Washington State Grape and Wine Research Program

FINAL REPORT

Mobile-App For Crop Estimation And Lag Phase Detection

Manoj Karkee, Center for Precision and Automated Agricultural Systems, Washington State University

Project Summary:

Automated crop estimation can be a highly valuable tool to wine grape growers (Read and Gamet, 2007). Precise prediction of wine grape yields helps in managing crops well to meet quality targets as well as in efficient handling of the grapes, and manage labor force and equipment requirement ahead of the harvest season. This information can also be useful in the study of long-term sustainability of the vineyards and the wine industries. Some of the existing methods for crop/yield estimation in wine grapes require information on number of vines per acre, number of clusters per vine and number of berries per cluster (Sabbatini, Dami and Howell, 2012). This task is tedious, labor intensive and time consuming and destructive in nature as it requires sampling, vine collection, manual counting of vines and weighing them. Labor-intensive tasks have been extensively challenging in modern agriculture industries due to the shortage of labor force and the increasing costs associated with it. Numerous studies have been reported in literature for automated crop/yield estimation in wine grapes using machine vision systems (Liu, Marden and Whitty, 2013; Bates, Grocholsky, Nuske, 2011; Nuske, Wilshusen, Achar et al., 2014). In addition, some private companies are also working on this technology. These systems were evaluated in laboratory environment (Liu, Marden and Whitty, 2013) and/or require some sophisticated technology such as very high-resolution imaging sensor mounted in an agriculture utility vehicle and, GPS unit that needs skilled manpower to operate. Moreover, these technologies have yet to meet the desired yield estimation accuracy. PI Karkee's team has successfully developed and tested a smartphone-based Software Application program (App) for yield estimation in apple orchards, which isolates computation of complex models in a distant server thus making the yield estimation easily available on smartphones and have shown up to 98% accuracy in estimating yield in a 0.25-acre block of a commercial orchard. Through funding from Washington State Wine Commission, PI Karkee's team expanded this Mobile-App to detect and count the number of clusters and number of berries in grape canopy images. This approach is a simple and low-cost yet practical design of a sensing system for crop estimation in wine grapes. The developed Mobile-App uses the sensors (cameras and GPS) of a mobile device (e.g., a smartphone) to acquire images in vineyards, which are uploaded to a cloud computing platform for image processing. The detection and count results (clusters and individual berries in clusters) are returned back to the App for display.

We are currently calibrating the estimated berry count and size against manual measurements and developing a correlation model to estimate cluster weights using the developed App. In addition, through PI Karkee's close collaboration with horticulturists and growers, it was identified that automated and in-field detection of lag-phase would be highly beneficial for efficient crop management. We have developed a model to detect lag-phase in vineyards, which could be integrated in the future for early season yield estimation. The model will also be integrated with the mobile App in the near future.

2. **Annual or Final Report**: Final report

3. **Project Title**: MOBILE-APP FOR CROP ESTIMATION AND LAG PHASE DETECTION

4. **Principal Investigator/Cooperator(s)**: Name, institutional affiliation, address, phone number and e-mail.

PI Name:	**MANOJ KARKEE**
Organization	CENTER FOR PRECISION AND AUTOMATED AGRICULTURAL SYSTEMS, WASHINGTON STATE UNIVERSITY (WSU CPAAS)
Address	24106 N BUNN ROAD, PROSSER, WA 99350
Telephone	509-786-9208
Email	MANOJ.KARKEE@WSU.EDU

Cooperator Name:	**NAIDU RAYAPATI AND MARKUS KELLER**	**Cooperator Name:**	**YUN ZHANG**
Organization	WSU	Organization	STE. MICHELLE WINE ESTATES
Description of participation:	SUPPORT WITH FIELD ACCESS AND VARIOUS PARAMETER ESTIMATION IN VINEYARDS	Description of participation:	GUIDANCE WITH INDUSTRY NEEDS AND ACCESS TO PLOTS FOR DATA COLLECTION AND SYSTEM EVALUATION

5. Objective(s) and Experiments Conducted to Meet Stated Objective(s):

1. Develop a spatial sampling and correlation model for block-level crop estimation
2. Develop the RGB -image-based detection model for lag-phase detection
3. Develop a viral symptom detection model
   Because of the challenges in hiring enough support staff at our lab and at the lab of our collaborators during the pandemic (first year of the project), collecting sufficient images for viral symptoms with lab-based calibration data was not possible. Based on the communication with the commission, this objective was dropped, and the budget was directed to perform additional field data collection and further validation of the block-level crop estimation under objective #1.

