Hyperspectral ZT and Food Safety Determination (Milestones 9) – Final Report
Milestone 9: Final Report

TABLE OF CONTENTS

1 EXECUTIVE SUMMARY

In this project a hyperspectral faeces and ingesta contaminant detection system is develop. The system uses data from a broad spectrum hyperspectral camera and a classification algorithm to detect contaminants. The system was designed to discriminate between clean and contaminated meat on a per pixel basis. As such, contaminants small as 1 mm^2 can be detected.

Hyperspectral data was collected on four separate occasions to design and refine the classification algorithm. In the first three experiments samples of beef and/or lamb were collected and contaminated manually with faeces and ingesta in a controlled manner. First, a preliminary data set was collected and scanned in a controlled laboratory environment to study the spectral characteristics of contaminated meat. Then, after the preliminary data was analyzed and a classification algorithm developed a, secondary data set was collected in a controlled abattoir environment to test and refine the classification algorithms. After the algorithm was confirmed to be robust, a contaminant detection rig suitable for wash down abattoir conditions was designed and constructed. To test the rig before abattoir trials, a tertiary set of data was collected on whole lamb carcasses to simulate abattoir conditions and to prepare for the final trials. Finally, the system was shipped to a lamb/mutton abattoir where the rig classified and collected data for two weeks.

In the spectral analysis of clean and contaminated samples, key wavelengths were identified and ranked according to their contribution to the classification algorithm. The analysis shows that a robust classifier can be constructed by using only 3 specific wavelength. However, best results were obtained when 4 or more wavelengths were used. Using only 3-4 wavelengths is desirable because conventional cameras with dedicated optical filters can be used instead of expensive and restrictive hyperspectral cameras.

The contaminant detection system developed throughout this research is capable of real time contaminant classification of whole carcasses, moving at a maximum speed of up to 1 m/s. The system was tested at an abattoir where more than 20,000 lamb and mutton carcasses were classified continuously over a period of two weeks. This trial demonstrated that visible contaminates can be reliably detected on the outer surface of lamb with the proposed system, and there is evidence to suggest that invisible levels of contamination can also be detected by the system.

2   INTRODUCTION

Hyperspectral imaging technology presents a unique opportunity to the red meat industry, particularly in the area of food safety. With the current growing need to drive down production costs and increase efficiency, the food industry is faced with a number of challenges, including maintenance of high quality standards and assurance of food safety while avoiding liability issues.

Prior research in real time faeces and ingesta detection exclusively relies on the fluorescence of chlorophyll found in digested plant matter (Lefcourt, 2005), (Ashby, 2003) and (Kim, 2003) . Such systems illuminate meat with high intensity ultra-violate light to make the chlorophyll fluoresce in the red spectra. The fluorescence in such systems is often very small and requires high intensity light for it to be visible and detectable on an imaging sensor. As such, the high intensity UV light can be potentially harmful especially because the harmful exposure may go unnoticed as UV is invisible to the human eye.

The objectives of this project are to analyze the spectral characteristics of contaminated meats using a hyperspectral camera, to develop a classification algorithm to discriminate between clean and contaminated meat and finally to design, construct and trial a contaminant detection system in an abattoir kill floor.

To meet these objectives the project was split up into a number of stages consisting of:

1. Preliminary data collection (Section 0).
   A controlled experiment was constructed to collect hyperspectral images of contaminated samples of meat. Faeces and ingesta contaminants were applied in small incremental quantities to samples of meat before capturing hyperspectral images of these samples for later analysis.
2. Preliminary data analysis (Section 0).
   Hyperspectral data was labeled and a number of classification models were tested to identify the best model and a range of "best" wavelengths that are useful for discriminating between clean and contaminated meat.
3. Software development – Data acquisition, classification and display (Section 6).
   Custom data acquisition, classification and display software was developed to acquire and process only 4 "best" wavelengths.
4. Software trial and data analysis (Section 7).
   The acquisition and classification software was tested in a controlled site trial, and the classification model was refined and evaluated.
5. System designed and development (Section 8).
   A contaminant detection system was designed for testing in an abattoir environment. The system consisted of new camera, lighting and software and was designed to be mobile and adjustable so that it can be positioned at a variety of stages at a lamb kill floor.
6. In-house system trial (Section 9).
   In-house trials were performed to test and refine the contaminant detection system on whole carcasses.
7. Onsite production trial (Section 10).
   Finally the system was shipped to an abattoir and trialed for two weeks.

The contaminant detection system developed through this research is able to reliably detect all visible traces of faeces and ingesta contaminants on the outer surface of lamb carcasses, and in some controlled cases is shown to detect invisible traces of contaminants.

Our approach is to illuminate meat with a broad spectrum light and take images of contaminated samples with a hyperspectral camera. Areas of the meat that are known to be contaminated will be labeled as such with an image mask. The hyperspectral data will then be analyzed manually and with machine learning algorithms to try and identify spectra that are most useful for discriminating between clean and contaminated meat.

Once a set of wavelengths are identified and an algorithm selected, an appropriate camera was purchased and a testing rig was designed and tested on site in an abattoir.

The classification algorithm developed in this research is a data driven process which requires accurately labeled data to train the algorithm. As such the algorithm can only accurately classify data similar to the data used for training. As such, a potential limitation of the system is that a particular variation of diet or animal species may have not been included in the training data set. Being a surface scan technology, the system requires even illumination to work with high reliability and this is difficult to achieve for a whole carcass, especially in the rib cage area where the light is often partially obstructed.

Another limitation of the system is it is sensitive to changes in lighting conditions. While all efforts have been made to make the illumination broad spectrum, intense, uniform, constant and diffused, ultimately the motion, non-uniform shape and size of the carcasses do cause variations in light intensity along the scanned surface. We compensate for this variation in lighting conditions by training the classification algorithm on a highly diverse data set so as to include many lighting variations. As such, our results from abattoir tests show that contaminants can be reliably detected despite the variations in lighting conditions.

3   PROJECT OBJECTIVES

The objectives of this phase are to utilize the existing Scott HSI camera and algorithm platform and:

1. Further refine Phase 1 algorithms.

2. Design and build a rig to enable the camera to be installed (temporarily) into a lamb slaughter floor/boning room (including waterproofing cabinet for camera and control box), and ascertain lighting requirements.

3. Evaluation of system at site whilst further refining of the algorithms.

4. Select and house high resolution camera.

5. Further refine algorithms in select regions and contaminants using a higher resolution of data.

6. Evaluation of refined algorithms at a lamb site.

4 PRELIMINARY DATA COLLECTION

4.1 Overview

In the initial data collection stage the focus was to collect hyperspectral images of both beef and lamb, contaminated with faeces and ingesta. Samples of meat and contaminant were collected from the JBS abattoir in Brooklyn Victoria. To capture a diverse set of data, the meat samples were cut into pieces containing equal amounts of fat and lean meat. First, clean samples were weighed, and scanned with two hyperspectral cameras, under both halogen an UV light. The experiment was repeated a multiple times with a small amount of contaminant incrementally added to the samples every time until the samples were saturated.

4.2 Data Collection Methodology

4.2.1 Sample collection

Samples were collected from the JBS abattoir in Brooklyn Victoria on 31 May 2017. Approximately 1.5kg of Beef loin and 1.5kg of Lamb rib cuts were collected. Faeces samples were collected from the stock yards and ingesta was collected from the offal room.

4.2.2 Sample preparation and acquisition procedure

The samples of beef and lamb were cut into 4 equally sized pieces (4 beef and 4 lamb) containing equal portions of lean and fat as shown in Figure 1. Two samples were contaminated with faeces and ingesta and scanned under halogen light, and the other two were similarly contaminated but scanned under UV light.

Figure 1: Example of uncontaminated beef sample.

For acquisition identification purposes every sample was labeled with a unique code before every acquisition. The code has the following format:

[Sample #].[Meat: B/L].[Contaminant: F/I].[Lighting: H/UVA].[Acquisition #]

For example, the unique code shown in Figure 1 describes the following acquisition:

Sample #:   1
Meat: Beef
Contaminant: Ingesta (not added yet as this is the first acquisition – the clean acquisition)
Lighting: Halogen
Acquisition #: 1

The next acquisition for the sample shown in Figure 1 is shown in Figure 2. Note that this sample has been contaminated and has the acquisition number "2" appended to the end.

Figure 2: Example of contaminated beef sample.

For reference RGB photos were taken of every sample prior to every acquisition.

4.2.3   Hyperspectral Equipment

Two Hyperspectral line scan cameras were used for data acquisition. A comparison of the two cameras is shown in Table 1, and their spectral range is illustrated in Figure 3.

	Camera A	Camera B
Spectral Bands (no. bands)	600-700	200-300
Spatial Resolution (no. pixels)	300-400	1000-1100
Maximum Frame Rate (Hz)	100	330
Minimum Focal Length (mm)	≈1000	≈200

Table 1: A comparison of the two cameras used in the experiments.
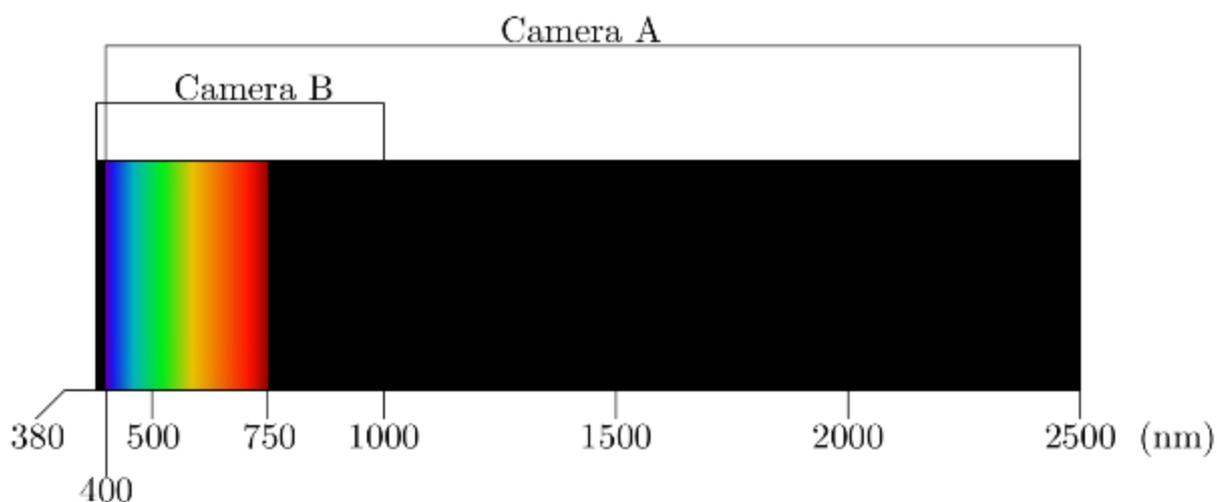
Figure 3: Spectral range of the two hyperspectral cameras compared to visible spectrum.

While the camera A covers a greater spectral range than camera B, we include data from camera B for its superior spatial resolution and its shorter focal length. The higher spatial resolution allows for more spectral samples to be taken and the shorter focal point allows the camera to be positioned closer to the samples, hence allowing for even more data to be collected and at a higher fidelity.

Figure 4 and Figure 5 show the camera, lighting and conveyor setup.

Figure 4: Camera A lighting and conveyor setup.
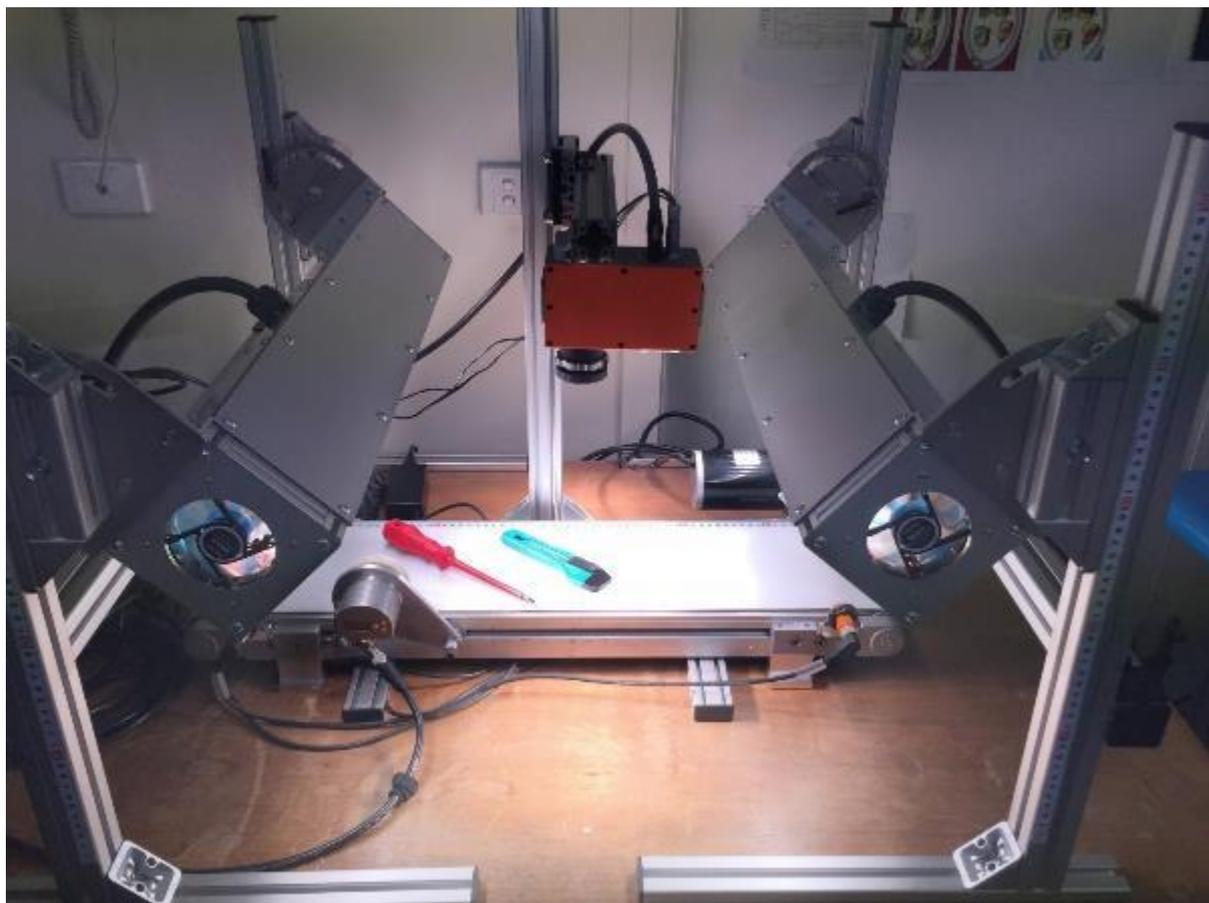
Figure 5: Camera B lighting and conveyor setup.

4.2.4 Camera Calibration

The following calibrations were performed on both cameras:

1) Height of camera above the sample was set to the minimum focal distance in order to use as much of the field of view as possible.
2) Focal length was adjusted to produce a sharp image at the sample height.
3) Exposure was adjusted by exposing a Spectralon[1] to the halogen or UV light source and setting the exposure time so that the brightest spectral reading was at approximately 80% of the sensors range. This ensures that most of the sensors dynamic range without saturation during acquisition.
4) Frame rate was adjusted to its maximum, which was either limited by the data rate on the camera link or the exposure time.
5) White and Dark reference images where captured and will be used to normalize the sensitivity of each pixel after acquisition.

---

[1] A fluoropolymer with the highest diffuse reflectance of any known material over the ultraviolet, visible and near-infrared regions.

4.2.5   Lighting

Two light sources were used to illuminate the samples for acquisition, a halogen light which produces a broad spectrum of light, and a UV light.

Despite the UV light having a spectral output below the detection capabilities of both hyperspectral cameras, it was suspected that the chloroform in faeces and ingest would flouress in the red to infrared spectrum, and hence be detected by the hyperspectral cameras.

4.3   Data Collection Summary

Table 2 summarizes all camera, lighting, meat and contaminant combinations that were used for the data acquired stage of this project. Each row represents a unique sample combination while the columns describe combination attributes:

a) **Cameras A and B**

b) **Light sources (Halogen and/or UV-A)**

c) **Meat (Beef and Lamb)**

d) **Contaminant (Faeces and Ingesta from respective animal)**

e) **Level of contamination (ranging from clean to saturated)**

Camera (a)	Light (b)	Meat (c)	Contaminant (d)	Level of contamination (e)					
				Clean	Light	Moderate	Heavy	Saturated	Just Contaminant
Cam A	Halogen	Beef	Faeces	✔	✔	✔	✔	✔	✔
			Ingesta	✔	✔	✔	✔	✔	✔
		Lamb	Faeces	✔	✔	✔	✔	✔	✔
			Ingesta	✔	✔	✔	✔	✔	✔
	UV-A	Beef	Faeces	✔	✔	✔	✔	✔	✔
			Ingesta	✔	✔	✔	✔	✔	✔
		Lamb	Faeces	✔	✔	✔	✔	✔	✔
			Ingesta	✔	✔	✔	✔	✔	✔
Cam B	Halogen	Beef	Faeces	✔	✔	✔	✔	✔	✔
			Ingesta	✔	✔	✔	✔	✔	✔
		Lamb	Faeces	✔	✔	✔	✔	✔	✔
			Ingesta	✔	✔	✔	✔	✔	✔

Table 2: Summary of data acquisition combinations performed in preliminary tests.

5   PRELIMINARY DATA ANALYSIS AND MODEL SELECTION

5.1   Overview

This section describes the hyperspectral data analysis on contaminated beef and lamb images. A Feed Forward Neural Network (FFNN) model for faeces and ingesta contaminant classification is identified as a suitable model. To illustrate the pros and cons of a broad range of available network topologies, two FFNN classification models from opposite ends of the model flexibility spectrum are compared. The lower complexity model uses only 3 wavelength bands and the high complexity model using 19 wavelength bands. It is shown that the 3-band model generalizes well on new data, while the 19-band model is able to fit the training data marginally better, but requires a highly diverse training set to generalize well. The 3-band model has the additional advantage of not being restricted to expensive hyperspectral camera technology, instead a much simpler, cheaper and more flexible camera technology can be used to effectively detect contaminants.

5.2   Methodology

5.2.1   Hyperspectral Data Set

The data set consists of hyperspectral images of beef and lamb contaminated with faeces and ingesta, taken with both hyperspectral cameras and illuminated with a halogen light source. The meat samples were specifically chosen to have approximately equal portions of fat and red meat to represent a broader range of variation in the meat. The Red-Green-Blue (RGB) representation of the camera A beef and lamb hyperspectral images are shown in Figure 6 and Figure 8 respectively, and camera B beef and lamb hyperspectral images are shown in Figure 7 and Figure 9 respectively.

The image are labeled according to the following format:

[image set ID]   [meat ID]   [contaminant]   [contamination %].

image set ID   **"A", "B" or "C" used to identify the image set. These sets are described in detail in the following section.**

meat ID   **number that identifies the unique cut of meat.**

contaminant   **"fae" or "ing" representing faeces and ingesta respectively.**

contamination %   **percentage indication of how much contaminant was applied or the level of dilution of the contaminant.**

For brevity, the camera name and meat type are not shown in the image labels. This information is recorded in figure captions.

The following is a description of the three image sets of beef and one of lamb that were collected over a period of 14 months and used for analysis in this report.

Set A: June 2017

This data was collection was described in the previous section. It contains beef and lamb images, contaminated with increasing amounts of faeces and ingesta. Set A images were captured with camera A and B cameras, and are shown in Figure 6 to Figure 9. Between scans, contaminants were incrementally added to the meat in small amounts (with a small paint brush) so as to record a range of contamination levels. The contamination percentage is only indicative of the level of contamination and does not reflect the actual percentage. The meats and contaminants were collected from JBS Brooklyn plant.

Set B: August 2016

This data set was captured with camera A and contains only beef samples, contaminated with varying levels of faeces and ingesta. The contaminant was prepared by diluting either faeces or ingesta in water and then dipping the top surface of the meat into the solution. In this set, the contamination percentage is an accurate indication of contamination level. Set B images are shown in Figure 6. The beef was purchased from a local butcher and the contaminant samples were collected from JBS Brooklyn plant.

Set C: April 2016

This data set was captured with camera A and contains only beef samples, contaminated with blotches of faeces and ingesta. The contaminant was applied unevenly to the meat in an uncontrolled manner. The percentage values are only indicative of the presence of contaminant and do not accurately reflect how much contaminant is present. The beef was purchased from a local butcher and the contaminant samples were collected from JBS Brooklyn plant.
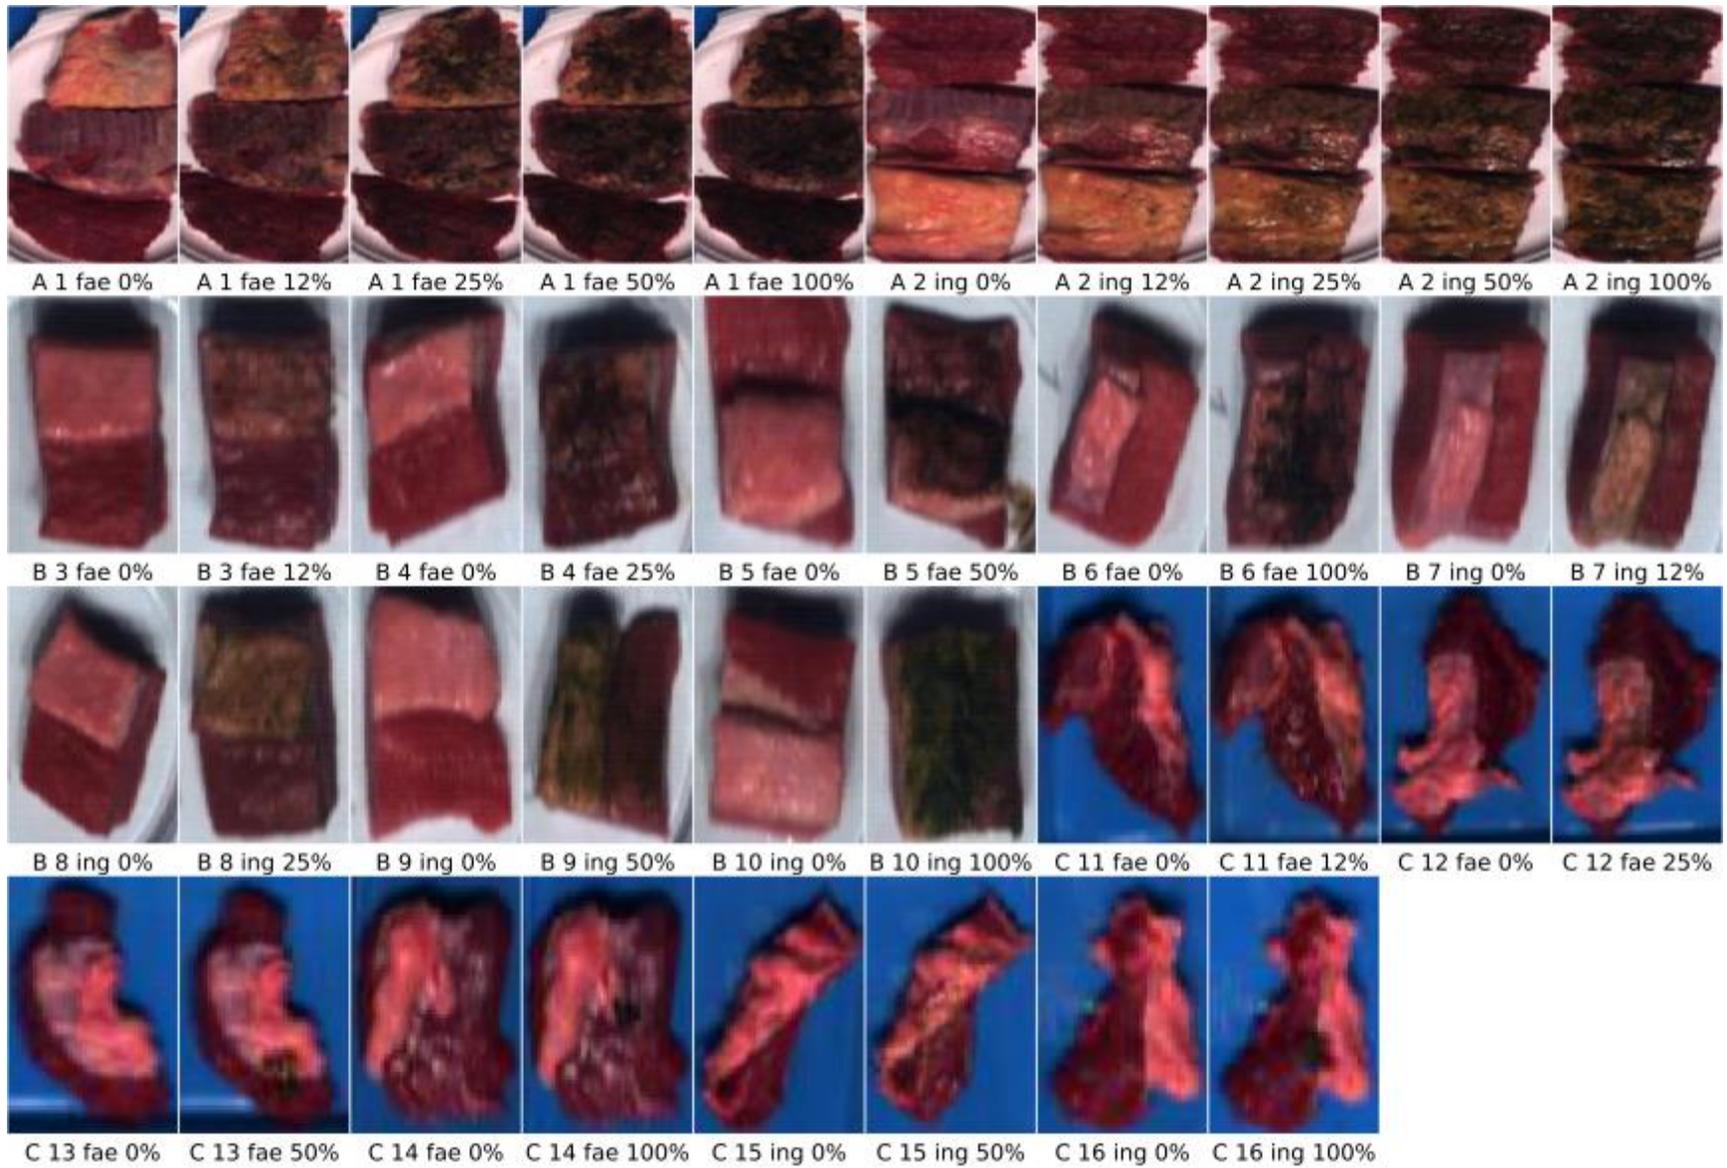
Figure 6: Camera A images of beef contaminated with varying levels of faeces and ingesta.

A 1 fae 0% | A 1 fae 12% | A 1 fae 25% | A 1 fae 50% | A 1 fae 100% | A 2 ing 0% | A 2 ing 12% | A 2 ing 25% | A 2 ing 50% | A 2 ing 100%
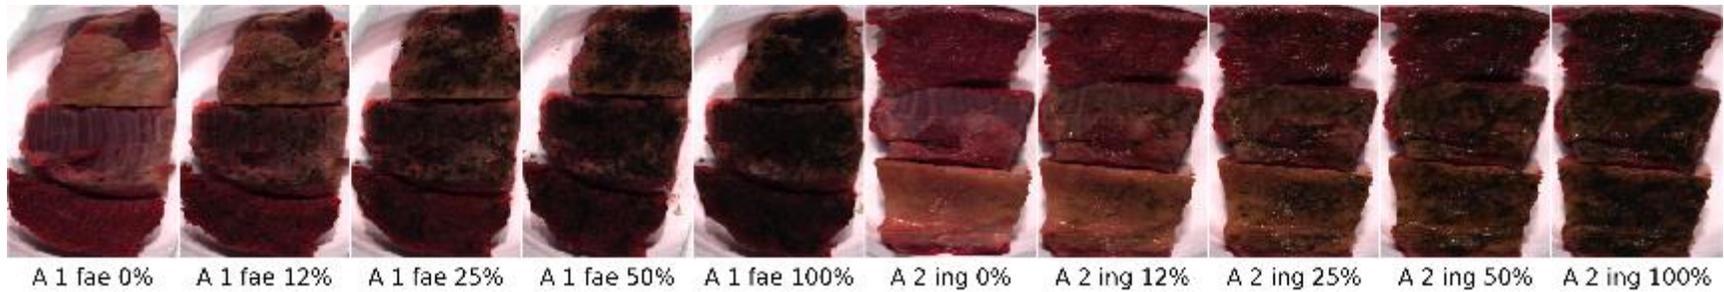
Figure 7: Camera B images of beef contaminated with varying levels of faeces and ingesta.

A 1 fae 0% | A 1 fae 12% | A 1 fae 25% | A 1 fae 50% | A 1 fae 100% | A 2 ing 0% | A 2 ing 12% | A 2 ing 25% | A 2 ing 50% | A 2 ing 100%
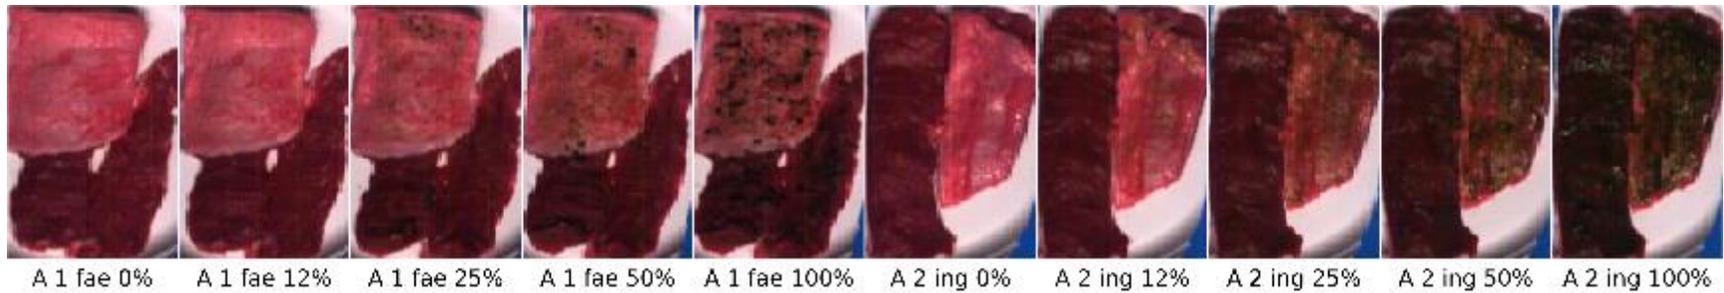
Figure 8: Camera A images of lamb contaminated with varying levels of faeces and ingesta.

