy'. In: *Computers and Electronics in Agriculture* 61.1. Emerging Technologies For Real-time and Integrated Agriculture Decisions, pp. 10–19. ISSN: 0168-1699. DOI: 10.1016/j.compag.2007.02.010. URL: https://www.sciencedirect.com/science/article/pii/S0168169907001639.

Cogliati, S., Sarti, F., Chiarantini, L., Cosi, M., Lorusso, R., Lopinto, E., Miglietta, F., Genesio, L., Guanter, L., Damm, A., Pérez-López, S., Scheffler, D., Tagliabue, G., Panigada, C., Rascher, U., Dowling, T.P.F., Giardino, C. and Colombo, R. (2021). 'The PRISMA imaging spectroscopy mission: overview and first performance analysis'. In: *Remote Sensing of Environment* 262, p. 112499. ISSN: 0034-4257. DOI: 10.1016/j.rse.2021.112499. URL: https://www.sciencedirect.com/science/article/pii/S0034425721002170.

Cohen, J. (1960). 'A Coefficient of Agreement for Nominal Scales'. In: *Educational and Psychological Measurement* 20.1. DOI: 10.1177/001316446002001.

Colaço, A.F., Trevisan, R.G., Karp, F.H.S. and Molin, J.P. (2020). 'Yield mapping methods for manually harvested crops'. In: *Computers and Electronics in Agriculture* 177, p. 105693.

ISSN: 0168-1699. DOI: 10.1016/j.compag.2020.105693. URL: https://www.sciencedirect.com/science/article/pii/S0168169920315453.

Colaço, A.F., Whelan, B.M., Bramley, R.G.V., Richetti, J., Fajardo, M., McCarthy, A.C., Perry, E.M., Bender, A., Leo, S., Fitzgerald, G.J. and Lawes, R.A. (2024). 'Digital strategies for nitrogen management in grain production systems: lessons from multi-method assessment using on-farm experimentation'. In: *Precision Agriculture* 25.2, pp. 983–1013. DOI: 10.1007/s11119-023-10102-z.

Cook, S.E. and Bramley, R.G.V. (1998). 'Precision agriculture—opportunities, benefits and pitfalls of site-specific crop management in Australia'. In: *Australian Journal of experimental agriculture* 38.7, pp. 753–763.

— (2001). 'Is agronomy being left behind by precision agriculture?' In: *Proceedings of the 10th Australian Agronomy Conference*. Australian Society of Agronomy.

Cook, S.E., Corner, R.J., Riethmuller, G., Mussel, G. and Maitland, M.D. (1996). 'Precision agriculture and risk analysis: An Australian example'. In: *Proceedings of the Third International Conference on Precision Agriculture*. Wiley Online Library, pp. 1123–1132.

Cotton Australia (2024). *Where is cotton grown?* URL: https://cottonaustralia.com.au/where-is-cotton-grown (visited on 16/12/2024).

Dang, Y.P., Dalal, R.C., Buck, S.R., Harms, B.R., Kelly, R., Hochman, Z., Schwenke, G.D., Biggs, A.J.W., Ferguson, N.J., Norrish, S., Routley R. McDonald, M., Hall, C., Singh, D.K., Daniells, I.G., Farquharson, R., Manning, W., Speirs, S., Grewal, H.S., Cornish, P., Bodapati, N. and Orange, D. (2010). 'Diagnosis, extent, impacts, and management of subsoil constraints in the northern grains cropping region of Australia'. In: *Soil Research* 48.2, pp. 105–119. DOI: 10.1071/SR09074.

Dang, Y.P., Menzies, N. and Dalal, R. (2022). *Soil Constraints to Crop Production*. Cambridge Scholars Publishing: Newcastle upon Tyne.

Dang, Y.P., Routley, R., McDonald, M., Dalal, R.C., Singh, D.K., Orange, D. and Mann, M. (2006). 'Subsoil constraints in Vertosols: crop water use, nutrient concentration, and grain yields of bread wheat, durum wheat, barley, chickpea, and canola'. In: *Australian Journal of Agricultural Research* 57.9, pp. 983–998. DOI: 10.1071/AR05268.

Davies, S., Page, K.L. and Armstrong, R. (2022). 'Australia'. In: *Soil Constraints to Crop Production*. Ed. by Dang, Y., Menzies, N. and Dalal, R. Cambridge Scholars Publishing: Newcastle upon Tyne, pp. 481–507.

Deloitte Access Economics Pty Ltd (2023). *Carbon border adjustment mechanisms: Implications for Australian agriculture*. Tech. rep. AgriFutures Australia.

Department of Agriculture, Fisheries and Forestry (2023). *Australian Government 2023 Report to the International Cotton Advisory Committee*. Tech. rep. Department of Agriculture, Fisheries and Forestry.

— (2024). *Land use and management*. URL: https://www.agriculture.gov.au/abares/aclump (visited on 16/12/2024).

Department of Climate Change Energy, the Environment and Water (2024). *Net Zero*. URL: https://www.dcceew.gov.au/climate-change/emissions-reduction/net-zero (visited on 16/12/2024).

Dodge, W.G. (2023). 'Digital phenotyping in cotton breeding using growth rate modelling based on visible light data collected with unmanned aerial systems'. Doctoral Dissertation. Texas AM University.

Donohue, R.J., Lawes, R.A., Mata, G., Gobbett, D. and Ouzman, J. (2018). 'Towards a national, remote-sensing-based model for predicting field-scale crop yield'. In: *Field Crops Research* 227, pp. 79–90. DOI: 10.1016/j.fcr.2018.08.005.

Durrence, J.S., Thomas, D.L., Perry, C.D. and Vellidis, G. (2005). 'Preliminary evaluation of commercial cotton yield monitors: the 1998 season in South Georgia.' In: *Proceedings of the Beltwide Cotton Conference 2005* (New Orleans, USA), pp. 366–372.

Eaton, F.M (1947). 'Nitrogen content of cotton in relation to other fiber properties'. In: *Textile Research Journal* 17.10, pp. 568–575. DOI: 10.1177/004051754701701004.

Edwards, J. and Roberts, K. (2008). 'Reproductive development'. In: *Wheat Growth and Development*. Ed. by White, J. and Edwards, J. NSW Department of Primary Industries, pp. 51–69.

Elms, M.K., Green, C.J. and Johnson, P.N. (2001). 'Variability of cotton yield and quality'. In: *Communications in Soil Science and Plant Analysis* 32.3–4, pp. 351–368. DOI: 10.1081/CSS-100103012.

Fajardo, M. and Whelan, B.M. (2021). 'Within-farm wheat yield forecasting incorporating off-farm information'. In: *Precision Agriculture* 22.2, pp. 569–585. DOI: 10.1007/s11119-020-09779-3.

Ferraciolli, M.A., Bocca, F.F. and Rodrigues, L.H.A. (2019). 'Neglecting spatial autocorrelation causes underestimation of the error of sugarcane yield models'. In: *Computers and electronics in agriculture* 161, pp. 233–240. DOI: 10.1016/j.compag.2018.09.003.

Fettell, N., Brill, R., Gardner, M. and McMullen, G. (2012). 'Yield and protein relationships in wheat'. In: *GRDC 2012 Grains Research Update – National*. Grains Research and Development Corporation (GRDC). URL: https://grdc.com.au/resources-and-publications/grdc-update-papers/tab-content/grdc-update-papers/2012/07/yield-and-protein-relationships-in-wheat.

Filippi, P., Cattle, S.R., Bishop, T.F.A., Odeh, I.O.A. and Pringle, M.J. (2018). 'Digital soil monitoring of top-and sub-soil pH with bivariate linear mixed models'. In: *Geoderma* 322, pp. 149–162. DOI: 10.1016/j.geoderma.2018.02.033.

Filippi, P., Han, S.Y. and Bishop, T.F.A. (2025). 'On crop yield modelling predicting, and forecasting and addressing the common issues in published studies'. In: *Precision Agriculture* 26 (8). DOI: 10.1007/s11119-024-10212-2.

Filippi, P., Jones, E.J. and Bishop, T.F.A. (2020a). 'Catchment-scale 3D mapping of depth to soil sodicity constraints through combining public and on-farm soil databases–A potential tool for on-farm management'. In: *Geoderma* 374, p. 114396. DOI: 10.1016/j.geoderma.2020.114396.

Filippi, P., Jones, E.J., Ginns, B.J., Whelan, B.M., Roth, G.W. and Bishop, T.F.A. (2019a). 'Mapping the depth-to-soil pH constraint, and the relationship with cotton and grain yield at the within-field scale'. In: *Agronomy* 9.5, p. 251. DOI: 10.3390/agronomy9050251.

Filippi, P., Jones, E.J., Wimalathunge, N.S., Somarathna, P.D.S.N., Pozza, L.E., Ugbaje, S.U., Jephcott, T.G., Paterson, S.E., Whelan, B.M. and Bishop, T.F.A. (2019b). 'An approach to forecast grain crop yield using multi-layered, multi-farm data sets and machine learning'. In: *Precision Agriculture* 20, pp. 1015–1029. DOI: 10.1007/s11119-018-09628-4.

Filippi, P., McPherson, T. and Bishop, T.F.A (2024a). 'Pulse vs . cereal crop variability: A comprehensive analysis with precision agriculture data across Australia'. In: *2024 Agronomy Australia Conference, 21-24 October 2024*. Australian Society of Agronomy.

Filippi, P., Whelan, B.M and Bishop, T.F.A. (2024b). 'Explainable Machine Learning to Map the Impact of Weather and Soil on Wheat Yield and Revenue Across the Eastern Australian Grain Belt'. In: *Agriculture* 14 (12). DOI: 10.3390/agriculture14122318.

Filippi, P., Whelan, B.M., Vervoort, W.R. and Bishop, T.F.A. (2020b). 'Mid-season empirical cotton yield forecasts at fine resolutions using large yield mapping datasets and diverse spatial covariates'. In: *Agricultural Systems* 184, p. 102894. DOI: 10.1016/j.agsy.2020.102894.

Food and United Nations, Agriculture Organization of the (2021). 'Emissions from agriculture and forest land. Global, regional and country trends 1990-2019'. In: *FAOSTAT Analytical Brief Series No 25*.

French, R.J. and Schultz, J.E. (1984). 'Water use efficiency of wheat in a Mediterranean-type environment. I. The relation between yield, water use and climate'. In: *Australian Journal of Agricultural Research* 35.6, pp. 743–764. DOI: 10.1071/AR9840743.

Fuhrer, L. (2022). 'Mapping of in-field cotton fiber quality utilizing John Deere's Harvest Identification System (HID)'. Master's thesis. University of Georgia.

Fuhrer, L., Porter, W., Barnes, E.M., Rains, G.C., Snider, J.L., Virk, S. and Ward, J.K. (2024). 'Utilizing John Deere's Harvest Identification System in Cotton Fiber Quality Mapping'. In: *Applied Engineering in Agriculture*. DOI: 10.13031/aea.15893.

Fuhrer, L., Porter, W.M., Rains, G., Snider, J. and Barnes, E. (2020). 'Mapping in-field Cotton Fiber Quality Utilizing John Deere Harvest Identification System'. In: *Proceedings of the Beltwide Cotton Conferences*. National Cotton Council of America, pp. 297–301.

Gallant, J., Wilson, N., Tickle, P.K., Dowling, T. and Read, A. (2009). *1 second SRTM Derived Digital Elevation Model (DEM)*. URL: https://pid.geoscience.gov.au/dataset/ga/72759.

Ge, Y., Thomasson, A.J. and Sui, R. (2007). 'Spatial variability of fiber quality in a dryland cotton field'. In: *Proceeding of the Beltwide Cotton Conferences* (Memphis, TN). National Cotton Council of America, pp. 929–937.

Ge, Y., Thomasson, A.J., Sui, R., Morgan, C.L., Searcy, S.W. and Parnell, C.B. (2008). 'Spatial variation of fiber quality and associated loan rate in a dryland cotton field'. In: *Precision Agriculture* 9, pp. 181–194. DOI: 10.1007/s11119-008-9064-2.

Gemtos, T.A., Markinos, A.T. and Nassiou, T. (2005). 'Cotton lint quality spatial variability and correlation with soil properties and yield'. In: *Precision Agriculture '05*. Ed. by Stafford, J. Wageningen Academic, pp. 361–368. DOI: 10.3920/978-90-8686-549-9_045.

Gitelson, A. and Merzlyak, M.N. (1994). 'Quantitative estimation of chlorophyll-a using reflectance spectra: Experiments with autumn chestnut and maple leaves'. In: *Journal of Photochemistry and Photobiology B: Biology* 22.3, pp. 247–252. DOI: 10.1016/1011-1344(93)06963-4.

Goovaerts, P. (2008). 'Kriging and Semivariogram Deconvolution in the Presence of Irregular Geographical Units'. In: *Mathematical Geosciences* 40, pp. 101–128. DOI: 10.1007/s11004-007-9129-1.

Gordon, S.G., Sluijs, M.H.J. van der and Prins, M.W. (2004). *Quality Issues for Australian Cotton from the Mill Perspective*. Tech. rep. Commonwealth Scientific and Industrial Research Organisation (CSIRO).

Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D. and Moore, R. (2017). 'Google Earth Engine: Planetary-scale geospatial analysis for everyone'. In: *Remote Sensing of Environment* 202, pp. 18–27. DOI: 0.1016/j.rse.2017.06.031.

Goss, M.J. and Oliver, M. (2023). *Encyclopedia of Soils in the Environment*. Ed. by Goss, M.J. and Oliver, M. 2nd ed. Elsevier.

Gourlot, J.P., Gérardeaux, E., Frydrych, E., Gawrysiak, G., Francalanci, P., Gozé, E., Dréan, J.Y. and Liu, R. (2005). 'Sampling issues for cotton fiber quality measurements. Part 2: Impact on cotton testing instruments results'. In: *Proceedings of the Beltwide Cotton Conference 2005* (New Orleans, USA).

Government, Queensland (2024). *SILO – Australian climate data from 1889 to yesterday*. URL: https://www.longpaddock.qld.gov.au/silo/gridded-data/ (visited on 17/12/2024).

Gozdowski, D., Leszczyńska, E., Stępień, M., Rozbicki, J. and Samborski, S. (2017). 'Within-field variability of winter wheat yield and grain quality versus soil properties'. In: *Communications in Soil Science and Plant Analysis* 48.9, pp. 1029–1041. DOI: 10.1080/00103624.2017.1323091.

Grains Australia (2024). *Wheat Classes*. URL: https://grainsaustralia.com.au/classification/wheat/wheat-process/wheat-classes-2 (visited on 17/12/2024).

Haan, S., Harianto, J., Butterworth, N. and Bishop, T. (2023). 'Geodata-Harvester: A Python package to jumpstart geospatial data extraction and analysis'. In: *Journal of Open Source Software* 8.89, p. 5205. DOI: 10.21105/joss.05205.

Habibi, L.N., Matsui, T. and Tanaka, T.S.T. (2024). 'Critical evaluation of the effects of a cross-validation strategy and machine learning optimization on the prediction accuracy and transferability of a soybean yield prediction model using UAV-based remote sensing'. In: *Journal of Agriculture and Food Research* 16, p. 101096. DOI: 10.1016/j.jafr.2024.101096.

Han, S.Y., Bishop, T.F.A. and Filippi, P. (2022). 'Data-driven, early-season forecasts of block sugarcane yield for precision agriculture'. In: *Field Crops Research* 276, p. 108360. DOI: 10.1016/j.fcr.2021.108360.

Hardin, I.V., Robert, G., Barnes, E.M., Delhom, C.D., Wanjura, J.D. and Ward, J.K. (2022). 'Internet of things: Cotton harvesting and processing'. In: *Computers and Electronics in Agriculture* 202, p. 107294. DOI: 10.1016/j.compag.2022.107294.

Harianto, J., Haan, S. and Butterworth, N. (2023). 'dataharvester: Download and Process Geospatial Data'. In: *R package version 0.1.2*. URL: sydney-informatics-hub.github.io/dataharvester/.

Hayman, P.T., O'Leary, G. and Meinke, H. (2019). 'Australian Agronomy in the Anthropocene: the challenges of climate'. In: *Australian Agriculture in 2020*. Ed. by Pratley, J. and Kirkegaard, J. Agronomy Australia. URL: https://api.semanticscholar.org/CorpusID:219720081.

He, L. and Mostovoy, G. (2019). 'Cotton yield estimate using Sentinel-2 data and an ecosystem model over the southern US'. In: *Remote Sensing* 11.17, p. 2000. DOI: 10.3390/rs11172000.

Herbert, A. (2012). 'Australian wheat production compares well to global competitors - an international benchmarking comparison'. In: *GRDC 2018 Grains Research Update – National*. Grains Research and Development Corporation (GRDC). URL: https://grdc.com.au/resources-and-publications/grdc-update-papers/tab-content/grdc-update-papers/2018/03/australian-wheat-production-compares-well-to-global-competitors-an-international-benchmarking-comparison.

Hernandez, C.M., Correndo, A., Kyveryga, P., Prestholt, A. and Ciampitti, I.A. (2023). 'On-farm soybean seed protein and oil prediction using satellite data'. In: *Computers and Electronics in Agriculture* 212, p. 108096. DOI: 10.1016/j.compag.2023.108096.

Himesh, S., Prakasa Rao, E.V.S., Gouda, K.C., Ramesh, K.V., Rakesh, V., Mohapatra, G.N., Rao, K.B., Sahoo, S.K. and Ajilesh, P. (2018). 'Digital revolution and Big Data: a new revolution in agriculture'. In: *CABI Reviews* 2018, pp. 1–7. DOI: 10.1079/PAVSNNR201813021.

Hochman, Z., Gobbett, D.L. and Horan, H. (2017). 'Climate trends account for stalled wheat yields in Australia since 1990'. In: *Global change biology* 23.5, pp. 2071–2081. DOI: 10.1111/gcb.13604.

Hochman, Z. and Horan, H. (2018). 'Causes of wheat yield gaps and opportunities to advance the water-limited yield frontier in Australia'. In: *Field Crops Research* 228, pp. 20–30. DOI: 10.1016/j.fcr.2018.08.023.

Hogan, B. (2024). 'Are rising input costs the biggest threat to farm profitability'. In: *GRDC 2024 Grains Research Update – South*. Grains Research and Development Corporation (GRDC). URL: https://grdc.com.au/resources-and-publications/grdc-update-papers/tab-content/grdc-update-papers/2024/02/are-rising-input-costs-the-biggest-threat-to-farm-profitability.

Holford, I.C.R., Doyle, A.D. and Leckie, C.C. (1992). 'Nitrogen response characteristics of wheat protein in relation to yield responses and their interactions with phosphorus'. In: *Australian Journal of Agricultural Research* 43.5, pp. 969–986. DOI: 10.1071/AR9920969.

Hu, M. and Huang, Y. (2020). 'atakrig: An R package for multivariate area-to-area and area-to-point kriging predictions'. In: *Computers and Geosciences* 139.104471. DOI: 10.1016/j.cageo.2020.104471.

Huete, A., Didan, K., Miura, T., Rodriguez, E.P., Gao, X. and Ferreira, L.G. (2002). 'Overview of the radiometric and biophysical performance of the MODIS vegetation indices'. In: *Remote sensing of environment* 83.1-2, pp. 195–213. DOI: 10.1016/S0034-4257(02)00096-2.

Ihuoma, S.O. and Madramootoo, C.A. (2017). 'Recent advances in crop water stress detection'. In: *Computers and Electronics in Agriculture* 141, pp. 267–275. ISSN: 0168-1699. DOI: -10.1016/j.compag.2017.07.026. URL: https://www.sciencedirect.com/science/article/pii/S0168169916310766.

Inc., Planet Labs (2019). *PLANET IMAGERY PRODUCT SPECIFICATIONS*.

Johnson, R.M., Downer, R.G., Bradow, J.M., Bauer, P.J. and Sadler, J.E. (2002). 'Variability in cotton fiber yield, fiber quality, and soil properties in a southeastern coastal plain'. In: *Agronomy Journal* 94.6, pp. 1305–1316. DOI: 10.2134/agronj2002.1305.

Jones, E.J., Bishop, T.F.A., Malone, B.P., Hulme, P.J., Whelan, B.M. and Filippi, P. (2022). 'Identifying causes of crop yield variability with interpretive machine learning'. In: *Computers and Electronics in Agriculture* 192, p. 106632. ISSN: 0168-1699. DOI: 10.1016/j.compag.2021.106632. URL: https://www.sciencedirect.com/science/article/pii/S0168169921006499.

Ju, C.H., Tian, Y.C., Yao, X., Cao, W.X., Zhu, Y. and Hannaway, D. (2010). 'Estimating Leaf Chlorophyll Content Using Red Edge Parameters'. In: *Pedosphere* 20.5, pp. 633–644. ISSN: 1002-0160. DOI: -10.1016/S1002-0160(10)60053-7. URL: https://www.sciencedirect.com/science/article/pii/S1002016010600537.

Kelly, R.M., Strong, W.M., Jensen, T.A. and Butler, D. (2004). 'Application of probability analysis to assess nitrogen supply to grain crops in northern Australia'. In: *Precision Agriculture* 5, pp. 95–110. DOI: 10.1023/B:PRAG.0000022356.01537.67.

Kelly, R.M., Strong, W.M., Jensen, T.A., Butler, D., Town, B. and Adams, M. (2003). 'Recurrence of yield and protein variation in the northern grains region'. In: *Solutions for a better environment: proceedings of the 11th Australian Agronomy Conference*. Australian Society of Agronomy.

Keogh, M. (2012). 'Factoring risk into farm decision making'. In: *GRDC 2012 Grains Research Update – North*. Grains Research and Development Corporation (GRDC). URL: https://grdc.com.au/resources-and-publications/grdc-update-papers/tab-content/grdc-update-papers/2012/11/factoring-risk-into-farm-decision-making.

Kerry, R., Goovaerts, P., Rawlins, B.G. and Marchant, B.P. (2012). 'Disaggregation of legacy soil data using area to point kriging for mapping soil organic carbon at the regional scale'. In: *Geoderma* 170, pp. 347–358. ISSN: 0016-7061. DOI: 10.1016/j.geoderma.2011.10.007. URL: https://www.sciencedirect.com/science/article/pii/S0016706111002928.

Kumhálová, J., Kumhála, F., Kroulík, M. and Matějková, S. (2011). 'The impact of topography on soil properties and yield and the effects of weather conditions'. In: *Precision Agriculture* 12, pp. 813–830. DOI: 10.1007/s11119-011-9221-x.

