AI on the Bog: Monitoring and Evaluating Cranberry Crop Risk

Peri Akiva*[1]   Benjamin Planche*[2]   Aditi Roy[2]   Kristin Dana[1]   Peter Oudemans[3]   Michael Mars[3]

[1]Department of Computer and Electrical Engineering, Rutgers University
[2]Siemens Corporate Technology
[3]Department of Plant Biology, Rutgers University

{peri.akiva, kristin.dana}@rutgers.edu   {benjamin.planche, aditi.roy}@siemens.com
{oudemans, mm2784}@njaes.rutgers.edu

Abstract

Machine vision for precision agriculture has attracted considerable research interest in recent years. The goal of this paper is to develop an end-to-end cranberry health monitoring system to enable and support real time cranberry over-heating assessment to facilitate informed decisions that may sustain the economic viability of the farm. Toward this goal, we propose two main deep learning-based modules for: 1) cranberry fruit segmentation to delineate the exact fruit regions in the cranberry field image that are exposed to sun, 2) prediction of cloud coverage conditions and sun irradiance to estimate the inner temperature of exposed cranberries. We develop drone-based field data and ground-based sky data collection systems to collect video imagery at multiple time points for use in crop health analysis. Extensive evaluation on the data set shows that it is possible to predict exposed fruit's inner temperature with high accuracy (0.02% MAPE). The sun irradiance prediction error was found to be 8.41-20.36% MAPE in the 5-20 minutes time horizon. With 62.54% mIoU for segmentation and 13.46 MAE for counting accuracies in exposed fruit identification, this system is capable of giving informed feedback to growers to take precautionary action (e.g., irrigation) in identified crop field regions with higher risk of sunburn in the near future. Though this novel system is applied for cranberry health monitoring, it represents a pioneering step forward for efficient farming and is useful in precision agriculture beyond the problem of cranberry overheating.

1. Introduction

Precision agriculture aims to improve farming methods and resource management decision making through computational approaches. The long term potential of precision agriculture is in optimizing yield by real-time monitoring to precisely tune irrigation, fertilization, and other crop man-

---
* Equally contributing co-primary authors

Figure 1. (Left) Aerial view of cranberry field. (Right) Imaging drone. We develop a pipeline for cranberry crop monitoring using drone and sky imagery. Our goal is a computational approach for cranberry growers based on machine learning and computer vision to quantitatively assess risk in order to inform agriculture resource management and enable precision agriculture methods.

agement. In this work, we focus on the issue of cranberry crop risk assessment to facilitate decisions on irrigation and other crop management methods. Sunburn, caused by direct solar radiation, is evidenced by irreversible damage to the tissues of the cranberry fruit [21]. Losses to fruit in the upper part of the cranberry canopy has become increasingly common over the past 8 years, and in some years, these losses can represent upwards of 30% of the entire crop in a particular field [3]. Croft *et al*. [10] reported that cranberry scald or sunburn occurs infrequently and is limited to areas in New Jersey. However, growing conditions have changed since that study was conducted. Most importantly, it is possible that the exposure of cranberry fruit to direct sunlight has increased as cranberry yields have increased and new, high yielding, large fruited cultivars are being planted. Preliminary research suggests that the overheating of cranberry fruit is driven by solar radiation and exposed fruits may experience internal temperatures 25°C above the canopy temperature [10] (see Figure 2). Exposed fruits initially soften and then begin to degrade, with fungal fruit rot ultimately developing within.

Linking sky conditions to berry internal temperature enables a risk assessment since the crop can survive high temperatures only for short time periods before failure. The
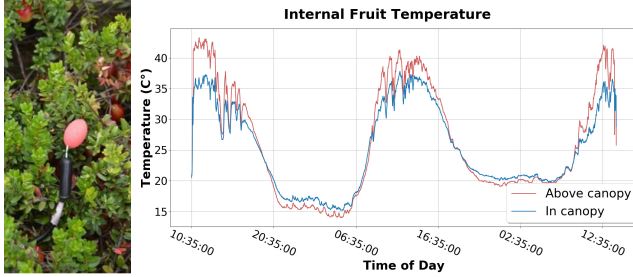
Figure 2. (Left) Synthetic cranberry temperature probe for measuring internal temperatures. (Right) Temperatures measured over time in and above cranberry canopy.

need for quick and accurate feedback is often critical to avoid loss of crop yield [47, 21, 35], which may reach up to 100% loss with sudden environmental condition changes [35]. This work shows that a combination of computer vision methods may provide an automated system to estimate yield loss and manage irrigation resources. Current yield estimation methods use previous years data or manual counting, while irrigation management relies on weather reports which are infrequent, insufficient and often inaccurate.

In this work, we develop a three part computer-vision based approach to 1) automatically segment and count exposed cranberries from down-facing drone imagery (see Figure 1) to provide near real-time estimates of crops with potential high risk of sun exposure and over-heating; 2) analyze up-facing sky imagery to determine cloud cover and cloud motion for future irradiance prediction and 3) predict internal berry temperature to assess crop risk of overheating. Segmenting and counting the cranberries enables estimation of the monetary crop value at a given time. By segmenting and counting visible cranberries, the method essentially detects cranberries with direct exposure to sunlight. The counts are then used to obtain density maps of exposed cranberries for an entire field, which may indicate to specific high risk regions that need to be attended during sudden environmental change. By predicting future irradiance with a time horizon of 15-20 minutes through sky monitoring, our method provides farmers a window of opportunity to mitigate over-heating events that may be irreversibly damaging to the fruit. To the best of our knowledge, this is the first approach linking sky observations with field-based fruit condition for its health monitoring.

2. Related Work

Precision agriculture using computer vision is growing rapidly in recent years. Drones play a significant role in providing an easy framework for efficiently imaging large-scale crop regions and current literature addresses general usage of drones in crop health monitoring [28, 30], drone models on the market [41], photogrammetric structure-from-motion point clouds of crops from drone imagery [20], automated drone navigation using vision [2]. Full 3D environment mapping methods for agriculture, including 3D

reconstruction of entire fields [14, 38, 39], often employ a collaboration of aerial and ground vehicles. While global view of fields have utility in long time-frame growth assessment, high resolution imagery to count individual cranberries can be prohibitively data-intensive over large regions. Toward the goal of real-time crop-monitoring, we adopt the efficient approach of stratified random sampling for representing large scale regions as in [40].

