Australian Government

Wine Australia

THE UNIVERSITY of ADELAIDE

VitiVisor: An Information, Prediction and Advisory Platform for Australian Viticulture



FINAL REPORT TO WINE AUSTRALIA

Project Number: **UA 1803 Project 1.5**

Principal Investigator: **PROFESSOR SETH WESTRA**

Authors: **Westra, S., Nott, B., Wheeler, S., Knowling, M., Zhang, H., Seidl, C., Tingey-Holyoak, J., Lu, T-F., Zhang, Z., Zhao, S., Collins, C., Grigg, D., Kiely, W., Bennett, B., Ostendorf, B., Shi, J., Tan, R., Lowe, A.**

Research Organisation: **UNIVERSITY OF ADELAIDE**

Date: **28th June 2022**

DOCUMENT STATUS RECORD

Project Title: VitiVisor: An Information, Prediction and Advisory Platform for Australian Viticulture
Client: Riverland Wine and Wine Australia
Document Number: 2022/15
File Name: VitiVisor_finalReport_DRAFT.docx

Issue No.	Date of Issue	Description	Signatures	
			Authors	Approved
1	28 June 2022	Draft report	SW, BN, SaW, MK, CS, JT-H, HZ, T-FL, ZZ, SZ, CC, DG, WK, BB, BO, JS, RT, AL	SW

Disclaimer:

1. The University of Adelaide has taken all reasonable steps to ensure that the information contained within this report is accurate at the time of production. In some cases, we may have relied upon the information supplied by the client.
2. This report has been prepared in accordance with good professional practice. No other warranty expressed or implied is made as to the professional advice given in this report.
3. The University of Adelaide maintains no responsibility for the misrepresentation of results due to incorrect usage of the information contained within this report.
4. This report should remain together and be read as a whole.
5. This report has been prepared solely for the benefit of the client listed above. No liability is accepted by the University of Adelaide with respect to the use of this report by third parties without prior written approval.

Contact Details:

Physical:
Engineering North Building
North Terrace Campus
University of Adelaide, Adelaide.

Postal:
University of Adelaide, Adelaide, 5005

Telephone: (08) 8313 1538

DRAFT

TABLE OF CONTENTS

Abstract 4

Executive Summary 5

1 Background and Project Context 7

1.1 What is VitiVisor? 7

1.2 Viticulture in an Era of Digital Transformation 7

1.3 Securing Value… For Whom? 9

1.4 'Control your Data, Control your Destiny' 10

2 Project Aims 11

2.1 Antecedents: Co-Design Prior to Project Inception 11

2.2 Guiding Design Principles 11

2.3 Project Aims and Alignment with Grower Priorities 12

3 Project Governance and Grower Engagement 14

3.1 Management Committee 15

3.2 Adoption and Commercialisation 15

3.3 Development Team 15

3.4 Grower Advisory Group 16

3.5 Support Sites 16

4 Method 17

4.1 Data Acquisition 17

4.1.1 Ground-based vineyard measures (ground truth and image acquisition) 17

4.1.2 'Internet of Things' sensor data 21

4.1.3 Third-party data 25

4.2 Ground-Based Vision System Development 27

4.2.1 System overview 27

4.2.2 VitiBox: A flexible image capture system 29

4.2.3 Algorithm development and post-processing subsystem 30

4.3 Viticulture modelling 34

4.3.1 Integration of established vine growth model 34

4.3.2 Building a vineyard 'digital twin' 36

4.3.3 Supporting grower decisions 37

4.4 Economic Assessment 39

4.4.1 Case studies and benchmarking 39

4.4.2 Net economic benefit assessment 41

4.5 Software Development 43
4.5.1 Software architecture overview 43
4.5.2 Business logic layer 44
4.5.3 Connection with established data providers 45
4.5.4 Integration of ground-based vision data 47
4.5.5 Ingestion of grower information 48
4.5.6 VineLOGIC integration 50
4.6 Adoption and Commercialisation 50
5 Results and Discussion 53
5.1 Ground-Based Vision System 53
5.1.1 Leaf Area Index Prediction 53
5.1.2 Inflorescence, Shoot Tips and Bud Counts 54
5.2 Viticulture Modelling 56
5.3 Economic Assessment 59
5.3.1 High-level assessment of Riverland grape growing farm financial characteristics 59
5.3.2 Identifying key factors associated with farm net return from an irrigator survey analysis – existing data 60
5.3.3 Identifying detailed variable and fixed cost and revenues for grape growing in the Riverland -Farm financial case studies – primary data collection 60
5.3.4 Understanding the net benefits of VitiVisor adoption for Growers - Financial BCA analysis 63
5.4 VitiVisor Dashboard 63
5.4.1 Landing page, property management and notification system 64
5.4.2 Sensors and irrigation system data 65
5.4.3 Spatial imagery 67
5.4.4 Finance and benchmarking 71
5.4.5 Scenario and prediction tool 72
5.5 Adoption and Commercialisation 73
5.5.1 Commercialisation through IP transfer 74
5.5.2 Commercialisation through establishment of a Social Enterprise 75
5.5.3 Industry-owned model 77
6 Outcome/Conclusion 79
7 Recommendations 81
8 References 85
Appendix 1: Technology Readiness Levels 87
Appendix 2: Communication 89

Communication with the project's Primary Stakeholders 89
Communication with Industry 90
Communication with the Scientific Community 91
Communication with the General Public 92
Appendix 3: Intellectual Property 93
Appendix 4: Project Team 95
Appendix 5: Acknowledgements 98

ABSTRACT

Initiated by Riverland Wine and co-developed with Riverland growers, VitiVisor is a digital platform that supports growers to manage production costs and enhance vineyard profitability. Key elements of the platform include integration of diverse data streams covering biophysical and economic information into a single web portal, the collection of new high-resolution data from ground-based camera systems, and state-of-the-art algorithms for predicting yield and providing guidance on management decisions. The tool has been developed to Technology Readiness Level 6 (pre-commercial demonstration), has been demonstrated on the Loxton Research Centre vineyard, and is available to the community through open source licensing.

EXECUTIVE SUMMARY

VitiVisor exists within a context of rapid change and digital disruption that has the potential to transform the Australian wine sector, presenting both challenges and opportunities to those within it. Recognising the need to be active participants in this transformation, Riverland Wine engaged the University of Adelaide in 2018 as part of a process of co-creating digital technologies that are specifically targeted to the unique contexts of Australia's inland winegrape growing regions.

This report describes the outcome of an important step in the journey of creating VitiVisor as a commercially viable digital platform that can help growers maximise the benefits of their data, with an initial focus on viticulture production costs and profitability. Specifically, this project focused on developing VitiVisor to a 'pre-commercial' technology readiness level, involving (1) the development of a prototype tested on real-world vineyards, (2) the development of a clearly identified value proposition, and (3) the identification and initial engagement of potential commercial partners.

This project was operational over the period from January 2020 to June 2022, and led to the creation of an open-source digital information, prediction and advisory platform that is capable of supporting grower decision making throughout a growing season. The platform has progressed through multiple design cycles, and includes the following functionality:

- A ground-based vision system that cost-effectively captures geo-located canopy imagery, converts the imagery into key vine health and development indices (e.g. plant area index, pruning weight, inflorescence number, shoot number, bud count and plant leaf temperature), and automatically uploads this information onto a digital platform for visualisation purposes and integration with other parts of the system. This system comprises a multi-camera box with flexible mount (e.g. to quad bike, tractor, ATV), and has been tested under a range of lighting conditions (i.e. cloud cover and direct sunlight).
- A data acquisition and storage system that collects real-time data from diverse sources including various 'Internet of Things' sensors on the vineyard (soil moisture, temperature, electrical conductivity, weather data, dendrometer readings), irrigation system data such as flow and pressure readings, airborne imagery such as NDVI and water stress, data from the ground-based vision system, water market information and financial accounting data.
- A library of labelled vine images comprising 170,867 separate images (geo-located, with PAI and canopy porosity) from 14 sites and across the five main varieties in the Riverland. Of these, 500 images have been annotated and ground-truthed using separate measures deemed to be decision-relevant to growers, and both the annotated and raw imagery have been deployed to train the ground-based vision system to Riverland conditions.
- A vine growth and yield prediction tool that builds on the VineLOGIC modelling platform, and that has been tailored to enable updating predictions (1) in real-time based on weather, soil moisture and canopy growth, and (2) for alternative management actions, particularly as they relate to watering strategy.
- Financial data collection and benchmarking capability that enables users to upload and track their farm costs and revenues, simplifying and optimising decision making. This tool includes integration with Xero and Know Your Numbers, as well as the outcomes of a financial benchmarking dataset developed based on a survey with 25 growers.

Each of the components have been integrated into a single operating environment, which has been refined through a series of feedback cycles with growers in relation both to core functionality and

user experience. The VitiVisor platform has been demonstrated on the Loxton Research Centre, with a lean version also demonstrated on three grower sites located throughout the Riverland.

A benefit-cost analysis of the adoption of VitiVisor was undertaken, assuming a range of fixed and variable costs over a five-year horizon, and benefits over the same horizon including a 10% reduction in electricity, fuel and labour cost and an estimated 3t/ha increase in yield. The expected net present value of adoption was estimated as $1,219 per hectare, or $69,749 per average farm enterprise. This led to a benefit cost ratio of 1.5 (indicating an expected net return on investment of $1.50 for every dollar of investment), with a sensitivity analysis suggesting the ratio could vary between 0.5 and 2.9 depending on a range of assumptions. The outcome of the analysis suggested that VitiVisor would yield positive net returns provided the overall cost of purchasing and using the tool (including accounting for grower's time) over a five-year period was no greater than $67,198.

All project objectives have been achieved; however, the broader direct and indirect sectoral benefits of the VitiVisor platform have yet to be realised. Indeed, the next phase is both the most important and arguably amongst the most difficult: taking the technology from pre-commercial to commercial, achieving impact at scale, and doing so in a manner that benefits the inland regions both through direct benefits in terms of reduced operating costs and increased profitability, and indirect benefits that might flow from greater end-user engagement, digital literacy and control of its digital future.

Although outside of the formal project scope, significant efforts were also made by the partnership to identify commercial models that aligned with Riverland Wine's requirements of open-source intellectual property and member involvement. This includes extensive consultation with 16 technology providers, as well as the development of a possible alternative 'social enterprise' model that would enable growers to have direct control over the future development of VitiVisor. At the time of writing this report, significant progress has been made in identifying a potential commercial model, although no firm decision has been made and several options continue to be investigated.

Learnings from commercialisation efforts suggest that there remains a significant gap in the transition of VitiVisor from Technology Readiness Level 6 ('pre-commercial') to TRL 9 ('commercial deployment'), as required to enable the benefits to be realised through widespread deployment across Australia's vineyards. All the recommendations from this report have therefore been designed with a single objective in mind: to describe potential next steps that help ensure VitiVisor or its component technologies are deployed at scale and create value for the sector. This will require further investment, as well as efforts to ensure strategic alignment between Wine Australia, representative bodies from the inland winegrape growing regions, and one or several third-party technology companies that might ultimately become custodians of the technology.

This work would not have been possible without the continual support of Wine Australia, Riverland Wine and their grower advisory group, Dr Paul Dalby for bringing researchers and growers together during the project formulation, Mr Hao Zhang and Mr Peter Zhou from AMI Fusion Technologies in providing software engineering support, A/Prof. Joanne Tingey-Holyoak of University of South Australia in providing accounting and financial advice, Ms Moira Were AM from Ethical Fields and Dr Ariella Helfgott from Collaborative Futures in providing commercialisation advice, staff at CSIRO for making VineLOGIC openly available and providing subsequent expert input, various data providers (Ceres, H_2OX, Central Irrigation Trust, Katalyst and Know Your Numbers, MEA and Green Brain, Opensensing), Oli Madgett and the Vineyard Open Standards project, as well as PIRSA in enabling the use of the Loxton Research Centre. VitiVisor only exists because of a co-design process initiated by Riverland Wine in 2018, and the continued passion, energy and commitment of its growers. It is only through that passion, energy and commitment that extraordinary ideas get turned into reality.

1 BACKGROUND AND PROJECT CONTEXT

1.1 WHAT IS VITIVISOR?

At its core, VitiVisor is a digital platform that integrates and builds on diverse data streams to generate value to support growers in addressing costs of production and enhancing profitability. Since its genesis in 2018 (Box 1), project priorities have been driven by Riverland growers, in partnership with Wine Australia and university researchers. Our mission: to support winegrape growers to maximise opportunities from the digital transformation that is now taking place.

This report describes the outcome of an important step in the journey to creating this platform. Using the technology maturity assessment framework developed by NASA, VitiVisor has been developed to 'Technology Readiness Level 6' out of a nine point scale. This level of development is described as 'pre-commercial technology demonstration'[1], and embeds the following attributes:

- Prototype capable of performing all the functions that will be required of it in the operational system;
- Prototype tested in an end-user site environment;
- A positive value proposition that has been refined; and
- Potential commercial partners identified and engaged.

"VitiVisor's mission is to support growers maximise opportunities from the digital transformation."

Working at this TRL level has represented a substantial shift from traditional University research, which commonly work at much lower TRLs (see Appendix 1 that describes the TRL scale). Indeed, within the 2.5 years of the VitiVisor project, we have been successful in taking a significant body of viticultural research—much dating back over many decades—from theory (or from within scientific papers) to a point where concepts can be demonstrated on real-world vineyards, and their value assessed through economic evaluation.

Most of the following pages showcase the technologies and associated concepts that have been the primary focus of the VitiVisor project. Yet VitiVisor is about more than just technology development; indeed, it asks the question—how does the Australian viticulture industry wish to be placed within the broader context of digital transformation? We touch on this broader question in the following sections, and return to it in Section 7 where we discuss possible future directions for VitiVisor.

1.2 VITICULTURE IN AN ERA OF DIGITAL TRANSFORMATION

VitiVisor exists within a context of rapid change and digital disruption that has the potential to reshape the wine industry. At the time of writing, the field of digital viticulture—like digital agriculture more broadly—can best be described as fragmented, unregulated and rapidly growing. Value propositions are still being identified and refined; multiple technologies and business models are vying for dominance; some businesses are merging or acquiring other businesses, while new enterprises are continually springing up; data standards are evolving; and issues associated with data ownership, privacy, trust and security are being actively debated. The end state of this transformative process remains unclear, but with exponential increases in capital flowing into digital agriculture (AgFunder, 2021), it is certain that the next decade will look very different to the last.

[1] The specific levels used in this report are based on the TRL frameworks of the US Department of Energy (2011) and US Department of Agriculture (2016). The TRL framework is summarised in Appendix 1.

Box 1: Responding to the South Australian Riverland challenge

VitiVisor was born in the Riverland.



At the invitation of Riverland Wine, researchers from the University of Adelaide visited a range of Riverland vineyards in mid-2018 and discussed with growers key pain points and future ambitions as they relate to digital viticulture. From these early meetings, a subsequent pilot project in early 2019 and continual engagement with growers both directly and through Riverland Wine as their representative body, a number of consistent messages emerged—messages that would form the basis of what has now become 'VitiVisor'.

First and foremost, growers highlighted concerns about lack of involvement in research and technology development funded through their levies. We responded by directly embedding growers and Riverland Wine into all key elements of project governance and decision making.

In terms of pain points, growers highlighted cost of production as their primary concern. We responded by developing software that would enable benchmarking and provision of advisory services, considering diverse resource categories but with a primary focus on water.

Growers also repeatedly commented on the inaccessibility of advanced viticulture research to support practical viticulture outcomes. We responded by viewing the digital platform as a vehicle to short-circuit the traditional research extension model, embedding research insights and reducing friction by placing that research both within context and at point-of-decision.

Growers identified digital fragmentation as a major friction point, with numerous stories of growers maintaining over a dozen separate 'apps' each with a narrowly defined capability. We responded by developing the capacity for seamless data integration into a single digital platform, and exploring new insights that arise through this convergence of data.

Lastly, growers highlighted the importance of an open source IP model, and the development of a business model that would retain value in the regions. We responded by making all project IP freely available to the world and exploring novel business models including the potential establishment of a cooperative structure run by Riverland growers.

Whilst not every aspiration expressed by growers has been met, much of it has been achieved, and the VitiVisor platform serves as a foundation for future development. Yet the next phase is both the most important and arguably amongst the most difficult: taking the technology from pre-commercial to commercial, achieving impact at scale, and doing so in a manner that directly benefits end users both immediately and over the longer term.

The foundation has been laid, but the real work has just begun.

There are several unique elements to digital disruption in agriculture that differentiate it from digital disruption in other sectors. Connectivity has come only recently to many rural communities, and many still suffer inadequate internet coverage. The workforce is aging, digital literacy is comparatively low, and many farms have only recently commenced their digital journey (McKinsey & Company, 2017). Farm budgets are often heavily squeezed, and available capital to invest in innovation at the farm level is limited (Wheeler et al, 2022). These factors mean that digitisation is still in its infancy, with low penetration of digital technologies compared to most other parts of the Australian economy (McKinsey & Company, 2017).

> "It is precisely at times of greatest uncertainty when our capacity to shape the future is at its greatest"

Each of these factors suggests that we are still in the early stage of the 'digital agriculture' business cycle, with this stage characterised both by high levels of uncertainty and significant potential for future growth[2]. At this stage, developing predictions of the future state of the industry is at its most uncertain, yet this is also the precise time when opportunity to positively shape that future is at its maximum. It is at this point that attention should be focused not only on seeking to develop scenarios of what digital viticulture *could* look like, but also provoking discussions of what it *should* look like, and ensuring that those with the greatest stake—the growers—have a direct opportunity to contribute to what it *will* look like.

1.3 Securing Value… For Whom?

The notion that growers will benefit from digital transformation is often made implicitly, but will not necessarily be the case. This becomes apparent when recognising that digital transformation is only the most recent incarnation of broader patterns of mechanisation and automation that have continuously driven change since the start of the industrial revolution, with attendant consequences such as the large-scale reduction in agricultural production as a proportion of GDP (Herrendorf et al, 2014), and the associated net migration of workers from rural to urban centres.

Of course, these changes may not necessarily be negative, with losses in parts of the economy offset by new value created elsewhere. Yet this report recognises there may be significant caveats about the capacity of digital technologies to create direct benefits to growers and their regional economies. Indeed if managed poorly, forces of information asymmetry, market concentration and lack of digital sovereignty that are often associated with sectoral digitisation (e.g., Regan, 2019; Zscheischler et al, 2022) could further entrench economic disadvantage within Australia's rural sector. These issues have been elaborated upon in a perspective paper (Box 2).

At this point it is important to stress that VitiVisor does not provide an answer to these questions, which are complex and require a broader strategic and sector-wide response. What it does do is provide an opportunity to reflect and ask the questions: What does a desirable future look like for digital viticulture (Figure 1)? And how can the technological and institutional foundations created through VitiVisor allow the Australian viticulture industry to create its own future?

The method of realising that future will be set by the VitiVisor adoption and commercialisation strategy, which has yet to be fully determined. Section 5.5 summarises some of the alternatives that have been identified, and no doubt others will emerge. Yet arguably the process for making strategic decisions is as important as the strategy itself, and this is now discussed.

[2] Using the model of the adaptive cycle, this phase might be termed the 'exploitation' phase that is characterised by both rapid growth and increasing connectedness.

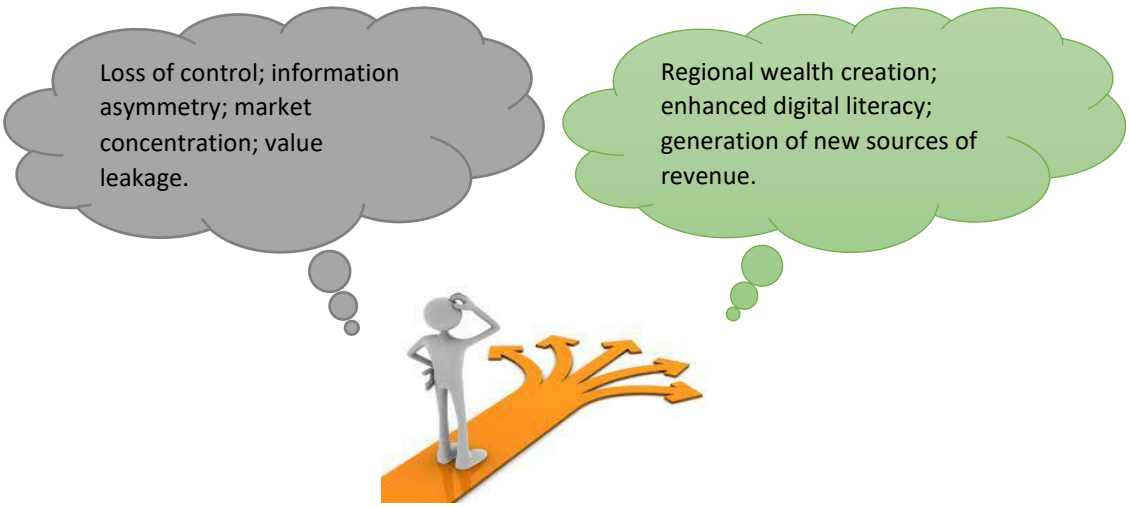


Figure 1: Stylized depiction of alternate future images associated with the digital transformation

Box 2: Research contributions on the digital agriculture transformation

A discussion on broad strategies to unlocking value from the digital transformation in agriculture have been articulated as part of this project in:

> Knowling, M.J., Westra, S., Wheeler, S., Helfgott, A., Lowe, A., Nott, B. Unlocking value from the digital agriculture transformation. Manuscript in preparation.

This perspective article brings concepts from the social sciences and economics literature to a more techno-centric digital agriculture audience. This allows us to tackle 'value' from a more expansive, humanist view, considering the conflict in what value may look like from different perspectives (e.g., public/private good, farmers and technology companies). We discuss how navigating different value perspectives as well as the forces of information asymmetry, market concentration and lack of digital sovereignty will be necessary to unlock the value of digital agriculture. We advocate that a balance be carefully sought between planned and emergent strategy development approaches to better manage and adapt digital agriculture into the future, drawing examples of when and where different actors such as governments can play a role in ensuring that digital agriculture creates value for all in the future.

1.4 'Control your Data, Control your Destiny'

This phrase became the theme of a workshop in the Riverland on 22nd June 2021 and is based on the idea that if Riverland growers are to navigate forces of digital disruption, any strategic decisions must be based on the principle of local empowerment and choice.

As such, grower involvement—both directly and through Riverland Wine—was the central principle of the project that affected all aspects of project design and development. This principle was practically realised in the co-design process prior to project inception (Section 2), and embedded in project governance and decision making at all levels (Section 3). Moreover, a significant portion of research was undertaken in operational vineyards (Section 4), and Riverland Wine played the key role in motivating the need to evaluate potential alternative adoption and commercialisation pathways (Section 5). Our recommendations on possible next steps (Section 7) were designed with this principle in mind.

We now turn to a description of the project aims and the journey that led to their formulation.

2 PROJECT AIMS

Project aims characterise the planned project activities, and are embedded within the formal agreement between Wine Australia, Riverland Wine and the University of Adelaide. We first summarise the co-design process that informed the development of the project aims, followed by a set of core design principles that were embedded throughout the project lifecycle. Lastly, the aims, as well as their alignment to grower concerns and the design principles, are summarised.

2.1 ANTECEDENTS: CO-DESIGN PRIOR TO PROJECT INCEPTION

As described in Section 1, the project mission is to support Riverland growers through the co-development of a digital platform that addresses issues primarily related to costs of production. As described in the preceding section, the project aims and objectives were formulated through a deliberate co-design process, spanning approximately 20 months from an initial meeting in April 2018 through to project inception in January 2020. The elements of co-design that preceded project inception are depicted in Figure 2.



Figure 2: Project co-design between Riverland Wine and University of Adelaide in the lead-up to the project

These structured interactions were interspersed with a variety of field trips and informal conversations with growers, Riverland Wine and Wine Australia, which helped shape both the guiding design principles described next, together with the formal project aims.

2.2 GUIDING DESIGN PRINCIPLES

The VitiVisor project has sought to realise a series of design principles that leverage the strengths of this unique partnership between growers, government (though Wine Australia) and the university sector. An early version of these design principles was first described in the pilot report (Westra et al, 2019), and have evolved over the project through a continual learning process. The key principles as experienced throughout the project as described in Figure 3 below:

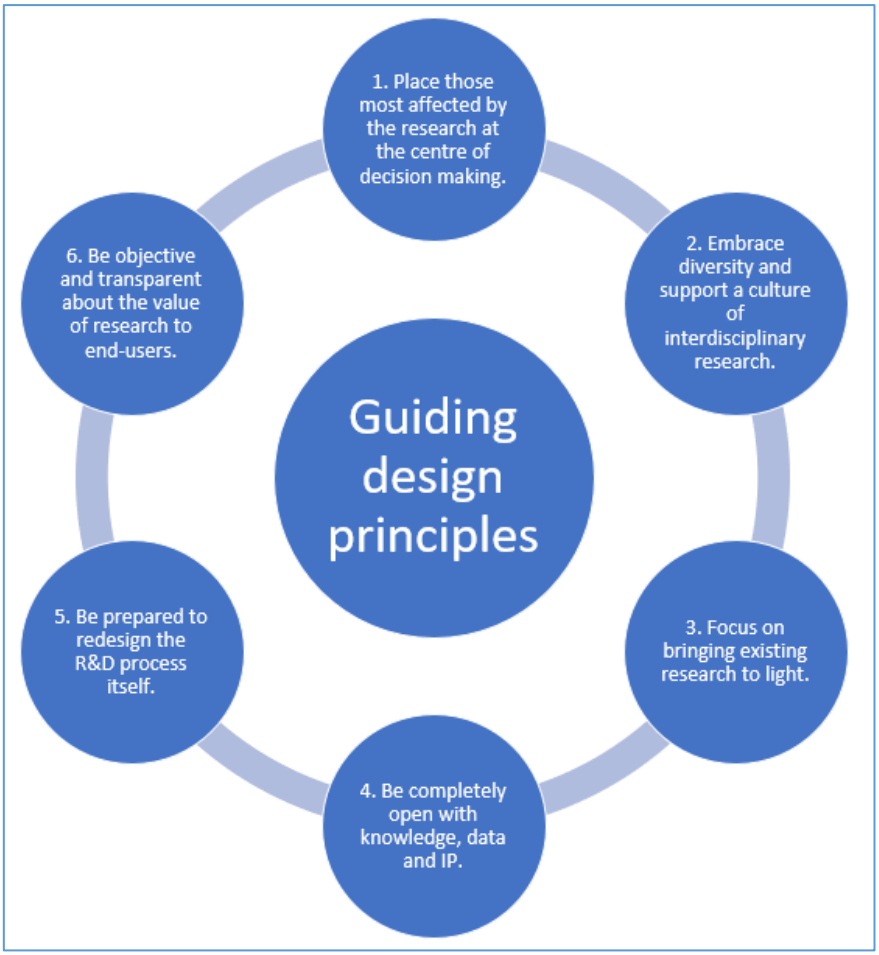


Figure 3: Guiding design principles that have informed the development of VitiVisor

2.3 Project Aims and Alignment with Grower Priorities

The grower priorities (Box 1) and design principles led to the following project aims.

Project Aim 1. A data collection and management system that captures, sends and stores data from vineyard 'Internet of Things' sensors, airborne, external third-party sensors and other data streams, creating a 'one-stop-shop' of all relevant datasets.

Project Aim 2. An operational ground-based vision system that cost-effectively captures and monitors canopy structure and development, as well as yield components, using existing and new vision technologies, and makes key labelled datasets openly accessible for community development.

Project Aim 3. A prediction and advisory system that focuses on relating management actions to vineyard performance measures (e.g. yield) particularly in the context of water management, using a combination of process-based and machine learning techniques.

Project Aim 4. On-going estimation of the net economic benefits of improving decision-making of Riverland viticultural producers to increase the overall sustainability of the industry, and to help steer and prioritise system development.

Project Aim 5. A dashboard (front end – 'VitiVisor') that will allow growers to visualise information, prediction and advisory services developed in the project, which will be developed in close consultation with Riverland growers who will help prioritise functionality and visualisation. In addition to the front-end, a backend software system that connects with the data management system and integrates all modules.

The alignment between project aims, grower priorities (Box 1) and design principles (Section 2.2) are given in Table 1.

Table 1: Alignment between project aims, grower concerns and research principles

Aim	Link to grower priorities and the design principles
1. Data collection and management system	Growers identified the fragmented nature of existing digital services and data providers as a constant source of friction, and also highlighted concerns around data ownership (Design Principle 4). By developing a platform that can collect and manage data from diverse sources, it becomes possible not only to provide a 'one-stop-shop' for data, but potentially also create a longer-term digital research environment to lower the cost of research 'on the data' (Design Principle 5) and share insights across disciplines (Design Principle 2).
2. Operational ground-based vision system	Ground-based vision was identified in the pilot project as a practical alternative to drone-based vision with multiple operational benefits. Although numerous private competitors have emerged since project commencement (e.g. BitWise and Green Atlas), there are likely to be competitive advantages associated with cost, customisation to the Riverland context, and scientific scrutiny of results. This aim helps accelerate the development of this technology (Design Principle 3) and places decisions related to future development in the hands of growers (Design Principle 1).
3. Prediction and advisory services	The VitiVisor project provides prediction and advisory services that focus first and foremost on what was identified as growers' primary concern: using digital technology to reduce costs of production. Many of the scientific insights that could enable prediction and advice were already available in the academic literature and in intellectual property developed over a period of multiple decades, although investment was needed to integrate and contextualise these insights into a product that could be used by growers (Design Principle 3).
4. Estimation of economic benefits	In parallel with the development of the VitiVisor platform, an objective, transparent and data-driven approach to assessing the economic value of digital technologies was completed (Design Principle 6). This helps create alignment between development work and grower value, and ensures a strong focus on addressing cost of production issues.
5. Dashboard 'front-end'	Early experience highlighted that the feedback cycle between growers and researchers needed to be shifted from conceptual discussions to the use of mock-ups or early versions of a potential end-product (Design Principles 1 and 3). This helped maximise alignment between growers and researchers on development priorities, and ensured that growers were closely involved in shaping how their levies were used to support research prioritisation.

3 PROJECT GOVERNANCE AND GROWER ENGAGEMENT

Whereas the preceding section focused on engagement activities in the lead-up to project initiation, this section focuses on approach to governance and engagement over the period from January 2020 to June 2022. The overarching approach to project governance and the interaction between the three key parties (Riverland Wine, Wine Australia and University of Adelaide) is depicted in Figure 4.

Core to the establishment of governance structures was the need to ensure that growers have a central role in shaping the project, contributing to key design decisions and importantly, participating in the management and decision-making structures throughout the project. This helps ensure both that the development of VitiVisor met grower needs, and to support subsequent adoption of the technology.

> "Riverland growers are core members of the design team throughout the design and development process. This ensures that the developed system meets key grower 'pain points' and maximises the likelihood of wide-spread adoption"

As a result, each of the project governance structures depicted in Figure 4 was designed to maximise the interface with growers at all levels. Further details on these structures are now provided.

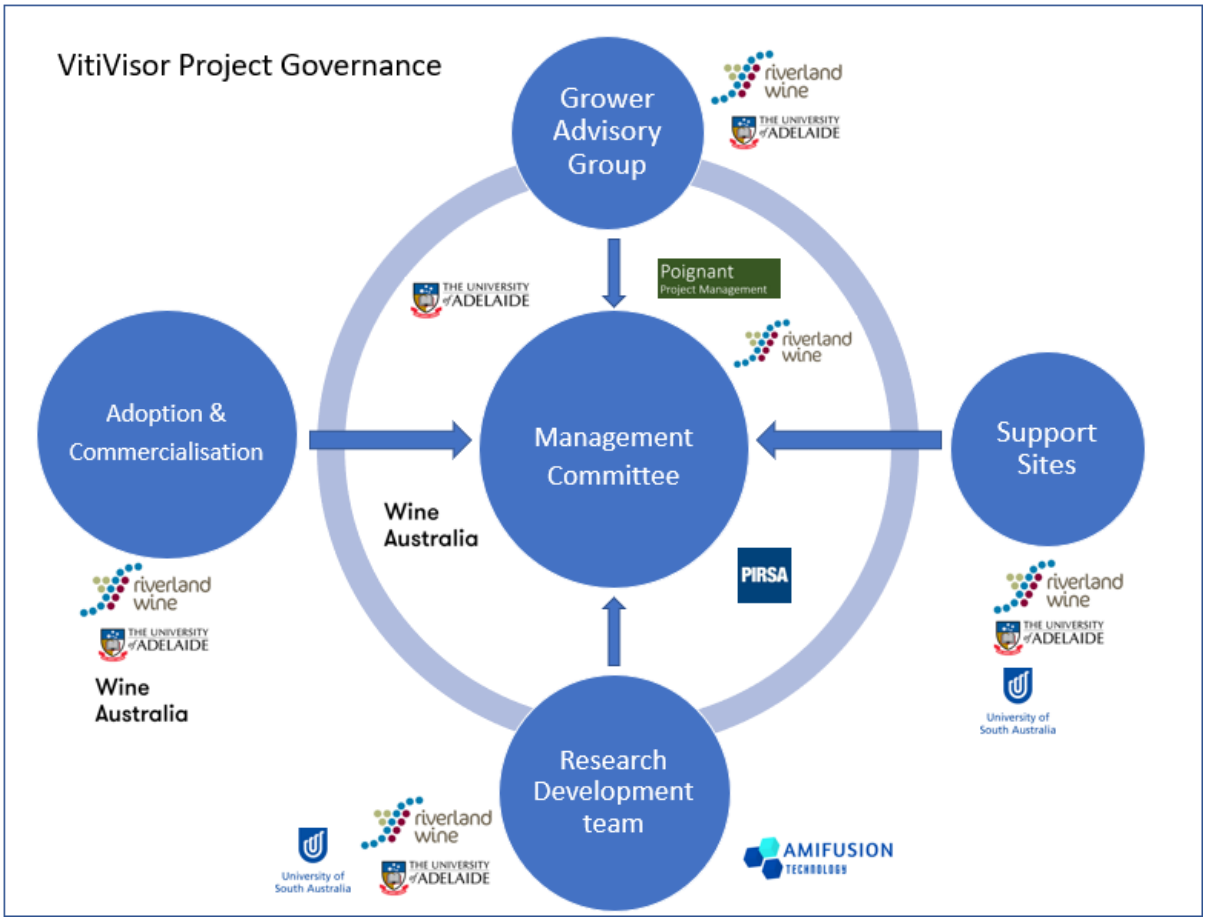


Figure 4: Approach to project governance

3.1 Management Committee

The management committee was the key governance committee for VitiVisor and was chaired by Professor Andy Lowe from the University of Adelaide. This committee was made up of the three primary parties of the agreement (RW, WA and UoA), as well as PIRSA given the central role of the Loxton Research Centre to this project. This committee initially met on a monthly basis during the project establishment phase, and then transitioned to quarterly.

