Final report

Deep Learning Vision for Label and Red Meat Cut Identification

Project code:	MLA Project Code = P.PSH.1504
	AMPC Project Code = 2024-1073
Prepared by:	Suvir Salins Retail Ready Operations Australia Pty. Ltd.
Date published:	30 May 2025

PUBLISHED BY
Meat & Livestock Australia Limited
PO Box 1961
NORTH SYDNEY NSW 2059

Australian Meat Processor Corporation
PO Box 6418
NORTH SYDNEY NSW 2059

This is a joint MLA Donor Company and AMPC funded project.

Meat & Livestock Australia acknowledges the matching funds provided by the Australian Government and contributions from the Australian Meat Processor Corporation to support the research and development detailed in this publication.

Abstract

This project has successfully developed the first off-the-shelf deep learning smart camera system for red meat label and meat cut and primal inspections. Utilising artificial intelligence, the example based deep learning algorithms are able to make more human-like decisions for difficult vision tasks that the traditional rule-based algorithms can't manage. This vision system has been named Deep Vision.

Around the core Deep Vision component, an automated reject system with locked reject bin and conveyor system was added, to automatically reject packs which are deemed to be failures.

A touch screen based interface was also developed to enable easy user selections for order management and Deep Vision management which are all fully audit logged.

Both the reject system and audit log make the system Quality Control Point (QCP) compliant.

The system's hardware design ensured ambient lighting variations do not affect the image quality. Internal lighting provides consistent image quality which uses a diffuse lighting model to capture images with even exposure so that high gloss labels or undulating packaging surfaces, such as with vacuum skin packaging (VSP), do not interact with the lighting in a way which obscures the label or meat cut being inspected.

The system is able to integrate with Enterprise Resource Planning (ERP) applications such as Innova via the use of common product codes and SQL databases. In this way production operators are able to make single order selections within the ERP which are then pushed down to production equipment, including Deep Vision. This centralised control alleviates mistakes caused when individual machines require floor operators to make selections.

Deep Vision has achieved impressive performance results with false positives at only 0.29% or lower, compared to at least 3% for the current vision systems employed at Coles, RROA. Additionally the accuracy of the system is exceptional with the true positives and false negatives being 100% and 0% respectively. Deep Vision is much less prone to errors, when compared to the current vision system, for instances where labels are slightly curled, rippled or looped as with VSP labels.

Deep Vision's ability to specify and reject labels based on position and orientation is both simple and powerful, allowing the business to tightly control the specification of what is acceptable.

Deep Vision's ability to inspect 2D labels, Chinese characters and primals makes the system future-proof, as the industry transitions to 2D codes, and to provide benefits for both exporters and primary processors, additional to those of secondary processors.

Finally Deep Vision's ability to also inspect secondary label characteristics, such as promotional pricing, and meat cut types, allows the business to control most aspects of product packaging to maximise quality and minimise the chance of product recalls.

The Deep Vision project has shown that significant performance can be achieved by a smart-camera based deep learning system above those of rule-based vision algorithms, which typically was only achievable by more expensive high-end PC-based vision systems. The Deep Vision alternative to traditional deep learning systems allows for simpler hardware at lower cost which is also less prone to cyberattacks. Additionally the Deep Vision achieves these performance benefits with less training images and the ability for in-house staff to perform the training rather than vendor experts, which again saves considerable time and cost.

Executive summary

Background

The focus of this research project was to address the question of whether a smart-camera based AI deep-learning vision system can outperform rule-based vision algorithms or achieve similar results to traditional high-end PC based deep learning vision systems.

Secondary and Primary processors, whether for domestic or export markets, require robust vision systems which can handle the various product packaging formats, labels and products being produced. All these target audiences aim to achieve the highest quality product for their customers while ensuring product recalls are minimised. With the constant pressure to add more information to labels, such as allergen information for compliance, or value-add information such as cooking instructions, or support for foreign languages or new coding formats such as 2D codes, the risk for labelling or packaging mistakes has risen.

With favourable results having been achieved for this research project, there will be significant benefits to the red meat industry in the form of an easily implemented vision system with more human-like capabilities to handle difficult visual inspection tasks in a format which is simpler and thus more affordable than traditional AI systems, and more easily managed with in-house staff, rather than vendor experts, while still achieving high performance levels.

Objectives

The following objectives were achieved at the conclusion of this project.

- Smart-camera based deep learning algorithms which inspect labels and meat cut types
- Automated reject system with soft fall to limit pack damage and locked bin for Quality Control Point compliance
- Centralised order management system with integration to ERP system
- Quality Control Point compliance application with full audit log of all relevant system selections, changes and product images via login access control for specific authorised users
- High performance levels achieved for false positive, true positive and false negative metrics when compared to traditional rule-based vision systems
- An AI vision system which can be managed with trained, in-house staff rather than expert vendor support for adding new products or fine-tuning existing products
- Local vendor support
- An AI vision system which doesn't require 1000's of training images to achieve exceptional performance levels, with only 5-12 trained images required for each product
- An AI vision system which is also capable of primal inspections, 2D code reading and Chinese character recognition
- A system which ensures high levels of quality for customers while reducing the risk of recalls

Methodology

The project was broken into 5 milestones with the following methodologies employed in each.

Milestone 1

- Develop the project plan, specification and design of the system

- Form the project steering group
- Develop the trial plans
- Procure the hardware
- Produce a milestone report of the project results

Milestone 2

- Based on the system specification develop deep-learning algorithms and vision software to inspect packaging labels, including main and secondary labels and beef and lamb cut types.
- Develop the algorithms such that processing and rejection signals can handle a pack rate of 120 packs per minute
- Trial the smart-camera using temporary mounting on a RROA production line for 1 week to determine the feasibility of the algorithms
- Produce a milestone report of the project results

Milestone 3

- Partner with a Primary processor to employ the meat cut deep-learning algorithms to inspect primals to determine their types
- Develop OCR and other software tools within the smart-camera to recognise Chinese characters for export label inspection
- Using the smart-camera with temporary mounting, trial the algorithms with static primals in a lab environment and on a production line to determine the feasibility of the algorithms
- Produce a milestone report of the project results

Milestone 4

- Develop the electrical and mechanical design of a production system to implement the smart-camera on a production line
- Build the production system
- Develop the software for the system interface
- Produce a milestone report of the project results

Milestone 5

- Design and develop the software for full quality control point compliance to record all system selections and changes against specific authorised personnel
- Design the integration methods between the vision system and RROA's ERP system Marel Innova to allow orders to be automatically pushed from Innova to the vision system
- Trial the systems capabilities using offline 2D codes (not on the production line) to ensure future capability when the industry transitions to 2D codes.
- Produce a milestone report of the project results

Milestone 6

- Combine the results from each milestone and report on the final results of the research and development project

Results/key findings

The result of the research and development project was the successful development of a production system which employed a Cognex smart-camera with built-in artificial intelligence deep-learning algorithms which outperformed traditional rule-based vision algorithms and performed as well as PC-based artificial intelligence algorithms, but without the need for vendor experts to add new products to the system, nor large data sets of images to train the system. The system's performance showed a low (<0.29%) false positive rate, while maintaining a 100% true positive and 0% false negative rate.

Additionally the system was developed with an interface that allowed for easy order management in a stand-alone mode as well as when integrated with an ERP system.

Finally the system delivered a quality control point compliant system, recording all system selections and changes against specific authorised users to maximise product quality while ensuring recalls and their associated losses and costs were minimised.

Benefits to industry

With favourable results having been achieved for this research project, there will be significant benefits to the red meat industry in the form of an easily implemented vision system with more human-like capabilities to handle difficult visual inspection tasks in a format which is simpler and thus more affordable than traditional AI systems, and more easily managed with in-house staff, rather than vendor experts, while still achieving high performance levels.

These benefits are applicable across the red meat industry from Primary to Secondary Processors, whether for domestic or export products, in label and meat inspection tasks.

Given the system's ability to take any type of image and learn, the benefits are only limited by the scope of the industry's imagination for applications that could benefit from the system's capabilities.

Future research and recommendations

This project recommends that the industry adopt this technology for the following applications.

- Primary processor primal identification for sorting primals or ensuring correct primals are packed in cartons/crates.
- Primary and secondary label verification to ensure the correct label is applied to product including product name (pre-printed or printed), price, best before date, 1D/2D codes and correct label position and orientation.
- Export label verification to ensure foreign language text is correct
- Secondary meat cut verification to ensure the meat cut packed, matches the label applied.

The project suggests the following future research projects.

- Meat specification verification to inspect for characteristics such as overall dimensions, fat thickness, and marbling.
- Meat grading – carcase, portions
- Packaging seal integrity looking for damage or contamination
- Allergen or cooking suggestion text verification
- Meat orientation and size verification for integration with pick and place automation for packaging
- Animal welfare monitoring – detect variations in posture, gait, injuries

Table of contents

1. Background

Current machine vision technology, focused on label inspection, utilises rule-based algorithms to determine whether an inspection passes or fails. Variations such as lighting, colour, material type, orientation, position, font type, smudges and others, can easily cause rule-based vision algorithms to generate false rejects. Artificial Intelligence (AI) based vision systems have long been touted as the solution to these limitations, giving machine vision a more human-like capability of discerning true failures from false. Up until now, AI-based deep learning vision systems have been extremely complex both in hardware as well as software configuration, requiring high levels of expertise to generate meaningful results.

This project will develop the first off-the-shelf deep learning smart camera for Beef and Lamb meat inspections to produce a truly holistic system which isn't limited to just standard label inspections, but can also recognise meat cuts. Utilising an off-the-shelf smart camera with built-in deep learning capabilities enhances the security of the system as it is not reliant on PC infrastructure which are prone to cyberattacks. Intrinsic deep learning functionality within the smart camera results in a very easy to configure system which can be taught to recognise any object with a fairly small dataset. This means that time to productive results is reduced, the number of people able to configure the system with limited specialised knowledge is increased, and the number of potential use cases is very large indeed.

Coles RROA will develop the pilot system on the Line 4 VSP/MAP product line which represents a challenging packaging format with C-wrap, D-wrap and DIGI labels. In particular, labels can undulate over the VSP pack causing warped images and lighting challenges. Even with the MAP format, the high gloss surfaces can cause challenging lighting conditions for traditional vision systems. This pilot system will also investigate the benefits for primal identification and QR code reading, which are both wider industry benefits for this technology.

2. Objectives

The overall objective is to evaluate the Cognex AI camera's ability to read labels and barcodes to determine errors and misprint while also assessing its capability to recognise and identify beef and lamb cuts and primals.

This project implemented on the Line 4 VSP line will deliver the following outcomes:

- The ability to identify and recognise Beef and Lamb meat cuts and primals.
- Label main text and date code reading via OCR without the requirement to train fonts
- 100% Label inspection determining if any labels do not conform to the selected order characteristics
- Barcode reading to determine if the barcode is legible for stores and that the barcode matches the selected order
- Security label inspection to ensure a security label is present
- Promo label inspection to ensure the correct promo labels is present for the selected order
- Historical log of images for all inspections for at least the period required for best before date compliance
- Historical log of the results of all inspections for the compliance period
- Live visual results of each inspection on the HMI
- Digital output indicating a failed inspection for the current pack being inspected
- The ability to track and reject packs that have failed inspections
- The ability to interface with the order management system
- The ability to log parameter changes for audit compliance requirements

3. Methodology

The following method and process steps will be applied:

1. Project planning, design and equipment supplied [Milestone 1]

Conduct start-up meeting with Coles RROA project team, MLA, AMPC, Form project steering group, Trial plans, Hardware rental. Submit Milestone Report, including trial plans to steering group for approval

2. Beef and Lamb meat cut recognition deep-learning development and software development [Milestone 2]

Label recognition (deep-learning), Date Code Reading (OCR), Barcode Reading, Promotional label (Deep-learning), Security Label Presence (Deep-learning), Beef and Lamb Portion Cut Recognition (Deep-learning), Speed requirements tested and satisfied (120 packs per minute). Submit Milestone Report, including trial plans to steering group for approval.

3. Beef and Lamb Primal recognition deep-learning development [Milestone 3]

In collaboration with agreed primary processing partner capability of technology to identify Beef and Lamb Meat Primals tested and associated software developed.

Submit Milestone Report, including trial plans to steering group for approval

4. Development of the reject handling system [Milestone 4]

Software development, Design and drafting, Mechanical and electrical works. Submit Milestone Report, including trial plans to steering group for approval

5. Development of the audit compliance module [Milestone 5]

Integration with MES for system parameter management, Integration with MES for order management. Submit Milestone Report, including trial plans to steering group for approval

6. Final report [Milestone 6]

A RROA confidential version of the final Report to include system performance testing, cost benefit and internal RROA business case and anticipated industry outcome.

RROA to provide a non-confidential version that contains aggregated and de-identifiable data into a case study report for wider industry dissemination.

A Coles RROA approved industry report to include potential applications & benefits. Final public report to provide recommendations to wider industry with video footage demonstrating the technology to share with industry members. Final Reports (confidential & public) submitted & approved by steering group.

4. Results

4.1 Project planning, design and equipment supplied [Milestone 1]

4.1.1 Start-up meeting

The project steering committee met on the 13th of June 2024 to kick off the project.

4.1.2 Trial plans

Please see Appendix 1 for the Timeline Chart for the project milestones.

This project is comprised of six milestones over a period of 10 months from June 2024 to April 2025. Phase 2 projects and overall end date scheduled for October 2025.

4.1.3 Hardware rental

Coles RROA are working with Automation Systems and Controls, the platinum distributor and integrator for Cognex vision systems, to provide the vision hardware for the pilot system. The system will be comprised of a conveyor, booth/enclosure, camera/lighting hardware, HMI/PLC.

4.1.4 Detailed vision specification

Coles RROA developed a detailed vision specification. This document details the label inspection, cut identification, throughput, image storage, reject handling, product selection, program creation, parameter fine-tuning, reporting functionality, external integration, line selection/location, network architecture, hardware and master data requirements and scope or works for suppliers and RROA.

4.1.5 Project KPIs

See Appendix 2 for the project KPI measures and baselines. At the completion of each Financial Year reporting half, these measures will be reported to and assessed by the project steering group.