For vision-based analysis of crop health, concentrating on a specific crop such as coffee [45], olive [15], grapes [31] and grains [34] has practical advantages since algorithms can be tuned to expected appearance. This work for cranberry crops develops a state-of-the-art unique pipeline to process cranberry crop imagery, sky images and weather data to quantify both the amount of exposed fruit and the level of risk at a given point in time. While we focus our results on cranberries, we expect that the overall framework is relevant to multiple types of crops using transfer learning, where deep learning network parameters can be fine-tuned for use in other agriculture imagery.

Deep learning Object Counting is fundamental for quantifying cranberry exposure, and current methods mainly divide into two approaches: counting by regression, and counting by detection. Regression approaches [24, 7] tend to be faster as they often explicitly learn count from images by encoding global features similar to classification methods [51, 22]. [7] shows that while counting by regression is faster and accurate in small count scenes, performance degrades with larger count scenes. On the other hand, counting by detection implicitly learns object counts by first detecting every object and its class in the image, and then summing occurrences per class. This implies that perfect detection results in perfect count. While best performing detection methods such as R-CNN and YOLO [17, 16, 43] use expensive [5] bounding box annotation, recent approaches relying on weak supervision have shown comparable results. LC-FCN (Localization based Counting Network) [23] leverages center points annotations to detect small regions used for counting, replacing full object detection with small blobs for each object in the scene. [44] also uses point supervision with a novel Weighted Hausdorff Distance loss to count through localization, relying on an encoder-decoder network for point localization prediction, which is concatenated with the network latent features and fed to a regressor for count prediction. Triple-S Network [1] builds on both [23, 44] by simultaneously predicting count and segmentation using a shape loss prior known about the objects. This work builds on Triple-S Network [1] and [44] to generate density maps of cranberry fields useful for risk assessments and resources management.

Cloud Monitoring has been done extensively in the research community over the past few decades for a wide variety of applications, such as nowcasting, to deliver ac-
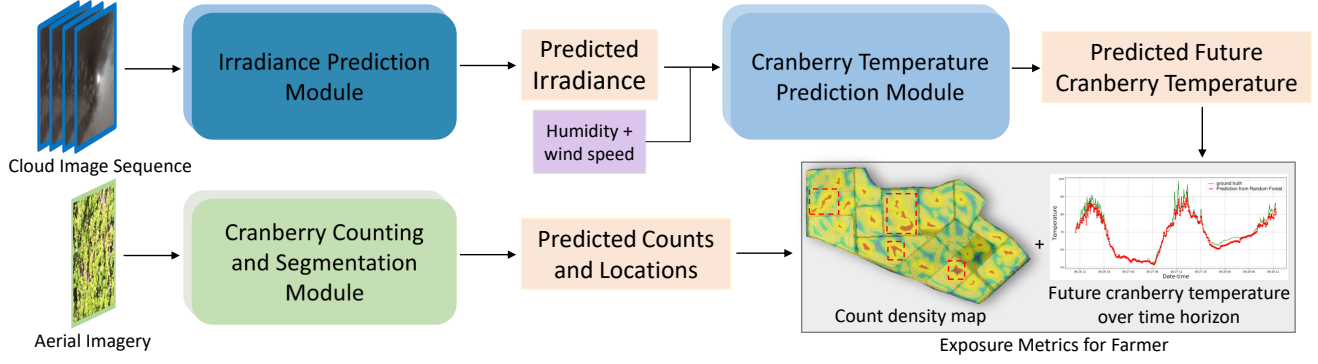
Figure 3. Pipeline overview. Cloud image sequence, humidity, and wind speed are used to predict future berry temperature over a time horizon to determine high risk time periods. Aerial imagery of cranberry crops are used to obtain count density maps to determine high risk regions. Exposure metrics (number of exposed cranberries with high berry internal temperature) are made available to the farmer in order to make resource decision such as crop irrigation. Red dashed boxes indicate high risk regions. Best viewed in color and zoomed.

curate weather forecasts [33], rainfall and satellite precipitation estimates [29], in the study of contrails [54], and various other day-to-day meteorological applications [9]. To monitor cloud dynamics in a local area and within a short time horizon of 20-30 minutes, a fish-eye camera-based system is the best choice as it provides high temporal and spatial resolution hemispherical information of the cloud cover by taking the sky pictures continuously every few seconds. From the time series of the images, it is possible to obtain a good estimate of the cloud trajectory and thus predict when and how much the sunlight would be occluded [53, 6]. To estimate cloud trajectory, we need to perform cloud segmentation and motion estimation, which are known to be closely related [6, 8, 42]. Most existing cloud segmentation techniques are based on color features [18, 11] and do not emphasize the generalization to the stochastic cloud appearance. Recent deep-learning-based popular segmentation algorithms (*e.g.*, U-Net++ [55], SegNet [4], *etc*.) are heavily dependent on finely labeled training data and seldom perform motion estimation jointly. To leverage temporal continuity of image sequences, few approaches [4, 8, 42, 56] consider optical flow prediction as an additional task apart from segmentation to make the network robust by synergy. Approaches that exploit the synergy between these two tasks have shown superior results on both tasks. Motivated by this observation, we employ a weakly supervised model for joint estimation of cloud motion and segmentation from sky image sequences.

3. Method

3.1. Problem Formulation and Pipeline

The application pipeline illustrated in figure 3 follows two branches: the berry temperature prediction branch, and the cranberry segmentation and counting branch. The input to the berry temperature prediction branch are sequences of up-facing (sky) images taken near the cranberry bog to predict future solar irradiance over that region. Irradiance

information is paired with current humidity and wind speed measurements and serves as input into the berry temperature prediction module, which provides predicted future internal berry temperature. As discussed in Section 1, internal berry temperature is a key element in determining berry spoilage, as they become irreversibly damaged beyond certain temperatures. This link between spoilage and temperature makes predicting future internal berry temperature highly valuable in preventing crop loss. Furthermore, cranberry counts provide a quantitative measure of the economic impact of the potential loss. The cranberry segmentation and counting branch also provides locations in which exposed cranberries are at present. Together, the two branches provide information to cranberry growers about high risk times and regions. This information is critical for resource decisions such as whether the cost of additional irrigation is justified by the value of the at-risk crop.

3.2. Cranberry Counting and Segmentation

This part of our application is built upon Triple-S Network's [1] count loss ablation study. Let X be an image with corresponding points P, containing cranberry (positive) and non-cranberry (negative) points noted as p_p and p_n respectively. The method feeds input X to an encoder-decoder architecture to obtain output $\tilde{Y}$ which is fed to three branches aiming to optimize segmentation, split, and count losses. The segmentation loss encourages the model towards accurate blob localization using the points as ground truth. This loss is defined as the cross entropy loss, where the target, Y, is a pixel-wise mask obtained by labeling the positive and negative points.

