Assessing produce freshness using hyperspectral imaging and machine learning

Riley D. Logan,[a] **Bryan Scherrer**,[a] **Jacob Senecal**,[b] **Neil S. Walton**,[b] **Amy Peerlinck**,[b] **John W. Sheppard**,[b] **and Joseph A. Shaw**[a,*]

[a]Montana State University, Electrical and Computer Engineering Department and Optical Technology Center, Bozeman, Montana, United States

[b]Montana State University, Gianforte School of Computing, Bozeman, Montana, United States

Abstract. A method of monitoring produce freshness with hyperspectral imaging and machine learning is described as a way to reduce food waste in grocery stores. The method relies on hyperspectral reflectance images in the visible–near-infrared spectral range from 387.12 to 1023.5 nm with a 2.12-nm spectral resolution. The images were recorded in a laboratory with the imager viewing produce samples illuminated by broadband halogen lights, but we also recorded and discussed the implications of the illumination spectrum of lights found in a variety of grocery stores. A convolutional neural network was used to perform freshness classification for potatoes, bananas, and green peppers. Additionally, a genetic algorithm (GA) was used to determine the wavelengths carrying the most useful information for age classification, with an eye toward a future multispectral imager. Hyperspectral images were processed to explore the use of RGB images, GA-selected multispectral images, and full-spectrum hyperspectral images. The GA-based feature selection method outperformed RGB images for all tested produce, outperformed hyperspectral imagery for bananas, and matched hyperspectral imagery performance for green peppers. This feature selection method is being used to develop a low-cost multispectral imager for use in monitoring produce in grocery stores. [DOI: 10.1117/1.JRS.15.034505]

Keywords: food waste; food quality; food safety; hyperspectral imaging; machine learning.

Paper 210136 received Mar. 8, 2021; accepted for publication Jul. 2, 2021; published online Jul. 15, 2021.

1 Introduction

Improved methods of monitoring produce freshness could help reduce the enormous problem of food waste. Recent estimates suggest that up to one third of food produced for human consumption is wasted,[1–3] at a cost of up to $198B annually[2,3] with corresponding increases of greenhouse gas emissions[4,5] and wasted consumption of freshwater and oil.[5] Much of the food waste in the United States occurs at retail stores and with consumers.[6] For example, unsold or discarded produce imposes a cost of approximately $15B annually at grocery stores, largely due to produce becoming overripe before being sold.[3] Local grocery store employees have told us that damaged or overripe produce is currently identified through visual inspection, a method which is not well suited to early prediction, which could enable retail strategies for rapidly selling produce before it is too ripe. In this paper, we present a method for detecting and predicting the ripening process of fruits and vegetables using hyperspectral imaging coupled with machine learning algorithms.

There is a substantial body of work in using optical sensing for non-destructive food safety and quality analysis.[7–11] Hyperspectral imaging is particularly useful because it provides both spatial and spectral information, the latter being especially well suited for monitoring produce ripening driven by chemical processes such as chlorophyll degradation, respiration changes, carotenoid biosynthesis, and ethylene production changes.[12] The addition of near-infrared

*Address all correspondence to Joseph A. Shaw, joseph.shaw@montana.edu

1931-3195/2021/$28.00

(NIR) wavelengths can improve ripeness monitoring by going beyond the visible pigmentation changes caused by varying chlorophyll content and carotenoid accumulation. For example, hyperspectral imaging has been used to monitor banana quality, with the most information provided by eight visible and NIR wavelengths.[13]

Machine learning algorithms have been applied to RGB images for classification and quality grading of fruits and vegetables[14–16] and ripeness monitoring of bell peppers[17] and gooseberries.[18] Convolutional neural networks (CNNs) have been employed for applications that include fruit and vegetable classification with low-cost cameras and Raspberry PI computers[19] and with smartphone images for detecting artificially ripened fruit.[20] Machine learning is particularly useful for analyzing the large quantities of spectral data generated by hyperspectral imagers or spectrometers. Applications include detecting fungi that cause citrus fruit to rot,[21] identifying mechanical damage to mangoes,[22] classifying produce type,[23] and evaluating tomato ripeness.[11,24]

Building on our prior work in agricultural remote sensing,[25,26] this paper describes a methodology for using hyperspectral imaging and machine learning for monitoring and classifying freshness of fruits and vegetables commonly available in grocery stores, focusing on bananas, potatoes, tomatoes, avocados, and bell peppers. A genetic algorithm (GA) was used to select optimal spectral channels in preparation for development of a low-cost multispectral imager.[27] The goal of our method is to determine the relative age of the produce as well as classify whether the produce is fresh, thus decreasing the amount of discarded food.

This paper is an expanded discussion of a work initially reported in an SPIE Proceedings conference paper.[28] In accordance with the SPIE publication policy,[29,30] we expanded the scope and depth but also reused figures and text from the conference paper with this acknowledgment. Compared with the earlier conference paper, we added to this paper experiments focused on predicting the age of the produce since purchase, a discussion of the different light source spectra found in grocery stores and the expected impacts of this on in-store measurements, and an expanded discussion of the GA used to select optimal wavelengths from hyperspectral images.

2 Methods

The methods used throughout this work are introduced by an overview of our experimental procedure, beginning with a broad overview of the processing steps. We then discuss the imaging system and light sources used to gather image data, the experimental setup and testing procedure, image preprocessing, and final processing steps.

2.1 *Process*

The experimental process used throughout this work is illustrated as a flowchart in Fig. 1. This processing flowchart serves as a map of our experimental methods, beginning with our imaging system and illumination techniques.

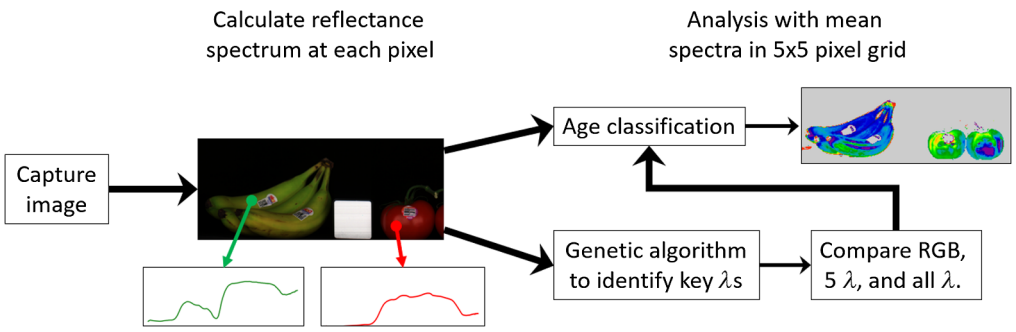

Fig. 1 Process flowchart used to outline our experimental methods.

2.2 *Imaging System and Light Sources*

All image data were collected using a Resonon, Inc., visible–near-infrared (VNIR) Pika L hyperspectral imaging system across a spectral range of 387.12 to 1023.50 nm with a spectral resolution of approximately 2.12 nm, producing 12-bit hyperspectral data cubes with 300 spectral channels. We assessed the polarization response of the Pika L imager, validated its spectral and pixel-to-pixel performance, and analyzed its image-to-image stability.