4.2 Beef and Lamb meat cut recognition deep-learning development and software development [Milestone 2]

4.2.1 Beef and Lamb Cut recognition deep-learning development and software development

The vendor trialled various strategies of deep-learning, OCR and pattern classifications to achieve the optimum results for the following inspection tasks.

- Main Label text recognition (deep-learning)
- Date Code Reading (OCR)
- Barcode Reading
- Promotional label (Deep-learning)
- Security Label Presence (Deep-learning)
- Beef and Lamb Portion Cut Recognition (Deep-learning)

The smart-camera deep learning algorithms built into the Cognex camera were fine-tuned to maximise the accuracy of recognition tasks while minimising false positives.

4.2.2 Main Label text recognition (deep-learning)

To avoid the issues with rule based vision, the Cognex camera is being trained on good images using the built-in deep learning algorithms. Coupled with very controlled lighting which creates a diffuse lighting environment with lighting coming from many angles, rather than a direct bar or spot light, the deep learning is able to recognise main label text very accurately.

Lighting artefacts which slightly obscure the text will not cause the algorithm to fail recognition nor mistake one label for another, as easily as rule based vision.

Label landmark used to locate the label region and then compared the main text to expected. The Cow is used as the landmark symbol in this example.

4.2.3 Date Code Reading (OCR)

Date code reading utilises standard Optical Character Recognition (OCR) tools which find individual character areas, through a process called segmentation, and then using pre-trained, in-built fonts, determine what each character is and outputs the string.

Deep learning is not appropriate for this task as the accuracy of each character must be checked and known and any obscuration must result in a failed inspection, as a legible date is a legal requirement.

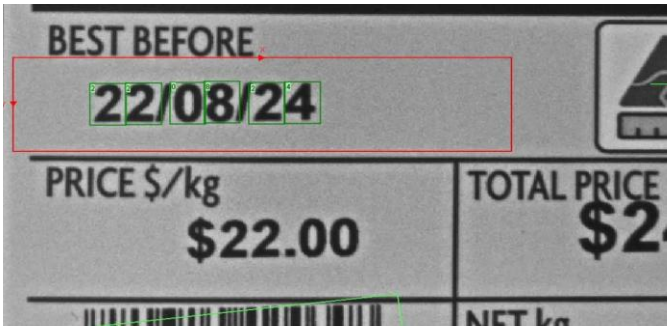

4.2.4 Barcode Reading

The barcode reading inspection tool is also a standard vision tool which doesn't require deep learning. The entire barcode string is read and output so that it can be compared to the expected data for each particular product. It tests both readability of the code and that it is as expected.

4.2.5 Promotional label (Deep-learning)

Promotional price labels have a legal requirement to meet any claims being made. As such the text must be read entirely and compared to the claim being made ie price value, whether it is a claim per pack ‘pk’ or each ‘ea’ or per kilogram ‘kg’ etc.

Each character must be found and compared to expected.

There are many varieties of font types, font sizes and character locations. To enable the system to provide easy product creation, different templates for each combination of these characteristics were created. Tools to locate a specific landmark character present on all, which in this case is the $ symbol, was used. Then each character nearby is segmented and recognised. Using pre-trained, in-built fonts allows the system to be configured without having to train every conceivable character manually. In this way any promotional label can be found and recognised and any new label meeting the combinations found in the templates can easily be added without requiring specialised assistance from the vendor.

Locating the $ symbol, segmenting and recognising each character in order to compare to expected.

Sample of Label Types

Locating Secondary Labels, checking their spatial location and type are correct

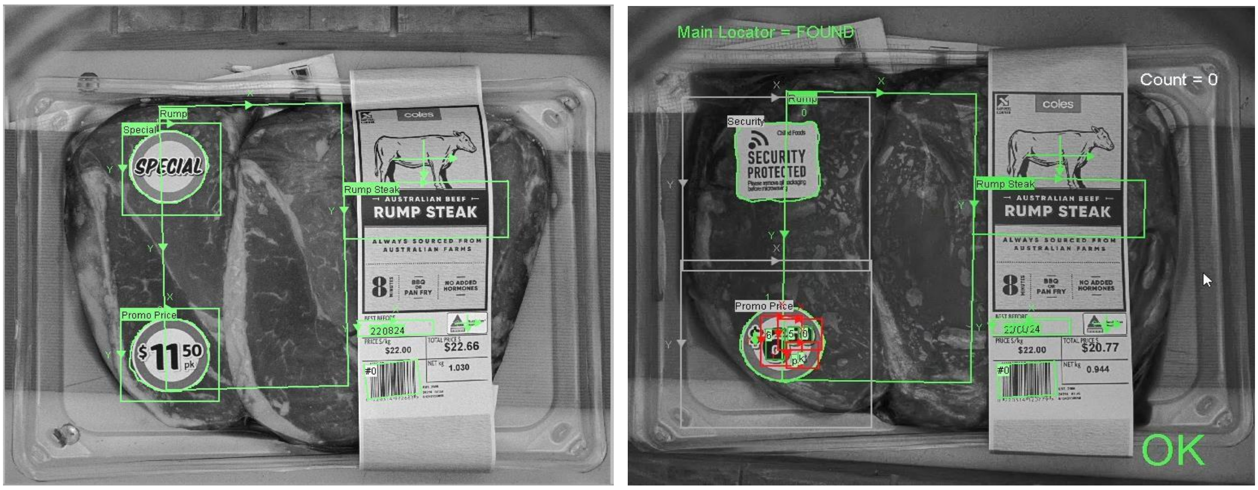

4.2.6 Security Label Presence (Deep-learning)

Check the Security Label is found and in the correct location, independent of the orientation

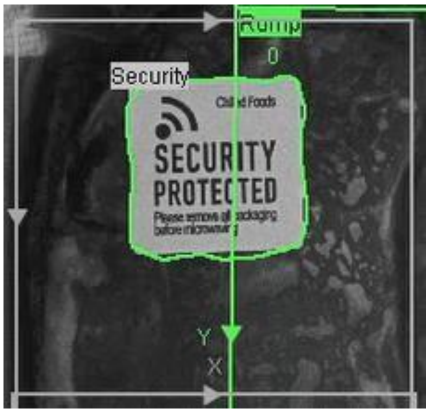

4.2.7 Beef and Lamb Portion Cut Recognition (Deep-learning)

One of the most challenging inspection tasks for rule based vision systems is recognising objects such as meat. Meat comes in many sizes, shapes and edge characteristics. Too many variables to accurately define with a few simple rules. The deep learning algorithm however, is uniquely capable of meeting such a challenge.

The camera is shown several images of each meat cut type. In this example the deep learning algorithm was trained using rump steak samples. The images are captured in various orientations, label locations, cut sizes, locations and number of portions. All are labelled as rump. In this way the algorithm is trained on a wide range of expected rump images. Once trained, the algorithm is able to recognise this meat cut regardless of many changing variables.

Rump expected, rum found

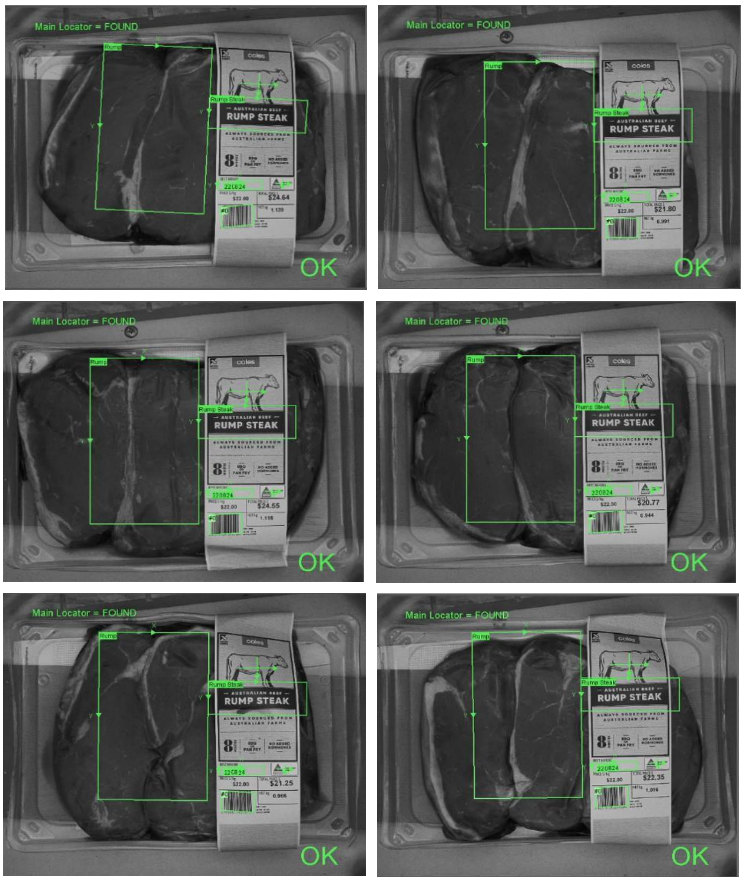

Rump expected, scotch fillet found

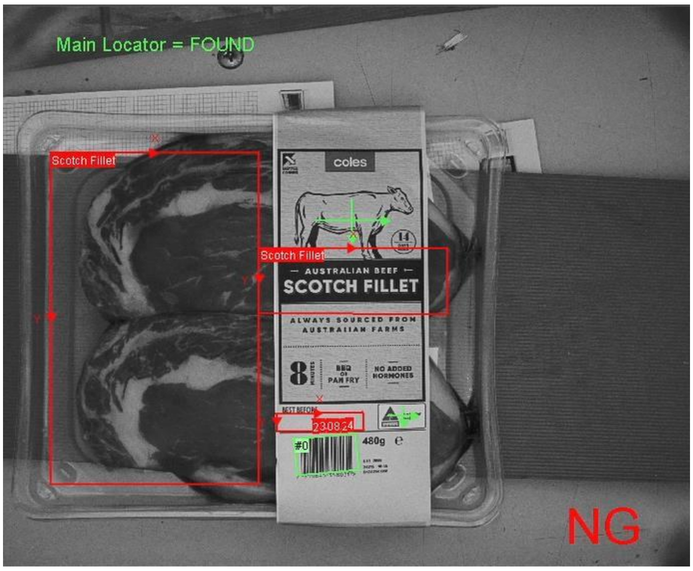

4.2.8 Speed requirements tested and satisfied (120 packs per minute)

The system must be capable of processing all required inspection tasks at a rate of 120 packs per minute to handle the fastest production rates in the RROA plant. This equates to a maximum 500ms processing time in total for all inspection tasks.

4.3 Beef and Lamb Primal recognition deep-learning development [Milestone 3]

4.3.1 Development of the Primal Recognition Deep-Learning Algorithms

The red meat industry would greatly benefit from the ability to automate the recognition of primal types. The ability to sort by type, aiding in automated pick and place, and to ensure quality control when packing boxes, whether automated or manual, or quality when receiving product, would all improve the efficiency and quality within this industry.

The Cognex deep-learning artificial intelligence smart cameras represent the possibility of this capability.

The following trials are thanks to RROA's partnership with Cognex and JBS Foods Australia.

Our partner JBS sort primals, which have been packed in vacuum bags, before placing them into cardboard boxes via robotic pick and place machines. There is a high cost associated with primals being sorted incorrectly and mixed in the same box.

This trial analysed the capability of improving this sorting process by utilising the Cognex D905C smart camera.

This particular sorting process requires the camera to look for every possible primal type that is going to be run on that line that day, in order to sort them to the correct carton lines. This means that each of the primal type's images must be added to what is called the model. So if three primals are being run on a particular line those three types must have their images added to a single model.

The following primals have been tested.

- Bolar Blade
- Navel End Brisket
- Cube Roll
- Eye Round
- Inside Skirt
- Oyster Blade
- Short Ribs
- Rump
- Tenderloin
- Topside

What was discovered is that the model health deteriorates as more primal types are added to the model. The model is expecting variations on a theme, but not images that are totally unique.

In this mode of operation the algorithm starts confusing primals that are similar if they are in the same model. However if the camera is just trained on one primal type and all else is supposed to be a fail, it achieves very high, 99%+ accuracy. This was discovered later.

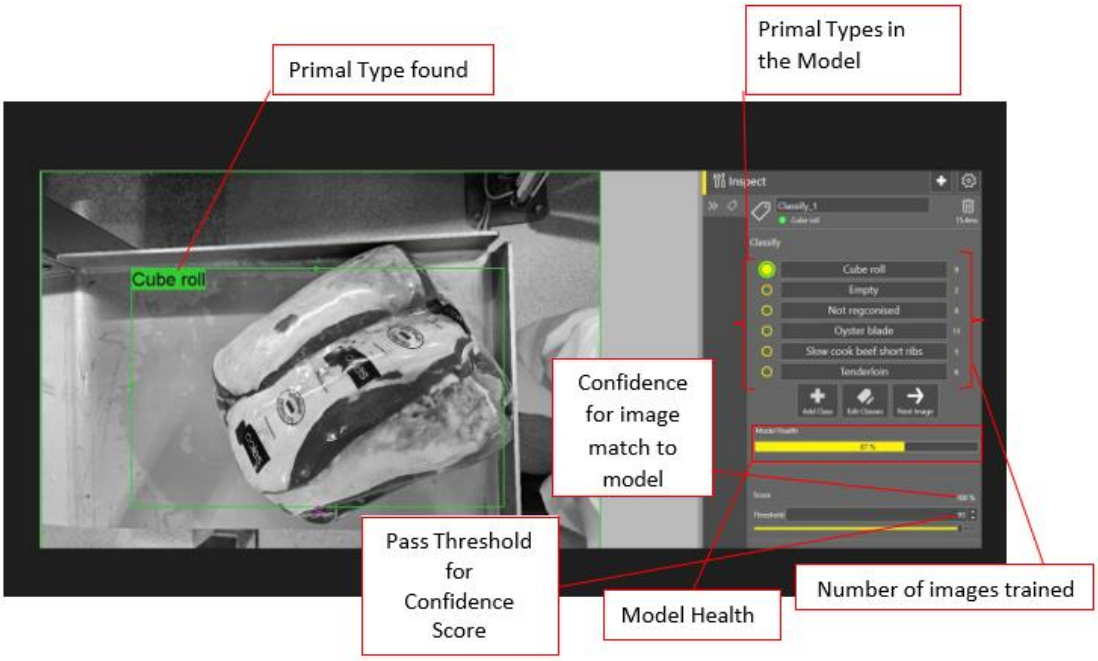

For the initial trial the following primals were selected - Brisket, Tenderloin and Short Rib, as these were the most common primals run on the test production line.