A 1 fae 0% | A 1 fae 12% | A 1 fae 25% | A 1 fae 50% | A 1 fae 100% | A 2 ing 0% | A 2 ing 12% | A 2 ing 25% | A 2 ing 50% | A 2 ing 100%

Figure 9: Camera B images of lamb contaminated with varying levels of faeces and ingesta.

5.2.2 Data labelling

The Machine Learning models used for classification in this analysis are data driven models that require the training data to be labeled. In this case the labels consist of clean meat, contaminant and background. The background is regarded as the set of all things that are not meat or contaminant.

The data was manually labeled by creating image masks for all of the training images. An example of such a mask for a contaminated piece of meat is shown in Figure 10. In this example, the white part of the mask is used to represent the area of interest for that mask category. Areas of an image that did not clearly fall into any of the three classification categories, were intentionally excluded from all the masks (labeled black) so as to create accurate training data.

| RGB image | contaminant mask | background mask |
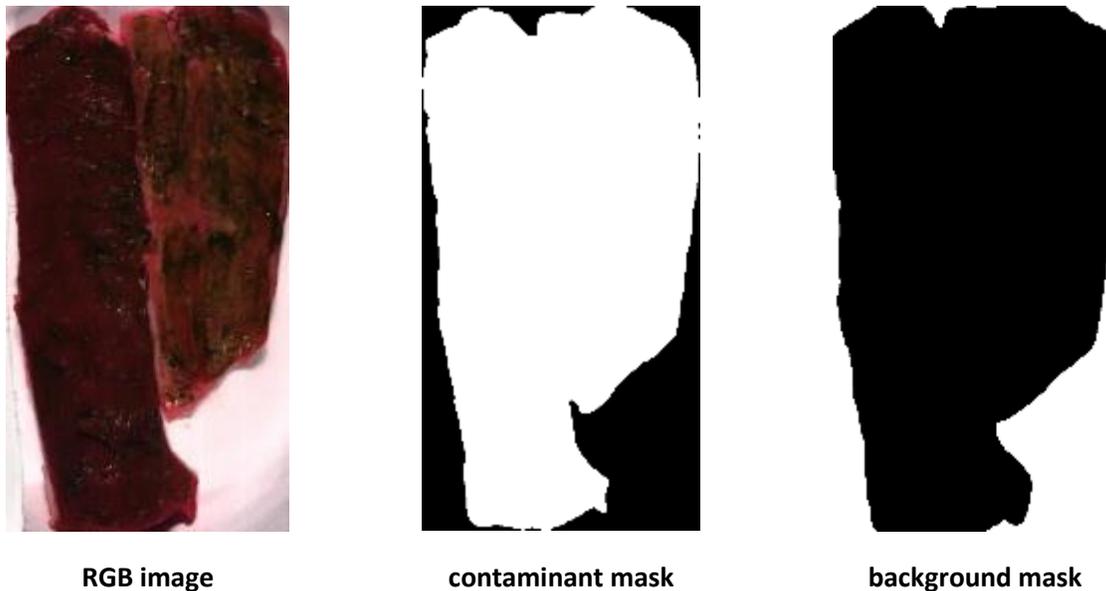
Figure 10: Example of image masks used to label contaminant and background areas of an image. White area represents the area of interest in each mask.

5.2.3 Classification Model Selection

The goal of this stage was to test a range of classification models to identify a range of model inputs, topology and parameters that would result in good classification performance. The input of these models are a set of specially selected spectral bands and the outputs are three numbers that represent the probability that the current pixel is either clean, contaminated or background as shown in Figure 11. In this example, the input spectrum of the current pixel is classified as clean meat with 80% confidence. Classification is performed on a per spectra basis, that is, the relevant spectra of one pixel is fed into the classifier at any one time and a percentage score is given to each pixel for all three categories.
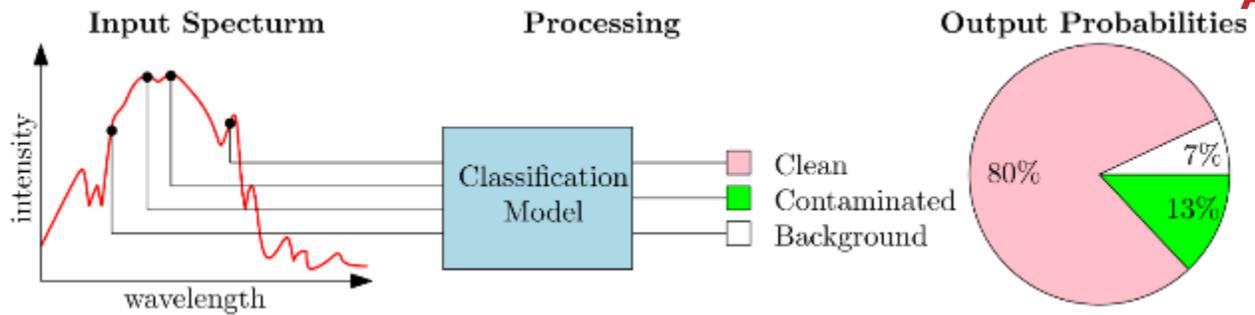
Figure 11: Information flow diagram showing the input spectral bands and out category probabilistic of the classification model.

Three classification models were proposed were tested, these include:

1. Feedforward Neural Networks (FFNN),
2. Support Vector Machines (SVM), and
3. Gaussian Process (GP).

All three classification models were tested and it was identified that FFNNs are the most suitable for this application.

In terms of prediction performance, all proposed methods performed well on small data sets. However, once the training data set was increased to more than 10,000 spectra, training the SVM and GP models became computationally impractical owing to their quadratic and cubic complexity respectively. In the larger training sets we used 100k – 1M spectra samples. Proceeding this much data with SVM and GP would require 100s of Gigabytes of random access memory (RAM) and many days of processing to train one model. For these reasons, SVM and GP models were rejected.

5.2.4   Training

FFNNs are parametric models which means they have many hyper parameters that need to be chosen to define the topology of the network. Once the hyper parameters are set, the model parameters can be trained using standard gradient decent algorithms to minimize the classification error.

The training process involves minimizing the classification error between predicted values and the target values by adjusting the model parameters. One of the difficulties of training a complex and high dimensional model is that the training process can converge on many "good" solutions, but it is unlikely to converge to the best solution. This is because the models used to train such models rely on algorithm based on gradient decent, which will converge to a local minimum and be completely ignorant of the fact that a global minimum may exist. To increase the chances of obtaining the best solution such models are often trained many thousands of times, at random parameter starting points in order to find a global minimum. Depending on the complexity of the model, the number of features and the number of samples, each training session may take up to 10 minutes.

High dimensional parameter spaces contain many local minima, which make it difficult to find the ideal solution. For this reason, the training process is repeated up to 10 times from random starting states, in order to increase the probability of finding a good solution.

The specific details of the hyper parameters is beyond the scope of this report, however, to find a set of good parameters, a search grid was created with a diverse range of hyper parameter combinations. Using k-fold cross validation as a performance metric, the hyper parameter combinations were reduced down to a small set.

5.2.5 Feature Selection

Feature selection is a technique used to reduce the input dimension size of the available data. Reducing the dimension size is an important step because it ultimately results in models that are easier to train, require less training data and generalize well to new data. In the context of this project, "features" refer to individual wavelength bands of the hyperspectral cameras. Choosing the "best" set of features can be regarded as hyper parameters selection problem, therefore, the feature set was chosen in conjunction with the other hyper parameters.

It should be noted, that alternative approaches to input dimension reduction can be performed by applying a transformations to the data (e.g. Principal Component Analysis) which projects the data into a lower dimensional space. Such transformations, can be very effective, however they require the use of all of the hyperspectral wavelength bands, which restricts the acquisition technology to a hyperspectral linescan camera. Such approaches were dismissed early in the analysis because it became clear that feature selection performed well and offered the possibility to move away from hyperspectral imaging.

Early in the analysis, features were selected visually, by studying the spectrum of clean and contaminated meats. Figure 12 shows the mean and 1 standard deviation of clean and 25% contaminated beef spectra. The shaded regions in these plots represent 68% of the data of the respective category. From this plot it is obvious that a single wavelength band cannot be used to discriminate clean from contaminated meat, instead, a higher dimensional space (more bands) is required. The wavelength range-of-interest was found by inspection and was within the range of the camera B, hence camera B data is also used in all experiments. In this range, the overlap between error bars is minimized and the general shape of the clean plot differs more significantly from the contaminated plot.
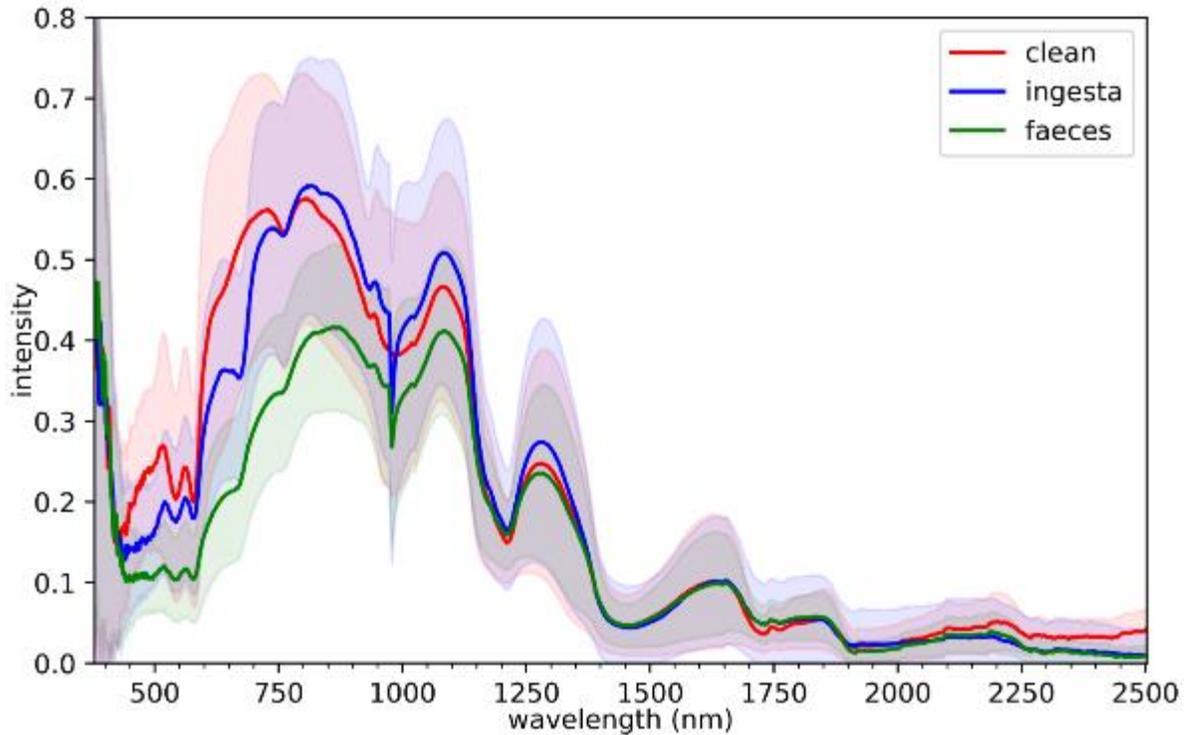
Figure 12: Spectrum mean and 1 standard deviation error bars of clean and 25% contaminated beef. 200 spectra samples were randomly selected from clean and contaminated images (A_1_fae_0%, A_1_fae_25% and A_2_ing_25%).

A genetic algorithm was used to find the set of best hyper parameters for the model topology. For every hyper parameter, the genetic algorithm generates multiple copies of the current model, each with hyper parameters randomly perturbed by a small amount. The resulting models are all trained and the model that performs best on a validation set is kept to generate the next generation of randomly perturbed models. This process was repeated until there was no more performance improvement.

A similar process was used to reduce the number of features. Features were removed one at a time, and those that resulted in negligible reduction to classification performance were permanently removed from the set. Using the above training procedure, the number of features were narrowed down to 3 wavelength bands. The scatterplot in Figure 13 shows the intensity values at the 3 "best" wavelengths after preprocessing the data with feature scaling and whitening (standard preprocessing procedures in machine learning). This plot shows that in this rather limited data set, there is a relatively clear separation boundary between clean and contaminated meat. However, as more diverse data is included, this separation boundary becomes obscured and highly non-linear requiring for a more complex model non-linear model to perform classification.
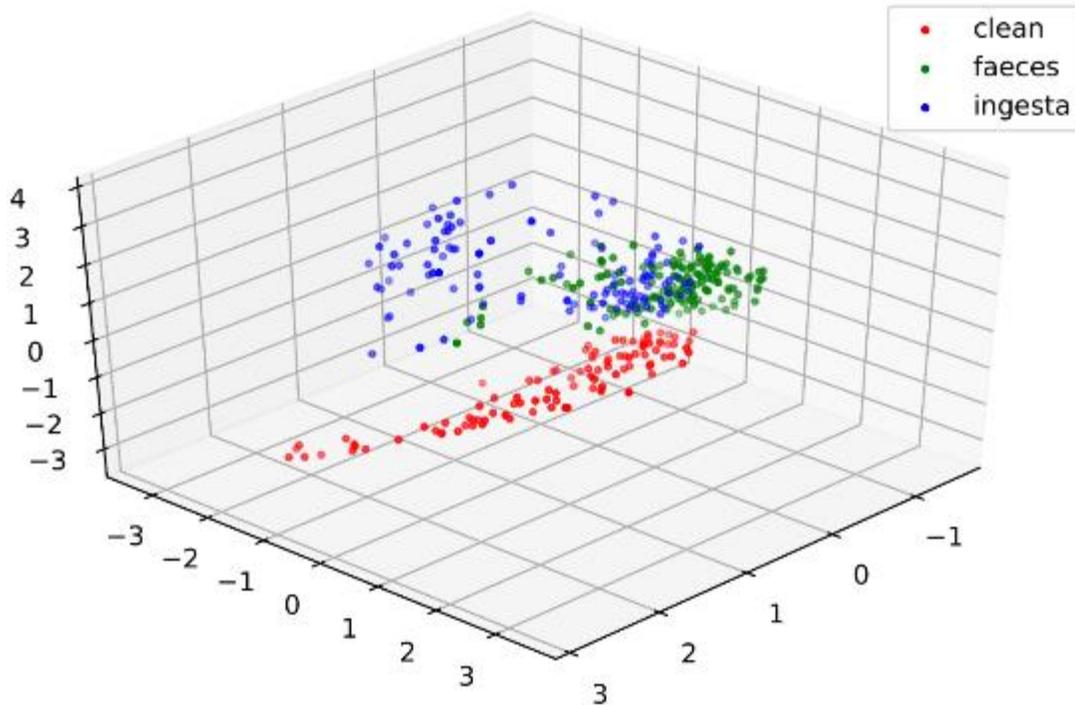
Figure 13: Scatterplot showing 3 'best' bands (after preprocessing) from the same data set as that used for Figure 12.

To verify that the best 3 wavelengths are indeed the best, a brute force approach was used to train and test all combination of a 3 wavelength model. This found the same set of 3 wavelengths as previously found by the more computationally efficient genetic algorithm. The brute force search was capped at a maximum of 3 wavelengths because the search space becomes too large requiring approximately 3 weeks of processing time to discover the best 4 wavelength model. The heat map in Figure 14 shows how the classification accuracy is affected by varying two of the three best wavelengths. The horizontal and vertical axis represent the index of the two wavelength bands and the heat map represents the accuracy of the classification model. Values for the lower right corner were not collected because the results would be identical to those in the upper left corner due to the symmetry of the search space.
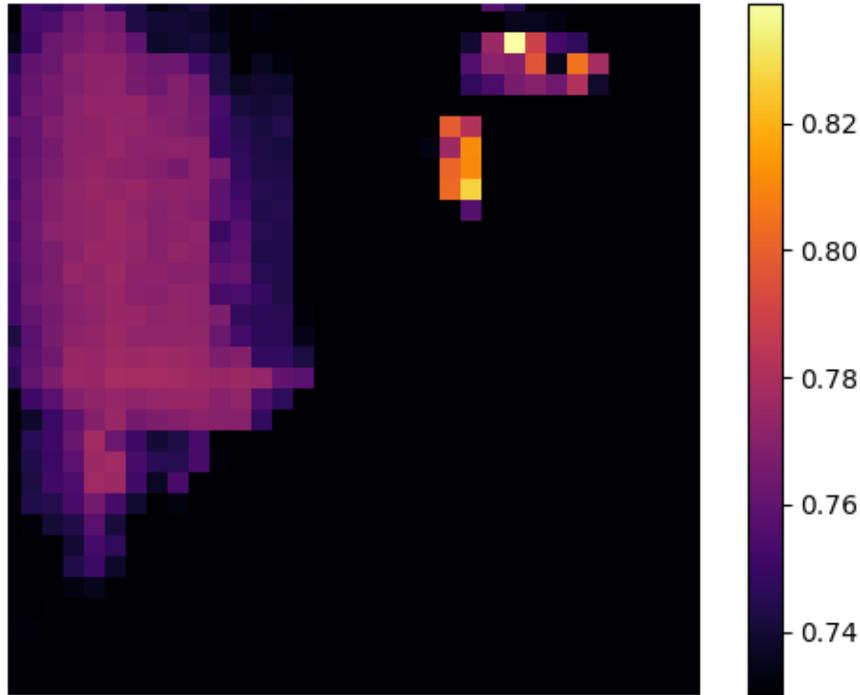
Figure 14: Heat map showing the classification accuracy when two of the three best wavelengths are varied.

5.3 Data Analysis Outcomes

In this section we briefly describe a set of experiments that are performed for the data analysis and summarize the results.

Classification results of two model topologies are shown, one using 3 and the other using 19 wavelength bands as inputs to the classification model. These models are chosen to represent the performance of the two ends of the model flexibility spectrum. The 3-band model represents the simplest possible solution that performs at a level that we regard as acceptable, and it generalizes well. The 19-band model represents a high complexity model which gives the best in-sample performance. However will be shown, the 19-band model does not generalize well when trained on a small data set. Using more than 19 bands resulted in negligible performance improvement.

In all of the experiments presented in this section the classification models are trained on only 30% of the spectra samples available in the chosen training set. We use only 30% of the data for the following reasons.

1. 30% of the data from each image is sufficient to represent the diversity of that image, using any more data results in a negligible performance improvement.
2. A goal of this analysis is to develop a model that generalizes well on spectra that were not included in the training set (the remaining 70% of the data).
3. A smaller data set reduces the computational complexity, allowing for more tests to be performed.

The results of two experiments are presented for both 3-band and 19-band models. These experiments are designed to illustrate the following two properties of the two classification models:

Experiment 1: Small Training Set

> To show how well the models generalize, that is to show how well they perform on spectra and images that were not included in the training set. This is achieved by training the model on only 30% of the spectra available in the training images, and also by using only four out of 48 beef (20 lamb) images in the training set. To further show how well the models generalize, the four training images are taken only from image set A (June 2017 batch), and are used to classify all other images of contaminated meat that where collected months apart, with different camera settings, with a different cameras, different lighting conditions and on meat and contaminant samples that were obtained from different sources.

Experiment 2: Large Training Set

> To show how well the models are expected to perform if a sufficiently large and diverse training set is available. This is achieved by training the models with 30% of the pixels from almost all of the available images.

In all of the results figures (Figure 15 to

Figure 24) the RGB representation of the hyperspectral images are overlaid with white for areas classified as background and green for areas classified as contaminant. Areas classified as meat remain unchanged so as to clearly see the meat part of the image.

To get an objective measure of performance, we show a confusion matrix for the classification results for images in set B (beef images only). Set B images are chosen because the meat samples from this set were toughly contaminated (beef was dipped into a diluted solution of the contaminant), which allows for the image to be accurately labeled. The confusion matrices are shown in Table 3 to Table 6. A confusion matrix show the percentage that a row category is classified as all other categories.

Note: Ultra Violate (UV) images were also taken as part of the data collection, however, early analysis of this data showed that UV light caused small amounts of fluorescence in very narrow wavelength bands, with all other bands containing only noise. Analysis of the fluorescent bands led to the conclusion that our UV setup cannot be used to discriminate between clean and contaminated meat. As such, the classification of the UV data will not be presented in this report.

5.3.1 Experiment 1: Small Training Set:

In this experiment the models are trained on four camera A images from set A, the classification test set includes all camera A images from set A, B and C and camera B images from set A. This test set represents a diverse set including variables such as meat and contaminant samples that were collected months apart and from different sources, and images that were collected with different cameras, lighting conditions and camera settings. With such a large test set diversity, this experiment demonstrates how well the models will perform when new - previously unseen piece of meats and/or contaminants are classified.

For training, one clean and one 100% contaminated image was chosen for each type of meat (beef or lamb) and for each contaminant (faeces or ingesta). The training images are highlighted with red fonts in Figure 15 to Figure 20.

5.3.1.1 3-band model with small training set

Figure 15 and Figure 16 show the beef classification results on images taken with camera A and B respectively, and Figure 17 shows the lamb classification results taken with both camera A (top row) and camera B (bottom row).

Given the small training set, the 3-band model (for both beef and lamb) performs well on images from both cameras. The only image that the classifier struggled with is 'B_3_fea_12%' in Figure 15, where the lower half of the image should have been classified as contaminant. As we shall see, this problem is overcome by using a more diverse training set. There are also some minor misclassification around the boarders of the clean images of set C in Figure 15, however this is due to the fact that a glass container was used to hold the meat for imaging. The glass caused irregular reflections which made it difficult for this simple model to learn all of the variation is the background. This issue is easily solved with the use of a uniformly colored opaque background as in set B. We also see some misclassification in set A, the edges of the paper plate are incorrectly classified as clean meat. Again, this issue is easily resolved with use of a uniformly colored opaque background.

The confusion matrix in Table 3 quantifies the out-of-sample performance of the 3-band classifier. It is out of sample because the classifier was trained on four images from set A and the confusion matrix only shows the classification results of set B. The 'clean' and 'background' category are quite accurately classified. However, 28% of the time contaminated pixels are incorrectly classified as 'clean'. While this misclassification may seem like a concern, recall that the classification matrix presents the classification rate on a pixel basis. So while it is true that 28% of contaminated will go undetected (under this model), this does not mean that there is a 28% chance that a contaminant on a whole piece of contaminated meat will go undetected. For example, if we assume that there is a cluster of 10 contaminated pixels. Each of these pixels has a probability of 0.28 of being undetected, therefore the probability that all 10 pixels are undetected is 0.28^{10} which is approximately equal to 1 in 338,000. It is also worth noting that clean meat is rarely misclassified as 'contaminated'.

The lamb image 'A_1_fae_12%' in Figure 17 had very light level of contamination. We see that the resolution of camera A (top row) is insufficient (approximately 2mm pixel size) to detect much of the contaminant. On the other hand, camera B was able to detect more of the contaminant.

		Predicted		
		Clean	**Contaminated**	**Background**
Actual	**Clean**	99.512	0.001	0.487
	Contaminated	28.124	71.836	0.04
	Background	0.002	0	99.998

Table 3: Confusion matrix showing the classification accuracy for test set B using the 3-band classifier, trained on four images shown in Figure 15.

5.3.1.2 19-band model with small training set

Figure 18 and Figure 19 show the beef classification results on images taken with camera A and B camera respectively, and Figure 20 shows the lamb classification results taken with both camera A (top row) and camera B (bottom row).

As you can see the 19-band model performs worse than the 3-band model on all images given a small training set. Even the clean images that were used in the training set are misclassified in some areas. This misclassification is due to the fact that the 19-band model is incredibly flexible and it fit the 30% training data very well, however, the remaining 70% of the data was inevitably slightly different and was incorrectly classified. This result clearly shows that a more complex model does not necessarily mean that it will perform better than a simpler model. On the other hand, if the model was trained on a very large data set which included all the diversities of meat and contaminant, then the more complex model would perform better than the simpler model. Evidence of this claim is presented in the next section where a larger training set is used.

The confusion matrix in Table 4 shows that the 19-band model does a rather poor job of classifying contaminants, with an accuracy of only 46.7%, and it appears that all other categories are fairly accurately classified. However, the actual performance for the 19-band model should actually be worse. For example, the confusion matrix shows that clean meat is accurately classified 98.9% of the time, however, the classification results of clean meat in Figure 18 do not agree. This discrepancy is due to way in which the images were labelled, we only labelled areas that we could classify with absolute certainty and so the dark regions at the top of the meat was not labelled as it is not clear if that part of the image is background or meat. Hence, the misclassifications in the clean images of set B in Figure 18 are not accurately reflected in the confusion matrix at all.

		Predicted		
		Clean	**Contaminated**	**Background**
Actual	**Clean**	98.951	0.001	1.048
	Contaminated	53.305	46.695	0
	Background	0	0	100