2.2.1 *Imager characterization*

As many imaging systems are partially sensitive to polarization[31] and glossy produce surfaces can create partially polarized reflected light, we measured the polarization response of the Pika L imager.[32] After imaging the uniform, randomly polarized output of an integrating sphere through a linear polarizer (Meadowlark Optics VLM-129-UV-C) rotated through 180 deg in steps of 2 deg, we calculated the degree of linear polarization and determined the polarization response to be relatively small, with a maximum value of 0.073 (7.3%) at 744 nm. This response, coupled with the weakly polarized light from the majority of the produce surface, convinced us that no polarization response compensation was necessary for this study.

A spectral calibration was performed by sweeping a monochromator (Acton Research Corporation SpectraPro-150) from 400 to 1000 nm in 10-nm steps, collecting both a hyperspectral data cube and a spectrometer measurement with an Ocean Optics USB4000 at each step. We found that the wavelengths from the monochromator, the spectrometer, and the Pika L all matched within a measurement uncertainty of 0.1 nm.

Finally, the uniformity of pixel response across the Pika L's detector array was characterized by imaging the spatially uniform integrating sphere output while scanning the imager across the aperture of the sphere. As each of the line scans captured across the aperture of the integrating sphere responded with the same pixels on the imager's detector, we averaged pixel values across the scanned spatial dimension, resulting in a 900 spatial pixel × 300 spectral pixel image. Since the integrating sphere did not have a flat spectral radiance output, each of the spectral channels captured by the Pika L had a unique response. Therefore, for each channel, we defined a non-uniformity deviation metric,

$$\text{Non-uniformity deviation} = \frac{\sigma}{\mu}, \tag{1}$$

where σ is the standard deviation and μ is the mean.

Using 36 images of the spatially uniform output of the integrating sphere, we calculated the mean and standard deviation of the 900 spatial pixels for each image across six spectral channels. From the resulting 36 means and standards deviations for each spectral channel, we calculated the average value of both the mean and standard deviation. Table 1 shows the spatial

Table 1 Spatial non-uniformity deviation of the Pika L hyperspectral imager.

Wavelength (nm)	Uncalibrated non-uniformity deviation %	Calibrated non-uniformity deviation %
449.6	1.49	0.15
550.0	1.42	0.11
650.3	1.52	0.11
750.4	1.62	0.11
850.2	2.03	0.16
949.6	2.67	0.33

non-uniformity deviation of the six measured spectral channels, with maximum values of 2.67% before application of a radiometric calibration and 0.33% after radiometric calibration (performed by relating the measured digital numbers to known brightness of an integrating sphere source with spatial uniformity of approximately 1% as reported by the manufacturer). The measurements in this paper were converted to reflectance after subtracting a dark signal from each image using the SpectrononPro software. The spatial non-uniformity establishes an upper bound on the uncertainty of reflectance values determined for the produce pixels in a ratio with the Spectralon panel pixels (see Fig. 1 and discussion in Sec. 2.4). The results in Table 1 show that implementing a radiometric calibration before computing reflectances could further reduce the uncertainty.

2.2.2 *Light sources*

Not all light sources produce light over the Pika L imager's spectral range of 400 to 1000 nm. For example, fluorescent and LED lighting commonly found in grocery stores provides only a small portion of this spectrum. To alleviate this issue for laboratory studies, we illuminated our samples with two 500-W halogen bulbs mounted behind diffusing screens. Each source was positioned at ± 30 deg from the imaging axis, resulting in approximately uniform illumination of the test samples (Fig. 2). We converted the raw digital numbers captured by the imaging system to reflectance by placing two 5 cm $\times$ 5 cm, 99% reflective Spectralon reference plates (Labsphere SRT-99-020) next to the produce. Further details are discussed in Sec. 2.4.

2.3 *Experimental Setup and Testing Procedure*

Measurements were made with produce placed on a low-reflectance "produce stage," beside two reflectance calibration targets and scanning the imager through a full angle of approximately 24 deg. We maintained consistent camera settings throughout the measurement cycle, with a gain of 1, frame rate of 30 fps, and exposure time of 10 ms using SpectrononPro software from Resonon, Inc. The low-reflectance stage reduced stray light and multiple reflections within the stage.

We began with fresh produce obtained from a local grocery store and ended each measurement cycle when the produce was visibly inedible (i.e., rot or mold present), taking one measurement every 24 hours. Store employees informed us that produce was restocked from distribution vehicles every 1 to 3 days, meaning fresh test produce was easily obtained. We selected test samples from the most commonly available items in produce departments. In this paper, we discuss measurements of bananas, bell peppers, potatoes, avocados, and tomatoes. We chose test produce by selecting the least-ripe options readily available on grocery store displays, judging by appearance, texture, firmness, shape, and smell. Once purchased, we placed the test produce on the produce stage in our temperature-controlled laboratory,

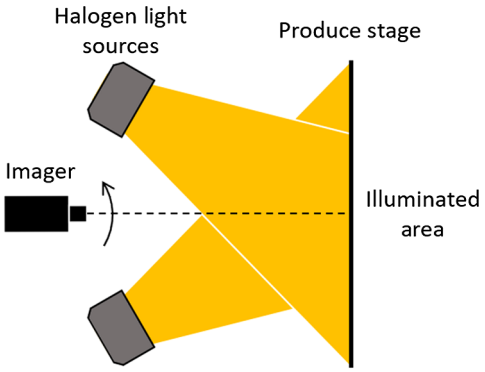

Fig. 2 Overview of experimental setup showing halogen light orientation, imager position and scanning direction, and illuminated stage on which produce and reflectance targets were placed.

held at an average ambient temperature of 21.7°C and average relative humidity of 27.5%, throughout the course of our measurements. Once positioned on the produce stage, the test produce was not allowed to shift position until the full measurement cycle was complete. By not moving the samples, we obtained measurements of the same area on the test produce throughout the entire observation period.

2.4 *Preprocessing*

The preprocessing pipeline used in this work involved removing the dark signal from the hyperspectral data cube, converting the data cube to reflectance, and isolating produce pixels in each data cube. After subtracting the mean dark frame, the raw data cube was converted to reflectance, found as the ratio of signals from the produce pixel signals and the Spectralon pixels. This required an assumption of spatial uniformity, which had been previously confirmed within 2.67%. Reflectance values were found at each produce pixel from

$$\rho_{\text{scene}} = \left(\frac{\text{DN}_{\text{scene}} - I_{\text{dark}}}{\text{DN}_{\text{target}} - I_{\text{dark}}}\right)\rho_{\text{target}}, \tag{2}$$

where ρ_{scene} is the reflectance at each pixel, ρ_{target} is the reflectance of the reference target, DN_{scene} is the digital number recorded at each pixel, $\text{DN}_{\text{target}}$ is the digital number recorded on the reflectance target, and I_{dark} is the dark signal. This calculation was performed at each spatial location, resulting in a reflectance spectrum from 400 to 1000 nm at each spatial location in the imaged scene (Fig. 3).

In the final preprocessing step, we manually identified and isolated the test produce in each data cube. By isolating the test produce, we created a training set for our classification and age analysis algorithms discussed in Sec. 2.5.

