

Joint 2D to 3D image registration workflow for comparing multiple slice photographs and CT scans of apple fruit with internal disorders

Dirk Elias Schut^{a,*}, Rachael Maree Wood^b, Anna Katharina Trull^c, Rob Schouten^d, Robert van Liere^{a,e}, Tristan van Leeuwen^{a,f}, Kees Joost Batenburg^{a,g}

^a Computational Imaging Group, Centrum Wiskunde en Informatica (CWI), Science Park 123, 1098 XG Amsterdam, the Netherlands

^b Horticulture and Product Physiology, Wageningen University and Research, Droevendaalsesteeg 1, 6708 PB Wageningen, the Netherlands

^c Greefa Machinebouw B.V., Langstraat 12, 4196 JB Tricht, the Netherlands

^d Wageningen Food and Biobased Research, Bornse Weiland 9, 6708 WG Wageningen, the Netherlands

^e Visualization Group, Eindhoven University of Technology, PO Box 513, 5600 MB Eindhoven, the Netherlands

^f Mathematisch Instituut, Utrecht University, Budapestlaan 6, 3584 CD Utrecht, the Netherlands

^g Leiden Institute of Advanced Computer Science (LIACS), Leiden University, Niels Bohrweg 1, 2333 CA Leiden, the Netherlands

ARTICLE INFO

Keywords:

Image registration
Internal browning
Non-destructive testing (NDT)
Transformation model
Automatic differentiation
Deep learning

ABSTRACT

A large percentage of apples are affected by internal disorders after long-term storage, which makes them unacceptable in the supply chain. CT imaging is a promising technique for in-line detection of these disorders. Therefore, it is crucial to understand how different disorders affect the image features that can be observed in CT scans. This paper presents a workflow for creating datasets of image pairs of photographs of apple slices and their corresponding CT slices. By having CT and photographic images of the same part of the apple, the complementary information in both images can be used to study the processes underlying internal disorders and how internal disorders can be measured in CT images. The workflow includes data acquisition, image segmentation, image registration, and validation methods. The image registration method aligns all available slices of an apple within a single optimization problem, assuming that the slices are parallel. This method outperformed optimizing the alignment separately for each slice. The workflow was applied to create a dataset of 1347 slice photographs and their corresponding CT slices. The dataset was acquired from 107 'Kanzi' apples that had been stored in controlled atmosphere (CA) storage for 8 months. In this dataset, the distance between annotations in the slice photograph and the matching CT slice was, on average, 1.47 ± 0.40 mm. Our workflow allows collecting large datasets of accurately aligned photo-CT image pairs, which can help distinguish internal disorders with a similar appearance on CT. With slight modifications, a similar workflow can be applied to other fruits or MRI instead of CT scans.

1. Introduction

Apples (*Malus × domestica* Borkh.) are often stored long-term to supply markets with a high-quality product throughout the year (Wood et al., 2022b). Long-term storage of apples is possible by utilizing cold storage under a controlled atmosphere (CA) where oxygen is lowered and carbon dioxide is increased (Neuwalde et al., 2021). However, these conditions involve a loss of fruit quality, allowing physiological disorders and decay to develop. Postharvest losses caused by physiological disorders and decay can range between 18–27% under commercial CA conditions (Argenta et al., 2021). Therefore, detecting and removing affected fruit from the supply chain is crucial for reducing food wastage.

However, disorders such as internal browning (IB) do not display any external disorders making in-line detection challenging. Furthermore, different patterns of IB occur in different regions of fruit tissue (Sidhu et al., 2023). IB can cause 34% of postharvest losses in sensitive cultivars even when stored under optimal CA conditions (Wood et al., 2022a). Accurately detecting internal disorders and defects is critical for maintaining consumer confidence.

X-ray-based imaging techniques are promising for the in-line detection of internal disorders. There are two main X-ray-based imaging modalities: radiography, which uses a single 2D X-ray image (radiograph), and computed tomography (CT), which combines the data of many radiographs of the same object to create a 3D representation of the

* Corresponding author.

E-mail address: dirk.schut@cw.nl (D.E. Schut).

<https://doi.org/10.1016/j.postharvbio.2024.112814>

Received 25 August 2023; Received in revised form 25 January 2024; Accepted 25 January 2024

Available online 15 February 2024

0925-5214/© 2024 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

object. CT scanning offers 3D information and higher contrast, but radiographs can be acquired much faster. While both modalities are promising, radiography currently has more potential for practical applications in in-line quality control of produce, because it has a lower acquisition time (Kotwaliwale et al., 2014). Nevertheless, the acquisition time of CT scanning can be lowered by optimizing the scanning setup for in-line use (De Schryver et al., 2016; Morton et al., 2009; Schut et al., 2022). Moreover, radiographs can be accurately simulated from CT scans (Van Aarle et al., 2016), and if labels or segmentation masks are available on the CT scan, these can also be transferred to the simulated radiograph. Zeegers et al. (2022) developed a workflow for foreign object detection on radiographs where a few CT scans with segmentation masks of the foreign objects were used to simulate many radiographs with segmentation masks to train a neural network. This method was extended and applied to predict the ripeness of avocados by Andriashen et al. (2023).

CT scanning of apples has been used for studying the processes underlying bruising (Diels et al., 2017), internal browning (Herremans et al., 2013; Chigwaya et al., 2021), watercore (Herremans et al., 2014) and bitterpit (Si and Sankaran, 2016; Jarolmasjed et al., 2016). However, several conditions can look similar in CT scans, so destructive methods may still be needed to classify specific disorders. Watercore in apples appears as high-intensity regions in CT images due to the flooding of pores in the water-soaked regions of affected fruit (Herremans et al., 2014). High-intensity regions are also reported to be associated with IB (Chigwaya et al., 2021). In some markets, apples with watercore are sold at a premium due to the sweet juicy taste (Herremans et al., 2014; Mink, 1973), whereas IB is associated with off-flavor and is an unwanted trait (Hatoum et al., 2014). Furthermore, watercore can dissipate during storage in mildly affected fruit (Upchurch and Throop, 1994), while IB will not. Similarly, bruises and bitter pit can appear as small low-intensity regions close to the fruit peel in CT scans (Diels et al., 2017; Jarolmasjed et al., 2016). Bitter pit progressively worsens with CA storage (Jarolmasjed et al., 2016), while fruit with mechanical damage or bruising can still remain healthy during long-term storage. Therefore, fruits must still be cut and visually examined for further development of detection algorithms for specific disorders.

When fruit are CT or MRI scanned for research, they are often also cut and photographed. It is helpful to present the photographic images alongside CT or MRI images, which requires aligning the photographs to the CT or MRI scan. Photographs have been aligned manually to CT or MRI scans for investigating disorders in apples (Chigwaya et al., 2021; Herremans et al., 2014; van Dael et al., 2019) and pears (Azadbakht et al., 2019; van Dael et al., 2017; Van De Looverbosch et al., 2021, 2020). However, there are significant downsides to manually aligning the photographs. Firstly, manually aligning photographs is tedious and time-consuming, which limits the dataset size. It is often unknown whether an apple has an internal defect before CT scanning, and there is natural variation in how internal defects affect the CT measurements. Therefore, acquiring large datasets is valuable because it increases the chance of capturing the variation. Moreover, having large and varied datasets helps with training machine learning methods. Secondly, manual alignment is challenging in slices without features that are clearly visible on both the CT slice and the photograph, such as the image pair displayed in Fig. 1. The apple core can sometimes be used as a shared visual feature because it is visible in CT and photographs but not in all slices. Thirdly, the alignment error is unknown. For small features (e.g. sepal or petal bundles), the alignment error may be larger than the feature size. In such cases, finding the corresponding position on the CT slice is challenging, and there is a risk of finding a false correspondence. Knowing the alignment error gives an indication of which size of features can reliably be matched.

