

PROCEEDINGS OF SPIE

SPIDigitalLibrary.org/conference-proceedings-of-spie

Hyperspectral imaging and machine learning for monitoring produce ripeness

Logan, Riley, Scherrer, Bryan, Senecal, Jacob, Walton, Neil, Peerlinck, Amy, et al.

Riley D. Logan, Bryan Scherrer, Jacob Senecal, Neil S. Walton, Amy Peerlinck, John W. Sheppard, Joseph A. Shaw, "Hyperspectral imaging and machine learning for monitoring produce ripeness," Proc. SPIE 11421, Sensing for Agriculture and Food Quality and Safety XII, 114210O (22 April 2020); doi: 10.1117/12.2560968

SPIE.

Event: SPIE Defense + Commercial Sensing, 2020, Online Only, California, United States

Hyperspectral imaging and machine learning for monitoring produce ripeness

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

^aElectrical and Computer Engineering Department, Montana State University, P.O. Box 173780, Bozeman, USA

^bGianforte School of Computing, Montana State University, P.O. Box 173880, Bozeman, USA

ABSTRACT

Hyperspectral imaging is a powerful remote sensing tool capable of capturing rich spectral and spatial information. Although the origins of hyperspectral imaging are in terrestrial remote sensing, new applications are emerging rapidly. Owing to its non-destructive nature, hyperspectral imaging has become a useful tool for monitoring produce ripeness. This paper describes the process that uses a visible near-infrared (VNIR) hyperspectral imager from Resonon, Inc., coupled with machine learning algorithms to assess the ripeness of various pieces of produce. The images were converted to reflectance across a spectral range of 387.12 nm to 1023.5 nm, with a spectral resolution of 2.12 nm. A convolutional neural network was used to perform age classification for potatoes, bananas, and green peppers. Additionally, a genetic algorithm was used to determine the wavelengths carrying the most useful information for age classification. Experiments were run using RGB images, full spectrum hyperspectral images, and the genetic algorithm feature selection method. Results showed that the genetic algorithm-based feature selection method outperforms RGB images for all tested produce, outperforms hyperspectral imagery for bananas, and matches hyperspectral imagery performance for green peppers. This feature selection method is being used to develop a low-cost multi-spectral imager for use in monitoring produce in grocery stores.

Keywords: Remote Sensing, Hyperspectral Imaging, Food Quality, Food Safety, Machine Learning

1. INTRODUCTION

In 2011, global food waste was estimated at 1.2 billion metric tons annually, meaning that approximately one-third of all food produced for human consumption was wasted.^{1,2} The United States doesn't fall far from the global average when it comes to food waste, with nearly 113 million metric tons, or roughly 29% of all food produced going to waste,² with some estimates placing the number as high as 40%.³ This not only costs businesses and consumers across the United States between \$165 and \$198 billion annually,^{2,3} but decomposition of wasted food releases greenhouse gasses that contribute to global climate change.^{4,5} Additionally, food waste can be tied to an increased consumption of over one-quarter of freshwater and nearly 300 million barrels of oil.⁵

According to the United States Department of Agriculture (USDA), 133 billion pounds of food was wasted at the retail and consumer level in 2010.⁶ The loss of fruits and vegetables in retail grocery stores, from either unsold or discarded produce, amounts to approximately \$15 billion annually. This loss is driven by many factors, such as overstocked product displays and damaged goods,³ but can also be attributed to produce becoming overripe on store shelves. Employees at our local grocery stores tell us that damaged or overripe produce is currently identified by employees who survey the display shelves throughout the day, removing produce that appears damaged or overripe. However, human observers are unable to predict when produce may be going bad, meaning produce is discarded once it has already passed the point of prime retail value. The ability to forecast the life cycle of produce would allow retail stores to implement strategies for rapidly selling

Further author information:

Joseph A. Shaw: E-mail: joseph.shaw@montana.edu, Telephone: 406-994-7261

Sensing for Agriculture and Food Quality and Safety XII, edited by Moon S. Kim, Bryan A. Chin, Young-Kwan Cho, Proc. of SPIE Vol. 11421, 114210O · © 2020 SPIE
CCC code: 0277-786X/20/\$21 · doi: 10.1117/12.2560968

Proc. of SPIE Vol. 11421 114210O-1

produce that is about to expire. In this paper, we present a promising method for detecting and predicting the ripening process of fruits and vegetables utilizing hyperspectral imaging coupled with machine learning algorithms.