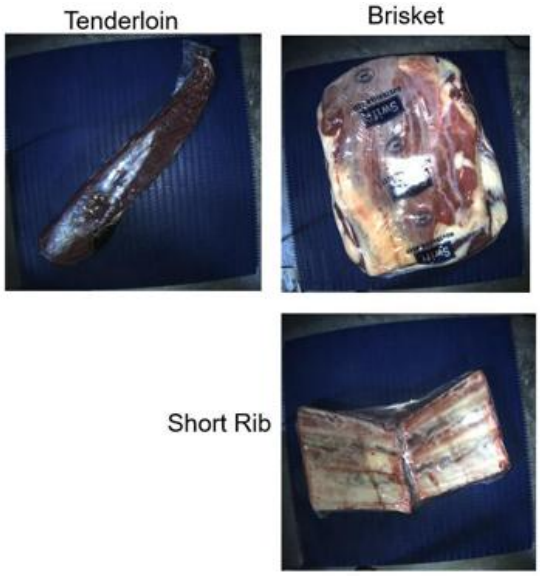

With minimal training of good images (4 to 7), the system was able to accurately recognise these three primal types despite variation in size, orientation, presentation and positioning, as shown below.

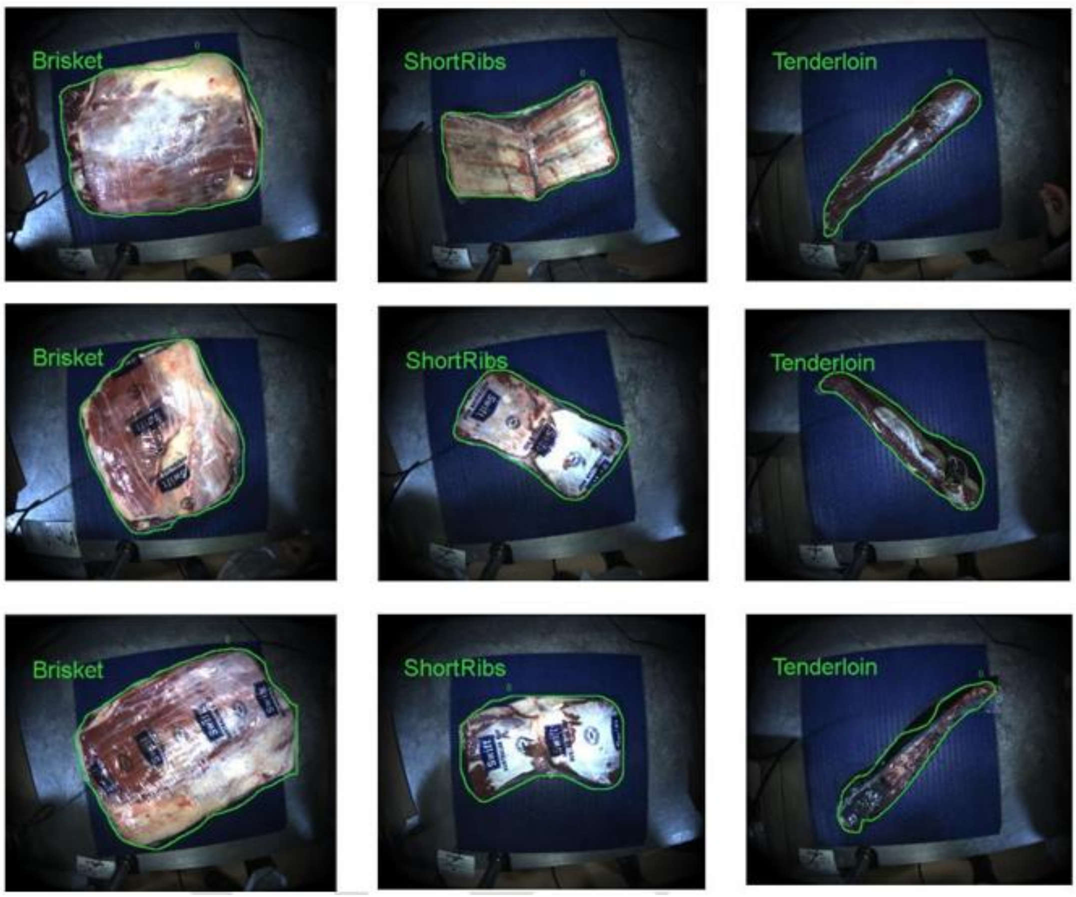

These primal cuts are correctly recognised as not being a brisket, short rib or tenderloin

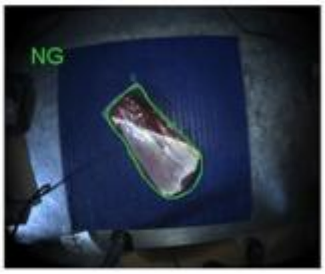

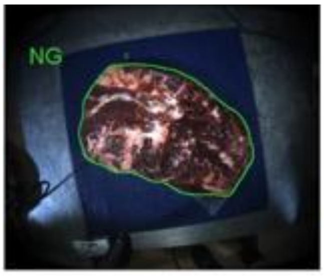

The system is capable of achieving 97%+ accuracy with these very different primal cuts in the one model. However if the system is trying to distinguish between cuts which are quite similar and the model is developed with all those primal types, the accuracy can drop to around 85% with a model health below 80%. A model health above 80% generally ensures a higher accuracy of recognition.

To solve this issue the camera should only be trained with one primal type. For example if the camera is trained on Brisket only, it can achieve 99%+ accuracy.

For a sorting process you would use one camera per primal type. At each sorting point the camera would send all failures to the next camera, while images which are a Pass ie matches the primal type it is expecting, would be send to the carton line for that primal type to be packed. All other primals are then sent to the next camera for further processing. This next camera would only be looking for, say Tenderloin. In this way the accuracy of the system can be increased dramatically. Additionally very similar primal types can be handled more accurately. However the drawback is the requirement of more cameras for this type of sorting process.

The process proposed at RROA is not a sorting processing but an inspection to verify a label is correct and matches the portions within the pack. As such the inspection is designed to only look for one SKU at a time while all other SKUs are deemed a fail. This method has been proven to achieve high accuracy compared to trying to model all SKUs together in one model.

For our partners JBS, inspecting the primals before they are vacuum packed may further improve accuracy. When multiple primals of the same type are packed within one bag they can be made to look like other primals which on their own look quite different.

4.3.2 Ability to recognise Chinese characters for export label inspection

China represents one of Australia's biggest trading partners for the red meat industry. As such the ability to inspect export product labels is essential in maintaining our quality standards for this market.

There are two main inspection methodologies the Cognex smart-camera can utilise to perform inspections on Chinese characters. One method is to train the camera to recognise specific characters as a graphic pattern and look for matches without understanding what the character means. The second method is to train the character set and perform OCR on label print and match to an expected meaning.

Each day at JBS, the QA department prints a single label for the range of SKUs which will be manufactured that day. They manually conduct a visual inspection of these labels to ensure the label contains all the correct information before it is released to Production.

The Cognex smart-camera can provide real-time feedback to a QA inspector as it validates a label's contents.

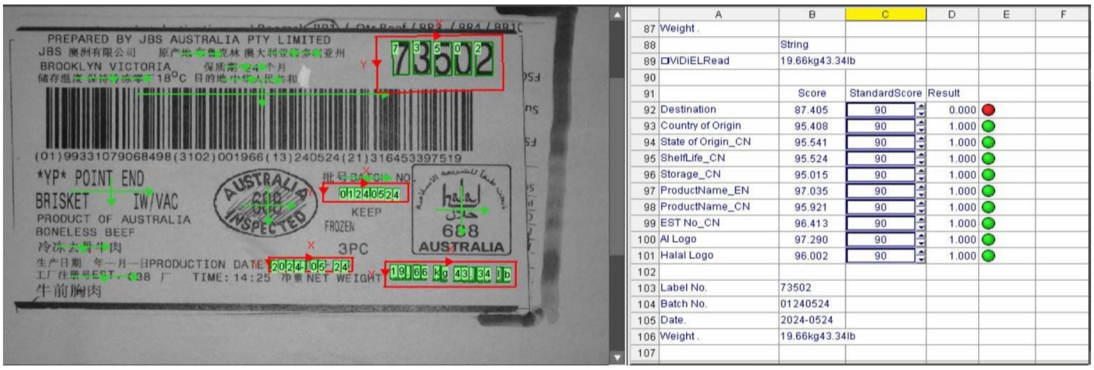

Vision tools to detect the presence or absence of key data to be validated were used. The pattern recognition tools were then employed within these key regions to match against the target pattern some of which represented Chinese or English text.

Over 100 label types are being inspected for different SKUs, countries and primals.

Below shows the second method of inspection where a whole character set used in the labels, has been trained for a particular font. Each character is trained so that when inspected the image is actually read and compared against an expected set of characters. The percentage match between read and expected is compared against a threshold to determine if it is a match or fail.

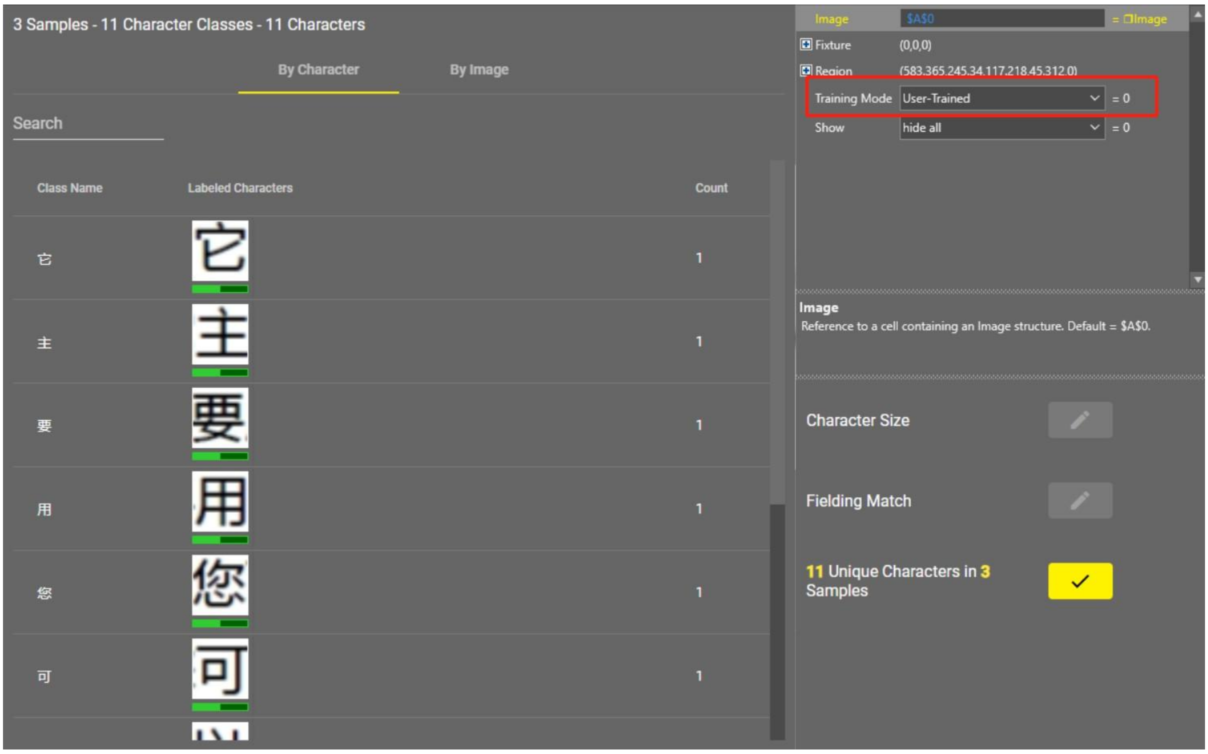

Both Pattern Recognition and OCR achieve near 100% recognition rates with confidence scores ranging from 70-100%.

The benefit of the OCR method is faster processing speeds and the ability to transfer the target characters from a string in a database rather than having to select character by character with the pattern recognition tool.

4.4 Development of the reject handling system [Milestone 4]

4.4.1 Software development

Reject Timing

The system outputs a digital signal via its PLC output every time a pack is tagged as failed by any of its inspection tools. The signal has an associated delay and hold timer. The period of time that the signal is delayed and/or held high is adjustable by a user with administrator rights. This allows the reject timing to be fine-tuned.

Reject Confirmation

The software monitoring the sensor to determine that a pack has been successfully rejected ensures that the sensor is not on for too long. If a pack gets jammed the sensor will see the pack not fall away. The software will stop the conveyor, signal to the host system that there is a fault and the user must acknowledge the alarm to restart the conveyor and clear the alarm.

Reject Bin Full

The software monitoring the sensor to determine that the reject bin is full will stop the conveyor and signal to the host system that there is a fault if the bin full sensor is covered for a sufficient length of time.

Reject Test Mode

There is a reject test mode which allows the reject mechanism to be disabled by admin level users. If a pack, which should be rejected, is run through the system and the reject mechanism, which has been disabled, fails to eject the pack, the conveyor will stop and an alarm displayed.

When the alarm is reset the conveyor and system will restart. Once the reject mechanism is re- enabled all packs which should be rejected will trigger the reject mechanism to push the pack into the reject bin.

Consecutive Failures

If there are 5 (can be altered by an admin user) consecutive failed packs the conveyor will stop, a signal to the host that there is a fault is sent, and alarm displayed on the display. The user must acknowledge the alarm to restart the conveyor.

