

Archie Jnr: A Robotic Platform for Autonomous Cane Pruning of Grapevines

Henry Williams^{1*}, David Smith¹, Jalil Shahabi¹, Trevor Gee¹, Ans Qureshi¹,
Ben McGuinness², Scott Harvey², Catherine Downes², Rahul Jangali², Kale Black², Hin Lim²,
Mike Duke², Bruce A. MacDonald¹

Abstract—Cane pruning grapevines is a complex manual task requiring expert vine assessment to determine which canes to prune. This paper presents Archie Jnr, which was developed to autonomously assess the structure of the vine and prune the lower-quality canes as an expert pruner would. The platform has been extensively evaluated in a real-world commercial vineyard using a three-cane pruning method. The results show the effectiveness of the vision system for generating accurate assessments of a vine's canes. The platform is also shown to be capable of successfully pruning 71.1% of the 311 total canes that required pruning across 25 vines.

I. INTRODUCTION

Aotearoa New Zealand's winegrape industry totals nearly \$2 billion annually and is rapidly growing [9]. Following similar global trends, there are growing challenges in maintaining skilled agricultural workers to prune the grapevines reliably to keep up with this expansion. Pruning is critical to maintaining the vine size and shape, which helps regulate crop load and balance. If grapes are not pruned yearly, they develop many unproductive shoots, reducing the quality and overall yield of the vineyard. A reliable means of pruning is crucial to maintain and lift growth targets in this sector.

In cane pruning, one to four new branches are selected each winter to be trained along the wires, and the rest are removed. Cane pruning can produce a higher yield in some varieties [2], [7] as you increase the number of fruit-bearing shoots if the right decisions are made during the pruning process. The wrong choices can severely reduce the plant yield for the current and future years. The choice of which canes are selected for a given year is a decision that requires extensive training. The difficulty in accessing and training a seasonal workforce that is large enough to keep up with the projected expansion of the industry has generated interest in automated platforms capable of assessing and pruning wine grapevines at the same quality as a trained human worker.

This paper presents an autonomous cane pruning robot called Archie Jnr in a real-world commercial vineyard, shown in figure 1. Prior work [8] has demonstrated the effectiveness of the vision system for modelling grapevines; this paper presents an additional vision system evaluation on a new vineyard, a novel automated pruning mechanism, and

autonomous pruning performance. A video demonstration of the platform can be found here¹.



Fig. 1: Archie Jnr in the vineyard driving over the row of vines on the left, and the internal view of the robot platform with the side panels removed on the right.

II. RELATED WORK

Automation in agriculture is a popular field of research, with numerous robotic platforms being developed for the task of pruning [12], [13], [6]. The crops that have seen the most work for automating pruning are cherries [12], apples [13], and grapevines [6]. In all cases, the robot aims to prune old material off the plant to make room for new productive growth. The different crops utilise various pruning methods, but most existing work is similar to spur pruning in grapevines.

Work conducted by [12] presents an automated system for cherry trees that share a similar structure to grapevines. The pruning of these cherry trees is similar to spur pruning in grapevines. The paper states the system detected 38 branches to cut but did not clarify how many cherry trees the system was tested across. Using standard off-the-shelf electric bypass pruners on the end of the platform arm, they showed a cutting success rate of 58% on the 38 detected branches. The primary limitation was the planning system's inability to place the cutting tool.

The most advanced example of automating winegrape pruning is presented in [6] as the Bumblebee, designed for spur pruning. Of the critical decision metrics as given in section III, only the node count is compared against a ground truth measurement, which presents an accuracy of 94.89%. The system could spur-prune a row of vines from both sides entirely in 213 sec/vine with a total pruning accuracy of 87%.

¹The authors are with the Centre for Automation and Robotic Engineering Science, The University of Auckland, New Zealand.

²The authors are with the School of Engineering, University of Waikato, Hamilton, New Zealand.

