

EXTRACTION OF VINEYARD MACROSTRUCTURE FROM SUB-OPTIMAL SEQUENCES OF AERIAL IMAGERY

Anthony Finn (corresponding author), Aaron Melville-Smith, Russell Brinkworth

School of Engineering, University of South Australia, Mawson Lakes, South Australia, SA 5095
{anthony.finn, aaron.melville-smith, russell.brinkworth}@unisa.edu.au

KEY WORDS: unmanned aerial vehicle, 3D point cloud, structure from motion, vineyard

ABSTRACT:

Remote sensing techniques can be used to identify and classify vine properties such as row width, height, cover-fraction, missing segments and leaf area density, providing opportunities to visualise plant vigour as a spatial function of vineyard geography. This information may then be integrated into decision support tools to improve vineyard management practices. An algorithm for identifying vines from a sequence of overlapping aerial images and then estimating their properties is described. The image stacks were obtained from visible and long wave infrared cameras carried by an unmanned aerial vehicle (UAV). Structure from motion (SfM) was used to create 3D thermal and coloured point clouds, from which the underlying topography of the surface terrain was extracted. The surface topographic model was obtained using bounded data query nearest neighbour calculations, which were reduced to computationally manageable levels using Kd-trees that recursively partitioned the point clouds by progressively separating them into binary trees. This allowed the point clouds to be classified in terms of their hue, saturation, surface temperature and height relative to surface topography using Lloyd's unsupervised k-means clustering. Individual samples were then associated using Gaussian probability density functions normalised by cluster statistics. The algorithm was evaluated against ground truth obtained using aerial data in terms of its accuracy and robustness using a combination of real world conditions that included high shadowing, poor contrast and UAV flight paths and camera settings that delivered sub-optimal SfM performance.

1. INTRODUCTION

Many factors affect grapevine productivity: climate, weather, soil properties, topography, grape variety, management practices, and pests and diseases, with spatial variations of such factors within and between vineyards impacting both grape quality and yield (Hall et al., 2002). As a result, optimal management practices are highly desirable, with decisions that depend on precise situational awareness of crop state ideally based on information gathered safely, efficiently and non-destructively.

In this context, remote sensing offers considerable opportunities, with visible (VIS), long wave infrared (LWIR) and near infrared (NIR) hyperspectral sensors all providing windows into the complex surface chemistry present in vineyards. Although potentially promising, satellite-based observations are not routinely exploited: imagery from several satellite systems is commercially available, but the spatial resolution of their VIS sensors (5 – 30m) is inadequate for mapping small vineyards (Weiss and Baret, 2017). Moreover, the superposition of signal returns from vines and inter-row material within a single pixel makes extraction of desired signatures difficult. Row structure and topography of underlying terrain must also be accounted for to avoid introduction of effects that depend upon the directionality of observations or signal reflectance (Zarco-Tejada et al., 2005, Meggio et al., 2008, López-Lozano et al., 2009, Holben and Justice, 1980, Shepherd and Dymond, 2003). Higher spatial resolution sensors (5 – 20cm) are therefore essential if the properties of inter-row returns are not to contaminate vine signatures, which are of prime importance to any decision support tool.

Furthermore, satellite orbital characteristics, together with sensor pose, drive temporal availability of such data sets, as does the prevalence of cloud cover.

Alternative, more operationally flexible remote sensing options, such as aircraft and Unmanned Aerial Vehicles (UAVs), are also available; and—relative to their satellite counterparts—the proximity of such platforms to vines intrinsically improves the ground resolution of the sensors carried. Although operation of such platforms is generally weather-dependent, planning flights close to particular stages of plant development is straightforward and several authors have examined mapping crop vigour using such techniques (Johnson et al., 2001, Dobrowski et al., 2003, Smit et al., 2010, Fiorillo et al., 2012). In the past few years increased availability of UAVs has also added new opportunities to acquisition of high spatial resolution imagery. As a result, there have been many studies into the acquisition and exploitation of multi- or hyperspectral data sets to assess vine vigour (Hernández-Clemente et al., 2012, Zarco-Tejada et al., 2013, Rey et al., 2013, Candiago et al., 2015), grape species (Lacar et al., 2001, Ferreiro-Armán et al., 2006), and water and nitrogen stress (Zarco-Tejada et al., 2012, Zarco-Tejada et al., 2013) against vegetation indices, such as the normalised difference vegetation index (Rouse et al., 1974).