4.4.2 Design and drafting

The vendor has designed a reject mechanism to reject packs off the line which fail any of the vision inspections. The maximum throughput rate of the system has been specified at 120 packs per minute.

The reject mechanism was designed to handle this maximum throughput with consecutive packs requiring to be rejected. The reject mechanism design must ensure that all pack sizes specified should be accurately rejected off the line without getting stuck at the entrance or within the reject bin.

The line control must provide tray gaps with a minimum of 150mm spacing between packs.

The bin to capture the rejects must provide a mechanism for soft fall so that packaging is less likely to be damaged especially with packaging materials becoming more brittle as recycled content increases. The reject mechanism must be a push mechanism and safely enclosed.

Initial concept drawing

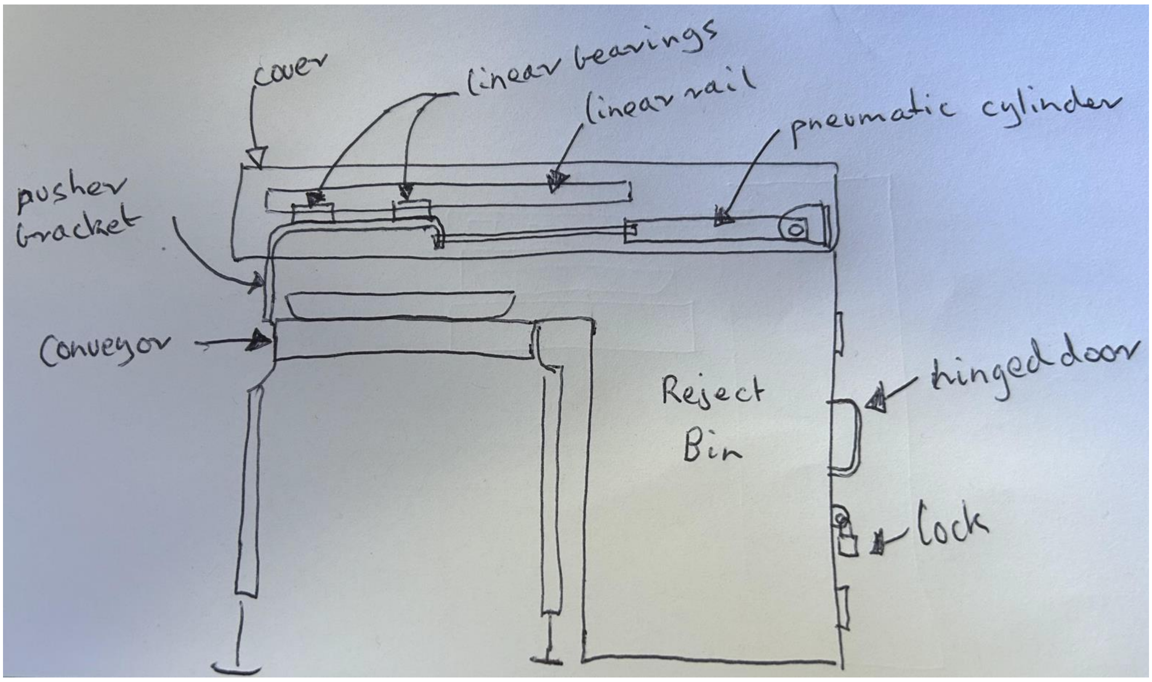

The initial concept was a pneumatically driven pull mechanism. However during trials caused jams on the return stroke. The design was changed to a push mechanism which was more forgiving with trays where the gap was not being maintained at the required 150mm distance.

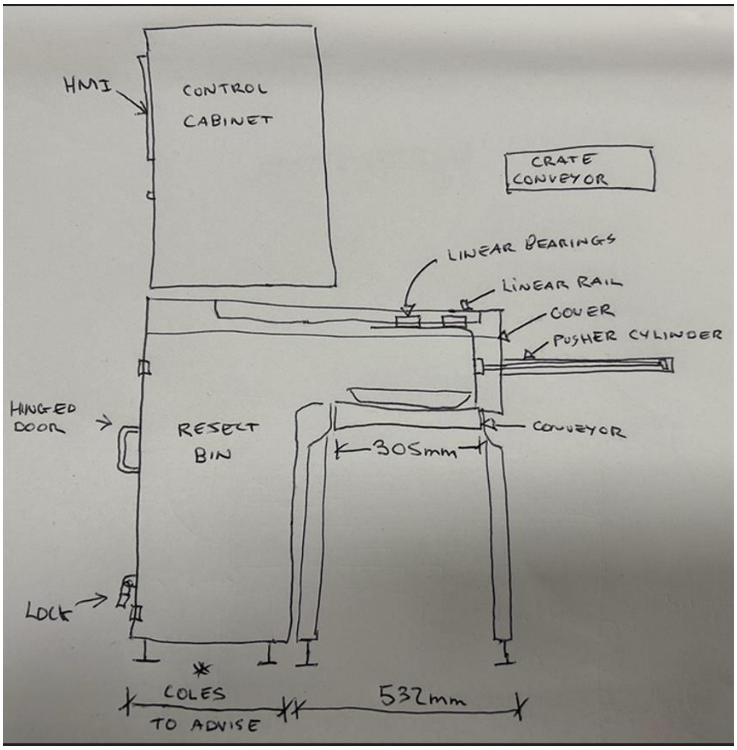

Soft Fall Mechanism Design Options

Option 1

Using a powered reject bin floor to start at the top of the bin and then move down as trays are rejected would ensure a soft fall for the rejected trays. This was seen as too complex with moving parts which could fail and requiring additional control.

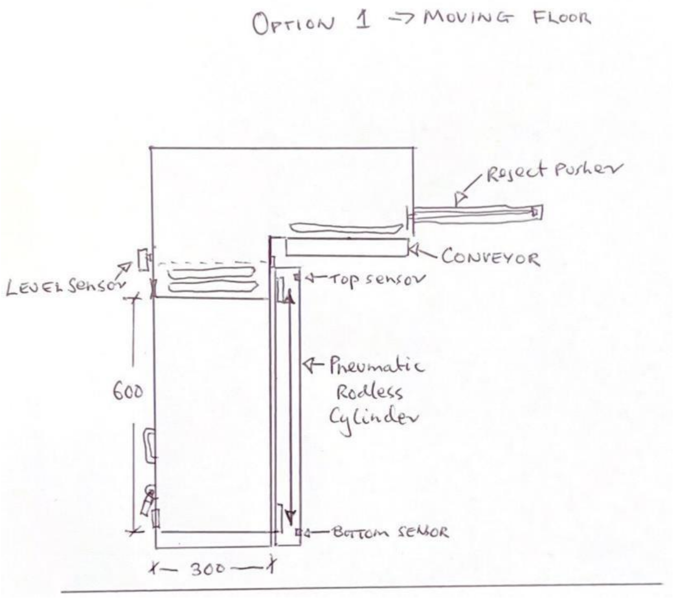

Option 2

Using a counter-balanced cylinder which pivots as it is loaded with more and more trays to lower the bin floor using precision pneumatics. Again seen as too complex mechanically.

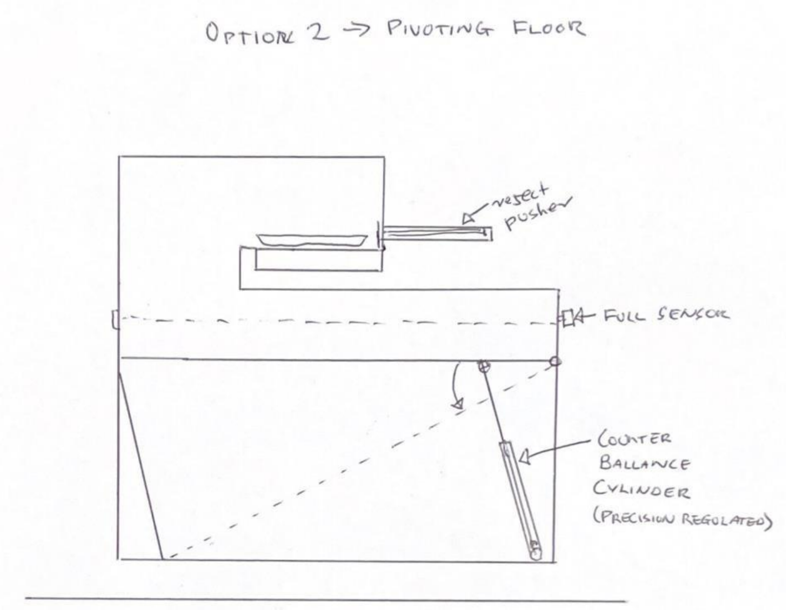

Option 3

The final design option utilises baffles angled and positioned to create a soft fall for rejected trays and thereby limiting the chance of trays being damaged during rejection. This was seen as a nice and simple, yet effective design.

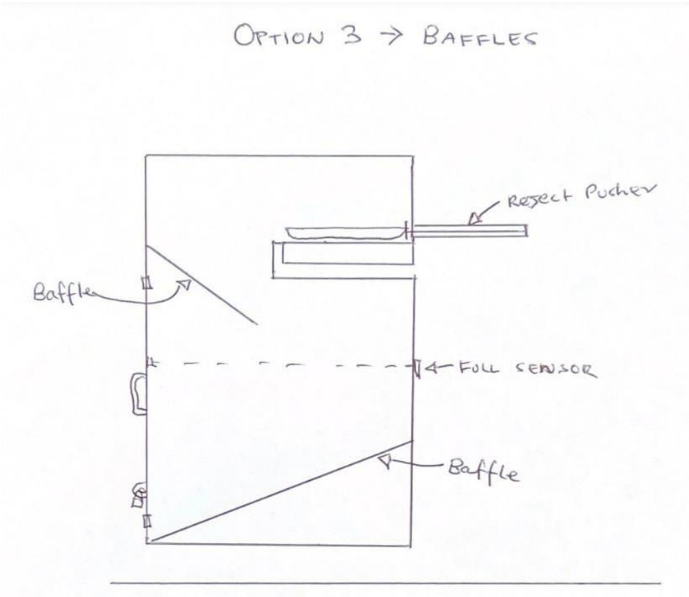

Design Drawing of the Reject Mechanism

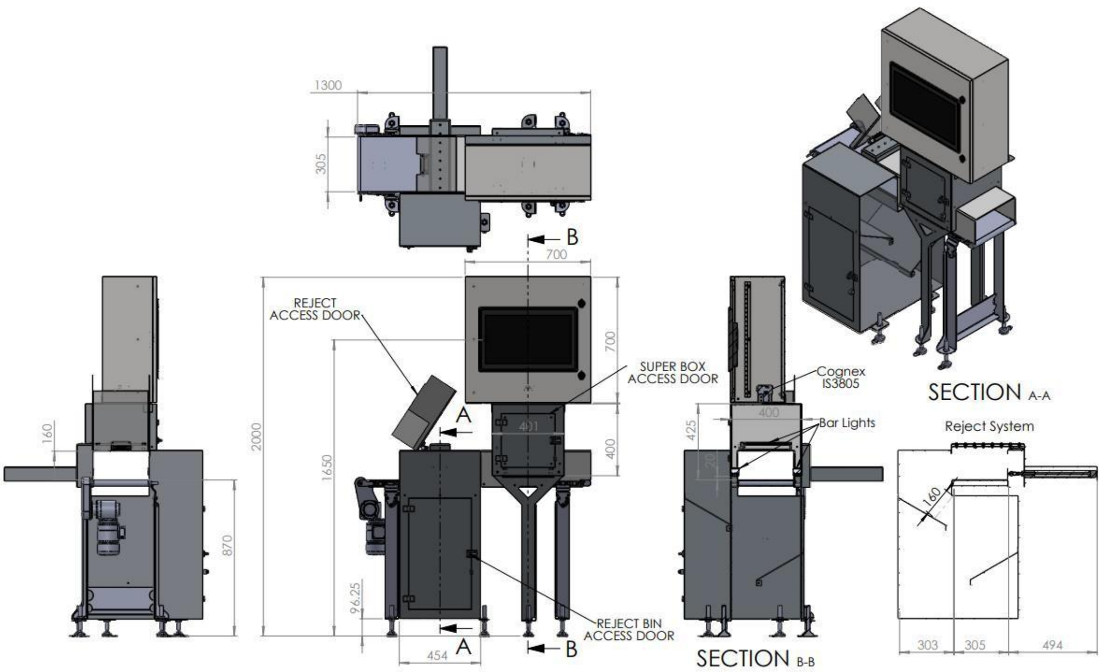

The final design is a pneumatically driven pusher mechanism which pushes failed trays into the reject bin. The opening of the reject bin is much larger than the largest tray width to ensure trays which are angled as they are rejected, will not easily jam.
Baffles are used within the reject bin to slow the fall of rejected trays to limit damaging the trays. This is particularly necessary moving forward as more plastic recycled content are used for packaging materials, which tend to be more brittle.

Sensors are placed at the top of the bin in line with the reject mechanism to confirm each reject has been successful and not jammed. Another sensor is placed slightly lower at the top of the bin to monitor for the bin full state. The line is covered with Perspex tunnel so people can't take packs off the line between inspection and reject to comply with Quality Control Point standards.

The bin, that the rejected packs are ejected into, is locked with a combination padlock on the lid which is on the front of the bin. The lid, which is hinged at the side, allows for easy and safe access to the packs within the bin i.e. side opening lid, not lifted up or down.

4.4.3 Mechanical and electrical works

Mechanical Works

Below is an exposed section of the reject mechanism. The pneumatic rod is enclosed in an enclosure behind the conveyor. The stainless steel pusher is extended by the rod half-way across the conveyor to push the tray with a design that spreads the load across the tray and limits deformation and damage to the tray. The pusher only extends half way across the conveyor to give an impulse to the tray to push it off with momentum, and allows the pusher to retract quickly back into position without interfering with the next tray coming up behind, at a minimum gap of 150mm.