Corresponding author: henry.williams@auckland.ac.nz*

¹<https://www.youtube.com/watch?v=GdVJnK1KZvc>

However, this evaluation was done on a pre-pruned vineyard to simplify the physical structure of the vines. Within this work, they point out that 32% of the cuts had to be conducted from the other side of the vine. This challenge is more significant in a cane-pruned system where canes can grow on either side of the structure as they are not limited to a cordon.

Extensive work has been explored for generating accurate 3D models of plants for phenotyping [4]. Most of this work operates on isolated individual plants in controlled conditions [3] where they can generate dense point clouds of the plant without gaps. These techniques struggle to handle movement (due to wind) and significant occlusion from complex plant geometries, such as overlapping canes on a grapevine. Recent work has attempted to extract the skeleton from complex grapevines for measuring pruning weight [5]. The approach shows promise for skeletonising dense grapevine structures but does not provide measures to indicate the performance of segmenting individual canes from the data.

Overall, robotic pruning has been extensively explored in the literature but has focused on the equivalents of spur pruning. Detecting the position of the vines is quite robust in literature at this stage, but extracting individual canes from the vine is still unsolved. However, generating a detailed 3D Model of the structure of the vine in a real-world vineyard is still a challenge.

III. REQUIREMENTS

Vines are pruned in winter when dormant and free of leaves. Vines are approximately 1.8m wide by 2.0m tall, although this varies from vineyard to vineyard. Cane-pruned vines are laid flat in a 2D structure and wrapped around the guide wires to maintain plant structure and stability. Through discussions with expert pruners - vineyard managers and trainers with many years of pruning experience - the key *quantifiable* metrics used to evaluate a given cane's quality can be defined as the DOLPHIN metrics, highlighted in figure 2.

- **Diameter** measures the cane's thickness and thus stored carbohydrates reserves necessary to support new growth. Typically measured between the second and third node.
- **Orientation** determines whether the cane is best suited for tying down on the left or right-hand side of the head, as bending too far in the opposite direction may break the cane.
- **Length** is a binary check that the cane is long enough to maximise this year's yield potential.
- **Position** determines if the cane is below the first structural wire. One of the objectives of pruning is to keep the growth of the vine's head below the first wire to maximise the potential growth year after year.
- **Health** of the cane is a visual inspection for any diseases on the plant; this specific measure was omitted for this paper as our data did not show any visible signs of disease or damage.

- **Internode length** determines the spacing of the new shoots along the cane. Typically measured between the first and second or the second and third nodes along the cane.
- **Node count** is the number of nodes along the cane, with the desired number required for producing new canes depending on the variety.

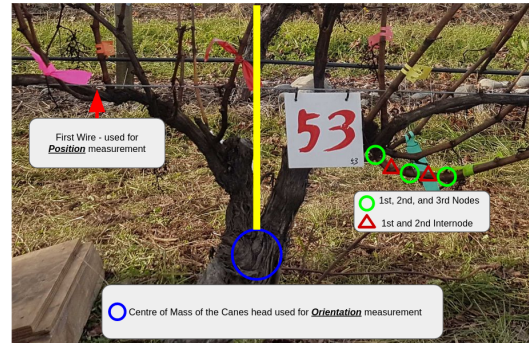


Fig. 2: An example of a vine with the first, second, and third nodes (buds) highlighted. An example of an internode connection is also shown along with the centre of mass of the vine's head.

IV. ROBOTIC PLATFORM

As seen in figure 1, Archie Jr is an experimental platform designed to automatically cane-prune grapevines. The platform encapsulates the vine by arching over the row at 2.3 m tall, 1.938 m wide, and 3.325 m long. On either side of the vine, a UR5 robotic arm is mounted with stereo cameras and the custom pruning device, the Barracuda, shown in figure 3 and described further in Section VI-A.




Fig. 3: The Barracuda: a custom end-effector for Archie Jr designed to compliantly cut canes wrapped around wires by pushing the wire into the recess along the anvil. The large red circle indicates the cane, and the small blue circle is the wire.