Larson, J.A., Roberts, R.K., English, B.C., Cochran, R.L. and Wilson, B.S. (2005). 'A computer decision aid for the cotton yield monitor investment decision'. In: *Computers and electronics in agriculture* 48.3, pp. 216–234. DOI: 10.1016/j.compag.2005.04.001.

Leo, S., De Antoni Migliorati, M. and Grace, P.R (2021). 'Predicting within-field cotton yields using publicly available datasets and machine learning'. In: *Agronomy Journal* 113.2, pp. 1150–1163. DOI: 10.1002/agj2.20543.

Leroux, C. and Tisseyre, B. (2019). 'How to measure and report within-field variability: a review of common indicators and their sensitivity'. In: *Precision Agriculture* 20.3, pp. 562–590.

Liakos, K.G., Busato, P., Moshou, D., Pearson, S. and Bochtis, D. (2018). 'Machine learning in agriculture: A review'. In: *Sensors* 18.8, p. 2674. DOI: 10.3390/s18082674.

Lin, L.I. (1989). 'A concordance correlation coefficient to evaluate reproducibility'. In: *Biometrics* 45.1, pp. 255–268. ISSN: 0006341X, 15410420. URL: http://www.jstor.org/stable/2532051 (visited on 12/12/2024).

Lin, Y., Zhu, Z., Guo, W., Sun, Y., Yang, X. and Kovalskyy, V. (2020). 'Continuous Monitoring of Cotton Stem Water Potential using Sentinel-2 Imagery'. In: *Remote Sensing* 12.7. ISSN: 2072-4292. DOI: 10.3390/rs12071176. URL: https://www.mdpi.com/2072-4292/12/7/1176.

Liu, H.Q. and Huete, A. (1995). 'A feedback based modification of the NDVI to minimize canopy background and atmospheric noise'. In: *Transactions on Geoscience and Remote Sensing* 33.2, pp. 457–465. DOI: 10.1109/TGRS.1995.8746027.

Liu, Y., Just, A. and Mayer, M. (2020). 'SHAPforxgboost: SHAP plots for 'XGBoost''. In: *R package version 0.1.0*. URL: https://cran.r-project.org/package=SHAPforxgboost.

Long, D.S. and Engel, R.E. (2011). 'Computing wheat nitrogen requirements from grain yield and protein maps'. In: *GIS Applications in Agriculture, Volume II: Nutrient Management for Energy Efficiency*. Ed. by Clay, David E. and Shanahan, John F. CRC Press, pp. 321–336. DOI: 10.1201/b10600-24.

Long, D.S., McCallum, J.D. and Scharf, P.A. (2013a). 'Optical-mechanical system for on-combine segregation of wheat by grain protein concentration'. In: *Agronomy Journal* 105.6, pp. 1529–1535. DOI: 10.2134/agronj2013.0206.

Long, R.L., Bange, M.P., Delhom, C.D., Church, J.S. and Constable, G.A. (2013b). 'An assessment of alternative cotton fibre quality attributes and their relationship with yarn strength'. In: *Crop and Pasture Science* 64.8, pp. 750–762. DOI: 10.1071/CP12382.

Long, R.L., Gordon, S. and McCarthy, C. (2020). 'Precision Management for Improved Cotton Quality'. In: *CRDC Final Reports*. Cotton Research and Development Corporation (CRDC). URL: https://www.insidecotton.com/precision-management-improved-cotton-quality.

Longchamps, L., Tisseyre, B., Taylor, J.A., Sagoo, L., Momin, A., Fountas, S., Manfrini, L., Ampatzidis, Y., Schueller, J.K. and Khosla, R. (2022). 'Yield sensing technologies for perennial and annual horticultural crops: a review'. In: *Precision Agriculture* 23.6, pp. 2407–2448. DOI: 10.1007/s11119-022-09906-2.

Magney, T.S., Eitel, J.U.H., Huggins, D.R. and Vierling, L.A. (2016). 'Proximal NDVI derived phenology improves in-season predictions of wheat quantity and quality'. In: *Agricultural and Forest Meteorology* 217, pp. 46–60. DOI: 10.1016/j.agrformet.2015.11.009.

Manning, B., Schulze, K. and McNee, T. (2008). 'Grain development'. In: *Wheat Growth and Development*. Ed. by White, J. and Edwards, J. NSW Department of Primary Industries, pp. 71–89.

Matthews, P., Hertel, K. and Jenkins, L. (2024). *Winter crop variety sowing guide 2024*. NSW Department of Primary Industries.

McBratney, A.B., Whelan, B.M. and Shatar, T.M. (1997). 'Variability and uncertainty in spatial, temporal and spatiotemporal crop-yield and related data'. In: *Ciba Foundation Symposium 210 – Precision Agriculture: Spatial and Temporal Variability of Environmental Quality*. Wiley Online Library, pp. 141–160. DOI: 10.1002/9780470515419.ch9.

McVeigh, M. (2017). *The impact of colour discounts to the Australian cotton industry. Project 1517*. Tech. rep. Accessed: 31 May 2024. Nuffield Australia.

Meloni, R., Cordero, E., Capo, L., Reyneri, A., Sacco, D. and Blandino, M. (2024). 'Optimizing nitrogen rates for winter wheat using in-season crop N status indicators'. In: *Field Crops Research* 318, p. 109545. DOI: 10.1016/j.fcr.2024.109545.

Menegat, S., Ledo, A. and Tirado, R. (2022). 'Greenhouse gas emissions from global production and use of nitrogen synthetic fertilisers in agriculture'. In: *Scientific Reports* 12.1, pp. 1–13.

Meredith, W.R. (1986). 'Fiber Quality Variation among USA Cotton Growing Regions'. In: *Proceedings: Beltwide Cotton Production Research Conferences* (Memphis, TN), pp. 105–106.

— (1994). 'Where does fiber quality come from?' In: *Proceedings of the 1994 Beltwide Cotton Conferences* (Memphis, TN), pp. 155–157.

Meteorology, Bureau of (2024a). *Monthly Rainfall Bindi Bindi*. URL: http://www.bom.gov.au/jsp/ncc/cdio/weatherData/av?p_nccObsCode=139&p_display_type=dataFile&p_startYear=&p_c=&p_stn_num=008202 (visited on 16/12/2024).

— (2024b). *Monthly Rainfall Irwin House*. URL: http://www.bom.gov.au/jsp/ncc/cdio/weatherData/av?p_nccObsCode=139&p_display_type=dataFile&p_startYear=&p_c=&p_stn_num=008276 (visited on 16/12/2024).

— (2024c). *Summary statistics Eneabba aero*. URL: http://www.bom.gov.au/climate/averages/tables/cw_008225.shtml (visited on 15/12/2024).

— (2024d). *Summary statistics Moree aero*. URL: http://www.bom.gov.au/climate/averages/tables/cw_053115.shtml (visited on 15/12/2024).

Miao, M.H. and Gordon, S. (2020). 'The effect of round module storage time and ambient conditions on cotton quality'. In: *Journal of Cotton Science* 24.4, pp. 211–228.

Miao, Y., Mulla, D.J. and Robert, P.C. (2006). 'Spatial variability of soil properties, corn quality and yield in two Illinois, USA fields: implications for precision corn management'. In: *Precision Agriculture* 7, pp. 5–20. DOI: 10.1007/s11119-005-6786-2.

Minty, B., Franklin, R., Milligan, P., Richardson, M. and Wilford, J. (2009). 'The Radiometric Map of Australia'. In: *Exploration Geophysics* 40, pp. 325–333. DOI: 10.1071/EG09025.

Monjardino, M., Hochman, Z. and Horan, H. (2019). 'Yield potential determines Australian wheat growers' capacity to close yield gaps while mitigating economic risk'. In: *Agronomy for Sustainable Development* 39.49. DOI: 10.1007/s13593-019-0595-x.

Morari, F., Zanella, V., Sartori, L., Visioli, G., Berzaghi, P. and Mosca, G. (2018). 'Optimising durum wheat cultivation in North Italy: understanding the effects of site-specific fertilization on yield and protein content'. In: *Precision Agriculture* 19, pp. 257–277. DOI: 10.1007/s11119-017-9515-8.

Muleke, A., Harrison, M.T., Yanotti, M. and Battaglia, M. (2022). 'Yield gains of irrigated crops in Australia have stalled: the dire need for adaptation to increasingly volatile weather

and market conditions'. In: *Current Research in Environmental Sustainability* 4, p. 100192. DOI: 10.1016/j.crsust.2022.100192.

Muschalik, M. (2024). *What Are Shapley Interactions, and Why Should You Care?* URL: http://maxmuschalik.com/blog/2024/shapley-interactions/#:~:text=Instead%20of%20providing%20a%20single,a%20value%20of%20an%20interaction. (visited on 20/12/2024).

Muschalik, M., Baniecki, H., Fumagalli, F., Kolpaczki, P., Hammer, B. and Hüllermeier, E. (2024). 'shapiq: Shapley interactions for machine learning'. In: *arXiv preprint arXiv:2410.01649*. DOI: 10.48550/arXiv.2410.01649.

National Farmers Federation (2019). *2030 Roadmap*. URL: https://nff.org.au/policies/roadmap/ (visited on 19/12/2024).

Nehbandani, A., Filippi, P., Alizadeh-Dehkordi, P., Dadrasi, A. and Soltani, A. (2023). 'Use of interpretive machine learning and a crop model to investigate the impact of environment and management on soybean yield gap'. In: *Crop and Pasture Science* 75.1. DOI: 10.1071/CP23032.

Nguyen, L.H., Zhu, J., Lin, Z., Du, H., Yang, Z., Guo, W. and Jin, F. (2019). 'Spatial-temporal multi-task learning for within-field cotton yield prediction'. In: *Advances in Knowledge Discovery and Data Mining: 23rd Pacific-Asia Conference, PAKDD 2019, Macau, China, April 14-17, 2019, Proceedings, Part I 23*. Springer, pp. 343–354. DOI: 10.1007/978-3-030-16148-4_27.

Nuttall, J.G., O'Leary, G.J., Panozzo, J.F., Walker, C.K., Barlow, K.M. and Fitzgerald, G.J. (2017). 'Models of grain quality in wheat – A review'. In: *Field Crops Research* 202. Modeling crops from genotype to phenotype in a changing climate, pp. 136–145. ISSN: 0378-4290. DOI: 10.1016/j.fcr.2015.12.011.

Nutter, F.W., Tylka, G.L., Guan, J., Moreira, A.J.D., Marett, C.C., Rosburg, T.R., Basart, J.P. and Chong, C.S. (2002). 'Use of remote sensing to detect soybean cyst nematode-induced plant stress'. In: *Journal of Nematology* 34.3, pp. 222–231.

Oliveira, R.P. de (2009). 'Contributions Towards Decision Support for Site-Specific Crop Management: a study of aspects influencing the development of knowledge-intensive differential management decisions'. Doctoral Dissertation. The University of Sydney.

Oliveira, R.P. de, Whelan, B., McBratney, A.B. and Taylor, J.A. (2007). 'Yield variability as an index supporting management decisions: YIELDex'. In: *Precision agriculture '07*. Ed. by Stafford, J. Wageningen Academic, pp. 281–288.

Oliver, Y., Wong, M., Robertson, M. and Wittwer, K. (2006). 'PAWC determines spatial variability in grain yield and nitrogen requirement by interacting with rainfall on northern WA sandplain'. In: *Ground-breaking Stuff: Proceedings of the 13th Australian Agronomy Conference*. Australian Society of Agronomy, pp. 10–14.

Oosterhuis, D.M., Coomer, T., Raper, T.B. and Espinoza, L. (2015). *Use of Remote Sensing in Cotton to Accurately Predict the Onset of Nutrient Stress for Foliar Alleviation for Optimizing Yield and Quality*. Tech. rep. Fluid Fertilizer Foundation. Fayetteville, NC, USA.

Orton, T.G., Mallawaarachchi, T., Pringle, M.J., Menzies, N.W., Dalal, R.C., Kopittke, P.M., Searle, R., Hochman, Z. and Dang, Y.P. (2018). 'Quantifying the economic impact of soil constraints on Australian agriculture: A case-study of wheat'. In: *Land Degradation & Development* 29.11, pp. 3866–3875. DOI: 0.1002/ldr.3130.

Orton, T.G., Pringle, M.J., Page, K.L., Dalal, R.C. and Bishop, T.F.A. (2014). 'Spatial prediction of soil organic carbon stock using a linear model of coregionalisation'. In: *Geoderma* 230, pp. 119–130. DOI: 10.1016/j.geoderma.2014.04.016.

Pallegedara Dewage, S.N.S, Minasny, B. and Malone, B. (2020). 'Disaggregating a regional-extent digital soil map using Bayesian area-to-point regression kriging for farm-scale soil carbon assessment'. In: *SOIL* 6.2, pp. 359–369. DOI: 10.5194/soil-6-359-2020. URL: https://soil.copernicus.org/articles/6/359/2020/.

Pan, W.L., Kidwell, K.K., McCracken, V.A., Bolton, R.P. and Allen, M. (2020). 'Economically optimal wheat yield, protein and nitrogen use component responses to varying N supply and genotype'. In: *Frontiers in plant science* 10, p. 1790.

Pannell, D.J., Marshall, G.R., Barr, N., Curtis, A., Vanclay, F. and Wilkinson, R. (2006). 'Understanding and promoting adoption of conservation practices by rural landholders'. In: *Australian Journal of Experimental Agriculture* 46.11, pp. 1407–1424. DOI: 10.1071/EA05037.

Pebesma, E.J. (2004). 'Multivariable geostatistics in S: the gstat package'. In: *Computers & geosciences* 30.7, pp. 683–691. DOI: 10.1016/j.cageo.2004.03.012.

Ping, J.L., Green, C.J., Zartman, R.E., Bronson, K.F. and Morris, T.F. (2007). 'Spatial Variability of Soil Properties, Cotton Yield, and Quality in a Production Field'. In: *Communications in Soil Science and Plant Analysis* 39.1–2, pp. 1–16. DOI: 10.1080/00103620701758840.

Poole, S., Bishop, T.F.A., Al-Shammari, D. and Filippi, P. (2024). 'Mapping within-field drivers of crop production in space and time using interpretive machine learning'. In: *Adaptive agronomy for a resilient future: Proceedings of the 21st Australian Agronomy Conference*. Australian Society of Agronomy.

Pracilio, G., Adams, M.L. and Smettem, K.R.J. (2003). 'Use of airborne gamma radiometric data for soil property and crop biomass assessment, northern dryland agricultural region, Western Australia'. In: *Precision Agriculture '03: Proceedings of the 4th European Conference*. Ed. by Stafford, J.V. and Werner, A. Wageningen Academic Publishers.

Pringle, M.J., McBratney, A.B., Whelan, B.M. and Taylor, J.A. (2003). 'A preliminary approach to assessing the opportunity for site-specific crop management in a field, using yield monitor data'. In: *Agricultural Systems* 76.1, pp. 273–292. DOI: 10.1016/S0308-521X(02)00005-7.

R Core Team (2021). 'R: A language and environment for statistical computing'. In: *R version 4.3.1*. URL: https://www.r-project.

Research, Cotton and Corporation, Development (2022). *2022 Grower Survey*.

Research, Cotton, Corporation, Development and CottonInfo (2023). *Australian Cotton Production Manual*. CRDC.

Research, Grains and Corporation, Development (2016). 'Wheat Section 15 – Marketing'. In: *GRDC Grownotes*.

— (2024). *Growing regions*. URL: https://grdc.com.au/about/our-industry/growing-regions (visited on 17/12/2024).

Roades, J.P., Beck, A.D. and Searcy, S.W. (2000). 'Cotton yield mapping: Texas experiences in 1999.' In: *Proceedings of the Beltwide Cotton Conferences*. National Cotton Council of America, pp. 404–407.

Robertson, M.J., Lyle, G. and Bowden, J.W. (2008). 'Within-field variability of wheat yield and economic implications for spatially variable nutrient management'. In: *Field Crops Research* 105.3, pp. 211–220. DOI: 10.1016/j.fcr.2007.10.005.

Rondeaux, G., Steven, M. and Baret, F. (1996). 'Optimization of soil-adjusted vegetation indices'. In: *Remote Sensing of Environment* 55.2, pp. 95–107. ISSN: 0034-4257. DOI: 10.1016/0034-4257(95)00186-7. URL: https://www.sciencedirect.com/science/article/pii/0034425795001867.

Rosenberg, C.W. von, Abbate, A., Drake, J. and Mayes, D.M. (2000). 'A Rugged Near-Infrared Spectrometer for the Real-time Measurement of Grains During Harvest-An NIR solid-state spectrometer can be used to measure commodity grains in the field in real'. In: *Spectroscopy* 15.6, pp. 34–39.

Roßbach, P. (2018). 'Neural networks vs. random forests–does it always have to be deep learning'. In: *Germany: Frankfurt School of Finance and Management*.

Roth, G., Harris, G., Gillies, M., Montgomery, J. and Wigginton, D. (2013). 'Water-use efficiency and productivity trends in Australian irrigated cotton: a review'. In: *Crop and Pasture Science* 64.12, pp. 1033–1048. DOI: 10.1071/CP13315.

Rouse, J.W., Haas, R.H., Schell, J.A. and Deering, D.W. (1974). 'Monitoring vegetation systems in the Great Plains with ERTS'. In: *NASA. Goddard Space Flight Center 3d ERTS-1 Symposium, Vol. 1*. Vol. 351. 1, p. 309.

Rudin, C. (2019). 'Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead'. In: *Nature Machine Intelligence* 1, pp. 206–2015. DOI: 10.1038/s42256-019-0048-x.

Ruß, G. and Brenning, A. (2010). 'Data mining in precision agriculture: management of spatial information'. In: *Computational Intelligence for Knowledge-Based Systems Design: 13th International Conference on Information Processing and Management of Uncertainty, IPMU 2010, Dortmund, Germany, June 28-July 2, 2010. Proceedings 13*. Springer, pp. 350–359. DOI: 10.1007/978-3-642-14049-5_36.

Schade, P., Struve, C. and Gralfs, R. (2023). 'Sensor Technology for Continuous Measurement of Constituents During Small Grain Harvest'. In: *ATZheavy duty worldwide* 16.2, pp. 10–15. DOI: 10.1007/s41321-023-1031-3.

Schielack, V.P., Thomasson, A.J., Sui, R and Ge, Y. (2016). 'Harvester-based sensing system for cotton fiber quality mapping.' In: *Journal of Cotton Science* 20.4, pp. 386–393. DOI: 10.56454/LNLB8251.

Scott, E. (2022). 'Protein mapping - getting more bang for your fertiliser buck'. In: *GRDC 2022 Grains Research Update – South*. Grains Research and Development Corporation (GRDC). URL: https://grdc.com.au/resources-and-publications/grdc-update-papers/tab-content/grdc-update-papers/2022/02/protein-mapping-getting-more-bang-for-your-fertiliser-buck.

Scott, E. and Clancy, M. (2012). 'Leveraging protein for profitability'. In: *GRDC 2024 Grains Research Update – South*. Grains Research and Development Corporation (GRDC). URL: https://grdc.com.au/resources-and-publications/grdc-update-papers/tab-content/grdc-update-papers/2024/07/leveraging-protein-for-profitability.

Segarra, J., Buchaillot, M.L., Araus, J.L. and Kefauver, S.C. (2020). 'Remote sensing for precision agriculture: Sentinel-2 improved features and applications'. In: *Agronomy* 10.5, p. 641. DOI: 10.3390/agronomy10050641.

Shackley, B., Power, S., Paynter, B., Troup, G., Seymour, M. and Dhammu, H. (2024). *2024 Western Australian Crop Sowing Guide*. Department of Primary Industries and Regional Development.

Simmonds, D.H (1989). *Wheat and wheat quality in Australia*. CSIRO Publishing. DOI: 10.1071/9780643101456.

Simmonds, N.W. (1995). 'The relation between yield and protein in cereal grain'. In: *Journal of the Science of Food and Agriculture* 67.3, pp. 309–315. DOI: 10.1002/jsfa.2740670306.

Skerritt, J.H., Adams, M.L., Cook, S.E. and Naglis, G. (2002). 'Within-field variation in wheat quality: implications for precision agricultural management'. In: *Australian Journal of Agricultural Research* 53.11, pp. 1229–1242. DOI: 10.1071/AR01204.

Sluijs, M.H.J. van der (2022). 'Effect of nitrogen application level on cotton fibre quality'. In: *Journal of Cotton Research* 5.9. DOI: 10.1186/s42397-022-00116-9.

Sluijs, M.H.J. van der, Delhom, C.D., Wanjura, J.D. and Holt, G.A. (2019). 'A preliminary investigation into the feasibility of gin blending'. In: *Journal of Cotton Science* 23, pp. 97–108. DOI: 10.56454/AMJD216.

Sluijs, M.H.J. van der, Long, R.L. and Bange, M.P. (2015). 'Comparing cotton fiber quality from conventional and round module harvesting methods'. In: *Textile Research Journal* 85.9, pp. 987–997. DOI: 10.1177/0040517514540770.

Smith, H.D., Dubeux, J.C.B., Zare, A. and Wilson, C.H. (2023). 'Assessing transferability of remote sensing pasture estimates using multiple machine learning algorithms and evaluation structures'. In: *Remote Sensing* 15.11, p. 2940. DOI: 10.3390/rs15112940.

Song, X., Yang, G., Yang, C., Wang, J. and Cui, B. (2017). 'Spatial variability analysis of within-field winter wheat nitrogen and grain quality using canopy fluorescence sensor measurements'. In: *Remote Sensing* 9.3, p. 237. DOI: 10.3390/rs9030237.

Song, Y., Zheng, X., Cheng, X., Xu, Q., Liu, X., Tian, Y., Zhu, Y., Cao, W. and Cao, Q. (2023). 'Improving the Prediction of Grain Protein Content in Winter Wheat at the County Level with Multisource Data: A Case Study in Jiangsu Province of China'. In: *Agronomy* 13.10, p. 2577. DOI: 10.3390/agronomy13102577.

Statistics, Australian Bureau of (2024). *Australian Agriculture: Broadacre Crops*. URL: https://www.abs.gov.au/statistics/industry/agriculture/australian-agriculture-broadacre-crops/latest-release (visited on 16/12/2024).

Stevens, A., Miralles, I. and Wesemael, B. van (2012). 'Soil organic carbon predictions by airborne imaging spectroscopy: Comparing cross-validation and validation'. In: *Soil Science Society of America Journal* 76.6, pp. 2174–2183. DOI: 10.2136/sssaj2012.0054.

