

Agtech Framework for Cranberry-Ripening Analysis Using Vision Foundation Models

Faith Johnson^{1,*}

Ryan Meegan¹

Jack Lowry¹

Peter Oudemans^{2,3}

Kristin Dana¹

Abstract

Agricultural domains are being transformed by recent advances in AI and computer vision that support quantitative visual evaluation. Using aerial and ground imaging over a time series, we develop a framework for characterizing the ripening process of cranberry crops, a crucial component for precision agriculture tasks such as comparing crop breeds (high-throughput phenotyping) and detecting disease. Using drone imaging, we capture images from 20 waypoints across multiple bogs, and using ground-based imaging (hand-held camera), we image same bog patch using fixed fiducial markers. Both imaging methods are repeated to gather a multi-week time series spanning the entire growing season. Aerial imaging provides multiple samples to compute a distribution of albedo values. Ground imaging enables tracking of individual berries for a detailed view of berry appearance changes. Using vision transformers (ViT) for feature detection after segmentation, we extract a high dimensional feature descriptor of berry appearance. Interpretability of appearance is critical for plant biologists and cranberry growers to support crop breeding decisions (e.g. comparison of berry varieties from breeding programs). For interpretability, we create a 2D manifold of cranberry appearance by using a UMAP dimensionality reduction on ViT features. This projection enables quantification of ripening paths and a useful metric of ripening rate. We demonstrate the comparison of four cranberry varieties based on our ripening assessments. This work is the first of its kind and has future impact for cranberries and for other crops including wine grapes, olives, blueberries, and maize. Aerial and ground datasets are made publicly available.

1. Introduction

Machine learning and computer vision methods play an increasingly vital role in facilitating agricultural advancement by giving real time, actionable crop feedback (27; 30; 51). These methods are enabling farming prac-



Figure 1. Cranberry bog at the measurement site. (Left) Cranberry harvesting. (Right) Drone at bog for in-field cranberry measurements during the growing season.

tices to adapt and evolve to keep up with changing conditions. Cranberry farmers are particularly poised to benefit from vision-based crop monitoring as they face numerous challenges related to fruit quality such as fruit rot and over heating (33; 36; 46; 48). As cranberries ripen and turn red, they become much more susceptible to overheating, partially because they lose their capacity for evaporative cooling (21; 37; 44). When this growth stage is reached, the cranberries exposed to direct sunlight can become unusable.

We develop a vision-based method for measuring in-field cranberry albedo to quantify ripening in order to predict when cranberries are nearing this vulnerable stage. Currently, cranberry growers quantify this ripening process manually using out-of-field albedo evaluation by imaging harvested cranberries over time (40). This approach is cumbersome and time-consuming, limiting its utility in larger-scale evaluations. For practical applications, only small numbers of berries can be harvested for out-of-field images. This is inefficient and incapable of painting a full picture of overall crop health as plants do not ripen uniformly. Berries on the top of the canopy with direct sun exposure have a high risk of overheating, while berries underneath the leafy

* Corresponding author: faith.johnson@rutgers.edu

¹ECE Department, Rutgers University

²Department of Plant Biology, Rutgers University

³ Philip E. Marucci Center for Blueberry and Cranberry Research and Extension, Rutgers University




Figure 2. An example segmentation of cranberry images using the Segment Anything Model SAM(39) without point-click prompts (automatic mask generation). Notice that in addition to cranberries, surrounding leaves and other structures are segmented.

canopy are generally well protected.

The conventional solution to overheating is increased crop irrigation during the growing season. These irrigation decisions must consider cost and efficient use of environmental resources. Furthermore, inadequate or poorly timed irrigation can lead to overheating whereas excessive irrigation encourages fungal fruit rot to develop (33). Since berries can overheat in a short period of time (21; 35), irrigation decisions should be coordinated with in-field albedo characterization, informing the grower on the number of vulnerable berries and enabling expedient decision making. Therefore, assessing the current albedo of the visible berries is directly relevant to irrigation decisions. For these reasons, in-field measurement of the ripening rate for a particular cranberry bed significantly informs crop management decisions.