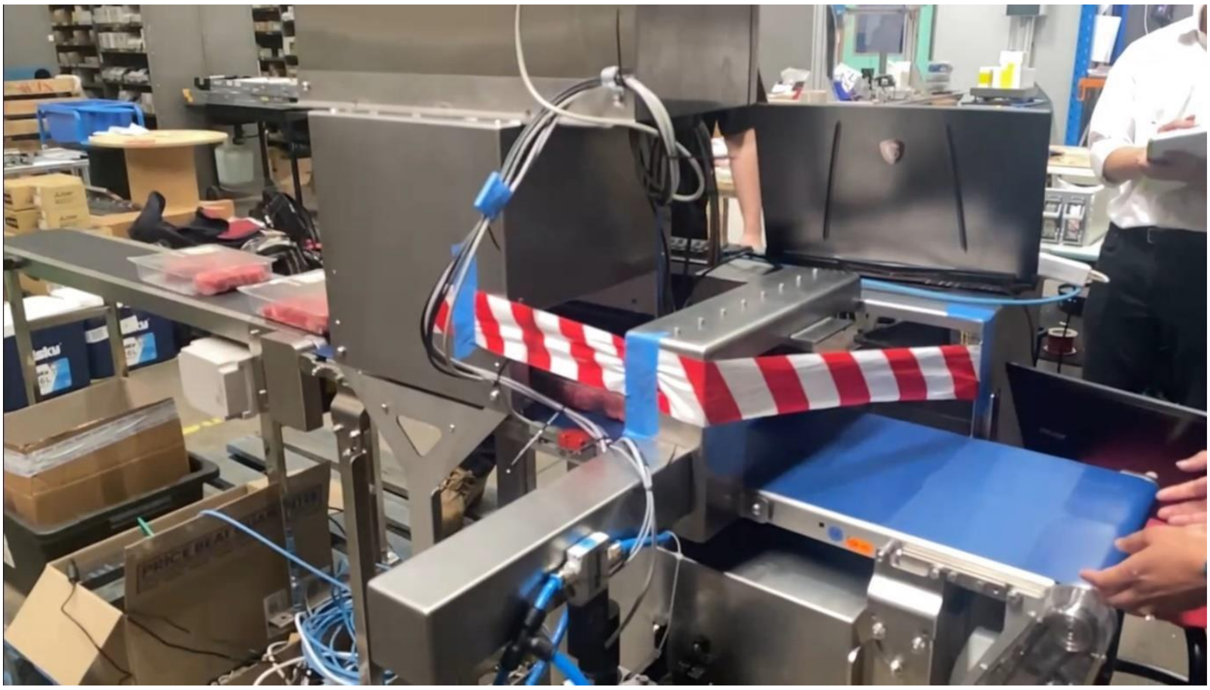

Reject pusher

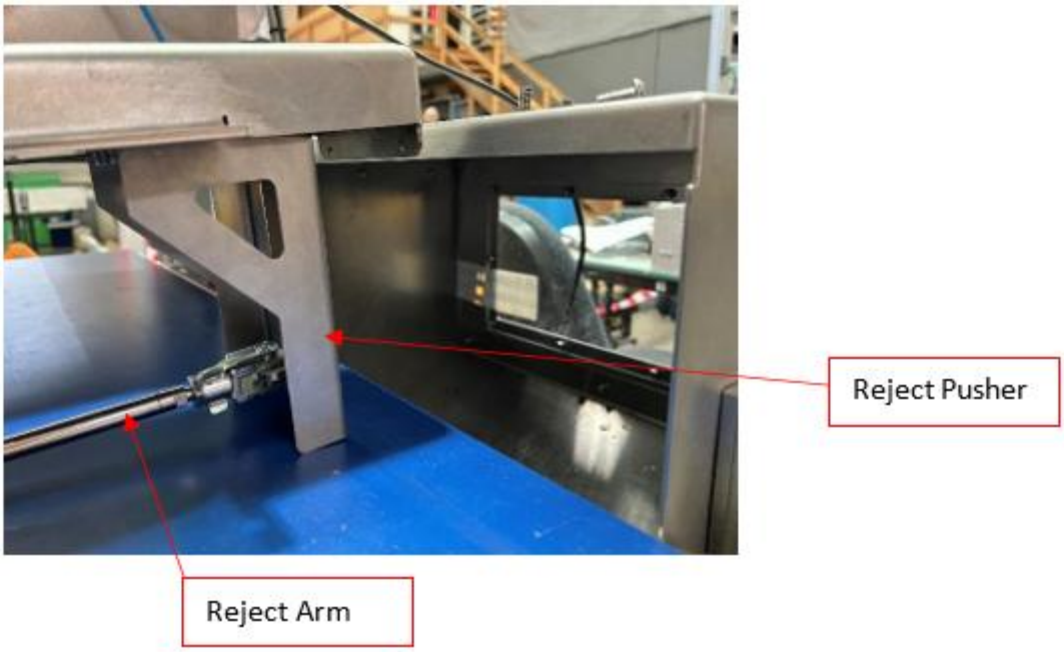

Reject bin

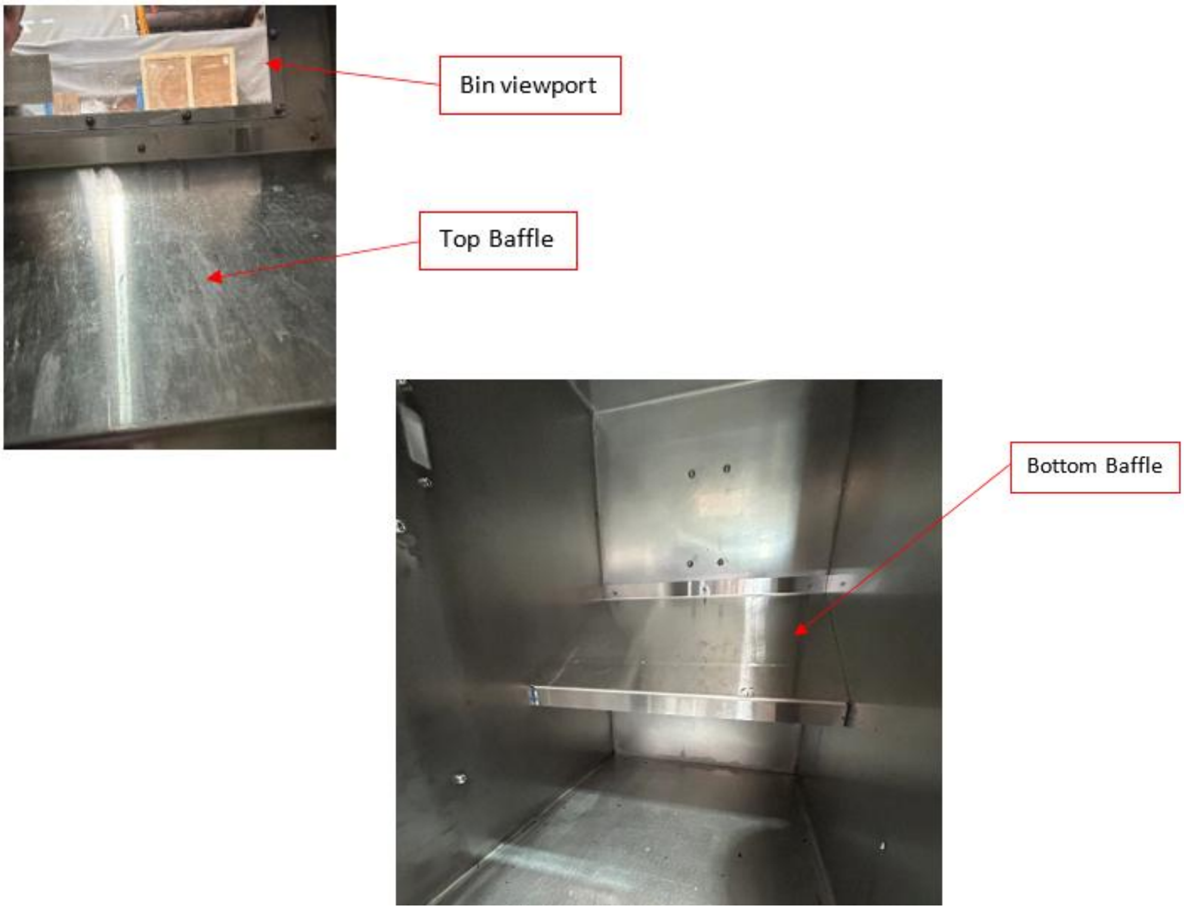

Quick soft push

Rejector arm completes motion half-way across belt to retract quickly before next tray comes into position

Maximum deflection of tray on push

Tray slides down baffle quickly before next tray in position

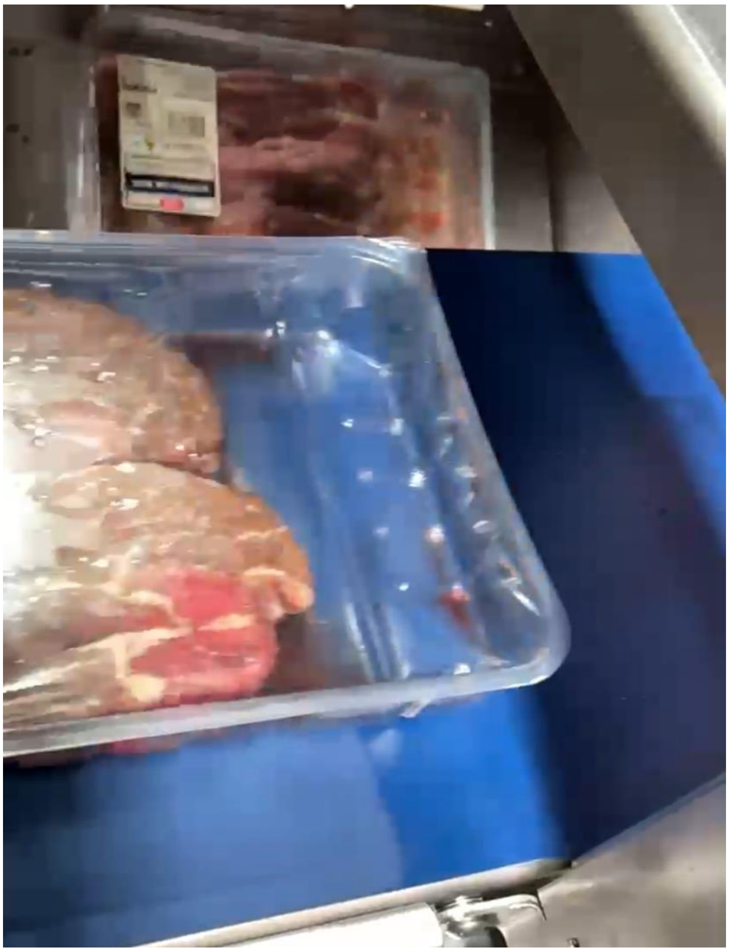

Electrical works

The electrical design involves the use of a Variable Speed Drive (VSD) which controls the speed of the conveyor. A Programmable Logic Controller (PLC) is used to control all signals including monitoring the reject sensors, control their timing, controlling the reject mechanism and controlling the speed of the conveyor.

The HMI screen is used to acknowledge faults such as too many consecutive failed inspections or failed rejections and reset the alarm and restart the conveyor.

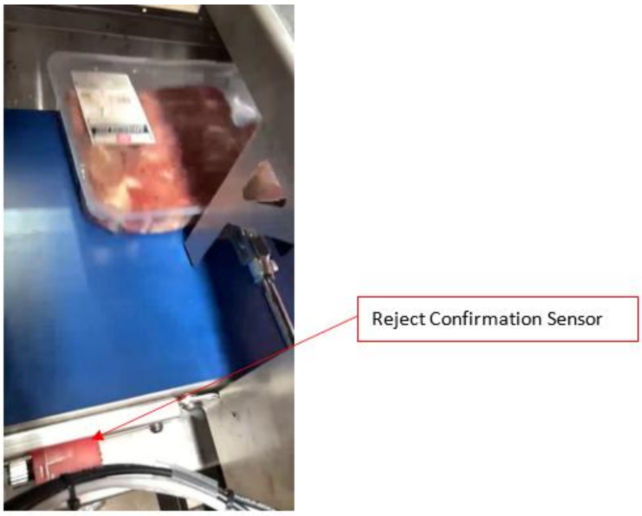

4.5 Development of the audit compliance module [Milestone 5]

It is a requirement for the Technical / Quality department to have audit logs that record critical production information related to food safety. The Deep Vision system has developed an audit compliance module to record all actions against the user who is logged in.

If there are any parameters which are altered which cause issues or even a recall, the date and time of these changes can be found in the audit log and the person associated with the change can be determined. Recalls can be managed most efficiently using this audit log as only packs within the timestamps of the change can be targeted.

The system will not function unless a user is logged in. The logins are based on specific user's names, not generic users. Each user is assigned to a category when first created. Each category has a specific access level for every function within the system.

General Function / Screen	Specific Function / Details	Admin	QA	Supervisor/Team Leader	Engineer	WPL Operator/QC
Navigation	Home Screen / Login	X	X	X	X	X
	Run Product / Product Selection	X	X	X	X	X
	Manage Products	X	X		X	
	Manage users	X	X		X	
	View audit trail	X	X	X	X	
	Tune job file	X				
	Exit application	X	X		X	
Run Product / Product Selection	Override automatic BBD		X			
	Production Selection / Load Product	X	X	X		X
Manage Products	Add new product	X				
	Delete product	X				
Edit users	Create user	X				
	Delete user	X				
Run screen	End run	X	X	X	X	X
	Acknowledge / reset alarm	X		X	X	X
	Start / stop conveyor	X		X	X	X
	Enable reject verification test mode	X	X	X	X	X
Key parameter change	Adjust key vision inspection values	X	X			

Every action which makes a change in the system is recorded in the audit log as shown below.

4.5.1 Integration with MES for system parameter management

The old vision system is stand alone and relies on production staff on the floor to make critical selections on the system on the floor. If many staff can make changes the risk of errors goes up exponentially. As such it is better to limit the number of people who can make critical changes and centralise this change mechanism.

The Deep Vision system has developed their own MES parameter management system and architected it so it can integrate with any external MES system through the use of a common ItemCode for each product.

