

Fine-Grain Prediction of Strawberry Freshness using Subsurface Scattering

Jeremy Klotz*, Vijay Rengarajan, and Aswin C. Sankaranarayanan
ECE Department, Carnegie Mellon University, Pittsburgh, PA, USA

*jklotz@andrew.cmu.edu

Abstract

Predicting fruit freshness before any visible decay is invaluable in the food distribution chain, spanning producers, retailers, and consumers. In this work, we leverage subsurface scattering signatures associated with strawberry tissue to perform long-term edibility predictions. Specifically, we implement various active illumination techniques with a projector-camera system to measure a strawberry’s subsurface scattering and predict the time when it is likely to be inedible. We propose a learning-based approach with captures under structured illumination to perform this prediction. We study the efficacy of our method by capturing a dataset of strawberries decaying naturally over time.

1. Introduction

Recent advances in agricultural imaging focus on specific types of fruit decay or mechanical injuries [9, 17], where we seek to understand which fruits among a collection are damaged or decayed. While this is invaluable in many post-harvest settings, various important applications benefit from a more nuanced and fine-grained prediction of fruit freshness; specifically, this prediction understands which fruits in a population are “fresher” and predicts *when* they would become inedible.

For example, predictions of how long fruits will remain edible would benefit a farmer by informing their harvest schedule, allowing them to optimize their yield for wholesalers’ demand. Wholesalers may also identify fruits with an extra day of freshness to transport to remote destinations without sacrificing food quality, thereby improving food equity for those in distant locations. These fine-grained predictions further inform consumers deciding which fruits to purchase and eat first. Precise freshness measurements for inedible fruits are also useful to estimate a decayed fruit’s exposure time to other products during transport and wholesale. Hence, it is not a stretch to suggest that a precise prediction of the freshness of a fruit in terms of time to decay benefits every actor in the food production, transportation, retail, and consumption chain.

To build a framework for predicting fruit freshness, we rely on measurements of how light scatters within fruit tissue. There are many physical signals that we can measure in this context. Perhaps the most informative amongst them is the bidirectional subsurface scattering reflectance distribution function (BSSRDF) [3], which measures light intensity transport between incoming and outgoing light ray pairs. Clearly, such a set would be unreasonably large with a long acquisition time. To simplify this capture, we can drop the angular dimension of light and simply measure the spatially-varying subsurface scattering kernel, which is measured by capturing a sequence of images with a single illuminated spot on the object. Noting that this kernel is often compact, we can speed up the acquisition by illuminating a point array on the object, measuring each kernel, and subsequently scanning the array across the object.

To simplify our measurements even further, we can drop the spatial nature of the illumination to obtain a very coarse measurement in the form of direct and global images. Here, we represent the light transport between a camera and projector with just two images: a direct image, comprised of single-bounce light paths, and a global image that includes the intensity of all other light paths, including those from subsurface scattering. Direct and global images can be rapidly measured [14], but they only provide a coarse understanding of the subsurface scattering functions. Clearly, each of these measurements describe subsurface scattering at a different level of precision and require setups with differing complexities and acquisition times.

This paper systematically studies the informativeness of various subsurface scattering measurements for fine-grained freshness predictions. We perform this study in the context of strawberries since they are widely consumed and readily prone to decay (or so has been the experience of the authors). We implement multiple imaging techniques to capture subsurface scattering measurements of varying spatial and scattering resolutions, each with differing acquisition times. We then compare the performance of single-shot and multi-shot measurements on a real dataset. Our idea can be adapted to other non-hard surface fruits such as apples, cherries, tomatoes, and grapes.

