Vision system for the automation of ovine carcass processing

Danny Cheng and **Chee Kit Wong** and **Patrick P. K. Lim**
Research & Technical Services
Callaghan Innovation
Auckland, New Zealand
{danny.cheng, kit.wong, patrick.lim}@callaghaninnovation.govt.nz

Abstract

A vision system is presented for analysing the 3D profile of ovine carcasses for the belly rip down cut using a Microsoft Kinect. The system is an improvement over previous work, because it is significantly more robust against changes in the environment, while offering a simpler and more cost-effective setup. The integration of the vision system with a 6-axis robot manipulator is presented, leading to a fully automated vision-guided robotic system. Experimental trials involving 100 ovine carcasses were carried out, and a success rate of 99% was achieved. The vision system has a resolution of approximately 2 mm and experimental results show that it correctly identified the critical insertion point for the belly rip down cut. Challenges with the insertion point localisation are discussed and future work where improvements can be made are proposed.

1 Introduction

Belly rip down is a task performed by a butcher as part of the ovine de-pelting process, where the pelt is cut from the brisket, along the belly of the carcass, to the cod region. In the production line, this highly skilled operation is performed with the carcass hung in the 'foreleg hang-up' position and the pelt needs to be cut without perforating the paunch [van Beurden *et al.*, 2014].

To automate this operation, the process is broken down into two tasks: (i) detect the insertion point with a vision system; and (ii) perform a belly rip down cut with a specially designed tool. This paper focuses on the former task.

To perform a successful belly rip operation, the device must be inserted into the insertion point, located between the pelt and the flesh, as shown in Figure 1. Vertically, the insertion point is just above the pelt that has been de-pelted from the forelegs and brisket areas.

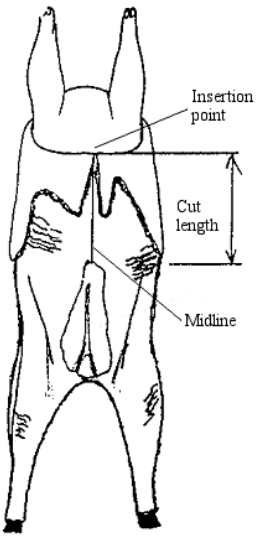

Figure 1: The belly rip down cut. The insertion point is above the flap/flesh interface.

However, pelt location varies depending on the butcher and in addition to that, the system has to cope with high degrees of variation resulting from carcass size, wool length, gender and colour. Hence, locating the insertion point is not a straightforward task.

An earlier development by Industrial Research Limited [Valkenburg and Preddey, 2000] featured a structured light system consisting of a projected laser stripe and Vision Components camera. Trial results showed the system successfully identified the insertion point in 95% of the samples, with a resolution of 1 mm [Taylor, 2000]. While the system provided good results, the field-of-view of the system was narrow and extensive screens were required to control ambient lighting. Occlusion was also an issue with the laser and camera units one metre apart.

Hurd [Hurd, 2004] explored alternative approaches, including distance, line-scan, capacitive, inductive and ultrasonic sensors to determine the height and depth of the insertion point on the carcass. Capacitive, inductive

and ultrasonic sensors were quickly eliminated as these were not suitable for the harsh production environment.

When scanning the pelt surface on a carcass, the line-scan sensor provided variable results. Single scans contained too much noise and were not sufficiently robust, however when averaging over ten scans, the insertion point can be readily identified. A single point distance sensor was also explored with it mounted on an industrial robot to perform a vertical scan. This approach required a sensor with a fast response time at a scan rate of 200 mm/s over the distance of 300 mm. The requirement took into account the speed of the production chain at a cycle time of 7.2 s per carcass allowing a maximum of 2 s for scanning. One main issue with the distance sensor is the need to perform a scan over a distance, which is time consuming. In addition, the approach also assumes that the carcass either does not swing or must be held in place so that the scan is performed down the midline.

