Final Report

Striploin camera upgrade

LEAP4Beef-Module L4B01 Project 2 – Striploin chining pre-production cell accuracy improvements

Project code
2025-1045

Prepared by
Vaughan Weatherall & Steve Maunsell

Date submitted
10/06/25

Published by
AMPC

Date published
10/06/25

1 Abstract

This project has focused on improving the accuracy of the striploin chine removal cell. The improvements include:

- The camera quality and geometry has improved.
- A camera pair has been added to capture the side of the product.
- An automated camera calibration mechanism has been developed with an associated robotic end-effector.
- A feedback mechanism has been created to create a dataset for CNN training.
- The robot path has added flexibility to allow curved paths.

These changes were implemented at the installed machine at the trial site. 714 striploin products were put through the machine to test these changes and generate data to train a CNN to generate cut co-ordinates for the robot bandsaw.

These changes significantly improved the accuracy of the machine from (Maunsell & McCrorie, 2024) and simplifies data collection to enable further improvement.

The development now indicates a minor yield improvement, relative to the manual operators. A concept has been proposed for a standalone machine with the capacity to remove two bandsaw operators and add one loader per shift. The adoption of the machine, including yield improvement and labour saving provides an estimated 2.3 year payback.

It is proposed that the development of a prototype machine is justified.

The development of a standalone beef chining machine would provide significant benefits to the Australian beef processing industry.

2 Executive summary

Purpose of the research

There is an opportunity to automate the removal of the chine bone from the beef striploin, targeting yield improvement, labour saving and improved operator safety.

This project has focused on improving the accuracy of the striploin chine removal cell installed in previous project (Maunsell, 2024).

Target Audience

The main target audience is the Australasian beef processing industry

Benefit of the results

The results of the research will form a business case for proposing a prototype machine development project. With the desired outcome being the development of a prototype which will contribute value, by way of yield, labour saving and improved safety, to the Australasian beef industry.

The objectives and delivery of the project were:

- To enhance the sensing of the L4B01 Module – Striploin Chine Bone Removal concept prototype cell currently installed at an Australian processor to predict the curvature of the vertebrae
- Assess improvement in cut accuracy and hence measure if this produces a sufficient improvement in yield.
- Provide the steering group including Scott with the information required to develop a “final” module design that will result in a pre-commercial machine build offer being made to an Australian processor.

Where the accuracy was enhanced from the previous project (Maunsell, 2024) and minor improvement relative to the manual process was determined. A concept has been proposed such that cycle rate, capital cost and associated payback has been estimated.

Methodology employed

The background was largely researched from the previous project (Maunsell, 2024). Particularly the performance metrics and evaluation method.

The key metric is the error distance of the resultant chine cut from the ideal cut surface.

The yield benefit can be estimated using a weight per mm of error model.

Results/key findings

The project has implemented additional sensing means, both enhanced the resolution of the end view.

The quantity of meat retained on the product (“yield”) is determined from the mean error. Where the smaller error determines in increase in yield.

The performance of the strategy is a trade off between minimising the mean error from the ideal cut surface and the percentage of production that is in the “No Go” region. Where the “No Go” region is defined by each boning room. For this project the “No Go” region is defined at -2mm from the ideal cut surface. At -2mm the “bone bridges” adjacent to the button bones are readily broken, without the need for hammering.

To enable valid comparison, the mean for each strategy is adjusted to constrain the two standard deviations value at the “No Go” region.

For this project it is assumed that 0-rib and 2-rib short loins are produced in approximately equal numbers and therefore the strategy can be evaluated by considering the average of the performance metrics.

With these assumptions, referring to Table 1, it can be stated that the strategy in (Maunsell & McCrorie, 2024) resulted in a meat per head loss of 219g and the strategy in this project resulted in a gain of 58g, relative to the data collected on the manual process.

In project (Maunsell & McCrorie, 2024) and limited trials performed in (Kennedy, Maunsell, & Brennan, 2019) the accuracy of the CT scanner and robotic bandsaw approach was also estimated as per Table 1.

It can be stated that the key parameter to enable tuning the mean error setting and associated yield improvement is the reduction of standard deviation.

A significant advantage of the auto system is that the mean position of the cut depth can be tuned to, on balance, give the most preferred results.

During the CNN tuning it has been noted that there was minor improvement in the output as more data was added. However, given the nature of the CNN approach, it cannot be estimated as to high much improvement will occur when production volumes are available to tune with.

If it can be assumed that 0 and 2 rib shortloins are processed in equal numbers, the determined improvement in the error measurement is 1mm towards the spine.

The yield improvement can be estimated, using 28.88g/mm and $25/kg (Maunsell & McCrorie, 2024) at $0.72 per side, or $1.44 per head.

Given the variation in processing rates across the Australian industry, an example case has been established and the payback calculations provided.

For the example case of 300,000 head processed per year, the forecast payback is 2.3 years.

According to (Maunsell, Kennedy, & Dickie, 2018), nine beef processing sites in Australia have annual production volumes exceeding the 300,000-unit benchmark. These sites are expected to achieve a better payback than the example case. Notably, two sites exceed this volume by more than double, indicating that even two machines at these locations would outperform the example case in terms of payback.

There is an opportunity, depending on layout and available cycle time, to incorporate the loading of the short loin into a standalone machine with an upstream process, such as the tenderloin removal. This “upside” has not been included in the payback calculation.

It can also be expected that when a prototype machine is in production, further refinement of the vision analysis, particularly the CNN training, would improve the standard deviation of the cut error.

It is recommended that the current test rig, with the upgraded cameras, is further operated to increase the trialled sample size and obtain more extensive feedback from the trial site boning room. Further trials will increase the accuracy of establishing the processor motivation to purchase a machine utilising the developed process.

It is recommended that the development of a manually loaded, standalone machine is viable and would provide significant benefit to the Australasian beef processor industry.

Benefits to industry

The results of the research will form a business case for proposing a prototype machine development project. With the desired outcome being the development of a prototype which will contribute value, by way of yield, labour saving and improved safety, to the Australasian beef industry.

Future research/extension/adoption and recommendations

It is recommended that the current test rig, with the upgraded cameras, is further operated to increase the trialled sample size and obtain more extensive feedback from the trial site boning room.

It is recommended that the development of a manually loaded, standalone machine is viable and would provide significant benefit to the Australasian beef processor industry.

It is recommended that a concept be further developed and a project proposal put forward.

3 Introduction

There is an opportunity to automate the removal of the chine bone from the beef striploin, targeting yield improvement, labour saving and improved operator safety.

This project has focused on improving the accuracy of the striploin chine removal cell installed in previous project (Maunsell, 2024).

The main question was to determine whether the striploin chine cut path, being driven of skeleton features, can be established from the exterior surfaces. And given the achievable accuracy for the required capital cost, is it possible to develop a viable machine.

The main target audience is the Australasian beef processing industry

The results of the research will form a business case for proposing a prototype machine development project. With the desired outcome being the development of a prototype which will contribute value, by way of yield, labour saving and improved safety, to the Australasian beef industry.

4 Project objectives

- To enhance the sensing of the L4B01 Module – Striploin Chine Bone Removal concept prototype cell currently installed at an Australian processor to predict the curvature of the vertebrae

- Assess improvement in cut accuracy and hence measure if this produces a sufficient improvement in yield.
- Provide the steering group including Scott with the information required to develop a “final” module design that will result in a pre-commercial machine build offer being made to an Australian processor.

5 Methodology

5.1 Background

In previous project the automation of the beef process has been broken up into the modules shown in Figure 1.