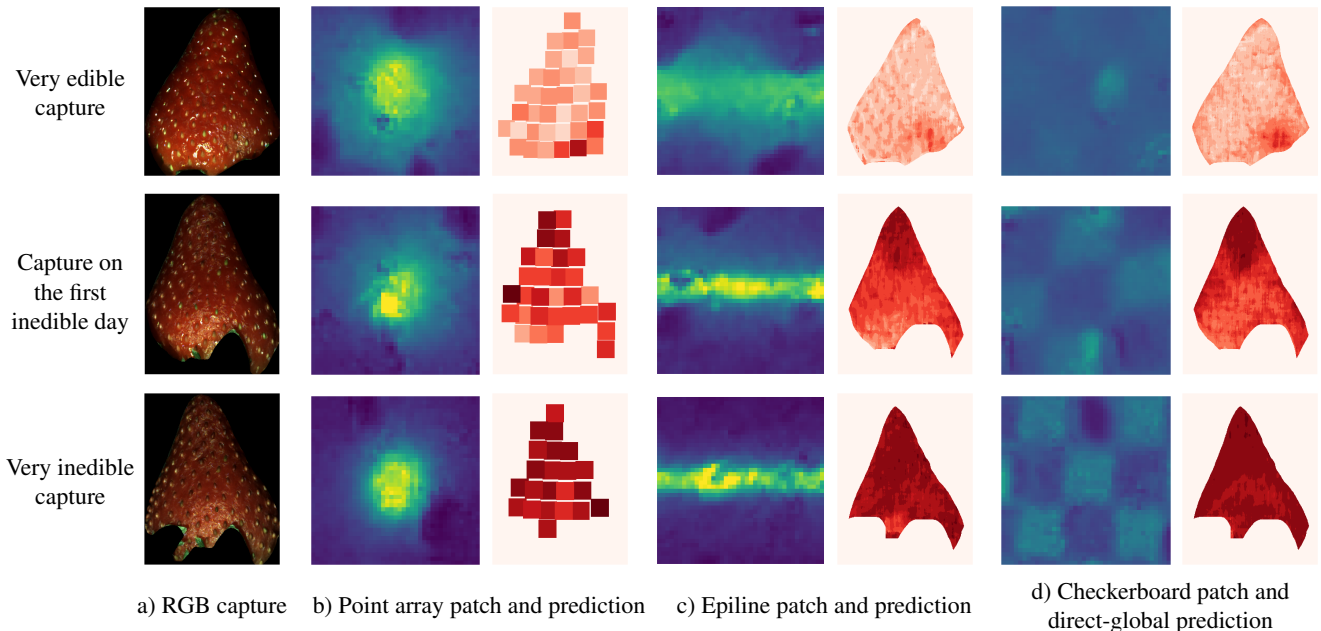


Figure 1: *Freshness predictions under different illuminations.* Each row represents the same fruit at different freshness stages, with the full fruit shown in a). Imaging the subsurface scattering around a single point-source in b) provides the highest resolution scattering measurement with the least spatial resolution. In c), epipolar measurements obtain scattering in one-dimension around a line and produce a dense prediction map after combining multiple captures. In d), direct-global separation uses checkerboard patterns to simultaneously measure scattering at many locations and produces a dense prediction map. While c) and d) improve spatial resolution, we find that dense spatial resolution is not necessary for fine-grained freshness predictions. Lighter red shades denote fresher patches, and darker red shades denote more inedible patches.

Contributions. Our main contributions are as follows:

1. We implement a low-cost camera-projector system to capture images of actively illuminated strawberries to predict their freshness.
2. We empirically validate the following methods' abilities to predict strawberry freshness: (i) point arrays, (ii) epipolar light transport, (iii) direct-global separation, (iv) structured-illumination reflectance imaging (SIRI), and (v) checkerboard patterns.
3. We capture a small dataset of 34 strawberries as they decay to evaluate each method's performance. Our dataset and code are available at https://github.com/Image-Science-Lab-cmu/Strawberry_ICCVW21.

Figure 1 shows three subsurface scattering measurements of a strawberry and their corresponding patch-wise prediction maps. As the fruit decays, point arrays and epilines show a clear drop in subsurface scattering around the directly-illuminated points, whereas checkerboard patterns exhibit global changes in intensity variation. We find that the higher resolution scattering measurements improve the fine-grained prediction accuracy, despite their coarse spatial

information. Specifically, the point array patches outperform every other method for fine-grained freshness predictions, achieving 33.6% accuracy using only four captures.

2. Related Works

We briefly discuss some of the key related works on fruit inspection, focusing broadly on visual techniques.

Visual signals for probing fruit decay. There are two primary optical measurements used to predict fruit freshness: light spectrum and scattering [21, 5, 20, 8, 15]. Hyperspectral imaging has been widely applied to non-invasive fruit inspection; specific examples of this include detecting defect in mangoes and citrus fruits [21, 5] as well as detecting fungus spores in dates [19]. Siedliska *et al.* [17] and Li *et al.* [6] use specific wavelengths in infrared spectra to detect fruit defects using discriminative analysis. Yet, despite spectrum being a powerful tool for analyzing material compositions, hyperspectral imaging can be extremely costly with long acquisition times, both of which pose a barrier to widespread use.

There have been many works that use subsurface scattering measurements to analyze produce. Laser-light backscat-

ter imaging (LLBI) is a fast, low-cost technique to measure subsurface scattering signatures of fruits. It captures the backscatter in highly-scattering media from a single point source, providing localized information about the tissue’s optical properties. Multiple works [10, 16, 11, 9, 8] have proposed techniques based on laser-light backscattering to detect fruit decay and identified optimal wavelengths for this prediction. While these techniques use features from subsurface scattering measurements, their predominant objective is to find the best laser wavelength that captures fruit decay as proposed by Lorente *et al.* [9, 8]. We implement backscatter imaging using a projector-camera pair to simultaneously acquire backscatter images at multiple locations. This enables high-resolution scattering measurements with rapid acquisition times to compare with other techniques with improved spatial resolution.

Spatial frequency domain imaging. A popular approach for probing subsurface scattering functions is spatial frequency domain imaging (SFDI), a low-cost, non-invasive technique that projects multiple sinusoids of different spatial frequencies and capturing the reflected image [2, 1]. This can be interpreted as probing the subsurface scattering functions in its frequency domain. However, the inverse problem of computing optical properties from captured images is computationally expensive.