Optical remote sensing has long been used for non-destructively analyzing food safety and quality,⁷⁻¹¹ especially hyperspectral imaging, owing to its rich spatial and spectral information. The spectral information available at each pixel is useful for monitoring the ripening of fruits and vegetables, which usually involves chemical processes such as chlorophyll degradation, changes in respiration, biosynthesis of carotenoids, and changes in ethylene production.¹² Many of the visible changes throughout the ripening process are attributed to changes in pigmentation induced by varying chlorophyll content and accumulation of carotenoids. Changes in pigmentation present in produce can be observed with the high spectral resolution of a hyperspectral imaging system, allowing for detection of the ripening cycle. In one such case, hyperspectral imaging was used to observe banana fruit quality and maturity stages at three different temperatures.¹³ Using quality parameters such as firmness, moisture content, and total soluble solids, they developed a relationship between quality and spectral data using partial least square analysis. Additionally, important wavelengths were determined using predicted residual error sum of squares, revealing eight visible and near-infrared wavelengths that carried the most information for determining quality.

Machine learning algorithms are commonly used in these applications to process large quantities of data. For example, machine learning methods have been used with machine vision or RGB cameras for classifying and grading the quality of fruits and vegetables¹⁴⁻¹⁶ and determining the ripeness of bell peppers¹⁷ and gooseberries.¹⁸ Convolutional neural networks have been employed to process images from low-cost RGB cameras coupled to a Raspberry PI computer to identify fruit and vegetable types¹⁹ and to process smartphone images to detect artificially ripened fruit.²⁰

Machine learning is particularly useful for analyzing the large quantities of spectral data generated by hyperspectral imagers or spectrometers. Some relevant 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}

In this paper, we build on our recent work in agricultural remote sensing^{25,26} to expand the analysis of fruits and vegetables using hyperspectral imaging and machine learning. In particular, we develop a framework for analyzing and predicting the ripeness of produce commonly available in the retail grocery store setting, focusing on bananas, potatoes, tomatoes, avocados, and bell peppers. We propose methods for analyzing hyperspectral images of commonly available fruits and vegetables using novel machine learning algorithms to predict remaining shelf life, thus decreasing the amount of discarded food.

2. METHODS

In this section, we introduce an overview of our experimental procedure, beginning with a broad overview of our 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 will serve as a road map of our experimental methods, beginning with our imaging system and illumination techniques.

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 nm to 1023.50 nm with a spectral resolution of approximately 2.12 nm, producing 12-bit hyperspectral data cubes with 300 spectral channels. We performed a radiometric calibration on our Pika L imager, assessed its polarization response, validated its spectral and pixel-to-pixel performance, and analyzed its image-to-image stability.




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

2.2.1 Imager Characterization

We performed the radiometric calibration with a variable-radiance integrating sphere (Labsphere USLR-V12F-NDNN). By capturing hyperspectral data cubes at 40 radiance levels, ranging from $192.7 \text{ W m}^{-2} \text{ sr}^{-1}$ to $0.62 \text{ W m}^{-2} \text{ sr}^{-1}$ within the Pika L's spectral range, we related the raw digital numbers captured by our imaging system to the known output of the integrating sphere. With this calibration, we obtained linear calibration curves across all of our spectral channels, mapping digital number to radiance.