Responsibilities of the management committee included:

- Approving strategy
- Monitoring progress relative to the annual Work plan milestones
- Evaluating risk and recommending mitigation
- Ratify decisions recommended from sub-committees
- Providing stakeholder / industry updates
- Biannual reporting to Wine Australia

The key decisions made in these meetings were recorded in a decision register, accessible through a collaborative, digital platform.

3.2 Adoption and Commercialisation

The adoption and commercialisation committee met over thirteen occasions through the project and provided the primary coordinating point for all discussions related to these topics. Made up of representatives from RW, WA and UoA, the committee's activities included:

- Developing strategy for approval by the management committee
- Coordinating and liaising with agTech companies to create awareness of VitiVisor, through presentations by team members
- Seeking expert advice
- Advise the team on IP pathways

The committee was joined on multiple occasions by Francois Duvenage, Commercial Manager, Innovations and Commercial Partner (on behalf of University of Adelaide) and also by Doug Adamson, DNA Innovation Pty Ltd (on behalf of Wine Australia). Technology companies contacted during the project timeline are summarised in Appendix 2.

3.3 Development Team

The development team was made up of all the project output leads together with representatives from Riverland Wine, and met up on a monthly basis to discuss progress updates, key design decisions and interfaces across the project. This team also provided the primary interface between Hao Zhang of AMI Fusion Pty Ltd who was responsible for software development, and Joanne Tingey-Holyoak from the University of South Australia who provided specialist accounting advice to the finance and economics team.

Project Output Leader	**Responsibilities**
Prof. Seth Westra	Chief Investigator, and Lead Researcher for Modelling and Prediction
Prof Javen Shi	Lead Researcher for Machine Learning and Dashboard
Dr Tien Fu Lu	Lead Researcher for the Ground-based Vision System
Prof. Sarah Wheeler	Lead Researcher for the Economic Benefit
Assoc. Prof Cass Collins	Lead Researcher for Viticulture
Mr Brad Nott	Project Manager

3.4 Grower Advisory Group

The VitiVisor user interface is an important aspect of the successful adoption of technology. A grower advisory group consisting of eight Riverland growers was established to consult with and test the proposed User Experience (UX) of the VitiVisor dashboard. The grower's immediate feedback on the look and feel of the dashboard proved invaluable, with suggested new functionality (e.g. display of growth stages and management tasks) being incorporated into the final design.

- Mr Ben Haslett, Woolenook Fruits
- Mr David Zadow, Paisley Hill
- Mrs Sheridan Alm, Starrs Reach
- Ms Nicole Pitman, Kingston Estate
- Mr Brian Caddy, Limestone Ridge
- Mr Chris Byrne, Gums n Roses
- Mr Brett Proud, Sherwood Estate
- Mr Brett Rosenzweig, Rosy Ridge Farms

3.5 Support Sites

Riverland Wine growers that represented a range of different growing conditions and viticultural practices were selected from each of the main regions across the Riverland: Waikerie, Barmera, Loxton and Renmark. Figure 2 illustrates the purpose and activities of the different types of support sites developed for the project. Researchers worked collaboratively with their stakeholders, establishing a control site at the Loxton Research Centre, connecting three growers to a VitiVisor dashboard proving the feasibility of the connection of their existing IoT vineyard systems. Moreover, viticulture data from 14 patches was captured across two seasons with twenty five growers being interviewed for economic case studies.

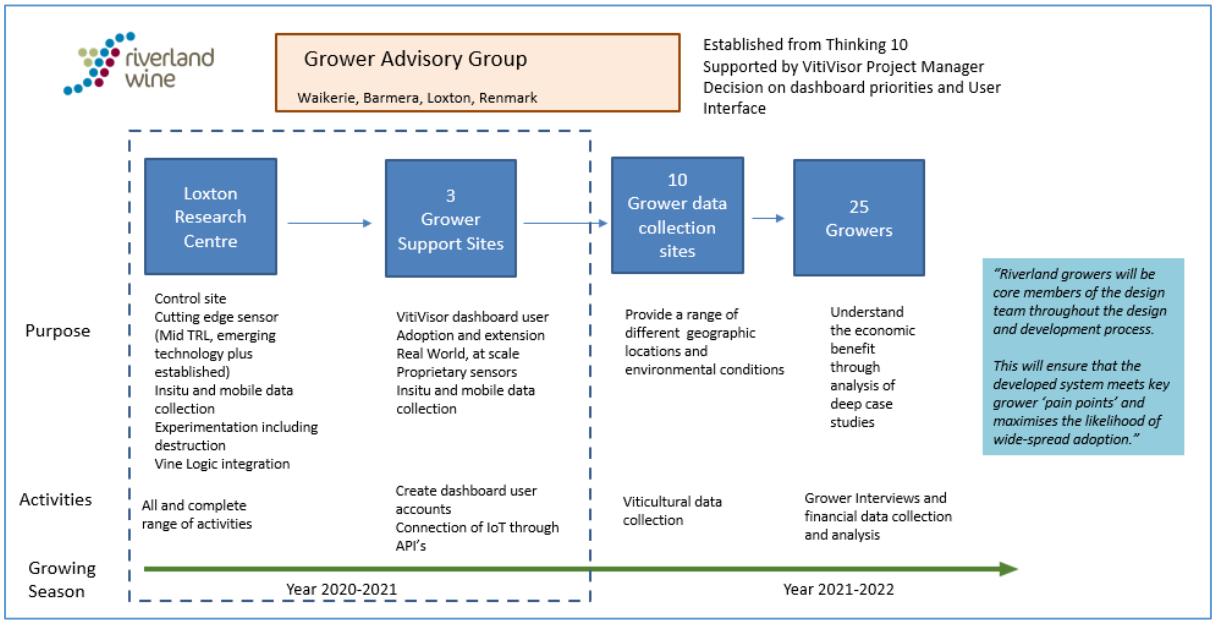


Figure 5: Grower engagement through support sites

4 METHOD

VitiVisor follows a 'data-to-decisions' pipeline, with the primary data categories collected within VitiVisor depicted in Figure 6. The approach to data acquisition for the majority of the data categories is summarised in Section 4.1, and the development of a new ground-based vision system that is capable of monitoring diverse measures of canopy development is then described in Section 4.2. The capacity to convert this data into forward-looking predictions is then described in Section 4.3, including the application of this capability for scenario-based decision making. The development of a detailed financial benchmarking dataset is then described in Section 4.4, as well as an evaluation of the potential return on investment to growers of using this technology. Finally, the approach to displaying all this information on a dashboard is presented in Section 4.5.

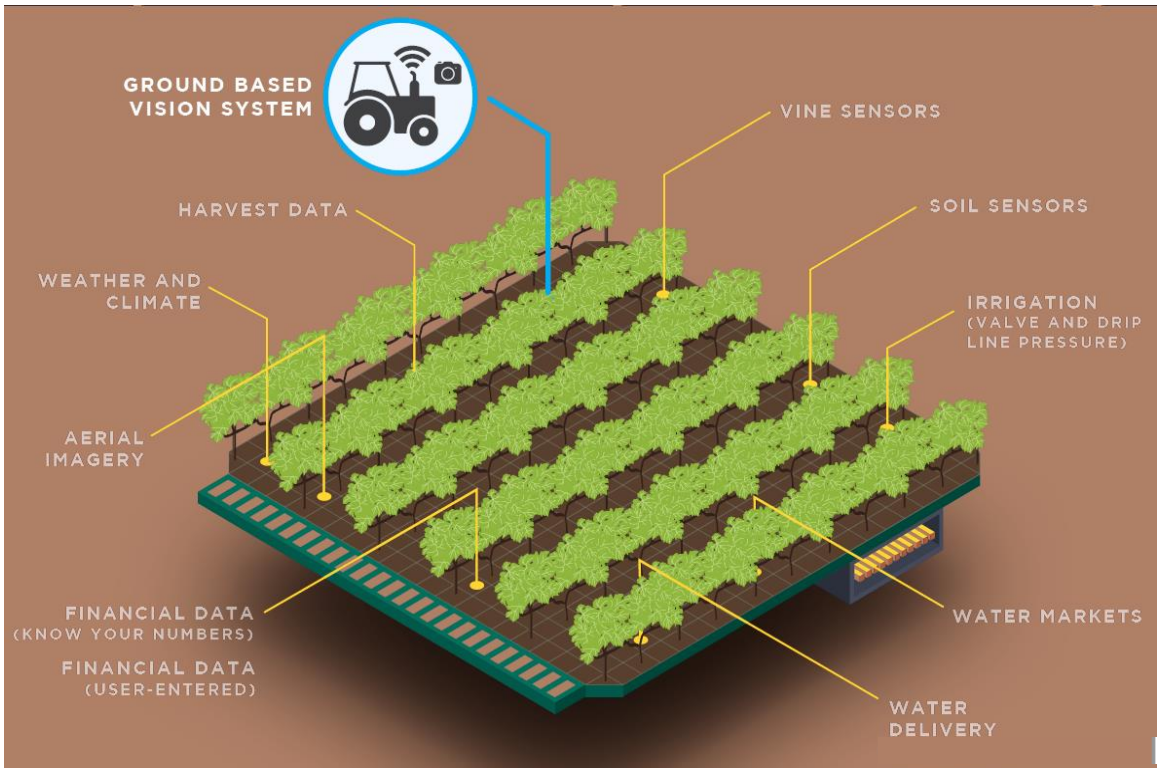


Figure 6: Primary data categories integrated into the VitiVisor platform.

4.1 DATA ACQUISITION

This section describes the process of data acquisition required to support the VitiVisor platform based on three primary data classes: (1) vineyard measures derived from the ground-based image acquisition system; (2) data from 'Internet of Things' sensors; and (3) third-party datasets. These are discussed in turn below.

4.1.1 Ground-based vineyard measures (ground truth and image acquisition)

4.1.1.1 Sampling strategy

Fourteen vineyard patches in the Riverland were selected for ground-based vineyard measures. Vineyards were selected based on the willingness of owners to participate in the study and to capture the variation in vineyard management across the region. Details of the selected varieties and sampling strategy are provided in Table 2.

Table 2: Variety and sampling frequency details for each patch.

Patch	Variety	Panel Frequency	Row Frequency	PAI image count / visit
Patch 01	Chardonnay	Every available	Every row	~ 460
Patch 02	Shiraz	Every available	Every row	~ 690
Patch 03	Shiraz	Every 2nd	Every row	~ 710
Patch 04	Shiraz	Every 2nd	Every row	~ 660
Patch 05	Chardonnay	Every 10 metres	Every row	~ 380
Patch 06	Cabernet Sauvignon	Every 2nd	Every 2nd row	~ 1,200
Patch 07	Chardonnay	Every 2nd (approx.)	Every row	~ 2,000
Patch 08	Cabernet Sauvignon	Every 2nd	Every 2nd row	~ 1,850
Patch 09	Shiraz	Every 2nd (approx.)	Every 2nd row	~ 1,200
Patch 10	Mataro	Every 2nd	Every 2nd row	~ 1,130
Patch 11	Cabernet Sauvignon	Every 2nd	Every 2nd row	~ 1,260
Patch 12	Chardonnay	Every 2nd (approx.)	Every 2nd row	~ 940
Patch 13	Pinot Gris	Every 2nd	Every 2nd row	~ 1,980
Patch 14	Shiraz	Every 2nd	Every 2nd row	~ 740

At each vineyard patch, images were collected at key stages of phenological development, and where applicable ground truth measurements were also recorded. Details of the frequency and resolution of data collection are shown in Table 3. Data collected from patches 01 and 02 had greater resolution and sampling frequency than other patches in the first season (2020/21). In the second season (2021/22), patches 01 to 05 had greater sampling frequency compared to the other patches; these patches were sampled at higher frequency due their size, geographic spread and as examples of contrasting management systems.

Table 3: Total patch visits and images taken, and total visits per season

Patch	Variety	2020/21 Season	2021/22 Season	Total Visits	Total Images
01	Chardonnay	37	33	70	31,558
02	Shiraz	37	30	67	44,638
03	Shiraz	5	30	35	21,920
04	Shiraz	5	30	35	18,684
05	Chardonnay	3	29	32	13,773
06	Cabernet Sauvignon	3	2	5	4,090
07	Chardonnay	3	3	6	5,358
08	Cabernet Sauvignon	4	3	7	4,340
09	Shiraz	6	4	10	4,406
10	Mataro	5	3	8	4,356
11	Cabernet Sauvignon	5	4	9	5,059
12	Chardonnay	5	3	8	3,786
13	Pinot Gris	5	4	9	5,876
14	Shiraz	4	3	7	3,032
TOTAL					**170,876**

At each site, 10 ‘sentinel frames’ were established within the vineyard at locations selected from vineyard visits during pruning in 2020 and via multi-season satellite NDVI images to capture vineyard

patch variation. Each sentinel frame represents the area captured by the data collection tool VitiCanopy utilising an inbuilt phone camera. Sentinel frames were set out as per the field of view of the cameras shown in Figure 7. The length of the frame on each cordon was measured and tagged to allow for repeatable image and measurement capture. Physical counts (i.e. bud numbers etc.) were made on a per 'frame' basis and relate to a corresponding image. Details of ground truth measurements and developmental stages of image data acquisition for each season and patch are outlined in Table 4.

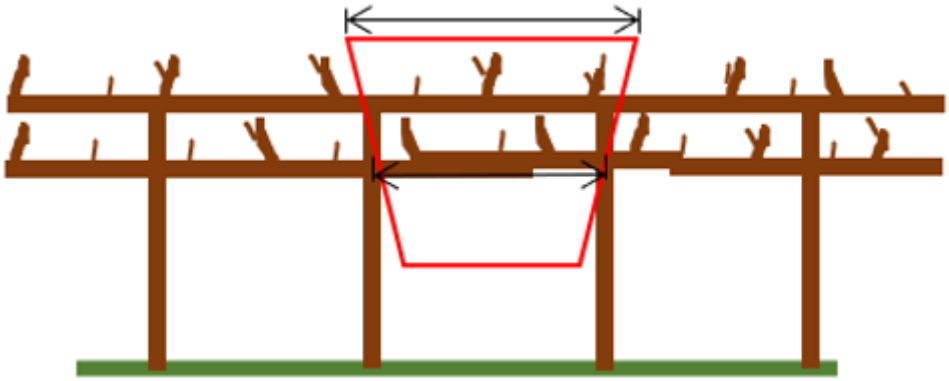


Figure 7: Frame capture representation, based on upward looking imagery captured with smartphone, the upper and lower bounds were tagged to allow repeatable imaging and set-out distance was standardised at each site.

Table 4: Ground truth measures and developmental stage of image acquisition captured during each season at each patch.

Patch	EL Stage Image acquisition 2020/21	Ground truth measures 2020/21	EL Stage Image acquisition 2021/22	Ground truth measures 2021/22
Patch 01	1, 2, 7, 9, 13, 15-18, 23, 25, 29, 31-38, 41, 43, 47	bud no., shoot no., inflorescence no., bunch no., bunch mass	1, 4, 11-12, 15-19, 23, 25, 26, 27, 29, 31-39, 41, 43, 45	pruning wt., bud no., shoot no., inflorescence no., bunch no., bunch mass
Patch 02	1, 2, 5, 11-15, 17-19, 26, 29, 31- 39, 41, 43, 47	bud no., shoot no., inflorescence no., bunch no., bunch mass	1, 4, 9, 13-14, 16-17, 18, 19, 25-27, 29, 30, 32-38, 41, 43	pruning wt., bud no., shoot no., inflorescence no., bunch no., bunch mass
Patch 03	1, 4, 18, 38, 47	bud no., shoot no., inflorescence no., bunch no., bunch mass	1, 11, 12, 14, 15, 17, 18, 19, 25-30, 32-39, 41, 43	pruning wt., bud no., shoot no., inflorescence no., bunch no., bunch mass
Patch 04	1, 2, 17, 38	bud no., shoot no., inflorescence no., bunch no., bunch mass	1, 3, 9, 11-13, 15, 17, 18, 19, 20, 25-30, 32-38, 41, 46	pruning wt., bud no., shoot no., inflorescence no., bunch no., bunch mass
Patch 05	1, 3, 13, 38, 47	bud no., shoot no., inflorescence no., bunch no., bunch mass	2, 11, 13, 14, 16, 19, 21, 23, 26-27, 29, 31-39, 41, 46	pruning wt., bud no., shoot no., inflorescence no., bunch no., bunch mass
Patch 06	1, 12, 19, 41, 47	pruning wt., shoot no., inflorescence no.	1, 2, 11, 16, 19, 38	pruning wt., bud no., shoot no., inflorescence no., bunch no., bunch mass
Patch 07	1, 4, 15, 38	pruning wt., shoot no., inflorescence no., bunch no., bunch mass	4, 13, 16, 38	bud no., shoot no., inflorescence no., bunch no., bunch mass
Patch 08	1, 7, 16-17, 38, 47	pruning wt., shoot no., inflorescence no., bunch no., bunch mass	1, 3, 14, 37	pruning wt., shoot no., inflorescence no., bunch no., bunch mass
Patch 09	1, 12, 17, 19, 38, 47	pruning wt., shoot no., inflorescence no., bunch no., bunch mass	1, 12, 17, 19, 38, 47	pruning wt., bud no., shoot no., inflorescence no., bunch no., bunch mass
Patch 10	1, 2, 21, 38	bud no., shoot no., inflorescence no., bunch no., bunch mass	1, 5, 16, 38	pruning wt., bud no., shoot no., inflorescence no., bunch no., bunch mass

Patch 11	1, 2, 14-16, 38, 47	pruning wt., bud no., shoot no., inflorescence no., bunch no., bunch mass	1, 2, 16, 38	pruning wt., shoot no., inflorescence no., bunch no., bunch mass
Patch 12	1, 2, 5, 17, 18, 38, 47	pruning wt., bud no., shoot no., inflorescence no., bunch no., bunch mass	4, 19, 38	pruning wt., bud no., shoot no., inflorescence no., bunch no., bunch mass
Patch 13	1, 2, 15-16, 38, 47	bud no., shoot no., inflorescence no., bunch no., bunch mass	1, 11, 16, 38	pruning wt., bud no., shoot no., inflorescence no., bunch no., bunch mass
Patch 14	1, 12, 18, 38, 47	pruning wt., bud no., shoot no., inflorescence no., bunch no., bunch mass	11, 19, 37	pruning wt., bud no., shoot no., inflorescence no., bunch no., bunch mass

4.1.1.2 Imagery acquisition and labelling

Images were acquired using an Apple smartphone (iPhone Model 7, 8 and 11) and processed using VitiCanopy, a smartphone application that uses images to calculate plant area index (PAI) and canopy porosity (De Bei et al., 2016). Image resolution and other details of phone-based image capture are provided in Table 5.

Table 5: Details of phone and associated camera for image capture

iPhone Model	Camera used	Image Format	Resolution	Date Used
7	Front	Landscape	3,088 x 2,320	Up to & including 23/11/2020
8	Front	Landscape	3,088 x 2,320	Only on 5/06/2020
11	Front	Landscape	4,032 x 3,024	2/12/2020 & after

Two scales of image collection were captured—the first being from each of the 10 sentinel capture sites at each vineyard patch and the second being the collection of images at a vineyard scale. At patches 1 and 2, images were captured from every unobscured panel in every row and for patches 3-5 from every second unobscured panel in every row. Images were also used to develop machine learning algorithms for bud, inflorescence and shoot counts. This has resulted in 170,876 images. Images processed have been mapped using their GPS coordinates within the VitiCanopy app. Imagery and ground truth data was captured by researchers on foot in each of the patches.

4.1.1.3 Ground truth measures

The computer vision algorithms described in Section 4.2.3 were trained on annotated imagery using ground truth measures that were deemed to be decision-relevant to growers, based on a combination of viticultural expertise and grower feedback. In total, eight measures were calculated and summarised below:

- **Pruning weight** was measured as it gives an indication of the vegetative growth of the vine during the season. Pruning weights were recorded during the winter dormancy period. These measures were collected by cutting all canes to two node spurs and weighing the total amount of material removed per sentinel capture.
- **Inflorescence** counts were recorded in spring before flowering in both seasons
- **Shoot counts** were recorded in spring before flowering in both seasons.
- **Bunch number** and **bunch mass** were two key yield components recorded just before commercial harvest at each vineyard patch.
- **Bud number** was included following pruning all buds on one year old wood were counted.

- **Cane weight** was then calculated (pruning weight/cane number) and gives a useful indication of shoot vigour.
- **Plant Area Index** was calculated at different developmental stages using the VitiCanopy App.

Unless indicated otherwise, measures were described as a count per sentinel frame.

For machine learning analysis, images from vines that were assessed for ground truth measures (buds, inflorescence and shoot tips) were manually annotated using 'labelImg' software (version 1.8.1). For each image, two qualified viticulturists manually annotated a subset of the same images to provide quality assurance, with the process illustrated in Figure 8 below. Annotated images were the same capture used for ground truth assessments of bud, inflorescence and shoot tip counts from sentinel vines at each patch over two growing seasons.

A total of 170,876 images were acquired to determine Plant Area Index. A selection of images were annotated for machine learning purposes. Ground truth measures were recorded for eight distinct measures identified by growers as important for vineyard management.

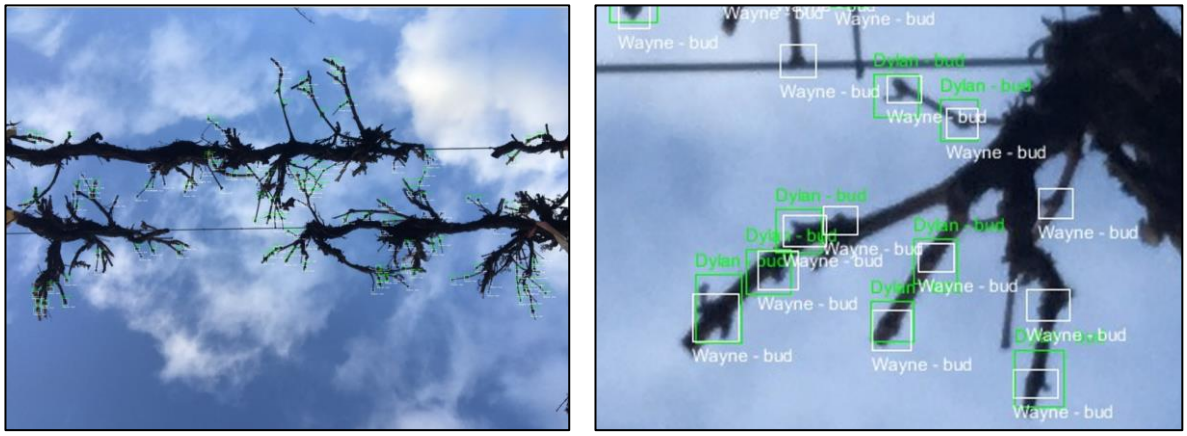


Figure 8: Approach to ground truthing, in which key measures were manually labelled using the 'labelImg' software. Left image shows a single sentinel image, and right image shows a zoomed-in version.

Imagery and ground truth measures were made available to the modelling and machine learning team members for further analysis. Data has also been made available in an open source way via Figshare (see www.vitivisor.com.au for links).

4.1.2 'Internet of Things' sensor data

Internet of Things (IoT) technology is an integral part of VitiVisor's overarching data acquisition and storage system that collects real-time data from sensors on the vineyard, connecting to the VitiVisor dashboard via LoRaWAN. The focus of this project was to develop a low-cost sensor network and supporting infrastructure for demonstration purposes, focusing on the Loxton Research Centre viticulture trail environment.

The approach taken in VitiVisor is illustrated in Figure 9, and comprises a LoRaWAN gateway, The Things Network (TTN) architecture and a series of sensors. The LoRaWAN gateway, which has been made available for public use, is described further in Box 3.

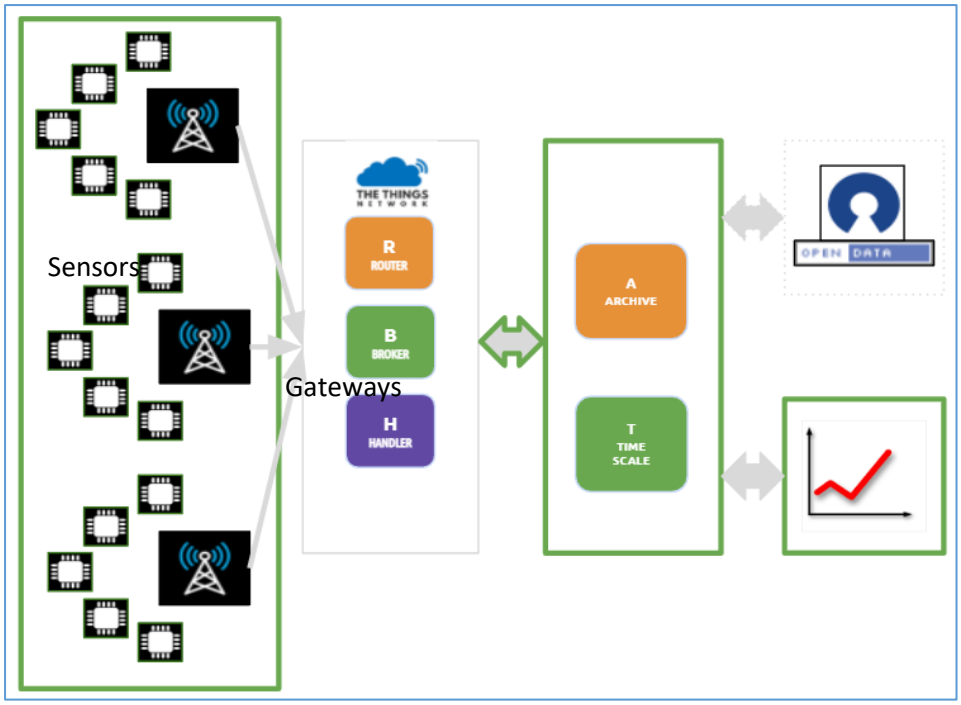


Figure 9: IoT via TTN project architecture

Box 3: Connecting to the Loxton LoRaWAN network

As part of the VitiVisor project, a low-power wide area networking (LoRaWAN) gateway has been installed, tested and validated in the Loxton Research Centre. The purpose of this infrastructure is to enable wireless connection of devices to the internet, and manage communication between end-node devices and network gateways.

The VitiVisor project has set up the network so that anyone within an approximate 45 km radius and with a LoRaWAN node can connect to LRC gateway. This low cost solution means that instead of using a telco sim card with monthly subscription, users could choose to opt for LoRaWAN.

To get connected to the public network, simply configure your LoRaWAN Node to AU915 within the coverage area and follow the instructions from the link below.

https://www.thethingsindustries.com/docs/devices/

In this project, various sensors were tested out using different LoRaWAN transmission devices focusing on ease-of-use and cost. As a result, a node from ICT International (Figure 10) was chosen due to the following characteristics:

- Capability of node to connect with various legacy sensors with SDI-12 connections;
- Good circuit design allowing for quick-connect;
- Flexibility to modify data output (e.g. frequency of data output); and
- Cost effectiveness.

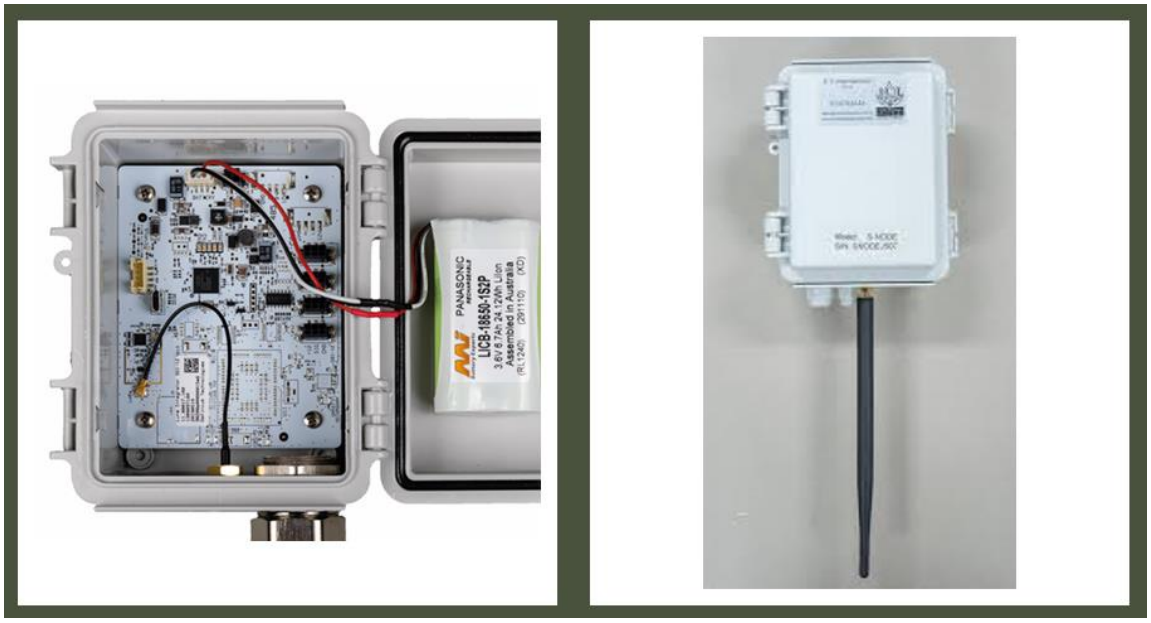


Figure 10: S-Node from ICT International

Table 6: Inventory of IoT sensors installed during the project

Data Type	Sensor type	Sampling frequency
Soil moisture	In-ground sensor	60 minutes
Soil temperature	In-ground sensor	60 minutes
Salinity	In-ground sensor	60 minutes
Vine water use	Dendrometer	15 minutes
Air temperature	Weather station, under canopy sensor	30 minutes
Atmospheric humidity	Weather station, under canopy sensor	30 minutes

The list of sensors that have been connected to the LoRaWAN gateway is summarised in Table 6, with the installation locations illustrated in Figure 11.

The data acquired from the installed sensors is transferred to TTN through LoRaWAN. Within TTN, data payload is decrypted, processed and stored on a MQTT layer ready for retrieval by the Vitivisor Dashboard. In-depth information of the data transmitted from the sensors and LoRaWAN Nodes (Figure 12) can be accessed and diagnosed. Each data packet transmitted by the node can be tracked as TTN, and has the capability to show the trail of data coming out from every sensos, including the strength of the LoRaWAN connection at the point of time during data transmission. Importantly, the validity of data source can be traced and erroneous data packets can be omitted.

The outcome of data transmission can be reflected onto the dashboard (Figure 13), where further farm management decisions can be made on a day-to-day basis.

The key idea of the IoT component of VitiVisor is to demonstrate how low-cost sensors can be connected to a public LoRaWAN network, and provide the enabling software and hardware so that the solution can be replicated by others.

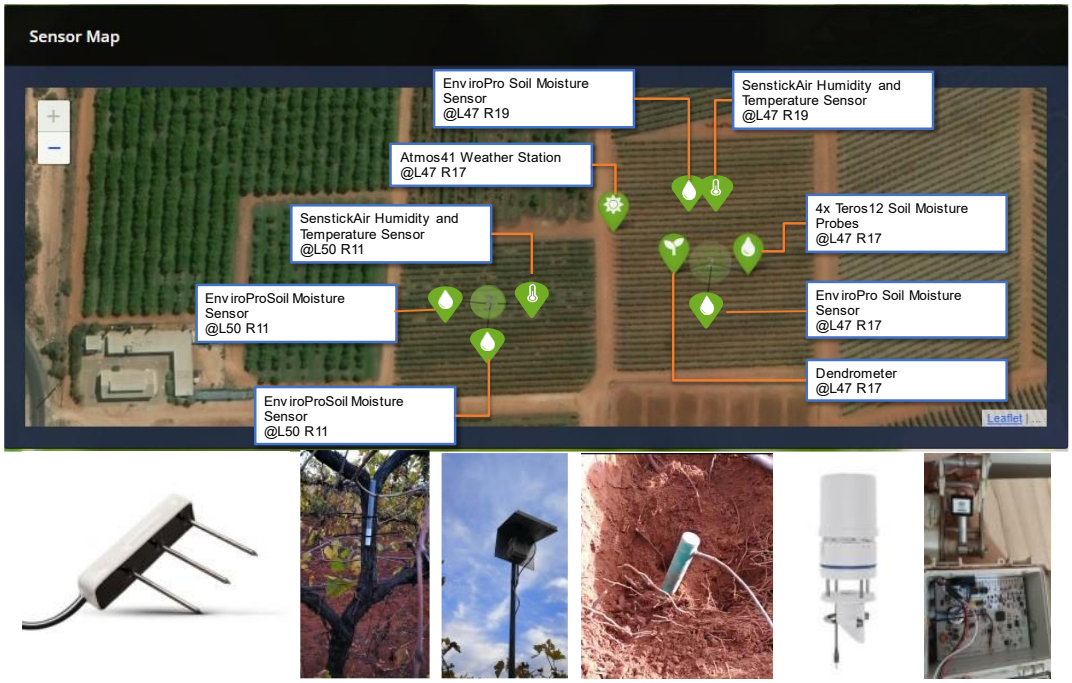


Figure 11: Sensors installed in LRC (Bottom left to right: Teros-12 Soil Moisture Probe; Senstick Air Humidity and Temperature Sensors, LoRaWAN S-Node with Solar panel Enviropro Soil Moisture Stick; ATMOS41 Micro Weather Station; Dendrometer).