Recent advancement in depth imaging with low-cost RGBD devices such as such the Microsoft Kinect [Microsoft, 2014] and Asus Xtion [Asus, 2014] is changing robotic perception. Specifically, these devices are capable of providing 3D data from its infrared projector and sensor, as well as RGB image from the colour sensor. The ability to align the two data easily and effectively makes these sensors very attractive for machine vision applications. As such, RGBD sensors have already been applied in various applications such as object detection [Lai *et al.*, 2011; Ren *et al.*, 2012] and tracking [Xia *et al.*, 2011; Han *et al.*, 2012; Spinello and Arras, 2011]; gesture recognition [Caputo *et al.*, 2012; Tang, 2011] and robot mapping and navigation [Meilland *et al.*, 2012; Sturm *et al.*, 2012].

The use of depth information also extends to other 3D sensing devices such as stereoscopic and time-of-flight cameras. For example, depth perception was also used for the localisation of spring onion crops for robotic manipulation in [Wong and Lim, 2012], Simultaneous Localisation And Mapping (SLAM) [Arbeiter *et al.*, 2010; Hochdorfer and Schlegel, 2010; May *et al.*, 2009] and obstacle avoidance in autonomous mobile robotics [Droeschel *et al.*, 2010; Yuan *et al.*, 2009] and 3D scene reconstruction [Feulner *et al.*, 2009; Swadzba *et al.*, 2007]. However, the fact that the RGBD devices are very low cost, accurate and reliable meant it was a straightforward decision for our application.

The paper is as arranged as follows: Section 2 presents the vision system developed for correctly detecting the insertion point on the carcass for the belly rip down cut using a Microsoft Kinect, Section 3 discusses the integration of the vision system with an industrial robotic system. The experimental setup is presented in Section 4, followed by the discussion of the results in Section 5.

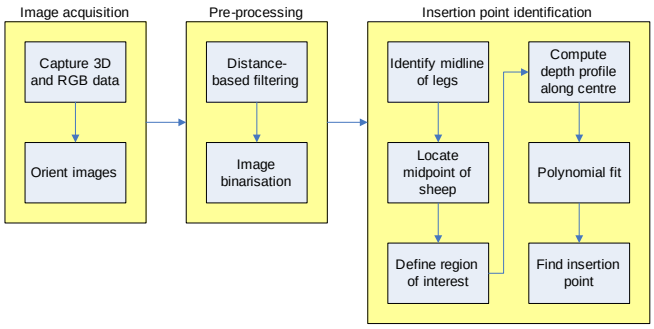

Figure 2: Flowchart of the image processing process.

Finally, conclusions are presented Section 6.

2 Vision system for insertion point detection

A novel vision system was developed for detecting the insertion point on ovine carcasses using the Microsoft Kinect. An overview of the vision system is shown in Figure 2.

2.1 Image acquisition

Using OpenNI drivers [OpenNI, 2014] for the Kinect in MathWorks MATLAB [MathWorks, 2014], first the required 3D data are acquired. The colour image is not used in this work, as the similarity in colour between the pelt and the flesh near the Region of Interest (ROI) meant a robust method for determining the insertion point was not possible. The colour image is recorded only when visual reference is required for checking the accuracy of the system.

2.2 Pre-processing

Prior to identifying the insertion point, the 3D data needs to be filtered to remove noise, as well as unwanted foreground and background. The sheep has a nominal distance of approximately 950 mm from the Kinect. Hence, a lower limit of 800 mm and an upper limit of 1150 mm was applied to the 3D data. Even though it is assumed that no objects occlude the view of the carcass, the Kinect assigns a value of zero to points which it cannot positively compute the distance. Therefore a lower limit is required to remove all 'zero' values from the dataset. Conversely, the background can be cluttered as shown in Figure 8a. If these objects are not removed in pre-processing, the task of identifying the insertion point becomes almost impossible, as it will be difficult to segment the sheep from the background. The filtered 3D data of a carcass is shown in Figure 3. Note some small objects exist in the image. If necessary, these can be easily removed by assuming the carcass is the largest object in the image, and filtering out the unwanted objects. However, this is not necessary for this application

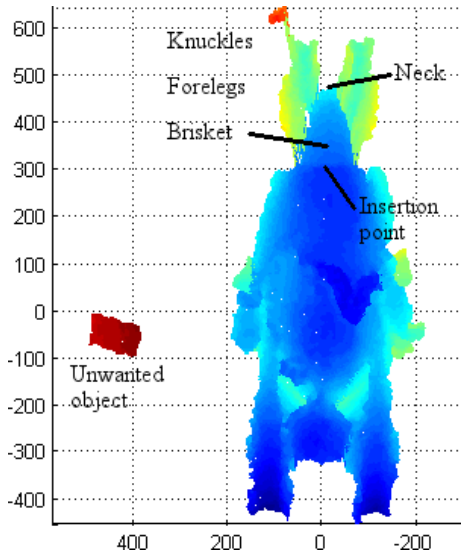

Figure 3: Mesh plot of 3D data after pre-processing.

as the ROI lies in the mid-region of the image. A binary image of the carcass is generated once the background and unwanted objects are filtered in the 3D data, by assigning a '1' to pixels that have valid depths and '0' to all other pixels.

