Australian Government
Wine Australia

CSIRO

New digital technologies for managing winegrape production and quality



FINAL REPORT TO WINE AUSTRALIA

Project Number: **CSA 1701-3.1**

Principal Investigators:
Everard J Edwards & Giacomo Betti

Research Organisation: **CSIRO Agriculture & Food**

Date: **30 June 2022**

New digital technologies for managing winegrape production and quality

CSA 1701 - 3.1

April 2020 to June 2022

Giacomo Betti, Ryan Lagerstrom, Luke Powell, Chad Hargrave, Rizwan Khokher, Qiyu Liao, Dadong Wang, Dan Smith, David Ahmedt Aristizabal, Stephen Gensemer, & Everard J Edwards

CSIRO
Locked Bag 2
Glen Osmond
SA 5064

COPYRIGHT

© 2023 CSIRO. To the extent permitted by law, all rights are reserved and no part of this publication covered by copyright may be reproduced or copied in any form or by any means except with the written permission of CSIRO.

DISCLAIMER

CSIRO Agriculture & Food advises that the information contained in this publication comprises general statements based on scientific research. The reader is advised and needs to be aware that such information may be incomplete or unable to be used in any specific situation. No reliance or actions must therefore be made on that information without seeking prior expert professional, scientific and technical advice.

To the extent permitted by law, CSIRO (including its employees and consultants) excludes all liability to any person for any consequences, including but not limited to all losses, damages, costs, expenses and any other compensation, arising directly or indirectly from using this publication (in part or in whole) and any information or material contained in it.

Contents

Abstract **1**

Executive summary **3**

Chapter 1 ..... **5**

1.1 Project Background 5

1.2 Project Aims and Performance targets 5

Chapter 2 - Yield estimation **7**

2.1 By machine vision 7

2.2 By RADAR 15

Chapter 3 – Vine water status monitoring **31**

3.1 Introduction 31

3.2 System overview 31

3.3 Data collection and field data collection 34

3.4 Conclusion 46

Chapter 4 – Non-contact determination of fruit composition **47**

4.1 Introduction 47

4.2 Laboratory work 47

4.3 Field data collection 51

4.4 Data analysis 54

4.5 Case study: in-field assessment of fruit maturity at Barossa Valley Estate 66

4.6 Conclusions 71

Chapter 5 – Estimation and mapping of canopy density and cover **73**

5.1 Background 73

5.2 Visualisation mapping of canopy data 73

5.3 Industry demonstration 74

Chapter 6 – Potential for commercialisation **79**

6.1 – Machine vision for yield estimation 79

6.2 - Radar for yield estimation 80

6.3 - Canopy temperature for vine water status monitoring 80

6.4 - NIRS for fruit and vine tissue composition 80

6.5 - 3D spinning LiDAR SLAM for canopy structure 81

Appendix 1: Communication ..... **i**

Refereed papers i

Papers under review i

Conference presentations i

Datasets i

Appendix 2: Intellectual property ..... **iii**

Appendix 3: References ..... **v**
Appendix 4 ..... **vii**
Project staff vii
Acknowledgements viii
Appendix 5: Other material ..... **ix**
Appendix 5.1: schematics of complete ML processing pipeline for yield estimation ix
Appendix 5.2: Planned Project Outputs xii
Appendix 6: Budget reconciliation ..... **xvi**

Abstract

Four areas of viticultural management were addressed through this project, which further developed technologies from predecessor projects CSA 1601 and CSA 1602.

- The viability of using machine vision for hands-off yield estimation at scale in a range of viticultural situations was clearly demonstrated with an accuracy >90%, but will require significant computational efficiencies to be commercially practicable. In contrast, utilising a radar approach to assess yield by 'seeing' through the canopy was not successful, despite examining a wide range of frequencies and methods.
- The 'ProxiCrop' prototype low-cost static sensor for irrigation management support was refined and tested, and proved to be capable of monitoring canopy temperature and subsequent estimation of vine water status.
- On-the-go field estimation and mapping of fruit maturation and composition was demonstrated using an NIR spectral sensor and calibration libraries that were extended and further validated.
- The use of dual-axis spinning LiDAR with CSIRO SLAM and 'raycloudtools' libraries to map canopy structure in three-dimensions was finalised and demonstrated in a trial with commercial growers.

This page intentionally left blank.

Executive summary

CSIRO undertook two projects with Wine Australia as part of the Rural R&D for Profit scheme, which ended in 2020. CSA 1601 developed 'New technologies for dynamic canopy and disease management' and CSA 1602 developed 'New non-destructive technologies for simultaneous yield, crop condition and quality estimation'. The work described in this report continued the testing and development of the technologies utilised in those projects. This encompassed five potential solutions to four industry problems, all five aimed at in-field use with four being on-the-go.

Two technologies addressed the need for accurate yield estimation. This is a key part of horticultural production, including winegrapes, with supply chains and fruit value significantly impacted when production volume does not match expectations. Current yield estimates rely significantly on the experience of the vineyard manager, are labour intensive and typically have an error rate of up to 30%. Yield can be split into three components, berry weight, berry number per bunch and bunch number per vine/vineyard area. Berry weight is only fixed at harvest, whereas berry number is determined by fruit-set early in the season and bunch number by budburst, at the start of the season.

Machine learning methods trialled in CSA 1602 for assessing *yield potential* by counting inflorescences in video data post-budburst were further refined, proving highly accurate (approximately 5% error) without any manual adjustment or additional calibration. Methods for estimating berry number and size were explored and data sets collected, but the capability available for pre-harvest *yield estimation* was primarily directed at estimating bunch mass from video in combination with bunch counts. This was undertaken in collaboration with the CSIRO Machine Learning and Artificial Intelligence Future Science Platform and produced a complete data pipeline demonstration for the determination of fruit yield from video data, again without manual adjustment or additional data collection.

Radar has the potential to 'see through' canopy to detect bunches that are occluded by leaves based on wood, fruit and leaves having a different radar cross section. Initial results from testing in CSA 1602 demonstrated differences in the radar returns depending on whether fruit were present in a canopy and between vine components. However, detailed testing in the current project of both transmission and reflection modes across a wide range of frequencies did not suggest that it is possible to obtain a quantitative relationship with fruit mass due to the scatter of signal within the canopy.

The second industry problem addressed was in-field determination of fruit maturity and quality. More than any other perennial crop, the commercial potential of winegrapes depends on composition. If objective assessment of fruit quality can be done in the field prior to harvest and transport, harvest timing and winery planning can be optimised and a number of potential points of friction between grower and winery avoided.

As part of CSA 1602, a lab based spectral classification tool was developed to allow grapevine parts (fruit, leaves, canes, other) to be identified using only their reflectance spectra. In addition, a calibration for prediction of key fruit composition parameters from their reflectance spectra was developed using four winegrape varieties. The current project expanded the calibration out to eight varieties, which allowed model validation by successfully predicting parameters for varieties excluded from the model. The field spectrometer-based instrument was further developed, and fitted with a lighting system that allowed it to be used in full daylight. Finally, a data processing method was trialled to allow the lab-based calibration to be used with the field instrument. The combination of these tools successfully predicted fruit sugar maturity when used on a continuously moving farm vehicle.

The third problem addressed was objective measures of canopy structure to allow consistent and repeatable canopy management. As part of CSA 1601, a dual-axis spinning LiDAR was successfully used with CSIRO's Simultaneous Localisation and Mapping (SLAM) tools and a novel technique using the 'ray-cloud' rather than just returned laser pulses to map canopy density in 3D. Through the current project this technique was packaged into an open-source library and released by CSIRO. In addition, a demonstration event was held in the Coonawarra, where two vineyards were scanned with the LiDAR system, the data processed overnight, and canopy density maps provided to growers at an information session the following morning.

Finally, in recent years there has been significant interest in using plant-based sensing in addition to soil moisture measurements for irrigation management. A component of Wine Australia-CSIRO project CSA 1701-1.5 developed such a sensor, with the aim of quantifying vine water status, based on canopy temperature. In contrast to existing solutions, which have a single point measurement of a few leaves or less, the sensor used a low cost thermal imaging sensor, fused with RGB imaging, to obtain canopy temperatures of large areas (approximately one panel of vines). This work was moved into the current project, with the sensor tested under commercial field conditions and developed accordingly. The resulting canopy temperature data were well correlated with direct measurements of vine water stress.

In conclusion, four of the five technologies further developed through this project show potential for use in commercial growing systems, although the technology readiness of these varies.

Chapter 1

1.1 Project Background

This project extended two prior projects (CSA 1601 and CSA 1602) for a little over two years to continue the technology development that those projects had begun. The prior projects trialled a range of new technologies to assess their potential to solve a number of important industry problems in the field; namely, early estimation of yield, crop condition and fruit composition, measurement tools to assist canopy management, early disease detection and monitoring of canopy condition. The subsequent step in this work was to take the most promising technologies and develop them to the point where they are practical and user friendly for adoption by industry.

The Wine Australia gap analysis of digital viticulture undertaken before the prior projects identified a range of priorities, foremost of which were 1) yield, crop condition and quality estimation and 2) canopy condition monitoring. The project reported here directly addressed these gaps, taking the results of our prior work and developing prototype systems that can be commercially assessed for future production.

Such technologies will assist the production of fruit to specification and maximising vineyard productivity by providing the industry with accurate, whole of vineyard assessments that reduce the current 30% error in yield estimation, accurate fruit condition assessments in the field rather than at the winery intake, and the detailed spatial information needed for accurate adjustments of canopy size and structure within-block. As a result of the final assessment of CSA 1601, it was decided not to pursue early disease detection in the current project.

In September 2020, a component of the work from CSA 1701-1.5, the development of a static canopy temperature sensor, was moved into the current project and the development and application of the sensor from this date is reported here.

1.2 Project Aims and Performance targets

The project objectives were:

- Develop the technologies trialled in CSA 1602 that are assessed as effective for estimation of yield and crop condition to the point of a functioning prototype unit, including data analytics software.
- Develop the technologies trialled in CSA 1602 that are assessed as effective for estimation of canopy structure and condition to the point of a functioning prototype unit, including data analytics software.

The project aims were:

3 year (2022):

- Prototype sensor system for at least one viticultural factor available for demonstration and for testing by identified industry commercial entities.
- Automated data analytics for at least one sensor system running on CSIRO servers and developed to a point where it could be transferred to a cloud computer commercial service to deliver results in an easy to understand format to assist industry decision-making.

10 year (aspirational outcomes):

Automated sensor systems, together with cloud-based data analytics, in commercial production and being used by growers for yield estimation, canopy management decision making, assessment of crop condition/composition and early disease detection.

Chapter 2 - Yield estimation

Yield prediction and estimation have long been a key target for viticultural research and high priorities for the wine industry. At present, yield estimation is typically a subjective assessment with very small sampling sizes compared to vineyard size. This commonly results in errors of up to 30%, resulting in an under or over supply to the winery. Knowing the yield accurately ahead of harvest allows a winery to better plan intake and adjust for over or under supply of a variety. Currently, yield estimation costs the industry around $16/ha and the high error rates enlarge this cost. A reduction in the yield estimate error from 33% to 20% would amount to a saving of $85 million per year (Clingeleffer et al. 2001). Further reductions would magnify the savings.

Predecessor project CSA 1602 took two approaches to developing yield estimation at scale. The first was by machine vision, with the development of this approach significantly expanded through the current project, and the second was using radar, with new methods tested here.

2.1 By machine vision

Grapevine yield can be split into three components as follows:

```
yield = bunch weight*bunch number
```

where:

```
bunch weight = berry weight*berry number per bunch
```

therefore:

```
yield = berry weight*berry number per bunch*bunch number
```

Bunch number accounts for approximately 60% of inter-seasonal yield variation (Clingeleffer et al. 1997), berry number for approximately 20-30% of inter-seasonal variation and berry weight approximately 10-20%. Bunch number is limited by the number of inflorescences produced by the vine, which in most production systems are limited to those in the primary shoot. The primordia for these are established in the previous season. Consequently, an early prediction of *yield potential* can made by counting the number of inflorescences produced at budburst.

The yield potential can be improved by estimating fruit set (and counting the bunches if thinning or an adverse event reduced the number of inflorescences) to generate a *yield prediction*. This prediction can be further improved to provide a *yield estimate* by determining average berry weight prior to harvest. An alternative method to generate a yield estimate is to count the bunches prior to harvest and make an estimate of average bunch weight at the same time.

The work built on the machine vision tools developed in CSA 1602 to improve and validate the method for yield potential and to develop a complete analysis pipeline for yield estimation. In addition, preliminary work was undertaken to estimate fruit set for yield predations.

2.1.1 Yield potential

The potential accuracy of estimating yield potential using inflorescence counts was determined by collecting manual counts of inflorescences and final yield in the field. Data were collected for each panel in a single row from three different vineyard blocks, each a different variety (Fiano, Shiraz and Mataro), over three seasons (2019/20 to 2021/22). The first season of this data was collected through project CSA 1602.

The proposed method relies on a strong relationship between the inflorescence number at E-L 12 and the bunch number at harvest. If such a relationship is not present, the yield potential calculation will not be useful. Across the three blocks and seasons there was indeed a strong relationship, with R^2 of this varying between 0.66 and 0.96, and a mean R^2 of 0.79 (Figure 2.1).

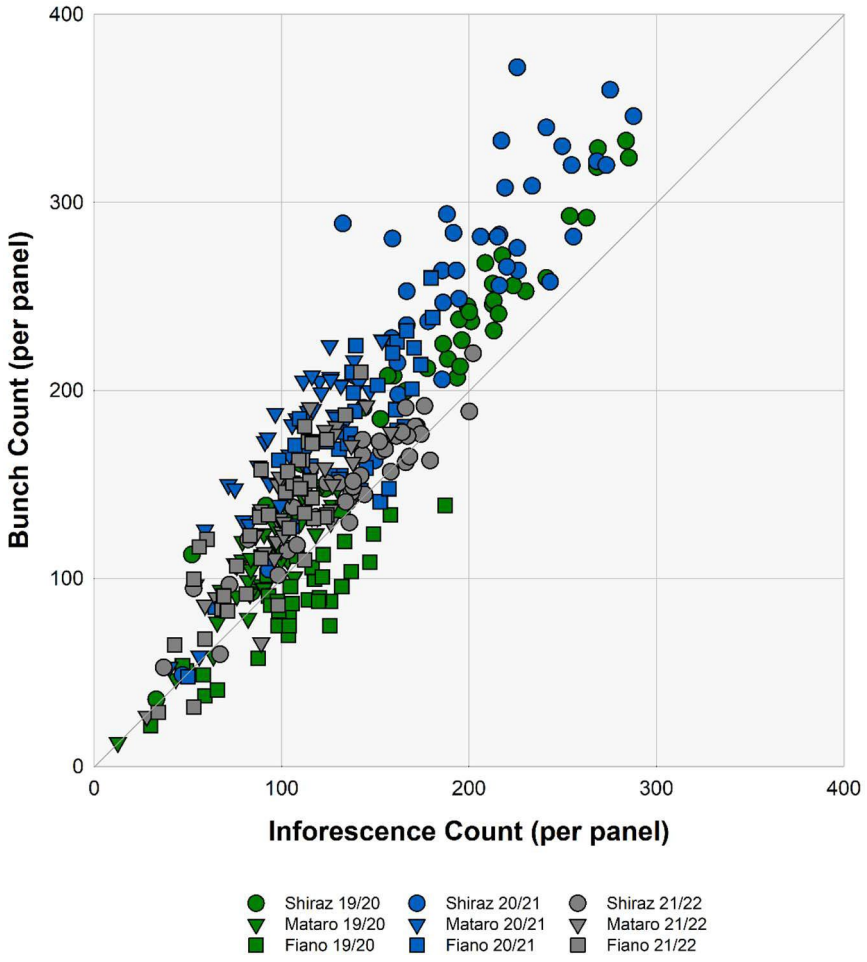


Figure 2.1 – relationship between inflorescence and bunch counts for three variety blocks across three growing seasons.

The data also demonstrated the scale of year-to-year variation in bunch number, with inflorescence and bunch counts in 2021/22 being almost half those of 2019/20 for example. The Fiano inflorescence counts in 2019/20 were higher than the bunch counts due to bunch thinning being carried out by the grower, in all other cases bunch counts were 10-15% higher than inflorescence counts, likely due to bunches on laterals formed later in the season or late shoots. The impact of timing of the count, relative to budburst, was examined in 2019/20 by doing hand counts of inflorescences in three blocks at McLaren Vale on three occasions, separated by approximately five days (only twice for Fiano). The results (Table 2.1) demonstrate that counting too early underestimates inflorescences and that hand counts are fairly stable after that, although visibility for imaging decreased.

Table 2.1 – mean number, standard deviation (SD), relative standard deviation (%RSD) and correlations between counts of inflorescences per panel, counted three times by hand over ten days.

	Shiraz			Mataro			Fiano		
	Count 1	Count 2	Count 3	Count 1	Count 2	Count 3	Count 1	Count 2	Count 3
Row mean #	129	175	185	58	84	91	n/a	108	106
SD	42	57	62	17	23	23	n/a	35	33
%RSD	32%	32%	34%	29%	27%	25%	n/a	32%	31%
r2 vs count 1	-	0.78	0.73	-	0.77	0.53	-	n/a	n/a
r2 vs count 2	-	-	**0.96**	-	-	**0.82**	-	-	**0.84**

Yield predictions were made per panel as:

$$Y_{pred} = C_{infl} \times \overline{\mu}_{BWt_h} \times \overline{\mu}_{C_{bun}/C_{infl}}$$

Where Y_{pred} is the prediction, C_{infl} is the inflorescence count, $\overline{\mu}_{BWt_h}$ is the multi-season mean bunch weight for the block, and $\overline{\mu}_{C_{bun}/C_{infl}}$ is the multi-season relationship between the inflorescence count at E-L 12 and the bunch count at harvest. The resulting predictions were compared to the harvest yield on both a panel and whole row (sum of all panels) basis (Figure 2.2). For the nine rows (three blocks x three seasons) this method predicted yield with an R^2 of 0.97. For the panels within a single row and season (n=39 in each case), R^2 values ranged from 0.07 to 0.83, with the lower values occurring when there was little panel to panel variation. This situation did not affect the estimation at scale (row), with the per panel data that produced the lowest R^2 still providing a per row prediction within 3% of actual yield. The mean % error between the predicted and actual yield across the nine datasets was 8%, whereas if yield was predicted using the three-year average yield for each block, the simplest form of yield prediction, the mean error was 28%. Clearly, inflorescence counts can be used to provide a better yield potential than simply using long-term yield averages.

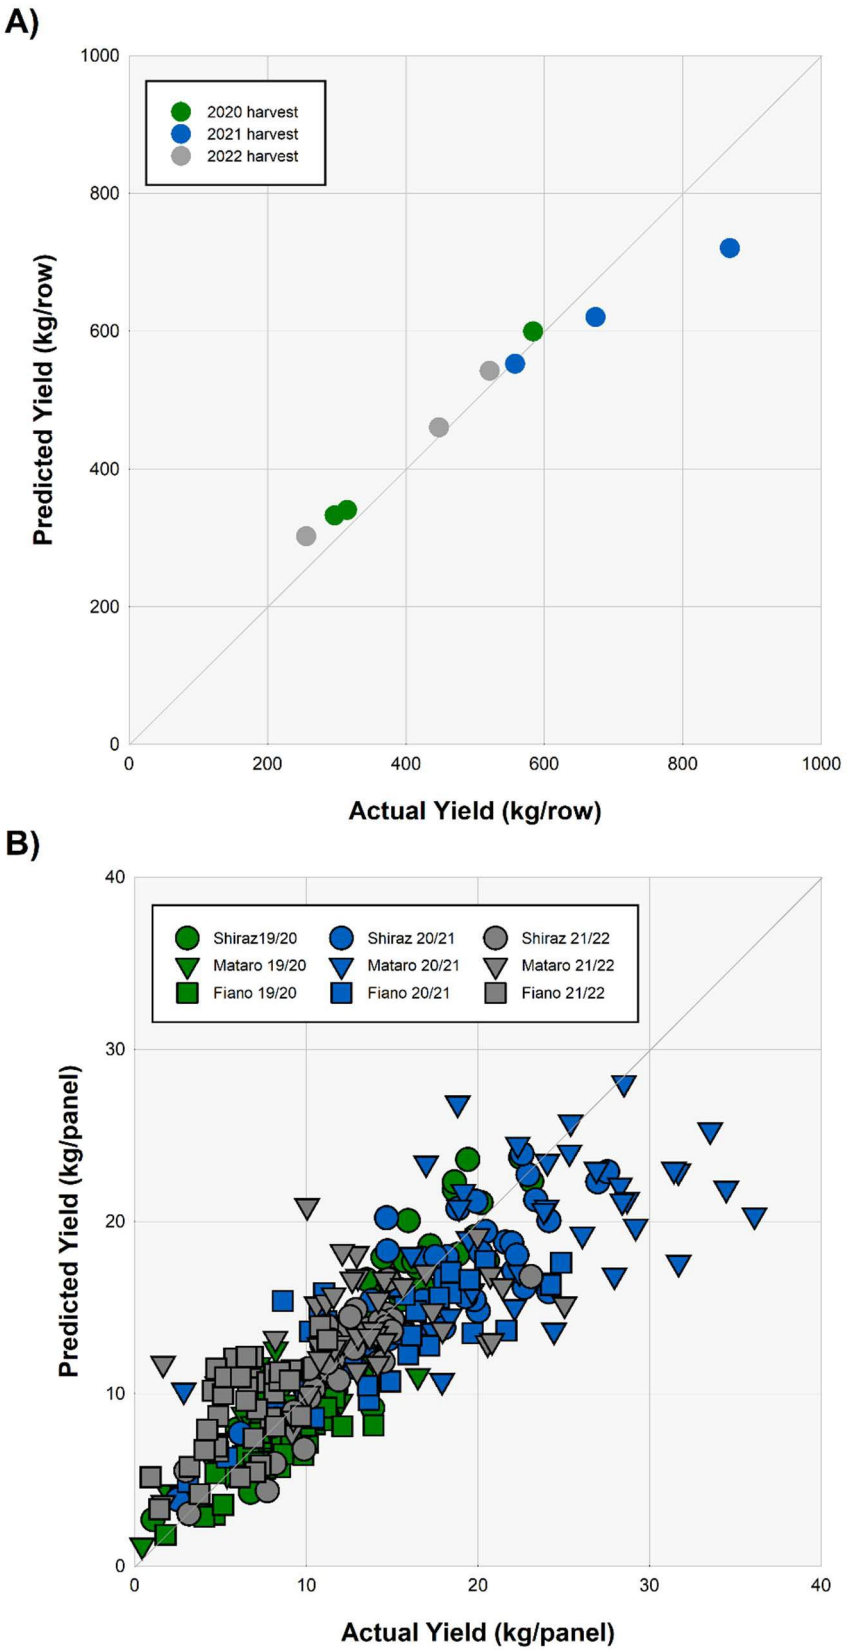


Figure 2.2 – relationship between actual yield and yield predicted from manual inflorescence counts and long-term mean bunch weight per row (A) and per panel (B) of three varieties over three seasons.

The detector and tracking algorithms utilised in CSA 1602 were further developed and improved during the current project, with particular attention applied to the tracking algorithm, which was critical to producing accurate counts from the input video data. Yield potential estimates were then generated for the same five rows of data described in the final report for CSA 1602 (Figure 2.3). At the Woodside vineyard, hand thinning of inflorescences post-budburst was employed by the grower, with the aim of reducing the number of inflorescences by 60%. For the rows at this vineyard an additional yield prediction was made, taking this into account, by simply multiplying the yield potential value by 0.4 (Figure 2.3). Historical bunch weights were not available for any of the rows utilised to demonstrate the yield potential algorithms, so data from the same season was used. For the McLaren Vale row, only a whole row harvest yield was available, but a bunch count was undertaken prior to harvest and an average bunch weight calculated. For the Woodside rows, an average bunch weight (and bunch count) across five rows was used.

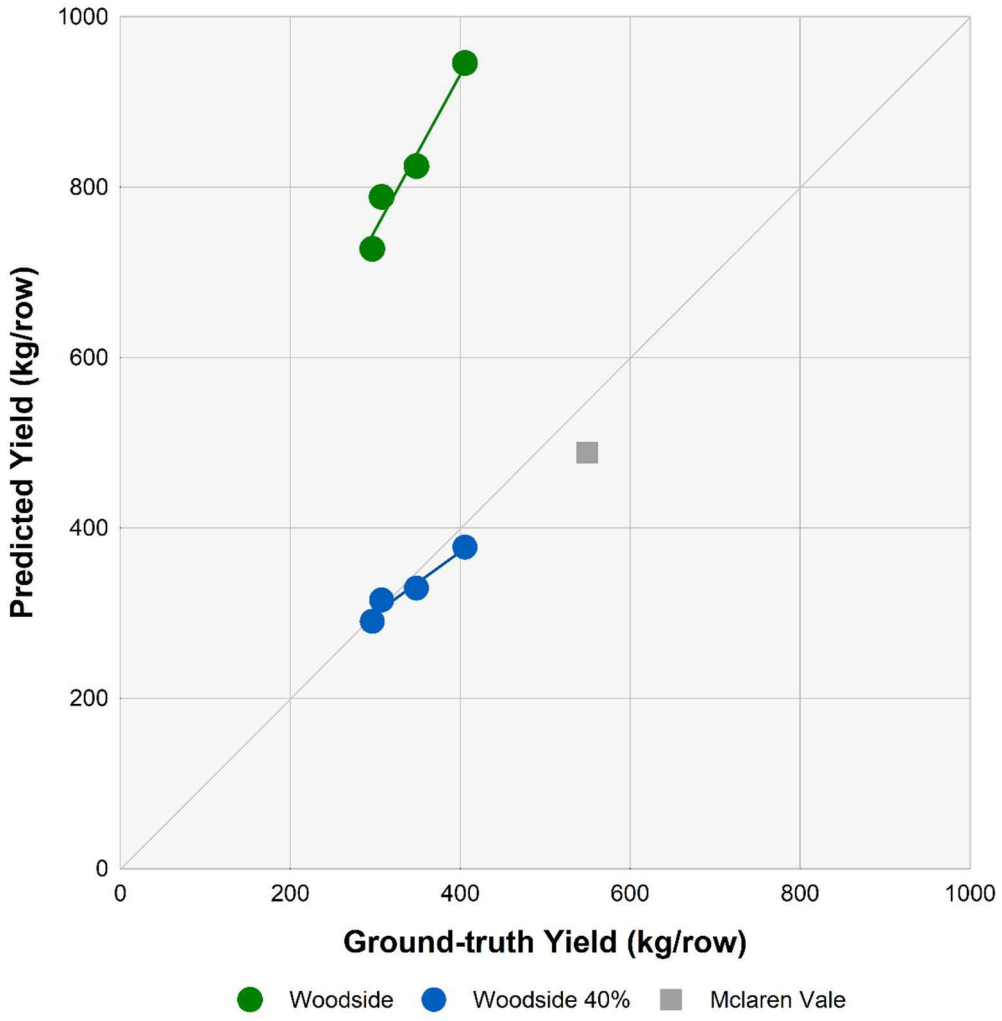


Figure 2.3 – relationship between the predicted yield from the machine vision approach and the actual yield at two sites, with and without an adjustment for manual thinning at Woodside.

The average error of the early yield prediction across the four Woodside rows, using the adjusted estimate, was 4% below harvest yield and the error of the McLaren Vale row was 11% below harvest yield. These values are probably better than would be achieved using historical bunch weight and count data, but still suggest that the machine vision method developed can be used to generate an accurate yield potential very early in the season.

2.1.2 Yield prediction

As described above, the accuracy of yield potential can be improved to give a prediction of yield following fruit-set. The wide remit of the work in this project meant that there were capacity constraints and, although video and still image data were collected in the 2021/22 season from the three McLaren Vale blocks described previously we were unable to undertake the image labelling and machine learning development required as a decision was made to concentrate this effort on the yield estimation component. The yield estimation work included development of berry counting algorithms and it is anticipated that these could be relatively easily adapted to work with bunches at an earlier stage of development.

2.1.3 Yield estimation

The principle used for pre-harvest yield estimation was similar to that used earlier in the season, except that the aim was to determine bunch size directly and thus mass (via an empirical calibration), rather than count bunches and use historical averages. Furthermore, in the majority of Australian vineyards there is significant occlusion of the fruit by the canopy, something that was able to be avoided with the early yield potential estimate.

To achieve this, a number of additional steps were required and the CSIRO Machine Learning and Artificial Intelligence Future Science Platform assisted the project by providing additional capability and know-how to build a proof-of-concept pipeline. In principle, the pipeline:

1. Detects each bunch in each frame of the video.
2. Segments each bunch from the image (identifies which pixels are bunch vs background/foreground).
3. Counts the number of bunches, by tracking each bunch through the video frames.
4. Estimates bunch mass using pre-determined relationship between bunch features and mass.

For each step, multiple methods were experimented with, including combining steps 1 & 2 in a single model. For step four, one of the bunch features examined was the use of berry number, a parameter that also necessitated experimentation with multiple machine learning models. Similarly, this step also required identification of bunches without significant occlusion from leaves and scaling of geometric bunch features (where used) according to distance from the camera, also requiring further model development.

A preliminary pipeline was used to establish the feasibility of this approach by using an Intel

RealSense D435i camera to simultaneously collect RGB video data and matching depth map. The video from this camera was a relatively low resolution and the depth map data can only be collected in an outdoor setting if moved very slowly, so it was only suitable for initial testing. Over 150 bunches were individually tagged to allow identification in the video and harvested and weighed after the video data was collected. Three frames per bunch were selected from the video data and the bunches segmented by adapting the bunch detector developed in the prior project CSA 1602. Figure 2.4 provides an example of a paired video frame and depth map together with two examples of frames with bunches segmented.

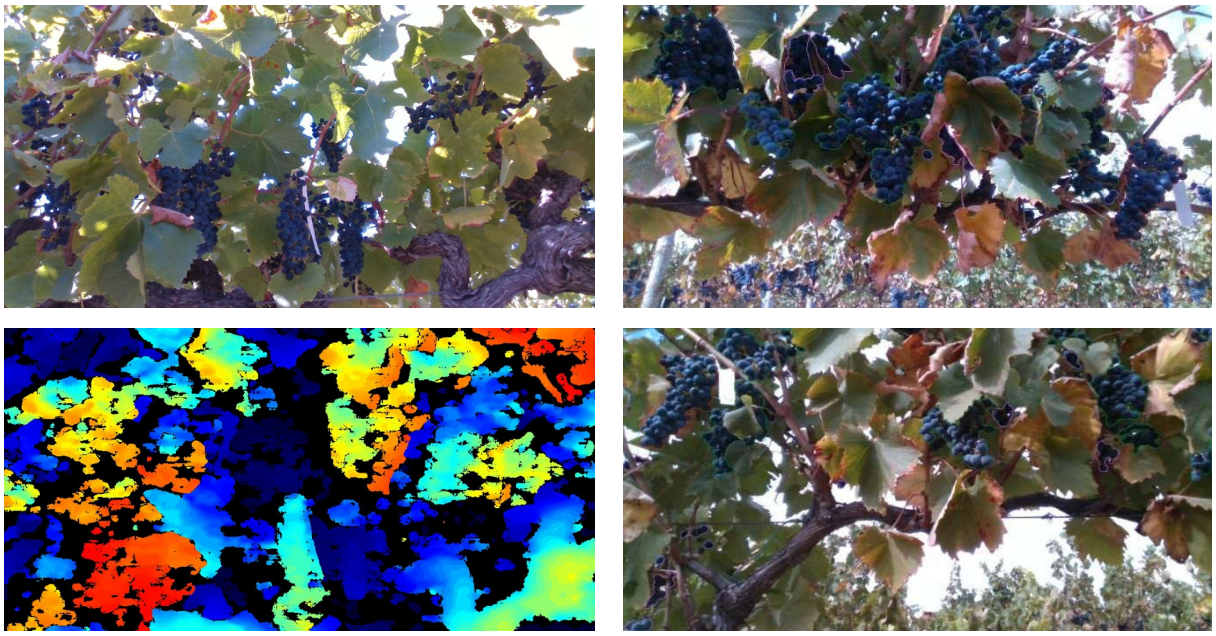


Figure 2.4 – two example frames with segmented bunches labelled (top) and an example frame with matching depth map (bottom).