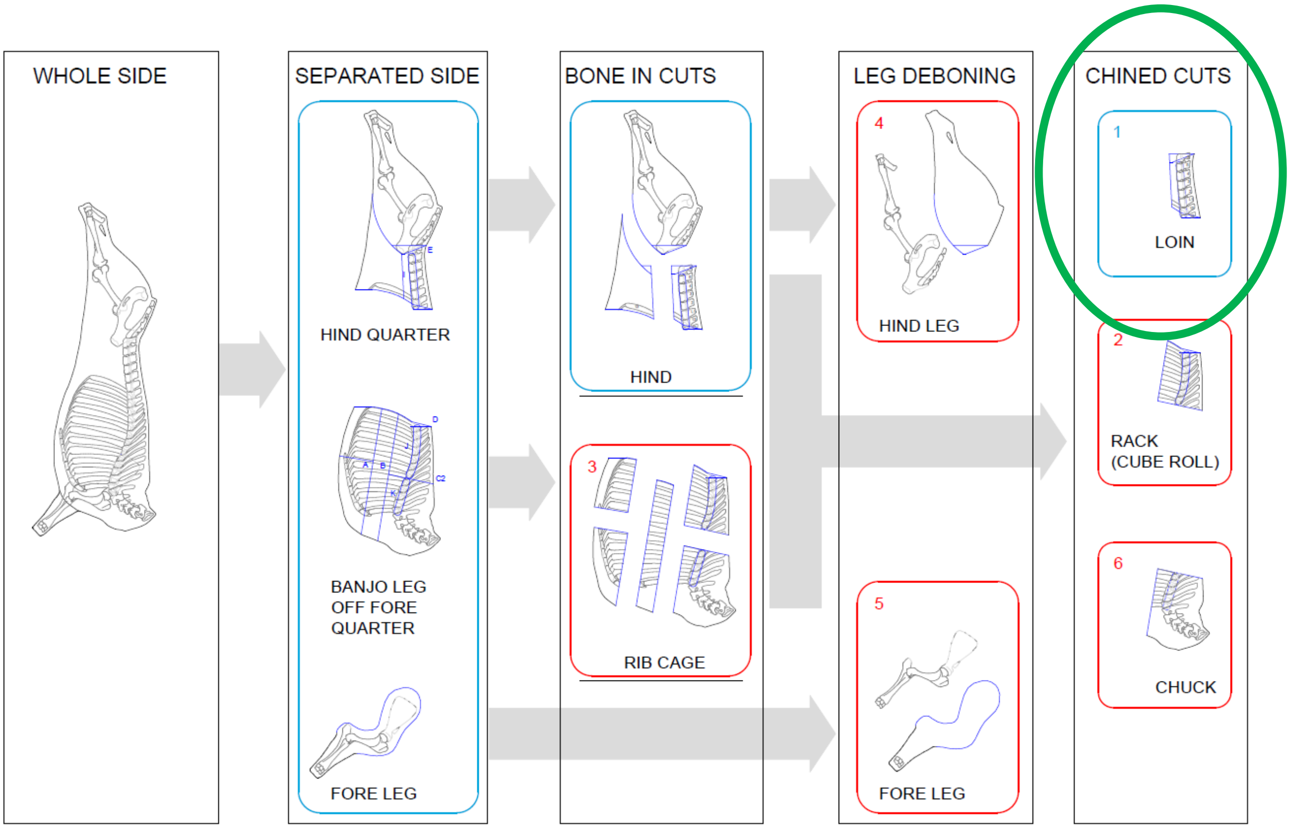

Figure 1: Automation of beef process modules

Where this project and its predecessor (Maunsell, 2024), is identified as CHINED CUTS, LOIN. (Marked in green)

The modules are:

- **Leap 4 Beef Module 01 Striploin chine bone removal (Current project)**
- Leap 4 Beef Module 02 Cube roll chine bone removal
- Leap 4 Beef Module 03 FQ Cuts processing (less banjo)

- Leap 4 Beef Module 04 HQ Leg bone removal from meat
- Leap 4 Beef Module 05 FQ Leg bone removal from meat
- Leap 4 Beef Module 06 Chuck chine removal

To assess the viability of a module it is necessary to understand

- the items produced
- the specifications that govern production of the produced items
- the production methods that are followed
- the value of produced items
- the possible benefits of automation in terms of labour requirements, yield and value

The key knowledge gap identified is as to whether measurement of the external surfaces of a beef short loin with an analysis system can be used to determine a chining cut path with beneficial accuracy.

To close this knowledge gap, the previous project, (Maunsell, 2024), was executed.

Project (Maunsell, 2024) developed a robotic chining test rig with vision and computing systems.

The test rig included (Figure 2):

- A developed clamp for securely hold the short loin spine, along the length of the product while leaving the working zone clear for the bandsaw.
- A bandsaw mounted on a robot where the path was able to be adapted from vision capture and analysis system. Where the vision system was two cameras viewing both ends of the short loin.

- Numerous trials were performed enabling measurement, training, development and final performance measurement for reporting.

Figure 2: (Maunsell, 2024) Robotic Chining Cell

5.1.1 The key metric for performance determination

The key metric for performance determination is a measurement from the cut plane to the “ideal” cut surface. This measurement is the error. The “ideal” cut surface is that surface where there is “zero” bridge (no meat and no bone) between buttons.

With this error metric, yield loss can be reasonably estimated. E.g. if you have 2mm valleys all the way across, you can multiply by a cut face area and density to determine mass of meat loss from the ideal.

The metrics recorded in the results table are:

Short loin ID	Number of ribs in short loin	Button# (For each of 7 buttons)		
		Split?	Dorsal Depth	Ventral Depth

5.1.1.1 Split metric definition

The split metric is a discrete measurement of cut quality. Defined in Figure 3 and Table 1