Though the Pika L imager was not designed for polarization imaging, many hyperspectral imaging systems are partially sensitive to polarization.²⁷ As many scenes commonly contain partially polarized light, ignoring the polarization-sensitivity of a system may lead to erroneous results. In the context of observing fruits and vegetables, light can become partially polarized when reflected from the glossy surface of produce. Therefore, to determine the severity of the polarization response of the Pika L imager, we imaged the uniform, randomly polarized output of an integrating sphere through a linear polarizer (Meadowlark Optics VLM-129-UV-C) mounted in a precision rotation stage. We rotated the test polarizer through 180° in steps of 2° , imaging at each step. After analysis, we determined the polarization response to be relatively small, with a maximum value of 0.073 (7.3%) at 744 nm. This relatively small response coupled with the weakly polarized light from the majority of the produce surface convinced us that no compensation was necessary for this current study. More details of the polarization response are reported elsewhere.²⁸

To verify that our imaging system accurately detected the spectral content of a scene, we performed a spectral calibration using a monochromator (Acton Research Corporation SpectraPro-150) and spectrometer (Ocean Optics USB4000-VIS-NIR). We swept the monochromator through a spectral range of 400 nm to 1000 nm in steps of 10 nm, collecting both a hyperspectral data cube and a spectrometer measurement at each step. We found that the wavelengths from the monochromator, the spectrometer, and the Pika-L all matched within measurement uncertainty.

To measure the uniformity of pixel response across the Pika L's detector array, we imaged the spatially uniform output of the integrating sphere by scanning the imager across the sphere's aperture. As each of the line-scans captured across the aperture of the integrating sphere responds with the same pixels on the imager's detector, we averaged pixel values across the scanned spatial dimension, resulting in a 900 spatial pixel x 300 spectral pixel image. Since the integrating sphere does not have a flat spectral radiance output, each of the spectral channels captured by the Pika L have a unique response. To report the spatial uniformity, we defined a non-uniformity deviation metric, calculated as

$$\text{Non-Uniformity Deviation} = \frac{\text{Standard Deviation}}{\text{Mean}}. \quad (1)$$

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 6 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 non-uniformity deviation across all measured spectral channels, with maximum values of 2.67% before application of the radiometric calibration and 0.33% after radiometric calibration, compared to the known integrating sphere spatial uniformity of approximately 1% as reported by the integrating sphere manufacturer. Additionally, before each image capture, we subtracted the dark signal from each image using the SpectronPro software.

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

2.2.2 Light Sources

Not all light sources produce light over the Pika L imager’s spectral range of 400 nm 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 watt halogen bulbs mounted behind diffusing screens. Each source was positioned at $\pm 30^\circ$ 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 x 5 cm, 99% reflective Spectralon reference plates (Labsphere SRT-99-020) next to the produce. Further detail are discussed in Sec. 2.4.




Figure 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.

2.3 Experimental Setup and Testing Procedure

Measurements were made with produce sitting on a specially designed low-reflectance produce stage, alongside the two reflectance calibration targets, and scanning the imager through a full angle of approximately 24° . 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 the SpectronPro software provided by Resonon, Inc.. The low-reflectance produce stage kept stray light from entering the test stage while also reducing multiple reflections between the walls of the produce stage and other pieces of produce or other objects in the setup.

We began with fresh produce obtained from a local grocery store and ended each measurement when the produce was visibly inedible (i.e. rot or mold present). Store employees informed us that produce was restocked from distribution vehicles every 1-3 days, meaning fresh test produce was easy to obtain. We selected test samples from the most commonly available items in produce departments. In this paper, we discuss measurements of bananas, bell peppers, 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, 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 ripening process 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 contains 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 0.33%. Reflectance values were found at each produce pixel from

$$\rho_{scene} = \left(\frac{DN_{scene} - I_{dark}}{DN_{target} - I_{dark}} \right) \rho_{target} \quad (2)$$

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, DN_{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, giving us a reflectance spectrum from 400 nm 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.

2.5 Processing

In this section, we discuss the development of machine learning algorithms to classify the type of produce under analysis, the relative age of the test produce measured from the time of purchase, and for selecting the most informative wavelengths for use in a low-cost multi-spectral imager.