For each segmented bunch, geometric features were defined (area, major axis, minor axis and perimeter) and a berry 'density' feature determined using the Otsu method (Figure 2.5). Four regression models were compared and validated to estimate bunch mass from the features for the individually weighed bunches. Figure 2.6 provides the output from the most successful of these and demonstrates the requirement to include the scaling by depth.

Finally, a bunch count was obtained for the entire row in which the labelled bunches were taken and the whole row harvested, coming in at 310 kg. A yield estimate for the row was made using the bunch count and the bunch weights distribution from the predicted weights (Figure 2.6 left), which produced an estimate of 444 kg. However, it was realised that the labelled bunches had been specifically selected to give an even spread of weights, whereas the natural bunch weight variation in the field was skewed (to smaller bunches), recalculating the estimate using the natural bunch weight distribution produced an estimate of 323 kg, within 5% of the actual harvest weight. Whilst this was a fairly artificial estimate that used

actual knowledge from the field, it demonstrated the feasibility of using bunch features from video images and the requirement to include distance in these.

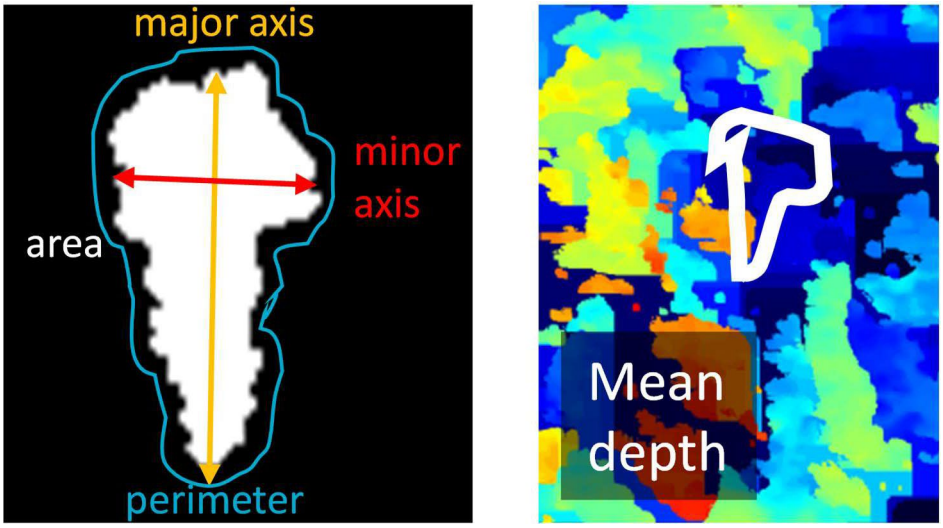


Figure 2.5 – example image showing the bunch features (left) and depth information (right) for a single bunch.

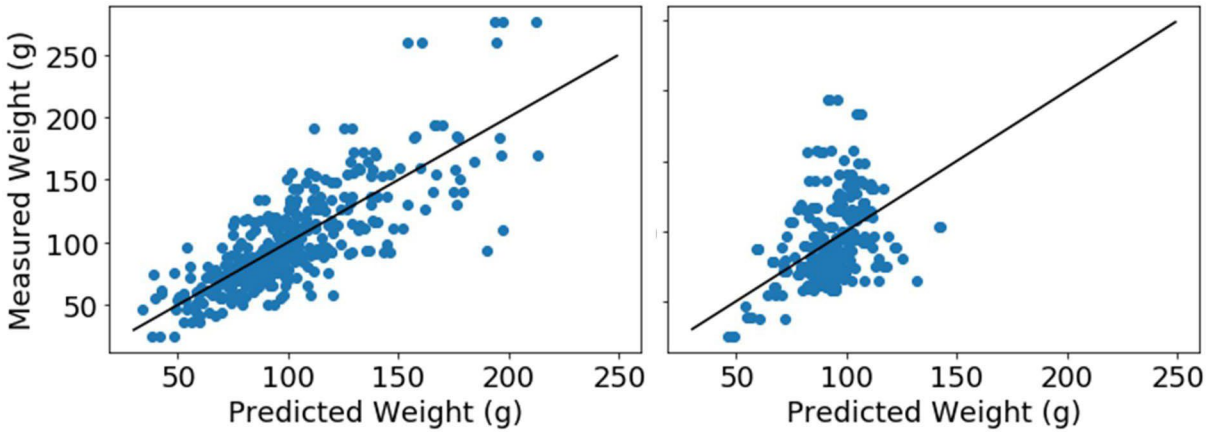


Figure 2.6 – the relationship actual bunch weight at harvest and weight predicted by the machine vision approach, taking into account the actual bunch weight distribution (left) or using only the distribution of the bunches selected for training (right).

The CSIRO MLAI FSP is continuing to work on the proof-of-concept pipeline at the time of writing. An update on this work is provided in appendix 5.1.

2.2 By RADAR

Although using digital images or video for yield estimation makes use of ubiquitous, low-cost sensors (cameras) visible light does not penetrate the canopy and thus occlusion of fruit when using RGB machine vision approaches will be a significant problem in the majority of Australian viticulture, thus necessitating work-arounds, such as those described above. The predecessor project, CSA 1602, demonstrated that longer wavelengths (e.g. microwaves) could penetrate the canopy and that reflected signals varied according to the presence of fruit. This work was continued, but no method to extract *quantitative* results from the sensors used in CSA 1602 was found. Consequently, during the 2021/22 season a comprehensive test system was set up to examine a much wider range of wavelengths and to look at both transmission and reflectance modes.

2.2.1 Data collection

In brief, a Keysight N9950A FieldFox Vector Network Analyser, with a range of 0.3 - 32 GHz was rented from ElectroRent (Penang, Malaysia), together with two Diamond Engineering DE0726 700MHz - 26GHz Broadband Reference Horn antenna (rented from The EMC Shop, Roseville, CA, USA) and set up using an aluminium frame to ensure all directions, angles and set-up was accurate and repeatable. All equipment and cabling were calibrated in the laboratory and in place in the field, as appropriate. A total of 15 test sets of data was collected from a vineyard in Lenswood (SA) over the course of three days. Each test set consisted of eights scans from 0.7-26 GHz with both transmission and reflection data, S21 and S11 respectively. As the propagation of the waves varied by frequency, with the higher frequencies having a much smaller field of view between antenna and receiver, the region of the vine row between the two antennas was broken up into three zones of interest, outer, middle and inner zones as depicted by the circles shown in Figure 2.7. A series of four scans was then undertaken, with progressive removal of fruit from each zone, starting with the outer (Table 2.2). The leaves were then removed from the area and the fruit manually hung from the vines. A second set of scans was then made, with fruit progressively removed as previously (Table 2.2). For each of the 15 sets of data, bunch counts and weights were recorded for each zone. Figure 2.8 provides images of the set-up and experimental procedure.

2.2.2 Data analysis approach

The objective of these tests was to examine as wide a range of frequencies as possibly in both reflectance (as in CSA 1602) and transmission modes to prove the best possible chance of finding a quantitative relationship between a signal property and fruit mass. This assumes that the objects (fruit, leaves, woody material, etc.) causing an effect on the scan energies are separable. In turn, this would mean that the energy signature and/or energy difference between scans with objects present and absent (e.g. fruit) was approximately all attributed

to the object. If the difference in energy either across the entire frequency range or a subset of frequencies and the volume of grapes has little correlation, this indicates that the energy contributions are not significantly separable. Further, for a commercially relevant on-the-go system, the collected scans would ideally provide an indication of the grape weight present in the scan without having to take reference scans in each position.



Figure 2.7 - diagrams to demonstrate the field of view of the radar system at three frequencies and the fruit removal steps to generate the scans detailed in Table 2.2.

Table 2.2 – presence or absence of fruit and foliage in the three zones for each of the scans a-h.

Scan – foliage present	Scan – foliage removed	Zones with grapes present
A	E	Outer, middle, inner
B	F	Middle, inner
C	G	Inner
D	H	No grapes present

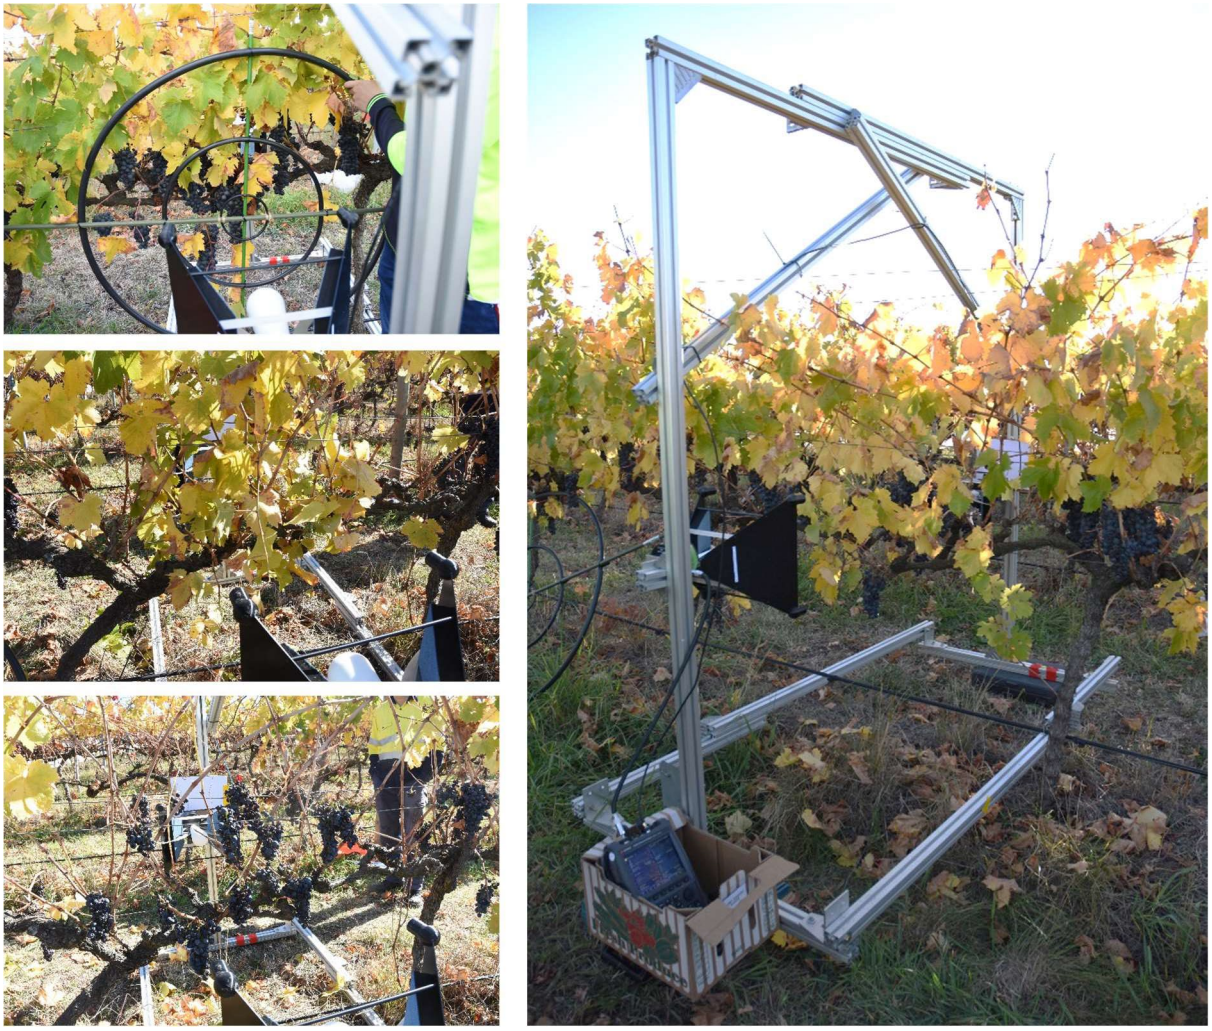


Figure 2.8 – photographs of the radar test setup and field of view zones.

For these tests, to assess the scan energy related to the grapes present in a scan, the scan was referenced to a series of different cases, namely: no grapes present; the previous scan (i.e. a change in grapes present in a single zone); open air; and raw data.

Referenced to the no grape case:

For each set, cases d and h, the scans where no grapes were present are subtracted from cases a-c and e-g respectively. Between the two scans, the data collection system position, foliage, and vines remain the same, the only difference is the grapes that were removed from the corresponding zones. The weights correspond to the zones listed in Table 2.2. This set of data is used to determine if the energy difference between two scans is correlated to the difference in grape weights.

Referenced to the previous case:

For each set, the case with one more zone removed is used as the reference case. Specifically, for the naturally hung grapes, case a-b, b-c and c-d are used, and the corresponding scans for the manually hung foliage removed cases undergo the same process. Likewise with referencing the no grape case, the only difference is the grapes removed between the two

scans. For test cases a and e, the outer zone weights will be taken into account, likewise test cases b and f will use the middle zone weights, and test cases c and g will use the inner zone weights.

Referenced to open air:

For each set, each case has the open air scan from the corresponding day used as the reference scan. Unlike the other cases, as the reference case is external to the four cases for each method, this means that test case d may also be used, and as such there is an extra plot. This results in plots of the test cases that are used to assess the correlation of the energy difference and the grape weights, and the ability to separate them from the foliage and vines.

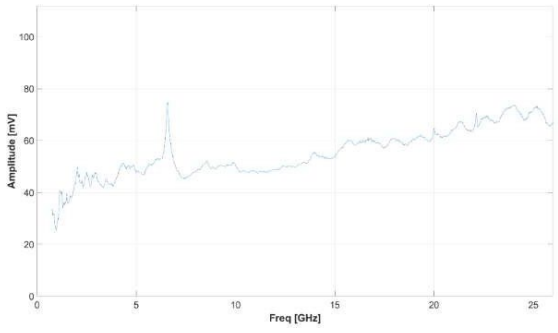
Raw data:

This set of data uses the raw data with no reference case subtracted from the test cases. This is similar to referenced to open air, in that there are four test cases per method, and the same factors are assessed for their energy contribution.

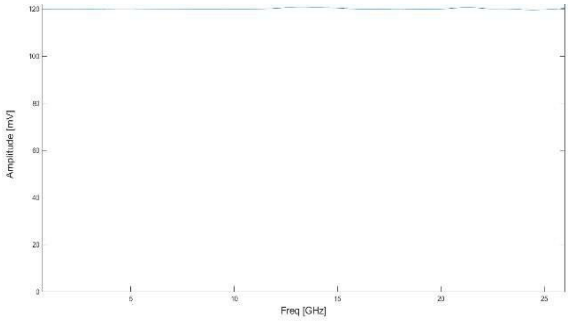
2.2.3 Radar scans

Atypical Scans:

Figure 2.9 shows two example scans, with Figure 2.9a showing a typical radar scan, and Figure 2.9b showing an atypical radar scan. The amplitude of the atypical radar scan is much larger than that of the typical scan, and the variations in the signal over the range of frequencies is significantly smaller in the atypical scan compared to the typical scan. Set 2 case a, and set 4 case c have been removed as they did not have a typical response, indicating something had gone wrong during the scan. For the later energy vs grape weight trend comparisons, the atypical test cases and any subsequent test cases that rely on these have been removed to eliminate their effect on the results. One concern that arises is that atypical scans need to be identified without relying on other scans or knowing the grape and vine content of the scan



(a) A typical radar scan of expected amplitude levels with variation over the range of frequencies.



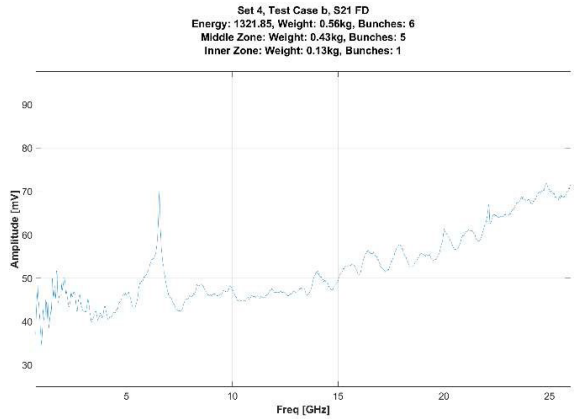
(b) An atypical radar scan of large amplitude with very little variation across the range of frequencies.

Figure 2.9: (a) Typical and (b) atypical radar scan examples.

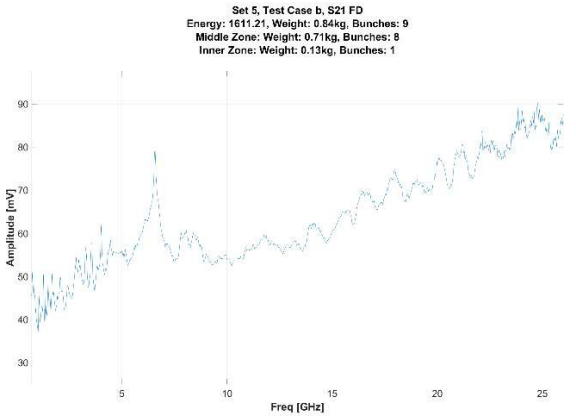
in advance, such that the scan may be retaken while the field data collection system is in position. The difference between the example typical and atypical scans is easily spotted and could be automatically identified in the field and have the user prompted to retake the scan.

Scan Data of Naturally Placed Grapes with Foliage Present:

Figure 2.10 shows the raw S21 scan data from sets 4 and 5 test case b. Set 4 had 0.56 kg of grapes present across six bunches in the middle and inner zones and set 5 had 0.84 kg of grapes present across nine bunches. The two scans have a similar profile and the scan from set 5 has a higher overall amplitude than that in set 4 by an energy increase of 22%. Figure 2.11 shows another two raw S21 scans from sets 10 and 11 test case b. Unlike Figure 2.10, there is an increase in scan energy of around 21%, yet the grape weight is 26% less than that of the scan from set 10. There are more variances between these two scan profiles than between the scans in Figure 2.10. Figure 2.12 shows two S11 scans from sets 1 and 2 test case b featuring a similar scan profile. The total reflectance energy in the set 1 scan is approximately the same as the energy in the set 2 scan, however the bunch weight is half that which was present in the scan from set 1. Figure 2.13 shows example S21 scan data from sets 10 and 12, test case c referenced to their associated scans with no grapes present, where the only difference between the two scans is the naturally hung grapes present in the inner zone. They contain similar overall energy levels but feature vastly different scan signatures and grape weight. There is no evident correlation between scan energy or signature and grape weight emerging across the presented scans. It is a possibility that the foliage present during the scan is interfering with the visibility of the grapes and causing sensitivity variations to grape weight.

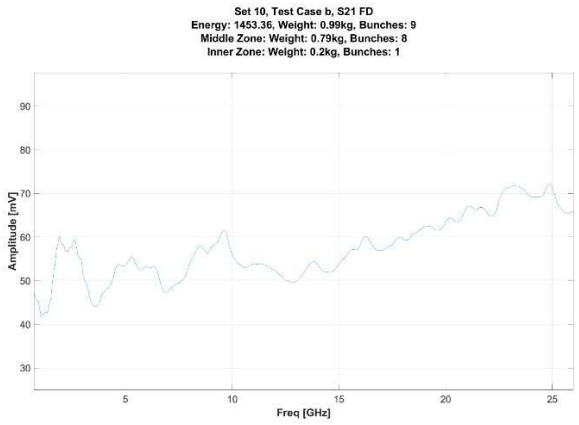


(a) Set 4, test case b raw S21 data scan.

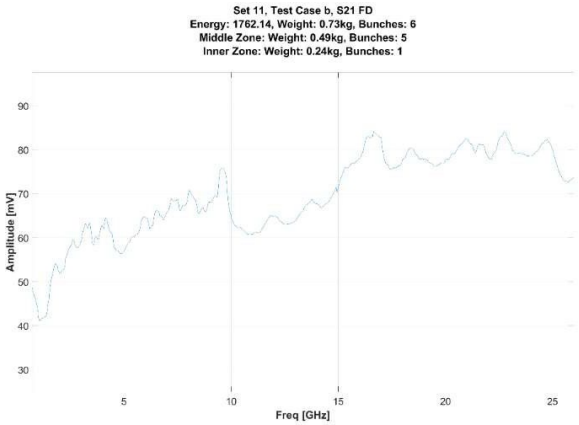


(b) Set 5, test case b raw S21 data scan.

Figure 2.10: S21 scan data from sets 4 and 5 showing an increase in amplitude along with an increase in grape weight.

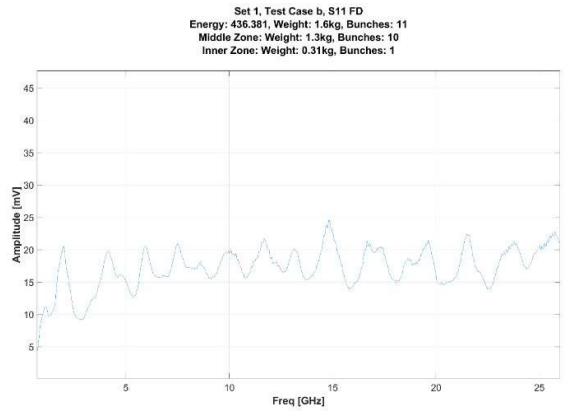


(a) Set 10, test case b raw S21 data scan.

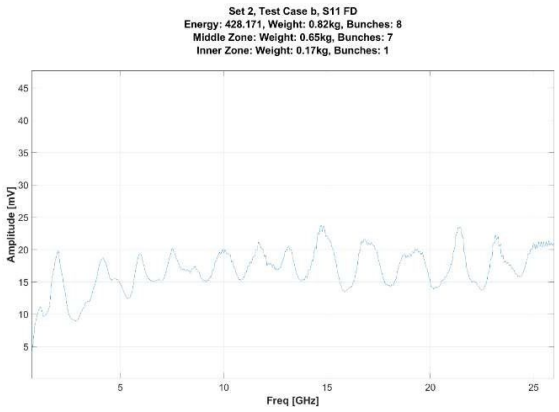


(b) Set 11, test case b raw S21 data scan.

Figure 2.11: S21 scan data from sets 10 and 11 showing an increase in amplitude along with a decrease in grape weight.

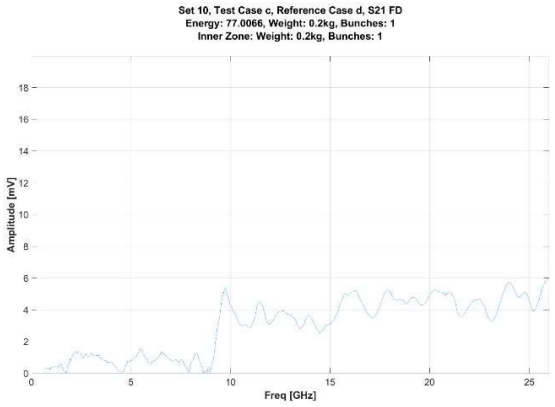


(a) Set 1, test case b raw S11 data scan.

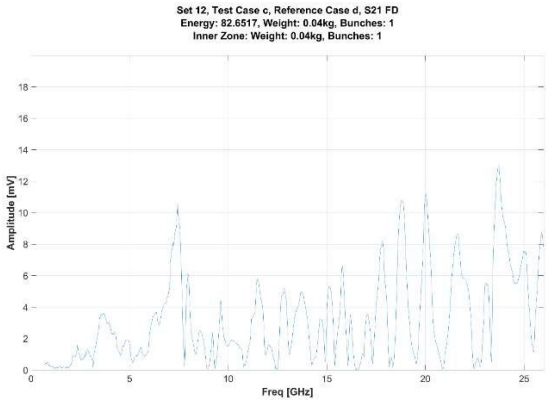


(b) Set 2, test case b raw S11 data scan.

Figure 2.12: S11 scan data from different sets with similar amplitude and scan signature along with a large difference in grape weight.



(a) Set 10, test case c referenced to no grape case S21 data scan.



(b) Set 12, test case c referenced to no grape case S21 data scan.

Figure 2.13: S21 scan data referenced to the associated no grape case from different sets with different scan signature and difference in grape weight.

Scan Data Weight Trends of Manually Placed Grapes:

Figure 2.14 and Figure 2.15 show S11 and S21 scan data for test case g referenced to the no grape case across a small range of sets. These show the difference in scan data attributed to grapes hung manually in the inner zone without foliage. This test case has been selected to minimise the ambiguity of grapes present, as case g has grapes only in the inner zone and case h has no grapes present at all. Additionally, the lack of foliage reduces the possibility of interference and provides the best chance at obtaining an ideal scan of grape bunches in the field. The grape weight present in the scans shown are as follows, set 1 had 0.28 kg of grape weight in the inner zone, sets 4, 2 and 7 had 0.14 kg, 0.15 kg and 0.16 kg respectively, and sets 3 and 6 each had 0.07 kg present. The selected sets show an example of variances in scans for different amounts of grape weight, an example of two scans of the same amount of grape weight, and a set of scans with minimal change in grape weight.

Sets 1 to 3 are in decreasing weight order and the energy present in the S21 scans are also in decreasing order, whether looking across the entire range of tested frequencies or in the 20-26 GHz range. While this may appear as a potential trend, looking at sets 4, 2, and 7 with weights in increasing order by 0.1 kg increments, their order by increasing energy is 7, 4, 2 which is neither positively nor negatively correlated to the grape weight present. At the lower end of the weight scale, sets 3 and 6 both have a grape weight of 0.07 kg. Sets 3 and 6 have similar S11 energies at 2.57 and 2.35 respectively, however their S21 energies are 76.80 and 32.10. In addition to set 3 having twice the S21 energy than set 6, the energy in set 3 is concentrated in the range of 10 to 20 GHz, while in set 6 the bulk of the energy is in the range of 5 to 10 GHz.

Additionally, there is a series of prominent peaks in the S11 scan data across the collected sets, however they are not consistent between sets and the range of grape weight present. There is a peak around 15 GHz in each S11 scan, though it does not correlate with the amount of grape weight present. The S11 energy weight plots show that there is less correlation in the K_u band (12 to 18 GHz) which contains the 15 GHz peak than the plot including the energy across the whole range of collected frequencies. The S21 energy weight plots show more correlation than that of the S11 data, however they are not sufficiently correlated, with an R^2 of 0.497.

Overall, in the clearest of scenarios without grape bunches present in adjacent zones or foliage present, there is no clear correlation between the total energy, subset of energies or scan signature and the bunch weight present in the inner zone. There are limited similarities between some sets, though this does not apply to all the collected sets and hold as a trend, nor cover the entire range of frequencies.

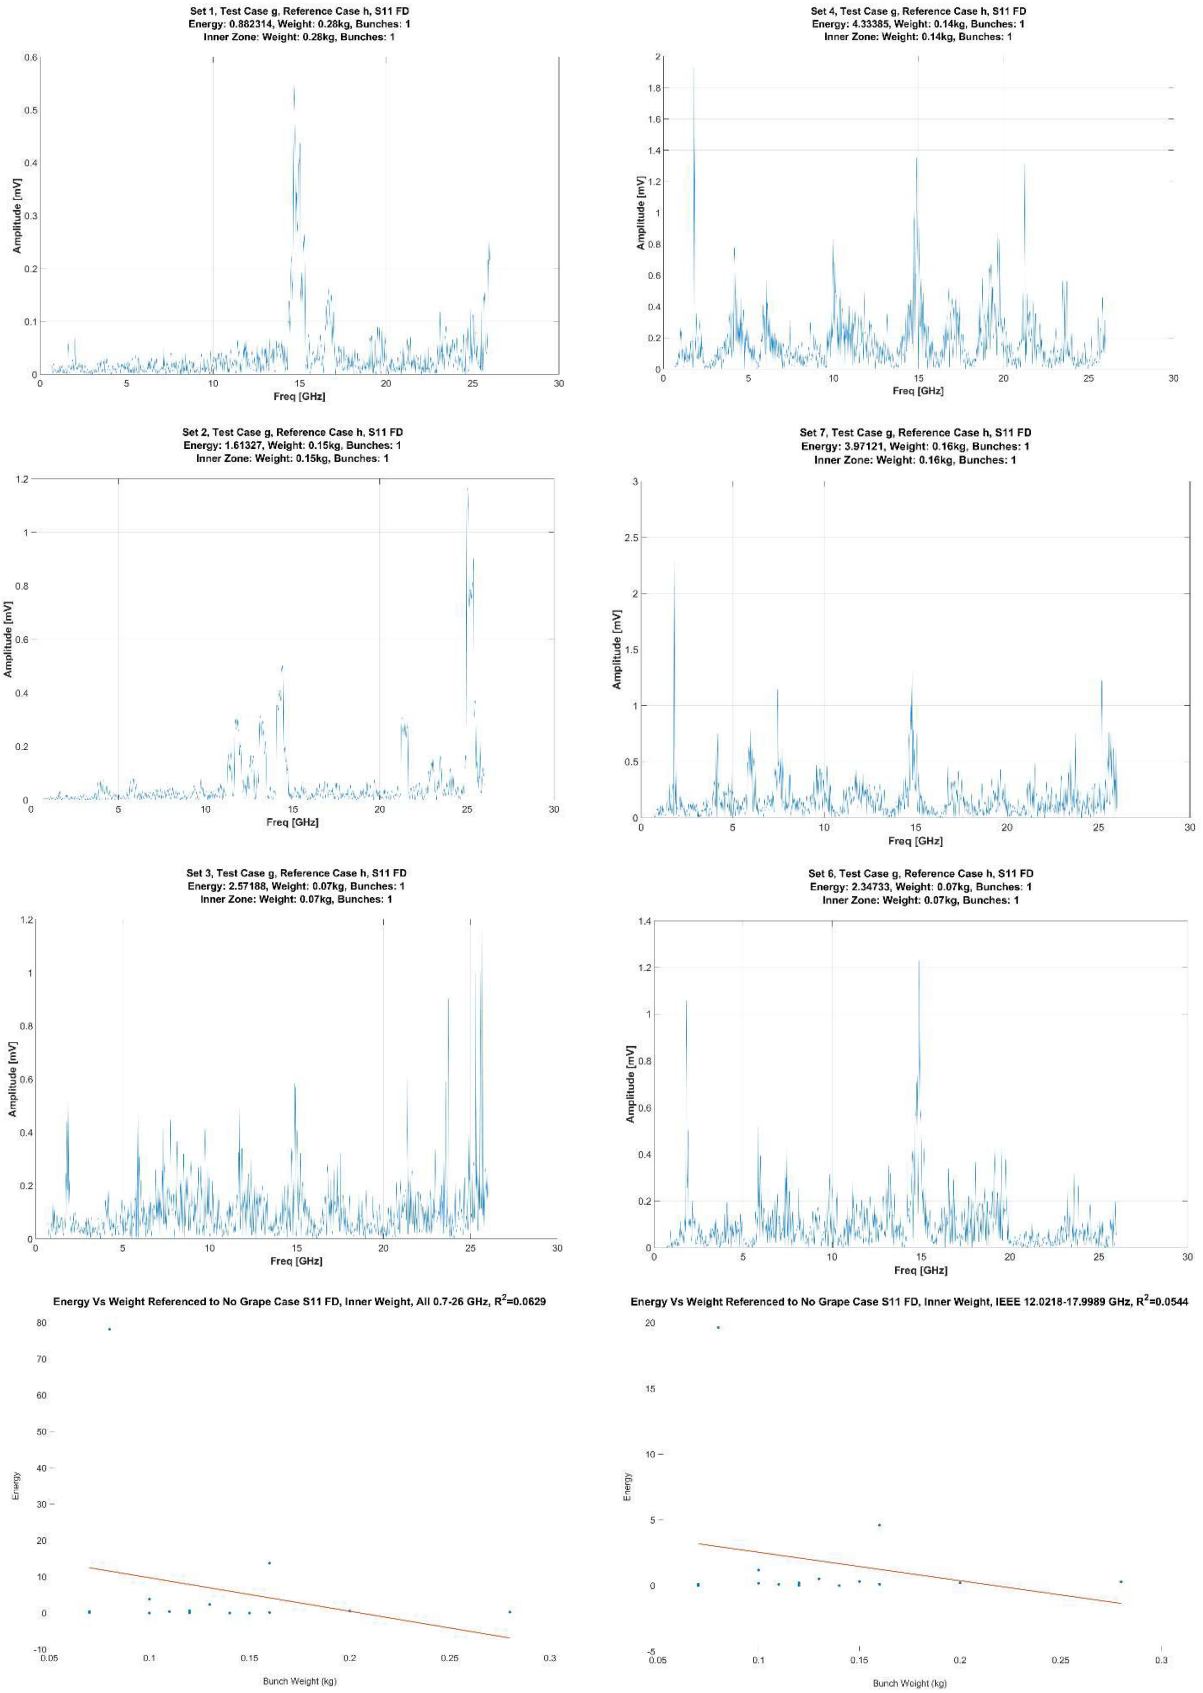


Figure 2.14: S11 scan data test case g referenced to the no grape case across a range of sets and summary plots.