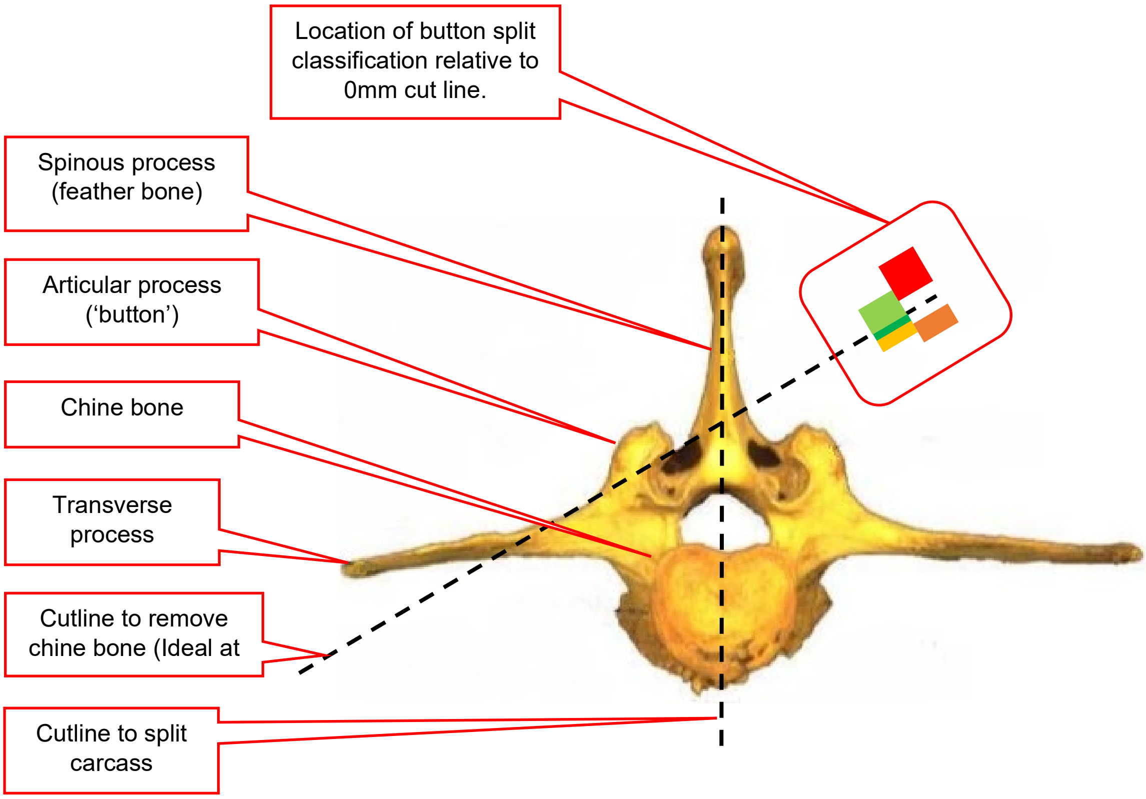

Figure 3: Split? metric definition

Button split?	Description	Min. (approx.)	Max. (approx.)	Yield
M	Missed – Button bone is not visible on the chine bone side of the bandsaw cut. There may be a streak of that connective tissue showing on the striploin side of the cut, but no bone so the blade has missed the button.	+10mm	+32mm	Lost, with not ability to recover as first cut is within eye meat.
Y	Yes – The button is clearly separated from the surrounding chine bone, feather bone and transverse process. Fatty, white connective tissue connecting the bones on the striploin side of the cut, whilst a very slight bridge showing on the chine side of the cut would typically indicate a 0mm measurement.	0mm	+18mm	Good to fair. Possible room for improvement.
PB	Partially broken – One side of the button may be separated, but the other side is still "bridged", but pressing on the connecting bridge with solid finger pressure will snap it. Typically, this type of bridge will indicate a -1mm or -2mm measurement.	-3mm	0mm	Maximum yield while button remains relatively easy to remove.
PBH	Partially broken (hammer) - One side of the button may be separated, but the other side is still "bridged", but pressing on the connecting bridge with solid finger pressure <u>will not</u> snap it. A good solid swing of the button hammer will break the bridge. Typically, this type of bridge will indicate a measurement between -3mm and -10mm.	-10mm	-3mm	Maximum yield, but buttons are difficult to remove and will require additional work. Could lead to lower throughput.
N	No – One or neither side of the button may be separated. Both or one side is still "bridged", but the hammer <u>will not</u> snap it.	-15mm	-5mm	Strip loin will need another cut, which may lead to further yield loss as the product no longer has the chine bone for rigidity.

Table 1: Split metric values definition

5.1.1.2 Dorsal and Ventral Depth metric definition (error measurement)

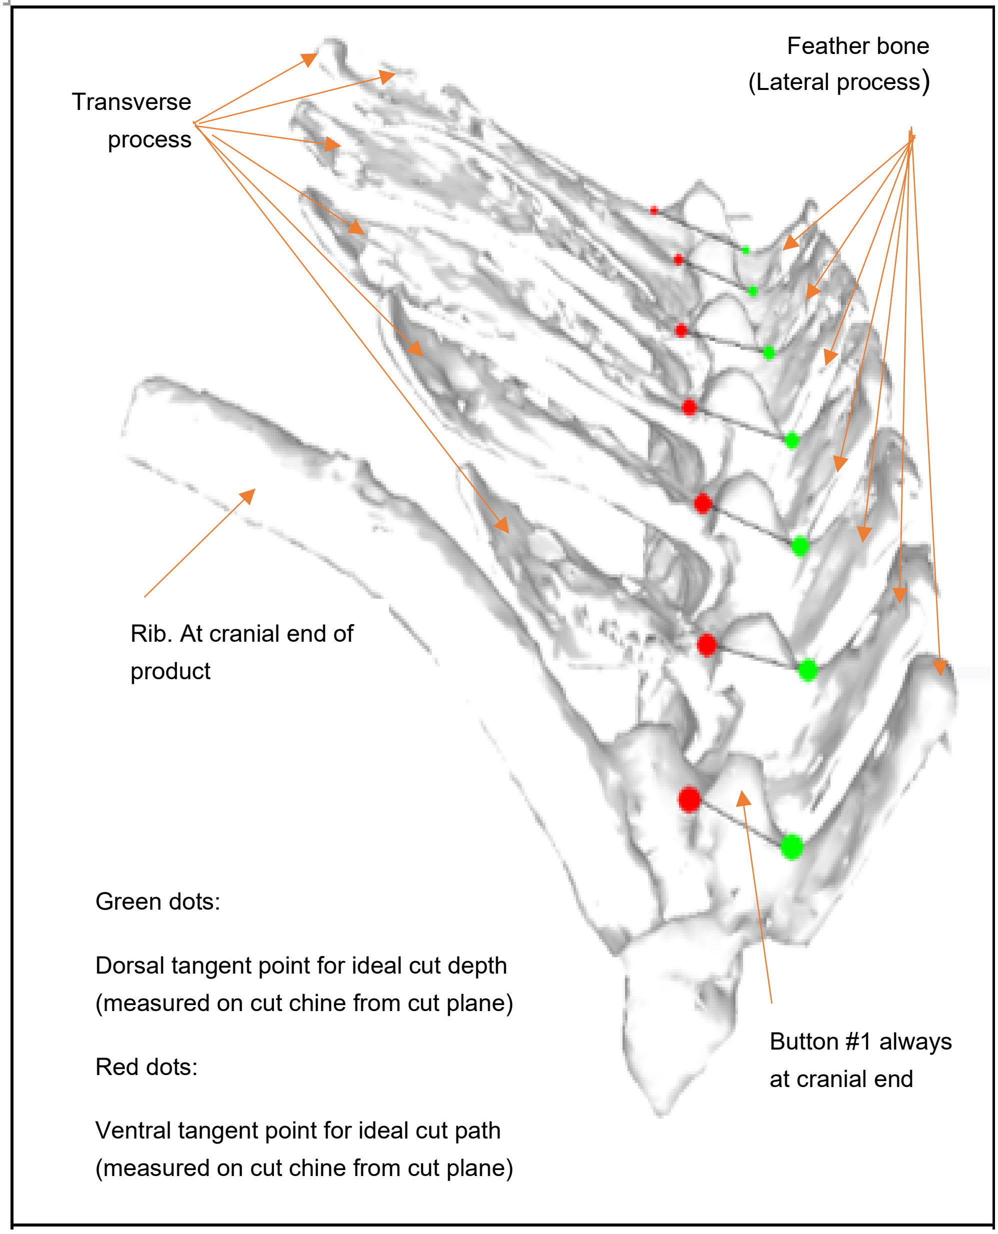

Figure 4: Striploin skeleton demonstrating the position for 0mm cut on dorsal and ventral sides

5.1.2 Strategy to realise benefit

The strategy to realise benefit is predominately from yield improvement. Labour saving is also targeted, but yield is the most significant gain and if yield is not improved, the labour-saving benefit is readily consumed.

The target is to achieve a standard deviation in cut accuracy that is less than or equal to the manual process. There is a significant product value differential across the cut, premium strip loin versus bone, therefore ideally with an improved standard deviation the cut mean position could be moved closer to the bone delivering a yield gain. The yield can be calculated from the manual process mean position minus the robotic position, multiplied by the nominal cross section, meat density and value per kg.

