Australian Government
Department of Agriculture, Fisheries and Forestry

Technical Report

Program and KPI: Sub-program 1.3 KPI 1.20

Report Title: Capability of integrated hyperspectral and 3D imaging cameras to determine subcutaneous fat depth and cover

Prepared by: Dr Alen Alempijevic, Dr Teresa Vidal Calleja, Dr Raphael Falque, Mr Jasprabhjit Mehami

University of Technology Sydney

Date published: <…>

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

Citation

Alempijevic, A., Vidal Calleja, T., Falque, R., Mehami, J. (2023). Capability of integrated hyperspectral and 3D imaging cameras to determine subcutaneous fat depth and cover. An *Advanced measurement technologies for globally competitive Australian meat* Sub-program 1.1 KPI 1.20. January, No. 1. 25 pp

Acknowledgements

This study was undertaken through the Advanced Livestock Measurement Technologies Project (ALMTech) and funded by the Department of Agriculture Rural Research and Development (R&D) for Profit program and Meat and Livestock Australia. Meat and Livestock Australia are thanked for their funding for data acquisition, and for access to commercial herds to compile the calibration datasets.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

Abstract

Hyperspectral imaging has the ability to provide information that goes beyond texture and colour of a surface, near and short-wave infrared spectrum penetrate into the subsurface of materials. Combining disparate sensing modalities, hyperspectral and 3D imaging cameras, could allow to reason about the thickness of the layer of fat over the surface of a carcass (known as subcutaneous fat). We present a learning-based method to estimate subcutaneous fat depth by leveraging hyperspectral data, modelling light sources and 3D surface shape information. In order to develop such a system, subcutaneous fat depth ground truth is recovered from computed tomography (CT) data via a novel systematic approach involving ray casting and non-rigidly aligning a 3D reconstruction captured by the 3D imaging camera. We utilise a collection of mid-loins supplied by a commercial abattoir to validate that subcutaneous fat is correlated to total fat with R^2 0.76. Our method estimates subcutaneous fat over the entire surface of mid-loins (patches of 1.5mm x 1.5mm) with R^2 0.73 and RMSE 1.56mm when compared to CT obtained ground truth. This paves a promising new direction to rapidly estimate information of carcass fat deposits that are complimentary to other LMY technologies which generally estimated whole carcass properties.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

Executive Summary

The overall objective of this work is to leverage the disparate sensing modalities, hyperspectral and 3D imaging cameras, to determine subcutaneous fat depth and cover in order to enable lamb carcass assessment.

To allow estimation of subcutaneous fat we present a hardware solution (cameras and lighting system) and calibration methods for the hardware. We present methodology to combine light source modelling with 3D shape reconstruction and produce a hyperspectral imaging overlay. This allows to remove light reflected from the surface (specular light) and focus only on the portion of light that has results from interaction of with subsurface fat tissue (diffuse light).

In order to develop machine learning methods for subcutaneous fat depth estimation, we propose a systematic approach to obtain “ground truth” subcutaneous fat depth via ray casting CT scans of lamb carcasses. We further transfer this information via non-rigid alignment to the 3D reconstruction obtained from our camera system, and therefore relate it to hyperspectral data. Our machine learning framework can then be trained to estimate subcutaneous fat over the entire surface of a carcass, each estimate relates to a surface area of approximately 1.5mm x 1.5mm

We evaluated the capability of our system on determining fat depth on a selection of mid-loins from a commercial abattoir, collected by Program 2, covering a representative distribution of C-site fat. We demonstrate that subcutaneous fat on mid-loins is correlated to total fat with R^2 0.76 whereas the single point C-site measure has R^2 0.34. While we cannot draw a conclusion on overall fat of the carcass, due to not having total carcass fat information and access to only mid-loins, this relationship provides us with confidence that the subcutaneous fat depth is correlated to overall fat.

We demonstrate ability of our approach to predict subcuntenout fat on previously unobserved mid-loins (fat depth range 1-17mm) with R^2 0.73 and RMSE 1.56mm. While the error tends to increase with fat depth, mean error is still below 2.5mm for areas of carcass with 15mm fat depth. The results indicate the approach with machine learning framework (MLP) can independently estimate fat depth for every patch on the surface. We belive the results could be further improved using a deep learning frameworks that exploits spatial information via convolution.

We conclude that the integrated 3D imaging and hyperspectral system shows good potential for measuring subcutaneous fat depth, and that the fat depth is correceted to total fat. This paves a promising new direction to rapidly estimate information of carcass fat deposits that are complimentary to other LMY technologies that generally deal with overall fat/lean. Further validation and training of across a more diverse phenotypic and genotypic range of lamb carcasses is warranted, as is predicting overall carcase measurements.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

Contents

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

1 Introduction

The overall objective of this work is evaluating the capability of integrated hyperspectral and 3D imaging cameras to determine subcutaneous fat depth and cover, as an objective technology for lamb carcass assessment.

Our earlier work using 3D information in lamb carcasses determined that shape information alone was not capable of capturing muscling related traits, such as Lean Meat Yield (LMY). In contrast to our work in beef cattle LMY, a lamb carcass has lack of visible internal bone structure which impacts extraction of muscle groups. Therefore, the approach to examine fat instead of lean muscle was deemed a possible alternative direction.

Two objective measures relating to fat/yield in lamb carcasses are already defined, C-site fat depth is a single point fat measurement, LMY is an overall meat yield value. Estimation of subcutaneous fat depth over the entire carcass (or parts thereof) could provide complimentary information of fat deposits that are not obtainable by either C-site of LMY technologies.

Robotics perception systems, in particular those comprising multiple sensing modalities, are key to providing autonomy and intelligence (de Medeiros Esper, From & Mason, 2021). Hyperspectral imaging captures information in the near-infrared (NIR) and short-wave infrared (SWIR) range of the electromagnetic spectrum that cannot be seen by conventional colour cameras. In particular, through hyperspectral imaging it is possible to gather internal (under the surface) information about certain meat or food products (Wendel, Underwood & Walsh, 2018; Yuwono & Saputro, 2021). Hyperspectral imaging has been used in red meat industry to estimate Intramuscular Fat (IMF. While a component of IMF could be related to depth of the fat approaches mostly distinguish between muscle and fat cells looking directly onto the meat cut at a fixed distance. Our preliminary work on estimating fat depth using hyperspectral imaging was on lamb short loin portions (Rahman, Quin, Walsh, Vidal-Calleja, McPhee, Toohey & Alempijevic, 2018) at a fixed distance. However, carcasses are curved and the distance between camera and the carcass surface varies. To allow estimation of subcutaneous fat in carcasses, the approach we are proposing is to combine 3D and hyperspectral data.

The main challenge of the task at hand is the lack of data available to evaluate, and more importantly, to develop machine learning methods to regress fat depth. Ground truth fat depth values can only be obtained via medical CT. Thus, we propose a new systematic approach to obtain subcutaneous fat depth by ray casting the CT scans of lamb carcasses and registering them via non-rigid alignment with a 3D reconstruction obtained from our 3D camera system.

Given our perception system and the ground truth data, we present a pipeline that leverages light source modelling, shape reconstruction using 3D, hyperspectral data to develop machine learning models for fat depth regression. We evaluate our work on mid-loins, estimating subcutaneous fat on areas of approximately 1.5x1.5mm. The mid-loins are from a commercial abattoir, collected by Program 2, covering a representative distribution of C-site fat.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

2 Methods

Our perception system developed to estimate subcutaneous fat is shown in Figure 1, consisting of a line-scan hyperspectral camera, a depth camera and a light source.