The split loss aims to discourage overlapping instances in segmentation prediction, and define appropriate regions for predicted blobs. The split loss branch uses the selective watershed algorithm, W_{select} [1], to obtain the set of positive and negative regions representing appropriate regions for cranberry and non-cranberry predictions. The loss function is constructed as the cross entropy loss with prediction
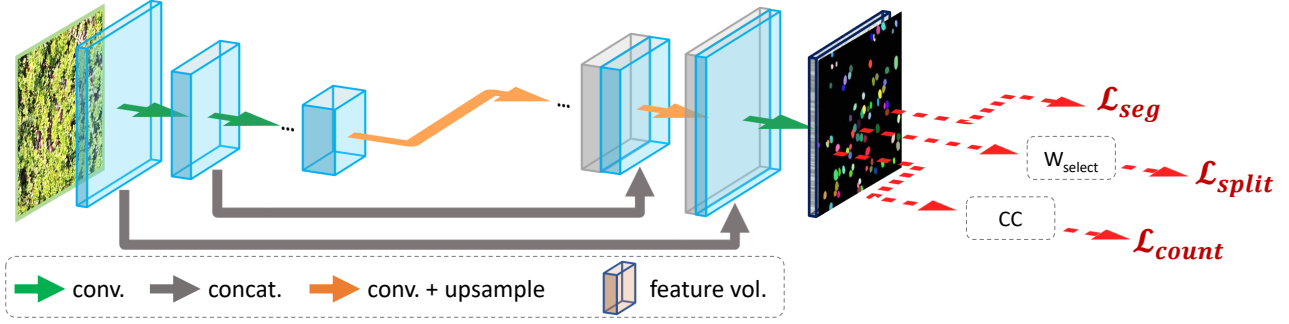
Figure 4. Triple-S Network architecture used in this application. Aerial image is input to a U-Net style network with output guided by segmentation loss, $\mathcal{L}_{seg}$, split loss, $\mathcal{L}_{split}$, and count loss, $\mathcal{L}_{count}$. W_{select} is the selective watershed algorithm introduced in [1], and CC is the connected components algorithm [13].

$W_{select}(\tilde{Y}, Y)$.

Lastly, the count loss branch aims to directly optimize training towards the correct number of instances in the input image. This branch is closely related to detection based counting but does not include a detection algorithm. Instead, it defines instances using the network's segmentation prediction, Y, to obtain separable blobs defined as instances. To that end, the connected components algorithm [13], noted as CC, is used on output Y, to obtain a set of separable blobs B. The predicted count $\tilde{C}$ is defined as the number of separable blobs B in Y, otherwise noted as the cardinality of B, $\tilde{C} = |B|$. The count loss is constructed as a Huber Loss [19] with ground truth $|p_p|$ and prediction $\tilde{C}$.

3.3. Irradiance Prediction

The goal of the second part of our system is to help predict local weather conditions that would affect the berry internal temperature, *e.g.*, solar irradiance. In this subsection, we therefore justify the importance and challenges of forecasting solar irradiance based on predicted cloud coverage and movement from local sky images.

Cloud Analysis from Sky Images. To predict current and near-future sun occlusion, cloud presence and motion should both be estimated from the sky images. We model the first task as a binary pixel-wise classification (identifying cloud and non-cloud pixels), and the second task as an optical flow (*i.e.*, 2D pixel motion map) regression. We extend a deep learning method for joint estimation of motion and segmentation from image sequences [42].

As shown in Figure 5, the baseline network [42] consists of two branches: a cloud motion estimation branch that is built on an unsupervised Siamese style recurrent spatial transformer network, and a fully-convolutional cloud segmentation branch. To leverage the synergy between their tasks, the two branches share a joint multi-scale feature encoder [37]. This encoder takes as input consecutive sky images $\{x_i \in \mathbb{R}^{H \times W \times 3}\}_{i=t}^{t+\tau}$ and samples them into one sequence of *new* frames $\{x_i\}_{i=t+1}^{t+\tau}$ and one sequence of their *preceding* frames $\{x_i\}_{i=t}^{t+\tau-1}$. Features from each consec-

utive pair (x_i, x_{i-1}) (*new* and *preceding*) are concatenated within the motion estimation branch, which then decodes the feature volume into the reverse optical flow $v_{i \to i-1}$ mapping x_i to x_{i-1}. Note that a recurrent neural layer is added to the module to refine the flow based on previous predictions. In parallel, the semantic segmentation estimates from the encoded representation of x_i its cloud pixel map m_i.

Additionally, we utilize the domain knowledge that, overall, only clouds should be moving from one frame to another in sky images, to increase the synergy between the two network branches. We multiply the predicted per-pixel optical flow values with the corresponding *softmax* cloud probabilities, virtually masking the optical flow for non-cloud pixels.

Since ground-truth semantic masks for cloud segmentation are tedious if not impossible to collect (due to fuzzy cloud contours, irregular cloud density, *etc.*), we adopt various schemes to weakly supervise the pipeline training. The joint encoder is learned by optimizing both branches simultaneously. This enables the weakly-supervised segmentation by taking advantage of features that are unsupervisedly learned in the motion estimation branch from a large amount of un-annotated data.