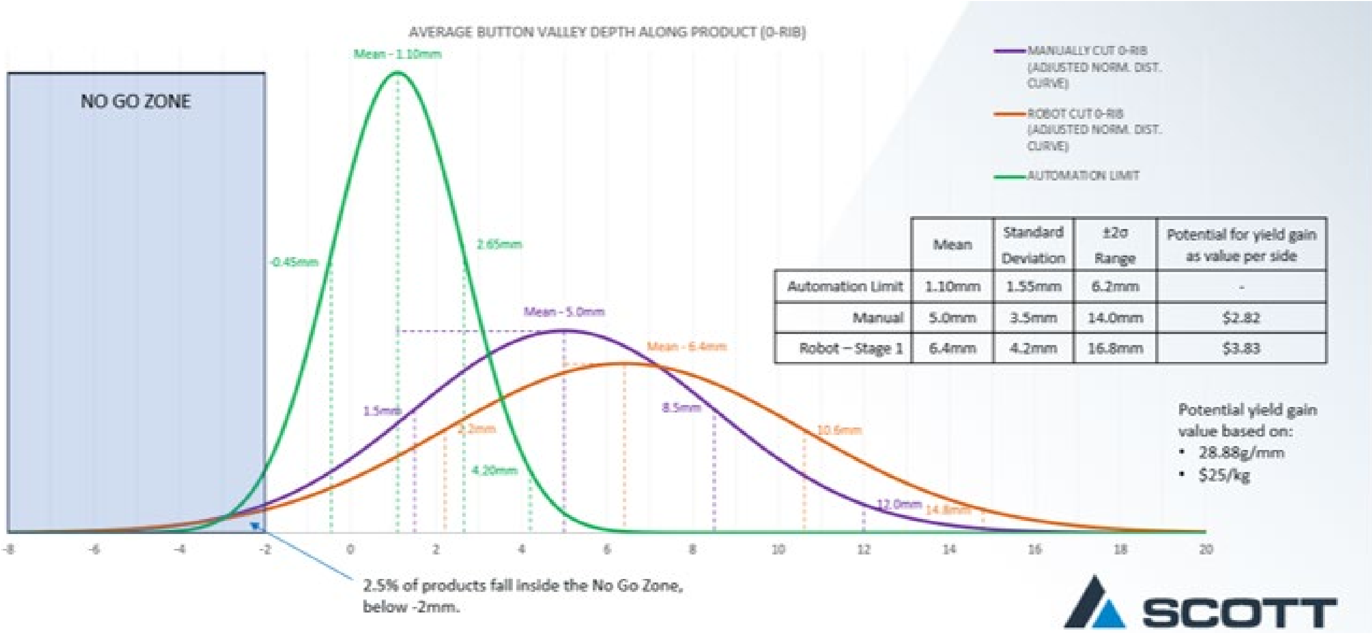

	Mean	Standard Deviation	±2σ Range	Potential for yield gain as value per side
Automation Limit	1.10mm	1.55mm	6.2mm	-
Manual	5.0mm	3.5mm	14.0mm	$2.82
Robot – Stage 1	6.4mm	4.2mm	16.8mm	$3.83

Figure 5: From (Maunsell 2024) demonstrating potential yield increases from increasing chining accuracy. Potential yield is calculated as $2.82 for a 0-rib product.

5.1.3 The results determined in (Maunsell, 2024) were:

- When aligned normal distribution curves for cut depth error are compared, the automatically cut product mean is between 1.0mm to 1.4mm more than the mean of manually cut product at an Australian processor.
- Using the current simple sensing means (cameras at each end), the mean of cut depth error is 5.3mm to 8.1mm more than the expected results from advanced sensing means (CT scanning).

5.1.4 The determined benefits to industry were:

Assuming the next stage development machine can improve on a manual cut product by 1mm for 0-rib product and 2mm for 2-rib product, this would result in a yield benefit of:

- $0.722 AUD per 0-rib striploin*
- $1.444 AUD per 2-rib striploin*

Using the above assumed benefits, the value of equipment utilising a 2-year payback would equate to: $991,744 AUD**

*	This benefit is plant and commercial solution dependant.
**	There may be further costs associated with equipment

5.2 Understand the past work and generate upgrade strategies

A brief internet search, review of (Maunsell, 2024) and the generation of strategies to provide the ultimate opportunity to demonstrate the determination of chine cut surface from external surface data.

5.3 Develop upgrade design

The development includes:

1. Vision hardware
2. Computer hardware
3. Image capture software
4. Image processing
5. Boning room, trials and associated feedback measurements
6. Trials and measuring

5.4 Yield measurement.

For the boneless short loin, the yield for both robotic and the manual boning process, could be measured and compared.

5.5 Propose commercial machine to establish an estimation of capital cost.

In concept form, propose an option for a commercial machine.

Using experience and comparisons with known machines, estimate a price of the proposal

6 Results

6.1 Understand the past work and generate upgrade strategies

In (Maunsell, 2024), the striploin chine machine was shown to work reliably using external sensing, but the cutting accuracy was notably worse than existing manual processing. Three main areas were determined to be important for improving this.

- Improving image data provided to CNN using better cameras and optics
- Improving CNN training and particularly data feedback
- In larger product, the cut surface tended to be more curved. In (Maunsell, 2024) the manual operators were determined to be aiming for the tenderloin attachment point on the striploin. To help with this, a side view was added and the robot cut path was modified to account for an additional point.

6.2 Develop upgrade design

6.2.1 Vision hardware

To improve vision data, the cameras and lenses were changed.

This improved sensor triples the total pixel count, with a 2.8x larger area. The longer lens focal length increases the usable area of the image. Combined with changing orientation, this increases the product area in view from 13% horizontal, 30% vertical the to 30% vertical, 45% horizontal for a large product. This further increases the useful information by approximately 3.5 times.

For this project, a primary goal was to improve the accuracy of the cut path.

The dataset for training the neural network was determined to be an important factor to improve. In (Maunsell, 2024), images and results were collected for a few hundred products. The images were used for input for the CNN, but the labelled cut positions from these images were from manual inspection instead of using the collected data. Using the measured feedback data is a much more accurate method of creating labels for images.

This information was generated for each product in (Maunsell, 2024). However, this was not set up to align these parts to ensure that the data generated matched the correct product. In particular, the feedback was written on paper. This required significant work to digitise the data and there was no way to ensure the images and cut positions matched with each other. Some products were analysed multiple times, feedback was recorded out of order or excluded from feedback.

An issue identified during boning room trials during (Maunsell, 2024) was there were sometimes several products lined up to be measured. This caused an issue that recorded feedback could be out of order compared with the cutting order if two chines were measured out of order.

To ensure data was easy to record and keep aligned, a data recording application was created. This was a web application hosted on the vision computer and accessed with a tablet which recorded the data for each product.

Algorithm 1: Method for recording data
1. On pressing button to start cut, a unique ID is generated. **2.** After cut, the feedback camera is swung around to take a picture of the cut product. A picture is taken by pressing a button on the tablet application. **3.** This adds an item to a queue on the tablet application to record the depths of the meat across the cut product. **4.** After recording data on the tablet, it is stored in the directory from the cut product.

The product ID is recorded in a database that associates a directory of images with the generated robot coordinates. The feedback is associated with the most recent analysed images. This method ensures that feedback data is digitally recorded and correctly associated with the ID from the same product.

The data recording application also shows the image taken from the feedback camera. This makes sure the cut product can be checked against the image to ensure the product being measured was correct. This was rarely necessary during testing.

Once these values are gathered, this can be used to generate corrected data. For a linear cut, two points are used. The positions of buttons are linearly interpolated. For an example of this calculation, a set of cut points are (500, 120) to (975, 100). The measured data is from the feedback tablet. Positive values mean the cut was too far away from the vertebrae, and vice versa[1].

Position	X (caudal-cranial) (mm)	Y (towards/away from chine)	Measured error (mm)	Adjusted Y (mm)
Start of cut [robot input 1]	500.0	120.0		115.2
Button 1 ventral	536.5	118.5	-2	116.5
Button 1 dorsal	573.1	116.9	-2	114.9
Button 2 ventral	609.6	115.4	-3	112.4
Button 2 dorsal	646.2	113.8	-6	107.8

[1] The bandsaw is a thin blade held at a constant angle. These Y positions are the position of the blade at Z=0. The Y values of the measured positions will be closer to the vertebra because they are located above the conveyor.

Button 3 ventral	682.7	112.3	-4	108.3
Button 3 dorsal	719.2	110.8	-5	105.8
Button 4 ventral	755.8	109.2	-5	104.2
Button 4 dorsal	792.3	107.7	-3	104.7
Button 5 ventral	828.8	106.2	2	108.2
Button 5 dorsal	865.4	104.6	1	105.6
Button 6 ventral	901.9	103.1	-3	100.1
Button 6 dorsal	938.5	101.5	2	103.5
End of cut [robot input 2]	975.0	100.0		100.0

Table 2: An example for calculating adjusted positions of button bridges. The Y values are adjusted by the measured error. N.B. In the actual system, the measurements are made perpendicular to the cut face. For simplification, these are treated as the error in Y instead.