Stewart, C.M., McBratney, A.B. and Skerritt, J.H. (2002). 'Site-specific durum wheat quality and its relationship to soil properties in a single field in northern New South Wales'. In: *Precision agriculture* 3, pp. 155–168. DOI: 10.1023/A:1013871519665.

Stoy, P.C., Khan, A.M., Wipf, A., Silverman, N. and Powell, S.L. (2022). 'The spatial variability of NDVI within a wheat field: Information content and implications for yield

and grain protein monitoring'. In: *PloS one* 17.3, e0265243. DOI: 10.1371/journal.pone.0265243.

Strobl, C., Boulesteix, A.L., Kneib, T., Augustin, T. and Zeileis, A. (2008). 'Conditional variable importance for random forests'. In: *BMC Bioinformatics* 9.307, pp. 1–11. DOI: 10.1186/1471-2105-9-307.

Szigeti, N., Sulyán, P., Labus, B., Földi, M., Hunyadi, E., Mikó, P., Milibák, F. and Drexler, D. (2024). 'Limitations and solutions for developing a grain yield and protein content forecasting model based on vegetation indices in organic wheat production–on-farm experimentation'. In: *Biological Agriculture & Horticulture* 40.3, pp. 1–15. DOI: 10.1080/01448765.2024.2364304.

Tan, C., Zhou, X., Zhang, P., Wang, Z., Wang, D., Guo, W. and Yun, F. (2020). 'Predicting grain protein content of field-grown winter wheat with satellite images and partial least square algorithm'. In: *PLoS One* 15.3, e0228500. DOI: 10.1371/journal.pone.0228500.

Taylor, J.A., McBratney, A.B. and Whelan, B.M. (2007a). 'Establishing management classes for broadacre agricultural production'. In: *Agronomy Journal* 99.5, pp. 1366–1376. DOI: 10.2134/agronj2007.0070.

Taylor, J.A. and Whelan, B.M. (2007b). 'On-the-go protein monitoring: A review'. In: *4th International Symposium on Precision Agriculture (SIAP07)*. Federal University of Viçosa.

Taylor, M.J., Smettem, K., Pracilio, G. and Verboom, W. (2002). 'Relationships between soil properties and high-resolution radiometrics, central eastern Wheatbelt, Western Australia'. In: *Exploration geophysics* 33.2, pp. 95–102. DOI: 10.1071/EG02095.

Terman, G.L., Ramig, R.E., Dreier, A.F. and Olson, R.A. (1969). 'Yield-Protein Relationships in Wheat Grain, as Affected by Nitrogen and Water 1'. In: *Agronomy Journal* 61.5, pp. 755–759. DOI: 10.2134/agronj1969.00021962006100050031x.

Thompson, Co.N., Guo, W., Sharma, B. and Ritchie, G.L. (2019). 'Using Normalized Difference Red Edge Index to Assess Maturity in Cotton'. In: *Crop Science* 59.5, pp. 2167–2177. DOI: 10.2135/cropsci2019.04.0227.

Tilse, M.J., Bishop, T.F.A. and Filippi, P. (2024a). 'Decoding the protein puzzle: understanding variability in grain protein content'. In: *2024 Australasian Precision Ag Symposium* (Dubbo, New South Wales).

— (2024b). 'Predicting and mapping grain protein content to better understand variability – utilising John Deere's new Harvestlab™ 3000 grain sensing system'. In: *GRDC 2024 Grains Research Update – North*. Grains Research and Development Corporation (GRDC). URL: https://grdc.com.au/resources-and-publications/grdc-update-papers/tab-content/grdc-update-papers/2024/02/predicting-and-mapping-grain-protein-content-to-better-understand-variability-utilising-john-deeres-new-harvestlab-3000-grain-sensing-system.

— (2024c). 'Predicting, mapping, and understanding the drivers of grain protein content variability – utilizing harvester mounted grain protein sensors'. In: *16th International Conference on Precision Agriculture* (Manhattan, Kansas USA), pp. 1–13.

— (2024d). 'Understanding the drivers of within-field grain protein content variability with spatial data layers'. In: *Adaptive agronomy for a resilient future: Proceedings of the 21st Australian Agronomy Conference* (Albany, Australia).

Tilse, M.J., Filippi, P., Whelan, B.M. and Bishop, T.F.A. (2024e). 'Downscaling crop production data to fine scale estimates with geostatistics and remote sensing: a case study in mapping cotton fibre quality'. In: *Precision Agriculture*, pp. 1–37. DOI: 10.1007/s11119-024-10161-w.

Tilse, M.J., Stockmann, U. and Filippi, P. (2023). 'Proximal soil sensing in the field'. In: *Encyclopedia of Soils in the Environment*. Ed. by Goss, Michael J. and Oliver, Margaret. 2nd ed. Vol. 4. Elsevier, pp. 579–590. DOI: 10.1016/B978-0-12-822974-3.00188-9.

Tilse, M.J., Whelan, B.M., Bishop, T.F.A. and Filippi, P. (2022). 'Soil costraint diagnosis and mapping'. In: *Soil Constraints to Crop Production*. Ed. by Dang, Y., Menzies, N. and Dalal, R. Cambridge Scholars Publishing: Newcastle upon Tyne, pp. 481–507.

Tucker, C.J. (1979). 'Red and photographic infrared linear combinations for monitoring vegetation'. In: *Remote Sensing of Environment* 8.2, pp. 127–150. ISSN: 0034-4257. DOI:

10.1016/0034-4257(79)90013-0. URL: https://www.sciencedirect.com/science/article/pii/0034425779900130.

Ulfa, F., Orton, T.G., Dang, Y.P. and Menzies, N.W. (2022). 'Developing and testing remote-sensing indices to represent within-field variation of wheat yields: Assessment of the variation explained by simple models'. In: *Agronomy* 12.2, p. 384. DOI: 10.3390/agronomy12020384.

United Nations (2015). *The Paris Agreement*. URL: https://unfccc.int/process-and-meetings/the-paris-agreement (visited on 17/12/2024).

Velandia, M.M., Rejesus, R.M., Segarra, E. and Bronson, K. (2008). 'Economics of management Zone delineation in cotton Precision agriculture'. In: *Journal of Cotton Science* 12, pp. 210–227. DOI: 10.22004/ag.econ.35387.

Wang, J. (2024). '3D maps take soils management to new levels'. In: *GRDC Groundcover Supplement: Grain Automate and Digital Agronomy May – June 2024*. Grains Research and Development Corporation (GRDC). URL: https://groundcover.grdc.com.au/innovation/precision-%20agriculture-and-machinery/3d-maps-take-soils-management-to-new-levels.

Wang, K., Huggins, D.R. and Tao, H. (2019). 'Rapid mapping of winter wheat yield, protein, and nitrogen uptake using remote and proximal sensing'. In: *International Journal of Applied Earth Observation and Geoinformation* 82, p. 101921. DOI: 10.1016/j.jag.2019.101921.

Wang, Q., Shi, W., Atkinson, P.M. and Zhao, Y. (2015). 'Downscaling MODIS images with area-to-point regression kriging'. In: *Remote Sensing of Environment* 166, pp. 191–204. ISSN: 0034-4257. DOI: 10.1016/j.rse.2015.06.003. URL: https://www.sciencedirect.com/science/article/pii/S003442571530033X.

Wang, R., Alex Thomasson, J., Cox, M.S., Sui, R. and Marley Hollingsworth, E.G. (2017). 'Cotton fiber-quality prediction based on spatial variability in soils'. In: *Journal of Cotton Science* 21, pp. 220–228. DOI: 10.5555/20193363291.

Wang, X., Zhang, L., Evers, J.B., Mao, L., Wei, S., Pan, X., Zhao, X., Werf, W. van der and Li, Z. (2014). 'Predicting the effects of environment and management on cotton fibre

growth and quality: a functional–structural plant modelling approach'. In: *AoB Plants* 6, plu040. DOI: 10.1093/aobpla/plu040.

Wang, Y., Peng, Y., Lin, J., Wang, L., Jia, Z. and Zhang, R. (2023). 'Optimal nitrogen management to achieve high wheat grain yield, grain protein content, and water productivity: A meta-analysis'. In: *Agricultural Water Management* 290, p. 108587. DOI: 10.1016/j.agwat.2023.108587.

Watcharaanantapong, P., Roberts, R.K., Lambert, D.M., Larson, J.A., Velandia, M., English, B.C., Rejesus, R.M. and Wang, C. (2014). 'Timing of precision agriculture technology adoption in US cotton production'. In: *Precision agriculture* 15, pp. 427–446.

Whelan, B.M. and McBratney, A.B. (2000). 'The "null hypothesis" of precision agriculture management'. In: *Precision Agriculture* 2, pp. 265–279. DOI: 10.1023/A:1011838806489.

Whelan, B.M. and Taylor, J.A. (2013). *Precision agriculture for grain production systems*. CSIRO publishing.

Whelan, B.M., Taylor, J.A. and Hassall, J.A. (2009). 'Site-specific variation in wheat grain protein concentration and wheat grain yield measured on an Australian farm using harvester-mounted on-the-go sensors'. In: *Crop and Pasture Science* 60.9, pp. 808–817. DOI: 10.1071/CP08343.

Whetton, R.L., Waine, T.W. and Mouazen, A.M. (2018). 'Evaluating management zone maps for variable rate fungicide application and selective harvest'. In: *Computers and Electronics in Agriculture* 153, pp. 202–212. DOI: 10.1016/j.compag.2018.08.004.

Wiggins, M.S., Leib, B.G., Mueller, T.C. and Main, C.L. (2014). 'Cotton growth, yield, and fiber quality response to irrigation and water deficit in soil of varying depth to a sand layer'. In: *Journal of Cotton Science* 18, pp. 145–152. DOI: 10.5555/20143341899.

Wilford, J. and Roberts, D. (2021). *Sentinel-2 Barest Earth imagery for soil and lithological mapping*. DOI: 10.11636/146125.

Wright, M.N. and Ziegler, A. (2017). 'ranger: A fast implementation of random forests for high dimensional data in C++ and R'. In: *Journal of Statistical Software* 77.1, pp. 1–17. DOI: 10.48550/arXiv.1508.04409.

Wright, M.N., Ziegler, A. and König, I.R. (2016). 'Do little interactions get lost in dark random forests?' In: *BMC Bioinformatics* 17.145, pp. 1–10. DOI: 10.1186/s12859-016-0995-8.

Xu, W., Chen, P., Zhan, Y., Chen, S., Zhang, L. and Lan, Y. (2021). 'Cotton yield estimation model based on machine learning using time series UAV remote sensing data'. In: *International Journal of Applied Earth Observation and Geoinformation* 104, p. 102511. ISSN: 1569–8432. DOI: 10.1016/j.jag.2021.102511. URL: https://www.sciencedirect.com/science/article/pii/S030324342100218X.

Xu, X., Teng, C., Zhao, Y., Du, Y., Zhao, C., Yang, G., Jin, X., Song, X., Gu, X., Casa, R., Chen, L. and Li, Z. (2020). 'Prediction of wheat grain protein by coupling multisource remote sensing imagery and ECMWF data'. In: *Remote Sensing* 12.8, p. 1349. DOI: 10.3390/rs12081349.

Zarco-Tejada, P.J., Ustin, S.L. and Whiting, M.L. (2005). 'Temporal and Spatial Relationships between Within-Field Yield Variability in Cotton and High-Spatial Hyperspectral Remote Sensing Imagery'. In: *Agronomy Journal* 97.3, pp. 641–653. DOI: 10.2134/agronj2003.0257.

Zhao, C., Liu, L., Wang, J., Huang, W., Song, X. and Li, C. (2005). 'Predicting grain protein content of winter wheat using remote sensing data based on nitrogen status and water stress'. In: *International Journal of Applied Earth Observation and Geoinformation* 7.1, pp. 1–9. DOI: 10.1016/j.jag.2004.10.002.

Zhao, H., Song, X., Yang, G., Li, Z., Zhang, D. and Feng, H. (2019). 'Monitoring of nitrogen and grain protein content in winter wheat based on Sentinel-2A data'. In: *Remote Sensing* 11.14, p. 1724. DOI: 10.3390/rs11141724.

Zhou, X., Kono, Y., Win, A., Matsui, T. and Tanaka, T.S.T. (2021). 'Predicting within-field variability in grain yield and protein content of winter wheat using UAV-based multispectral imagery and machine learning approaches'. In: *Plant Production Science* 24.2, pp. 137–151. DOI: 10.1080/1343943X.2020.1819165.

Žilić, S., Barać, M., Pešić, M., Dodig, D. and Ignjatović-Micić, D. (2011). 'Characterization of proteins from grain of different bread and durum wheat genotypes'. In: *International journal of molecular sciences* 12.9, pp. 5878–5894. DOI: 10.3390/ijms12095878.

APPENDIX A

Soil Constraint Diagnosis and Mapping

Mikaela J. Tilse, Brett Whelan, Thomas Bishop, Patrick Filippi

Sydney Institute of Agriculture, School of Life and Environmental Sciences, Faculty of Science, The University of Sydney, Sydney, New South Wales, Australia.

Abstract

Soil constraints significantly impact crop growth and yields. The ability to accurately measure and map these constraints has the potential to greatly improve on-farm productivity. Soil constraints have traditionally been assessed using farmer observations or conventional soil sampling and analysis techniques. Technological advancements have enabled the assessment of constraints using more formal sampling strategies, and proximal and remote sensing technologies. While regional and national soil databases and products contain valuable information about the soil environment, many lack information specific to soil constraints and have limited use on-farm due to their broad extent. The collection of field-scale data, via point-based or spatial techniques, provides an abundance of raw soil data. However, this information is often not presented in a useful or meaningful way for applied decision-making on-farm. Digital soil mapping (DSM) presents an opportunity to integrate this plethora of grower-sourced and publicly available soil data into useful decision-making tools. As techniques used to measure and map soil constraints continue to improve, the translation of this information into useful decision-making tools remains imperative.

Key words: Digital soil mapping, soil sampling, soil constraints

Introduction

Soil constraints have the potential to significantly impact crop growth and yields. Data on the soil landscape is becoming increasingly more accessible across a range of spatial and temporal extents. However, information specific to soil constraints and their impact on crops is either lacking, or not portrayed in a way that is useful for guiding informed management decisions in the field. Growers must be able to measure and quantify soil constraints on farm to ameliorate or manage them accordingly.

This chapter will explore the range of soil sampling strategies and techniques available to understand soil physical, chemical, and biological constraints in both the surface and/or subsoil. The University of Sydney farm "L'lara" will be used as a case study throughout to demonstrate the range of soil sampling and mapping techniques available (Figures A1 – A7). A range of sampling strategies, from arbitrary and ad hoc approaches, through to more formal systematic and stratified sampling methods will be discussed. Sample extraction techniques, including profile pits, augers, and cores, and analysis techniques such as proximal and remote sensing will be explored. Digital soil mapping (DSM) has been developing in recent decades and is an opportunity to integrate grower-sourced and publicly available data. The application of DSM to better understand the nature and extent of soil constraints will be discussed, and future directions will be explored.

Identifying a soil constraint to crop production

There are a suite of ways to identify whether a soil constraint to crop production exists in a particular area. These include informal approaches, such as farmer experience and knowledge, as well as more systematic approaches that use yield data or satellite imagery. In the absence of more temporary constraints (pests, diseases, drought, frost), areas of poor and variable crop growth and yields over multiple seasons often indicate the presence of soil constraints (Lobell *et al*. 2007; Dang *et al*. 2011). While farmer observations play a key role in the initial

identification of lower yielding areas, these visual assessments are often subjective, and a more formal assessment of this variability is increasingly being sought.

Many growers have access to yield monitor data and produce maps of yield variability (Bramley *et al*. 2019). This may complement farmer observations and can be used to formally identify lower yielding areas within and between fields. Remote sensing, using satellite or airborne sensors and imagery, also presents an additional source of information to assess crop variability. For example, Normalised Difference Vegetation Index (NDVI) maps can be used to assess variability in crop vigour, canopy health, and may be used as an indirect indicator of soil condition. This may complement existing yield maps and on-the-ground observations. Multiple years of remotely sensed surrogate yield data may also be used to identify areas affected by soil constraints, where consistently low-performing yields over multiple seasons may infer the presence of a soil constraint when the impacts of more temporary constraints are filtered out (Lobell *et al*. 2007; Dang *et al*. 2011).

However, while yield maps may suggest the presence of at least one soil constraint, they cannot identify the nature of the constraint or separate the interactions and relationships between multiple soil attributes. Ground-truthing of visual observations, yield maps, and satellite imagery through the collection of soil data is thus required to identify the exact nature and extent of soil constraints. Attributing crop variability to a single constraint is often difficult due to the interactions between the physical and chemical properties of soil (Adcock *et al*. 2007). Interrelationships between soil properties also make it difficult to identify the relative effect of individual or multiple constraints, and antagonistic interactions between soil properties may result in yield or growth-limiting conditions, even if each individual soil property does not exceed defined threshold values (Adcock *et al*. 2007).

Soil sampling and analysis for ground-truthing constraints

Soil sampling strategies

Soil analysis is important for ground-truthing observations of suspected soil constraints. The accuracy with which a soil sample represents the broader soil landscape depends on the number of sampling units, the sampling technique, and the inherent variability of the soil (Cline 1944). Thus, it is important to know both where and how to sample to identify soil constraints.

Since the early 20th century, the inherent spatial variability of soils has been increasingly recognised and is now a key consideration when identifying where to sample. Over time, sampling approaches have become increasingly sophisticated, supported by emerging academic research, extension publications, and agronomic advice (Lawrence *et al*. 2020). Soil sampling approaches range from ad hoc or arbitrary techniques, through to more formal approaches including systematic and stratified sampling schemes. The ad hoc selection of sampling points "at random" over an area is not considered an effective method of randomisation as it may introduce strong personal bias and fail to effectively capture variability of the soil (Cline 1944). Prior observations and existing data, such as farmer knowledge, yield maps, or satellite imagery, may be used to arbitrarily guide the location of sampling points to capture sufficient variability. Ad hoc or arbitrary sampling approaches are commonplace due to their convenience and ease of use and depending on the experience and knowledge of the surveyor, may be time and cost-effective, particularly for exploratory investigations. However, unintended bias remains a concern.

Systematic sampling strategies select sample sites based on organised grids, patterns or transects and aim to objectively capture as much variability as possible across the study area and reduce personal bias. However, in practice, systematic sampling approaches using arbitrarily located transects or zig-zag patterns, for example, are likely to produce biased results as they fail to capture a truly randomised sample (Lawrence *et al*. 2020). The advent of Global Positioning System (GPS) and Geographic Information System (GIS) technologies

have facilitated the use of grid-based sampling strategies to geo-locate soil information and produce digital soil maps. Grid soil sampling aims to capture soil variability most-objectively by sampling based on a regular pattern of cells (Figure A1). However, soils are highly variable through space and some parts of a field may be more inherently variable than others. Grid sampling ignores this variability and has a uniform sampling density irrespective of the underlying variation present. While grid-based sampling can capture sufficient detail required for DSM, the process is time consuming and expensive (Dang *et al*. 2010). The variability of soil parameters, time, and financial constraints are important considerations when deciding upon grid sizes (Flowers *et al*. 2005). While smaller grid cells are able to capture more variability (Flowers *et al*. 2005), grid cell sizes may range anywhere from 10's to 100's of metres depending on logistical constraints. For inexperienced or non-technically trained individuals, grid sampling is generally preferred over sampling at random as it can limit potential bias (Cline 1944).

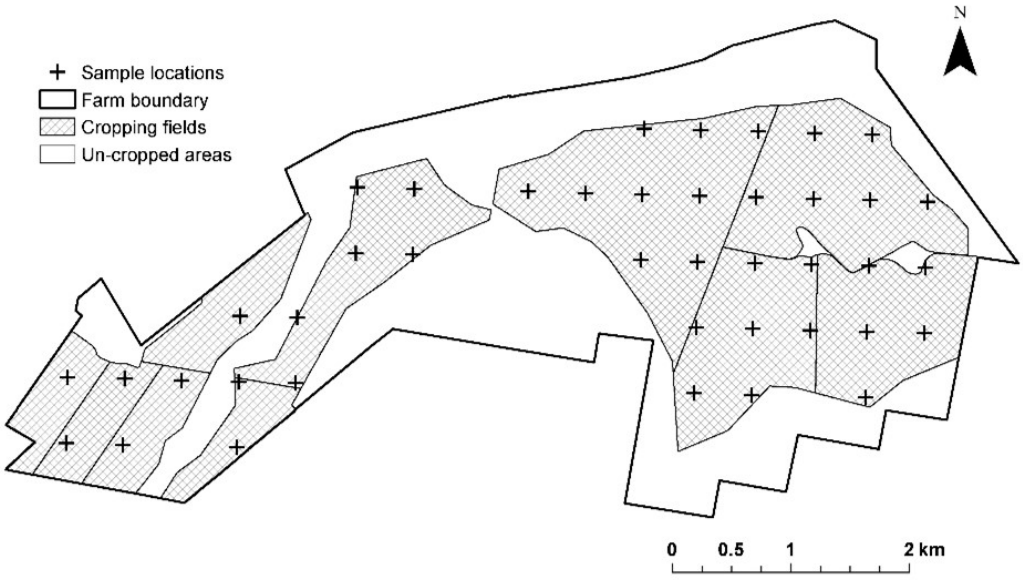

FIGURE A1. The distribution of sampling locations based on a regular-spaced grid pattern across "L'lara".

Statistical or stratified sampling separates the study area into sub-areas (i.e. strata). Strata may be derived by separating the study region into equal-area strata or by subdivision into more homogenous regions through a range of statistical methodologies including Latin hypercube

sampling (Minasny *et al*. 2006) or k-means clustering (Filippi *et al*. 2019; Taylor *et al*. 2003) (Figure A2). Depending on the complexity of the stratification technique, heterogeneity of the study area, and number of desired strata, information used to delineate strata may include farmer observations (Fleming *et al*. 2000), yield data (Flowers *et al*. 2005), satellite imagery (Filippi *et al*. 2019), or a combination of this information alongside other ancillary data. However, there is a trade-off between increased precision and increasing cost and complexity that must be considered when choosing the optimum number of strata.