2.3 Insertion point identification

Using the binary image, first the locations of the two forelegs are identified. The knuckles of the forelegs are used as these have shown to provide the most consistent results regardless of the size, gender and type of sheep involved. A horizontal scan along the knuckles results in four edges, two for each leg. The midline of the two legs can then be computed by taking the middle value of the two edges. Using the midlines of the forelegs, the midline of the carcass is located.

Given the location of the midline of the carcass, the location of the neck can be assumed to be the first '1' pixel from the top of the binary image.The ROI used for the remainder of the process is then defined as a function of the midline of the carcass through the neck location. A similar approach was also applied to the hind legs to determine the location of the bottom of the carcass. An algorithm for identifying the centre of the carcass is advantageous over the previous approach of using trigger sensors [Taylor, 2000], as it provides a generic approach suitable for any production chain, as additional sensors are not required for identifying the centre of the carcass.

The profile of the carcass along the midline is generated using the y- and z-components of the point cloud data, as shown in Figure 4. A 20^{th} order polynomial is fitted to the profile (1), with the first (2) and second (3) derivatives of the polynomial providing the gradient and curvature for each point of the profile, respectively

$$f(p) = c_1 p^n + \cdots + c_{n-1} p^2 + c_n p + c_{n+1} \quad (1)$$

$$f'(p) = n c_1 p^{n-1} + \cdots + 2 c_{n-1} p + c_n \quad (2)$$

$$f''(p) = n(n-1) c_1 p^{n-2} + \cdots + 2 c_{n-1} \quad (3)$$

Where $n = 20$, the order of the polynomial, $c_1 \ldots c_n$ are the computed coefficients for each carcass, and p the depth of the points along the depth curve.

Using (3) a set of local minimas is then computed by finding all $f'(p) = 0, f''(p) > 0$. The local minimas are of particular interest as they represent potential insertion points.

The local minimas close to the neck are then removed and the insertion point is then

$$\max(f''(p)), f'(p) = 0 \quad (4)$$

Whilst a 20^{th} order approximation seems excessive, variations between carcasses and the need for accurate localisation of the insertion point meant that a high order was a necessity. This was confirmed during the experimental phase.

3 Vision-guided robotic system for ovine carcass processing

Once the insertion point is identified by the vision system, the robot needs to move the specially designed tool to the insertion point to perform the belly rip down cut. The integration and calibration of the vision system with the robotic system is discussed in this section.

3.1 System setup

In order to perform the belly rip down cut along the belly of the carcass, a vision-guided robotic system is used and is shown in Figure 4. It shows a static frame where a carcass is hung for experimental trials before the robot is placed on the production chain. It avoids disruptions to the production line during development of the vision-guided robotic system. The carcass is hung in the 'foreleg hang-up' position, where the forelegs are mounted on hooks. The Kinect is mounted on a custom-made tripod, which can move in front of the carcass, and moved out of the way when the robotic system performs the cut. A fixed stop for the camera tripod ensures the camera returns to the same location every time, avoiding the need for recalibration. In addition, a bull's eye level was also used to ensure the Kinect camera is consistently mounted when removed from the tripod. The robotic system used is a Stäubli TX200 [Stäubli, 2014].