2.5.1 Age Analysis

The main part of the analysis performed here was with respect to determining the age of the produce that had been imaged. For this step, we employed a convolutional neural network architecture that we call "Spectrum-Net." The details of this network are described in Senecal,²⁹ with a specific description of its application to related multispectral classification problems in Senecal *et al.*³⁰

The basic architecture of SpectrumNet defines a set of "spectral" modules that correspond to squeeze layers and expand layers. The squeeze layers define the bulk of the network, except for a convolutional layer placed at the beginning and end of the network. In addition, depthwise convolution and batch normalization are

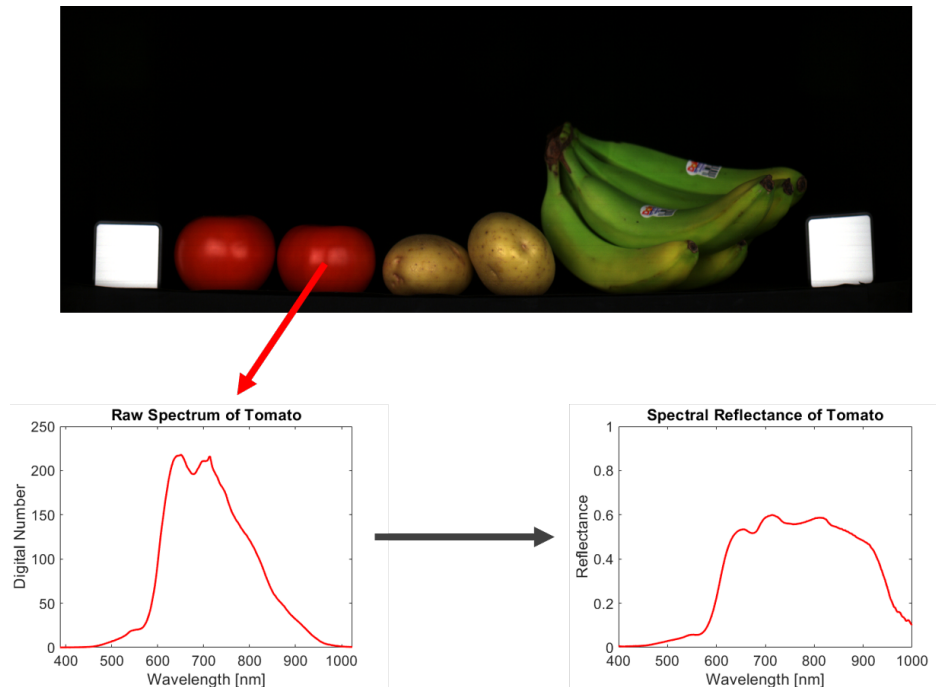


Figure 3. Example of converting from digital number to reflectance using the known spectral reflectance of the Spectralon target (white squares). Here, the raw spectrum and corresponding reflectance spectrum are shown for a tomato pixel.

incorporated into the squeeze and expand layers, thereby significantly reducing the computational burden of the network. The resulting network consisted of an input layer, a convolutional layer, three spectral layers (combining squeeze and expand), a max-pooling layer, six more spectral layers, a final convolutional layer, and an average pooling layer. A basic feedforward prediction layer was then placed on the top with classes corresponding to *fresh* and *non fresh*.

2.5.2 Wavelength Selection

Given that our goal is to deploy a low-cost imager in a grocery store setting, it is evident that full hyperspectral imaging is unlikely to be cost-effective. As a result, we have been investigating alternative methods for selecting small subsets of wavelengths from the spectrum for purposes of designing low-cost, multi-spectral imagers tailored to the produced being monitored. With this in mind, we developed a method using a genetic algorithm (GA) that simultaneously selects a subset of informative wavelengths and identifies effective filter bandwidths for such an imager.