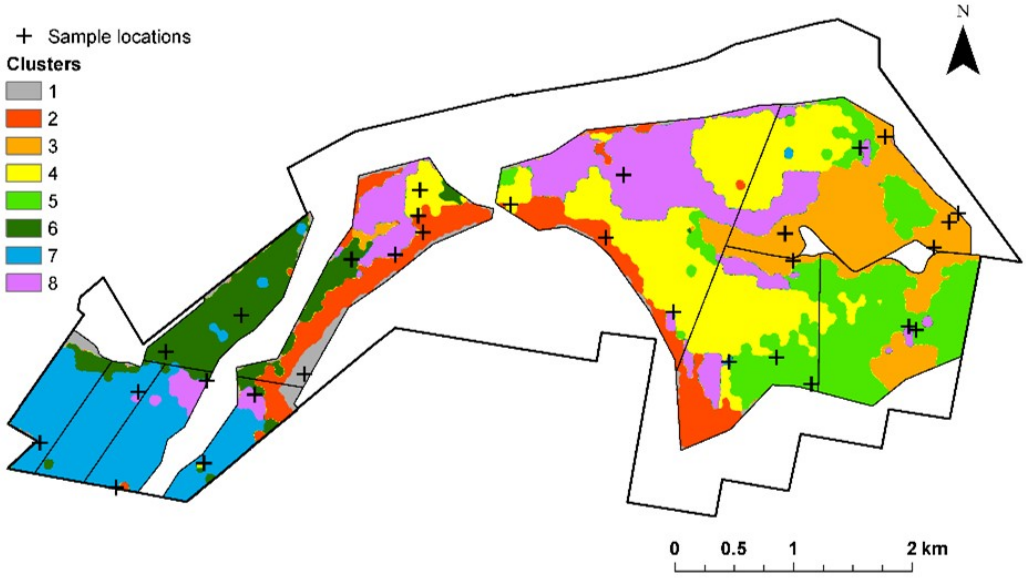

FIGURE A2. The distribution of eight clusters derived from a K-means clustering analysis, and the locations soil sampling sites randomly selected within each cluster across "L'lara" (Filippi *et al*. 2019).

Soil extraction techniques

Traditional soil sampling approaches treat fields or land management areas as homogenous regions (Dawson *et al*. 2018). As such, bulk or composite samples are created to represent the area of interest. Composite soil sampling is still widely used by farmers and agronomists in the field and is an approach best suited to more homogenous areas where there is existing knowledge about the soil condition (Dawson *et al*. 2018). Composite sampling operates under the assumption that if each individual sample was analysed and averaged, it would

give the same result as for the larger composite sample (Tan 2005). The number of soil samples recommended to capture a sufficiently representative composite sample varies for different soil properties measured, although logistical constraints often determine the number of samples that can feasibly be collected (Lawrence *et al*. 2020). A general rule of thumb, the more homogenous an area is, the larger the area the composite sample can represent (Dawson *et al*. 2018). Composite sampling is limited by its inability to capture the inherent spatial variability of soils, however, and sampling error will likely increase as sampling areas become more heterogenous (Tan 2005).

Point-based sampling strategies, including digging profile pits, hand augering, and soil core samples, can be used to capture spatially referenced soil information. Soil profile pits can provide detailed information about a soil's physical, chemical, and biological properties. When using ad hoc or arbitrary sampling approaches, soil profile pits may be dug based on observations of suspected constraints, historical yield maps, vegetation indices, or pre-existing soil surveys (Taylor *et al*. 2003). However, the time and labour required to dig and analyse soil profile pits means that they are rarely used on-farm (Earl *et al*. 2003), they are impractical for spatial analysis and cannot be used to capture the inherent spatial variability of soils (Taylor *et al*. 2003).

Soil augers and cores offer cheaper, faster alternatives to soil pits that are highly suited to systematic or statistical sampling strategies. Hand auger sampling is typically used for the assessment of surface-soil or shallow subsurface properties, while mechanised soil cores enable the extraction of a high density of soil samples to depths typically up to 2 m. The mechanised extraction of soil cores is less time and labour intensive compared to hand augering (Earl *et al*. 2003), however, the vehicle access to fields required for mechanical sampling may be limited at certain crop growth stages and due to agronomic concerns, such as compaction. By geo-referencing the locations of samples, soil maps can be created through DSM procedures.

Analysis of soil samples

Conventional soil analysis techniques continue to be a key contributor to our understanding of the soils physical, chemical, and biological condition. Soil samples may be analysed in situ or be transferred to a laboratory for analysis. Conventional analysis methods include the use of indicator dyes via a colorimetric method such as the Raupach *et al*. (1959) method for the in-field assessment of soil pH, the use of electrodes to measure pH and EC in soil solutions (Tan 2005) or a cylindrical core sampler to measure bulk density (Hao *et al*. 2008). There is an extensive body of literature available detailing local and national guidelines and recommendations for conventional soil analysis techniques (e.g. Tan 2005).

Proximal soil sensing may also be used to compliment, or replace, conventional techniques by providing quantitative results via portable and/or handheld sensors. Viscarra Rossel *et al*. (2011) offer a comprehensive review of proximal soil sensing. Here we explore their application for the analysis of soil constraints. Proximal soil sensors can facilitate the (near instantaneous) collection of vast amounts of data using simpler, less time and labour intensive techniques than conventional approaches (Viscarra Rossel *et al*. 2011). For example, spectral sensors, such as Portable X-Ray Fluorescence (pXRF), use elemental data as a proxy for the analysis of various soil attributes and have been used to assess soil pH (Sharma *et al*. 2014), CEC (Sharma *et al*. 2015), and salinity (Aldabaa *et al*. 2015). Neutron probes, nuclear density gauges, and low-activity nuclear density gauges have also been used to measure soil moisture and bulk density (Dep *et al*. 2022). Similarly, visible-near-infrared (NIR) spectroscopy has been used to measure a suite of soil properties including salinity and soil organic carbon (Ahmadi *et al*. 2021). Today, proximal soil sensing techniques are increasingly shifting from the research and development phase into practical application.

However, while handheld or portable proximal soil sensors have been used with some success, the output from individual sensors is not as accurate as output derived from a combination of multiple sensors and ancillary data (Ji *et al*. 2019; Aldabaa *et al*. 2015). Proximal soil sensors do not directly measure soil constraint or attribute values. Rather, sensor outputs serve as a proxy that must be correlated via mathematical processing and modelling. As correlations are

often made with calibration samples unique to a particular location or soil type, the use of proximal sensors is restricted by the soil condition(s) it has been calibrated for (Stenberg *et al*. 2010). In addition, interrelationships between soil properties requires multivariate calibrations to account for the non-specificity of proximal soil sensor output for individual constraints or attributes stenberg2010diffuse.

Mapping soil constraints

Changes in soil sampling and analysis techniques, from point-based measurements in situ to the use of spatial sensors, has rapidly expanded the volume and complexity of information collected about the soil landscape. In its raw form, this data gives minimal insight into the extent and distribution of soil constraints, and must be converted into a mapping format before it can provide information about the spatial distribution of soil properties and the nature and extent of constraints across the landscape in a readily interpretable, easy to use format.

Conventional soil maps typically rely on qualitative inferences of soil information, including field observations and aerial imagery, to classify soil properties and features. While these maps served their purpose throughout much of the 20th century, conventional soil surveys and maps have been criticised for their subjectivity and non-replicability (Arrouays *et al*. 2020). The rise of precision agriculture and increasingly intensive broadacre cropping systems is generating more and more quantitative information and spatial data. DSM has emerged in recent decades following rapid growth and advancements in computational and information technologies (Arrouays *et al*. 2020). Early DSM simply represented point observations without spatial interpolation (e.g. Webster *et al*. 1979), or involved the digitisation of conventional soil maps (Tomlinson 1978). Modern DSM techniques use mathematical or statistical techniques to derive relationships between soil observations and ancillary environmental information for the quantitative prediction of soil properties across a landscape (Arrouays *et al*. 2020). McBratney *et al*. (2003) proposed the scorpan framework as a methodology to derive these quantitative soil-landscape relationships and produce digital soil maps. While this is not a standardised procedure, this framework has been widely adopted (Arrouays *et al*. 2020). In

the last decade, DSM has shifted from a research and development phase into operational use, emphasising the need for DSM outputs to present information that is relevant to end-users (Carré *et al*. 2007).

Publicly available digital soil maps

Access to relevant and accurate soil information is important for making on-the-ground management decisions. Globally, two DSM projects have emerged in response to a demand for up-to-date, spatially referenced soil information with global coverage – SoilGrids (Hengl *et al*. 2014; Hengl *et al*. 2017) and GlobalSoilMap (Arrouays *et al*. 2014a). SoilGrids (Hengl *et al*. 2014; Hengl *et al*. 2017) provides complete, globally consistent gridded maps of soil properties. Released in 2014, the first SoilGrids product mapped soil properties on a 1 km grid using a top-down approach, which uses a global model to predict properties using data derived from legacy soil surveys. Although the coarse spatial resolution fails to capture sufficient variability of soil properties at the local level, the first SoilGrids product demonstrated the ability to map soil properties at a global scale using legacy soil data (Arrouays *et al*. 2017). In using a top-down approach, SoilGrids can be used to highlight regions where data may be sparse or missing, thus potentially encouraging ongoing collaborative efforts to fill these gaps through bottom-up approaches (Arrouays *et al*. 2017). An update to SoilGrids was released in 2017, which mapped soil properties across a 250 m grid (Hengl *et al*. 2017).

The GlobalSoilMap project aims to provide a collaborative, standardised framework for the provision of publicly accessible, continuous global soil information at a high spatial resolution (90 m) (Arrouays *et al*. 2014b). In using a bottom-up approach (i.e. from country to globe), the framework specifies that soil information is included for a minimum of twelve key soil attributes (depth to rock, plant exploitable (effective) depth, organic carbon, pH, clay, silt, sand, coarse fragments, CEC, bulk density, available water capacity, and electrical conductivity) at six standard depth intervals (i.e. 0 – 0.05, 0.05 – 0.15, 0.15 – 0.30, 0.30 – 0.60, 0.60 – 1.00, 1.00 – 2.00 m), mapped at a 3 arcsecond (roughly 90 m) spatial resolution with estimates of uncertainty (5^{th} and 95^{th} percentile prediction limits) (Arrouays *et al*. 2014b). Numerous contributions have been made to the GlobalSoilMap project already at the continental and

national level, including for Sub-Saharan Africa (Hengl *et al*. 2015), Australia (Grundy *et al*. 2015), France (Mulder *et al*. 2016), and Nigeria (Akpa *et al*. 2014).

While digital soil maps from SoilGrids and the GlobalSoilMap project may be useful for regional soil assessments and other land-use applications, the large spatial resolutions of these maps (i.e. 90 m, 250 m, 1 km) is generally too coarse for use at the within-field-scale. As such, these maps are rarely used by growers on farm. Also, understanding the depth at which a constraint is reached in the soil profile is important for guiding management decisions and agronomic practices. However, the vertical resolution of the standard depth intervals of the GlobalSoilMap increases in coarseness with depth (i.e. 0.30 – 0.60, 0.60 – 1.00 m), which makes it difficult to accurately determine the depth to a constraint.

While several agronomically important soil constraints are included in the GlobalSoilMap project, including EC for salinity, pH for acidity/alkalinity (Figure A3), and bulk density for compaction, other attributes such as ESP for sodicity are not. Other soil attributes such as Available Water Capacity (AWC) are estimated via the use of pedotransfer functions (PTFs). Region-specific PTFs are being developed and applied to GlobalSoilMap projects across the globe. However, error derived from the accuracy of PTF input variables and PTF coefficients may result in uncertainty in AWC estimates (Román Dobarco *et al*. 2019). The incorporation of uncertainty estimates into GlobalSoilMap products and predictions is important for the application of this information in ecological and agricultural modelling and decision making (Román Dobarco *et al*. 2019).

By presenting soil attribute information at six depth intervals, end-users are provided with an abundance of soil information. However, these soil maps do not indicate the presence or absence of a constraint. Rather, they only represent the spatial distribution of attribute values at specified depth intervals. This abundance of data may overwhelm end-users, and additional processing may be required to understand the nature and extent of a constraint with respect to the application of this data.

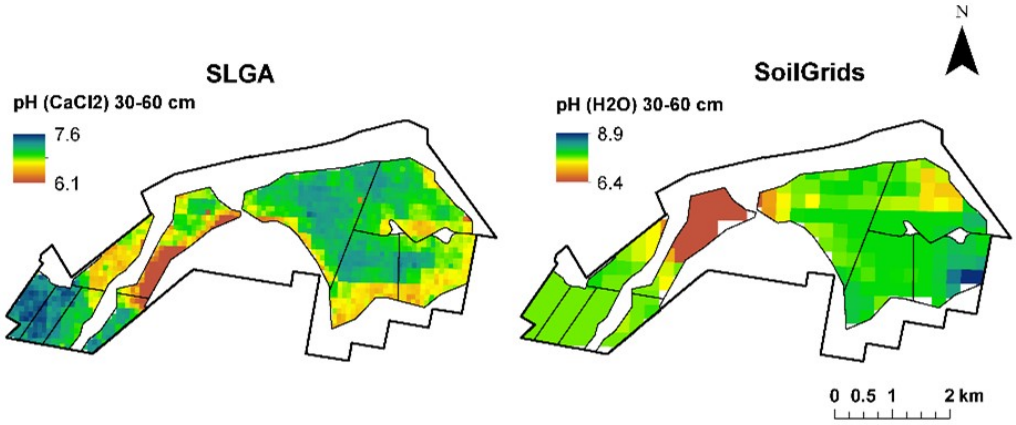

FIGURE A3. Digital soil map soil pH products for the 30–60 cm clay layer at their provided resolution from the Soil and Landscape Grids of Australia (SLGA) (90 m) and SoilGrids (250 m) across "L'lara".

Proximal spatial sensing

Logistical constraints limit the volume of data that can be obtained via conventional soil sampling strategies and handheld/portable proximal soil sensors. Proximal spatial sensing enables the measurements of soil information at high densities, thus filling in these gaps and facilitating soil sampling for high-resolution soil mapping (De Gruijter *et al*. 2010).

Output from proximal spatial sensors (PSS), such as electromagnetic induction (EMI) instruments and ground-based gamma ray spectrometers (GRS), often have good relationships with soil constraints, but are indirect measures. Both EMI and GRS surveys are commonly used PSS for research purposes and commercial application. By sampling above the soil surface, EMI and GRS surveys measure apparent soil electrical conductivity (EC¬a) and broad elemental radioactive emissions, respectively, at depth within the soil profile. Measurements are taken at fixed intervals across the study area by sampling on-the-go in swathes, forming a regular transect pattern which is then used to create a map of soil variability (De Gruijter *et al*. 2010) (Figure A4). Early uses of EMI surveys were for the assessment of soil salinity (Rhoades *et al*. 1981). Today, EMI surveys have been used for the assessment of numerous soil constraints, including sodicity (Paz *et al*. 2020), pH (Dunn *et al*. 2007), and bulk density

(Al-Gaadi 2012), and have been used at the field (Triantafilis *et al*. 2002) and landscape scale (Paz *et al*. 2020).

Other PSS have also been developed for the assessment of soil constraints, including mechanical and electrochemical sensors. Tine-based sensors have been used to measure soil strength to estimate compaction (Lapen *et al*. 2002), and the Veris® pH meter is a commercially available proximal spatial sensor that has been used with considerable accuracy to produce maps for the variable rate application of lime to manage soil acidity (Trengove *et al*. 2019).

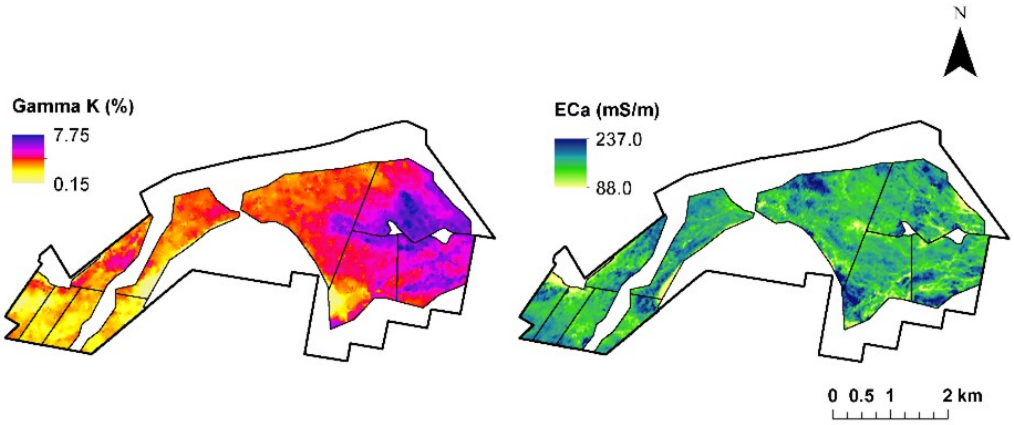

FIGURE A4. Gamma radiometrics K and Soil ECa (0 – 3.0 m) maps derived from GRS and a Dual EM-21S, respectively, across "L'lara".

The management of soil constraints involves complex decision-making processes which require multiple layers of (complex) information (Viscarra Rossel *et al*. 2010). However, a single sensor can only measure certain soil properties (Ji *et al*. 2019). A selection of complementary sensors is often required to provide a concurrent quantitative measurement of multiple soil constraints (Grunwald *et al*. 2015). Output from PSS can be used with DSM techniques to create soil maps to inform management decisions. However, few growers actually collect proximally sensed soil data. This limited adoption may be attributed to a poor understanding and perception of the use of these technologies (Bramley *et al*. 2019), or a lack of integration with other data to enable meaningful interpretation.

Overall, the application of PSS for the assessment of soil constraints is limited in much the same ways as for handheld or portable proximal soil sensors. Readings from PSS are generally

an indirect measure of the soil constraint being examined, interrelationships between soil properties limit the number of constraints that can be measured with a single sensor (Ji *et al*. 2019), and calibration models and equations tend to be site-specific and time-dependent (Doolittle *et al*. 2014).

Publicly available data layers

Publicly available data layers, including remotely sensed satellite imagery, terrain attributes, and other environmental covariates, can also be used in conjunction with DSM techniques to create maps of soil constraints. Remote sensing, via satellite or aerial sensing platforms, enables the collection of space-time dense data in a cost and time-effective manner (Grunwald *et al*. 2015). Over time, more and more data layers obtained via remote sensing platforms have been made publicly available as technologies and spatial, temporal, and spectral resolutions have improved. Databases of freely available, medium resolution (i.e. $\geq$ 10 m pixel) data layers are available, including Landsat, Shuttle Radar Topography Mission (SRTM), Sentinel (from the European Space Agency), and Advanced Very High Resolution Radiometer (AVHRR). Other databases include higher spatial resolution data layers (i.e. $\leq$ 5 m) available at a cost, such as RapidEye and GeoEye-1. The recent proliferation of open-access databases has resulted in an increase in the number of studies utilising satellite remote sensing data (Khanal *et al*. 2020). Both Mulla (2013) and Khanal *et al*. (2020) provide comprehensive summaries of the applications of satellite remote sensing platforms to (precision) agriculture for the last 20 – 25 years. Here, we will focus on the application of publicly available data layers to measure and map soil constraints.

Many important soil constraints have been mapped using DSM approaches that utilise environmental covariates obtained from publicly accessible databases. Environmental covariates include a range of terrain attributes such as elevation and slope, and reflectance data in a range of forms such as direct single-band imagery (e.g. red-band) or multi-band processed imagery (e.g. NDVI) (Figure A5). By building relationships between available soil data and publicly available environmental covariates, soil constraints including pH, sodicity, and others have been mapped at the national (Roudier *et al*. 2020), catchment (Filippi *et al*. 2020) and farm

scale (Dang *et al*. 2011) respectively. In addition, publicly available data layers may also be used to directly measure and map the distribution and impact of soil constraints such as soil salinity (Abbas *et al*. 2013). There is increasing recognition of the value in combining publicly available and farm-sourced data to create "bespoke digital soil maps" unique to fields, farms, and regions (Filippi *et al*. 2020). However, not all soil constraints can be accurately identified from publicly available data layers. Satellite or aerial remote sensing-based studies of soil compaction have yielded poor results, for example (Wells *et al*. 1998), and some remotely sensed indices, such as NDVI, are sensitive to the effects of canopy shading, soil colour, and brightness. Post-hoc adjustments of remote sensing products may be required to more accurately assess some soil or vegetation attributes (Xue *et al*. 2017).

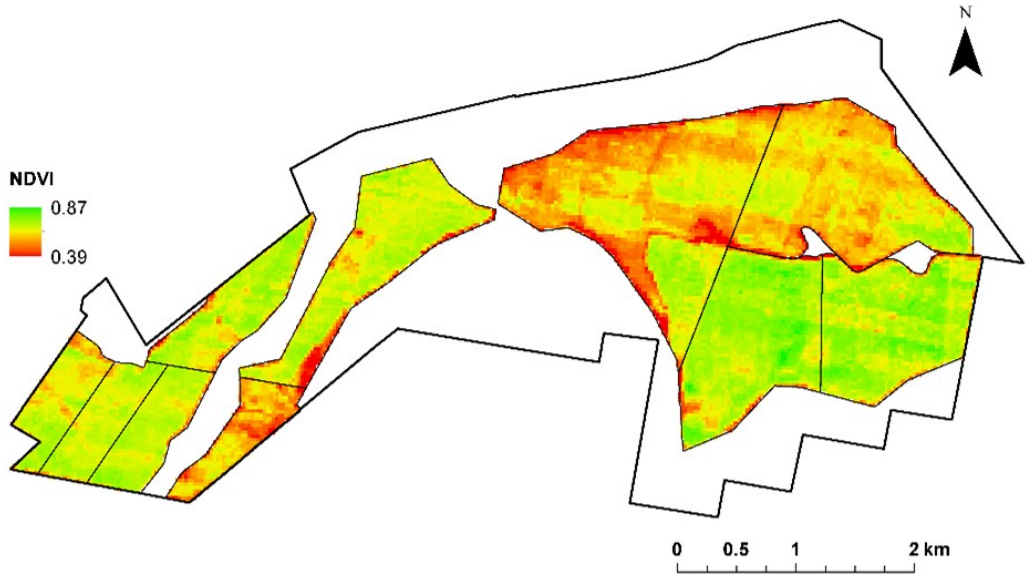

FIGURE A5. Average 95th percentile NDVI imagery for 2000 – 2018 from Landsat 7 across "L'lara".