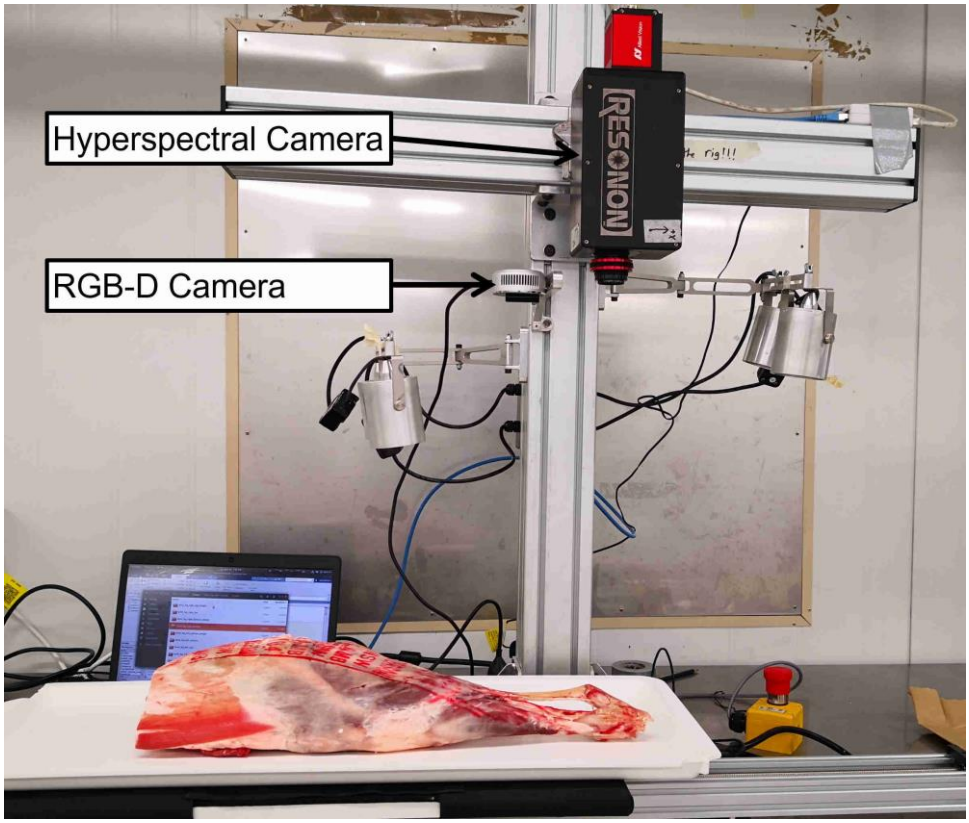

Figure 1 – Acquisition rig, rigidly-mounted camera system comprised of a Resonon Pika NIR-320 hyperspectral line-scan camera and an Intel Realsense L515 RGB-D camera. The specimen (lamb leg in the image above, and in other experiments is a mid-loin), moves linearly under the camera system. The surface of the specimen changes distance and angling with respect to light source and camera(s) given it is not a flat surface.

2.1 Hyperspectral Reflectance

The reflectance estimation method involves the light (LS), the hyperspectral camera (HS), the RGB-D image (F), and the specimen observed by the cameras (P) which is the lamb carcass, as noted in Figure 2. The method aims at decomposing the components of the image observed by the hyperspectral camera in order to isolate the reflectance of the specimen (carcass) which could be described as the true colour of the medium without the influence of the light. The camera system is fixed, and calibrated using the active method outlined by (Mehami, Vidal-Calleja & Alempijevic, 2020) providing spatial transformations between the RGB-D camera towards the hyperspectral camera and the RGB-D camera towards the Light source respectively. In practice this would always hold, the carcass could move, however the camera system and light source would be a single fixed unit. The non-isotropic distribution of irradiance produced by the light source is modelled using a Gaussian Process (GP) following the methods of (Mehami, J., Falque, R., Vidal-Calleja, T., & Alempijevic, A., 2022)

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

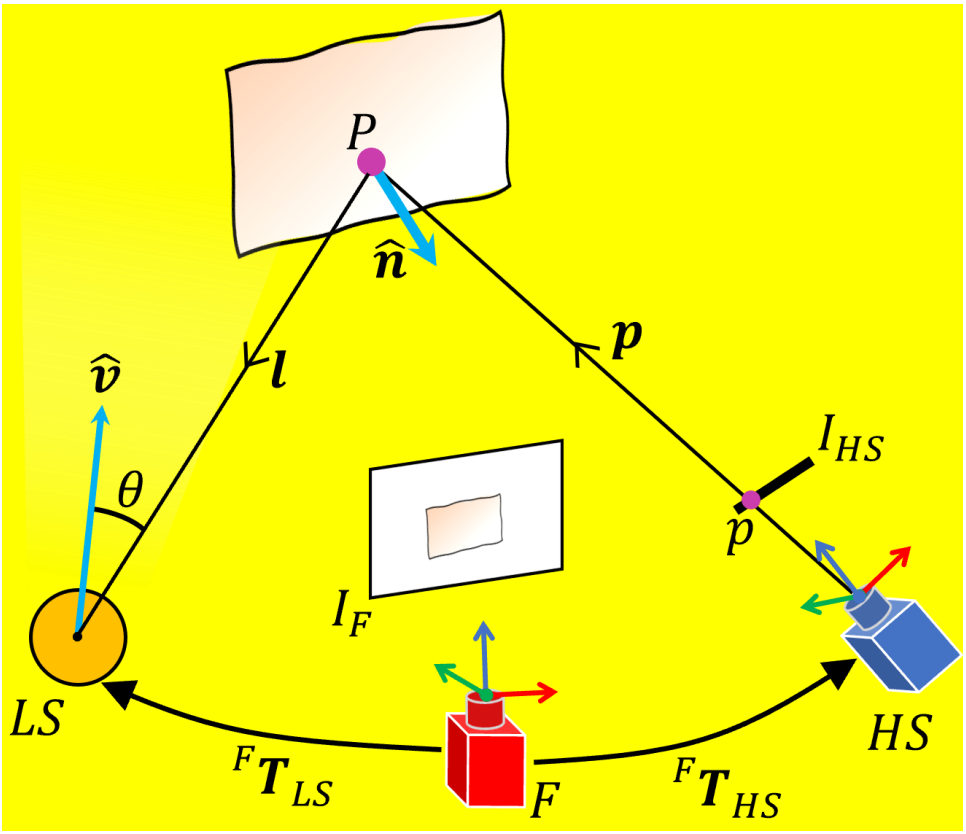

Figure 2 – The RGB-D camera (F), line-scan hyperspectral camera (HS), and light source (LS). In explaining the phenomena between the cameras, light source and surface: light carries radiant flux, which is reflected off the surface and measured by the hyperspectral camera.

2.1.1 Hyperspectral Ray Casting

The first step of the reflectance estimation consists of estimating the surface normals and associating an estimated normal to each hyperspectral measurement. This is achieved by analysing the depth information capture by the RGB-D camera. The depth image from the RGB-D camera is converted into a 3D occupancy grid map, which can then be used for ray casting as shown in Figure 3. The normal is then computed using the depth image from the RGB-D camera.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

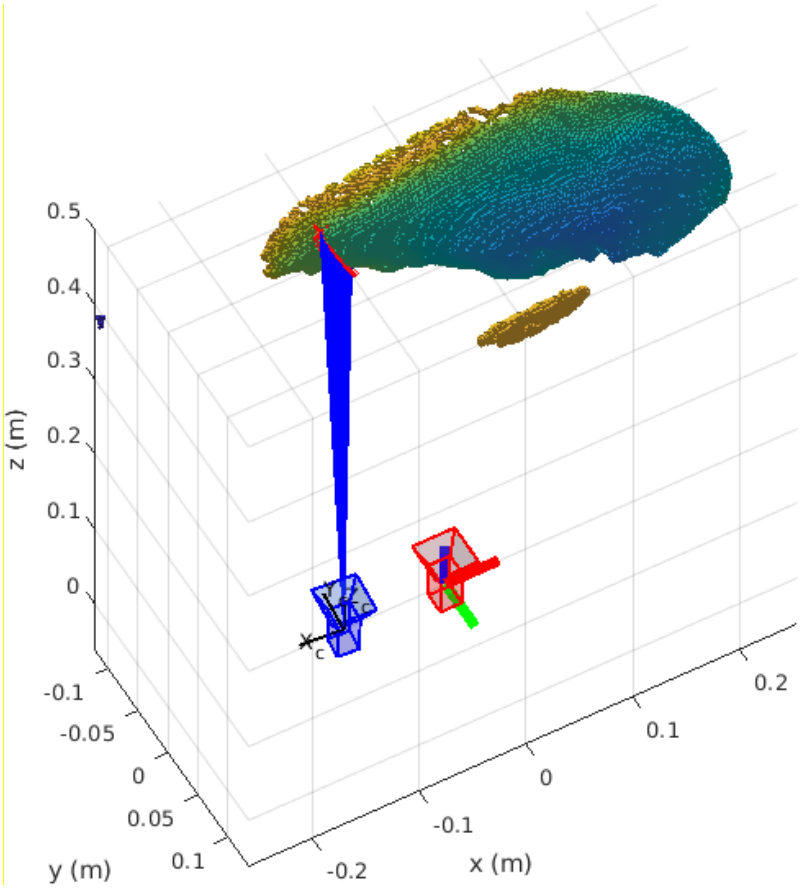

Figure 3 - Ray casting 3D occupancy grid map of a lamb cut. The RGB-D camera in red is used to generate the 3D representation. The hyperspectral camera origin shown in blue is used for ray casting. The rays originating from the hyperspectral are shown in blue and the intersection points on the cut are shown as red crosses.