Li *et al.* applied the concepts of SFDI to fruit inspection and proposed structured-illumination reflectance imaging (SIRI), a system for agriculture imaging that does not directly compute optical properties. Instead of computing optical properties, they project spatially-modulated sinusoids described in [1] and perform training and inference on a set of demodulated images derived from the captures [12, 13]. The demodulated images I_{DC} and I_{AC} correspond to the DC bias and amplitude envelope of the reflected photon density standing wave at the given spatial frequency, respectively. They found that I_{AC} , the amplitude envelope, increased contrast for bruised apple tissue before any visible signs [12]. Sun *et al.* [18] extended this imaging technique as a multi-spectral SIRI system to detect specific fungal infections in peaches using learning-based classifications technique. We follow a similar approach by measuring subsurface scattering and proposing a learning-based technique for patch-wise classification, but our ground-truth labels are not exclusive to any single type of decay.

Direct-global separation. This technique separates light transport into direct and global components from captures under high-frequency illumination [14]. The direct component of a scene contains single-bounce light and specular highlights, and effects such as interreflections and subsurface scattering comprise the global component. The direct-global separation method described in [14] is particularly

useful due to its simple implementation and rapid acquisition. Because the global component is simultaneously measured for many points on a scene, this method achieves poor scattering resolution as backscatter from adjacent locations may interfere. Despite this shortfall, we implement this method as it provides the highest spatial resolution.

Epipolar light transport. Epipolar light transport describes the light transport between a projector’s epipolar line and a camera pixel. This method provides improved scattering resolution over direct-global separation as light scattering orthogonal to an epipolar line exhibits less interference. Kubo *et al.* describe and capture this epipolar light transport using a rectified projector-camera pair. They exploit the camera’s rolling shutter and laser projector scanline to rapidly acquire measurements [4]. Liu *et al.* [7] extend this work to compute a target’s optical properties with a similar setup. This method provides improved scattering resolution over direct-global separation as light scattering orthogonal to an epipolar line exhibits less interference.

3. Problem Setup

Our goal is to classify a strawberry image into a freshness class, $C \in \mathbb{Z}$. $C = 0$ is the first day a strawberry is inedible. A negative C denotes an edible class, where its magnitude represents the number of days until the strawberry becomes inedible. A positive C denotes an inedible class and represents the number of days since the strawberry became inedible. For example, $C = -5$ indicates five days until becoming inedible and $C = 2$ suggests that the strawberry had become inedible two days prior.

Our idea is to rapidly capture the subsurface signatures of strawberries using a projector-camera system to predict strawberry freshness. At a high level, we implement multiple approaches for obtaining projections of the subsurface scattering function and train classifiers for each approach to predict the time to decay.

3.1. Imaging Setup

Our imaging setup uses a FLIR FL3-U3-13E4C-C RGB camera and an IncrediSonic Pico DMD projector. The strawberry is imaged in a partially enclosed diffuse black box to minimize interreflections. We cycle through different patterns on the projector and simultaneously capture images with the camera. Figure 2(a) shows the photostage with the camera and projector. Each strawberry is imaged everyday at two rotations separated by approximately 180° , and the initial rotation is chosen arbitrarily. To minimize condensation and its effects on the captured images, we store each strawberry in the refrigerator until it is ready to be imaged.

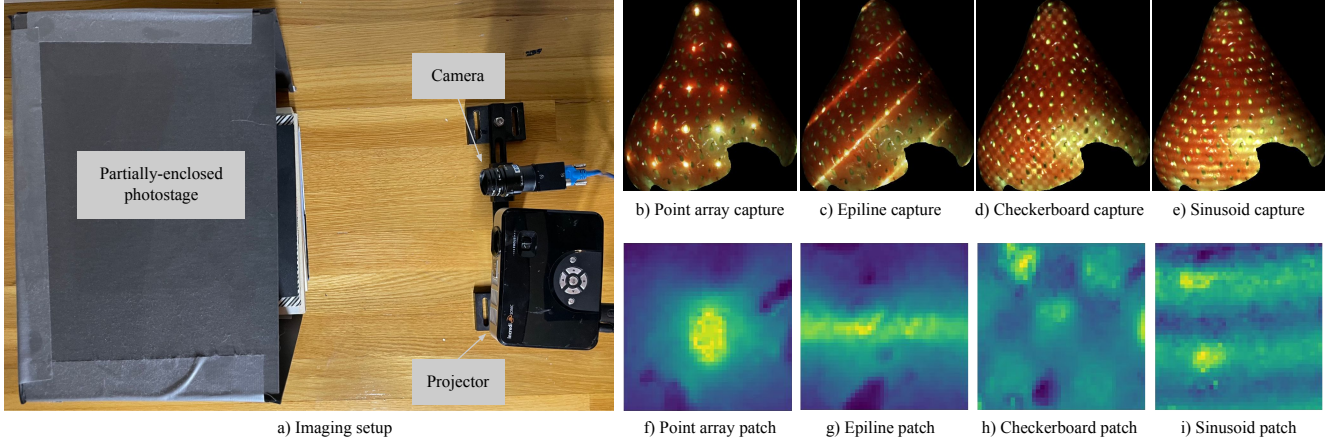


Figure 2: *Strawberry imaging setup and captures.* The camera-projector pair (left) images strawberries under different projections. The right images show example captures and patches of a single strawberry for point arrays, epipolar lines, checkerboards, and sinusoids. The single-channel patches are printed in false color to improve contrast.

3.2. Projector Patterns

We discuss the various projections that we implement in this subsection. Figure 2 shows an example capture and patch from each projection type described below.