In this paper, we present a workflow for creating datasets of slice photographs and their corresponding CT slices, which solves the three issues mentioned above. Image registration (automatic alignment) was used instead of manual alignment, making it possible to create large




Fig. 1. Registered image pair of a slice photograph and a CT slice generated using our method.

datasets with little extra effort. The alignment quality is not affected if a slice has little shared visual information inside of the apple (Fig. 1), because the image registration only uses the outer shape of the apple. Moreover, the alignment error can be calculated by using the position of the apple core, because the apple core was not used for the image registration. This paper is the first work where image registration is applied to 3D scans and photographs of sliced fruit. We expect that with slight modifications, a similar workflow could be applied to other fruits or MRI scans instead of CT scans.

There are two additional scientific contributions in this paper. Firstly, the problem formulation of the image registration method is new. It uses a task-specific transformation model to describe the alignment of multiple parallel 2D slices to a 3D volume, which is optimized jointly by using automatic differentiation. In this approach, information is shared between multiple slices, which improves the overall registration accuracy. Secondly, the presented workflow has been applied to create a dataset of 1347 slice photographs acquired from 107 ‘Kanzi’ apples with registered CT slices, which was made public for further research.

2. Materials and methods

2.1. Apple samples

In 2022, 120 ‘Kanzi’ apples that had been stored under CA conditions (4 °C, 1 kPa O₂, 1.5 kPa CO₂) for 8–9 months were obtained from FruitMasters, The Netherlands. The fruit was grown in orchards surrounding Geldermalsen, the Netherlands, and harvested at physiological maturity in 2021.

2.2. Workflow overview

Image registration methods use shared information between images to find a transformation function T that describes the relative alignment between the images. In this paper, the image registration problem is finding a transformation function from the slice photographs to the CT scan in order to sample the corresponding CT slice for each photograph. CT scans and photographs have relatively little shared information (Fig. 1), but the outer shape of the apple is visible on both CT and photographic images. Therefore, the image registration method was designed to use the outer shape to find the transformation function. However, by only using the outer shape, relatively little information is available on each slice photograph. To overcome this, the image registration method optimizes the alignment of all slices jointly. Apples can be sliced into parallel slices, causing all slices to be rotated equally. Therefore, only one set of rotation parameters is sufficient to describe the rotation of all slices relative to the CT scan. The image registration method optimizes these parameters so that the outer shape matches well with the CT scan on all slices. This means that the combined information

from all slices is used to find the rotation parameters and other parameters that are shared between all slices.

The complete workflow to obtain registered image pairs of apples consists of four steps: Data acquisition, image segmentation, image registration, and validation (Fig. 2). In the data acquisition step, the fruit is first CT scanned and then sliced into parallel slices and photographed. In the image segmentation step, the outer shape of the apple is determined in both the CT scan and the slice photographs. In the image registration step, the transformation function from the slice photographs to the CT scan is optimized based on the outer shape of the apple. In the validation step, the registration error is measured from annotations of the apple's core. All steps are implemented as code in the Python programming language.

2.3. Step 1: Data acquisition

The goal of this step was to acquire CT scans and slice photographs of the apple. We first CT scanned the apples. After that, we used a setup to slice and photograph the apple. This setup ensured that the slices were parallel and equally spaced and that the apple did not rotate during slicing.

2.3.1. X-ray computed tomography(CT) acquisition and reconstruction

The apples were scanned at the Flex-ray laboratory (Coban et al., 2020) using a custom scanner developed by Tescan-XRE, Gent, Belgium. A cone beam geometry with a circular trajectory was used to acquire 1440 projection images at an exposure time of 100 ms, a tube peak voltage of 90 kV, and a current of 550 μ A. Volumes were reconstructed with the FDK algorithm (Feldkamp et al., 1984). To limit dataset size and scan time, 2 times detector pixel binning was used, resulting in a voxel size of 129.3 μ m. Beam hardening correction was used from the Flexbox package (Kostenko et al., 2020). All apples were scanned with the stem side on top. Moreover, a line was drawn on all apples from the stem to the calyx. The apples were put in the CT scanner so that the line was facing the X-ray source.

2.3.2. Slicing and photograph acquisition

One day after CT scanning, the apples were sliced using a modified meat-slicing machine (CaterChef, EMGA, Mijdrecht, The Netherlands) (Fig. 3). The sliding surface of the meat-slicing machine was replaced by a transparent acrylic sheet. A camera was placed behind the sliding surface, facing the apple. While in the machine, each apple was kept in place by a suction cup so it could not rotate during the slicing. All apples were sliced from the stem end to the calyx end with a slice thickness of approximately 4 mm. Every time before slicing, a picture was taken of the remaining part of the apple through the transparent sliding surface.




Fig. 3. Setup used for slicing and photographing the apples.

The jpeg files from the camera were used without color corrections. To ensure that the slice photographs were already roughly aligned with the CT scans, the apples were oriented in the slicing machine so that the line drawn earlier was on top. After each apple was sliced, the slices were visually inspected for brown tissue. At a later time, the type of browning was also noted for each fruit based on the photographs.

Sometimes, parts of the apple would break off while slicing the last few slices. The corresponding photographs were excluded. For three apples, the cutting damage already started halfway through the apple, and these apples were therefore excluded completely. Two apples were excluded because they had severe decay. The first eight apples were excluded because the background was too bright due to direct sunlight. In the remainder of the photographs the direct sunlight was blocked using a piece of carton. All included apples were photographed on the same day in the same location, resulting in similar lighting conditions. After the exclusions, 107 apples remained.

2.4. Step 2: Image segmentation

The goal of this step was to acquire a segmentation mask for each photograph and the CT scan. A segmentation mask is a black-and-white image of the same size as the original image, where a value of 1 indicates this part of the image shows the inside of the apple, and a value of 0 indicates this part of the image shows the outside of the apple. The segmentation masks were used in the image registration step to represent the outer shape of the apple.




Fig. 2. Overview of the workflow. For every step, one CT slice and photograph are shown to illustrate the results from that step. Note: In the actual workflow the CT scan is 3D data and there are multiple slice photographs.

2.4.1. CT segmentation

Thresholding was used to segment the shape of the apple from the CT scan. The threshold was determined by applying Otsu's method (Otsu, 1979) to all CT scans individually and taking the average of the results. This average Otsu threshold was applied to all CT scans. Morphological closing, selection of the largest connected component, and hole filling were used as post-processing.

2.4.2. Deep learning based photo segmentation

Segmenting the apple photographs could not be achieved with a single threshold. Only the outer shape of the apple should be segmented, which in the pictures is the part where the apple was sliced. However, the colors of the inside of the apple were very similar to the background or the outside of the apple in some photographs (Fig. 4). Deep learning has shown excellent results on various segmentation tasks (Wang et al., 2022), so it was used for segmentation.