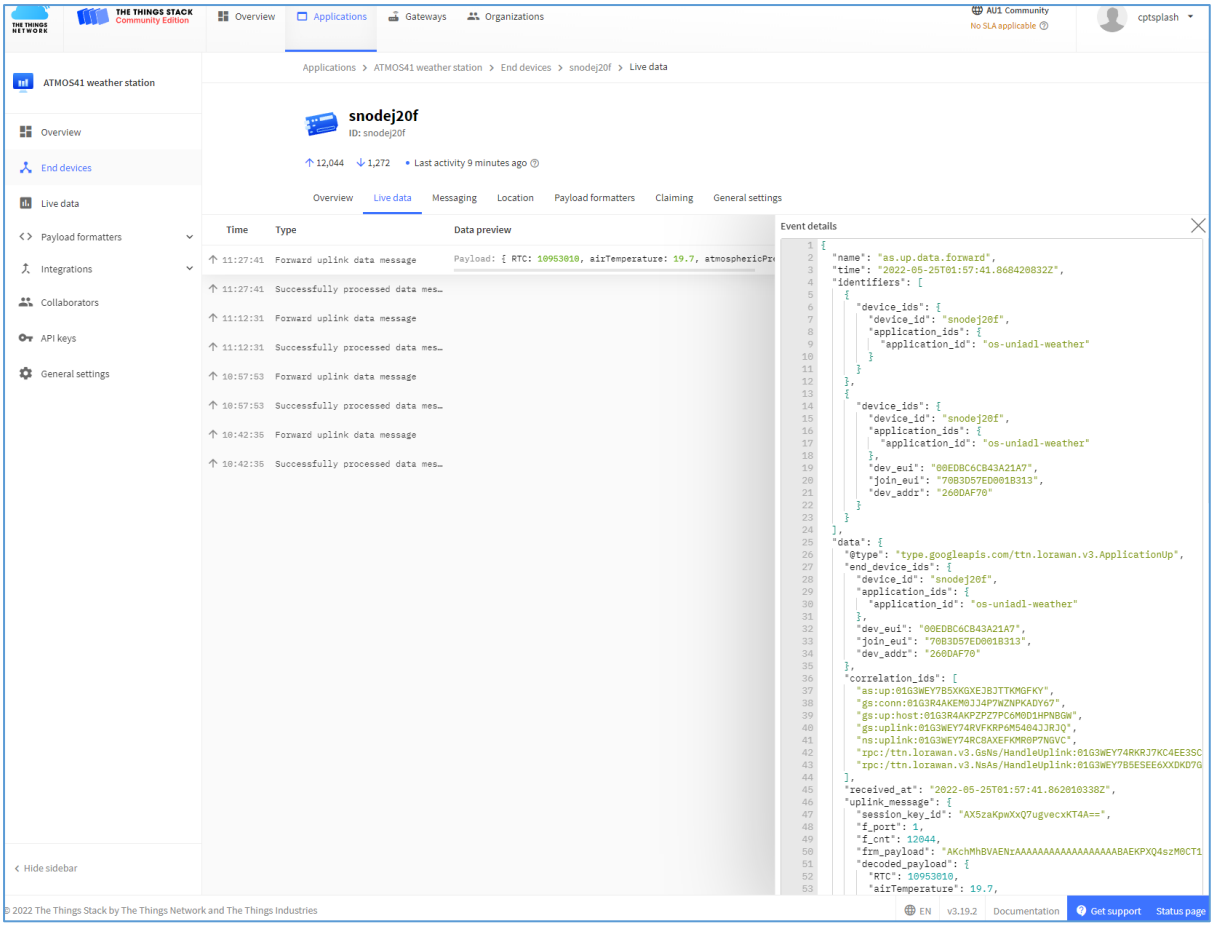


Figure 12: A typical overview of The Things Network (TTN)

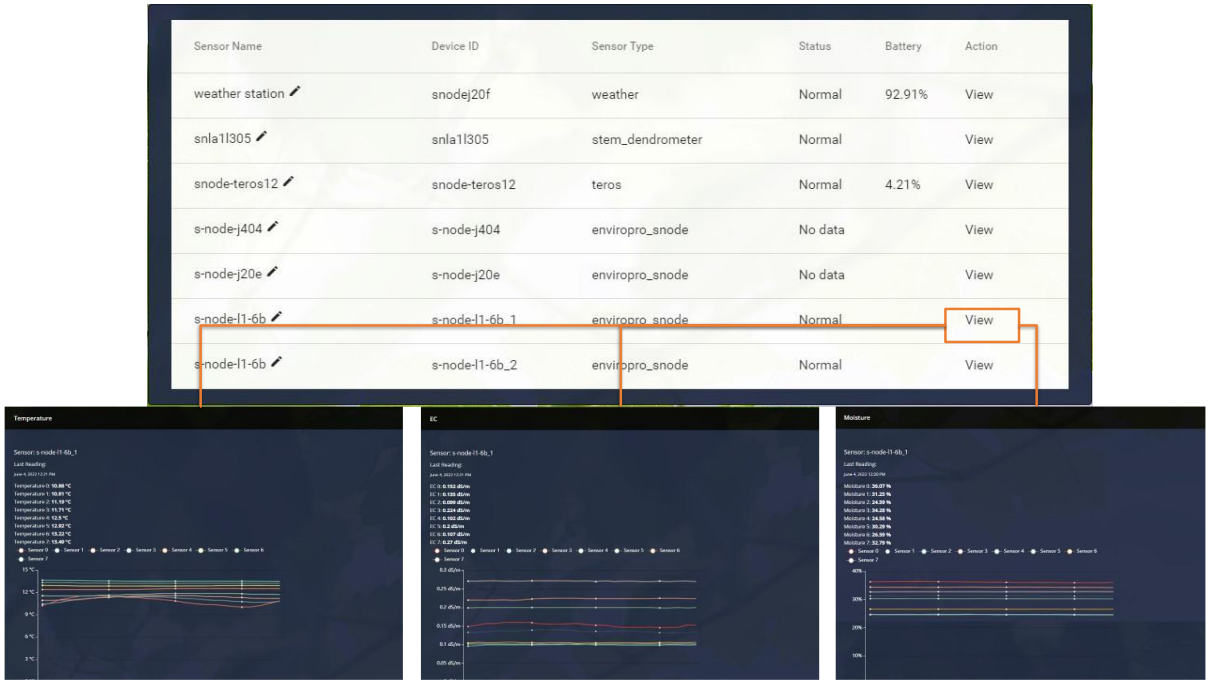


Figure 13: A typical overview of the Vitivisor dashboard on installed IoTs

4.1.3 Third-party data

The initial consultation with Riverland growers indicated the use of multiple applications on their phones, whereby the third-party data was used as part of their vineyard decision making. A key grower request was to have all this third-party data available to them on a single digital platform.

In developing the design of the VitiVisor dashboard a number of key requirements from Riverland Wine guided the selection of third party data and their providers. This was predicated on a number of factors, with two overarching key principles:

- all of the necessary data used for making decision in the vineyard should be available in one place, both through display on a dashboard and to support algorithms that translate this into viticultural advice; and
- service providers participating in VitiVisor act in a transparent, participatory and collaboratively way.

This section describes the key data categories required to support VitiVisor, the effort made for the selection, and connection of third party data providers to the dashboard. This is summarised in Table 7, together with the key drivers for each data category.

Table 7: Primary data categories, and the third-party providers selected in the current version of VitiVisor.

Category	**Third party**	**Driver**
Property management and aerial imagery	Ceres Imagery	Enables high-resolution aerial imagery including NDVI, water stress and thermal imagery
Sensor management	Opensensing Greenbrain	IoT sensors including soil moisture sensors and a dendrometer provide real-time information on key vineyard parameters

Irrigation system	Central and Renmark Irrigation Trusts	Information on water consumption.
VitiCanopy	University of Adelaide	Interim data collection solution during the construction of VitiBox, providing information on key vine development parameters including plant area index, canopy porosity and canopy cover.
Financial	Katalyst (and Xero) Bentley Charted Accountants Water market data – H2OX	Information on financial information, used to support scenario-based prediction tool.

4.1.3.1 Property management

Ceres Imagery Pty Ltd is located within the Loxton Research Centre (LRC) providing high resolution aerial imagery services to a range of agricultural sectors. Through a series of meetings and subsequent negotiation with Ceres, they provided imagery of the LRC as an in-kind contribution to the project, including the two trial patches (Block 47 and 50) used as a primary support site for the project. The imagery used in the VitiVisor dashboard includes NDVI, water stress, and thermal imagery.

A potential future pathway could be that Ceres would engage with Riverland Wine or the CCW Cooperative, providing them with timely seasonal imagery. As a member benefit of one of these organisations, growers would then have access to the Ceres imagery data, which would be made available on VitiVisor via an API that they would create and maintain. This, in-turn, would allow the imagery from Ceres to be combined with the other datasets collected through VitiVisor to provide value-added information, prediction and advisory services.

4.1.3.2 Sensor Management

The pilot project prior to VitiVisor (Integrated Vineyard Precision Control System, University of Adelaide, 2019) established the use of a range of commercial and near commercial *in situ* sensors (e.g. soil moisture and dendrometers) operating in the open source environment of The Things Network (TTN) using LoraWAN connectivity. This was provided by OpenSensing, a South Australian service provider of sensors and expertise in open source networks. On this basis, the VitiVisor project engaged Opensensing for the provision and maintenance of a range of hardware to be installed at the LRC. Further information on this can be found in Section 4.1.2 above.

During the project, an additional connection was made with Greenbrain from MEA, as this is a commercial platform that integrates *in situ* soil moisture sensors with its web based dashboard displaying data such as soil moisture tension and providing irrigation advice. In developing the LRC trial sites, the project team was made aware of the installation of Greenbrain by MEA in the viticultural blocks, and this information was therefore also integrated into the VitiVisor platform.

4.1.3.3 Irrigation System

Water management was a key research priority for Riverland growers as it is both a significant viticultural management lever and a major production expense; moreover, water in the Riverland is a tradeable asset through the national water markets. Therefore, it was necessary for the project to

engage with the two main water delivery providers: Central Irrigation Trust (CIT) and Renmark Irrigation Trust (RIT).

In the case of CIT, they outsourced the electronic monitoring and management of water delivery to Outpost Central. In negotiation with both these organisations and through the consent of selected growers (their customers), the project was able to connect an API to the dashboard displaying water delivery data for a vineyard operation. In the case of RIT, initial contact was made to undertake a discovery exercise to understand their framework of monitoring and managing water delivery through a telemetry system; however, the planned investigation did not proceed further due to the COVID travel restrictions at the time.

4.1.3.4 VitiCanopy

The collection of viticultural data was undertaken by the use of the VitiCanopy application over the project. This technology was selected as the operational data collection system within VitiVisor, as the VitiBox data acquisition system was only completed towards the end of the final growing season. VitiCanopy was selected because it mature technology, easy to use, highly portable and able to provide results available in real time.

Through negotiation with the VitiCanopy IP owners, the connection of an API to VitiVisor provided the perfect application for its use, demonstrating that a grower already using VitiCanopy could connect to the VitiVisor dashboard displaying key vine parameters including plant area index (PAI), canopy porosity and canopy cover.

4.1.3.5 Financial

Know your Numbers (KY#) is a web-based tool for managing and displaying a vineyard's fixed and variable expenses. The tool was conceived by Riverland Wine, with the aim of increasing the financial literacy of its growers. At the start of the VitiVisor project, Riverland Wine directed its inclusion in the VitiVisor project as the key platform for collecting and providing financial information to the VitiVisor system. KY# was created by the Adelaide software development company Katalyst. The VitiVisor economics team worked very closely with Katalyst and helped supply funds in order to:

- Upgrade the KY# system to a current version
- Connect and enable syncing of KY# with the Xero accounting software
- Connect and enable syncing of KY# with the VitiVisor dashboard

In addition, meetings were held with Bentleys Chartered Accountants, to gain a perspective on how this tool may be used by financial advisors.

4.2 Ground-Based Vision System Development

4.2.1 System overview

The ground-based vision system collects and analyses imagery from camera systems that can be mounted on any agricultural vehicle and can interpret these images through a set of sophisticated computer vision algorithms to generate insights for growers. Through the process of an iterative design cycle and benchmarking against alternatives in the marketplace, the following design principles were adopted:

- **Affordable**: start from basic elements including the use of an inexpensive camera module, parts, and the associated technologies. The hardware is light, small, low cost, and widely available.
- **Deployable**: can be mounted on any agricultural vehicle as a 'collect as you go' image capture device, to be applied either as an incidental capture system or for dedicated capture missions.
- **Flexible**: prototyped to demonstrate that VitiBox can be expanded to connect to more RGB cameras, thermal cameras, and other data collection devices.
- **User-friendly**: provides the use of touch panel GUI showing control buttons, live image preview window, status of the connected devices as well as activate/deactivate image capturing. Images are saved into a local USB flash disk and also uploaded to cloud automatically once having access to Wi-Fi.
- **Fast and accurate:** captures, time stamps, and geo-tags high-resolution RGB images every one second with location accuracy of 2-3 centimetres.
- **Robust:** works under all day-time lighting conditions (i.e. cloud cover and direct sunlight).
- **Insightful**: applies state-of-the-art computer vision algorithms to provide information on key measures relevant to grower decisions, including pruning weight, bud number, inflorescence number, shoot number, bunch number and bunch mass.
- **Open source:** hardware schematic diagrams, CAD files, software, and instructions will be made available online. The device can be replicated by anyone with basic computer skills.

The system was built as an end-to-end solution that involves image capture, processing, and presentation back to the grower. The workflow is illustrated in Figure 14 and shows the setup, captured and processed images, and the extracted information display on dashboard. Note that the sample image shown in Figure 14 was captured by an earlier version of VitiBox (VitiBox 2.1).

The ground-based vision system is designed as two separate sub-systems: the image collection device (VitiBox) and the post-processing sub-system, and each is discussed further below.

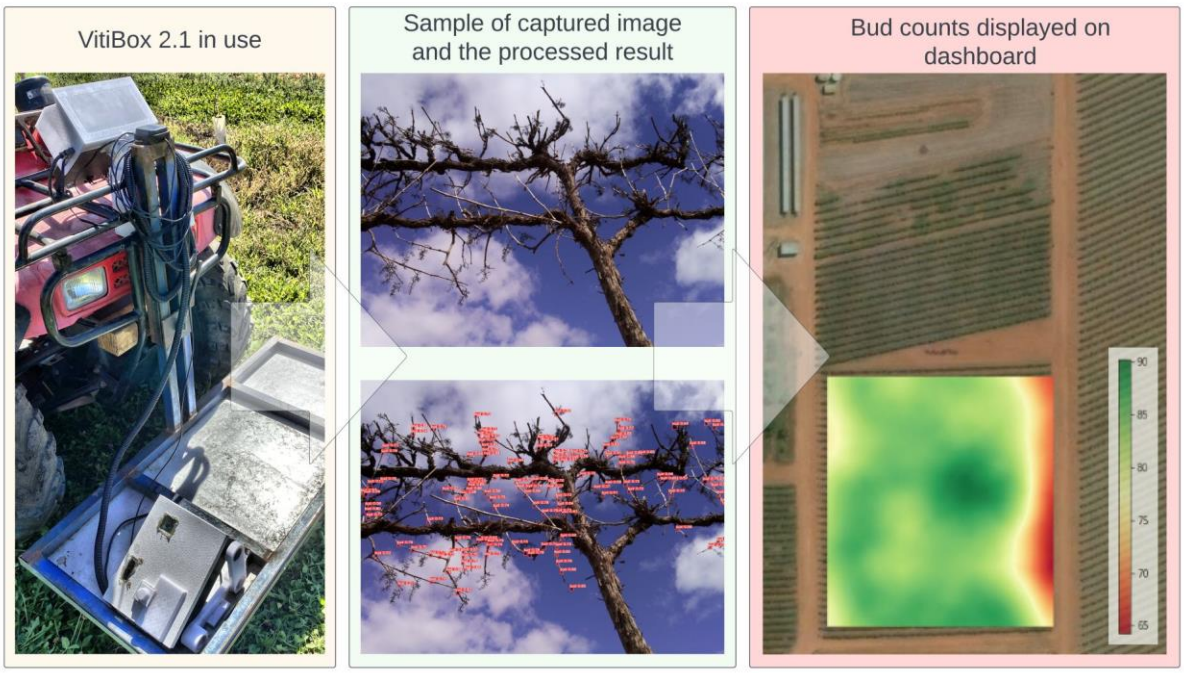


Figure 14: Ground-based system workflow

4.2.2 VitiBox: A flexible image capture system

VitiBox is designed as a low-cost and flexible ground-based image capturing system. Compared to off-the-shelf image capture solutions (e.g. a GoPro camera or equivalent), VitiBox provides significant capacity to tailor the image capture process (e.g. frame rate), interactive user feedback for image quality through a user display, improved geopositioning (several centimetres using RTK GPS, rather than several metres for standard GPS solutions), seamless transfer to the cloud, and flexibility to easily add additional camera systems (e.g. thermal). This not only provides a more complete user experience, but also leads directly to higher quality imagery with anticipated benefits in terms of predictive skill for the machine learning algorithms.

Compared to existing integrated hardware systems, VitiBox has focused on providing a flexible, low-cost (both capital and operating) and open-source solution that can be tailored to a broad range of image capture missions. Figure 15 illustrates the system design architecture.

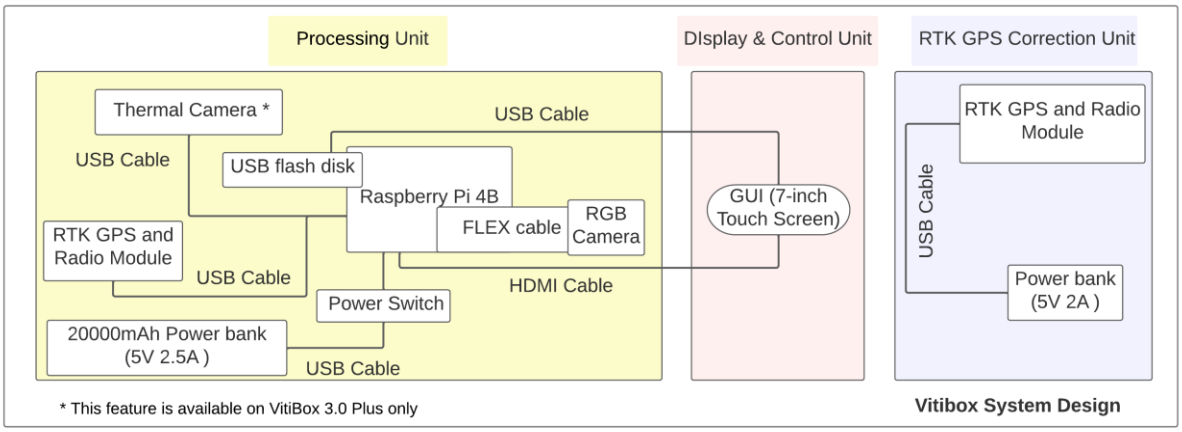


Figure 15: VitiBox 3 system design

Throughout the project, the VitiBox system has gone through a series of design cycles, summarised briefly below:

- **VitiBox 1.0:** This version has an RGB camera and a long-wave infrared (LWIR) thermal camera for fixed location image capturing. The surface dripping irrigation system was tested at Waite campus only. RGB and thermal images are taken every three minutes. The resolutions for the RGB and thermal images are 640 x 480 and 160 x 120, respectively.
- **VitiBox 2.x:** Based on feedback from growers, a ground-based vision system that captures high-resolution, geo-located canopy images at a high frequency was designed and developed. Four versions were evolved from VitiBox 2.0 to VitiBox 2.3. The image resolution was eventually increased to 3280 x 2464. In addition, real-time timestamp in milliseconds and RTK GPS geo-tag in centimetres were added as part of the metadata. The VitiBox 2.x series were tested at the Waite Campus of the University of Adelaide and Loxton Research Centre at Loxton.
- **VitiBox 3.0.** In this version, a post-processing function is created to add rich metadata to images which runs on the Windows platform. VitiBox 3.0 evolves to have some variations:
    - **VitiBox 3.0 Cloud**: uploads all images and corresponding timestamps, geo-tags automatically to the Cloud when it connects to the Internet. The post-processing sub-system starts the processing once image is received and the extracted information will be display on the dashboard.
    - **VitiBox 3.0 USB**: allows users to view the images on a local computer and upload images to the Cloud selectively.

- **VitiBox 3.0 Plus**: can connect to thermal camera to capture both RGB and thermal images at the same time. The image resolution is limited to 2560 x 1920 to pair with the thermal image resolution.

A summary of the VitiBox 3.0 technical specifications are listed in Table 8:

Table 8: Technical specifications of VitiBox 3.

Item	Details	Comments
Processor	Raspberry Pi 4B 8GB	Small size, low power demand, and low cost
RGB Camera	Raspberry Pi V2 3280 × 2464 pixels @ 1Hz*	Light, High resolution, and low cost
Thermal Camera**	FLIR Lepton 3.5 160 × 120 pixel @ 1Hz	Light, radiometric temperature accuracy ∓ 5℃, low cost, and good support
Geolocation	ZED-F9P GPS-RTK2 Boards	Widely available, improved accuracy, and good support
Camera Mounting	3D printed enclosure with 1 DOF adjustment on camera shooting angle	Light, reliable, easy to use

* RGB image resolution is set as 2560 x 1920 for VitiBox 3.0 Plus
** This feature is available on VitiBox 3.0 Plus only.

Throughout the design cycles of the VitiBox series, many trials have been taking place involving eight vineyard patches at the Waite Campus of the University of Adelaide and at Loxton Research Centre. More than 50 separate trails and testings in total were conducted by attaching Vitibox to agriculture vehicles travelling over 40 km of vineyard rows. These trials and testings finally led to over 25,000 images captured. Valuable feedback and comments through trials and testings were received from farmers and users for the improvements and evolutions of the VitiBox series. The captured images were proven to have very good image quality and processed by the AI based computer vision algorithms to produce satisfactory results as the example shown in Figure 14.

Importantly, VitiBox has been made available open-source with associated instructions that would allow someone with reasonable computer literacy to reproduce the technology (Box 4).

Box 4: Build your own VitiBox

VitiBox has been designed as an open-source solution that is available for anyone to reproduce. The details are provided on Hackster (see www.vitivisor.com.au for links) and includes a complete inventory of hardware components, software services, 3D printing parts, associated code and schematics, and detailed instructions regarding reproduction copy.

It is envisaged that assembly of VitiBox could be achieved by someone with reasonable computer literacy, and with development of associated computer skills (Linux and 3D printing), the assembly could be taught over several hours to senior high school students and/or TAFE.

4.2.3 Algorithm development and post-processing subsystem

The ground-based vision system's post-processing subsystem takes the images from the image collection device (VitiBox), and converts this into important vegetation structural variables for vines, including leaf area index (LAI), pruning weight, bunch count, bud number, shoot counts and counts of the inflorescences. The conversion is conducted using powerful machine learning algorithms that

can produce machine learning models from a relatively small number of samples with high accuracy. These predictions are then displayed on the dashboard (Section 4.5) so that the growers can view the distribution of these variables over the whole vineyard.

The machine learning algorithms are more advanced than many established systems in three ways:

- estimates and detects not only canopy development, but also many other key indices such as pruning weight, bud count, and inflorescence and shoot number;
- uses cutting edge deep learning techniques thus gaining more robustness to the variations of the images captured;
- utilises triplet loss function and domain knowledge to turn insufficient small amount of annotated data virtually to medium to large amount of data, thus significantly boosting the accuracy.

Further detail on the development, training and testing of these algorithms is provided in the sections below.

4.2.3.1 Performance Evaluation Protocol

To evaluate the performance of the collected labelled data, data have been randomly separated into three disjoint sets:

- A **training set**, to train our models using differentiable loss functions;
- A **validation set**, validation set are used to tune the hyper-parameters of our models; and
- A **testing set**, which was kept away from the model and only used at the final stage to give an indication of model performance.

These sets were usually divided by the ratio of 0.8:0.1:0.1, respectively, in-line with disciplinary conventions. The specific error metric adopted is the mean absolute error, allowing error metrics to be presented in the same unit as the indicator of interest. The error rate is presented on an individual image basis; this is likely to over-estimate the error for decision-relevant measures such as patch-level metrics and spatial variation, due to the potential for errors to cancel out on aggregation. Moreover, the dataset, the machine learning algorithms as well as the performance evaluation protocol have been made publicly available to allow for independent scrutiny. The therefore supports transparency and is consistent with industry best-practice.

We now turn to the specific algorithm implementation for some of the key indices used in this project.

4.2.3.2 LAI Prediction

The leaf/plant area index (LAI, PAI) is a critical vegetation structural variable and is essential in the feedback of vegetation to the climate system. An accurate LAI monitoring procedure can provide valuable information for a vineyard. While traditional direct method of LAI measurement requires destructive sampling procedure as it requires harvesting vegetation leaves, indirect optical methods can provide LAI estimation from images of vines without destructive sampling. However, not only the leaves are contained in the images of vines, but also shoots, cordons, bunches, and dead leaves. Thus, it is necessary to provide a robust algorithm that can efficiently reduce the influence of unwanted patterns in the images.

Recently, the deep neural networks have been shown to have great advantage over traditional methods in various computer vision tasks including segmentation and detection. Thus, we designed a data-driven LAI prediction method using deep neural networks (LeCun et al, 2015). One main

drawback for deep neural networks is that it often requires a large amount of labelled data to train a deep neural network to attain good performance. Particularly in our task, deep neural networks require accurate image/LAI value pairs, which is impractical as it is consumes significant resources.

To solve this problem, we designed an algorithm that can make use of unlabelled data. Our basic idea is to dig the rank information between unlabelled data. Due to the natural order, the LAI of vines in October must be higher than the one in September (See Figure 16 for example). Using this information, even though we do not have the accurate LAI for images, we will be able to obtain a partial order between the LAI of unlabelled images. Such partial order can be used to construct a rank-based loss (Zheng et al, 2018) for LAI prediction such that the model can benefit from a huge amount of unlabelled data. Using this idea, we are able to obtain a machine learning system for LAI prediction as shown in Figure 17. As our system can make use of huge amount of unlabelled data (a subset of the collected 170,867 images), the performance of our model can be improved continuously as additional raw image data are collected over time.



Figure 16 Example images captures for LAI prediction. ***Left:*** *image captured in September 2020,* ***Right:*** *Image captured in October 2020. Obviously, the LAI in September should be less than the LAI in October.*

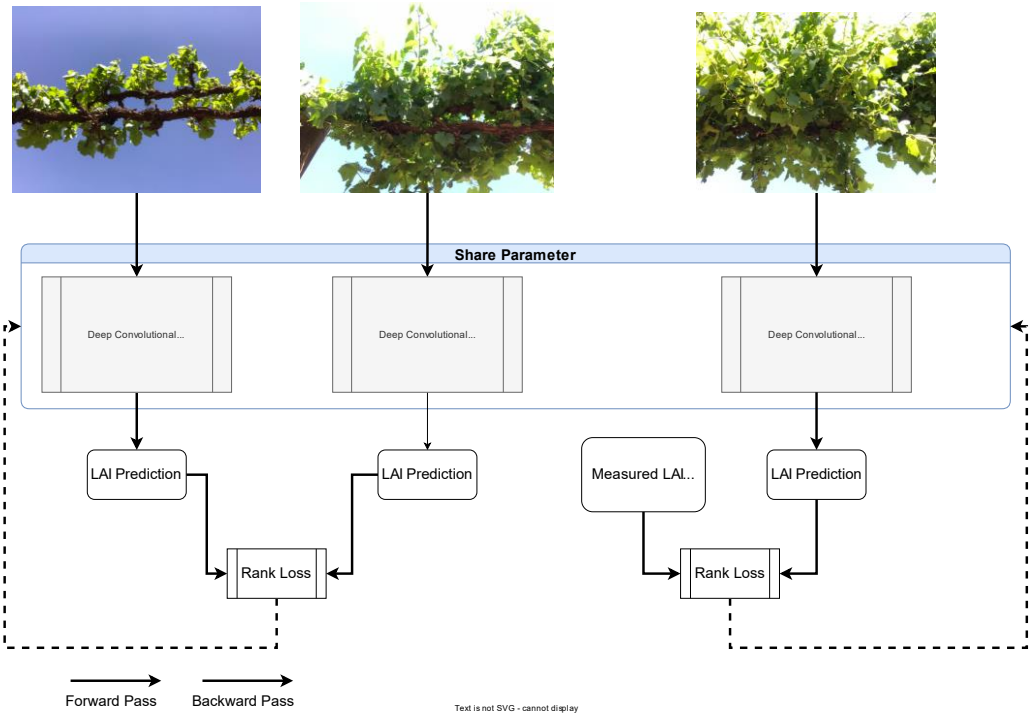


Figure 17: High-level structure of our LAI prediction module training procedure. In our model training procedure, we not only use a small amount of labelled data, where the ground truth is obtained from destructive sampling, but also use a large amount of unlabelled image. For the unlabelled data, as we are able to obtain the rank relation between LAI of images from different times, we can use this information to obtain a rank-based loss.

4.2.3.3 Pruning Weight Prediction

From the machine learning's perspective, the pruning weight problem is very similar to the LAI prediction task. The input of the machine learning algorithm is an image, and the output is an estimation of pruning weight. Again, we are facing the problem of low number of labelled images, given the significant time and expense required for labelling as discussed in Section 4.1.1. However, unlike the LAI prediction problem, we are not able to obtain the partial order between unlabelled data. Thus, we use a different way to enlarge the dataset. In our model, instead of directly predicting the pruning weight from one image, we predict the ratio of pruning weight of an image against a template image. By using this approach, assume that we have n labelled images, we will be able to obtain $n^2 - n$ images pairs where one image in the pair serves as template. Using this idea, the small data set are extended to a large dataset and thus the deep learning-based approaches can attain relatively high accuracy.

4.2.3.4 Bunch Counting

The bunch counting problem can be viewed as a regression problem. However, the images of bunch counting are very similar to the ones for LAI prediction, thus we finetune the model of bunch counting from the LAI prediction model, rather than training the model scratch. Using this technique, the knowledge learned in LAI prediction can be transferred to the bunch counting model to attain better performance.

4.2.3.5 Buds/inflorescence/shoot-tips Detection and Counting

The buds/inflorescence/shoot-tips detections and counting problem can be viewed as an object detection problem from the machine learning's perspective (Zhao et al, 2018). However, in typical objective detection problems the size of object is relatively large, while in our problem the size of object is usually relatively small compared to the whole image size (example shown in Figure 18). In this case, directly using the large image to train an objective detection model would results in very low accuracy (Tong et al, 2020). As a result, in our implementation for every image we randomly sample 200 patches with resolution 640x480 and in such patches the buds/Inflorescence/shoot-tips becomes larger such that richer features are obtained to achieve a higher accuracy. In our implementation, we use the yolo[3] network to train a detector.

[3] https://github.com/ultralytics/yolov5

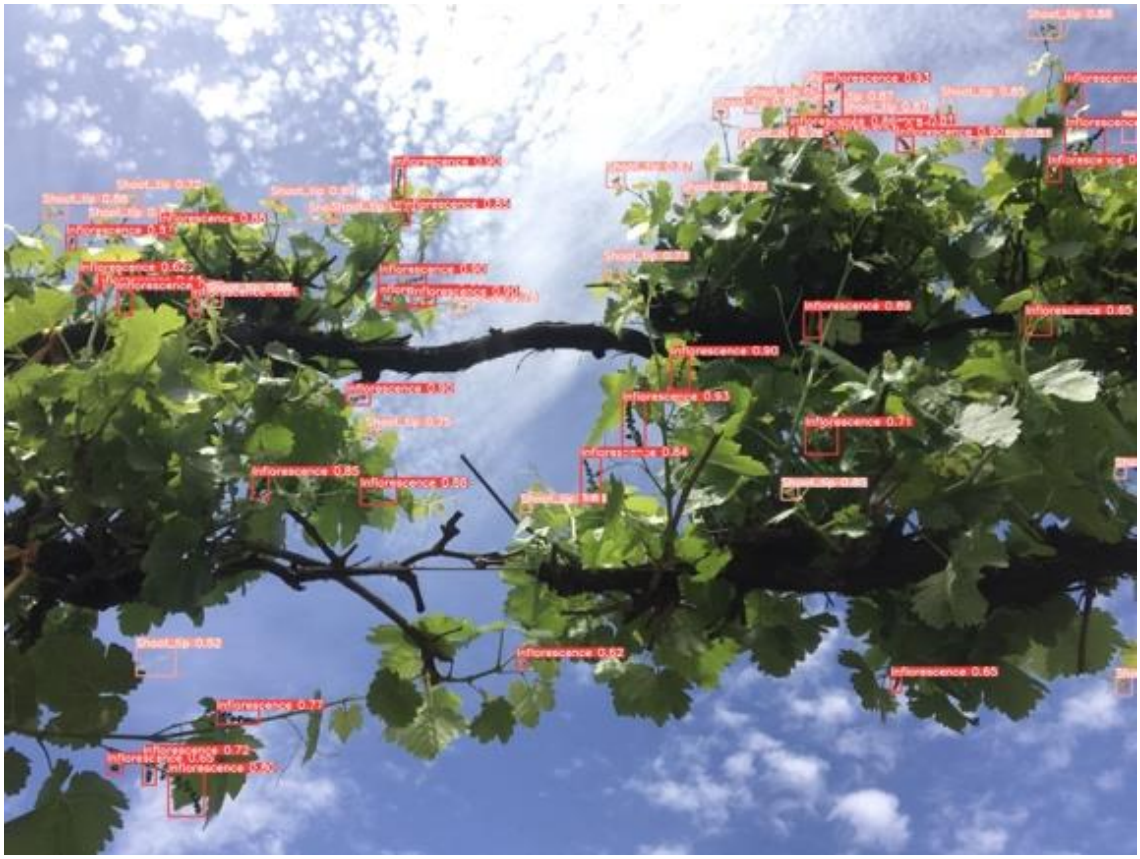


Figure 18: Typical inflorescence and shoot-tips detection task

4.3 VITICULTURE MODELLING

The preceding sections all described the collation of various farm-related data streams—some commonly collected on operational vineyards and some still relatively novel—as part of the VitiVisor digital platform. Yet for data to become of value to growers and their advisors, it must be converted into a form that supports decision making. This section describes techniques to achieve this purpose.

> "Data is of limited value by itself. It needs to be combined with the capacity to predict the future in order to inform decision making."

4.3.1 Integration of established vine growth model

A key principle underpinning the project is to 'connect' decisions and outcomes (Figure 19). The relationships between environmental factors (e.g. climate, soil), operational decisions (e.g., irrigation) and vineyard outcomes (e.g. yield), are complex and dependent on many factors. To this end, we use a process-based viticulture model as this class of model is designed to reflect contemporary scientific understanding of the key biophysical processes governing grapevine growth and development in mathematical terms, thereby enabling prediction of future vineyard status and outcomes under different management decisions and environmental factors.