Experiments Conducted

Experiment 1: Obj #1 - Data Collection

Two distinct datasets were acquired during the wine grape harvest seasons of 2021 (September) and 2022 (October). In the 2021 dataset, a total of 14 vines and 42 grape clusters were labeled, and their corresponding images were captured prior to the harvest. Subsequently, post-harvest assessments were conducted to record the yield of each vine, the number of clusters per vine, the weight of the imaged clusters, and the total number of berries within each cluster. The image dataset for this study was acquired using a Samsung S9 camera in a Cabernet Sauvignon commercial vineyard. All the data collection during the harvest such as yield per vine, cluster per vine count and cluster weight measurement was performed manually.

In 2022 harvest season, a larger sample size was used, consisting of 60 labeled vines each of Chardonnay and Syrah from the Columbia Crest Vineyards, Paterson, WA. Similar to the 2021 dataset, RGB images were acquired prior to the harvest, and additional information including the yield per vine and the number of clusters per vine was recorded. Additionally, the total yield of 60 Chardonnay vines from the Washington State University (WSU) experimental vineyard located at Prosser, WA, was collected. Among these 60 vines, a subset of 15 vines had their vine canopy images captured using the same camera system, and crop yield and number of clusters per vine were recorded. This subset of data was utilized for testing the yield estimation model.

Experiment 2: Obj #1 - Smartphone Application Development

An Android application (App) was developed to acquire and upload images to Cloud for processing. The App facilitates users to acquire image through the smartphone camera as well select from the gallery. Once the images are processed, the results are received and displayed by the App to the end-users. Another App which is supported by iOS platform is also being developed.

Experiment 3: Obj #1 - Cluster number and Cluster weight Estimation Model Development

Using the dataset collected in 2021 growing season, a linear relationship between the total berry count in each cluster (manually counted) and corresponding cluster weight was established, which achieved an R^2 of 0.96. Because of this strong relationship between the berry count and cluster weight, the proposed study utilized image-based berry counting as the basis for yield estimation as presented in Fig. 1. To automate the overall yield estimation process, the number of berries visible in cluster images was counted using the berry detection method, which was then correlated to the total weight of the corresponding cluster. The total yield per vine was then estimated using the total number of clusters in each vine predicted with the visible number of clusters detected in canopy images.

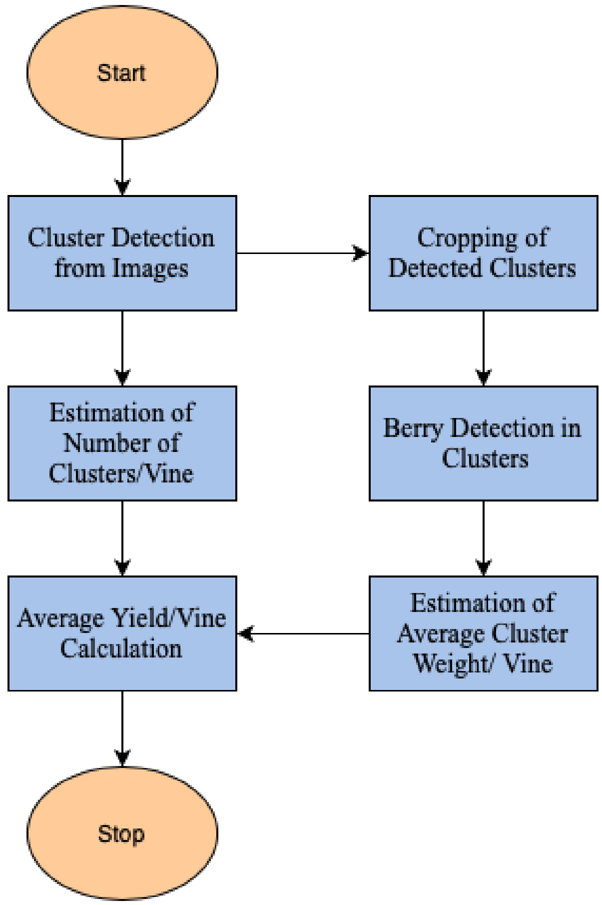

Fig. 1. Flowchart showing yield estimation in a vine level. The methodology includes image acquisition, cluster and berry detection, average cluster weight and number estimation to give average yield /vine estimation.