For in-field measurements, we use drone imaging (see Figure 1 and ground-based imaging with standard RGB cameras (digital SLR). Our ripening assessment framework uses cranberry image segmentation to evaluate albedo variation over time to compare cranberry varieties. We conduct this albedo analysis on two spatial scales: broad area (bog-wide) and on the individual berry level. Ripening rates vary among cranberry varieties, and ones that ripen early are at the greatest risk. While surveying ripening rates over an entire bog can provide ripening statistics over the group, the appearance changes of an individual berry allows for a more detailed assessment of traits and variations of the berry ripening patterns.

In prior work (4; 6), neural networks for segmentation have been used for yield estimation through counting. In our work, we use these cranberry-tuned segmentation networks for bog-wide albedo analysis. For the individual berry analysis, we use recent foundation model segmentation networks (SAM (39)) to isolate individual berries over time to find temporal patterns in cranberry albedo across varieties. The foundation models are useful when point-



Figure 3. Example of breeding plots (drone view) that are typically evaluated manually. Planting design permits approx 3500 plots/ha, and this entire block is approximately 2 ha. Convenient quantitative evaluation can be supported by our vision-based ripening assessment framework. *Location removed for blind review.*

click prompts can be given for individual berries, while the cranberry-tuned methods are useful for large areas/bogs since they require no point-click prompts. Without point-click prompts, SAM-based segmentation segments both cranberries and leaves as shown in Figure 2.

Using the framework of imaging, photometric calibration, and cranberry segmentation, and albedo analysis for estimation of a ripening metric: we present a ripening comparison of four cranberry varieties over a two month span and show clear timelines of albedo change indicating when each variety becomes at great risk for overheating.

1.1. Impact for Crop Breeding

Screening for the heritability of novel genotypes requires high through-put phenotyping (HTP) methods to discover desirable genetic traits (8; 14). In crop breeding, there may be hundreds to thousands of progeny/offspring to evaluate, and HTP methods make this evaluation practical. Computer vision algorithms for segmentation and calibrated albedo measurements enable quantitative comparisons. The methodology we present in this paper fits those requirements well, and HTP is an application domain for this work.

The rate of color development is a crop trait that can affect the quality of cranberries at harvest. For consumer appeal, the timing and uniformity of ripening is critical, i.e. asynchronous ripening is a problem. For breeding, uniformity is desirable so HTP is used to look at multiple genotypes. For example, our related current work (unpublished) evaluates 300-400 genotypes planted in small plots (e.g. 3.3 sq.\m.) where a ripening evaluation is done out-of-field and only a few times (1-2) per season depending on time and labor. To illustrate the scale of these studies, consider that



Figure 4. Drone images from multi-temporal drone-scout imaging. Weekly inspection of multiple cranberry bogs over the late July/September growing season for four varieties (Mullica Queen, Stevens, Crimson Queen, Haines). (Left to Right) Imaging Dates for 2022: 7/27, 8/2, 8/16, 8/25, 8/31, 9/9.

they include 0.5 acre plots with 350 individual small plots and 0.2 hectare (ha) plots with 7000 individual plots (see drone image shown in Figure 3).

2. Related Work

2.1. Precision Agriculture

Precision agriculture is revolutionizing farming and challenging traditional methods. Future farms will integrate multiple technical advances such as soil sensors (2), plant wearables (50), drone aerodynamics (38), and remote sensing (43). Advances in machine learning and AI have been particularly impactful, enabling significant breakthroughs in agricultural applications in recent years (9; 28; 42; 49). Computer vision is capable of giving real time, high fidelity feedback to farmers about yield estimation (11; 18; 34; 45), phenotype identification (23; 24; 25), and crop health assessment (3; 13; 19) while also being useful for larger scale applications like farm automation (17).