A key issue in all such characterisations is isolation of pixels belonging to the vine rows from those belonging to the spaces in between and elsewhere. Difference in spectral properties between these domains were first attempted by applying a threshold to vegetation indices computed from

multi-sensor images (Candiago et al., 2015, Puletti et al., 2014). Other researchers have applied Fourier transforms to the red band of the red-green-blue (RGB) classification of images (Wassenaar et al., 2002, Chanussot et al., 2005, Delenne et al., 2010). Unfortunately, complex lighting conditions, such as poor contrast, shadow, and glare—as well as the need for spectral contrast between the green of vine rows and the presence of grass between rows—has limited the performance of such methods.

To overcome such shortfalls, a photogrammetric technique known as Structure from Motion (SfM) (Longuet-Higgins, 1981, Furukawa and Ponce, 2010, Hartley, 1997) is often used to obtain dense point clouds (DPC) from the overlapping nature of multiple VIS/RGB images obtained by a UAV. This has been applied to vineyards (Mathews and Jensen, 2013) and other crops (Bendig et al., 2013) to obtain 3D models of terrain and extract fractions of vegetation cover; and to estimate vineyard structure (Weiss and Baret, 2017). Unfortunately, while these approaches are able to discriminate between vines and grass between rows, existing algorithms are sensitive to even slight illumination variations common in real world settings. Additionally, in order to obtain good performance, current techniques need flight configurations and camera settings that are optimised (Weiss and Baret, 2017).

In the final analysis, the value proposition for decision support tools based on remote sensing approaches derives not from the capabilities of any given technology or sensor per se, but from a system’s holistic potential to improve efficiencies and reduce fieldwork. In other words, when developing aids to support mapping and early identification of vineyard characteristics like disease, biophysical stress and equipment failures—for practical reasons—successful algorithms should not critically rely on specialist users who are expert in the design or operation of the relevant software, mathematical techniques or manipulation of information sets. Nor should algorithms make use of machine learning approaches that require training on unique or hard-to-obtain data sets. They should intrinsically cope with sub-optimal equipment configurations and reasonable performance degradation in the processing chain. Ideally, users should need to supply only a few intuitive parameters.

This paper offers an algorithm that can be used for estimating the structural characteristics of vineyards from remotely sensed DPC. It requires input of only a single parameter, readily determined by an untrained operator. The algorithm does not require training on specific data sets and is insensitive to complex colour, contrast and lighting variations typically found in the real world. It also copes with noisy and imperfect point clouds generated by sub-optimal UAV flight and camera configurations, as well as the uneven spatial performance of SfM algorithms. It is evaluated under intentionally complex and sub-optimal conditions over a vineyard where green grass is sometimes present between vines. Its performance strongly suggests it

is suitable for application in vineyard decision support tools.

2. METHODOLOGY FOR CHARACTERISING VINEYARD MACROSTRUCTURE

Simple UAV flight paths were executed with aerial observations of the test site made in both the RGB and LWIR components of the electromagnetic spectrum. 3D DPC were generated using SfM and the underlying topography of the surface terrain identified and excised. To enhance computational speed, a Kd-tree (Bentley, 1975) partitioned the DPC by recursively splitting it into binary clusters. This allowed a nearest neighbour search to be quickly executed using a bounded data query (Elseberg et al., 2012), from which the surface topography was statistically determined. This topographic model was then used to normalise the DPC to a reference surface, whereupon Lloyd’s classification algorithm (MacQueen, 1967) was applied to a combination of height, hue, saturation and irradiance (the latter measured by the LWIR sensor). The clusters were associated using Gaussian probability density functions (PDFs). Finally, the macrostructure and spatial properties of the vineyard (vine vs. inter-row, row height, width, cover fraction and missing row segments) were mapped as spatial functions of vineyard geography.

An overview of the algorithm that computes vine structure is shown in Figure 1. As all mathematical techniques involved are well understood, wherever possible, only verbal descriptions of their function and references to source material is provided in this paper. This is simply to ensure concise narrative.

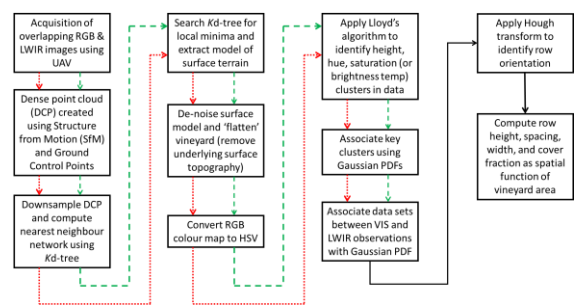


Figure 1: Algorithmic overview of vine macrostructure characterisation algorithm. Red (dotted) and green (dashed) lines indicate processes executed sequentially on both VIS and LWIR imagery. The black (continuous) line indicates processes executed on the joint imagery

Image Acquisition: Field experiments took place at the Jacob’s Creek vineyards in the Barossa Valley, South Australia on 8th February 2018 (Figure 2). Visibility was clear, the sky cloudless and air temperatures ranged from temperate (25°C) at about 9.00am to very hot (> 43°C) around 3pm local time. These conditions provide regions of dark shadowing, as well as zones of high irradiance and directionally dependent illumination and reflectance: a

combination of circumstances that challenge many image processing techniques. Moreover, by judicious selection of observation sets, the thermal conditions were able to offer both a challenging test set for the LWIR data and sound ground truth.