Table 4: Confusion matrix showing the classification accuracy for test set B using 19-band classifier, trained on four images shown in Figure 18.
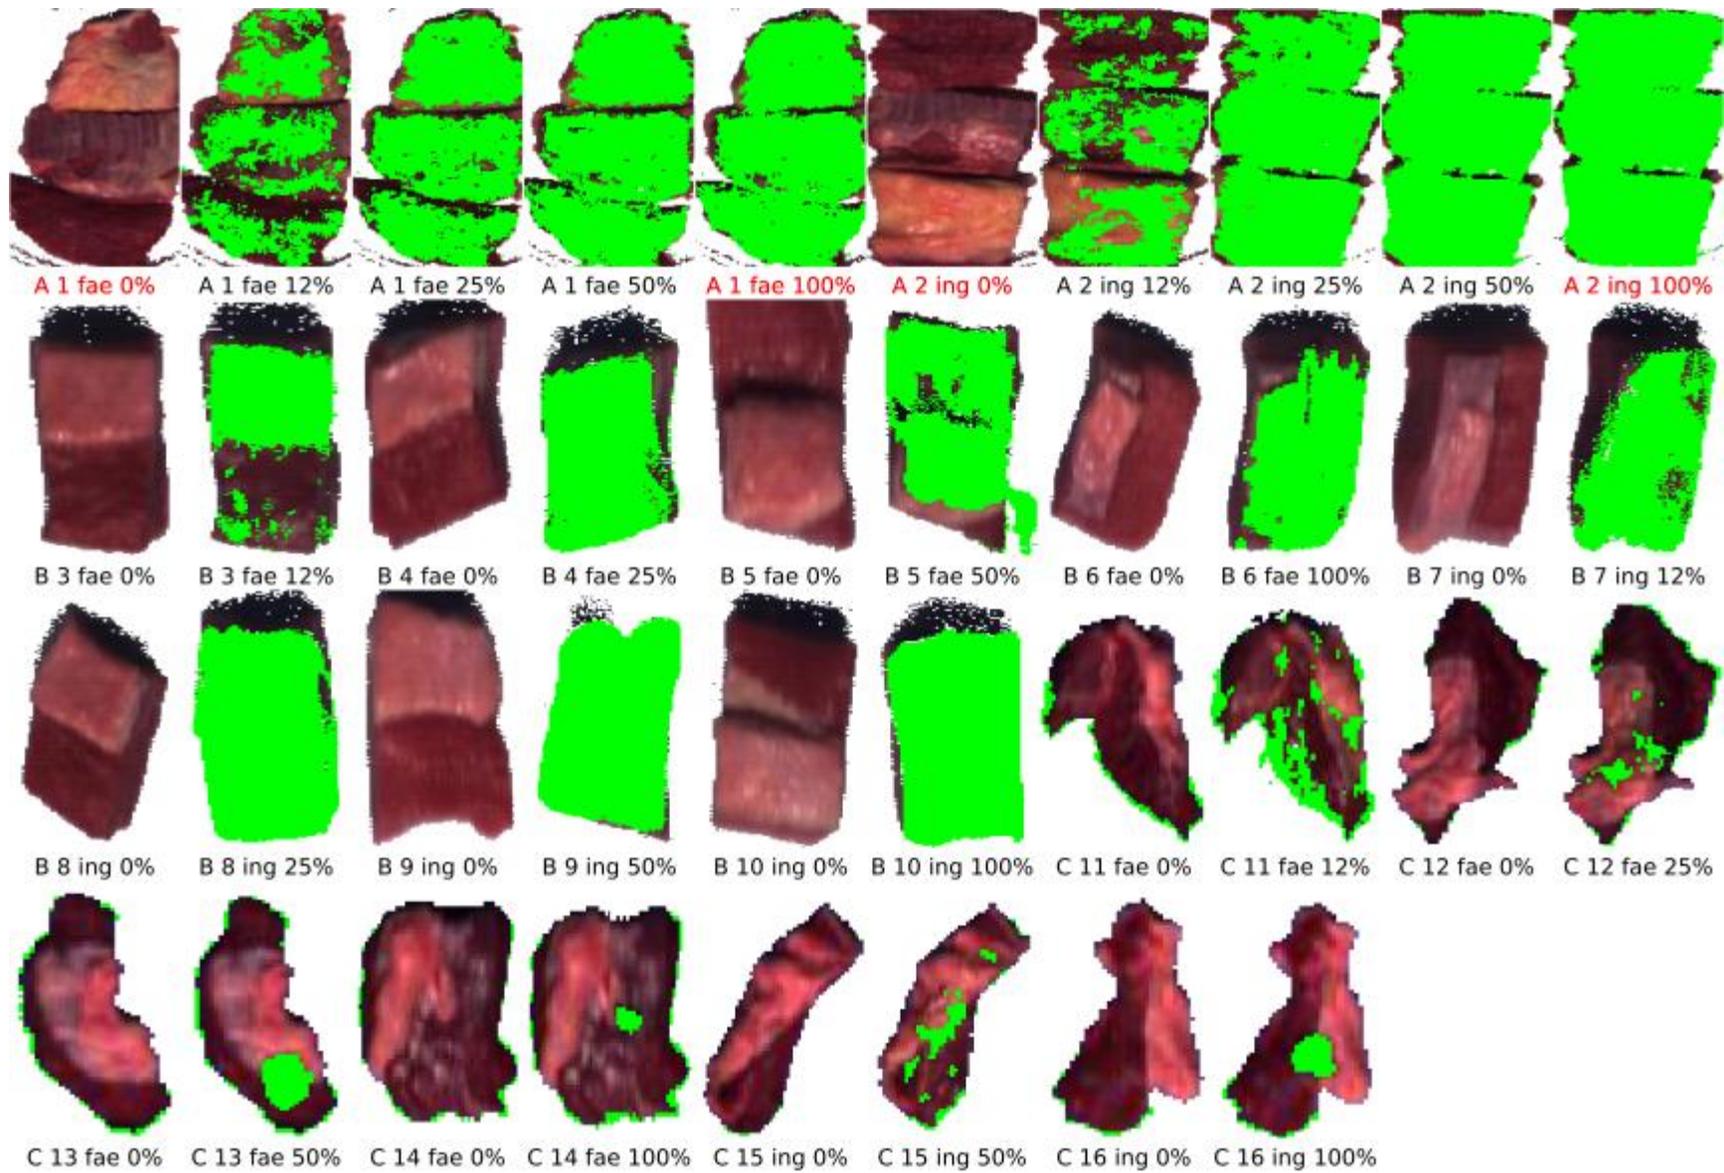
Figure 15: Beef classification results using 3 wavelength bands from camera A. Classification model was trained on the four camera A images highlighted in red.
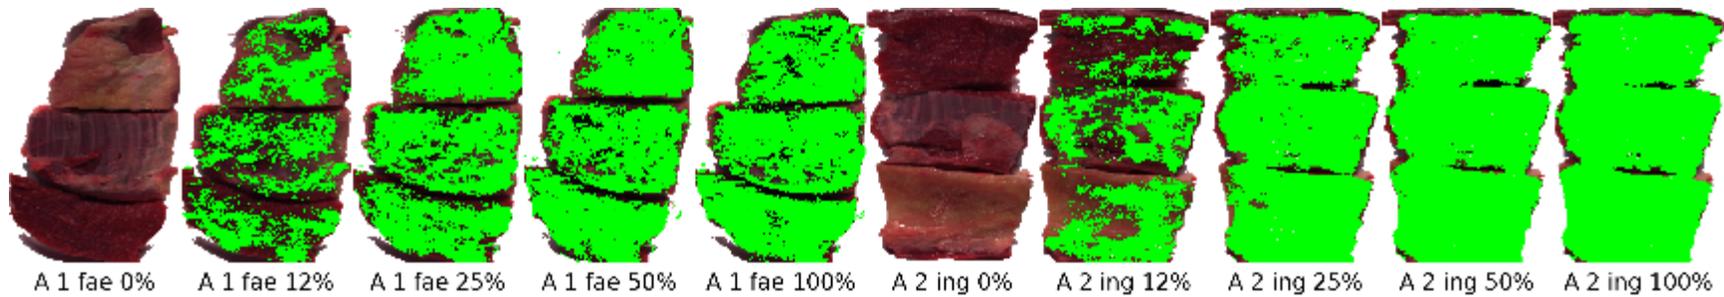
Figure 16: Beef classification results using 3 wavelength bands from camera B. Classification model was trained on four highlighted camera A images from Figure 15.
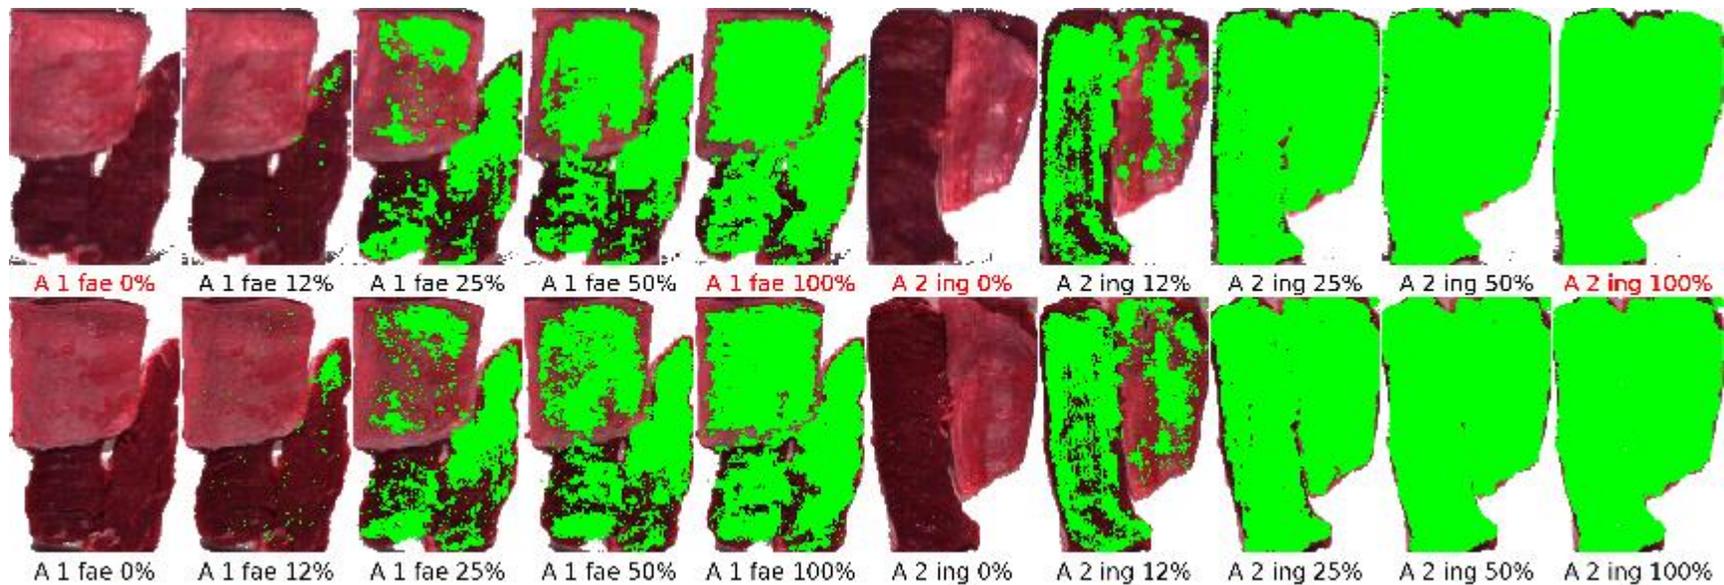
Figure 17: Lamb classification results using 3 wavelength bands from camera A (top) and camera B (bottom). Classification model was trained on the four camera A images highlighted in red.
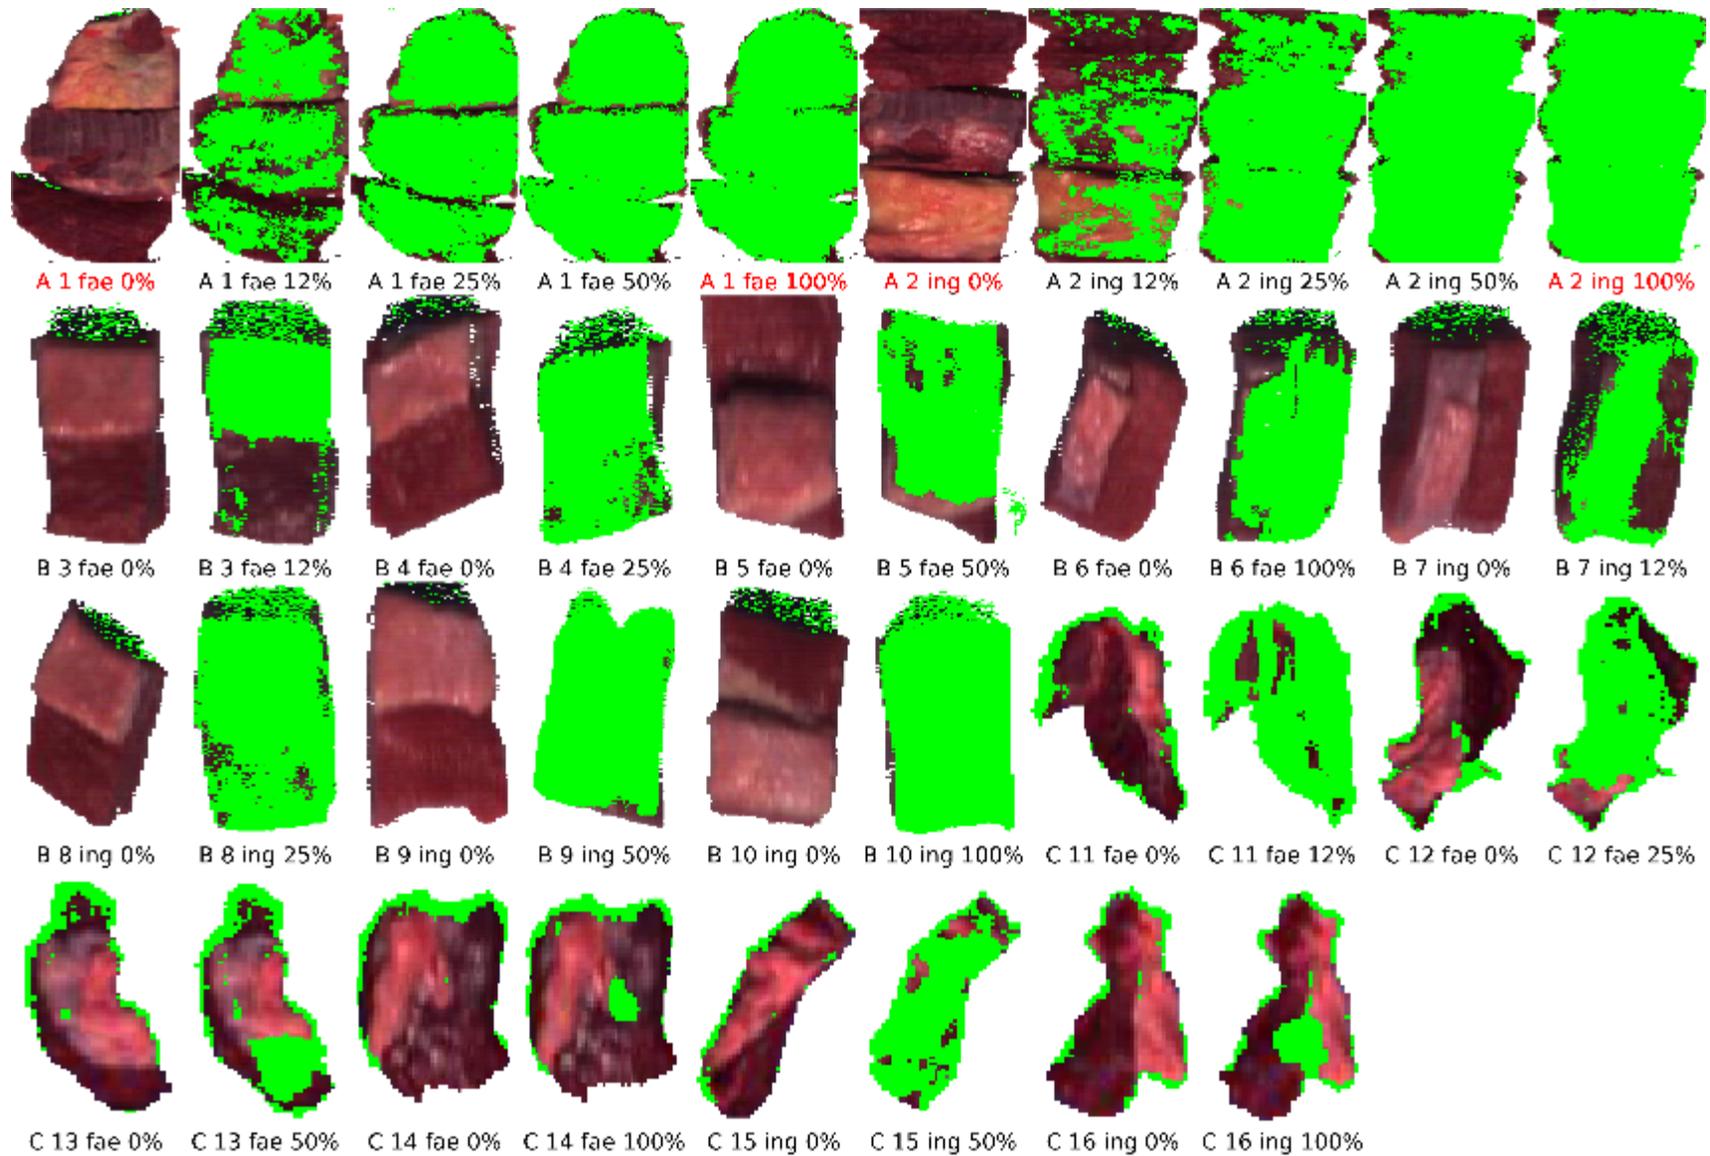
Figure 18: Beef classification results using 19 wavelength bands from camera A. Classification model was trained on the four camera A images highlighted in red.
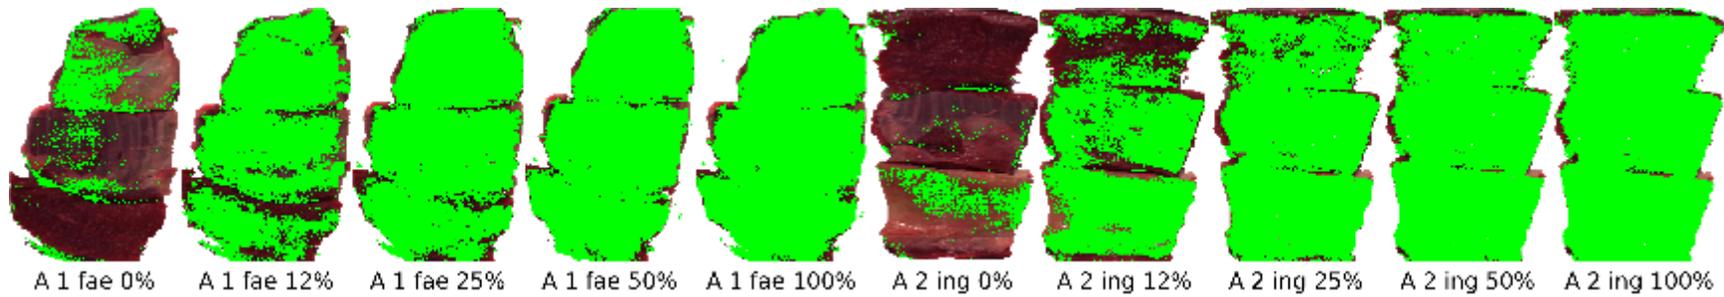
Figure 19: Beef classification results using 19 wavelength bands from camera B. Classification model was trained on four highlighted camera B images from Figure 18.
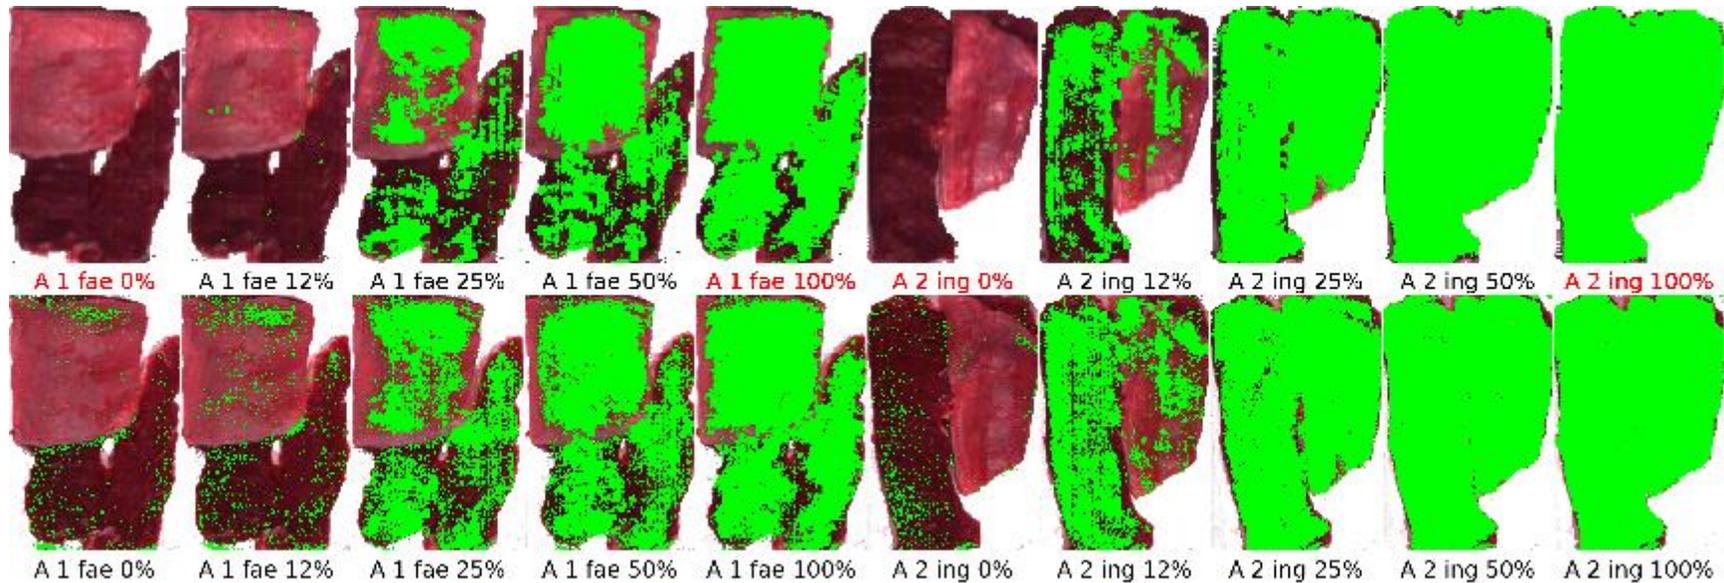
Figure 20: Lamb classification results using 19 wavelength bands from camera A (top) and camera B (bottom). Classification model was trained on the four camera A images highlighted in red.

5.3.2 Experiment 2: Large training set

In this experiment, the classifier is trained on 30% of the pixels in all clean and thoroughly contaminated beef images. Lamb was not included in this experiment because only set A data is available for lamb. The 26 training images from camera A are highlighted with red font in Figure 21 and Figure 23 for 3-band and 19-band model respectively. Figure 22 and

Figure 24 show the camera B classification results for the 3-band and 19-band models respectively. This experiment is designed to illustrate how well the two models would perform if a sufficient amount of training data can be made available.

5.3.2.1 3-band model with large training set

Figure 21 and Figure 22 show beef classification results on images taken with camera A and camera B respectively. These results show that the 3 band model performs very well when trained on a diverse data set. Most notable is the improved contaminant classification performance in 'B_3_fae_12%' in Figure 21, compared the small training set performance in Figure 15. Again the model struggles near the edges of the meat to accurately discriminate between the irregular backgrounds. However, as we pointed out earlier this problem is easily resolved with the use of a uniformly colored opaque background.

In some clean images, there are some small contaminant misclassifications. These occur because much of the model is dedicated to accurately classifying the various types of backgrounds that are present in the training data. By using a constant background, more of the models recourses would be used to discriminate between meat and contaminant. Additionally, these speckled misclassifications can be removed with some post processing as contaminated meat is expected to consist of larger clusters.

The confusion matrix in Table 5 confirms our observation that a large training set has significantly improved the performance of the 3-band model, increasing the contaminant detection rate from 72% to 92%.

		Predicted		
		Clean	**Contaminated**	**Background**
Actual	**Clean**	99.972	0.025	0.003
	Contaminated	7.288	92.704	0.008
	Background	0.007	0	99.993

Table 5: Confusion matrix showing the classification accuracy for test set B using 3-band classifier, trained on images shown in Figure 21.

5.3.2.2　19-band model with large training set

Figure 23 and

Figure 24 show beef classification results on images taken with the camera A and B respectively. Figure 23 shows that the 19-band model can perform exceptionally well when a large training set is used, accurately classifying all of the camera A images. However, once again the 19-band model performs poorly out of sample in

Figure 24 on camera B data, where it is used to classify images that were not included in the training set. These results show that a more complex model is able to perform very well, however, it requires a very large and diverse training set.

The confusion matrix in Table 6 confirms our observations a large training set produces good classification performance on the set B image set, however, it does not reflect the models poor out-of-sample performance (as illustrated in Figure 24).

		Predicted		
		Clean	**Contaminated**	**Background**
Actual	**Clean**	99.998	0.001	0.001
	Contaminated	0.104	99.896	0
	Background	0.136	0.002	99.863

Table 6: Confusion matrix showing the classification accuracy for test set B using 19-band classifier, trained on images shown in Figure 23.

5.4　Data Analysis Discussion

Overall, the results show that if a highly diverse training sample of meat and contaminant samples cannot be collected at the installation site, then a low complexity (3-band) model should be chosen. On the other hand, a higher complexity model should be chosen if access to a large and diverse training set can be guaranteed.

Given that spectral properties of meats and contaminants may change over time due to factors such as diet, age and breading, we strongly recommend to proceed with a simpler 3 to 5-band model that can generalize well. A low complexity model also has the advantage that it does not require the use of a hyperspectral camera. A camera capable of capturing a narrow bands centered at a few specific wavelengths will suffice. Additionally, a 3 to 5-band system does not restrict the camera technology to a linescan camera (as all hyperspectral cameras are), instead it is possible to use an area scan camera capable of capturing the entire area of interest in one shot.

We believe that the performance of a 3 to 5-band models can be further improved by using a more uniform light source (removing shadows), and by using a constant color opaque background. For example, the training data used in this analysis was collected months apart, by different people and so the backgrounds of the hyperspectral images consist of paper plates, plastic plates, glass containers, and a blue conveyor. As a consequence, some of the flexibility of the 3-band model has gone into accurately classifying these objects into the background category. By reducing shadows and using a uniform opaque

background, it is expected that there would be less variation in the data, which would therefore increase the separation boundary between the classification categories, making it easier for the model to fit the data.