Each camera system in this architecture is connected to a central server. Only control room staff or those with high levels of access can make critical selections on the system. The control room staff remote access each floor system to select the product to be run and any critical parameters such as what secondary labels are to be run. The QA staff on the floor confirm the selections made are correct and then the production staff have permission to start the run.

All master data within the existing MES is able to be synchronised with the Deep Vision master data based on the ItemCode which is a primary key, unique code, related to each product. Once synchronised, the Deep Vision system manages all selections and parameters internally without needing further interaction with the MES application.

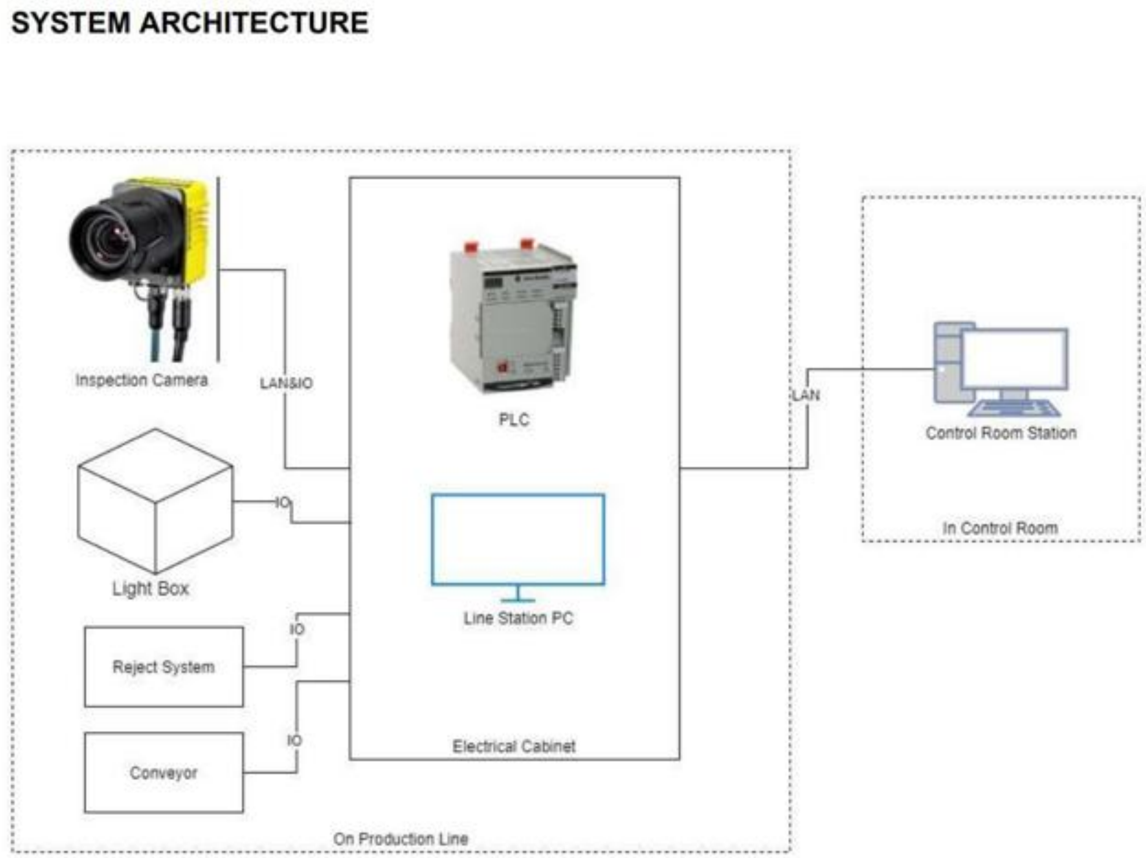

This architecture is not connected to the internet and so there is no attack surface available for external attack. Even if the system is connected to the internet, the camera runs the vision applications without constant need of the embedded PC or control room server.

All selections and control can be made directly on the camera with the appropriate user access level. In this way the system retains its hardness to cyberattack.

As detailed in the previous section, the client server architecture allows for central control of critical selections and parameter assignment by expert users. However should the internal network between the server and camera client go down, or for some other reason no server control is available, local control is available for users with the appropriate access.

4.5.2 Integration with MES for order management

Once master data between the MES is sychronised with the Deep Vision system, all order management can be handled internally.

4.5.3 Assess and report on how this system will be able to integrate with GS1 2D barcodes

GS1 2D barcodes will become the industry standard in 2027. Industry has set a date to make the transition to accepting 2D barcodes at point-of-sale (POS) or point-of-care (POC) — referred to as Sunrise 2027. By the end of 2027, retailers would need to ensure their POS systems are equipped with scanners capable of reading both traditional barcodes and 2D barcodes.

2D barcodes can carry additional data, such as expiration date, batch/lot number, serial number and more. Having this additional data in the barcode adds value by allowing the information to be automatically captured and acted on.

Some 2D barcodes, like a QR Code powered by GS1 Digital Link, can carry additional data while connecting consumers and other users to online resources and experiences. Having this data in the barcode adds value by allowing the information to be automatically captured and acted on. In addition to carrying more data, 2D barcodes are likely to be smaller than their 1D counterparts and include features, like built-in error correction that add to their reliability.

The Deep Vision has 2D barcode reading capabilities as can be seen below.

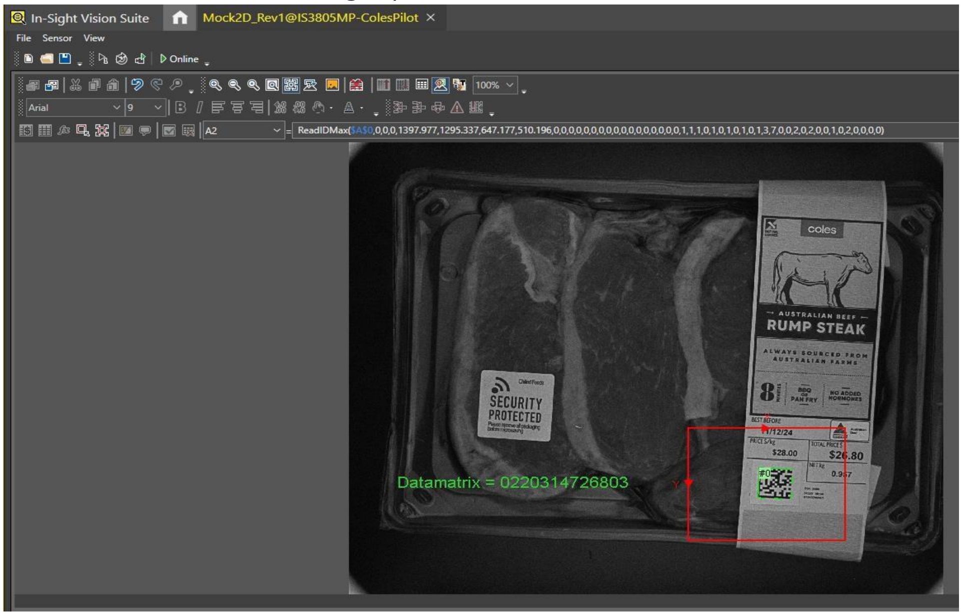

4.6 Final report [Milestone 6]

4.6.1 System performance testing and comparison against current vision system

In this section, the new Line 4 Cognex Deep Learning AI vision system, which has been named ***Deep Vision,*** will be compared against the vision system employed on the other lines. For each category a performance metric for Deep Vision will be detailed.

4.6.1.1 Mechanical performance

The new vision system reduces pack jams at the rejector when multiple packs in a row have to be rejected or when large packs, on an angle, have to be rejected. By utilising a larger opening into the reject bin and using two solenoids, one for push and the other for retract, for faster reject action, both these causes of ***jams have been reduced to less than 0.2% down from 3.3%.*** Additionally the rejection mechanism isn't required to stroke fully, only enough to push the pack into the bin with a fast nudge rather than a full push, which again reduces the risk of the next pack hitting the rejector arm as it retracts.

The current vision systems rely on single acting solenoid to push reject packs off the line, through a fairly tight opening, with a full extension push. When several packs in a row require rejection the rejector's speed is not always sufficient to retract back in time before the next pack moves into position causing a jam. With larger packs the small gap into the reject bin can cause jams if the pack is on an angle.

The Reject bin design was also a significant enhancement on the current design. Using *baffles* within the bin, the *fall of rejected packs is softened*, limiting the chance of damage to rejected trays. As packaging material continues to increase in recycled content, they tend to increase in brittleness which can cause damage when dropped from a great height. The baffles reduce the fall height, from surface to surface, until the pack is safely deposited at the base of the bin.

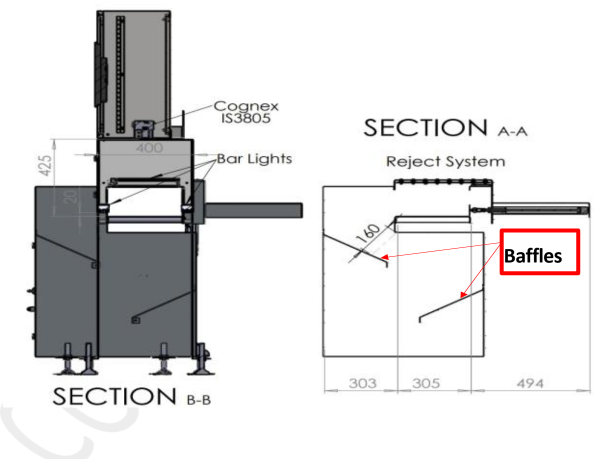

A side hinged access door on the reject bin, unlike the current top hinged doors, removes any risk of doors swinging shut on operator’s hands, thereby providing a safety enhancement.

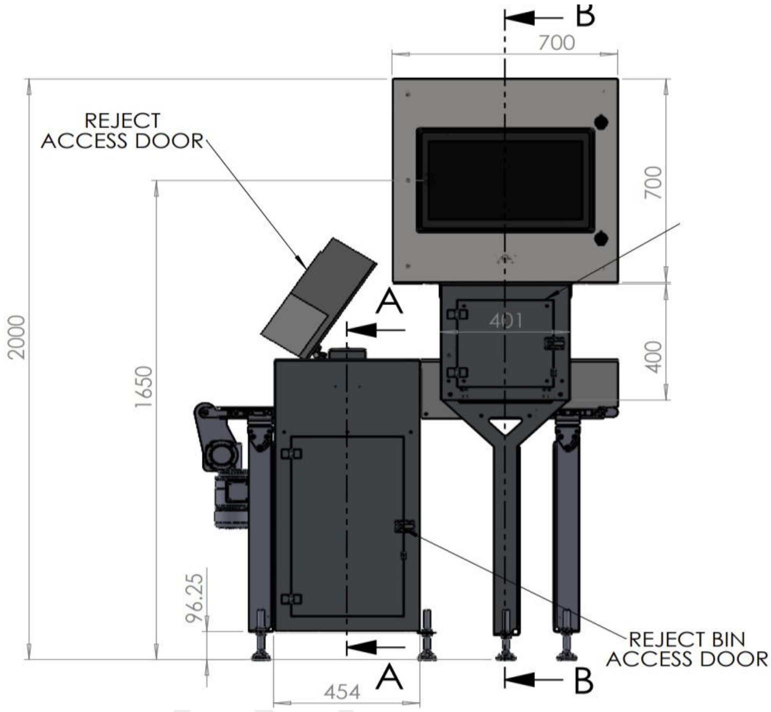

4.6.1.2 Electrical performance

The current vision system conveyors are complicated servo drives with embedded controllers which are out-dated. The current vision system conveyors are a simple VSD controlled motor, more easily integrated with any PLC or SCADA system.

The PLC was selected to match the current site standard of Allen-Bradley PLCs, again for easy integration with the rest of the plant systems, and to standardize spare parts which ensures quicker recovery times, should parts fail or need to be replaced.

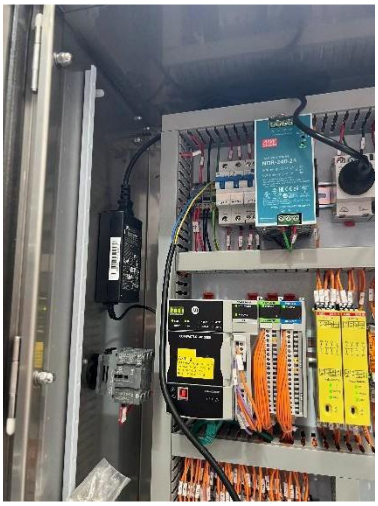

4.6.1.3 Vision performance

4.6.1.3.1 Lighting

One of the key improvements implemented in the new vision system is a more controlled lighting environment. The light box that the pack enters is more enclosed, which limits the variations in lighting that can be present in the external environment. Additionally the lighting box uses indirect lighting with LED lights pointed upwards into the white Perspex box so that light bounces down onto the pack from all directions, creating very even lighting.

External lighting sources can cause bright spots on glossy labels in the current vision system, which can obscure key elements of the label that are being inspected. This will result in labels failing for no reason other than poor lighting. The current system loses up to 5% of packs due to this lighting issue. The current system loses 0% of packs due to lighting issues.