The vine was encapsulated to eliminate wind, which shakes the canopy and disrupts the alignment of the vine mapping process by introducing a complex registration problem, an issue also seen in [6]. This also has the benefit of enabling the system to operate at night with internal lighting. Finally, similar to the work in [6], not all the canes are physically accessible or visible from a single side of the canopy. An actuator on either side of the vine allows us to see and cut canes from either side.

The Basler acA2440-35uc USB 3.0 cameras were used with a baseline of 100 mm. Kowa lenses (Kowa F1.8, 2/3" format, focal length of 5 mm) were attached to provide a

working distance of 300 mm to 600 mm suitable for scanning the width of the cane and directing the cutting tool.

V. VISION SYSTEM

A grapevine consists of primary canes growing off the vine's head. Along the canes are a series of nodes interconnected by cane growth, referred to as internodes. Additional shoots (new canes) grow from these nodes, making up the structure of the vine. The vine can be considered as a *3D tree structure* between all the nodes, with the head of the vine acting as the root connecting all the canes. The vision system attempts to extract the 3D tree structure of the vine and subsequent DOLPHIN metrics from the graph. The specific approach presented in this paper relies on a naive node-matching approach that leverages the physical morphology of the vine. The main challenges to achieving an accurate tree structure are differentiating overlapping canes and determining where sub-canecan branch from the primary canes. The full details of the vision system have previously been described in [8]; this section provides an overview of the vision system and highlights new features of the fully integrated system presented in this work.

A. Scanning Process

For this field trial, the cameras were moved approximately 0.4 m through a plane in front of the vine, measured from the guide wires. The arms moved in 0.1 m steps -0.5 m to 0.5 m across, and 0.0 m to 0.6 m up the vine centred on the head. The stereo data on a single side of the platform is aligned geometrically based on hand-eye calibration, which relies on the precision of the UR5s movements and calibration between the arm and the stereo pair. The stereo data between each side of the platform is aligned geometrically using a Charuco board on the platform's roof to provide a common reference point for each side.

The method of generating the 3D Model while the platform scans operates with the following sub-systems: *vine detection* to extract cane information from the images, *stereo depth inference* to generate depth information, *vine skeletonisation* to find the internode connections, and *cane extraction* to extract individual canes from the 3D tree structure. A flow diagram of this approach is shown in figure 4.

B. Vine Detection

This step utilises Detectron 2²[10] as a Panoptic Segmentation approach to extract the cane information from the 2D images. The model has been trained to detect *canes*, *nodes*, the *vine-head* and *wires*. Figure 4 shows an example of this detection at step two. HSMnet [11] was chosen to generate depth information from the high-resolution stereo images as it showed fewer gaps and outliers in the depth data compared to alternatives on the Middlebury Stereo Evaluation v3 database (<https://vision.middlebury.edu/stereo/eval3/>).

²<https://github.com/facebookresearch/detectron2>

C. Keypoint Skeletonisation

The first step uses panoptic segmentation to extract the cane, head, wire, and node locations in the 2D image, as seen in steps one and two of figure 4. The average height of the wires in the canopy structure is measured throughout the scan to determine the relative Position of the canes. The structure of the head is captured throughout the scan, with the centre of mass being used to determine the canes' Orientation.

The next step is to find the cane skeleton within the 2D image by linking the nodes along the cane segmentation. The internode growth typically occurs in straight vectors between the nodes, a natural property of the canes' growth. Treating the nodes as keypoints, the skeleton of the cane is found by linking the nodes with straight lines.

Straight lines are imposed on the 2D image between all detected node instances within 150 mm of each other. The cane segmentation and the depth information are then used to determine the valid connections. Lines that have 95% of their length fitting within the cane segmentation are kept, with the rest being filtered. The remaining canes are checked for depth continuity; the depth along the line is contiguous in the 3D space without a change more significant than 2 mm.

Crossing canes cause discontinuities in the depth data along the line, as shown in figure 5. If a discontinuity is detected through the line, the system checks if the original depth continues over the gap. If the original depth within 2 mm continues, the node connection is still valid; otherwise, it is removed. The remaining lines are the most likely internodes providing the vine's skeleton in the immediate scan data. Step three in figure 4 shows the result of this skeletonisation process.