Data layers obtained via remote sensing platforms have been critiqued as having spatial and temporal resolutions that are too coarse for application at the field scale for in-season use (Omran *et al*. 2020). While this may limit the application of data from some platforms such as Landsat (with a spatial resolution of 30 m pixels and a 16-day observation cycle), technological advancements have meant that sensors can now measure with sub-meter accuracy with relatively short re-observation times. Satellite imagery obtained from high spatial-resolution

platforms, including Pleiades (with a 0.5 m spatial resolution and daily observation cycle) and IKONOS (~1 m resolution and a revisit time of 3 days), have been used to assess soil salinity for example (Laiskhanov *et al*. 2016; Allbed *et al*. 2014). While significant improvements in the spatial and temporal resolution of satellite remote sensing platforms and publicly available data layers requires higher capacities of data storage and processing requirements (Mulla 2013), the ability to capture a vast amount of high-resolution data has significantly improved soil constraint assessment capabilities.

Techniques used to map soil constraints at the field, farm, and regional scales

Traditional DSM has generally mapped soil attributes at different depth increments in an attempt to capture both the horizontal and vertical variability of soil. Conventional soil sampling approaches divide the soil profile into horizons. These horizons vary in their number and position within the profile which may differ between sampling locations. Other sampling strategies sample at specified depth increments down the profile. These increments often vary between different publications, however, more and more studies are producing maps at depths in line with the standard depth intervals of the GlobalSoilMap project (i.e. 0 – 0.05, 0.05 – 0.15, 0.15 – 0.30, 0.30 – 0.60, 0.60 – 1.00) (Arrouays *et al*. 2014b) (Figure A6). In the past, inconsistencies in sampling depths between new and/or legacy soil surveys meant that it was often difficult to use this soil data to map soil properties down the profile. Equal-area quadratic smoothing splines (EAQSS) were proposed to overcome this challenge and make horizon data more continuous (Bishop, McBratney, and Laslett 1999). Today, EAQSS are used to map soil properties at a range of depth layers, ranging from 1 cm increments (Filippi et al. 2019) through to coarser intervals (Akpa et al. 2014). While mapping soil constraints at different depth intervals offers a wealth of information about the soil landscape, this may also overwhelm end-users and over-complicate decision making. Attempts to condense this information include simplifying map layers to only the top and/or subsoil (Filippi et al. 2018), or by producing fewer maps at coarser depth intervals. However, these options may fail to

capture the true extent of the variability of soil constraints depending on the thickness of each depth layer.

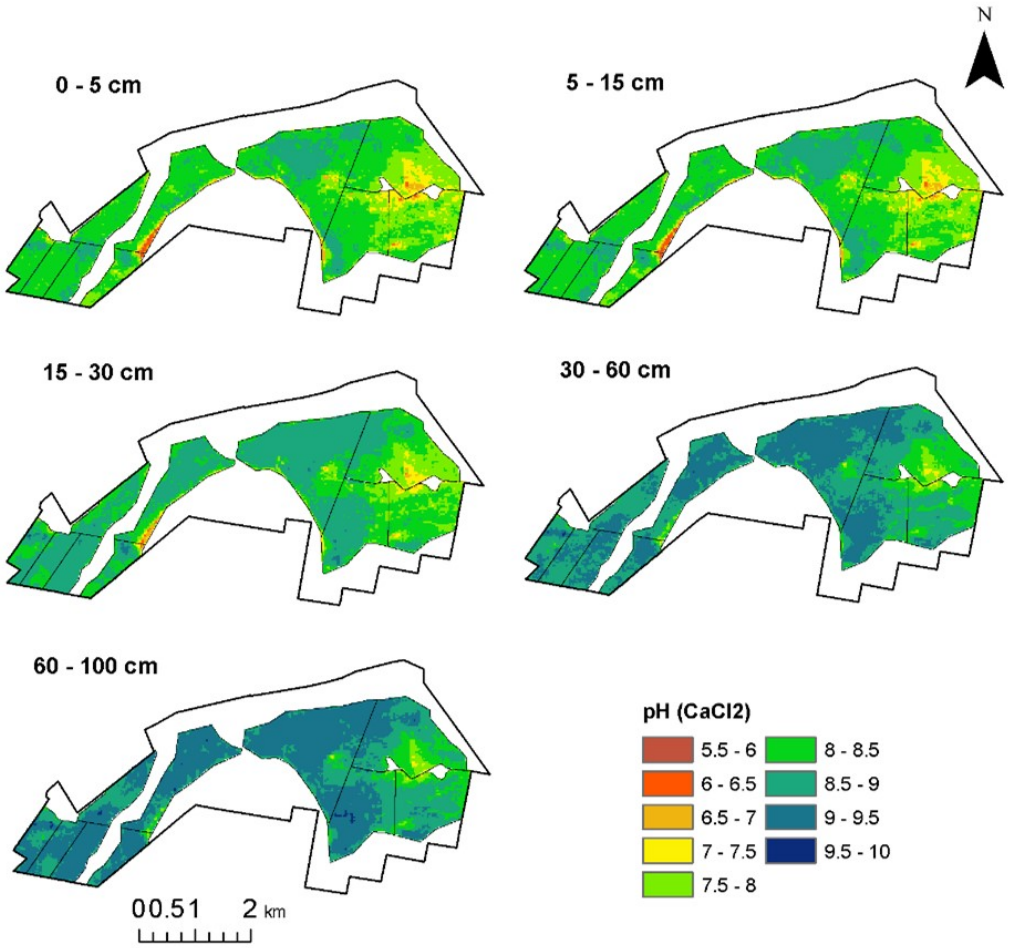

FIGURE A6. Digital soil map products for pH ($CaCl_2$) at GlobalSoilMap standard depth intervals for "L'lara".

Although maps of soil properties may be useful for identifying relationships between crop yields and soil condition, they are unable to identify the presence or absence of a constraint. The threshold values at which a soil property becomes constraining differs between crop types, growth stages, and other environmental factors. There are published thresholds for the effect of soil properties on crops, usually based on reductions in crop growth and yield. Filippi *et al*. (2019) proposed the production of a depth-to constraint map to simplify the plethora of information describing the variability of soils (Figure A7). By mapping the depth-to a specific constraint threshold value at a 1 cm vertical resolution, multiple layers of soil information

can be integrated into a single map. In recognising both the impact of constraints on crop growth and yields, and mapping constraints at a high vertical resolution, these maps may be of significant benefit for informing management decisions and assessing the efficacy of potential amelioration strategies via a simple and informative decision-making tool.

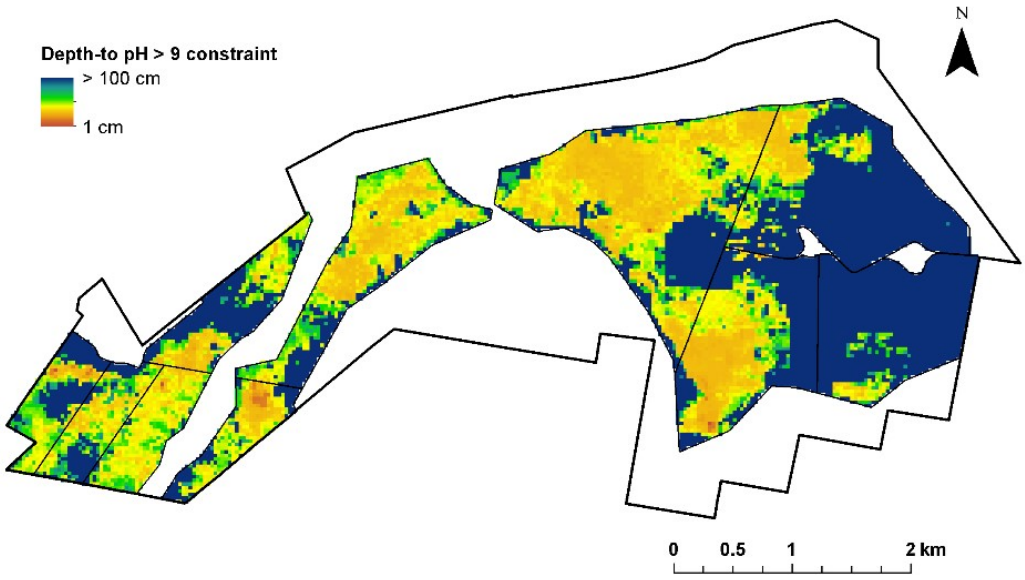

FIGURE A7. Depth-to pH >9 constraint map for "L'lara".

Soils are highly variable, however, and more than one constraint may be present in the soil environment at a particular location. This may make it difficult to attribute crop growth or yield responses to a particular constraint. Also, several maps showing the distribution of different constraints may present an excess of information that is difficult to interpret. Quantifying the impact of constraints may highlight those that most-significantly impact crop growth and yields, or identify constraints requiring additional amelioration or management strategies to better optimise crop production. These impacts have been considered in terms of estimating potential yield gaps (Hajkowicz *et al*. 2005), economic impacts (Orton *et al*. 2018), or the effect of constraints on soil available water capacity (Sadras *et al*. 2003). However, most efforts focussed on spatially mapping the impacts of constraints have focused on potential yield gaps.

In using the boundary-line analysis method, Shatar *et al*. (2004) modelled relationships between crop yields and soil data to identify soil constraints that resulted in the largest yield

gaps. These were identified to be the most-limiting constraints to yield and were presented on a patchwork-styled map at the field-scale. Mapping the spatial distribution of most-limiting constraints enables potential amelioration or management strategies to be directed towards those constraints that result in the greatest quantifiable impact on crop growth and yields for a particular area. However, although these maps may be easy to interpret, these may also over-simplify soil-plant relationships and ignore the antagonistic interactions and interrelationships between soil constraints and other properties (Adcock *et al*. 2007).

Future directions

The measurement and mapping of soil constraints relies on the collection of useful, appropriate, and accurate data. A plethora of data is currently available to growers and farm managers for the on-farm assessment of soil constraints. Moving forward, several key hurdles must be overcome to continue to support informed decision making into the future.

Firstly, the ongoing collection of additional soil data through on-farm soil surveys should be encouraged. There is significant value in using on-farm data in conjunction with publicly available digital soil maps and data layers to create bespoke field, farm, and region-specific maps of soil properties (Filippi *et al*. 2020). Through the standardisation of sampling procedures, this farm-sourced data can be integrated with existing soil databases to increase the information available and improve mapping capabilities at a range of scales. However, the return on investment of soil sampling surveys must be realised. The translation of soil data into useful decision-making tools and frameworks, such as maps of soil constraints, an understanding of yield gaps (Dang *et al*. 2016), and a calculation of the return on investment from soil constraint amelioration (Bennett *et al*. 2021), can highlight the importance of data collection and investment in these strategies. As such, the collection of soil data should not be goal of technological developments for measuring and mapping soil constraints. Rather, this data should be collected with the purposes of guiding informed decision making. Through the effective translation and communication of this information to end-users, the ongoing sustainability and productivity of agricultural systems can be ensured moving into the future.

Secondly, the benefits of collecting additional soil data cannot be appreciated if the quality and accuracy of this information is poor. Direct sampling methods are time-consuming and expensive (Dang *et al*. 2010), while proximal and remote sensing options are indirect measures that rely on proxies to reflect soil information. Improving the accuracy of on-the-go sampling and analysis methods and improving direct measurement techniques may improve the accuracy and ease with which data is collected. Further, the ongoing calibration of proximal and remote sensing technologies to different soil types and regions will improve the accuracy of proxy-based sensing systems. Improvements in uncertainty reporting of DSM products is also important for communicating the reliability of this information decision-making by a range of end-users including farmers and consultants, policy makers, and the broader community (Kidd *et al*. 2020).

Finally, improving the understanding of the values at which soil properties become constraining to crop growth and yields should be considered. Describing the presence or absence of a soil constraint via the use of arbitrarily defined threshold values can be beneficial (e.g. Filippi *et al*. 2019). However, crop growth and yield responses to soil constraints are not defined at a single set value. Rather, the magnitude of impact may vary between different soil types, crops, seasons, and regions. Also, antagonistic interactions between multiple soil properties may promote crop growth and yield-limiting conditions, even if each individual property does not exceed a defined threshold limit. Given that there is abundant yield monitor data relative to soil constraint data, further work should examine using this data source for refining constraint thresholds that are season, site and crop-specific.

References

Abbas, A., Khan, S., Hussain, N., Hanjra, M.A. and Akbar, S. (2013). 'Characterizing soil salinity in irrigated agriculture using a remote sensing approach'. In: *Physics and chemistry of the Earth, Parts A/B/C* 55–57, pp. 43–52. DOI: 10.1016/j.pce.2010.12.004.

Adcock, D., McNeill, A.M., McDonald, G.K. and Armstrong, R.D. (2007). 'Subsoil constraints to crop production on neutral and alkaline soils in south-eastern Australia: a review

of current knowledge and management strategies'. In: *Australian Journal of Experimental Agriculture* 47.11, pp. 1245–1261. DOI: 10.1071/EA06250.

Ahmadi, A., Emami, M., Daccache, A. and He, L. (2021). 'Soil properties prediction for precision agriculture using visible and near-infrared spectroscopy: A systematic review and meta-analysis'. In: *Agronomy* 11.3, p. 433. DOI: 10.3390/agronomy11030433.

Akpa, S.I.C., Odeh, I.O.A., Bishop, T.F.A. and Hartemink, A.E. (2014). 'Digital mapping of soil particle-size fractions for Nigeria'. In: *Soil Science Society of America Journal* 78.6, pp. 1953–1966. DOI: 10.2136/sssaj2014.05.0202.

Aldabaa, A.A.A., Weindorf, D.C., Chakraborty, S., Sharma, A. and Li, B. (2015). 'Combination of proximal and remote sensing methods for rapid soil salinity quantification'. In: *Geoderma* 239, pp. 34–46. DOI: 10.1016/j.geoderma.2014.09.011.

Allbed, A., Kumar, L. and Aldakheel, Y.Y. (2014). 'Assessing soil salinity using soil salinity and vegetation indices derived from IKONOS high-spatial resolution imageries: Applications in a date palm dominated region'. In: *Geoderma* 230, pp. 1–8. DOI: 10.1016/j.geoderma.2014.03.025.

Arrouays, D., Grundy, M.G., Hartemink, A.E., Hempel, J.W., Heuvelink, G.B.M., Hong, S.Y., Lagacherie, P., Lelyk, G., McBratney, A.B., McKenzie, N.J., Mendonca-Santos, M.d.L., Minasny, B., Montanarella, L., Odeh, I.O.A., Sanchez, P.A., Thompson, J.A. and Zhang, G-L. (2014a). 'Chapter Three – GlobalSoilMap: Toward a fine-resolution global grid of soil properties'. In: *Advances in agronomy* 125, pp. 93–134. DOI: 10.1016/B978-0-12-800137-0.00003-0.

Arrouays, D., Leenaars, J.G.B., Richer-de-Forges, A.C., Adhikari, K., Ballabio, C., Greve, M., Grundy, M., Guerrero, E., Hempel, J., Hengl, T., Heuvelink, G., Batjes, N., E., Carvalho., Hartemink, A., Hewitt, A., Hong, S-Y., Krasilnikov, P., Lagacherie, P., Lelyk, G., Libohova, Z. and Rodriguez, D. (2017). 'Soil legacy data rescue via GlobalSoilMap and other international and national initiatives'. In: *GeoResJ* 14, pp. 1–19. DOI: 10.1016/j.grj.2017.06.001.

Arrouays, D., McBratney, A.B., Bouma, J., Libohova, Z., Richer-de-Forges, A.C., Morgan, C.L.S., Roudier, P., Poggio, L. and Mulder, V.L. (2020). 'Impressions of digital soil maps:

The good, the not so good, and making them ever better'. In: *Geoderma Regional* 20, e00255. DOI: 10.1016/j.geodrs.2020.e00255.

Arrouays, D., McKenzie, N., Hempel, J., Forges, A.R. de and McBratney, A.B. (2014b). *GlobalSoilMap: basis of the global spatial soil information system*. London: CRC press. ISBN: 978-1-138-00119-0. DOI: 10.1201/b16500.

Bennett, J.McL., Roberton, S.D., Ghahramani, A. and McKenzie, D.C. (2021). 'Operationalising soil security by making soil data useful: Digital soil mapping, assessment and return-on-investment'. In: *Soil Security* 4, p. 100010. DOI: 10.1016/j.soisec.2021.100010.

Bramley, R.G.V. and Ouzman, J. (2019). 'Farmer attitudes to the use of sensors and automation in fertilizer decision-making: Nitrogen fertilization in the Australian grains sector'. In: *Precision agriculture* 20, pp. 157–175. DOI: 10.1007/s11119-018-9589-y.

Carré, F., McBratney, A.B., Mayr, T. and Montanarella, L. (2007). 'Digital soil assessments: Beyond DSM'. In: *Geoderma* 142.1-2, pp. 69–79. DOI: 10.1016/j.geoderma.2007.08.015.

Cline, M.G. (1944). 'Principles of soil sampling'. In: *Soil Science* 58.4, pp. 275–288.

Dang, Y.P., Dalal, R.C., Buck, S.R., Harms, B.R., Kelly, R., Hochman, Z., Schwenke, G.D., Biggs, A.J.W., Ferguson, N.J., Norrish, S., Routley R. McDonald, M., Hall, C., Singh, D.K., Daniells, I.G., Farquharson, R., Manning, W., Speirs, S., Grewal, H.S., Cornish, P., Bodapati, N. and Orange, D. (2010). 'Diagnosis, extent, impacts, and management of subsoil constraints in the northern grains cropping region of Australia'. In: *Soil Research* 48.2, pp. 105–119. DOI: 10.1071/SR09074.

Dang, Y.P. and Moody, P.W. (2016). 'Quantifying the costs of soil constraints to Australian agriculture: a case study of wheat in north-eastern Australia'. In: *Soil Research* 54.6, pp. 700–707. DOI: 10.1071/SR15007.

Dang, Y.P., Pringle, M.J., Schmidt, M., Dalal, R.C. and Apan, A. (2011). 'Identifying the spatial variability of soil constraints using multi-year remote sensing'. In: *Field Crops Research* 123.3, pp. 248–258. DOI: 10.1016/j.fcr.2011.05.021.

Dawson, A. and Knowles, O. (2018). *To grid or not to grid–a review of soil sampling strategies*. Tech. rep. Fertilizer and Lime Research Centre, Massey University, Palmerston North. URL: http://flrc.massey.ac.nz/publications.html.

De Gruijter, J.J., McBratney, A.B. and Taylor, J.A. (2010). 'Sampling for high-resolution soil mapping'. In: *Proximal soil sensing*. Ed. by Viscarra Rossel, R., McBratney, A.B and Minasny, B., pp. 3–14. DOI: 10.1007/978-90-481-8859-8_1.

Dep, L., Troxler, R.E., Sawyer, J. and Mwimba, S. (2022). 'A new low-activity nuclear density gauge for soil density measurements'. In: *Journal of Testing and Evaluation* 50.1, pp. 44–65. DOI: 10.1520/JTE20200483.

Doolittle, J.A. and Brevik, E.C. (2014). 'The use of electromagnetic induction techniques in soils studies'. In: *Geoderma* 223–225, pp. 33–45. ISSN: 0016-7061. DOI: 10.1016/j.geoderma.2014.01.027. URL: https://www.sciencedirect.com/science/article/pii/S0016706114000548.

Dunn, B.W. and Beecher, H.G. (2007). 'Using electro-magnetic induction technology to identify sampling sites for soil acidity assessment and to determine spatial variability of soil acidity in rice fields'. In: *Australian Journal of Experimental Agriculture* 47.2, pp. 208–214. DOI: 10.1071/EA05102.

Earl, R., Taylor, J.C., Wood, G.A., Bradley, I., James, I.T., Waine, T., Welsh, J.P., Godwin, R.J. and Knight, S.M. (2003). 'Soil factors and their influence on within-field crop variability, part I: field observation of soil variation'. In: *Biosystems engineering* 84.4, pp. 425–440. DOI: 10.1016/S1537-5110(03)00004-7.

Filippi, P., Jones, E.J. and Bishop, T.F.A. (2020). 'Catchment-scale 3D mapping of depth to soil sodicity constraints through combining public and on-farm soil databases–A potential tool for on-farm management'. In: *Geoderma* 374, p. 114396. DOI: 10.1016/j.geoderma.2020.114396.

Filippi, P., Jones, E.J., Ginns, B.J., Whelan, B.M., Roth, G.W. and Bishop, T.F.A. (2019). 'Mapping the depth-to-soil pH constraint, and the relationship with cotton and grain yield at the within-field scale'. In: *Agronomy* 9.5, p. 251. DOI: 10.3390/agronomy9050251.

Fleming, K.L., Westfall, D.G., Wiens, D.W. and Brodahl, M.C. (2000). 'Evaluating farmer defined management zone maps for variable rate fertilizer application'. In: *Precision Agriculture* 2, pp. 201–215. DOI: 10.1023/A:1011481832064.

Flowers, M., Weisz, R. and White, J.G. (2005). 'Yield-based management zones and grid sampling strategies: Describing soil test and nutrient variability'. In: *Agronomy Journal* 97.3, pp. 968–982. DOI: 10.2134/agronj2004.0224.

Al-Gaadi, K.A. (2012). 'Employing electromagnetic induction technique for the assessment of soil compaction'. In: *American Journal of Agricultural and Biological Sciences* 7.4, pp. 425–434. DOI: 10.3844/ajabssp.2012.425.434.

Grundy, M.J., Viscarra Rossel, R.A., Searle, R.D., Wilson, P.L., Chen, C. and Gregory, L.J. (2015). 'Soil and landscape grid of Australia'. In: *Soil Research* 53.8, pp. 835–844. DOI: 10.1071/SR15191.

Grunwald, S., Vasques, G.M. and Rivero, R.G. (2015). 'Fusion of soil and remote sensing data to model soil properties'. In: *Advances in agronomy* 131, pp. 1–109. DOI: 10.1016/bs.agron.2014.12.004.

Hajkowicz, S. and Young, M. (2005). 'Costing yield loss from acidity, sodicity and dryland salinity to Australian agriculture'. In: *Land Degradation & Development* 16.5, pp. 417–433. DOI: 10.1002/ldr.670.

Hao, X., Ball, B.C., Culley, J.L.B., Carter, M.R. and Parkin, G.W. (2008). 'Soil density and porosity'. In: *Soil Sampling and Methods of Analysis*. Ed. by Carter, M.R. and Gregorich, E.G. Vol. 2. CRC Press Boca Raton, FL, pp. 743–759.