Experiment 4: Obj #2 - Data Collection

For tracking berry growth and developing and testing the lag-phase detection model, RGB images of grape clusters were acquired using the same mobile device and the field as mentioned before (Experiment 1). These images were collected over a growing season that included a period before, during and after the actual lag-phase occurred in the field. The images were used for detecting berries and estimating their size (diameter) over the growing season to develop a temporal growth pattern for berries captured in images from randomly selected clusters. For this experiment, 20 grape clusters from 20 different vines (10 from Chardonnay (white) variety and 10 from Merlot (red) variety) were randomly selected. Altogether 500 images (25 images per cluster) were collected throughout the growing season from July to September of the year 2022. The sample vines were selected from 4 different vine rows (5 vines per row) in the WSU experimental vineyard. A portion of checkerboard was also placed next to each cluster, which was used as a calibration reference for estimating physical berry size (mm). For evaluating the performance of the diameter/size estimation algorithm, the diameters of the 20 randomly selected berries were measured using a digital Micrometer Screw Gauge during the lag-phase period. Two distinct datasets have been utilized for accomplishing two different tasks: berry segmentation and lag-phase modeling.

Experiment 5: Obj #2 - Berry segmentation and diameter measurement algorithm development

A Mask-R-CNN (He et al., 2020) instance segmentation network was used to detect and isolate berries within grape clusters to estimate their size. The model was trained using images of individual grape clusters and a calibration object. This allowed the model to detect berries within the clusters and generate bounding boxes and instance masks for each berry. The instance mask was then used to estimate the berry diameters in pixels. In the 2022 dataset, a calibration object in the form of a checkerboard was placed alongside the grape clusters. Canny Edge detection was employed to automatically segment the checkerboard squares, which provided distinct edges. This eliminated the need for manual interaction and the use of a measuring tape, resulting in more accurate and reliable measurements of berry size. The checkerboard squares with a known size of 25 mm served as a reference for calibration, providing a consistent calibration factor in mm/pixel across different images. This approach outperformed the previous method using a measuring tape and the GrabCut algorithm (Rother et al., 2004) for estimating berry size, as reported in 2021.

Calibration factor (mm/pixel) = known physical distance (mm)/object width (in pixels)

The output of the Mask-R-CNN-based berry detection model was used to estimate the diameter of individual berries as well as average diameter of berries within a cluster, which were essential for tracking berry growth and estimating lag-phase. The model provided masks and corresponding bounding boxes for each berry. The height and width of the bounding boxes were determined, and the wider side was used as the pixel diameter of the corresponding berries. The diameter of each berry in the physical dimension (mm) was then estimated using the pixel to physical dimension calibration factor estimated using the calibration object (checkerboard for 2022 dataset), which is given by:

Average berry diameter (mm) = Diameter in pixels * Calibration factor (mm/pixel)

Experiment 6: Obj #2 - Berry growth trend tracking and Lag-phase Detection

To track berry growth, the average diameter of berries within each cluster was used, aligning with growers' practice to mitigate variations in individual berry growth. Average berry size estimation per cluster was repeated over the growing season, and the resulting diameters were plotted against sampling time to create growth trends. A polynomial curve was fitted to the scatter plot, with different degrees of complexity tested and evaluated using statistical measures like R-squared value. The polynomial model offering the best fit represented the estimated berry growth, providing insights into temporal dynamics and enabling future predictions.

From the growth trends, the lag-phase period can be observed qualitatively, which is the region with negligible or slow berry growth. Before entering the lag-phase, berries initially grow at a constant rate, which then slows down, forming a leading edge to the lag-phase. To detect the region of decreasing growth leading to no-growth or slow growth, derivatives of the berry growth model (polynomial) were calculated for each tracked cluster. The slope (first derivative) decreased after the initial constant growth, indicating slower growth near the lag phase. A change in slope direction was detected by the zero value of the second derivative of the model. The day of this change and a narrow future window (slope change: 0 to 0.1 mm/day) were used to find the day of minimum slope change, determining the Lag-phase start point. For each variety and season, statistical measures (mean and range) of the lag-phase start date were estimated and compared with ground truth data provided by the viticulture team. As growers use lag-phase information at the field level, the average berry growth pattern across all tracked grape clusters was used to estimate the lag-phase start date, which was then compared against the actual start dates.

6. Summary of Major Research Accomplishments and Results by Objective

*6.1 **Objective#1:** Develop a spatial sampling and correlation model for block-level crop estimation*

The relationship between the total number of berries detected in a cluster image and the weight of the corresponding cluster, as well as the relationship between the number of clusters observed in a vine canopy image and the total number of clusters in a vine (manually counted), were modeled using the image and ground truth data collected in both 2021 and 2022 harvest season. The R^2 value, which indicates the strength of the relationship, was calculated to be 0.22 between the number of berries and cluster weight and 0.64 for the number of visible clusters and total clusters for a WSU research vineyard (2021 Season) (see Fig. 2). This relatively low R^2, especially for the number of berries to berry weight, suggests a weak correlation between the image count and the actual count, primarily attributed to the variable shape (3D) of the clusters and different degree of occlusion of berry clusters within the vines. When this model was implemented to estimate total yield in the experimental rows used in our study, a total yield of 431.0 kg was estimated whereas the actual yield recorded in the field was 479.0 kg, which is a deviation of 10% of the actual yield. This result is promising because the measured deviation meets the deviation requirement of less than 15%. Similar experiment was conducted in a commercial vineyard, which resulted in poorer correlation models, particularly between the number of clusters seen in the images, primarily attributed to much bigger and denser vine canopies in this plot (Fig. 3). We are currently exploring the potential of using shoot length and canopy volume as additional parameters to develop more reliable correlation models.