The network architecture was a Mixed-scale dense convolutional neural network (MSD-net) (Pelt and Sethian, 2018) implemented in PyTorch using the msd_pytorch library (Hendriksen, 2019) using 200 layers. The normalized RGB values were used as the three input channels of the network. The network had one output channel to which a sigmoid activation function was applied. The binary cross entropy was used as a loss function, and the ADAM optimizer (Kingma and Ba, 2014) was used to minimize the loss function. For training the neural network, 175 photos were manually segmented. Data augmentation was applied to the training data using the ColorJitter, ShiftScaleRotate and GaussNoise functions of the Albumentations package (Buslaev et al., 2020). Cross-validation was used to train five neural network instances using 80% of the data as the training set and 20% of the data as the validation set. The validation set was used for early stopping to reduce overfitting without compromising model accuracy. The outputs of the five neural network instances were converted into a segmentation mask by taking the average and applying a threshold of 0.5. Selection of the largest connected component and hole filling were used as post-processing.

2.5. Step 3: Image registration

The goal of this step was to find a transformation function T that describes the relative alignment of the slice photographs to the CT scan. A task-specific transformation model was developed with shared parameters between all slices because the slices are parallel and equally spaced. The parameters of the transformation model were optimized based on the outer shape of the apple. This optimization was done in two steps: an initialization step and a full optimization step. Using the optimized transformation function, CT slices were sampled from the CT volume corresponding to the photographs.




Fig. 4. (Left) Slice photograph showing both the inside and the outside of the apple. (Right) Segmentation mask from the training set indicating the inside of the apple.

2.5.1. Transformation model

Transformation functions can include different transformations such as rotating, shifting, scaling, and shearing. The transformation model defines which transformations are included and how they are parameterized. A transformation model was developed specifically for this image registration task. It describes how the parallel 2D slice photographs are aligned relative to each other and to the 3D CT scan. Most transformation parameters are shared between all photographs. These are: a scalar “scaling” to match the photograph pixel size to the CT voxel size; a scalar “spacing” specifying the z-axis distance between adjacent photographs; a scalar “offset_z” specifying the z-axis offset of the whole apple; and three scalars “rotation_x”, “rotation_y”, “rotation_z” describing the rotation as Euler angles. The non-shared parameters are two scalars “offset_{x,i}” and “offset_{y,i}” for each photograph specifying the offset along the x and y-axes. The transformation model is illustrated in Fig. 5. The order of applying the transformations is as follows for a given coordinate vector c in the i -th slice photograph:

$$T(c, \theta, i) = R_{\theta} \left(\begin{bmatrix} \text{offset}_{x,i} \\ \text{offset}_{y,i} \\ \text{offset}_z + \text{spacing} \cdot i \end{bmatrix} + \text{scaling} \cdot c \right). \quad (1)$$

All transformation parameters are contained in vector $\theta = \{ \text{rotation}_x, \text{rotation}_y, \text{rotation}_z, \text{scaling}, \text{spacing}, \text{offset}_{x,1}, \text{offset}_{y,1}, \text{offset}_{x,2}, \text{offset}_{y,2}, \dots \}$. R_{θ} is the rotation matrix based on the rotation_x, rotation_y and rotation_z parameters in θ .

2.5.2. Initialization of the transformation parameters

Before performing the full optimization of the transformation parameters, the parameters are initialized. Initializing the parameters avoids converging to the wrong local optimum and saves some computation time.

The rotation parameters were initialized to zero. Because of the way the position of the stem and the drawn line were used during acquisition (Sections 2.3.1 & 2.3.2) the rotation was already roughly aligned.

The “scaling”, “offset_z”, and “spacing” parameters were initialized by matching the area of the segmentations in horizontal slices between the CT and the photographs. We call this the area profile. The area profile of the photographs was defined as the set of points p_i with one point for each slice photograph, where i stands for the index of the photograph and $A_{\text{photo}}(i)$ for the segmentation area of the photograph at index i :

$$p_i = \begin{bmatrix} \sqrt{A_{\text{photo}}(i) \cdot \text{scaling}} \\ \text{offset}_z + \text{spacing} \cdot i \end{bmatrix}. \quad (2)$$




Fig. 5. Illustration of the parameters of the parallel slice transformation model. Only “offset_{x,i}” and “offset_{y,i}” are unique for each slice. The other parameters are shared between all slices of the same apple.

The area profile of the CT scan was defined as a set of points s_j with one point for each horizontal slice of the CT scan, where j stands for the index of the horizontal CT slice and $A_{CT}(j)$ for the segmentation area of slice j .

$$s_j = \begin{bmatrix} \sqrt{A_{CT}(j)} \\ j \end{bmatrix}. \quad (3)$$

The mean over the squared distances from every point in the photo area profile to the closest point in the CT area profile was used as a cost function. This cost function was optimized using gradient descent with momentum (Rumelhart et al., 1985; Goh, 2017), and the gradient was automatically derived using PyTorch (Paszke et al., 2017). This approach is similar to the SGD-ICP method (Maken et al., 2019). The learning rates were manually tuned for each parameter to be as large as possible without causing exploding gradients. A momentum of 0.6 was used on all parameters. A fixed number of 10000 iterations was used.

The “offset_{x,i}” and “offset_{y,i}” parameters were initialized so that the center of mass of each slice photograph segmentation was in the same position as the center of mass of the CT segmentation slice that was sampled based on the current “scaling”, “offset_z” and “spacing” estimates.

2.5.3. Joint optimization of the transformation parameters

After initializing the transformation parameters, a more precise optimization was performed to get to the final image registration result. This step uses the mean square error (MSE) between the segmentations of the CT scans and slice photographs as a cost function. To evaluate the cost function, the positions of all pixels in the photograph segmentation masks were transformed into CT space using the current transformation parameters, and lookups were done in the CT segmentation mask using trilinear interpolation.

The cost function was optimized using gradient descent with momentum (Rumelhart et al., 1985; Goh, 2017). The gradient was automatically derived using PyTorch (Paszke et al., 2017), and all calculations were performed on the GPU. The learning rates were manually tuned for each parameter to be as large as possible without causing exploding gradients, and a momentum of 0.75 was used on all parameters. The stopping criterion was that over the last 1000 iterations, the minimum and maximum values of the cost function were less than 0.00001 apart. After the optimization finished, for each photograph a corresponding CT slice was sampled from the CT volume using trilinear interpolation.

2.6. Step 4: Validation

The goal of this step was to measure the registration error. The segmentation error was also measured because the registration was optimized based on segmentation masks.

2.6.1. In plane core endpoint distance (IPCED)

The core of the apple is visible on both the CT scans and the slice photographs and was not used for optimizing the registration parameters. Therefore, the core could be used to measure the registration error. The outer corner points of the apple core were selected manually on the photographs. Later, the closest corresponding points on the registered CT slice were manually selected (Fig. 6). The in-plane core endpoint distance (IPCED) was defined as the average distance between the annotations in the slice photograph and the registered CT slice.