5.5 Data Analysis Conclusions

This analysis shows that faeces and ingesta contaminants can be reliably detected on beef and lamb, using data from a hyperspectral camera. The results compared the performance of two classification models on a broad range of contaminated beef and lamb meat samples. The results show that the lower complexity 3-band model is able to generalize well from a small training set, as it is able to accurately classifying previously unseen contaminated meats. The more complex 19-band model performs very well on the training set, but requires much more diversity in the training set to generalize well. As training on a diverse range of meats and contaminants on site may be an issue, it was decided that the project proceed with the lower complexity 3 to 5-band classification model. An additional advantage of the 3 to 5-band model is that image acquisition is not restricted to a hyperspectral linescan camera. Instead, cheaper, simpler and more flexible camera technology can be used.
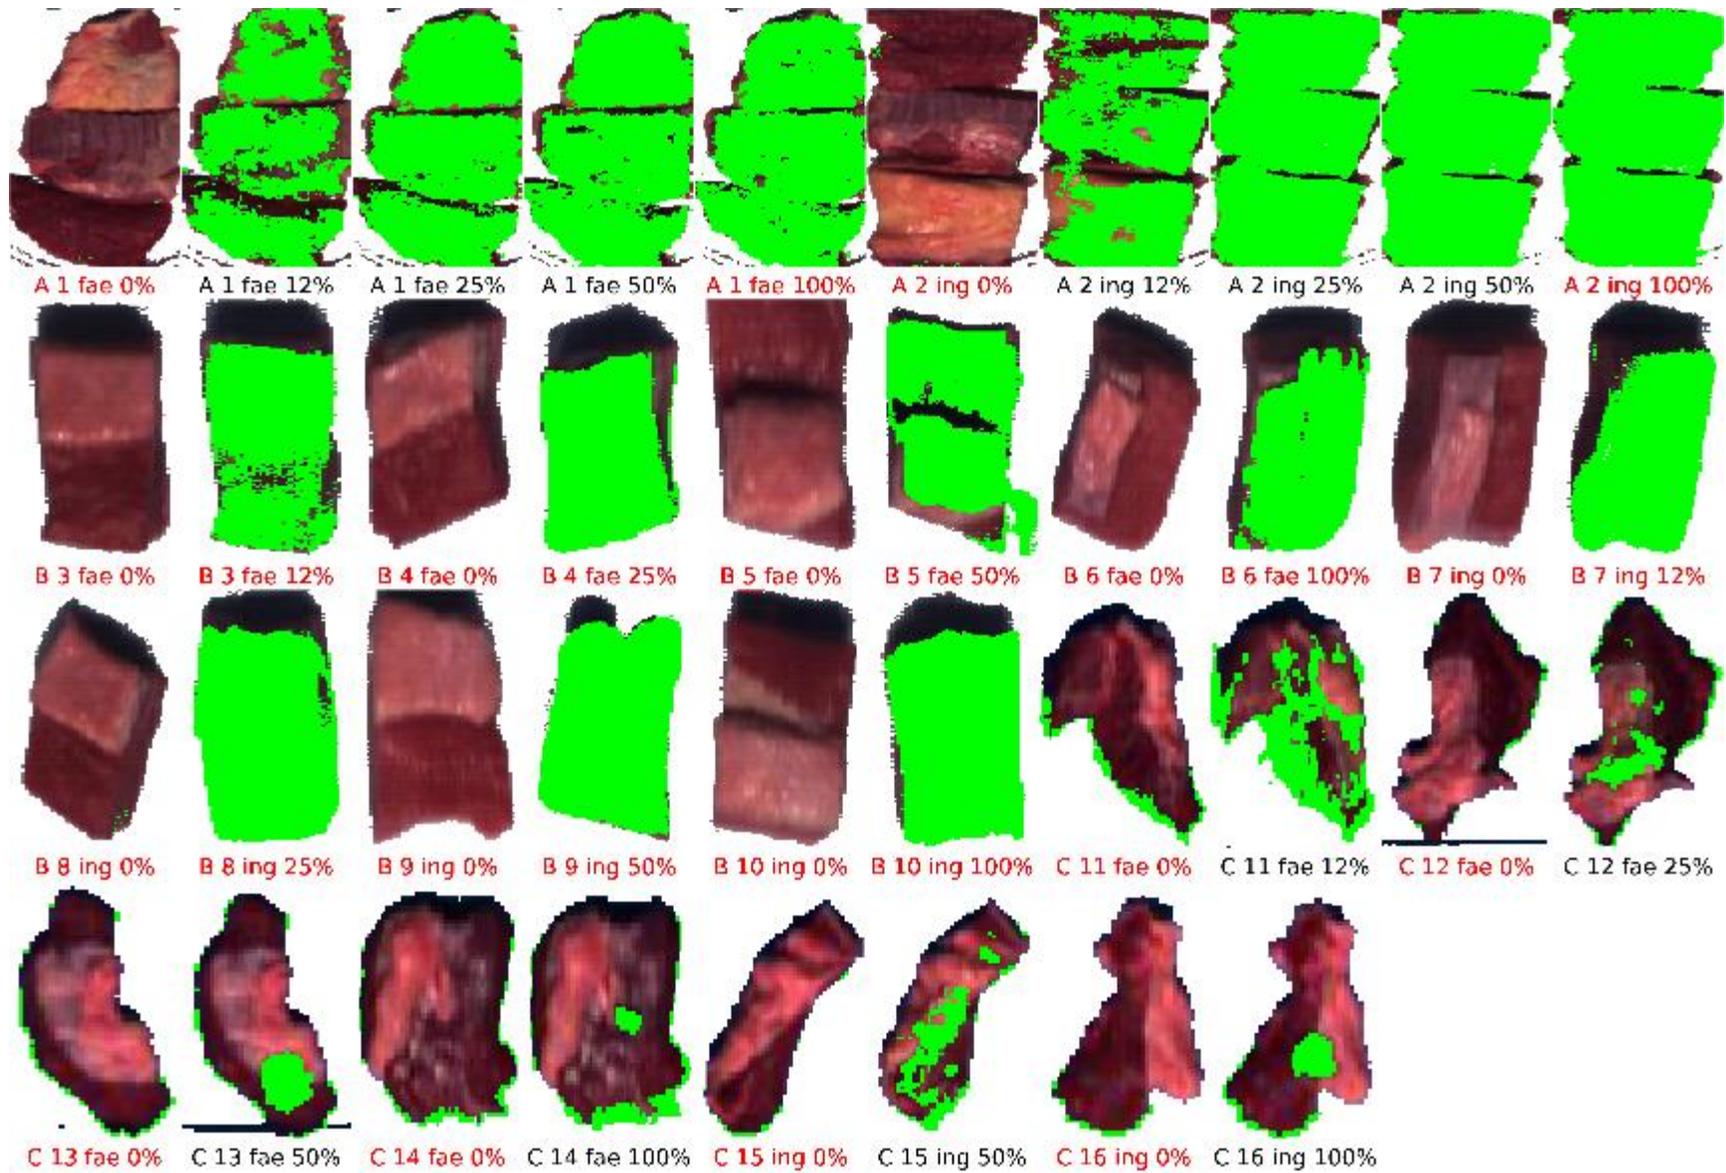
Figure 21: Beef classification results using 3 wavelength bands from camera A. Classification model was trained on images highlighted in red.
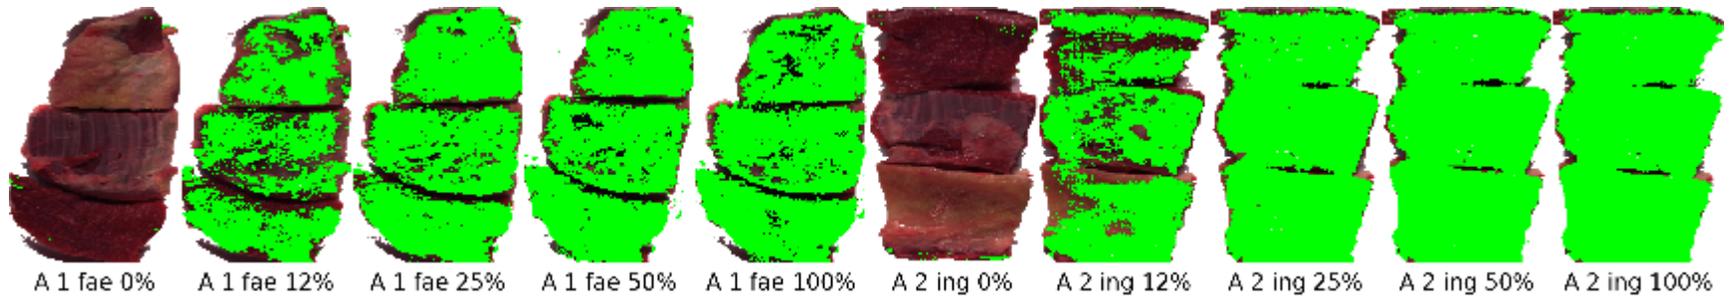
Figure 22: Beef classification results using 3 wavelength bands from camera B. Classification model was trained the highlighted camera A images from Figure 21.
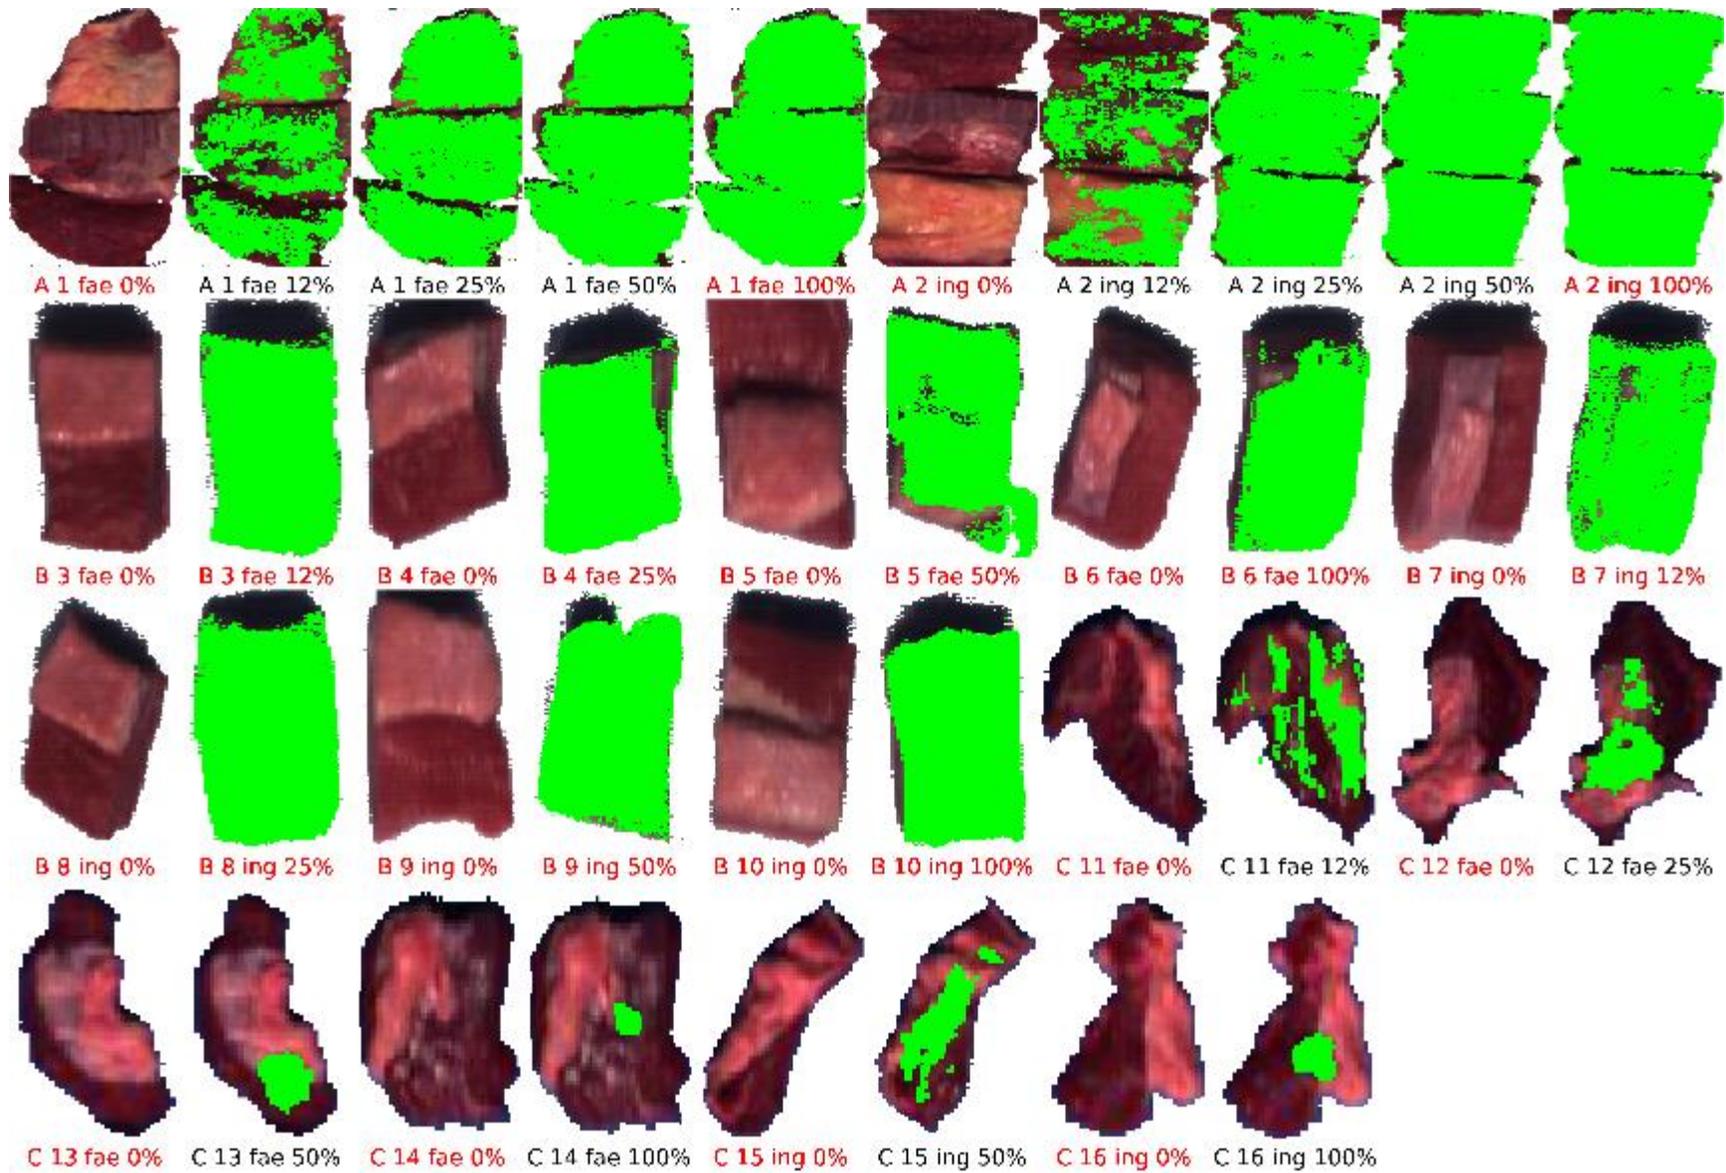
Figure 23: Beef classification results using 19 wavelength bands from camera A. Classification model was trained on images highlighted in red.
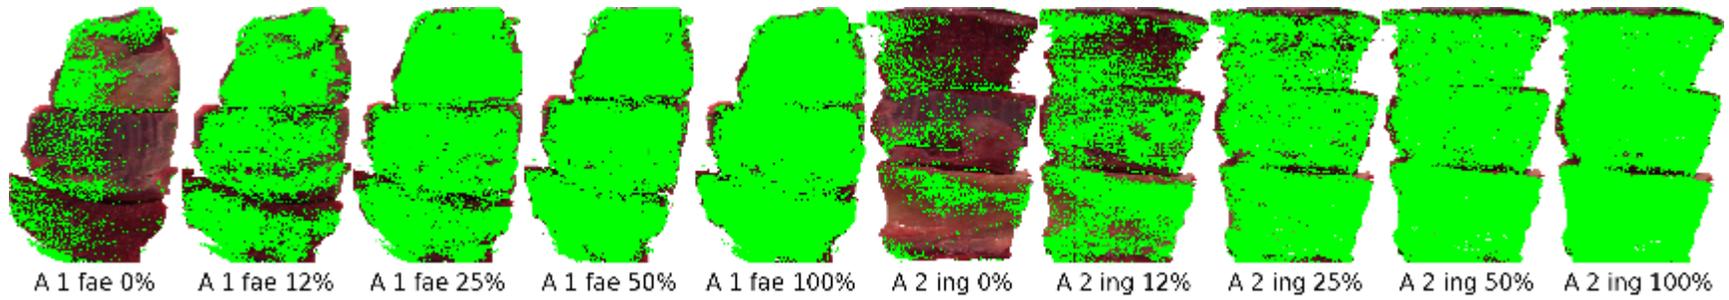
Figure 24: Beef classification results using 19 wavelength bands from camera B. Classification model was trained the highlighted camera A images from Figure 23.

6 SOFTWARE DEVELOPMENT

6.1 Overview

This section describes construction of the test rig and software development. The system consists of an aluminum frame housing the Hyperspectral camera, high intensity halogen lights, conveyor, electrical cabinet and Computer. The trial system is designed to scan samples of meat that are placed on the conveyor and display an RGB image of this scan on a screen together with an overlay mask showing areas of interest (clean meat, contaminants and background). The scanning and classification process is designed to be completely automated and only requires the operator to place meat samples onto the moving conveyor.

This report also discusses some of the physical limitations of camera A, and explains why a water proof enclosure was not constructed at this stage of the project. This decision was made with consultation and approval from AMPC.

Figure 25: Scott CT room

6.2 METHODOLOGY

6.2.1 Physical limitations of Camera A

Camera A is a state of the art hyperspectral camera which is capable of imagine a very wide spectral range (from 380nm to 2500nm). Owing to this wide spectral range we have been able to identify a set of wavelengths that are very well suited for the detection of faeces and ingesta in meat. Camera A has proven to be an exceptional laboratory instrument. However, in order to achieve such a wide spectral range, the design of the camera compromises in some of its other optical characteristics, and as a consequence this camera is poorly suited for industrial applications. For example, the focal length and aperture size of the camera are fixed because at such a large spectral range variation in these properties would result in severe chromatic aberration. In an effort to improve the sensitivity of the camera, the optics were designed with a wide aperture of f/2.4 to allow for a lot of light to reach the sensor. Consequently a wide aperture also restricts the depth of field of the camera resulting in a depth of approximately 20mm at a distance of 1 meter from the camera. This means that if there is more than 20mm variation in the thickness of a carcass, than those parts of a carcass will be out of focus.

The depth of field can be increased by moving the camera away from the carcass. To obtain a sharp image of a lamb carcass the camera would need to be 3-5 meters away for the carcass for the depth of field to be sufficiently. Reserving such a large amount of clear space at a processing plant is impractical for a trial. Setting the camera back so far would result in a pixel size of approximately 5-10mm on camera A, making it impossible to detect small contaminants. Furthermore a waterproof enclosure would require the camera to be behind a Perspex window which will also cause unwanted chromatic aberrations.

6.3 3.3 Testing and validation

6.3.1 Mechanical Requirements for the Trial Rig

Given the limitations of camera A, the hyperspectral trial rig was designed to operate in a controlled environment. We have sought and have been granted permission from JBS Brooklyn to test and demonstrate the system in their CT room (as shown in Figure 25). This room provides a clean and dry environment to accommodate the physical limitations of the camera A. As contaminants will need to be brought into the CT room, a thorough scenario testing document was provided to JBS which outlines procedures and precautions that were taken while handling meat and contaminants in a clean environment.

The scan rig was designed to enable consistent scanning of trial samples. The rig is designed to be portable and to house all of the necessary equipment. The following are a list of requirements that rig was designed for:

1. Portability: the rig was designed with wheels so that it can be easily transported and maneuvered on premises.
2. Size: the rig was designed to have a maximum width of 800mm so that it can comfortable fit though a standard door way.
3. Due to optical focal point of this camera, the camera had to be mounted at least 800mm above the inspection surface, otherwise the image would be out of focus.
4. The height of the camera needs to be adjustable.
5. High intensity halogen lights should be mounted on either end of the conveyor belt to provide constant and uniform illumination.

6. The angle and vertical position of the lights should be adjustable to control the intensity and of the light.
7. Conveyor: the conveyor needs to be speed controlled for testing purposes.
8. A photoelectric sensor will be used to trigger the camera to start and stop scanning.
9. The needs to have enough space beneath the conveyor to house an electrical cabinet (for I/O module, power supply, variable speed drive) and inspection computer.
10. Work bench for keyboard, mouse and monitor.

Figure 26 and Figure 27 show the CAD model of the scan rig that was designed to meet the above requirements.
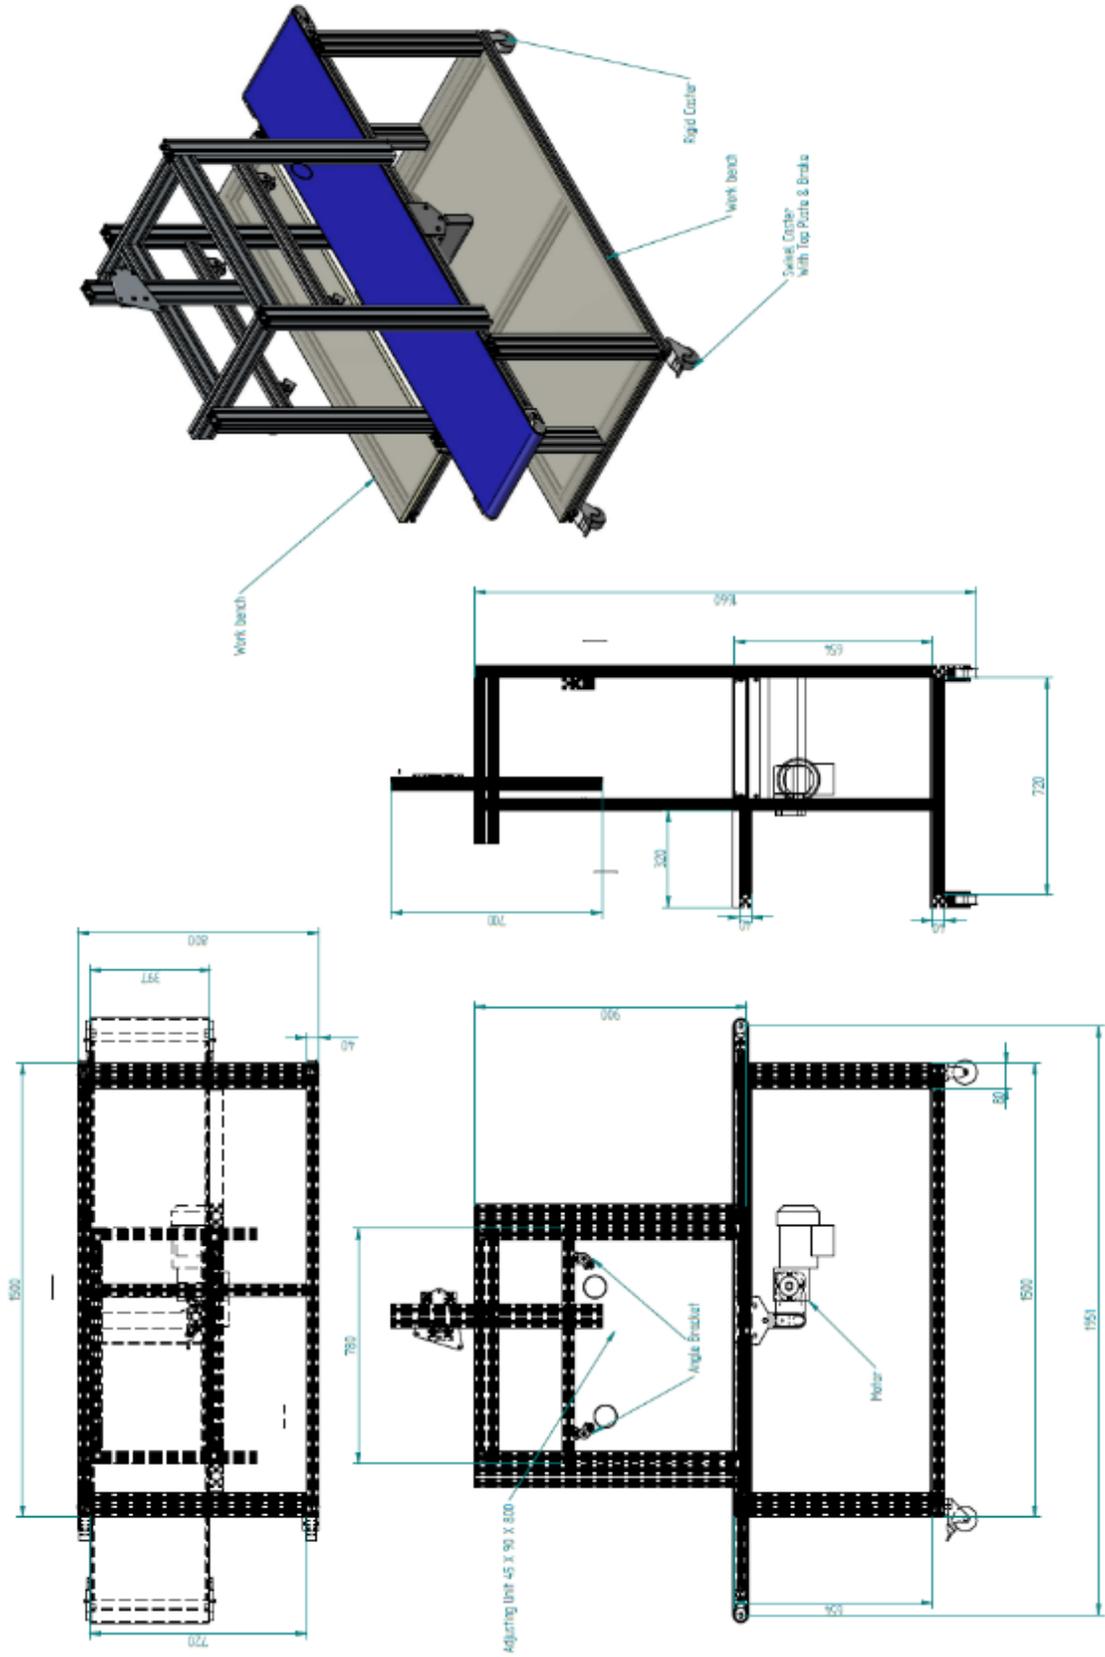
Figure 26: Projection view of the proposed scanning rig.
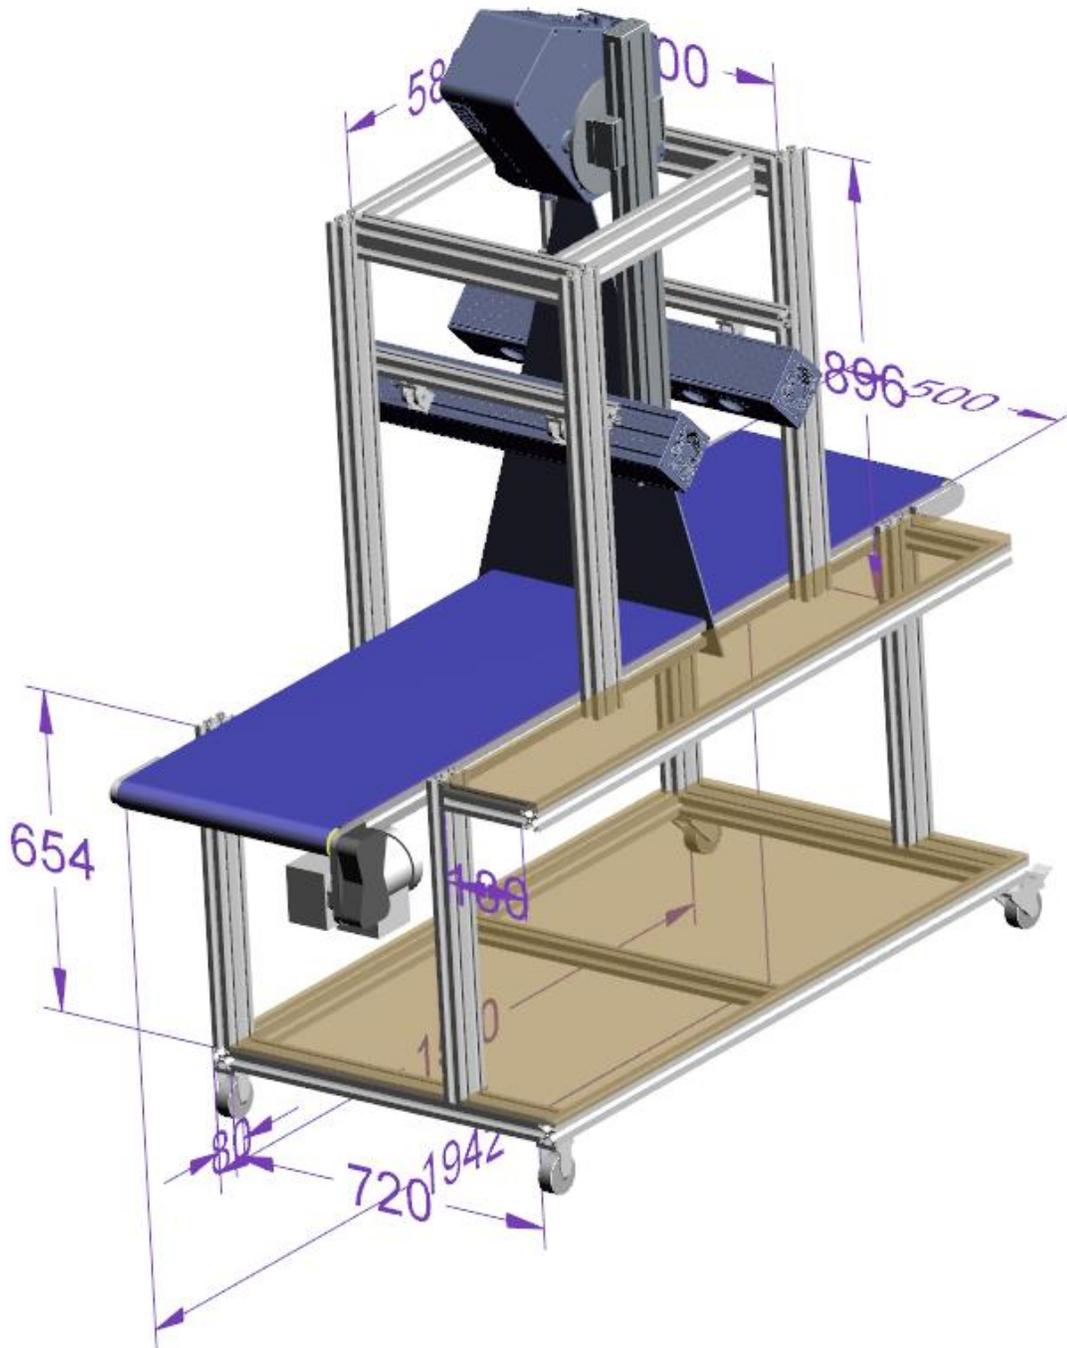
Figure 27: Three dimensional render of the proposed scanning rig.

6.4 STAGE SUMMARY

Owing to the limitations of camera A (described in the previous section), the scanning rig was designed to operate in a dry and controlled environment. Figure 28 shows the constructed scanning rig with the labels in the image identify the following:

1. Camera A.

2. High intensity halogen lights.

3. Conveyor belt.

4. Work bench.

5. Motor for conveyor belt.

6. Electrical cabinet.

7. Acquisition computer.

8. Wheels.

Due to the cameras optical limitations the camera had to be mounted 800mm above the conveyor. At this height, each pixel represents approximately 1mm wide. The high intensity halogen lights illuminate the scan line from both ends of the conveyor to provide uniform lighting at the scan line. The electrical cabinet is located beneath the conveyor and houses the variable speed drive for the conveyor. The computer is also located beneath the conveyor and is used for all of the application processing including frame grabbing from the camera, image processing and HMI control.

Once powered up, the conveyor runs at a constant speed and the camera constantly acquires data and streams it to the HMI where the data is displayed and classified in real time. As the data is streamed from the camera, the HMI shows a scrolling RGB image as the data is streamed from the camera and clearly displays the presence of contaminants as an overlay mask (as shown in Figure 29).
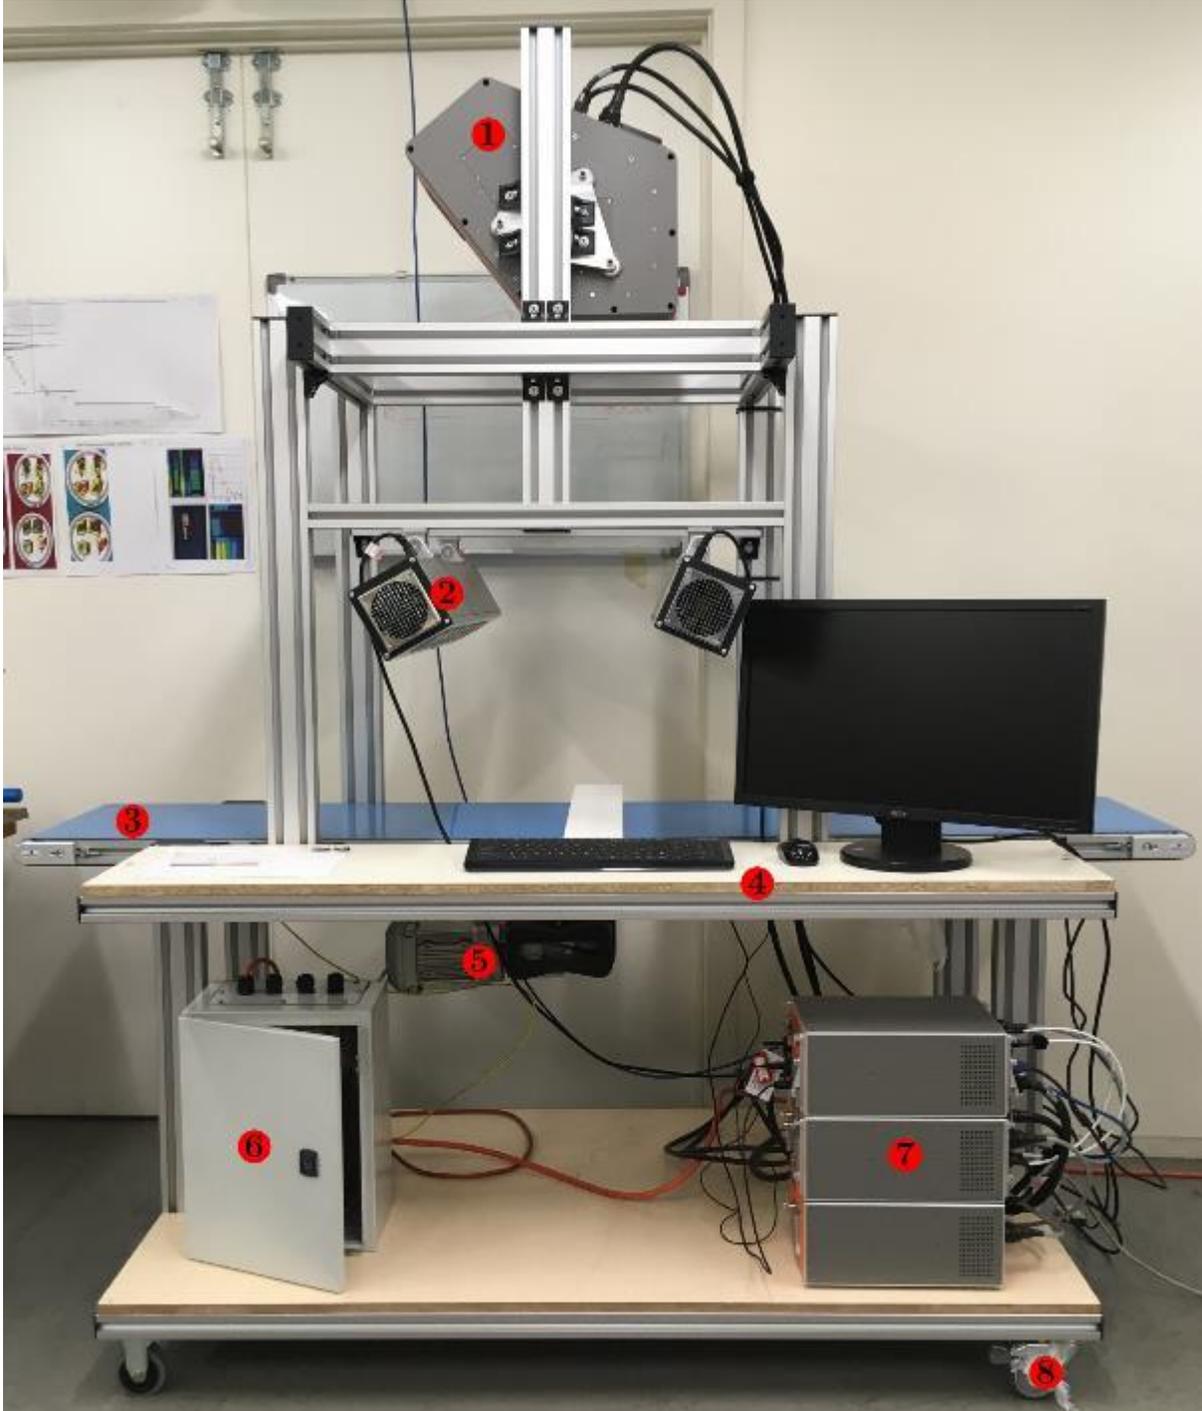
Figure 28: Photo of scanning rig.
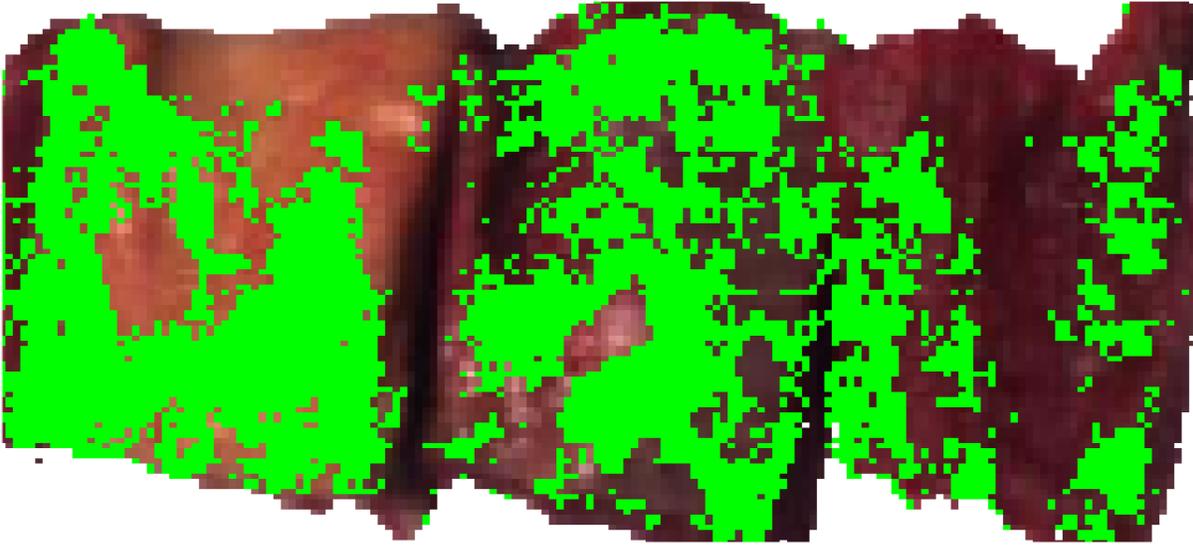
Figure 29: Example of HMI display, clearly identifying contaminants in green.