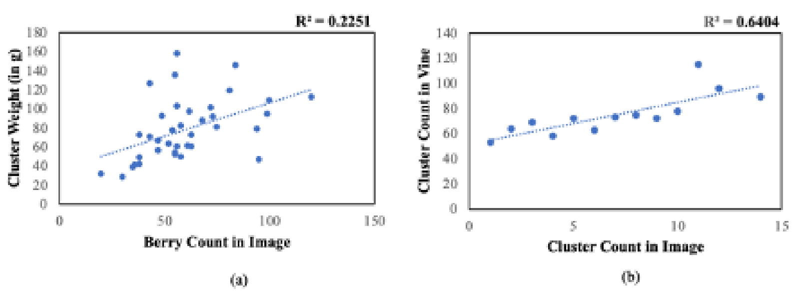

Fig. 2. Relationship between; (a) cluster weight and berry number in image; and (b) total number of clusters in a vine and clusters detected in an image. Dataset used in these models were collected in 2021 for Cabernet Sauvignon in a WSU research field.

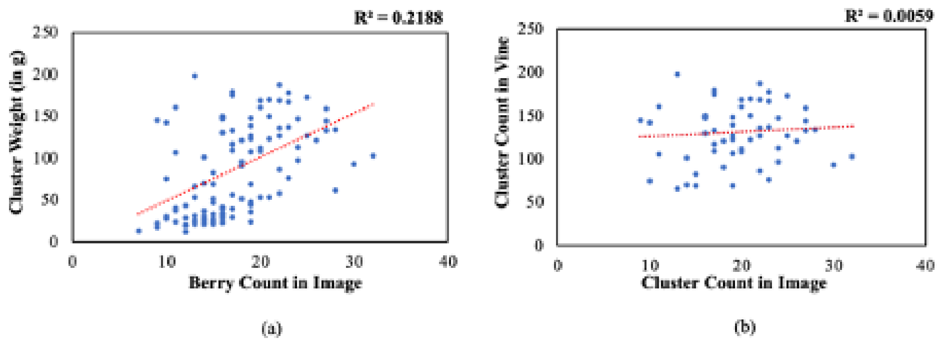

Fig. 3. Relationship between; (a) cluster weight and berry number in image; and (b) total number of clusters in a vine and clusters detected in an image. Dataset used in these models were collected in 2022 for Chardonnay in a commercial vineyard.

6.2 Objective#2: *Investigate the application of mobile App-based model for Lag-phase detection*

<u>Diameter Measurement Algorithm Evaluation:</u>

For testing the algorithm used, the ground truth data was compared to the measurements obtained from the algorithm. In the 2022 season, a checkerboard was used as the calibration object, enabling automated segmentation with a Canny edge detector, resulting in improved accuracy. These improvements yielded a reduced RMSE of 0.47 mm and a higher R^2 of 0.84 compared to the 2021 dataset (Fig. 4.). It is also noted that the berries in grape clusters are tightly packed together creating a 3-D structure with only the berries on the outer surface being completely visible. The calibration surface was placed next to the most visible berries that were utilized in berry size estimation. However, inconsistency (checkerboard not aligned to the same plane the detected berries were present) in the placement of the checkerboard next to the cluster was noted, which can reasonably be attributed to human error. Additionally, variations may occur when the camera placement is not parallel to the image plane being captured, which might have contributed, partly, to the inaccuracy in the berry size estimation.