To minimize the effect of the out-of-plane registration error, the points were only selected on one slice photograph close to the center of the apple core because, at its center, the edges of the core are roughly perpendicular to the slicing plane. Moreover, very thin core structures or those with an unclear endpoint were not annotated. To avoid bias, the photographs were selected and the annotations on the photographs were done without looking at the CT scans. The CT point annotations were




Fig. 6. Example of the annotations of the core endpoints used to measure the registration error between a slice photograph (left) and the registered CT slice (right). When the endpoint of a part of the core was unclear (circular region) or when the core structure was very thin (diamond region) the endpoints were not annotated.

placed with the annotated photograph and the registered CT scan side-by-side in a view similar to Fig. 6.

2.6.2. Segmentation error

The segmentation error was evaluated to understand how much of the registration error was caused by the quality of the segmentation of the photographs. For this goal, a test set was created of 20 manually segmented slice photographs of 20 different apples that had not been used for training the segmentation neural network. The five neural network instances that were trained using different cross-validation splits were applied to these photographs. The postprocessing was applied to the mean of the five networks like in Section 2.4.2 and to each network output separately. True positives (TP) represent pixels correctly segmented as being inside the apple; true negatives (TN) represent pixels correctly segmented as being outside the apple; false positives (FP) represent pixels incorrectly segmented as inside the apple; and false negatives (FN) represent pixels incorrectly segmented as being outside the apple. The following error metrics were calculated: accuracy $(TP + TN)/(TP + TN + FP + FN)$, precision $TP/(TP + FN)$, recall $TP/(TP + FP)$, and distance from the edge pixels in the neural network segmentation to the closest edge pixels in the manual segmentation.

2.7. Validation of registration method

The registration method (Section 2.5) has not been used before; thus, two additional experiments were performed to measure the added value of this new method.

2.7.1. Number of slices experiment

The slicing method (Section 2.3.2) was supposed to result in parallel slices, but slight deviations may occur that could accumulate when many slices are used. However, using many slices may also improve results by including more information in the optimization. The goal of this experiment was to explore how the number of slices used in the joint registration method related to the registration error. To vary the number of slices, not all available slices were used.

For every apple, slices at the top and the bottom of the apple were excluded to make subsets with a specific number of slices. To be able to measure the IPCED, the annotated center slice was always included. Additionally, the three slices above and two slices below the annotated slice were included one by one, creating a new subset with every included slice, resulting in six subsets for each apple. For 30 Apples, these slices were not all available, so those fruit were excluded from this experiment. On the subsets with fewer slices on some apples, the registration error was too large to do the annotations necessary for calculating the IPCED. To allow comparison between subset sizes these fruit were excluded from calculating the average IPCED on all subset sizes.

2.7.2. Separate slice registration intersection experiment

Because the previous experiment (Section 2.7.1) used the IPCED, it only calculated the error at one central slice. This experiment was performed to get insight into the registration robustness over the whole apple. Each apple's photographs were acquired from slices from top to bottom. Therefore, it is physically impossible that slices intersected each other, or that the position of later photographs was before the position of earlier photographs. Our joint registration method enforces that the slices are parallel and in the correct order, but in this experiment, each slice was registered separately to test what would occur if these properties were not enforced.

The same method was used as for joint registration, with a minor modification. In the transformation model, the offset_z and spacing parameters would offer the same degree of freedom when just one slice is used, so the spacing parameter was kept constant. The initialization method from Section 2.5.2, and the full joint optimization method from Section 2.5.3 were used to initialize the transformation model parameters.

An intersection test was developed to test if the slices were parallel and in the correct order. First, the coordinates of every pixel inside each photo slice segmentation mask were transformed into CT space. A pixel belonging to slice *i* should not be inside either of the two convex hulls formed by slices with indices larger and smaller than *i*. If any pixel in a slice was inside either of these convex hulls, that entire slice was marked as intersecting. A distinction was made between apples where three or fewer adjacent slices were intersecting and where more slices were intersecting. Fig. 7 illustrates the intersection test.

3. Results

3.1. Registration and segmentation errors

The full workflow was applied to obtain the slice photographs and registered CT slices in the 'Kanzi' data set. The average IPCED for all 107 apples was 1.47 ± 0.40 mm.

The calculated segmentation metrics are displayed in Table 1. The *combined* row represents the segmentations that were used for registration. The average segmentation edge distance was 0.195 ± 0.104 mm.

3.2. Validation of registration method

The results of the number of slices experiment are displayed in Table 2. The results show that the average IPCED decreased with every increase in the subset size. Moreover, only when a subset size of three or lower was used there were cases where the registration error was too large to do the annotations necessary for calculating the IPCED.

The results of the intersection experiment are displayed in Table 3. The results show that intersections occur when registering the slices separately. Using the full joint registration method as an initialization instead of the normal profile initialization method decreased the number of intersections, but intersections remained in 32 of the apples.

Table 1

Segmentation quality metrics of the neural network instances trained on different training and validation data splits.

	Accuracy (%)	Precision (%)	Recall (%)	Edge distance (mm)
instance 0	99.37 ± 0.30	99.05 ± 1.07	99.41 ± 0.63	0.234 ± 0.139
instance 1	99.26 ± 0.50	99.77 ± 0.27	98.45 ± 1.57	0.259 ± 0.180
instance 2	99.37 ± 0.26	99.10 ± 0.88	99.39 ± 0.65	0.241 ± 0.133
instance 3	99.33 ± 0.44	99.61 ± 0.36	98.78 ± 1.41	0.247 ± 0.206
instance 4	99.39 ± 0.33	99.52 ± 0.47	99.04 ± 1.05	0.222 ± 0.134
combined	99.43 ± 0.30	99.54 ± 0.42	99.12 ± 0.94	0.195 ± 0.104

Table 2

Registration results when using different numbers of slices in the joint registration.

No. of slices in subset	No. of fruit where IPCED could not be calculated	IPCED (mm)
1	3	1.97 ± 0.89
2	7	1.83 ± 0.80
3	3	1.83 ± 0.82
4	0	1.75 ± 0.81
5	0	1.74 ± 0.83
6	0	1.71 ± 0.86
All	0	1.52 ± 0.39

Table 3

Results of the separate slice registration intersection experiment.

Initialization	0 Intersections	≤ 3 Adjacent intersections	More intersections
Profile (Sec. 2.5.2)	42 (39.2%)	42 (39.2%)	23 (21.5%)
Joint (Sec. 2.5.3)	75 (70.0%)	24 (22.4%)	8 (7.5%)

3.3. Disorder evaluation

Of the 107 'Kanzi' apples in the dataset, 45 were healthy after storage, and 62 contained brown regions of flesh with varying levels of severity. Senescent breakdown, radial flesh, diffuse flesh, and core browning were all identified. An example of each browning type is presented in Fig. 8. We found that dark regions in CT scans corresponded to regions of brown flesh in 'Kanzi' fruit. High-intensity regions occurred on the border of dark regions in CT scans for fruit with core browning (Fig. 8C). Fruit with decay displayed areas of high and low-intensity regions in CT scans (Fig. 8F). The wet, decayed flesh had a high voxel intensity, and the remaining flesh resembled senescence browning (Fig. 8D).

During the browning classification, several fruit marked as healthy




Fig. 7. Illustration of the intersection test. The red lines indicate intersecting slices.