2.2. Albedo Characterization over Time

As cranberries ripen, their risk of spoilage increases due to overheating (35) caused by a decrease in evapotranspiration. This ripening corresponds with visual changes in the berry albedo, which allows us to use albedo characterization over time to predict when a cranberry bog is most at risk. This same phenomena is also found in apples (37) and grapes (44). Ripening patterns can also indicate the presence of viruses as occurs in wine grapes (7; 10). Despite the importance of quantifying color development, automated methods for albedo characterization have received limited attention in the literature. Most existing studies of ripening, do out-of-field measurements that rely on harvested berries for evaluating ripening (20; 47). These methods are time-consuming and do not scale to large evaluations or real-time assessments. The framework of this paper is an important step for using computer vision methods (foundation models, deep learning networks, and classic vision methods) as a tool in agriculture.



Figure 5. Ground-based imaging (hand-held DSLR camera) of the same region in a cranberry bog was done over 27 sessions (almost daily during the growing season) to show the appearance of individual cranberries over time. The semi-permanent PVC frame enabled identification of the imaged region over time.

3. Methods

3.1. Cranberry Bog Drone Imaging

We introduce CRAID-4, a new dataset from bog monitoring with drones. We combine the CRAID-1 dataset (4) with our new drone-based cranberry bog images that we call CRAID-4. In this work, the term CRAID+ dataset refers to the combination of CRAID-1 and CRAID-4. In total, this dataset contains four different cranberry varieties over seven bogs. The cranberry variety names are *Mullica Queen*, *Stevens*, *Haines*, and *Crimson Queen*. The images include three beds of Mullica Queen, one bed of Stevens, two beds of Haines, and one bed of Crimson Queen cranberries, as shown in Figure 4. The images were taken by drone in weekly increments between the months of July and September. We calibrate each drone image and crop each into 72, non-overlapping 456×608 sub-images used for training the cranberry segmentation network (6). A selection of 220 crops representative of the diverse berry appearances in the entire growing season were manually labelled with point-wise annotations for all berries in the image. This data was combined with the labeled dataset of 2368 images from (4) to create a new training dataset comprised of 2588 total images. We train on this combined dataset of 2588 images (CRAID+). The resulting segmentation is high quality (visually assessed) and requires no point-clicks after training.

3.2. Individual Cranberry Imaging

For single berry tracking, our goal is to image *the same berry* over a time sequence to characterize ripening and appearance in a more precise manner. To this end, in addition to drone imaging, we manually captured the same 12×12 region of the cranberry bog for 27 days sampled



Figure 6. Images for photometric calibration. The drone imaging protocol includes images of the Macbeth card over multiple days. Although the same camera is used for imaging, camera parameters change. Notice slight variations in card appearance from different days, which we remove through photometric calibration.

between August 7th, 2023 to September 22nd, 2023. A semi-permanent fiducial marker was constructed as a square of PVC pipes (see Figure 5) to denote the area for image capture. The images (of size 1361×907) of this region over time are registered using SIFT (26), specifically the FLANN based feature matcher with 20 trees and 200 checks (31). We align each image in this time series to the first image of the sequence using a homography estimated with RANSAC (16). To derive the cranberry segmentations, we use SAM 2 (39) where point clicks are passed into this SAM model’s image predictor class. These point clicks are necessary for this task, instead of using the automatic mask generation, to avoid segmentation of crop leaves and shaded regions (see Figure 2). Using this approach we obtain the first time-lapse cranberry imaging series, key for computing ripening metrics for crop assessment and variety comparisons.

3.3. Photometric Calibration

The images in CRAID-4 were first photometrically calibrated using the Macbeth Color Checker card (shown in Figure 6) and a well-established approach of estimating the optimal radiometric correction using measurements of the card under uniform illumination (1; 12; 22). Radiometric or photometric calibration is needed to account for the effects of the changing camera parameters and sun angle between imaging sessions. For an invariant albedo measurement, raw pixel values from images are insufficient since they depend on camera parameters and environment conditions. Reference images of the card were taken from every bog for each day of data collection using the drone camera. For each reference image we extracted intensity values for the 6 grey scale squares on the Macbeth Color Checker. The measured values were used to find a linear transformation to recover the radiometric correction parameters, and the images were corrected accordingly.