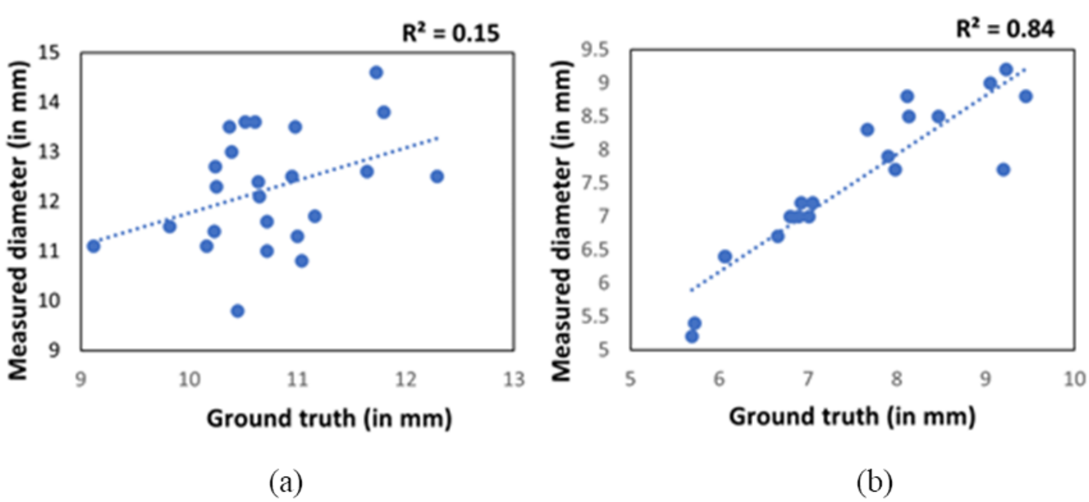

Fig. 4. Measured vs ground truth berry diameter; (a) 2021 dataset with measuring tape as the calibration surface and manual/ GrabCut segmentation; and (b) 2022 dataset with checkerboard as the calibration object with automated segmentation. Ground truth measurements were acquired using a Micrometer Screw Gauge.

Berry segmentation and Diameter measurement

A Mask-R-CNN model was developed and trained using grape canopy images collected during the 2019 growing season. It was employed for berry detection in grape clusters (Fig. 5). The model achieved a commendable Mean Average Precision (mAP) of 0.93 and Mean Average Recall (mAR) of 0.97. This indicates a reliable segmentation of berries in the images. These results indicate a low occurrence of false positives and a high accuracy in detecting the desired objects, i.e., berries, during the detection process. Similarly, calibration object was also segmented (Fig. 6) and the calibration factor was calculated.

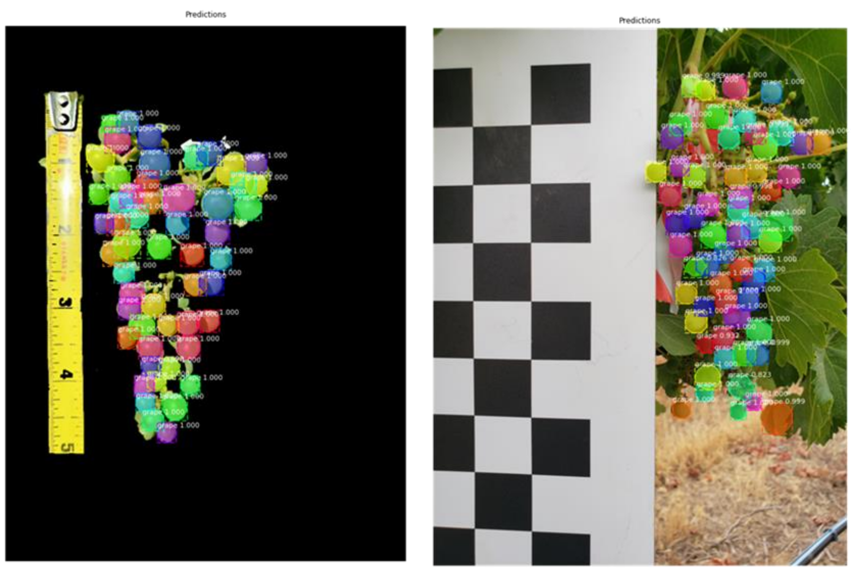

(a) (b)

Fig. 5. Example berries in a sample cluster detected using a Mask-R-CNN model; (a) 2021 dataset; and (b) 2022 dataset.

Fig. 6. Example results showing the segmentation of calibration object using Canny Edge Detection; (a) Sample image of a berry cluster with a checkerboard; (b) Segmented square from checkerboard with its physical dimension. The white boxes of the checkerboard (dimensions 25mm) are segmented out and used for calibrating the diameter of the berries from pixel dimension to physical dimension (mm).

Berry Growth Modelling

A fourth-degree polynomial provided the best fit, achieving relatively higher R^2 values for the sample clusters than when fitted with other curves. The growth trends of individual berries within the same cluster also exhibited similar patterns. Initially, there was steady growth, followed by stunted growth and finally, another period of rapid growth. This was reported in previous work (Merlot 2021 dataset). The fitness evaluation of third and fourth-degree polynomials for all clusters in each variety and year consistently favored the fourth-degree polynomial (Table 1). Moreover, the fourth-degree polynomials were better aligned qualitatively with the observed growth trends. Given the growth pattern of the berries, a more complex model beyond the fifth- degree polynomial was unnecessary and not tested.