Point array captures. We coarsely sample the tissue’s BSSRDF by projecting point-sources and imaging the backscatter. This accurately measures local subsurface scattering around each projected point. To avoid the long acquisition time described above, we project a grid of points with fixed spacing to spatially multiplex the measurements. We then shift the points by large step sizes to increase the spatial sampling period. Each fruit is a collection of non-overlapping patches centered around a point-source that contain subsurface scattering signatures. We compare multi-shot captures using four shifted point arrays with a single point array capture. These features are denoted “PA-4” and “PA-1” to denote the multi-shot and single-shot point array, respectively. We note that the spatial multiplexing requires sufficient spacing between the projected points. There must be a minimum separation equal to the tissue’s scattering distance to avoid interference in the measured radiance due to adjacent projections.

Epipolar line captures. While point array captures provide radial signatures around directly illuminated points, it limits the number of patches in a single capture. Based on the epipolar light transport proposed in [4], we trade scattering resolution in one direction for increased patch density by projecting an epipolar line instead of a point. A projected epipolar line provides the scattering signature orthogonal to the line, and the epipolar geometry eliminates any distortion due to the three-dimensional shape. Similar to the point ar-

rays, we spatially multiplex epipolar lines to simultaneously scan multiple lines in a single image.

Formally, we obtain an image stack I_e of $N \times H \times W$ captures, where each $H \times W$ capture contains spatially-multiplexed epipolar lines that are scanned across the scene. We find the capture index i where pixel $p = (r, c)$ is directly illuminated by solving

$$i = \arg \max_n I_e(n, r, c) \quad (1)$$

Each pixel maps to a scattering profile, an M -length vector with intensity values that vary with the distance to the projected line, where capture i corresponds to the capture with the line directly on top of the measured pixel. Each image patch is then a collection of M -length scattering profiles for each pixel. We use the term “EPI” for the scenario where we obtain pixel-wise scattering profiles from multiple captures. We also compare results using the scattering around epipolar lines from only a single capture. This single-shot feature is denoted “EPI-1.”

Checkerboard captures for direct-global separation.

To implement direct-global separation, we project multiple checkerboard patterns shifted in the x and y directions. We apply the direct-global separation algorithm to the capture stack to obtain direct and global images [14]. These direct and global images comprise the “DG” feature, a multi-shot scattering measurement. We also train on the captures under checkerboard projections before the direct-global separation. These captures, denoted “CHK-1,” provide unprocessed subsurface scattering measurements in a single capture. With an optimal checkerboard size, a learning-based approach should be able to adapt to these shifted features.

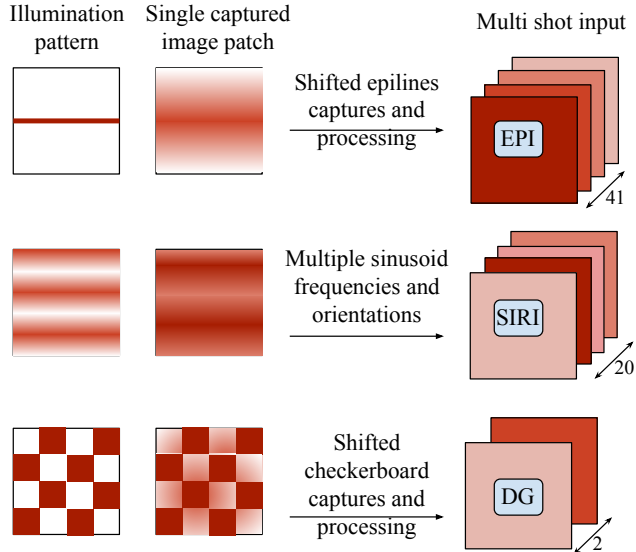


Figure 3: *Illustration of multi-shot illumination patterns.* Epipolar, SIRI, and direct-global images require multiple images to be captured under the illumination of multiple shifts and orientations of a specific pattern.

SIRI. We implement SIRI measurements, which capture images under spatially-modulated sinusoids to probe the spatial modulation transfer function [12]. For a highly scattering medium illuminated by a spatially modulated sinusoid, the captured image is modeled as:

$$I(x, y) = I_{DC} + I_{AC} \cos(wx + \phi) \quad (2)$$

If the images I_0, I_1, I_2 are captured under a sinusoid with phase shifts $\phi_i \in [-120^\circ, 0^\circ, 120^\circ]$, then I_{DC} and I_{AC} are recovered by three-phase demodulation [12]:

$$I_{DC} = \frac{1}{3} (I_0 + I_1 + I_2) \quad (3)$$

$$I_{AC} = \frac{\sqrt{2}}{3} \sqrt{(I_1 - I_0)^2 + (I_2 - I_1)^2 + (I_2 - I_0)^2} \quad (4)$$

Given that the illuminant phases are separated by exactly 120° , three-phase demodulation provides a simple pixel-wise demodulation scheme. Lu *et al.* note that I_{DC} is equivalent to the capture under uniform illumination [12].

We implement SIRI by projecting sinusoids at five spatial frequencies. Each spatial frequency includes captures at three phase shifts separated by 120° modulated in both the horizontal and vertical direction. The set of AC and DC images for each spatial frequency and orientation is denoted as the ‘‘SIRI’’ feature. An illustration of multi-shot illumination techniques is shown in Figure 3.