A small (1.4kg), bespoke-manufactured quad-rotor UAV overflew the vineyard at a constant forward horizontal velocity of 10m/s and altitude of 120m in a sequence of opposing parallel lines, oriented along vine row direction (roughly SE-NW). This pattern, colloquially known as a ‘lawn mower’, is shown in Figure 2 as set of continuous white lines. UAV position was determined using standard positioning service GPS (accuracy 3-5m), with altitude maintained using barometric pressure. Separation between each parallel path was about 70m. Images were collected at a rate of 10Hz, but—in order to deliver sub-optimal SfM performance—only one in ten images was used, i.e. an effective update rate of 1Hz. Sub-optimal SfM flight paths and camera settings were preferred because, while several UAV-SfM integrated/automated flight path planners exist, they typically require specialised user knowledge (i.e. skill sets outside those of a vineyard manager) and often result in time-consuming post processing computations; and the overarching aim is to avoid a need for users to become expert in the design of UAV flight paths and SfM optimisation. Thus, the trials demonstrate the algorithm’s robustness against imperfectly acquired real world data sets, and its insensitivity to noisy, imperfect and uneven performance of SfM.



Figure 2: Orthomosaic view of the VIS dense point cloud superimposed onto a Google Earth extract of the Jacobs Creek site at 34° 34' 04" South, 138° 56' 00" East. The white lines depict the approximate ‘lawn mower’ flight.

The UAV carried a gimbal-mounted Odroid digital camera (nominally oriented normal to the surface of the vineyard) with a 35mm lens. This translates to a roughly 60° field of view, providing about 85% and 50% overlap between images in the along- and across-track directions, respectively. This is sub-optimal for SfM but allows efficient coverage of the roughly 10Ha vineyard in about 3 minutes. UAV and camera orientation was not recorded (and hence remained unused) by the SfM.

The resolution of the camera provided a ground sampling distance of about 10cm. The vine canopies were approximately 1.5m high. They are not pruned and terrain undulates slightly over the roughly 10Ha area of the test site, varying by around 12m from lowest to highest points.

The UAV also carried an ICI 8640P long wave infrared (LWIR) thermal imaging camera, which has a 640 x 480 14-bit Vanadium Oxide radiometric imager, update rate of 10fps, and accuracy of 1°C. The observed irradiances (Figure 3) are converted to surface brightness temperature using the standard techniques for estimating such properties (Usamentiaga et al., 2014), i.e. the surface temperatures are not corrected for emissivity < 1.0, reflected temperature or atmospheric transmissivity.

As with the VIS imager, to ensure sub-optimal SfM performance, only every tenth image was used in the generation of DPC. Moreover, observations used in the generation of the DPC were taken at about 10.00am. This provides only modest contrast between leaf and ground irradiances and, as the sun has not yet reached its zenith, the vines create significant shadowing. Both circumstances are designed to challenge the vine segmentation algorithms operating on the data sensed by the LWIR payload.




Figure 3: Orthomosaic view of the dense point cloud generated from LWIR thermal imaging camera superimposed onto a Google Earth extract of the Jacobs Creek site. The DPC is colour-coded according to the temperature scale on the right of the image

Structure from motion (SfM) is the process of estimating the 3-D structure of a scene from a set of 2D images. It is used in many applications, such as robot navigation, autonomous driving, and augmented reality. SfM has been extensively described in the literature (Longuet-Higgins, 1981, Furukawa and Ponce, 2010, Hartley, 1997).

The process consists of two main components: camera motion estimation and DPC reconstruction. Initially, a sparse set of points are matched across the image stack to find correspondences. Such features are typically extracted using algorithms like SIFT (scale invariant feature transformation) (Lowe, 1999) and SURF (speeded up robust features) (Bay et al., 2008). The sequence of views are then iteratively processed to track a denser set of points

across the views so that the pose of the camera can be established for each image set and—after the relevant coordinate transforms have been accommodated—a dense 3D reconstruction of the scene made. The process of estimating camera motion (and hence DPC) is generally improved if camera pose is recorded during image capture and applied during SfM computations, which in our case it was not. If a suitable selection of ground control points (GCP) are applied to the resulting DPC, absolute scale and orientation may also be determined. This was achieved through matching the reconstructed DPC to a Google Earth view of the same area.