The clusters from each variety and year were averaged to obtain a mean plot. The Merlot-2021 dataset had the lowest fit (R^2 = 0.87) due to higher error and variability in berry size estimation. In the 2022 datasets, the 4th-degree polynomial exhibited an excellent fit with an average growth-based R^2 of 0.96 for both Merlot and Chardonnay. Fitting with a third-degree polynomial resulted in lower R^2 values (0.83, 0.92, and 0.93) for Merlot-2021, Merlot-2022, and Chardonnay-2022, respectively.

Table 1. R-squared values of 3rd and 4th degree polynomials in fitting the berry growth trend in two different varieties.

	Merlot-2021		Merlot-2022		Chardonnay-2022	
Clusters	3^{rd} degree	4^{th} degree	3^{rd} degree	4^{th} degree	3^{rd} degree	4^{th} degree
1	0.61	0.62	0.84	0.85	0.79	0.80
2	0.72	0.81	0.55	0.65	0.61	0.66
3	0.61	0.62	0.76	0.86	0.83	0.83
4	0.46	0.58	0.53	0.53	0.85	0.87
5	0.56	0.62	0.86	0.86	0.66	0.70
6	0.06	0.17	0.55	0.55	0.85	0.85
7	0.27	0.32	0.67	0.74	0.55	0.61
8	0.81	0.82	0.52	0.53	0.68	0.69
9	0.62	0.62	0.68	0.69	0.80	0.80
10	0.84	0.84	0.79	0.87	0.64	0.76
Average plot R^2	0.83	0.87	0.92	0.96	0.93	0.96

Fig. 7 (a), (b), and (c) depict polynomials plots of distinct clusters within the same variety and growing season, illustrating variability in growth trends. In Fig. 7(a), steady berry growth was observed until mid-July, followed by a slowdown and resumption in early August. Similarly, Fig. 7(b) showed growth halting in early to mid-August and resuming in late August. Deviations from the general trend were present due to estimation error and biological variability, but an average estimation was used to develop a field-specific lag-phase detection method as it minimized the impact of the variability and represented the general trend.

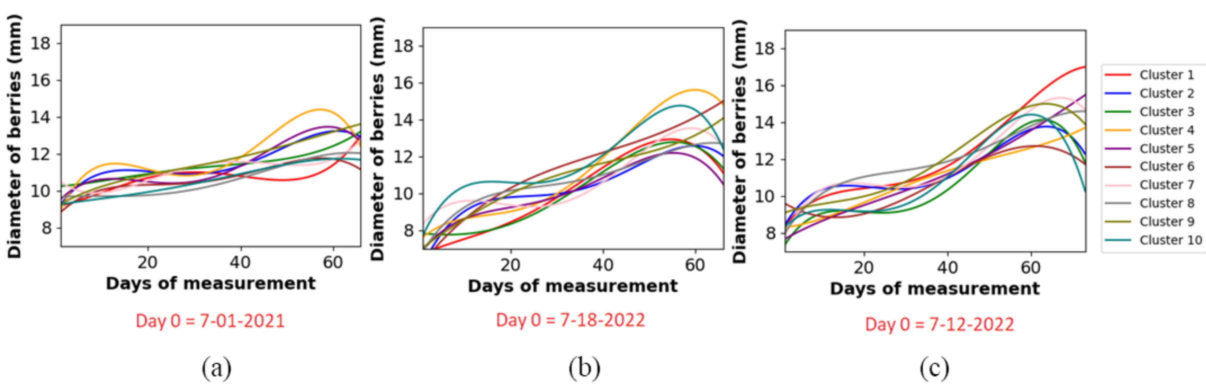

Fig. 7. (a) and (b) Growth trend of individual berries from an example cluster (Merlot-2021); (c) Growth trend of the same cluster represented by the average berry diameter.

Table 1 presents R^2 for the 4th degree polynomial modeling. The Merlot-2021 dataset had the lowest fit ($R^2 = 0.87$) due to higher error and variability in berry size estimation. In the 2022 datasets, the 4th-degree polynomial exhibited an excellent fit with an average growth-based R^2 of 0.96 for both Merlot and Chardonnay. Fitting with a third-degree polynomial resulted in lower R^2 values (0.83, 0.92, and 0.93) for Merlot-2021, Merlot-2022, and Chardonnay-2022, respectively.

Lag phase Detection

When the size of individual berries on each observation date were averaged together for all the clusters for a given year and variety as plotted in Fig. 8, the 4th-degree polynomial model represented the growth with very high R^2 of 0.96 for both Merlot-2022 and Chardonnay-2022 datasets.