2.1.2 Estimating the Irradiance

The irradiance arriving on the surface of the sample is assumed to only come from the light source and is estimated using a modelled Gaussian process (GP) distribution of the light source. This again is possible in practise, as light sources can be either isolated, or the contribution from light sources very far away from the camera system could be ignored.

The GP's mean function is modified to fit more accurately the behaviour of the light following the method described by Mehami (Mehami, Falque, Vidal-Calleja & Alempijevic, 2022). The irradiance arriving at all points on the specimen can be estimated from this GP model and projected into the 2D coordinate system of the RGB-D camera. A property that is proportional to the shading can be acquired, where the shading describes the interaction of the irradiance with the surface geometry; the surface normal, light source direction and shading shown in Figure 4.

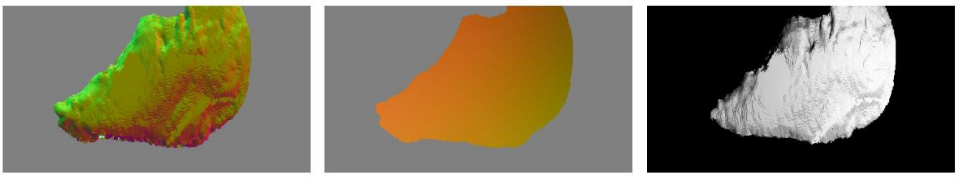

Figure 4 - The variation in shape and lighting for a lamb cut sample – lamb leg used in this example Left: surface normals, middle: the light source direction map, right: the resulting dot product which relates to the shading.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

2.2 Fat Depth Ground Truth

The ground truth of subcutaneous fat depth is obtained through Computed tomography (CT) scans. All specimens are moved to a CT scanner once scanning with the camera system containing Hyperspectral camera is finished. A CT measures the X-rays attenuation and is an ideal modality for distinguishing muscle from fat components (Borkan, Gerzof, Robbins, Hults, Silbert & Silbert, 1982; Goodpaster, Thaete & Kelley, 2000). The CT measurements are in the form of a series of sequential 2D slices. They are stacked into a 3D tensor to build a 3D grid of the X-ray absorption. The voxels have a size of 0.56 × 0.56 × 0.7 millimetres and are stored as DICOM images.

2.2.1 CT Reconstructions

To acquire the subcutaneous fat depth of the lamb specimens and align CT scans to the data of the camera system shown earlier, 3D reconstructed meshes were created from the 2D CT scan slices. This is achieved by first thresholding the X-ray absorption tensors to isolate the fat, muscles, and outer contour of the lamb samples and then using marching cubes for reconstructing the surface of each component (Adrian, S., 2020). Three meshes are created which isolate the fat, muscle and outer contour of the lamb specimen. The fat and muscle are shown in the left part of Figure 5.

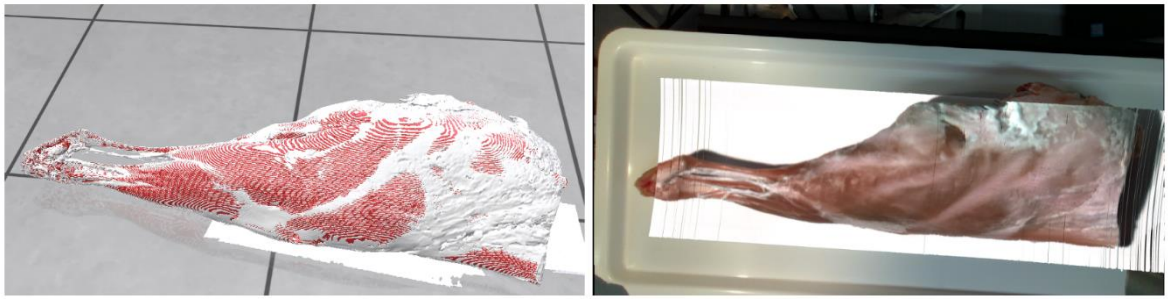

Figure 5 – Left: An example of surface reconstruction of the muscle (in red) and fat (in white) on a leg of lamb obtained from the computed tomography (CT) scans by thresholding the CT data and applying marching cubes (Adrian, S., 2020). Right: Reprojected hyperspectral composite image overlaid onto the RGB image of same leg of lamb.

2.2.2 CT and Depth Alignment

As the lamb specimens scanned with CT scanner were also acquired using the camera scanning rig in Figure 1 this results in two 3D reconstructions. The CT scans are reconstructed into meshes and the RGB-D camera allows us to create a 3D reconstruction. Ideally, these two separate 3D reconstructions of the specimen should perfectly align in a rigid manner. However, this is not what happens in practice, as slight deformations occur when transporting the sample between the camera scanning rig (which is setup in a chiller) and the CT scanner (generally located elsewhere). In our study the acquisition with cameras was at chillers at UNE whereas the CT scanner was in Armidale Clinic located ~5km away. Moving the specimens between the two locations involved loading/unloading crates in a refrigerated ute.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

To solve this problem, the outer contour is aligned to the depth reconstruction in a non-rigid manner by using embedded deformation (Sumner, Schmid & Pauly, 2007) bootstrapped assuming rigid alignment or manual annotation. The deformation is then applied to the fat and muscle meshes to align them to the 3D camera reconstruction. The alignment of the CT fat mesh to the camera 3D reconstruction using the deformation is shown in Figure 6.

Figure 6 - Non-rigid alignment of CT mesh to depth reconstruction using embedded deformation, example shows a leg of lamb. (left) misaligned CT fat mesh shown in yellow and depth reconstruction, (middle) deformation graph of CT fat mesh, and (right) aligned CT fat mesh to depth reconstruction.

2.2.3 Subcutaneous Fat Depth via Ray Casting

The subcutaneous fat depth must be acquired for all hyperspectral imaging measurements of the carcass to allow development and evaluation. This is estimated by ray casting from the position of camera into two aligned meshes of fat and muscle (obtained from CT). If the intersections are valid, with both muscle and fat intersections occurring and $d_f < d_m$, the fat depth is calculated as shown in Figure 7. Given slice size of 0.7mm from CT scanner we can only measure subcutaneous fat depth in increments of 0.7mm.

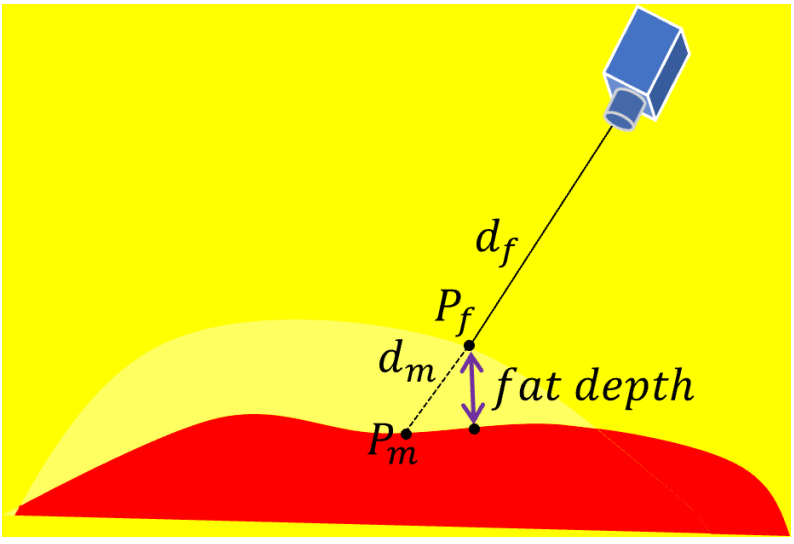

Figure 7 - Ray casting from the hyperspectral camera to the subcutaneous fat (yellow) and muscle (red) meshes (obtained from CT). The figure shows the intersection points (Pf – point on fat and Pm – point on muscle) and their corresponding distances dm. The actual subcutaneous fat depth is the distance highlighted by the purple arrow.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

2.3 Statistical Analysis

The subcutaneous fat depth regression is performed using machine learning approach. Performance is measured using a few metrics. First, we consider root mean square error RMSE and R^2 for all patches across all cuts. Then we examined mean absolute error (MAE) over the subcutaneous fat depth range, which quantifies the overall prediction error by averaging the error for each prediction of a fat depth. Then we examine spatially (over entire surface or via virtual slices) the comparison between actual and predicted fat depth.