Hengl, T., De Jesus, J.M., MacMillan, R.A., Batjes, N.H., Heuvelink, G.B.M., Ribeiro, E., Samuel-Rosa, A., Kempen, B., Leenaars, J.G.B., Walsh, M.G. and Gonzales, M.R. (2014). 'SoilGrids1km – global soil information based on automated mapping'. In: *PloS one* 9.8, e105992. DOI: 10.1371/journal.pone.0105992.

Hengl, T., Heuvelink, G.B.M., Kempen, B., Leenaars, J.G.B., Walsh, M.G., Shepherd, K.D., Sila, A., MacMillan, R.A., Mendes de Jesus, J., Tamene, L. and Tondoh, J.E. (2015). 'Mapping soil properties of Africa at 250 m resolution: Random forests significantly improve current predictions'. In: *PloS one* 10.6, e0125814. DOI: 10.1371/journal.pone.0125814.

Hengl, T., Mendes de Jesus, J., Heuvelink, G.B.M., Ruiperez Gonzalez, M., Kilibarda, M., Blagotić, A., Shangguan, W., Wright, M.N., Geng, X., Bauer-Marschallinger, B., Guevara, M.A., Vargas, R., MacMillan, R.A., Batjes, N.H., Leenaars, J.G.B., Ribeiro, E., Wheeler, I., Mantel, S. and Kempen, B. (2017). 'SoilGrids250m: Global gridded soil information based on machine learning'. In: *PLoS one* 12.2, e0169748. DOI: 10.1371/journal.pone.0169748.

Ji, W., Adamchuk, V.I., Chen, S., Su, A.S.M., Ismail, A., Gan, Q., Shi, Z. and Biswas, A. (2019). 'Simultaneous measurement of multiple soil properties through proximal sensor data fusion: A case study'. In: *Geoderma* 341, pp. 111–128. DOI: 10.1016/j.geoderma.2019.01.006.

Khanal, S., Kc, K., Fulton, J.P., Shearer, S. and Ozkan, E. (2020). 'Remote sensing in agriculture – accomplishments, limitations, and opportunities'. In: *Remote Sensing* 12.22, p. 3783. DOI: 10.3390/rs12223783.

Kidd, D., Searle, R., Grundy, M., McBratney, A.B., Robinson, N., O'Brien, L., Zund, P., Arrouays, D., Thomas, M., Padarian, J., Jones, E., Bennett, J.M. andMinasny. B., Holmes, K., Malone, B.P., Liddicoat, C., Meier, E.A., Stockmann, U., Wilson, P., Wilford, J, Payne, J., Ringrose-Voase, A., B., Slater., Odgers, N., Gray, J., Gool, D. van, Andrews, K., Harms, B., Stower, L. and Triantafilis, J. (2020). 'Operationalising digital soil mapping–Lessons from Australia'. In: *Geoderma Regional* 23, e00335. DOI: 10.1016/j.geodrs.2020.e00335.

Laiskhanov, S.U., Otarov, A., Savin, I.Y., Tanirbergenov, S.I., Mamutov, Z.U., Duisekov, S.N. and Zhogolev, A. (2016). 'Dynamics of Soil Salinity in Irrigation Areas in South Kazakhstan'. In: *Polish Journal of Environmental Studies* 25.6. DOI: 10.15244/pjoes/61629.

Lapen, D.R., Hayhoe, H.N., Topp, G.C., McLaughlin, N.B., Gregorich, E.G. and Curnoe, W.E. (2002). 'Measurements of mouldboard plow draft: II. Draft-soil-crop and yield-draft associations'. In: *Precision Agriculture* 3, pp. 237–257. DOI: 10.1023/A:1015519408578.

Lawrence, P.G., Roper, W., Morris, T.F. and Guillard, K. (2020). 'Guiding soil sampling strategies using classical and spatial statistics: A review'. In: *Agronomy Journal* 112.1, pp. 493–510. DOI: 10.1002/agj2.20048.

Lobell, D.B., Ortiz-Monasterio, J.I., Gurrola, F.C. and Valenzuela, L. (2007). 'Identification of saline soils with multiyear remote sensing of crop yields'. In: *Soil Science Society of America Journal* 71.3, pp. 777–783. DOI: 10.2136/sssaj2006.0306.

McBratney, A.B., Santos, M.L.M. and Minasny, B. (2003). 'On digital soil mapping'. In: *Geoderma* 117.1-2, pp. 3–52. DOI: 10.1016/S0016-7061(03)00223-4.

Minasny, B. and McBratney, A.B. (2006). 'A conditioned Latin hypercube method for sampling in the presence of ancillary information'. In: *Computers & geosciences* 32.9, pp. 1378–1388. DOI: 10.1016/j.cageo.2005.12.009.

Mulder, V.L., Lacoste, M., Richer-de-Forges, A.C. and Arrouays, D. (2016). 'GlobalSoilMap France: High-resolution spatial modelling the soils of France up to two meter depth'. In: *Science of the Total Environment* 573, pp. 1352–1369. DOI: 10.1016/j.scitotenv.2016.07.066.

Mulla, D.J. (2013). 'Twenty five years of remote sensing in precision agriculture: Key advances and remaining knowledge gaps'. In: *Biosystems engineering* 114.4, pp. 358–371. DOI: 10.1016/j.biosystemseng.2012.08.009.

Omran, E-S.E. and Negm, A.M. (2020). 'Smart Sensing System for Precision Agriculture'. In: *Technological and Modern Irrigation Environment in Egypt: Best Management Practices & Evaluation*, pp. 77–105. DOI: 10.1007/978-3-030-30375-4_5.

Orton, T.G., Mallawaarachchi, T., Pringle, M.J., Menzies, N.W., Dalal, R.C., Kopittke, P.M., Searle, R., Hochman, Z. and Dang, Y.P. (2018). 'Quantifying the economic impact of soil constraints on Australian agriculture: A case-study of wheat'. In: *Land Degradation & Development* 29.11, pp. 3866–3875. DOI: 0.1002/ldr.3130.

Paz, A.M., Castanheira, N., Farzamian, M., Paz, M.C., Gonçalves, M.C., Santos, F.A.M. and Triantafilis, J. (2020). 'Prediction of soil salinity and sodicity using electromagnetic conductivity imaging'. In: *Geoderma* 361, p. 114086. DOI: 10.1016/j.geoderma.2019.114086.

Raupach, M. and Tucker, B.M. (1959). 'The field determination of soil reaction'. In: *Journal of the Australian Institute of Agricultural Science* 25.2, pp. 129–133.

Rhoades, J.D. and Corwin, D.L. (1981). 'Determining soil electrical conductivity-depth relations using an inductive electromagnetic soil conductivity meter'. In: *Soil Science Society of America Journal* 45.2, pp. 255–260. DOI: 10.2136/sssaj1981.03615995004500020006x.

Román Dobarco, M., Bourennane, H., Arrouays, D., Saby, N.P.A., Cousin, I. and Martin, M.P. (2019). 'Uncertainty assessment of GlobalSoilMap soil available water capacity products: A French case study'. In: *Geoderma* 344, pp. 14–30. DOI: 10.1016/j.geoderma.2019.02.036.

Roudier, P., Burge, O.R., Richardson, S.J., McCarthy, J.K., Grealish, G.J. and Ausseil, A-G. (2020). 'National scale 3D mapping of soil pH using a data augmentation approach'. In: *Remote Sensing* 12.18, p. 2872. DOI: 10.3390/rs12182872.

Sadras, V., Baldock, J., Roget, D. and Rodriguez, D. (2003). 'Measuring and modelling yield and water budget components of wheat crops in coarse-textured soils with chemical constraints'. In: *Field Crops Research* 84.3, pp. 241–260. DOI: 10.1016/S0378-4290(03)00093-5.

Sharma, A., Weindorf, D.C., Man, T., Aldabaa, A.A.A. and Chakraborty, S. (2014). 'Characterizing soils via portable X-ray fluorescence spectrometer: 3. Soil reaction (pH)'. In: *Geoderma* 232, pp. 141–147. DOI: 10.1016/j.geoderma.2014.05.005.

Sharma, A., Weindorf, D.C., Wang, D. and Chakraborty, S. (2015). 'Characterizing soils via portable X-ray fluorescence spectrometer: 4. Cation exchange capacity (CEC)'. In: *Geoderma* 239, pp. 130–134. DOI: 10.1016/j.geoderma.2014.10.001.

Shatar, T.M. and McBratney, A.B. (2004). 'Boundary-line analysis of field-scale yield response to soil properties'. In: *The Journal of Agricultural Science* 142.5, pp. 553–560. DOI: 10.1017/S0021859604004642.

Stenberg, B. and Viscarra Rossel, R.A. (2010). 'Diffuse Reflectance Spectroscopy for High-Resolution Soil Sensing'. In: *Proximal Soil Sensing*. Ed. by Viscarra Rossel, R.A., McBratney, A.B. and Minasny, B. Dordrecht: Springer Netherlands, pp. 29–47. ISBN: 978-90-481-8859-8. DOI: 10.1007/978-90-481-8859-8_3. URL: https://doi.org/10.1007/978-90-481-8859-8_3.

Tan, K.H. (2005). *Soil sampling, preparation, and analysis*. CRC press. DOI: 10.1201/9781482274769.

Taylor, J.C., Wood, G.A., Earl, R. and Godwin, R.J. (2003). 'Soil factors and their influence on within-field crop variability, part II: spatial analysis and determination of management zones'. In: *Biosystems engineering* 84.4, pp. 441–453. DOI: 10.1016/S1537-5110(03)00005-9.

Tomlinson, R.F. (1978). 'Design considerations for digital soil map systems.' In: *Transactions, 11th International Congress of Soil Science* Transactions, 11th International Congress of Soil Science, pp. 191–207.

Trengove, S. and Sherriff, S. (2019). 'pH Mapping and variable Rate Lime Application'. In: *GRDC 2019 Grains Research Update – South*. Grains Research and Development Corporation (GRDC). URL: https://grdc.com.au/resources-and-publications/grdc-update-papers/tab-content/grdc-update-papers/2019/02/ph-mapping-and-variable-rate-lime-application.

Triantafilis, J., Ahmed, M.F. and Odeh, I.O.A. (2002). 'Application of a mobile electromagnetic sensing system (MESS) to assess cause and management of soil salinization in an irrigated cotton-growing field'. In: *Soil Use and Management* 18.4, pp. 330–339. DOI: 10.1111/j.1475-2743.2002.tb00249.x.

Viscarra Rossel, R.A., Adamchuk, V.I., Sudduth, K.A., McKenzie, N.J. and Lobsey, C. (2011). 'Chapter Five - Proximal Soil Sensing: An Effective Approach for Soil Measurements in Space and Time'. In: *Advances in Agronomy* 113. Ed. by Sparks, D.L., pp. 243–291. ISSN: 0065-2113. DOI: 10.1016/B978-0-12-386473-4.00005-1. URL: https://www.sciencedirect.com/science/article/pii/B9780123864734000051.

Viscarra Rossel, R.A., McBratney, A.B. and Minasny, B. (2010). *Proximal soil sensing*. Ed. by Viscarra Rossel, R.A., McBratney, A.B. and Minasny, B. Springer Science & Business Media. DOI: 10.1007/978-90-481-8859-8.

Webster, R., Harrod, T.R., Staines, S.J. and Hogan, D.V. (1979). *Grid sampling and computer mapping of the Ivybridge area, Devon.* Tech. rep. Rothamsted Experimental Station, Harpenden, Herts AL5 2JQ, UK.

Wells, L.G., Stombaugh, T.S. and Shearer, S.A. (1998). 'Application and assessment of precision deep tillage'. In: *2001 ASAE annual meeting*. American Society of Agricultural and Biological Engineers, p. 1. DOI: 10.13031/2013.6258.

Xue, J. and Su, B. (2017). 'Significant remote sensing vegetation indices: A review of developments and applications'. In: *Journal of sensors* 2017.1, p. 1353691. DOI: 10.1155/2017/1353691.

APPENDIX B

Proximal soil sensing in the field

Mikaela J. Tilse[ac], Uta Stockmann[b], Patrick Filippi[ac]

[a]Sydney Institute of Agriculture, The University of Sydney, Camperdown, NSW, Australia.

[b]Commonwealth Scientific and Industrial Research Organisation (CSIRO), Agriculture and Food Business Unit, Canberra, ACT, Australia.

[c]School of Life and Environmental Sciences, The University of Sydney, Camperdown, NSW, Australia.

Abstract

Proximal soil sensing in the field is the application of handheld, ground, or vehicle-based sensors to obtain data when the sensor is placed in contact, or within 1–2 m proximity, of the soil. Proximal soil sensors operate as on-the-go, stationary, or stop-and-go systems, and cover a broad range of the electromagnetic spectrum. Field-based proximal soil sensors are a cheaper, less time and labor-intensive alternative to conventional laboratory analyzes, but are an indirect method that often requires site and time-specific spectral libraries and calibration models for meaningful interpretation. Multi-sensor data fusion and multi-sensor approaches are explored, alongside future directions.

Key points

- Proximal soil sensors are defined and classified according to their model of operation, with particular focus on applications in the field.
- Advantages and disadvantages of proximal soil sensing is discussed.

- Types and applications of on the-go, point-based, and stop-and-go field-based proximal soil sensors are described.
- The development and importance of calibration models is highlighted.
- Multi-sensor data fusion and multi-sensor approaches are discussed.
- Future directions of proximal soil sensing are explored, including shifting technologies from the research to commercialization phases.

Introduction

Soil varies spatially across landscapes and with depth, and temporally both within, and between years. Understanding this spatial and temporal variation is crucial for a range of stakeholders, including land managers, farmers, and policy makers. Timely and accurate information describing soil variability is needed at increasingly finer (spatial and temporal) scales, but obtaining a representative sample of a particular soil landscape is an ongoing challenge (Adamchuk *et al*. 2004). While conventional soil sampling and laboratory analysis methods have contributed greatly to our scientific understanding of the soil landscape, they are often laborious, time-consuming, and expensive (Viscarra Rossel *et al*. 2011). Sensor-based technologies, which can be used in the field can overcome these shortcomings and facilitate the acquisition of high-resolution soil data, allow better understanding of the spatial and temporal dynamics of soil.

Proximal soil sensing (PSS) in the field is the application of handheld, ground, or vehicle-based sensors to obtain data on (soil) properties when the sensor is placed in contact, or within 1–2 m proximity, of the soil (Viscarra Rossel *et al*. 1998; Viscarra Rossel *et al*. 2010). The signals detected can be related to specific soil properties directly, or are proxies which require calibration. This definition of PSS excludes remote sensing via airborne or satellite platforms, and laboratory measurements with benchtop instruments.

Proximal soil sensors can rapidly provide users with cost-effective, detailed and timely information about soil physical, chemical, and biological properties (Viscarra Rossel *et al.*, 2011). While some proximal sensors may not provide results that are as accurate or precise as

conventional analyzes, per individual measurement, the collection of large amounts of (spatial and temporal) data via PSS offers users with more information than ever before about the soil landscape. Technological advances, improved accessibility, and reduced costs of PSS tools are contributing to an increased shift of PSS from the research-phase, to being implemented in more applied and commercial settings. The list of available PSS systems is expanding year-on-year, although the rate varies for different sensors, and in different areas globally.

This chapter focuses on PSS in the field and does not cover the use of proximal sensors in a laboratory setting. Types of proximal soil sensors, recent advances in commercialization, and their applications will be described. Multi-sensor approaches, sensor calibration and prediction will also be explored, including the future of PSS. Numerous reviews on the principles and applications of PSS for agriculture and natural resource management can be found elsewhere and are referred to throughout.

Sensor classification

Viscarra Rossel *et al*. (2011) classify proximal soil sensors into four main categories according to:

- The manner in which they measure (invasive [in situ or ex situ] or non-invasive);
- Their energy source (active or passive);
- How they operate (stationary or mobile); and
- The inference used in the measurement of the target soil property (direct or indirect).

Additional characterization criteria from Gebbers (2019) also classify proximal soil sensors according to:

- Their temporal responsiveness (the time taken for measurement, processing, and subsequent decision-making);
- Their selectivity (the extent to which a sensor can determine a particular property without interference); and

- Scope of the electromagnetic (EM) spectrum that they encompass (ranging from gamma-rays to mechanical methods).

Table B1 describes the classifications of common proximal soil sensors within the EM spectrum used in the field, and their applications to a range of soil properties.

Invasive sensors require a degree of sample preparation or direct contact (e.g., insertion into or extraction of the soil). Invasive in situ sensors measure soil properties in place, usually in direct contact with the soil (e.g., soil moisture sensors), whereas invasive ex situ systems require sample extraction and placement of the sample in contact with the sensing element for analysis (e.g., ion-selective electrodes) (Adamchuk *et al*. 2018). Non-invasive sensors measure from a proximal distance and do not require any sample preparation or processing (e.g., electromagnetic induction, EMI, sensors). Passive sensors measure ambient sources of energy (e.g., natural radioactive decay), whereas active sensor systems provide their own source of energy (e.g., ground penetrating radar (GPR)) (Viscarra Rossel *et al*. 2011).

Proximal soil sensors operate as either on-the-go (mobile), stationary (point-based), or stop-and-go sensors (see section on types of sensors and their applications). On-the-go (mobile) sensors continuously collect data as they move across an area (e.g., EMI; Gebbers 2019. Stationary sensors may be handheld (e.g., handheld infrared spectrometer) or permanently fixed sensors (e.g., a soil moisture probe buried in the soil), and must remain stationary during sample collection and measurement. Stop-and-go sensors can be moved from place-to-place, but need to remain stationary during sample collection and operate automatically (e.g., Veris; see section on B). Sensors that can provide data immediately (e.g., soil moisture or pH sensors) can be used for real-time management decisions (e.g., irrigation scheduling). If data processing, including cleaning and calibration, is required, sensors can be used for asynchronous applications (e.g., Digital Soil Mapping, DSM) (Gebbers 2019).

Sensor selectivity indicates how well a sensor can determine a particular property without interference (Gebbers 2019). High selectivity sensors (e.g., pH electrodes) are preferred, however there is a trade-off between sensor selectivity, cost, and robustness (Gebbers 2019). Non-selective sensors such as EMI sensors, which measure electrical conductivity (EC), are

widely used to assess a range of soil properties. While sensors that measure a physical process (i.e., direct measures) are most desirable, these are generally more expensive, less developed, and more technically demanding (Viscarra Rossel *et al*. 2016). Indirect methods, which rely on a proxy, are more readily available, more mature, and are cheaper to own and operate (Viscarra Rossel *et al*. 2016). However, these may be less accurate as they require site-specific calibration, are inherently more subjective (as they rely on proxies), and may be subject to inference from other soil constituents (Viscarra Rossel *et al*. 2016).

The most commonly used proximal soil sensors in research, applied, and commercial settings are electrical and EM sensors (e.g., EMI), X-ray, optical and radiometric sensors (e.g., Gamma-rays, Near and Mid-Infrared spectroscopy), and capacitance probes (e.g., soil moisture sensors) (Table B1). Other sensor types not covered in detail in this chapter include acoustic, pneumatic, and mechanical sensors.

TABLE B1. Commonly used proximal soil sensors in the field and their relative position on the electromagnetic spectrum.

		Sensor type											
		10^{-12}		10^{-10}		10^{-8} - 10^{-4}		10^{-2}		10^{1} - 10^{6}			
		Gamma ray		X-ray		Optical		Microwave		Radio wave			
Classification	Method	TNM	Spectroscopy	XRF	Visible	NIR	MIR	TDR	FDR/Capacitance	GPR	NMR	EMI	EC/ER
	Energy[a]	A	A/P	A	A/P	A/P	A	A	A	A	A	A	A
	Interaction[b]	I	N	N	I/N	I/N	I/N	I	I	N	N	N	I
	Operation[c]	S/M	S/M	S	S/M	S/M	S	S	S/M	S/M	S/M	S/M	M
	Development status[d]	R/C	C	C	C	R/C	R	C	C	C	R	C	C
Soil property	Chemical properties												
	Organic matter content					X		X		X			
	Total Carbon	X		X		X							
	Organic Carbon	X		X		X							
	Inorganic Carbon	X				X							
	Nitrogen	X		X		X		X		X			
	Phosphorous	X		X		X							
	Potassium	X		X		X							
	Micronutrients, elements	X		X		X							
	Heavy metals	X		X		X							
	CEC	X				X				X			
	Soil pH	X				X		X		X			
	Salinity							X		X			
	Physical properties												
	Colour			X									
	Water content	X		X		X		X		X			
	Soil texture (clay, silt, and sand)	X		X		X				X			
	Clay minerals	X		X		X							
	Bulk density	X		X		X				X			

Note: TNM: thermalized neutron methods; XRF: X-ray fluorescence; NIR: near infrared; MIR: mid infrared; TDR: time-domain reflectometry; FDR: frequency-domain reflectometry; GPR: ground-penetrating radar; NMR: nuclear magnetic resonance; EMI: electromagnetic induction; ER: electrical resistivity.

[a] Proximal sensors can be classified by their energy source (active sensors (A) provide their own source of energy, passive sensors (P) rely on ambient or emitted energy).

[b] Proximal sensors can be classified by their method of interaction with the soil (invasive sensors (I) rely on direct contact with soil, non-invasive sensors (N) are operated without any soil distortion).

[c] Proximal sensors can be classified by their mode of operation (stationary operation (S) requires placing the sensor in a specific geographic location at a fixed or variable depth, mobile operation (M) allows on-the-go soil sensing).

[d] Proximal sensors can be classified by their development status (either remaining in the research phase (R) or has been commercialized (C)).

Adapted from Adamchuk VI, Hummel JW, Morgan MT, Upadhyaya SK (2004) On-the-go soil sensors for precision agriculture. *Computers and Electronics in Agriculture* **44**: 71–91. https://doi.org/10.1016/j.compag.2004.03.002; Viscarra Rossel RA, Adamchuk VI, Sudduth KA, McKenzie NJ, Lobsey C (2011) Proximal soil sensing: An effective approach for soil measurements in space and time. *Advances in Agronomy* **113**: 243–291. https://doi.org/10.1016/B978-0-12-386473-4.00005-1; Kuang B, Mahmood HS, Quraishi MZ, Hoogmoed WB, Mouazen AM, van Henten EJ (2012) Sensing Soil Properties in the Laboratory, in situ, and on-line: A review. *Advances in Agronomy*, pp. 155–223, Elsevier Inc. https://doi.org/10.1016/B978-0-12-394275-3.00003-1; Viscarra Rossel RA and Lobsey C (2016) *Scoping review of proximal soil sensors for grain growing*. 52. https://doi.org/10.13140/RG.2.2.34785.51049.