The overall motion estimation branch is trained in a self-supervised manner. For each pair of training images, the next frame x_i is warped by the predicted flow $v_{i \to i-1}$, and the resulting image x_i^w is compared to the preceding target frame x_{i-1} *c.f.* loss $\mathcal{L}_x^w$ in Figure 5. We also augment the baseline solution with an additional loss $\mathcal{L}_x^{w_{rev}} = \sum_{i=t+1}^{t+\tau} |x_i - x_{i-1}^{w_{rev}}|$, with $x_{i-1}^{w_{rev}}$ obtained by warping x_{i-1} with the forward flow $v_{i-1 \to i}$ predicted by switching the feature concatenation order in the motion estimation branch. This loss ensures backward/forward consistency of predicted optical flows. The segmentation branch can be trained via weak supervision (*c.f.* loss $\mathcal{L}_m^w$), as well as direct supervision if target semantic masks $\{\widetilde{m_i}\}_{i=t}^{t+\tau}$ are available (*c.f.* loss $\mathcal{L}_m$). Enforcing temporal consistency and compensating for scarce/noisy target labels, $\mathcal{L}_m^w$ com-
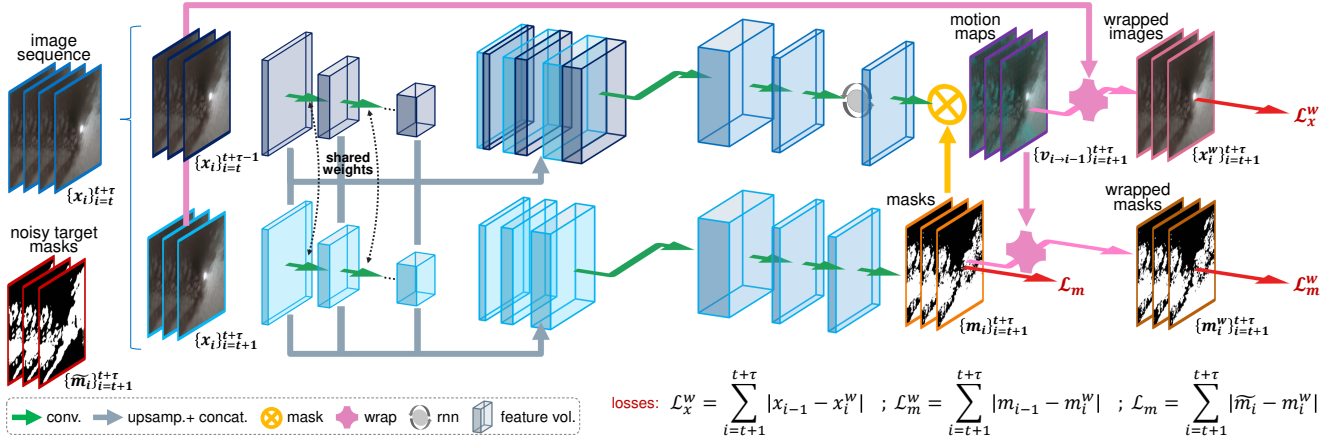
Figure 5. Training of baseline method for joint semantic segmentation and optical flow regression [42] with additional self-supervision.

$$\text{losses: } \mathcal{L}_x^w = \sum_{i=t+1}^{t+\tau} |x_{i-1} - x_i^w| \quad ; \quad \mathcal{L}_m^w = \sum_{i=t+1}^{t+\tau} |m_{i-1} - m_i^w| \quad ; \quad \mathcal{L}_m = \sum_{i=t+1}^{t+\tau} |\widetilde{m}_i - m_i^w|$$

pares m_i^w, obtained by warping predicted masks m_i using $v_{i \rightarrow i-1}$, to m_{i-1}.

Solar Irradiance Prediction. Using the cloud-related predictions, we infer the sun occlusion, and from there the irradiance, for a time interval of 1-20 minutes. To that end, we define a prediction zone in the sky images that contains the clouds, if any, that will possibly move in to occlude the sun in the considered time horizon. Like [52], we model this region of interest as a 2D band starting at the depicted sun position and extending in the opposite direction of the approximate global cloud motion, computed over the predicted motion maps (see purple rectangles in Figures 7-10); and we then apply forward extrapolation in time to estimate the occlusion probability.

The sun position s_i in 2D image space is obtained by projecting the current sun position in 3D sky coordinates obtained from "The Astronomical Almanac" [32], using the camera's extrinsic and intrinsic parameters. The global cloud motion vector $\vec{V} \in \mathbb{R}^2$ has the advantage to average out estimation errors over individual pixels and cloud deformation (which is not true motion), thus achieving a more stable estimate. We compute $\vec{V_i} = \sum_{x,y=1}^{H,W} \gamma_i(x,y) \cdot v_{i \rightarrow i+1}(x,y)$ where $\gamma_i = h(m_i) \cdot f(\|d_i\|_2) \cdot g(v_{i \rightarrow i+1}, d_i)$, with $d_i(x,y) = (s_{i,x} - x, s_{i,y} - y)$ and h, f, and g functions of choices (e.g., exponential or polynomial). Finally, we determine the clear sky irradiation curve from past observations and use it as the base for computing the irradiance from cloud probability map in the prediction zone.

3.4. Cranberry Internal Temperature Prediction

We explore a feature selection and regression framework to infer the berry internal temperature from predicted and measured data samples. The main drivers of berry internal temperature are environmental condition such as solar irradiance, humidity, wind speed and wetness [10]. In order to first rank these features w.r.t. their impact on the berry temperature, we train a random forest classifier [26] and con-figure it to return feature importance scores. From the collected dataset and results presented in Section 4, we observe that two of the top-3 features are solar irradiance and humidity. This shows the importance of predicting irradiance, since it varies frequently due to cloud movement, unlike the relatively constant humidity feature. Assuming constant humidity in the near-term (10-30 min), we can then predict future berry temperature with reasonable accuracy from the predicted future irradiance values. Finally, from this study, we complete our framework with a random forest [26] and a multilayer perceptron (MLP) [48, 49] to predict berry temperature from selected features.

4. Results

4.1. Cranberry Counting and Segmentation Results

Qualitative and quantitative results are shown in Figure 6 and Table 1 respectively for the cranberry counting and segmentation module used in this application. Quantitative results are evaluated using mean intersection over union (mIoU) for segmentation performance, and mean absolute error (MAE) for counting performance. We use the Cranberry Aerial Imagery Dataset (CRAID) [1] for training, validation and testing. The dataset provides 1022 images with center point annotations for training, and 231 images with pixel-wise annotation for evaluation. It can be seen that cranberries are correctly localized and separated into instances for accurate counting. Table 1 reports the metrics over the testing dataset for this method. In practice, each count is tied to an image location. Using a sequence of images covering an entire cranberry bog, we can generate a count density map for that bog. This provides farmers with essential information pertaining to exposed, high risk regions containing many exposed cranberries.

4.2. Irradiance Prediction Results

Cloud Data set. Rigorous evaluation of any segmentation (and motion estimation) algorithm requires annotated image data sets. Most research groups working on sky/cloud