2.3.1 Fat Depth Regression Methods

A small portion of the observed fat depth is shallow (approximately 1–2 mm) and only a few regions exhibiting deeper fat content (up to 14 mm). A multilayer perceptron (MLP) is used as this provides higher capabilities to model the nonlinearities of the subcutaneous fat. The MLP consists of three hidden layers with output sizes 16 × 10 × 1. Each layer employs rectified linear unit (ReLU) as an activation function. The MLP will learn from pixel reflectance inputs where each sample is a normalised feature vector with the reflectance from all the different frequencies captured by the hyperspectral camera.

2.3.2 Dimensionality Reduction

Due to the high dimensionality of hyperspectral data, it is common to apply band selection or reduction techniques. These methods not only compress the data and reduce its overall size but can also remove redundancies that are detrimental when training machine learning algorithms. In line with the work from Pena et. Al. (Peña, F., Mehami, J., Guenot-Falque, R., Patten, T., Alempijevic, A., & Vidal Calleja, T., 2020), this work applies factor analysis to reduce the dimensionality of data (number of hyperspectral channels). As per the wotk reported, performance gain beyond 20 components is small, thus we use this value for all experiments as it is a good trade-off between performance and data size.

2.4 Lamb Carcass - Dataset

In November 2021 and November 2022, data collection was conducted at the University of New England in Armidale, New South Wales to study the subcutaneous fat-depth of lamb cuts using a hyperspectral camera. The purpose of these trials was to examine the feasibility of estimating the subcutaneous fat-depth of lamb cuts using the hyperspectral camera.

The first data collection, conducted in November 2021, involved scanning cuts from a single lamb carcass. The cuts that were scanned included one saddle and two leg cuts (left and right). The aim of the trial was to estimate the subcutaneous fat-depth on the individual cuts using the hyperspectral camera where a diversity of fat depth was present over entire surface of the cut. The trial involved the rig scanning each cut with a total of 599 images captured for each camera. The cameras used in the trial were hyperspectral and RGB-D, and they were calibrated with the help of a reflective hemisphere on a planar fiducial board. In addition to the camera images, the trial also involved the capture of CT datasets to measure the ground-truth subcutaneous fat depth. The CT datasets included DICOM images of the bone and soft tissue for each cut, with approximately 1000 image slices captured for each cut. This data analysis is reported in (Peña, F., Mehami, J., Guenot-Falque, R., Patten, T., Alempijevic, A., & Vidal Calleja, T., 2022), given small number of samples is not part of this report.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

The second data collection, conducted in November 2022, involved scanning 30 mid loins from different lamb carcasses supplied by a commercial abattoir. The mid-loins were selected from a total of 147 available. The selection was based on covering the distribution of C-site values from the larger population of 147 (as demonstrated in Figure 8) and within constrains of using the CT scanner. were individually wrapped and on foam trays as shown in Figure 9. The aim of the trial was to show the statistical significance of subcutaneous fat depth models using a larger sample size and to prove generalization. The previous results from the first trial were from a single lamb carcass, so the second trial aimed to demonstrate that the results could be generalized to a larger sample size.

The trial involved 15 scans, with 2 saddle cuts per scan (refer Figure 10), and a total of 250 images captured for each camera. The cameras used in the trial were the same as in the first trial, hyperspectral and RGB, and same calibration process applied (illustrated in Figure 11). The trial also involved five CT scans, each with 6 saddle cuts lined up (refer Figure 10), with a total of 400 image slices captured for each cut. The CT datasets included DICOM images of the bone and soft tissue for each cut.

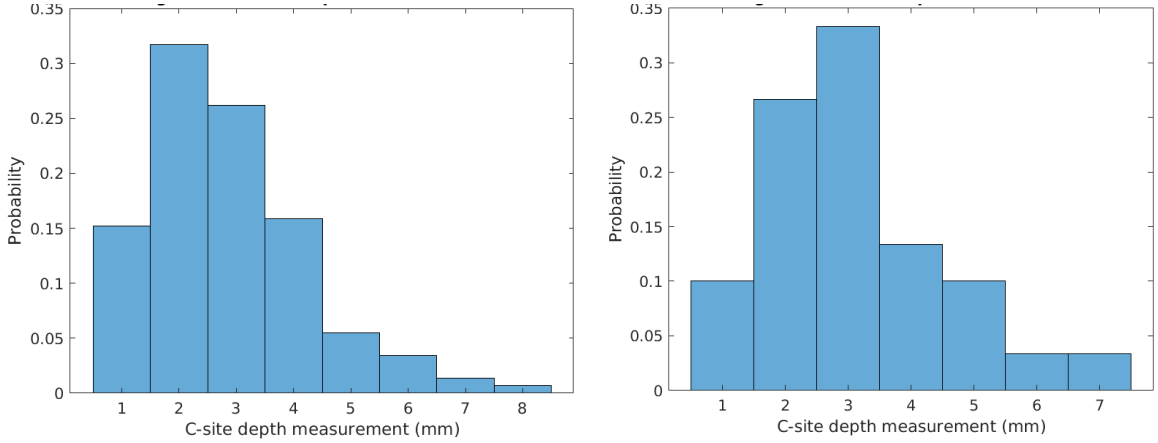

Figure 8 - Histogram of C-site fat depth. Left: for all of the 147 mid loins available. Right: for the 30 that have been selected to be part of the hyperspectral camera and CT scanner trial

Figure 9 - A view of the selection of 30 mid-loins, individually wrapped on foam trays for easier moving between the chiller and CT scanner. From total sample size supplied by commercial abattoir the 30 mid-loins were selected to cover C-site fat depth variation.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

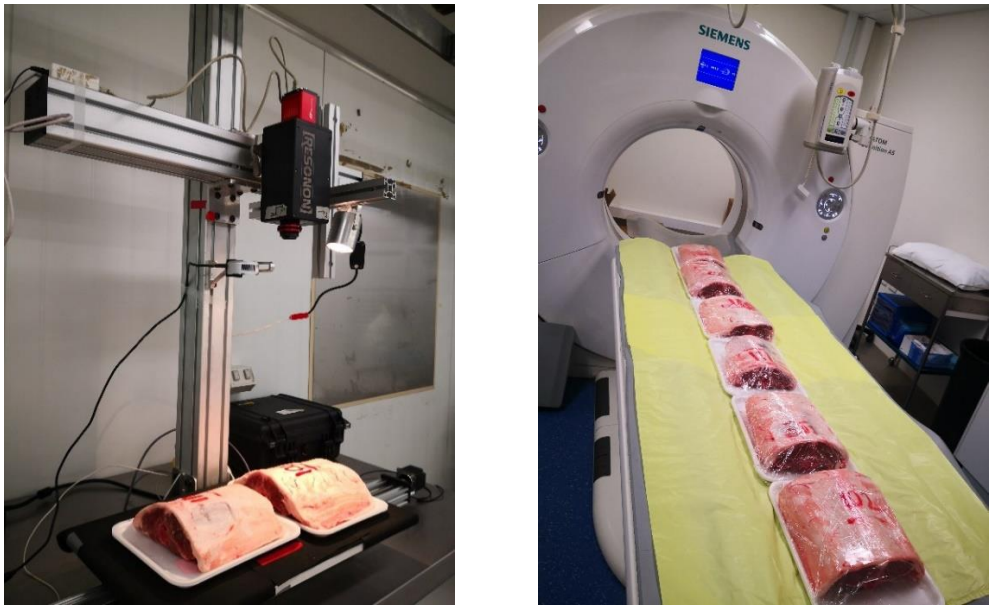

Figure 10 – Process of camera scanning and CT scanning of the mid-loins. Right: camera system scanning the mid-loins (two at a time), protective film is removed during scanning with the cameras, Left: CT scanner scanning 5 mid-loins at a time at a Medical Clinic in Armidale.