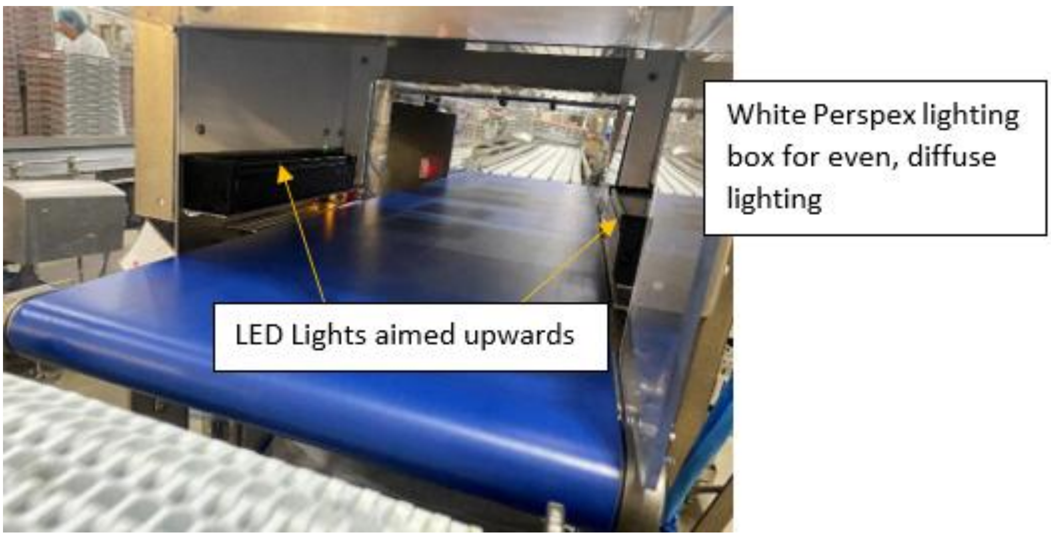

Deep Vision - Even lighting across Vacuum Skin Pack surfaces and their glossy labels

Current vision image quality issues

Without the lighting box and using ambient lighting we have very poor lighting quality

Current vision system

Under-exposed but still experiencing glare off the glossy surfaces

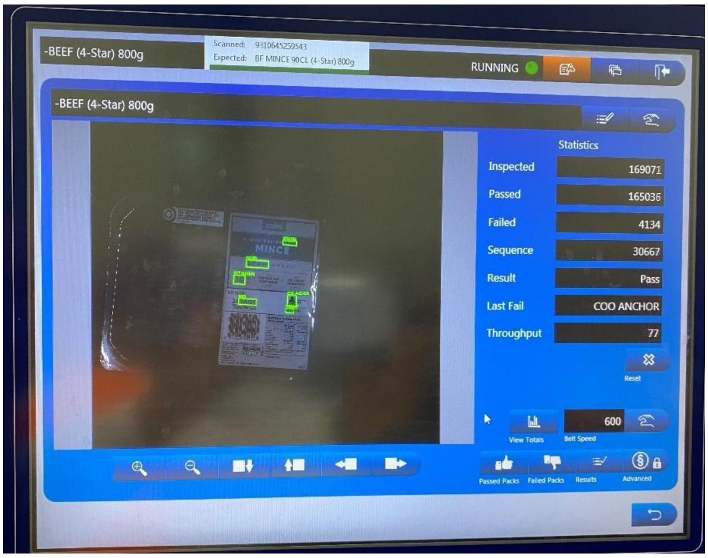

Current vision system

No OCR, barcode or full product description inspection

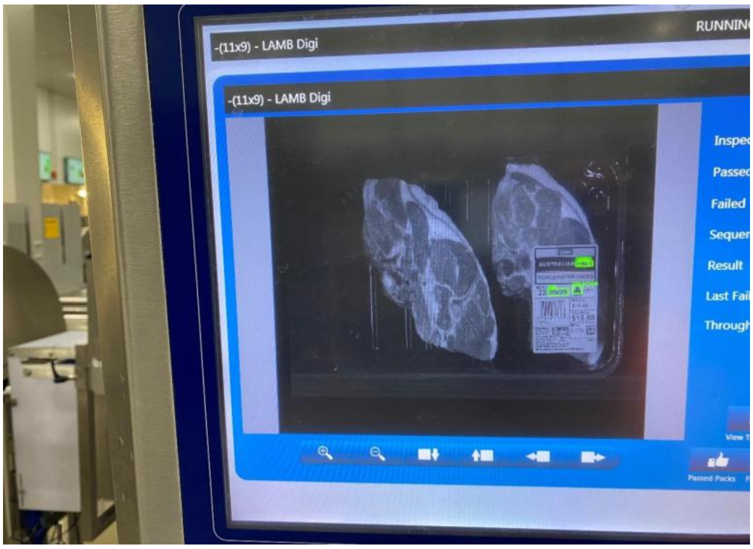

4.6.1.3.2 Deep learning

The current vision systems are all rule-based algorithms which utilise simple pixel pattern matching to look for the main label text. The tool doesn't read the actual text as in Optical Character Recognition (OCR). As such, the Best Before Date is not actually compared to expected, but only checked for presence. The main label is not read, but its pattern is checked against the expected pattern to ensure a match. However not every aspect of the main label's text is checked, and as such, small variations can trick the algorithm i.e. a B can be mistaken for a P due to the large degree of pixel matching.

The Deep Vision's Deep Learning Artificial Intelligence system however, uses example based learning to teach the system what to expect for each Main label description type. Using more human-like learning, the system can fully match the main label's description and not be tricked by similar letter shapes.

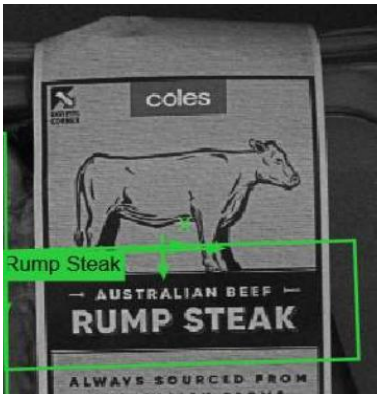

Furthermore Deep Vision's Cognex smart camera has full ***OCR*** capabilities with built-in font recognition, so the font is instantly recognizable without any special character training and the characters are actually read. In this way the Best Before Date can be read and compared against the actually expected date.

The ***Deep Learning Artificial Intelligence*** system, uses example based learning to teach the system what to expect for each ***Meat type***. With example based learning, various images of each meat cut are trained. With very few, typically under 10, training images required, the deep-learning AI can recognise many variations of each meat cut despite orientation, fat content variations and size of portions.

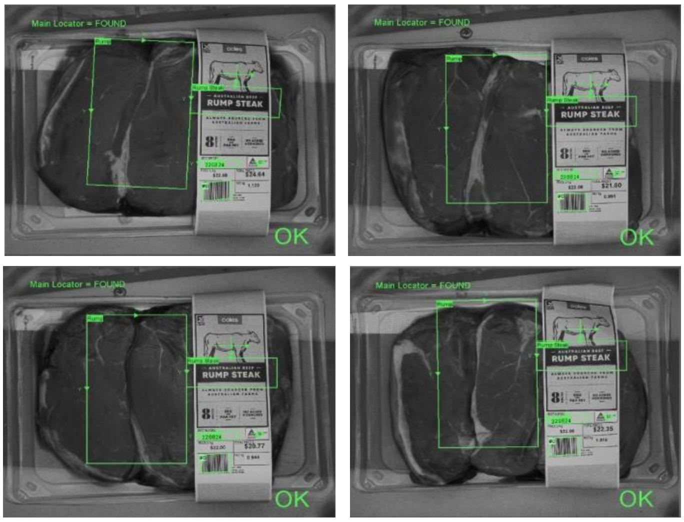

4.6.1.3.3 Barcode

Deep Vision's Cognex barcode reading tool reads the full barcode and compares it against the expected code. The current vision systems only check for barcode readability not whether the code is correct for the selected order. Even with the readability check, the current vision system isn't as sensitive as the Cognex tool. The Cognex barcode tool can detect when 1 pixel in the printer head has died and has removed a critical part of a black bar and changed the barcode value or damaged it so it is unreadable in stores.

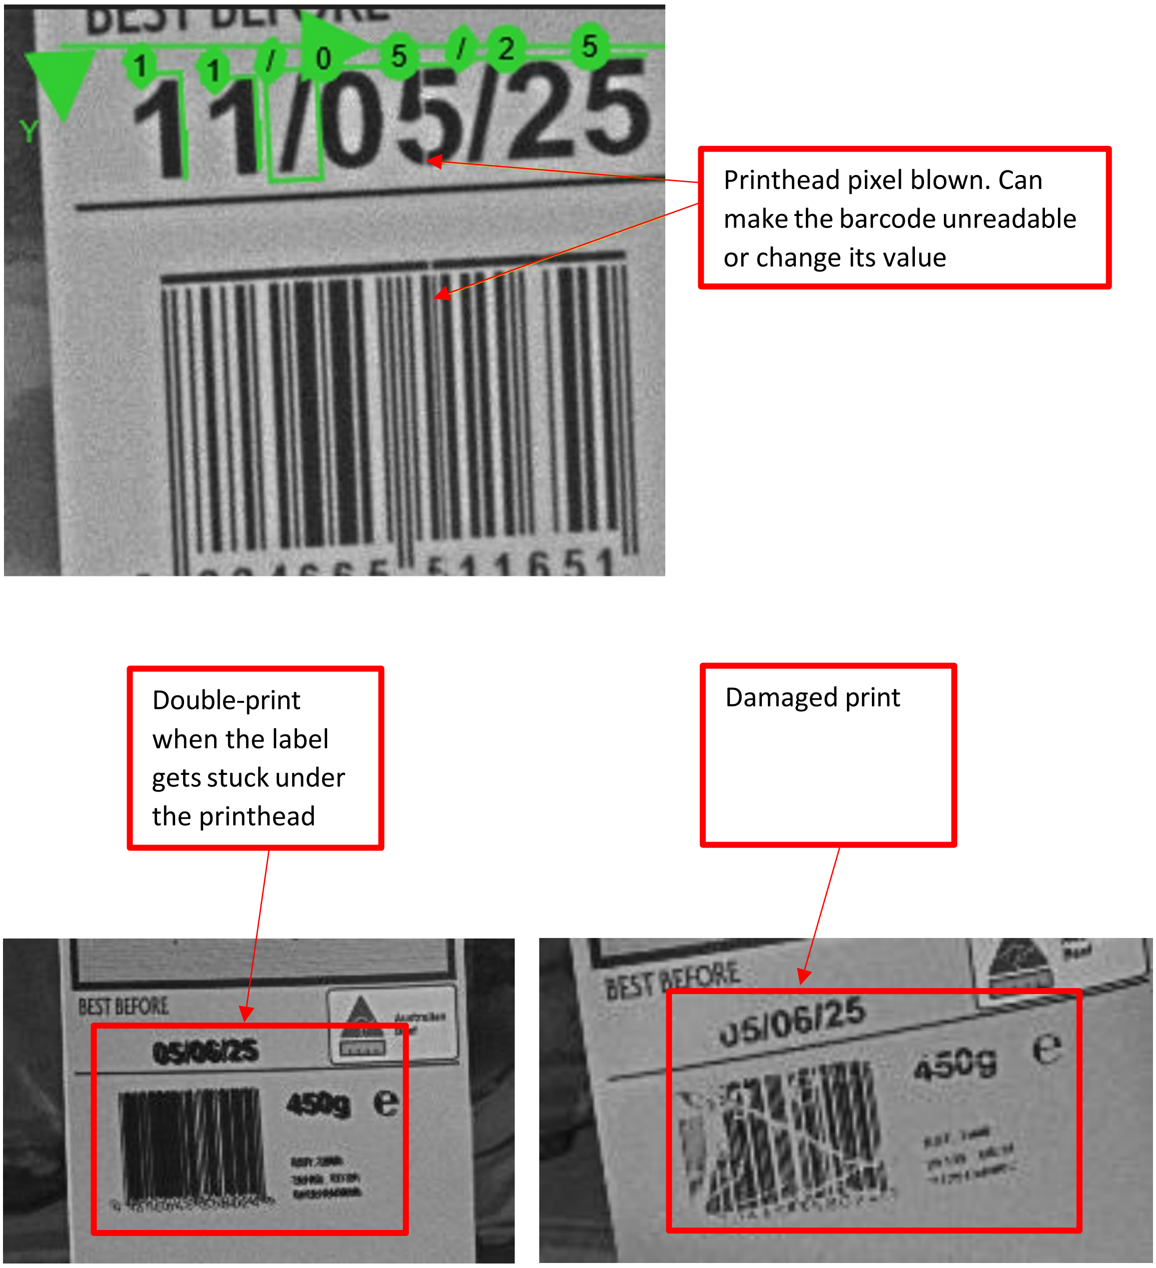

4.6.1.3.4 Secondary labels

The current vision system doesn't inspect secondary labels at all. So whether the label is correct for the order or in the correct location is not enforced. The Deep Vision system on the other hand, can inspect both the label accuracy and position. For promotional price labels this is particularly important as they are a legally binding offer to the customer and placing the wrong promotional label on a pack, may result in a recall, as would the wrong best before date or main label if it doesn't match the contents of the pack.

Can check the label is $6.50 pk as opposed to kg, ea or anything else which was not selected

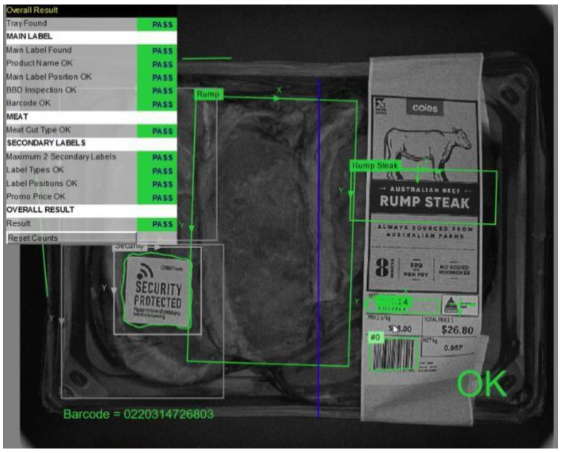

4.6.1.3.5 Label presentation

The current vision system doesn't manage label presentation well. The current system can handle a certain amount of label positional or rotational variation but these can't be enforced in a flexible manner. The Deep Vision system on the other hand, can specify the exact rotational amount that is acceptable, while anything in excess will be rejected and various positions of main and secondary label positions can be specified and enforced.