Once the skeleton has been estimated, the Diameter and Length of each internode are measured. The length is measured as the distance between the centre point of each node. A perpendicular line is drawn at the centre of the two nodes to calculate the internode diameter, as shown in step four in figure 4. Two edge points are then found by following this line until either a depth change occurs or the edge of the segmentation has been reached. These two points are then transformed into 3D positions, and the distance between the two points is taken as the internode diameter.

The fifth step merges the immediate scan data into the full scan of the vine, linking the canes into the complete vine structure. This is required to find the complete cane information to determine the cane's Orientation, Length, and Node count. This is achieved by matching all nodes in the global reference frame within a Euclidean threshold of 10 mm. The internode measurements between associated nodes are averaged between the immediate snap-shot and prior measurements. These five steps are repeated until the scanning phase has finished.

D. Cane Extraction

Step six takes the merged skeletonisation data and validates the vine as a tree structure, removing invalid internode connections. Invalid internode connections are detected as loops in the tree structure. As this is physically impossible

before the robot operated on them, as would be the case for a human pruner. The goal was to evaluate the performance of the vision system to produce accurate DOLPHIN metrics and its ability to prune the vine autonomously. Experiments were conducted over one week with the robot operating in various lighting and wind conditions. Experiments were not conducted during periods of rain.

For health and safety reasons, the platform was manually driven over a vine to place the scanning arms centred on the vine's head. Future work will fully automate the vehicle's driving. The scanning process was run to capture the 3D skeleton of the vine. Once the model was produced, the human operators manually selected all the canes on the head for the system to prune. All other canes that might need to be pruned as part of the cane pruning process were also selected based on a general expert rule set. In cases where the 3D model was incomplete, an internode was manually added to the model for the robot to attempt to prune the cane.

The robot then autonomously attempted to prune the canes by placing the cutting tool between the first or second internode connection. As this is an active commercial vineyard, the robot did not complete the cut to avoid damage to the vines. The Barracuda's cutting performance was evaluated manually using the device to cut various canes throughout the vineyard, similarly to the robot. The pruning attempt was successful if the tool was correctly placed around the cane such that a cut could be executed, as shown in figure 8. To avoid damage to the vine or robot, canes that could not physically be attempted were not attempted and were marked as uncuttable. Once completed, the robot was driven to the next vine in the row. This process was carried out on a total of 25 vines.




Fig. 8: The top three images show the successful placement of the cutting tool. The bottom three images from left to right show failures - a collision with a neighbouring cane, a collision with the wire structure, and an uncuttable cane due to neighbouring canes.

Ground truth measurements of the DOLPHIN metrics for each vine were taken and recorded to compare against the metrics extracted by the platform. The DOLPHIN metrics for ten canes on each vine were measured manually using callipers and a measuring tape as accurately as practically

measurable by hand.

VIII. RESULTS

The performance of the pruning system was evaluated in four parts. First was skeletonisation, the vision system's ability to generate an accurate tree structure of the vine. Second, the ability to accurately measure the DOLPHIN metrics for every cane. Third, the platform's ability to successfully prune canes on the vines. Finally, the cutting performance of the Barracuda design.

A. 3D Skeletonisation

The 3D skeletonisation is evaluated based on how many candidate canes were correctly extracted from the scanning process. Furthermore, the completeness of the extracted cane was measured by determining if the cane was extracted as one single cane or split into multiple parts due to failures in the skeletonisation process. An example of a 3D Model with a whole cane and split cane is shown in figure 4. The results of the skeletonisation process are given in table I. These results are based on the total number of canes across the 25 scanned vines.

TABLE I: Skeletonisation Results

Parts	Canes 2022 [8]	Canes 2023 (latest)
1	51.45%	45.3%
2	35.51%	19.7%
3+	13.04%	35.0%

In total 45.3% of the canes were fully captured and extracted by the skeletonisation process. An additional 19.7% only had a single internode connection missing along the length of the cane. The missing connection was typically seen around the head of the cane, where the canes are bunched closer together and frequently overlap. The primary reason for missed connections across 35.0% of the other canes was from crossing canes, where the depth data was insufficient to differentiate them.