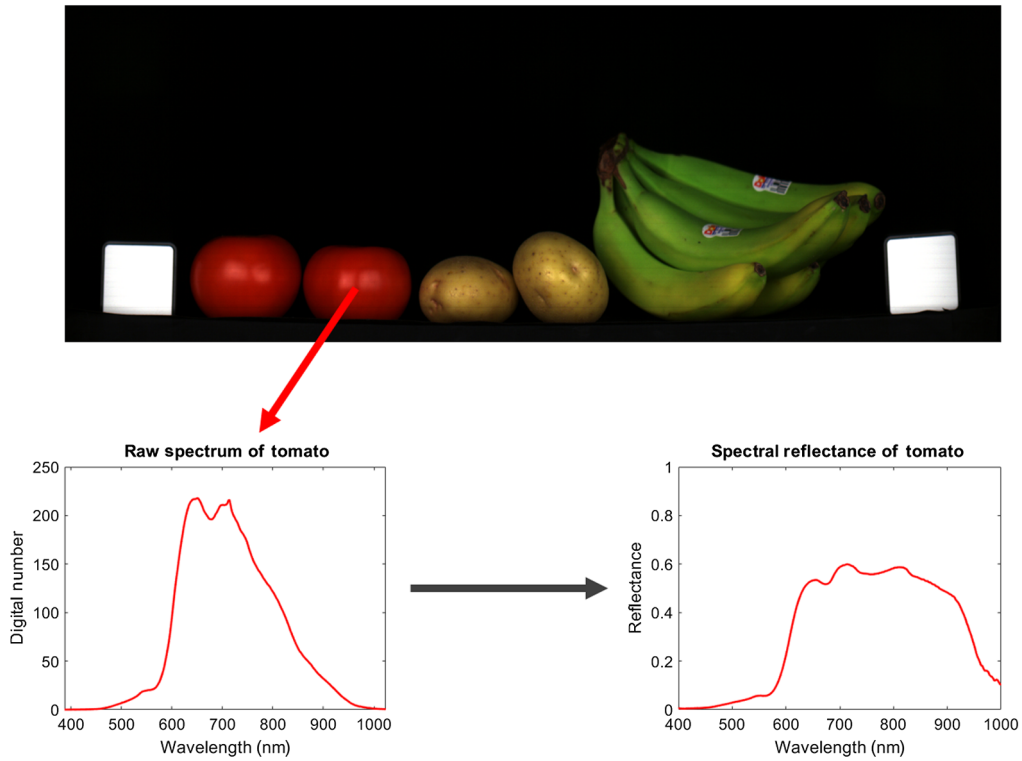

Fig. 3 Example of converting from digital number to reflectance using the known spectral reflectance of the Spectralon target (white squares). Here, the raw spectrum and corresponding reflectance spectrum are shown for a tomato pixel.

2.5 *Processing*

In this section, we discuss the development of machine learning algorithms to estimate the relative age of the test produce measured from the time of purchase, classify whether the produce was fresh, and select the most informative wavelengths for use in a low-cost multispectral imager.

2.5.1 *Convolutional neural network*

The key to the analysis described here is a novel CNN, previously designed to reduce the number of parameters needed while maintaining high performance.[33] The reduced parameterization was particularly important due to the curse of dimensionality.[34] Specifically, the more parameters that needed to be learned, the more data were required to prevent overfitting; therefore, by reducing the number of parameters, it was possible to develop accurate models with fewer data.

The CNN that we developed used both a macroarchitecture and a microarchitecture. Specifically, the macroarchitecture consisted of several spectral modules based on the SqueezeNet architecture and its associated "fire modules."[35] The microarchitecture then further reduced the number of parameters from SqueezeNet by reducing the number of 1×1 convolutions, thereby also reducing the number of 3×3 convolutions. The resulting network then applied 64×64 input patches (spatial pixels) for each channel. The architecture of the resulting SpectrumNet network is shown in Table 2 (note that the variable x in the table indicates the number of wavelengths used).

2.5.2 *Freshness classification*

The goal of the analysis performed here was to determine whether the produce that had been imaged could be classified as "fresh" or "non-fresh." For this step, we employed the CNN architecture described in Sec. 2.5.1 (i.e., SpectrumNet) with a classification "layer" appended to the output of the CNN. In particular, a simple feedforward prediction layer was placed on the top with two output nodes for classes corresponding to fresh and non-fresh. Three different versions

Table 2 SpectrumNet convolutional neural network.

Layer	Filter size	Filter/stride	1×1 Squeeze	1×1 Expand	3×3 Expand
Input	$64 \times 64 \times x$				
Conv1	$32 \times 32 \times 96$	$2 \times 2/1$			
Spectral2	$32 \times 32 \times 128$		16	96	32
Spectral3	$32 \times 32 \times 128$		16	96	32
Spectral4	$32 \times 32 \times 256$		32	192	64
Maxpool4	$16 \times 16 \times 256$	$2 \times 2/2$			
Spectral5	$16 \times 16 \times 256$		32	192	64
Spectral6	$16 \times 16 \times 384$		48	288	96
Spectral7	$16 \times 16 \times 384$		48	288	96
Spectral8	$16 \times 16 \times 512$		64	385	128
Maxpool8	$8 \times 8 \times 512$	$2 \times 2/2$			
Spectral9	$8 \times 8 \times 512$		64	385	128
Conv10	$8 \times 8 \times 10$	$1 \times 1/1$			
Avepool10	$1 \times 1 \times 10$	$8 \times 8/1$			

of this network were trained, corresponding to $x = 3$ (for RGB images), $x = 300$ (for full-spectrum images), and $x = 5$ [for multispectral images (MSIs)]. RGB images were created using single spectral channels from the hyperspectral data cubes at wavelengths of 641 nm (red), 551 nm (green), and 460 nm (blue). Details for selecting the five wavelengths for the MSIs are given in Sec. 2.5.4. All analyses presented here were performed with single 2.1-nm spectral bands. The training labels of fresh and non-fresh were assigned based on the result of a subjective evaluation made by the authors based on the visual appearance of the produce.

2.5.3 *Age analysis*

In addition to learning the classification of fresh versus non-fresh produce, we also explored estimating the age of the produce. For this analysis, we used the purchase date of the produce as the baseline age under the assumption that it would serve as a reasonable proxy for the produce being maximally fresh upon delivery to the store. Given this, we attempted to purchase produce from the grocery store as close to its delivery date as possible.

Using the purchase date as the baseline, we then trained a CNN based on SpectrumNet to solve a corresponding regression problem of estimating the number of days since purchase. This utilized the same architecture as described in Sec. 2.5.1 except that the prediction layer consisted of a single linear output node. Once again, three different networks were trained corresponding to $x = 3$, $x = 300$, and $x = 5$ for RGB, full spectrum, and multispectral, respectively.

2.5.4 *Wavelength selection*

One of the primary goals of this work was to design a low-cost imager for a grocery store setting. Given the potentially high expense of hyperspectral imaging, we investigated alternative methods for selecting small subsets of wavelengths from the spectrum for purposes of designing low-cost, multispectral imagers tailored to the produce being monitored. With this in mind, we developed a method using a GA that simultaneously selected a subset of informative wavelengths and identified effective filter bandwidths for such an imager.