Advantages of proximal sensing

Proximal soil sensors can facilitate the (near-instantaneous) collection of high-resolution spatial (or temporal) data using simpler, less time and labor-intensive techniques compared to conventional laboratory methods (Viscarra Rossel *et al*. 2011). While data from proximal sensors may not be as precise as laboratory measurements (Viscarra Rossel *et al*. 2011), users are offered more information than ever before about the soils condition to characterize the (horizontal and vertical) spatial and temporal variability of soil properties (Viscarra Rossel *et al*. 2016) with minimal disturbance.

Disadvantages of proximal sensing

Many proximal sensors are indirect measures that rely on an inference of the relationship between the sensor signal (a proxy) and the properties of interest for meaningful interpretation. These calibrations and empirical relationships are, in many cases, specific to certain soil and environmental conditions that may change over time (Adamchuk *et al*. 2011). Thus, the use of proximal sensors is restricted to the soil conditions it has been calibrated for. Complex interactions and interrelationships between multiple soil properties limits the number of attributes that can be measured with a single sensor, and multivariate calibrations to account for these interactions and the non-specificity of some proximal sensor outputs may be required (Stenberg *et al*. 2010). The performance of different sensors varies considerably due to different soil and environmental factors, including soil parent material, water content, and temperature (Kuang *et al*. 2012), and the deployment of proximal sensors is limited by access to the area of interest. This is a particular issue on waterlogged soils (e.g., rice paddies) or in densely forested areas.

Types of sensors and their applications

Proximal soil sensors range from simple handheld, stationary systems, through to more complex, autonomous robotic sensor platforms and are in various stages of development

and commercialization (Gebbers 2019). Systems available "off the shelf" can be purchased directly and may be ready to use, or may require some additional development or calibration before use. Others are in mature stages of research where additional testing for implementation is needed, or remain in the 'proof of concept' phase. Advancements in sensor capabilities, and information and communication technologies (e.g., telemetry systems and cloud-based services) have aided the development of PSS systems and are contributing to their increased adoption in-the-field.

There are issues and opportunities associated with on-the-go, stationary, or stop-and-go systems that cover a broad range of the EM spectrum (Table B1).

On-the-go sensors

Logistical constraints limit the number of samples that can be obtained via conventional sampling strategies or point-based proximal soil sensors. Improvements in the quality, accuracy, and adoption of global positioning systems (GPS) and geographical information systems (GIS) within agriculture and landscape management as well as ever increasing computing power has supported more intensive soil sampling strategies and the rise of spatial on-the-go sensing and mapping. On-the-go sensing technologies enable the collection of high-resolution soil information and can facilitate the interpretation of soil spatial patterns using a more time and cost-effective method compared to conventional analysis techniques.

On-the-go (mobile) sampling presents the "ultimate solution" (Adamchuk *et al*. 2018) for collecting large amounts of data for spatial interpolation and decision making. On-the-go proximal soil sensors are typically mounted to tractors, all-terrain vehicles, or other farm machinery and measurements are recorded while continuously moving across a field (Figure B1). Most commercial on-the-go PSS tools measure soil properties indirectly from on or above the soil surface, although sensor operation varies for different systems. The information collected is often used to derive maps of soil properties, which are then used as covariates for DSM or spatial interpolation methods.

Electromagnetic induction (EMI)

Electromagnetic induction is a noninvasive, ex situ soil sampling method, and is one of the most widely-used on-the-go PSS tools worldwide. Electromagnetic induction sensors measure the apparent electrical conductivity (EC_a) through the transmission and reflection of EM waves between two or more electric coils (solenoids) (Gebbers 2019), effectively turning the soil into an EM circuit. Soil properties that (directly or indirectly) influence conductivity can be inferred from EC_a measurements (via sampling and calibration), and the most common soil properties that are correlated with EC_a are soil clay content (texture), soil moisture, and salinity (Table B2; Whelan *et al*. 2013). Calibrations are site and time-specific, which limits the application of EMI to the soil conditions it has been calibrated for. Readings are influenced by the soil's water and salt content, temperature, and mineral type and composition, making universal calibrations complex and difficult to obtain (Whelan *et al*. 2013; Doolittle *et al*. 2014).

FIGURE B1. On-the-go Dual EM-21S electromagnetic induction (EMI) meter (Dualem Inc., Ontario, Canada) and RSX-1 gamma-ray spectrometer (Radiation Solutions Inc., Ontario, Canada) can be used as part of a multi-sensor platform to map soil moisture content and salinity (via EMI), and soil type and parent material (via Gamma). Together, these complementary sensors can be used to map patterns of soil variability, alongside a high-precision GPS unit mounted to the vehicle.

Compared to conventional sampling, EMI is a non-invasive, relatively quick, easy-to-use, cost-effective and data-rich measure of soil variability (Doolittle *et al*. 2014). When coupled with data logging and GPS technologies, a high density of spatial data can be gathered on-the-go by surveying the area along transects or swathes (Figures B1 and B2). Data is logged at regular (1–3 s) intervals from (typically) ground-based vehicles at (potentially) multiple depths down to 3 or 6 m for the most commonly used EMIs in soil science applications

TABLE B2. General impact of soil properties on soil apparent electrical conductivity (EC_a) readings.

	Material	Soil texture	Clay type	Moisture potential	CEC	Salinity	Nutrients
Higher ECa	Soil	Clay	Smectite	Saturated	High CEC	High EC	High N, P, K, Ca, Mg, Na
	Stony soil	Silt	Illite	Field capacity			
Lower ECa	Rock	Sand	Kaolin	Wilting point	Low CEC	Low EC	Low N, P, K, Ca, Mg, Na

Adapted from Whelan BM and Taylor J (2013) Precision agriculture for grain production systems. *International Journal of Digital Earth*. https://doi.org/10.1080/17538947.2013.817183.

(refer to Figure B1 for an example) (Viscarra Rossel *et al*. 2011), with effective sensing depth largely determined by coil spacing, orientation, frequency, and height above the ground (Gebbers 2019). Numerous EMI sensors are commercially available for use in the field, including the DUALEM (Dualem, Inc., Milton, Ontario, Canada) and EM38 (Geonics Limited, Mississauga, Ontario, Canada) meters. Each commercial sensor has operational advantages and disadvantages depending on the depths and soil properties of interest, the preferred data collection mode, and soil condition at the time of sampling (Sudduth *et al*. 2003).

Gamma-rays

Mobile gamma-ray spectrometry (GRS) is a noninvasive, ex situ soil sampling method used to estimate the concentrations of naturally in the soil occurring radioelements (Potassium [40K], Uranium [238U], Thorium [232Th], and their ratios) below the earth's surface (Reinhardt *et al*. 2019). This is done by measuring the gamma rays of radioactive isotopes emitted during radioactive decay. Passive and active gamma ray sensors are available, although passive gamma-ray sensors are preferred due to radiation and safety concerns associated with active sensors, which provide their own source of radiation (Adamchuk *et al*. 2018).

Simple gamma-ray sensors detect bulk radiation (total count), while others can distinguish different energy levels (Gebbers 2019). Three key regions of interest (ROIs) in the gamma spectra (40K, 238U, and 232Th) are commonly used to (directly and indirectly) infer information for understanding land forming processes and soil attribute distributions; primarily

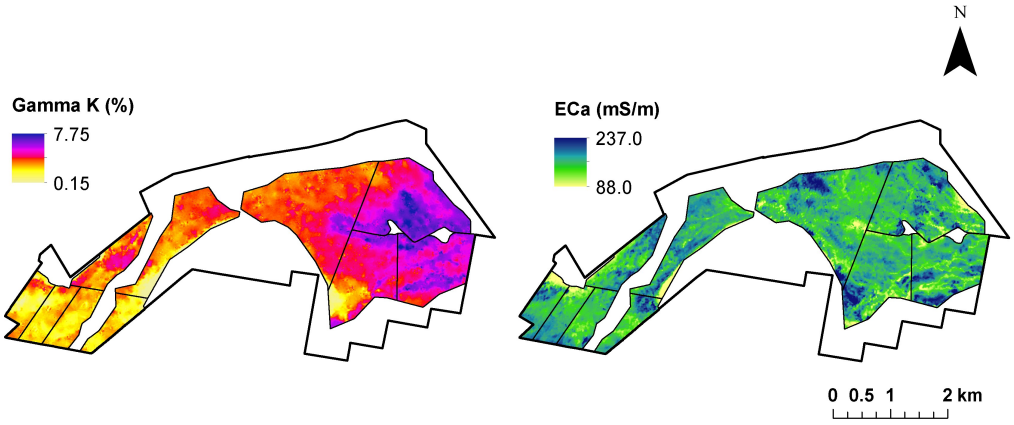

FIGURE B2. Gamma radiometrics K ROI and Soil EC_a (0–3 m) maps derived from a RSX-1 gamma-ray spectrometer (Radiation Solutions Inc., Ontario, Canada) and a Dual EM-21S electromagnetic induction (EMI) meter (Dualem Inc., Ontario, Canada), respectively, across "L'lara", a property close to Narrabri in North West NSW, Australia. Gamma radiometrics primarily relate to changes in parent material and soil type, with high Gamma K values related to soils with a low soil moisture potential (i.e., wilting point), high cation exchange capacity and K content, and a shallow depth to rock/gravel. EMI relates to soil texture, moisture, and salinity, and high EC_a values generally relate to soils with a high moisture potential (saturated), high clay (smectite) content with a high nutrient (N, P, K, Ca, Mg, Na) content and EC. Together, these on-the-go soil measurements can be used to describe soil variability across a landscape.

for parent material characterization and soil type differentiation (Table B3; Whelan *et al.* 2013; Reinhardt *et al.* 2019). Direct, standalone total count readings may be sufficient for delineating basic patterns of soil variability, but a more detailed interpretation and relation to specific soil properties requires sophisticated data processing (Pätzold *et al.* 2020).

TABLE B3. General impact of soil properties on soil gamma emission (total count; TC).

	Material	Soil texture	Clay type	Moisture potential	CEC	Depth to rock or gravel	Nutrients
Higher TC	Soil	Clay	Smectite	Wilting point	High CEC	Shallow	High K
	Stony soil	Silt	Ilite	Field capacity			
Lower TC	Soil	Sand	Kaolin	Saturated	Low CEC	Deep	Low K

Adapted from Whelan *et al.* (2013) Precision agriculture for grain production systems. *International Journal of Digital Earth*. https://doi.org/10.1080/17538947.2013.817183.

Gamma counts can be logged continuously on-the-go while driving over an area (Figures B1 and B2). Approximately 90% of all aboveground gamma radiation originates from the upper 30–50 cm of the soil profile (Gebbers 2019), although effective depth largely depends on bulk density and moisture content (Reinhardt *et al*. 2019). As for EMI, GRS enables the fast collection of dense spatial data and has a relatively well-established theoretical background (Gebbers 2019). However, careful site- and time-specific calibrations are required to account for different geopedological conditions (Pätzold *et al*. 2020). The most common commercially available gamma-ray sensors for use in the field include the RS-X (Radiation Solutions Inc., Ontario, Canada) and the Medusa (Medusa Radiometrics BV, Groningen, The Netherlands) gamma-ray spectrometer series.

While EMI and Gamma-ray spectrometer readings can be used individually to create soil property maps (e.g., soil EC_a, cation exchange capacity (CEC), soil texture), these maps are often difficult for users to understand and are not useful for making management decisions. Instead, on-the-go EMI and Gamma-ray sensors are often used in combination to map variability across the soil landscape and diagnose potential constraints and limitations (Figure B2). Gamma-ray spectrometer readings are largely impacted by shallow rocky materials, and total-count values are high in soils with a large clay (smectite), CEC and K contents, and a low soil moisture potential (i.e., wilting point). Soil EC_a readings (from EMI) are primarily related to soil texture, moisture, and salinity. High EC_a values relate to soils with high moisture potential (saturated), and large EC, clay (smectite), and nutrient (N, P, K, Ca, Mg, Na) content. The inverse response of these two soil measurements to soil moisture content, salinity status, and rocks/ stones means that they can be used together to identify the causal factors of high EC_a readings (whether that be soil moisture, salinity, clay, or CEC), and identify patterns of variability. For example, low gamma counts are associated with sandy soils, and EMI can be used to determine if these sands are saline or not. These instruments are often used in tandem as part of a multi-sensor platform (Figure B1).

COSMOS

Cosmic-ray Neutron Sensing (CRNS; Zreda *et al*. 2008) is a passive, non-invasive, indirect soil sensing tool used to estimate soil moisture based on the inverse relationship between the intensity of naturally occurring cosmic-ray generated neutrons, and the amount of water present in the surrounding environment (Desilets *et al*. 2010). While CRNS probes are used within agricultural fields and as part of larger soil moisture monitoring programs (i.e., the COsmic-ray Soil Moisture Observing System, COSMOS; Zreda *et al*. 2012), this stationary data does not capture the inherent spatial heterogeneity of soil moisture at a range of scales. A ground-based, mobile CRNS, called the cosmic-ray rover (Desilets *et al*. 2010), has been used for spatial soil moisture surveys over large areas (Figure B3). However, there are few cosmic-ray rovers across the world, and they are not available commercially. As for all indirect soil sensing methods, site and time-specific calibrations are required to account for local conditions at the time of sampling. The calculation of volumetric soil moisture content requires information on the spatial variation of bulk density, soil organic matter, and lattice water (i.e., the amount of hydrogen held within the lattice structure of the soil minerals), which requires spatial data at an appropriate scale or additional spatial sampling to collect the necessary information. Vegetation cover and roads can also influence soil moisture readings from mobile cosmic-ray sensors (McJannet *et al*. 2017).

Time-domain reflectometry (TDR) and ground-penetrating radar (GPR)

Time-domain reflectometry (TDR) and ground penetrating radar (GPR) are active sensors that operate based on the soil's dielectric constant (or the soil's ability to store electrical charge), and can be related to soil moisture. Both tdr and GPR rely on the same principle in that they measure the travel times and amplitudes of an EM pulse (Corwin 2008; Gebbers 2019). However, TDR uses two (or more) buried electrodes that are in direct contact with the soil, while GPR operates indirectly above the soil surface (Gebbers 2019). TDR sensors are mainly used for stationary measurements, but prototype stop-and-go mobile TDR sensors have been developed (Thomsen *et al*. 2007) and more work is needed to further develop and commercialize continuous mobile TDR measurements (Gebbers 2019). Mobile GPR

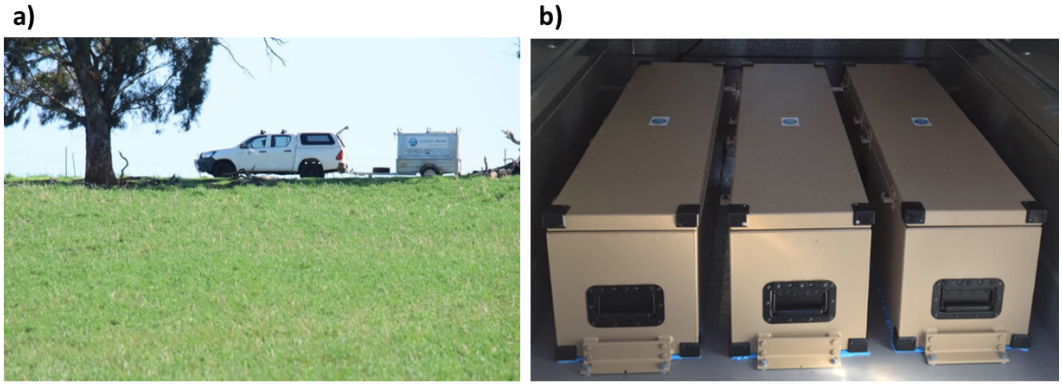

FIGURE B3. Example of a mobile cosmic ray neutron sensing (CRNS) system, the 'CosmOz Rover'–CSIRO (https://cosmoz.csiro.au/about), operating in the field (a) and a view of the cosmic-ray neutron sensors inside the trailer. The CosmOz Rover has a counting rate equal to about 20 standard stationary cosmic-ray probes, which enables the collection of the spatial variation in neutron counting rates while driving the trailer through the landscape.

sensors are commercially available and are predominately used for geophysical exploration and hydrology (Knight 2001).

Soil compaction and strength

In addition to proximal soil sensors that operate within the EM spectrum, soil compaction, impedance or strength can be measured on-the-go or at a single point via mechanical approaches. Mechanical sensors either measure the amount of energy required to pull an implement through the soil (e.g., tine), or measure the resistance of the soil to the insertion of a probe (penetration resistance) (Adamchuk *et al*. 2018). Broadly, any soil tillage tool (e.g., plows, load cells) can be used to measure the spatial variability of soil strength and compaction at single or multiple depths, and are relatively inexpensive, robust, and easy-to-use. Acoustic and pneumatic sensors may also be used to assess soil structure and compaction.

Point-based sensors

Point-based, or stationary, proximal sensing systems can be operated in the field to make a single measurement, produce several measurements at different depths at a given location, or continuously monitor soil parameters when installed at a site for a period of time (Adamchuk

et al. 2011). These may be handheld sensors, which can be moved from site-to-site (e.g., handheld infrared spectrometers), or be permanently fixed sensors, and must remain stationary during sample collection and analysis. Although point-based proximal sensors have a smaller spatial sampling density compared to on-the-go and stop-and-go systems, the ability to continuously monitor soil properties and understand their temporal dynamics (Adamchuk *et al*. 2018) means that they are an important component of proximal sensing repertoire for farmers and land managers.

Visible, near and mid-infrared diffuse reflectance spectroscopy

Diffuse reflectance infrared spectroscopy (DRS) is a proven method for rapid soil characterization under soil laboratory conditions (Stenberg *et al*. 2010) and with the development of field portable systems is slowly emerging for the rapid and cost-effective analysis of soil physical, chemical, and biological properties in field settings (Stockmann *et al*. 2018). Infrared spectroscopy is based on the reflection, absorbance, and scattering of light (photons) to measure soil properties, based on characteristic peaks or valleys at certain wavelengths in the visible (vis; 390–700 nm), near-infrared (NIR; 700–2500 nm), and mid-infrared, MIR; 2500–25,000 nm) parts of the electromagnetic spectrum. Spectral reference libraries, which are collections of known relationships between spectra and soil properties, are key to DRS. Calibration models are developed from these reference libraries to produce predictions of soil properties at varying levels of accuracy (Soriano-Disla *et al*. 2014). Additional pre-processing transformations may also be required to reduce noise, address the effects of field condition scanning such as soil moisture (e.g., external parameter orthogonalization, EPO; Minasny *et al*. 2011), and enhance more chemically-relevant peaks (Stenberg *et al*. 2010). Applications of DRS in research and commercial settings is increasing because it is a rapid, non-destructive, less labor-intensive, and more cost-effective alternative to conventional laboratory analyzes alone. There is a growing number of portable, handheld NIR spectrometers being developed and made available commercially for soil sensing applications (Tang *et al*. 2020), including from Hone Carbon (https://www.honecarbon.com/) and Yard Stick (https://www.useyardstick.com/), which can both be used to measure soil carbon levels in the field (B4). However, developing representative and diverse spectral

libraries for these technologies remains a challenge at present for widespread commercial uptake.

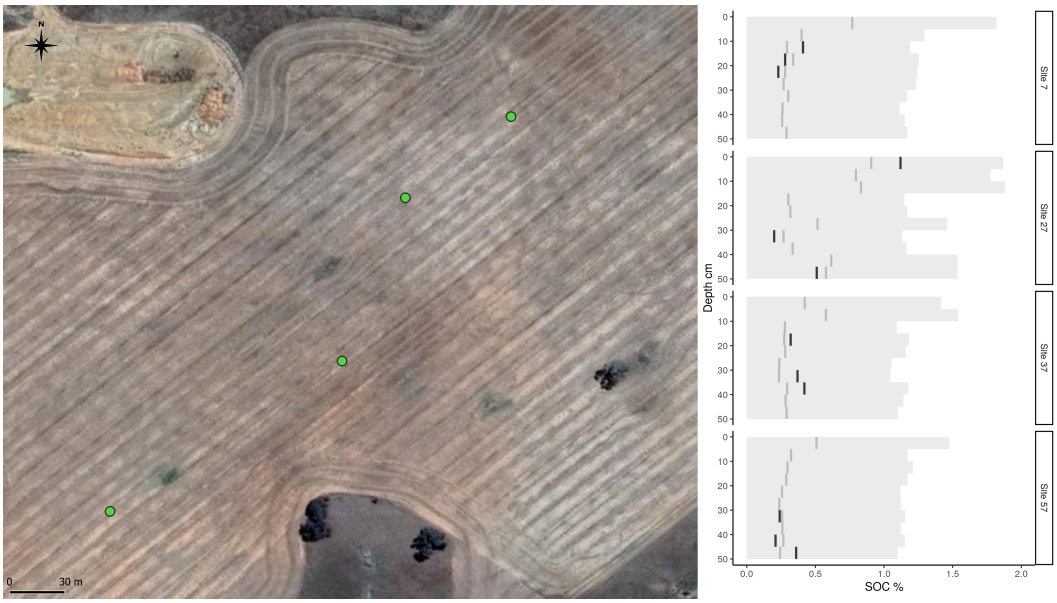

FIGURE B4. Predictions of Soil Organic Carbon (SOC) % using the Hone Lab Red sensor from 0 to 50 cm at 5 cm increments (dark grey) in four locations across the CSIRO Boorowa Agricultural Research Station long-term trial site in NSW, Australia, using a local calibration model with 50 samples produced through an ensemble cubist model with 50 bags. The prediction interval (light grey) is also shown and observed SOC analytical values are presented in black. Adapted from Moloney J, Stockmann U, Karunaratne S, Malone B, Wheeler S, Brown C, Gray B, Howard S and Shiley D (2022) Australia's national soil spectral library empowering rapid soil organic carbon measurement. In: *Feasibility Study Report for the National Soil Carbon Innovation Challenge Feasibility Study Grants Project SCMICF000073*. CSIRO: Australia.