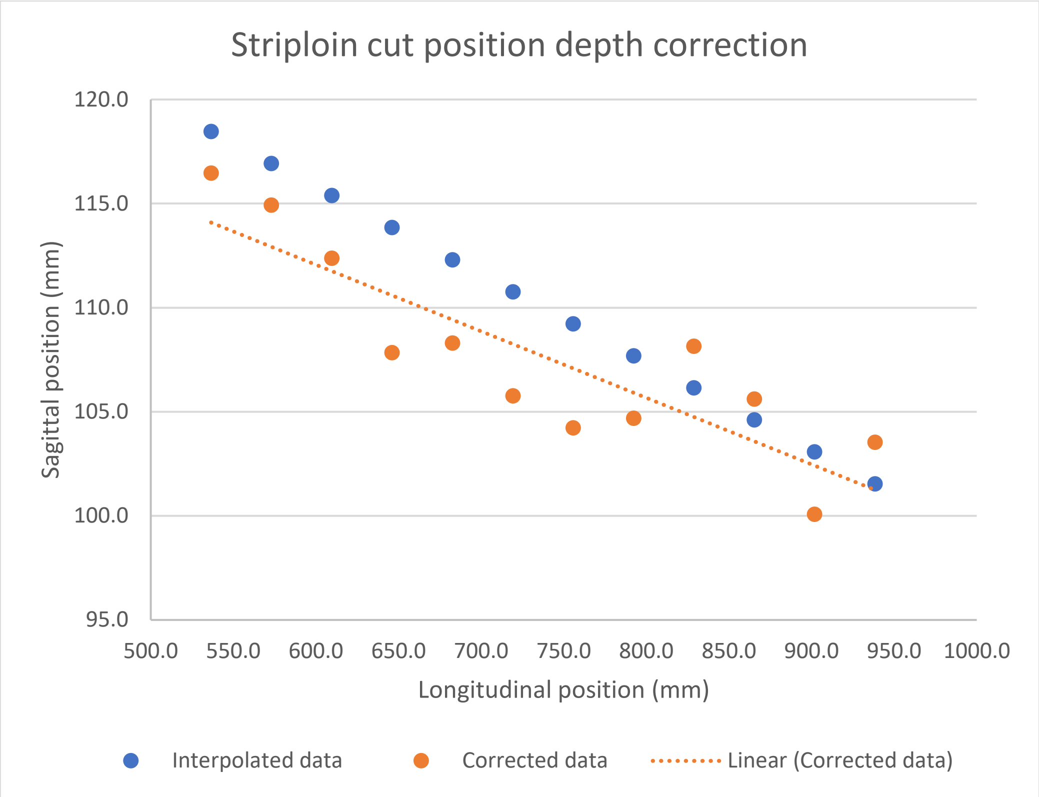

Figure 6: Original button points and points adjusted from feedback with line of best fit.

This generates a new set of cut co-ordinates (500, 115.2) and (975, 100.0). This feedback mechanism means if the original cut was inaccurate, the data provided to the CNN will be corrected to give a better cut informed by the feedback.

6.2.1.1 Computer hardware

The upgraded cameras use GMSL connections which are primarily used for vision systems in cars and robotics. This is designed for reliable connections up to 15m duplex connections transporting power and data. To facilitate these connections, the camera provider sells industrial PCs based on Nvidia Jetson which support 4 x ZED One connections. To support 6 ZED One cameras, two of these were used. Nvidia Jetson is an ARM based Linux computer. These factors significantly increased the development complexity.

6.2.1.2 Robot path

To accommodate a curved path, the robot cut code was rewritten to include a curve through the product. This presented difficulty due to the conveyor belt being used as an external axis and the complexity of ensuring the bandsaw path correctly aligned with the direction of the bandsaw blade.

6.2.1.3 Calibration

Converting between camera pixels and robot co-ordinates is dependent on intrinsic and extrinsic parameters. The intrinsic parameters are dependent on the camera sensor and the lens used. The extrinsic parameters depend on the relative position and orientation of the camera and the robot frames. Both are necessary to convert from camera co-ordinates to robot co-ordinates. The intrinsic parameters should be nearly identical for cameras with the same sensor and lens. The extrinsic parameters change if the camera or robot frame moves.

These calibration parameters can be derived from a series of camera co-ordinates paired with known points in the robot co-ordinate system. To facilitate this, a calibration end-effector was created, shown in Figure 7. This creates a target to be found at different robot positions. The process for finding the parameters are described in Algorithm 2.

Figure 7: Calibration end-effector in image

Algorithm 2: Automated calibration image capture **Input:** Robot positions **Output:** Images with ball end-effector at each position
1. For each position in a grid: a. Robot moves to new position b. The robot sends its position in robot co-ordinates and requests an image c. The vision PC takes images d. The vision PC sends a signal for the robot to proceed to the next position

6.3 Boning room, trials and associated feedback measurements

The CNN transforms a set of images into robot coordinates for a cut path. In (Maunsell, 2024) images were taken, and a cut line was manually marked on the images to estimate where to cut the product. This produced results in the correct region but was likely to be very limited. To address this, a method of easily connecting the measurements for the cut product was developed. In total 714 product were cut and results measured. The initial 300 were cut by choosing the cut lines on images, and the remaining 414 were cut using a CNN trained on the trialled images.