3.4. Bog Albedo Analysis

The industry standard for ripeness defines five classes of cranberries based on albedo (40). These distinct stages of ripeness are hand-defined by field experts from the periodic collections of the cranberries throughout the season (40). We use a similar approach for classifying albedo change in the CRAID data, but opt for using in-the-field images

of the cranberries in the bog instead. In the visual analysis of cranberry images, each class is defined by k-means clustering the RGB pixel values of a collection of randomly sampled cranberry detections sampled over the entirety of the growing season ($k=5$). We then map those clusters to the 5 “common classes”, spanning from green to red, that best align with the industry standard. This human-in-the-loop mapping combines domain knowledge of the industry-standard classes with the automated clustering results.

Once the classes have been determined, we match each berry to its corresponding color class by matching each pixel belonging to a single berry to its closest color cluster. The cluster belonging to the majority of the pixel values is chosen as the label for the berry. We repeat this process for each cranberry and count the number of detections in each image. The change in class density is plotted, as in Figure 7, and clearly shows patterns in the cranberry albedo. From the progression of these plots, we also pinpoint when each cranberry variety becomes most at risk of overheating.

3.5. Individual Berry Albedo Analysis

Fourteen individual berries from one bog were tracked over a time series consisting of 27 time points (over 6 weeks). Once the individual berries were tracked using alignment and segmentation, as described in Section 3.2, visual features were extracted from each individual berry image. We employ and compare the feature vector quality of four different off-the-shelf feature extractors: DinoV2 Giant (32), Google ViT (Vision Transformer) Huge (15), SAM 2 Hiera Huge (39), and Laion CLIP Big G¹ (41). The resulting features vectors are projected to a 2D manifold using UMAP (29) (as shown in Figure 8) with the cranberry images rendered at their UMAP locations. Figure 9 shows the ViT features projected in UMAP since this feature is selected for the ripeness metric, and an individual berry’s trajectory over time is shown. This projection provides an interpretable representation useful for growers and plant biologists. Knowing where the crop currently resides in this manifold in real-time allows for expedient decision-making that can increase crop yield and overall health.

4. Results

4.1. Broad Area (Bog) Ripening Metrics

The cranberry segmentation network has a mean intersection over union (mIOU) of 62.54% and a mean absolute error (MAE) of 13.46 as reported in (5). Training on the larger CRAID+ dataset produces smaller predicted cranberry blobs than training on only the 2022 season data in CRAID-4. This may be due to a scale mismatch between the datasets. The CRAID-1 dataset contained images of cranberries taken by drone from a higher elevation than

| Bog | Ripeness Ratio | | | | | |
|-----|----------------|-------|-------|-------|-------|------|
| | 8/2 | 8/16 | 8/25 | 8/31 | 9/9 | 9/14 |
| A5 | 0.007 | 0.082 | 0.331 | 0.497 | 0.902 | 1 |
| I15 | 0.001 | 0.108 | 0.167 | 0.409 | 0.874 | 1 |
| J12 | 0.002 | 0.088 | 0.419 | 0.609 | 0.968 | 1 |
| K4 | 0.012 | 0.151 | 0.339 | 0.433 | 0.872 | 1 |
| A4 | 0.127 | 0.453 | 0.926 | 1.118 | 0.808 | 1 |
| B7 | 0.035 | 0.217 | 0.622 | 0.798 | 1.119 | 1 |
| I3 | 0.010 | 0.079 | 0.347 | 0.678 | 1.121 | 1 |

Table 1. Ripeness ratio for each bog of cranberries over time. We define the ripeness ratio for a bog of cranberries to be the percentage of red berries at the current time over the percentage of red berries on the final collection date. Bog key indicates the following cranberry types: A5 Mullica Queen, I5 Mullica Queen, J12 Mullica Queen, K4 Stevens, A4 Crimson Queen, B7 Haines, I3 Haines.

in the CRAID-4 data. Another deficiency of the model trained on the CRAID+ data is that it misses detections of greener berries. Because the CRAID-1 dataset contains mostly red berries, it is unable to make the color invariant predictions necessary to accurately segment green berries in the CRAID-4 dataset. For this reason, we use the segmentation network trained on CRAID-4 for our albedo characterization.