3.2 Calibration

A vision-guided robot has two coordinate systems for the vision system and the robotic system. These coordinate

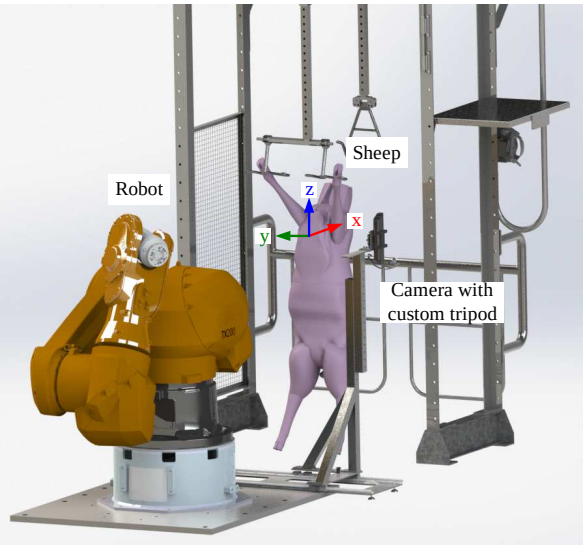

Figure 4: Vision-guided robotic system for automated belly rip down.

systems are calibrated and aligned correctly to ensure that the robot can move the tool to the location provided by the vision system accurately and repeatedly.

To align the vision coordinate system with the robot coordinate system, a calibration rig is used, as shown in Figure 5. The calibration rig is a horizontal bar, mounted between the legs of the static frame, with a number of randomly placed points at different locations in 3D. For optimal calibration, the points should cover the entire image of the vision system or the workspace of the robotic system, whichever is smaller. Two sets of calibration points were used in this work, achieved by mounting the calibration rig at different heights. This was used to achieve better coverage of the workspace. In theory, only a set of three non-degenerate points are required, however, in practice more points improve the accuracy of the calibration and are almost always used. Eight calibration points are shown in Figure 5.

To calibrate the vision system with the robotic system, first the colour image and 3D data are acquired using the Kinect. For each pixel in the colour image, a corresponding 3D point exists. By selecting the pixels which correspond to the calibration points, the relevant 3D data can be obtained. The same points are then measured by the robot, achieved by manually moving the tool centre point of the belly rip tool to the same points and recording the points' coordinates. The mapping of points from one coordinate system can be described as a function of the rotation matrix $\mathbf{R}$ and translation vector $\mathbf{t}$

$$\mathbf{x}' = \mathbf{R}\mathbf{x} + \mathbf{t} \tag{5}$$

Figure 5: Calibration rig used to calibrate the vision coordinate system with the robotic coordinate system.

Where $\mathbf{x}$ and $\mathbf{x}'$ are Cartesian coordinates of the calibration points in the Kinect and robot coordinates, respectively. Given n pairs of matching points $\mathbf{x_i}$ and $\mathbf{x'_i}$, a least squares approach is used to compute the rotation matrix and translation vector. This is obtained by first defining an error function

$$E = \sum_{i=1}^{n} \left|\mathbf{x'_i} - \mathbf{R}\mathbf{x_i} - \mathbf{t}\right|^2 \tag{6}$$

The error function is then minimised and its first order derivative calculated

$$\frac{\partial E}{\partial t} = -2\sum_{i=1}^{n} \left(\mathbf{x'_i} - \mathbf{R}\mathbf{x_i} - \mathbf{t}\right) \tag{7}$$

By rearranging (7), $\mathbf{t} = \mathbf{x'_0} - \mathbf{R}\mathbf{x_0}$, where $\mathbf{x_0}$and $\mathbf{x'_0}$ are the centroids of the two sets of points. The centred points $\mathbf{y_i} = \mathbf{x_i} - \mathbf{x_0}$ and similarly $\mathbf{y'_i} = \mathbf{x'_i} - \mathbf{x'_0}$ can be defined, resulting in

$$E = \sum_{i=1}^{n} \left|\mathbf{y'_i} - \mathbf{R}\mathbf{y_i}\right|^2 \tag{8}$$

By first solving for $\mathbf{R}$, the translation vector $\mathbf{t}$ can also be found [Forsyth and Ponce, 2011].

Experiments conducted show that a set of eight point pairs produce a mapping accuracy of approximately 1 mm, sufficient for the requirements of this work. More point pairs can increase the accuracy, however at the cost of a longer setup and calibration time.

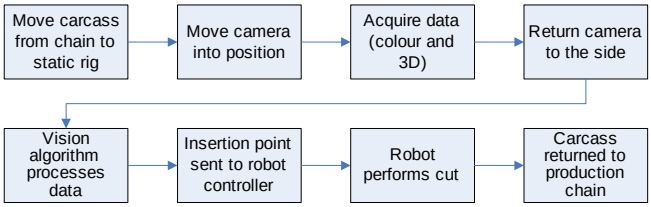

Figure 6: Flowchart of the trial process.