Fig. 8. Examples images of a healthy apple and apples with different types of disorders. A: healthy, B: radial browning, C: core browning, D: senescence browning, E: diffuse flesh browning, F: decay. The combined view shows both images in the same plot and can be used to see the distance between edges. The combined view is green when the photo is brighter, purple when the CT slice is brighter, and a shade of gray when both images have similar intensities.

during slicing displayed slightly darker areas in the photo images. However, it was difficult to determine with the human eye whether these areas were brown tissue or shadows. In six apples a slightly darker area in the photo images coincided with a dark region in the CT scan (Fig. 9). Based on the CT scan and photo together, we expect that these apples had mild browning.

3.4. Data and code

The CT scans, slice photographs, and browning status are publicly available on Zenodo (Schut et al., 2023a). All data used in the results section, including registration and segmentation results, are publicly




Fig. 9. Examples of 'Kanzi' apples that were classified as healthy after cutting, but contained regions of flesh that appeared brown in CT images. The arrows indicate regions of suspected brown tissue, which were confirmed using the registered CT slice. The combined view shows both images in the same plot and can be used to see the distance between edges. The combined view is green when the photo is brighter, purple when the CT slice is brighter, and a shade of gray when both images have similar intensities.

available as a separate dataset (Schut et al., 2023b). This dataset also includes side-by-side views of all photographs and their corresponding CT slices similar to Fig. 8, saved as .png files so they can be viewed in any image viewer. The Python code is available on Github (Schut, 2023).

4. Discussion

4.1. Physiological disorders

Most physiological disorders have affected regions much larger than the average IPCED of 1.47 mm. Therefore, the provided photo-CT image pairs should be suitable for disorder evaluation. The photographs can be used as ground truth to understand the appearance of disorders in CT. Registered image pairs can also provide complementary information, which could be explored in future work. Combining information from registered CT, MRI, or PET scans has been used in medical imaging to diagnose diseases such as cancer (Blodgett et al., 2007; Du et al., 2016).

Our finding that dark regions in CT scans corresponded to regions of brown flesh is consistent with previous reports (Herremans et al., 2013; van Dael et al., 2019). Additionally, high-intensity regions surrounding the dark regions in fruit with core browning, were also found by Lamertyn et al. (2003) study on pears. The dark regions in affected fruit are due to the structural collapse of cells and the moment of water away from affected flesh, reducing voxel intensity (Herremans et al., 2013; Van De Looverbosch et al., 2020). Herremans et al. (2013) suggest that the available water will diffuse towards the fruit peel, where it is lost to the environment. Water-rich tissue has a high voxel intensity (Diels et al., 2017); thus, the high-intensity regions surrounding the dark regions in fruit with core browning are likely caused by the movement of water away from the affected flesh and towards the fruit's surface. On six apples, browning was detected on the CT scan but not during visual inspection, which suggests that CT can potentially be used for earlier or more reliable detection of browning than visual inspection. Early non-destructive detection of browning would be commercially valuable. Wood et al. (2024) demonstrated that the level of browning in an apple can be determined from the percentage of dark spots in a single CT slice.

The high-intensity regions of decayed flesh in our CT scans are similar to those obtained from fruit with watercore (Herremans et al., 2013). Watercore is a physiological disorder associated with the flooding of air-filled pores with aqueous fluid, increasing voxel intensity in CT scans (Herremans et al., 2013). Herremans et al. (2013) study

successfully identified the water-soaked regions typical for fruit affected with watercore using CT with up to 89% accuracy. As watercore and decay appear similar under CT, fruit with decayed flesh may be successfully identified using CT and image thresholding. However, distinguishing between regions of flesh suffering from decay and watercore might be challenging due to their similar appearance. Although, the separation of decay and watercore may only be an issue in growing regions where apples with watercore are a desired quality trait and are sold at a premium.

4.2. Data acquisition and image segmentation

After using the workflow on the ‘Kanzi’ dataset, several areas of improvement have been identified for acquiring sets of slice photographs to make image registration to CT or MRI scans simpler and more accurate.

Firstly, camera calibration would improve the usability of the photographs. An approach that involves photographing a flat checkerboard test image in multiple orientations (Zhang, 2000) is available in OpenCV and Matlab. After calibration, lens distortions are compensated for. By also photographing the checkerboard test image when placed against the clear sliding surface of the slicing setup, the external calibration parameters from that image could be used to compensate for the perspective projection in the photographs. Moreover, the physical size of the apple slices could be derived directly from the images instead of being fitted in the registration step, making the registration simpler and more accurate.

Secondly, the photo segmentation could be simplified if the contrast could be improved between the inside of the apple and both the peel and the background. If the contrast can be increased sufficiently, classical methods for segmentation might be used instead of a neural network. This would free up the time required for manually labeling images and save the costs for hardware and electricity of neural network training. Moreover, more accurate segmentation masks may be achieved. It may even be possible to remove the segmentation step completely by using a multi-modality cost function such as mutual information during the image registration (Maes et al., 1997).

The contrast with the background can easily be improved by placing an evenly colored surface of a contrasting color (e.g. blue or black) behind the machine. The contrast between the peel and the inside of the apple could be improved by painting the apple in a contrasting color before cutting. Another option would be to cut the apple through the middle at the widest point and then put the apple in the machine in two parts with the cutting plane facing the camera. This way the apple would get narrower further away from the camera so that the peel would not be visible in the photographs. The middle slice would be photographed twice, and by registration between these two images, the rotation difference between the sets of photographs of the two halves of the apple could be compensated. The contrast of the whole image could be improved by fine-tuning the lighting so that it is uniform and uses the whole dynamic range of the sensor. Improving the lighting would also benefit visual inspection of the photographs.

4.3. Image registration

The idea of segmentation-based registration has been applied before by Museyko et al. (2010), which showed that the results for image registration between two CT scans of the human femur or lumbar spine were improved by segmenting the bone structures before doing image registration. Segmentations were also used to overcome the differences in appearance while registering a 3D ultrasound scan and a CT scan of the liver (Schut, 2018). Pears have a more irregular shape than apples, so performing segmentation-based registration on pears may be more accurate or require less data.

Task-specific transformation models have been used in medical image registration of the heart (Heyde et al., 2013) and liver (Wein et al.,

2008) to allow for specific motions with less risk of overfitting than using a more general transformation model. The most common approach for image registration is to use a toolkit such as Elastix (Klein et al., 2009) or ITK (Avants et al., 2014), where users can combine high-quality implementations of commonly used transformation models, cost functions, and optimizers. Many optimizers use the gradient, so gradient computation code has to be provided for each cost function and transformation model. While toolkits make it easy to prototype combinations of existing methods, the interactions with existing C++ code and the necessity of manually writing gradient code introduce a high barrier to adding custom transformation models or cost functions. The use of PyTorch for automatic differentiation (Paszke et al., 2017) in this work lowered this barrier considerably. AirLab (Sandkühler et al., 2018) is a recently developed image registration toolkit that uses PyTorch automatic differentiation internally. In future work on custom transformation models, the existing AirLab implementations of optimizers and cost functions could be used.