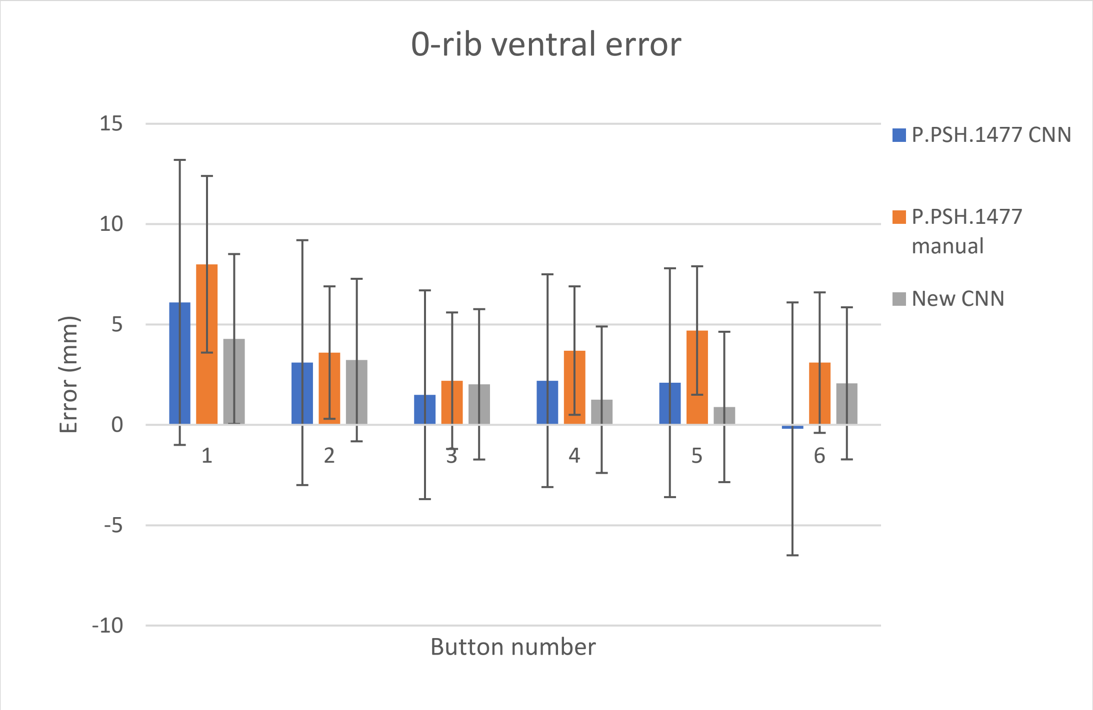

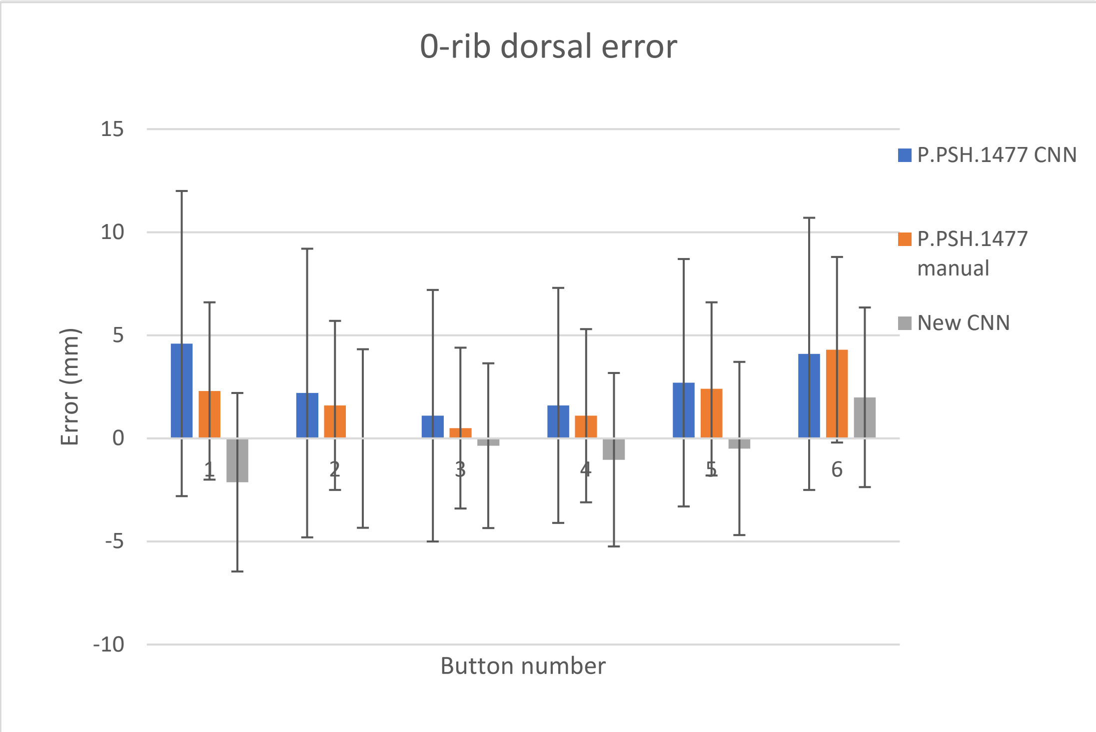

Figure 8: Measured error for 0-rib product for ventral and dorsal buttons. The error bars show 1 standard deviation from the mean. Assuming a normal distribution, 68% of the cut product will fall within these lines.

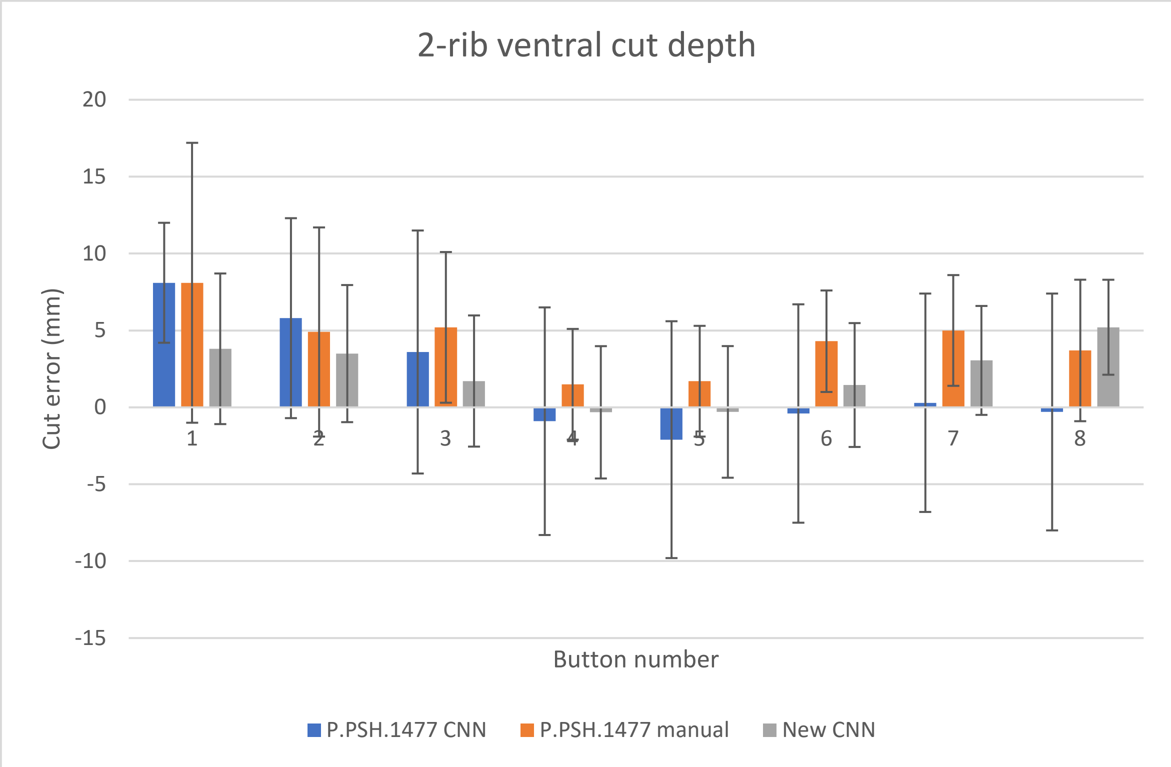

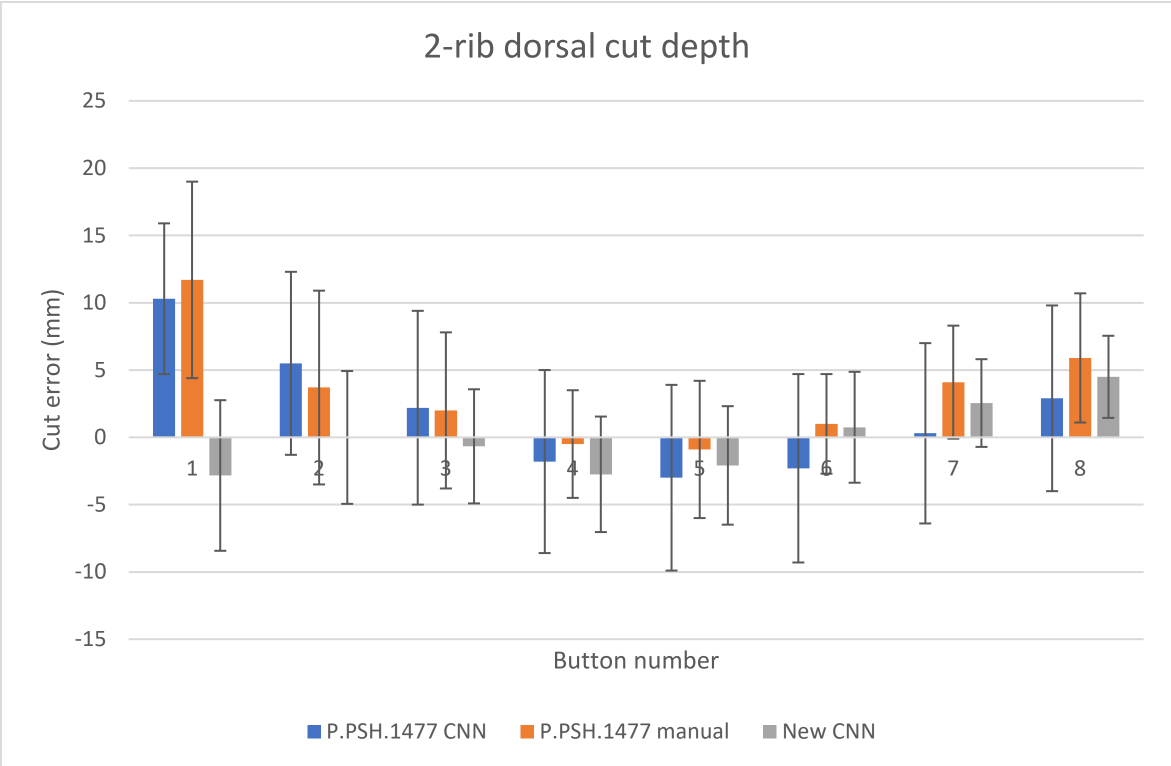

Figure 9: Measured error for 2-rib product for ventral and dorsal buttons. The error bars show 1 standard deviation from the mean. Assuming a normal distribution, 68% of the cut product will fall within these lines.