A typical GA follows an iterative, stochastic search process in which one initializes a "population" of solutions and then proceeds through a loop of selecting members of the population based on their quality (or fitness).[36] The selected individuals then reproduce, generating offspring via a process of crossover and mutation. The offspring are then re-introduced into the population. From one generation to the next, the population of individuals is expected to improve in fitness by applying a process inspired by "survival of the fittest" in natural selection. Then the solution to the optimization problem is selected from the final population.

Details of our GA-based approach are discussed by Walton et al.[27,37] To summarize the key steps, our approach proceeded from a common evolutionary feature-selection method[38] whereby we represented the features (i.e., wavelengths) as an integer vector in which the integers identified the index corresponding to a selected wavelength. The specific GA design consisted of applying tournament selection, binomial crossover, and generational replacement. The mutation operator consisted of selecting a wavelength index, perturbing the index by adding a value selected from neighboring integers according to $U(-3, +3)$, and truncating if we exceeded the bounds of possible wavelengths. The fitness evaluation then cast our approach as a "wrapper" method using the selected wavelengths to train a decision tree and evaluated fitness based on classification performance.[39] For the wrapper evaluation, we randomly selected 20 patches of 5×5 pixels out of the data cube corresponding to each of the particular produce types (e.g., banana) using all 300 wavelengths for each of the 25 pixels in the patch.

Given the GA as the foundation for selecting wavelengths, our method proceeded in four steps once the GA was terminated. First, we considered the population as if it had clustered wavelengths. In other words, rather than returning a single individual in the population with its selected wavelengths, we constructed a histogram of the wavelengths represented in the population. Empirically, we found that these histograms appeared to follow an underlying Gaussian mixture model (GMM); therefore, the second step was to fit such a model to the population data via expectation maximization. Third, once the Gaussians had been identified in the GMM, we extracted the means of the component Gaussians as the selected wavelengths and extracted the

standard deviations as the basis for setting the filter bandwidths. Note, however, that we did not use the derived bandwidths in the experiments reported here. We plan to discuss the effectiveness of using filters derived from these bandwidths in a future paper in which we will transform the input hyperspectral data to mimic what a multispectral imager would have provided based on these filter specifications. We refer to this method as histogram-assisted genetic algorithm for reduction in dimensionality (HAGRID). Details of the HAGRID approach and experimental results utilizing the procedure can be found in the work by Walton et al.[27]

2.6 *Grocery Store Lighting*

Not all lighting systems are created equal. As discussed in Sec. 2.2.2, the spectral output of lighting elements can vary significantly. From our survey of six local grocery stores, we found that the most common lighting elements in produce sections were white LEDs and fluorescent bulbs. Unfortunately, the spectral outputs of LEDs and fluorescent bulbs are relatively sparse in comparison with the broadband spectral output of a halogen bulb used in our laboratory studies (Fig. 4). The spectral output of light sources becomes important when small spectral regions are required for analysis, such as the chlorophyll absorption line near 680 nm. If the lighting system found in a grocery store has no output in the spectral region of interest, a supplemental light source may be required for analysis.

To address the issue of varying spectral outputs in lighting systems, work is currently underway to simulate measurements taken in a grocery store setting. These analyses incorporate the spectral outputs of various grocery store lighting elements into the spectra measured from pieces of produce in a laboratory setting. More details are presented in Sec. 3.4.

3 Results and Discussion

In this section, we discuss the results obtained for the analysis of bananas, avocados, bell peppers, tomatoes, and potatoes. We begin with an overview of the reflectance image data obtained throughout the measurement process and then move to a discussion of the performance of the machine learning algorithms developed for freshness classification, age analysis, and wavelength selection. Finally, we discuss the possible implementation of our methodology into a grocery store setting.

3.1 *Image Analysis of Bananas, Avocados, Bell Peppers, Tomatoes, and Potatoes*

Changes in the chemical makeup of produce during ripening (e.g., decreasing chlorophyll content) result in changes in the reflectance spectrum.[9,13,40,41] This process is illustrated for a banana in Fig. 5. When the banana was in the early stages of ripening, high chlorophyll content gave it a characteristic green color.[7] Chlorophyll absorption was evident as a large trough near 680 nm in the reflectance spectrum. As the banana ripened and chlorophyll content decreased, the visible color of the banana shifted from green to yellow. This phase of ripening was represented in days 3 to 7, during which the chlorophyll absorption trough decreased and spectral features between approximately 550 and 900 nm became smoother. When the banana was brown and soft at the

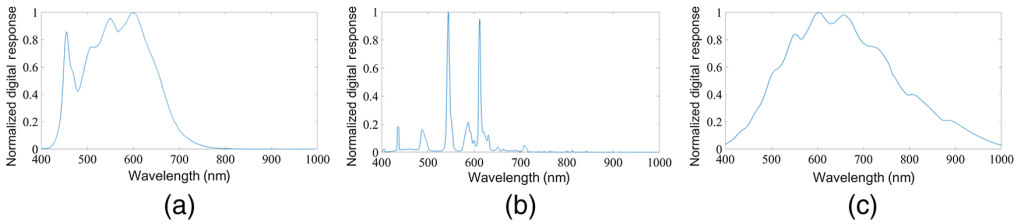

Fig. 4 Comparison of spectral output for various light sources: (a) white LED, (b) fluorescent bulb, and (c) halogen bulb.

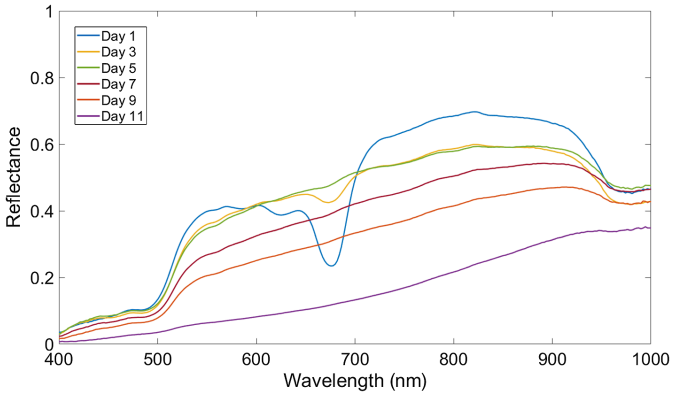

Fig. 5 Temporal variation in the spectral reflectance of a banana. Each line represents a measurement of the same banana taken on different days.

end of its ripeness cycle, the spectral reflectance features were almost entirely smooth (days 9 to 11).

The reflectance spectra of other types of produce also changed systematically during the ripening process. For example, though selecting a ripe avocado can be a bewildering task for some shoppers, our measurements of avocados showed relatively direct aging activity (Fig. 6). Avocado ripening is a complex process, with water content, ethylene content, and temperature all playing a role.[42] When not yet ripe, the avocado took on a firm, green appearance, represented by the increased reflectance near 550 nm and high reflectance at the knee of the red edge just beyond 700 nm (days 1 to 3). As the avocado ripened, its skin became darker and its reflectance decreased significantly beyond 700 nm (days 5 to 11). When the avocado became too ripe to eat and shriveled due to decreased water content, the reflectance remained below 10% until approximately 800 nm, where an increase in reflectance occurred (days 13 to 19).