Our approach involves fitting a multivariate Gaussian mixture model to a histogram of the GA population at convergence, selecting the wavelengths associated with the peaks of the distributions as our solution. Using this method, we are also able to specify filter bandwidths by calculating the standard deviations of the Gaussian distributions and computing the full-width at half-maximum values. In our experiments, we find that this novel histogram-based method for feature selection is effective when compared to both the standard GA and partial least squares discriminant analysis. Details on this approach can be found in Walton *et al.*³¹ and Walton.³²

3. RESULTS AND DISCUSSION

In this section, we discuss the results obtained for the analysis of bananas, tomatoes, avocados, and bell peppers. We begin with an overview of the reflectance image data obtained throughout the measurement process then move to a discussion of the performance of the machine learning algorithms developed for 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, Tomatoes, Avocados, and Bell Peppers

Changes in the chemical makeup of produce during ripening (e.g. decreasing chlorophyll content) result in changes in the reflectance spectrum.^{9,13,33,34} This process is illustrated for a banana in Fig. 4. When a banana is in the early stages of ripening, high chlorophyll content gives it a characteristic green color.⁷ The chlorophyll absorption is evident as the 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 is represented in Days 3 - 7, during which the chlorophyll absorption trough decreased and spectral features between approximately 550 nm to 900 nm became smoother. When the banana was brown and soft at the end of its ripeness cycle, the spectral reflectance features were almost entirely smooth (days 9 - 11).




Figure 4. Temporal variation in the spectral reflectance of a banana. Each line represents a measurement of the same banana taken on different days.

The reflectance spectra of other types of produce also change systematically during the ripening process. For example, though selecting a ripe avocado is can be a bewildering task for some shoppers, our measurements of avocados showed relatively direct aging activity (Fig. 5). Avocado ripening is a complex process, with water content, ethylene content, and temperature all playing a role.³⁵ 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 - 3). As the avocado ripened, its skin became darker and its reflectance decreased significantly beyond 700 nm (days 5 - 11). When the avocado became too ripe to eat and shriveling occurred because of decreased in water content, the reflectance remained below 10% until approximately 800 nm, where an increase in reflectance occurred (days 13 - 19).

Though bananas and avocados exhibit clear ripening processes, not all test produce 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. 6). Though unique spectral curves are helpful for classification, they create a challenge for determining consistent measures for aging 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 arises because of both chlorophyll and carotenoids that have observable absorption features in our spectral range,³⁶ our measurements showed there is no simple change in the spectral reflectance as the pepper ages




Figure 5. Temporal variation in the spectral reflectance of avocados. Each colored curve represents a new measurement day on the same avocado.

(Fig. 7). Though there may be some detectable change in the relative heights of the two large reflectance peaks or the slope of the reflectance ramp from approximately 600 nm to 700 nm, often denoted as the red edge, the ripening process in peppers is much more difficult to detect.

Tomatoes present a similar challenge. Tomato ripening is characterized by varying lycopene content in the skin, the breakdown of chlorophyll, and the build-up of carotenes.³⁷ 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. 8). While 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 incident light for an extended period, an increase in chlorophyll production occurs in a process known as greening.³⁸ 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. 9). 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 opposite 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.




Figure 6. Variation in the spectral reflectance of bell peppers. Each colored curve represents a different variety of bell pepper.




Figure 7. Temporal variation in the spectral reflectance of a green bell pepper. Each line represents a measurement of the same pepper taken on different days.




Figure 8. Temporal variation in the spectral reflectance of a tomato. Each line represents a measurement of the same tomato taken on different days.




Figure 9. 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

For our experiments, the networks with standard 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 an Nvidia GTX 1080TI GPU. We present the results presented of 10 fold cross-validation experiments.