Once the berries are successfully segmented, their constituent pixels are matched to the closest color cluster, and each berry is labeled with the cluster that appears most frequently. We compute the classes of all the berries and plot the percentage of berries in each class for a particular collection date and bog in Figure 7. The top three rows are the Mullica Queen variety. The next row is the Stevens variety followed by the Crimson Queen variety. The final two rows are the Haines variety. Each column of graphs is made from berries imaged on a specific day. From left to right, the columns were imaged on the following dates: 8/2, 8/16, 8/25, 8/31, 9/9, and 9/14; the ripening weeks for cranberry bogs. As the berries redden, the cranberry bog enters the high risk category and the the ripeness ratio, as shown in Table 1, can be used to determine a ripeness threshold (e.g. approximately 0.6) as an indicator. This ripeness ratio is measured as the percentage of red berries (class 4 and 5) on a collection date divided by the percentage of red berries on the final collection date.

The Mullica Queen variety has a relatively low risk of overheating based on its albedo class distribution for the first four collection dates. On the fifth collection date of 9/9, the number of red berries significantly increases, indicating that the berries’ overheating risk is now high. This pattern is observed with slight variations over all three Mullica Queen cranberry beds. The Stevens variety has a majority of green berries for a significant portion of the collection period. However, by 9/9 it begins to cross over into the cat-

¹implementations from <https://huggingface.co/>



Figure 7. Plots comparing albedo over time for four cranberry varieties. Histograms of pixels in the five main color classes are shown. The four varieties are: Mullica Queen (top three rows), Stevens, Crimson Queen, and Haines (bottom two rows). Residual green pixels at the later dates are artifacts due to some misclassifications of background leaf pixels.

egory for a high risk of overheating.

The Crimson Queen variety crosses into the high risk category by 8/25. (Green albedo values in the late-season graphs are an artifact due to mis-classification of some leaf pixels as berries.) The Haines variety crosses into the high risk category on 8/25 (in the sixth row) or on 8/31 (in the seventh row).

From Figure 7, we see that the Haines variety ripens the

fastest. The next fastest ripening cranberry variety is Crimson Queen, followed by Mullica Queen. The Stevens variety is the slowest to ripen. These dates indicate a rough timeline indicating when cranberry farmers will need to monitor their crop more closely. These dates also serve as markers for when to focus more heavily on crop irrigation to mitigate overheating concerns.




Figure 8. We use four different off-the-shelf feature extractors: Laion CLIP Big G (41), SAM 2 Hiera Huge (39), Google ViT Huge (15), and DinoV2 Giant (32). Fourteen individual berries were imaged for 27 time points and the images were aligned and segmented. Feature extraction of the berry images using these four feature extractors. The resulting features vectors are projected to an interpretable representation using UMAP resulting in the embeddings illustrated here. (All 14×27 berry images are shown here at their corresponding UMAP coordinate). ViT features led to the most useful 2D embedding, showing ripeness progression in a well-distributed path close to a line.




Figure 9. (a) Segmented cranberry images projected to a 2D manifold using UMAP on ViT features. Berry images are shown in the location of the projected UMAP features. (b) An individual berry tracked over 27 timepoints (from 6 weeks) and projected to the learned 2D appearance manifold revealing its ripening path.




Figure 10. Ripeness measure for fourteen individual berries tracked and image over the growing season (27 measurements over 6 weeks). The ripeness ratio is the result of our multi-step framework consisting of: imaging, alignment, segmentation, ViT feature extraction, manifold projection.

4.2. Individual Berry Ripening Metrics

For individual berries, we follow the steps described in Sections 3.2 and 3.5. When comparing the four feature extractors shown in Figure 8, we see that ViT features led to the most useful 2D embedding showing ripeness progression in a well-distributed path close to a line. The other feature embeddings show a less clear progression from green to red, making them less easily interpretable. A ripeness metric is directly obtained by fitting the ViT-UMAP coordinates to a line and mapping initial values to zero ripeness and final values to unity. Plots of the ripeness values for the fourteen tracked individual berries are shown in Figure 10. These plots show the quantification of individual berry ripeness over time, and statistics of these ripeness metrics (e.g. mean and variance) over the bog show patterns for the cranberry variety.