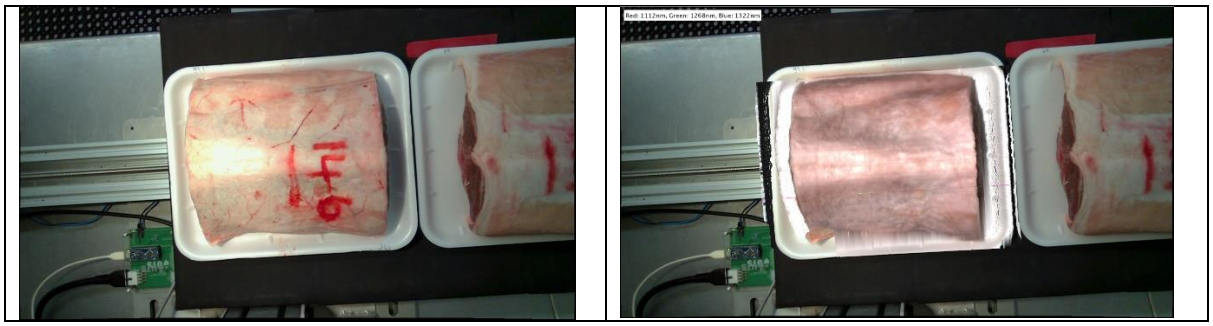

Figure 11 – The steps of registering Hyperspectral data against RGB data. Right: RGB image of a mid-loin at the centre of the image acquired with RGB camera clearly indicating a cut number, the reflection of light on the cut surface is also evident. Left: The RGB image has been altered such that the hyperspectral data appears on the mid-loin, 3 channels have been selected and superimposed. The cut number is not visible and the specular component of light (the one that bounces of the surface) is no longer present. This is the data that is subsequently used for estimation of subcutaneous fat.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

3 Results

Using the dataset (30 mid-loins) a relationship between the total fat of the mid-loin (obtained by CT thresholding fat) and the subcutaneous fat (via ra-tracing) or C-site fat depth is determined. We can see that the subcutaneous fat and C-site fat depth have a R^2 of 0.76 and 0.34 against total fat in Figure 12.

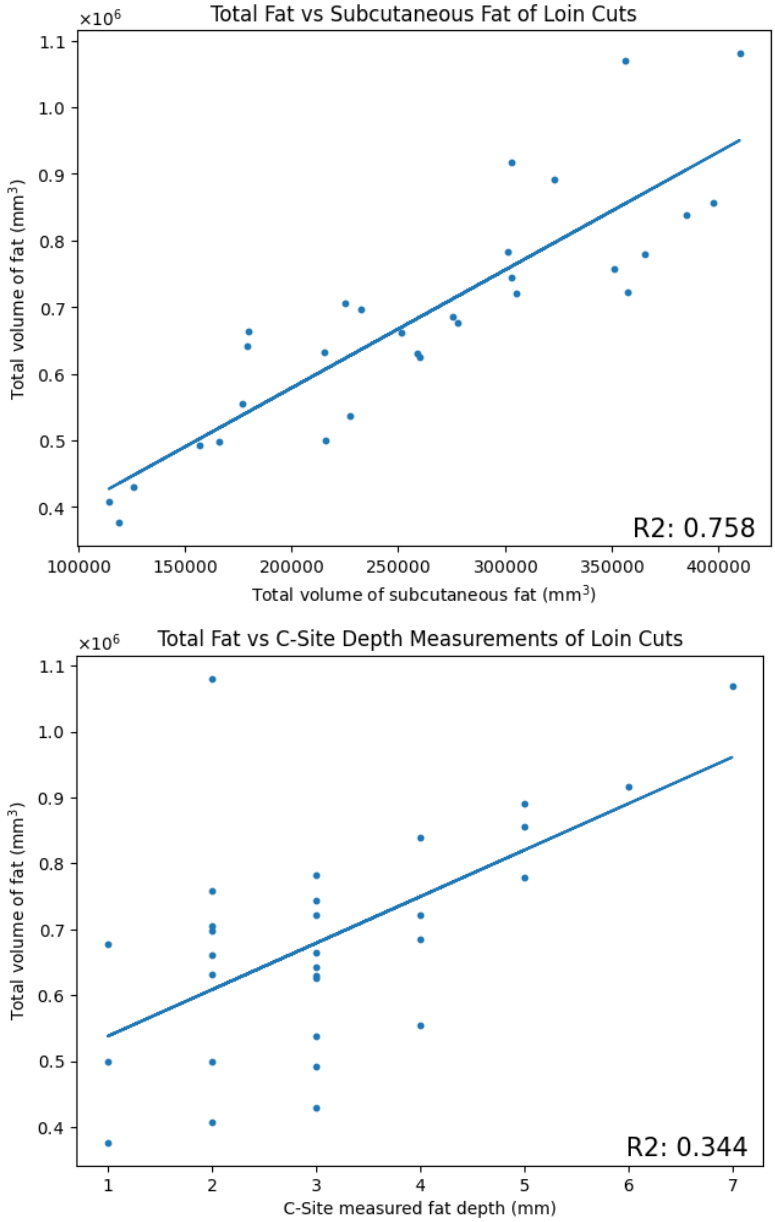

Figure 12 – Analysis on 30 cuts, relationship between total fat of the mid-loin and subcutaneous fat (top) and C-site fat (botom).

The dataset (30 mid-loins) were divided randomly into 80% training and 20% testing split. As we estimate subcutaneous fat depth of patches a histogram for fat depth of the patches is also presented, indicating training and test cuts have a similar distribution Figure 13. This

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

has resulted in a total of 172,543 patches across the 5 mid-loins that are used for testing, each patch being 1.5mmx1.5mm.

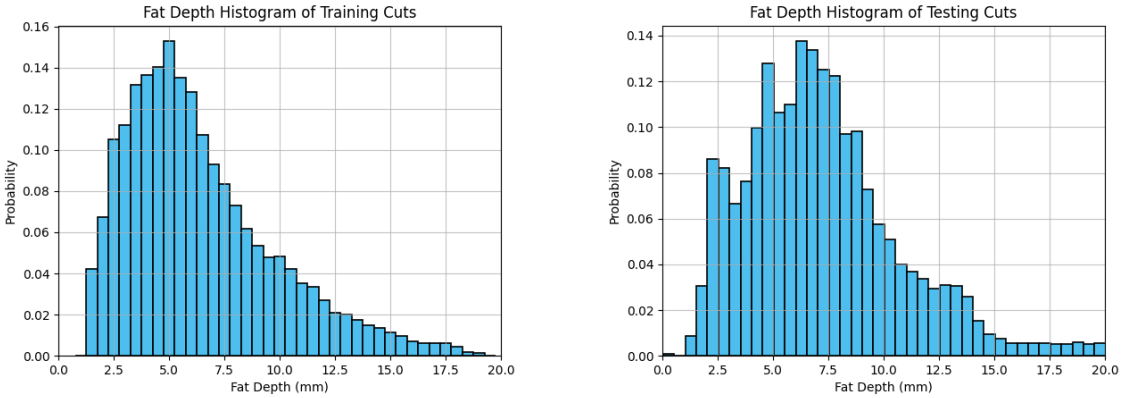

Figure 13 - Fat depth distribution for patches of 1.5mmx1.5mm on the 30 samples. These are separated to a training test and testing set from the data collection undertaken at UNE Armidale. (left) the training set (right) the testing set.

We present results in a few levels. Firstly, at the level of every patch the results of estimating the subcutaneous fat are presented in Figure 14, RMSE is 1.56mm and R^2 0.73. The MAE across the range of subcutaneous fat is also presented in Figure 14, with one standard deviation annotated in shaded blue. There are no points across the mid-loin with fat less than 1mm (complete fat cover), the resolution of fat used for ground truth from CT scanner is 0.7mm. Noting that MAE is under 2.5mm for subcutaneous fat depth under 15mm.