The last two tables give an overall picture of the results. For 0-rib, compared with the CNN developed for (Maunsell, 2024) on 0-rub product, the new CNN improved by an average of 11.9%, with particularly strong improvement on the second and final buttons. However, it remained worse than the manually cut benchmark by an average of 9.5%.

The improvement was more pronounced on 2-rib product. This was an average of 3.6% than the manual benchmark and was better by an average of 7.8% if cuts requiring a hammer are considered acceptable. However, this is heavily skewed by the final button. Excluding the final button lowers this to -13% and -1%. Another caveat is the number of manual 2-rib product sampled was very low.

Comparing the 0-rib CNN results and manual benchmark, the CNN typically had similar or lower rates of missing the buttons entirely, but much higher rates of cutting too deeply and either requiring a hammer or rework. This indicates the cut is biased towards the spine compared with the manual process.

This is further shown in Figure 8; the CNN mean dorsal depths are very close to zero or slightly negative. A cut depth between 0mm to -2mm is the ideal range, but if the mean is -2mm, half of the product will cut too deeply into the spine and require a hammer or rework.

As explained in (Maunsell, 2024), the mean distance can be directly changed with offsets. By adding a 4mm offset to the cut, the dorsal mean for the first button would increase from -2mm to 2mm. Assuming a normal distribution, this would decrease the number of product with more than 2mm bone on the first button to approximately 20%. This is a trade-off because this decreases the yield.

Controlling offsets provides a significant advantage over manual cutting because it allows the operator to move the cut closer or further from the spine, determining the balance between yield and rework. This means the important factor for comparing with manual cutting is how the data is distributed from the mean.

On 0-rib product, the automated cuts had a similar standard deviation to manually processed striploin on the dorsal side, and 0.4mm worse on the ventral side. On 2-rib product, the automated processing was significantly better. This has caveats that the sample size was small and the manual standard deviation is highly swayed by the first two buttons. Omitting the first two buttons lowers the 2-rib manual standard deviations to 4.0 and 4.7 which is still worse than the automated system.

Overall, this suggests the capability of the automated processing has significantly improved from the trials in (Maunsell, 2024), and now has a similar, but slightly worse precision as manual processing in 0-rib product, and slightly better for 2-rib product. After tuning offsets to minimise rework, the results should be similar.

6.3.1.1 Further at desk, CNN optimisation

After the boning room trials had been performed, the dataset was further analysed and refined. This involved looking through each image to check for irregularities such as image artifacts or incorrect marking.

A CNN to find a function that minimises a value referred to as *loss*. This depends on the system, but in general the loss value should be lower if the network outputs values close to the answer, and higher if the output is further away. The choice of loss function is an important factor in determining what the network will converge towards.

For this machine, the three sides of the product have separate CNNs that each take a single image and produce a cut co-ordinate. Since the cameras are in stereo pairs, the results from these are averaged. The cut co-ordinates are created from making a line of best fit for minimising the distance to estimated co-ordinates of the button valleys. By minimising the distance to the ideal cut points, this will make the cut more accurate.

The loss used for this system is the mean square error. This is a common metric with a similar formula to the standard deviation. To get a measure of how the amount of data affects the accuracy of the CNN, a CNN was trained using different numbers of input points, shown in Figure 10. As the number of data points increases, the loss decreases. If the system is close to its maximum accuracy, the CNN loss will start to plateau, and the accuracy will no longer increase. On this graph the line is still decreasing at the end side of the graph, this implies the accuracy of the system will improve with additional data.

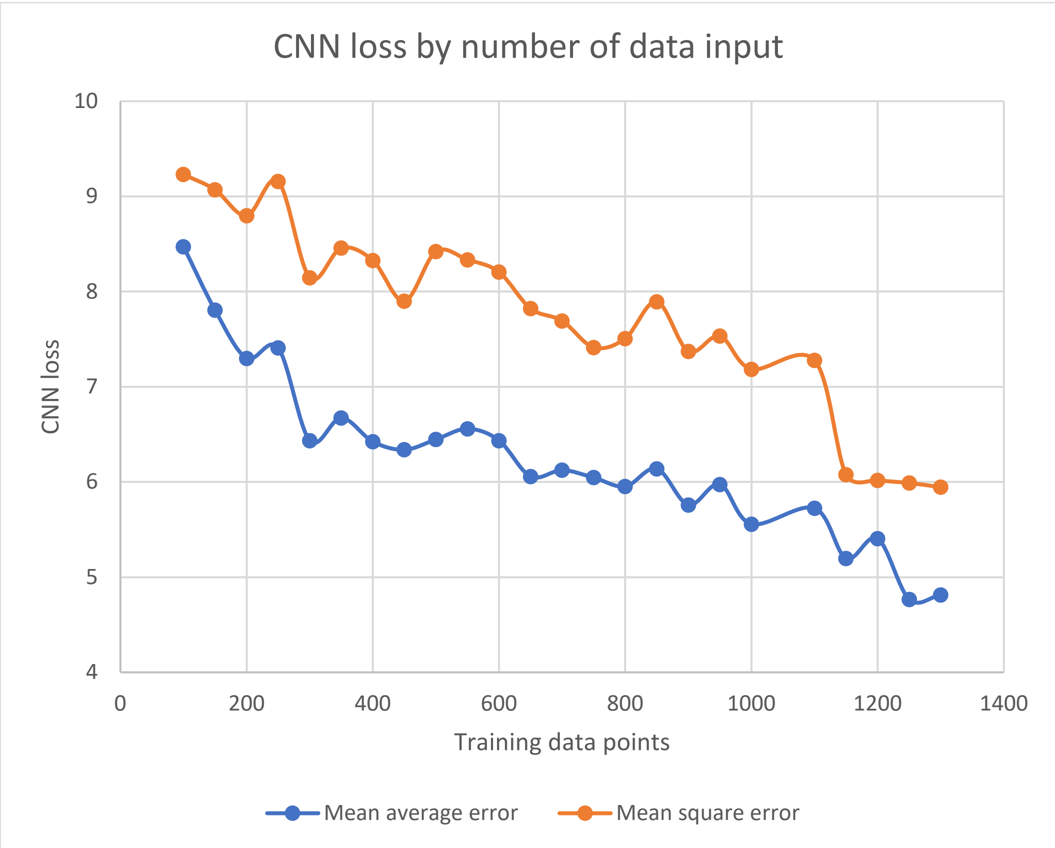

Figure 10: Comparison of mean average error and mean square error for different numbers of data points. The error values are calculated using the data points that have been taken outside the input data.

6.3.2 Australian processor yield measurement.

For the boneless short loin, the yield for both robotic and the manual boning process, was not performed in this project, but ideally would be monitored by the room in subsequent production.

6.3.3 Proposed commercial machine and estimation of capital cost.

It is proposed that the preferred concept for a commercial machine would cost $1.5M – $2M.

6.3.4 Determination of yield gains and other benefits for the return on investment calculation

The error improvement is -1mm for the average of 0 or 2 rib short loins.

The yield improvement can be estimated, using 28.88g/mm and $25/kg (Maunsell, 2024) at $0.72 per side, $1.44 per head.

For the example case of 300,000 head processed per year, the forecast payback is 2.3 years.

There is an opportunity, depending on layout and available cycle time, to incorporate the loading of the short loin into a standalone machine with an upstream process, such as the tenderloin removal. This “upside” has not been included in the payback calculation.

7 Discussion

The images captured from the new system are of higher resolution that the system in (Maunsell, 2024).