Though bananas and avocados exhibited clear ripening processes, not all test produce that we analyzed followed this theme. To expand our analysis of produce commonly available at grocery stores, we analyzed green, yellow, orange, and red bell peppers. Freshly purchased bell peppers exhibited a unique spectral curve for each color (Fig. 7). Though unique spectral curves were helpful for distinguishing between types of peppers, they created a challenge for determining consistent measures for age analysis across all varieties of peppers. An additional challenge was considering the temporal variation in the reflectance spectra of peppers. Despite the unique spectrum for each pepper color, which arose because both chlorophyll and carotenoids have observable absorption features in our spectral range,[43] our measurements showed that there was no simple change in the spectral reflectance as the pepper aged (Fig. 8). Though there may have

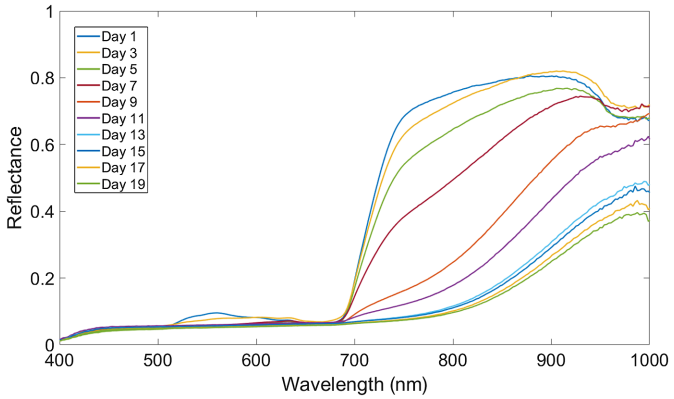

Fig. 6 Temporal variation in the spectral reflectance of avocados. Each colored curve represents a new measurement day on the same avocado.

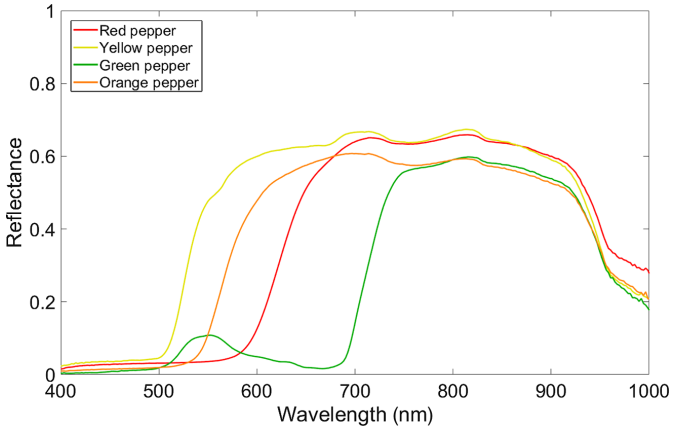

Fig. 7 Variation in the spectral reflectance of bell peppers. Each colored curve represents a different variety of bell pepper.

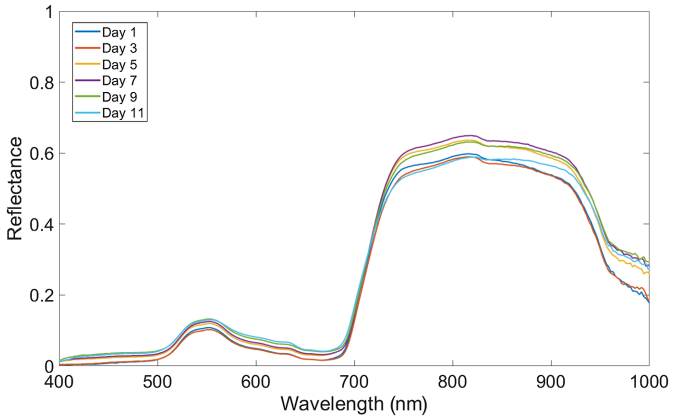

Fig. 8 Temporal variation in the spectral reflectance of a green bell pepper. Each line represents a measurement of the same pepper taken on different days.

been some detectable change in the relative heights of the two large reflectance peaks or the slope of the reflectance ramp from approximately 600 to 700 nm, often denoted as the red edge, the ripening process in peppers was much more difficult to detect visibly.

Tomatoes presented a similar challenge. Tomato ripening is characterized by varying lycopene content in the skin, the breakdown of chlorophyll, and the build-up of carotenes.[44] Though absorption features for chlorophyll and lycopene are well-known, our measurements showed no clear temporal variation of the reflectance spectra as the tomato aged (Fig. 9). Although the tomato was still fresh, there was some indication of a chlorophyll absorption feature near 680 nm, but this feature was subtle and only present during very early stages of the ripening process.

The final piece of produce presented in this section is the Yukon Gold potato. When potatoes are exposed to light for an extended period, an increase in chlorophyll production occurs in a process known as greening.[45] Greening is undesirable due to its adverse effect on potato marketability and the development of the toxic group of plant compounds, glycoalkaloids. The greening process was evident when observing the reflectance spectra of a Yukon Gold potato over the course of 17 days (Fig. 10). As the potato aged and chlorophyll production increased, an absorption feature developed near 680 nm. It is interesting to note that this aging process is nearly the opposite of that of bananas, with a chlorophyll absorption feature developing as the potato aged.

To explore the detection and use of the spectral features outlined above, we turned to machine learning algorithms. For brevity, we present results for a select subset of the produce analyzed, including potatoes, bananas, and green peppers.

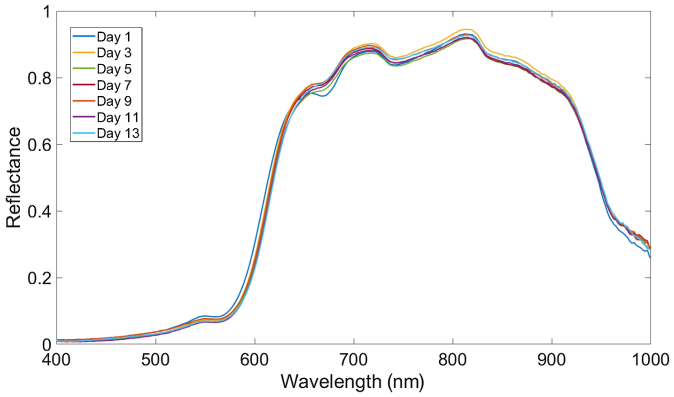

Fig. 9 Temporal variation in the spectral reflectance of a tomato. Each line represents a measurement of the same tomato taken on different days.

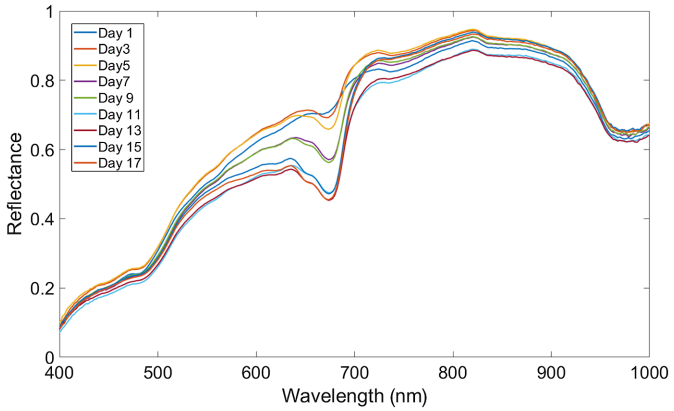

Fig. 10 Temporal variation in the spectral reflectance of a Yukon Gold potato. Each line represents a measurement of the same potato taken on different days.