5. Conclusion and Discussion

We develop an Agtech framework for evaluating cranberry ripening patterns over multiple bogs using time-series imaging (both aerial and ground-based). We show the effectiveness of using foundation models for combined with

medium scale segmentation models tuned for cranberry segmentation. Classic computer vision methods for image registration (SIFT and homography estimation) are applied for aligning the time-series images. Classic computer vision methods for basic radiometric/photometric calibration using gray-card measurements are also part of the workflow. This framework characterizes color development over time for cranberries and provides key insight into berry overheating risk and crop health. We create a timeline of albedo change that gives farmers the tools to make more informed irrigation choices to prevent crop rot and conserve resources. The resulting temporal signatures give important predictive power to the growers enabling choices among cranberry crop varieties and implications of those choices in best agriculture practices. The methodology can be automated for large scale crop evaluation to support new methods of high throughput phenotyping.

6. Acknowledgments

This project was sponsored by the USDA NIFA AFRI Award Number: 2019-67022-29922 and the NSF NRT-FW-HTF: Socially Cognizant Robotics for a Technology Enhanced Society (SOCRATES) No. 2021628.

References

- [1] Radiometric self calibration. In *Proceedings. 1999 IEEE computer society conference on computer vision and pattern recognition (Cat. No PR00149)*, volume 1, pages 374–380. IEEE, 1999. 4
- [2] Alireza Abdollahi, Karim Rejeb, Abderahman Rejeb, Mohamed M Mostafa, and Suhaiza Zailani. Wireless sensor networks in agriculture: Insights from bibliometric analysis. *Sustainability*, 13(21):12011, 2021. 3
- [3] Aanis Ahmad, Dharmendra Saraswat, and Aly El Gamal. A survey on using deep learning techniques for plant disease diagnosis and recommendations for development of appropriate tools. *Smart Agricultural Technology*, page 100083, 2022. 3
- [4] Peri Akiva, Kristin Dana, Peter Oudemans, and Michael Mars. Finding berries: Segmentation and counting of cranberries using point supervision and shape priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 50–51, 2020. 2, 4
- [5] Peri Akiva, Benjamin Planche, Aditi Roy, Kristin Dana, Peter Oudemans, and Michael Mars. Ai on the bog: Monitoring and evaluating cranberry crop risk. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2493–2502, 2021. 5
- [6] Peri Akiva, Benjamin Planche, Aditi Roy, Peter Oudemans, and Kristin Dana. Vision on the bog: Cranberry crop risk evaluation with deep learning. *Computers and Electronics in Agriculture*, 203:107444, 2022. 2, 4
- [7] Olufemi J Alabi, L Federico Casassa, Linga R Gutha, Richard C Larsen, Thomas Henick-Kling, James F Harbertson, and Rayapati A Naidu. Impacts of grapevine leafroll disease on fruit yield and grape and wine chemistry in a wine grape (*vitis vinifera* l.) cultivar. *PLoS One*, 11(2):e0149666, 2016. 3
- [8] José Luis Araus and Jill E Cairns. Field high-throughput phenotyping: the new crop breeding frontier. *Trends in plant science*, 19(1):52–61, 2014. 2
- [9] Lefteris Benos, Aristotelis C Tagarakis, Georgios Dolias, Remigio Berruto, Dimitrios Kateris, and Dionysis Bochtis. Machine learning in agriculture: A comprehensive updated review. *Sensors*, 21(11):3758, 2021. 3