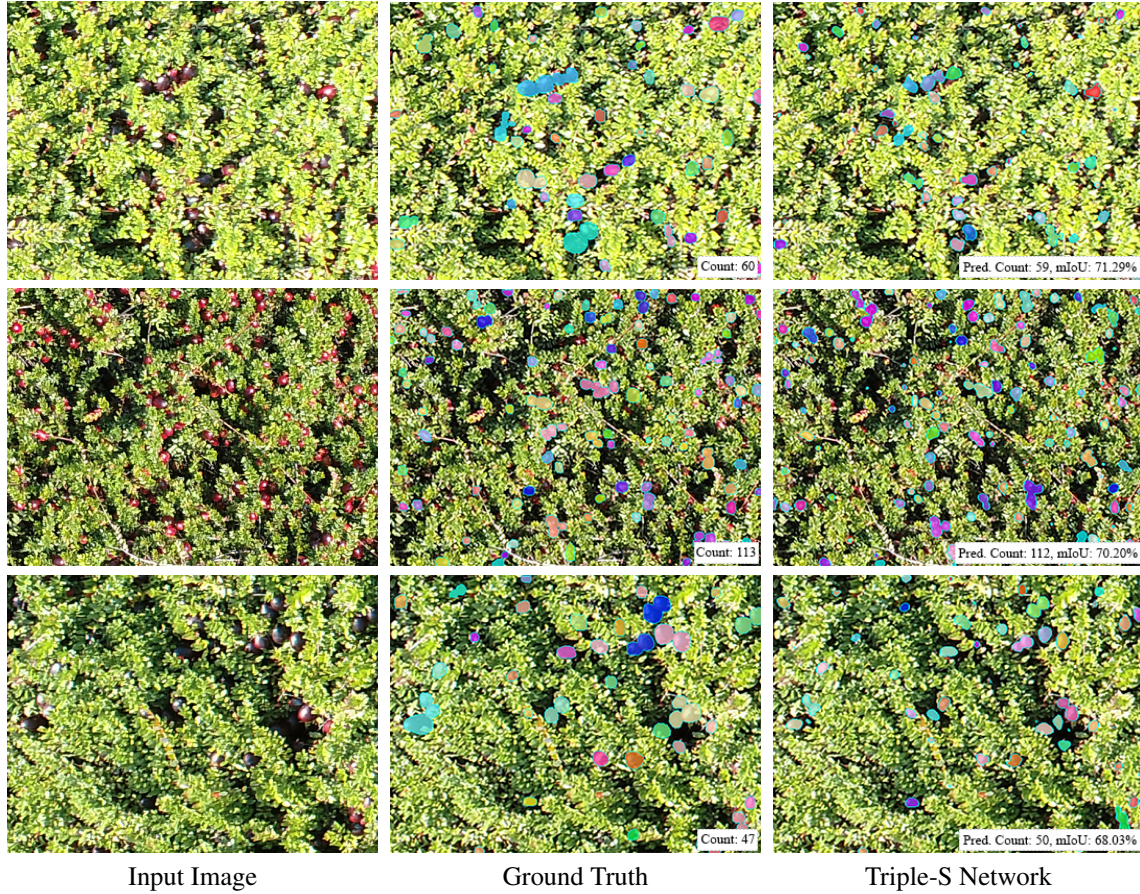| | Input Image | Ground Truth | Triple-S Network |

Figure 6. Results of cranberry counting and segmentation using Triple-S method which is integrated in this application. Best viewed in color and zoomed.

segmentation use their own proprietary images that are not publicly available. To date, only two publicly available annotated sky image data sets are available. The HYTA database [25] consists of 32 distinct images and Singapore Whole Sky IMaging SEGmentation Data set (SWIMSEG) [12] consists of 1,013 images of captured over a period of 22 months. However, both datasets are too small for training any deep learning algorithms and do not provide time-series data for motion estimation.

We collected our own dataset of hemispheric sky images ($180°/360°$) at an interval of 5 seconds spread over ~12 hrs. per day with Mobotix MX-Q24 fish eye camera. Corresponding irradiance measurements were collected using Kipp & Zonen SMP11 pyranometer in unit of $watt/m^2$. The camera is calibrated using the OcamCalib Toolbox [50] with 11 checkerboard images. The results of the calibration are camera intrinsic matrix and fisheye camera model that are used for projecting the original fisheye camera image in spherical coordinate to Cartesian coordinate for further image processing. For generating the segmentation label map, we apply standard color-based segmentation algorithm [6] which creates a noisy label data for training deep learning algorithms.

Method	MAE$^\downarrow$	mIoU (%)$^\uparrow$
U-Net [46]	18.67	60.61
LC-FCN [23]	17.46	61.97
Triple-S Network ($\mathcal{L}_{seg} + \mathcal{L}_{split} + \mathcal{L}_{count}$)	**13.46**	**62.54**

Table 1. Mean Intersection over Union (mIoU) and Mean Absolute Error (MAE) accuracies on Cranberry Aerial Imagery Dataset [1] (values: "↓" = lower is better; "↑" = higher is better).

Results. We tested our framework on a data set featuring different time of the day and month of the year. The training set was composed of image sequences and their noisy labels (~10,590 pairs) captured over the months of June and September; whereas the test set was made of sequences from May and August. Random image augmentations were performed to virtually increase the variability of the training data (random cropping, image flipping, brightness and contrast jitters) [36]. The network was trained until convergence using Adam optimizer with a learning rate of 0.001.

Figure 8 shows qualitative results on the semantic cloud segmentation task. It can be observed how, despite being trained on noisy target labels, the pipeline achieves more consistent results, thanks to the additional supervisions (through motion estimation and sequence consistency). Fig-
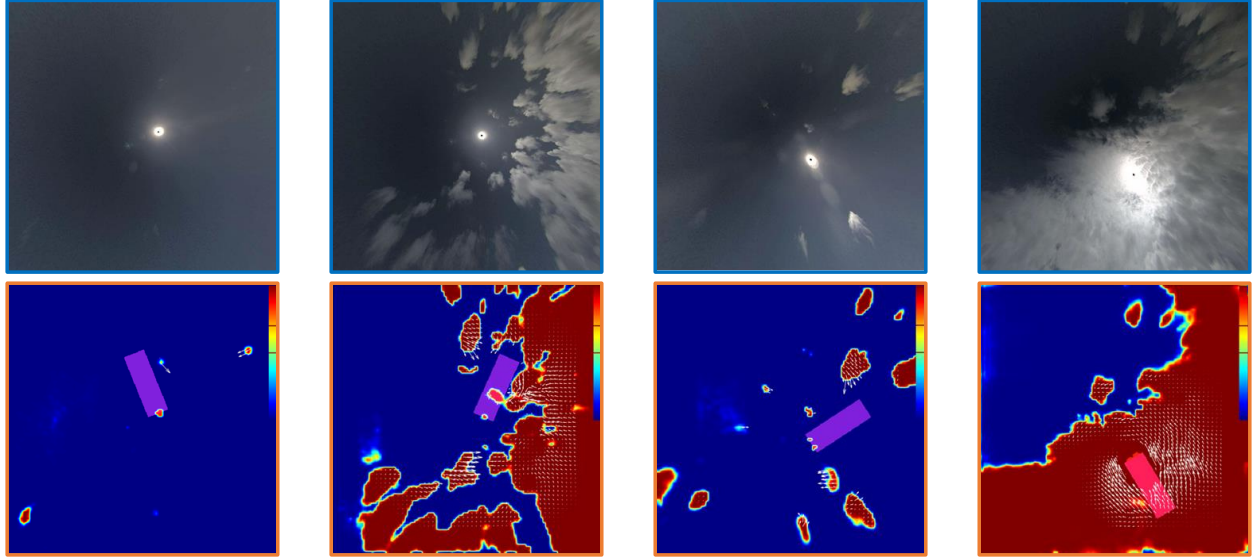
Figure 7. Results of cloud segmentation and motion estimation. Top row: original image; bottom row: output of our framework. The color gradient represents the predicted cloud probability, the arrows the predicted cloud motion, and the purple rectangle the region of interest for 20-min irradiance forecast based on sun position and cloud motion.

ure 7 shows additional results of joint segmentation and motion estimation.