7   SOFTWARE CONTROLLED TRIAL

7.1   Overview

In this section the trial hyperspectral system is tested in a controlled environment. The goal of this trial was to test our data acquisition and classification software, and to demonstrate that acquisition, display and classification can be performed in real time. Unlike the preliminary trial which used the standard data capture that captured all available hyperspectral data, this trial uses specially written acquisition software that captures only four specific wavelength, classifies, displays and saves the data in real. A graphical user interface was also developed to displays a scrolling image of meat moving down a conveyer and to highlight any detected contaminants. This trial confirms that a low complexity model, only using 4 wavelengths are suitable for real time contaminant detection. This trial was performed in the Scott CT room at the JBS abattoirs in Brooklyn (shown in Figure 30).

Figure 30: Hyperspectral contaminant detection system in Scott CT room at JBS Brooklyn.

7.2 Methodology

This section briefly describes the hyperspectral acquisition program, contaminant detection program, and describes the samples set and contamination process.

7.2.1 Acquisition program

To develop the acquisition program a C++ software development kit was purchased which provided full application programming interface of hyperspectral cameras. This was used to develop a custom acquisition server that streams several specific wavelength channels from the camera, over a TCP socket, to a client program running the contaminant detection program. The server was designed to accept a range of messages to start and stop acquisition, and to adjust a range of camera parameters.

7.2.2 Contaminant detection program

A contaminant detection program is made up of 3 parts:

1. A client socket connection to send/receive messages to/from the acquisition server.
2. A contaminant detection model, that has been trained to classify pixels of in image into three categories: clean, contaminated and background.
3. A Graphical User Interface GUI to display a scrolling live feed from the acquisition server and highlight any contaminants in real time. This GUI also allows users to adjust camera parameters, capture white/dark reference images and record the raw image data.

A screen shot of the contaminant detection GUI is shown in Figure 31. This screenshot is of a scrolling image (scrolling right to left) of a contaminated meat sample as it moved down the conveyor of the test rig. Mid-way though the scan the "Contaminant mask" checkbox was enabled to display a green overlay to highlight any detected contaminants.

7.2.3 Meat and contaminant samples

Lamb samples were collected from the boning room floor. A total of 8 pieces of lamb were used for this experiment with approximately equal amounts of lean and fat. The samples fit onto two large trays as shown in Figure 32. Faeces samples were collected from the stockyard and ingest was collected from offal room.
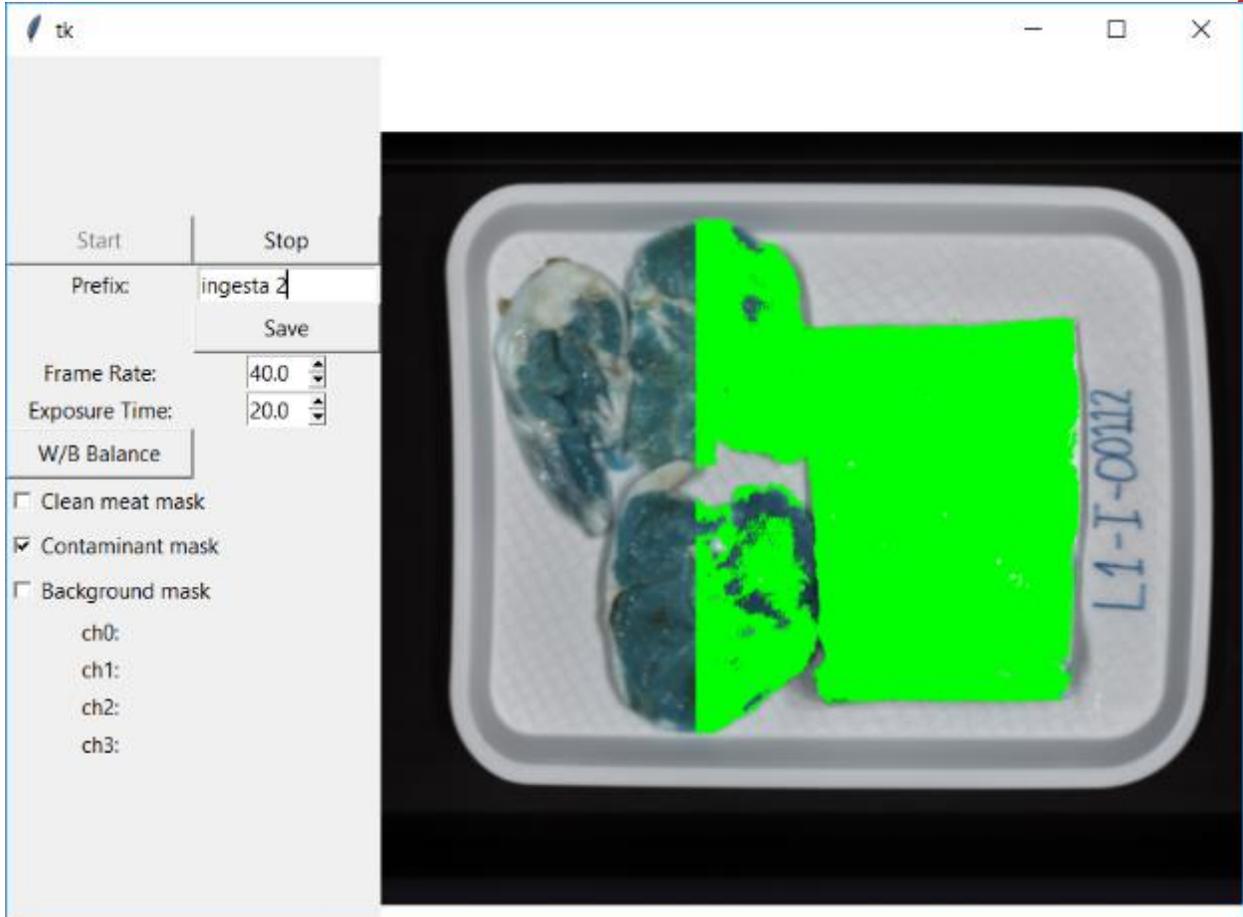
Figure 31: Screen shot of GUI showing real time contaminant detection. Contaminant mask was enabled mid-way through the scan to illustrate the mask overlay.

7.2.4 Contamination

The two trays of clean meat samples were contaminated with faeces and ingesta (respectively) in small incremental amounts. First a control sample of clean meat was scanned as shown in the two left images of Figure 32, then the meat was incrementally contaminated by painting a small amount of contaminant onto the meat. The first application of contaminant was very minimal, and almost indistinguishable from clean meat. All up there were 7 levels of contamination for both faeces and ingesta, and each contaminated sample was scanned at two orientations, for a total of 28 tray sized images. Color photos of faeces and ingesta images for one orientation are shown in Figure 32 and the corresponding hyperspectral scanned are shown in Figure 33. Contamination level is indicated by a number ranging from 0 to 7 representing no contamination to maximum contamination. The hyperspectral images have false colors as the four wavelength channels in these images do not correspond to red, blue or green wavelengths.
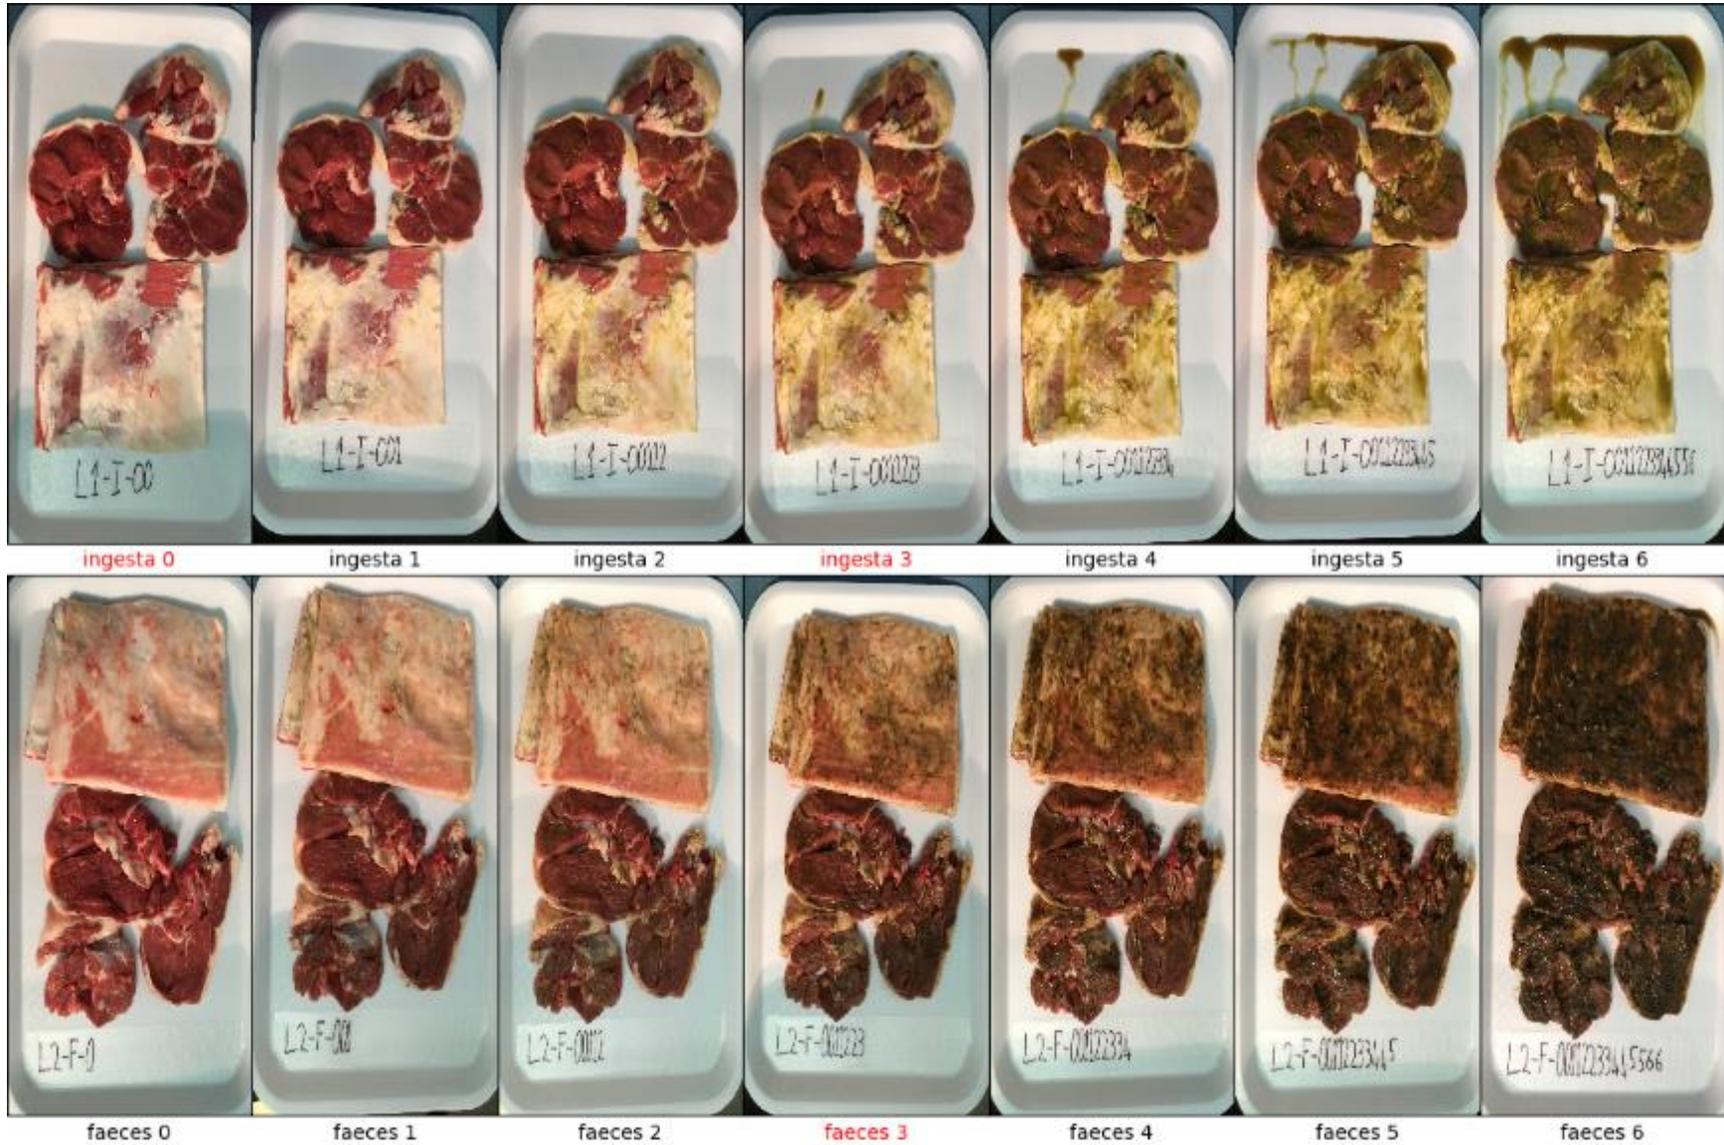
Figure 32: Colour photos of the ingesta (top) and faeces (bottom) contamination samples.
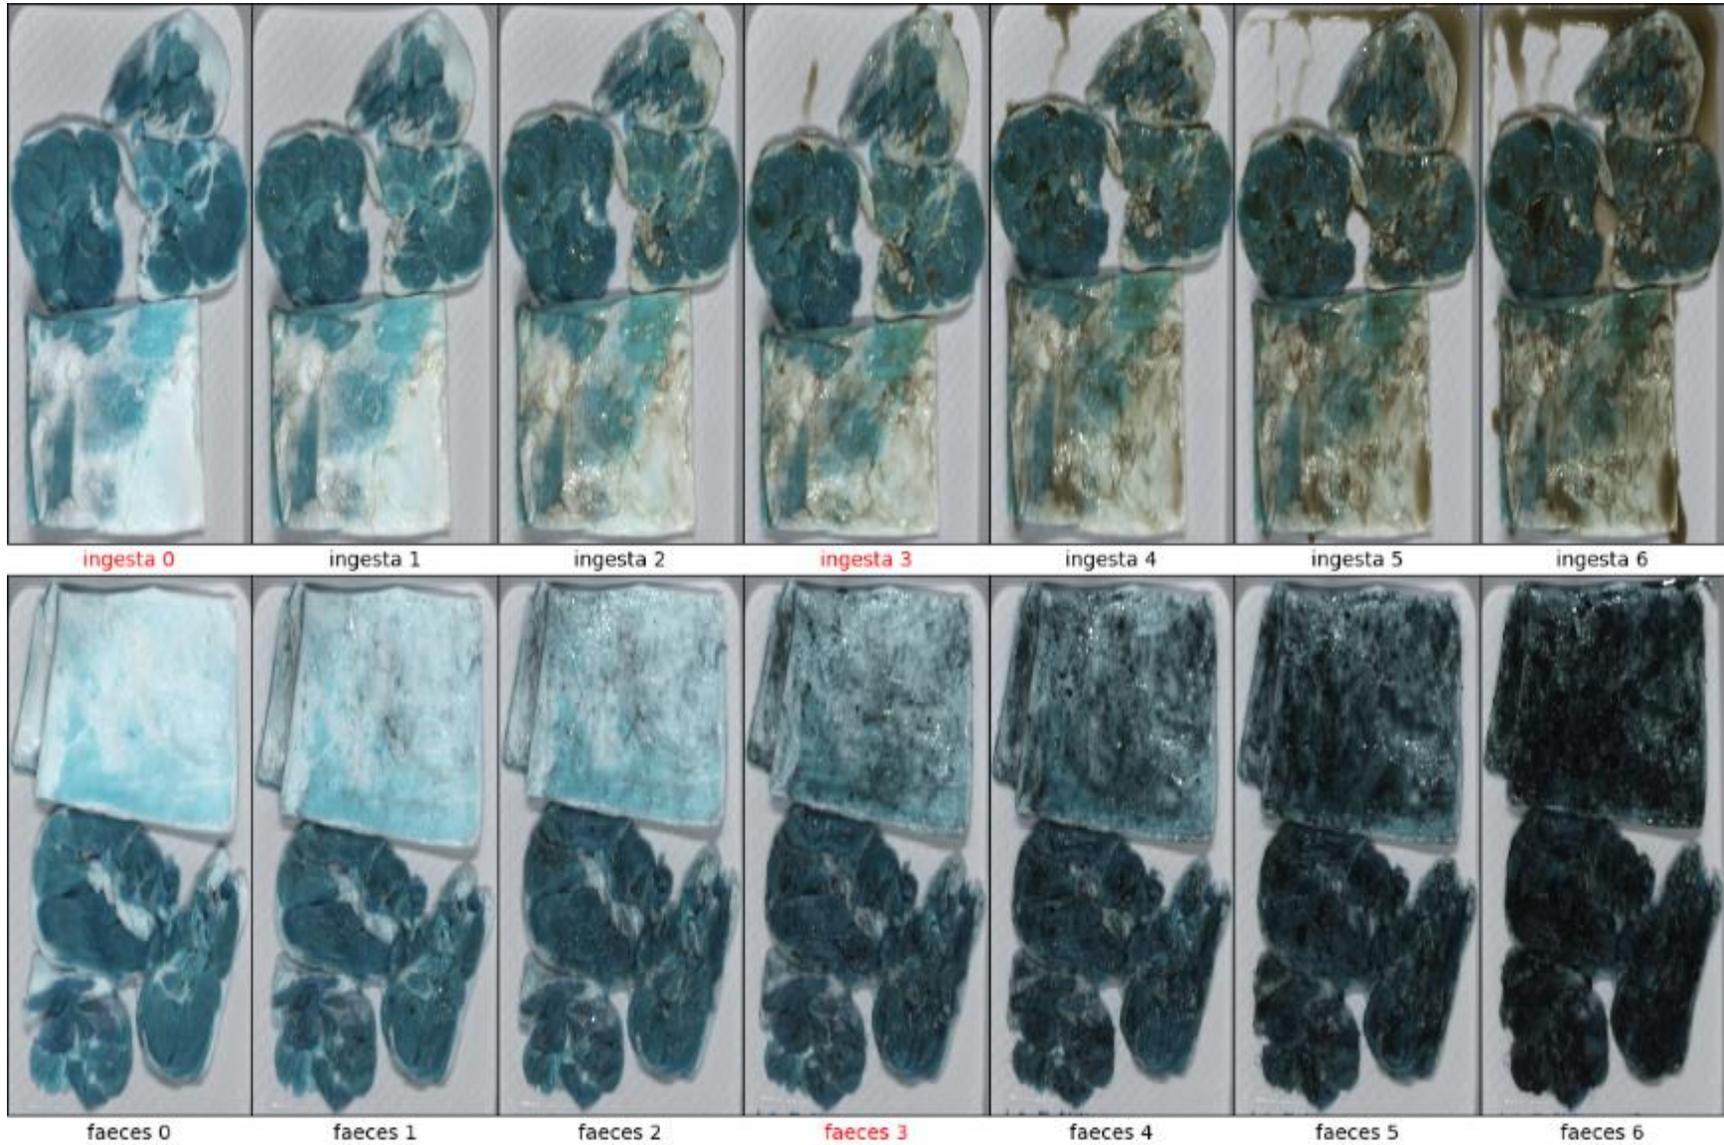
Figure 33: False color images of of the ingesta (top) and faeces (bottom) contamination samples.

7.3   STAGE SUMMARY

In this section the performance of the contaminant classifier is evaluated on a small subset of the hyperspectral images.

To illustrate how well the prediction algorithm works on new data (data that was not used in the training stage), the model is only trained on 20% of the pixels of the following 3 images:

- ingesta 0: Clean lamb
- ingesta 3: moderate ingesta contamination
- faeces 3: moderate faeces contamination

These images and corresponding results are labeled with red text in Figures Figure 32, Figure 33 and Figure 34 and in Table 7.

Moderately contaminated images were chosen for training to demonstrate how well the algorithm generalizes to images of lower or greater contamination. Additionally, by training on only 20% of the pixels in only 3 images, we again show that the algorithm generalizes well because it shows that the models does not over fit the training data and does indeed work well on data that the algorithm has never seen before.

In the contaminant detection results shown in Figure 34 the false RGB images are overlaid with white for areas classified as background and green for areas classified as contaminant. Areas classified as meat remain unchanged so as to clearly see the meat part of the image.

To get an objective measure of performance, we show a confusion matrix for the classification results of every image in Table 3. A confusion matrix show the percentage that a row category is classified as all other categories.
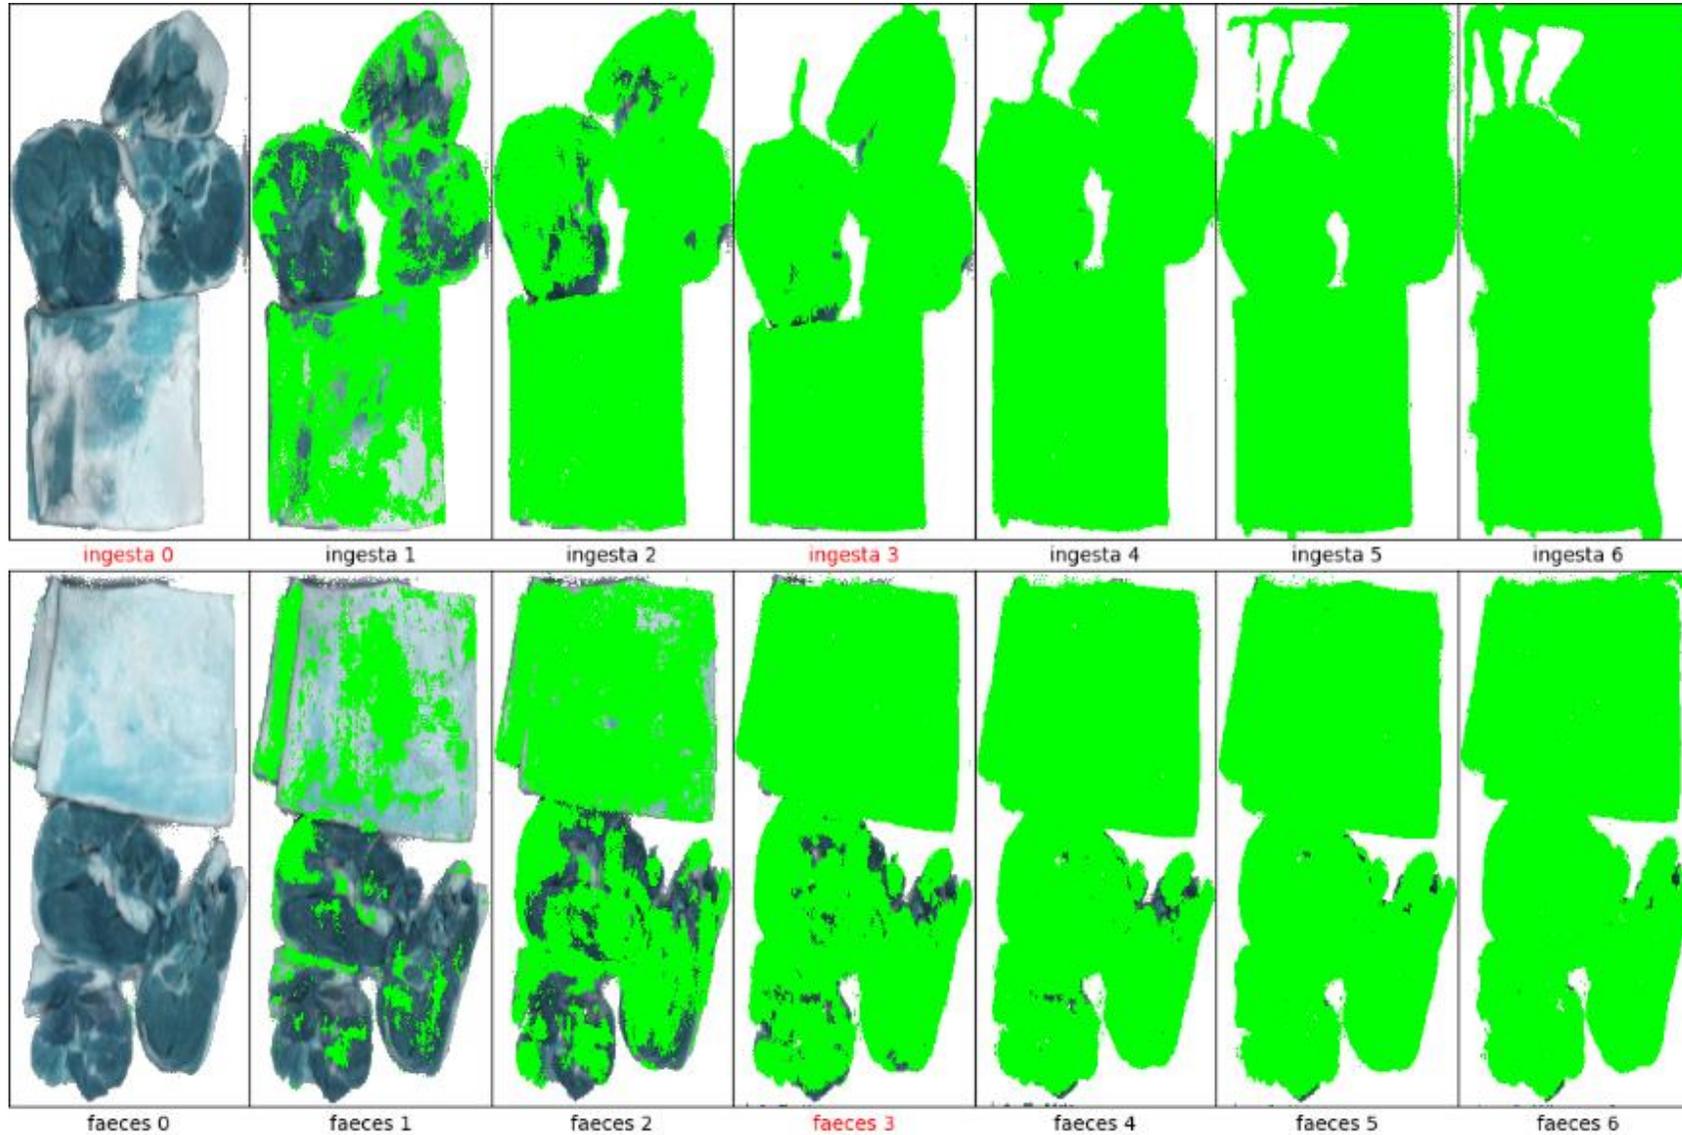
Figure 34: Contaminant detection results. Contaminants are highlighted with a green overlay.

			Predicted																				
			0			1			2			3			4			5			6		
			clean	contaminated	background	Clean	Contaminated	Background	Clean	Contaminated	Background	Clean	Contaminated	Background	Clean	Contaminated	Background	Clean	Contaminated	Background	Clean	Contaminated	Background
Actual	Ingesta	Clean	100	0.0	0.0	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
		Contaminated	-	-	-	42.8	56.9	0.3	5.5	94.3	0.1	0.8	99.2	0.0	0.0	100	0.0	0.0	100	0.0	0.0	100	0.0
		Background	2.8	0.1	97.2	0.8	0.3	98.8	0.3	0.4	99.3	0.1	0.9	98.9	0.0	1.0	99.0	0.0	1.5	98.5	0.0	1.8	98.2
	Faeces	Clean	100	0.0	0.0	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
		Contaminated	-	-	-	69.7	30.3	0.0	21.9	78.1	0.0	4.5	95.4	0.1	1.3	98.6	0.1	0.3	99.7	0.0	0.2	99.7	0.1
		Background	5.2	0.2	94.6	2.6	0.2	97.2	1.6	0.6	97.8	0.6	0.4	99.1	0.4	1.0	98.5	0.0	0.3	99.6	0.1	0.2	99.7

Table 7: Confusion matrix of all sample image.