Figure 2.15: S21 scan data test case g referenced to the no grape case across a range of sets and summary plots.

2.2.4 Zone Energy Level Summaries

All Frequencies:

Figure 2.16 shows the S21 total energy level in each zone across all sets, case e referring to energy in the outer zone, case f to the middle zone, and case g for the inner zone. These cases are where the foliage has been removed and the grapes are manually placed. Figure 2.17 also shows the S21 total energy level in the associated group of zones across all sets. Case e shows the energy for the weight in all zones, case f for the middle to inner zones, and case g for the inner zone.

Figure 2.16 show changes in weight in the outer zone shows very little energy contribution compared to the changes observed due to weight changes in the middle and inner zones. It is notable that in most sets there is significantly more grape weight in the outer zone than in the middle and inner zones, yet there is very little energy compared to the amount seen in the middle and inner zones. The lower energy level does not indicate increased transmission attenuation as the energy levels presented are the absolute difference between two scans, and if there were high levels of attenuation corresponding to the large grape weight difference there would be a larger energy response.

Figure 2.17 follows this trend of varying levels of sensitivity throughout the zones, although it also highlights that the energy levels do not consistently correlate with grape weight. This lack of consistency is shown by the order from least to most energy of cases e to g varying from set to set.

Figure 2.18 to Figure 2.20 shows the S21 energy vs grape weight plots for the referenced to previous case data for the manually placed grapes in the inner, middle and outer zones respectively. When using the energy across the whole collected frequency spectrum these figures show that there is very little correlation between the S21 scan energy level and the grape weight present. The most promising of the three zones is the inner zone, however its low R^2 of 0.497 indicates a low-quality fit.

The exhibited decreasing sensitivity to grapes further away from the centre may be responsible for some level of decreasing correlation due to the grape removal approach. During the trials, the way grapes were harvested was that if the grapes were not entirely contained within the zone of interest, they were removed. A grape bunch could be partially inside the middle zone, but predominantly in the inner zone, yet would be pruned with the middle grapes - this has the potential to skew the weight reading attributed to a zone. In this scenario, the grapes attributed to the middle zone that hang in the more sensitive inner zone would cause an increase in attributed energy and alter the correlation between energy and grape weight. This would explain a small number of outliers, however due to the large spread of data seen across the outer and middle zone plots, this issue does not explain all the variance exhibited.

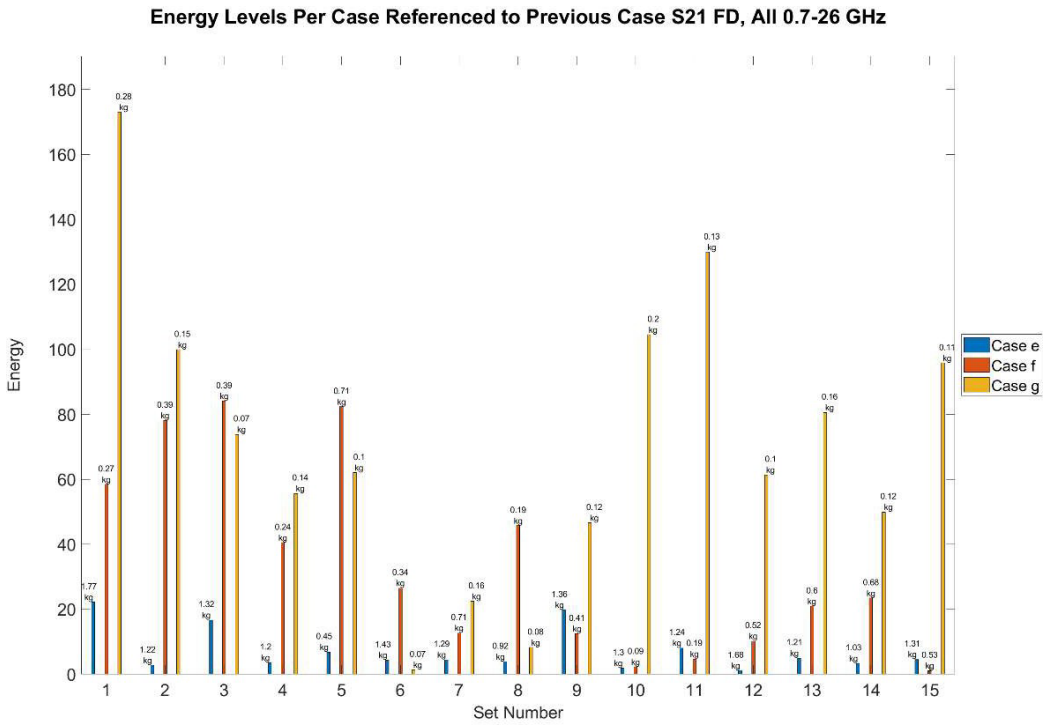


Figure 2.16. Energy for each test case referenced to the previous case's transmission data, manually placed grape bunches with foliage removed.

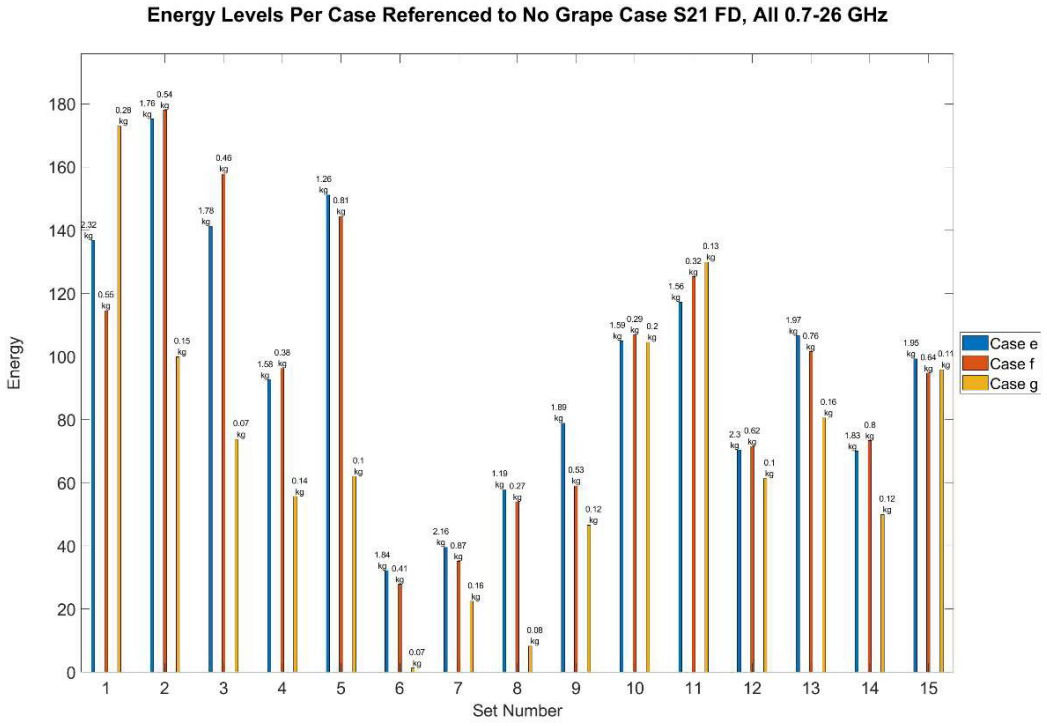


Figure 2.17. Energy for each test case referenced to the previous case's transmission data, manually placed grape bunches with foliage removed.

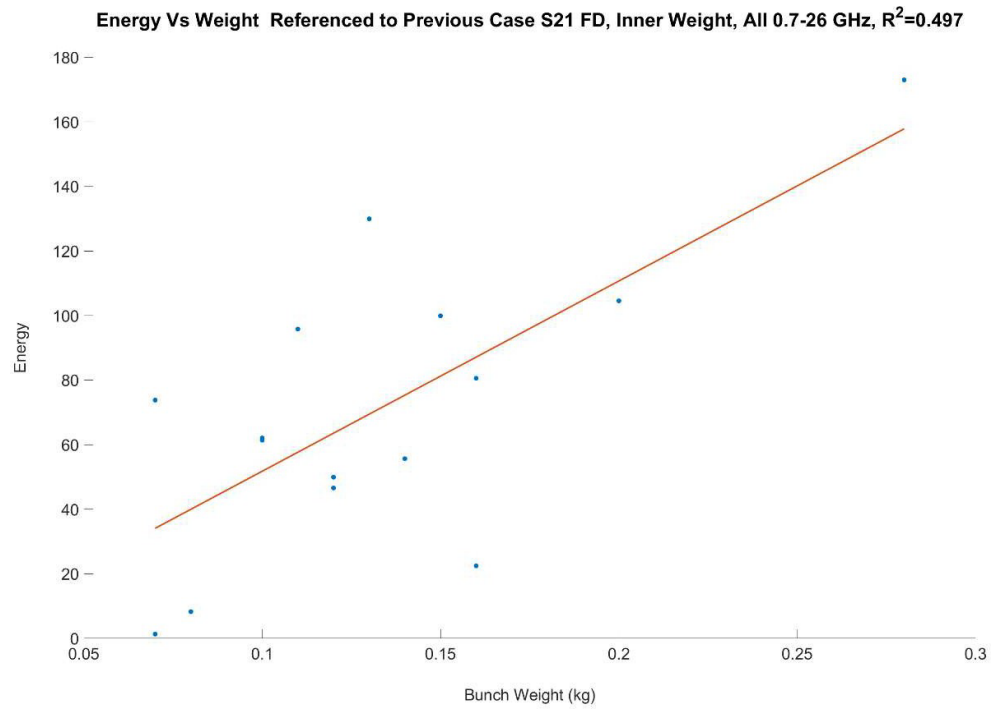


Figure 2.18. Energy vs weight for each test case referenced to the previous case's transmission data, manually placed grape bunches present in the inner zone with foliage removed.

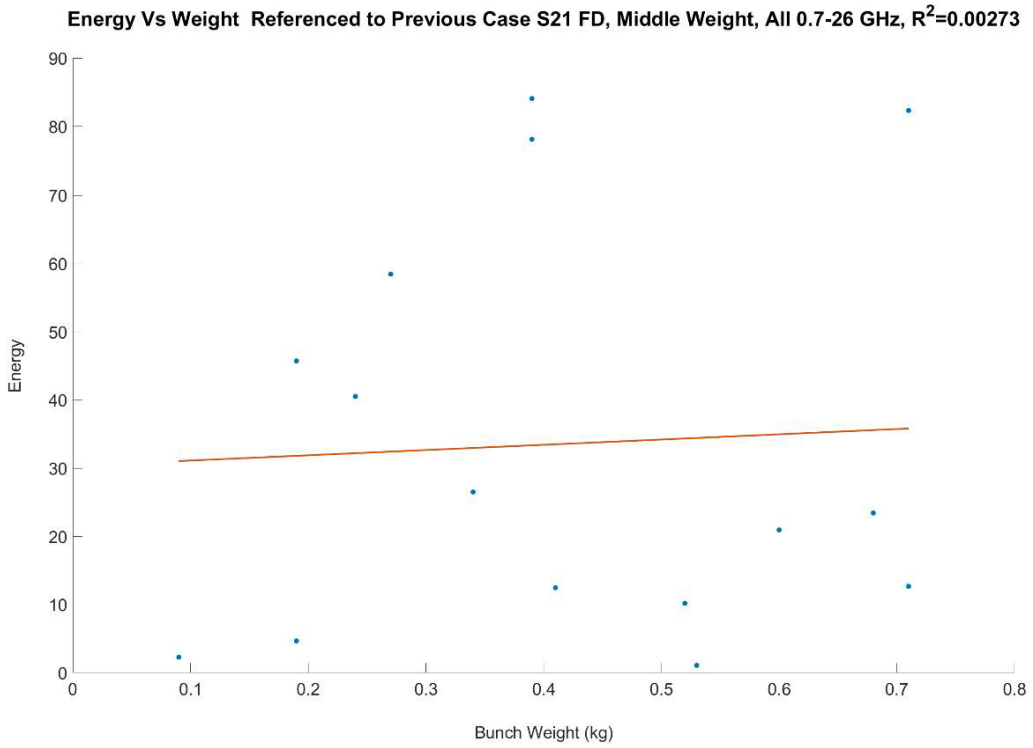


Figure 2.19. Energy vs weight for each test case referenced to the previous case's transmission data, manually placed grape bunches present in the middle zone with foliage removed.

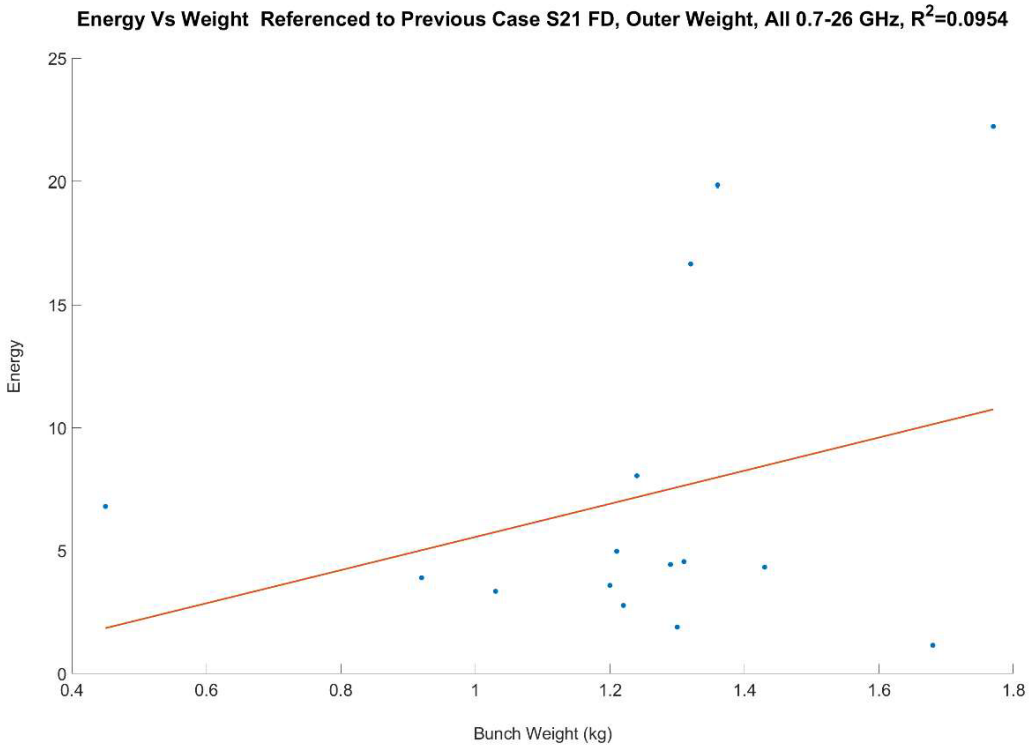


Figure 2.20. Energy vs weight for each test case referenced to the previous case's transmission data, manually placed grape bunches present in the outer zone with foliage removed.

Frequency Band Correlation

Due to the previously exhibited varying system response sensitivity between a target at the system's centre and those radially further out, the energy vs weight plots are shown for each step of increasing radial distance. For each of the steps, for both naturally and manually placed grape bunches, the frequency band with the highest R^2 value is presented below in Figure 2.21 to Figure 2.26. It was found that the frequency range with the best R^2 value in the inner zone was 18-26 GHz, a majority of the K band, for the combination of the middle and inner zones it was the S band 2-4 GHz, and for the outer to inner zones it was the C band, 4-8 GHz. For both the naturally hung and the manually hung grape bunches, the highest R^2 value was found in the same band, this indicates that the foliage present does not significantly affect the relative energy profile across the bands.

While these bands had the highest R^2 value for the region of interest, they did not reveal a high level of correlation between grape weight and band energy level. Except for the inner zone, the selected bands provided a better fit in terms of R^2 value than those that used the energy across the whole scanned range. Figure 2.27 and Figure 2.28 show the energy from the whole range of collected frequencies and have a higher R^2 value than those of their leading band counterparts. However, they do not have significantly higher correlation indicating a predictable grape weight based on scan energy.

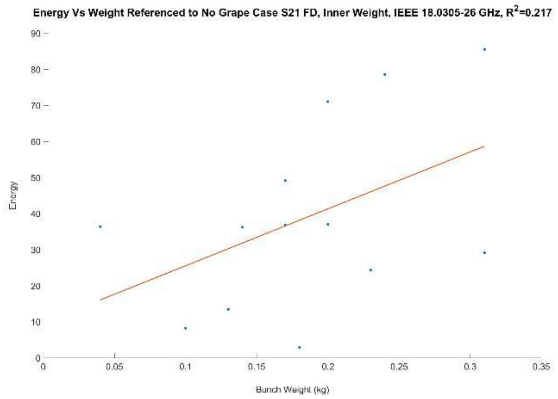


Figure 2.21. Energy vs weight for each set containing grapes exclusively in the inner zone, referenced to the no grape case transmission data, naturally placed with foliage present.

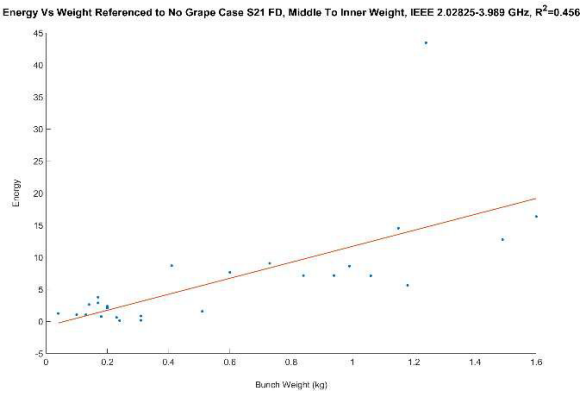


Figure 2.22. Energy vs weight for each set containing grapes exclusively in the middle to inner zones, referenced to the no grape case transmission data, naturally placed with foliage present.

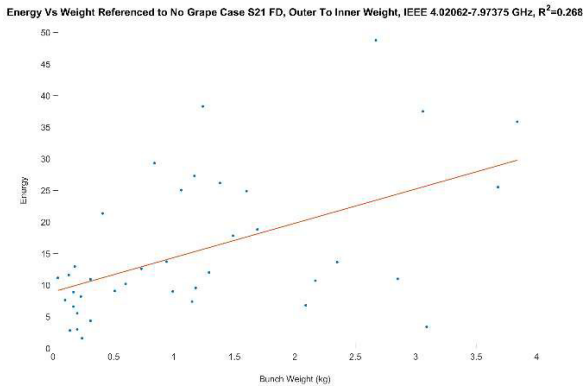


Figure 2.23. Energy vs weight for each set containing grapes in all zones, referenced to the no grape case transmission data, naturally placed with foliage present.

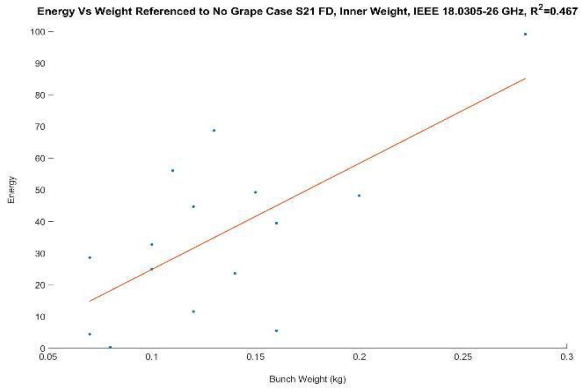


Figure 2.24. Energy vs weight for each set containing grapes exclusively in the inner zone, referenced to the no grape case transmission data, manually placed with no foliage present.

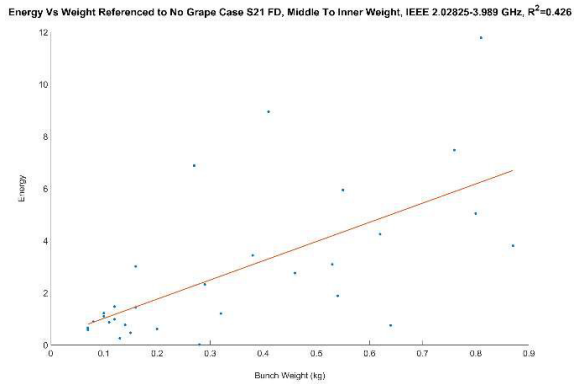


Figure 2.25. Energy vs weight for each set containing grapes exclusively in the middle to inner zones, referenced to the no grape case transmission data, manually placed with no foliage present.

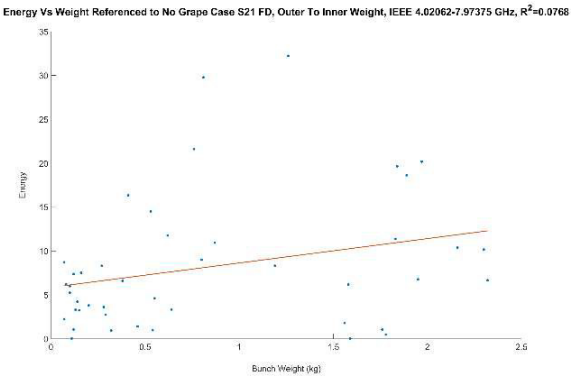


Figure 2.26. Energy vs weight for each set containing grapes in all zones, referenced to the no grape case transmission data, manually placed with no foliage present.

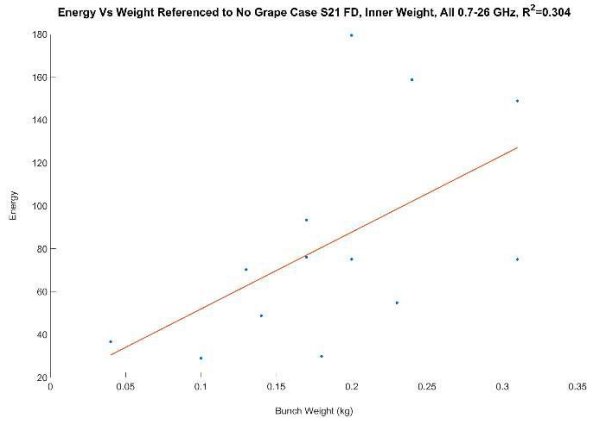


Figure 2.27. Energy across all frequencies vs weight for each set containing grapes exclusively in the inner zone, referenced to the no grape case transmission data, naturally placed with foliage present.

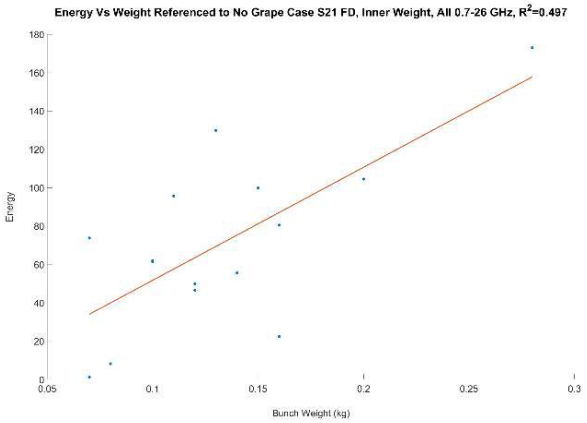


Figure 2.28. Energy across all frequencies vs weight for each set containing grapes exclusively in the inner zone, referenced to the no grape case transmission data, manually placed with no foliage present.

2.2.5 Summary

Across the range of scans and tests conducted, there are limited instances where the beginning of a trend emerges among a subset of the scans collected, however there are still many variances between scans, whether looking at relative or raw scan data. The positional sensitivity of the system exacerbates the variances in scan energy for a given grape weight. While the inner zone had the strongest response, there was not significant consistency in energy level in any given frequency band compared to the amount of grape weight present during the scan to constitute a predictable correlation.

It is likely that a single pair of opposing antennas is not sufficient to determine the grape contents of a scan due to the sensitivity variances and missed diverted energy. However, an alternative approach to overcome this issue would be to use an antenna array in order to isolate regions of grape bunches and capture the redirected energy that would otherwise be missed.

This page intentionally left blank.

Chapter 3 – Vine water status monitoring

3.1 Introduction

Water loss through transpiration represents more than 90% of plant water use. This process results in cooling of the leaf, relative to the environment. Consequently, accurately estimating canopy and environment temperatures offers a means to estimate vine transpiration at that point in time. Further, transpiration is under active control by the vine, being reduced when the vine is under water stress. As a result, the difference between actual vine transpiration and the transpiration that would be expected when water is not limiting, based on environmental conditions, is typically correlated to vine water status.

Previous attempts to use temperature for plant-based sensing have been limited to single point measurements, either small and unlikely to represent the vineyard block, or larger and likely to provide a plant signal mixed with background. Alternatively, high-resolution thermal imaging allows for larger section of canopy to be assessed while still giving the ability to isolate the pixels belonging to leaves from those of the background. However, the high cost of this technology still represents the main constraint to the development of a commercial product that uses thermography.

ProxiCrop represents a low-cost solution for thermography-based sensors. The cost of the device is significantly reduced by using low-resolution thermal camera paired with a low-cost RGB camera. A dedicated automated segmentation algorithm is then used to identify the pixels belonging to the canopy in the high resolution RGB image. From this, a mask is created and overlaid on the thermal image in order to obtain an average temperature from only canopy pixels.

3.2 System overview

The schematic representation of the sensor architecture in Figure 3.2a shows the main components in the ProxiCrop, which are all commercially available. The core of the device is a single-board computer, a Raspberry pi model 2B (Figure 3.2b), where a Raspberry pi RGB camera and a Lepton thermal image sensor are mounted (Figure 3.2c). The RGB and thermal cameras produce images with sizes of 820x616 pixels and 160x120 pixels respectively, both at resolution of 72dpi. A GPS receiver allows for the recording of the device's geographic location, while the data transmission is obtained via a 3G mobile dongle. Power supply comes from a 12,800 mAh USB 'power bank' with 5V output. The battery is recharged daily via a 10 W portable solar pane (Figure 3.2a). All components are mounted in a 162x122x90 mm polycarbonate enclosure with at total weight of less than 1 kg, making the device very portable and easy to install in the field.

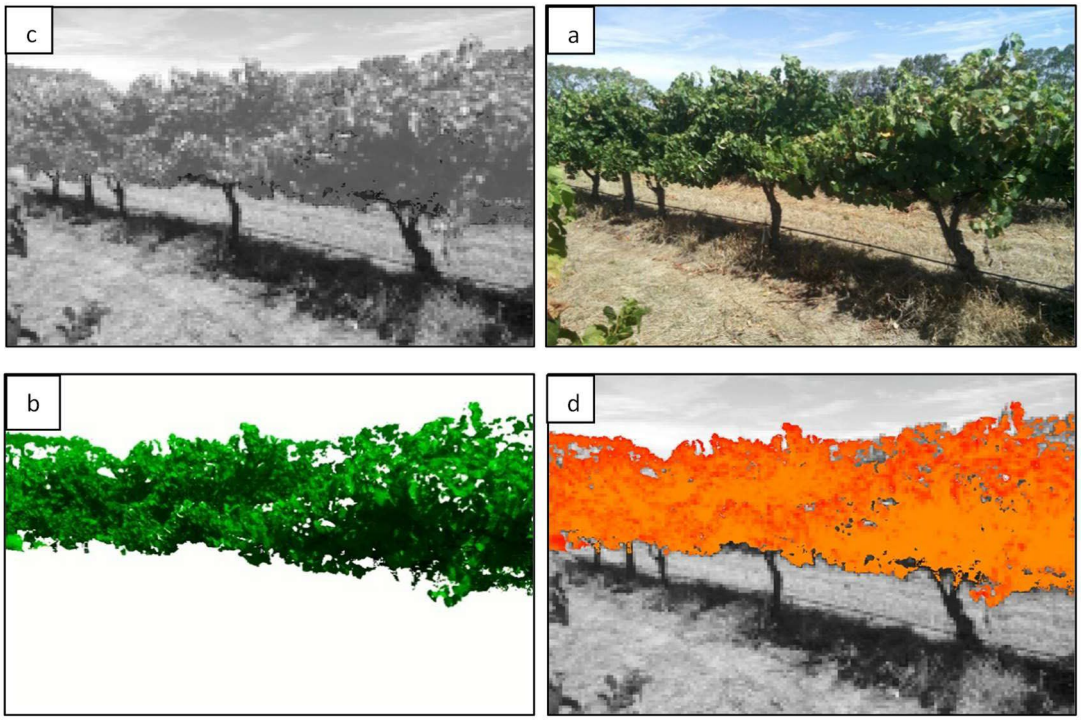


Figure 3.1 - Representation of the image processing using the ProxiCrop dual camera system. High resolution RGB images (a) are used to identify the canopy area (b) and create a mask overlaid on the thermal images (c) in order to select the thermal data from only the pixels of the canopy

The sensor takes an RGB photo (.jpg) and a thermal image (.pgm) every 5 minutes (users can change the interval according to their needs). The thermal camera configuration is stored as metadata in the .pgm files and a log file records when photos are captured. Each ProxiCrop transfers the data to an FTP server once a day, approximately around 12:00 pm, to minimise the power requirements for the cellular connection. Once the FTP connection is established, the sensor downloads the latest version of the firmware from the FTP server to itself and then uploads all images (.jpg and .pgm) as well as the logfile to the FTP server. When all files are copied successfully, the sensor archives all image files on the device and disconnects from the ftp server.

On the server side, the images are further processed once a day and average thermal data are created. In summary:

- Update firmware (from server to ftp server)
- Load images from ftp server to processing server
- Generate the leaf mask from the RGB image (Figure 3.2b, this happens daily and is based on an image around midday with good brightness).
- Scale the mask to the thermal image size.
- Apply leaf mask to the thermal image and calculate average leaf temperature (Figure 3.2d).
- Load BOM weather data.
- Run notification scripts that emails daily results.

The average canopy temperatures at 5 min intervals and selected weather information from the Bureau of Meteorology can them be displayed via a dedicated dashboard from which the raw data can be downloaded in .csv format (Figure 3.3).