The VIS DPC of the vineyard generated by SfM in this way contains around 35×10^9 points, which are randomly, but approximately homogeneously, arranged throughout the 10Ha vineyard [Note: computation may be significantly enhanced, without great loss of accuracy, if the DPC is downsampled using interpolation based on grid averaging].

The LWIR DPC was also generated using SfM, with the image stack colour-coded according to temperature. The lower resolution of the LWIR imager results in a noisier and more distorted DPC as—in addition to generating camera location and orientation—the SfM process also estimates distortion created by the camera lens, which is more difficult for an LWIR sensor (Nistér, 2005).

We have used proprietary software (written in MATLAB), as well as Agisoft Photoscan and Reality Capture for the SfM computations.

Compute model of surface terrain (nearest neighbour network using Kd-tree): As the underlying topography of the terrain varies in altitude by approximately 12m over the vineyard, in order to use height as a discriminator for vine vs. inter-row material (both of which can be similar in colour), the surface topography must first be identified and vine height computed relative to this. Furthermore, as SfM can be an imperfect, noisy technique, vineyards with more subtle topographic variations generally represent greater challenges for height identification algorithms. This is because noise or inaccuracies in the DPC represent larger errors in relative terms, presenting greater challenges for clustering and association algorithms later in the processing chain. Consequently, a vineyard with shallow undulations was selected as the test site (Figure 4, left).

To ensure computational tractability, data is down-sampled by a factor of 100 and a Kd-tree (Zhou et al., 2008) used to partition the DPC by recursively splitting it into clusters. This allows a k-nearest neighbour (k-NN) search to be quickly evaluated on sub-blocks of the data using a query bounded by a circle of radius 5m. This in turn enables the approximate surface topography to be determined by finding the local minimum for each point queried in the k-NN search. The influence of outliers is minimised using a technique devised by (Holland and Welsh, 1977).




Figure 4: Variation of height for Jacob’s Creek test site (left) prior to removal of surface topography and (right) afterwards. Note change of scale for images

After resampling back to original ground resolution using 2D spline interpolation, with Golay filtering (Chen et al., 2004) applied to reduce noise, the complex 3D reference ground plane is subtracted from the original DPC height data to normalise (‘flatten’) the point cloud (Figure 4 right). Points lower than 0m are set to this value. It is recognised that grass growing between the rows may result in an elevated reference plane, and thus lower estimates of vine heights in some regions of the vineyard.

Cluster vineyard properties using Lloyd’s algorithm: In addition to normalising vineyard height, prior to clustering the colour properties of the VIS DPC are transformed from RGB to hue-saturation-value (HSV). The 2D circular scale of hue and saturation better enables colour properties to be associated than the linear RGB representation. The value component of the HSV scale is ignored as the degree of illumination can vary dramatically across a strongly sunlit vineyard. In the case of the LWIR sensor data, surface temperature is colour coded from minimum to maximum and thus clustered in terms of its HSV representation (with its value component again ignored).

Lloyd’s algorithm, also known as k-means clustering, is then applied separately to the normalised height, irradiance and HS values of the two DPCs. Lloyd’s algorithm is an iterative, data-partitioning (sometimes also referred to as unsupervised machine learning) technique that assigns observations to one of k clusters defined by centroids, where k is chosen before the algorithm starts. We typically used $k = 7$ for each category, finding results to be largely insensitive to values higher than this.

Lloyd’s algorithm operates as follows: (i) point-to-cluster-centroid distances for all observations to each centroid are computed; (ii) each observation is associated to the cluster with the closest centroid [note: observations can be individually assigned to a different centroids if the reassignment decreases the sum of the within-cluster, sum-of-squares point-to-cluster-centroid distances]; (iii) the mean of the observations in each cluster is computed to obtain k new centroid locations; (iv) steps (i) through (iii) are repeated until cluster assignments do not change, or the maximum number of iterations is reached. Figure 5 shows the outcome of the algorithm, colour-coded by cluster.




Figure 5: Outcome of Lloyd's classification, colour-coded (left) by HSV and (right) height ($k = 7$). Each colour represents elements of the DPC identified with each separate cluster. Each of the seven colours in the two images show points associated with each of the clusters.

Associate clusters using Gaussian PDFs: It is clear from the above images that, on its own, k -means clustering does not successfully identify vineyard properties in terms of either height or colour. We thus apply a priori knowledge to the height data, assuming that vines are less than 2m tall. The results of the height clustering are therefore ordered from smallest to tallest and the most significant cluster with a Lloyd's centroid less than 2m assumed to be that most likely to contain the bulk of the points containing vine foliage. Based on Bayesian association, the colour (irradiance) cluster most associated with this height cluster is then declared the most likely vine colour (irradiance).