3.2 *Age Analysis Results*

For our experiments, the networks with standard convolutions (we use the term "standard" convolutions to distinguish them from "depthwise separable" convolutions) were trained using stochastic gradient descent with Nesterov momentum set to 0.9 and weight decay (L2 penalty) set to 5×10^{-4}. With the depthwise separable convolutions, due to the small number of parameters in the network, weight decay was not necessary. The initial learning rate was set to 0.001 and was reduced by 25% every 10 epochs. We used a batch size of 64 and ran all experiments on a Nvidia GTX 1080TI GPU. We present the results of 10-fold cross-validation experiments.

We ran experiments with RGB, hyperspectral, and MSIs. Results for RGB images and full-hyperspectral image classification of fresh versus non-fresh for potatoes, bananas, and green peppers are shown in Table 3. Particularly surprising was the relatively low freshness classification performance of RGB images on bananas, given the distinct browning as they aged. The freshness classification performance on green peppers was also surprising, given less discernible visual differences as they aged. Previous results using SpectrumNet on MSIs when compared with competing methods, such as ResNet-50[46] and DenseNet-161,[47] show the relative superiority of our approach, both in terms of accuracy and computational complexity.[33]

Similarly, the results of our age analysis for RGB and hyperspectral images are shown in Table 4. Overall, we were pleasantly surprised at the age prediction in all three cases (RGB, MSI, and full spectrum), given the fact that the baseline for determining age was the point of purchase. We hypothesize that much of the variability exhibited in these results could be due to the variability in the age at the point of purchase. Furthermore, it appeared that variability is higher in

Table 3 SpectrumNet produce freshness classification (mean % accuracy ± one standard deviation). Bold values indicate statistically significantly different from RGB at $\alpha = 0.05$ using a paired *t*-test.

Produce	RGB	Full spectrum	HAGRID wavelengths (5)
Potatoes	92.5 ± 5.12	$\mathbf{98.0 \pm 3.32}$	92.0 ± 4.36
Bananas	83.4 ± 9.27	$\mathbf{95.4 \pm 4.38}$	$\mathbf{96.0 \pm 2.61}$
Green peppers	98.3 ± 5.00	$\mathbf{100.0 \pm 0.00}$	$\mathbf{100.0 \pm 0.00}$

Table 4 SpectrumNet produce age estimation in days (mean % absolute error ± one standard deviation). Bold values indicate statistically significantly different from RGB at $\alpha = 0.05$ using a paired *t*-test.

Produce	RGB	Full spectrum	HAGRID wavelengths (5)
Potatoes	2.82 ± 0.77	$\mathbf{1.41 \pm 0.52}$	$\mathbf{1.51 \pm 0.60}$
Bananas	1.34 ± 0.34	$\mathbf{0.84 \pm 0.10}$	1.14 ± 0.17
Green peppers	0.94 ± 0.12	$\mathbf{0.18 \pm 0.06}$	$\mathbf{0.49 \pm 0.13}$

the classification case, due to the freshness labels being based on human inspection. For the age prediction, although there was variability in age at the point of purchase, under the assumption produce was delivered in a timely manner to the store, this measure is arguably less arbitrary.

3.3 *Genetic Algorithm*

For our experiments, we set the population size to 1000 and ran the algorithms for 300 generations. The crossover rate, mutation rate, and tournament size were tuned using a grid search, as were the parameters for the neural network. Since the GA only changed the way in which the solution was selected from the final population of the GA, one complete run of the GA was used for training the neural network.

The freshness classification results for the GA-based feature selection are shown in the rightmost column of Table 3. For the potato, we saw no appreciable improvement over the RGB images, but this was not surprising given the relative ease of seeing potatoes aging visually. However, we found a significant improvement for banana age classification over RGB and even some improvement (albeit not statistically significant) when compared with the full spectrum. The results on the green peppers were surprising in that the performance was perfect, as with the full spectrum.

The age prediction results for the GA-based feature selection are shown in the rightmost column of Table 4. We observed in this case that, with the exception of potatoes using RGB images, performance was strong, only erring on average less than one day. Even so, a 2 to 3 day error on potatoes was not bad, given their general longevity. We note that there was no appreciable difference between the multispectral-based prediction and RGB classification on bananas. In contrast to the freshness classification results, this result was not surprising, given the ability to observe bananas aging visually. Of interest was the low mean absolute error on green peppers, again given the less discernible visual differences as they aged.

Experimental results in both freshness classification and age prediction revealed that the GA-based method was capable of finding the relevant wavelengths on difficult classification problems. Given that the main goal of wavelength selection was to reduce the computational burden (resulting from the reduction in network complexity) without sacrificing classification performance, our results indicated that this goal was achieved. In particular, we note that there was no statistically significant difference in freshness classification between HAGRID and the full spectrum on either bananas or green peppers. Although there was a significant difference between

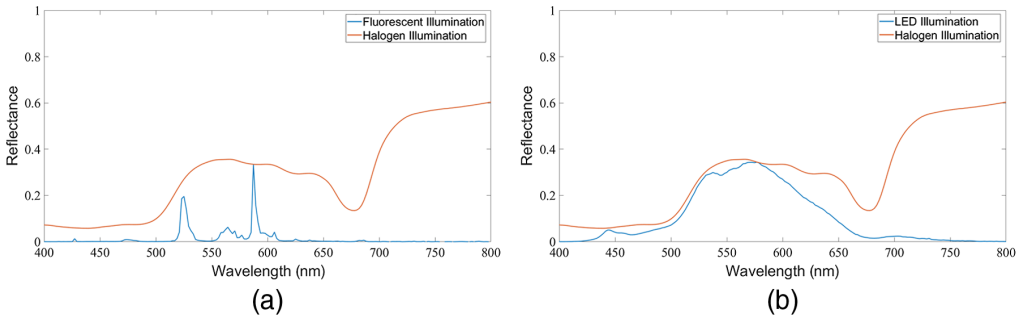

Fig. 11 Banana spectra before and after illumination overlay with (a) fluorescent bulb and (b) LED.

HAGRID and the full spectrum on potatoes, we found that the HAGRID-based results were statistically equivalent to RGB classification. With respect to age prediction, we found no statistical difference between HAGRID and the full spectrum on potatoes. Although there was a significant difference between HAGRID and the full spectrum on bananas and green peppers, the overall prediction performance of both was so strong as to render the difference practically insignificant.