Combining information from multiple slices has been done in medical histology image registration (Pichat et al., 2018; Ferrante and Paragios, 2017). In medical histology, a tissue sample is sliced into thin sections, which are put on glass slides, stained, and photographed under a microscope. Like apples in this paper, histological samples are sometimes scanned in a CT or MRI scanner before being sliced. Registration between histological slides and CT or MRI scans is therefore similar to the apple registration problem described in this paper. However, there are also differences. Large full-object deformations may be present in registration between histology and CT or MRI, so a deformable transformation model is used. To fit the many parameters of a deformable model, combining information from multiple histological photographs can be beneficial (Ceritoglu et al., 2010). The processes of slicing, mounting the slice on a glass slide, and staining can introduce artifacts and deformations in the slicing plane (Pichat et al., 2018). This makes it challenging to find the accurate relative positions of adjacent slices, which is called 3D reconstruction. Joint optimization of the 3D reconstruction and image registration problems can compensate for these effects. Joint optimization has been done by alternating between registering the slices individually to the CT or MRI scan to improve the 3D reconstruction and registering the whole reconstructed volume to the CT or MRI scan to improve the whole object registration (Adler et al., 2014; Goubran et al., 2013; Malandain et al., 2004; Yang et al., 2012) and Section 5.1.4 in Pichat et al. (2018). The method in this paper differs from those earlier methods in that the 3D reconstruction and registration problems are optimized within the same optimization problem.

4.4. Validation

The IPCED metric defined in Section 2.6.1 is easy to calculate for any apple or pear with only a few manual annotations. However, it is limited in that it only considers the in-plane error and can only be calculated at the core of the apple. For medical histology applications, needle tracks have been used as a marker for evaluation (Pichat et al., 2018). Inserting and removing multiple needles before data acquisition introduces straight holes in the apple. These are likely easy to segment in the CT scan and easy to detect in the photographs. The registration error can be calculated as the distance from the holes in the photograph to the nearest hole in the CT scan. Moreover, these features could be used directly for image registration. However, this metric would require damaging the apples and would introduce an extra sample preparation step and an extra detection step to the workflow.

5. Conclusions

In this paper, we presented a workflow for creating datasets of registered image pairs of slice photographs and CT slices. Part of the workflow is a joint 2D-3D image registration method (Section 2.5), which uses a task-specific transformation model to register all slices of

one apple within a single optimization problem. It was shown that this joint approach is both more precise and more robust to large registration errors compared to image registration of separate slices (Section 3.2). Using this workflow, a dataset was created of 1347 slice photographs acquired from 107 ‘Kanzi’ apples, which has been published alongside this paper. The image registration error of this dataset according to the in-plane core endpoint distance (IPCED) metric was 1.47 ± 0.40 mm (Section 3.1). The appearance of disorders in this dataset mostly matched the descriptions from earlier works on CT scanning of disordered apples (Section 4.1). However, six apples that looked healthy directly after slicing showed dark regions on the CT scans associated with browning (Section 3.3), indicating that CT may have the potential for detecting the early stages of browning.

Funding

This work was funded by the Dutch Research Council (NWO) through the UTOPIA project (ENWSS.2018.003).

CRediT authorship contribution statement

Batenburg Kees Joost: Conceptualization, Funding acquisition, Project administration, Supervision, Writing – original draft, Writing – review & editing. **van Liere Robert:** Conceptualization, Project administration, Supervision, Writing – review & editing. **van Leeuwen Tristan:** Conceptualization, Project administration, Supervision. **Trull Anna Katharina:** Conceptualization, Data curation, Investigation, Methodology, Project administration, Resources. **Schouten Rob:** Conceptualization, Writing – review & editing. **Schut Dirk Elias:** Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Wood Rachael Maree:** Conceptualization, Data curation, Investigation, Validation, Visualization, Writing – original draft, Writing – review & editing.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The data is publicly available as two datasets on Zenodo (links in article). The code is available on Github (link in article).

Acknowledgements

The authors would like to acknowledge Michiel Couvée for technical assistance with ‘Kanzi’ fruit slicing. The authors also acknowledge TESCAN-XRE NV for their collaboration and support of the Flex-ray laboratory.

References

- Adler, D.H., Pluta, J., Kadivar, S., Craigie, C., Gee, J.C., Avants, B.B., Yushkevich, P.A., 2014. Histology-derived volumetric annotation of the human hippocampal subfields in postmortem MRI. *Neuroimage* 84, 505–523. <https://doi.org/10.1016/j.neuroimage.2013.08.067>.
- Andriashen, V., van Liere, R., van Leeuwen, T., Batenburg, K.J., 2023. Ct-based data generation for foreign object detection on a single X-ray projection. *Sci. Rep.* 13 (1), 1881. <https://doi.org/10.1038/s41598-023-29079-w>.
- Argenta, L.C., de Freitas, S.T., Mattheis, J.P., Vieira, M.J., Ogoshi, C., 2021. Characterization and quantification of postharvest losses of apple fruit stored under commercial conditions. *HortScience* 56 (5), 608–616. <https://doi.org/10.21273/HORTSCI5771-21>.
- Avants, B.B., Tustison, N.J., Stauffer, M., Song, G., Wu, B., Gee, J.C., 2014. The Insight Toolkit image registration framework. *Front. Neuroinform.* 8, 44. <https://doi.org/10.3389/fninf.2014.00044>.