Compared to the previous results on the vineyard in [8] where the system fully extracted 51.45% of the canes fully and 35.51% missing only one internode, the performance is poorer. Qualitatively, the more mature vines in the 2023 vineyard had a higher growth density around the head of the vine, leading to more overlapping canes. These overlaps are a significant challenge to disambiguate with the accuracy of the current depth system. Future work needs to focus on improving the depth information; one means of doing this is specifically training HSMnet or alternatives specifically for vineyard environments.

B. DOLPHIN Metrics

The DOLPHIN metrics for each cane are estimated from the 3D Model and compared against the ground truth manually measured on the physical vine. The canes were manually selected from the 3D Model to match the cane's location in the physical world. If the cane is in multiple parts, the part starting closest to the head is selected for generating the DOLPHIN metrics. This means the DOLPHIN metric results

depend on the skeletonisation process’s accuracy. The results of this comparison are shown in table II.

TABLE II: DOLPHIN Metric Results

	Accuracy
Identification Rate	83.2%
Diameter (mm)	±2.88
Orientation	95.2%
Length	50.2%
Position	66.2%
Internode Length (mm)	±20.3
Node	±7

From the skeletonisation process, 83.2% of the canes had the correct first or second internode connection extracted to measure the DOLPHIN metrics. However, given the skeletonisation results in table I, the DOLPHIN metric accuracy varied regarding whether the cane was extracted in a single part. This is particularly noticeable in the low length and node accuracy, with only portions of the cane being extracted 35.0% of the time. A manual investigation of the estimates of the DOLPHIN metrics on the vine demonstrates sufficient accuracy can be captured from the model.

C. Barracuda Performance

The cutting performance of the Barracuda was done by manually using the device on various vines throughout the vineyard. The cuts must be as clean as possible to promote healthy cane growth for the following year. A clean cut is characterised by a smooth surface without any loose fibres. Otherwise, the vine will dry out beneath the cut surface, and nearby buds or the sap flow will suffer damage.

The Barracuda was first tested by cutting a variety of canes where no wire was involved to determine a baseline cutting performance. Cuts were then made on sections of the cane where the cane was wrapped around the wire, and unless moved, the wire would obstruct a clean cut through. Cane and shoot sizes ranged from 2 mm to 20 mm off the wire, and 6 mm to 20 mm on the wire. Alignment of the secateurs with the cane was done by referencing the position of the secateurs to the cane and wire, specifically the position of the Barracuda teeth. As seen in Figure 9, testing was conducted over four different placements of the wire and cane with the Barracuda: 0°, 90°, 180° and 270°. The Barracuda was placed as accurately as feasible by hand at these relative angles.




Fig. 9: Pruning tests were conducted over four different positions presented in the image: A 0°, B 180°, C 90°, and 270°.

Performance was graded based on how well the secateurs cut through the cane, with the results shown in table III. Five

TABLE III: Barracuda cutting performance results.

Success Type	No Wires		Wire 0°		Wire 90°		Wire 270°	
Clean Cut	76	74.5%	229	59.8%	23	24.0%	36	29.0%
Acceptable Cut	21	20.6%	133	34.7%	50	52.1%	61	49.2%
Incomplete Cut	2	2.0%	21	5.5%	20	20.8%	26	21.0%
Damage	0	0.0%	0	0.0%	0	0.0%	0	0.0%
Failed Cut	3	2.9%	0	0.0%	3	3.1%	1	0.8%
Total Cuts	102		383		96		124	

qualitative measures were used to categorise the performance of the Barracudas from least desirable to most desirable:

- **Failed Cut:** no incision, or minimal incision
- **Damage:** damage to the cane from the cutting process, typically seen as the crushing of the cane or a minimal incision
- **Incomplete Cut:** over 2 mm of the cane has not been cut, leaving a partial cut that is not easily removed by machinery or hand.
- **Acceptable Cut:** under 2 mm of the cane has not been cut; this is considered acceptable by commercial partners as it is still easily removable by machinery or hand without causing damage to the canes.
- **Clean Cut:** A complete cut of the vine with a clean separation between the two pieces of the cane.