3.4 *Grocery Store Lighting Analysis*

The analyses presented above are for data obtained with a broadband lighting system; however, grocery stores often use fluorescent or LED lighting systems. To extend our analyses to include a more realistic grocery store setting, we began by visiting six grocery stores in the Bozeman, Montana, area and measured the spectral output of the general and produce-specific lighting systems using a VNIR spectrometer (Ocean Optics USB4000). The measured spectra were normalized to their maximum value and multiplied by the daily reflectance spectra measured previously for each piece of produce under broadband lighting. This simulated the reflectance spectra that would have been measured in a grocery store setting for each type of lighting system. The spectra in Fig. 11 show an example of the illumination correction process using a typical fresh banana spectrum. The spectrum under diffuse halogen lighting is shown in red, and the illumination-corrected reflectance spectrum is shown in blue for fluorescent lighting [Fig. 11(a)] and LED lighting [Fig. 11(b)].

It is evident from each of the illumination-corrected reflectance spectra that significant spectral information is lost under both fluorescent and LED illuminations. Work is currently underway to understand how this loss of spectral information affects the performance of the algorithms introduced in Secs. 3.2 and 3.3.

4 Conclusion

In this paper, we presented a promising method for characterizing the ripening process of fruits and vegetables through age analysis and freshness classification using a calibrated VNIR hyperspectral imager coupled with machine learning algorithms. Several calibrations and analyses were performed on the hyperspectral imaging system, including a conversion to reflectance spectra, pixel uniformity testing, image stability analysis, spectral performance testing, and an assessment of polarization response. Hyperspectral data cubes were obtained with produce placed in a specially designed, low-reflectance stage. Measurements were made once a day throughout the life cycle of the produce being tested, which included bananas, tomatoes, potatoes, avocados, and bell peppers. For analysis, all measurements were converted into reflectance spectra using two Spectralon panels placed on the produce stage.

Using the reflectance spectra, machine learning algorithms were used to analyze the ripeness of fruits and vegetables. Age analysis experiments were performed with a novel CNN called SpectrumNet trained using stochastic gradient descent. These experiments considered two prediction tasks—freshness classification and age prediction (since purchase). Experiments were

run using both RGB images and hyperspectral images to compare performance. Results showed that hyperspectral images outperformed RGB images for freshness classification and age prediction on all tested produce, including potatoes, bananas, and green peppers.

Feature selection was performed using a GA, with hyperparameters tuned using a grid search. Using the GA, we selected five salient wavelengths from the full-hyperspectral spectrum and compared freshness classification and age prediction performance against RGB images and hyperspectral imagery. The GA-based feature selection method typically outperformed RGB images on freshness classification, with a significant improvement on classification of bananas and green peppers. Similarly, the GA-based feature selection method outperformed RGB images on age prediction, with significant improvement on potatoes and green peppers. When compared with hyperspectral imagery, the GA-based method used with freshness classification showed modest improvement in classifying bananas and matched performance for green peppers. However, hyperspectral imagery outperformed all methods for classification of potatoes. In the case of age prediction, the GA-based method was statistically significantly worse than hyperspectral images on both bananas and green peppers, but the overall low mean absolute error indicated that this difference in performance was not significant in the practical sense of the word. There was no significant difference (statistically or practically) in the performance on potatoes.

The GA-based feature selection method presented in this paper provided a method for determining salient wavelengths for a given application. Currently, this method is being used to develop and test a low-cost multispectral imager for use in age classification of bell peppers. For future work, further efforts in the development and testing of a wide range of multispectral imagers for age classification of produce in grocery stores could be explored.

Acknowledgments

This work was funded initially by the Intel Corp. under Agreement Montana-Intel SRA 10-18-17. This material is also based on work supported in part by the National Science Foundation EPSCoR Cooperative Agreement OIA-1757351. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. Disclosures: This paper is an expanded discussion of work initially reported in an SPIE Proceedings conference paper.[28] In accordance with the SPIE publication policy,[29,30] we have expanded the text but also reused figures and text from the conference paper with this acknowledgment. Compared with the earlier conference paper, we added to this paper experiments focused on predicting the age of the produce since purchase, a discussion of the different light source spectra found in grocery stores and the expected impacts of this on in-store measurements, and an expanded discussion of the genetic algorithm used to select optimal wavelengths from hyperspectral images.

References

1. J. Gustavsson et al., "Global food losses and food waste—extent, causes and prevention," in *Int. Congr. Save Food*, Food and Agriculture Organization of the United Nations, Rome (2011).
2. K. Venkat, "The climate change and economic impacts of food waste in the United States," *Int. J. Food Syst. Dyn.* **2**, 431–446 (2011).
3. D. Gunders, "Wasted: how America is losing up to 40 percent of its food from farm to fork to landfill," Technical Report 12-06-B, Natural Resources Defense Council (2012).
4. M. Melikoglu, C. S. K. Lin, and C. Webb, "Analysing global food waste problem: pinpointing the facts and estimating the energy content," *Central Eur. J. Eng.* **3**, 157–164 (2013).
5. K. D. Hall et al., "The progressive increase of food waste in America and its environmental impact," *PLoS One* **4**, e7940 (2009).
6. J. C. Buzby, H. F. Wells, and J. Hyman, "The estimated amount, value, and calories of postharvest food losses at the retail and consumer levels in the United States," Technical Report 121, USDA (2014).