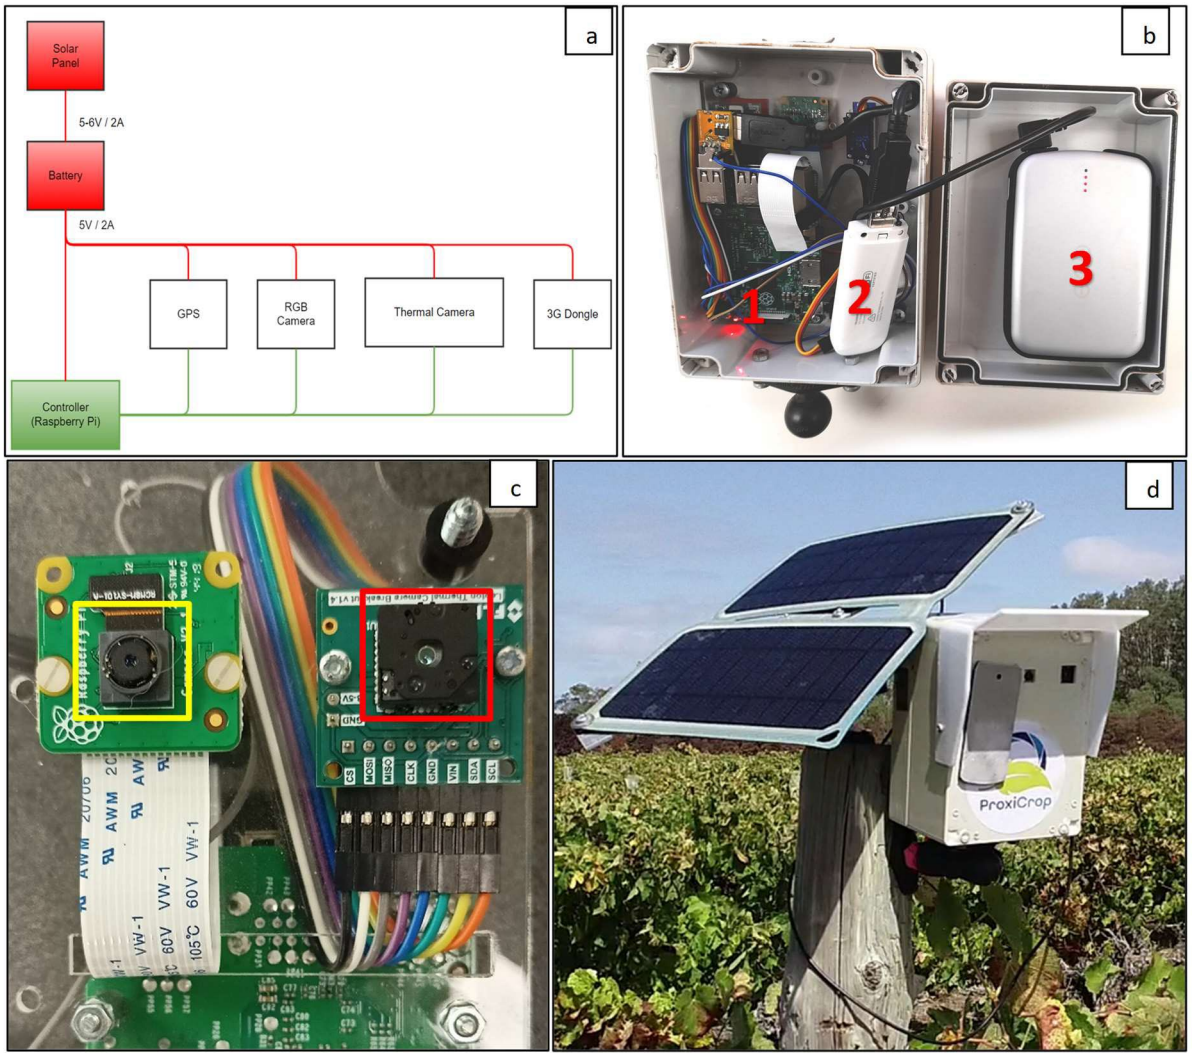


Figure 3.2 - a) Schematic representation of the sensor architecture. b) Internal view on the device showing the Raspberry pi board (1), the 3G dongle and the battery pack (3). c) Close-up image showing the Raspberry pi RGB camera (yellow box) and the Lepton thermal camera (red box). d) Typical field setup of a ProxiCrop.

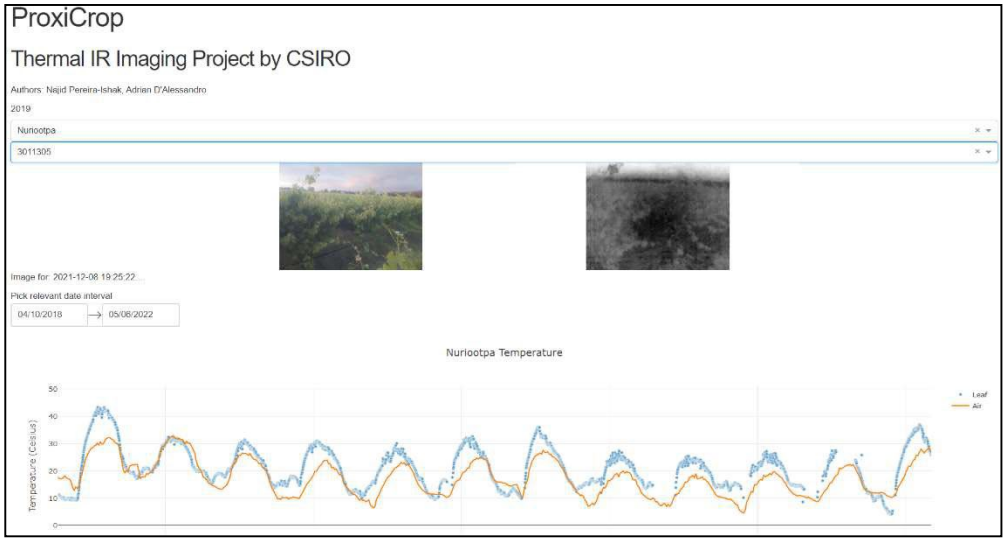


Figure 3.3 -ProxiCrop dashboard.

3.3 Data collection and field data collection

Following the preliminary field tests at the SARDI Nuriootpa Research Centre (NRC) during the 2019/20 season, originally through CSA 1701-1.5, further field testing was conducted in the 2020/21 and 2021/22 growing seasons. In both seasons, the main goals were: *1)* evaluating the thermographic data under different irrigation management, *2)* assessing the hardware components under the typical field conditions of a commercial vineyard and *3*) validating the thermographic data against selected plant physiological data.

3.3.1 Evaluation of the thermographic data under different irrigation management

During the 2020/21 and 2021/22, the ProxiCrop cameras were installed in commercial vineyards where different management of soil moisture was present. This allowed for the comparison of the thermographic data captured from canopies that, in principle, would have experienced different levels of water stress. In most vineyards, water stress is often monitored using a single soil moisture sensor that must be representative of a whole vineyard block (or multiple blocks) for the irrigation scheduling. This is a consequence of the high cost of installation and maintenance of such sensors. The low cost and high portability of the ProxiCrop cameras means that more sensors for water stress monitoring or irrigation scheduling could be installed in different areas of the vineyard to better capture block variability.

Below, we present the results from case studies at different commercial vineyards where the ProxiCrop cameras were tested to evaluate how thermographic data can be used to inform growers on canopy water stress under different irrigation/soil moisture management.

Case study 1 – 2020-21 season, Pernod Ricard (Rowland Flat):

In 2020, in collaboration with SWAN Systems and Pernod Ricard (PR), a site was chosen near Rowland Flat. In addition to the typical canopy and agronomic management, the site was chosen due to the presence of a soil moisture sensor (continuous data logger) operated by SWAN and a PR field trial where different practices for water management were under evaluation, including:

- On row straw mulch application for the reduction of water loss due to evaporation
- Under-vine sprinklers
- Standard drip irrigation (control)

A map of the location, showing the position of the different water management treatments, is presented in Figure 3.4.

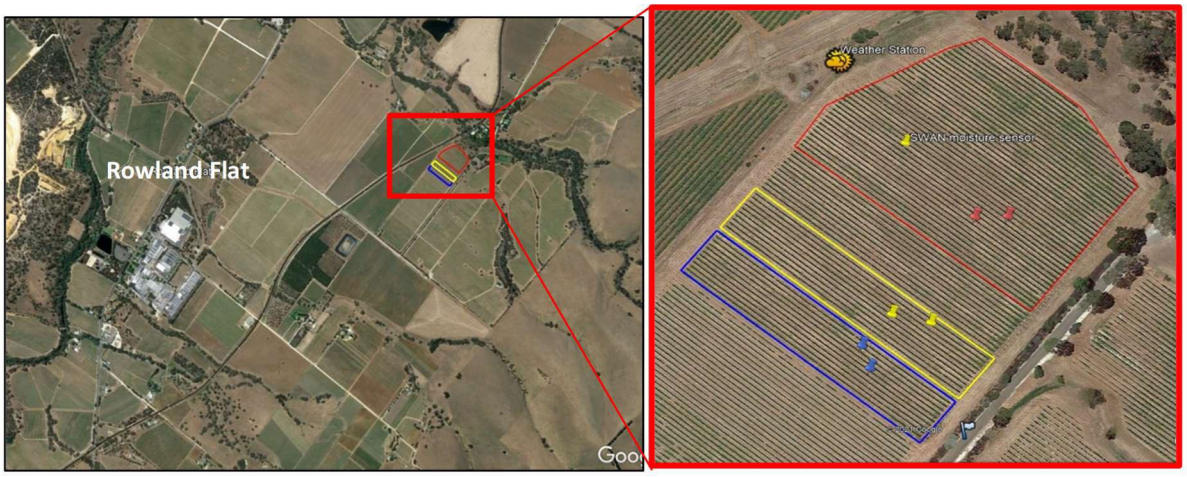


Figure 3.4 - Location and trial layout at the Roland Flat site. On the right, the blue, yellow and red polygons represent the areas with straw mulch applications, sprinkler irrigation and standard drip irrigation respectively. The coloured pins represent the approximate location of the ProxiCrop sensors and the soil moisture sensors used by SWAN Systems (BLUE= Straw mulch; YELLOW= Under-vine sprinklers; RED=Control).

The soil moisture sensor was in the north-west corner of the block, within the area under the standard drip irrigation (*control*). Two ProxiCrop cameras per treatment (six in total) were installed in November 2020 and kept in the field until February 2021 (Figure 3.4 right).

In addition to the 5 minute canopy temperature data, a *Temperature-Time Threshold* (TTT) index was calculated daily. The TTT index accumulates the time that a crop canopy is above a specified critical temperature and was developed by the USDA for automatic irrigation scheduling in different crops. In this case study, the critical canopy temperature for grapevine was set at 40° C.

The 2020/21 season at Rowland Flat was below average in terms of air temperature. Nevertheless, in January 2021 the site experienced multiple days with maximum temperature exceeding 30° C.

The data collected by the ProxiCrop sensors showed that the canopy temperature was affected by the different water management (Figure 3.5).

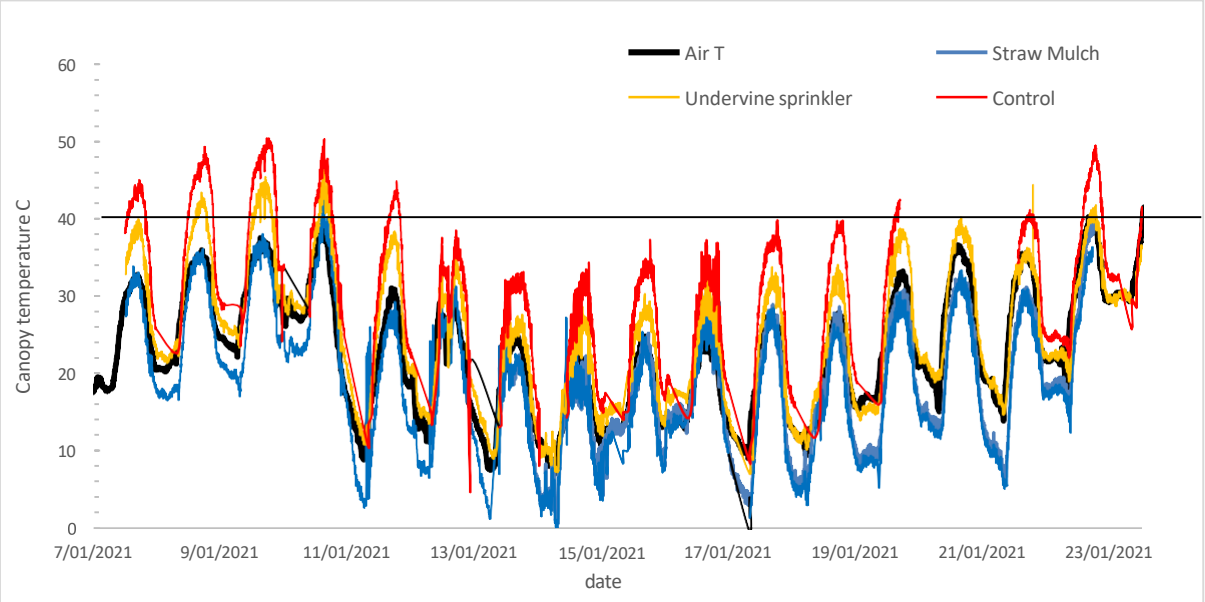


Figure 3.5 - Average canopy temperatures (at 5 min interval) from the two ProxiCrop sensors in the three different soil water management treatments (7-23 January 2021) at the site in Rowland Flat (black dashed line represents the critical canopy temperature at 40° C).

Between 7-23 January 2021, the canopy temperature in the *control* treatment exceeded the critical temperature of 40° C on multiple occasions and, on average, the canopy temperature at midday was more than 10° C above the air temperature measured by a nearby wheatear station. The treatment with the *undervine sprinklers* followed a similar pattern to the *control* but on average the canopy temperature was lower. On the other hand, the canopy temperatures in the *straw-mulch* treatment were significantly lower and closer to the air temperature, denoting a much smaller level of canopy stress.

This was also indicated by the TTT index calculated from the ProxiCrop data (Figure 3.6). The canopy in the *control* treatment experienced more than 77 hours above the critical 40° C threshold, in comparison to only 4 h and 1 h in the *undervine sprinklers* and *straw-mulch* treatments, respectively. Clearly the canopy temperature sensors were recording differences associated with the treatments, providing information on vine performance in the absence of soil moisture sensors.

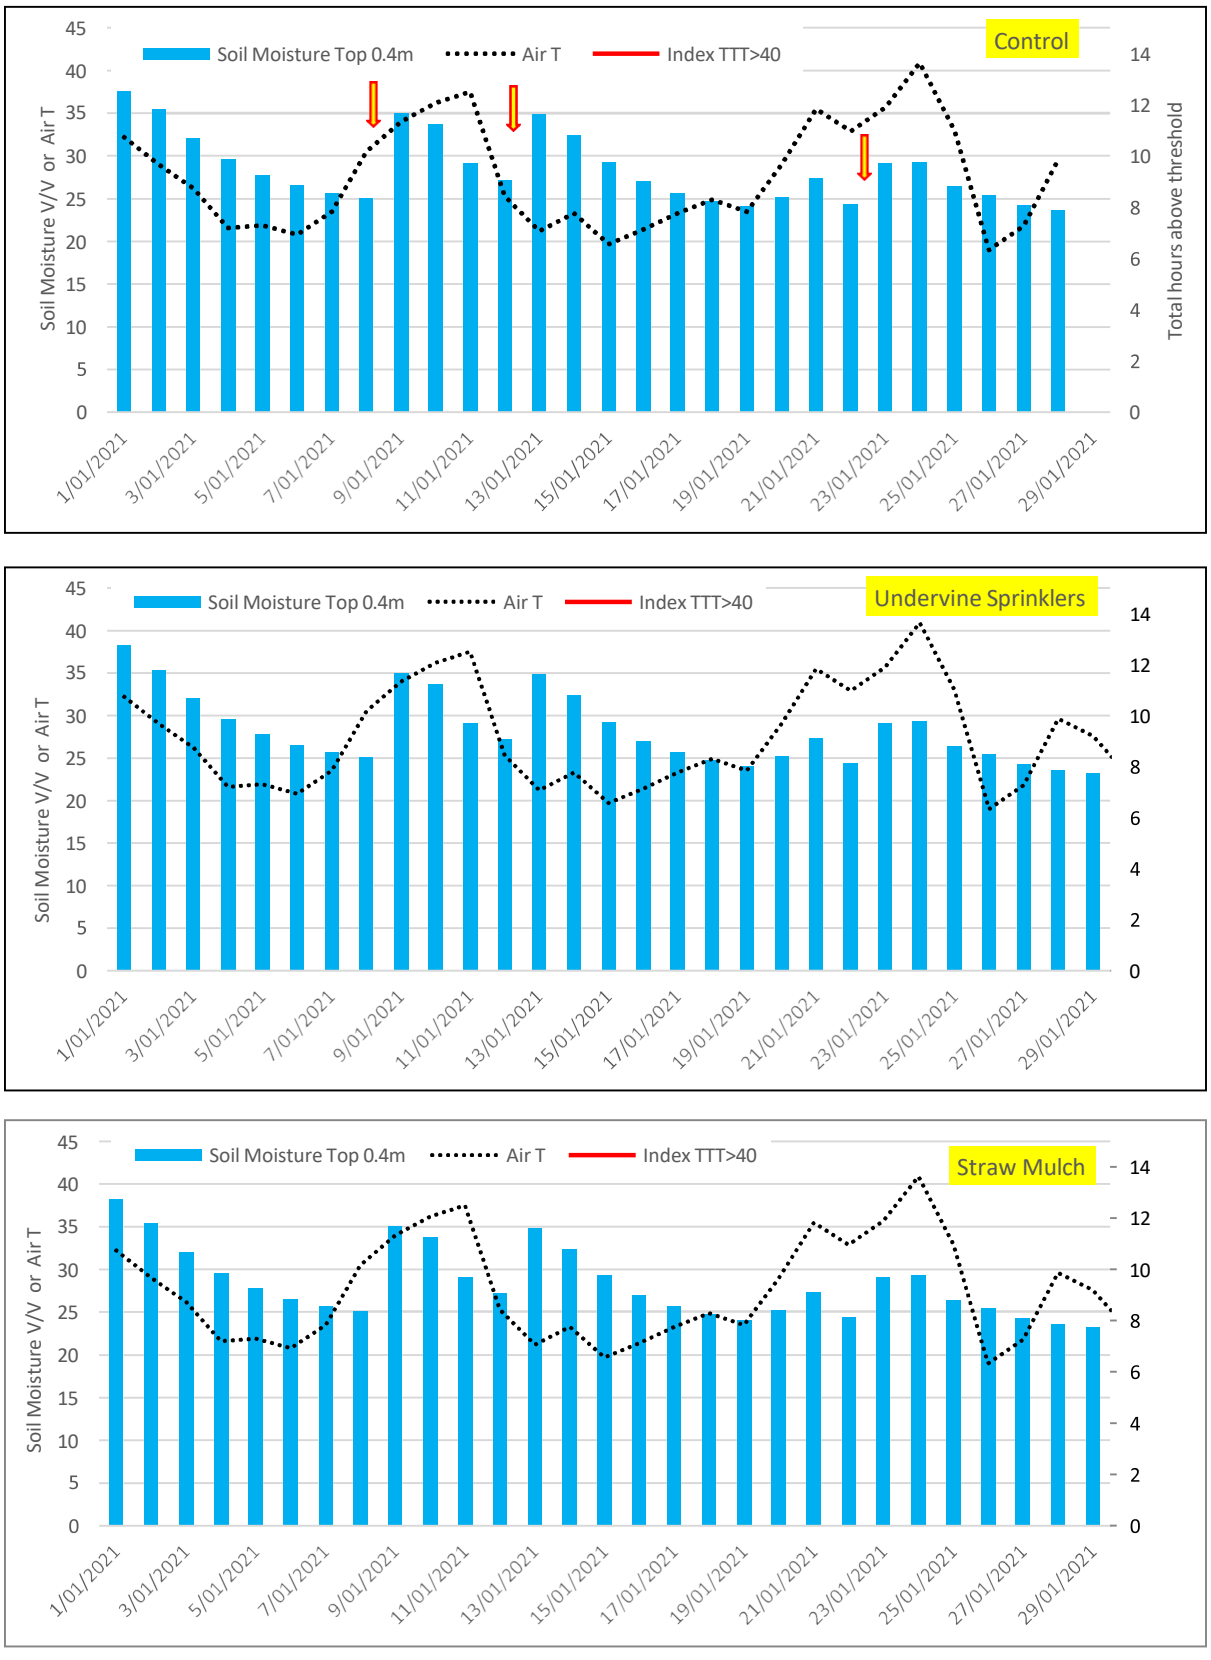


Figure 3.6 - Daily soil water content in the top 0.4m of soil within the control treatment and daily Temperature-Time Thresholds (>40° C) index in the three treatments at the site near Rowland Flat in the 2020/21 season. The yellow arrows mark the days when the block was irrigated.

The field testing in the 2021/22 season continued the devices' testing based on the lessons learnt in the 2020/21 season. Again, the focus was on improving the hardware functionality under the typical field conditions of commercial vineyards. Moreover, the designs of the field trials were such that the data of the ProxiCrop devices could be compared to the data from other continuously logging field sensors, such as soil moisture sensors and micro-tensiometers (Flora Pulse).

Two locations were chosen in the Langhorne Creek viticultural region. The criteria for the selection of the locations were:

- Uniformity of the block, with the same grapevine variety under two different irrigation managements and/or two different heat damage reduction management (i.e. with or without overhead sprinklers)
- Possibility to install soil moisture sensors, micro-tensiometers (vine trunk size dependent) and weather station.

Based on the above criteria, the two chosen sites were (Figure 3.7):

- Pernod Ricard (Manager: Tim McCarthy)
- Jess Farm (Manager: Mark Hall)

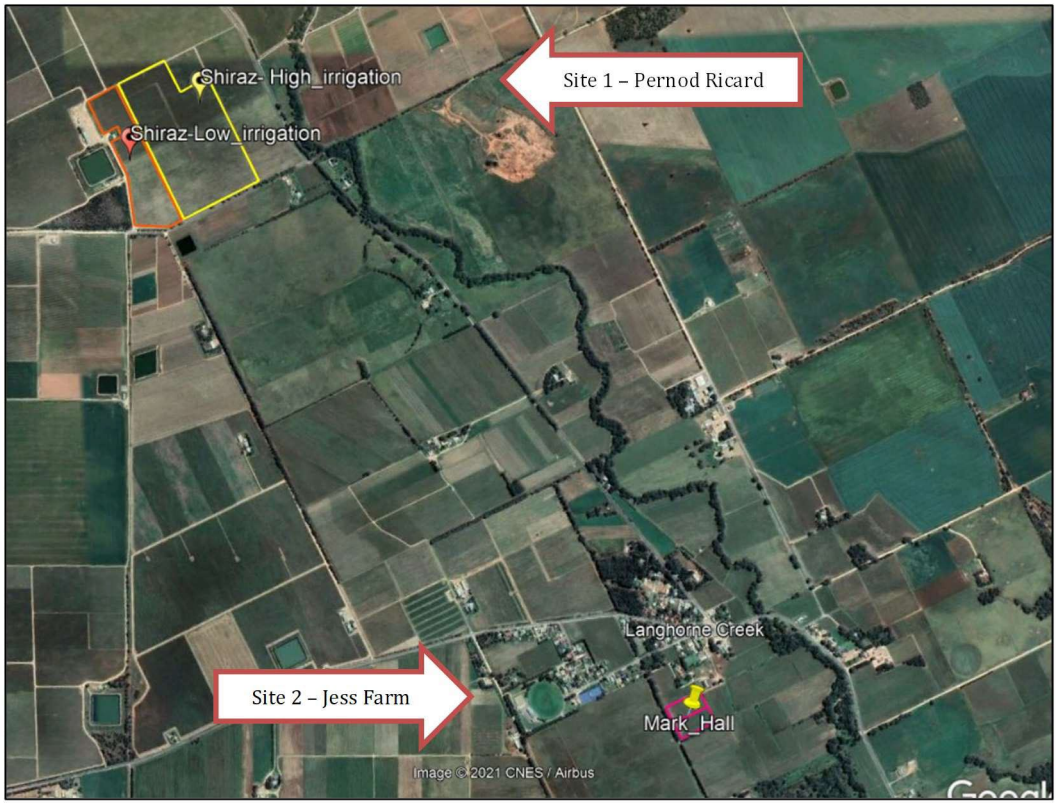


Figure 3.7 - Trial locations at Langhorne Creek.

Case study 2 – Pernod Ricard (Langhorne Creek):

The main goal at this site was the test of ProxiCrop camera as a supporting tool for irrigation scheduling under two different irrigation management. This site was characterised by having two adjacent Shiraz blocks at high (Block D) and low irrigation (Block C) regimes (Figure 3.8).

Two ProxiCrop sensors were installed along two rows on both Block C and D and one soil sensor was installed between the sensors in each block.

In addition to the soil sensors, two *Flora Pulse* micro-tensiometer (one in each block) were installed in two vines within the field of view of the ProxiCrop cameras. The micro-tensiometer directly measures the plant water status as water potential. Because the measurement is taken directly inside the water-carrying tissue, this sensor is considered to be one of the most accurate ways to inform irrigation scheduling. However, the sensor can be expensive and quite destructive to the plants, as it needs to be embedded into the woody tissue of the vine trunk. Therefore, the ProxiCrop sensors could represent a non-destructive and cheaper alternative to the micro-tensiometers.



Figure 3.8 - Trial layout and location of the sensors at Pernod Ricard in Langhorne Creek in 2021/22.

As previously, the logged canopy temperature data demonstrated the impact of management treatments, in this case irrigation (Figure 3.9). As expected, the canopy temperature differences between treatments were highest when the stem water potential differences

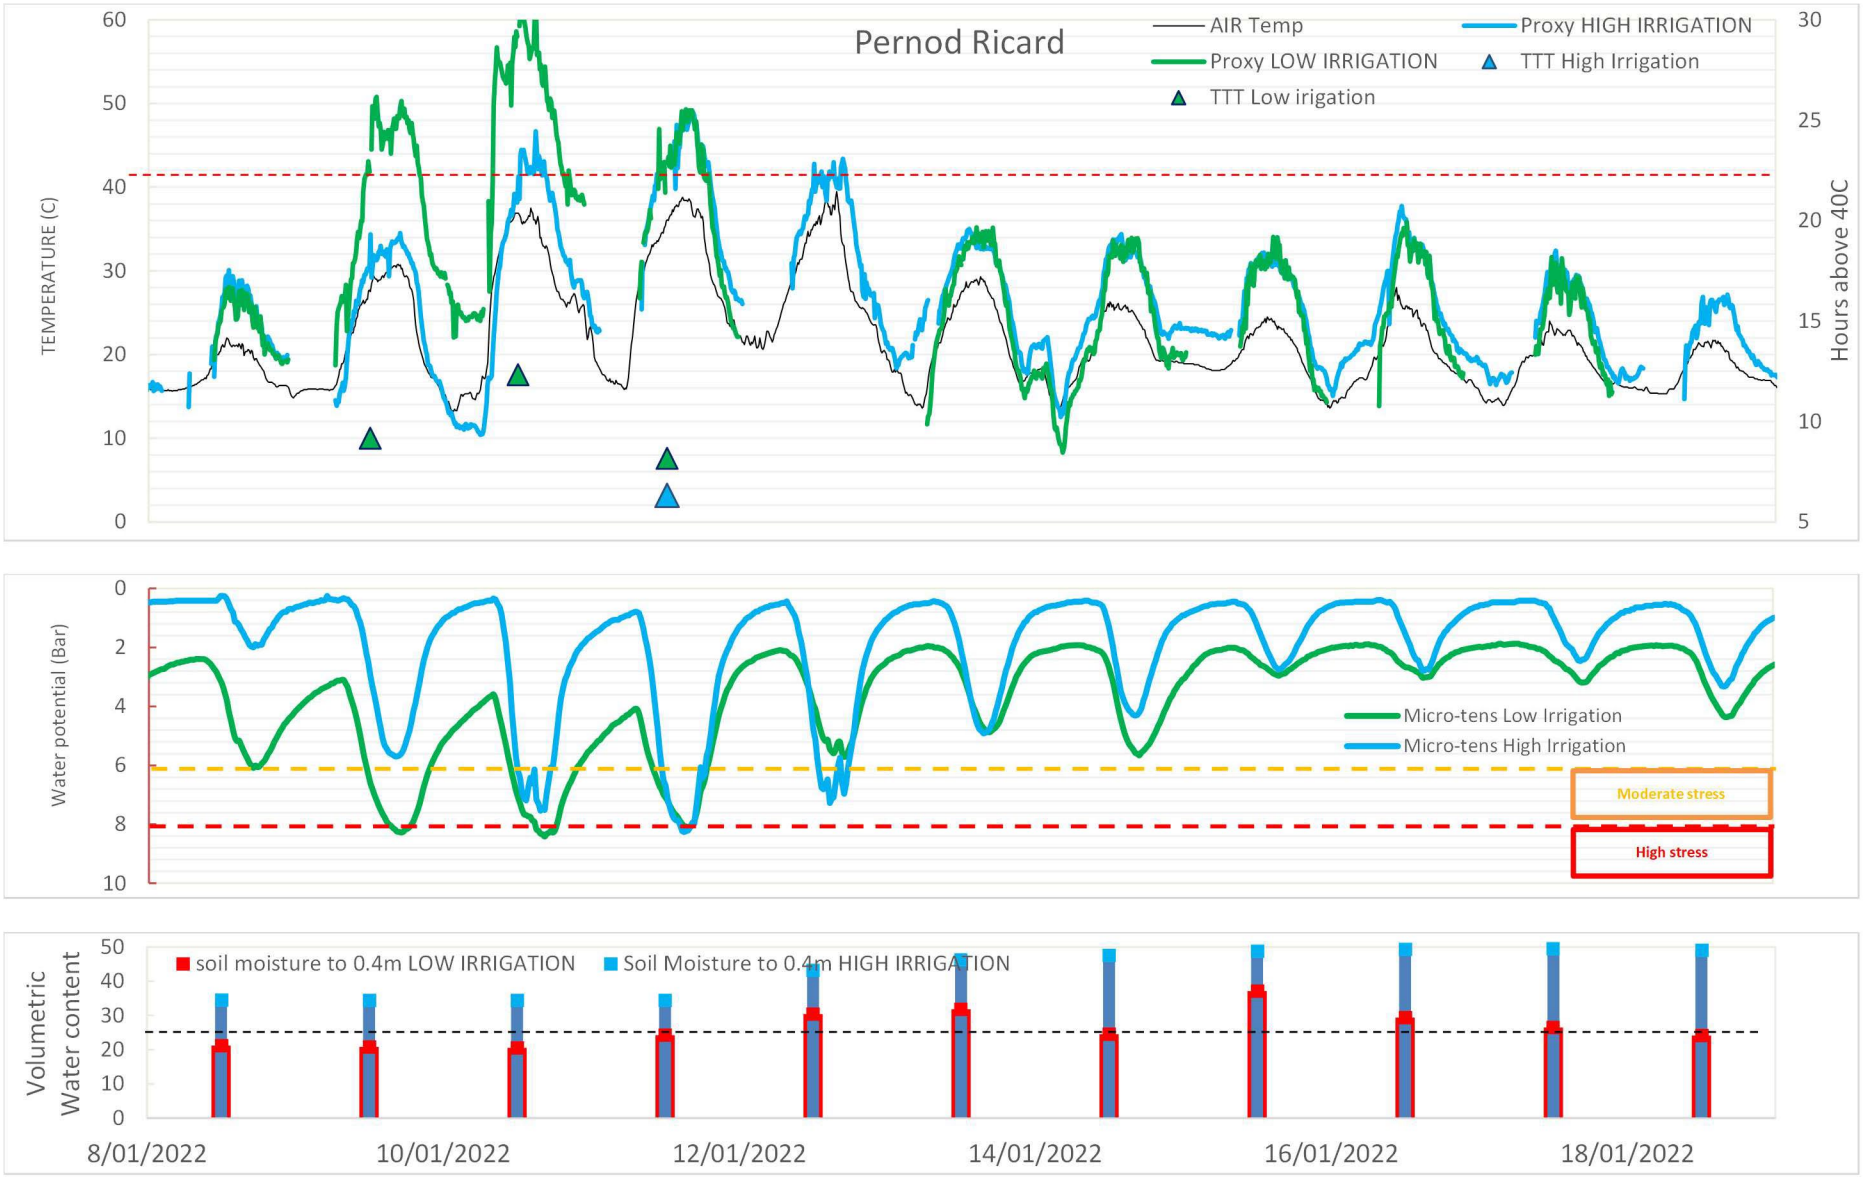

Pernod Ricard
AIR Temp
Proxy HIGH IRRIGATION
Proxy LOW IRRIGATION
TTT High Irrigation
TTT Low irigation
TEMPERATURE (C)
Hours above 40C
Water potential (Bar)
Micro-tens Low Irrigation
Micro-tens High Irrigation
Moderate stress
High stress
Volumetric Water content
soil moisture to 0.4m LOW IRRIGATION
Soil Moisture to 0.4m HIGH IRRIGATION
8/01/2022
10/01/2022
12/01/2022
14/01/2022
16/01/2022
18/01/2022

Figure 3.9 - (previous page)– Thermal data from ProxiCrop sensors and temperature time threshold index from the Pernod Ricard site at Langhorne Creek in 2021/22(top graph) compared to plant water potential measure by the Flora Pulse micro-tensiometers (centre graph) and average soil moisture content (bottom graph). Orange and red dotted lines in the centre graph represent water potential (bar) upper limits associated to moderate and high plant water stress respectively, according to the Flora Pulse system.

were highest, with high canopy temperatures where air temperature was high and stem water potential most negative. Later in the time-course provided in Figure 3.9, air temperatures were low and the stem water potentials indicated no stress, there was no difference in canopy temperature.