Portable X-ray fluorescence spectrometry

Portable X-ray fluorescence (pXRF) spectroscopy is a widely used, non-destructive, indirect proximal sensing tool used to determine the elemental composition of a sample, based on the excitation and release of energy from electrons. X-ray fluorescence spectrometers emit X-rays which are absorbed by electrons and shift them into an excited energy state, resulting in the release of X-ray energy which is referred to as X-ray fluorescence. This measured energy spectrum (fluorescence) contains peaks that are characteristic for each element, and this X-ray absorption can be used to identify and quantity different elements in a sample (Weindorf

et al. 2014; Gebbers 2019). Robust pXRF instruments can quantify up to 42 elements (Gebbers 2019), but the detectable range is greater for more sophisticated instruments. XRF spectroscopy is an established laboratory method in earth sciences, and field-based, handheld pXRF instruments are increasingly becoming available at decreasing prices (Gebbers 2019). Field pXRF spectrometers enable the rapid, relatively easy, in- and ex situ analysis of solid soil samples with minimal sample preparation compared to conventional laboratory analysis and are commonly used for land contamination assessments (Pozza *et al*. 2020). However, invasive site-preparation is still necessary to prepare soil pits or extract cores for scanning (Weindorf *et al*. 2014). As for other indirect methods, there are several sources of interference on pXRF signals mainly soil matrix heterogeneity, but also soil moisture content, and thus the degree of sample preparation (Weindorf *et al*. 2014), although pXRF is less impacted by soil moisture than is DRS.

Soil moisture sensors

Proximal soil moisture sensors may directly or indirectly measure soil moisture content from above the soil surface or via direct contact. Stationary soil moisture sensors can be used to take single measurements, multi-depth measurements, or continuously monitor soil moisture at a given location. Linking numerous moisture sensors together via wireless networks can also be used to monitor the spatial and temporal dynamics of soilmoisture (Adamchuk *et al*. 2018). Many soil moisture sensors are relatively mature technologies, and commonly used proximal soil moisture sensors in the field include capacitance probes and frequency domain reflectometry (FDR) sensors, matric potential sensors, active neutron probes and gamma-ray sensors, and passive cosmic-ray sensors.

Capacitance probes and FDR sensors are both active, invasive sensors that measure soil moisture content based on the soils dielectric constant by propagating an EM wave through the soil (i.e., using it as a dielectric medium). Matric potential sensors (e.g., tensiometers, gypsum blocks, and granular matrix sensors) measure the amount of force (suction) required to extract water from the soil (i.e., the soils matric potential), rather than the soil moisture content (Hardie 2020). Capacitance, FDR, and matric potential sensors all require direct

soil-to-sensor contact to obtain accurate readings, and the measurement area (i.e., the volume of inference around the sensor) is relatively small (Hardie 2020).

Other field-based, stationary soil moisture sensors include active neutron probes and gamma-ray sensors, and passive cosmic-ray sensors. These noninvasive sensors all facilitate the fast, accurate, and continuous monitoring of soil moisture at a fixed point, and are particularly well-suited to hard-set, vertic (shrink-swell) soils where the installation of other soil moisture sensor types is difficult (Hardie 2020). However, they are costly, require complex and precise calibration methods, and the measurement footprint of cosmic ray neutron sensors in particular (i.e., 200–600 m) means they are poorly suited to short-range precision agriculture treatments such as variable rate irrigation (Hardie 2020). Further, application of active neutron and gamma-ray sensors in the field is limited due to radiation safety concerns.

Stop-and-go sensors

Stop-and-go sensors can be moved from place-to-place, but the sensor platform needs to remain stationary during sample collection. While sampling locations and intervals are typically defined by the user, the sample collection itself is often automated into the sensor platform once the vehicle has stopped. Stop-and-go platforms usually have a smaller sampling density compared to on-the-go systems but are greater than conventional sampling methods or point-based proximal sensors. Multiple sensors may be incorporated onto a single stop-and-go sensor platform for the collection of several measurements at once, and modular systems which enable the customization of sensor platforms are increasingly being adopted.

The uptake of single-vendor, modular sensing systems is increasing in applied and commercial settings. A range of on-the-go, stop-and-go, and point-based proximal soil sensors are commercially available and marketed by Veris® Technologies (Salina, USA; https://veristech.com/). Depending on the platform and its configuration, modular sensing systems like Veris® integrate a range of proximal sensors, such as EMI, infrared spectrometers, and penetrometers, to measure a range of soil properties in the field using a multi-sensor approach, including soil pH, EC, and soil organic matter. Sensors may be mounted on existing

farm equipment, embedded within a complete sensing system, or be individual sensors that can be used in combination within a customized setup. For all Veris® sensors, sample uptake, analysis, and sensing is automated on-the-go or within a stop-and-go system. Single vendor, modular systems are widely used in research and commercial applications as they require little-to-no user calibration and configuration, provide a high-density of information to map soil properties, and are relatively easy to use and interpret.

Other automated stop-and-go proximal sensing systems are incorporating spectrometers into multi-sensor platforms. The Soil Condition Analysis System (SCANS) is an automated soil core sensing system which measures a range of soil properties (e.g., SOC, bulk density, pH, available water capacity) of an intact soil core under field condition using a complementary set of sensors covering different portions of the EM spectrum (i.e., a gamma-ray attenuation densitometer, digital camera, and a vis-NIR spectrometer) (Viscarra Rossel *et al*. 2017). For field deployment, soil cores are extracted and loaded manually into the SCANS platform which is mounted within a specialized field trailer, and the sample is scanned automatically along the core (Figure B5). The modularity of SCANS can also facilitate the integration of other sensors (e.g., portable XRF and mid-IR spectrometers) as they become available (Viscarra Rossel *et al*. 2017).

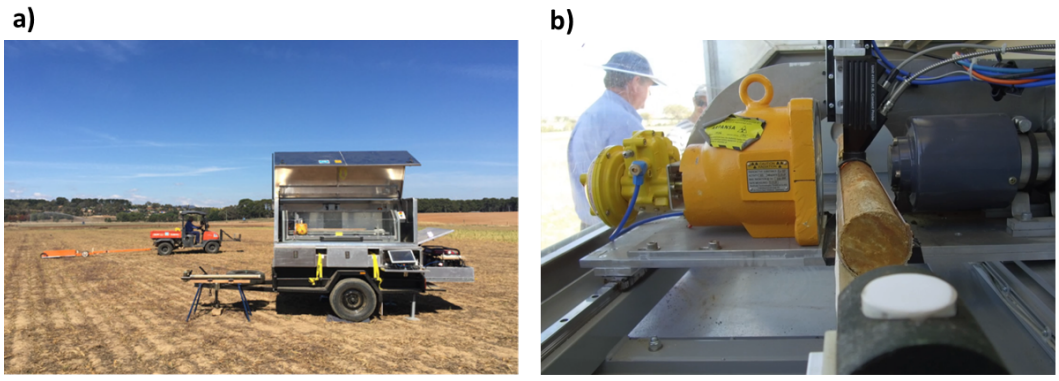

FIGURE B5. (a) The trailer-mounted SCANS automated proximal sensing system at a field site. (b) The soil core sensing system collecting data down the length of a soil core through a gamma densitometer, vis-NIR spectrometer, and high-resolution RGB camera.

Predicting soil properties

Calibration and spectral libraries

Many proximal sensors are indirect measures and parameter values cannot be directly deciphered from sensor signals or spectra. Instead, spectra need to be related to known reference samples, collectively known as a spectral library, through calibration of a prediction model (Stenberg *et al*. 2010). These reference samples must be representative of the range of soil types and conditions that the model and sensor is intended for. Commonly used calibration models include stepwise multiple linear regression, principal component analysis, partial least squares regression, and machine learning methods (e.g., regression trees). Also, additional pre-processing using mathematical functions is often used to correct for variation (e.g., normalizing spectra in respect to each other using pre-processing methods such as the Standard Normal Variate (SNV) method) and reduce noise (Stenberg *et al*. 2010).

Spectral libraries and calibration models enable predictions on unknown samples using only spectral data, saving both time and money compared to conventional laboratory assessments. However, the construction of a representative and diverse soil spectral library still requires conventional soil sampling and analysis which is both site- and time-specific. While fewer calibrations are better in terms of logistics, cost, and error, there is a trade-off between the accuracy of this calibration data and generalization capacity (Stenberg *et al*. 2010). Further, sensor signals respond to multiple factors that depend on environmental conditions (e.g., soil moisture) and diverse reference samples or libraries are needed to account or correct for this variability.

Multi-sensor data fusion

No single sensor can measure all soil properties, and sensors can respond to more than one soil or landscape attribute. Using a selection of complementary sensing techniques and integrating this data, through multi-sensor data fusion or via a multi-sensor approach, presents an opportunity to provide more information about soil properties (Adamchuk *et al*. 2011).

Multi-sensor data fusion, or the use of two-or-more sensors and their subsequent data, can improve the quality and diversity of information that would otherwise not be available, or not be as accurate, from individual sensors (Adamchuk *et al*. 2011). Spectral reference libraries and calibration models are used to define relationships between sensor signals and soil attributes, including multivariate models or machine learning algorithms (Adamchuk *et al*. 2018). Different combinations of sensor inputs and models may be necessary depending on the target soil property, including the concatenation of spectra or model averaging of predictions from sensors across different areas of the EM spectrum (Pozza *et al*. 2020).

Multi-sensor approaches involve the use of multiple sensors in tandem for unique applications and output. Multi-sensor approaches can increase the robustness and validity of measurements, capture a broader range of soil attributes, and increase the density of data collected (Adamchuk *et al*. 2011) by using multiple sensors across the same area at the same time, often in a single pass for on-the-go or stop-and-go sensors. Within these multi-sensor platforms, sensor configuration may be complementary (i.e., information from two-or-more sensors can give a more complete image of the environment when used together), competitive (i.e., multiple sensors deliver independent measurements of the same property, often used to increase robustness or reduce error), or cooperative (i.e., use two-or-more sensors to derive information that would otherwise be unavailable from a single sensor as part of multi-sensor data fusion) (Durrant-Whyte 1990; Mitchell 2012). Modular multi-sensor systems like Veris® (refer to section on B) are an example of a multi-sensor approach that is used in commercial settings to gain a more complete understanding of the soil landscape.

Future directions of proximal sensing

In practice, proximal soil sensors are not a standalone tool. Instead, they form part of a complex decision-making process that draws information from a range of sources (Bellingham 2011). Data integration from across multiple PSS systems and other on-farm (e.g., yield monitor data), remote sensing (e.g., LANDSAT, MODIS), and publicly available (e.g., ELVIS—Elevation and Depth—Foundation Spatial Data) data platforms can provide users

with a thorough understanding of the soil landscape. Maps derived from PSS tools, such as soil gamma radiometric or EC_a maps, are often used as covariates for design-based soil sampling, DSM and spatial interpolation methods, alongside this additional auxiliary information. The acquisition, integration, and presentation of information from various sources can aid in the assessment of soils at a range of spatial and temporal scales, and can help to assess the uncertainty of proximal sensing predictions. Beyond this, PSS tools can also be used to guide future soil sampling strategies, such as delineating agricultural management zones or areas of land contamination (Bellingham 2011). However, further research on the complementary use of multi-sensor information and spatial (up and down) scaling of this data is needed (Viscarra Rossel *et al*. 2016).

For several soil attributes, appropriate proximal sensors remain in the research phase. However, greater commercialization potential is being realized as the cost, accuracy, and ease of use of these sensors continues to improve. To increase the uptake and application of PSS tools, better access to commercialized and integrated PSS systems and services is required (Viscarra Rossel *et al*. 2016). Further, integrating this information with other on-farm, industry, and publicly available data to enable meaningful interpretation needs to be supported. It should be emphasized that investment, research, and development into new and improved PSS tools should not be for the sole purpose of collecting more data. These tools should be used to collect information that is relevant to farmers, land managers, policy makers and other stakeholders that need data to support well-informed management and policy decisions. The translation of existing information into accessible tools for all end-users is an important first step to better understanding the soil landscape.

References

Adamchuk, V.I., Hummel, J.W., Morgan, M.T. and Upadhyaya, S.K. (2004). 'On-the-go soil sensors for precision agriculture'. In: *Computers and Electronics in Agriculture* 44.1, pp. 71–91. ISSN: 0168-1699. DOI: 10.1016/j.compag.2004.03.002.

URL: https://www.sciencedirect.com/science/article/pii/S0168169904000444.

Adamchuk, V.I., Ji, W., Viscarra Rossel, R.A., Gebbers, R. and Tremblay, N. (2018). 'Proximal soil and plant sensing'. In: *Precision agriculture basics*. Wiley Online Library, pp. 119–140. DOI: 10.2134/precisionagbasics.2016.0093.

Adamchuk, V.I., Viscarra Rossel, R.A., Sudduth, K.A. and Schulze Lammers, P. (2011). 'Sensor Fusion for Precision Agriculture'. In: *Sensor Fusion - Foundation and Applications*. Ed. by Thomas, Ciza. Rijeka: IntechOpen. Chap. 2. DOI: 10.5772/19983. URL: https://doi.org/10.5772/19983.

Bellingham, B.K. (2011). 'Proximal soil sensing'. In: *Vadose Zone Journal*. Ed. by Viscarra Rossel, R.A., McBratney, A.B. and Minasny, B. The Soil Science Society of America, Inc. DOI: 10.2136/vzj2011.0105br.

Corwin, D.L. (2008). 'Past, present, and future trends of soil electrical conductivity measurement using geophysical methods'. In: *Handbook of agricultural geophysics*. 1st ed. CRC Press, Taylor Francis Group: New York, NY, USA, pp. 17–44.

Desilets, D., Zreda, M. and Ferré, T.P.A. (2010). 'Nature's neutron probe: Land surface hydrology at an elusive scale with cosmic rays'. In: *Water Resources Research* 46.11. DOI: 10.1029/2009WR008726.

Doolittle, J.A. and Brevik, E.C. (2014). 'The use of electromagnetic induction techniques in soils studies'. In: *Geoderma* 223–225, pp. 33–45. ISSN: 0016-7061. DOI: 10.1016/j.geoderma.2014.01.027. URL: https://www.sciencedirect.com/science/article/pii/S0016706114000548.

Durrant-Whyte, H.F. (1990). 'Sensor Models and Multisensor Integration'. In: *Autonomous Robot Vehicles*. Ed. by Cox, I.J. and Wilfong, G.T. New York, NY: Springer New York, pp. 73–89. ISBN: 978-1-4613-8997-2. DOI: 10.1007/978-1-4613-8997-2_7. URL: https://doi.org/10.1007/978-1-4613-8997-2_7.

Gebbers, R. (2019). 'Proximal soil surveying and monitoring techniques'. In: *Precision agriculture for sustainability*. Ed. by Stafford, J. Burleigh Dodds Science Publishing, pp. 49–98. DOI: 10.1201/9781351114592.

Hardie, M. (2020). 'Review of Novel and Emerging Proximal Soil Moisture Sensors for Use in Agriculture'. In: *Sensors* 20.23. ISSN: 1424-8220. DOI: 10.3390/s20236934. URL: https://www.mdpi.com/1424-8220/20/23/6934.

Knight, R. (2001). 'Ground penetrating radar for environmental applications'. In: *Annual Review of Earth and Planetary Sciences* 29.1, pp. 229–255.

Kuang, B., Mahmood, H.S., Quraishi, M.Z., Hoogmoed, W.B., Mouazen, A.M. and Henten, E.J. van (2012). 'Chapter four - Sensing Soil Properties in the Laboratory, In Situ, and On-Line: A Review'. In: *Advances in Agronomy*. Ed. by Sparks, D.L. Vol. 114. Advances in Agronomy. Academic Press, pp. 155–223. DOI: 10.1016/B978-0-12-394275-3.00003-1. URL: https://www.sciencedirect.com/science/article/pii/B9780123942753000031.

McJannet, D., Hawdon, A., Baker, B., Renzullo, L. and Searle, R. (2017). 'Multiscale soil moisture estimates using static and roving cosmic-ray soil moisture sensors'. In: *Hydrology and Earth System Sciences* 21.12, pp. 6049–6067. DOI: 10.5194/hess-21-6049-2017. URL: https://hess.copernicus.org/articles/21/6049/2017/.

Minasny, B., McBratney, A.B., Bellon-Maurel, V., Roger, J.M., Gobrecht, A., Ferrand, L. and Joalland, S. (2011). 'Removing the effect of soil moisture from NIR diffuse reflectance spectra for the prediction of soil organic carbon'. In: *Geoderma* 167–168, pp. 118–124. ISSN: 0016-7061. DOI: 10.1016/j.geoderma.2011.09.008. URL: https://www.sciencedirect.com/science/article/pii/S0016706111002758.

Mitchell, H.B. (2012). *Data fusion: concepts and ideas*. 2nd ed. Springer Science & Business Media. DOI: 10.1007/978-3-642-27222-6.

Pätzold, S., Leenen, M. and Heggemann, T.W. (2020). 'Proximal Mobile Gamma Spectrometry as Tool for Precision Farming and Field Experimentation'. In: *Soil Systems* 4.2. ISSN: 2571-8789. DOI: 10.3390/soilsystems4020031. URL: https://www.mdpi.com/2571-8789/4/2/31.

Pozza, L.E., Bishop, T.F.A., Stockmann, U. and Birch, G.F. (2020). 'Integration of vis-NIR and pXRF spectroscopy for rapid measurement of soil lead concentrations'. In: *Soil Research* 58.3, pp. 247–257. DOI: 10.1071/SR19174.

Reinhardt, N. and Herrmann, L. (2019). 'Gamma-ray spectrometry as versatile tool in soil science: A critical review'. In: *Journal of Plant Nutrition and Soil Science* 182.1, pp. 9–27. DOI: 10.1002/jpln.201700447.

Soriano-Disla, J.M., Janik, L.J., Viscarra Rossel, R.A., Macdonald, L.M. and McLaughlin, M.J. (2014). 'The Performance of Visible, Near-, and Mid-Infrared Reflectance Spectroscopy for Prediction of Soil Physical, Chemical, and Biological Properties'. In: *Applied Spectroscopy Reviews* 49.2, pp. 139–186. DOI: 10.1080/05704928.2013.811081. eprint: https://doi.org/10.1080/05704928.2013.811081. URL: https://doi.org/10.1080/05704928.2013.811081.

Stenberg, B. and Viscarra Rossel, R.A. (2010). 'Diffuse Reflectance Spectroscopy for High-Resolution Soil Sensing'. In: *Proximal Soil Sensing*. Ed. by Viscarra Rossel, R.A., McBratney, A.B. and Minasny, B. Dordrecht: Springer Netherlands, pp. 29–47. ISBN: 978-90-481-8859-8. DOI: 10.1007/978-90-481-8859-8_3. URL: https://doi.org/10.1007/978-90-481-8859-8_3.

Stockmann, U., Jones, E.J., Odeh, I.O.A. and McBratney, A.B. (2018). 'Pedometric Treatment of Soil Attributes'. In: *Pedometrics*. Ed. by McBratney, A.B., Minasny, B. and Stockmann, U. Cham: Springer International Publishing, pp. 115–153. ISBN: 978-3-319-63439-5. DOI: 10.1007/978-3-319-63439-5_5. URL: https://doi.org/10.1007/978-3-319-63439-5_5.

Sudduth, K.A., Kitchen, N.R., Bollero, G.A., Bullock, D.G. and Wiebold, W.J. (2003). 'Comparison of electromagnetic induction and direct sensing of soil electrical conductivity'. In: *Agronomy Journal* 95.3, pp. 472–482. DOI: 10.2134/agronj2003.4720.

Tang, Y., Jones, E.J. and Minasny, B. (2020). 'Evaluating low-cost portable near infrared sensors for rapid analysis of soils from South Eastern Australia'. In: *Geoderma Regional* 20, e00240. ISSN: 2352-0094. DOI: 10.1016/j.geodrs.2019.e00240. URL: https://www.sciencedirect.com/science/article/pii/S2352009419302391.

Thomsen, A., Schelde, K., Drøscher, P. and Steffensen, F. (2007). 'Mobile TDR for geo-referenced measurement of soil water content and electrical conductivity'. In: *Precision Agriculture* 8, pp. 213–223. DOI: 10.1007/s11119-007-9041-1.

Viscarra Rossel, R.A., Adamchuk, V.I., Sudduth, K.A., McKenzie, N.J. and Lobsey, C. (2011). 'Chapter Five - Proximal Soil Sensing: An Effective Approach for Soil Measurements in Space and Time'. In: *Advances in Agronomy* 113. Ed. by Sparks, D.L., pp. 243–291. ISSN: 0065-2113. DOI: 10.1016/B978-0-12-386473-4.00005-1. URL: https://www.sciencedirect.com/science/article/pii/B9780123864734000051.

Viscarra Rossel, R.A. and Lobsey, C. (2016). *Scoping review of proximal soil sensors for grain growing*. Tech. rep. CSIRO. DOI: 10.4225/08/5953fcda5ab78.

Viscarra Rossel, R.A., Lobsey, C.R., Sharman, C., Flick, P. and McLachlan, G. (2017). 'Novel Proximal Sensing for Monitoring Soil Organic C Stocks and Condition'. In: *Environmental Science & Technology* 51.10, pp. 5630–5641. DOI: 10.1021/acs.est.7b00889.

Viscarra Rossel, R.A. and McBratney, A.B. (1998). 'Laboratory evaluation of a proximal sensing technique for simultaneous measurement of soil clay and water content'. In: *Geoderma* 85.1, pp. 19–39. DOI: 10.1016/S0016-7061(98)00023-8.

Viscarra Rossel, R.A., McBratney, A.B. and Minasny, B. (2010). *Proximal soil sensing*. Ed. by Viscarra Rossel, R.A., McBratney, A.B. and Minasny, B. Springer Science & Business Media. DOI: 10.1007/978-90-481-8859-8.

Weindorf, David C., Bakr, Noura and Zhu, Yuanda (2014). 'Chapter One – Advances in Portable X-ray Fluorescence (PXRF) for Environmental, Pedological, and Agronomic Applications'. In: *Advances in Agronomy* 128. Ed. by Sparks, D.L., pp. 1–45. ISSN: 0065-2113. DOI: 10.1007/s11119-013-9338-1. URL: https://www.sciencedirect.com/science/article/pii/B9780128021392000019.

Whelan, B.M. and Taylor, J.A. (2013). *Precision agriculture for grain production systems*. CSIRO publishing.

Zreda, M., Desilets, D., Ferré, T.P.A. and Scott, R.L. (2008). 'Measuring soil moisture content non-invasively at intermediate spatial scale using cosmic-ray neutrons'. In: *Geophysical research letters* 35.21. DOI: 10.1029/2008GL035655.

Zreda, M., Shuttleworth, W.J., Zeng, X., Zweck, C., Desilets, D., Franz, T. and Rosolem, R. (2012). 'COSMOS: the COsmic-ray Soil Moisture Observing System'. In: *Hydrology and Earth System Sciences* 16.11, pp. 4079–4099. DOI: 10.5194/hess-16-4079-2012. URL: https://hess.copernicus.org/articles/16/4079/2012/.