Due to the lack of exact ground-truth labels for the cloud segmentation task, meaningful quantitative results can only be obtained through indirect metrics. Synchronized pyrometer data captured with the sky data was used as ground-truth irradiance to evaluate the cloud segmentation-based irradiance prediction method. Three different metrics were used to evaluate the accuracy of the prediction over a time horizon of 20 minutes, namely MAPE (Mean Absolute Percentage Error), R-squared (R^2), and Fréchet distance. The ground truth and predicted series, resp. $\widetilde{Y} = \{\widetilde{y}_i\}_{i=t}^{t+\tau}$ and $Y = \{y_i\}_{i=t}^{t+\tau}$, are normalized before comparison based on the maximum and minimum values of the series over the time period. The MAPE is defined as $\text{MAPE} = \frac{1}{\tau+1} \sum_{i=t}^{t+\tau} |\frac{\widetilde{y}_i - y_i}{\widetilde{y}_i}|$. R-squared (explained variance) is defined as $R^2 = 1 - \frac{\|\widetilde{Y} - Y\|_2}{\|\widetilde{Y} - \overline{y}\|_2}$ with $\overline{y} = \frac{1}{\tau+1} \sum_{i=t}^{t+\tau} \widetilde{y}_i$ average of ground-truth series. The Fréchet distance is a measure of similarity between curves that takes into account the location and ordering of the points along the curves. It is defined as the minimum cord-length sufficient to join a point traveling forward along predicted curve and one traveling forward along ground truth curve, although the rate of travel for either point may not necessarily be uniform. If the Fréchet distance between two curves is small, it indicates that the curves are similar. Table 2 demonstrates performance of the proposed system using these three metrics in several scenarios, with an accuracy increase attributable to our contributions w.r.t. network self-supervision. Prediction results on sampled sequences of different sky conditions are also shown in Figure 10. We can observe how cloud motion prediction (even if locally challenging due to lack of tex-
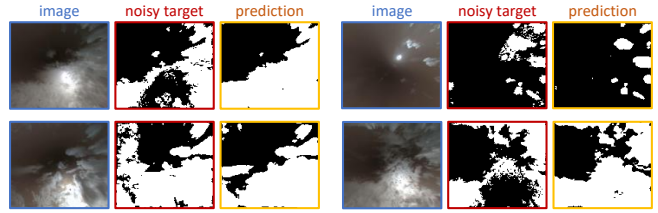

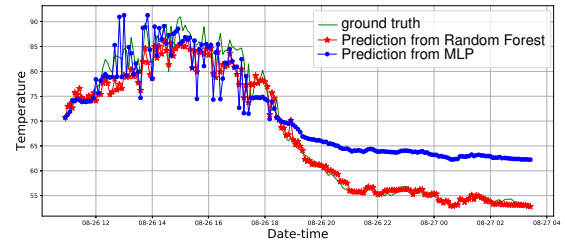Figure 8. Segmentation results on test image sequences.

Figure 9. Berry temperature prediction results using Random Forest and MLP.

ture and complex dynamics) effectively guides the irradiance forecasting in varied sun occlusion scenarios. By extracting a region of interest within each sky image based on the sun position and overall cloud motion, only detected clouds that may occlude the sun in the time horizon are considered by the irradiance prediction module as mentioned in Section 3.3.

4.3. Berry Temperature Prediction Results

We synchronously collected berry temperature and weather data over a period of 2.5 months with 5-minute sampling rate. Collected weather data contains information about ambient temperature (°F), wind speed (mile/hr), gust speed (mile/hr), wind direction (0-360°), relative humid-

Components of Pipeline for Sky Image Analysis				Irr. Forecast – 5min			Irr. Forecast – 10min			Irr. Forecast – 20min		
baseline [42]	$\mathcal{L}_x^{wrev}$	$v \odot m$	flow ref. [27]	MAPE$\downarrow$	$R^2\uparrow$	Fréchet$\downarrow$	MAPE$\downarrow$	$R^2\uparrow$	Fréchet$\downarrow$	MAPE$\downarrow$	$R^2\uparrow$	Fréchet$\downarrow$
✓				10.01	0.77	0.501	16.92	0.79	0.634	29.18	0.67	0.765
✓	✓			9.13	0.75	**0.477**	14.53	0.72	0.593	21.23	0.79	0.664
✓	✓		✓	8.86	0.92	0.479	12.89	0.76	0.602	**20.36**	0.81	**0.564**
✓		✓		8.88	0.93	0.485	13.76	**0.84**	0.596	22.52	0.84	0.655
✓		✓	✓	**8.41**	**0.95**	0.507	**12.32**	0.84	0.573	22.84	**0.86**	0.592

Table 2. Ablation study over irradiance prediction results. We compare the accuracy on the end task of irradiance forecasting for the baseline [42] and our extensions, *i.e.*, the backward/forward consistency loss $\mathcal{L}_x^{wrev}$ and the masking of motion maps v with cloud probability maps m. We also consider the refinement of motion maps using Lucas-Kanade method [27] (values: "$\downarrow$" = lower is better; "$\uparrow$" = higher is better).
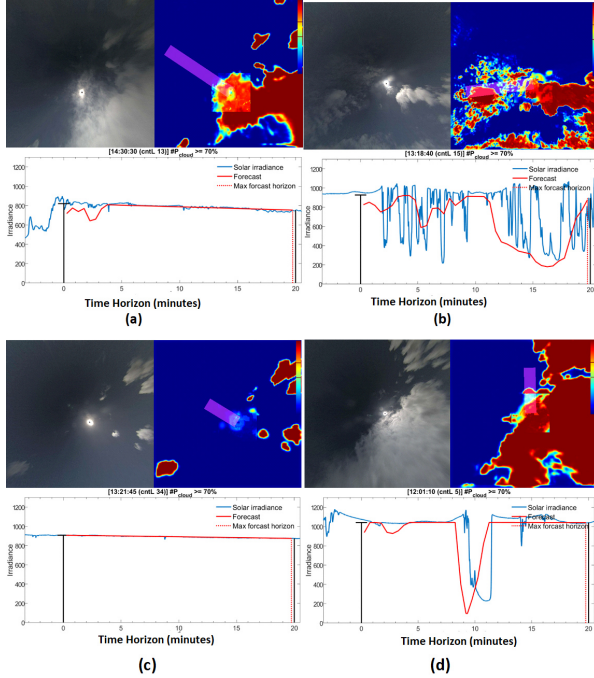