Case study 3 – Jess Farm (Langhorne Creek):

At this site, overhead sprinklers were installed over a Malbec block as a heat stress mitigating tool in an area prone to seasonal heat waves. Currently, the decision making on when to activate the sprinklers is not based on direct information from the vine canopy, but only on indirect data such as weather forecast and air temperature. Therefore, the main goal at this site was the test of the ProxiCrop sensors as a supporting decision-make tool for the mitigation of seasonal heat stress due to heat waves. Precisely, the continuous monitoring of canopy temperature by the ProxiCrop sensors could be used to better inform the growers on when to activate heat stress mitigating tools (like overhead sprinklers) or even automatically switch on and off such mitigating tools.



Figure 3.10 - Trial layout and location of the sensors at the Jess Farm in Langhorne creek in 2021/22.

Four ProxiCrop sensors were installed along two rows of a Malbec block and one soil sensor was installed between the sensors in each row. A small weather station was installed in the vicinities of the block to monitor the local weather conditions (Figure 3.10).

Unfortunately, season 2021/22 had below average temperatures at Langhorne Creek, with no occurrence of significant heatwaves. However, on 31 December 2021, the forecast for the air temperature at the site was expected to exceed 40° C and the overhead sprinklers were switched on at approximately 10:48 in the morning to prevent any damage to the crop because of heat stress.

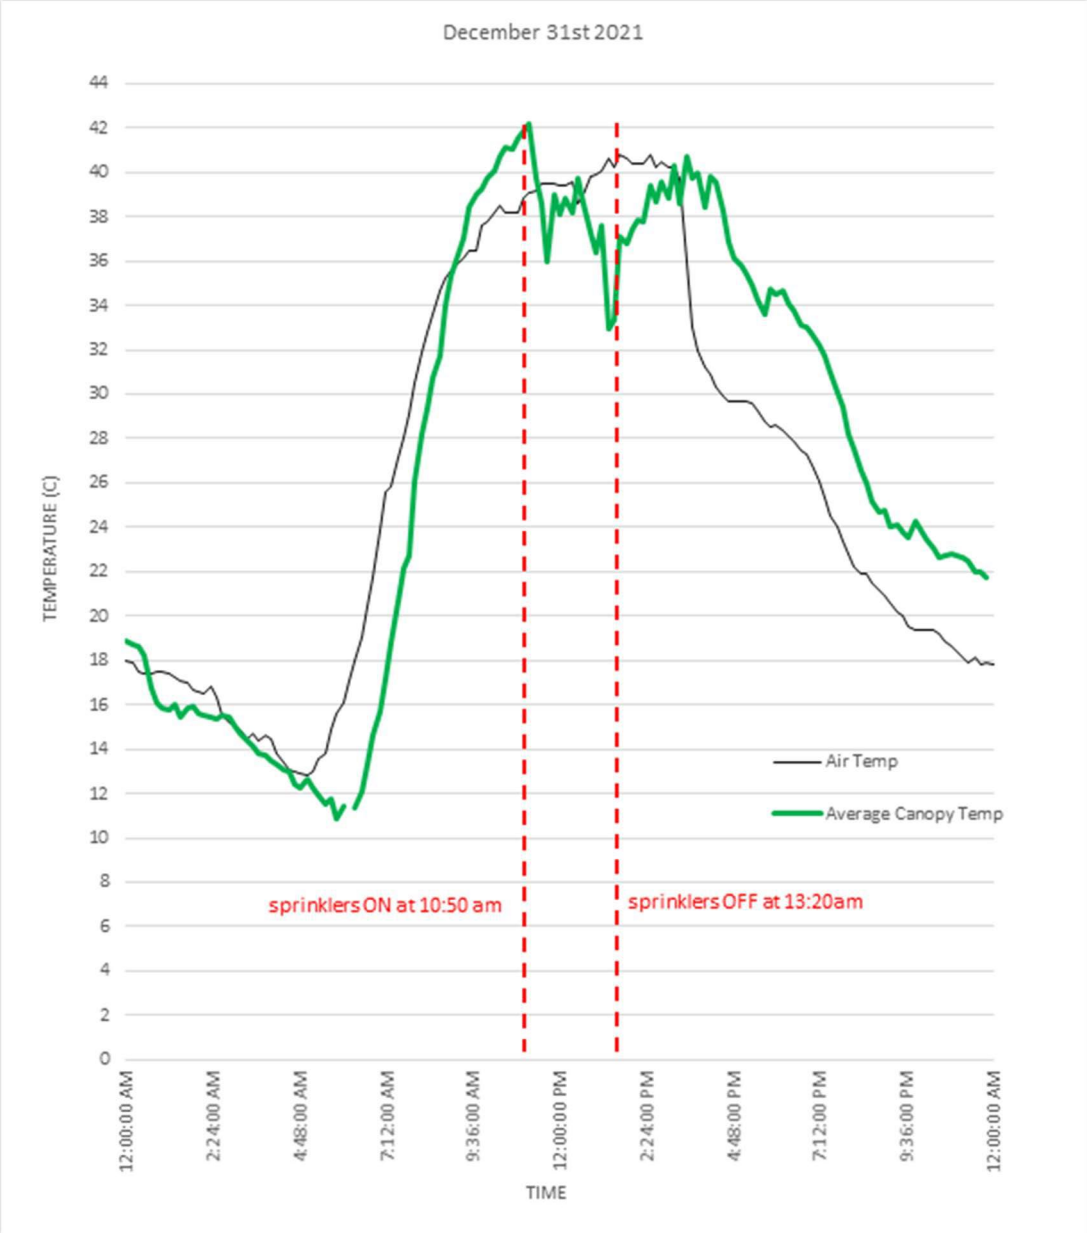


Figure 3.11 - Average canopy temperature from the Malbec block and air temperature on 31 December 2021. The red dotted lines represent the times at which the overhead sprinklers were switched on and then off to reduce canopy heat stress.

The sprinklers were switched off approximately two and a half hours later at 13:20 (Figure 3.11). During this time, the ProxiCrop sensors were able to capture the changes in canopy temperature as affected first by the local weather conditions and then by the activation of the overhead sprinklers.

As shown in Figure 3.11, the canopy temperature prior to the activation of the sprinklers exceeded 40° C for approximately 15 min. After that, the cooling effect of the sprinklers reduced the canopy temperature to below 34° C. However, once the sprinklers were off again, the canopy temperature rose again to above 40° C within 2 hours.

Note that the actual decision to switch on and off the sprinklers was based on the air temperature and the grower's experience. However, the results demonstrated that the ProxiCrop cameras could be an effective way to directly monitor canopy heat stress and automatically switch on and off the overhead sprinklers based on a critical canopy temperature threshold rather than subjective decisions by the grower.

3.3.2 Assessment of the hardware components and design development

A key component of the ProxiCrop development was the assessment of the hardware (design and functionality) under the typical conditions of a commercial vineyard. In particular, the functionality of the thermal cameras was assessed against mechanical operations in the vineyard such as chemical spraying.

During the 2021/22 season at Rowland Flat, two ProxiCrop prototypes (Figure 3.12 centre and right), were developed to prevent spray build up on the shutter mechanism of the thermal camera and compared to the original design (Figure 3.12 left). One of these used protective windows installed directly in the custom-made aperture of the PVC box (Prototype 1, P1, Figure 3.11 centre). Thermal imaging uses long-wave infrared frequencies, requiring a window transparent to this part of the spectrum. Such windows are made of materials such as Zinc selenide and attract a significant cost (circa $300).

The second prototype (Prototype 2, P2, Figure 3.11 right) had a custom-made external shutter that covered the ProxiCrop lenses between image capturing. The shutter was made of aluminium and the opening/closing movement was done by an SG90 micro servo motor. Prior to the field season, a stress test of the shutter was also conducted to check the durability of this system during a typical season of data collection. The SG90 micro servo performed well over 55,000 open/close cycles (equivalent to three seasons) with no issues.

Both prototypes performed well during the 2020/21 season at Rowland Flat, in terms of protection from chemical spraying and wet weather. For 2021/22 the P2 prototype with shutter and servo was used as this was a significantly cheaper solution than the P1 prototype with IR transparent window.

A further design improvement in the 2021/22 for P2 was the addition of a custom-made cover, similar to a 'roof', that increased weather protection of the system and reduced potential glare in the images due to direct sunlight.

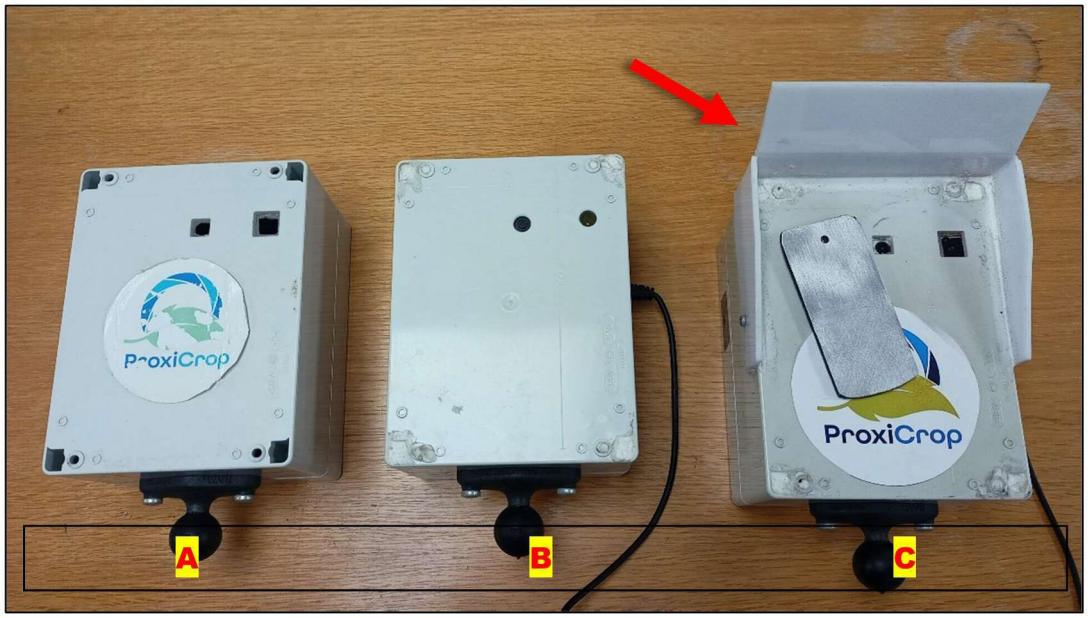


Figure 3.12 - Evolution of the design of the ProxiCrop sensors. A) original design. B) Prototype 1 with protective lenses. C) Prototype 2 with mechanical shutter

Additional development for the 2021/22 season targeted image quality assurance. The aim was to remove 'bad' thermal images that would have affected the calculations of the thermographic data. During field testing it was noticed that, on some occasions, the thermal images from the ProxiCrop cameras were either corrupted or too noisy. In most cases, this was due to water covering the lenses, either from drops falling while the shutter was open or from excess moisture due to morning dew. We considered two approaches to filtering bad data:

Histogram: This method is mainly based on the spread of brightness values by looking at the standard deviation of a curve fitted over the histogram of the image. As shown in Figure 3.13, a good image has a wider spread than a typical bad image. When we plot the standard deviation for a test set of good and bad images, we get the graph below. The x-axis represents the image number. Green are good images; red are bad images. The y-axis shows the standard deviation.

Machine Learning (ML): This method trains an ML system to classify good and bad images. In this method, we attempted to train a Convolutional Neural Network (CNN) to distinguish between good and bad photos. The initial tests only had a success rate of 55%, probably because of the different types of bad images with different features. The proposed way forward with this approach would be to use data from the existing network to identify all the bad images classified by the CNN. These images can then be removed from the dataset and the CNN can then be trained to detect a new type of bad images. Doing this iteratively, a

multi-layered set of CNNs can then be used to identify the different types of bad images.

Alternatively, once a group of bad images is identified, they can be labelled in the training set as a specific type of bad image. The neural network can then be expanded to a multi-class neural net that can classify three groups of images (good, bad type 1, other bad) instead of 2 (good, bad). Doing this iteratively we can expand the number of classes (1 good + multiple bad) that our network can classify.

After exploring both options, we concluded that both methods can work well. However, the histogram method is recommended as the implementation is simple and easy to understand or debug.

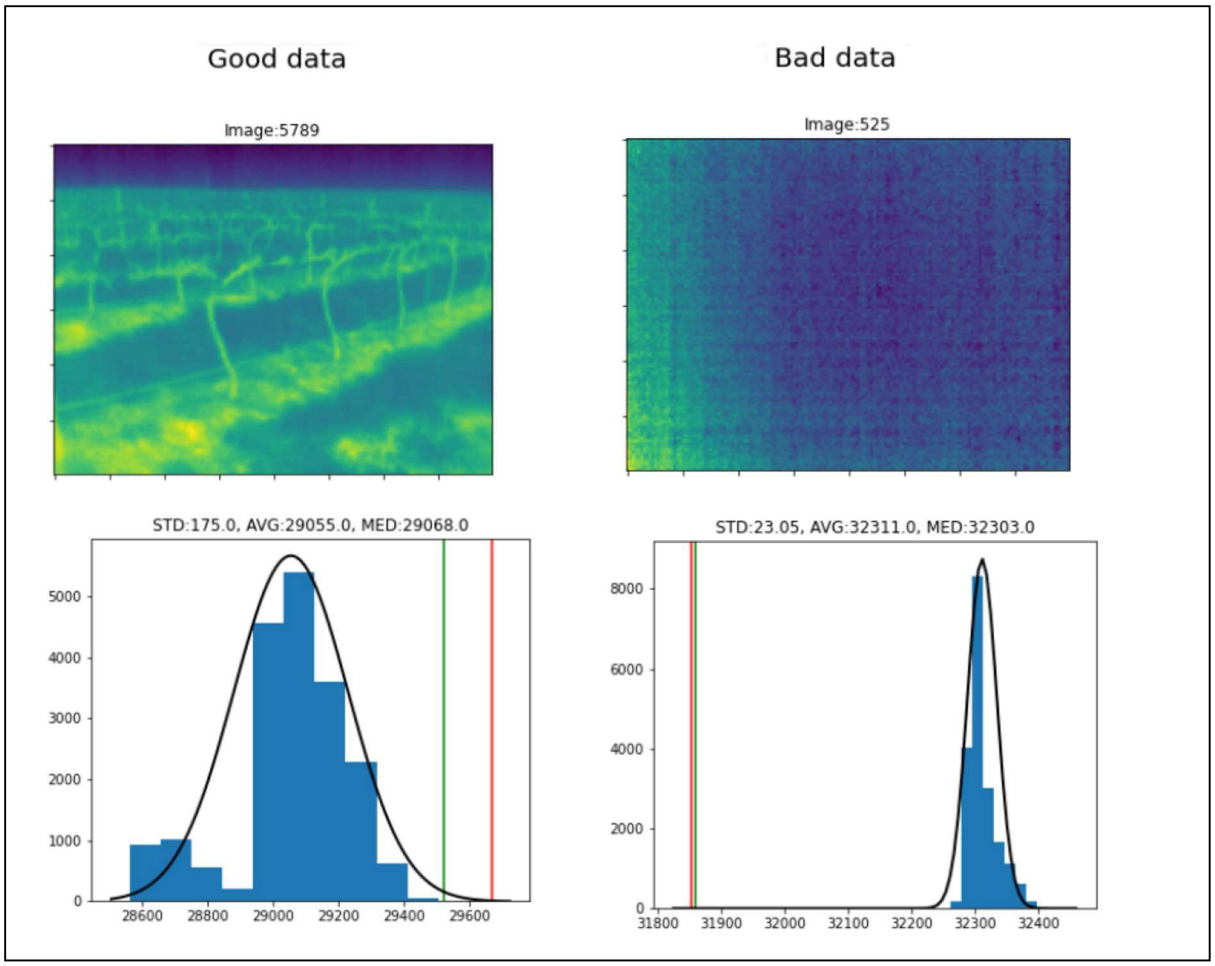


Figure 3.13 – examples of effective and poor thermal images (top) and the histograms for those images (bottom).

3.4 Conclusion

At least ten ProxiCrop devices were deployed in commercial vineyards in each of the 2020/21 and 2021/22 seasons. This allowed for developments to improve reliability as well as testing against soil moisture, atmospheric and plant physiological data to demonstrate the ability of such thermal data to capture changes in plant water status. Although only limited physiological data could be collected due to capacity constraints, the data set could be used as the basis for the development of novel quantitative indices of vine water status.

Chapter 4 – Non-contact determination of fruit composition

4.1 Introduction

Fruit quality and maturation estimation in viticulture is typically labour intensive, prone to error and the information may not be available in a timely manner for decision making. This chapter reports on the potential for near infrared reflectance spectroscopy (NIRS) to be used to non-destructively estimate parameters associated with berry quality that are currently measured via chemical analysis. In particular, we focus on the estimation of berry sugar and its potential to be estimated on-the-go, in the vineyard. We discuss what has been achieved in the lab environment, a lower-cost (non-imaging) spectrometer-based system for deployment in the vineyard, translation of lab models to field models to allow calibrations to be built rapidly and easily and an assessment of these estimation methodologies.

4.2 Laboratory work

In our previous work, near infrared (NIR) hyperspectral images of individual berries paired with ground truth measures of composition were collected over multiple seasons for multiple genotypes (Edwards et al. 2020b). Models were developed from these data to predict fruit composition parameters from the NIR spectra. Edwards et al. (2020) details the berry sampling regime used to form the data sets, the controlled environment imaging system used to collect data, chemical measure ground-truthing, statistical modelling of the data and the results of estimation methodologies up to 2019.

Subsequent to the above-mentioned report, further data have been collected using the same methods to assess model performance across significantly more varieties. An overview of all the data collected is shown in Table 4.1 where for each variety, the number of samples and the type of ground truthing data available for analysis is reported. The OenoFoss measures included TSS, BAUME, pH, Total Acid etc. while the Anth/Phen measures are anthocyanin and phenolics.

The previous report discussed performance in terms of R^2 and RMSE for single variety models and models both tested and trained across the three varieties available at the time (Shiraz and two glasshouses grown varieties).

To assess the robustness of estimation we investigated the performance of models where a variety was omitted from model training and then used exclusively as the test set. This simulates a situation where an unknown variety is exposed to the model and is particularly challenging compared with the single variety models reported on earlier. Table 4.2 and Table 4.3 detail the performance of these models for each of the available varieties over various composition measures. It is clear that models for TSS and BAUME are quite robust to varietal differences with R^2 values greater than 0.8 over all varieties. Models for Total acid, malic acid and anthocyanin also show reasonable robustness to variety with R^2 values greater than

0.5 apart from Total Acid for estimating the Muscat Hamburg variety where R^2 is -1.20.

Table 4.1 - Details of all data collected where imagery and ground truth was available. Shown is the variety, number of berries and if OenoFoss and/or Anthocyanin/Phenolics measurements formed the ground truth.

Variety	# Samples	OenoFoss	Anth/Phen
Shiraz	1440	Yes	Yes
Sauv Blanc	1080	Yes	No
White GH	720	Yes	No
Chardonnay	360	Yes	No
Carignan	360	Yes	Yes
Semillon	360	Yes	No
MuscatHam	360	Yes	Yes
Taminga	360	Yes	No
Mataro	185	No	Yes
Pinot	184	No	Yes
Tempranillo	186	Yes	Yes
Cab Sauv	298	Yes	Yes

Table 4.2 - Performance for estimating OenoFoss based grape composition measures where variety is omitted from training and used only for testing. The values in each cell are R^2 (RMSE).

	Shiraz	Sauv Blanc	White GH	Chardonnay	Carignan	Semillon	MuscatHam	Taminga
TSS	0.97(1.26)	0.94(1.88)	0.82(2.24)	0.92(1.64)	0.95(1.23)	0.87(1.37)	0.92(0.92)	0.91(1.64)
Total Acid	0.67(4.20)	0.72(3.35)	0.44(4.47)	0.87(2.66)	0.62(4.22)	0.66(3.88)	-1.20(6.03)	0.76(3.43)
pH	0.80(0.38)	0.82(0.38)	0.57(0.26)	0.17(0.311)	0.18(0.20)	-1.81(0.36)	-0.43(0.25)	-0.84(0.31)
Malic Acid	0.84(2.86)	0.81(4.04)	0.39(6.11)	0.79(1.87)	0.59(2.82)	0.22(3.54)	-0.62(2.90)	0.67(3.01)
BAUME	0.96(0.84)	0.94(1.08)	0.81(1.24)	0.94(0.88)	0.94(0.71)	0.86(0.87)	0.94(0.57)	0.94(1.0)

When a negative value of R^2 is observed, the interpretation is that the mean of the data is a better predictor than the model. With this in mind, estimation of pH is good for some varieties (Shiraz, Sauv Blanc) and poor for others. An observation in a number of cases where R^2 is low or close to zero is that the observed and estimated values are reasonably correlated but shifted from the one-to-one correspondence i.e. have a large bias.

Table 4.3 - Performance for estimating anthocyanin related composition measures where variety is omitted from training and used only for testing. The values in each cell are R^2 (RMSE).

	Shiraz	Carignan	MuscatHam	Mataro	Pinot	Tempranillo	Cab Sauv
Anthocyanin	0.66(0.32)	0.75(0.22)	0.84(0.14)	0.65(0.20)	0.59(0.29)	0.74(0.27)	0.73(0.31)
Phenolics	0.11(0.27)	0.49(0.18)	0.45(0.11)	0.21(0.15)	0.20(0.23)	0.48(0.19)	0.56(0.22)

This method of assessment is likely penalising the performance given that in practice we would use all varieties in a training model and new data may well come from the varieties we have studied.

For the remainder of this report, the lab models discussed refer to models trained on balanced samples across all varieties available. Figure 4.1 shows observed versus estimated values for both TSS and Anthocyanin using LASSO models across all available berries. These models maximised the use of our database. Specifically, for TSS we used 360 berries from each variety to train the model where 360 was chosen as it is the size of the smallest variety classes (Taminga, Semillon, MuscatHam, Carignan and Chardonnay). The size of the training sample is 2880 and the plot shows all 5040 points available, which include the training sample. The corresponding R^2 and RMSE for the TSS model is 0.94 and 1.50 respectively, however these values should be treated with caution given they include the training sample. Similarly for anthocyanin, 185 points were sampled from each varietal class as the smallest class is 185, giving us 1295 training samples. The plot shows all 2470 samples from our database with corresponding R^2 and RMSE of 0.80 and 0.25. Again, these values are calculated using all available data, including the training sample, and should be interpreted correspondingly. R^2 and RMSE should be calculated on data not seen in the training process, but we have elected to show these values for our final models as omitting the training data removes more than half of the grape varieties.

These models represent our final models based on all data available to us. They are deployable on new data captured with the lab hyperspectral system and form the basis for field models described in the remainder of this document.

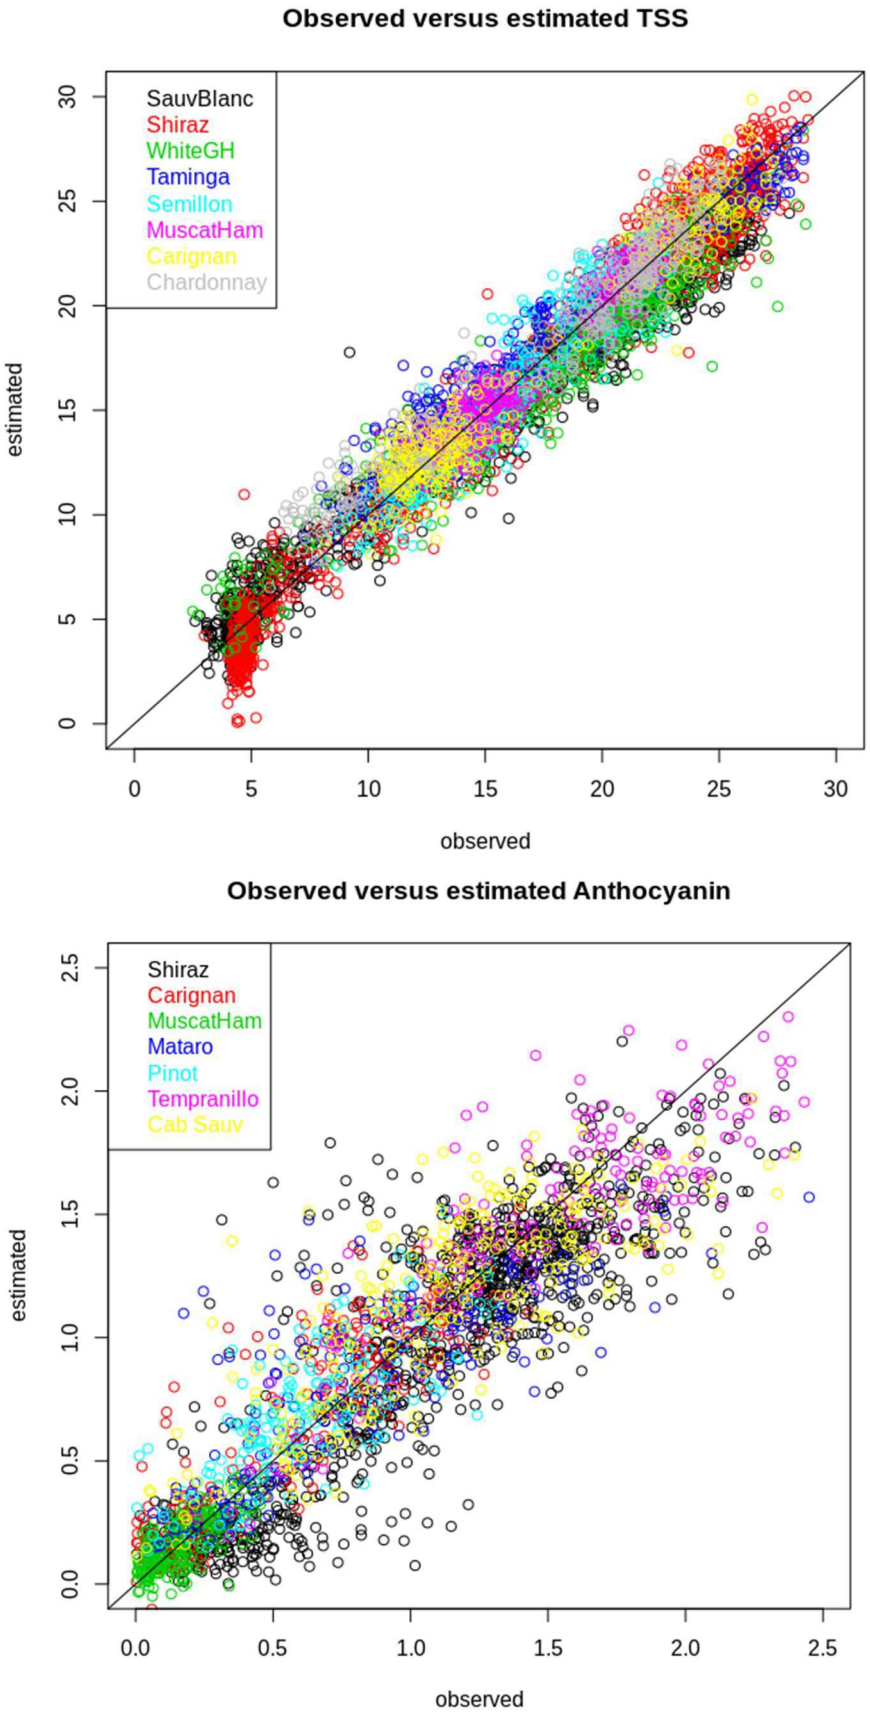


Figure 4.1 - observed versus estimated TSS and Anthocyanin from models using all available varieties

4.3 Field data collection

Data collection in the field was carried out using a dual point spectrometer system. The point spectrometers were Hamamatsu C13053MA and C14486GA covering wavelength ranges 500-1100 nm and 900-1700 nm respectively. Combined, the two spectrometers cover the wavelength range of the lab imaging system and also have a higher spectral resolution (<5 nm). Illumination was provided by two 500W halogen lights during the 2021 season. The lighting was chosen to minimise the impact of sun on the captured spectra. The lighting was changed to two 1000W halogen lights in the 2022 season, as it was determined that more illumination was required. The system is shown in Figure 4.2 attached to a Kubota.

The remainder of this section details data collections obtained in the field. Over two seasons, three rows from Accolade's McLaren Vale vineyard were utilised for collecting data. Each row corresponded to a different variety, namely Shiraz, Mataro and Fiano. Data collection was carried out approximately weekly during the growing season.



Figure 4.2 - Point spec system attached to a Kubota showing lighting and point spectrometer optic.

4.3.1 Field bunch scan

At each time point, individual bunches were tagged in an approximately even spacing along each of the rows, corresponding to every 8th panel. The point spec system was moved to each tagged bunch, where the bunch was then removed from the vine and scanned by the system by slowly moving and rotating the bunch while the spectrometers acquired data over an approximately 10 second period. The bunches were then bagged for further analysis in the lab. An example tagged bunch is shown in Figure 4.3. White and dark references were taken prior to each scan.

Over the course of two seasons (2020/21 and 2021/22), scans of 360 bunches were captured in this manner. Capture was carried out for three rows from the Shiraz, Mataro and Fiano varieties at approximately weekly time points.



Figure 4.3 - An RGB image of a tagged bunch while intact on the vine.

4.3.2 Field row scan

After the individual field bunch scans were collected, full row scans were performed as the Kubota mounted point spec system was moved along the rows. The point spec system was focussed into the bunch zone. The system collected data at approximately 30 points per second. For each row, the number of points collected was in the order of 10000. White and dark references were captured before and after each row scan. Weekly scans of six rows were captured in this fashion. Three of the six rows corresponded directly to those where field bunch scans were carried out, while the remaining three were neighbouring rows of the same variety. Figure 4.4 shows a single frame from a GoPro video that was captured concurrently with the point spec scanning. The range of the lighting illustrates where in the bunch zone the point spec is capturing data.

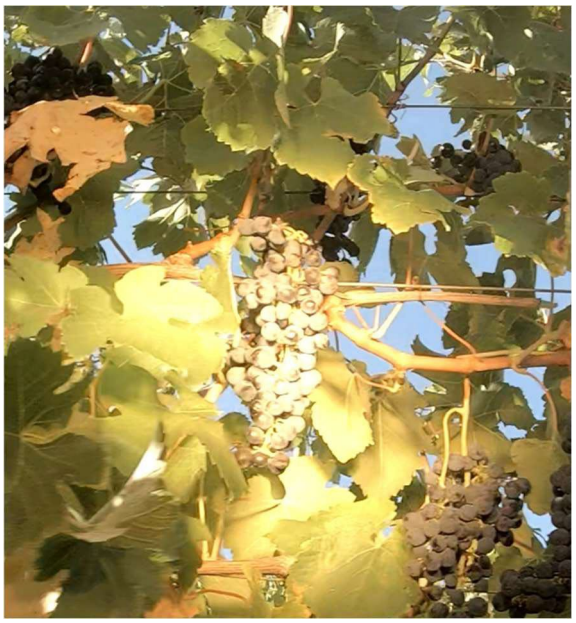


Figure 4.4 - Frame from GoPro video captured concurrently with point spec field row scan.

4.3.3 Laboratory bunch scan

Each of the bunches that were tagged and scanned with the point spec system were also scanned with the lab imaging system. These bunches were then subjected to the OenoFoss system to give us TSS measurements. This was done for the 360 bunches that were scanned in the field. Figure 4.5 shows a single wavelength band of a bunch that has been tagged in the field, scanned by the point spec system, scanned in the lab with the imaging system and finally measured for TSS using the OenoFoss.

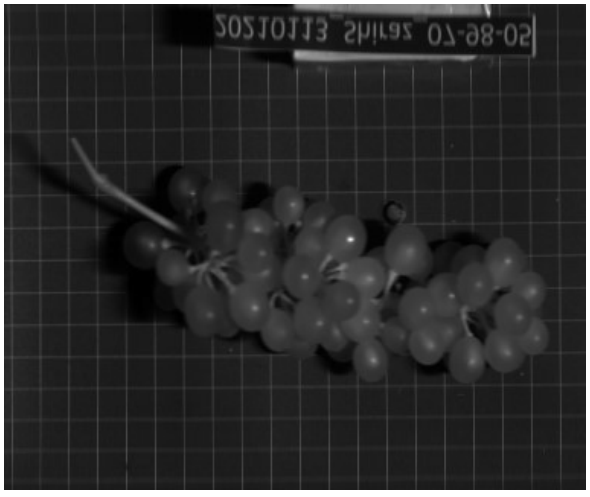


Figure 4.5 - A single band from a hyperspectral image of bunch.