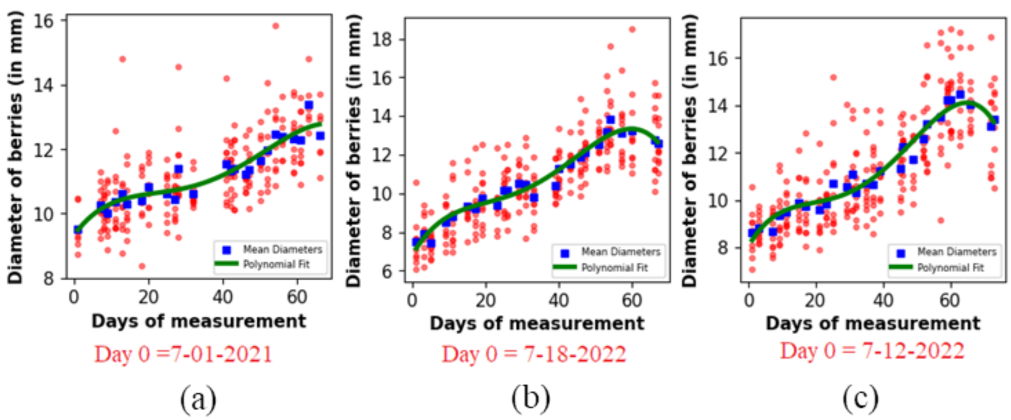

Fig. 8. Scatter plots of individual berry cluster and average berry size/diameter for; (a) Merlot-2021; (b) Merlot-2022; and (c) Chardonnay-2022. Individual berry diameters were averaged for each day the images were collected over the growing season.

Using the first and second derivatives of the polynomial model (Fig. 9), the lag-phase start date for Merlot-2022 was estimated to be August 10, slightly later than the ground truth date of August 2, 8 days earlier than the date provided by viticulturists through manual measurements. For Chardonnay-2022, it was August 2, eleven days earlier compared to the ground truth dates. With addition of Merlot 2021 data (estimated- July 26 and ground truth- July 24), the model showed a mean absolute error (MAE) of 7 days for these three dates. The lag-phase period in 2022 exhibited a shift compared to Merlot in 2021, indicating a delayed fruit set that affected both Chardonnay and Merlot varieties.

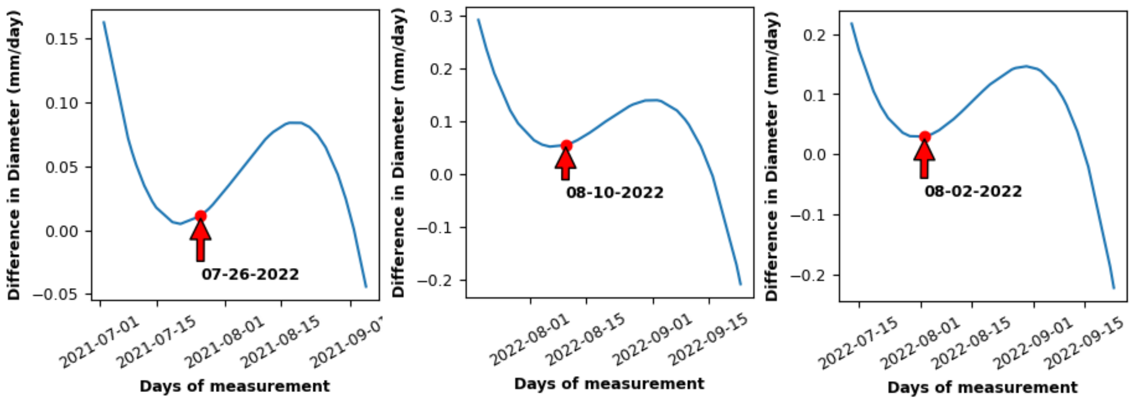

Fig. 9. Comparison of change in berry size (first derivative of the polynomial model) over the growing season; (a) 2021 Merlot; (b) 2022 Chardonnay; and (c) 2022 Merlot. The arrows on these plots show the points where the minimum slope change occurred denoting the lag-phase start date.