- [10] Barbara Blanco-Ulate, Helene Hopfer, Rosa Figueroa-Balderas, Zirou Ye, Rosa M Rivero, Alfonso Albacete, Francisco Pérez-Alfocea, Renata Koyama, Michael M Anderson, Rhonda J Smith, et al. Red blotch disease alters grape berry development and metabolism by interfering with the transcriptional and hormonal regulation of ripening. *Journal of experimental botany*, 68(5):1225–1238, 2017. 3
- [11] Bini Darwin, Pamela Dharmaraj, Shajin Prince, Daniela Elena Popescu, and Duraisamy Jude Hemanth. Recognition of bloom/yield in crop images using deep learning models for smart agriculture: a review. *Agronomy*, 11(4):646, 2021. 3
- [12] Paul E Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In *ACM SIGGRAPH 2008 classes*, pages 1–10. 2008. 4
- [13] Vijaypal Singh Dhaka, Sangeeta Vaibhav Meena, Geeta Rani, Deepak Sinwar, Muhammad Fazal Ijaz, and Marcin Woźniak. A survey of deep convolutional neural networks applied for prediction of plant leaf diseases. *Sensors*, 21(14):4749, 2021. 3
- [14] Luis Diaz-Garcia, Brandon Schlautman, Giovanni Covarrubias-Pazaran, Andrew Maule, Jennifer Johnson-Cicalese, Edward Grygleski, Nicholi Vorsa, and Juan Zalapa. Massive phenotyping of multiple cranberry populations reveals novel qtls for fruit anthocyanin content and other important chemical traits. *Molecular Genetics and Genomics*, 293:1379–1392, 2018. 2
- [15] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 5, 7
- [16] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. 4
- [17] Othmane Friha, Mohamed Amine Ferrag, Lei Shu, Leandros Maglaras, and Xiaochan Wang. Internet of things for the future of smart agriculture: A comprehensive survey of emerging technologies. *IEEE/CAA Journal of Automatica Sinica*, 8(4):718–752, 2021. 3
- [18] Leilei He, Wentai Fang, Guanao Zhao, Zhenchao Wu, Longsheng Fu, Rui Li, Yaqoob Majeed, and Jaspreet Dhupia. Fruit yield prediction and estimation in orchards: A state-of-the-art comprehensive review for both direct and indirect methods. *Computers and Electronics in Agriculture*, 195:106812, 2022. 3
- [19] Teja Kattenborn, Jens Leitloff, Felix Schiefer, and Stefan Hinz. Review on convolutional neural networks (cnn) in vegetation remote sensing. *ISPRS journal of photogrammetry and remote sensing*, 173:24–49, 2021. 3
- [20] Markus Keller. Managing grapevines to optimise fruit development in a challenging environment: a climate change primer for viticulturists. *Australian Journal of Grape and Wine Research*, 16:56–69, 2010. 3
- [21] R Kerry, P Goovaerts, Daniel Gimenez, and PV Oudemans. Investigating temporal and spatial patterns of cranberry yield in new jersey fields. *Precision Agriculture*, 18:507–524, 2017. 1, 2
- [22] Seon Joo Kim and Marc Pollefeys. Robust radiometric calibration and vignetting correction. *IEEE transactions on pattern analysis and machine intelligence*, 30(4):562–576, 2008. 4
- [23] Shrikrishna Kolhar and Jayant Jagtap. Plant trait estimation and classification studies in plant phenotyping using machine vision—a review. *Information Processing in Agriculture*, 10(1):114–135, 2023. 3
- [24] Zhenbo Li, Ruohao Guo, Meng Li, Yaru Chen, and Guangyao Li. A review of computer vision technologies for plant phenotyping. *Computers and Electronics in Agriculture*, 176:105672, 2020. 3
- [25] Guoxu Liu, Zengtian Hou, Hongtao Liu, Jun Liu, Wenjie Zhao, and Kun Li. Tomatodet: Anchor-free detector for tomato detection. *Frontiers in Plant Science*, 13:942875, 2022. 3
- [26] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60:91–110, 2004. 4