A different spatial view of the prediction results across 5 cuts is presented in Figure 14. We show the estimated and predicted fat depth over the surface of the mid loin, the general trend across the surface is consistent with the ground truth. We can see the areas with high fat depth (over 15mm) are at the edges of the mid-loin and are underestimated.

Finally, a spatial representation consisting of virtual slices across each mid-loin is presented in Figure 15. We show that the fat depth predicted in each slice follows the ground truth and that the fat depth increases symmetrically going away from the spine to the belly. The subcutaneous fat depth across the entire mid-loin is seldom at the C-cite fat depth, the two cuts have C-site fat depth of 3mm and 2mm respectively while loin eye width was 51mm an 53mm.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

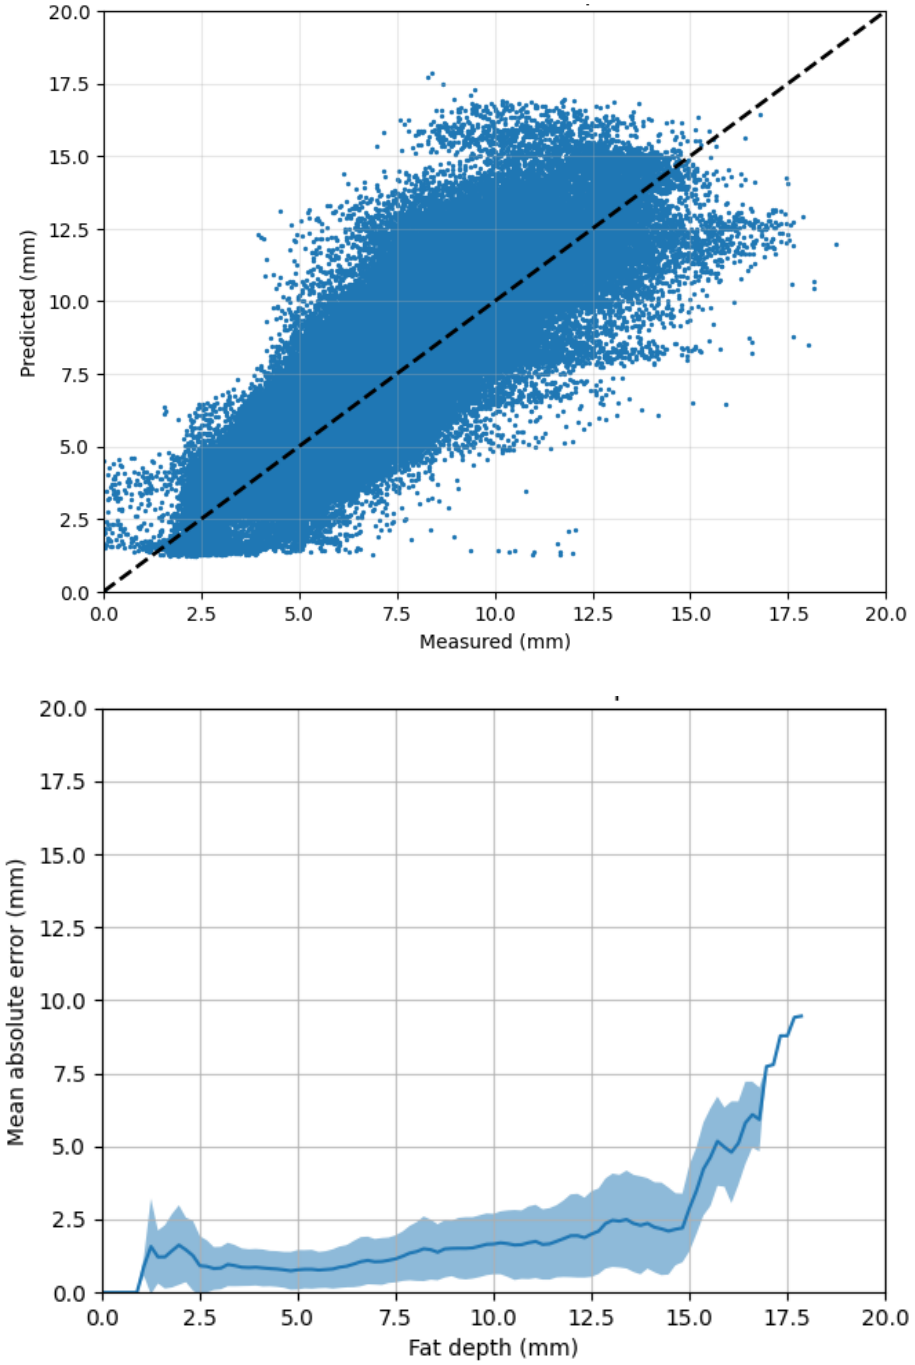

Figure 14 – Top: Measured vs predicted subcutaneous fat depth for each patch of 1.5mmx15.mm across all cuts in training data set. There are 172,453 patches in the cuts Bottom: the MAE vs the subcutaneous fat depth for cuts in the testing set, the MAE is under 2.5mm until the fat depth reaches 15mm. Thereafter the error increases exponentially.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

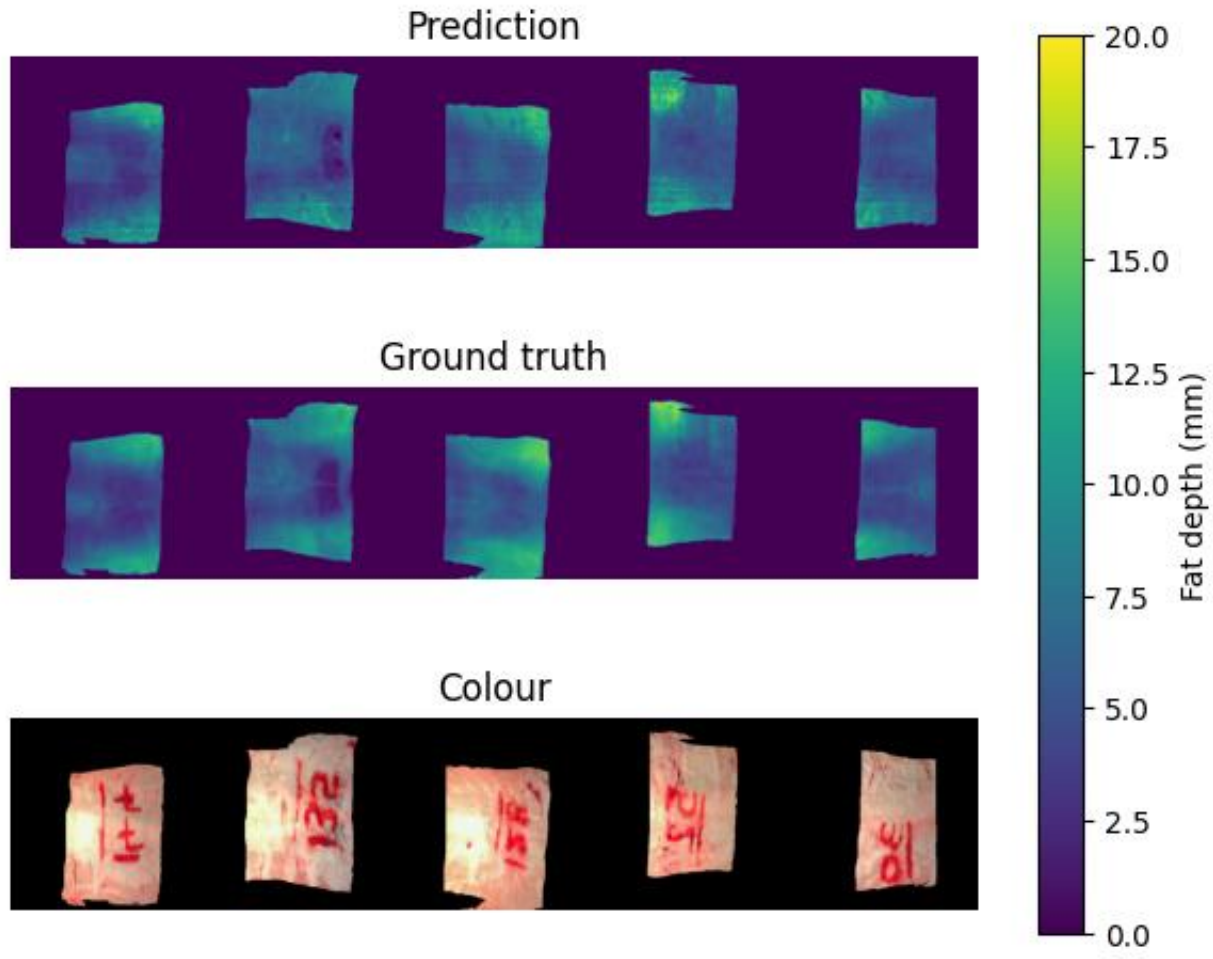

Figure 15 - 2D reprojection of the fat-depth prediction (top) compared to the ground truth value (middle). It can be seen that the red marking on the cuts (bottom) does not affect the prediction of the fat depth