Figure 10. Irradiation prediction results in different cloud scenarios. Proposed framework could correctly match ground truth irradiation during (a) sun occlusion by thin clouds and (c) clear sky without any cloud coverage of the sun. (b) When sun occlusion happens by presence of frequent irregular clouds, our approach could closely predict the overall ground truth irradiation trend. However, minor variation in irradiation is not predicted. (d) Sudden pop-up cloud appearance was identified by the present method with limited time-difference.

ity, dew point (°F), rain (in), wetness (%), solar irradiance (w/m^2). Our current temperature prediction uses ground truth solar irradiance (measured by a pyranometer) as input, since our sky image dataset was collected from a separate location. Ground truth internal temperature data for training is obtained with a synthetic cranberry temperature probe as shown in Figure 2. Training and testing samples for temperature prediction were split 70-30 over temporally-disjoint periods, performing 5-fold cross validation. Predicted berry temperature values against ground truth using the two models are shown in Figure 9. Comparative analysis of the two machine learning models, namely random forest and MLP, using three metrics (MAE, MAPE, and R^2 defined in previ-

Method	MAPE$\downarrow$	MAE$\downarrow$	$\mathbf{R^2}\uparrow$
Random Forest	0.02	1	0.98
MLP	0.04	3	0.91

Table 3. Method comparison for berry temperature prediction.

ous sections) is shown in Table 3. It can be observed from the table as well as from the figure, that the random forest based model prediction is better than the MLP model. The features selected by the random forest method in order of importance are: solar irradiance, relative humidity, dew point, gust speed, wind speed, wind direction and wetness.

5. Discussion and Conclusions

The proposed crop-monitoring technology for small field farming achieves advantages of low cost, high efficiency and precision. By quantifying crop exposure and risk in near real-time, the proposed framework provides growers the ability to optimize resources and sustain the economic viability of the farm.

There are several major challenges. Variations in cloud appearance, cloud motion, and abrupt deformation makes cloud segmentation a challenging task. The presence of anomalies covering the camera field of view such as trees, birds, and debris, can interfere with cloud coverage prediction. For temperature prediction, training data from the synthetic cranberry temperature probe is representative of real cranberry temperature but not exact due to the probe position located at a different height from the imaged cranberries. Cranberry segmentation and counting are affected by change in berry appearance (shape, size, and albedo) and partial occlusions. To expand the technology into new fruit application areas in the future, we need to handle technological issues related to different fruit properties and acquisition of the corresponding large-scale datasets to make the system generalizable. While challenges remain, the proposed technology is a significant step towards real-time robust automatic fruit health monitoring, thus promoting the development of agricultural automation equipment and precision agriculture systems.

References

[1] 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] Bilal Hazim Younus Alsalam, Kye Morton, Duncan Campbell, and Felipe Gonzalez. Autonomous uav with vision based on-board decision making for remote sensing and precision agriculture. In *2017 IEEE Aerospace Conference*, pages 1–12. IEEE, 2017.

[3] Anonymous. Personal communication with cranberry growers, name removed for anonymous review, 2019.

[4] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 39(12):2481–2495, 2017.

[5] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What's the point: Semantic segmentation with point supervision. In *European conference on computer vision*, pages 549–565. Springer, 2016.

[6] Ti-chiun Chang, Jan Ernst, Jeremy-Ralph Wiles, Joachim Bamberger, and Andrei Szabo. Short term cloud forecast, improved cloud recognition and prediction and uncertainty index estimation, May 2019.

[7] Prithvijit Chattopadhyay, Ramakrishna Vedantam, Ramprasaath R Selvaraju, Dhruv Batra, and Devi Parikh. Counting everyday objects in everyday scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1135–1144, 2017.

[8] Jingchun Cheng, Yi-Hsuan Tsai, Shengjin Wang, and Ming-Hsuan Yang. SegFlow: Joint Learning for Video Object Segmentation and Optical Flow. *arXiv:1709.06750 [cs]*, Sept. 2017. arXiv: 1709.06750.

[9] C.I. Christodoulou, S.C. Michaelides, and C.S. Pattichis. Multifeature texture analysis for the classification of clouds in satellite imagery. *IEEE Trans. Geosci. Remote Sensing*, 41(11):2662–2668, Nov. 2003.

[10] Paul J Croft. Field conditions associated with cranberry scald. *HortScience*, 30(3):627–627, 1995.

[11] Soumyabrata Dev, Yee Hui Lee, and Stefan Winkler. Color-based Segmentation of Sky/Cloud Images From Ground-based Cameras. *arXiv:1606.03669 [cs]*, June 2016. arXiv: 1606.03669.

[12] Soumyabrata Dev, Yee Hui Lee, and Stefan Winkler. Color-based segmentation of sky/cloud images from ground-based cameras. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 10(1):231–242, 2016.

[13] Michael B Dillencourt, Hanan Samet, and Markku Tamminen. A general approach to connected-component labeling for arbitrary image representations. *Journal of the ACM (JACM)*, 39(2):253–280, 1992.

[14] Jing Dong, John Gary Burnham, Byron Boots, Glen Rains, and Frank Dellaert. 4d crop monitoring: Spatio-temporal reconstruction for agriculture. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3878–3885. IEEE, 2017.

[15] Gabriel Gatica, Stanley Best, José Ceroni, and Gastón Lefranc. Olive fruits recognition using neural networks. *Procedia Computer Science*, 17:412–419, 2013.

[16] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015.

[17] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014.

[18] A. Heinle, A. Macke, and A. Srivastav. Automatic cloud classification of whole sky images. *Atmos. Meas. Tech.*, 3(3):557–567, May 2010.

[19] Peter J Huber. Robust estimation of a location parameter. In *Breakthroughs in statistics*, pages 492–518. Springer, 1992.

[20] E Raymond Hunt Jr and Craig ST Daughtry. What good are unmanned aircraft systems for agricultural remote sensing and precision agriculture? *International journal of remote sensing*, 39(15-16):5345–5376, 2018.

[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(4):507–524, 2017.

[22] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.