- Azadbakht, M., Vahedi Torshizi, M., Mahmoodi, M.J., 2019. The relation of pear volume and it's bruised volume by CT scan imaging. *J. Food Meas. Charact.* 13 (2), 1089–1099. <https://doi.org/10.1007/s11694-018-00024-0>.
- Blodgett, T.M., Meltzer, C.C., Townsend, D.W., 2007. PET/CT: form and function. In: *Radiology*, 242, pp. 360–385. <https://doi.org/10.1148/radiol.2422051113>.
- Buslaev, A., Iglovikov, V.I., Khvedchenya, E., Parinov, A., Druzhinin, M., Kalinin, A.A., 2020. Alumentations: fast and flexible image augmentations. *Information* 11 (2). <https://doi.org/10.3390/info11020125>.
- Ceritoglu, C., Wang, L., Selemón, L.D., Csernansky, J.G., Miller, M.I., Ratnanather, J.T., 2010. Large deformation diffeomorphic metric mapping registration of reconstructed 3D histological section images and in vivo MR images. *Front. Hum. Neurosci.* 43. <https://doi.org/10.3389/fnhum.2010.00043>.
- Chigwaya, K., Karuppanapandian, T., Schoeman, L., Viljoen, D.W., Crouch, I.J., Nugraha, B., Verboven, P., Nicolai, B.M., Crouch, E.M., 2021. X-ray CT and porosity mapping to determine the effect of ‘Fuji’ apple morphological and microstructural properties on the incidence of CO₂ induced internal browning. *Postharvest Biol. Technol.* 174, 111464. <https://doi.org/10.1016/j.postharvbio.2021.111464>.
- Chigwaya, K., duPlessis, A., Viljoen, D.W., Crouch, I.J., Crouch, E.M., 2021. Use of X-ray computed tomography and 3D image analysis to characterize internal browning in ‘Fuji’ apples after exposure to CO₂ stress. *Sci. Hortic.* 277, 109840. <https://doi.org/10.1016/j.scienta.2020.109840>.
- Coban, S.B., Lucka, F., Palenstijn, W.J., Van Loo, D., Batenburg, K.J., 2020. Explorative imaging and its implementation at the Flex-ray Laboratory. *J. Imaging* 6 (4), 18. <https://doi.org/10.3390/jimaging6040018>.
- De Schryver, T., Dhaene, J., Dierick, M., Boone, M.N., Janssens, E., Sijbers, J., van Dael, M., Verboven, P., Nicolai, B., Van Hoorebeke, L., 2016. In-line NDT with X-Ray CT combining sample rotation and translation. *NDT E Int.* 84, 89–98. <https://doi.org/10.1016/j.ndteint.2016.09.001>.
- Diels, E., van Dael, M., Keresztes, J., Vanmaercke, S., Verboven, P., Nicolai, B., Saeys, W., Ramon, H., Smeets, B., 2017. Assessment of bruise volumes in apples using X-ray computed tomography. *Postharvest Biol. Technol.* 128, 24–32. <https://doi.org/10.1016/j.postharvbio.2017.01.013>.
- Du, J., Li, W., Lu, K., Xiao, B., 2016. An overview of multi-modal medical image fusion. *Neurocomputing* 215, 3–20. <https://doi.org/10.1016/j.neucom.2015.07.160>.
- Feldkamp, L.A., Davis, L.C., Kress, J.W., 1984. Practical cone-beam algorithm. *JOSA A* 1 (6), 612–619. <https://doi.org/10.1364/JOSAA.1.000612>.
- Ferrante, E., Paragios, N., 2017. Slice-to-volume medical image registration: a survey. *Med. Image Anal.* 39, 101–123. <https://doi.org/10.1016/j.media.2017.04.010>.
- Goh, G., 2017. Why momentum really works. *Distill.* <https://doi.org/10.23915/distill.00006>.
- Goubran, M., Crukley, C., De Ribaupierre, S., Peters, T.M., Khan, A.R., 2013. Image registration of ex-vivo MRI to sparsely sectioned histology of hippocampal and neocortical temporal lobe specimens. *Neuroimage* 83, 770–781. <https://doi.org/10.1016/j.neuroimage.2013.07.053>.
- Hatoum, D., Buts, K., Hertog, M., Geeraerd, A., Schenk, A., Vercammen, J., Nicolai, B., 2014. Effects of pre-and postharvest factors on browning in Braeburn. *Hortic. Sci.* 41 (1), 19–26. <https://doi.org/10.17221/180/2013-HORTSCI>.
- Hendriksen, A.A., 2019. ahendriksen/msd_pytorch: v0.7.2. Version v0.7.2.10.5281/zenodo.3560114.
- Herremans, E., Verboven, P., Bongaers, E., Estrade, P., Verlinden, B.E., Wevers, M., Hertog, M.L., Nicolai, B.M., 2013. Characterisation of ‘Braeburn’ browning disorder by means of X-ray micro-CT. *Postharvest Biol. Technol.* 75, 114–124. <https://doi.org/10.1016/j.postharvbio.2012.08.008>.
- Herremans, E., Melado-Herreros, A., Defraeye, T., Verlinden, B., Hertog, M., Verboven, P., Val, J., Fernández-Valle, M.E., Bongaers, E., Estrade, P., Wevers, M., Barreiro, P., Nicolai, B.M., 2014. Comparison of X-ray CT and MRI of watercore disorder of different apple cultivars. *Postharvest Biol. Technol.* 87, 42–50. <https://doi.org/10.1016/j.postharvbio.2013.08.008>.
- Heyde, B., Barbosa, D., Claus, P., Maes, F., and D’hooge, J., 2013. Three-dimensional cardiac motion estimation based on non-rigid image registration using a novel transformation model adapted to the heart. In: *Statistical Atlases and Computational Models of the Heart. Imaging and Modelling Challenges: Third International Workshop, STACOM 2012, Held in Conjunction with MICCAI 2012, Nice, France, October 5, 2012, Revised Selected Papers* 3.142–150. Springer. 10.1007/978-3-642-36961-2.17.
- Jarolmasjed, S., Espinoza, C.Z., Sankaran, S., Khot, L.R., 2016. Postharvest bitter pit detection and progression evaluation in ‘honeycrisp’ apples using computed tomography images. *Postharvest Biol. Technol.* 118, 35–42. <https://doi.org/10.1016/j.postharvbio.2016.03.014>.
- Kingma, D.P., Ba, J., 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*. 10.48550/arXiv.1412.6980.
- Klein, S., Staring, M., Murphy, K., Viergever, M.A., Pluim, J.P., 2009. Elastix: a toolbox for intensity-based medical image registration. *IEEE Trans. Med. Imaging* 29 (1), 196–205. <https://doi.org/10.1109/TMI.2009.2035616>.
- Kostenko, A., Palenstijn, W.J., Coban, S.B., Hendriksen, A.A., van Liere, R., Batenburg, K.J., 2020. Prototyping X-ray tomographic reconstruction pipelines with Flexbox. *SoftwareX* 11, 100364. <https://doi.org/10.1016/j.softx.2019.100364>.
- Kotwaliwale, N., Singh, K., Kalne, A., Jha, S.N., Seth, N., Kar, A., 2014. X-ray imaging methods for internal quality evaluation of agricultural produce. *J. Food Sci. Technol.* 51, 1–15. <https://doi.org/10.1007/s13197-011-0485-y>.
- Lammertyn, J., Dresselaers, T., Van Hecke, P., Jancsó, P., Wevers, M., Nicolai, B., 2003. MRI and X-ray CT study of spatial distribution of core breakdown in ‘Conference’ pears. *Magn. Reson. Imaging* 21 (7), 805–815. [https://doi.org/10.1016/S0730-725X\(03\)00105-X](https://doi.org/10.1016/S0730-725X(03)00105-X).

- Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., Suetens, P., 1997. Multimodality image registration by maximization of mutual information. *IEEE Trans. Med. Imaging* 16 (2), 187–198. <https://doi.org/10.1109/42.563664>.
- Maken, F.A., Ramos, F., and Ott, L., 2019. Speeding up iterative closest point using stochastic gradient descent. In: 2019 International Conference on Robotics and Automation (ICRA).IEEE, 6395–6401.10.1109/ICRA.2019.8794011.
- Malandain, G., Bardinet, E., Nelissen, K., Vanduffel, W., 2004. Fusion of autoradiographs with an MR volume using 2-D and 3-D linear transformations. *NeuroImage* 23 (1), 111–127. <https://doi.org/10.1016/j.neuroimage.2004.04.038>.
- Mink, G., 1973. The apple industry in Japan. *HortScience* 8 (2), 81–86. <https://doi.org/10.21273/HORTSCI.8.2.81>.
- Morton, E., Mann, K., Berman, A., Knaup, M., Kachelrieß, M., 2009. Ultrafast 3D reconstruction for X-ray real-time tomography (RTT). In: 2009 IEEE Nuclear Science Symposium Conference Record (NSS/MIC).IEEE, 4077–4080.10.1109/NSSMIC.2009.5402393.
- Museyko, O., Eisa, F., Hess, A., Schett, G., Kalender, W.A., Engelke, K., 2010. Binary segmentation masks can improve intrasubject registration accuracy of bone structures in CT images. *Ann. Biomed. Eng.* 38, 2464–2472. <https://doi.org/10.1007/s10439-010-9981-x>.
- Neuwald, D., Wood, R., Saquet, A., Pilla, R., Büchele, F., Kessler, M., Pansera-Espindola, B., Schmidt, D., and Wünsche, J., 2021. Impact of postharvest hot water and calcium treatments on fruit quality and decay incidence of ‘Topaz’ apple under various storage conditions.IX International Symposium on Mineral Nutrition of Fruit Crops 1333, 291–298.10.17660/ActaHortic.2022.1333.38.
- Otsu, N., 1979. A threshold selection method from gray-level histograms. *IEEE Trans. Syst., Man, Cybern.* 9 (1), 62–66.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A., 2017. Automatic differentiation in PyTorch. NIPS-W.
- Pelt, D.M., Sethian, J.A., 2018. A mixed-scale dense convolutional neural network for image analysis. *Proc. Natl. Acad. Sci.* 115 (2), 254–259. <https://doi.org/10.1073/pnas.1715832114>.
- Pichat, J., Iglesias, J.E., Yousry, T.A., Ourselin, S., Modat, M., 2018. A survey of methods for 3D histology reconstruction. *Med. Image Anal.* 46, 73–105. <https://doi.org/10.1016/j.media.2018.02.004>.
- Rumelhart, D.E., Hinton, G.E., and Williams, R.J., 1985. Learning internal representations by error propagation.Tech. rep. California Univ San Diego La Jolla Inst for Cognitive Science.
- Sandkühler, R., Jud, C., Andermatt, S., Cattin, P.C.2018. Airlab: autograd image registration laboratory.arXiv preprint arXiv:1806.09907. 10.48550/arXiv.1806.09907.
- Schut, D.E., 2018. Automatic Initialization for 3D Ultrasound CT Registration During Liver Tumor Ablations.MSc thesis. Delft University of Technology.
- Schut, D.E.2023. apple_photo_ct_workflow GitHub repository. (https://github.com/D1rk123/apple_photo_ct_workflow).
- Schut, D.E., Trull, A.K., Couvée, M., 2023b. Dataset with results of joint 2D to 3D image registration workflow for comparing multiple slice photographs and CT scans of apple fruit with internal disorders.Zenodo. (<https://zenodo.org/records/8275793>).
- Schut, D.E., Trull, A.K., Couvée, M., 2023a. Dataset of CT scans, slice photographs, and visual browning scores of 120 ‘Kanzi’ apples.Zenodo. 10.5281/zenodo.8167285.
- Schut, D.E., Batenburg, K.J., van Liere, R., van Leeuwen, T., 2022. TOP-CT: trajectory with overlapping projections X-ray computed tomography. *IEEE Trans. Comput. Imaging* 8, 598–608. <https://doi.org/10.1109/TCI.20223192125>.
- Si, Y., Sankaran, S., 2016. Computed tomography imaging-based bitter pit evaluation in apples. *Biosyst. Eng.* 151, 9–16. <https://doi.org/10.1016/j.biosystemseng.2016.08.008>.
- Sidhu, R.S., Bound, S.A., Swarts, N.D., 2023. Internal flesh browning in apple and its predisposing factors—a review. *Physiologia* 3 (2), 145–172. <https://doi.org/10.3390/physiologia3020012>.
- Upchurch, B., Throop, J., 1994. Effects of storage duration on detecting watercore in apples using machine vision. *Trans. ASAE* 37 (2), 483–486. <https://doi.org/10.13031/2013.28100>.
- Van Aarle, W., Palenstijn, W.J., Cant, J., Janssens, E., Bleichrodt, F., Dabrvolski, A., De Beenhouwer, J., Batenburg, K.J., Sijbers, J., 2016. Fast and flexible X-ray tomography using the ASTRA toolbox. *Opt. Express* 24 (22), 25129–25147. <https://doi.org/10.1364/OE.24.025129>.
- van Dael, M., Verboven, P., Dhaene, J., Van Hoorebeke, L., Sijbers, J., Nicolai, B., 2017. Multisensor X-ray inspection of internal defects in horticultural products. *Postharvest Biol. Technol.* 128, 33–43. <https://doi.org/10.1016/j.postharvbio.2017.02.002>.
- van Dael, M., Verboven, P., Zanella, A., Sijbers, J., Nicolai, B., 2019. Combination of shape and X-ray inspection for apple internal quality control: in silico analysis of the methodology based on X-ray computed tomography. *Postharvest Biol. Technol.* 148, 218–227. <https://doi.org/10.1016/j.postharvbio.2018.05.020>.
- Van De Looverbosch, T., RahmanBhuiyan, M.H., Verboven, P., Dierick, M., Van Loo, D., De Beenhouwer, J., Sijbers, J., Nicolai, B., 2020. Nondestructive internal quality inspection of pear fruit by X-ray CT using machine learning. *Food Control* 113, 107170. <https://doi.org/10.1016/j.foodcont.2020.107170>.
- Van De Looverbosch, T., Raeymaekers, E., Verboven, P., Sijbers, J., Nicolai, B., 2021. Non-destructive internal disorder detection of Conference pears by semantic segmentation of X-ray CT scans using deep learning. *Expert Syst. Appl.* 176, 114925. <https://doi.org/10.1016/j.eswa.2021.114925>.
- Wang, R., Lei, T., Cui, R., Zhang, B., Meng, H., Nandi, A.K., 2022. Medical image segmentation using deep learning: a survey. *IET Image Process.* 16 (5), 1243–1267. <https://doi.org/10.1049/ipr2.12419>.
- Wein, W., Brunke, S., Khamene, A., Callstrom, M.R., Navab, N., 2008. Automatic CT-ultrasound registration for diagnostic imaging and image-guided intervention. *Med. Image Anal.* 12 (5), 577–585. <https://doi.org/10.1016/j.media.2008.06.006>.
- Wood, R.M., Proske, M., de Freitas, S.T., Scheer, C., Vögele, R.T., Neuwald, D.A., 2022a. Seasonal variation in calcium and ascorbic acid content at harvest related to internal browning in ‘Braeburn’ apple during controlled atmosphere storage. *Sci. Hortic.* 297, 110943. <https://doi.org/10.1016/j.scienta.2022.110943>.
- Wood, R.M., Thewes, F.R., Reynaud, M., Kitemann, D., Sautter, C.K., Wünsche, J.N., Neuwald, D.A., 2022b. Apple fruit recovery from anoxia under controlled atmosphere storage. *Food Chem.* 371, 131152. <https://doi.org/10.1016/j.foodchem.2021.131152>.
- Wood, R.M., Schut, D.E., Trull, A.K., Marcelis, L.F., Schouten, R.E., 2024. Detecting internal browning in apple tissue as determined by a single CT slice in intact fruit. *Postharvest Biol. Technol.*
- Yang, Z., Richards, K., Kurniawan, N.D., Petrou, S., Reutens, D.C., 2012. MRI-guided volume reconstruction of mouse brain from histological sections. *J. Neurosci. Methods* 211 (2), 210–217. <https://doi.org/10.1016/j.jneumeth.2012.08.021>.
- Zeegers, M.T., van Leeuwen, T., Pelt, D.M., Coban, S.B., van Liere, R., Batenburg, K.J., 2022. A tomographic workflow to enable deep learning for X-ray based foreign object detection. *Expert Syst. Appl.* 206, 117768. <https://doi.org/10.1016/j.eswa.2022.117768>.
- Zhang, Z., 2000. A flexible new technique for camera calibration. *IEEE Trans. Pattern Anal. Mach. Intell.* 22 (11), 1330–1334.