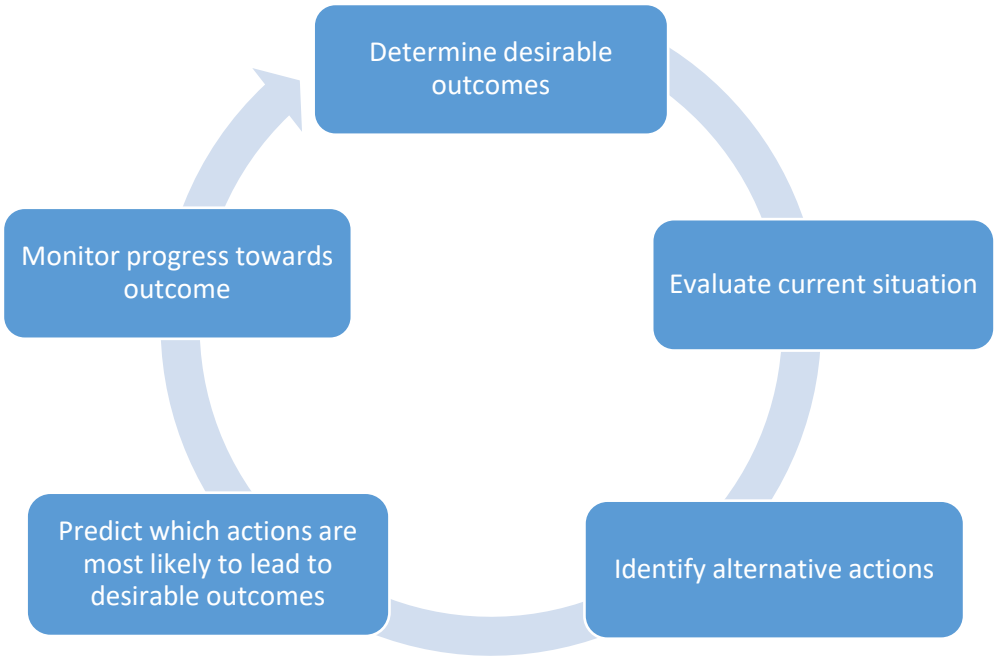


Figure 19: The approach to combining models and data in this project support grower decision making through improved situational awareness, predictive capabilities and the capacity to monitor progress towards goals

To determine which of the currently available viticulture models are best suited to provide decision support in the Riverland, we conducted a global review (Knowling et al., 2021). This review highlighted numerous potential advantages of the model VineLOGIC—developed over the period since the early 1990s—as the basis of supporting VitiVisor. In particular, VineLOGIC was identified as being well suited to the Riverland context where irrigation application and yield represent key performance indicators.

Table 9: Summary of the capability and suitability of process-based models to support decisions across different vineyard management contexts. Italics indicate that capability has been reported, but not yet demonstrated. Table adapted from Knowling et al (2021)

	Outcomes	
Decision Levers	**Yield**	**Berry composition attributes**
Water management	APSIM STICS SWAP **VineLOGIC** VitiSim	*STICS* ***VineLOGIC***
Canopy management	APSIM IVINE STICS **VineLOGIC** VitiSim	*IVINE* *STICS* ***VineLOGIC***
Nutrient management	*APSIM* *STICS*	*STICS*

Our choice of VineLOGIC also supports another key principle underpinning the project—to build on existing Australian research and further unlock its value. VineLOGIC's long development history in Australia reflects a significant investment from CSIRO, AWRI and Wine Australia. This project was able to revitalise the IP from this partnership by focusing on four key pillars:

1) Negotiating with CSIRO and AWRI to make the license open source for both non-commercial and commercial uses. Prior to this, restrictions were placed on access to the software, ultimately meaning that at the beginning of the project, it was only used for teaching purposes for a single course at University of Adelaide.
2) Modernising elements of the code, including:
    a) streamlining input and output handling (e.g. widely used JSON and CSV file structures);
    b) translating a large portion of the code base that the user interacts with to the popular object-oriented R programming language;
    c) shifting variables that were previously 'hard-coded' to the input-output interface to allow for more control of key variables; and
    d) supporting additional decision variables that enable greater flexibility around exploration of canopy management practices (e.g. within-season trimming).
3) Improved documentation of VineLOGIC through contributions to the development of demonstration tool VineLOGIC View and through preparation of a journal article that serves as an overview of VineLOGIC capability, with a particular focus on its unique strengths relative to other models.
4) Broadening the application of the model by making VineLOGIC a real-time system that ingests real-time data feeds and is capable of supporting operational decisions.

Implementation of this last pillar (pillar 4) represents the most significant modelling innovation in this project and is discussed further below.

4.3.2 Building a vineyard 'digital twin'

A key principle of VitiVisor is the integration of diverse data and knowledge streams 'in one place'. Simply collating these diverse streams is already an important project outcome, as indicated through early discussions with growers. However, to fully leverage the benefits of data it is necessary to combine what is measured in the field with model predictions both to 'fill in the blanks' (in other words, provide estimates of key parameters in the vineyard for which data is not available), and to provide future predictions.

'Data assimilation' algorithms are the primary tool for combining diverse data streams (e.g. soil moisture, or canopy measures such as leaf area index) with a viticultural model (e.g. VineLOGIC) to produce a unified digital environment ('digital twin'). Key benefits of this approach include:

- **Situational awareness.** By blending observations with a numerical model of vine growth, it becomes possible to estimate variables that are difficult to observe (e.g. carbon stores), or interpolate information in space (e.g. soil moisture across a vineyard or at different soil depths, even if the actual soil moisture data is limited to a single probe at a single depth) and time (e.g. between physical measurements).
- **Causal reasoning**. The modelling environment can provide information about why certain outcomes occur within the vineyard.
- **Predictions.** The physical vine development processes represented by the model enable forward predictions of within-season measures (e.g. canopy measures) and end-of-season outcomes (e.g. yield), updated on a day-to-day basis as new data is collected.

- **Interventions.** Often, the primary purpose of having a digital twin is to enable users to change the digital version of the vineyard to see what might happen, before making those changes in the real world. This allows for different options to be tested before making a decision on the preferred approach.

"Given the importance of water, we focused on the relationship between watering strategies and end-of-season yield."

To achieve these outcomes, our DA system optimally blends diverse data streams (e.g. approximately weekly LAI data from VitiCanopy, together with daily weather station data) and our understanding of grapevine growth and development processes (i.e. expressed mathematically in VineLOGIC). This involves continually updating VineLOGIC on a daily basis as new data became available. In other words, VineLOGIC 'learns' from the data by adjusting its internal state variables, leading to improved forecasts as more and more data is incorporated. For example, as LAI data are assimilated into the model, forecasts of key end-of-season outcomes such as yield will be updated, depending how the observed rate of canopy development compares to that predicted in the model and the relationship between canopy development and yield.

Our DA system is a first for the viticulture domain and is innovative in several ways. First, our system predicts to the end-of-season every day, so that outcomes such as yield can be tracked continually as weather or observation data are updated. Second, we update both dynamic and static aspects of the model each day, enhancing the flexibility and reliability of our system. These innovations are the subject of Knowling et al. (under review).

We implemented DA using a recently developed software suite PESTPP-DA (Alzraiee et al., 2022). PESTPP-DA is unique in that it is data- and model-agnostic and can scale to address complex problems with many dimensions. PESTPP-DA is a publicly available and open-source tool available at: https://github.com/usgs/pestpp.

4.3.3 Supporting grower decisions

As stated above, a key focus of the project was to explore and characterise relationships between key decisions and outcomes. We focused on the relationship between water application and yield, since this relationship is a strong driver of Riverland growers' gross margin, including the role of different watering strategies at different stages of phenological development.

To better understand how grapevines respond to water application (often encapsulated within 'water production' functions), we used VineLOGIC to predict end-of-season yield under many different irrigation scenarios (e.g. changes in irrigation rates during different growth stages). We also used VineLOGIC to assess the impact of different genetic and environmental (climate and soil) factors on the relationship between water applied and yield. Both of these relationships are critical for effective irrigation decision making.

As highlighted in our global review (Knowling et al., 2021), VineLOGIC has a relatively detailed representation of soil water dynamics and the influence of water deficit on grapevine growth including berry development—both key to its selection for VitiVisor. VineLOGIC also supports different means to represent irrigation—through direct daily time series input or through daily simulation using soil moisture deficit thresholds. Figure 20 shows how irrigation decisions influence yield. This figure highlights the process-based nature of VineLOGIC—its outputs are a result of a network of causal relationships that represent our understanding of grapevine growth and

development processes. It is for these reasons that when combining VineLOGIC with data feeds from the vineyard, it represents a 'digital twin'.

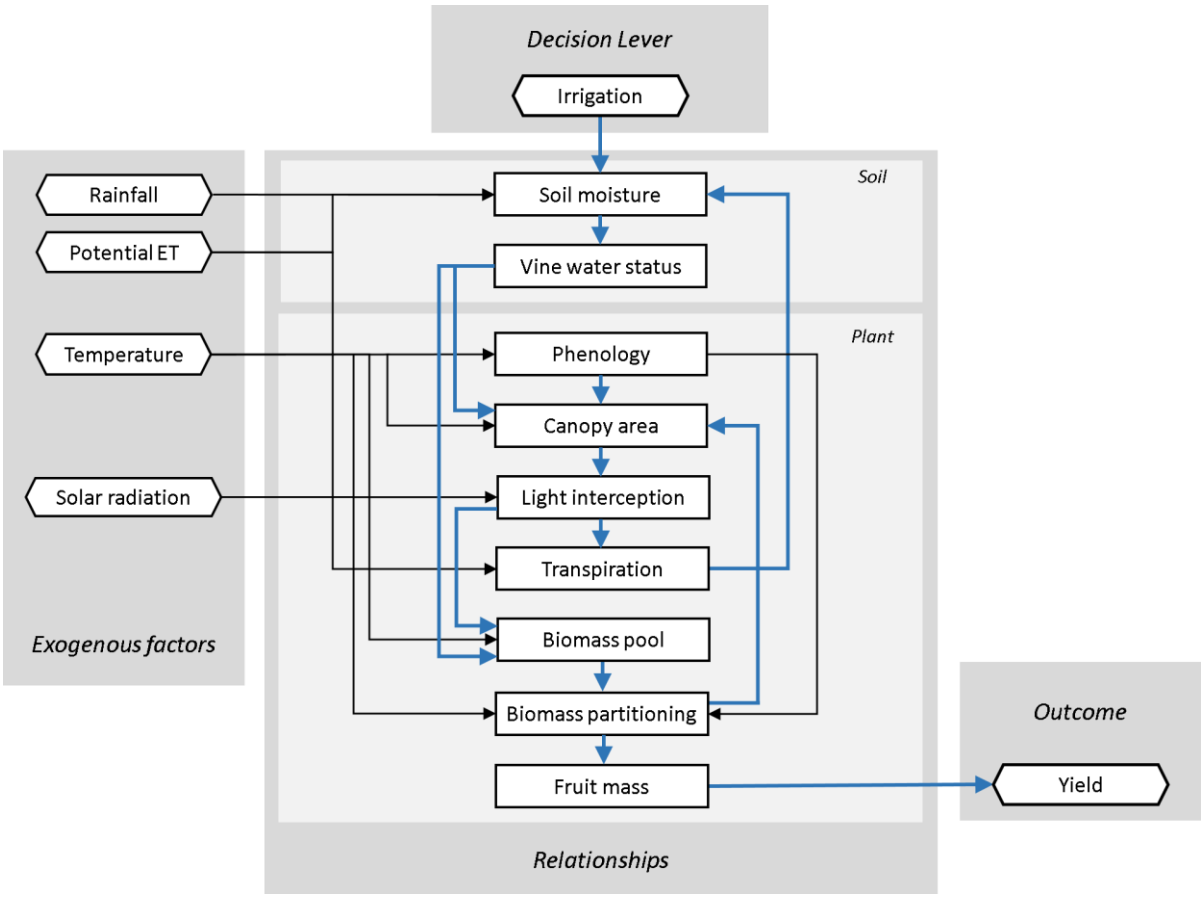


Figure 20: *Schematic of how the effects of water deficit on grapevine yield are represented in VineLOGIC (blue arrows). Water deficit stress indices associated with vine water status serves to reduce the biomass carbon assimilate pool for growth both directly through reducing photosynthetic activity, and indirectly through reducing the canopy area for light interception (From Knowling et al. [under review])*

Box 5: Research contributions on selection of viticulture modelling platform

A review of process-based models to support grower decisions in the viticulture section is available in:

> *Knowling et al, 2021, Bridging the gap between data and decisions: A review of process-based models for viticulture, Agricultural Systems, 193, 13 pages.*

Included in this review is articulation of the advantages and disadvantages of alternative modelling platforms, which formed the basis for selecting VineLOGIC for this project.

Further details on the VineLOGIC model is provided in:

> *Walker et al, 2020. VineLOGIC View V1. CSIRO Service Collection.*

VineLOGIC's capability to simulate the implications of different irrigation strategies on crop yield is described in:

> *Knowling et al, under review at Agricultural Water Management, Generalised water production relations through process-based modelling: A viticulture example.*

This manuscript also discusses the use of process-based models more generally for characterising complex, cause-and-effect relationships. As described in the manuscript: "… the relationship between applied water and end-of-season crop yield ('water production relations') … is often partial and lacks generalisability, owing to its many determining factors, especially for woody perennial crops such as grapevines. Process-based models are a way in which to represent these relationships in a manner that is both generalisable and interpretable. Here we conduct numerical experiments using a process-based crop model to evaluate water production relations for grapevines, and investigate how these relations are influenced by genetic and environmental factors and irrigation timing decisions."

4.4 Economic Assessment

Throughout the project co-design process (Section 2.1), growers emphasised the need for VitiVisor to focus directly on issues associated with costs of production. To this end, a rigorous economic analysis was conducted to address two key questions:

- To what extent can financial and economic data be used within the VitiVisor platform to support growers to manage costs of production and make decisions?
- To what extent can the VitiVisor platform provide a reasonable return on investment for Riverland growers if adopted?

A detailed stand-alone report (Wheeler *et al,* 2022) has been produced to address these questions, and a brief synopsis of the methodology is provided below.

4.4.1 Case studies and benchmarking

The aim of the economic work conducted was to provide a comprehensive overview of the economics—and returns—of growing grapes in the Riverland. In particular we focussed on benchmarking research, driven largely by the recognition that comparison of one organisation against other organisations is an integral part business performance analysis and improvement processes (e.g. Isoraite, 2004).

A thorough literature review was conducted to understand what historical benchmarking activities had been conducted, and to develop guidance for the key financial measures. Following this literature review, benchmarking was undertaken on three main levels:

1. using Australian tax office data on all grape growing enterprises in the Riverland;
2. using University of Adelaide unique irrigation farm survey data; and
3. undertaking new case studies of 25 Riverland typical grape growers.

The detailed data from the case study then allowed for the creation of a benchmarking tab in VitiVisor, supporting growers to evaluate their key production costs relative to other growers in the Riverland, as well as provide objectives measures of industry best practice.

The long-term vision is that widespread adoption of VitiVisor within and beyond the Riverland would enable the platform to continually add to its collection of data at a more finite levels (e.g., by variety or location) over time, and thus provide consistently updated information on costs and sources of revenue. The case studies and benchmarking analysis described here provides a solid foundation to achieve this broader vision, using a diverse set of information including a thorough literature review of all benchmarking relevant issues for the Riverland grape growing industry, as well as the collection of primary and secondary data.

The primary dataset to inform the benchmarking analysis originated from detailed interviews with 25 grape growers in the Riverland (University of Adelaide ethics approval number: H-2019-086). This was required as there was very little recent information of detailed up-to-date variable and fixed costs of grape growing businesses in this region. The grape growers were selected by contacts, information and criteria provided by Riverland Wine and the CCW Co-operative. Interviews with the research team and 25 growers were undertaken on-farm between July 2020 and April 2021, and subsequently updated in 2022, with questions aimed at collecting detailed information about:

- yield and varieties;
- growing practices;
- water use and trading strategies;
- farm income; and
- costs.

Interview and farm financial records were combined to create an in-depth assessment of growers' income and cost positions, with detailed findings especially for the financial year of 2019-2020. In addition, information on previous tax years was collected wherever it was available, along with information on the 2020/21 financial year once the tax information become available.

As discussed, this primary dataset was supplemented by a range of secondary data sources that were collected throughout the project. This additional information was utilised to understand the historical picture of grape growing and profitability in the Riverland over time, and exit and agglomeration of businesses. A summary of key data sources is provided in below.

Table 10: Secondary data sources to support economic analysis

Data source	Description	Role in analysis
Australian Bureau of Statistics (ABS) DataLab (2020) microdata	Australian Tax Office (ATO) data on individual business, with annual data available from 2000-01 to 2017-18 as of 2022. In order to identify grape growing businesses, the ANZSIC code 0131 for Grape Growing was used (ABS 2014)). According to Wine Australia's Riverland GI region and Australia	Enabled an analysis of ATO data including information on farm assets (cash, machinery, water entitlements, farmland, buildings and bearing plants such as vines)

	Post's postcode map, the Riverland GI region encompasses the postcodes 5311, 5320, 5321, 5322, 5330, 5331, 5332, 5333, 5340, 5341, 5342, 5343, 5344, 5345, 5346, 5357, and 5417. In total, 929 businesses were identified in this dataset in 2001/02, decreasing to 628 in 2017/18 (this compares to 1213 and 963 businesses, respectively registered through the Wine Australia dataset).	and liabilities (farm debt, or bills yet to be paid).
SA Auditor General (2020) data	This dataset comprises statutory property valuations of Riverland lots classified as 'Vines'; 'Vines – irrigated'; 'Vines and others'; 'Vines and others – irrigated'; and 'Vines and Stock – irrigated' from 1985/85 – 2017/18. The dataset comprises lot sizes, site values and capital values, both nominal and in 2017/18 prices.	Allowed comparison between rates data and Wine Australia data
Wine Australia (1999-2020) data	Data on volume of grapes by variety produced across Australia's wine growing regions from 2015/16 – 2020/21. Data on South Australian wine regions go back as early as 1999/00 – 2020/21 and also contains farm size and farm numbers for each year, separating growers into five different sizes: <10 ha; 10-24 ha; 25-49 ha; 50-99 ha; and >100 ha.	Allowed comparison between other datasets and detailed price information
University of Adelaide irrigation farm surveys from 2010-11 and 2015-16	This dataset provided a cross-sectional dataset of 593 southern Murray-Darling Basin horticultural irrigators, of which 258 were SA Riverland grape farms.	Allowed regression analysis to be applied for identification of the factors associated with greater farm net returns (include net returns per hectare) for Riverland grape farms
ABS agricultural census information	The ABS agricultural census dataset for 2015-16 was used to estimate number of grape growers in the Riverland.	Allowed comparison between other datasets
H20X water market current data	An agreement was sought with H_20X to provide current and up-to-date water market data (prices and volumes and water allocations to various entitlements) into the dashboard	Allows growers to see current water market prices, and to be able to use them for current decision-making very easily

All analysis was conducted using standard economic and accounting terminology, functions and definitions.

4.4.2 Net economic benefit assessment

In addition to benchmarking analysis as part of the VitiVisor platform to enable growers to evaluate and manage their costs of production, a separate analysis was conducted to evaluate the extent to which the VitiVisor platform—if adopted by growers—will represent a reasonable return on investment to growers.

This analysis was undertaken using a financial benefit cost analysis (BCA). BCA is a widely used tool for evaluating alternative courses of action that enables costs and benefits to be monetised to enable direct comparison across the full suite of options under evaluation, and standard assumptions were applied.

The data to inform the BCA was based on ABS/AWRI data, interview findings from farm case studies conducted in 2019-20 and consultations with the project focus group, and a number of assumptions

were made (see Wheeler et al, 2022, for more details). Table 11 provides an overview of the main assumptions and figures used. As a broad summary:

1. Benefits and costs were assumed to accrue over a five-year lifespan of VitiVisor dashboard technology.
2. The base case was the continuation of the standard financial results, without adopting the VitiVisor dashboard.
3. Average per-hectare benefits and costs were calculated using data on average costs incurred by the average viticultural farm in financial year 2019-20.
4. Members of the VitiVisor team provided estimates of the potential costs of the Dashboard and software. Other information on other software packages was also collected, and we allowed for greater potential fixed adoption costs.
5. Savings of 10% on electricity, fuel and, labour costs were assumed as a result of better-informed and timely decision-making. Some case studies identified improvements of up to 50% (e.g. in electricity), however, we used a conservative estimate of 10% in this analysis.
6. It was assumed that savings in chemical costs would be offset by additional costs of harvesting and freight associated with increased yields. Savings in water costs were also entirely possible, but we conservatively only valued the average additional cost of extra watering to maximise yields.
7. Other potential benefits of adoption not quantified in this analysis, include reduced operator labour time on farms, increases in farm revenue from improvements in price received due to improvements in grape quality, increased soil carbon levels, improvements in vine nutrition, soil quality and water quality. In addition, no social benefits of industry knowing predicted yield, or optimising harvest times etc, were quantified.

Table 11: Parameters, parameter values and data sources used to quantify costs and benefits of adoption of Dashboard Technology

Parameter	Value
Period of analysis (years)	5
Discount rate	7%
Total grape area (ha)	57.2
Fixed expenses per hectare ($/ha)	$3,815
Costs of adoption of Vitivisor Dashboard Technology	
Initial cost required to develop Vineyard Manager's skills (hours valued at $50/hour)	$500
Annual cost of operating the dashboard estimated at 4 hours a week	$10,400
Initial cost of dashboard software packages (soil moisture sensors/micro weather station/flow meter/ground-based vision system)	$6,000 (up to $10,500)
Ongoing annual maintenance cost (soil moisture sensors/micro weather station/flow meter/ground-based vision system)	$6,000
Cost of increasing watering by 1 ML/ha ($/ML/ha)	$298
Benefits of adoption of Vitivisor Dashboard Technology	
Yield gains in tonnes from increasing watering by 1 ML/ha (t/ha)	3
Additional revenue from increasing watering by 1 ML/ha ($/ha)	$1,970
10% Electricity cost saving ($/ha)	$34
10% Fuel cost saving ($/ha)	$25
10% Labour cost saving ($/ha)	$38

A variety of sensitivity analyses was also conducted where key adoption benefit variables were varied for an: a) optimistic scenario where greater benefits were achieved; and b) pessimistic

scenario where benefits were less than those predicted. We also conducted a threshold analysis, to establish the point at which costs of technology adoption equalled benefits.

4.5 Software Development

4.5.1 Software architecture overview

Back-end and front-end software was developed for VitiVisor to connect each of the data streams described in the preceding section with a set of state-of-the-art models and algorithms, and a user interface that displays both the data and associated insights. The conceptual relationship between the key data streams is provided in Figure 21. The dashboard as part of the software platform is designed to demonstrate the core elements of project IP and, importantly, provide a vehicle for grower engagement and feedback elicitation.

As suggested by the logic of Figure 21, the project focused firstly on the water related management actions (canopy and irrigation) as per the initial direction and priority from the growers. Subsequent priorities were derived from multiple meetings with the grower advisory group, with progress managed and tracked using a Trello dashboard.

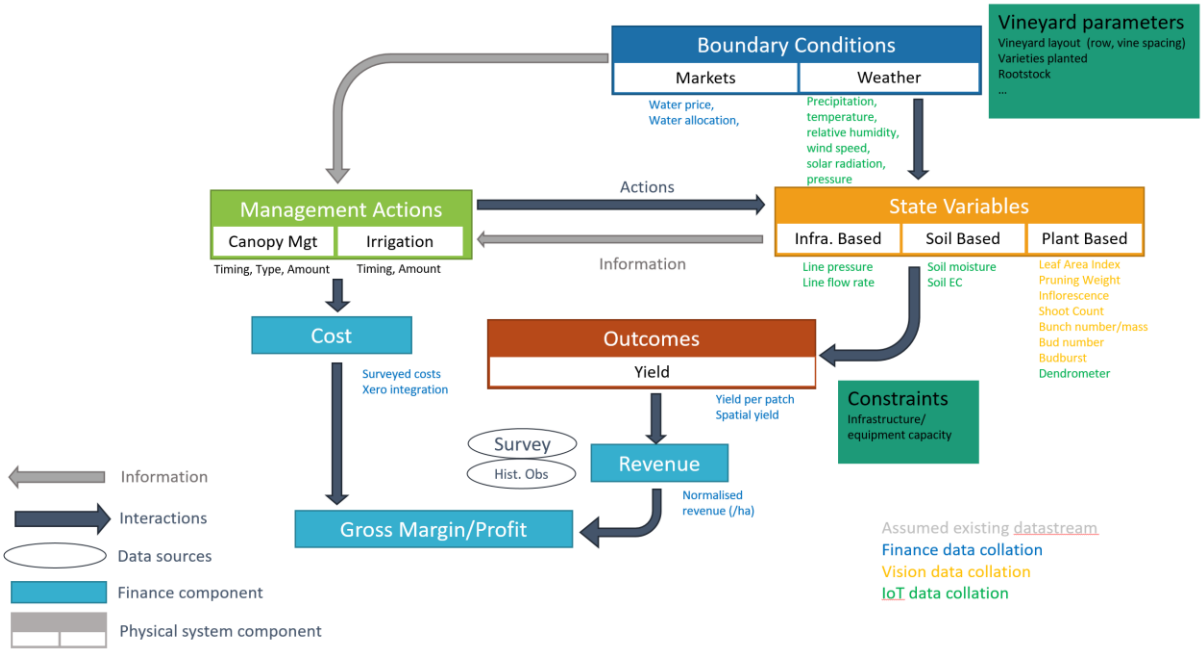


Figure 21: Relational schematic of key data flows associated with VitiVisor

The implementation of this logic is achieved through a software architecture as illustrated in Figure 22, and consists of the following components:

- **User interface (UI),** which is the visual representation of the Vitivisor web application. The system allows users to access the system via web browsers. The content of the system is shown in web pages, where users then interact via a primary and sub menu system to access relevant information about their vineyard, provided by the system.
- **Business logic layer**, which encodes the real-world business rules that determine how data can be created, stored, and changed. It is contrasted with the remainder of the software that is concerned with lower-level details of managing a database or displaying the user interface, system infrastructure, or generally connecting various parts of the program.

- **Database & file storage systems,** which is responsible for storing files into a file system and recording data into the database.
- **Connections with external data sources,** which are used to create and maintain connections to multiple external data sources using different protocols. The data sources range from finance, weather, aerial image and property through to IoT and ground-based vison systems.
- **Processors for VineLOGIC and the machine learning models**, which have been designed to operate as a standalone to the business logic layer either running periodically or triggered by the receipt of new data using the 'data & file' in the data storage component. This functionality stores the generated files and records back to the file system and database.

These components are elaborated upon in the sections below.

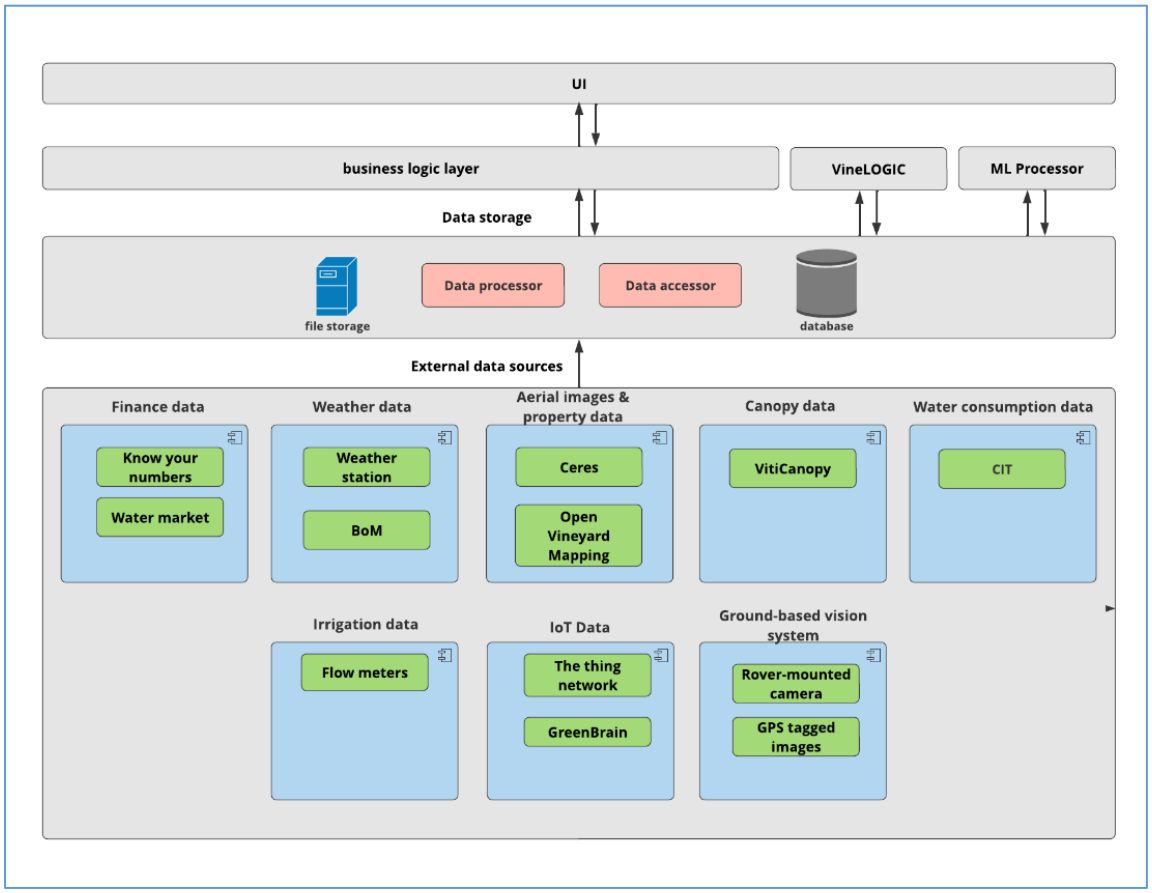


Figure 22: VitiVisor software architecture

4.5.2 Business logic layer

The business logic layer consists of many small components that define the business rules in the VitiVisor system. The composition of the business layer is detailed in Figure 23, with each subcomponent being responsible for related business functions. For example, the user account system defines users and provide APIs for user registration and login.

The business logic layer also connects to two external data collection systems, Know your Numbers (financial data) and VitiCanopy (viticultural data). Importantly, the finance component allows users to input and update their financial records, and undertake benchmarking with the regional average, using the built-in financial models.

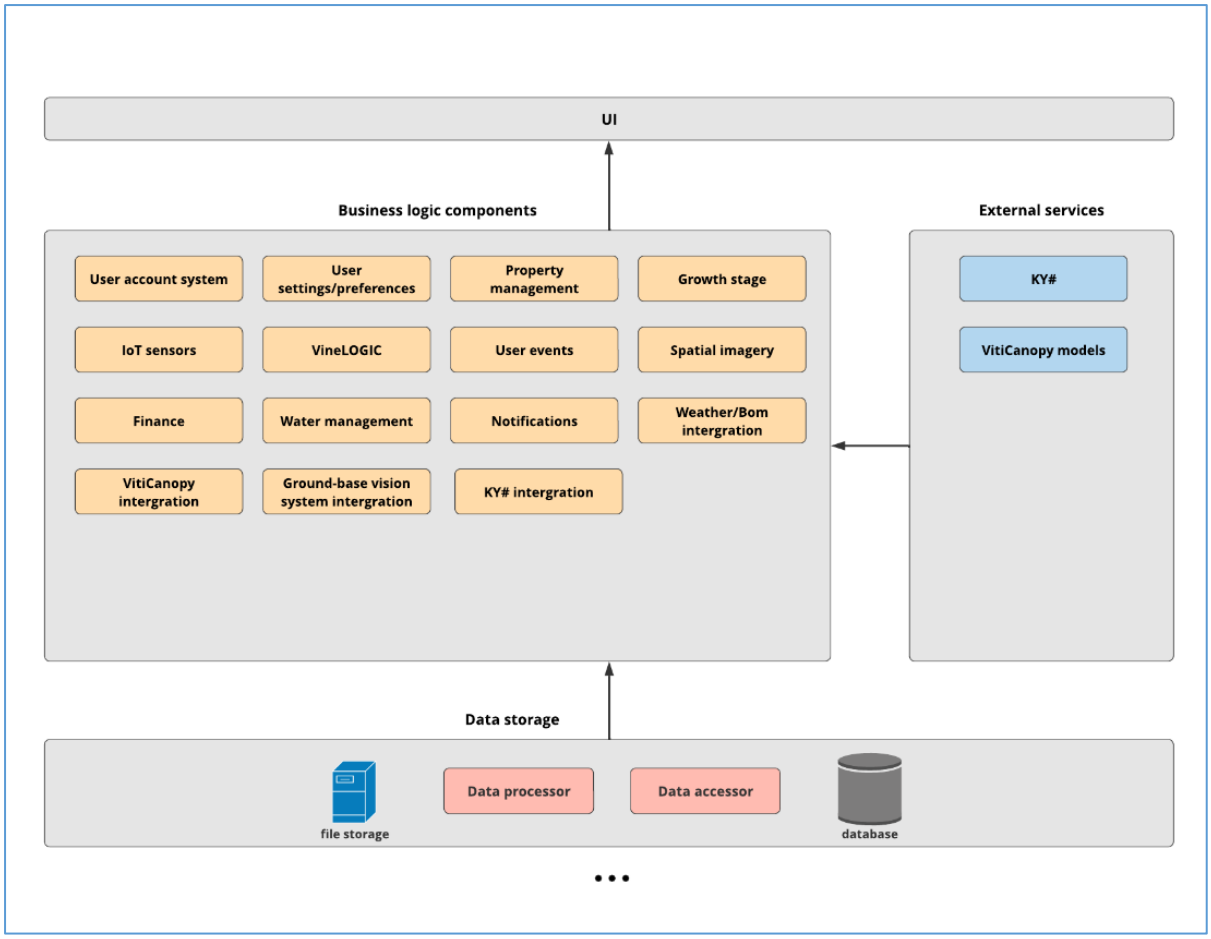


Figure 23: VitiVisor business logic layer

4.5.3 Connection with established data providers

In order to demonstrate the capacity to develop algorithms, scenario analysis tools and decision support capability 'on the top' of data, it was necessary to establish a foundation whereby the VitiVisor platform was able to connect to existing and widely available data streams.

The VitiVisor dashboard development team used the Ruby on Rails framework for the backend development, connecting with the various data providers' Application Programming Interface (API) using the Ruby language.

One of the key requirements from Riverland Wine is that data service providers participating in VitiVisor act in an open-source manner. In contrast to proprietary software, the data providers were encouraged to act in a transparent, participatory and collaboratively way. With this in mind, VitiVisor has also been designed to be compatible with the Open Vineyard Standard (Collabriculture), allowing property data to be parsed directly into the dashboard.

The data providers in

Table 12 are categorised to demonstrate capability, rather than providing comprehensiveness.