Uniform illumination. Unlike the methods described above, uniform illumination does not extract any subsurface scattering features. We capture an RGB image under

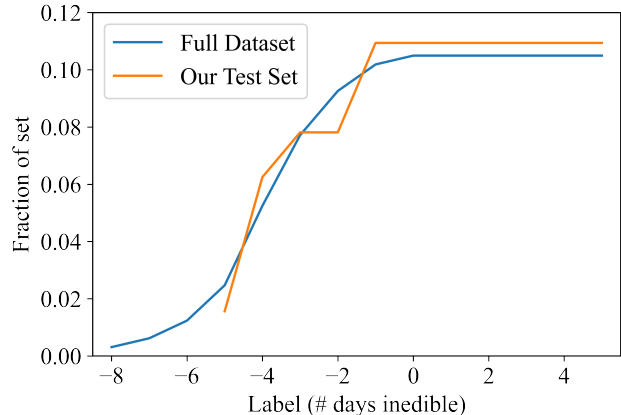


Figure 4: Freshness label distribution of our full and test datasets.

uniform illumination for baseline comparison. This feature is denoted ‘‘UNIF.’’

4. Dataset Acquisition

Fruit preparation and labeling. Two boxes of non-organic strawberries were purchased from a local grocery store on January 14, 2021. Every strawberry is rinsed exactly once immediately after the purchase. We manually choose 34 strawberries with flat calyxes to comprise the dataset.¹ Strawberries are stored in egg cartons labeled $s_i \in [0, 34]$ in a refrigerator.

During each capture session, we manually inspect the freshness of each strawberry and note any obvious bruising or decay. Every strawberry freshness label is the subjective judgment of one author. If the author deems a strawberry inedible, that date is noted as ‘‘Day 0’’ for the strawberry, the first day the strawberry is no longer edible. While multiple annotators would have improved the ground-truth labeling, the single annotator attempted to follow the same criteria when labeling each strawberry.

The strawberry edibility is independent of rotation. While captures of the same strawberry at the two rotations are considered independently, they correspond to the same label despite the location of any decay.

Dataset statistics. Figure 4 shows the distribution of strawberry labels for the strawberry dataset collected over 14 days along with the fixed test set described below in section 5.2. The distribution is biased toward inedible strawberries as a result of the initial distribution of strawberry freshness at the time of purchase.

¹Strawberry calyx must be relatively flat for the berry to sit upside down on the imaging stage.

Table 1: Strawberry inedibility reasons.

Inedibility Reason	Full Dataset	Test Set
Shrinkage due to water loss	26	6
Bruising or physical damage	7	1
Suspected mold	1	0

The dataset is composed of strawberries that decayed due to water loss, localized bruising, and one instance of suspected mold. Table 1 shows the inedibility reasons for the strawberries in the full dataset and the test set.

Image preprocessing. Masks are generated in the normalized RGB colorspace to remove the background and leaves. Each capture corresponds to reference captures under uniform illumination, I_{white} , and no illumination, I_{dark} . To correct for non-uniform illumination, we compute the corrected image I_{corr} from the capture I_c :

$$I_{corr} = \frac{I_c - I_{dark}}{I_{white} - I_{dark}} \quad (5)$$

Epiline and point array patches are further normalized to the mean intensity around the directly illuminated points.

5. Experiments

In this section, we first describe the neural network architecture followed by a description of training and testing data partitioning and fruit classification results with various subsurface scattering measurements. In addition to freshness predictions on every strawberry, we examine results on edible and inedible strawberries separately. Finally, we comment on capture and inference times of each projection type.

5.1. Neural Network Architecture and Training

We employ a convolutional neural network architecture that takes an image stack $I : N \times 40 \times 40$ as the input and outputs a 14-length vector corresponding to edibility class probabilities. The image patch’s fixed width and height correspond to the scattering distance of the strawberry tissue. The network consists of three convolutional layers each with 512 channels followed by a final convolutional layer with 1024 channels. Each convolutional layer is followed by a ReLU non-linearity, and a 2×2 max-pooling down-samples the feature maps after the last three convolutional layers. The convolutional layers are followed by a fully-connected layer that converts the flattened 25600×1 feature map into a 14-length vector, followed by a soft-max operator to output a probability distribution of single-day freshness predictions.

To account for uncertainty in ground-truth labeling, each patch corresponds to a ground-truth normalized Gaussian

function with a standard deviation of 1.5 centered around the true label instead of a one-hot vector. We optimize the mean squared error of the predicted and ground-truth probability distributions during training.

Data augmentations on the training patches include small, random rotations, random flips across the vertical axis, and random vertical jitters for epilene patches. The network is implemented in PyTorch and trained on four Nvidia GTX Titan X GPU’s. Each model is trained for at least 2000 iterations of 512-length mini-batches with the Adam optimizer and a fixed learning rate of 10^{-3} . Training takes approximately 2 hours per model. A network is trained individually for each of the projection types. Every network is trained and tested on captures with integer labels to learn freshness with 1-day precision.