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

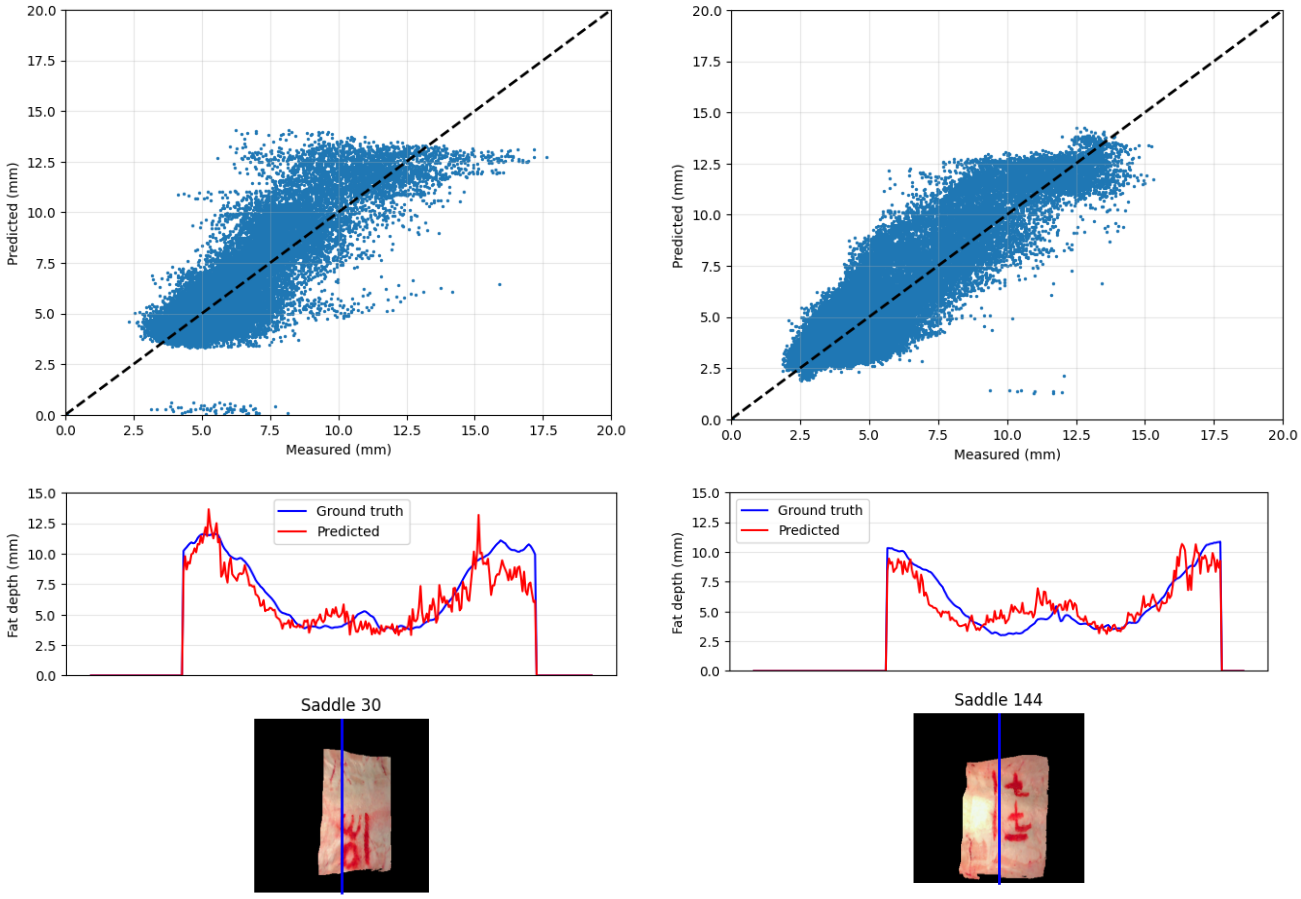

Figure 16 - Visualization of prediction of the subcutaneous fat depth. Top row: measured vs predicted for entire mid-loin. Middle row: of a 2D slice, ground truth is shown in blue and the prediction is plotted in red. Bottom row, the visualisation of the actual cut. Cut labelled 30 (on left) had c-site fat depth of 3mm loin eye width 51mm. Cut labelled 144 (on right) had c-site fat depth of 2mm loin eye width 53mm.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

4 Discussion

Using all 30 cuts at our disposal we demonstrate that for mid-loins the subcutaneous fat is correlated to total fat of the cut R^2 0.73, which is expected measure accounting for vastly more fat deposits (such a subcutaneous fat cover over the whole mid-loin). This is significantly higher than correlation to the single point C-site measure R^2 0.34. While we cannot draw a conclusion for overall fat of the carcass, due to not having total fat information, this relationship provides us with confidence that the subcutaneous fat depth is related to overall fat.

We demonstrate ability to train a model for subcuntenous fat from a subset of mid-loins supplied by a commercial abattoir and predict on previously unobserved mid-loins. The prediction of subcutaneous fat on patches of 1.5mmx1.5mm over the entire surface reports R^2 0.73 and RMSE 1.56mm (fat depth range 1-17mm, total of 172,543 independent patches of 5 mid loins). Our quantisation of the "ground truth" is 0.7mm, related to the resolution of CT scanner and the methodology for ray tracing into CT data to obtain subcutaneous fat depth.

Investigating the error via MAE across the range of subcutaneous fat indicates that the MAE is under 1mm until a depth of 7.5mm and slowly increases to 2.5m at fat depth of 15mm. Thereafter there is an exponential increase of error with fat depths over 15mm. This is related to absorption of IR/SWIR by the fat tissue, and the lack of light penetration further into the tissue as its depth increases. Analysis of the prediction over the entire mid-loin, indicates unsurprisingly that areas very far from the central spine have higher subcutaneous and thereby higher error in prediction.

Examining data presented as virtual slices though the mid-loin indicates the predicted subcuntenous fat depth clarly follows the trend of actual measurements. We posit that this clearly indicates the approach and machine learning framework (MLP) can independently estimate fat depth for each patch. There is a level of noise present when looking at neighbouring estimates of fat depth on this virtual line; exploiting spatial information over the surface to produce more consistency could be accomplished with a slightly "deeper" learning framework. Deeper frameworks exploit spatial knowledge through convolutions and typically need more data to train, having a higher number of internal parameters towards prediction.

We acknowledge that while the mid-loins are from different animals, and have been cut to same specifications, they could belong to a similar cohort and that tests would need to be run across mid-loins ensuring different breeds.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

5 Conclusion

The overall objective of this work is evaluating the capability of integrated hyperspectral and 3D imaging cameras to determine subcutaneous fat depth and cover, as an objective technology for lamb carcass assessment.

To allow estimation of subcutaneous fat we present hardware (cameras and lighting system), calibration methods and methodology that leverages light source modelling with 3D shape reconstruction and hyperspectral data to develop machine learning models for fat depth regression. The methodology separates total light visible on a surface of the carcass to the portion reflected from the surface itself (specular light) and that from the interaction of light with subsurface tissue (diffuse light).

In order to develop machine learning methods and have a “gold standard” for validation we proposed a systematic approach to obtain subcutaneous fat depth by ray casting the CT scans of lamb carcasses and registering them via non-rigid alignment with a 3D reconstruction obtained from our camera system.