An important assumption when performing trials on the static rig is that the carcass stays relatively stationary when acquiring data and when the robot performs the cut. In a production environment, the chain does not stop and the carcass is tracked continuously by an encoder sensor from where data is acquired, to the workspace where the robot performs the cut. It is vital that the carcass position and orientation remain unchanged else the vision data becomes invalid, and risks a failed cut or damaging the carcass.

4 Experimental Setup

In order to verify the proposed algorithm, various trials involving a total of 100 ovine carcasses were carried out over a two month period. The aim is to identify the possible issues and further improve the approach. A flowchart of the experimental process is shown in Figure 6.

The trial was carried out at Auckland Meat Processors Limited (AMP), on a static rig as shown in Figure 4. The experimental setup at AMP is a semi-controlled environment with no significant changes in the lighting condition. Small variations in the illumination level and shadows were observed, however the algorithm was sufficiently robust and was not affected by these changes.

For each carcass, first it is taken off the production chain and moved onto the static rig. This ensures minimal disruption to the production chain, and minimises variables such as movement or a change of pose of the carcass throughout the duration of the trial, which could affect the results. The camera on the custom tripod is moved and locked in place, ensuring the same relative position to the robot every time. Once in position, the 3D data and colour image are acquired, and the camera moved out of the way of the robot. The vision system processes the data acquired, resulting in an insertion point for the tool in the vision coordinate system. This insertion point is transformed to the robot coordinate system and sent to the robot controller. A visual inspection of the location of the insertion point is carried out to ensure that the insertion point is correct, before the cut is carried out by the robot. The tool is sanitised before the next sheep is put on the static rig to eliminate cross-contamination.

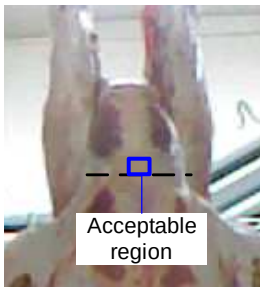

Figure 7: Acceptable region for the insertion point.

Number of carcasses	Success	Failure
100	99	1

Table 1: Summary of results from trials involving a total of 100 carcasses.

5 Results and discussions

The results of the trials are summarised in Table 1. A success is defined as when the vision system is able to identify the insertion point correctly, within the region 20 mm above the flap, but not below, as shown in Figure 7. A failure is when the algorithm is not able to identify the insertion point. A set of representative results are shown in Figure 8.

Of the 100 ovine carcasses, the vision system achieved a success rate of 99%. Furthermore, the robot was also able to move to the location of the correct insertion point as indicated by the vision system, demonstrating that the calibration approach is working correctly. The resolution of the Kinect at the distance involved, and therefore the insertion point detection, is estimated to be approximately 2 mm. This is ample for the work, as the tolerance for the insertion point is significantly larger than the resolution of the Kinect based on results from the trials.

One of the main advantages of the proposed approach is its ability to cope with naturally occurring variations in its subjects. Trials were carried out over a two month period, and hence the system was exposed to different batches of sheep. Results from the trials prove the algorithm was successful when applied to carcasses of varying: (i) size; (ii) wool length; (iii) gender; (iv) shape; and (v) colour. Further investigation should be carried out for different sheep breeds. This was not possible at the time of trial due to limitations in the breeds available.

Furthermore, it is also worth noting that the operating environment was not altered or constrained in any form or manner, for example the lighting condition or the geometry of the hooks used to hold the carcasses. This ensured that the trials were carried out in a realistic environment, and can be easily transferred to production