Table 12: Inventory of external data providers

Data Category	Data Provider	Nature of connection
Weather	BoM	Weather data via API and click through for detailed view
Farm Finance	Know Your Numbers / Xero Accounting	Fixed and variable expenses per patch per property The Know your Numbers software (KY#) is connected to Xero for the enterprises chart of accounts, and VitiVisor is connected to KY#.
Water market data	H20X	Live pricing data Live price is not stored in the database, but the last valid data will be cached in case of invalid data coming through Trade data Historical trade data is stored in the database, variables stored are instrument, price, date, time and quantity
Water Delivery (infrastructure)	Central Irrigation Trust (CIT)	Data service through API, providing property water delivery data
Geospatial Imagery	CERES Imaging Pty Ltd	Geographic Parsing geojson data and store geoboundary data for each patch Spatial Imagery Spatial imagery data of three types (NDVI, water stress, thermal) for a particular property in a particular date is provided by Ceres Imagery and stored in the Vitivisor system
Geospatial viticultural data	VitiCanopy	Plant Area Index (PAI) PAI maps per patch at a particular date is provided by VitiCanopy and stored in the Vitivisor system
IoT Sensor	Commercially available soil moisture sensors (EnviroPro, and ICT international) Green Brain	IoT data from soil moisture sensor via The Things Network (TTN) gateway IoT data from soil moisture sensors via their API

4.5.4 Integration of ground-based vision data

In this project, a VitiBox is configured to be associated with a VitiVisor user account. Once a user opens the VitiBox side software, they will be able to select which patch of the vineyard property they are working on. The VitiBox can be easily mounted onto a moving vehicle such as an ATV or tractor, and by adjusting the angle of the camera, start capturing images whilst driving along the rows. Once capturing starts, images will be stored into a local disk (which could be a SD card or external drive), and records of these images with their associated GPS locations will be stored in a local database. After image capture, the data will be uploaded to the cloud automatically, once VitiBox is connected to the internet.

After uploading to the cloud, the backend processing is undertaken whereby an image will be stored in an AWS S3 bucket, and a record with GPS information will be stored into the cloud database. It will then be placed in a queue, for processing by the machine learning algorithms. The processing results of these images will be stored in the cloud database.

Once all the images of the same patch captured in the same day are processed, the system will start to generate heatmaps (Figure 24) using the predicted or detected results. The heatmaps will then be stored into S3 and the metadata will be stored into the cloud database, which can be accessed by the dashboard for displaying purposes.

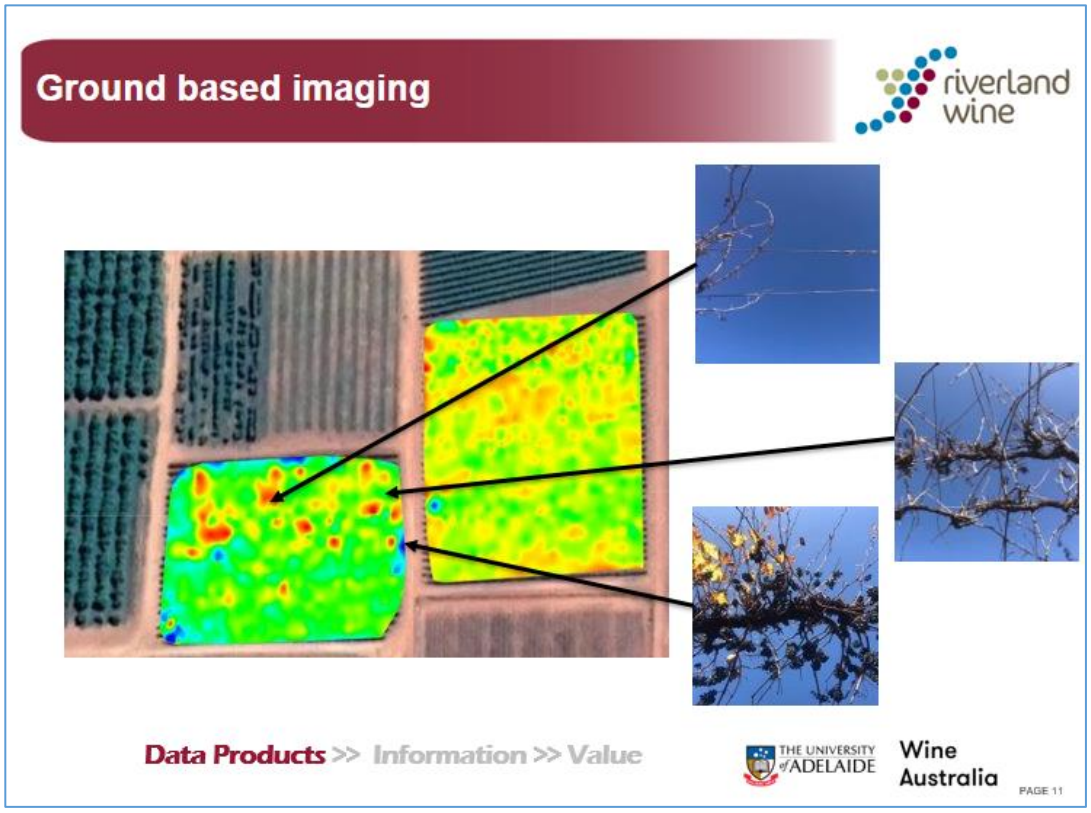


Figure 24: Image of a vine to a heat map

Each of the indices described in Section 4.2.3 are presented on the dashboard using heatmaps, with explanatory information provided to help growers interpret this information and support decision making.

4.5.5 Ingestion of grower information

Several data sets are required to be provided by the user to display information in the dashboard menus of Property Management and Finance. Moreover, to support a number of the features of VitiVisor, it is necessary to have access to a set of key vineyard parameters (e.g. row spacing, variety). As such, several mechanisms have been designed to support ingestion of vineyard data.

In property management the vineyard patch area can be derived from a combination of three data sources (see Figure 25):

- Data from geospatial imagery
- Know Your Numbers (finance module) or
- Custom – manually entered.



Figure 25: Display of the vineyard parameters

The VitiVisor platform is integrated with the Riverland Wine developed financial software called Know your Numbers (KY#, as described in Box 6) or, if users do not want to use this functionality, they can enter their data manually through the Finance tab on the main menu, as illustrated in Figure 26.

Box 6: Know Your Numbers (KY#)

The Know Your Numbers (KY#) Program was developed in 2014 to support growers in their decisions on profitability of their whole farm, varieties and individual patches. The program focuses on a simple input process of readily available information within accepted farm profitability conventions to deliver a meaningful farm profit (or loss) figure. It has been designed to deliver true farm profit (or loss), not distorted by cash flow or tax considerations, and contains the costs and income that relate to a specific crop year.

The system was also designed to capture regional data to enable enterprises of similar scale to be benchmarked, and has recently been updated to enable syncing with Xero to reduce friction associated with data entry and duplication.

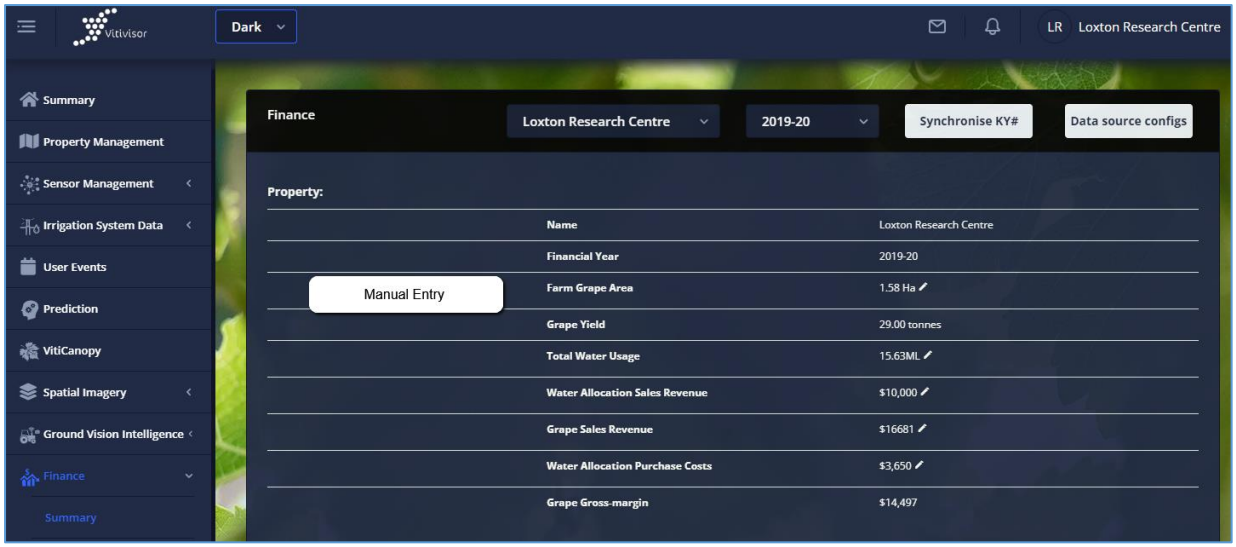


Figure 26: Manual entry of data without KY#

In the case of the KY# integration with VitiVisor, the user will need to input the enterprise's data directly into the key areas of KY# as per Figure 27.

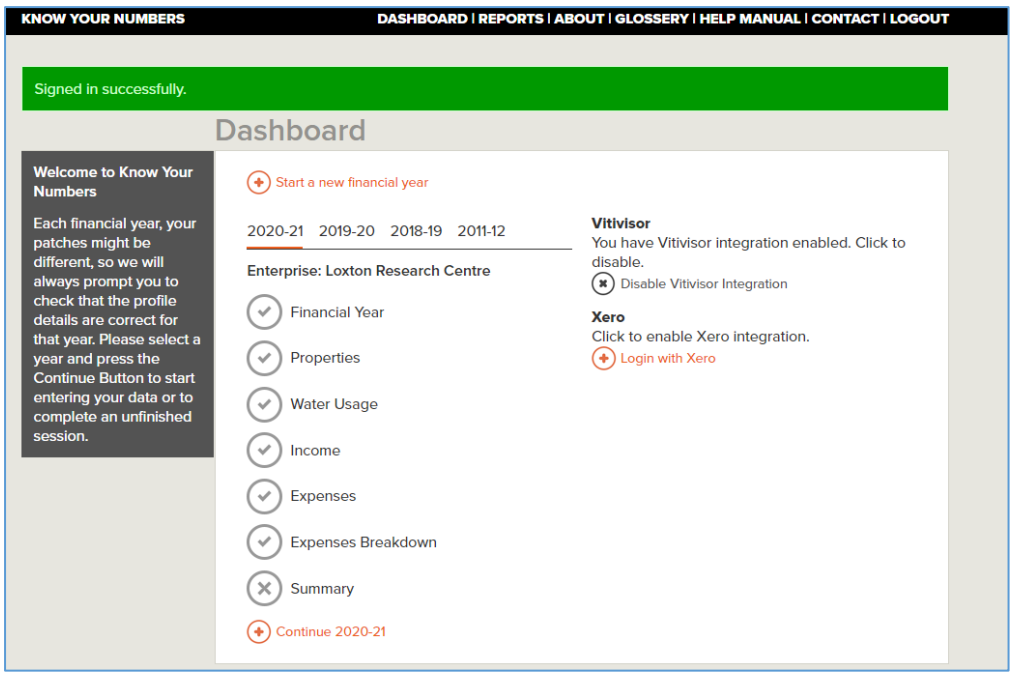


Figure 27: KY# parameter setup

Lastly, Figure 28 displays the vineyard parameters at a patch level that will be uploaded to the VitiVisor dashboard when it's synchronized on a daily basis, or when manually selected on the dashboard.

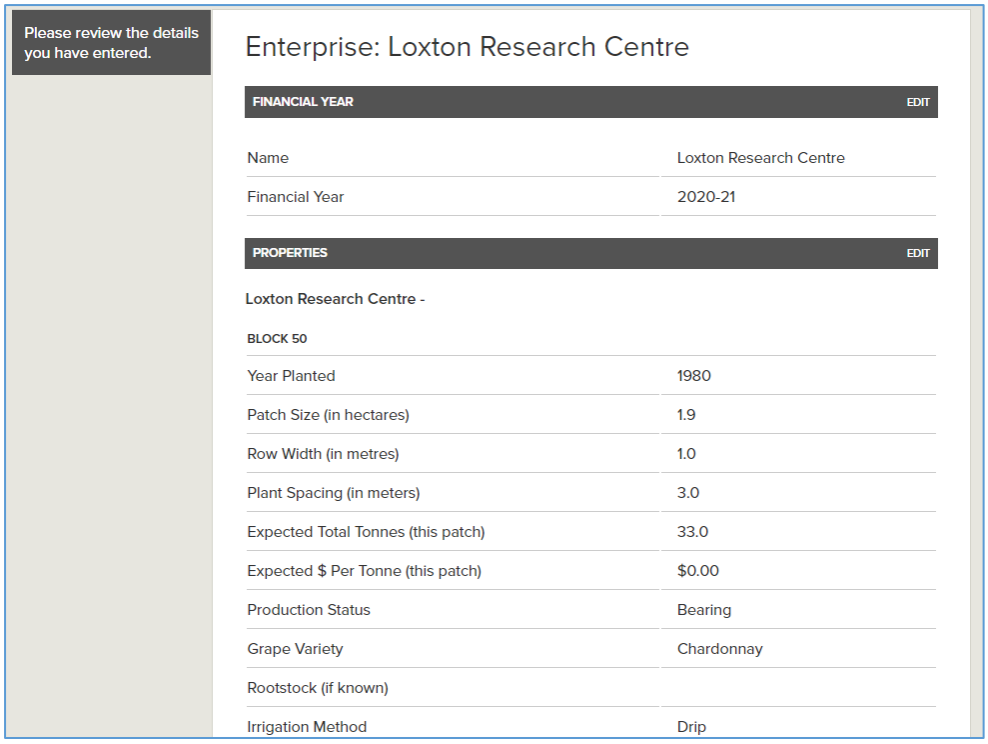


Figure 28: KY# Property data setup.

4.5.6 VineLOGIC integration

The VineLOGIC model runs as a standalone processor. The VitiVisor backend is collecting data as input for VineLOGIC, and periodically sends data into VineLOGIC to generate predictive output, such as predictive canopy measures, water consumption and yield. The generated output is processed by the VitiVisor backend and stored in the database, which is associated with a user account. The running of the VineLOGIC algorithm is also triggered by multiple types of events. For example, when there is a new VitiCanopy heatmap generated, or the user changed the scenario configurations, the VineLOGIC algorithm will rerun based on the new input and configurations.

There is also a module in the VitiVisor backend to handle data processing and displaying in the VitiVisor dashboard, for example by combining the (past) measured data from sensors with (future) predictive data of the same variable (e.g. soil moisture, leaf area index) and displaying them in the same chart. It will also handle user-triggered events in the dashboard, such as changing scenario configurations.

4.6 Adoption and Commercialisation

The adoption and commercialisation strategy of VitiVisor has been a collaborative process between Wine Australia, Riverland Wine and University of Adelaide over the course of the project and led through the VitiVisor Adoption and Commercialisation Committee (Section 3.2). The strategy represents a process of continual learning and testing ideas with different stakeholders and can be divided into three distinct phases.

In the first phase, from approximately 2019-2021, the strategy centred on transferring the Intellectual Property developed by the project and their related value propositions, either in whole or in part, to interested third party technology organisations operating in the Riverland region.

This stage involved establishing the framework in Figure 29 with the initial objective of establishing a list of industry partners to connect with and understand if there was a strategic fit between their organisation and VitiVisor's value proposition. A short survey was used to collect more information and used to both inform the different VitiVisor presentations to interested organisations and the development of an IP Summary document.

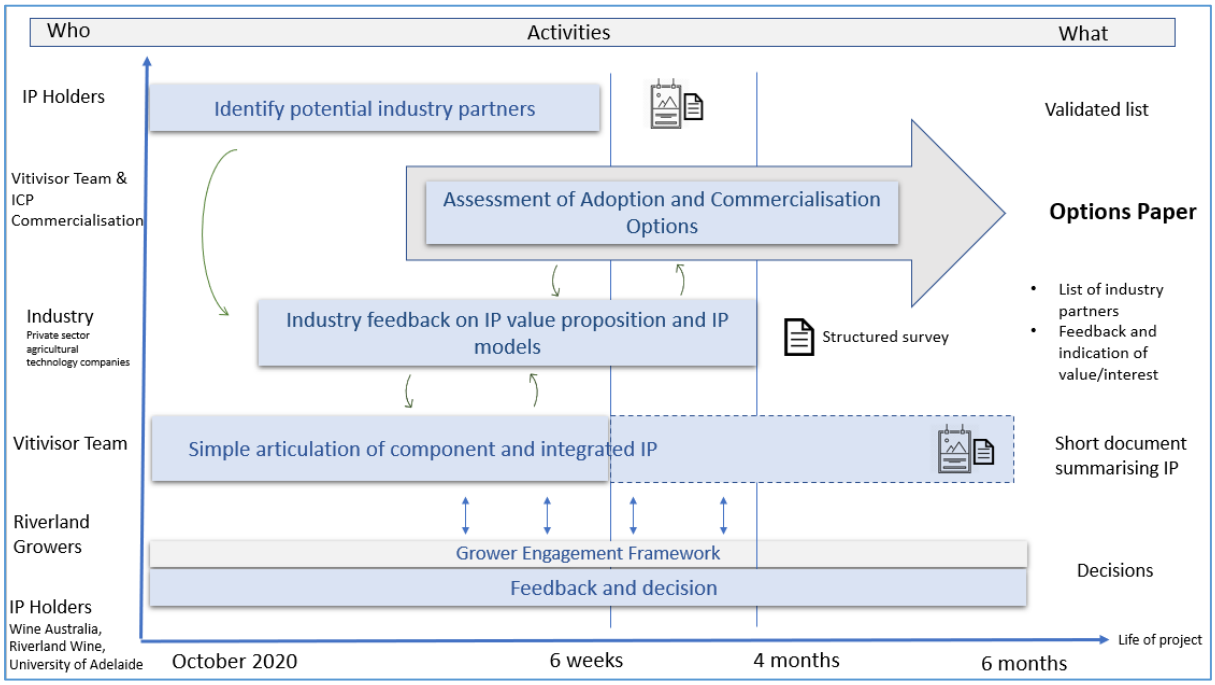


Figure 29: Adoption & Commercialisation framework October 2020

The second phase commenced in January 2021 when, at the request of Riverland Wine, the traditional commercialisation approach was halted and an alternative pathway was explored—referred to henceforth as the 'social enterprise' model[4]. This model was proposed as a result of perceptions by Riverland Wine that the 'traditional model' was unlikely to be successful in achieving broader outcomes such as enhanced digital literacy, data sovereignty and community wealth building, and that alternative options should be explored. An organisation with expertise in social enterprises (Ethical Fields) was brought in as an advisor and facilitator of discussions with Riverland Growers, with work comprising interviews with twenty eight stakeholders, a workshop on the 30th August 2021 and a subsequent report. This report was considered by Riverland Wine in September 2021, and additional funding was provided and a governance structure developed to action the recommendations of that report.

The third phase commenced in late November 2021, as governance and management changes at Riverland Wine led to a halt in discussions associated with the social enterprise model. These changes affected the co-design process for a scheduled 2nd workshop in December 2021 that was consequently cancelled. This led to a reversion to the original commercialisation model described in the first phase, with Wine Australia committing additional resources in the form of Doug Adamson from DNA Innovation Pty Ltd to both facilitate conversations with a variety of technology companies

[4] A social enterprise is commercial organisation that in addition to achieving commercial goals, seeks to achieve broader social, cultural and/or environmental outcomes. It is a broad term that can include both for-profit and not-for-profit entities, and may take a diversity of legal forms (e.g. cooperative, mutual or traditional company structure).

and to work with the project manager in developing an adoption and commercialisation pathway. Some of these conversations are on-going.

At the time of writing, no single model has yet emerged as a preferred model, with discussions continuing to refine various options. Nevertheless, there have been significant learnings throughout this process which are discussed in Section 5.5, and recommendations for further developments are summarised in Section 7.

5 RESULTS AND DISCUSSION

5.1 GROUND-BASED VISION SYSTEM

The ground-based vision system was tested by focusing on the predictive performance of a set of key indices that were identified by growers as being decision-relevant. Each of these are discussed in the following sections, including descriptors of performance both at the level of individual images, and at different levels of aggregation (including patch-level). Visualisations of the heatmaps produced from the computer vision algorithms are described in Section 5.4.3, and also can be accessed on www.VitiVisor.com.au.

5.1.1 Leaf Area Index Prediction

Using a small set of images with the ground truth leaf area index (LAI) value, together with a large set of images without ground truth LAI value, we can obtain an LAI prediction model with relatively high accuracy. Traditionally, transfer learning is often used to solve this kind of problem; that is to first train the model on a large number of labelled images from a variety of problem domains (using ImageNet, which comprises a library of approximately one million images) and then fine-tune the model on the smaller target dataset. We followed the same protocol to start from a pre-trained model from ImageNet, and then we use two different strategies to fine-tune the model. In the first strategy, we directly fine-tuned the model on our labelled but relatively small LAI dataset, while in the second strategy, we fine-tuned the model using both the labelled and unlabelled LAI dataset using our rank-based loss.

In Figure 30 we show a performance comparison of the proposed rank-loss based approach with naive machine learning based methods. Here we applied different random seeds to retrain the model five times to obtain an average performance comparison with error bar. We can see that our algorithm achieves much lower absolute error and much higher correlation coefficients compared to traditional deep learning approaches. In this application, we used 182 images of different varieties of vines with ground truth leaf area[5] information obtained from the viticulture team. Here we use 142 images for training and validation, and 40 images for testing. When splitting between the training, validation and testing set, we strictly enforce that the images in testing set do not belong to the training and validation datasets to avoid any data leakage.

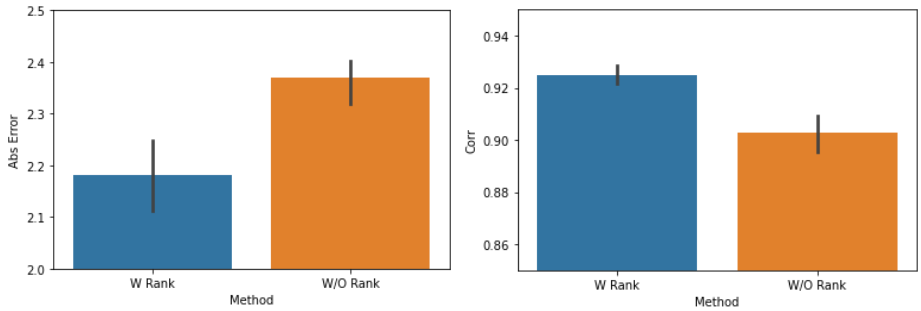


Figure 30 Comparison of the machine learning algorithms with/without rank loss. If we only use the labelled data where the exact LAI value is known, the final performance is significantly poorer than the model trained with both labelled and unlabelled data. For both the absolute error and correlation coefficients, models were trained five times with different random seeds to get an average result with less uncertainty. The errors and correlation coefficients were measured on the test set with 40 images, and the range of ground truth leaf area varies from 0 m^2 to 23.14 m^2.

[5] For this part of the analysis, the focus was on replicating the raw leaf area within an image, rather than the leaf area index which includes an additional normalisation factor.

As the range of the ground truth leaf area varies from 0 m^2 to 23.14 m^2, the relative error for our approaches is less than 10%. However, in practice we not only rely on the leaf area value of a single image, but also use a set of images to generate a heat map of leaf area value for a vine patch. Here the leaf area value in the heatmap can be derived as an average of several images taken in the vineyard. In this case, the error is much lower than the ones obtained from single images. As shown in Figure 31, the distribution of errors of our model are very close to Gaussian distribution, and thus by central limit theorem, the error of the average of several predictions will converge to a Gaussian distribution with zero mean. Here if we consider the average of 40 images in the testing set, the error for the average leaf area values becomes 0.35 m^2, which is very small.

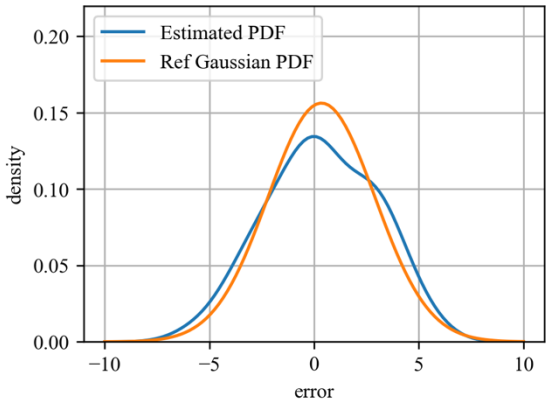


Figure 31: Distribution of error of our leaf area prediction model. The pdf of the distribution is estimated using the kernel density estimation (KDE) using the 40 images in the testing set, and the bandwidth for KDE is set to 1.2. From the result we can see that the distribution of the errors is very close to reference Gaussian distribution, where the mean and variance of the reference Gaussian distribution are obtained from the mean and variance on the errors from the testing set.

Similarly for the pruning weight prediction and bunch counting prediction, the error rates for a single image are 7.01% and 2.47% respectively. However for a patch in the vineyard with multiple sampled images, the accuracy can be much higher.

5.1.2 Inflorescence, Shoot Tips and Bud Counts

We first detect the inflorescences, shoot tips and buds, and then count the number of inflorescences and shoot tips. One challenge for inflorescence, shoot tip, and bud detection is that the resolution of the image must be high to maintain a proper visual feature for the detection. In this case, including the full image for training will cost a very large amount of GPU memory so that the training batch size must be very small. However, modern deep learning-based approaches heavily depend upon techniques such as batch normalization (Bjorck et al., 2018) to attain acceptable performance, which requires sufficient large batch size to work. Thus as described in the methods section, we the sample small patches from the full images to form a large batch for the training detector.

For inflorescence and shoot tips detection, we use 127 images for training, 27 images for validation and 20 images for testing. During the training procedure, for each 3088x2320 image, we randomly sample 100 640x480 patches from that image and then use the patch to form a training set of 12,700 images. The trained detector is then applied on the testing set to count the number of inflorescences and shoot tips, and the accuracy for a single image is shown in Table 13.

Table 13: Performance of inflorescences and shoot tips counting on a single image.

	Relative error rate	**Correlation coefficients**
Count of inflorescences	18.7%	92.13%
Count of shoot tips	17.4%	87.66%



Figure 32: Typical Inflorescences and shoot tip detection results. The left is the generated heat map of one patch in Loxton Research Centre, where green means more inflorescences, and red means less inflorescences; the middle is the typical detection results from the green zone, and the right is the typical results from the red zone. The results also demonstrate the good correlation between our predicted inflorescences/shoot tips counting and ground truth counting.

In practice, we will use a set of images to generate a heat map of number of inflorescences and shoot tips as shown in Figure 32. In this case, the value on the heatmap will be very close to the ground truth. The distribution of the counting error is very close to Gaussian distribution as shown in Figure 33.

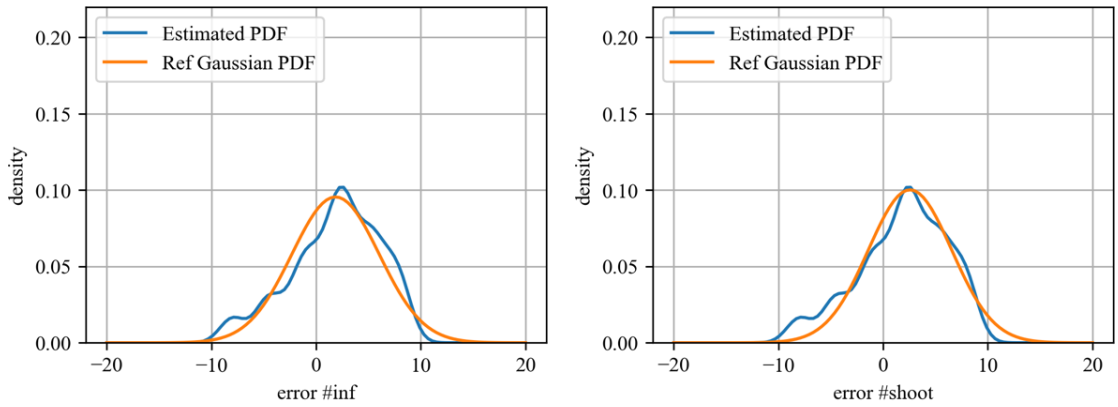


Figure 33: Distribution of our inflorescence and shoot tips counting results. The p.d.f. of the distribution is estimated using the Kernel Density Estimation (KDE) using the 20 images in the testing set, and the bandwidth for KDE is set to 1.2. From the result we can see that the distribution of the errors is very close to reference Gaussian distribution, where the mean and variance of the reference Gaussian distribution are obtained from the mean and variance on the errors from the testing set.

Once again, the counting accuracy of the whole patch can be much higher than the counting accuracy of a single image. For example, if we consider a small patch consists of 20 images in our test set, the overall counting error of inflorescences is 7.6%. In practice we taken 600-1000 images for the whole patch and in that case the overall counting error can become very close to zero. Similarly, for buds detection the average for single image is 10.3%, but the overall counting error over 12 test images is 0.7%, which indicates that the accuracy of the heatmap is very high.

The detection time of a single image on an NVIDIA GTX 1030 video card is about 1.1 seconds. Since GTX1030 is a quite weak video card with only 1,127GFLOPS, it will be possible to port our method to embedding systems such as NVIDIA Jeston AGS with 275 TOPS. While currently we use an embedding system to capture images and send the images to the server for object detection, in future work it is possible to replace the current embedding system with stronger system such as NVIDIA Jeston AGS and port our detector to the system to achieve real-time inflorescence/shoot tip/bud counting.

5.2 VITICULTURE MODELLING

An important challenge facing growers and agronomists is to make best use of the diverse data being collected, to try to improve vineyard outcomes and ultimately business gross margin. For example, how do canopy size data—being collected in this project through VitiCanopy (Section 4.1.3) and computer vision technologies (Section 4.2.3)—inform decisions pertaining to water application? The answer is not straightforward due to the complex relationship between biophysical variables being measured (canopy size), operational decisions (water application rates) and farm outcomes (yield and gross margin). This highlights that to fully leverage the benefits of data, it is necessary to combine what is measured in the field with model predictions (Section 4.3).

We demonstrate the prediction system's capability to provide an unprecedented level of situational awareness of vineyard status and likely future development by blending diverse data streams with model predictions in real-time throughout the season. Specifically, we demonstrate how predictions of key end-of-season outcomes—namely yield, which can ultimately inform irrigation decisions—are continually updated and improved by learning from data being collected in the field. In this section we provide a high-level summary of these results; links to scientific publications describing these results are provided in Box 7.

Box 7: Research contributions on viticulture modelling

The first real-time application of VineLOGIC, including integration of the VineLOGIC model with real-time data feeds on canopy volume, is contained in:

> *Knowling, M.J., White, J.T., Grigg, D., Collins, C., Westra, S., Walker, R.R., Pellegrino, A., Ostendorf, B., Bennett, B., Alzraiee, A. Operationalising crop model data assimilation for improved on-farm situational awareness and decision making. Under review and Agricultural and Forest Meteorology.*

This article not only provides a demonstration of the technology in viticulture, it also offers guiding principles to operationalise data assimilation (DA) technology for on-farm decision support. An excerpt from the paper, describing its contribution to the viticulture as well as wider agricultural science community, is as follows:

"Despite routine adoption in other fields, there is a lag in uptake of DA in agriculture due to the need for crop models to express end-of-season outcomes such as yield, to map insights to forecasts of these end-of-season outcomes throughout the growing season, and to enhance forecast reliability. To overcome these challenges, three guiding principles are introduced, providing a means to operationalize crop model DA for robust on-farm decision support. We apply the guiding principles using a South Australian viticulture case study. Our case study involves application of an iterative form of a widely used DA algorithm (ensemble Kalman filter) to dynamically update both static parameters and states associated with a grapevine simulation model. … It is shown how crop model DA can lead to not only significant improvements in forecasts of LAI but also to forecasts of end-of-season yield. … This study highlights the critical role that formal crop model DA can play in agricultural decision support through enhancing situational awareness, generating insights and supporting decisions in real time."

Figure 34 shows how the end-of-season yield prediction is updated as the season progresses (from left-to-right), as more data become available. The *x*-axis represents 'real time'; that is, the date at which the prediction of end-of-season yield is made, based on all available data at that time (even if there are none). An end-of-season yield prediction—made using all data up available at the time of prediction—is shown on the summary page of the prediction menu of the dashboard (Section 5.4.5).

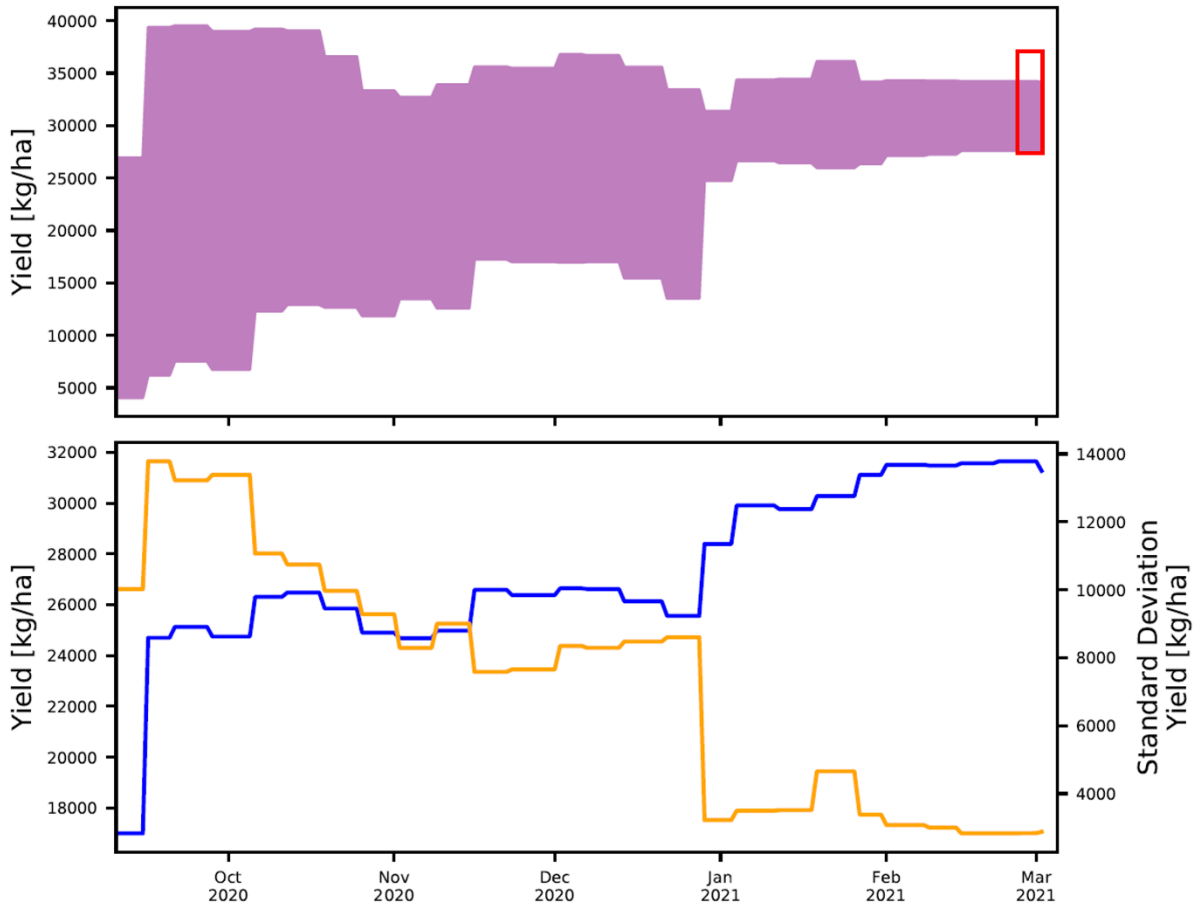


Figure 34: End-of-season yield ensemble (purple) and mean (blue) and standard deviation (orange) as a function of time. The field estimate range is depicted by the red box (95th percentile)

The uncertainty in end-of-season yield (the height of the purple bar in the upper plot and the orange line in the lower plot) is significantly reduced as the season progresses as more data are assimilated. This is expected—predictions are expected to become more reliable when based on more data that are relevant to the prediction. The considerable convergence of yield predictions in late-December coincides with the simulation and assimilation of canopy intervention.