The lag-phase typically spans a duration of 2-3 weeks. Given the calculated MAE of 7 days in predicting the lag-phase start date, the result could be practically useful to growers as they can confidently assume the berries to be in the lag-phase about one week after the model detected date and complete any lag-phase-based sampling around that time for effective and efficient farming operations such as yield estimation. To observe the variation in growth patterns and lag-phase start date at individual plant level contributing to the overall estimate discussed in the previous paragraphs, lag-phase detection method was applied to all 10 clusters of each dataset (Merlot-2021, Merlot-2022, and Chardonnay-2022) individually (Table 2). The lag-phase start date for the entire field calculated by averaging individual cluster-level dates was found to be July 24 (standard deviation, sd–4 days), August 13 (sd–12 days) and August 18 (sd–7 days) respectively for Merlot-2021, Merlot-2022, and Chardonnay 2022, which were 7, 4 and 9 days off of the ground truth dates respectively.

7. Outreach and Education Efforts - Presentations of Research:

Strategy for communicating research results to end-users and stakeholders:

PI Karkee's group presented the results and findings in various conferences such as 15th ICPA (International Conference on Precision Agriculture) and Washington State Grape Society meeting in November 2022. Our lab also hosted growers, policymakers and other stakeholders from WA, OR, CA in our lab when we demonstrated the App developed and got feedback in terms of its applicability. Over the next several months, we plan on presenting our results and findings in various conferences such as the annual ASABE meeting in Lincoln, Nebraska. We also have submitted a research article in a peer-reviewed journal and published a conference paper in ICPA conference.

8. Research Success Statements: In a few sentences, describe in detail how your research program benefits the viticulture industry. For example, “This research has provided vintners/growers with the essential tools to control…” ***These statements are different from the project summary in that they specifically indicate how the research benefits industry rather than summarizing progress.***

In this research project, a smartphone application (App)-based system was developed to detect, locate, and count clusters in wine grape canopies. In addition, the App can provide number of berries in an image of grape cluster. The App has just been released for Beta testing by interested stakeholders. The count values obtained were then utilized to obtain the yield estimation at the field/block level. Other features such as lag-phase detection will also be added in the smartphone App soon. The results from the research show that this simple, easy to use App can provide approximate date when the lag-phase will occur in the grapes, which can be a practical tool for growers for crop-estimation in wine grapes. The App is also being expanded to estimate total yield at the field scale and to add capability for shoot detection and growth tracking.

9. Funds Status: Include a general summary of how funds were spent. (Copy of budget tables is acceptable if accurate.)

Approved By: **Samantha Bridger** Date: 01/22/2021	**% Time on Project**	**Year 1 FY**	**Year 2 FY**
		Jul 2021 – Jun 2022	Jul 2022 – Jun 2023
Personnel: Priyanka Upadhyaya	100%		
Salaries		17,870	18,585
Benefits			14,629
Wages		5,957	6,195
Benefits		596	620
Total (include all costs)		**$24,413**	**$40,029**
Total Requested		**$24,413**	**$40,029**

Total Project Request (entire project duration): $64,442

Total Project Costs (all costs, winemaking): $64,5442 *(no winemaking in this project)*

Current Year Request: $40,029

Literature Cited:

Bates, T., Grochalsky, B., & Nuske, S. (2011). Automating measurements of canopy and fruit to map crop load in commercial vineyards.

Bhattarai, U., Member, S., Bhusal, S., Zhang, Q., & Karkee, M. (n.d.). *Regression-based Network for Simplified Flower / Fruit Distribution and Count Estimation in*. 1–9.

Sabbatini, P., Dami, I., & Howell, G. S. (2012). *Predicting harvest yield in juice and wine grape vineyards*. Michigan State University Extension.

He, K., Gkioxari, G., Dollár, P., & Girshick, R. (2017). Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision* (pp. 2961-2969).

Liu, S., Marden, S., & Whitty, M. (2013, December). Towards automated yield estimation in viticulture. In *Proceedings of the Australasian Conference on Robotics and Automation*, Sydney, Australia (Vol. 24, pp. 2-6).

Interactive foreground extraction using GrabCut algorithm. OpenCV. (2021). Retrieved January 28, 2022, from https://docs.opencv.org/3.4/d8/d83/tutorial_py_grabcut.html

Nuske, S., Wilshusen, K., Achar, S., Yoder, L., Narasimhan, S., & Singh, S. (2014). Automated visual yield estimation in vineyards. *Journal of Field Robotics*, 31(5), 837-860.

Read, Paul E., and Gamet, Stephen. (2007). Crop Estimation. Lincoln: Department of Agronomy and Horticulture, Viticulture Program, University of Nebraska-Lincoln. Accessed 27 09, 2020. https://viticulture.unl.edu/publication/CROP%20ESTIMATION.pdf

Rother, C., Kolmogorov, V., & Blake, A. (2004). “GrabCut” - Interactive foreground extraction using iterated graph cuts. *ACM Transactions on Graphics*, *23*(3), 309–314. https://doi.org/10.1145/1015706.1015720