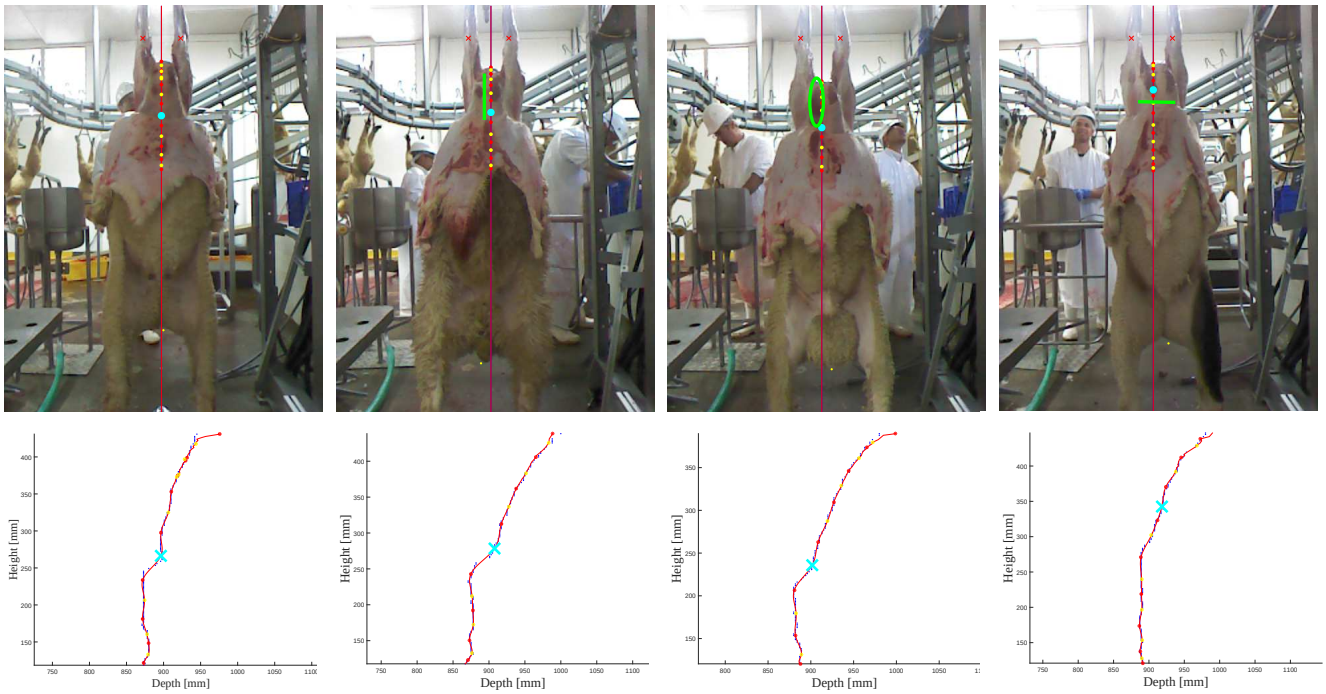

(a) Good.

(b) Rotation about the longitudinal axis. The actual midline is shown in green.

(c) Excess fat removed. The exposed flesh region is enclosed in green.

(d) Too high. The required height is shown in green.

Figure 8: Graphs of carcasses with the depth line shown and the detected insertion point marked in blue.

chains without significant further work.

5.1 Rotation about the longitudinal axis

Of the 99 successful carcasses, 13% were observed to be rotated about the longitudinal axis, or skewed. An example of this is shown in Figure 8b. Regardless of the rotation of the carcass, the vision system is able to successfully detect the height of the insertion point for all the carcasses which were skewed. However, trials have shown that if the tool is not inserted into approximately the midline of the carcass, the risk of a failed belly rip down cut significantly increases.

The challenge with skewed carcasses is that the rotation of the carcass significantly deviates the insertion point away from the true midline of the carcass. Due to natural variation of the carcass, this rotation is not easily quantifiable for the system to compensate for. As the belly rip down cut requires the insertion to be as close to the midline of the carcass as possible, an off-centre insertion will result in the tool penetrating the carcass or withdrawing from the pelt. Preliminary studies show that it may be possible to detect carcass rotation based on the depth profile around the foreleg shoulders and brisket region, but further investigation is needed.

5.2 Excess fat removed

Earlier trials with a handheld version of the tool showed that the tool needs be inserted into the fat regions and not flesh. One of the issues noted during the trial was that on some carcasses, excess fat would be removed from the carcass, resulting in a very small or no region left on the carcass for the tool to make contact, as shown in Figure 8c. This is a result of variation in the human operation.

When small amounts of fat are left on the carcass, it is desirable to shift the insertion point onto a fat region if it is not far from the centre of the carcass. If no fat remains, then it is not possible for the tool to correctly perform the cut and should therefore be aborted.