A series of Gaussian PDFs are then established such that $p(v_{VIS}) = e^{-\alpha^2/\mu_H^2} e^{-\beta^2/\mu_C^2}$, where $p(v_{VIS})$ is the probability that any given VIS point is a vine, α and β are the point-to-centroid distances in the height and colour clusters, and μ_H and μ_C are the mean point-to-centroid distances obtained for the height and colour clusters declared in the previous paragraph. The process is then repeated for the LWIR DPC and $p(v_{LWIR})$ computed.

The VIS and LWIR clusters, which are co-registered to the same coordinate system, are fused using a joint PDF, $p(v_{Joint}) = p(v_{VIS}) + p(v_{LWIR})e^{-(X_{VIS}-X_{LWIR})^2/\sigma_d^2}$, where X_{VIS} , X_{LWIR} are coordinates of the points identified as vines using the VIS and LWIR data sets, respectively, and σ_d the root mean square value of the vine width derived from the VIS data.

A threshold, $p(v_{Joint})_T$, is set, such that joint probability estimates, $p(v_{Joint}) > p(v_{Joint})_T$, are declared vine material. This is the only parameter that need be set by a user, and which is readily adjusted on the basis of objectively viewing the degree of contamination by the background, i.e. roads, trees, bushes, etc.

3. ASSESSMENT OF SEGMENTATION PERFORMANCE

The results of the approach when applied only to the VIS DPC are shown in Figure 6, with an expanded view (Region 1) of an area where an irrigation pipe has ruptured shown in Figure 6, right. This region has a combination of significant green inter-row material, much more vigorous vine growth, shadowing caused by extended vine growth,

and varying hue from strong sunlight. Despite this, the algorithm has correctly identified only vine material.




Figure 6: Result of vine identification algorithm using VIS data alone (left) and (right) close up view (Region 1). The same area is shown in Figure 8.




Figure 7: Expanded view of the raw imagery (left) and DPC (right) of Region 2 (Figure 6)

There are also patches of vineyard where the algorithm has failed to correctly identify vine material, e.g. Region 2. While the algorithm suggests a complete absence of vine material, a better interpretation is that the plant density has fallen below a threshold ground resolution required by the SfM. This is confirmed by examination of the raw imagery (Figure 7), which shows the sparsity to be largely due to poor performance of the SfM (noting it was intentionally provided imagery gathered under sub-optimal conditions to show performance under these circumstances). The image on the left of Figure 7 shows an expanded view of the raw VIS imagery observed onboard the UAV. This shows the vine material in this region, albeit thinner than in other areas of the vineyard. The image on the right of Figure 7 shows SfM output, where distortions in the resultant DPC have obscured the sparser vine material and resulted in a much flatter 3D structure than truly exists. As a result, these sparser segments of the vineyard have been identified as ground material.




Figure 8: Result of vine identification algorithm based on LWIR data alone (left) for the entire vineyard and (right) for an expanded area (Region 1)

The results of LWIR-only segmentation are shown in . As with VIS-only outcomes, due a combination of poor height determination by the SfM and slightly cooler surface

temperatures in the more densely vegetated areas, performance is imperfect in areas such as Region 1.




Figure 9: Result of vine identification algorithm using VIS and LWIR data (left) and (right) expanded view of Region 1




Figure 10: Superposition of vine segmentation and LWIR observations with a selection of regions that are expanded below for more detailed scrutiny of performance

A more holistic validation of the segmentation algorithm was evaluated over the entire vineyard by contrasting the identified vine material against thermal imagery. The ground truth data was observed on a day for which air temperatures reached about 43°C. As ground brightness temperatures typically reach up to 60°C, but the vine leaves transpire (and their surface temperatures only reach about 40°C), the vines are clearly visible against the background. Unfortunately, regions in which irrigation infrastructure has broken, areas shaded from the direct effects of the sun and areas with sparse vine density but irrigated also have surface brightness temperatures around 40°C. A data set observed close to local noon (when the sun was near its zenith) was selected (a) to minimise the effects of shadowing and maximise the effects of contrast and (b) because this was prior to the cycle of irrigation, minimising any evapotranspiration effects in sparse vine regions.

The vine segmentation and ground truth DPCs were then co-registered and superimposed onto one another for visual examination (Figure 10). There are small registration errors apparent in the super-position, particularly in the

along row direction. Nevertheless, it is clear, even from this busy image, that the vine segmentation algorithm has successfully identified only vines over the vast majority of the vineyard.