5.2. Results

Fixed training and testing sets. The dataset is randomly split into 27 training strawberries and 7 testing strawberries. Non-overlapping patches are extracted from each training strawberry to comprise the training dataset. We perform 10-fold cross-validation to train patch-wise classifiers for each subsurface scattering feature.

Prior to data augmentations, there are 16,075 training patches for point arrays. The epipolar, direct-global, SIRI, and uniform illumination features contain 7,638 training patches because they initially draw larger non-overlapping patches to enable random rotations. There are 322,270 training epilene patches and 871,120 training checkerboard patches since each strawberry provides 40 captures under these two projections.

A single strawberry prediction is obtained by performing forward pass on its patches. The patch-wise probability distributions are averaged together to obtain a global prediction.

Due to the dataset’s class imbalance described in Section 4, we report the balanced accuracy, the average recall for each class. Formally, given N class labels $i \in C$, each with N_i samples, we let g_{ij} and p_{ij} be the ground-truth and predicted label of the j ’th sample in class i , respectively. We obtain the balanced accuracy by averaging each class-wise recall:

$$\frac{1}{N} \sum_{i \in C} \frac{1}{N_i} \sum_j \mathbb{1}(g_{ij} = p_{ij}) \quad (6)$$

We also define the average k -accuracy, the balanced accuracy for predictions to $\pm k$ days from the true class:

$$\frac{1}{N} \sum_{i \in C} \frac{1}{N_i} \sum_j \mathbb{1}(|g_{ij} - p_{ij}| \leq k) \quad (7)$$

Figure 5 summarizes the test strawberry accuracies for each subsurface scattering measurement in predicting

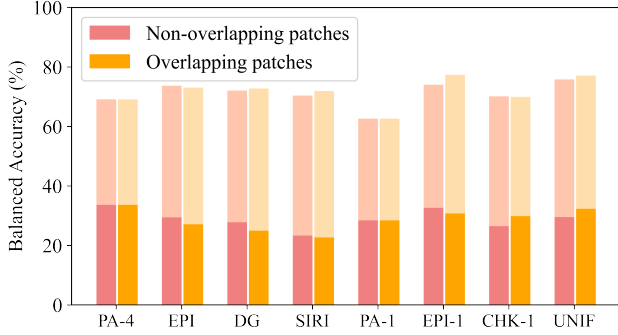


Figure 5: Test strawberry balanced accuracies with non-overlapping and overlapping patches. The lower bars (darker shade) show the accuracy with single-day precision and the upper bars (lighter shade) show the accuracy to ± 1 day of the true label.

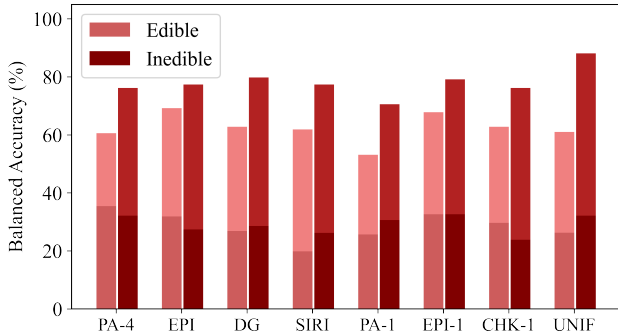


Figure 6: Test strawberry balanced accuracies on edible and inedible strawberries separately. The lower bars (lighter shades) show the accuracy with single-day precision and the upper bars (darker shades) shows the accuracy to ± 1 day of the true label.

freshness with non-overlapping and overlapping patches. For fine-grained predictions with single-day precision, the multi-shot point arrays (PA-4) outperform every other feature with a 33.6% prediction accuracy using non-overlapping patches. For predictions to ± 1 day of the true freshness, the uniform illumination captures (UNIF) achieve the highest accuracy with non-overlapping patches, and the single-shot epilines (EPI-1) achieve the highest accuracy with overlapping patches.

Figure 6 shows the same feature accuracies on edible and inedible strawberries separately using non-overlapping patches. For the fine-grained freshness predictions, the multi-shot point arrays (PA-4) are best for edible strawberries, and the single-shot epilines (EPI-1) are best for inedible strawberries. For predictions to ± 1 day, the multi-shot epipolar feature (EPI) is best for edible strawberries while the uniform illumination (UNIF) feature outperforms every




Figure 7: *Dataset permutations.* Validation strawberry balanced accuracy mean (blue), maximum (green), and minimum (orange) over the dataset permutations. Error bars represent unit standard deviation.

Table 2: Point array accuracies and patch coverage using a varying number of captures.

Feature	Predict Accuracy (%)	Patch Coverage (%)
PA-1	28.4	22.4
PA-2	32.3	49.0
PA-3	33.6	69.4
PA-4	33.6	96.0

other scattering measurement on inedible strawberries.

Dataset permutations. Due to the small test set, we train on permutations of the training and validation sets without a test set. The dataset is randomly partitioned 20 times into 24 training and 10 validation strawberries, and we train a single model for each partition.

Figure 7 shows the validation balanced accuracy distribution for each feature over the 20 permutations. Figure 8 shows the mean validation balanced accuracies for edible and inedible strawberries separately.

Point arrays consistently outperform all other scattering features, and the single-shot point array exhibits only a small loss in accuracy. Although uniform illumination captures show improved accuracies for inedible strawberries in Figure 8, the point arrays remain best. Finally, we note that every feature is more accurate on inedible strawberries.