D-Wrap Label - Porterhouse

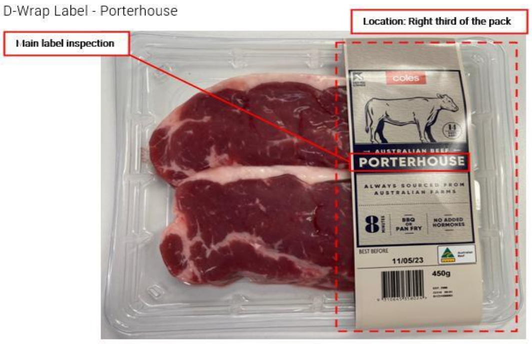

C-Wrap Label – Loin Chops

Main label position defined within the right 1/3 of the pack

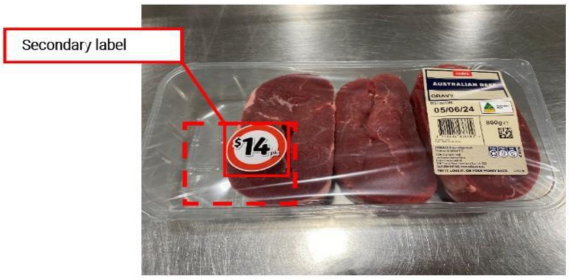

Secondary label positions defined within specific regions

Too Rotated

OK

4.6.1.3.6 Artwork quality inspection

The current vision system will fail if the locator artwork varies too much. Deep Vision can handle this variation while also being capable of alarming if the quality exceeds specified tolerances. The locator tool is used to find the location of all the elements of the label required to be inspected as well as measuring the angle of the label relative to the tray edges.

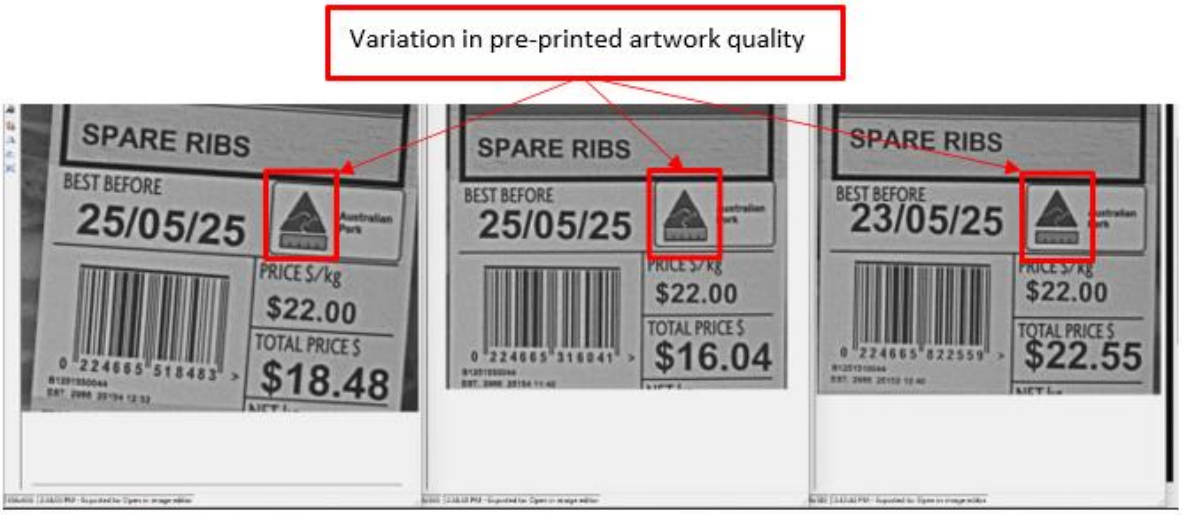

4.6.1.3.7 Inspection performance

There are several metrics that are typically used to measure the performance of a vision system. The first metric is the percentage of **False Positives**/Rejects which is where a pack is incorrectly rejected but was actually ok (ideal is 0%). The second metric is the percentage of **True Positives**/Rejects which is where a pack is correctly rejected because it was a failure (ideal is 100%). The third metric is the percentage of **False Negatives**/Accepted which is where a pack is incorrectly accepted but should have been rejected (ideal is 0%).

The Deep Vision system is in the initial phase of determining these metrics. Statistically significant volumes of data must be analysed to ensure the metrics are accurate. However initial results are as follows:

Deep Vision		**Current Vision**
False Positives	**<0.29%**	**3-5%**
True Positives	**100%**	**>95%**
False Negatives	**0%**	**5-7%**

4.6.2 Cost benefit and internal RROA business case

4.6.2.1 Brand reputation

There are several financial benefits gained from the Deep Vision system. However the primary benefit for the business is one of brand reputation, ensuring high levels of quality packaging, labelling and product are received by our customers. Each 1% loss in market share could translate to hundreds of millions of dollars in lost revenue annually.

4.6.2.2 Recall fines

If labelling is incorrect, relating to the Best Before Date, Promotional Label or Main Label, and these packs are delivered to stores, it may trigger a product recall and fine. Each recall incident attracts a $200K fine.

4.6.2.3 Product rework

If there are labelling issues initially missed on the production line, which are captured in downstream processes, there are rework costs associated with Quality Assurance labour, OutBound logistics labour, material rework costs and opportunity costs associated with the line running reworked product rather than new product.

4.6.2.4 Quality control labour

Quality control (QC) personnel are required for intermittent label inspections on each line. With

Deep Vision’s Quality Control Point compliant system, there is less need for as many inspections by QC staff. As such one QC could be used across two lines rather than on each line. This would provide a labour saving benefit to the business.

4.6.2.5 Waste and markdown

Each recall results in loss of product which represents a cost to the business.

4.6.2.6 Comparison with competitor technology

When reviewing the Cognex Smart Camera Deep Learning technology, several leading competitors were assessed. The criteria used for the comparison of the technologies was as follows.

- Use of Artificial Intelligence to make decisions more human-like and also future proof the system
- Ability to inspect labels as well as meat cut types
- Not prone to adverse effects from ambient lighting conditions
- Minimise risks from cyberattacks
- Ability for on-site staff to maintain and add products to the system with minimal or no help from the vendor
- Not require long periods of time or external vendor support for training of new products
- Easy installation within the current production constraints of space and line integration
- Automated reject system
- Ability to integrate with the current order management system
- Ability to centralise the product configuration and selection from the Control Room
- Ability to handle both VSP and MAP products of varying sizes and heights
- Limit damage to packs which are rejected from the line
- Handle inspections and rejections of packs up to a production rate of 120 packs per minutes (ppm)
- Manufacturing materials to comply with food grade production environment requirements

- Support remote access for vendor to manage updates or issue resolution

Cognex and their partners were selected as the preferred vendor to deliver a solution which satisfied the above criteria.

4.6.2.7 Quality control point compliance

The system developed in conjunction with the Cognex camera technology is one which achieves a high level of Quality Control Point (QCP) compliance.

Each selection or parameter change, is fully audit logged against individuals. The reject bin is locked and only authorised personnel can remove rejected packs. In this way if there are changes which result in a recall, the affected packs can be determined through datetime stamps and recalled efficiently with less waste. However, prevention is better than cure, and as such, the QCP features of the system ensure changes to the system are only made by qualified staff to minimise the risk of recalls altogether.

4.6.2.8 Cost benefit analysis

A Return On Investment (ROI) calculator was used to take the costs detailed in the previous sections, coupled with the development cost of the system, to determine the ROI of the Deep Vision system for one production line.

The annual savings per line, which has the Deep Vision system installed, is significant and results in a relatively quick payback period.

It should be noted that it is expected that subsequent iterations of the Deep Vision system may be lower in price, given this initial system has several development costs associated with it, which will not be required for the next units. This will result in an even shorter payback period for those units.

It should also be noted that these ROI calculations don’t take into account the costs of brand reputational damage which result in a loss of market share. Given each 1% can result in hundreds of millions of dollars, the ROI detailed above is dwarfed by the brand reputation consideration.

4.6.3 Anticipated industry outcomes

The Cognex Deep Vision system is anticipated to have a very positive impact on the red meat industry. Artificial Intelligence (AI) is making its way into every facet of life. It promises more human-like capabilities in realms, such as vision, that computers and machines have typically lagged behind humans.

The area of machine vision has long sought the benefits of AI but it has failed to deliver on these promises. Due to very costly computer infrastructure required to host the AI software, the high levels of time, and thus money, required by experts to train the system, the large amount of data required to supply the system in order to train it effectively and the inability for in-house staff to be able to make additions to the system, have kept AI from achieving mass adoption within general industrial environments, including the red meat industry.

However, with the advent of the world’s first smart camera with built-in AI, Cognex have finally brought a system to market which overcomes these technology hurdles to a point which makes it accessible to most general industries, including red meat.

As such, most primary and secondary red meat plants which rely on human inspections, which are prone to high levels of mistakes, can’t keep up with the speeds of production lines, are episodic rather than continuous and are prone to inconsistency, can adopt a Deep Vision inspection system. The industry will be able to achieve high levels of quality control and thus reduced recalls, lower labour costs, lower packaging rework, lower waste and markdowns, higher line utilisation, higher levels of consistency and increased levels of efficiency.

This technology is expected to have a much faster time to implementation and reduced development costs given the smart camera is much simpler to setup and configure compared to a computer-based, custom designed AI system. Additionally because the system needs only 1 to 20 training images, on average, to produce very accurate and consistent results, the time to a working solution is reduced.

4.6.4 Applications and benefits of the project and technology

The following are some of the key benefits of the Cognex Deep-Learning Artificial Intelligence Smart Camera and associated Deep Vision system built around this technology.

- Camera has built-in deep-learning algorithms – small form-factor, less prone to cyberattacks, no need for high-end PC.
- OCR with built-in fonts – no need for time consuming font training for standard fonts, can compare text against expected, fast processing time.
- Deep-learning algorithms – low number of example images required to train (1 to 20 on average) and achieve high levels of performance with human-like image recognition.
- SCADA interface which has been developed for ease of use – simple order selection, label types and values, barcode values, best before date auto calculations (ability for Quality department to override on deviations)
- SCADA audit log – allows all actions to be logged against specific users to secure against unauthorised changes and to track any recalls to specific time ranges thereby limiting the potential waste. Key feature in this system achieving Quality Control Point compliance.
- High-speed and robust reject system which is Quality Control Point compliant
- Ability to bypass the vision system and run the conveyor in Bypass Mode should the screen or camera fail.

The following are some of the key applications this system can be applied to, some of which have yet to be implemented but the system is capable of delivering.

- 100% inspection, not sample-based
- Up to 120 inspections per minute
- Automated rejection into a locked reject bin
- Main label inspections – description (pre-printed or printed), best before date, barcode, rotation and location, price, artwork elements
- Secondary label inspections – description, orientation and location
- Meat cut type inspections – primals or packs for quality or sorting applications
- Ability to check meat specifications such as external fat thickness and portion orientation in the pack
- Carcase inspections
- Grading meat
- Accurate pick and place data for automation equipment such as robotics

- Contamination detection
- Seal integrity

4.7Key findings

The result of the research and development project was the successful development of a production system which employed a Cognex smart-camera with built-in artificial intelligence deep-learning algorithms which outperformed traditional rule-based vision algorithms and performed as well as PC-based artificial intelligence algorithms, but without the need for vendor experts to add new products to the system, nor large data sets of images to train the system. The system's performance showed a low (<0.29%) false positive rate, while maintaining a 100% true positive and 0% false negative rate.

Additionally the system was developed with an interface that allowed for easy order management in a stand-alone mode as well as when integrated with an ERP system.

Finally the system delivered a quality control point compliant system, recording all system selections and changes against specific authorised users to maximise product quality while ensuring recalls and their associated losses and costs were minimised.

4.8 Benefits to industry

With favourable results having been achieved for this research project, there will be significant benefits to the red meat industry in the form of an easily implemented vision system with more human-like capabilities to handle difficult visual inspection tasks in a format which is simpler and thus more affordable than traditional AI systems, and more easily managed with in-house staff, rather than vendor experts, while still achieving high performance levels.

These benefits are applicable across the red meat industry from Primary to Secondary Processors, whether for domestic or export products, in label and meat inspection tasks.

There are benefits to the wider red meat industry for meat grading, defect detection, meat score analysis, animal welfare monitoring and automating pick and place meat packing automation applications. Given the system's ability to take any type of image and learn, the benefits are only limited by the scope of the industry's imagination for applications that could benefit from the system's capabilities.

5. Conclusion

With favourable results having been achieved for this research project, there will be significant benefits to the red meat industry in the form of an easily implemented vision system with more human-like capabilities to handle difficult visual inspection tasks in a format which is simpler and thus more affordable than traditional AI systems, and more easily managed with in-house staff, rather than vendor experts, while still achieving high performance levels.

The ability to teach the smart-camera what 'good' looks like with specific examples and for it to then recognise the class of 'good' objects from objects it has not seen before is truly remarkable. This opens up many possibilities for red meat applications such as label and meat verification, meat grading, meat scoring, animal welfare monitoring among many others.

The system works very well with a small set of example images to learn from. Having too many images in the learning set can reduce system performance due to the longer amounts of processing time required.

Cognex have a team of experts who provided assistance to understand any limitations in the system such as processing time constraints and how best to train the deep-learning system to achieve maximum results while maintaining efficiency.

The production system developed and built for this research project can be used as-is for any Secondary processor with similar label and meat cut verification requirements. Other Primary or Secondary requirements, apart from those shown in this project, can easily be added to this system with additional software development by the Cognex team.

6. Future research and recommendations

This project recommends that the industry adopt this technology for the following applications.

- Primary processor primal identification for sorting primals or ensuring correct primals are packed in cartons/crates. Could utilise the software system developed in this project with a modified mechanical structure for primals, cartons or crates

- Primary and secondary label verification to ensure the correct label is applied to product including product name (pre-printed or printed), price, best before date, 1D/2D codes and correct label position and orientation. Can utilise the production system developed through this project and teach the specific labels of the processor.

- Export label verification to ensure foreign language text is correct. Could utilise the software system developed in this project with a modified mechanical structure for cartons or crates.

- Secondary meat cut verification to ensure the meat cut packed, matches the label applied. Can utilise the production system developed through this project and teach the specific meat cuts of the processor.

The project suggests the following future research projects.

- Meat specification verification to inspect for characteristics such as overall dimensions, fat thickness, and marbling.

- Meat grading – carcase, portions

- Packaging seal integrity looking for damage or contamination

- Allergen or cooking suggestion text verification

- Meat orientation and size verification for integration with pick and place automation for packaging

- Animal welfare monitoring – detect variations in posture, gait, injurie