7. M. Li, D. C. Slaughter, and J. F. Thompson, "Optical chlorophyll sensing system for banana ripening," *Postharv. Biol. Technol.* **12**, 273–283 (1997).
8. J. A. Abbott, "Quality measurement of fruits and vegetables," *Postharv. Biol. Technol.* **15**, 207–225 (1999).
9. A. Peirs et al., "Prediction of the optimal picking date of different apple cultivars by means of VIS/NIR-spectroscopy," *Postharv. Biol. Technol.* **21**, 189–199 (2000).
10. M. S. Kim, Y. R. Chen, and P. M. Mehl, "Hyperspectral reflectance and fluorescence imaging system for food quality and safety," *Trans. ASAE* **44**(3), 721–729 (2001).
11. G. Polder, G. W. A. M. Van Der Heijden, and I. T. Young, "Spectral image analysis for measuring ripeness of tomatoes," *Trans. ASAE* **45**(4), 1155–1161 (2002).
12. V. Prasanna, T. N. Prabha, and R. N. Tharanathan, "Fruit ripening phenomena—an overview," *Crit. Rev. Food Sci. Nutr.* **47**, 1–19 (2007).
13. P. Rajkumar et al., "Studies on banana fruit quality and maturity stages using hyperspectral imaging," *J. Food Eng.* **108**, 194–200 (2012).
14. S. Cubero et al., "Advances in machine vision applications for automatic inspection and quality evaluation of fruits and vegetables," *Food Bioprocess Technol.* **4**, 487–504 (2011).
15. R. Pandey, S. Naik, and R. Marfatia, "Image processing and machine learning for automated fruit grading system: a technical review," *Int. J. Comput. Appl.* **81**(16), 29–39 (2013).
16. S. Naik and B. Patel, "Image processing and machine learning for automated fruit grading system: a technical review," *Int. J. Comput. Appl.* **170**(9), 22–34 (2017).
17. E. Elhariri et al., "Bell pepper ripeness classification based on support vector machine," in *Int. Conf. Eng. and Technol.*, pp. 1–6 (2014).
18. W. Castro et al., "Classification of cape gooseberry fruit according to its level of ripeness using machine learning techniques and different color spaces," *IEEE Access* **7**, 27389–27400 (2019).
19. F. Femling, A. Olsson, and F. Alonso-Fernandez, "Fruit and vegetable identification using machine learning for retail applications," in *14th Int. Conf. Signal-Image Technology & Internet-Based Systems (SITIS)* (2018).
20. H. Vaviya et al., "Identification of artificially ripened fruits using machine learning," in *2nd Int. Conf. Adv. in Sci. & Technol.*, pp. 1–6 (2019).
21. J. Gomez-Sanchis et al., "Detecting rottenness caused by *Penicillium* genus fungi in citrus fruits using machine learning techniques," *Expert Syst. Appl.* **39**(1), 780–785 (2012).
22. N. V. Rivera et al., "Early detection of mechanical damage in mango using NIR hyperspectral images and machine learning," *Biosyst. Eng.* **122**, 91–98 (2014).
23. O. Gupta et al., "Machine learning approaches for large scale classification of produce," *Sci. Rep.* **8**(1), 5226 (2019).
24. N. El-Bendary et al., "Using machine learning techniques for evaluating tomato ripeness," *Expert Syst. Appl.* **42**(4), 1892–1905 (2015).
25. P. W. Nugent et al., "Discrimination of herbicide-resistant kochia with hyperspectral imaging," *J. Appl. Remote Sens.* **12**(1), 016037 (2018).
26. B. Scherrer et al., "Hyperspectral imaging and neural networks to classify herbicide-resistant weeds," *J. Appl. Remote Sens.* **13**(4), 044516 (2019).
27. N. S. Walton, J. W. Sheppard, and J. A. Shaw, "Using a genetic algorithm with histogram-based feature selection in hyperspectral image classification," in *Proc. ACM Genetic and Evol. Comput. Conf.* (2019).
28. R. D. Logan et al., "Hyperspectral imaging and machine learning for monitoring produce ripeness," *Proc. SPIE* **11421**, 114210O (2020).
29. SPIE, "Revision of conference proceedings manuscripts for journal submission," https://www.spie.org/conferences-and-exhibitions/authors-and-presenters/spie-proceedings-to-journal-pathway?SSO=1 (2021).
30. M. T. Eismann, "Natural synergy of conferences and journals," *Opt. Eng.* **53**(12), 120101 (2016).
31. J. S. Tyo et al., "Review of passive imaging polarimetry for remote sensing applications," *Appl. Opt.* **45**(22), 5453–5469 (2006).
32. R. D. Logan and J. A. Shaw, "Measuring the polarization response of a VNIR hyperspectral imager," *Proc. SPIE* **11412**, 113–119 (2020).

33. J. Senecal, J. Sheppard, and J. Shaw, "Efficient convolutional neural networks for multi-spectral image classification," in *Proc. IEEE Int. Joint Conf. Neural Networks* (2019).
34. R. Bellman, *Dynamic Programming*, Princeton University Press (1957).
35. F. Iandola, "Exploring the design space of deep convolutional neural networks at large scale," PhD Thesis, University of California, Berkeley, California (2016).
36. M. Mitchell, *An Introduction to Genetic Algorithms*, The MIT Press, Cambridge, Massachusetts (1996).
37. N. Walton, "Utilizing distributions of variable influence for feature selection in hyperspectral images," Master's Thesis, Montana State University, Bozeman, Montana (2019).
38. J. Yang and V. Honavar, *Feature Subset Selection Using a Genetic Algorithm*, pp. 117–136, Springer US, Boston, Massachusetts (1998).
39. R. Kohavi and G. John, "Wrappers for feature subset selection," *Artif. Intell.* **97**, 273–324 (1997).
40. D. C. Slaughter and J. A. Abbott, *Analysis of Fruits and Vegetables*, pp. 377–398, John Wiley & Sons Ltd. (2004).
41. D. Laval-Martin, J. Quennemet, and R. Monéger, "Pigment evolution in *Lycopersicon esculentum* fruits during growth and ripening," *Phytochemistry* **14**(11), 2357–2362 (1975).
42. J. P. Bower and J. G. Cutting, *Avocado Fruit Development and Ripening Physiology*, pp. 229–271, John Wiley & Sons, Ltd. (1988).
43. P. A. Luning et al., "Gas chromatography, mass spectrometry, and sniffing port analyses of volatile compounds of fresh bell peppers (*Capsicum annuum*) at different ripening stages," *J. Agric. Food Chem.* **42**, 977–983 (1994).
44. G. Polder, G. W. A. M. van der Heijden, and I. T. Young, "Tomato sorting using independent component analysis on spectral images," *Real-Time Imaging* **9**, 253–259 (2003).
45. L. Grunenfelder, L. K. Hiller, and N. R. Knowles, "Color indices for the assessment of chlorophyll development and greening of fresh market potatoes," *Postharv. Biol. Technol.* **40**, 73–81 (2006).
46. K. He et al., "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, pp. 770–778 (2017).
47. G. Huang et al., "Densely connected convolutional networks," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, pp. 2261–2269 (2017).

Riley D. Logan received his MS degree in optics and photonics from Montana State University. He is a PhD student at Montana State University working with the Optical Remote Sensor Laboratory. His current research includes optical polarization and remote sensing with a focus on developing airborne hyperspectral imaging systems and techniques for environmental monitoring.

Bryan Scherrer received his MS degrees in physics and electrical engineering from Montana State University. He is a research engineer at the Optical Remote Sensor Laboratory, Montana State University, Bozeman, Montana. His current work focuses on the development of hyperspectral imaging and lidar systems for use in precision agriculture, ecology, and food processing.

Jacob Senecal received his MS degree in computer science and his BS degree in mechanical engineering from Montana State University. He is currently a senior software engineer at Figure Technologies, building technology to transform financial services through blockchain, AI, and analytics.

Neil S. Walton received his MS degree in computer science from Montana State University and his BA degree in computer science and economics from Colgate University. He is currently employed as a software engineer at Goldman Sachs, where he works to automate regulatory reporting processes.

Amy Peerlinck received her MS degree in computer science from Montana State University, her BA degree in applied linguistics from the University of Antwerp, and her BS degree in information science from Karel De Grote College/University. She is currently working toward her PhD in computer science at Montana State University, where she is a research assistant

on a precision agriculture grant, optimizing profit for farmers through machine learning techniques.

John W. Sheppard received his PhD and MS degree in computer science from Johns Hopkins University and his BS degree in computer science from Southern Methodist University. He is a fellow of the Institute of Electrical and Electronics Engineers and a Distinguished Professor of Computer Science at Montana State University. His research interests include deep learning, probabilistic graphical models, and evolutionary optimization applied to electronic prognostics and health management, precision agriculture, and medical diagnostics.

Joseph A. Shaw received his PhD and MS degree in optical sciences from the University of Arizona, his MS degree in electrical engineering from the University of Utah, and his BS degree in electrical engineering from the University of Alaska. He is the director of the Optical Technology Center and Distinguished Professor at Montana State University in Bozeman, Montana, USA, where he develops optical remote sensing systems for environmental science. He is a fellow of OSA and SPIE and was the 2019 recipient of the G. G. Stokes Award from SPIE.