In the control experiments without the wire, the Barracuda successfully cut 97.0% of the canes, with 76.0% being clean. The two rows of teeth effectively double the contact area the pressure is applied to over the cane as the blade slices through it. This enabled it to successfully cut 76.0% of the canes, with 21.0% leaving only a minor and acceptable connection along the cane. The remaining but acceptable connections were primarily caused by the bottom segment of the cane falling into the recesses and avoiding the blade. The three failed cuts were on canes less than 3 mm in diameter. They were pushed entirely into the recesses at this size, avoiding the blade. Similarly, the two failed cuts with greater than 2 mm remaining were less than 3.5 mm in diameter and mostly avoided blade contact by falling into the recess. Reducing the recess size may reduce these limitations, but it would likely obstruct the wire, potentially causing damage to the wire or Barracuda. Overall, the Barracuda did not substantially reduce the acceptable quality of typical cuts with no wires.

The Barracuda showed a negligible 0.6% drop in successful cuts with the wire in the ideal 0° position while maintaining a 59.8% clean cut rate. A success rate of 76.1% and 78.2% was also still achievable at the extreme positions of 90° and 270°. The rounded tooth profile effectively allowed the wire to fall into the recesses during cutting. These results demonstrate that the Barracuda is suitable for the automated platform and provides a reasonable tolerance level in its positioning around the wire and cane. The Barracuda was ineffective when the wire was at the 180° position, leading to immediate damage to the blade or wire. Cuts at this position were abandoned to avoid further damage to the Barracuda or the vineyard. This means the automated platform will be required to detect the wire and place the anvil of the Barracuda on the opposite side.

D. Automated Pruning

Table IV presents the outcome of the platform attempting to cut 270 canes out of a total of 311 over 25 vines in the vineyard. Of the cuttable canes, the platform successfully placed the end-effector at 81.9% of them. The similar success rate per vine at 81.4% shows that the platform is consistent from vine to vine. However, incorporating the uncuttable canes, the cut success rate falls to 71.1%. There was an almost 50% split on which side the canes had to be cut on at 53.5% and 46.5%. This reinforces the requirement for the pruner to operate on both sides of the vine.

The primary cause of failures to place the end-effector was caused by collisions of the end-effector with neighbouring canes, the vine head, or the wire structure. Examples of these failure cases are shown in figure 8. The most prevalent failure case when attempting to cut was the end-effector catching on close neighbouring canes. Reducing the physical profile of the cutting tool would enable more precise placement of the cutting head. However, this must be carefully considered to not compromise the ability to cut the larger canes.

TABLE IV: Cutting Results.

	Total	Average Per Vine
Number of Vines	25	25
Candidate Canes	311	12.44
Cuttable Canes	270	10.8
Uncuttable Canes	41	1.64
Success Cuttable	81.9%	81.4%
Success Overall	71.1%	71.0%

IX. DISCUSSION

Archie Jnr can extract accurate DOLPHIN metrics if a cane is sufficiently extracted - that is, in one complete part. The Diameter and internode Length are within a reasonable tolerance to the precision of the measured ground truth of the canes. The Orientation, Length, and Position accuracy depends explicitly on whether the individual canes' primary node is successfully detected. If the primary node is detected, the system can differentiate these metrics accurately.

The skeletonisation process requires further improvements to associate the missing internode connections. Failure to link the correct nodes in the skeletonisation process is primarily caused by ambiguities in the depth data. These ambiguities prevent the canes from being distinguished in 3D space, as shown in figure 5. Incorrectly resolving the connection errors shown in figure 6 also leads to incorrect cane paths and disconnected canes. Observing the 2D images, it is clear to a human which cane the nodes belong to by tracing the path of the canes between the nodes. This feature could be exploited through machine learning-based keypoint detection for extracting the internode connections [1].