Figure 11: Expanded views of “Region 1” in Figure 10 (left) and the corresponding VIS DPC (right). In the left hand image the superimposed black dots indicate segments identified as vines, whereas yellow and red colouring below these dots indicate higher (likely between row material) and lower (likely vine material) ground temperatures, respectively




Figure 12: Expanded views of “Region 3” (left) and “Region 4” (right) of Figure 10

Closer inspection of an expanded region (Figure 11) confirms the technique’s capacity to discriminate between the dense vine foliage where irrigation has been damaged and grass is growing strongly between rows (Region 1, Figure 11). Similarly, the segmentation has successfully dealt with shadowing in both the VIS and LWIR data. Figure 12 (left) shows an expanded view of Region 2 from Figure 10: an area that appears to be only sparsely populated with vines based on the VIS DPC. Figure 12 (right) shows another expanded view of a section of vineyard where the algorithm has had to deal with complex lighting and foliage density.

Hough Transform: Using the identified vine locations, some useful parameters such as those identified by (Weiss and Baret, 2017) may now be computed, e.g. row width, height, spacing, cover fraction and missing segments. In order to do this, however, the exact orientation of the vine rows must first be identified. This can be achieved manually, by visually determining row orientation, or automatically, using a Hough transform (Ballard, 1981). This transform uses a parametric representation of a line, $\rho = x \cos \theta + y \sin \theta$, where ρ is the distance from the origin of the DPC to the line along a vector perpendicular to the line, and θ is the angle between the x-axis and this vector. The Hough transform is used extensively in image processing to identify continuous

point sets that constitute lines in data. Consequently, as vine rows are typically arranged as a set of parallel lines, the method readily identifies row orientation.

In general, however, a Hough Transform will not successfully determine orientation where vines are planted along iso-altitude curves. Under such circumstances rotation angles should be computed manually or the location of the vine rows computed as a function of vineyard geography using an alternative technique. One simple approach is to identify a small section of vine (in the along-row direction) at the left or right extreme of a vineyard, compute the across-row median and row separation, and perform the property identification calculations below for a short vine segment; and then repeat, having advanced to the next segment in the along-row direction (note: this increases the computational cost of the procedure).

Crop row width: Crop row width, $W_{Cij} = |x_i - x_{M_j}|$, may be computed for each point in the DPC (, upper) [Note: the calculations of row width associated with the vineyards NW and SE of the main 10Ha test site show erroneous values in . This is because the row centroids, x_{M_j} , are computed for the 10Ha vineyard]. The elements of the DPC associated with each row may then be sorted into ascending order and along-row separation distances computed (, lower). The more significant breaks in the vine rows are clearly visible. The enlarged view of the data in the right hand image show gaps in the vines less readily visible in the entire DPC.

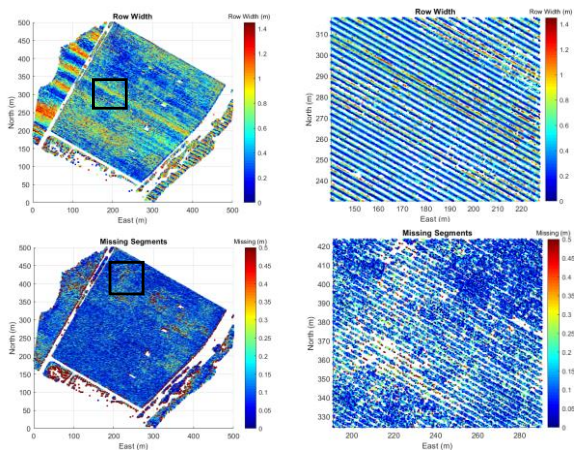


Figure 13: Spatial distribution of vine row width (top) and missing row segments (bottom) over vineyard (enlarged views on are shown on right)

4. CONCLUSIONS

A technique for identifying vines and discriminating them from inter-row material was developed and tested. Although the algorithm can be used solely on dense point clouds generated from aerial RGB imagery obtained from a UAV, better results derive from a combination of RGB and LWIR sensed data. The technique was evaluated against aerial measurements and showed good performance, despite the test data being obtained under

conditions that included harsh shadowing, poor contrast, and UAV flight paths and camera settings that delivered sub-optimal performance from the SfM. Unlike existing classification techniques, the approach in this paper is not critically reliant on specialist users who are expert in the algorithm’s design, mathematical techniques or manipulation of information: they need supply only a single, intuitive parameter. Moreover, the algorithm readily copes with degraded performance of upstream processing. Finally, the approach does not rely on supervised machine learning approaches that require training data sets. Characteristics such as canopy leaf area density and plant vigour can be readily evaluated as a function of vineyard geography, making the algorithm suitable for inclusion in decision support tools that aim to reduce fieldwork and improve vineyard management practices.