Importantly, yield predictions appear to converge towards field estimates of yield from bunch analysis. These estimates are derived via completely independent means; the model does not 'see' yield data, it is only using weather and canopy size data together with biophysical processes represented in VineLOGIC to arrive at yield predictions. This result therefore provides an important validation step. It is noteworthy that the yield predictions encompass the field estimates of yield throughout the majority of the season. While this may be advantageous in some contexts, it does reflect a degree of variability about the yield predictions that may be considered overly conservative. This could be addressed through further refinement of prior uncertainty specifications or imposing stronger penalties on yield in accordance with expert and/or grower knowledge into the assimilation

procedure (we are currently investigating both of these), and through further 'training' of the system (i.e. assimilating data from more field sites and more growing seasons).

It is important to understand the basis for the yield prediction updates (e.g. Why did yield variability reduce late-December?). As discussed, assimilation of canopy size data involves updating the model's predictions of the canopy size to better align with observed canopy size data. By updating the model's internal state variables (e.g. LAI) to better reflect observed conditions, other predictions such as yield may also be informed—to the extent that these predictions are sensitive to updated variables—as shown above.

Figure 35 shows how the predicted trajectory of canopy size (in terms of LAI) is updated as the season progresses as more data become available. Reading from top-to-bottom are LAI time series predicted at different times during the growing season. From left-to-right (*x*-axes) is LAI time series at each of the prediction times. An LAI time series—predicted using all data up to present day—is shown via the 'see more details' button in the prediction menu of the dashboard (Section 5.4.5).

It can be seen that as more weather and canopy data become available and are assimilated into the prediction system, the predicted trajectory of canopy size is updated, with the variability in canopy size through the season substantially reduced. The simulated effect of canopy intervention (simulated in the model when LAI reaches a threshold) in December serves to reduce the variability of the canopy size trajectory, and also reduce canopy size values overall. Moreover, VitiCanopy data suggested canopy development rates during October were larger than expected based on the model alone, and therefore the model used this information to increase its predicted LAI values, which in-turn led to an increased end-of-season yield prediction at this time.

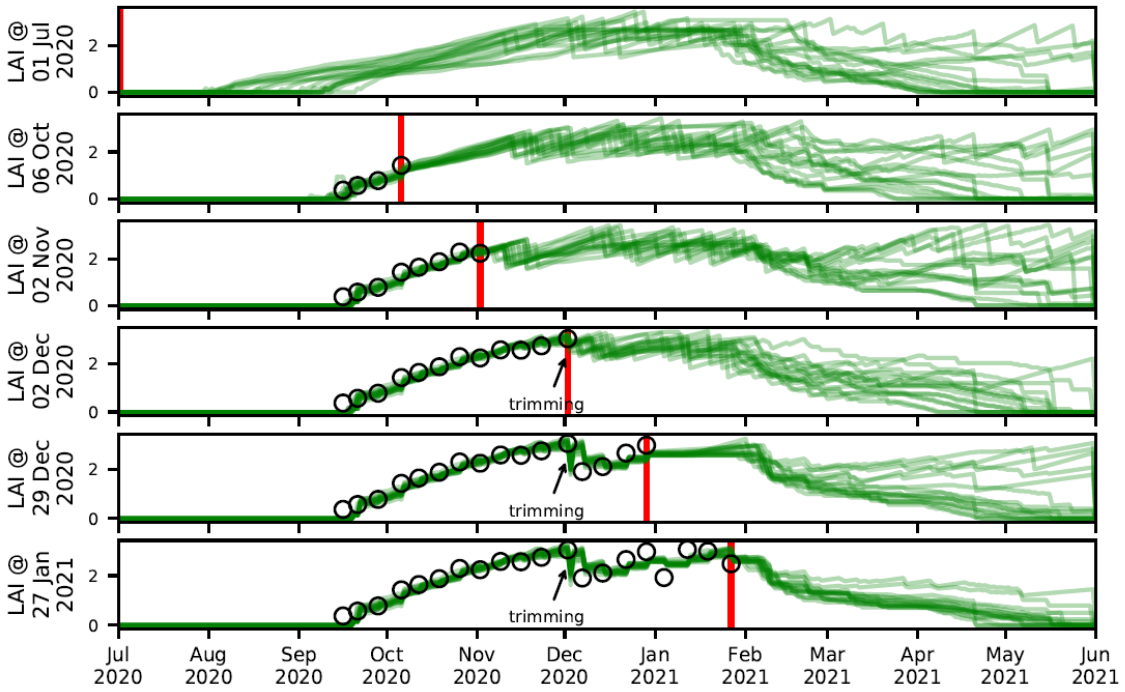


Figure 35: Ensemble LAI time series predicted at different stages throughout the growing season (depicted by red line). Observed LAI values are depicted by black circles. Canopy intervention events (i.e., trimming) are labelled

In developing the VineLOGIC prediction system, it is worth noting that there is considerable flexibility to augment the model's functionality and modify the contexts in which it is able to operate. In particular:

- Since the model underpinning the prediction system is process-based (i.e. driven by biophysical principles), the prediction system is simulating many more than the two key biophysical variables that are displayed on the VitiVisor dashboard. As such, it would be straightforward to evaluate and expose these variables on future iterations of a dashboard with low computational cost. It is suggested that variables such as Brix (or Baume) and carbohydrate reserves be evaluated in future.
- The prediction system is flexible in its ability to 'handle' different data availability situations. That is, the system will provide predictions of end-of-season outcomes (e.g. yield) as well as state variables (e.g. LAI) regardless of whether or not canopy data is available. Removing the requirement to assimilate canopy data would still produce predictions, although they would be most likely be less precise. Extension of the prediction system to include new data types is relatively straightforward from a computational standpoint, but will require validation with regard to how representative the data are of certain field variables, and how these relate to variables being simulated in the model.

In summary, we have demonstrated predictions of key vineyard outcomes throughout the growing season, and how these can be informed by data being collected in the field, ultimately leading to enhanced situational awareness. This can be combined with supporting grower operational decisions such as water application rates by exploring how changes in those decisions lead to different outcomes. Finally, the VineLOGIC prediction system has been combined with outcomes from the finance and economic assessment (Section 5.3) to develop estimates of gross margin.

5.3 Economic Assessment

As discussed in Section 4.4, the economic analysis conducted as part of the VitiVisor project focused on two distinct issues: (1) the use of financial, economic and benchmarking data to greater understand growers' cost of production and financial viability issues; and (2) an evaluation of the return on investment associated with the VitiVisor tool itself. Sections 5.3.1 to 5.3.3 focus on the first of these issues, with the benefit-cost analysis of VitiVisor is provided in Section 5.3.4. Note that a larger report titled "The economics and financial benchmarking of Riverland Grape production, and potential benefits of Vitivisor technology", independent to this one, characterises in more details the economics research associated with this project.

5.3.1 High-level assessment of Riverland grape growing farm financial characteristics

A detailed review of commonly cited benchmarks and indicators used in farm business and grape growing was undertaken. Utilising benchmarking best practice, this report uses indicator results over multiple years, rather than relying too heavily on snapshot indicators. The benchmarks presented here focus on whole-of-farm characteristics to provide a general overview of business viability and potential problem areas, to provide high-level context for more detailed analysis in subsequent sections.

There is limited published benchmarking grape growing information available that focuses on the South Australian Riverland. The studies that exist show that whilst cost of production seems to have been falling from 2009/10 – 2015/16, this has coincided with low grape prices over the same period; as a result, a substantial number of grape farms have had operative losses from 2009/10 onwards (Wheeler et al., 2022). This project addresses this knowledge gap by undertaking benchmarking analysis of the most recent available

Riverland grape farms were found to hold too few assets compared to liabilities, indicating the potential for business solvency issues.

secondary data sets for Riverland grape production based on sources including the ABS DataLab (2020) microdata which contains information on all grape growing businesses that lodged a tax return, as well as data from the SA Auditor General (2020) and Wine Australia (1999 - 2020). Key findings from this analysis are as follows:

- Average farm assets for medium-sized Riverland grape farms rose from $660,000 in 2001/02 to $930,000 in 2017/18.
- Capital values of SA grape properties fell consistently from 2005/06 to 2017/18. However, this does not include water entitlement value, which increased over the same period.
- Riverland grape farms hold too few assets compared to liabilities, as indicated by a low equity to asset ratio: the ratio was 40%, 35% and 32% for small, medium, and large grape farms respectively, and these ratios indicate business solvency issues.

Issues associated with solvency validate the concerns of growers highlighted during the early co-design meetings for VitiVisor about the need to focus on technologies that support growers to enhance profit margins, with diagnostic information to support action provided below.

5.3.2 Identifying key factors associated with farm net return from an irrigator survey analysis – existing data

As commented earlier in the methods section, regression analysis was applied to two representative telephone surveys of 1,946 southern MDB irrigators from 2010/11 (946) and 2015/16 (1,000), leading to a final cross-sectional dataset of 593 southern MDB horticultural irrigators, of which 258 were SA Riverland grape farms. It was found that farm size and water entitlement ownership have the most important influence on horticultural and grape irrigator net farm income (net farm income is gross revenue take gross variable costs). While larger farms and farms with more water entitlement ownership have higher net farm income in SA and for MDB grape farms overall, farm size has a negative impact on net farm income for grape farms outside SA, and negatively impacts net farm income per hectare on SA grape farms.

This points towards the conclusion that large—and likely more automated—farms have larger net farm income than small family farms. The negative influence of farm size on net farm income per hectare might indicate that there is a size ceiling for grape farms, above which further size increases will have a negative marginal income effect. There is also a negative influence of water consumption rate on net farm income which can be explained by the savings of water costs associated with efficient irrigation. The negative influence of a high debt to equity ratio on net farm income is not surprising, given that businesses with high debt levels tend to be less financially viable, and often have to spend considerable resources on debt servicing.

5.3.3 Identifying detailed variable and fixed cost and revenues for grape growing in the Riverland - Farm financial case studies – primary data collection

The case study analysis of twenty-five grape growers represents the primary dataset for the analysis, and can be used to understand in more detail the exact breakdown of various costs by grape growing farms in the Riverland. It also allows for a detailed assessment of the potential benefits that may be derived from VitiVisor dashboard adoption by growers. Interview and farm financial records were combined to create an in-depth assessment of growers' income and cost positions, with findings including that for the financial year of 2019-2020:

- Growers had a total average grape area of 57 ha, comprising an average of 30 hectares of red grapes and 26 hectares of white grapes.

- Average water use per hectare was 7.61 ML, and average grape yield was 21.8 tonnes per hectare for red varieties and 23.4 tonnes per hectare for white varieties.
- Growers received on average a grape gross margin of $ 6,268 /ha. Grape revenue is on average $ 13,197 /ha, roughly double of average variable expenses per ha.
- Profit is on average $3,883 /ha and operating margin averages at 12%. Interestingly, operating margin seems not structurally associated with farm size, as there are small farms with high operating and large farms with negative operating margins.
- Average total variable grape cost per hectare for case study irrigators are $6,928 /ha and average total fixed costs per hectare are $3,815 /ha.
- Average water costs are the largest component of total variable grape costs per hectare with 28% and $2,221 /ha, followed by contract and labour costs with 16% and $1063 /ha. Variable cost analysis reveals that contract & labour, and water costs are very different across growers, mirroring the different water ownership, labour and trading strategies.

Figure 36 provides an overview of the distribution of various average variable costs per hectare for 2019-20. The important points to note from this figure is that water costs are highly variable: some farmers had very large water variable costs, while others had minimal water costs. The next largest variable expenses were contract and labour, followed by 'other' expenses. This variability is shown by the water costs considerably decreasing as a percentage of total variable costs in 2020-21.

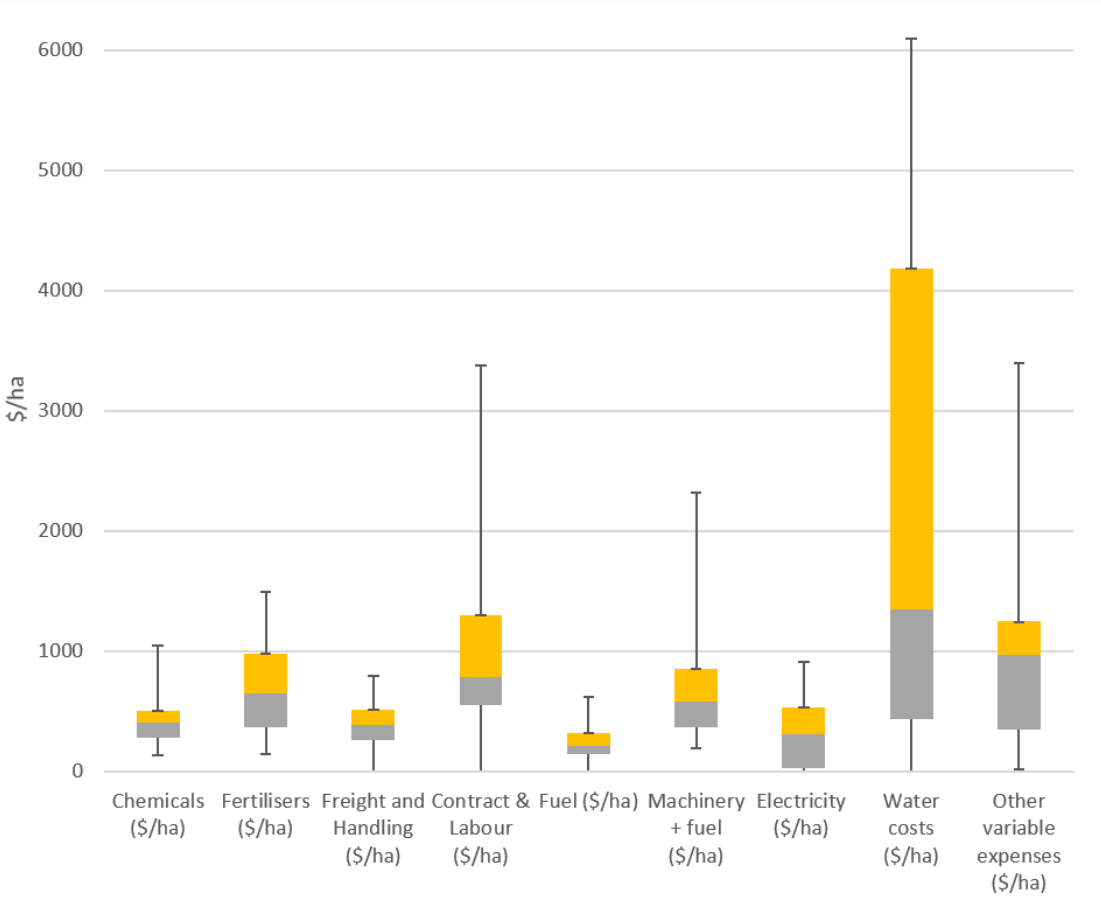


Figure 36: Boxplot of case study variable viticultural costs of per ha for 2019/20

Case study participants were invited to provide their historical farm financial records for the years 2016/17- 2018/19, and also updated data for 2020-21 was sought. Some findings included:

- Historical median grape gross margin per ha was $3,618 /ha and $495 /ML. Median grape revenue was 390 $/tonne or $9,158 /ha, with variable grape expenses per hectare roughly 53% of revenue, at $4,891 /ha. Median fixed expenses per hectare amounted to $4,842 /ha.
- The largest single variable grape cost position for the period 2016/17 – 2018/19 was water costs per hectare at a median of $450 /ha and 17% of average total variable costs. Median other variable grape expenses also made up 23% of total variable grape costs at a median of $1,148 /ha consistent across time periods.

Grape gross margin per hectare and grape revenue per hectare are associated with the level of new technologies irrigators' use. Case study irrigators who use more (sophisticated) technology tend to achieve higher grape gross margins per hectare, though the relationship is not strong, and the causal relationship is still in question. A journal publication describing this relationship and others is described in Box 8:

Box 8: Decision-making and resilience in agriculture: improving awareness of the role of accounting

An exploration of farmer decision-making and resilience, including the challenges with soft information such as accounting data and its formal role in strategic decisions, has been provided in the following manuscript:

> Tingey-Holyoak, J., Wheeler, S. & Seidl, C., 2022, "Decision-making and resilience in agriculture: improving awareness of the role of accounting", submitted to *Meditari Accountancy Research*.

Australian agriculture is facing increasingly uncertain weather patterns which is impacting financial performance, exacerbated by worsening terms of trade and a decline in commodity prices. Increasing the resilience and adaptive capacity of the primary production sector is of key importance. Governments and farmers' groups often depict technology adoption as the salvation of farming, often ignoring the importance of decision-making processes and soft information skills and needs. Drawing on a strategic choice perspective, we explore the links between farmer characteristics, attitudes, technology orientation, decision-making and financial performance, to explore how growers could be better supported with accounting information and tools in their strategic choices and adaptive capacity. Results show that farmers with low operating profit margins struggle with minimising their variable costs, especially water costs, in addition to spending double the time making decisions. These growers were also likely to demonstrate a lack of resilience, and indeed, this was also shown by lower profit growers being less likely to perceive climate change as a threat. The results highlight the potential for accountants to make more use of technological advances and for this data to be utilised for enhanced prediction that can provide information for important on-farm strategic decision-making. The paper recommends that simply packaged biophysical and financial data could better support farmers struggling to make a profit in their strategic choices and adaptation.

All of this work provides a strong foundation for how to define, depict and describe the economics of grape growing on the VitiVisor platform. The overall aim is to provide comparisons and snapshots to growers of a) their own financial situation, and b) how they compare to other growers in the region. This comparison and provision of information is the first step in allowing growers to recognise that

perhaps there is different aspects of their own management that could be better optimised. For an example of how this is displayed in the dashboard refer to Section 5.4.4.

5.3.4 Understanding the net benefits of VitiVisor adoption for Growers - Financial BCA analysis

As mentioned previously, a financial benefit cost analysis (BCA) was conducted to estimate the net benefit of adopting Vitivisor. BCA is a widely used tool for evaluating alternative courses of action that enables costs and benefits to be monetised to enable direct comparison across the full suite of options under evaluation.

Utilising very conservative estimates of the potential benefits of the technology, the expected net present value of adoption over the five-year period was estimated as $1,219 per hectare, or $69,749 per average farm enterprise. The benefit cost ratio was calculated as 1.5, indicating an expected net return on investment of $1.50 for every dollar of investment. If more optimistic assumptions were used, the return on investment was 2.9.

The findings in the report show that the threshold initial cost of VitiVisor dashboard software packages beyond which the expected incremental net return from the investment would be negative was calculated for the average farm enterprise as $67,198 over the five-year lifespan. There were some limitations noted in the report but overall analysis shows potentially reduced decision-making times overall for some farmers, including significantly greater estimated savings in electricity, fuel and labour costs and greater revenue from maximising grape yield. In addition, it is highly likely that our estimates are conservative for the following reasons:

- Reduced operator labour time on farms was not included and the time involved for using the Dashboard may be overestimated;
- Significantly greater savings in electricity, fuel and labour costs could be achieved than the 10% value used in this analysis, plus some farms will achieve significant savings in water that were not costed;
- The salvage value from the technology investment was not included and the lifespan of the technology may be longer than five years;
- Improvements in health and safety operating issues, identified as potential benefit of adoption, and reduced chemical application, were not included;
- Improvements in grape, soil, water and vine quality were not included; and
- Community benefits of the technology (e.g. such as reduced chemical use; predicting yield for the season, improved logistics etc) were not quantified.

Overall, it is important to note that these assumptions are based on predictions only, and a greater understanding of the actual adoption and efficiency benefits of the Dashboard may be warranted, to encourage the maximum social benefits that can be derived. The findings demonstrate that the Riverland region's growers need improved farm outcomes including profitability, and that the adoption of VitiVisor dashboard technology can be expected to be beneficial for growers for improved on farm decision-making.

5.4 VitiVisor Dashboard

In this section, key elements of the front-end functionality of VitiVisor is presented and briefly described. Access to the VitiVisor dashboard is also available at www.VitiVisor.com.au.

5.4.1 Landing page, property management and notification system

The summary page (Figure 37) is the default landing page that users will see after logging into the VitiVisor dashboard. This page provides widgets including notifications, growth stage tasks, weather, water allocation price, water trade and finance expenses. Users can customise what information is displayed by clicking the settings icon on the right top corner of this page.

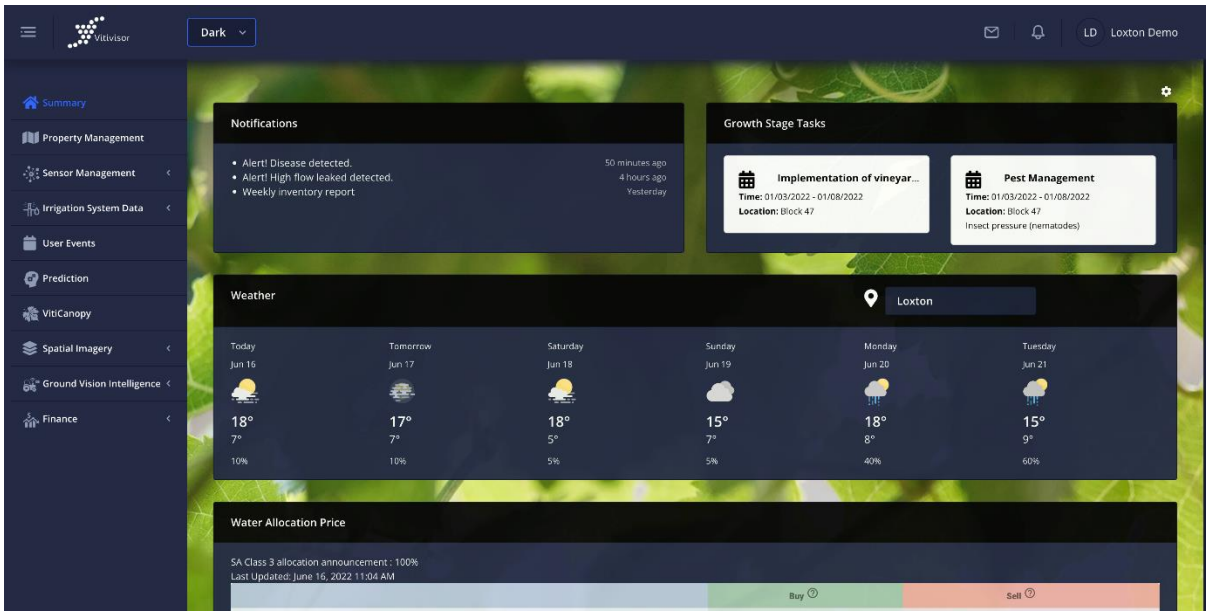


Figure 37: VitiVisor dashboard landing page

The property page (Figure 38) displays all the vineyard patches of a user's property on a map view. Additionally, it provides more information in a table view, where users can access information about the patch by clicking the patch name in the table. A unique feature of this page is the display of the current vine growth stage, whereby users can also go to the growth stage page by clicking the growth stage name displayed in the table view.

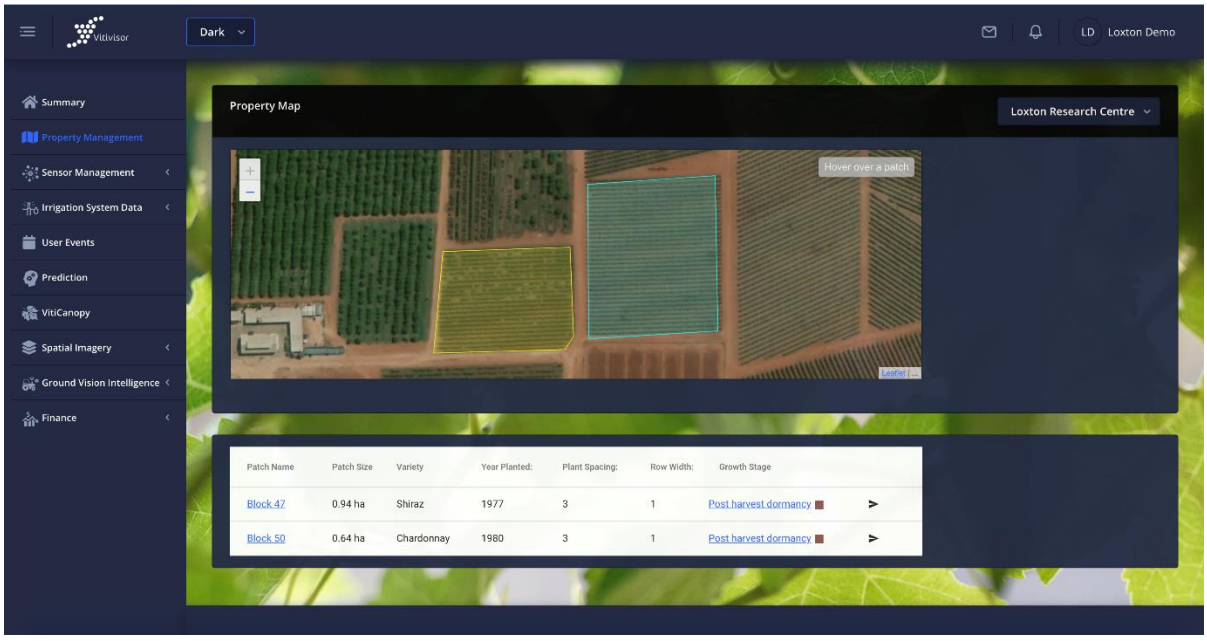


Figure 38: VitiVisor dashboard property page

In growth stage page (Figure 39), users can see all the growth stages details for a patch. The system has default values for these growth stages, and users can edit them any time to fit their own needs. For example, they can edit the start and end date of a growth stage and edit tasks for this growth stage. Importantly, for each growth stage there are customisable management stages with corresponding tasks (Figure 40), providing a starting point for a customizable guidance system.

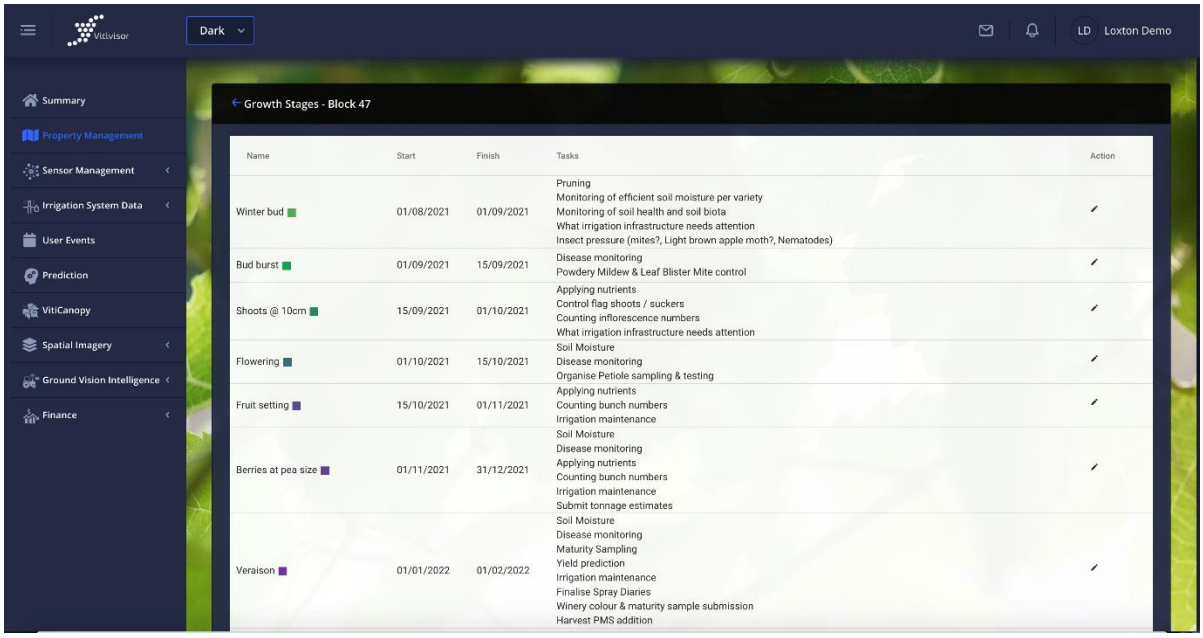


Figure 39: VitiVisor dashboard growth stage page

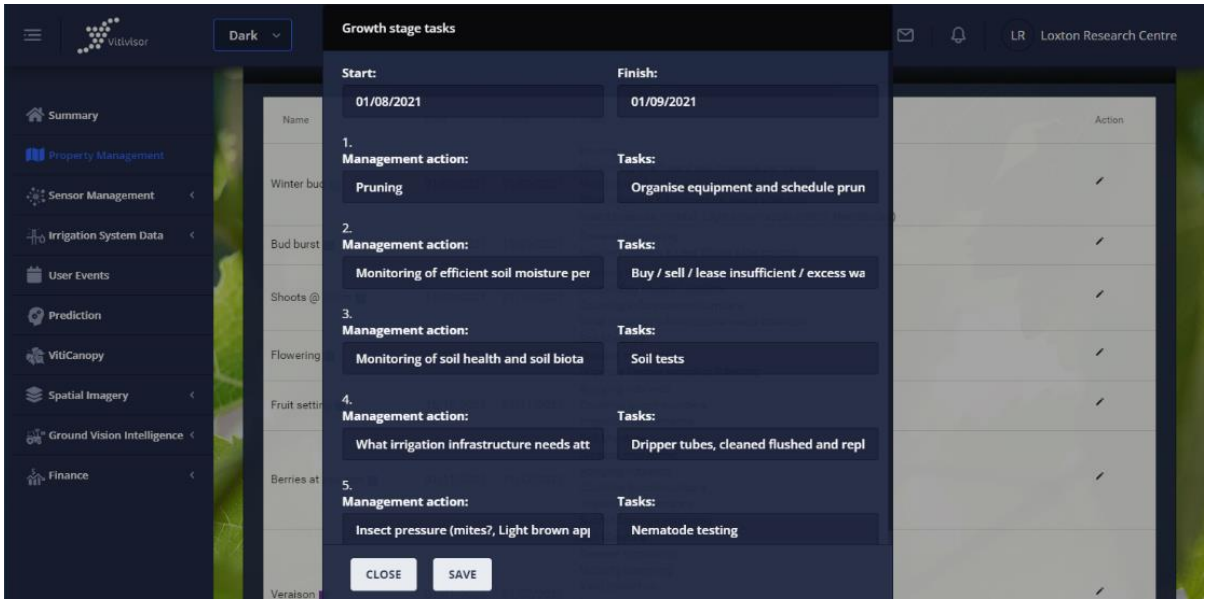


Figure 40: VitiVisor dashboard management actions by growth stage

5.4.2 Sensors and irrigation system data

The Internet of Things (IoT) is an integral part of VitiVisor's overarching data acquisition and storage system that collects real-time data from sensors on the vineyard. Multiple types of sensors such as soil moisture and dendrometers were installed by the Vitivisor Infrastructure team at the Loxton Research Centre and are connected to The Things Network (TTN) using LoRaWAN gateways. The

data is streamed to the Vitivisor dashboard and displayed in the TTTN page. More detail is provided in Section 4.1 of this report.

On this page, the sensors are displayed as markers on a map view and in table view. A sensor is shown as a green marker if it is working correctly and shown as a red marker if it is faulty. Users can click the view button on the Action column of the table view to see details of the sensor data in another page.

The VitiVisor dashboard is also connected with a third-party commercial product called GreenBrain from MEA that connects to its sensors, shown as markers in a map and table views (Figure 41). Users can click the view button on the Action column of the table view to see details of the sensor data in another page.



Figure 41: VitiVisor dashboard IoT sensor page

Water for irrigation of the vineyard is delivered by the Central Irrigation Trust (CIT) infrastructure. The VitiVisor dashboard is connected with CIT's API, providing irrigation meter readings on a property. The CIT irrigation outlets are shown as markers in a map and table views (Figure 42). Users can click the view button on the Action column of the table view to see details of the meter reading data in another page.