The results clearly show that the contaminant detection system is capable of detecting both faeces and ingesta contaminants at a wide range of contamination levels.

The confusion matrix in Table 7 shows that very low levels of contaminants (images ingesta 1 and faeces 1) are also detected but at a reduced rate of approximately 30-57% per pixel. There are a number of plausible reasons why the detection rate of these images is so low. First, contaminants were manually applied with a paint brush, and with such low contamination levels, areas of the meat may not have been sufficiently contaminated. This is visible in images "faeces 3" (a moderately contaminated sample) in Figures Figure 32 and Figure 33 where uncontaminated regions are clearly visible. Secondly, the faeces collected from the stockyard was quite dry and difficult to apply onto the sample, this may have led to areas that were poorly or not-at-all contaminated and consequently the low detection rate of 30 in "faeces 1" image. Finally, the ingesta samples were extremely watery and were invisible to the naked eye when applied at very low concentration, thus possibly resulting in uncontaminated regions or undetectable levels of contamination.

8   HYPERSPECTRAL RIG DESIGN

8.1   Overview

This section describes the mechanical concept design of the hyperspectral contaminant detection rig suitable for onsite abattoir testing. The rig is an adjustable supporting structure for a camera enclosure and two light towers. The system was designed to be portable and flexible. The height and angle of the line scan camera enclosure is adjustable and is suitable for scanning upper or lower areas of beef/lamb carcass. The angle and horizontal position of the light towers is also adjustable. The light towers are positioned on either side of the camera's scanline and are large enough to illuminate up to 1200mm of carcass. This section also describes the line scan camera chosen for this enclosure and presents some preliminary results.

8.2   Methodology

This section briefly describes the concept design of the camera, enclosure, and the rig.

8.2.1   Camera Selection

Initially, it was planned to replace the hyperspectral camera with a 4 camera system, which would use 4 ordinary grayscale line scan cameras with optical filters placed over each lens to capture the four wavelengths of interest. This option was attractive because it would significantly reduce the costs, increase the resolution and allow for more lens options. However, after purchasing the cameras and commencing programming, this option was identified as an unacceptable risk due to the time required to write custom software to correct for the parallax error (to align the 4 images due their different positions). The difficulty was that parallax is usually corrected by using identical cameras that view the image in the same set of wavelengths, however, in this case each camera had a different optical filter on it, which makes it difficult to find similar features in the 4 images (a stage required for parallax correction).

To reduce risk, camera B is an industrial grade line scan camera and was specifically purchased for this application.

Camera B provides the following advantages over camera A that was used during the analysis phase:

1. It is designed for industrial applications.
2. It has a spatial resolution of 1024 pixels, which is 3 times greater than camera A (used for analysis in earlier sections).
3. It uses the same Software Development Kit (SDK) as our existing camera, so the existing code can be used with only small modifications.
4. It has a GigE interface, so specialized capture card is not require.
5. The camera came with a number of lens options, allowing us to choose the most suitable lens for this application.
6. It has a spectral range from below1000 nm, which includes the 4 wavelengths that we identified in earlier sections.

8.2.2 Camera Enclosure

The camera enclosure is a water tight enclosure suitable for abattoir wash down. The enclosure is a standard cabinet that has been modified to allow for a camera mount and plastic window as shown in Figure 35 . The camera mount is designed to allow for precise adjustments to the camera position and orientation shown in Figure 36.

The plastic window is an anti-reflection high transmission plastic window. It was chosen because it has an optical transmission greater than 97% over the wavelengths of interest and it is a shatter proof material suitable for use in food production.

The enclosure has a door at the top to allow easy access to the camera.

Figure 35: Camera enclosure showing hinged door and plastic window.

Figure 36: Inside view of camera enclosure showing mounting brackets.

8.2.3 Halogen Lights

Illumination is provided by 60 of 35W halogen lights shown in Figure 37. The lights have a beam angle of 10 degrees to focus the light onto the area of interest, this small angle focuses 10 times more light onto the subject compared to our laboratory 35 degree lighting system. High intensity lighting is required to reduce the exposure time on the camera and therefore allowing for faster frame rates.

Figure 37: 35W 10 degree 12V halogen lamp.

Halogen lights are required for this application (as opposed to LED, incandescent or HID lights) because they produce fairly uniformly distributed and broad spectrum light as shown in Figure 38.

The use of halogen lighting produces significant design challenges due to the large amount of heat they produce. All together these lights produce approximately 2100 Watts of heat. This heat must be dissipated efficiently to extend the life of the lamps and to make the enclosure safe to handle.

These lights have an intensity half-life of 5000 hours. In the future we intend to upgrade the lighting to a hyperspectral LED light source to extend life span, power and potentially cost, however this technology is in its infancy and the lights are not sufficiently intense at this point in time.
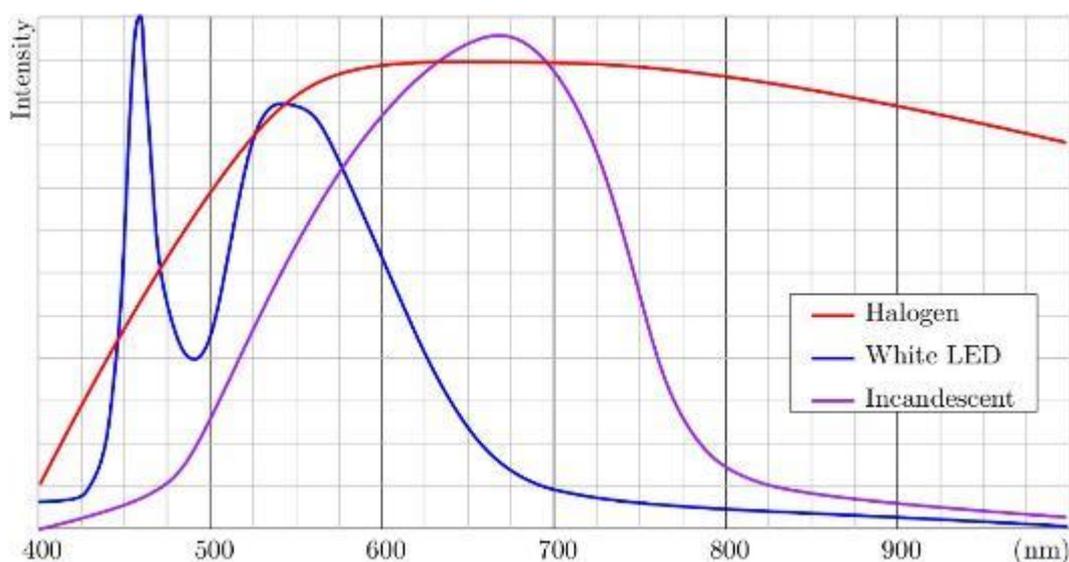
Figure 38: Spectral intensity of Halogen, LED and Incandescent lights.

8.2.4 Light Towers and Frame

The light towers are positioned on either side of the camera lens as shown in Figure 39, focusing intense halogen light on the area of interest. The towers are designed so that the height, horizontal angle and position are adjustable. A flexible design was chosen to allow for the contaminant detection rig to be used no matter the position of the carcass in an abattoir.
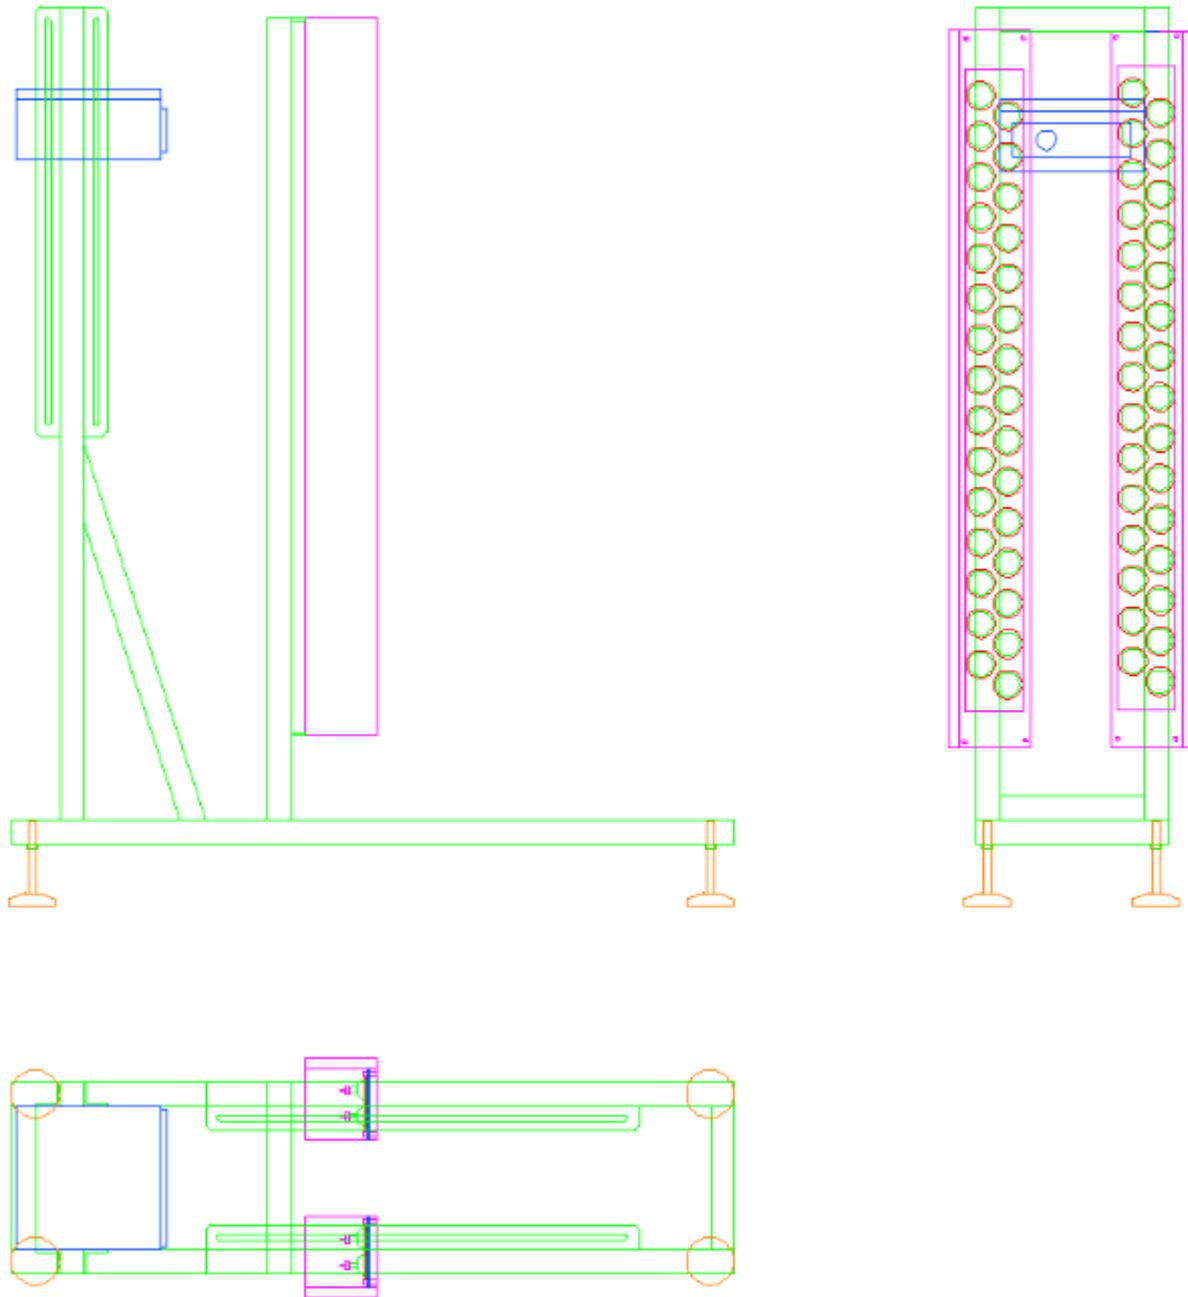
Figure 39: Hyperspectral contaminant detection rig showing frame (green), camera (blue) and light towers (magenta).

The light towers are made of a stainless steel enclosure with a tempered glass window for the light source. Tempered glass was chosen because anti-reflection high transmission plastic cannot be

manufactured in the size required for this application. Guards have been attached to the towers to protect the glass from any physical contact and to protect operators from intense halogen light.
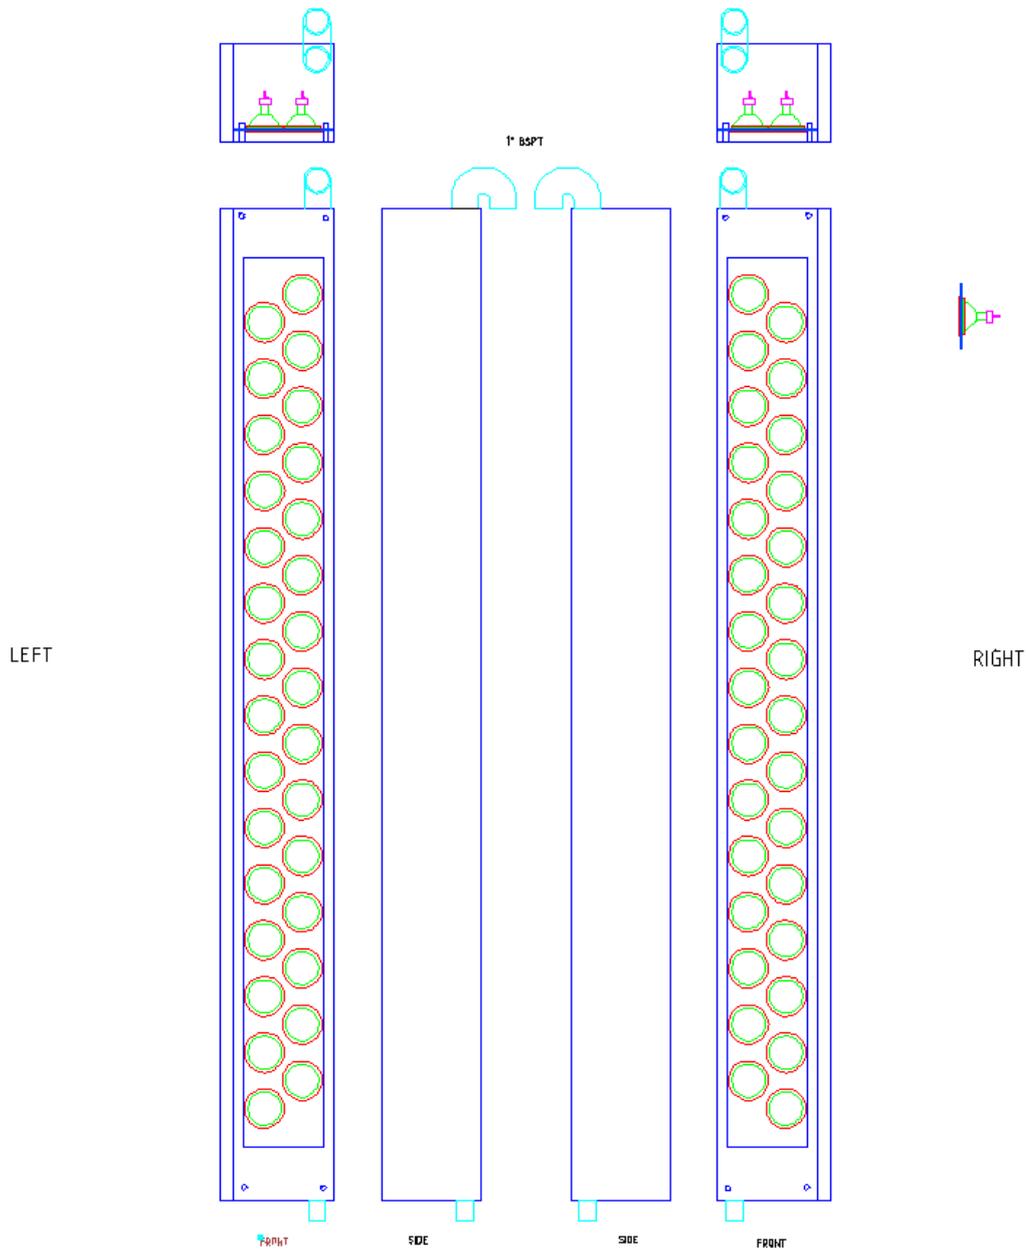
Figure 40: Drawing of light towers.

As mentioned earlier, cooling the lights in the enclosure is particularly difficult if the system is to be portable and self-contained. To meet these requirements, the enclosure was designed with open air vents. The vents (shown cyan color in Figure 40) are pointing down to minimize the possibility of material entering the enclosure. Compressed was pumped into the bottom of the enclosures to circulate air through the towers, replacing all of the air in the tower at least once per second. This cooling extends the life of the halogen lights and keeps the enclosure/glass cool for handling.

Since this is a portable system and the light towers are not 100% water tight, the system will only be operated under the supervision of a Scott engineer. Furthermore, a procedure has been put in place to clean the system after use and cover appropriately for wash down.

8.2.5 Power and Computer Interface

Power supplies for the lights and camera are mounted in a stainless steel enclosure at the bottom for the frame (Figure 41). The camera is connected to a laptop via a Gig-E Ethernet connection. All image acquisition and processing was performed on a standard laptop.

Figure 41: Electrical cabinet showing 12 VDC, 3 kW power supply, distribution blocks and safety switches.

8.3 Summary

This contaminant detection system was designed to be portable and temporarily stationed on site. For this reason, the designee was focused on keeping the system as self-contained and compact as possible. To meet this requirements we had to design the light enclosures with open air vents. To reduce the chance of materials falling into the enclosure, the air vents were designed to point down. To meet the wash down requirements, a strict shutdown procedure was put in place that involved manually washing and covering the lights towers with a protective cover.

9 INHOUSE WHOLE CARCASS DATA COLLECTION AND ANALYSIS

9.1 Overview

This section describes how the hyperspectral contaminant detection system was tested and retrained on lamb carcasses in preparation for abattoir trials. In this in-house trial, the focus was to test the newly built contaminant detection rig, to simulate abattoir conditions and to collect more accurate data.

In this test it was decided to focus only on lamb/mutton because the processor had kindly granted permission for us to perform the onsite trials at their abattoir. As such, four rear-halves of lamb were purchased, and contaminants collected from the processing plant. To simulate an overhead conveyor, the lamb carcasses we suspended from a robot arm and moved past the contaminant detection rig for data acquisition. As in all prior tests, the meat was incrementally contaminated with faeces and ingesta, the data was captured and used to train a new classification model.

9.2 Methodology

9.2.1 Experiment setup

The in-house trial was performed at the Scott warehouse in Melbourne and consisted of suspending lamb carcasses from a pelletizing robot (Figure 42) and scanning them with the hyperspectral system as they moves past the camera. The linear velocity of the robot was set 350mm/sec – the approximate speed of the overhead conveyor at the processing plant. The camera was positioned at a distance of approximately 1.5 m from the carcass which gave a field of view of approximately 1000 mm. At this distance the vertical height of each pixel was approximately 1 mm. To achieve an equal aspect ratio, the frame rate of the linescan camera was set to 350 fps, hence giving a horizontal pixel width of 1 mm.

Figure 42: Photo showing a carcass suspended from a pelletizing robot, moving past the hyperspectral rig to simulate the movement of an abattoir overhead conveyor.

9.2.2 Data Collection

To diversify the data set, both sides of the carcasses were scanned at 4 levels of contamination (clean, light, moderate and heavy). Each scan was also repeated 5 times to test for repeatability, therefore a total of 160 images were collected for this trial.

Figure 43 shows false RGB images of the four carcasses. Each row shows front and back images of one carcasses contaminated with increasing amounts of contaminant. These images appear in a different colour compared to earlier results (Figure 33 for example), because it was decided to change the order of the wavelengths used to create the false RGB images and make them appear as close to natural colour as possible. In this new representation, contaminants should appear brown/maroon colour.

It should be noted that images labeled as "ingesta 1" look almost completely clean. This is attributed to the fact that the ingesta obtained for this experiment was extremely watery and unlike any of the prior samples. If left for a few second in the jar the ingesta would separate into layers of what appears to be undigested grass, seeds and water. Earlier samples of ingesta had a very similar spectra to faeces.
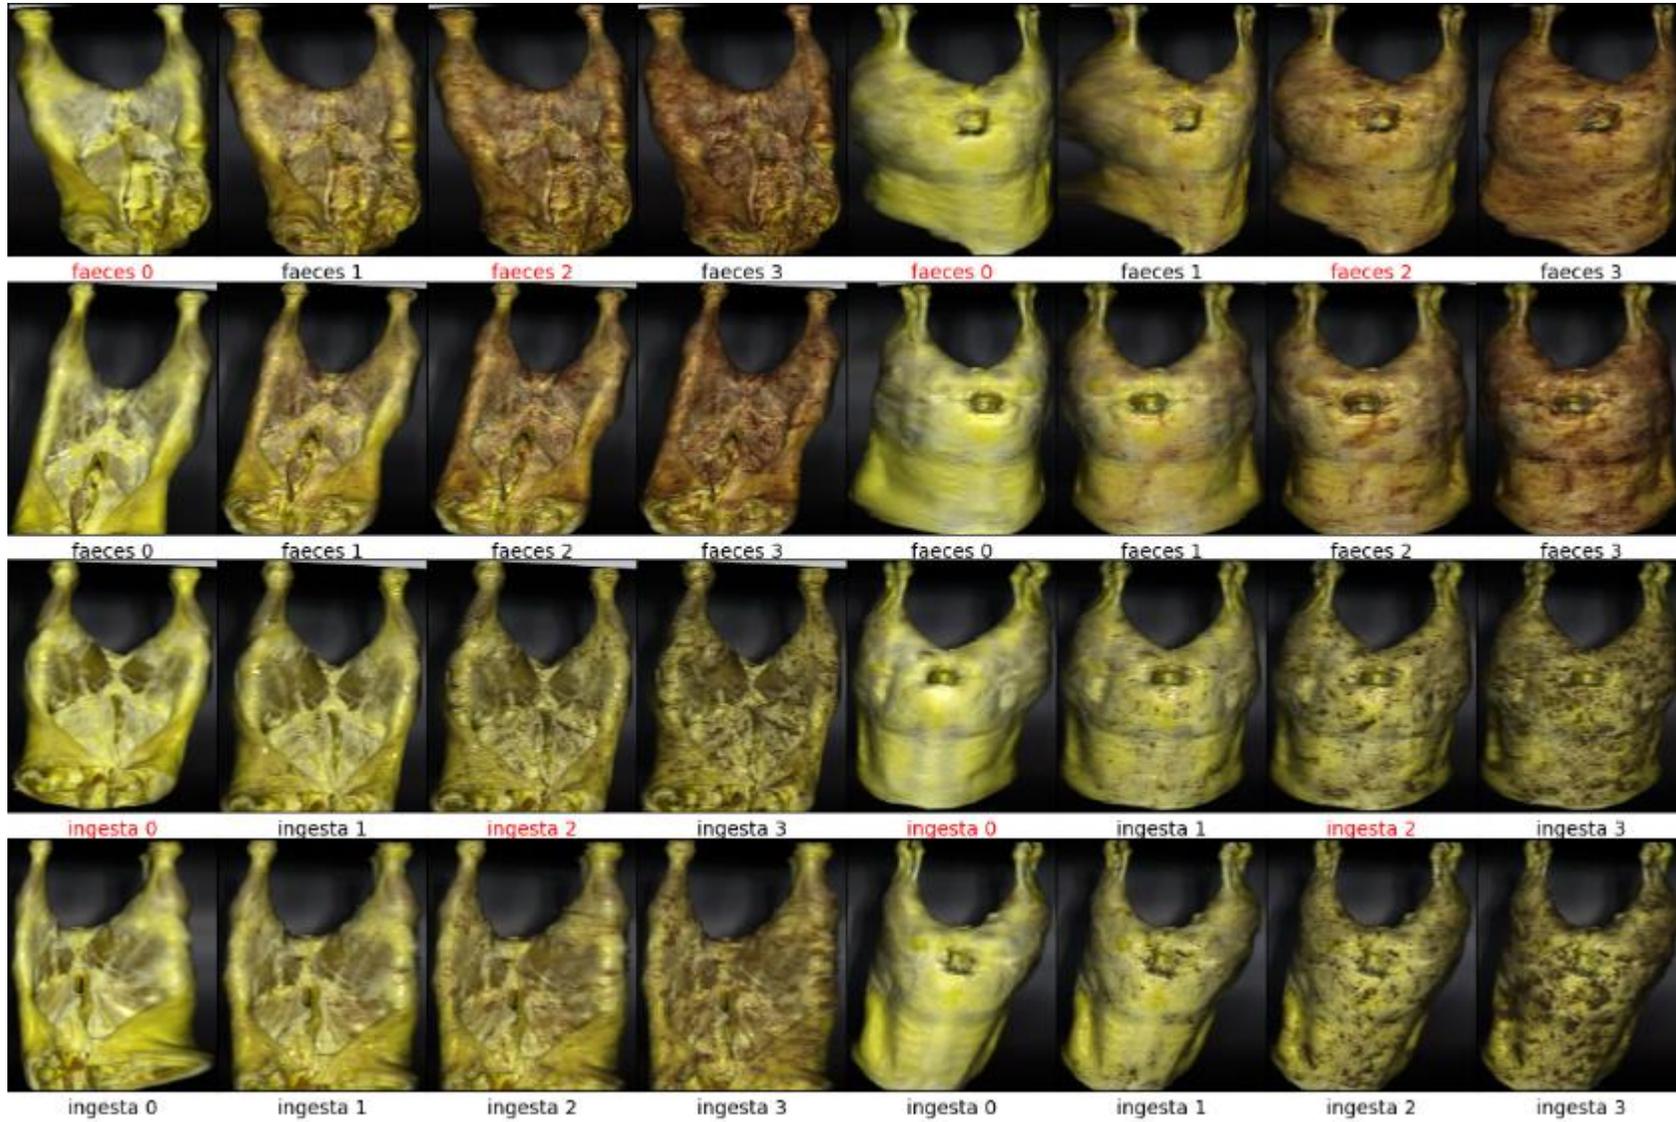
Figure 43: Front and back hyperspectral images of four carcasses of varying levels of contamination. Contaminant type and level is labeled with a caption beneath each image. Red captions represent images that were used for model training, with all other images used for validation.

During the data collection process the previously trained model was used to display contaminants on the HMI as the robot moved the carcass past the camera. Figure 44 shows a phot of the HMI while a carcasses is being scanned and displayed in real time. In this photo, all of the masks have been enabled in the HMI, and so clean meat is shown in red, contaminants are green and the background is white. Given that the model used in this example was trained in a highly controlled environment on a rather small set of data, it was surprising and reassuring that the classifier worked as well as it did in this photo.
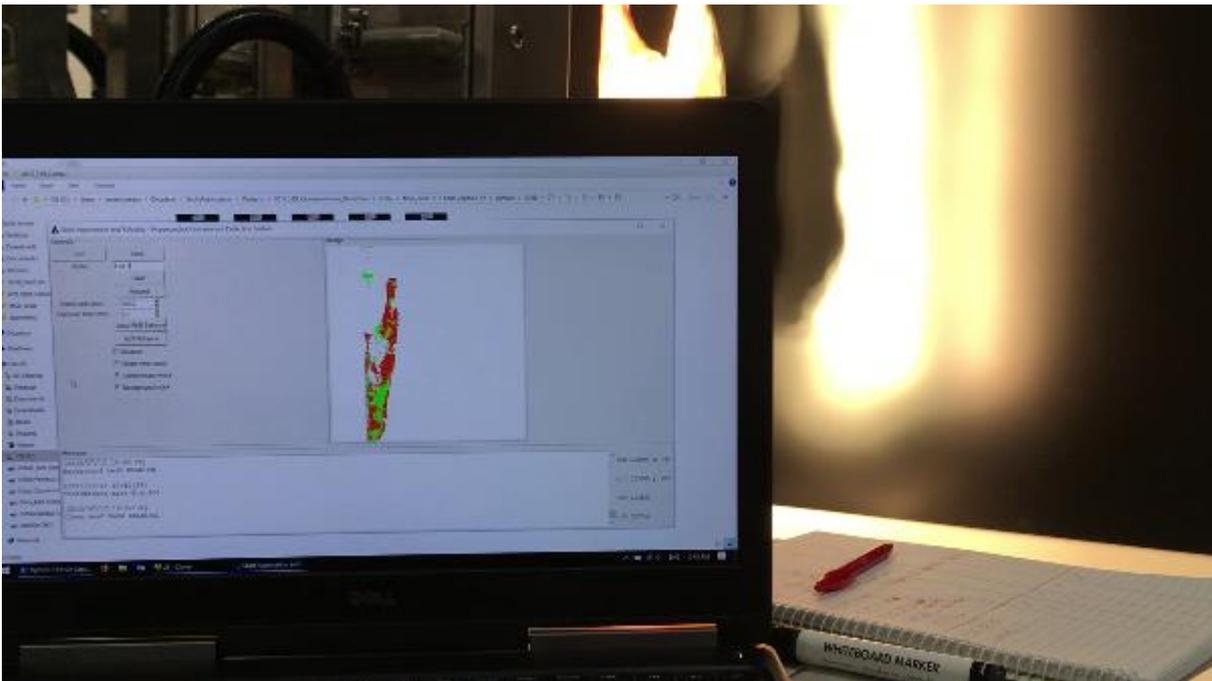
Figure 44: Photo of the carcass being classified in real time as the pelletizing robot moves it past the camera.