5.3 Incorrect height

The vision algorithm detects the insertion point by identifying the local minima in the fitted polynomial line with the greatest curvature. Figures 8d show that the local minimas in the polynomial lines can appear very similar depending on the profile of the carcass. Due to these similarities, the algorithm could mistake another part of the carcass as the insertion point which would either be too high or too low.

A possible solution is to approximate the likely location of the flap based on the length of the carcass. Since

the top and bottom of the carcass are identifiable, the length can be easily computed. Whilst a correlation has been visually observed, further investigations will be required to determine if a statistical correlation exists for this approach to work.

6 Conclusions

A novel vision system for detecting the insertion point on the pelt/flesh interface of ovine carcasses is presented. The vision system relies only on the 3D data from the Microsoft Kinect, where the depth profile of the carcass is analysed and the insertion point identified. The integration of the vision system with a robotic system to form a vision-guided robotic system is also presented. A calibration accuracy of approximately 1 mm was achieved for mapping points from the vision system to the robotic system.

Results from trials of 100 carcasses show that the presented vision system is capable of handling carcasses of different sizes, wool lengths, genders and colours, with a success rate of 99%, an improvement over the structured lighting approach used previously [Taylor, 2000]. The vision system has a resolution of 2 mm, which is sufficient based on the results from trials. Practical issues were identified which could affect the belly rip down cut using the robotic system, including rotation of the carcass about the longitudinal axis, and excess fat removed from the flesh. While the vision system correctly identified the correct height for the insertion point in all the cases, it is not possible for the belly rip down tool to correctly insert and perform a cut in these instances. Changes in illumination levels were consistent with those expected in a typical ovine carcass processing environment and did not affect the results. Possible improvements were discussed and trials will be conducted as part of the future work for improving the success rate of the system.

Acknowledgements

This research was supported by the Ovine Automation Ltd (OAL) consortium and Auckland Meat Processors Limited (AMP), who provided the carcasses and facilitated our vision and robotic systems on the slaughter floor for experimental trials.

References

[Arbeiter *et al.*, 2010] Georg Arbeiter, Jan Fischer, and Alexander Verl. 3-d-environment reconstruction for mobile robots using fast- slam and feature extraction. In *Robotics (ISR), 2010 41st International Symposium on and 2010 6th German Conference on Robotics (ROBOTIK)*, pages 1–5, June 2010.

[Asus, 2014] Asus. Xtion pro. [online]: http://www.asus.com/Multimedia/Xtion_PRO/, 2014. Accessed 2014.

[Caputo *et al.*, 2012] Manuel Caputo, Klaus Denker, Benjamin Dums, and Georg Umlauf. 3d hand gesture recognition based on sensor fusion of commodity hardware. In Harald Reiterer and Oliver Deussen, editors, *Mensch & Computer 2012: interaktiv informiert allgegenwrtig und allumfassend!?*, pages 293–302, Mnchen, 2012. Oldenbourg Verlag.

[Droeschel *et al.*, 2010] David Droeschel, Dirk Holz, Jrg Stckler, and Sven Behnke. Using time-of-flight cameras with active gaze control for 3d collision avoidance. In *In Proc. of the IEEE International Conference on Robotics and Automation (ICRA*, 2010.

[Feulner *et al.*, 2009] J. Feulner, J. Penne, E. Kollorz, and J. Hornegger. Robust real-time 3d modeling of static scenes using solely a time-of-flight sensor. In *Computer Vision and Pattern Recognition Workshops, 2009. CVPR Workshops 2009. IEEE Computer Society Conference on*, pages 74–81, June 2009.

[Forsyth and Ponce, 2011] D. A. Forsyth and J. Ponce. *Computer vision, a modern approach*. Prentice Hall, 2nd edition, 2011.

[Han *et al.*, 2012] Jungong Han, E.J. Pauwels, P.M. de Zeeuw, and P.H.N. de With. Employing a rgb-d sensor for real-time tracking of humans across multiple re-entries in a smart environment. *Consumer Electronics, IEEE Transactions on*, 58(2):255–263, May 2012.

[Hochdorfer and Schlegel, 2010] S. Hochdorfer and C. Schlegel. 6 dof slam using a tof camera: The challenge of a continuously growing number of landmarks. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pages 3981–3986, Oct 2010.