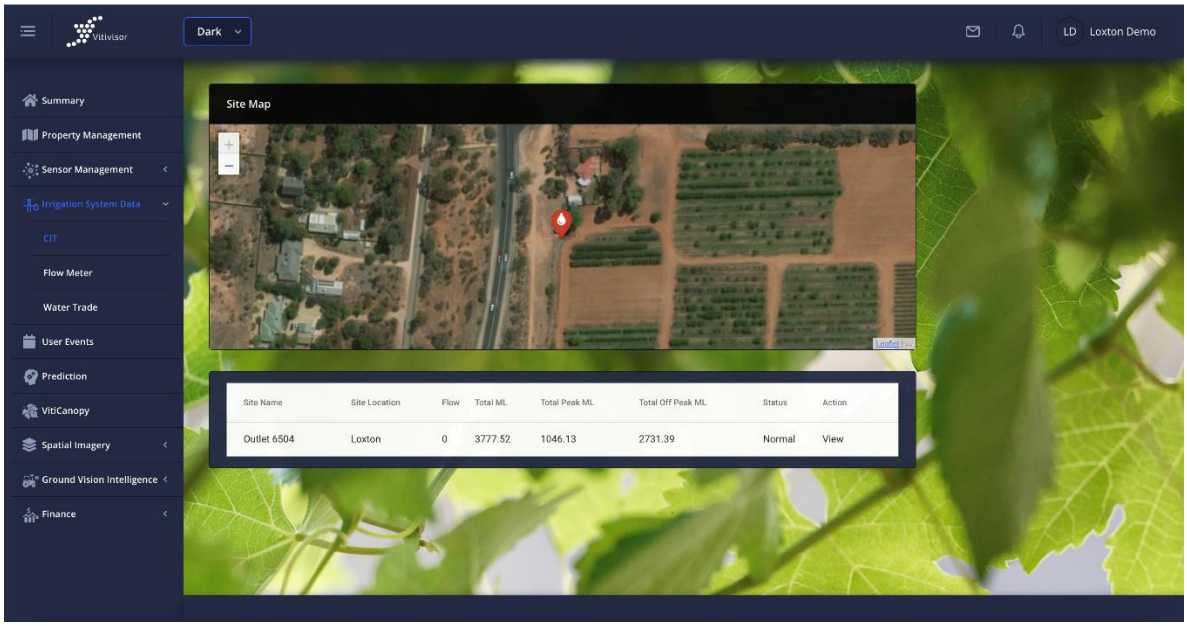


Figure 42: VitiVisor dashboard water delivery page

Another key piece of the water infrastructure is a flow meter that measures the flow rate and pressure of the mainline of the irrigation infrastructure. The sensor data is transmitted to the VitiVisor backend and displayed on the dashboard, shown as a marker in a map and table views (Figure 43). Users can click the view button on the Action column of the table view to see details of the meter reading data in another page.

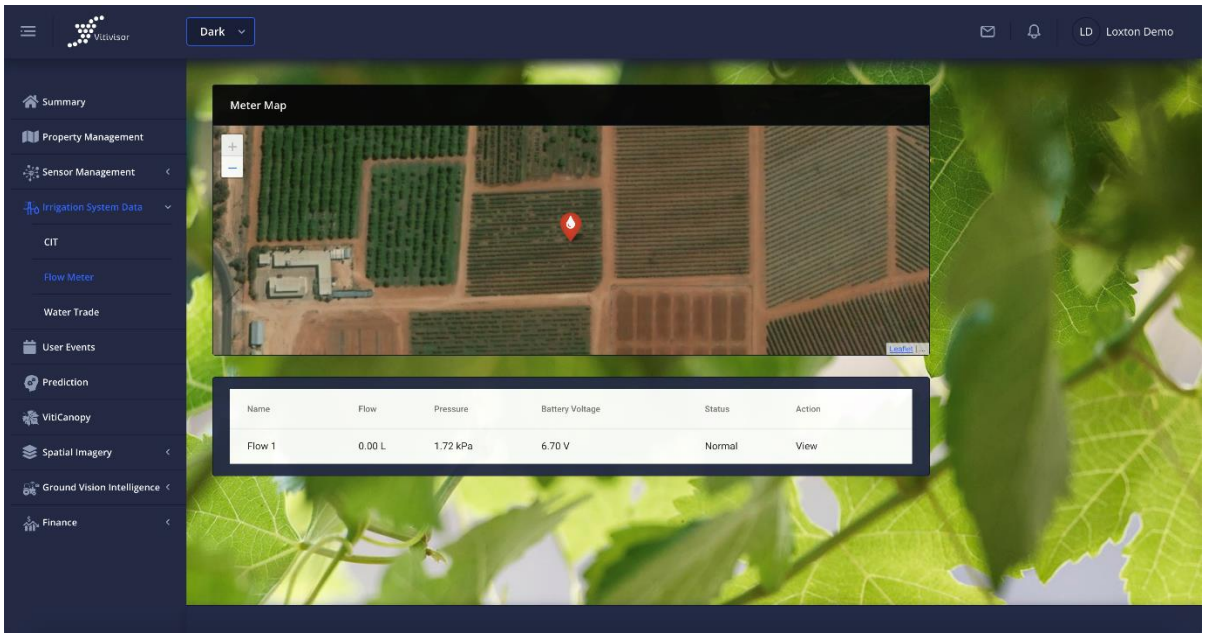


Figure 43: VitiVisor dashboard flow meter page

5.4.3 Spatial imagery

The spatial imagery pages display the data from VitiCanopy, Ceres Imaging, and the Ground-based Vision System sources that are described in more detail in sections 4.1 and 4.2 of the report.

5.4.3.1 VitiCanopy

VitiVisor can be configured to connect with a user's existing VitiCanopy account. Users can create a location with the same name of a patch in VitiCanopy. The heatmap generated by VitiCanopy will be displayed in the dashboard. Users can switch between different patches and date to view the heatmaps. A time series chart of patch-average PAI is also displayed adjacent to the heatmap, assisting the user by indicating the average PAI changes of a patch over the time period.



Figure 44: VitiVisor dashboard VitiCanopy page

5.4.3.2 Ceres Imaging Pty Ltd

Ceres imaging combines high resolution aerial imagery with advanced analytics to provide growers with decision making information. Their data for the Loxton Research Centre is used to generate a maps for:

- The NDVI, which allows identification of areas of variable canopy growth which may be the source of reduced production.
- The water stress index, which can be used to identify areas of irrigation deficit for growers to check or make changes to their irrigation.
- The core thermal index, which can be used to identify irrigation leaks and other features signified by changes in the canopy temperature.

These maps can be used by the grower to compare previous time periods and use them as a basis for further investigation. An illustration of the NDVI map is provided in Figure 45.

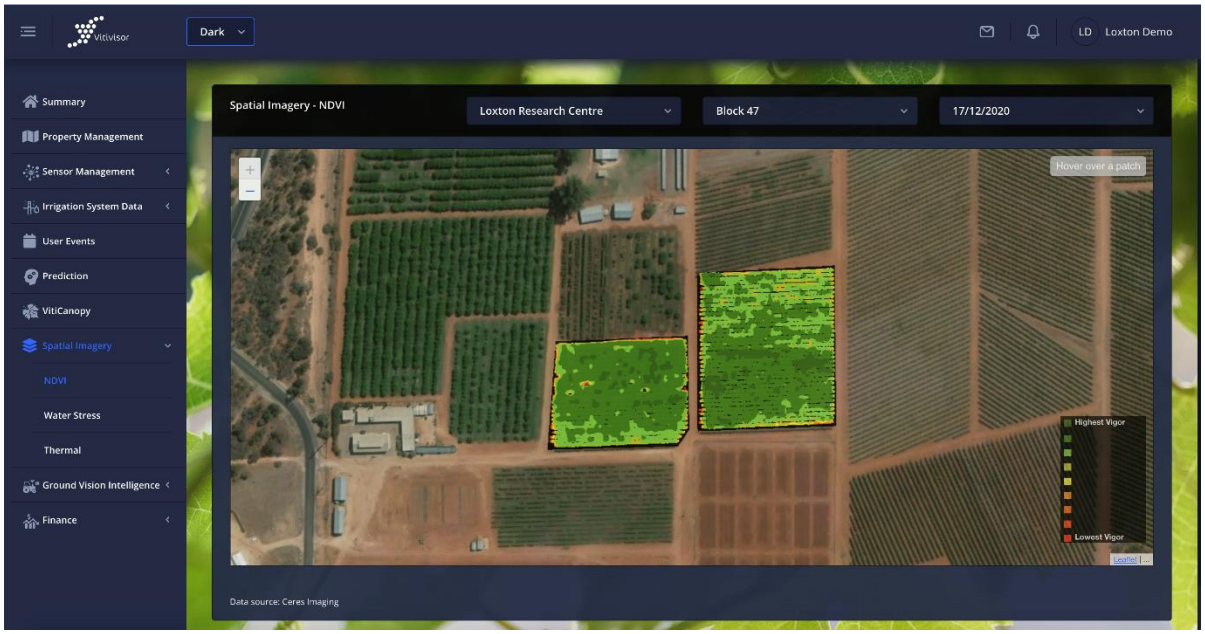


Figure 45: VitiVisor dashboard display of Ceres imagery (in this case, NDVI)

5.4.3.3 Ground-based Vision Intelligence

The ground-based vision intelligence pages display the output of the ground-based vision system (GBVS) designed and created by the VitiVisor team, to collect grapevine images in the field throughout the growing season. These images are processed in the cloud to generate user-friendly heatmaps, generated by the machine learning algorithm in the VitiVisor backend, as described in Section 4.2.

A range of key vine indices are presented using heat maps, which are displayed on the dashboard in a map view only, correlated with a legend ranging from the lowest to highest value. Users can select different patches and date for the relevant heat maps:

- Leaf Area Index (LAI) prediction
- Pruning weight prediction
- Inflorescence detection
- Shoot tip detection
- Bud detection

An example of the display of one index (inflorescence) is presented in Figure 46. Explanatory information accompanies each heat map, including information about the specific index and its relevance for grower decisions. Samples of the actual detection results (images and associated index values) are also shown on this page. Users can view the enlarged image for details by click the image on the dashboard.

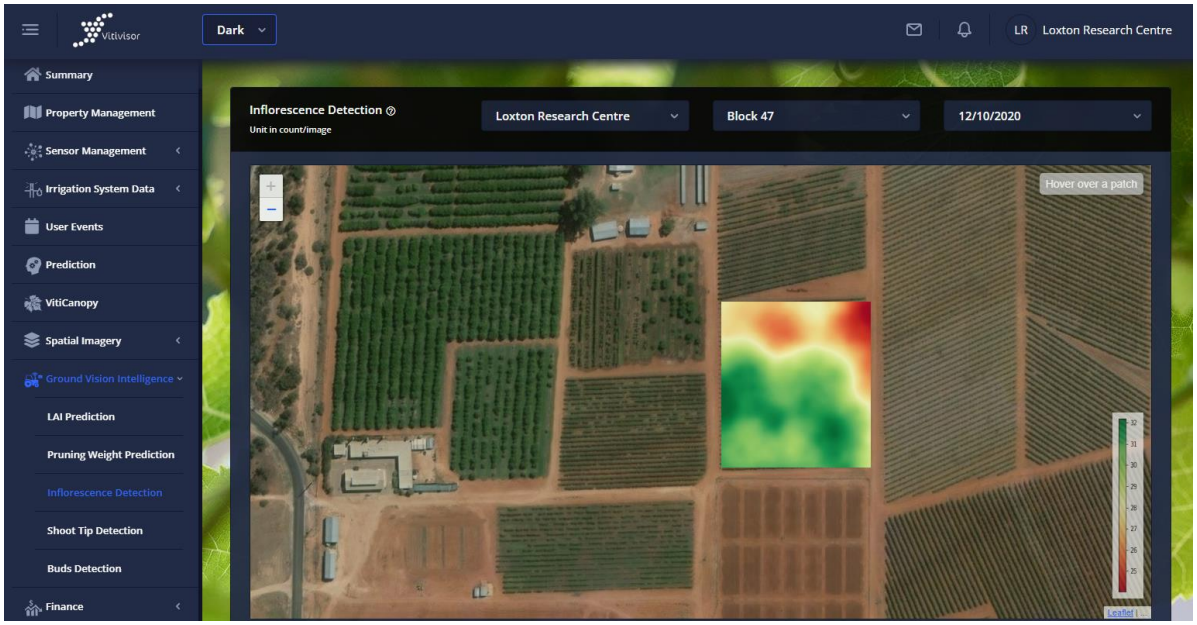


Figure 46: VitiVisor dashboard display of ground-based vision system indices, including explanatory information and example images corresponding to different index values

The inflorescence or early flowering stage images were collected via mobile phone camera and manually counted to ground truth. Later a machine learning algorithm was applied to automatically detect and locate all the inflorescence visible in vineyard wide images (reference to images below). These detection results generate a heatmap that can be displayed on the dashboard. Samples of the actual detection results (images) are also shown on this page as per Figure 47. Users can view the enlarged image for details by clicking the image on the dashboard, as per Figure 48.

The heat map displaying predicted inflorescence number is relevant for grower decision making as it can be used for early yield prediction and to assess the variability of inflorescence or potential bunch numbers across a vineyard block. For example, in the heat map shown above areas in red may indicate an irrigation problem or other issue that may need to be addressed to improve the uniformity in that vineyard patch. Creating a history of the patch also allows a grower to monitor yield fluctuations and better manage seasonal and patch variations, directing appropriate resources such as water and nutrients to where they are needed most.



Figure 47: Samples of the actual detection results

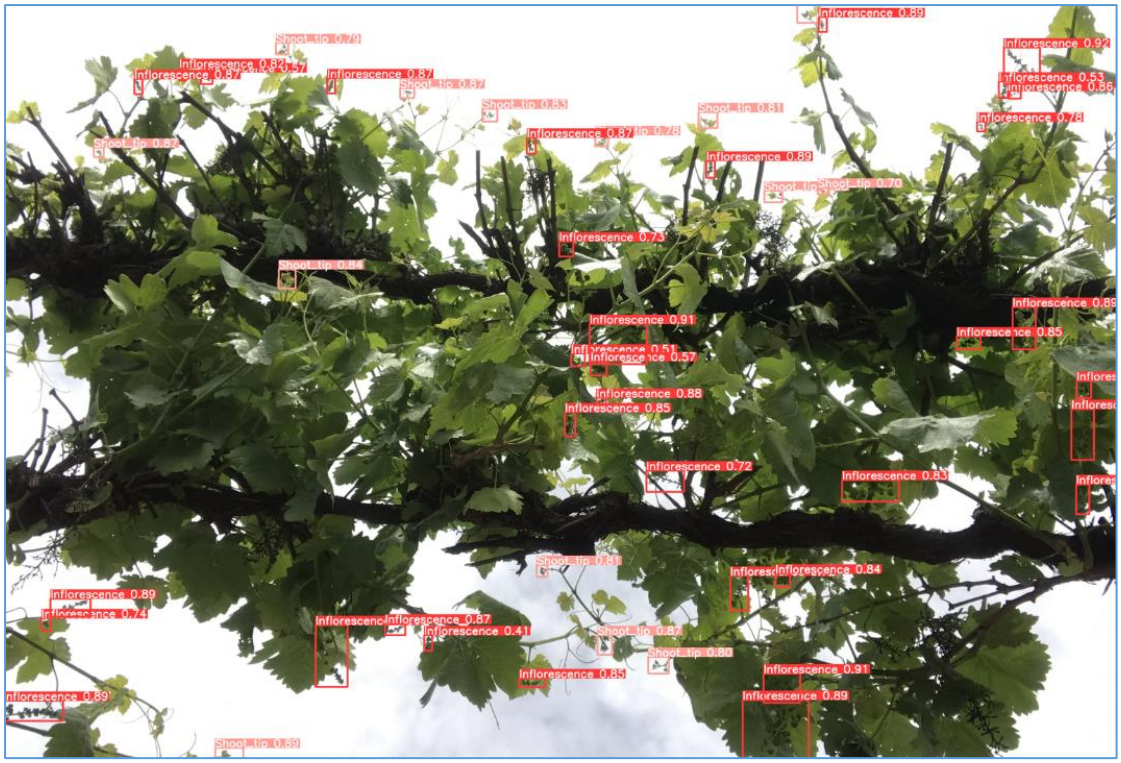


Figure 48: Enlarged image of actual detection results

5.4.4 Finance and benchmarking

After retrieving data from KY# or entering manually on the Finance summary page, users are able to view their expenses breakdowns in the Finance expenses page (Figure 49). Variable and fixed expenses are displayed in pie charts and tables in this page.

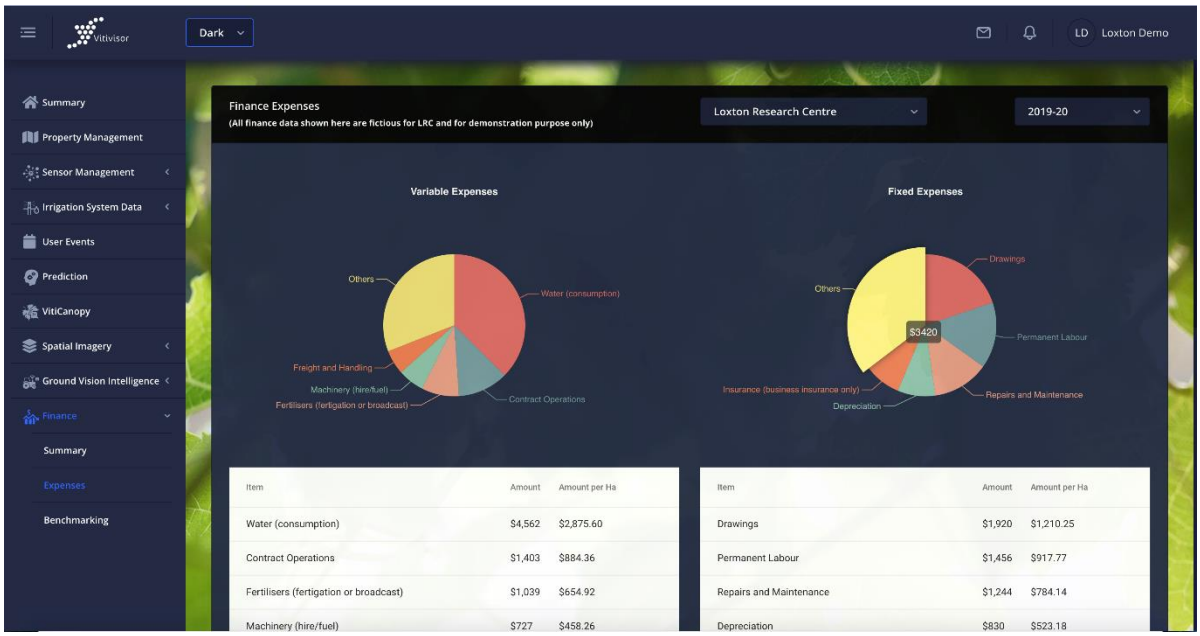


Figure 49: VitiVisor dashboard display of financial information

Users can compare their financial data with the regional average data in the benchmarking page (Figure 50). The regional average data is calculated based on all the financial data collected in the VitiVisor application.



Figure 50: VitiVisor dashboard depicting the grower's situation against the regional average benchmark

Name	My Farm	Regional Average	% Difference
Yield / ha	45.38 tonnes	24.55 tonnes	84.86% ↑
Water use / ha	9.85 ML	7.6 ML	29.63% ↑
Grape gross margin / ha	$2,834.63	$12,892	78.01% ↓
Variable cost / ha	$7,680.03	$5,000	53.6% ↑
Fertilizer cost / ha	$678.87	$620	9.5% ↑
Chemical cost / ha	$491.66	$545.2	9.82% ↓
Casual labour cost / ha	$1,149.1	$1,777	35.33% ↓
Contract operation cost / ha	$884.36	$1,000	11.56% ↓
Fuel cost / ha	$266.63	$490	45.59% ↓
Electricity cost / ha	$237.64	$402	40.89% ↓
Water consumption cost / ha	$3,101.26	$6,421.5	51.71% ↓

Figure 51: VitiVisor dashboard snapshot comparisons of the differences in grower's situation against the regional average benchmark in VitiVisor

5.4.5 Scenario and prediction tool

The VineLOGIC summary page provides users with the predictions that are generated by the VineLOGIC model (Figure 52). The predicted yield is shown in a horizontal bar chart with estimated ranges. Patch level grape prediction, water consumption prediction and finance information is also displayed in this page. Users could edit water price and grape crush price to explore different gross margin scenarios based on different watering strategies and the yield implications that result from those strategies.

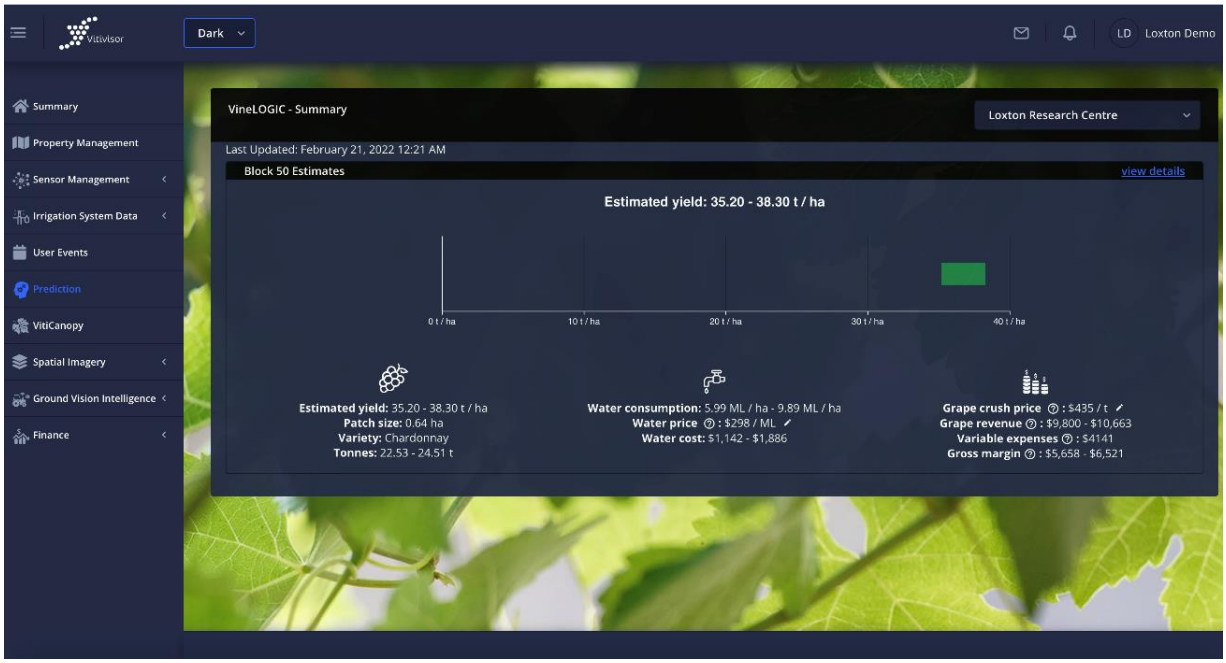


Figure 52: VitiVisor dashboard VineLOGIC summary page

The VineLOGIC details page shows irrigation, soil moisture, LAI and available carbohydrates predictions (Figure 53). Historical soil moisture and LAI data is also shown on the charts as dot points. Users could go into the scenario configurations page by clicking the button on the right top corner of this page, where users could configure “maximum daily irrigation” and “days between irrigation applications” parameters for the scenario.



Figure 53: VitiVisor dashboard VineLOGIC details page

5.5 ADOPTION AND COMMERCIALISATION

As discussed in Section 4.6, there has been significant investment by the partnership in the exploration of various adoption and commercialisation models for VitiVisor. The partnership between Riverland Wine, Wine Australia and University of Adelaide has the unified position that the

primary objective is to accelerate adoption of VitiVisor by growers, and that there are likely to be multiple alternative pathways to achieving this outcome.

From a broad range of interactions over the project lifetime, three primary options have emerged: 1) commercialisation through IP transfer; 2) commercialisation through establishment of a social enterprise; and 3) an industry-owned model. A high-level overview of each is provided below. At the time of writing, no single model has yet emerged as a preferred model, with on-going discussions continuing to refine various options. Additional work is required to further develop these models and identify a preferred direction, as discussed in Section 7 (Recommendations).

5.5.1 Commercialisation through IP transfer

The traditional approach to commercialisation involves transfer of intellectual property from the VitiVisor partnership to one or several third-party commercial entities, who would take on full responsibility for further product development, pricing and servicing models, and associated initiatives that drive adoption outcomes. This approach was investigated in the latter half of 2020 and revisited in the first half of 2022 after discussions with Riverland Wine about establishing a social enterprise were placed on hold (see Section 5.5.2). In total, presentations were delivered to thirteen companies (see Appendix 2), and discussions are ongoing with several of these companies about possible commercialisation options.

Feedback from industry has been that, given established agricultural technology companies already have a digital platform, it is likely that IP transfer will occur in the form of harvesting specific algorithms from VitiVisor, rather making VitiVisor available to growers as an integrated product. For example, multiple companies expressed interest both in the VineLOGIC vine simulation model, and in the computer vision algorithms for monitoring vine development. Interestingly, no companies raised the potential of an open-source IP position as being a significant barrier, mitigating one of the primary risks to achieving adoption.

Table 14: Advantages, disadvantages, risks and next steps associated with the traditional commercialisation option

Advantages • Existing path to market • Capacity to raise capital from capital markets to support continued product development • Likely to be the fastest approach to adoption if accompanied by adequate capital and with an established customer base	**Risks** • Adoption would depend on success of commercial entity in creating scale and thus achieving widespread adoption • IP 'lock-in' ordinarily would be a risk to traditional commercialisation approaches, but is mitigated through VitiVisor's open-source IP position
Disadvantages • Loss of control of subsequent development and pricing decisions. • Profits unlikely to flow directly to growers (although growers would potentially benefit from value provided by the commercial entity). • May be difficult to ensure alignment with broader sector strategies around digital literacy and extension.	**Possible next steps** • Commit resources to continue and accelerate engagements with private markets and identify business opportunities • Demonstrate a sufficient market for the capabilities embedded within VitiVisor through formal market testing • Identify funding vehicles (e.g. commercialisation grants) to support and reduce commercial risk of

• RDCs and grower groups may not be constrained in preferencing a single commercial entity in achieving scale. • VitiVisor likely to be 'carved up' as third-party technology companies generally already have an existing platform and may only be interested in specific areas of IP.	integration of VitiVisor IP into existing technology stacks

5.5.2 Commercialisation through establishment of a Social Enterprise

The social enterprise model centred on the establishment of a new commercial entity that would be governed by—and for the interests of—its members. Foundational work for the development of the social enterprise model was outsourced to Ethical Fields and Collaborative Futures, with further details provided in Section 4.6.

Ethical Fields and Collaborative Futures conducted 28 separate one-on-one interviews with growers in the Riverland, with a number of key tensions and challenges identified and summarised in their report (Ethical Fields and Collaborative Futures, 2021). Through this process, interviewees identified five key desirable outcomes that they hoped could be achieved through the business (Figure 54). The consultants also presented growers with overseas case studies of similar initiatives (Figure 55), including JoinData (The Netherlands; https://join-data.nl/en/) and Agrigate (New Zealand; https://agrigate.co.nz/).



Figure 54: Desired outcomes identified by growers in interviews conducted by Ethical Fields and Collaborative Futures (2021)

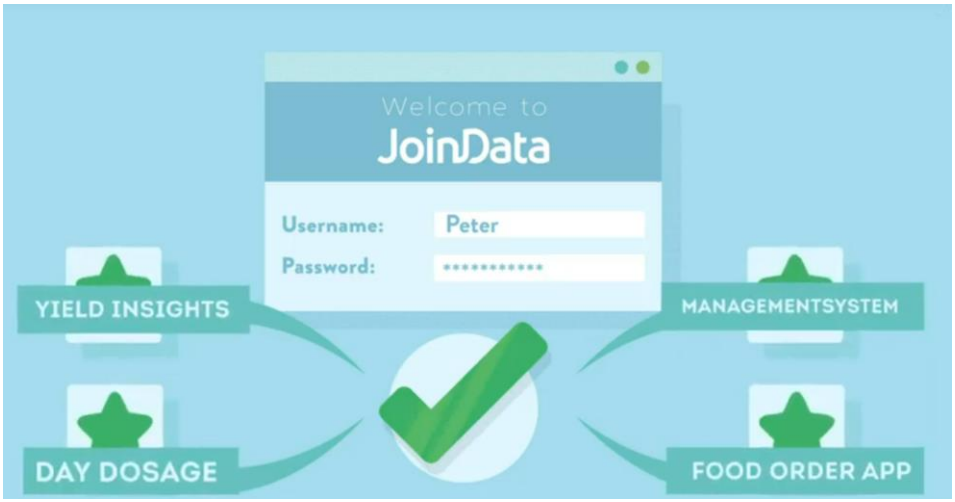


Figure 55: JoinData (https://join-data.nl/en/) from The Netherlands (pictured) was one of several case studies of successful social enterprises presented to growers as a potential model

Following the interviews, a workshop was subsequently held on 30th August 2021 and the outcomes of this process suggested overwhelming support for the concept of a social enterprise, Ethical Fields concluding that '*the majority of workshop participants, which included all the growers, supported the advancing of the idea of a social enterprise that would be able to generate community wealth*.' It is noted that this workshop did not cover topics including detailed business model, legal entity (e.g. company, cooperative, mutual), and so forth. Instead, Ethical Fields and Collaborative Futures (2021) identified the following areas of development as next steps:

- **Governance and business development**, including governance structure, legal entity, strategic capacity, funding and revenue streams, financing and resource model
- **Engagement and adoption**, focusing on fostering conditions for adoption, building literacy, confidence and communication to stakeholders
- **Data and technology**, including development approach, identification of 'minimum viable product', and consideration of issues such as data use and data sovereignty.

> "The majority of workshop participants… supported the advancing of the idea of a social enterprise…"
>
> Ethical Fields and Collaborative Futures report (2021)

Working groups to further develop each of these work streams were to have been formed in the workshop scheduled for December 2021, but as discussed in Section 4.6, this process was halted in November 2021 due to the leadership change at Riverland Wine and associated failure to find a local champion to drive the process further.

Table 15: Advantages, disadvantages, risks and next steps associated with the social enterprise model

Advantages	**Risks**
• Regenerative business model that ensures a significant portion of the profits are retained within the region • Control of future research, development, extension and capacity building activities	• Significant leadership commitment required to realise enterprise • Difficulty in attracting private capital and reliance on grants may lead to unacceptably slow development timelines

• Greater control of data ownership and potential for greater comfort by end-users to share private data • Enhanced potential for government grants, given aligned purpose as well as potential alignment between membership and levee payers • Maintain integrity of VitiVisor as a single cohesive unit	
Disadvantages • Requires establishment of new entity (e.g. company, cooperative or mutual) which would take time and effort • Requires skillsets that may not exist within the region, although these can be brought in or outsourced • Private capital is likely to be more difficult to attract, as it may not be possible to own equity in the company (although this depends on nature of the legal entity and associated ownership structure that is established).	**Possible next steps** • Unlikely to be viable unless leadership emerges from a core membership group to further develop the social enterprise model

5.5.3 Industry-owned model

An alternative model that has not been broadly explored is that the technology could be retained by bodies that represent industry interests (e.g. Wine Australia) as enabling 'research infrastructure'. This is aligned with many of the values of the social enterprise model, but without the requirement for growers to stand up a new business. Instead, the industry body would seek to operate and maintain VitiVisor to achieve industry-wide benefits, with several possible areas for further investigation identified below:

- **Reducing the costs and increasing the impact and velocity of research.** VitiVisor could help dramatically lower the cost of research, as well as increase the speed at which research can find commercial outcomes, by allowing future digital technologies (sensors, algorithms, etc) to integrate with existing technologies. This will therefore ensure that the efforts in building the VitiVisor architecture do not need to be replicated in future R&D efforts, whilst ensuring that any research can be demonstrated to a broad set of commercialisation partners. This is enabled by the intellectual property arrangements of VitiVisor, which means that:
    - VitiVisor could serve as infrastructure to the wider research community, and thus not be limited to the current research partnership.
    - By working in an open environment whilst still providing advanced TRL support, VItiVisor could reduce the requirement for public investment to 'pick winners' or distort the market by investing in individual private sector technology providers.
- **Building trust.** VitiVisor could facilitate benchmarking of performance of various commercial algorithms related to computer vision, yield prediction and so forth, through provision of transparent benchmarking tools and datasets. This has the potential for building broader trust in the sector, the lack of which anecdotally has been an important barrier to technology adoption at present.

- **Accelerating extension activities.** As an advisory tool that has significant information on grower context, VitiVisor could embed scientific advances within a context-aware environment, so that the immediacy and relevance of extension activities can be enhanced. For example, by knowledge of a grower's variety and stage in phenological development, it would be possible to provide notifications and insights specifically related to what a grower needs at that point in time. This will help increase velocity and reduce friction associated with current extension efforts.
- **Facilitating reporting requirements by growers and across the value chain**. Regular reporting requirements represents a common friction point but can also become a touch-point with growers through two-way flow of information. Building reporting capability into VitiVisor as a 'lean' or 'minimum viable product' version of the software could be an opportunity to rapidly build the scale required for higher-value activities.

It is noted that this option is not necessarily mutually exclusive to the other commercialisation options, as for example a commercial provider could be commissioned to provide commercial solutions to Wine Australia or other grower group on behalf of a collective of growers.

Table 16: Advantages, disadvantages, risks and next steps associated with a Wine Australia-owned model

Advantages • Control of future research, development, extension and capacity building activities • Capacity to achieve rapid scale through servicing sector-wide needs • Capacity to leverage broader value streams as articulated in the main text • Capacity to use this as a 'digital vineyard' and thus increase the impact and velocity and lower the cost of future research • Maintain integrity of VitiVisor as a single cohesive unit	**Risks** • Potential for market distortion if the focus is not carefully targeted to public-good outcomes.
Disadvantages • May be perceived as competing with private sector operators • Lack of experience in hosting software platforms	**Possible next steps** • Requires contextualisation within a Wine Australia digital strategy

6 OUTCOME/CONCLUSION

This report documented the development of VitiVisor over the period from January 2020 to June 2022, implementing the recommendations of a pilot project summarised in Westra *et al* (2019). The overall idea and call for collaboration was initiated by Riverland Wine in 2018, and the focus of the current work has been to develop the concepts to Technology Readiness Level 6 ('pre-commercial technology demonstration'), as a foundation to generating economic and environmental sustainability outcomes for the Riverland.

The co-development process between researchers and industry represents a unique and fundamental component of the VitiVisor process, with the approach to governance and decision making described in Section 3 of this report. The outcome has been a technology platform that brings together a unique set of capabilities, including:

- A cost-effective ground-based vision system that combines comparatively low-cost camera technology with state-of-the-art computer vision algorithms to track canopy development throughout the growing season;
- A data acquisition and storage system that collects diverse data streams, and thus enables the development of algorithms and insights build on a strong empirical foundation;
- A vine growth and yield prediction tool that updates in real time based on the latest data on weather, soil moisture and canopy development, and allows predictions to be aligned to grower actions such as watering strategies; and
- Financial data collection and data collection capabilities, that enable users to upload and track their farm costs and revenues, simplifying and optimising decision making.

Importantly, each of these components are closely interconnected and have been integrated into a single operating environment. This integration enables insights from diverse data streams to be directed to achieve grower value through the provision of information to growers associated with production costs and enhancing profitability. Moreover, by making the technology open source, it is hoped that elements of the VitiVisor technology base can be deployed to achieve sector-wide benefits through supporting the broader ag-tech industry.