4.4 Data analysis

This section reports on the analysis of the field point spec data in the context of estimating sugar (TSS). We concentrated on this aspect as the results of the lab work show TSS to be the best estimated compositional attribute.

4.4.1 TSS estimation from lab bunch scans

All 360 lab bunch scans were pre-processed to extract a single representative spectrum for each bunch. Each hyperspectral image was treated as outlined in Edwards et al. (2020), pixels identified as berry tissue were extracted and a median at each wavelength calculated. The representative spectra were then passed to the LASSO TSS estimation algorithm discussed at the conclusion of the Laboratory work section above. Figure 4.6 shows plots of observed and estimated TSS values for 2021 on the left and 2022 on the right. The R^2 values are 0.92 and 0.90 while RMSE is 1.6 and 1.8 for 2021 and 2022 respectively. These values are similar to those in Table 2, noting that neither Fiano nor Mataro were varieties used to build the lab model. There does, however, appear to be a slight bias in the estimated values for Fiano with those values being consistently below the observed.

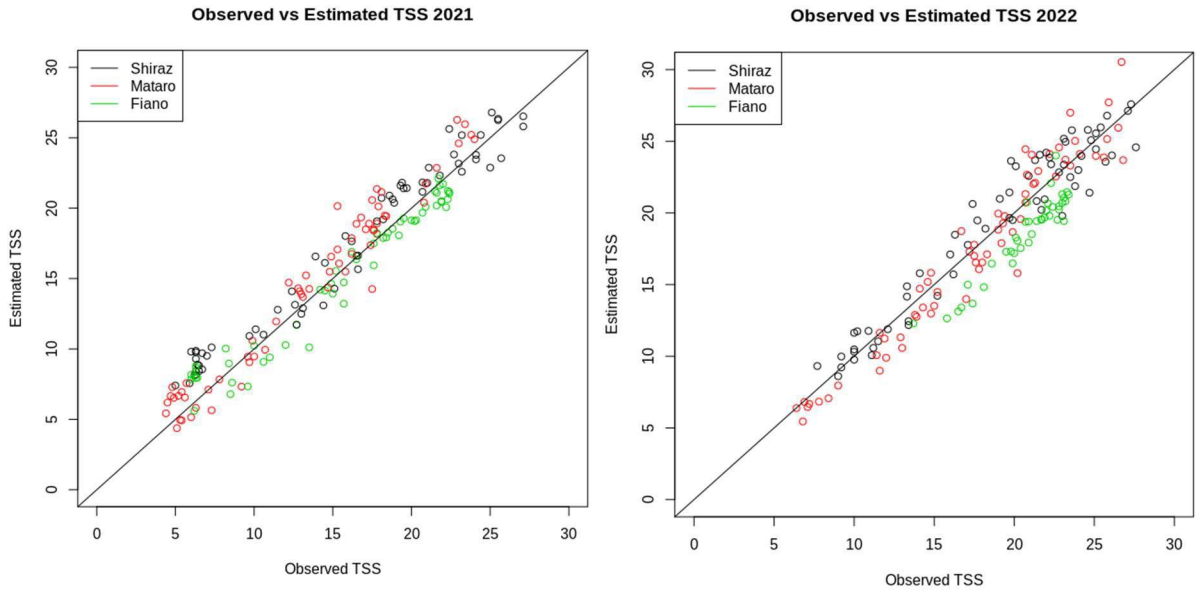


Figure 4.6 - Observed versus estimated TSS for lab bunch scans in 2021 and 2022.

4.4.2 Transform from field spec to lab spec

To make use of our rich lab data set we needed to establish the correspondence between our field bunch scans and lab bunch scans, so we aimed to produce a representative spectrum from both the lab and field for each of the bunches.

Over the 10 second period the field bunches were exposed to the point spec setup, the bunches were rotated, so at some points the spectra produced were not berries. Figure 4.7 shows three example spectra from a single 10 sec scan. Our aim is to find the spectra that are from berries and take their median to produce a representative spectrum for that bunch.

This was done by clustering the berries into three groups and determining the cluster that is berry spectra (they have the shape of the black spectra in Figure 4.7).

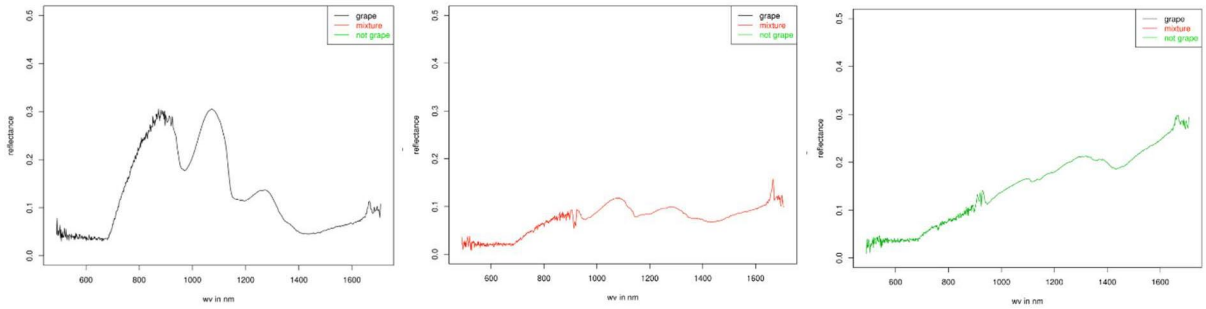


Figure 4.7 - Three example spectra produced in a single 10 sec period when producing field bunch data.

An example of all the grape spectra for a single bunch is shown in Figure 4.8. Here the variation in reflectance due to illumination gives reason to standardise the spectra before taking the median to find the representative spectrum for this bunch).

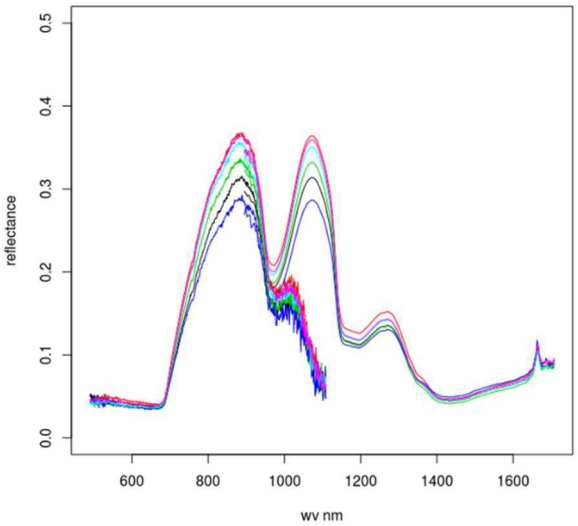


Figure 4.8 - Spectra identified as grape from a single field bunch scan.

As with the field bunch spectra, we need to create a representative spectrum for each of the bunches imaged in the lab. The process for identifying grape spectra in an image was established in through the preceding CSA 1602 project (Edwards et al. 2020). Those pixels identified as belonging to grapes can then have their median spectrum taken over each wavelength to form the representative spectrum.

Examples of representative spectra for a single bunch from both the field and lab scans are shown in Figure 4.9. The panels show reflectance for representative spectra from the field in

grey, and the lab in black, for the two wavelength ranges corresponding to the point spectrometers. It is clear that there is considerable difference between lab and field spectra.

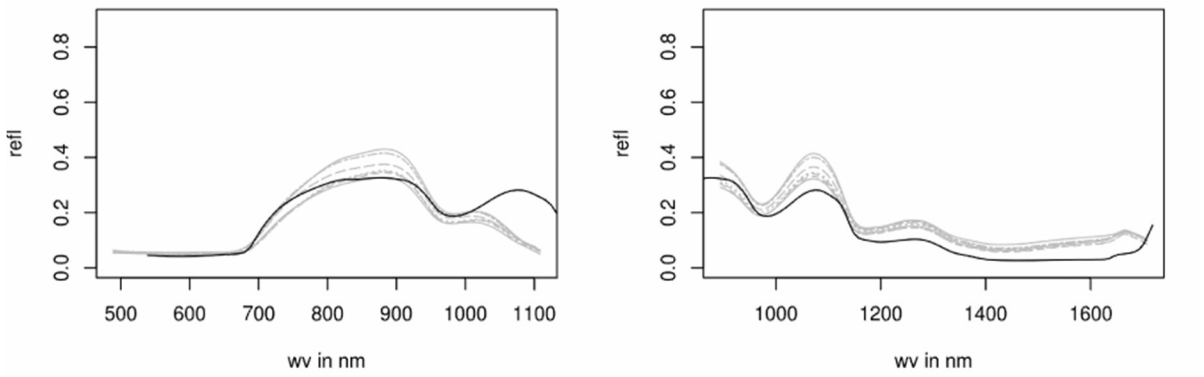


Figure 4.9 – Reflectance spectra from the two spectrometers (left and right). Field bunch spectra are in grey, while lab bunch spectra are in black.

Our lab setup covers the wavelength range 550-1650 nm while our two field spectrometers cover the ranges 500-1100 nm and 900-1700 nm. Recall also, that our lab-based models are based on standardised spectra over their wavelength range (to constrain illumination variation, while retaining shape). Because the spectrometers operate independently, this can not be done reliably with the field data. Instead, we standardise the data from each spectrometer separately. We can re-standardise the lab data so it has similar properties. This is done by splitting the lab spectra at 1000 nm and scaling, independently in the regions above and below 1000 nm. The splitting point, 1000 nm, was chosen as it is the middle of the overlap of the point spectrometer ranges. Figure 4.10 illustrates this idea. On the left is an example spectrum where its reflectance was standardised using the lab wavelength range while on the right, the standardisation was applied independently across the wavelength ranges of the two point spectrometers. This means the lab models need to be rebuilt on the re-standardised lab data. After testing, very similar estimation performance can be achieved.

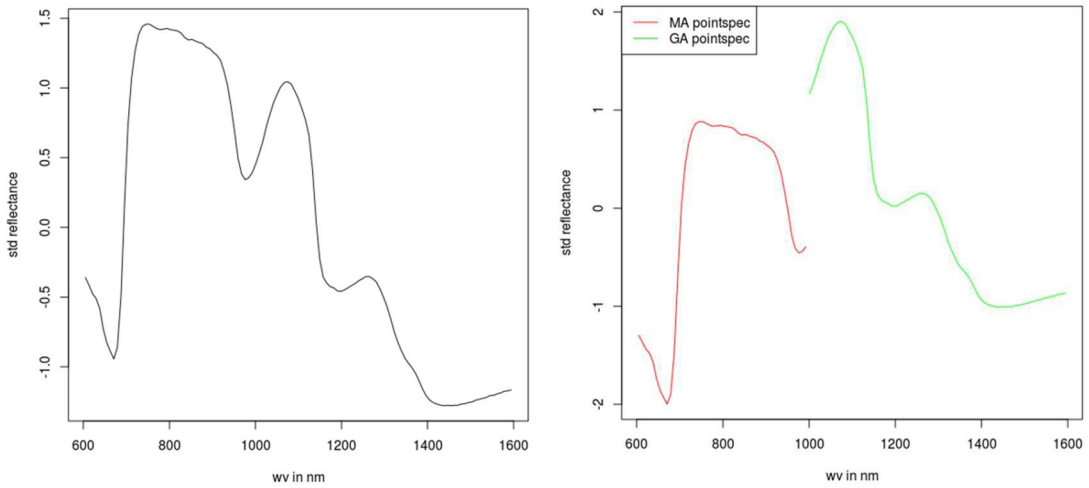


Figure 4.10 - Lab spectra standardised in the normal fashion (left) and standardised by splitting and independent standardisation (right).

The independent standardisation process was then applied to our representative field and lab bunch spectra. Figure 4.11 shows the results of the standardisation applied to the example in Figure 4.9. The correspondence is much tighter and would be enough to drive a classification type process based on a model built from our lab samples. However, the correspondence is not good enough to use for estimating sugar.

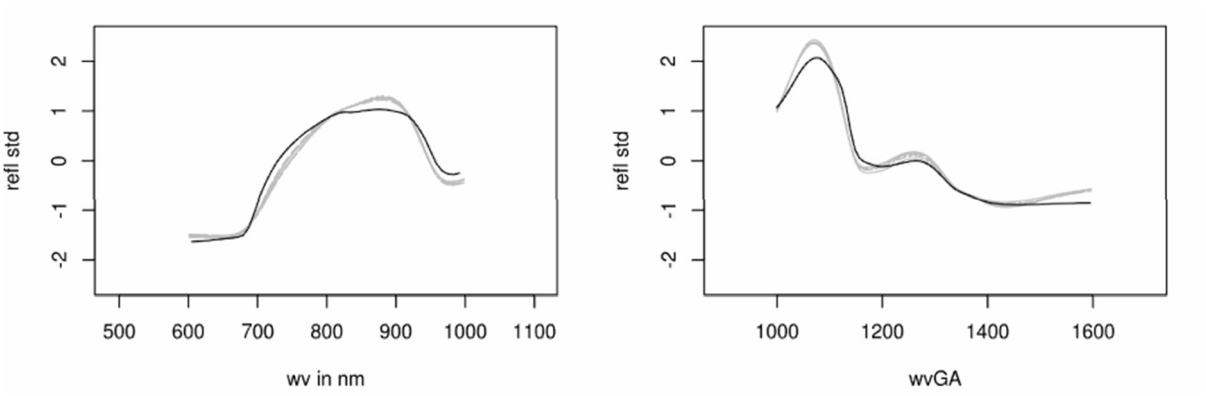


Figure 4.11 – Standardised reflectance spectra from the two spectrometers (left and right). Field bunch spectra are in grey, while lab bunch spectra are in black.

To arrive at spectra with a better correspondence between field and lab bunch spectra we investigated techniques for transforming the standardised field spectra. We looked at a simple univariate linear regression at each wavelength and two multivariate random forest modelling approaches. The random forest approaches both used the values in the field spectra at specific wavelengths from the field spectra as predictors for each of the

wavelengths for the lab spectra. The predictor wavelengths were 604, 679, 745, 885, 984, 1074, 1198, 1264, 1421 and 1545 nm, which are turning points in typical berry spectra. The two random forest models differed in that one model also used the current wavelength in the model i.e. when modelling lab spectra at wavelength 907 nm, the same wavelength of the field spectrum was also included along with the above-mentioned wavelengths.

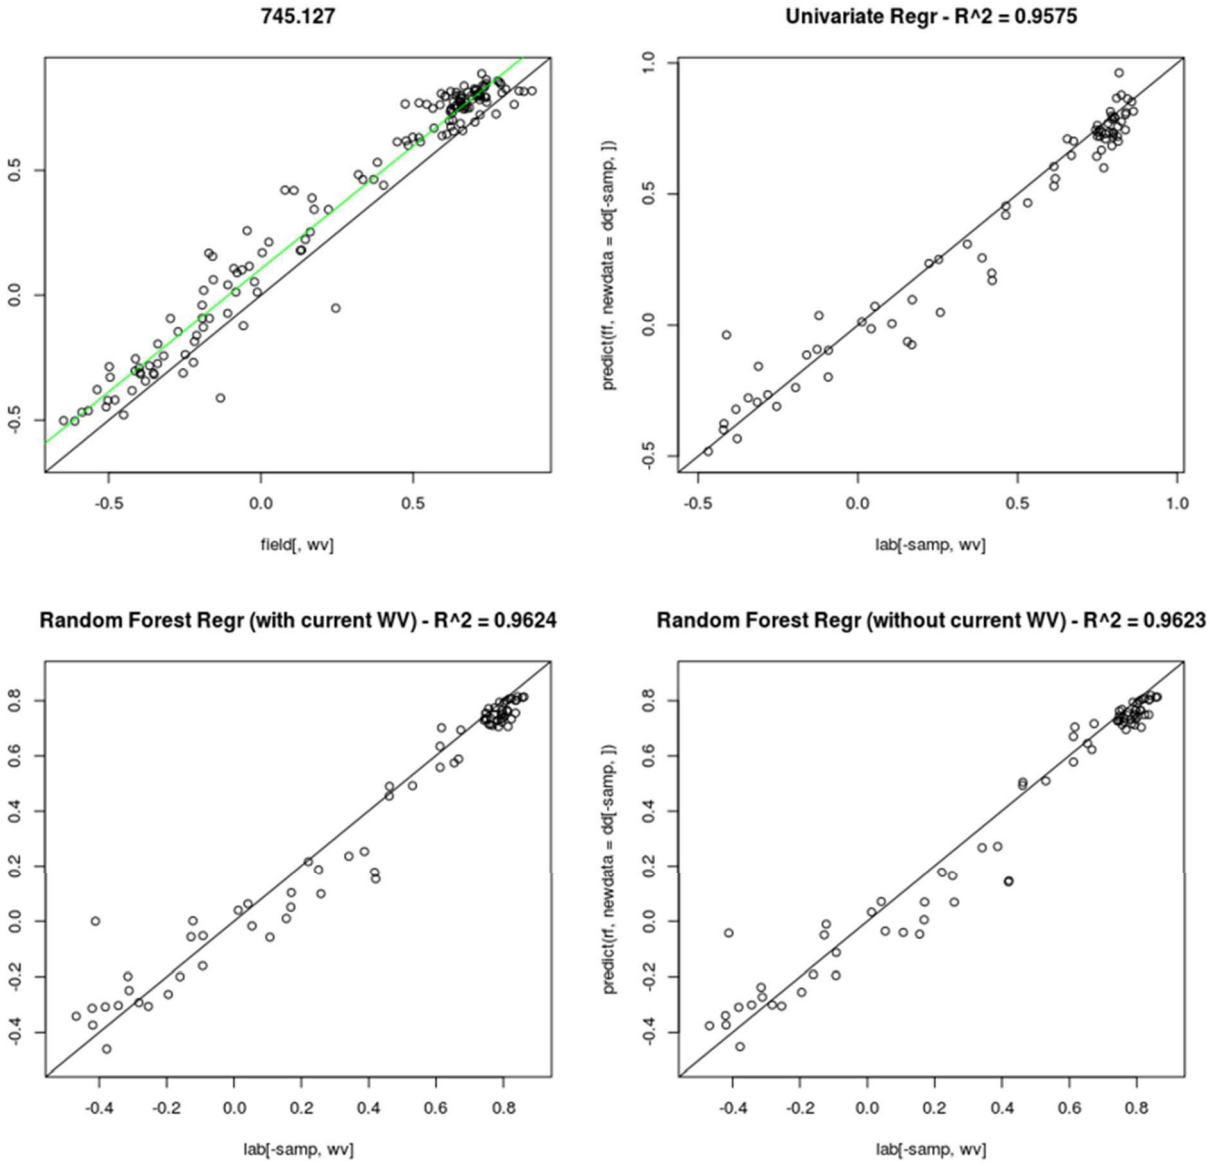


Figure 4.12 – Plots illustrating transform modelling at wavelength 745 nm

At the time of writing, the transform models were only explored on the 2021 data. A number of early season representative spectra were also omitted from the models as the data was generated while the capture process was being fined tuned and hence were considered unreliable. The number of representative spectra used was 142. Figure 4.12 shows four plots to illustrate the modelling concept. The top left panel plots the field vs lab values of representative spectra at wavelength 745.127 nm. The green line is a line of best fit showing

a consistent shift across all representative spectra at that wavelength. This shift can be seen in the left panel of Figure 4.11 where the grey spectra sit consistently below the black for that example. The remaining panels in Figure 4.12 show the results of the three transform models where the x-axes are what we observe in the lab at 745.127 nm and what we estimate it to be from the models. In this example, the univariate regression and random forest models all perform well which is expected given the strong linear relationship between the field and lab representative spectra. Figure 4.13 shows the equivalent set of plots for 1132.7 nm. Here the line of best fit is almost horizontal and the univariate regression essentially estimates the mean. The random forest approaches, however, offer more information and produce estimates that improve the relationship between observed and estimated.

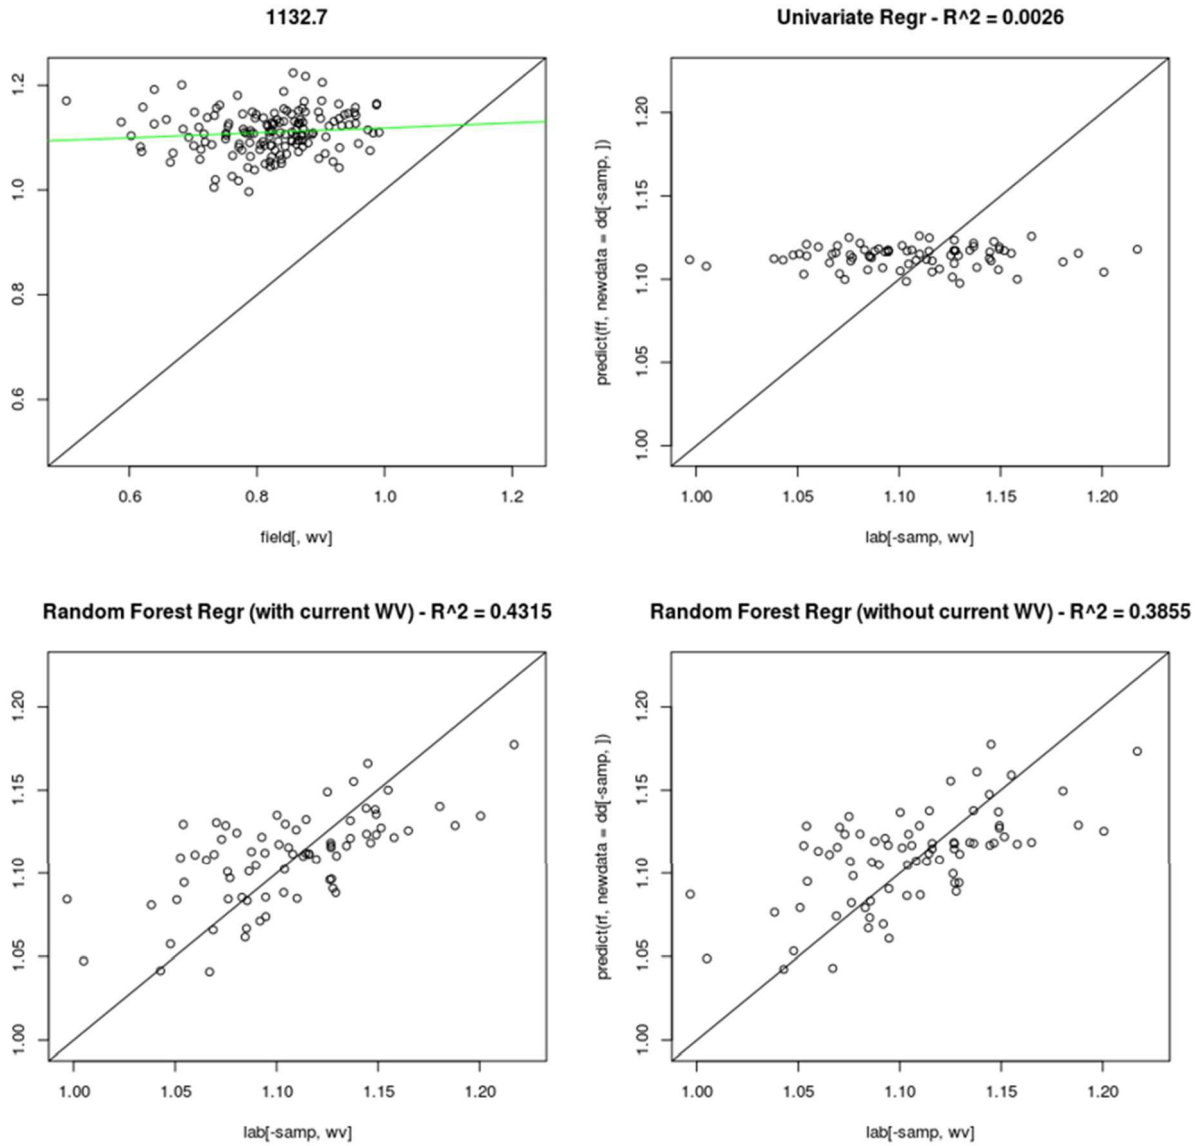


Figure 4.13 - Plots illustrating transform modelling at wavelength 1132 nm

The transform models can now be applied at every wavelength of field spectral data. Figure 4.14 shows two examples of field spectra that have been subjected to the three transform models. The green curves show the input field spectra while the black show our target lab spectra. The blue is the resulting spectra for the univariate regression approach while the orange and red show the random forest results. The left panel shows a result where the correspondence between the lab and both the random forest approaches is particularly good while the right panel shows an example where the random forest approaches produce reasonable but somewhat noisy results. In both cases the univariate regression approach does not yield good results, particularly in the upper wavelength ranges.

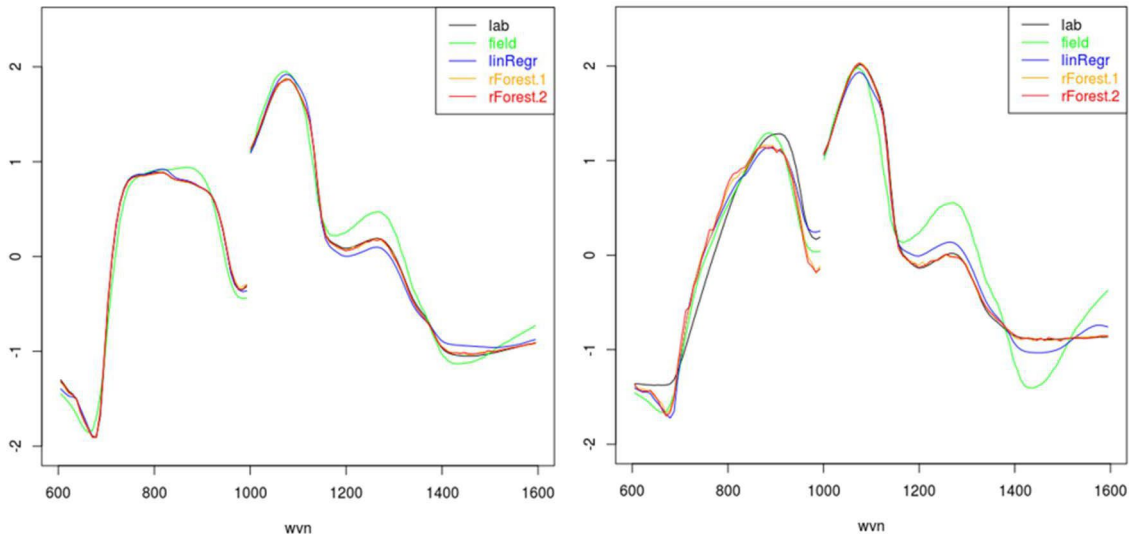


Figure 4.14 – Examples of transformed field spectra using the 3 models

As discussed, at the time of writing, only the 2021 season's data have been used to develop these models. It is noted that early season data are under-represented due to experimentation with the point spec system where data points were removed. Further to this, the spectral shapes for Mataro and Shiraz (red varieties) follow similar patterns compared to Fiano (white varieties), so white grapes may not be represented well. This is coupled with the fact Fiano was not collected at all in the first few weeks of the 2021 season collection.

Further work needs to be carried out to integrate the 2022 data and to try to mitigate Fiano/white under-representation.

4.4.3 Estimation of TSS for bunch scans

The field to lab transformation can now be applied to the representative bunch spectra captured in the field and in turn the LASSO model for estimating TSS built from our lab database. For the 2021 data, instead of using all 142 data points to train the lab transform model, we used only 71 points to avoid over-fitting when assessing the performance of estimating bunch TSS for 2021. Figure 4.15 shows plots of observed versus estimated TSS with

the left panel corresponding to the 2021 data. Colours correspond to the three varieties. All 142 points are plotted, with triangles representing points included in training the transform model, while circles are points that were not in the training set. The R^2 for the points outside of the training was 0.82 with an RMSE of 3.3. The right panel of Figure 4.15 corresponds to the 2022 data. Here, no data from 2022 were used in training the transform model. The most obvious issue is the estimation of TSS for the Fiano variety. In the previous section we discussed Fiano's under-representation in the transform models and consider this to be the cause of its poor result. If Fiano is included in calculations of R^2 for the 2022 data, $R^2$is 0.23, if it is excluded, the value is 0.79. For both seasons the results tend to overestimate TSS when the actual values range from 15-20, while for 2022, TSS is underestimated when the actual TSS is high, above 24. However, given the results for 2022 were generated from transform models trained on the previous season and estimation models based on lab data, this result is encouraging.

Further work needs to be carried out to improve the modelling. Including the 2022 data in the transform models should be investigated to determine whether it improves performance for Fiano and the under- and over-estimation issues discussed above.

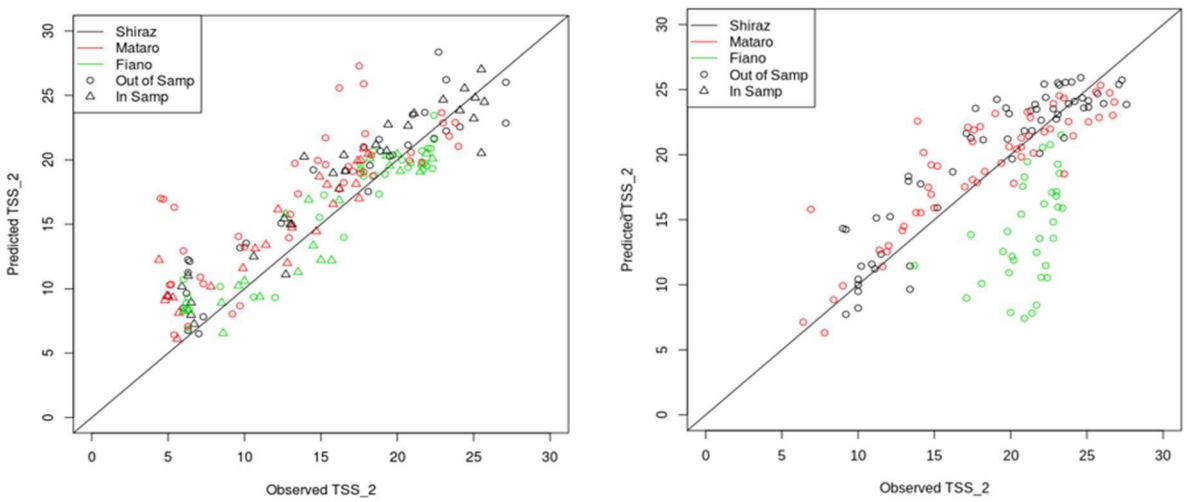


Figure 4.15 – Observed vs estimated TSS from field transformed spectra. 2021 is on the left, 2022 on the right

4.4.4 Estimation of TSS for row scans

To assess the performance of estimating TSS across whole rows, pre-processing is required before passing the spectra to the field to lab transform. The pre-processing has two aims, firstly to decide which white target to use for transforming to reflectance given a white target scan is taken before and after each full row scan. The second aim is classification to decide whether a spectrum is from a berry or some other material i.e. leaves.

For classification, spectral angle mapping (SAM) is used. SAM can be thought of as a similarity score between two spectra. The 2021 field bunch spectra, in standardised reflectance, were clustered into five groups using a hierarchical clustering method and the group centres calculated. Each spectrum from a full row scan was converted to standardised reflectance twice, once each for the two white targets. SAM was then calculated for reach of the five cluster centres. If the maximum score was over 0.995, the spectrum was deemed to be of a berry. The maximum score was determined empirically. Using the maximum also determines the appropriate white target as it suggests we achieve the most berry-like spectrum. Two examples of classified spectra are shown in Figure 4.16 with the right panel showing a successful berry classification.

This classification could be further refined by introducing the 2022 season's data into the clustering process, which may address issues with Fiano's under-representation along with estimation performance at high TSS values.