We ran experiments with both RGB images and HSI images. We present the results for experiments with RGB images and the full spectrum for potatoes, bananas, and green peppers in Table 2. Particularly surprising is the relatively low performance of RGB images on bananas given their distinct browning as they age. The performance on green peppers was also surprising given less discernible visual differences as they age. Previous results using SpectrumNet on multi-spectral images when compared to competing methods such as ResNet-50³⁹ and DenseNet-161⁴⁰ show the relative superiority of our approach, both in terms of accuracy and computational complexity.³⁰

Table 2. SpectrumNet Produce Age Classification (Mean % Accuracy \pm One Standard Deviation)

Produce	RGB	Full Spectrum	HAGRID Wavelengths (5)
Potatoes	92.5 \pm 5.12	98.0 \pm 3.32	92.0 \pm 4.36
Bananas	83.4 \pm 9.27	95.4 \pm 4.38	96.0 \pm 2.61
Green Peppers	98.3 \pm 5.00	100.0 \pm 0.00	100 \pm 0.00

3.3 Genetic Algorithm

For our experiments, we set the population size to 1,000 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 changes the way in which the solution is selected from the final population of the GA, one complete run of the GA is used for the training the neural network.

The results for the GA-based feature selection are shown in the rightmost column of Table 2. For the potato, we see no appreciable improvement over the RGB images, but this is not surprising given the relative ease of seeing potatoes aging visually. We do find a significant improvement for banana age classification over RGB, and even some improvement (albeit not statistically significant) when compared to the full spectrum. The results on the green peppers was surprising in that the performance was perfect, as it was with the full spectrum. This reveals that the GA-based method is capable of finding the relevant wavelengths on difficult classification problems.

4. CONCLUSION

In this paper, we presented a promising method for detecting and predicting the ripening process of fruits and vegetables using a calibrated VNIR hyperspectral imager coupled with machine learning algorithms. Several calibrations and analyses were performed on the hyperspectral imaging system, including a radiometric calibration, conversion to reflectance spectra, pixel uniformity testing, image stability analysis, spectral performance testing, and an assessment of polarization response. Once calibrated, 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.

*We use the term "standard" convolutions to distinguish them from "depthwise separable" convolutions.

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 convolutional neural network called SpectrumNet trained using stochastic gradient descent. Experiments were run using both RGB images and hyperspectral images to compare performance. Results show that hyperspectral images outperform RGB images for age classification on all tested produce, including potatoes, bananas, and green peppers.

Feature selection was performed using a genetic algorithm (GA), with hyperparameters tuned using a grid search. Using the GA, we selected 5 salient wavelengths from the full hyperspectral spectrum and compared age classification performance against RGB images and hyperspectral imagery. The GA-based feature selection method typically outperformed RGB images, with a significant improvement on classification of bananas. When compared to hyperspectral imagery, the GA-based method showed modest improvement in classifying bananas and matched performance for green peppers. However, hyperspectral imagery outperformed all methods for classification of potatoes.

The GA-based feature selection method presented in this paper provides 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 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 by the Intel Corp. under Agreement Montana-Intel SRA 10-18-17. This material is also based upon 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.