5.3. Point Array Ablation Study

We perform an ablation study of point arrays on the fixed test set by performing inference on a subset of the 4 shifted captures to limit the fruit’s spatial sampling. From Table 2, the marginal benefit of adding another shifted capture to increase the patch coverage quickly diminishes after two captures. Clearly, dense spatial information is not the sole requirement for accurate fine-grained freshness predictions.

5.4. Capture and Inference Times

Table 3 summarizes each feature’s captures and performance in fine-grained predictions with single day precision

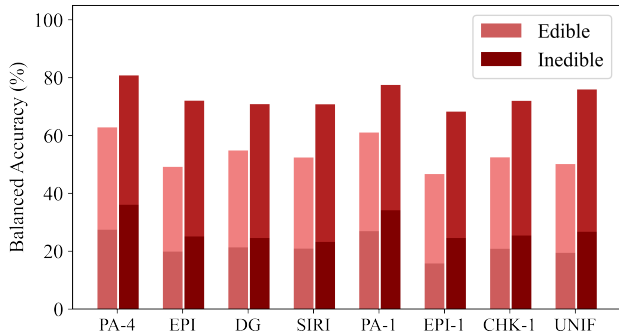


Figure 8: Mean validation strawberry balanced accuracies for edible and inedible strawberries separately over the dataset permutations. The lower bars (darker shades) show the accuracy with single-day precision and the upper bars (lighter shades) show the accuracy to ± 1 day of the true label.

using non-overlapping patches on the fixed test set. Each method requires two reference images to compute I_{corr} in eq. (5) that are not included in the capture count.

The multi-shot point array captures provide the highest accuracy with a relatively small number of captures and input channels. While the epipolar feature requires all 40 captures to align pixel-wise scattering profiles in eq. (1), the epilines measurements provide improved accuracy using only one capture from that set.

6. Discussion

The point array captures outperformed all other features in fine-grained freshness predictions with the exception of coarse predictions on inedible strawberries. Point arrays provided the sparsest spatial sampling; however, most of the measured strawberries exhibited global scattering changes as they decayed due to water loss. If water content is the most discriminative feature, higher frequency spatial sampling may not be necessary for freshness predictions.

Inference is performed both on non-overlapping and overlapping patches. We find no significant accuracy gain from overlapping patches to justify the increased inference time. This aligns with the suspected advantages of point arrays; high spatial resolution is unnecessary for fine-grained freshness predictions on our dataset.

The low-cost imaging system posed multiple limitations. Primarily, the measurements are captured on the visible spectrum. An infrared camera and projector may have probed deeper subsurface scattering. Secondly, the projector’s narrow depth of field made the projection sensitive to scene depth. In-focus projections are essential to ensure any captured signal is due to subsurface scattering and not defocus blur. Finally, specular highlights diminish the robust-

Table 3: Summary of captures.

Label	Feature	Pred. Acc. (%)	Input Chan.	# Captures excl. two ref.
PA-4	Point Arrays	33.6	1	4
EPI	Epipolar	29.4	41	40
DG	Direct-Global	27.8	2	40
SIRI	SIRI	23.3	20	30
PA-1	PA single shot	28.4	1	1
EPI-1	EPI single shot	32.6	1	1
CHK-1	Checkerboard	26.5	1	1
UNIF	Uniform illum.	29.5	1	1

ness of the feature preprocessing, particularly for epilines preprocessing.

Our dataset is limited by its size and label imbalance. The large variance in cross-validation results is likely due to the small dataset’s inability to capture all strawberry decay. There exists a wide variety of fungal infections and decay that are not measured in this dataset; a large dataset would help identify specific types of decay. In addition, fusing multiple illumination patterns at the input or feature level of neural network inference may improve the accuracy of our approach.

7. Conclusion

We have proposed and validated a learning-based approach to predict strawberry freshness from subsurface scattering measurements. We have found that subsurface scattering measurements with point arrays and epilines are best for fine-grained freshness predictions on our dataset. These captures that trade spatial resolution to measure the highest resolution scattering signatures improve freshness predictions, particularly for edible fruits. The uniform illumination captures outperformed the scattering measurements only for coarser predictions to ± 1 day, an effect particularly seen on inedible fruits. All of these results suggest that subsurface scattering measurements provide a promising avenue for a fine-grain prediction of freshness that goes beyond the traditional framing of edible vs. inedible fruits.

8. Acknowledgments

The authors acknowledge support via the National Science Foundation (NSF) under the following awards: CAREER CCF-1652569, the Expeditions in Computing Award 1730147, and the SaTC award 1801382.