9.2.3 Labeling the data

To train the classification model, the data was labeled. The labels we generated as before, image masks was manually drawn for every clean and contaminated image. Since the images were uniformly contaminated, this labeling process involved drawing a silhouette of the carcass. The labeled pixels were then used to retrain the classification model. An example of a "clean" image mask is shown in Figure 45. The white part of the image indicate the area of interest (clean meat in this case).
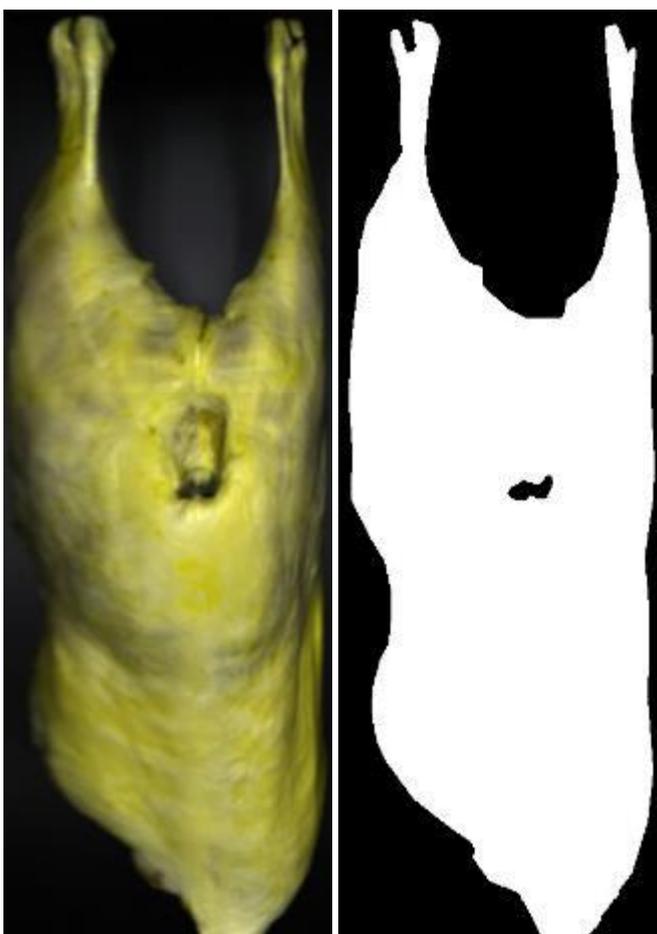
Figure 45: An example of a "clean" image mask.

9.3 Retraining the Classification Model

Of the 32 labelled images shown in Figure 43, only 8 have been used during the training process. These 8 images are labelled with a red caption. Of these training images only 30% of the pixels were used as part of the training set and all remaining data was used for validation. The data was split in this manner to demonstrate that the classifier generalises well to previously unseen images and contamination levels. Furthermore, the use of mode data only marginally improves the classification accuracy on this data set. The classification model that was developed in earlier stages was retrained on the new data - the topology of the model was not modified at all.

9.4 Classification Results

The classification results of the in-house trial are shown in Figure 46. These results show that faeces contamination is detected quire accurately with an average accuracy of approximately 92%. However, as mentioned earlier the ingesta sample was very watery and this resulted in poor classification on images with low contamination, achieving an average classification accuracy of 65% across the 16 ingesta images.

Despite the unusually watery sample of ingesta these results show that all levels of faeces and moderate to high levels of ingesta can be reliably detected. This retrained model was later used to test the classification algorithm at the processing plant as described in the next section.
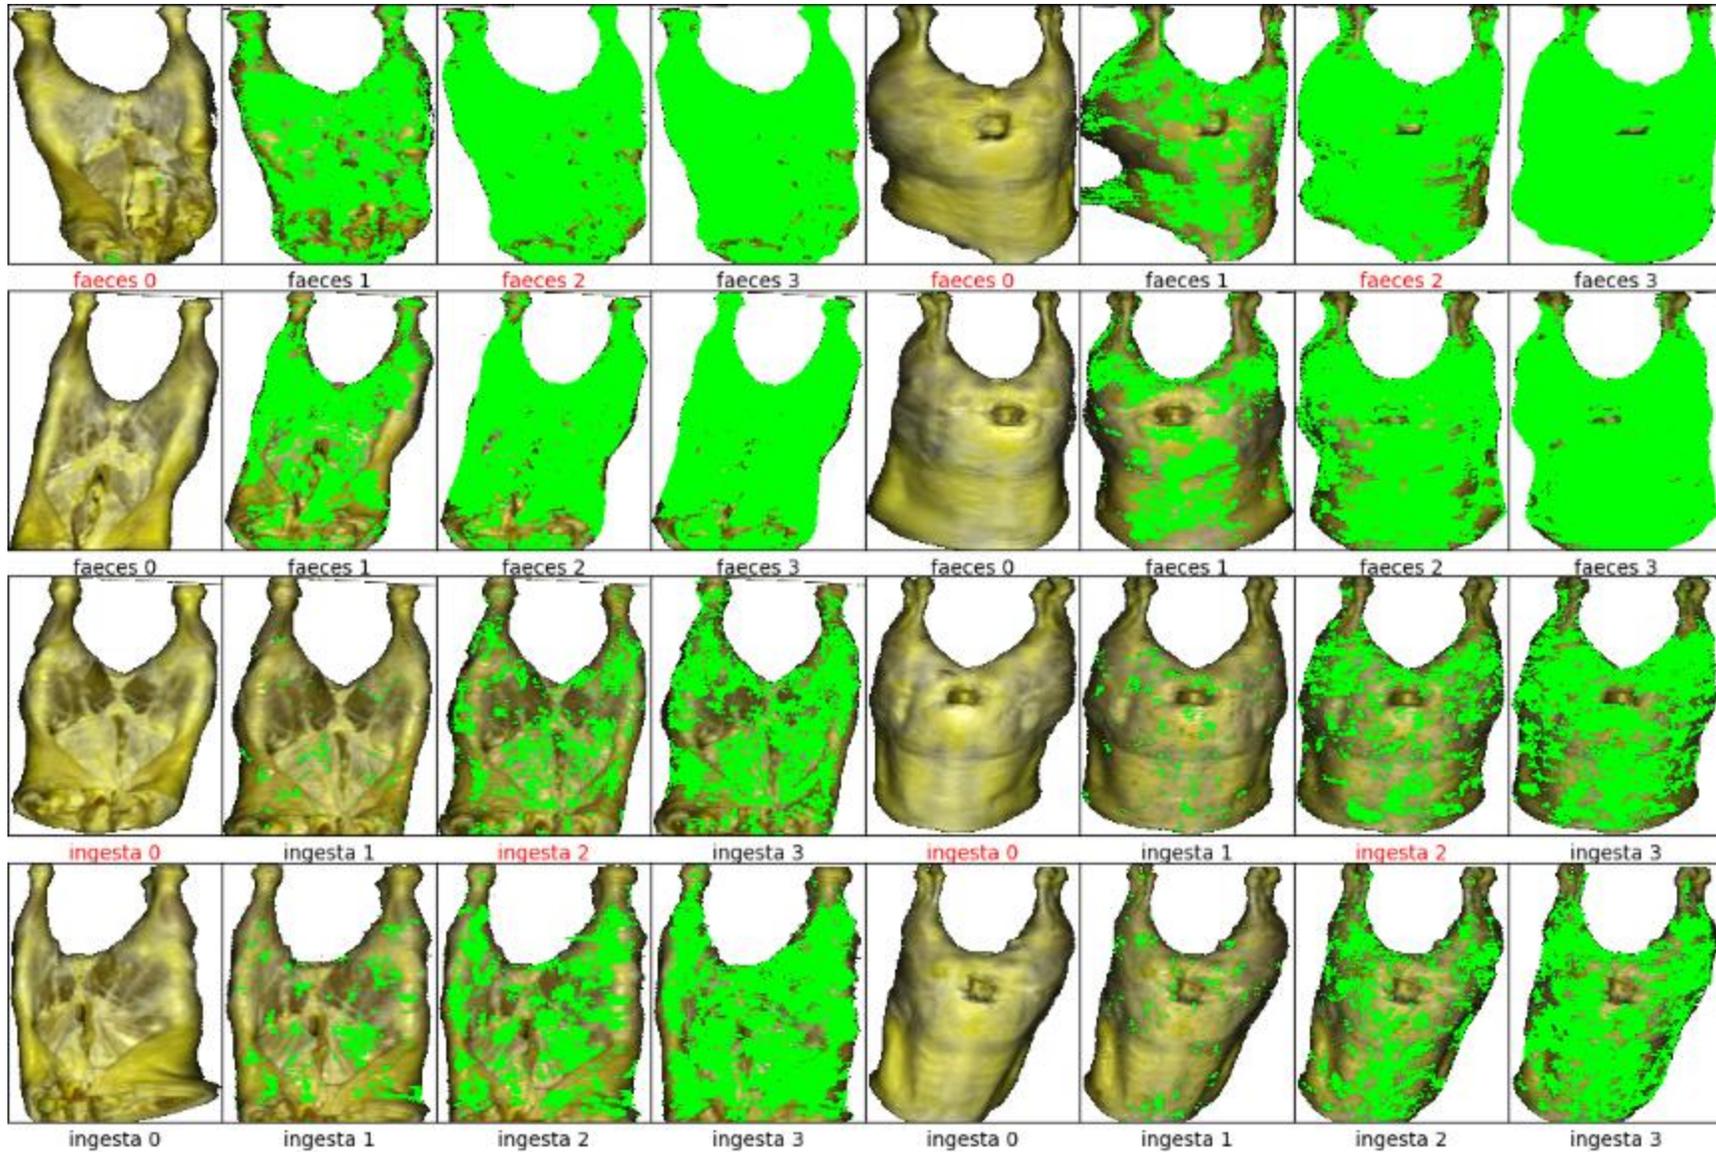
Figure 46: Contaminant classification results of in-house half carcass trials. Green mask shows the detected contaminants. Original images are shown in Figure 43

10 ONSITE PRODUCTION TRIALS

10.1 Overview

In this section the Scott Hyperspectral Contaminant Detection system was refined and tested onsite for two weeks at a lamb/mutton kill floor. The contaminant classification model was refined using data collected from an in house trial on lamb carcasses, then the hyperspectral rig was set up at a processing plant in Victoria where data was collected for two weeks and the system was demonstrated to acquire, classify and record data in real time. A Scott Engineer was present at all times to ensure safe operation of the rig, perform any necessary calibrations and to label images that had visible contaminants. More than 20,000 images were taken during this two week period.

Following the onsite trial, all images were classified and manually browsed to identify all images with visible contamination. These images were used to retrain the model and improve the classifier. A comparison of the current and previous iteration of the classification algorithm shows that the new classifier performs better on the diversities of the images from the abattoir environment.

The results show that the Scott Contaminant Detection System effectively detected all visible contaminants on the outer surface of the carcasses and there is strong evidence to support that invisible levels of contaminant are also being detected.

10.2 METHODOLOGY

In this section we describe the system preparation, installation and final data analysis.

10.2.1 On site installation

The hyperspectral contaminant detection system was installed at a lamb and mutton abattoir. The system was installed on the slaughter floor just before the carcasses are washed. Figure 47 show a range of onsite photos of the system.

Figure 47: Onsite photos of the Contaminant Detection rig.

As a safety measure for the surrounding workers, aluminum screens were positioned around the lights and behind the carcass, to avoid direct exposure to the halogen lights. These black screens are clearly seen in Figure 47 and a bird's eye view of the rig is shown in Figure 48. The width and angle of the screens was designed to protect surrounding workers from direct light exposure.
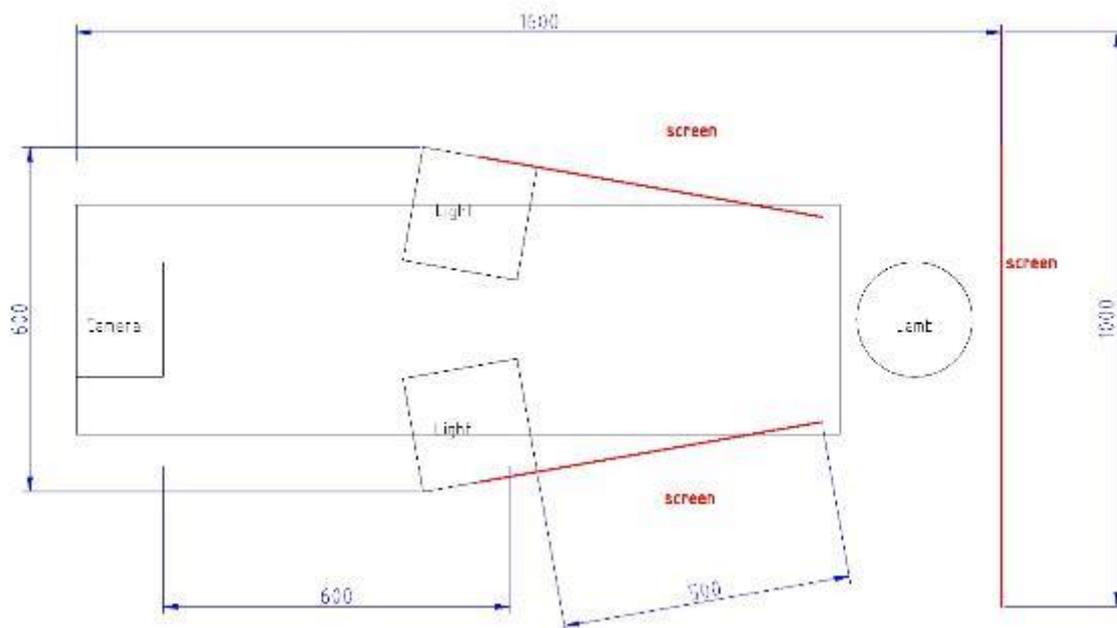
Figure 48: Top view of the Contaminant Detection rig showing protective light screens.

10.2.2 On-site data acquisition

10.2.2.1 Camera setup

The camera B hyperspectral camera used in the system is a linescan camera positioned so that it captures vertical lines. The vertical position of the camera was adjusted so that the field of view was approximately 1m. At a resolution of 1024 pixels each pixel represented approximately 1mm (vertically) of carcass.

The frame rate of the camera determines the effective width of each pixel. As such, the frame rate was chosen to match the speed of the overhead conveyor, so that each pixel captures approximately 1 mm x 1 mm at the distance of the carcass. The overhead conveyor was moving at approximately 350 mm/sec, therefore the frame rate was set to 350 lines per second. At this frame rate, the maximum exposure time is approximately 2.86 ms. The camera has a specified setup time between exposures, and when that is taken into account, the maximum exposure was restricted to 2.3 ms.

To calibrate the camera, white and black reference images were taken every time the lighting conditions of the system change significantly. The average minimum and maximum values of these reference images at each pixel were used to standardize the raw image and effectively illuminate the effects due to the different lighting conditions.

To perform this calibration, a 1.5 m spectralon (white material with high reflectance from 250-2500 nm) was placed in front of the camera at the position of the carcase, then 100 lines were collected with the aperture open and another 100 with the aperture closed to get the white and black reference images respectively.

Despite having 60 halogen lights focused on the spectralon, we were unable to saturate the image sensor during calibration. The maximum white value that was recorded used up only 30% of the sensors dynamic range. At this level of exposure each pixel only used up about 10 bits/pixel of the available 12 bits/pixel sensor. Hence, each pixel represented approximately 1024 values. This is a sufficient dynamic range, however a range closer to 4096 values is achievable if higher intensity lighting is used or if the exposure is increased (although this would require the overhead conveyor to slow down).

10.2.2.2 Real time Acquisition and Classification

To demonstrate that data acquisition and classification can be performed in real time on a standard laptop PC, the classifier was permanently enabled during the whole two week acquisition period. A screen shot of the classifier operating in real time is show in Figure 49. This screen shot was taken from a video recording of the classification process. In this example, the clean (red), contaminated (green) and background (white) masks have been enabled to show that the system is indeed classifying the data. Also shown is the frame rate and exposure settings, and in the messages window a log is displaying the destination folder of the most recently acquired image.
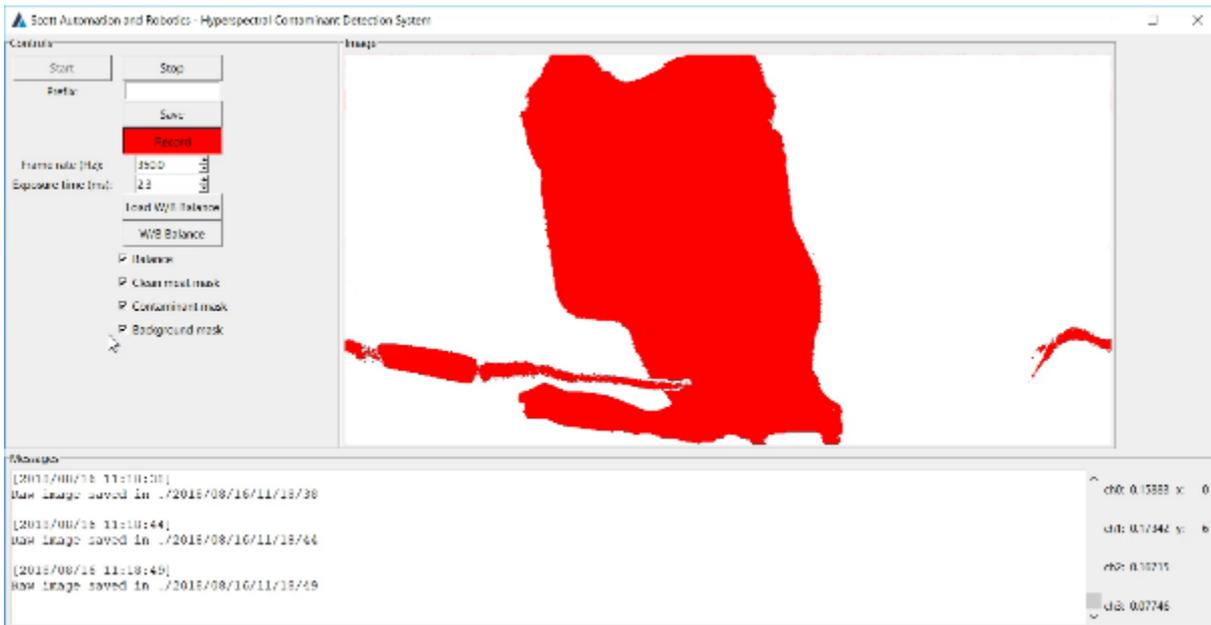
Figure 49: Screen shot of the Hyperspectral Contaminant Detection Program as the system is saving and classifying data in real time. Image shows the silhouette of a lamb carcass and a hand stabilizing the carcass.

10.2.2.3 Overhead Conveyor Data

To keep the data acquisition as simple as possible, the system was set to constantly record. Every time the acquisition program captured 2000 lines of data, it would save that data into a time stamped folder, where the 4 channels of each images were saved as separate 16 bit/pixel PNG files. Each image represents approximately 2 m of overhead conveyor movement. A panoramic view of 6 adjacent images stitched together is shown in Figure 50. For reference, Figure 51 shows the same panoramic image, but with the contaminant mask enabled. This contamination mask clearly shows the presence of contaminants in the second carcase.

Figure 50: A view of 6 adjacent 2000 line images stiched together. Image shows a hand stabalizing the carcase as it is being captured. The 3 "U" shaped markings after the second carcase are a hand being waved in front of the camera to indicate that the previous carcase was visibly contaminated.

Figure 51: A view of the images from Figure 50 with the contamination mask enabled. The contaminant detection algorithm used to here was trained only on data collected from the in house trials.

The panoramic images also show a hand stabilising the carcase as it moves past the camera. The hand appears unusually warped because often it was stationary in front of the camera preparing to grab the carcass.

The "U" shaped markings on the panoramic images are a hand being waved in front of the camera to indicate that the previous carcase was visibly contaminated, and indeed the green contaminant mask in Figure 51 shows that the second carcase was heavily contaminated.

Over the course of two weeks, more than 20,000 images were captured (or approximately 40 km of overhead conveyor loaded with carcasses).

It should be noted that the hyperspectral camera is only acquiring 4 very specific wave lengths of light, and these wavelengths do not represent a true RGB image. Hence, the images shown here are a false RGB image, where only 3 out of the 4 channels are being used to display the image. As such, contaminants that are usually a dark green colour, appear as a maroon/brown colour.

10.2.3 Data Analysis

10.2.3.1 Locating the data

In order to analyze the data, all contaminated images need to be identified. The identification process was performed by first processing all the 20k+ images with the contaminant detection classifier and saving the classified images (those overlaid with a green contaminant mask) to disk. Then the images were browsed manually to locate those containing blotches of green contaminant mask. As a secondary detection mechanism, contaminated images were also detected by looking out for a waving hand (as in Figure 50 and Figure 51) – a sign that a contaminant was clearly visible to the Scott engineer. A total of 190 images were identified where contaminants were clearly visible.

10.2.3.2 Labeling the Data

To further improve the classification algorithm, the model needs to be retrained on the processor data set, and as before, the data needs to be labeled for training. A small selection of 20 images were chosen to retraining the classifier. While the number of images seems small, 20 images are made up of $20\times2000\times1024 \approx 41$ M pixels. Since the classification is performed on a per pixel basis, 41 million pixels is more than enough to sufficiently train a model. Another reason for using only 20 images is to illustrate how well the algorithm generalizes to images that were not used in training. Contaminated image were chosen to include heavy and light levels of contamination. Only areas where clean or contaminated pixels can be visually confirmed were labeled and used for training.

Clean images were selected to include images at varying brightness and at a range of orientations (front, back and side).

10.2.3.3 Retraining the System

To train the classifier, the labeled pixels were randomly split into a training and validation set, with 70% of the pixels used for training and 30% used for validation.

A script was written to train a model hundreds of times, with each time changing the topology of the model. After 3 days of processing the model with the highest accuracy was selected. This model had a validation accuracy of 98.8%. While this accuracy appears very high, it should be stressed that the data was labeled by masking only those areas that appeared contaminated, therefore since we could identify the contaminants by eye, it is expected that the model should be able to accurately detect those same pixels.

To truly assess the accuracy of the model the classification results would need to be compared against swab samples, which is beyond the scope of this trial.

10.3 SUMMARY

In this section we present and compare the contaminant detection results of two classification models – one trained on in-house data and the other trained on onsite data. Both models are tested on 56 contaminated images that were collected from the onsite processor trials. The contaminated test images are shown in Figure 52.
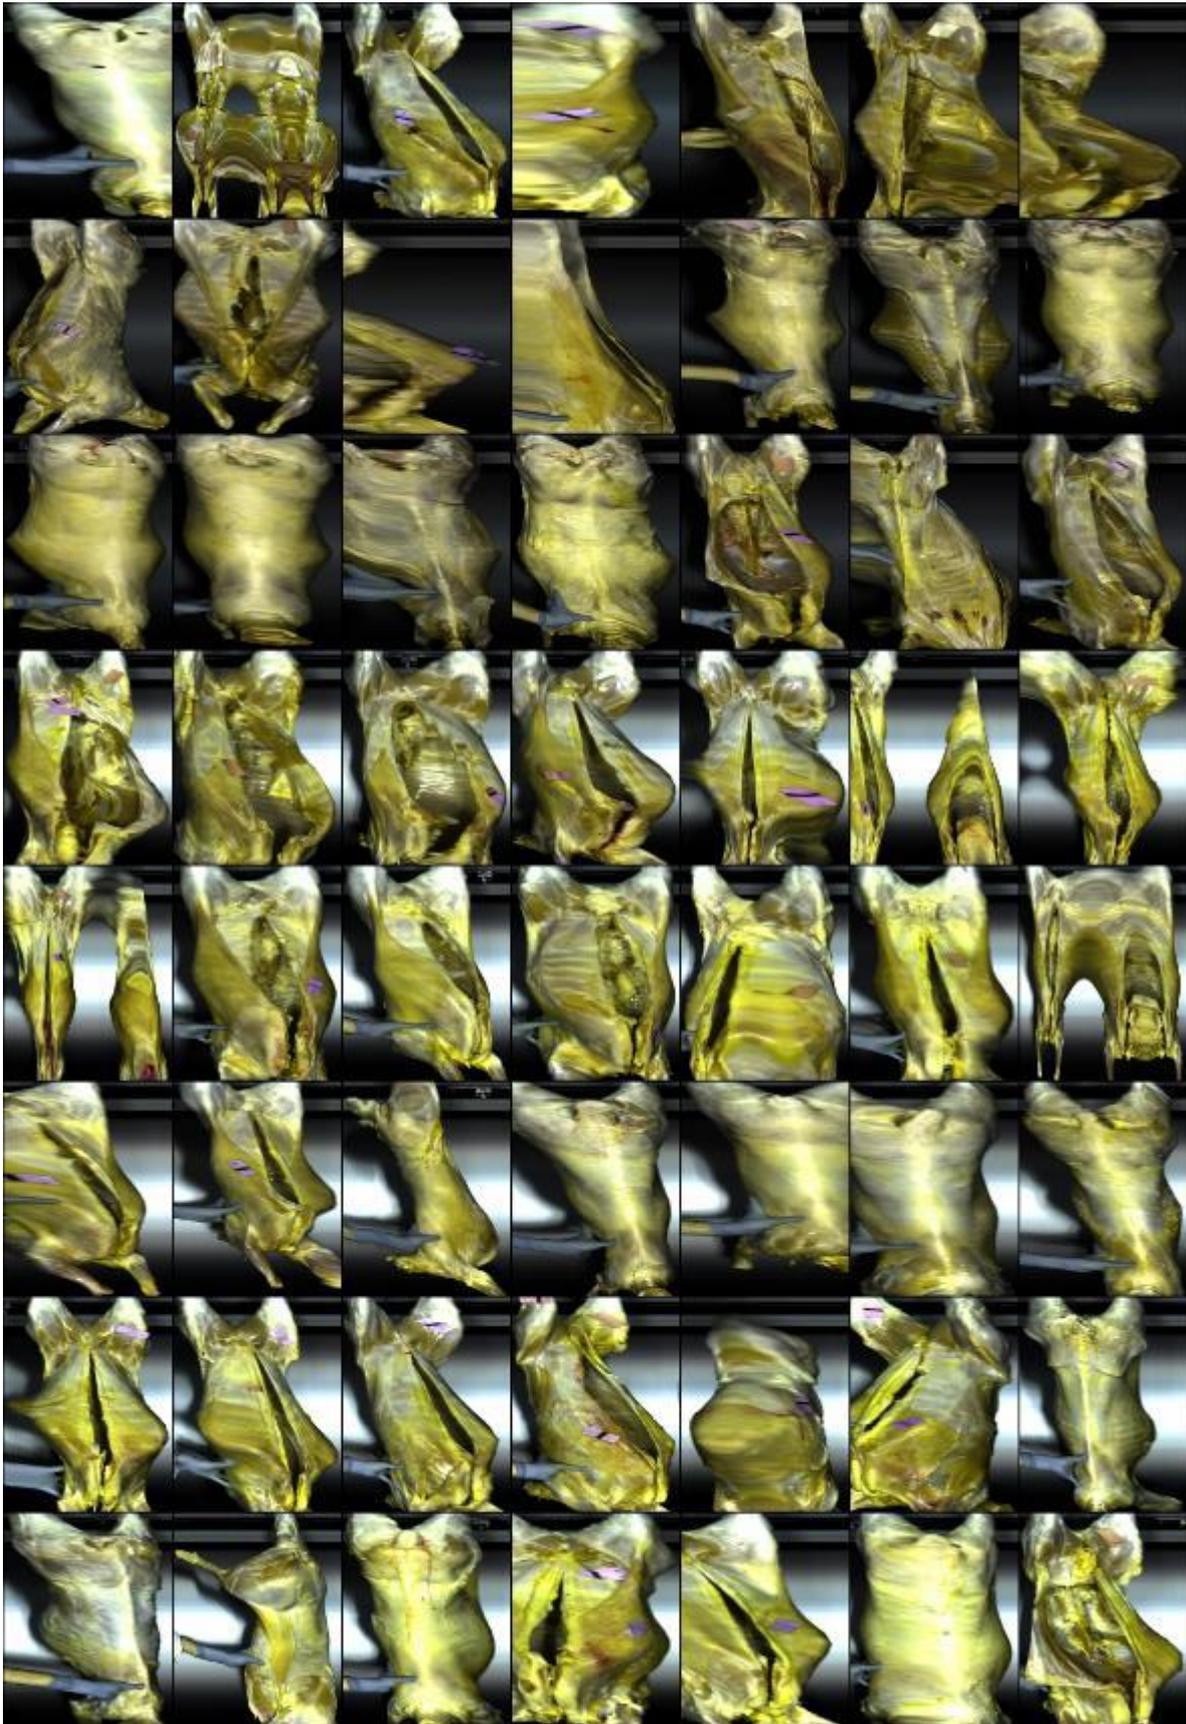
Figure 52: A sample of contaminated images collected onsite.

10.3.1 Model trained on In-house data

Figure 53 shows the classification result of the model trained on in-house data (the same model used while the system was trialed on site).