- [27] Zifei Luo, Wenzhu Yang, Yunfeng Yuan, Ruru Gou, and Xiaonan Li. Semantic segmentation of agricultural images: A survey. *Information Processing in Agriculture*, 2023. **1**
- [28] Efthimia Mavridou, Eleni Vrochidou, George A Papakostas, Theodore Pachidis, and Vassilis G Kaburlasos. Machine vision systems in precision agriculture for crop farming. *Journal of Imaging*, 5(12):89, 2019. **3**
- [29] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018. **5**
- [30] Vishal Meshram, Kailas Patil, Vidula Meshram, Dinesh Hanchate, and SD Ramkteke. Machine learning in agriculture domain: A state-of-art survey. *Artificial Intelligence in the Life Sciences*, 1:100010, 2021. **1**
- [31] Marius Muja and David G Lowe. Fast approximate nearest neighbors with automatic algorithm configuration. *VISAPP (1)*, 2(331-340):2, 2009. **4**
- [32] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. **5, 7**
- [33] Peter V Oudemans, Frank L Caruso, and Allan W Stretch. Cranberry fruit rot in the northeast: a complex disease. *Plant Disease*, 82(11):1176–1184, 1998. **1, 2**
- [34] Fernando Palacios, Maria P Diago, Pedro Melo-Pinto, and Javier Tardaguila. Early yield prediction in different grapevine varieties using computer vision and machine learning. *Precision Agriculture*, 24(2):407–435, 2023. **3**
- [35] V Pelletier, S Pepin, J Gallichand, and J Caron. Reducing cranberry heat stress and midday depression with evaporative cooling. *Scientia horticultrae*, 198:445–453, 2016. **2, 3**
- [36] JJ Polashock, FL Caruso, PV Oudemans, PS McManus, and JA Crouch. The north american cranberry fruit rot fungal community: a systematic overview using morphological and phylogenetic affinities. *Plant Pathology*, 58(6):1116–1127, 2009. **1**
- [37] J Racsco and LE Schrader. Sunburn of apple fruit: Historical background, recent advances and future perspectives. *Critical reviews in plant sciences*, 31(6):455–504, 2012. **1, 3**
- [38] Panagiotis Radoglou-Grammatikis, Panagiotis Sarigiannidis, Thomas Lagkas, and Ioannis Moscholios. A compilation of uav applications for precision agriculture. *Computer Networks*, 172:107148, 2020. **3**
- [39] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. **2, 4, 5, 7**
- [40] Ocean Spray Rodney Serras. Private Communication, 2022. **1, 4**
- [41] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. **5, 7**
- [42] Abhinav Sharma, Arpit Jain, Prateek Gupta, and Vinay Chowdary. Machine learning applications for precision agriculture: A comprehensive review. *IEEE Access*, 9:4843–4873, 2020. **3**
- [43] Rajendra P Sishodia, Ram L Ray, and Sudhir K Singh. Applications of remote sensing in precision agriculture: A review. *Remote Sensing*, 12(19):3136, 2020. **3**
- [44] Richard E Smart and Thomas R Sinclair. Solar heating of grape berries and other spherical fruits. *Agricultural Meteorology*, 17(4):241–259, 1976. **1, 3**
- [45] Thomas Van Klompenburg, Ayalew Kassahun, and Cagatay Catal. Crop yield prediction using machine learning: A systematic literature review. *Computers and Electronics in Agriculture*, 177:105709, 2020. **3**
- [46] Nicholi Vorsa and Jennifer Johnson-Cicalese. American cranberry. *Fruit breeding*, pages 191–223, 2012. **1**
- [47] Nicholi Vorsa, Jennifer Johnson-Cicalese, Kim Patten, Cassie Bouska, and Robert Donaldson. Performance of welker, haines and other advanced selections in regional trials. 2017. **3**
- [48] Nicholi Vorsa and Juan Zalapa. Domestication, genetics, and genomics of the american cranberry. *Plant Breeding Reviews*, 43:279–315, 2019. **1**
- [49] Dashuai Wang, Wujing Cao, Fan Zhang, Zhuolin Li, Sheng Xu, and Xinyu Wu. A review of deep learning in multiscale agricultural sensing. *Remote Sensing*, 14(3):559, 2022. **3**
- [50] Heyu Yin, Yunteng Cao, Benedetto Marelli, Xiangqun Zeng, Andrew J Mason, and Changyong Cao. Soil sensors and plant wearables for smart and precision agriculture. *Advanced Materials*, 33(20):2007764, 2021. **3**
- [51] Hua Yin, Wenlong Yi, and Dianming Hu. Computer vision and machine learning applied in the mushroom industry: A critical review. *Computers and Electronics in Agriculture*, 198:107015, 2022. **1**