[Hurd, 2004] S. Hurd. Automated ripdown: sensing review. Industrial Research Limited, 2004.

[Lai *et al.*, 2011] Kevin Lai, Liefeng Bo, Xiaofeng Ren, and Dieter Fox. Sparse distance learning for object recognition combining rgb and depth information. In *International Conference on Robotics and Automation*, pages 4007–4013, 2011.

[MathWorks, 2014] MathWorks. Matlab - the language of technical computing. [online]: http://www.mathworks.com.au/products/matlab/, 2014. Accessed 2014.

[May *et al.*, 2009] Stefan May, Stefan Fuchs, David Droeschel, Dirk Holz, and Andreas Nchter. Robust 3d-mapping with time-of-flight cameras. In *In IROS*, pages 1673–1678, 2009.

[Meilland *et al.*, 2012] M. Meilland, P Rives, and A. I. Comport. Dense RGB-D mapping for real-time localisation and navigation. In *IV12 Workshop on Navigation Positionnig and Mapping*, Alcalá de Henares, Spain., June 3 2012.

[Microsoft, 2014] Microsoft. Kinect for windows. [online] http://www.microsoft.com/en-us/kinectforwindows/, 2014. Accessed 2014.

[OpenNI, 2014] OpenNI. Openni/openni github. [online]: https://github.com/OpenNI/OpenNI, 2014. Accessed 2014.

[Ren *et al.*, 2012] X. Ren, L. Bo, and D. Fox. Rgb-(d) scene labeling: Features and algorithms. In *Computer Vision and Pattern Recognition*, pages 2759–2766, 2012.

[Spinello and Arras, 2011] Luciano Spinello and Kai O. Arras. People detection in rgb-d data. In *In IEEE/RSJ Int. Conf. on*, 2011.

[Stäubli, 2014] Stäubli. Tx200 6-axis heavy payload robot. [online] http://www.staubli.com/en/robotics/6-axis-scara-industrial-robot/heavy-payload-robot/6-axis-industrial-robot-tx200/, 2014. Accessed 2014.

[Sturm *et al.*, 2012] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers. A benchmark for the evaluation of rgb-d slam systems. In *Proc. of the International Conference on Intelligent Robot Systems (IROS)*, Oct. 2012.

[Swadzba *et al.*, 2007] Agnes Swadzba, Bing Liu, Jochen Penne, Oliver Jesorsky, and Ralf Kompe. A comprehensive system for 3d modeling from range images acquired from a 3d tof sensor. In *International Conference on Computer Vision Systems*, International Conference on Computer Vision Systems. University Library of Bielefeld, 2007.

[Tang, 2011] M. Tang. Recognizing hand gestures with microsofts kinect. Technical report, Department of Electrical Engineering, Stanford University, 2011.

[Taylor, 2000] M. G. Taylor. Automated ripdown project at alliance stockburn. Industrial Research Limited, 2000.

[Valkenburg and Preddey, 2000] R. J. Valkenburg and J. T. Preddey. Robust 3d feature location using structured light. Industrial Research Limited, 2000.

[van Beurden *et al.*, 2014] J. van Beurden, D. Cheng, K. Wong, C. Cheam, and K. Heffel. Oal integrated belly rip down. Industrial Research Limited, 2014.

[Wong and Lim, 2012] Chee Kit Wong and P.P.K. Lim. Processing of point cloud data from tof camera for the localisation of ground-based crop. In *Mechatronics and Machine Vision in Practice (M2VIP), 2012 19th International Conference*, pages 184–189, Nov 2012.

[Xia *et al.*, 2011] Lu Xia, Chia chih Chen, and J. K. Aggarwal. Human detection using depth information by kinect. In *International Workshop on Human Activity Understanding from 3D Data in conjunction with CVPR (HAU3D)*, 2011.

[Yuan *et al.*, 2009] Fang Yuan, Agnes Swadzba, Roland Philippsen, Orhan Engin, Marc Hanheide, and Sven Wachsmuth. Laser-based navigation enhanced with 3d time-of-flight data. In *IEEE International Conference on Robotics and Automation*, International Conference on Robotics and Automation, pages 2844–2850. IEEE, 2009.