We evaluated the capability of our system on a selection of mid-loins from a commercial abattoir, collected by Program 2, covering a representative distribution of C-site fat. Our approach estimating subcutaneous fat over the entire surface of the mid loins, with each estimate of fat corresponding to a surface of approximately 1.5mm x 1.5mm.

Using all 30 cuts at our disposal we demonstrate that for mid-loins the subcutaneous fat is correlated to total fat of the cut R^2 0.73, which is expected measure accounting for vastly more fat deposits (such a subcutaneous fat cover over the whole mid-loin). This is significantly higher than correlation to the single point C-site measure R^2 0.34. While we cannot draw a conclusion for overall fat of the carcass, due to not having total fat information, this relationship provides us with confidence that the subcutaneous fat depth is related to overall fat

We demonstrate ability to predict subcuntenout fat on previously unobserved mid-loins (fat depth range 1-17mm) with R^2 0.73 and RMSE 1.56mm. The MAE across the range of subcutaneous fat indicates errors are below 1mm until a depth of 7.5mm and slowly increases to 2.5m at fat depth of 15mm. Examining data presented as virtual slices though the mid-loin indicates the predicted subcuntenous fat depth clearly follows the trend of actual ground truth measurements. Again, we point out each estimate is independent, and there is no spatial knowledge (relationship between neighbouring measurements) exploited. We posit that this clearly indicates the approach and machine learning framework (MLP) can independently estimate fat depth for each patch. We belive the results could be further improved using a learning frameworks that exploits spatial information via convolution.

We conclude that the integrated 3D imaging and hyperspectral system shows good potential for measuring subcutaneous fat depth to potentially allow predict overall carcase measurements. Further validation and training of across a more diverse phenotypic and genotypic range of lamb carcasses is warranted.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

6 References

Adrian, S. (2020) DicomToMesh. https://github.com/AOT-AG/DicomToMesh.

Alevizos, E. (2020). A combined machine learning and residual analysis approach for improved retrieval of shallow bathymetry from hyperspectral imagery and sparse ground truth data. Remote Sensing, 12(21), 3489.

Alonso, I., Riazuelo, L., & Murillo, A. C. (2020). Mininet: An efficient semantic segmentation convnet for real-time robotic applications. IEEE Transactions on Robotics, 36(4), 1340-1347.

Borkan, G. A., Gerzof, S. G., Robbins, A. H., Hults, D. E., Silbert, C. K., & Silbert, J. E. (1982). Assessment of abdominal fat content by computed tomography. The American journal of clinical nutrition, 36(1), 172-177.

de Medeiros Esper, I., From, P. J., & Mason, A. (2021). Robotisation and intelligent systems in abattoirs. Trends in Food Science & Technology, 108, 214-222.

Eddy, P. R., Smith, A. M., Hill, B. D., Peddle, D. R., Coburn, C. A., & Blackshaw, R. E. (2014). Weed and crop discrimination using hyperspectral image data and reduced bandsets. Canadian Journal of Remote Sensing, 39(6), 481-490.

Goodpaster, B. H., Thaete, F. L., & Kelley, D. E. (2000). Composition of skeletal muscle evaluated with computed tomography. Annals of the New York Academy of Sciences, 904(1), 18-24.

Krebs, A., Benezeth, Y., & Marzani, F. (2020). Intrinsic RGB and multispectral images recovery by independent quadratic programming. PeerJ Computer Science, 6, e256.

Li, S., Song, W., Fang, L., Chen, Y., Ghamisi, P., & Benediktsson, J. A. (2019). Deep learning for hyperspectral image classification: An overview. IEEE Transactions on Geoscience and Remote Sensing, 57(9), 6690-6709.

Ma, L., Liu, J., Pei, X., Hu, Y., & Sun, F. (2019). Calibration of position and orientation for point light source synchronously with single image in photometric stereo. Optics express, 27(4), 4024-4033.

Maier, P. M., & Keller, S. (2018, September). Machine learning regression on hyperspectral data to estimate multiple water parameters. In 2018 9th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS) (pp. 1-5). IEEE.

Mehami, J., Vidal-Calleja, T., & Alempijevic, A. (2020). Observability driven multi-modal line-scan camera calibration. In 2020 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI) (pp. 285-290). IEEE.

Mehami, J., Falque, R., Vidal-Calleja, T., & Alempijevic, A. (2022). Multi-Modal Non-Isotropic Light Source Modelling for Reflectance Estimation in Hyperspectral Imaging. IEEE Robotics and Automation Letters, 7(4), 10336-10343.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

Paoletti, M. E., Haut, J. M., Plaza, J., & Plaza, A. (2019). Deep learning classifiers for hyperspectral imaging: A review. ISPRS Journal of Photogrammetry and Remote Sensing, 158, 279-317.

Peña, F., Mehami, J., Guenot-Falque, R., Patten, T., Alempijevic, A., & Vidal Calleja, T. (2022). Subcutaneous Fat Depth Regression Using Hyperspectral and Depth Imaging. In *Australasian Conference on Robotics and Automation*. ARAA, pp 1-8.

Pyo, J., Duan, H., Baek, S., Kim, M.S., Jeon, T., Kwon, Y.S., Lee, H. & Cho, K. H. (2019). A convolutional neural network regression for quantifying cyanobacteria using hyperspectral imagery. Remote Sensing of Environment, 233, 111350.

Rahman, S., Quin, P., Walsh, T., Vidal-Calleja, T., McPhee, M. J., Toohey, E., & Alempijevic, A. (2018). Preliminary estimation of fat depth in the lamb short loin using a hyperspectral camera. Animal Production Science, 58(8), 1488-1496.

Robles-Kelly, A., & Huynh, C. P. (2012). Imaging spectroscopy for scene analysis. Springer Science & Business Media.

Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18 (pp. 234-241). Springer International Publishing.

Sumner, R. W., Schmid, J., & Pauly, M. (2007). Embedded deformation for shape manipulation. In ACM siggraph 2007 papers (pp. 80-es).

Torres, I., & Amigo, J. M. (2019). An overview of regression methods in hyperspectral and multispectral imaging. Data Handling in Science and Technology, 32, 205-230.

Trajanovski, S., Shan, C., Weijtmans, P. J., de Koning, S. G. B., & Ruers, T. J. (2020). Tongue tumor detection in hyperspectral images using deep learning semantic segmentation. IEEE transactions on biomedical engineering, 68(4), 1330-1340.

Weijtmans, P. J. C., Shan, C., Tan, T., De Koning, S. B., & Ruers, T. J. M. (2019, April). A dual stream network for tumor detection in hyperspectral images. In 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019) (pp. 1256-1259). IEEE.

Wendel, A., & Underwood, J. (2016, May). Self-supervised weed detection in vegetable crops using ground based hyperspectral imaging. In 2016 IEEE international conference on robotics and automation (ICRA) (pp. 5128-5135). IEEE.

Wendel, A., Underwood, J., & Walsh, K. (2018). Maturity estimation of mangoes using hyperspectral imaging from a ground based mobile platform. Computers and Electronics in Agriculture, 155, 298-313.

Xin, Z., Jun, S., Yan, T., Quansheng, C., Xiaohong, W., & Yingying, H. (2020). A deep learning based regression method on hyperspectral data for rapid prediction of cadmium residue in lettuce leaves. Chemometrics and Intelligent Laboratory Systems, 200, 103996.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.

Yang, Y., Soyemi, O. O., Landry, M. R., & Soller, B. R. (2005). Influence of a fat layer on the near infrared spectra of human muscle: quantitative analysis based on two-layered Monte Carlo simulations and phantom experiments. Optics express, 13(5), 1570-1579.

Yuwono, H. A., & Saputro, A. H. (2021, October). Hyperspectral and Deep Learning-based Regression Model to Estimate Moisture Content in Sea Cucumbers. In 2021 8th International Conference on Electrical Engineering, Computer Science and Informatics (EECSI) (pp. 283-287). IEEE.

This project is supported by funding from the Australian Government Department of Agriculture, Water and the Environment as part of its Rural R&D for Profit programme in partnership with Research & Development Corporations, Commercial Companies, State Departments & Universities.