[23] Issam H Laradji, Negar Rostamzadeh, Pedro O Pinheiro, David Vazquez, and Mark Schmidt. Where are the blobs: Counting by localization with point supervision. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 547–562, 2018.

[24] Victor Lempitsky and Andrew Zisserman. Learning to count objects in images. In *Advances in neural information processing systems*, pages 1324–1332, 2010.

[25] Qingyong Li, Weitao Lu, and Jun Yang. A hybrid thresholding algorithm for cloud detection on ground-based color images. *Journal of atmospheric and oceanic technology*, 28(10):1286–1296, 2011.

[26] Andy Liaw, Matthew Wiener, et al. Classification and regression by randomforest. *R news*, 2(3):18–22, 2002.

[27] Bruce D Lucas, Takeo Kanade, et al. An iterative image registration technique with an application to stereo vision, 1981.

[28] Wouter H Maes and Kathy Steppe. Perspectives for remote sensing with unmanned aerial vehicles in precision agriculture. *Trends in plant science*, 24(2):152–164, 2019.

[29] Majid Mahrooghy, Nicolas H. Younan, Valentine G. Anantharaj, James Aanstoos, and Shantia Yarahmadian. On the Use of a Cluster Ensemble Cloud Classification Technique in Satellite Precipitation Estimation. *IEEE J. Sel. Top. Appl. Earth Observations Remote Sensing*, 5(5):1356–1363, Oct. 2012.

[30] UM Rao Mogili and BBVL Deepak. Review on application of drone systems in precision agriculture. *Procedia computer science*, 133:502–509, 2018.

[31] Stephen Nuske, Supreeth Achar, Terry Bates, Srinivasa Narasimhan, and Sanjiv Singh. Yield estimation in vineyards by visual grape detection. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2352–2358. IEEE, 2011.

[32] US Naval Observatory. Astronomical almanac, 1985.

[33] C. Papin, P. Bouthemy, and G. Rochard. Unsupervised segmentation of low clouds from infrared METEOSAT images based on a contextual spatio-temporal labeling approach. *IEEE Trans. Geosci. Remote Sensing*, 40(1):104–114, Jan. 2002.

[34] Diego Inácio Patrício and Rafael Rieder. Computer vision and artificial intelligence in precision agriculture for grain crops: A systematic review. *Computers and electronics in agriculture*, 153:69–81, 2018.

[35] V Pelletier, S Pepin, J Gallichand, and J Caron. Reducing cranberry heat stress and midday depression with evaporative cooling. *Scientia horticulturae*, 198:445–453, 2016.

[36] B Planche and E Andres. Hands-on computer vision with tensorflow 2, 2019.

[37] Benjamin Planche, Sergey Zakharov, Ziyan Wu, Andreas Hutter, Harald Kosch, and Slobodan Ilic. Seeing beyond appearance-mapping real images into geometrical domains for unsupervised cad-based recognition. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2579–2586. IEEE, 2019.

[38] Ciro Potena, Raghav Khanna, Juan Nieto, Daniele Nardi, and Alberto Pretto. Collaborative uav-ugv environment reconstruction in precision agriculture. In *Proceedings of the IEEE/RSJ IROS Workshop" Vision-based Drones: What's Next*, 2018.

[39] Ciro Potena, Raghav Khanna, Juan Nieto, Roland Siegwart, Daniele Nardi, and Alberto Pretto. Agricolmap: Aerial-ground collaborative 3d mapping for precision farming. *IEEE Robotics and Automation Letters*, 4(2):1085–1092, 2019.

[40] Larisa Pozdnyakova, Daniel Giménez, and Peter V Oudemans. Spatial analysis of cranberry yield at three scales. *Agronomy Journal*, 97(1):49–57, 2005.

[41] Vikram Puri, Anand Nayyar, and Linesh Raja. Agriculture drones: A modern breakthrough in precision agriculture. *Journal of Statistics and Management Systems*, 20(4):507–518, 2017.

[42] Chen Qin, Wenjia Bai, Jo Schlemper, Steffen E. Petersen, Stefan K. Piechnik, Stefan Neubauer, and Daniel Rueckert. Joint Learning of Motion Estimation and Segmentation for Cardiac MR Image Sequences. *arXiv:1806.04066 [cs]*, June 2018. arXiv: 1806.04066.

[43] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.

[44] Javier Ribera, David Guera, Yuhao Chen, and Edward J Delp. Locating objects without bounding boxes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6479–6489, 2019.

[45] Jhonn Pablo Rodríguez, David Camilo Corrales, Jean-Noël Aubertot, and Carlos Corrales. A computer vision system for automatic cherry beans detection on coffee trees. *Pattern Recognition Letters*, 2020.

[46] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.

[47] Teryl Roper. The physiology of cranberry yield. *Wisconsin Cranberry Crop Management Newsletter Vol. XIX*, 2006.

[48] Frank Rosenblatt. Principles of neurodynamics. perceptrons and the theory of brain mechanisms. Technical report, Cornell Aeronautical Lab Inc Buffalo NY, 1961.

[49] Dennis W Ruck, Steven K Rogers, and Matthew Kabrisky. Feature selection using a multilayer perceptron. *Journal of Neural Network Computing*, 2(2):40–48, 1990.

[50] Davide Scaramuzza, Agostino Martinelli, and Roland Siegwart. A toolbox for easily calibrating omnidirectional cameras. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5695–5701. IEEE, 2006.

[51] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[52] Shanhui Sun, Jan Ernst, Joachim Bamberger, and Jeremy Ralph Wiles. Sun location prediction in image space with astronomical almanac-based calibration using ground based camera, Oct. 22 2015. US Patent App. 14/711,002.

[53] Shanhui Sun, Jan Ernst, Archana Sapkota, Eberhard Ritzhaupt-Kleissl, Jeremy Wiles, Joachim Bamberger, and Terrence Chen. Short term cloud coverage prediction using ground based all sky imager. In *2014 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pages 121–126, Nov. 2014.

[54] J.M. Weiss, S.A. Christopher, and R.M. Welch. Automatic contrail detection and segmentation. *IEEE Trans. Geosci. Remote Sensing*, 36(5):1609–1619, Sept. 1998.

[55] Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. Unet++: A nested u-net architecture for medical image segmentation. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, pages 3–11. Springer, 2018.

[56] Yi Zhu, Karan Sapra, Fitsum A. Reda, Kevin J. Shih, Shawn Newsam, Andrew Tao, and Bryan Catanzaro. Improving Semantic Segmentation via Video Propagation and Label Relaxation. *arXiv:1812.01593 [cs]*, July 2019. arXiv: 1812.01593.