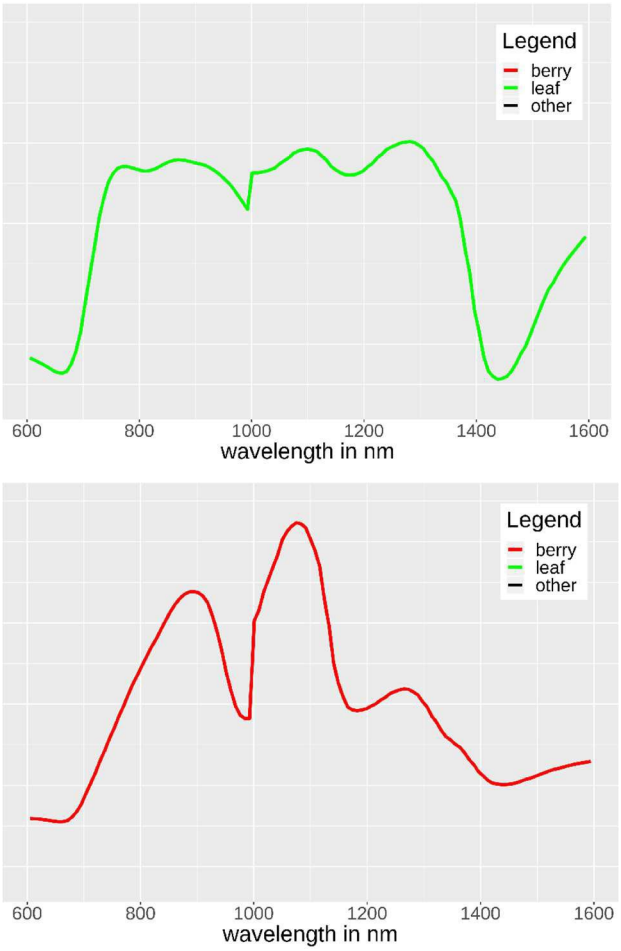


Figure 4.16 – Plots illustration spectral classification for finding berry spectra in row data

All standardised reflectance spectra from the row that have been classified as berry can then be passed on to the field to lab transform and further to the LASSO model for estimating TSS.

Figure 4.17 shows estimated TSS for the Mataro variety row monitored across the 2021 season. The vertical panels in the plot correspond to the five time points where row data was captured. The time points are labelled in green, so for example, 20210127 is 27 January 2021. Within each panel, the black circles correspond to estimated TSS for spectra classified as berry (in the order of 60 points per time point). The red squares are the observed TSS measurements for bunches collected and analysed in the lab. It should be noted, there is no direct correspondence between the position of the red squares relative to the black circles as we have not tried to identify tagged bunches with the stream of row data. They should be interpreted as a guide to how the estimation of TSS within a time point is tracking. Generally, estimated TSS is increasing across the season from time point to time point. TSS is however being overestimated for the first four time points, which is consistent with our observation on the field bunch data for TSS values between 15 and 20. It is also noted that observed TSS, the red squares, is largely constant between 3 and 10 February, and 16 February apart from a single high value, is also similar. Another consideration is that the row data represents the first trial of the system, with lower powered lighting compared with 2022. The observed weekly increase in estimated TSS in a plausible range is a positive outcome despite the bias towards overestimation.

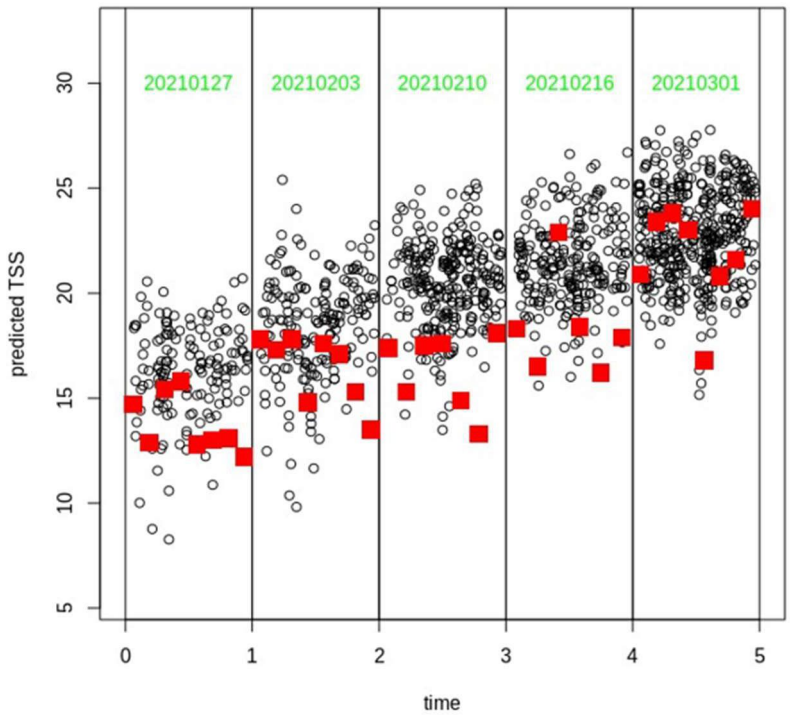


Figure 4.17 – Estimated TSS from row data, Mataro 2021 in black, Red represents ground truth

Figure 4.18 shows the equivalent plot of estimated TSS for Mataro across the 2022 season. Again, estimated TSS increases week on week during the season. There is also generally better correspondence between the estimated TSS values and the observed field bunch TSS. Encouragingly, the observed values of TSS fall within the range of the estimated TSS values at all time points apart from the last time point, 15 March, where we underestimate TSS. The range of estimated TSS on 8 February is large and at this stage an explanation is unavailable, however the mean within this time point is similar to that for the observed values. Overall, this plot demonstrates the ability of the system to estimate TSS across a season, noting that no data collected in 2022 were used in any aspect of training the algorithms.

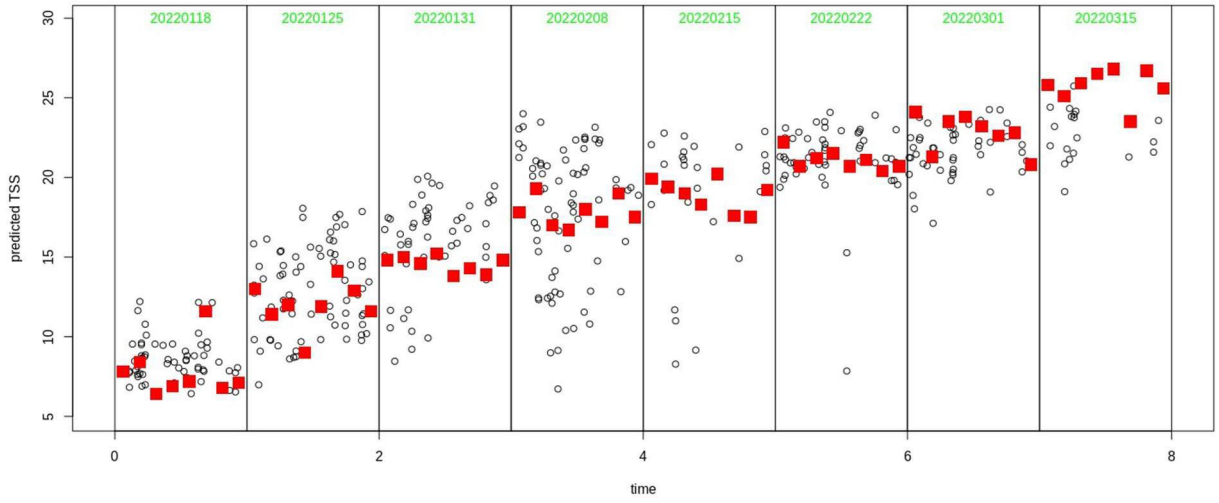


Figure 4.18 - Estimated TSS from row data, Mataro 2022 in black, Red represents ground truth

Figure 4.19 shows estimated TSS for the Shiraz variety row monitored across the 2021 season. Only four time points were captured during this season. Similarly, to the Mataro results for the 2021 season, there is a tendency toward overestimation for TSS in the 15-20 range. Again, encouragingly, estimated TSS is increasing over the four time points in a plausible fashion with the observed TSS falling within the range of the estimated TSS within each time point.

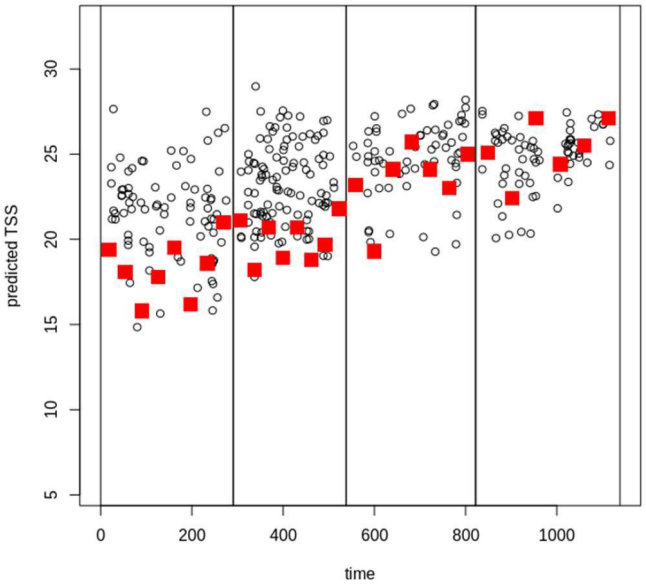


Figure 4.19 - Estimated TSS from row data, Shiraz 2021 in black, Red represents ground truth

Figure 4.20 shows the equivalent plot of Shiraz for the 2022 season. Again, estimated TSS increases week on week during the season. There is also good correspondence between the estimated TSS values and the observed field bunch TSS compared with the 2021 results. For 2022, there were eight separate time points covering more of the early season. Interestingly, 31 January has a large variation of estimated TSS and is likely timed with veraison which makes the variation plausible. The TSS estimates at the final time point seem to be biased down which was similarly noted for the field bunch data. This plot is a largely convincing demonstration of estimating TSS across a whole season. Again we note that no data collected in 2022 were used to train any components of the algorithms used to estimate TSS.

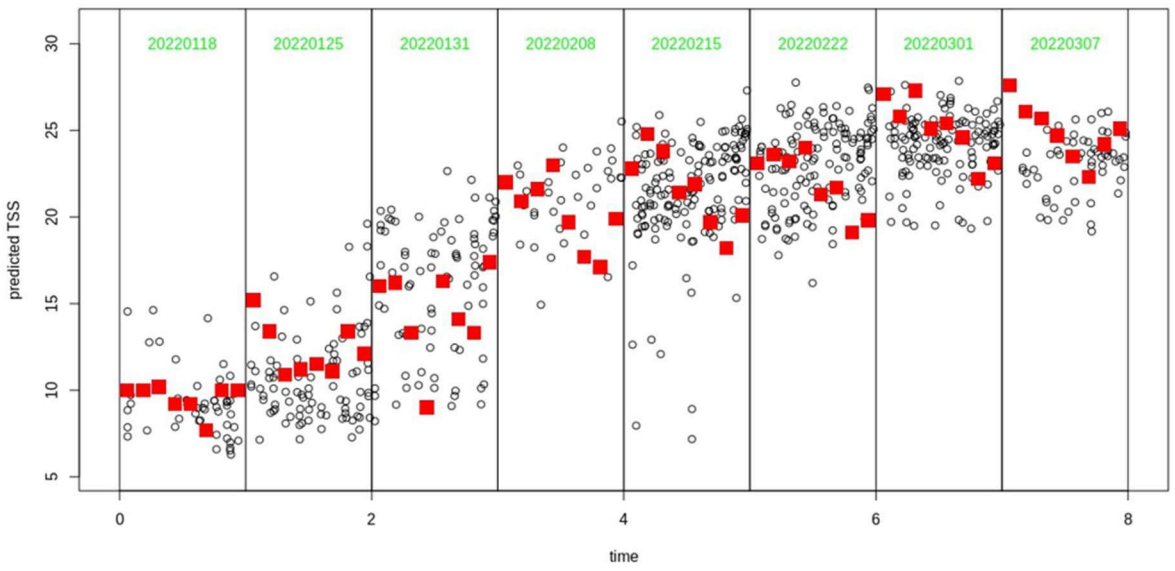


Figure 4.20 - Estimated TSS from row data, Shiraz 2022 in black, Red represents ground truth

The results for the Fiano data sets were generally poor and are omitted from detailed discussion. Investigation into the cause for this is yet to be carried out, but it is reasonable to conclude that Fiano is under-represented in the 2021 data and that including the 2022 data is likely to improve estimation performance.

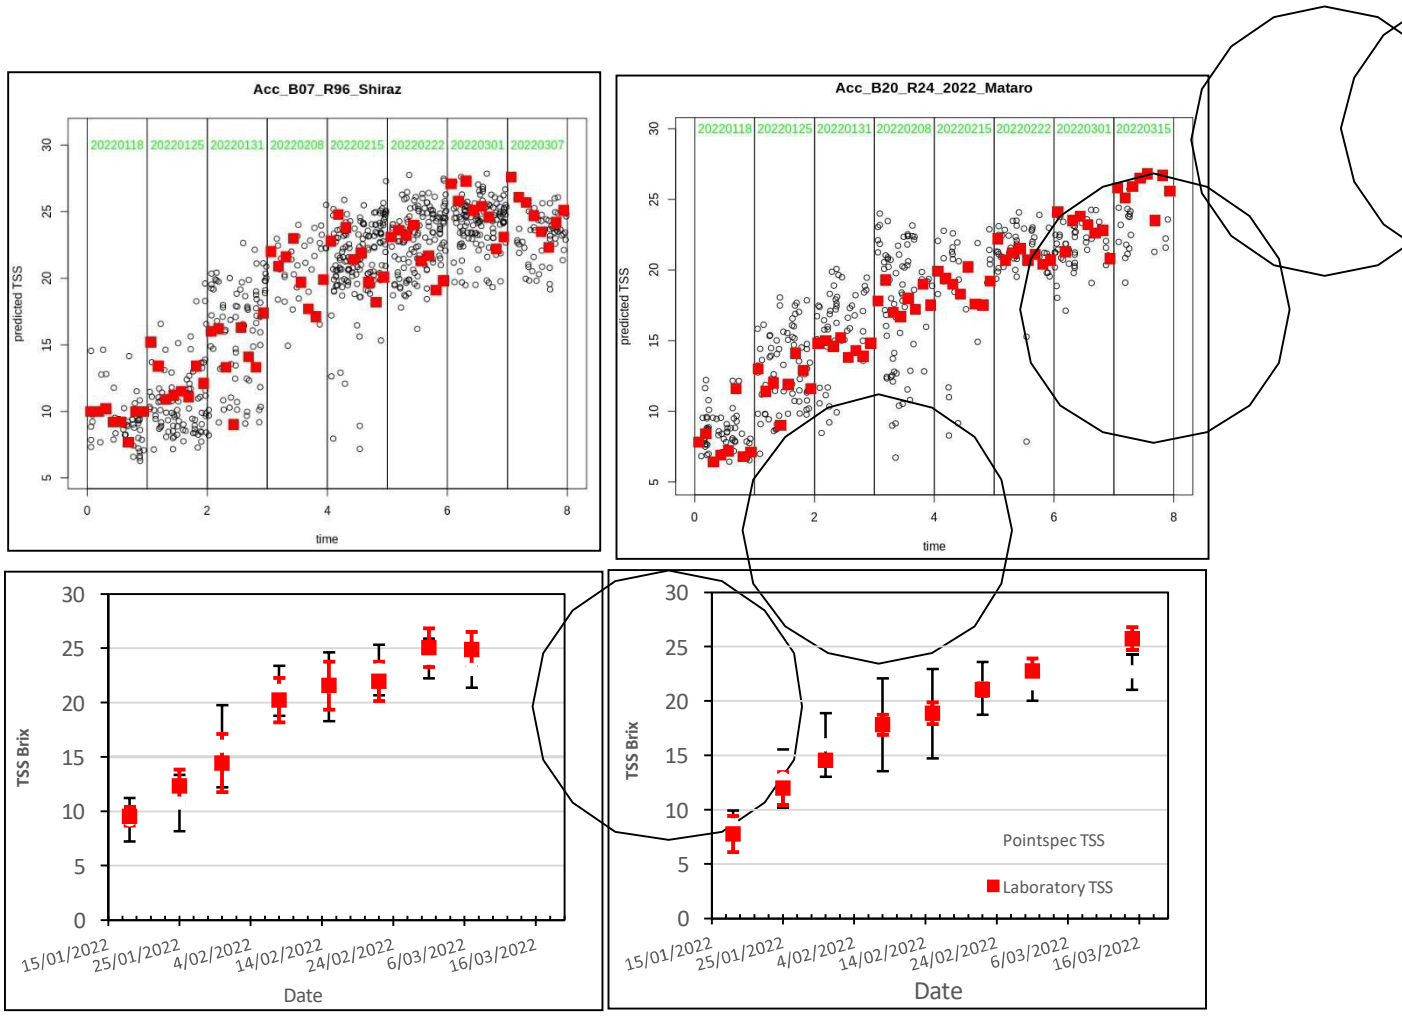


Figure 4.21 – Top row: estimated TSS from row data, Shiraz and Mataro 2022 in black, Red represents ground truth. Bottom row: The same data represented as box and whisker plots.

4.5 Case study: in-field assessment of fruit maturity at Barossa Valley Estate

4.5.1 Overview

Effective scheduling of harvesting and winery intake depends on accurate knowledge of fruit maturity. At best, this is assessed by a manual sampling of a few bunches at a few locations in a block, with significant variation across a typical block. Further, in-field assessment is normally restricted to sugar maturity, and other measures require laboratory analysis.

As part of the Yield Estimation Project in collaboration with the Barossa Valley Estate winery, CSIRO has tested the capability of the point spectrometer for the assessment of fruit maturity prior to harvest at the Krondorf Road vineyard (KRR) in the Barossa, SA.

4.5.2 Data Collection at the BVE site

CSIRO Point Spectrometer data

On 23 March 2022, the point spectrometer system (PS) was mounted on a Kubota RTV vehicle and run in six selected rows in each of the two experimental blocks chosen for the BVE project: KRR-KRR-A-SH -SH and KRR-KRR-D-SH -SH (Figure 4.22). In both blocks, the grapevines were of Shiraz variety trained on a Scott-Henry system.

For the data collection, rows 18, 21, 22, 24, 26 were selected from Block KRR-A-S, while rows 21, 24, 25, 27, 29 and 30 were selected from block KRR-D-SH. The PS runs in each row in the west-east direction, following the same order of the row panels numbering. As the PS was mounted on the left-hand side of the RTV, all spectral data were collected from the south-facing side of the rows. During each row run, the PS was paired with and GoPro camera providing GPS location of the PS data at any point and assisting with a visual of the area of data collection. This additional information was used in conjunction with the spectral data from the PS to estimate the mean total soluble solids (TSS) from each panel within the selected rows (as ^{o}Brix, subsequently converted to ^{o}Baume).

Moreover, it was estimated in that the PS detected 30 bunches per panel on average.

BVE fruit collection and analysis

The day after the data collection by CSIRO, BVE technical staff collected all the grape bunches from single vines representing selected row panels. These vines were in panels 7, 15, 23, 31 and 39 from rows 21 and 22 in Block KRR-A-SH and panels 7, 15, 23, 31, 39 and 47 from rows 24 and 25 in Block KRR-D-SH. Average sugar content (^{o}Baume) from each vine was therefore measured after the bunches were fermented for 24 hours.

Furthermore, 24h after harvest two samples from the bulk crushed and fermented fruit were also collected by BVE staff for the analysis of the mean sugar content of blocks KRR-A-SH and KRR-D-SH (^{o}Baume).



Figure 4.22 - Overview map of the KRR vineyard in the Barossa, showing the locations of the two experimental blocks: KRR-A-SH and KRR-D-SH

4.5.3 Results

Average sugar content in blocks KRR-A-SH and KRR-D-SH

Results from the sugar contents analyses and estimates (°Baume) at the KRR vineyard are presented in Table 1.

The average sugar contents from the bulk analysis were high, with 15.4 and 15.2 °Baume from the KRR-A-SH and KRR-D-SH blocks, respectively. However, the estimate of the average sugar content from both blocks was lower when only the bunches from the selected panels were used for the analysis.

The average sugar content of each block was substantially lower when estimated from the PS data (Table 4.4), with little difference between using the whole dataset from the six rows or just the data from the selected panels (same panels as per BVE bunch collection).

Table 4.4 - Average sugar content in blocks KRR-A-SH and KRR-D-SH measured from chemical analyses of collected grapevine bunches (a and b) and estimated from the spectral data using the CSIRO point spectrometer (c and d). Values in brackets represent the standard deviation of the means.

	a	*b*		*c*		*d*	
Block	**BVE bulk block analysis**	**BVE selected panels**	**% diff with bulk analysis**	**CSIRO 6 rows**	**% diff with bulk analysis**	**CSIRO selected panels**	**% diff with bulk analysis**
KRR-A-SH	**15.4(n.a)**	**14.5(0.8)**	-6%	**13(0.5)**	-16%	**12.8(0.3)**	-17%
KRR-D-SH	**15.2(n.a.)**	**14.7(0.6)**	-3%	**13.2(0.5)**	-13%	**13.4(0.2)**	-12%

Correlation between measured and estimated data

There was no correlation between the estimated (PS) and measured sugar contents of individually selected panels, as shown in Figure 4.23 left. The correlation was also low (Figure 4.23 right) when the panel data were averaged to provide the mean sugar contents from rows 21 and 22 (KRR-A-SH) and from rows 24 and 25 (KRR-D-SH). Typically, the result from the point spectrometer analysis was lower than the BVE sampling.

The lack of significant correlation between the PS data and the measured data from individual panels can also be seen in Figure 4.24. However, the bulk analysis by BE was from the harvested and processed fruit in the winery, with the manual sampling by BE lower than this value (Table 4.4). Nevertheless, while the PS derived results were reasonable, they were lower than the direct analysis.

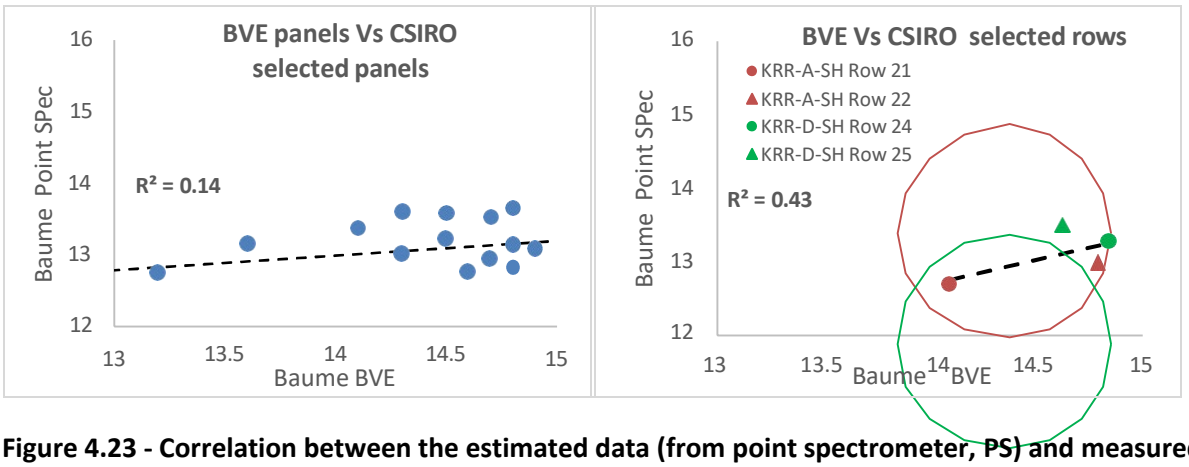


Figure 4.23 - Correlation between the estimated data (from point spectrometer, PS) and measured data (laboratory analysis by BVE) of the individual selected panels (left) and selected rows (right)

There are a number of potential explanations for the discrepancy in absolute values. Firstly, the sensor translation method may need to be improved (it was not trained on the full data set, only 2020/21 season results), and secondly, the fruit spectra classification algorithm may be excluding spectra of shrivelled berries, which are being included in the manual measurements (particularly the bulk analysis by BE) and thirdly, the lab based calibration may not include enough spectra from fruit with high (above 24 ºBrix) sugar content. All of which could be readily rectified and tested in future work.

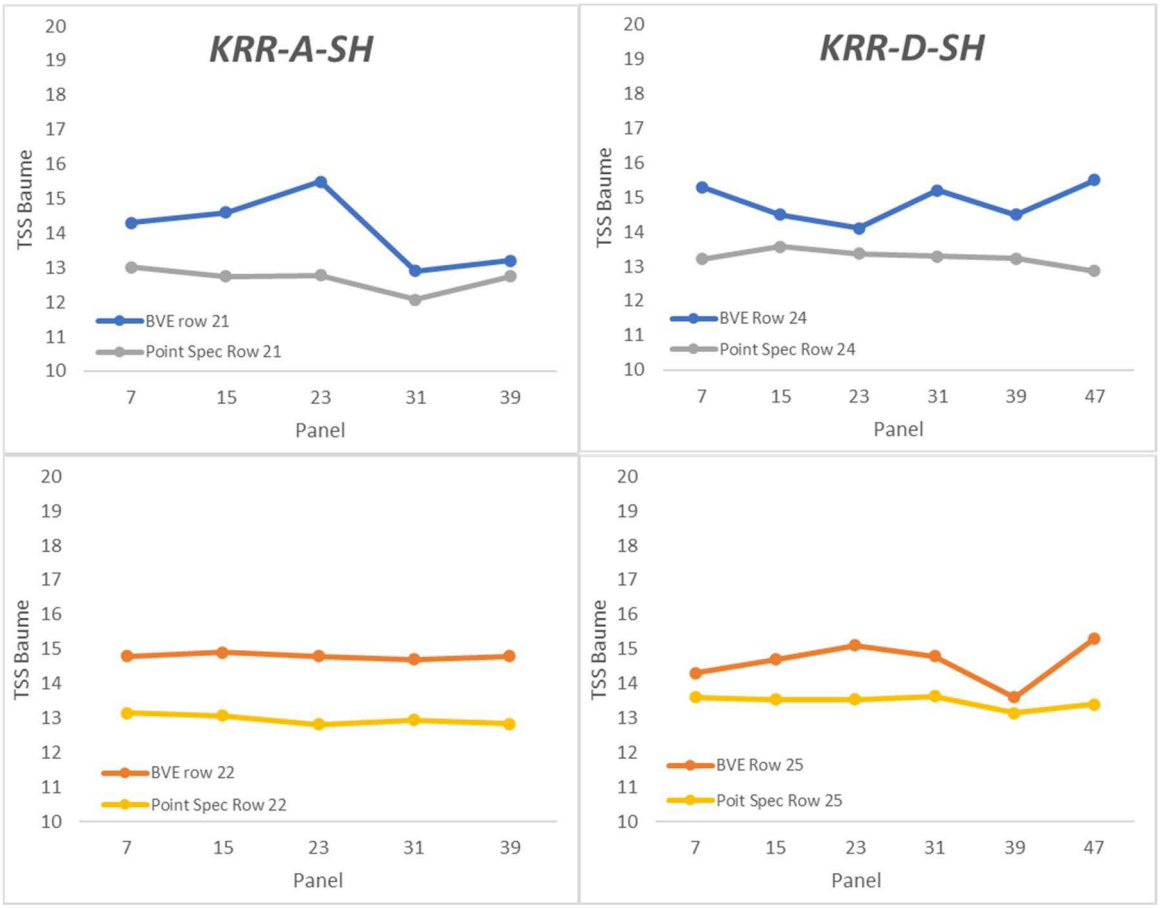


Figure 4.24 - Comparison of the mean sugar content (°Baume) from individual panels measured from harvested grapevine bunches (BVE) and estimated from point spectrometer data.

4.5.4 Spatial variability across the vineyard blocks of estimated TSS

The maps in Figure 4.25 and Figure 4.26 represent the spatial distribution of grapevine sugar content (ºBaume) within the areas surveyed using the point spectrometer sensor. Both maps were obtained by interpolating the PS data from the individual panels of the six selected rows in blocks KRR-A-SH and KRR-D-SH using the Kriging method in QGIS.

Both blocks showed a clear spatial variability of the sugar content within the areas where the data were collected. In block KRR-A-SH, an area of low mean sugar content is evident across all rows and close to the centre of the rows.

In block KRR-D-SH, the spatial variability showed a gradual change in the mean sugar content between the rows, rather than within the row, with the lowest values found on the south-facing side of the surveyed area.

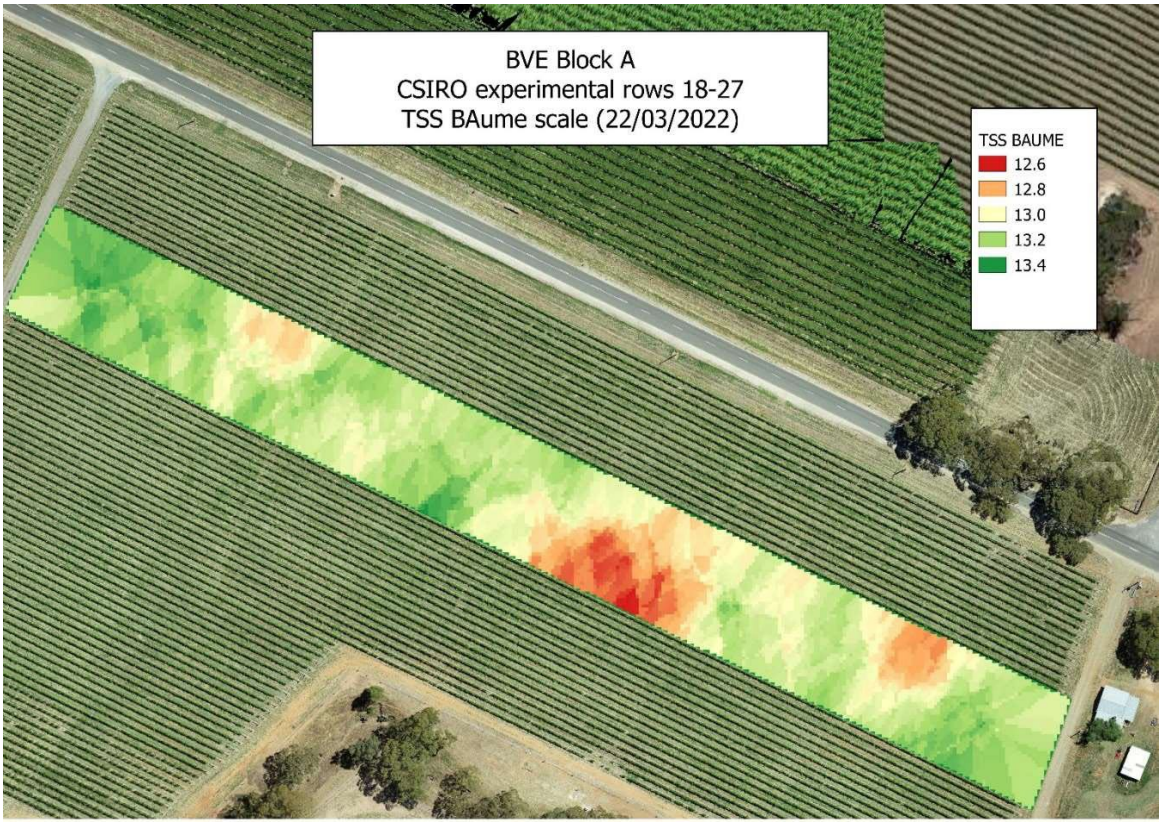


Figure 4.25 - Interpolation map representing the spatial variability of the grapevine sugar content within the area surveyed using the point spectrometer at block KRR-A-SH

4.6 Conclusions

The approach taken, whereby the point spectrometer based instrument was used to collect data at scale in the field, but matched with lab based calibrations that could be much more quickly developed, allowed for a general purpose instrument that could be used for many quality parameters. Further, although not tested in this project, the trial calibration for macro-nutrients generated in CSA 1601 suggests that it can be used for other viticultural purposes as well.