The automated pruning performance demonstrated the platform's ability to prune 71.1% of the 311 possible canes. Reducing the profile of the cutting tool and end-effector would help reduce failures caused by interference from neighbouring canes. However, we consider the uncuttable canes unreachable without some mechanism to shift them

out of the way as a human would. Significantly increasing the complexity of the end-effector. The evaluation of the Barracuda design demonstrates its suitability for pruning the canes. Future evaluations will fully automate the cutting step.

X. CONCLUSIONS

This paper presents Archie Jnr, the autonomous cane pruning robot. The platform has been extensively evaluated in a real-world commercial vineyard. The results have shown the effectiveness of the vision system for generating accurate DOLPHIN metrics for making quality pruning decisions. The platform was also able to successfully prune 71.1% of the 311 canes that required pruning across 25 vines. Future work will focus on improving the skeletonisation process and refining the custom cutting tool for improved positioning within the complex vine structure.

ACKNOWLEDGEMENTS

This research was supported by the New Zealand Ministry for Business, Innovation and Employment (MBIE) on contract UOAX1810.

REFERENCES

- [1] Qi Dang, Jianqin Yin, Bin Wang, and Wenqing Zheng. Deep learning based 2d human pose estimation: A survey. *Tsinghua Science and Technology*, 24(6):663–676, 2019.
- [2] Tremain A Hatch, Mizuho Nita, and Tony K Wolf. Vegetative and reproductive responses of mature cabernet sauvignon grapevines converted from spur pruning to cane pruning at five-foot in-row vine spacing. *Catalyst: Discovery into Practice*, 3(1):8–16, 2019.
- [3] Kenta Itakura, Itchoku Kamakura, and Fumiki Hosoi. Three-dimensional monitoring of plant structural parameters and chlorophyll distribution. *Sensors*, 19(2):413, 2019.
- [4] Stefan Paulus. Measuring crops in 3d: using geometry for plant phenotyping. *Plant methods*, 15(1):1–13, 2019.
- [5] Eric Schneider, Sushanth Jayanth, Abhishesh Silwal, and George Kantor. 3d skeletonization of complex grapevines for robotic pruning. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3278–3283. IEEE, 2023.
- [6] Abhishesh Silwal, Francisco Yandun, Anjana Nellithimaru, Terry Bates, and George Kantor. Bumblebee: A path towards fully autonomous robotic vine pruning. *Field Robotics*, 2:1661–1696, 2022.
- [7] Rachael White, Annie Vogel, John Scaduto, and Cain Hickey. Evaluation of canopy division and cane pruning to retrofit spur-pruned, vertical shoot-positioned petit manseng. *Catalyst: Discovery into Practice*, 4(1):21–32, 2020.
- [8] Henry Williams, David Smith, Jalil Shahabi, Trevor Gee, Mahla Nejati, Ben McGuinness, Kale Black, Jonathan Tobias, Rahul Jangali, Hin Lim, et al. Modelling wine grapevines for autonomous robotic cane pruning. *biosystems engineering*, 235:31–49, 2023.
- [9] NZ Wine. 2022 new zealand winegrowers inc annual report, 2023.
- [10] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. visited 2022-10-01.
- [11] Gengshan Yang, Joshua Manela, Michael Happold, and Deva Ramanan. Hierarchical deep stereo matching on high-resolution images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5515–5524, 2019.
- [12] Alexander You, Nidhi Parayil, Josyula Gopala Krishna, Udhav Bhattarai, Ranjan Sapkota, Dawood Ahmed, Matthew Whiting, Manoj Karkee, Cindy M Grimm, and Joseph R Davidson. An autonomous robot for pruning modern, planar fruit trees. *arXiv preprint arXiv:2206.07201*, 2022.
- [13] Azlan Zahid, Md Sultan Mahmud, Long He, Paul Heinemann, Daeun Choi, and James Schupp. Technological advancements towards developing a robotic pruner for apple trees: A review. *Computers and Electronics in Agriculture*, 189:106383, 2021.