ACKNOWLEDGMENTS

The authors are grateful to: Treasury Wine Estates, Pernod Ricard and Barton Vale Technologies—in particular Dr Catherine Kidman, Tim McCarthy and Peter Smith—for their collaboration, our UniSA colleagues Joshua Meade, Jarrod Skinner, Kevin Rogers, and Amir Zargarian for trials support and Steve Andriolo of EyeSky for UAV operations. **Funding:** This research was funded by Wine Australia, grant number USA1601, “Use of Unmanned Air Vehicles for early, real time detection of extreme weather events in vineyards”

REFERENCES

- Ballard, D. H. (1981) Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recognition* **13**(2):Pages 111-122.
- Bay, H., Ess, A., Tuytelaars, T. & Van Gool, L. (2008) SURF: Speeded Up Robust Features. *Computer Vision and Image Understanding* **110**(3):346-359.
- Bendig, J., Bolten, A. & Bareth, G. (2013) UAV-based imaging for multi-temporal, very high resolution crop surface models to monitor crop growth variability. *PGF - Journal of Photogrammetry, Remote Sensing and Geoinformation Science* **6**:551–562.
- Bentley, J. L. (1975) Multidimensional binary search trees used for associative searching. *Communications of the ACM* **18**(9).
- Candiago, S., Remondino, F., De Giglio, M., Dubbini, M. & Gatelli, M. (2015) Evaluating multispectral images and vegetation indices for precision farming applications from UAV images *Remote Sensing* **7**:4026–4047.
- Chanussot, J., Bas, P. & Bombrun, L. (2005) Airborne remote sensing of vineyards for the detection of dead vine trees In *Proceedings of IEEE International Geoscience and Remote Sensing Symposium*, pp. 3090–3093.
- Chen, J., Jönsson, P., Tamura, M., Gu, Z., Matsushita, B. & Eklundh, L. (2004) A simple method for reconstructing a high-quality NDVI time-series data set based on the Savitzky–Golay filter. *Remote Sensing of Environment* **91**(3–4):332-344.
- Delenne, C., Durrieu, S., Rabatel, G. & Deshayes, M. (2010) From pixel to vine parcel: A complete methodology for vineyard delineation and characterization using remote-sensing data. *Computers & Electronics in Agriculture*:78–83.