The number of pixels for the product in view has increased to approximately 1000 x 1300 from 300 x 400. This is 11x the information for the product in view.

An additional pair of cameras have been implemented to give the side view:

Figure 11: Side view cropped from previous image

These aspects have improved the overall quality of data produced.

Having better and greater quantity vision data combined with a more complex robot path has significantly increased the accuracy of the CNN system. Using the error metric that has been establish for both (Maunsell, 2024) and this report the summary comparison is:

		Previous manual	Previous auto	New auto
0-rib	Product count	72	218	325
	Good split	81.9%	60.6%	72.5%
	Including PBH	90.3%	75.9%	85.1%
	Mean error	3.1mm	2.5mm	1.0mm
	Standard deviation	3.9mm	6.2mm	4.1mm
	Mean absolute error	3.9mm	5.1mm	3.4mm
2-rib	Product count	15	64	115
	Good split	65.0%	47.1%	61.4%
	Including PBH	74.2%	63.6%	81.9%
	Mean error	3.8mm	1.8mm	1.1mm
	Standard deviation	5.4mm	6.9mm	4.2mm
	Mean absolute error	5.0mm	6.3mm	3.9mm

Table 3: Comparison of metrics for cut quality for (Maunsell, 2024) and new tests. The good split percentage is the number of splits that are either 'Y' or 'PB'. A separate measure is given to include 'PBH' splits. Mean error is the combined dorsal and ventral average depth from a 0mm cut including negative values if there is bone left on the chine. Mean absolute error is the same measure except leaving bone is also counted as a positive error. The mean error and standard deviation give a distribution of where the cut sits relative to a 0mm cut. The absolute error shows the average distance away from a 0mm cut.

For 0-rib:

The accuracy of buttons that have been correctly split has increased by an average of 12% compared with (Maunsell, 2024) and the standard deviation decreased by 34%. This is a significant improvement in the main metrics measured.

Compared with the manual benchmark, the percentage of well split buttons was 9% lower or 5% if PBH is included. The standard deviation of the upgraded system was very similar to the manual process.

The automated system cut an average of 2.1mm closer to the chine than the manual process and 1.5mm closer than the previous CNN. This is especially pronounced at the caudal end. This would leave greater yield on the chine but require significant rework due to bone left on the product. Because this can be tuned using offsets, this is less important to compare than the standard deviations.

Including an offset to bias the automated cut away from the chine will increase the percentage of buttons that are correctly split. From the model in Section 6.2.1, tuning offsets on the first button was projected to increase the split accuracy by 5%. This would raise the accuracy to near the manual benchmark.

For 2-rib:

The 2-rib data is a similar picture. The percentage of correctly split product has improved by 14% and the standard deviation has decreased by 39% compared with (Maunsell, 2024). This is a significant improvement in the main metrics measured.

Compared with the manual benchmark, there is a caveat that the sample size measured for a benchmark was only 15 2-rib product measured. The percentage of well split buttons was 3.5% lower but 6.7% more accurate if PBH is included. The standard deviation of the upgraded system was very similar to the manual process.

The automated system cut an average of 2.7mm closer to the chine than the manual process and 0.7mm closer than the previous CNN.

Based on these values, the automated system performs similarly or better than the manual benchmark taken from (Maunsell, 2024). Based on these results, the automated system should produce higher yield and similar accuracy. Without any offset the automated system is comparatively biased towards the chine, so more products will require rework. As with 0-rib, tuning offsets should further increase accuracy.

8 Conclusions

The project has implemented additional sensing means; both enhanced the resolution of the end view cameras and added a side view camera. The cameras are implemented as a stereo pair.

A CNN strategy was implemented.

The quantity of meat retained on the product (“yield”) is determined from the mean error. Where the smaller error determines in increase in yield.

The performance of the strategy is a trade-off between minimising the mean error from the ideal cut surface and the percentage of production that is in the “No Go” region. Where the “No Go” region is defined by each boning room. For this project the “No Go” region is defined at -2mm from the ideal cut surface. At -2mm the “bone bridges” adjacent to the button bones are readily broken, without the need for hammering.

To enable valid comparison, the mean for each strategy is adjusted to constrain the two standard deviations value at the “No Go” region.

For this project it is assumed that 0-rib and 2-rib short loins are produced in approximately equal numbers and therefore the strategy can be evaluated by considering the average of the performance metrics.

With these assumptions, it can be stated that the strategy in (Maunsell & McCrorie, 2024) resulted in a meat per head loss of 219g and the strategy in this project resulted in a gain of 58g, relative to the data collected on the manual process.

In project (Maunsell & McCrorie, 2024) and limited trials performed in (Kennedy, Maunsell, & Brennan, 2019) the accuracy of the CT scanner and robotic bandsaw approach was also estimated.

It can be stated that the key parameter to enable tuning the mean error setting and associated yield improvement is the reduction of standard deviation.

A significant advantage of the auto system is that the mean position of the cut depth can be tuned to, on balance, give the most preferred results.

During the CNN tuning it has been noted that there was minor improvement in the output as more data was added. However, given the nature of the CNN approach, it cannot be estimated as to high much improvement will occur when production volumes are available to tune with.

The determined improvement in the error measurement is 1mm towards the spine.

The yield improvement can be estimated, using 28.88g/mm and $25/kg (Maunsell & McCrorie, 2024) at $0.72 per side, or $1.44 per head.

Given the variation in processing rates across the Australian industry, an example case has been established and the payback calculations provided.

For the example case of 300,000 head processed per year, the forecast payback is 2.3 years.

According to (Maunsell, Kennedy, & Dickie, 2018), nine beef processing sites in Australia have annual production volumes exceeding the 300,000-unit benchmark. These sites are expected to achieve a better payback than the example case. Notably, two sites exceed this volume by more than double, indicating that even two machines at these locations would outperform the example case in terms of payback.

There is an opportunity, depending on layout and available cycle time, to incorporate the loading of the short loin into a standalone machine with an upstream process, such as the tenderloin removal. This “upside” has not been included in the payback calculation.

It can also be expected that when a prototype machine is in production, further refinement of the vision analysis, particularly the CNN training, would improve the standard deviation of the cut error.

It is recommended that the current test rig, with the upgraded cameras, is further operated to increase the trialled sample size and obtain more extensive feedback from the trial site boning room. Further trials will increase the accuracy of establishing the processor motivation to purchase a machine utilising the developed process.

It is recommended that the development of a manually loaded, standalone machine is viable and would provide significant benefit to the Australasian beef processor industry.

9 Recommendations

It is recommended that further work is undertaken to establish market demand for a machine utilising the developed process.

There is an opportunity to process more products through the as developed test rig and formalise processor benefit from the Australian processor point of view.

It is recommended that the development of a manually loaded, standalone machine is viable and would provide significant benefit to the Australasian beef processor industry.

It is recommended that a concept be developed and a project proposal put forward.

10 Project outputs

Outputs (tangible deliverables) delivered during the project include:

- Technical reports for milestone 1, summarising options evaluated, results from trials and recommended future paths of research
- Data has been collected for the various research, alternative solution evaluation and experimental activities and analysis performed. Presented in the technical reports in various tabular and graphical formats.
- The upgraded cameras are on the test rig at an Australian processor and the rig can be run when required.

11 Bibliography

AUS-MEAT. (2005). *Handbook of Australian Meat* (7th Edition ed.). Tingalpa DC: AUS-MEAT.

Kennedy, G., Maunsell, S., & Brennan, J. (2019). *P.PSH.0893 LEAP 4 Beef Sub project 1 - Automated Chine, button and fat trim proof of concept for the Striploin and Cube Roll.* Sydney: MLA.

Maunsell, S., & McCrorie, K. (2024). *P.PSH.1477 - LEAP 4 Beef Automated Cutting Systems - Scott pre-production module 1 prototypes.* Sydney: MLA.

Thomson, A. (2024). *2024-1088 Leap4Beef Project 0 (stage1) for various modules.* Sydney: AMPC.

12 Appendices

-