A benefit-cost analysis of the adoption of VitiVisor was undertaken, assuming a range of fixed and variable costs over a five year horizon, and benefits over the same horizon including a 10% reduction in electricity, fuel and labour cost and an estimated 3t/ha increase in yield. The expected net present value of adoption was conservatively estimated as $1,219 per hectare, or $69,749 per average farm enterprise. This led to a benefit cost ratio of 1.5 (indicating an expected net return on investment of $1.50 for every dollar of investment), with a sensitivity analysis suggesting the ratio could vary between 0.5 and 2.9 depending on a range of assumptions. The outcome of the analysis suggested that VitiVisor would yield positive net returns provided the overall cost of purchasing and using the tool (including accounting for grower's time) over a five year period was no greater than $67,198.

Although outside of the formal project scope, significant in-kind efforts were also made by the partnership to identify commercial models that aligned with Riverland Wine's requirements of open source intellectual property and member involvement. This includes extensive consultation with 16 technology providers, as well as the development of a possible alternative 'social enterprise' model that would enable growers to have direct control over the future development of VitiVisor. At the time of writing this report, significant progress has been made on identifying a potential commercial model, although no firm decision has been made and several options continue to be investigated.

Whilst all project objectives have been achieved, it is clear that the next phase will be both the most important and arguably amongst the most difficult: taking the technology from pre-commercial to commercial, achieving impact at scale, and doing so in a manner that directly benefits end users both immediately and over the longer term. We finish this report by describing a series of recommendations that collectively might facilitate the process of achieving this broader outcome.

7 RECOMMENDATIONS

VitiVisor has been built on the idea that if Riverland growers are to navigate the forces of digital disruption, any strategic decisions must be based on the principles of co-design, local empowerment and capacity building. Each of the recommendations below therefore have been designed to ensure that:

- the $5.5m of investment into VitiVisor by the partnership achieves direct and indirect benefits for growers and the broader wine-grape value chain in the Riverland and beyond; and
- any adoption and commercialisation pathways are consistent with the principles of co-design, local empowerment and digital capacity building.

Recommendation 1. That as a priority, a dedicated resource be identified and invested in to support the realisation of one or several preferred adoption and commercialisation pathways, to ensure the significant investments into VitiVisor by the partnership achieve impact and scale.

Commercialisation efforts were outside of the original scope of VitiVisor, and have been resourced largely through additional in-kind efforts of the project partners. Moreover, in addition to resource constraints, efforts at commercialisation were affected by:

- a significant shift in strategic direction away from Riverland Wine's preferred model of a social enterprise model in late 2021; and
- difficulties in showcasing tangible capabilities during early commercialisation discussions, with a working version of the dashboard only available for broader presentation in the final months of the project.

Therefore, it is highly likely that there are significant untapped and underexplored opportunities for commercialisation and adoption.

As part of this approach, it is recommended that any adoption and commercialisation strategy take an impact-centred approach, recognising that different commercialisation strategies are likely to impact on growers in different ways—particularly over the medium- to long-term. As highlighted in Section 1.3, it is not a given that commercialisation of digital technology will automatically yield benefits to growers or regional communities. Therefore, a deliberate approach by growers and their representative bodies (e.g. Riverland Wine, Wine Australia) will be required to ensure benefits are appropriately shared.

In considering this option, it should be noted that considerable investment is likely to be required to shift VitiVisor from TRL 6 through to TRL 7 or 8, with this shift representing a 'grey area' in terms of alignment with private capital. In particular, it is widely recognised that it can be difficult to attract private equity in the 'valley of death' generally defined as the zone between TRL 4 and 7 (a zone in which VitiVisor still operates). As such, it is likely that investment in business development activities will only achieve desired outcomes when combined with other strategies described herein, and a focus exclusively on business development to the exclusion of other strategies is likely to put the significant investment already placed into VitiVisor at risk.

Recommendation 2. That efforts at commercialisation and adoption of VitiVisor be placed within a broader sectoral digital strategy that takes an expansive view of the value of digital technologies to growers and across the wine value chain.

A sectoral digital strategy would provide broader context related to desired outcomes and the role of different actors across the digital value chain in securing these outcomes. Possible themes that could be included in such a digital strategy include:

- value creation and the different sources of value;
- accelerating digital literacy in regional communities;
- digital opportunities to enhance research extension experiences;
- equity and value retention within regional communities;
- issues associated with regulation and generation of trust, including product benchmarking as well as issues of data integrity and security;
- design of vibrant R&D ecosystems including the potential role of digital infrastructure as research enabler; and
- the role of grower bodies and funding agencies to support a favourable environment for the adoption of digital technologies.

> Strategic foresighting does not '*aim to predict the future but rather to make it possible to imagine multiple futures in creative ways to that heighten our ability to sense, shape and adapt to what happens in the years ahead*.'
>
> Peter Scoblic

Given the high levels of uncertainty and rapid pace of change associated with the digital transformation (see Section 1.2), it is likely that the strategy would benefit from the institutionalisation of strategic foresighting capabilities within the sector (e.g. Scoblic, 2020), thereby creating a clear and transparent framework to support decisions related to commercialisation and adoption, and provide direction for future investment priorities.

Recommendation 3. Recognising that lack of digital literacy and low rates of adoption are pervasive issues in the field of digital agricultural technology, any adoption strategy should better integrate scientific and technological research with strategies associated with extension and training.

Lack of digital literacy and low technological adoption rates are frequently raised as significant barriers that limit the capacity of the sector to realise value from digital technologies (e.g. Leonard et al, 2017). As such, any research and development in digital technologies such as those embedded within VitiVisor will likely need to be more closely integrated with investments associated with training and extension, rather than consider these issues in isolation.

Importantly, the scientific rigour associated with the VitiVisor platform (much of which has, or will be, subject to the international peer review process), combined with the open source nature of the tool, provide significant advantages relative to proprietary systems in terms of using the platform for training purposes. Thus, with the right investment VitiVisor may have unique competitive advantages as a pedagogical tool to support enhancement of digital literacy in the viticulture sector.

Recommendation 4. Reinvigorate the co-design process through the adoption of a lean approach to technology development, with a specific focus on enabling the transition from TRL 6 to 7.

Experience from adoption and commercialisation efforts during the VitiVisor project suggests that additional investment is required make VitiVisor attractive to private investors. It is therefore recommended that a lean approach (Ries, 2011) is supported in order to explore potential commercialisation avenues through the creation of accelerated feedback loops that build on what has already been created. It is suggested that investment comprise the following elements:

- Rapidly deploying VitiVisor in several commercial vineyards (and/or a lean 'VitiVisor Lite' or 'Freemium' version deployed to a larger number of growers), using vineyard partners who are willing and able to trial integrating insights from VitiVisor into their business operations;
- Leveraging the VitiVisor dashboard and associated technologies that have already been developed, but deploying these in multiple different ways through a 'lean' process with an accelerated design cycle;
- Incorporating one or several commercial partners (e.g. technology companies) into the co-design process, whilst recognising that the current TRL of VitiVisor will require some level of de-risking; and
- Consideration (and rapidly testing) of multiple potential sources of value beyond the emphasis on production costs considered in this current project phase, including areas such as:
    - objective grape quality measures;
    - access to sustainability markets (e.g. carbon markets);
    - minimising regulatory burden (e.g. spray diaries, Sustainable Wine Australia certification); and
    - logistics including real-time prediction of harvest dates and better integration between vineyard and winery.

Given the software architecture and associated capabilities that were developed in this iteration of VitiVisor, many of the above sources of value would be comparatively quick to implement. For example, berry quality predictions could be informed by a berry quality module within VineLOGIC that exists but has not been activated or sufficiently tested. Similarly, much of the data for sustainability certification is already embedded within VitiVisor. Each of these would require more detailed scoping, but the VitiVisor infrastructure is specifically designed to shorten design cycles, and thus the marginal costs of implementing any of these measures would be substantially reduced compared to other means of funding the same endeavour.

Recommendation 5. Consider broader sector-wide benefits of VitiVisor by recognising the tool's potential to serve as research infrastructure

VitiVisor is a digital platform that has the potential to serve as digital infrastructure for the broader grape and wine sector. Thus, whilst Recommendations 1 and 4 focus largely on using commercialisation as a means of obtaining direct benefit for growers, it is likely that a number of indirect benefits can be realised when aligned to a more strategic approach to digital technologies for the sector (Recommendation 2). Examples include:

- **Building trust.** VitiVisor could facilitate benchmarking of performance of various commercial algorithms related to computer vision, yield prediction and so forth, through provision of transparent benchmarking tools and datasets. This has the potential for building broader trust in the sector, which anecdotally is an important barrier to technology adoption at present.
- **Reducing the costs of research.** VitiVisor could help dramatically lower the cost of research, by allowing future digital technologies (sensors, algorithms, etc) to integrate with existing technologies and thus ensure that the efforts in building the VitiVisor architecture do not need to be replicated in future R&D efforts. This is enabled by the unique intellectual property arrangements of VitiVisor, which means that:
    - VitiVisor could serve as infrastructure to the wider research community, and thus not be limited to the current research partnership.

    - By working in an open environment whilst still providing advanced TRL support, VItiVisor could reduce the requirement for public investment to 'pick winners' or distort the market by investing in individual private sector technology providers.
- **Accelerating extension activities.** As an advisory tool that has significant information on grower context, VitiVisor could embed scientific advances within a context-aware environment, so that relevance of extension activities can be enhanced. For example, by knowledge of a grower's variety and stage in phenological development (functionality that already exists within VitiVisor), it would be possible to provide notifications and insights specifically related to what a grower needs at that point in time. Development of this option would need further co-design with extension providers, but may potentially assist in increasing velocity and efficacy of extension efforts.

8 REFERENCES

ABS DataLab, 2020, Business Longitudinal Analysis Data Environment (BLADE), Australian Bureau of Statistics, Canberra

AgFunder, 2021, AgFunder Farm Tech Investment Report, https://agfunder.com/research/2021-farm-tech-investment-report/ 54pp

Alzraiee, A.H., White, J.T., Knowling, M.J., Hunt, R.J., Fienen, M.N., 2022, A scalable model-independent iterative data assimilation tool for sequential and batch estimation of high dimensional model parameters and states, *Environmental Modelling and Software*, 150, 105284.

CRC for Optimising Resource Extraction https://www.crcore.org.au/index.php/technology-readiness-levels

De Bei, R., Fuentes, S., Gilliham, M., Tyerman, S., Edwards, E., Bianchini, N., ... & Collins, C. (2016). VitiCanopy: A free computer App to estimate canopy vigor and porosity for grapevine. Sensors, 16(4), 585. https://doi.org/10.3390/s16040585

Bjorck, N., Gomes C., Selman, B. &* Weinberger, K.Q., 2018, Understanding batch normalization, Advances in neural information processing systems, 31.

Ethical Fields and Collaborative Futures, 2021, Preliminary exploration of a social enterprise approach for adoption and commercialisation of the VitiVisor platform, 10 September 2021, 35pp.

Herrendorf, B., Rogerson, R. & Valentinyi, A., 2014, Growth and Structural Transformation, *Handbook of Economic Growth*, Volume 2B.

Isoraite, M 2004, 'Theoretical Aspects of Benchmarking Theory', *Viesoji Politika ir Administravimas*, 9.

LeCun, Y., Bengio, Y. & Hinton, G., 2015, Deep Learning, *Nature*, 521, 436-444.

Leonard, E. (Ed), Rainbow, R. (Ed), Trindall, J. (Ed), Baker, I., Barry, S., Darragh, L., Darnell, R., George, A., Heath, R., Jakku, E., Laurie, A., Lamb, D., Llewellyn, R., Perrett, E., Sanderson, J., Skinner, A., Stollery, T., Wiseman, L., Wood, G. and Zhang, A. (2017). Accelerating precision agriculture to decision agriculture: Enabling digital agriculture in Australia. Cotton Research and Development Corporation, Australia.

McKinsey & Company, 2017, Digital Australia: Seizing the opportunity for the Fourth Industrial Revolution, 164pp.

Regan, A., 2019. 'Smart farming' in Ireland: A risk perception study with key governance actors. NJAS – Wageningen Journal of Life Sciences 90-91, 100292. https://doi.org/10.1016/j.njas.2019.02.003

Ries, E., 2011 The Lean Start-up: How Constant Innovation Creates Radically Successful Businesses, Penguin publishers, 336pp.

SA Auditor General, 2020, Statutory Property Valuations - specialised data request, South Australian Auditor General's Office, Adelaide

Scoblic, P.J., 2020, Decision Making and Problem Solving: Learning from the Future, Harvard Business Review (Jul-Aug 2020).

Tong, Kang, Yiquan Wu, and Fei Zhou. "Recent advances in small object detection based on deep learning: A review." *Image and Vision Computing,* 97 (2020): 103910.

US Department of Energy, 2011, Technology Readiness Assessment Guide, downloaded from https://www2.lbl.gov/DIR/assets/docs/TRL%20guide.pdf (1 May 2019)

US Department of Agriculture, 2016, Crop Research Technology Readiness Levels, accessed from https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=1&ved=2ahUKEwi5476G8_nhAhVDfisKHZaECaEQFjAAegQIAhAC&url=https%3A%2F%2Fnifa.usda.gov%2Fsites%2Fdefault%2Ffiles%2Fresources%2FCrop%2520Research%2520Technology%2520Readiness%2520Level.docx&usg=AOvVaw2Rvr0qzIFPmoXNW08fp-zj accessed 1 May 2019

Wine Australia, 1999 - 2020, *South Australian Winegrape Crush Survey*, Vine Health Australia, Adelaide.

Wheeler, S., Seidl, C., Zuo, A., Tingey-Holyoak, J., Xu, Y. & Kandulu, J., 2022, The economics and financial benchmarking of Riverland Grape production, and potential benefits of VitiVisor technology, report prepared for Wine Australia, 108pp.

Zscheischler, J., Brunsch, R., Rogga, S., Scholz, R.W., 2022. Perceived risks and vulnerabilities of employing digitalization and digital data in agriculture – Socially robust orientations from a transdisciplinary process. Journal of Cleaner Production 358, 132034. https://doi.org/10.1016/j.jclepro.2022.132034

Zhao, Z-Q., Zheng, P., Xu, S-T & Wu, X., 2018, Object detection with deep learning: A review, *IEEE transactions on neural networks and learning systems,* 3212-3232.

Zheng, X., Ji, R., Sun, X., Wu, Y., Huang, F. & Yang, Y., 2018 "Centralized Ranking Loss with Weakly Supervised Localization for Fine-Grained Object Retrieval." International Joint Conference on Artificial Intelligence

APPENDIX 1: TECHNOLOGY READINESS LEVELS

The concept of technology readiness levels was first developed by NASA in the 1970s, an adapted version of the TRL scale is provided in Table 17. This framework is very useful both to describe the level of maturity of technology, and also to gain insight into key activities that are required to transition from one level to the next. This framework has been used extensively in this report as a theoretical basis for key investment decisions related to the development of VitiVisor.

Table 17: Description of technology readiness levels. Adapted from US Department of Energy (2011), the CRC for Optimising Resource Extraction, and the US Department of Agriculture (2016).

Relative Level of Technology Development	**Technology Readiness Level**	**TRL definition**	**Description**
Commercial Deployment	TRL 9	Actual system proven reliable through operation	The technology is in its final form and operated under the full range of operating (environmental) conditions). Final bugs are fixed and the technology is routinely implemented.
System Commissioning	TRL 8	Commercially relevant system deployed on end-user site with proven value proposition	The technology has been proven to work in is final form and under expected conditions. This TRL usually represents the end of system development, and can include integration of new technology into an existing system. Outcome is often the commercial manufacture and site uptake of the technology.
	TRL 7	Full-scale prototypical system demonstrated on end-user site	This represents a major step up from TRL 6, requiring the demonstration of an actual system prototype at an end-user site. The prototype should be near or at scale of planned operations, and have commercial partners involved.
Technology Demonstration (Pre-Commercial)	TRL 6	System model tested on end user site with refinement of positive value proposition	Engineering-scale models or prototypes are tested in an end-user site environment. Potential commercial partners will have been identified or already engaged. TRL 6 begins true engineering development of the technology as an operational system. The prototype should be capable of performing all the functions that will be required of the operational system.
Technology Development (Pre-Commercial)	TRL 5	System model tested in simulated or realistic environment	The basic technology components are integrated so that the system configuration is similar to the final application in almost all respects. The major difference between TRL 4 and 5 is the increase in the fidelity of the system and environment to the actual application. The system tested is almost prototypical.
	TRL 4	Component and/or system validation in	The basic technological components are integrated to establish that the pieces will work together. This is relatively 'low fidelity' compared

		laboratory environment, and initial value proposition	with the eventual system. TRL 4-6 represents the bridge from scientific research to engineering. TRL 4 is the first step in determining whether the individual components will work together as a system. The system will probably be a mix of on-hand equipment and a few special purpose components that may require special handling, calibration, or alignment to get them to function.
Experimental testing	TRL 3	Analytical and experimental critical function and/or characteristic 'proof of concept'	Active research and development is initiated. This includes analytical studies and laboratory scale studies to physically validate the analytical predictions of separate elements of the technology. At TRL 3 work has moved beyond the paper phase to 'proof-of-concept' experimental work that verifies that the concept works as expected. Components of the technology are validated, but there is no attempt to integrate the components into a complete system.
Preliminary Technology Solution Evaluation	TRL 2	Technology concept and/or application formulated	Once basic physical principles are observed, the next level involves identification or invention of practical application of those principles. This can include estimation of the value of the technology relative to existing technologies.
	TRL 1	Basic principles observed and reported	This is the lowest level of technology readiness. Scientific research begins to be translated into applied R&D. Examples include identification of industry challenges the scientific research is capable of addressing, the potential for the science to modify industry practice. Supporting information includes published research or other references that identify the basic principles that underlie the technology.

APPENDIX 2: COMMUNICATION

An extensive communication strategy was adopted throughout the project, with separate efforts directed to (1) the primary stakeholder group, comprising winegrape growers of the Riverland; (2) the agricultural technology industry; (3) the scientific community; and (4) the general public. Each of these are discussed in turn below.

COMMUNICATION WITH THE PROJECT'S PRIMARY STAKEHOLDERS

Communication with stakeholders and industry occurred throughout the project through a combination of face-to-face meetings, workshops and industry events. Numerous informal meets and conversations with growers occurred throughout the project, and there were also significant touch-points with the finance and economics team through interviews with 25 growers. In addition to these interactions, Table 18 lists the key communication events held throughout the project lifecycle from January 2020 to June 2022.

Table 18: Key stakeholder communication events throughout the VitiVisor project

Stakeholder Communication	**When**	**Attendees**	**Purpose**
Grower workshop and Engagement Strategy	25th June 2020 Loxton Research Centre	Riverland Wine Thinking Ten (T10) growers VitiVisor Output Leaders PIRSA	To reconnect T10 with the project from their previous involvement in the pilot project To establish the Grower Engagement strategy, and to establish project support sites across the Riverland
Grower Advisory Group	Monthly Zoom meeting	Key Riverland wine grape growers	To consult with growers on the ongoing development of the VitiVisor Dashboard
Riverland Wine Management Committee (MC)	Quarterly committee meetings	Project Manager and RW MC	To provide an update directly to the MC and facilitate deep discussion on the projects progress.
Hands Off Ha Showcase	24th June 2021 Loxton Community Theatre Loxton High School	Riverland Wine members Wine Australia Australian Wine Research Institute PIRSA	This showcase presented the project outputs to the key stakeholders and broader community including students studying agriculture and technology.
Social Enterprise	30th August Loxton Hotel	Riverland Wine members Wine Australia Australian Wine Research Institute	A workshop to explore a social enterprise approach to the adoption and commercialisation of the Vitivisor Platform.

		PIRSA	
Advance Ag	3rd November 2020 Loxton research Centre	18 demonstrators from the SA AgTech ecosystem attended by 55 people from the Riverland region.	The Vitivisor team provided information on the project through both a demonstration stand and in vineyard talk.
Project Closure Presentation	9th June 2022	Riverland Wine Thinking Ten (T10) growers VitiVisor Output Leaders PIRSA Wine Australia CCW Co-op.	The VitiVisor team provided presentation that connected the projects objectives with the research outcomes, highlighting the key achievements.
18th Australian Wine Industry Technical conference	27th to 29th June 2022	Wine industry professionals, including both developers and users of technology.	To provide broader profile for the project and discuss the intellectual property that has been developed with potential industry partners.

COMMUNICATION WITH INDUSTRY

Communication with industry has been facilitated by the VitiVisor project manager through both formal and informal channels:

- Regular updates to the Riverland region were provided through their website, socials and newsletter to their stakeholders
- Attendance at the quarterly South Australian AgTech meetup
- Active participation in the Open Standards for Vineyard mapping (Collabriculture)

Over the project lifecycle, meetings were also held with representatives from:

- Almond Board of Australia
- Horticultural Innovation Australia
- Accolade and CCW Pty Ltd
- Central Irrigation Trust (CIT)

Finally, as part of the Adoption and Commercialisation Framework, the VitiVisor project team frequently engaged with a range of potential commercial partners by demonstrating the features and functionality of VitiVisor.

- The table below lists the organisations that the VitiVisor team has engaged with.

Table 19: Technology organisations engaged with throughout the project lifecycle

Organisation	**Key Contact**
Swan Systems	John Pargeter
Deep Planet	David Carter
Consilium	Sebastien Wong
MEA	Dominic Skinner
PlatFarm	Oli Madgett
BitWise Agronomy	Fiona Turner
Sentek	Peter Buss

Green Atlas	James Underwood
Nio Labs	Doug Standley
Amazon Web Services	Rob Absalom
Australian Wine Research Institute (AWRI)	Tony Robinson
Sage Automation	Antonio Pisctelli
Axel Space Corp - Japan	Yasu Yamazaki
TerraView	Prateek Srivastava
Ceres Imagery	Scott Gillett
Eartrumpet Consulting	Dan Fischl
Opensensing	Leo Gaggl
Sensand	Davi La Ferla

The project team also supported the Vineyard Open Standards project led by Oli Madgett from PlatFarm, by participating in four workshops throughout 2021, and by adopting the open standards with the VitiVisor dashboard software architecture.

COMMUNICATION WITH THE SCIENTIFIC COMMUNITY

The applied 'mid-TRL' nature of the VitiVisor project meant that scientific publications were not a primary emphasis of the project team. Nevertheless, a number of publications in high-quality scientific journals was produced during the project, with articles that have been either published or submitted for peer review summarised in Table 20; several other publications are in the process of being completed, but have not been listed here.

Table 20: Scientific publications generated during the VitiVisor project

Title	**Authorship**	**Journal or other publication vehicle**	**Status (as of 30 June 2022)**
Bridging the gap between data and decisions in viticulture: A review of process-based numerical models	**Matthew Knowling**, Bree Bennett, Bertram Ostendorf, Seth Westra, Rob Walker, Anne Pellegrino, Everard Edwards, Cassandra Collins, Vinay Pagay, Dylan Grigg	Agricultural Systems (IF = 5.4)	Published
Generalised water production relations through process-based Modelling: A viticulture example	**Matthew Knowling**, Rob Walker, Anne Pellegrino, Everard Edwards, Cassandra Collins, Seth Westra, Bertram Ostendorf, Bree Bennett	Agricultural Water Management (IF = 4.5)	Submitted for peer review
Operationalising crop model data assimilation for improved on-farm situational awareness and decision making	**Matthew Knowling**, Jeremy White, Dylan Grigg, Cassandra Collins, Seth Westra, Rob Walker, Anne Pellegrino, Bertram Ostendorf, Bree Bennett, Ayman Alzariee	Agricultural and Forest Meteorology (IF = 5.7)	Submitted for peer review
Overcoming small data sets in viticulture: Case Study for canopy growth assessments	**Zhen Zhang**, Javen Shi, Dylan Grigg, Cassandra Collins		Submitted for peer review
Decision-making and resilience in agriculture: improving awareness of the role of accounting	**Joanne Tingey-Holyoak,** Sarah Wheeler, Constantin Seidl	Meditari Accountancy Research	Submitted for peer review

COMMUNICATION WITH THE GENERAL PUBLIC

Broader communications have occurred throughout the project, both to the general public and outlets to wine industry professionals. This has been assisted by Vitivisor being featured in a range of media releases from eNewsletters to a technology segment featured on the BBC. Key media coverage is summarised in Table 21.

Table 21: Media coverage during the VitiVisor project

Date	Organisation	Media Link
15/01/2020	Wine Australia	https://www.wineaustralia.com/news/media-releases/technology-to-transform-vineyard-management
01/04/2020	Food and Drink Business	https://www.foodanddrinkbusiness.com.au/news/wine-industry-collab-develops-open-source-platform
9/07/2021	Wine Australia	https://www.wineaustralia.com/news/articles/vitivisor-showcases-the-digital-future
11/11/2021	The Australian	https://www.theaustralian.com.au/special-reports/future-adelaide/wine-benefits-from-hitech-skills/news-story/f030a9f636130080ea2e33463a71d9a1
18/11/2020	BBC	https://www.bbc.co.uk/news/technology-57850284
27/05/2022	Winetitles Media	https://winetitles.com.au/smart-technology-for-vineyards/
22/06/2022	Winetitles Media	https://winetitles.com.au/new-canopy-surveillance-systems-are-set-to-ring-in-the-age-of-precision-viticulture/
23/6/2022	Winetitles Media	https://winetitles.com.au/viticulturists-one-step-closer-to-advanced-computer-vision/
29/06/2022	Environment Institute, University of Adelaide	https://blogs.adelaide.edu.au/environment/2022/06/29/new-canopy-surveillance-systems-are-set-to-ring-in-the-age-of-precision-viticulture/
29/06/2022	NASAA Organic Insights	https://cld.bz/4FWhga
29/06/2022	Murray Pioneer, Riverland	“5M vineyard project’s finale” “Better canopy data poised to give growers the edge”

APPENDIX 3: INTELLECTUAL PROPERTY

The intellectual property both used in the project (Background IP) and produced during the project (Project IP) is summarised in the tables below. Where possible, Project IP has been made available on the VitiVisor website through an open-source license, and can be accessed at www.VitiVisor.com.au.

Background IP

Area	Item	Description	License	Limitation/ restrictions	Access
Prediction and Modelling	VineLOGIC	VineLOGIC (bio-physical model) source code	GNU General Public License version 3.[6]	As per license	Persistent link
		VineLOGIC associated datasets	CC BY 4.0[7]	As per license	Persistent link
		VineLOGIC View (web app to interact with model)	CC BY 4.0	As per license	Persistent link
		PEST++	CC0 1.0 Universal Public Domain Dedication.[8]	As per license	Link
Ground-based Vision System	Annotation tool for detection.	Used to label the bounding boxes of the inflorescences for later training the models.	MIT License.[9]	As per license	Link
	Deep Learning tool for object detection	Used to train on the annotated data and then detect inflorescence / shoot tips/ buds in	GNU General Public License version 3.	As per license	Link

[6] This license is a free, copyleft license for software and other kinds of works. Here, "free" refers to freedom with respect to rights associated with the software not price.
[7] This license means that a user is "free" to share (copy and redistribute in any medium or format) as well as adapt (remix, transform, and build upon for any purpose, even commercially).
[8] This license means that one can copy, modify, distribute and perform the work, even for commercial purposes, all without asking permission.
[9] This license means that one can copy, modify, distribute and perform the work, even for commercial purposes, all without asking permission.

		images in this project.			
	RTK high precision GPS tool (~4cm precision)	Used to accurately localize the point of interest in the vineyard			TBA
	Sequential modelling tool (LSTM/GRU).	Used for vine growth modelling.	BSD License 10	As per license	Link
	Sequential modelling tool (BERT).	Used for vine growth modelling.	Apach License 2.011	As per license	Link

Project Output IP

Area	**Item**	**Description**	**License**	**Limitation/ restrictions**	**Access**[12]
Prediction and Modelling	VineLOGIC Wrapper	Integration with VitiVisor dashboard	Free to access	As per Bit Bucket terms and conditions	Bit Bucket
Viticulture	Vitidata	Viticultural data	Free to access	As per Figshare terms and conditions	Figshare
Ground-based Vision System	Machine Learning algorithms				GitHub
	VitiBox		BSD 3-Clause"New" or "Revised" License	As per Hackster terms and conditions	Hackster
VitiVisor Platform	Dashboard	Allows visualisation	TBA	TBA	TBA

[10] This license means that one can copy, modify, distribute and perform the work, even for commercial purposes, all without asking permission.
[11] This license means that one can copy, modify, distribute and perform the work, even for commercial purposes, all without asking permission. However, the patent (possibily) in the code are not automatically granted to the user.
[12] The specific links for project IP are provided in the VitiVisor site (www.VitiVisor.com.au).

APPENDIX 4: PROJECT TEAM

A collaborative team was formed and maintained throughout the project lifecycle, and comprised representatives from the University of Adelaide, Riverland Wine and Wine Australia. A summary of key team members is provided in the tables below.

Client	Name	Role
Wine Australia	Dr Paul Smith	(Grantor)
	Mr Alex Sas	Representative

Customer	Name	Role
Riverland Wine	Ms Lyndall Rowe	Executive Officer
	Mr Brett Rosenzweig	Grower representative
	Mr Ian McRae	Grower representative (CCW)
	Mr Chris Byrne	Former Executive Chair

Grower Advisory Group (Riverland Wine members)	Name
	Mr Ben Haslett
	Mr Chris Byrne
	Mr Brett Proud
	Mr Andrew Proud
	Mr Brian Caddy
	Mr Dave Zadow
	Ms Sheridan Alm
	Mr Dave Zadow
	Mr Brett Rosenzweig

Project Management	Name	Role
UoA	Prof Andrew Lowe	Project Director
	Prof Seth Westra	Chief Investigator
Poignant Pty Ltd	Mr Brad Nott	Project Manager

Project Delivery Team

Output	Name	Role	University of Adelaide
Data management	Prof Javen Shi	Output Lead	Australian Institute for Machine Learning, University of Adelaide
	Mr Hao Zhang	Systems Engineer	AMI Fusion Pty Ltd
	Mr Peter Zhou	Software Engineer	AMI Fusion Pty Ltd
Ground-based vision systems and Irrigation Infrastructure management	Dr Tien Fu Lu	Output Lead	School of Mechanical Engineering, Faculty of Sciences, Engineering and Technology

	Dr Shi Zao	Post Doc Researcher	School of Mechanical Engineering, Faculty of Sciences, Engineering and Technology
	Mr Ryan Kuan Tan	Researcher / Engineer	School of Mechanical Engineering, Faculty of Sciences, Engineering and Technology
Prediction and Modelling	Prof Seth Westra	Chief Investigator	School of Civil, Environmental and Mining Engineering, Faculty of Sciences, Engineering and Technology
	Dr Matthew Knowling	Post Doc Researcher	School of Civil, Environmental and Mining Engineering, Faculty of Sciences, Engineering and Technology
	Dr Bree Bennett	Investigator	School of Civil, Environmental and Mining Engineering, Faculty of Sciences, Engineering and Technology
	A/Prof Bertram Ostendorf	Investigator	School of Biological Sciences, Faculty of Sciences, Engineering and Technology
Vision Systems (Machine Learning	Prof Javen Shi	Output Lead	Australian Institute for Machine Learning, University of Adelaide
	Dr Zhen Zhang	Post Doc Researcher	Australian Institute for Machine Learning, University of Adelaide
	Assoc Prof Cassandra Collins	Output Lead	School of Agriculture, Food and Wine, Faculty of Sciences, Engineering and Technology
	Dr Dylan Grigg	Post Doc Researcher	School of Agriculture, Food and Wine, Faculty of Sciences, Engineering and Technology s
	Mr Wayne Kiely	Field Assistant and data collection	School of Agriculture, Food and Wine, Faculty of Sciences, Engineering and Technology

Economic Benefits	Prof Sarah Wheeler	Supervision/oversight; benefits estimation; literature review, benchmarking and analysis; case study analysis	School of Economics and Public Policy, Faculty of Arts, Business, Law and Economics
	Assoc Prof Joanne Tingey-Holyoak	Assistance and support	University of South Australia Business School
	Dr Alec Zuo	Benchmarking and analysis; benefits estimation, survey analysis	Centre for Global Food and Resources, Faculty of Arts, Business, Law and Economics
	Dr Ying Xu	Benchmarking and analysis; benefits estimation, survey analysis	School of Economics and Public Policy, Faculty of Arts, Business, Law and Economics
	Dr Constantin Seidl	Assistance and support; literature review; case study analysis	School of Economics and Public Policy, Faculty of Arts, Business, Law and Economics
Dashboard	Prof Javen Shi	Output Lead	Australian Institute for Machine Learning, University of Adelaide
	Mr Hao Zhang	Systems Engineer	AMI Fusion Pty Ltd
	Mr Peter Zhou	Software Engineer	AMI Fusion Pty Ltd
	Mr Bradley Nott	Grower Engagement	Poignant Pty Ltd

APPENDIX 5: ACKNOWLEDGEMENTS

The research team would like to acknowledge the valuable support from the following people and businesses over the duration of the project:

Riverland Wine

Mr Chris Byrne (formerly, Executive Chair of Riverland Wine), for seeding the original idea and for his unrelenting passion to focus the team and drive the project.

Riverland Wine Grower Advisory Group

To the members of this group that are winegrape growers who share in the vision and continuously volunteer their time to support and contribute to the project's success.

Department of Primary Industries and Regions South Australia (PIRSA)

Mr Peter McGee (Farm Manager, Loxton Research Centre), for providing support and technical know-how to the research team at the trial viticulture patches used in the research project.

School of Agriculture, Food and Wine, University of Adelaide

Dr Vinay Pagay, for using his viticultural knowledge in assisting the project team to scoping the project.

Mr Benjamin Pike (Vineyard Manager at Waite Campus), for providing access and support for the viticultural and IoT team, to the viticultural research vineyard.

Central Irrigation Trust

Mr Greg McCarron, for enabling and supporting the API connection (via Outpost Central) to the VitiVisor dashboard for water delivery data.

Wine Australia

Dr Paul Smith, for his support and guidance through regular, positive feedback to the research team throughout the duration of the project.

South Australian Agtech Community

Mr Oli Madgett, PlatFarm Pty Ltd, for his unwavering support and promotion of VitiVisor to the South Australian Agtech community.

Agtech Companies

- Ceres Imaging Pty Ltd, for their ongoing support providing aerial imaging data for the trial patches at the Loxton Research Centre.
- H_20X, for developing and enabling the API connection to the VitiVisor dashboard for water market and trading data.
- Outpost Central, for enabling the API connection to the VitiVisor dashboard for water delivery data.
- Katalyst, for their upgrade to Know Your Numbers and connecting API's to Xero and VitiVisor.