References

- [1] David J Cuccia, Frederic Bevilacqua, Anthony J Durkin, Frederick R Ayers, and Bruce J Tromberg. Quantitation and mapping of tissue optical properties using modulated imag-

- ing. *Journal of Biomedical Optics*, 14(2):024012–0240113, 2009. 3
- [2] David J. Cuccia, Frederic Bevilacqua, Anthony J. Durkin, and Bruce J. Tromberg. Modulated imaging: Quantitative analysis and tomography of turbid media in the spatial-frequency domain. *Optics Letters*, 30(11):1354–1356, Jun 2005. 3
- [3] Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. A practical model for subsurface light transport. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '01*, page 511–518, New York, NY, USA, 2001. Association for Computing Machinery. 1
- [4] Hiroyuki Kubo, Suren Jayasuriya, Takafumi Iwaguchi, Takuya Funatomi, Yasuhiro Mukaigawa, and Srinivasa G Narasimhan. Programmable non-epipolar indirect light transport: Capture and analysis. *IEEE Transactions on Visualization and Computer Graphics*, 27(4):2421–2436, 2021. 3, 4
- [5] Jiangbo Li, Wenqian Huang, Xi Tian, Chaopeng Wang, Shuxiang Fan, and Chunjiang Zhao. Fast detection and visualization of early decay in citrus using Vis-NIR hyperspectral imaging. *Computers and Electronics in Agriculture*, 127:582–592, 2016. 2
- [6] Jiangbo Li, Wei Luo, Zheli Wang, and Shuxiang Fan. Early detection of decay on apples using hyperspectral reflectance imaging combining both principal component analysis and improved watershed segmentation method. *Postharvest Biology and Technology*, 149:235–246, 2019. 2
- [7] Chao Liu, Akash K. Maity, Artur W. Dubrawski, Ashutosh Sabharwal, and Srinivasa G. Narasimhan. High resolution diffuse optical tomography using short range indirect subsurface imaging. In *IEEE International Conference on Computational Photography (ICCP)*, pages 1–12, 2020. 3
- [8] D Lorente, M Zude, C Idler, J Gómez-Sanchis, and J Blasco. Laser-light backscattering imaging for early decay detection in citrus fruit using both a statistical and a physical model. *Journal of Food Engineering*, 154:76–85, 2015. 2, 3
- [9] D Lorente, M Zude, C Regen, L Palou, J Gómez-Sanchis, and J Blasco. Early decay detection in citrus fruit using laser-light backscattering imaging. *Postharvest Biology and Technology*, 86:424–430, 2013. 1, 3
- [10] Renfu Lu. Multispectral imaging for predicting firmness and soluble solids content of apple fruit. *Postharvest Biology and Technology*, 31(2):147–157, 2004. 3
- [11] Renfu Lu, Haiyan Cen, Min Huang, and Diwan P Ariana. Spectral absorption and scattering properties of normal and bruised apple tissue. *Transactions of the ASABE*, 53(1):263–269, 2010. 3
- [12] Yuzhen Lu, Richard Li, and Renfu Lu. Detection of fresh bruises in apples by structured-illumination reflectance imaging. In *Proceedings of SPIE - The International Society for Optical Engineering*, volume 9864, pages 986406–986406–12. SPIE, 2016. 3, 5
- [13] Yuzhen Lu and Renfu Lu. Structured-illumination reflectance imaging for the detection of defects in fruit: Analysis of resolution, contrast and depth-resolving features. *Biosystems Engineering*, 180:1–15, 2019. 3
- [14] Shree Nayar, Gurunandan Krishnan, Michael Grossberg, and Ramesh Raskar. Fast separation of direct and global components of a scene using high frequency illumination. *ACM Transactions on Graphics*, 25(3):935–944, 2006. 1, 3, 4
- [15] Zhaoshen Qing, Baoping Ji, and Manuela Zude. Non-destructive analyses of apple quality parameters by means of laser-induced light backscattering imaging. *Postharvest Biology and Technology*, 48(2):215–222, 2008. 2
- [16] Giuseppe Romano, László Baranyai, Klaus Gottschalk, and Manuela Zude. An approach for monitoring the moisture content changes of drying banana slices with laser light backscattering imaging. *Food and Bioprocess Technology*, 1(4):410–414, 2008. 3
- [17] Anna Siedliska, Piotr Baranowski, Monika Zubik, Wojciech Mazurek, and Bożena Sosnowska. Detection of fungal infections in strawberry fruit by VNIR/SWIR hyperspectral imaging. *Postharvest Biology and Technology*, 139:115–126, 2018. 1, 2
- [18] Ye Sun, Renfu Lu, Yuzhen Lu, Kang Tu, and Leiqing Pan. Detection of early decay in peaches by structured-illumination reflectance imaging. *Postharvest Biology and Technology*, 151:68–78, 2019. 3
- [19] M.A Teena, A Manickavasagan, L Ravikanth, and D.S Jayas. Near infrared (NIR) hyperspectral imaging to classify fungal infected date fruits. *Journal of Stored Products Research*, 59:306–313, 2014. 2
- [20] Xi Tian, Shuxiang Fan, Wenqian Huang, Zheli Wang, and Jiangbo Li. Detection of early decay on citrus using hyperspectral transmittance imaging technology coupled with principal component analysis and improved watershed segmentation algorithms. *Postharvest Biology and Technology*, 161:111071–, 2020. 2
- [21] Nayeli Vélez Rivera, Juan Gómez-Sanchis, Jorge Chanona-Pérez, Juan José Carrasco, Mónica Millán-Giraldo, Delia Lorente, Sergio Cubero, and José Blasco. Early detection of mechanical damage in mango using NIR hyperspectral images and machine learning. *Biosystems Engineering*, 122:91–98, 2014. 2