The results show that the classifier is able to detect all contaminants, however, it also occasionally falsely classifies pixels that are expected to be clean. This misclassification can be attributed to the fact that the in-house model was trained on a small sample of lambs and in a very controlled lighting environment. It is possible but unlikely that the classifier is actually detecting invisible contaminants. This is unlikely because the test data has far more diversity in terms of lighting, carcass size, shadowing and variation in contaminants, and so it is unrealistic to assume that a model trained only on a restricted data set could accurately classify this more diverse set.
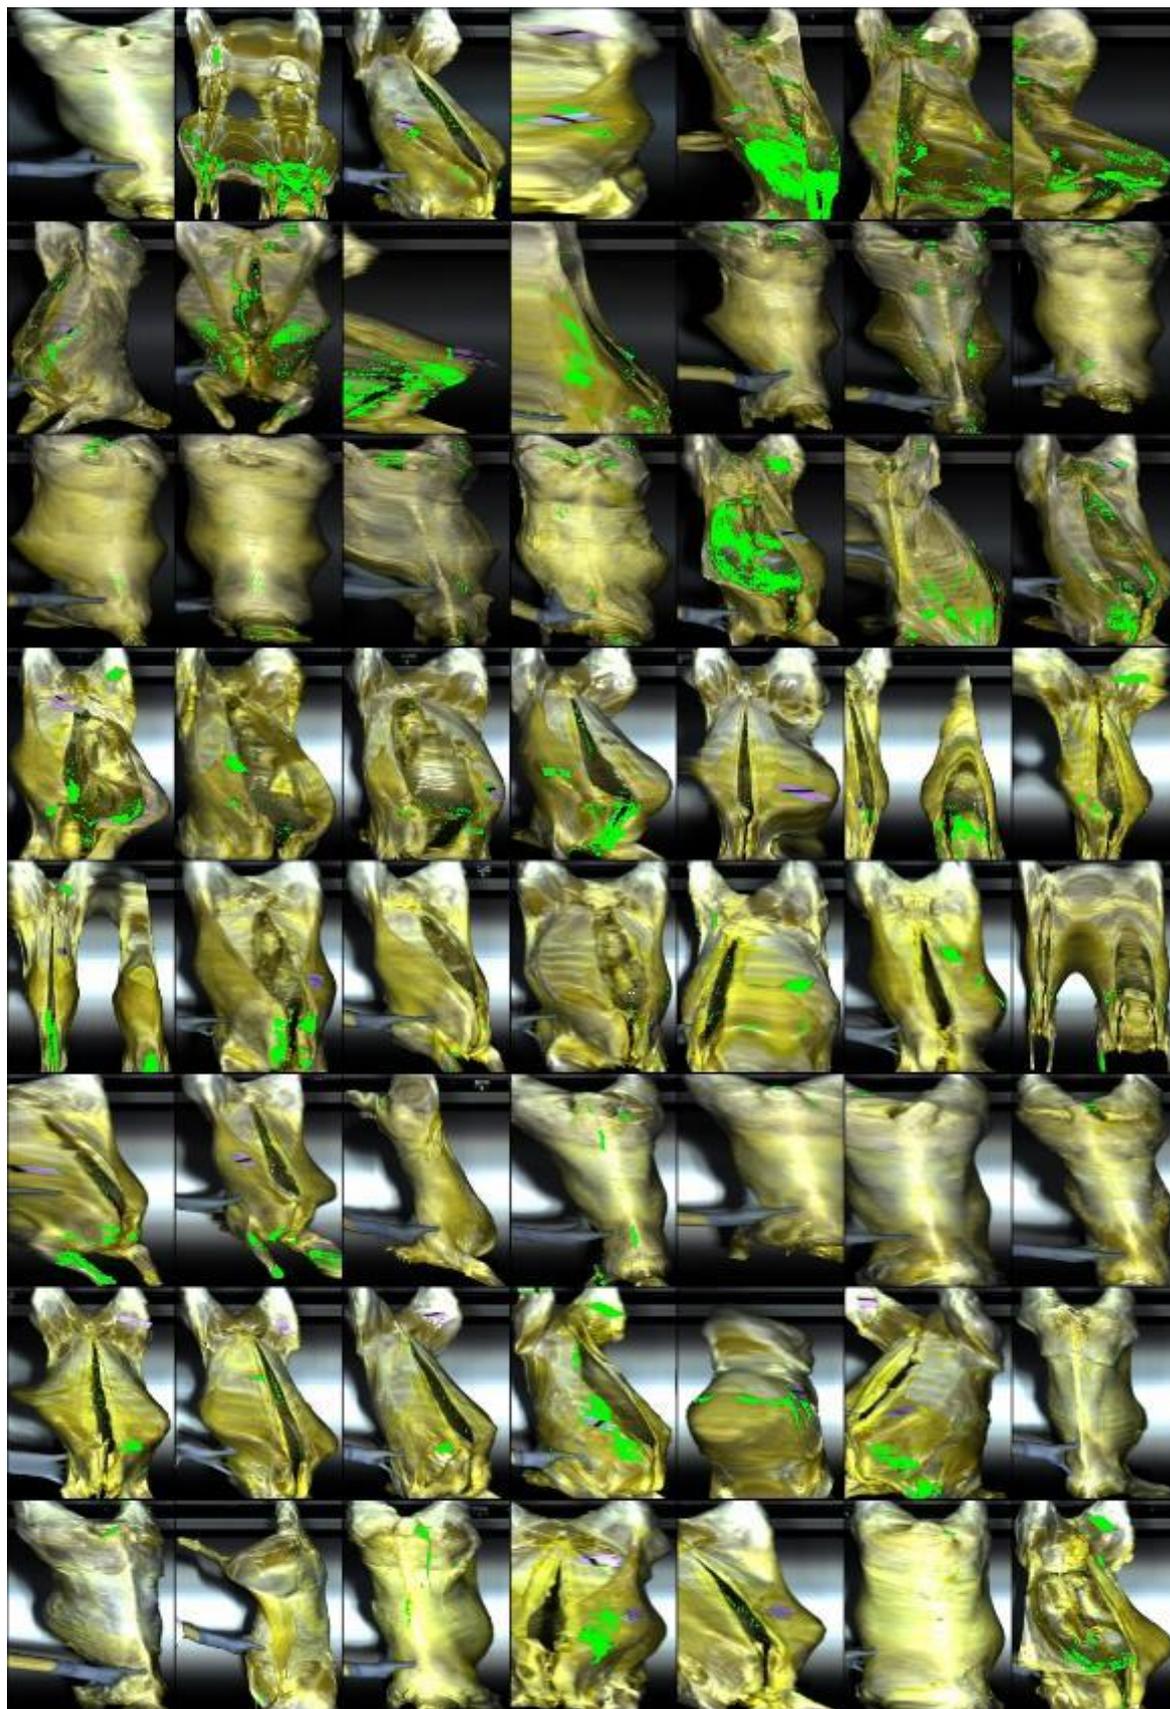
Figure 53: Classified images using model trained on in-house data.

The prevalence of these false positives can be mitigated by performing additional image processing. Specifically, we note that such potential false positives are being detected in sparse and pixelated patches. These sparsely spaced pixels can be filtered out using a simple digital filter. For example, Figure 54 shows a contaminated carcass with and without the speckle noise filter applied to the contaminant mask.
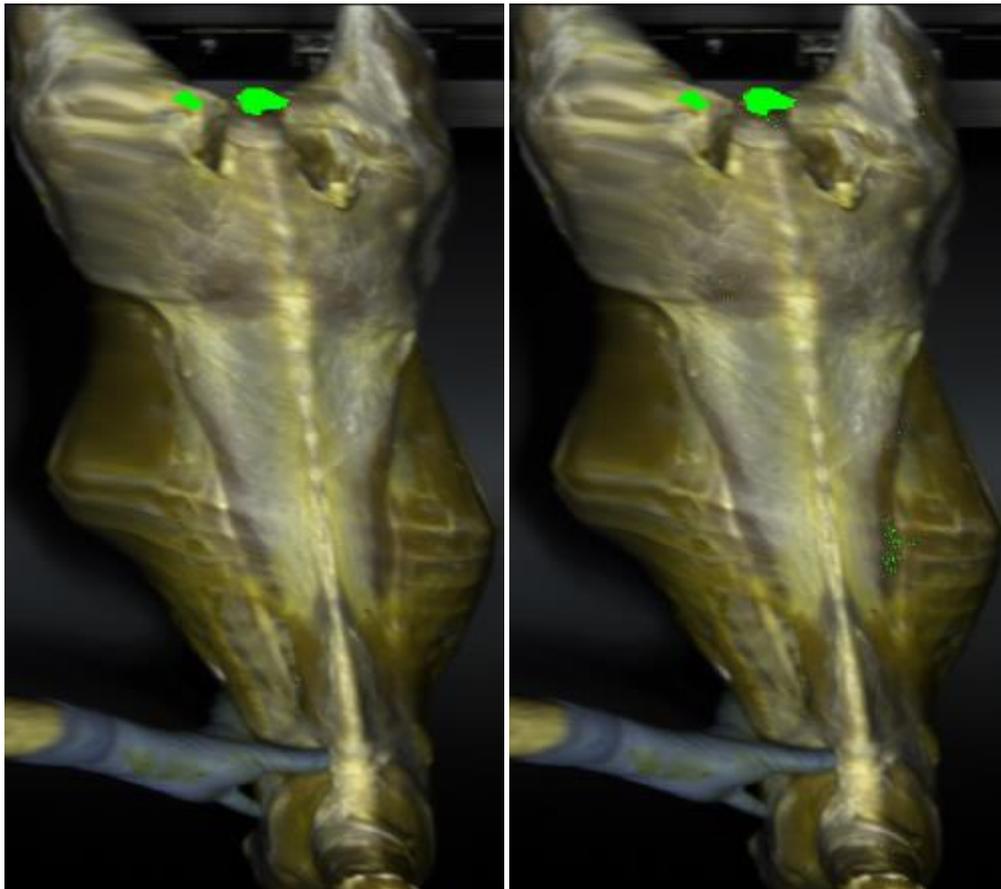
Figure 54: Contaminant mask with and without speckle noise removal.

10.3.2 Model Trained on On-site data

Figure 55 shows the classification result of the model trained on 20 images from data collected during the onsite trial. The results clearly show that the onsite model makes fewer false positive classifications, while still accurately detecting all contaminants compared to the in-house model.
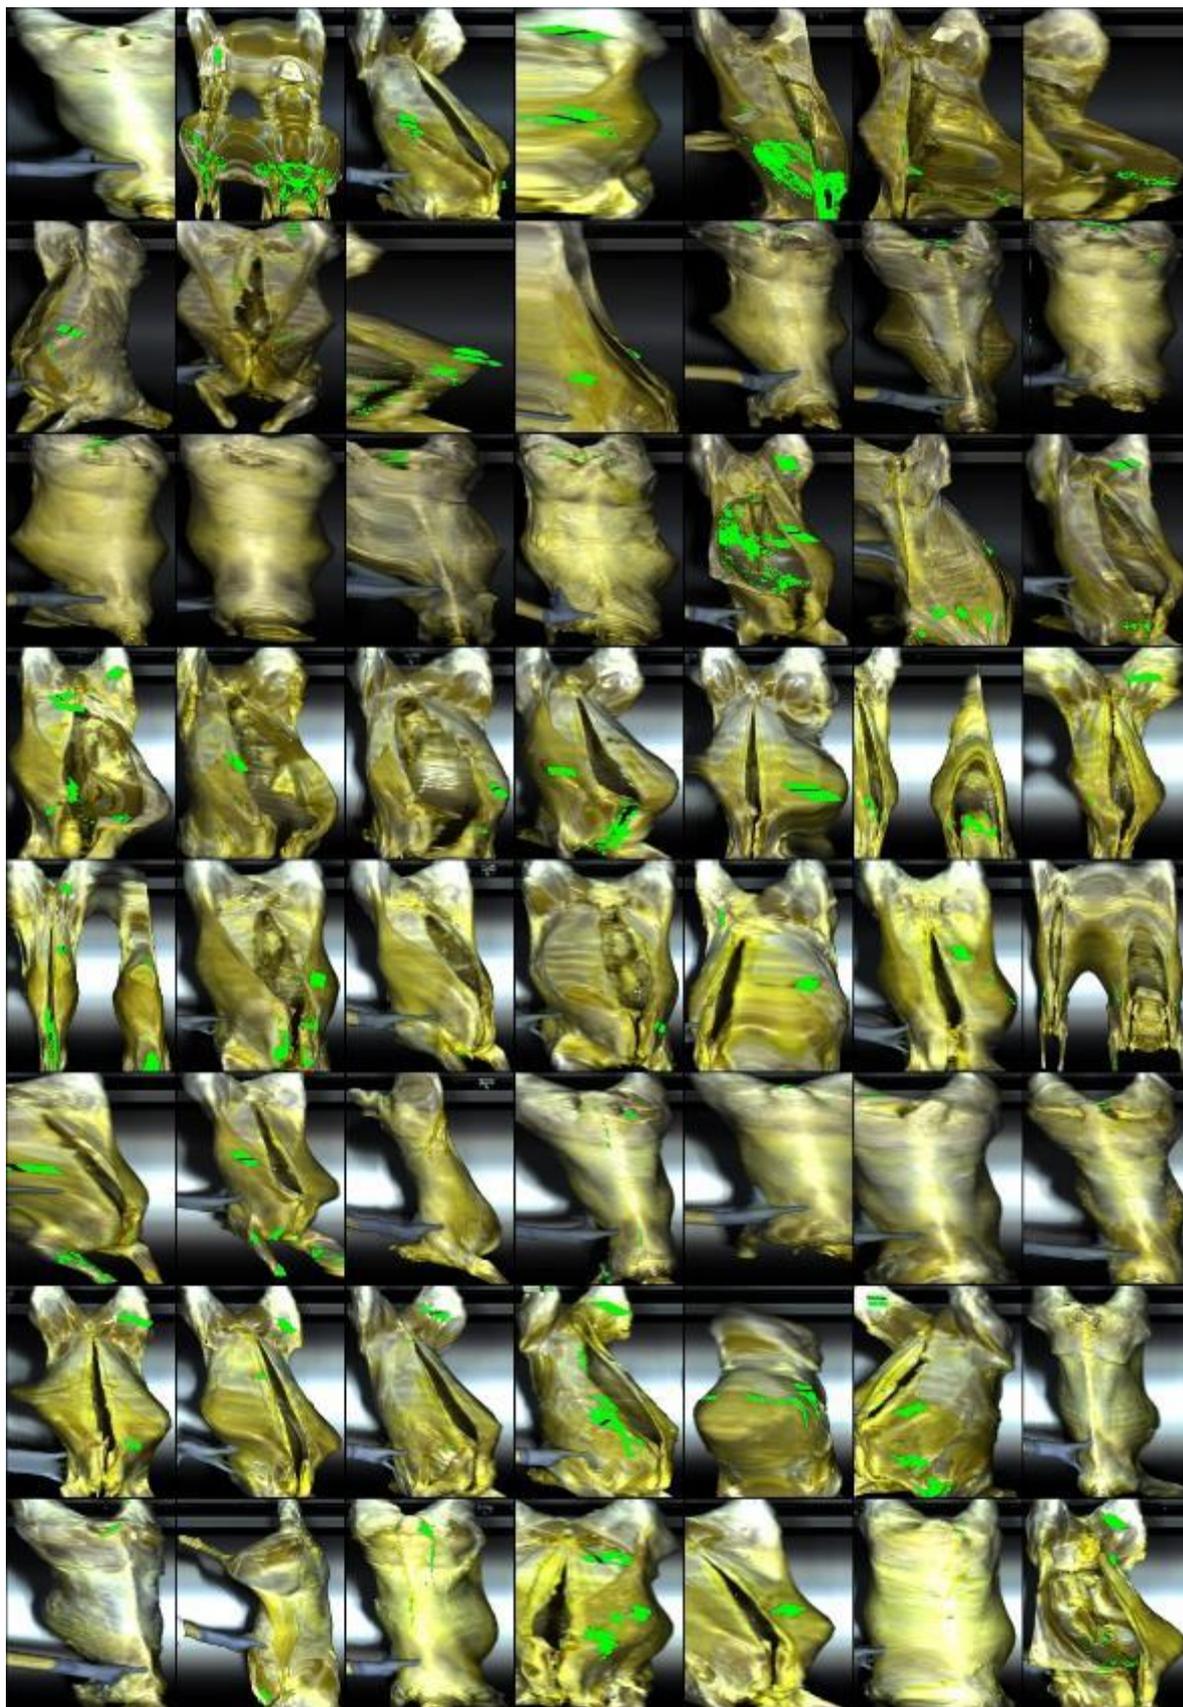
Figure 55: Classified images using model trained on onsite data.

10.3.3 Offcut Data

With approval from the

from the offal room. A photo of these offcuts was taken with an RGB camera and then they were waved in front of the hyperspectral system. This data was collected so that it could be determined if all visible contaminants can be detected. An example of such data is shown in Figure 56. Of particular interest in this example is the fact that there are some areas where it is difficult (if not impossible) to see contaminants with the naked eye, yet the classifier is detecting some invisible contamination. On the other hand, when a clean offcut is tested, there are no traces of visible contamination as shown in Figure 57.
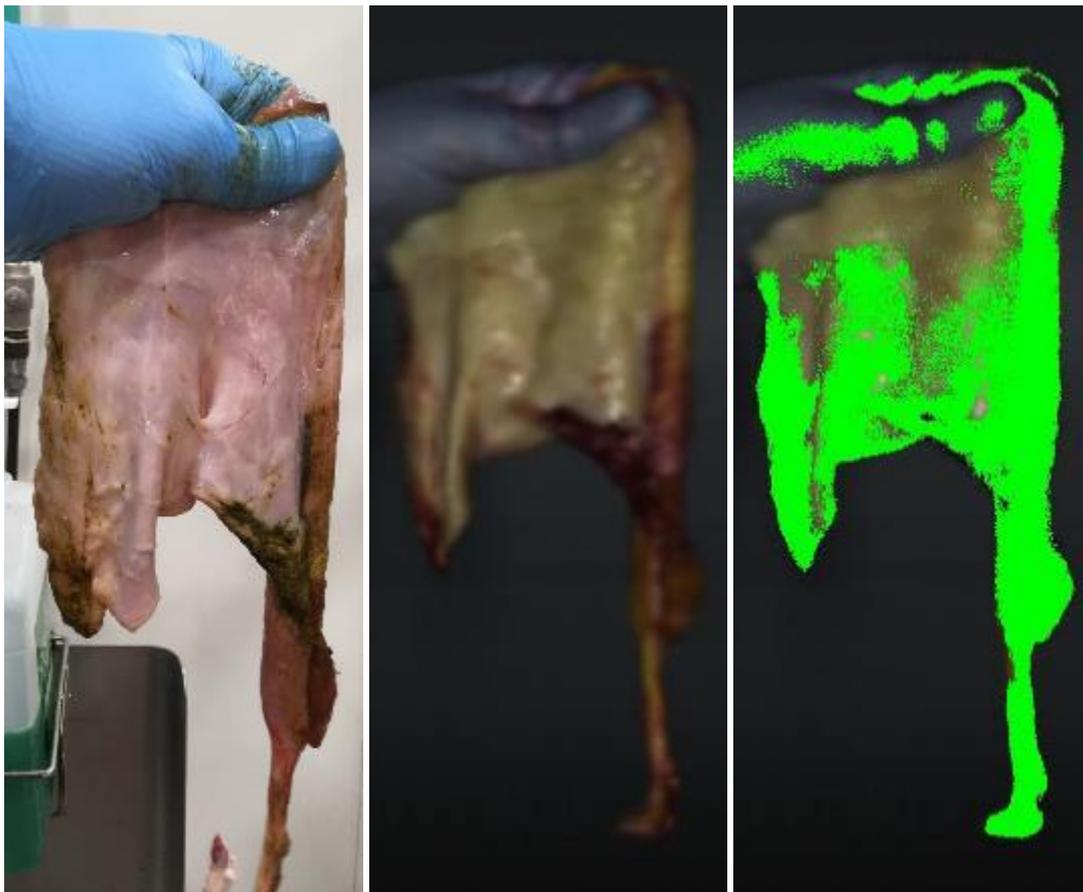
Figure 56: Example of a contaminated offcut showing a photo, hyperspectral and classified image.
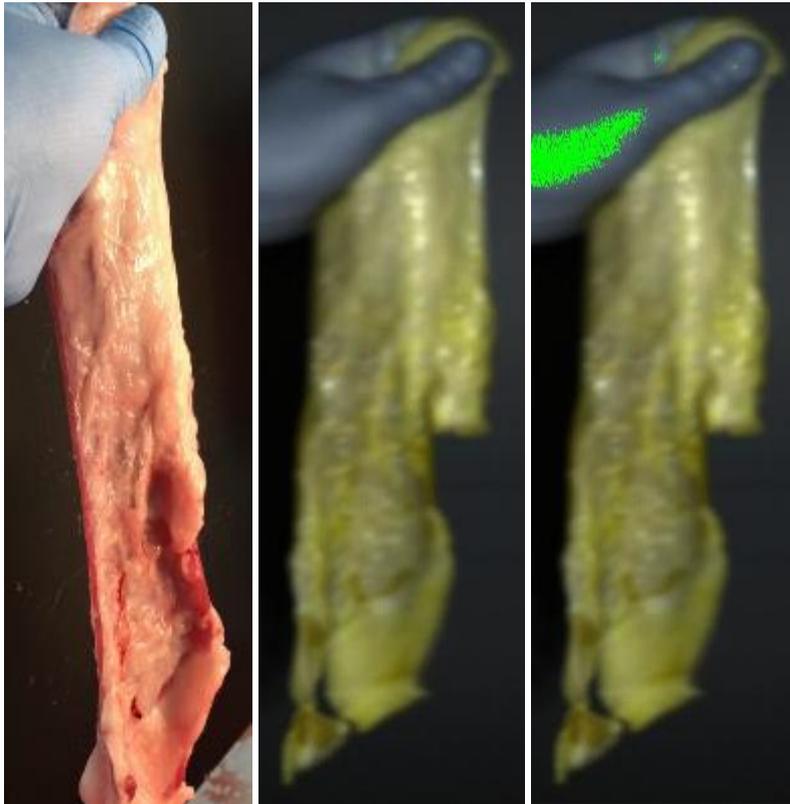
Figure 57: Clean offcut showing no signs of contamination.

11 DISCUSSION

The contaminant detection system detailed in this report performs contaminant detection on a spectral level and as such the accuracy is heavily dependent on the lighting conditions. Ideally, the lighting needs to meet the following criteria:

1) Broad spectrum – To study the spectrum of contaminants the light source needs to provide even illumination across the broad spectrum that is detectable by the hyperspectral camera.

2) High intensity – The individual pixels in the hyperspectral imaging sensor are exposed to a very narrow spectrum of light with an average Full Width Half Maximum value of ranging between 2-4 nm. As such, a light in such a narrow spectrum may contain less than 1% of the total energy being emitted at the corresponding point on the sample. Therefore, to sufficiently expose the pixels, a very high intensity of light is required.

3) Constant – For repeatability over time, it is required that the light intensity does not change over time. Unfortunately the intensity of all lights change over time, so it is important to calibrate the camera on a regular basis.

4) Uniform – For the results to be repeatable for all pixels on the hyperspectral camera, then the subject needs to be uniformly illuminated. In other words, the intensity of the light at any two points illuminated on the subject need to be the same.

5) Diffused – Since the carcass is of varying shape and distance from the camera and light source, it is paramount that the light is emitted from a diffused source so that the light remains close to uniform across the non-uniform carcass surface.

The design of the lighting met the first 3 criteria well. Halogen lamps were used to produce broad spectrum light, 60 lamps were used to achieve high intensity, and the system was calibrated regularly to maintain constant intensity (after image acquisition). Despite this high power the images were still noticeably affected by the abattoir lighting which caused false positives contaminant detections. This was mitigated by attaching a screen directly above the carcasses to block the overhead LED light. The overhead screen can be seen in Figure 58.

The uniform and diffused criteria go hand in hand. In a sense, uniform light was achieved, but that is assuming that the surface of the carcass is perpendicular to the light source. However, given the non-uniform nature of a carcass and variation in size, a perfectly uniform illumination could not be achieved.

A more diffused light source would have aided in providing a more uniform light across the non-uniform landscape of the carcasses, however a highly diffused light was not used for practical reasons. The system in its current state, is already drawing 2.1kW of power through 60 halogen lamps and requires cooling. A sufficiently diffused light source would require up to 10 times the number of lights distributed evenly around the carcass. Therefore, due to the cost, power, heat and space requirements, this type of design was abandoned.

Through its flexibility, the data driven classification algorithm developed in this project was able to

compensate for the variations of light intensity that were expected in our design. However, as illustrated in the results, for the system to be robust it must be trained with a diverse set of data – data that includes will illuminated areas and poorly illuminated crevices. Often it is difficult to supply a sufficiently diverse training set in the training process. For future designs it is therefore recommended that the system allow for the classification algorithm to be updated on site. Typically this would require a user interface that would allow an operator to review the scanned data and correct any misclassifications.

Figure 58: Contaminant detection system with overhead screen shown above the carcass.

Controlled laboratory experiments, confirmed that faeces and ingesta contaminants can be detected in lamb carcasses with a pixel classification accuracy of approximately 95%. In these experiments, the meat was cut into flat pieces, allowing it to be illuminated evenly, and contaminates were applied by hand so that an objective accuracy measurement could be calculated. Under these conditions it was shown that the system was able to reliably detect very low levels of contaminants, even when it was not visible to the eye.

The system accuracy in the production trials is more difficult to evaluate because contaminated areas were identified visually, and so every pixel of the data could not be labeled with absolute accuracy. Instead, entire images were labeled as containing contaminants or not. In all cases where visual contaminant was identified, the system did indeed detect contamination. Similarly, contaminants were

not detected in images where contaminants were not visually present.

The project as a whole presented many challenges including technological, optical, electrical and software challenges, most of these were overcome through innovative design. However, the biggest challenge and one that remains a potential risk to the robustness to the project is lighting. Fundamentally, this system detects contaminants on a per pixel basis by analyzing reflected light from the carcass. The spectral characteristics of the reflected light is heavily dependent on the intensity and uniformity of the incident light source, and the regularity of the illuminated carcass and the carcass distance from light source and camera. Minimizing the variation in these dependencies is difficult to achieve in an abattoir environment because of the carcass movement, variation in thickness and irregular shape.

Despite this potential lighting risk, the results presented in this report demonstrate that contaminant detection can still be performed quite reliably. A level of robustness was achieved by training the classification model on a diverse data set, one that included a variety of illumination variations, and hence the system was able to "learn" how to cope (to some degree) with varying illumination. In our laboratory experiments, uniform illumination was easy to achieve because samples of meat were cut into flat pieces of constant thickness.

An additional lighting difficulty was achieving the necessary intensity. Unlike ordinary grey scale or RGB cameras whose pixels are respond to a broad range of wave lengths of light, the individual pixels of a hyperspectral camera are illuminated by a very narrow band of radiation, thus requiring significantly higher light intensity for sufficient exposure.

Due to the absence of more efficient lighting systems on the market, we had to rely on halogen lights to supply the broad spectrum illumination necessary for hyperspectral applications. The power demands of the halogen lights presented their own set of challenges. For example, the lighting system required a 2.1kW 12VDC power supply which supplied a 175A of current to 60 halogen lamps. As such, the light enclosures were designed to facilitate active cooling via compressed air.

12 CONCLUSIONS/RECOMMENDATIONS

We recommend that future research in the hyperspectral space focus on developing hyperspectral LED technology. With LED illumination it would be more practical to construct a more uniform and intense light source, while minimizing many of the technical challenges that come with halogen lamps (heating, power consumption and low intensity). Hyperspectral LED technology would make future research and development of hyperspectral systems more accessible and longer term more economical system.

Through the research and development of hyperspectral contaminant detection systems, we collected and analyzed a large amount of data in four separate experiments ranging from highly controlled laboratory experiments to onsite abattoir tests. Through these tests, we have been able to identify key wavelengths that are particularly useful for discriminating between clean and contaminated meat. The identification of these key wavelength creates opportunities to reduce the cost and complexity of the camera technology to something more readily available, more reliable and more flexible.

To perform the classification, a data driven process was used to "train" a classification model by showing it examples of clean and contaminated pixels. Using this approach is advantageous because it allows the algorithm to be easily updated upon receiving new data and it does not require any expert knowledge of the specific spectral characteristics of faeces and ingesta on meat and how these characteristics may change under different lighting conditions. As such, this sophisticated classification algorithm is more robust to lighting variations compared to traditional approaches which require near perfect illumination.

Custom image acquisition and user interface (UI) software was developed to acquire data, control camera display, classify and save hyperspectral data in real time. The UI displays a scrolling image on the screen as it is being acquired on the camera, it allows for the operator to enable various image masks to easily identify clean or contaminated meat, it allows the white balance of the camera to be calibrated and allows for camera parameters to be adjusted.

To perform abattoir tests, a new camera was purchased and rig was constructed to house the camera and illuminate the carcasses. This hyperspectral contaminant detection system was designed to be portable and flexible, so that it could be installed at any abattoir and at virtually any stage of the production process. The position and angle of the lights and camera were made adjustable and allowed for fine tuning while on site. Due to the high intensity lighting, screens were fixed around the lights and behind the carcass to protect the workers from direct exposure.

The abattoir and laboratory tests showed that the Scott Hyperspectral Contaminant Detection system can effectively detect contaminants on the outer surface of lamb carcasses. Despite the variations in illumination caused by the motion and irregular shape of the carcasses, the classifier was shown to generalize well across all of these diversities. The system was demonstrated to be computationally efficient as it was able to acquire, classify and save the image data in real time for a period of two weeks while onsite.

A limitation of this technology is that it is a surface scan technology and so only areas that are obstructed from the cameras view will not be tested for contaminants. A particularly risky area is the inner rib cage region which has a flap of skin that block the cameras view. It is therefore recommended

that the system is better suited for a beef plant as the carcasses are cut vertically though the middle which removes all obstructions from the outer surface of the two halves.

Lighting was identified as one of the most significant risks to the reliability of the system. To reduce this risk it is recommended that further research focus on developing high intensity, diffused, hyperspectral LED technology. LED technology would significantly reduce many of the power and heat issues that were experienced with halogen lights. Improving the efficiency would allow for more LED lights to be used, which would increase intensity and allow for a more diffused source of light to provide more uniform illumination over the non-uniform carcasses surface. Hyperspectral LED technology would make further research and development of hyperspectral systems more accessible.

Finally, we recommend that further research investigate the use of ultraviolet light for contaminant detection. We expect that the chlorophyll in faeces and ingest will fluoresce under UV light and become more detectable by the hyperspectral camera.

13 BIBLIOGRAPHY

Ashby, K. D. a. W. J. a. C. P. a. C. T. A. a. R. M. A. a. P. J. W., 2003. Fluorescence of dietary porphyrins as a basis for real-time detection of fecal contamination on meat. *Journal of agricultural and food chemistry,* pp. 3502-3507.

Kim, M. S. a. L. A. M. a. C. Y.-R., 2003. Optimal fluorescence excitation and emission bands for detection of fecal contamination. *Journal of food protection,* pp. 1198--1207.

Lefcourt, A. M. a. K. M. S. a. C. Y.-R., 2005. A transportable fluorescence imagining system for detecting fecal contaminants. *Computers and Electronics in Agriculture,* pp. 63-74.