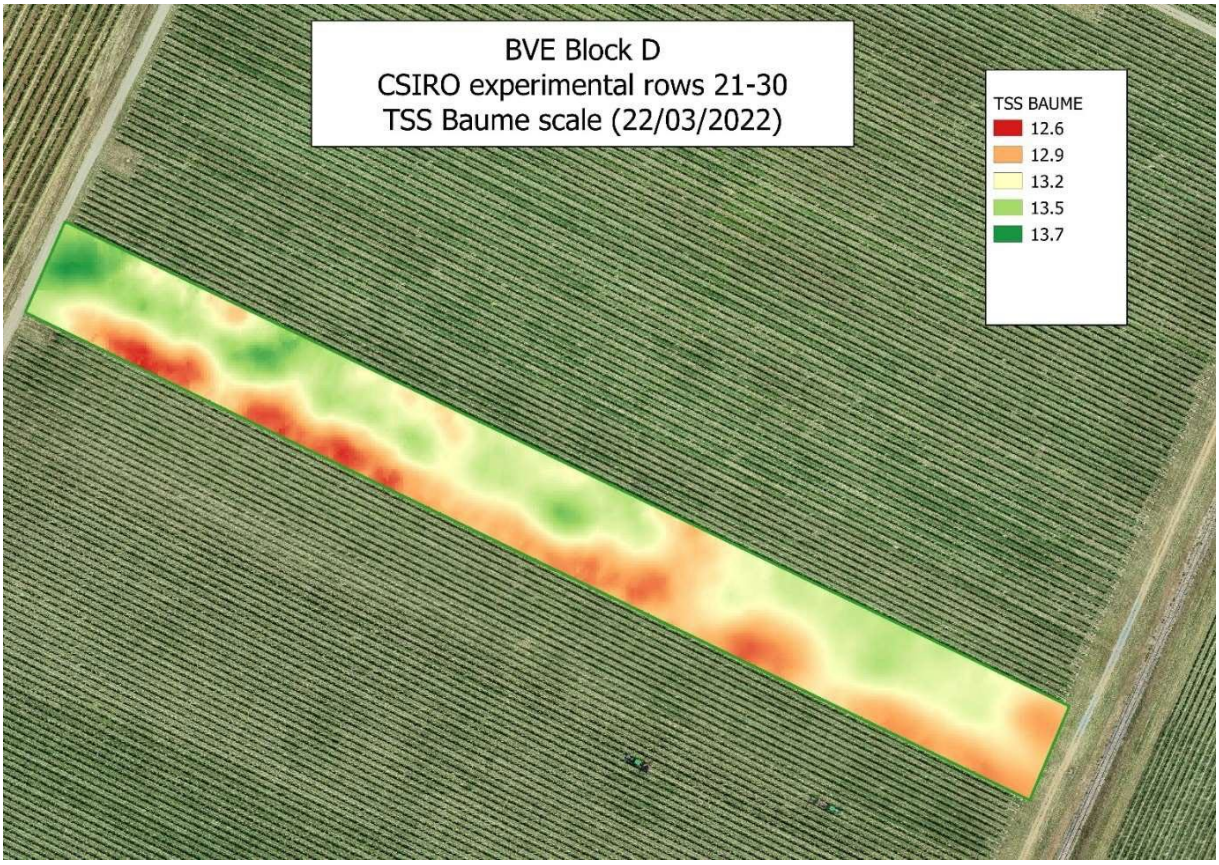


Figure 4.26 - Interpolation map representing the spatial variability of the grapevine sugar content within the area surveyed using the point spectrometer at block KRR-D-SH.

Chapter 5 – Estimation and mapping of canopy density and cover

5.1 Background

Project CSA 1601 developed a method for using a 3D spinning LiDAR SLAM (Simultaneous Localisation and Mapping) method to estimate canopy density over a vineyard block (Edwards et al. 2020a). Using a mobile spinning LiDAR, rather than a push-broom (line scan) system allows data from within the canopy to be obtained, even in vigorous plantings. The method uses SLAM to register a 'ray cloud', that is the origin, vector and contact point of each laser pulse from the LiDAR. The volume of the vineyard is then split into 3D spaces, voxels, and the ray cloud segmented, so that each ray that enters a given voxel is known. The penetration of the rays that enter a voxel are then used to estimate the single sided leaf area within that voxel, giving a canopy density. Figure 5.1 provides a diagrammatic view of this process.

In conjunction with CSIRO Data61, a ray cloud tools software library was developed and released (available at https://github.com/csiro-robotics/raycloudtools/). The method was also summarised in Lowe et al. (2021).

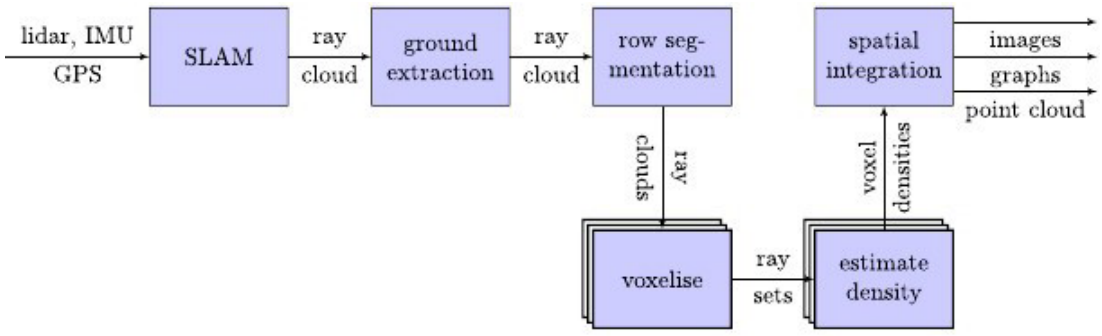


Figure 5.1 – workflow and data forms for processing LiDAR data into 3D canopy density data.

The photogrammetry based method for determining canopy cover from drone data using a digital surface map (DSM) and digital terrain map (DTM) developed through CSA 1601 was also used during the current project and the workflow further optimised.

5.2 Visualisation mapping of canopy data

Irrespective of data source (LiDAR or drone), the mapping and display was undertaken using the QGIS graphical modeller (QGIS Development Team 2022). An initial geolocated base map is made of the individual rows and vines, ideally using RTK GPS accuracy (Figure 5.2 left). This only needs to be undertaken once during the lifetime of the vineyard block. For the Penley Estate blocks used in the Coonawarra demonstration (see below), this was kindly supplied by

Hans Loder (pers comm). A QGIS Graphical Model, developed as part of this project, calculates an 'area of influence' for each vine (Figure 5.2 right), effectively mid-row to mid-row and mid-vine to mid-vine. This area can then be used for any spatial data pertaining to the vine.

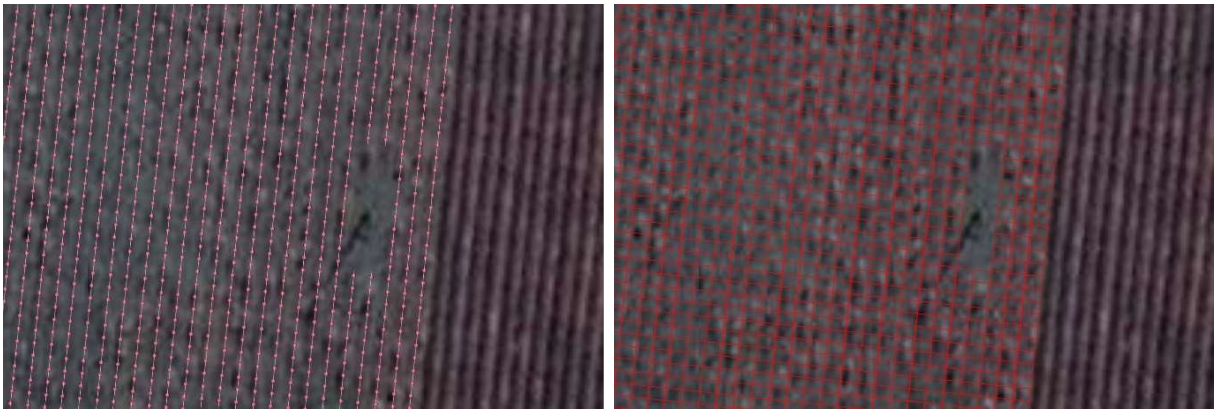


Figure 5.2 – base maps with individual rows and vines marked (left) and individual 'areas of influence' per vine (right).

Additional QGIS Graphical Models were developed to allow the LiDAR data or drone data (from either photogrammetry or pixel colour) to be displayed over the base map and for zonal statistics (e.g. canopy density or cover per vine) to be calculated and displayed.

5.3 Industry demonstration

To demonstrate the practicality and potential value of the system, a field test was arranged with growers in the Coonawarra for early December 2021. The aim was to run the system in a commercial vineyard, process the data overnight and present the results, including mapping of canopy density, to the growers the following morning. In addition, drone imagery was to be taken to utilise the improved pipeline for canopy cover.

After discussion with local industry, including Ockert Le Roux of the Coonawarra Vignerons, it was decided to run the test on two blocks at Penley Estate and run an information session open to all in the local industry, where the results would be presented, at Rymill Wines.

The LiDAR was set up on a Kubota RTV (Figure 5.3 left) and a workstation PC with the photogrammetry software set up in the site office (Figure 5.3 right). All data capture was undertaken on the afternoon of 9 December 2021. Multiple LiDAR runs were made, focussing on four rows in each block, with different configurations to also provide data for future analysis on the most efficient configuration for taking the LiDAR data. Drone flights were made of both blocks using a DJI Phantom 3 with its standard RGB camera.

The workflows for both the LiDAR and drone data were run overnight and the mapping prepared the following morning ready for the industry presentation at 10:30 a.m. Figure 5.4 provides the first visualisation of the LiDAR data indicating successful processing, with the left image visualising a point cloud from the SLAM process and the right image visualising a top

down view of the canopy density results from raycloudtools.



Figure 5.3 – set up for on site data collection and processing at Penley Estate.

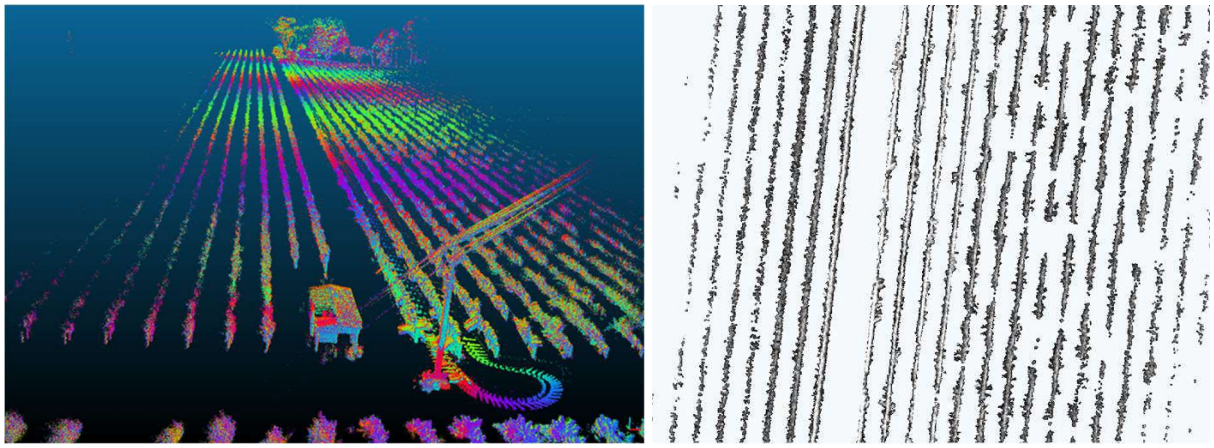


Figure 5.4 – initial imagery from LiDAR workflow indicating successful processing.

Block 6 was chosen to illustrate a LiDAR derived canopy density output (Figure 5.5 per vine, three categories) and Block 18 the drone output (per vine, proportion of cover). As noted in the report for CSA 1601, there are many ways that the canopy density can potentially be presented and used, but it was decided that a per vine map was the simplest. Feedback from the growers present was very positive with interest in when this might be available and what the cost point might be.



Figure 5.5 -visualisation of canopy density results from LiDAR scan processing. Five complete rows (top) and a small area of vineyard with and without per vine data classification overlay (bottom).



Figure 5.6 – visualisation of canopy cover results from drone imagery processing. Vineyard block, with (right) and without (left) per vine data overlay.

This page intentionally left blank.

Chapter 6 – Potential for commercialisation

This project developed five potentially technologies that had the potential to be commercialised into services for use in viticulture in Australia.

1. Machine vision for yield estimation
2. Radar for yield estimation
3. Canopy temperature for vine water status monitoring
4. NIRS for composition
5. 3D spinning LiDAR SLAM for canopy structure

All but canopy temperature monitoring are continuous measurements taken from a moving vineyard vehicle, providing data at scale over the whole vineyard. Of those four, all but the LiDAR system have the potential to be truly on-the-go, providing results during the data collection runs themselves.

6.1 – Machine vision for yield estimation

Since the start of the CSIRO/Wine Australia work on the use of machine vision with machine learning for estimation of winegrape fruit yield through CSA 1602 in late 2016, there has been significant effort to develop such systems in a number of winegrape growing countries around the world, including Australia. Typically these rely on imaging at night, use in highly managed vines where the majority of fruit are fully exposed, additional hand collected data in excess of that used for traditional yield estimation or a combination of these. To the best of our knowledge, no other group has demonstrated a completely hands off workflow for estimation of actual yield in 'sprawl' vines from continuous video alone.

Consequently, the prototype pre-harvest yield estimation tool referred to here offers a unique opportunity for commercial development of a system applicable across the majority of the Australian viticultural sector and beyond. However, significant effort is required to move the prototype workflow to a reliable software tool ready for broad-scale testing, let alone a commercially ready tool. Furthermore, the current workflow relies on access to significant high-performance computing capacity that would not be economically viable if based on a commercially available system. As a result, a concerted effort to optimise the computing requirements would also be required to bring a tool based on the work here to market.

It is likely that improving the ML components and the workflow would take 1-2 FTE data scientists 1- 2 years and that the optimisation of the computing requirements would require several suitable software engineers for several months. Both of these areas of work would require ongoing support and development once a minimum viable product was prepared. These requirements for further development mean the investment required is likely to be out of reach of a start-up and only be viable for an established commercial entity willing to

make the investment.

6.2 - Radar for yield estimation

The detailed analyses undertaken as part of this project indicate that, while the concept of using radar for yield estimation is an attractive one, the physics involved and the complexity of the vineyard environment mean that it is an intractable problem with current technology and understanding. As a result, commercial development is not possible in the foreseeable future.

6.3 - Canopy temperature for vine water status monitoring

The working ProxiCrop prototype device offers a system for collecting canopy temperature data at a realistic scale (multiple vines vs. one to a few leaves for existing start-up products) that has been successfully demonstrated over a full growing season with 10+ deployed in multiple commercial vineyards. The ability for the software backend to adjust to changes in canopy size during the season also provides additional flexibility in its use. The next step for commercialisation of the hardware prototype would be industrial design and manufacturing, which would allow it to be used for any irrigation management that uses canopy temperature as an input.

Initial conversations with providers and industry, together with an informal examination of the market, suggest that a more likely avenue would be to market ProxiCrop as a standalone system that provides an index or other quantifiable measure of vine water status. While the work presented in this report demonstrates that ProxiCrop and the data it provides matches vine water status, it has not reached the stage of a validated water status index. It is likely that a further R&D project to build and validate such an index is required prior to full commercialisation. Such a project would likely need 2 FTE (a combination of field technicians and both plant physiology and data scientists) for two growing seasons.

6.4 - NIRS for fruit and vine tissue composition

Across this project, CSA 1601 and CSA 1602, a significant amount of work has been undertaken to collect NIR spectra between 550 and 1650 nm matched to berry composition parameters, with over 7,000 individual samples collected. Both the data set and the calibration could potentially be licensed to a start-up utilising a hyperspectral device within this range. However, as spectra from different sensors rarely give exactly the same spectral shape even after radiance correction, techniques such as the one described in this report to match spectra between devices and allow calibrations to be portable would likely be required to make this realistic.

The project has also demonstrated a field device mountable to vineyard vehicles that can be used with such a sensor translation technique and our lab based calibration to provide timely

on-the-go collected data on fruit composition at scale to harvest decision making. There is clearly the potential to develop such a system and the back-end calibrations into a commercial offering. This could be done with start-up investment or in conjunction with an existing entity that operates with full spectral range sensors, but would need funding for further development and validation. The requirements for this would be similar to those described for the ProxiCrop. It should be noted that CSA 1601 developed a proof-of-concept calibration for macro-nutrients in leaves, so it is very likely that the same field instrumentation could be used for assessing vine nutrition early in the season.

6.5 - 3D spinning LiDAR SLAM for canopy structure

The 3D canopy structure tool currently relies on a CSIRO SLAM implementation (protected IP) and implementing the CSIRO developed raycloudtools library (open source) as well as CSIRO developed dual-axis LiDAR with IMU (not protected in concept, but not commercially available). In combination with the computing requirements of these systems, commercialisation of the tool is probably only realistic as a service (rather than commercially available units with all required software). Without additional market research it is not clear whether a large enough fraction of the industry desires canopy knowledge at this level of detail and whether they would be willing to pay enough for the service to make it commercially viable and this would be recommended prior to further development of this tool.

Note that CSIRO did trial a dual-axis LiDAR with a drone and the above software tools to generate a similar map of canopy density, but the detail and accuracy of this has not been compared to the ground based data collection as yet. The drone-based LiDAR offers fewer laser pulses per volume of canopy and a smaller range of angles of penetration, both of which will affect the output. However, a drone-based system may be more economically viable due to the ability to more quickly assess larger areas of vineyard.

This page intentionally left blank.

Appendix 1: Communication

Refereed papers

Lowe, T., Moghadam, P., Edwards, E., & Williams, J. (2021). Canopy density estimation in perennial horticulture crops using 3D spinning lidar SLAM. *Journal of Field Robotics*, **38**(4), 598–618. https://doi.org/10.1002/rob.22006

Papers under review

Khokher, M.R., Liao, Q., Smith, A., Sun, C., Mackenzie, D., Thomas2, M.R., Wang, D., Edwards, E.J. Automatic Early Yield Estimation in Viticulture based on Grapevine Inflorescence Detection and Counting in Videos

Conference presentations

Edwards, E.J. (2020). 'On the go' sensors to monitor fruit and canopy development. *AWITC Workshop*, June 2020.

Datasets

Lowe, T., Moghadam, P., Edwards, E., Williams, J., Brosnan, S., Haddon, D., Dungavell, R. (2021). AgScan3D Viticulture Datasets. CSIRO Data Collection. https://doi.org/10.25919/gbvy-yz23

This page intentionally left blank.

Appendix 2: Intellectual property

No patents were taken out during the term of this project.

Background IP and development of background IP remain the property of CSIRO.

Project IP is jointly owned by CSIRO and Wine Australia.

Project data are jointly owned by CSIRO and Wine Australia.

This page intentionally left blank.

Appendix 3: References

Clingeleffer, P.R.; Sommer, K.J.; Krstic, M.P.; Small, G.; Welsh, M. (1997) Winegrape crop prediction and management. *Australian and New Zealand Wine Industry Journal* **12**: 354–359.

Clingeleffer, P.; Dunn, G.; Krstic, M.; Martin, S. (2001) Crop development, crop estimation and crop control to secure quality and production of major wine grape varieties: a national approach. Final report to Wine Australia.

Edwards, E.J., Thomas, M.R., Gensemer, S., Smith, A., Lowe, T., Moghadam, P., Lagerstrom, R., Knauer, U., Seiffert, U. (2020) New technologies for dynamic canopy and disease management. Final report to Wine Australia.

Edwards, E.J., Thomas, M.R., Gensemer, S., Lagerstrom, R., Khokher, R., Liao, Q., Sun, C., Wang, D., Hargrave, C. Ralston, J. (2020) New non-destructive technologies for simultaneous yield, crop condition and quality estimation. Final report to Wine Australia.

QGIS Development Team (2022). QGIS Geographic Information System. Open Source Geospatial Foundation Project. http://qgis.osgeo.org

This page intentionally left blank.

Appendix 4

Project staff

Chief Investigator

Everard Edwards

Principal Investigator

Giacomo Betti

Administration

Shane Wynn

CSIRO A&F

Adam Smith

Najid Pereira-Ishak

Jack Conroy

Don Mackenzie

Chaoyong Huang

Alex Lawler

Adelle Semmler

CSIRO Manufacturing

Stephen Gensemer

CSIRO Data61

Peyman Moghadam

Tom Lowe

Dadong Wang

Ryan Lagerstrom

Rizwan Khokker

Qiyu Liao

Xun Li

CSIRO Mineral Resources

Chad Hargrave

Jonathon Ralston

Paul McPhee

Luke Powelll

MLAI FSP

Daniel Smith

David Ahmedt Aristizabal

Acknowledgements

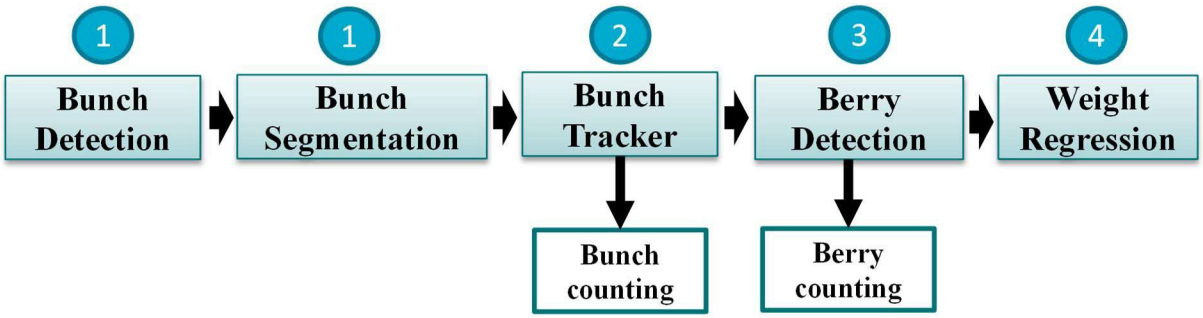
We gratefully thank the wine companies and their staff that gave us access to their vineyards and supported our staff in undertaking their work, including Accolade Wines (particularly Warren Birchmore & Glen Cheel), Adam Loveys & James Thorpe, Barossa Valley Estate (particularly Alan Hoey), Penley Estate (particularly Hans Loder), Pernot Ricard (particularly Tim McCarthy), Mark Hall at Jess Farm, and Rymill Wines. We also thank the other staff at CSIRO and these wine companies that assisted us at various times.

This research was supported by funding from Wine Australia. Wine Australia invests in and manages research, development and extension on behalf of Australia's grape growers and winemakers and the Australian Government.

Appendix 5: Other material

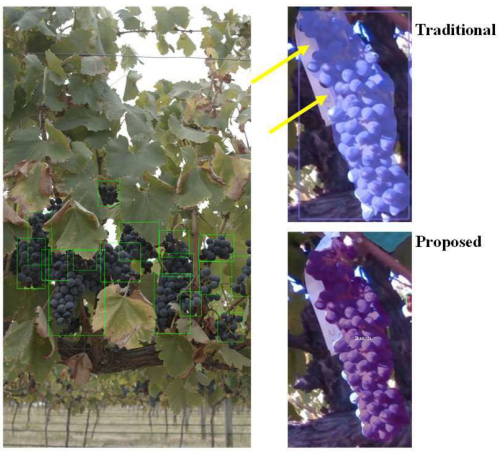
Appendix 5.1: schematics of complete ML processing pipeline for yield estimation

Overview:



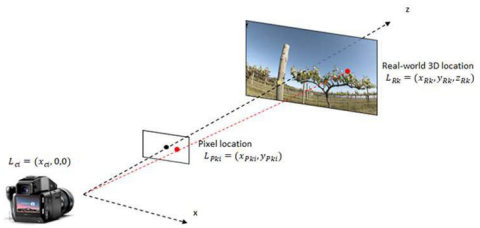
Bunch detection and segmentation



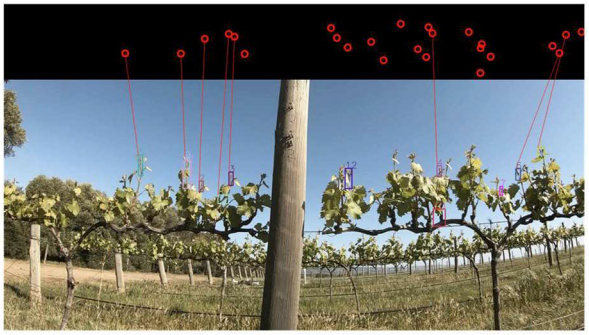


Red bounding-boxes and polygons represent ground-truth annotations and blue bounding-boxes and filled regions represent the detected and segmentation

Bunch tracking



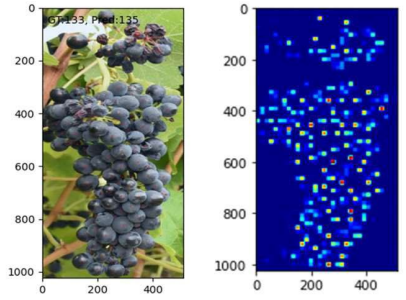
Count Inflorescences/Bunches

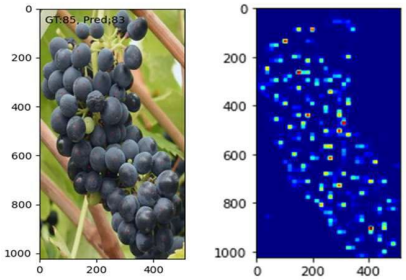


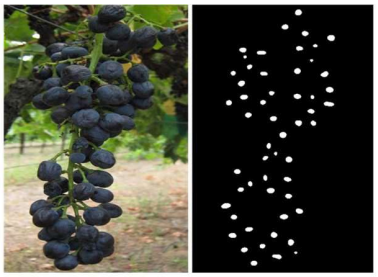
Black Magic sequence



Berry detection



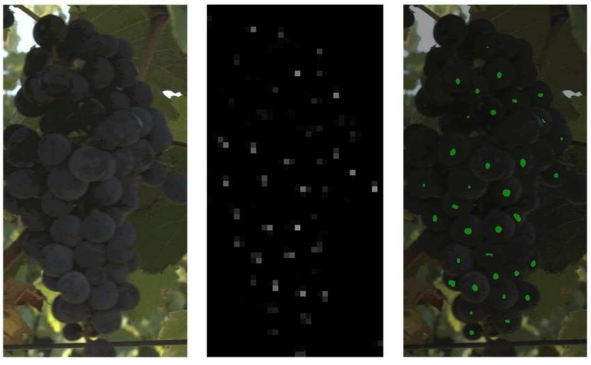




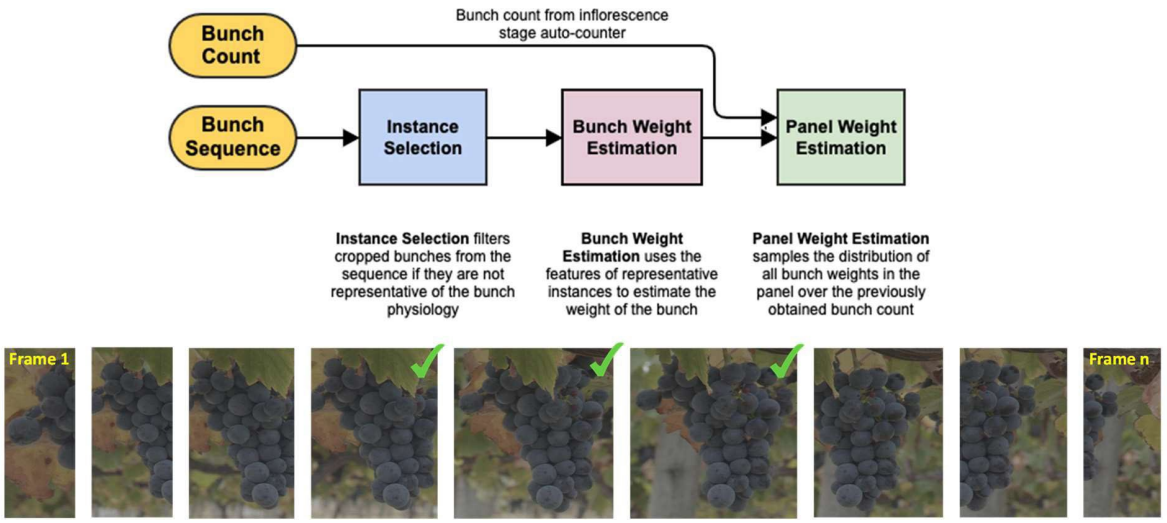


Density map

Segmentation



Weight regression



Appendix 5.2: Planned Project Outputs

Outputs and activities from the annual operating plans.

Outputs and Activities 2019–20

Year 1 ran from the end of projects CSA 1601 and CSA 1602, encompassing only three months and did not have any formal outputs beyond those reported in CSA 1601 and CSA 1602.

Outputs and Activities 2020-21

Year 2	Output	Target Date	Activities
	Future technology directions of the project determined.	30/11/2020	One-page technology assessment made for each of the technologies developed in CSA 1601 and CSA 1602 and provided to Wine Australia and industry reference group. Meeting with Wine Australia and reference group to discuss technologies and determine directions towards commercialisation, further development or to put on hold for each technology.
	Machine learning algorithms for analysis of RGB video imagery further developed, improved and validated.	30/6/2021	Image labelling requirements for machine learning tools determined and undertaken. Machine learning tools compared ground truth and previous iterations.
	Yield estimation/prediction technology further developed.	30/6/2021	Field collection of data (imagery, other sensors as appropriate and manual assessment) for inflorescence counting, bunch counting and bunch volume assessment undertaken on at least three occasions at multiple commercial vineyards. At least one technology for yield prediction/estimation further developed using field data. Steps towards commercialisation undertaken as determined by activity a.

	Canopy management decision support technology further developed.	30/6/2021	Field collection of data (imagery, other sensors as appropriate and manual assessment) for canopy size and assessment undertaken on at least three occasions at multiple commercial vineyards. At least one technology for yield prediction/estimation further developed using field data. Steps towards commercialisation undertaken as determined by activity a.
	ProxiCrop thermal sensors and analytics further developed.	30/6/2021	Sensor data analysis improvements based on improvements in colour segmentation algorithms. Processed thermal image data compared with vine and climate data and assessed for development of novel or existing vine stress index to assist irrigation management. Continued software (and hardware if suitable) improvements to improve collection and processing of data. Further assessment of commercialisation potential and pathway.

Outputs and Activities 2021–22

Year 3	Output	Target Date	Activities
	Commercialisation plan for each technology developed by the project.	30/6/2022	Written assessment of commercialisation potential for each technology. Commercialisation plan available for each technology assessed as suitable for further commercialisation. EOIs for interaction with CSIRO on video+ML yield estimation and ProxiCrop published prior to beginning

			of 21/22 growing season. Responses followed up in a timely manner. Demonstration of LiDAR based canopy management decision support technology organised and provided to industry.
	Yield estimation/prediction technology further developed.	30/6/2022	Field collection of data (RGB imagery and radar sensors as appropriate and manual assessment) for inflorescence counting, bunch counting, berry counting and bunch volume assessment undertaken at multiple commercial vineyard blocks. Machine learning algorithms for processing of previously collected UWB radar trialled. Transmission radar system tested in lab and prototype deployed for field testing. Machine learning algorithms for analysis of RGB video imagery further developed, improved and validated. Machine learning tool predictions compared with ground truth and previous iterations, including bunch mass. Interaction with CSIRO MLAI Future Science Platform to provide additional ML technology testing where practical. Interaction with commercial entities in response to RGB video+ML EOI if suitable responses.
	Fruit composition technology further developed.	30/6/2022	All spectral and lab-based ground truth data collected during 20/21 FY incorporated into new/updated calibration models and assessed. Point-spec based calibration developed using existing lab-based calibration data and accuracy demonstrated. Field collection of point-spec instrument expanded, at least three commercial

			vineyards scanned on-the-go together with hand-scanning of bunches and collection of samples for lab analysis on at least three time points. Second point-spec instrument built to expand data collection if required. Prototype data analysis pipeline developed and deployed on HPC system.
	Canopy management decision support technology further developed.	30/6/2022	Field collection of drone imagery, together with manual assessment (LAI) undertaken at multiple commercial vineyards. Canopy cover data analysis pipeline further developed and compared with ground truth and, if practical, with LiDAR assessments. Processing pipeline to convert drone data into maps compatible with a commercially available application for growers, e.g. PlatFarm. Developed and demonstrated.

Appendix 6: Budget reconciliation

A budget reconciliation will be provided separately.