REFERENCES

- [1] Gustavsson, J., Cederberg, C., Sonesson, U., Van Otterdijk, R., and Meybeck, A., "Global Food Losses and Food Waste - Extent, causes and prevention," in [*International Congress Save Food!*], Food and Agriculture Organization of the United Nations, Rome (2011).
- [2] Venkat, K., "The Climate Change and Economic Impacts of Food Waste in the United States," *International Journal on Food System Dynamics* **2**, 431–446 (dec 2011).
- [3] Gunders, D., "Wasted: How America Is Losing Up to 40 Percent of Its Food from Farm to Fork to Landfill," *NRDC Issue Papers* (12-06-B) (2012).
- [4] Melikoglu, M., Lin, C. S. K., and Webb, C., "Analysing global food waste problem: Pinpointing the facts and estimating the energy content," *Central European Journal of Engineering* **3**, 157–164 (jun 2013).
- [5] Hall, K. D., Guo, J., Dore, M., and Chow, C. C., "The Progressive Increase of Food Waste in America and Its Environmental Impact," *PLoS ONE* **4**, e7940 (nov 2009).
- [6] Buzby, J. C., Wells, H. F., and Hyman, J., "The Estimated Amount, Value, and Calories of Postharvest Food Losses at the Retail and Consumer Levels in the United States," *USDA-ERS Economic Information Bulletin* (121) (2014).
- [7] Li, M., Slaughter, D. C., and Thompson, J. F., "Optical chlorophyll sensing system for banana ripening," *Postharvest Biology and Technology* **12**, 273–283 (dec 1997).
- [8] Abbott, J. A., "Quality measurement of fruits and vegetables," *Postharvest Biology and Technology* **15**, 207–225 (1999).
- [9] Peirs, A., Lammertyn, J., Ooms, K., and Nicolai, B. M., "Prediction of the Optimal Picking Date of Different Apple Cultivars by Means of VIS/NIR-Spectroscopy," *Postharvest Biology and Technology* **21**, 189–199 (2000).
- [10] Kim, M. S., Chen, Y. R., and Mehl, P. M., "Hyperspectral Reflectance and Fluorescence Imaging System for Food Quality and Safety," *Transactions of the ASAE* **44**(3), 721–729 (2001).
- [11] Polder, G., Van Der Heijden, G. W. A. M., and Young, I. T., "Spectral Image Analysis For Measuring Ripeness of Tomatoes," *Transactions of the ASAE* **45**(4), 1155–1161 (2002).

- [12] Prasanna, V., Prabha, T. N., and Tharanathan, R. N., "Fruit Ripening Phenomena—An Overview," *Critical Reviews in Food Science and Nutrition* **47**, 1–19 (jan 2007).
- [13] Rajkumar, P., Wang, N., Elmasry, G., Raghavan, G. S. V., and Garipey, Y., "Studies on Banana Fruit Quality and Maturity Stages Using Hyperspectral Imaging," *Journal of Food Engineering* **108**, 194–200 (2012).
- [14] Cubero, S., Aleixos, N., Molto, E., Gomez-Sanchis, J., and Blasco, J., "Advances in machine vision applications for automatic inspection and quality evaluation of fruits and vegetables," *Food Bioprocess Technology* **4**, 487–504 (2011).
- [15] Pandey, R., Naik, S., and Marfatia, R., "Image processing and machine learning for automated fruit grading system: A technical review," *International Journal of Computer Applications* **81**(16), 29–39 (2013).
- [16] Naik, S. and Patel, B., "Image processing and machine learning for automated fruit grading system: A technical review," *International Journal of Computer Applications* **170**(9), 22–34 (2017).
- [17] Elhariri, E., El-Bendary, N., Hussein, A. M. M., Hassanien, A. E., and Badr, A., "Bell pepper ripeness classification based on support vector machinedetecting rottenness caused by pennicillium genus fungi in citrus fruits using machine learning techniques," *ICET 2014 - 2nd International Conference on Engineering and Technology* (2014).
- [18] Castro, W., Oblitas, J., De-La-Torre, M., Cotrina, C., Bazan, K., and Avila-George, H., "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] Femling, F., Olsson, A., and Alonso-Fernandez, F., "Fruit and vegetable identification using machine learning for retail applications," (2018).
- [20] Vaviya, H., Yadav, A., Vishwakarma, V., and Shah, N., "Identification of artificially ripened fruits using machine learning," *2nd International Conference on Advances in Science Technology (ICAST)* (2019).
- [21] Gomez-Sanchis, J., Martin-Guerrero, J. D., Soria-Olivas, E., Martinez-Sober, M., Magdalena-Benedito, R., and Blasco, J., "Detecting rottenness caused by pennicillium genus fungi in citrus fruits using machine learning techniques," *Expert Systems with Applications* **39**(1), 780–785 (2012).
- [22] Rivera, N. V., Gómez-Sanchis, J., Chanona-Pérez, J., Carrasco, J. J., Millán-Giraldo, M., Lorente, D., Cubero, S., and Blasco, J., "Early detection of mechanical damage in mango using nir hyperspectral images and machine learning," *Biosystems Engineering* **122**, 91 – 98 (2014).
- [23] Gupta, O., Das, A. J., Hellerstein, J., and Raskar, R., "Machine learning approaches for large scale classification of produce," *Scientific Reports* **8**(1) (2019).
- [24] El-Bendary, N., Hariri, E. E., Hassanien, A. E., and Badr, A., "Using machine learning techniques for evaluating tomato ripeness," *Expert Systems with Applications* **42**(4), 1892 – 1905 (2015).
- [25] Nugent, P. W., Shaw, J. A., Jha, P., Scherrer, B., Donelick, A., and Kumar, V., "Discrimination of herbicide-resistant kochia with hyperspectral imaging," *Journal of Applied Remote Sensing* **12**(1), 1 – 10 (2018).
- [26] Scherrer, B., Sheppard, J., Jha, P., and Shaw, J. A., "Hyperspectral imaging and neural networks to classify herbicide-resistant weeds," *Journal of Applied Remote Sensing* **13**(4), 1 – 15 (2019).
- [27] Tyo, J. S., Goldstein, D. L., Chenault, D. B., and Shaw, J. A., "Review of passive imaging polarimetry for remote sensing applications," *Applied Optics* **45**(22), 5453–5469 (2006).
- [28] Logan, R. D. and Shaw, J. A., "Polarization Response of the Pika L Hyperspectral Imaging System," *SPIE Proceedings* (2020).
- [29] Senecal, J., *Convolutional Neural Networks for Multi- and Hyper-Spectral Image Classification*, Master's thesis, Montana State University, Bozeman, MT (2019).
- [30] Senecal, J., Sheppard, J., and Shaw, J., "Efficient convolutional neural networks for multi-spectral image classification," in [*Proceedings of the IEEE International Joint Conference on Neural Networks (IJCNN)*], (2019).
- [31] Walton, N. S., Sheppard, J. W., and Shaw, J. A., "Using a genetic algorithm with histogram-based feature selection in hyperspectral image classification," in [*Proceedings of the ACM Genetic and Evolutionary Computation Conference (GECCO)*], (2019).
- [32] Walton, N., *Utilizing Distributions of Variable Influence for Feature Selection in Hyperspectral Images*, Master's thesis, Montana State University, Bozeman, MT (2019).
- [33] Slaughter, D. C. and Abbott, J. A., "Applications in Analysis of Fruits and Vegetables," *Near-Infrared Spectroscopy in Agri- culture, Agronomy Monograph no. 44* (44), 1–12 (2004).