- Dobrowski, S. Z., Ustin, S. L. & Wolpert, J. A. (2003) Grapevine dormant pruning weight prediction using remotely sensed data. *Aust. J. Grape Wine Res.* **9**:177–182.
- Elseberg, J., Magnenat, S., Siegwart, R. & Nuchter, A. (2012) Comparison of nearest-neighbor-search strategies and implementations for efficient shape registration. *Journal of Software Engineering for Robotics* **3**(1):2-12.
- Ferreiro-Armán, M., Da Costa, J. P., Homayouni, S. & Martín-Herrero, J. (2006) *Hyperspectral image analysis for precision viticulture*. Berlin/Heidelberg, Germany, Springer.
- Fiorillo, E., Crisci, A., De Filippis, T., Di Gennaro, S. F., Di Blasi, S., Matese, A., Primicerio, J., Vaccari, F. P. & Genesio, L. (2012) Airborne high-resolution images for grape classification: Changes in correlation between technological and late maturity in a sangiovese vineyard in central Italy. *Australian Journal of Grape & Wine Research* **18**:80–90.
- Furukawa, Y. & Ponce, J. (2010) Accurate, dense, and robust multiview stereopsis. *IEEE Transactions in Pattern Analysis & Machine Intelligence* **32**:1362–1376.
- Hall, A., Lamb, D. W., Holzapfel, B. & Louis, J. (2002) Optical remote sensing applications in viticulture—A review *Australian Journal of Grape Wine Research* **8**:36–47.
- Hartley, R. (1997) In defense of the eight-point algorithm. *IEEE Transactions of Pattern Analysis & Machine Intelligence* **19**:580–593.
- Hernández-Clemente, R., Navarro-Cerrillo, R. M. & Zarco-Tejada, P. J. (2012) Carotenoid content estimation in a heterogeneous conifer forest using narrow-band indices and PROSPECT + DART simulations. *Remote Sensing & Environment* **127**:298–315.
- Holben, B. N. & Justice, C. O. (1980) The topographic effect on spectral response from nadir-pointing sensors. *Photogrammetry, Engineering & Remote Sensing* **46**:1191–1200.
- Holland, P. W. & Welsch, R. E. (1977) Robust regression using iteratively reweighted least-squares. *Communications & Statistical Theory & Methods* **6**:813-827.
- Johnson, L., Bosch, D., Williams, D. & Lobitz, B. (2001) Remote sensing of vineyard management zones: Implications for wine quality. *Applied Engineering in Agriculture* **17**:557–560.
- Lacar, F. M., Lewis, M. M. & Grierson, I. T. (2001) Use of hyperspectral imagery for mapping grape varieties in the Barossa Valley, South Australia In *Proceedings of IEEE International Geoscience and Remote Sensing Symposium*, pp. 2875–2877.
- Longuet-Higgins, H. C. (1981) A computer algorithm for reconstructing a scene from two images. *Nature*.
- López-Lozano, R., Baret, F., García De Cortázar-Atauri, I., Bertrand, N. & Casterad, M. A. (2009) Optimal geometric configuration and algorithms for lai indirect estimates under row canopies: The case of vineyards *Agriculture, Forestry & Meteorology* **149**:1307–1316.
- Lowe, D. G. (1999) Object recognition from local scale-invariant features In *Proceedings of IEEE International Conference on Computer Vision* vol. 2, pp. 1150-1157.
- Macqueen, J. (1967) Some methods for classification and analysis of multivariate observations In *Proceedings of Berkeley Symposium on Mathematical Statistics and Probability*. University of California Press Berkeley Symposium on Mathematical Statistics and Probability vol. 1, pp. 281.
- Mathews, A. & Jensen, J. (2013) Visualizing and quantifying vineyard canopy LAI using an unmanned aerial vehicle (UAV) collected high density structure from motion point cloud. *Remote Sensing* **5**.
- Meggio, F., Zarco-Tejada, P. J., Miller, J. R., Martín, P., González, M. R. & Berjón, A. (2008) Row orientation and viewing geometry effects on row-structured vine crops for chlorophyll content estimation. *Canadian Journal of Remote Sensing* **34**:220–234.
- Nistér, D. (2005) Preemptive RANSAC for live structure and motion estimation. *Machine Vision and Applications* **16**(5):321-329.
- Puletti, N., Perria, R. & Storchi, P. (2014) Unsupervised classification of very high remotely sensed images for grapevine rows detection. *European Journal of Remote Sensing* **47**:45–54.
- Rey, C., Martín, M. P., Lobo, A., Luna, I., Diago, M. P., Millan, B. & Tardáguila, J. (2013) Multispectral imagery acquired from a UAV to assess the spatial variability of a tempranillo vineyard. *Precision Agriculture* **13**:617–624.
- Rouse, J. W., Haas, R. H., Schell, J. A., Deering, D. W. & Harlan, J. C. (1974) *Monitoring the Vernal Advancement of Retrogradation of Natural Vegetation* Texas A&M University, College Station, TX, USA.
- Shepherd, J. D. & Dymond, J. R. (2003) Correcting satellite imagery for the variance of reflectance and illumination with topography. *International Journal of Remote Sensing* **24**:3503–3514.
- Smit, J. L., Sithole, G. & Strever, A. E. (2010) Vine signal extraction—An application of remote sensing in precision viticulture. *South African Journal of Enology & Viticulture* **31**:65–74.
- Usamentiaga, R., Venegas, P., Guerediaga, J., Vega, L., Molleda, J. & Bulnes, F. G. (2014) Infrared Thermography for Temperature Measurement and Non-Destructive Testing. *Sensors* **14**(7):12305-12348.
- Wassenaar, T., Robbez-Masson, J. M., Andrieux, P. & Baret, F. (2002) Vineyard identification and description of spatial crop structure by per-field frequency analysis. *Int. J. Remote Sens* **23**:3311–3325.
- Weiss, M. & Baret, F. (2017) Using 3D Point Clouds Derived from UAV RGB Imagery to Describe Vineyard 3D Macro-Structure. *Remote Sensing* **9**(1).
- Zarco-Tejada, P. J., Berjón, A., López-Lozano, R., Miller, J. R., Martín, P., Cachorro, V., González, M. R. & De Frutos, A. (2005) Assessing vineyard condition with hyperspectral indices: Leaf and canopy reflectance simulation in a row-structured discontinuous canopy. *Remote Sensing & Environment* **99**:271–287.
- Zarco-Tejada, P. J., González-Dugo, V. & Berni, J. a. J. (2012) Fluorescence, temperature and narrow-band indices acquired from a UAV platform for water stress detection using a micro-hyperspectral imager and a thermal camera. *Remote Sensing & Environment* **117**:322–337.
- Zarco-Tejada, P. J., Guillén-Climent, M. L., Hernández-Clemente, R., Catalina, A., González, M. R. & Martín, P. (2013) Estimating leaf carotenoid content in vineyards using high resolution hyperspectral imagery acquired from an unmanned aerial vehicle (UAV). *Agriculture Forestry & Meteorology* **171–172**:281–294.
- Zhou, K., Hou, Q., Wang, R. & Guo, B. (2008) Real-time KD-tree construction on graphics hardware. *ACM Transactions on Graphics (TOG)* **27** (5).