- [34] Laval-Martin, D., Quennemet, J., and Monéger, R., "Pigment Evolution in *Lycopersicon Esculentum* Fruits During Growth and Ripening," *Phytochemistry* **14**(11), 2357–2362 (1975).
- [35] Bower, J. P. and Cutting, J. G., "Avocado fruit development and ripening physiology," tech. rep. (1988).
- [36] Luning, P. A., De Rijk, T., Wichers, H. J., and Roozen, J. P., "Gas Chromatography, Mass Spectrometry, and Sniffing Port Analyses of Volatile Compounds of Fresh Bell Peppers (*Capsicum annuum*) at Different Ripening Stages," *Journal of Agricultural and Food Chemistry* **42**, 977–983 (1994).
- [37] Polder, G., van der Heijden, G. W. A. M., and Young, I. T., "Tomato Sorting Using Independent Component Analysis on Spectral Images," *Real-Time Imaging* **9**, 253–259 (2003).
- [38] Grunenfelder, L., Hiller, L. K., and Knowles, N. R., "Color indices for the assessment of chlorophyll development and greening of fresh market potatoes," *Postharvest Biology and Technology* **40**, 73–81 (apr 2006).
- [39] He, K., Zhang, X., Ren, S., and Sun, J., "Deep residual learning for image recognition," in [*Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*], 770–778 (2017).
- [40] Huang, G., Liu, Z., Maaten, L. V. D., and Weinberger, K. Q., "Densely connected convolutional networks," in [*Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*], 2261–2269 (2017).