Meat Science 96 (2014) 366–372

Contents lists available at ScienceDirect

Meat Science

journal homepage: www.elsevier.com/locate/meatsci

Vision-based method for tracking meat cuts in slaughterhouses

Anders Boesen Lindbo Larsen [a,*], Marchen Sonja Hviid [b], Mikkel Engbo Jørgensen [b], Rasmus Larsen [a], Anders Lindbjerg Dahl [a]

[a] Technical University of Denmark, Lyngby, Denmark
[b] Danish Meat Research Institute, Roskilde, Denmark

ARTICLE INFO

Article history:
Received 16 April 2013
Received in revised form 10 July 2013
Accepted 16 July 2013

Keywords:
Tracking
Traceability
Computer vision
Image processing
Object recognition

ABSTRACT

Meat traceability is important for linking process and quality parameters from the individual meat cuts back to the production data from the farmer that produced the animal. Current tracking systems rely on physical tagging, which is too intrusive for individual meat cuts in a slaughterhouse environment. In this article, we demonstrate a computer vision system for recognizing meat cuts at different points along a slaughterhouse production line. More specifically, we show that 211 pig loins can be identified correctly between two photo sessions. The pig loins undergo various perturbation scenarios (hanging, rough treatment and incorrect trimming) and our method is able to handle these perturbations gracefully. This study shows that the suggested vision-based approach to tracking is a promising alternative to the more intrusive methods currently available.

1. Introduction

In recent years, traceability has become an increasingly important aspect of the meat industry. For consumers, meat safety and quality is a persistent concern strengthened by reoccurring food recalls and scandals as well as increased animal welfare awareness (Vanhonacker, Verbeke, Poucke, & Tuyttens, 2008). In Western markets, this public concern has lead to legislations and regulations regarding food traceability to ensure quality and safety standards (Trienekens & Zuurbier, 2008). For producers, traceability adds extra value to their end products (Wang & Li, 2006). Demand for traceability information is on the rise yielding a competitive advantage to the producers who can deliver better guarantees of origin and handling (Buhr, 2003; Carriquiry & Babcock, 2007; Pouliot & Sumner, 2008).

In industrial abattoirs individual meat cuts become hard to trace after having cut up the carcass. Today most tracking systems are based on secondary systems like boxes or Christmas trees with RFID technology or conveyor belts. These systems offer only batch-level tracking of meat cuts because the secondary devices cannot be attached to the products individually.

In this work we propose a new technology for enabling meat traceability of individual meat cuts in slaughterhouse environments. Our approach is based on modern methods from the field of computer vision and image processing. Instead of attaching identification information to an object in order to track it we capture an image of the object and can identify the same object at a later point by capturing a new image. That is, we extract a *description* of an object from its appearance and use it as identifier for that object. We believe that this approach offers attractive advantages compared to current technology. While our experiments are limited to tracking pork loins, the method is sufficiently generic to be applied in other domains where the objects exhibit adequate diversity in appearance like the meat cuts considered in this work.

1.1. Related work

Food traceability has been approached from many angles with different applications in mind. This has led to a diverse literature with a limited agreement on how to implement food traceability. For an overview of food traceability literature, we refer to Karlsen, Dreyer, Olsen, and Elvevoll (2013).

In this article, we focus on a single aspect of traceability in the meat industry; the technology that enables object tracking along a production line. In recent literature, the use of RFID tags as underlying food tracking technology is dominating (Cimino & Marcelloni, 2012; Lefebvre, Castro, & Lefebvre, 2011; Regattieri, Gamberi, & Manzini, 2007). However, RFID tagging of meat in a slaughterhouse environment has drawbacks for mainly one reason: Tags may disappear into the meat product and turn up on the consumer's plate. This is a very critical point with the consequence that slaughterhouses avoid tagging meat cuts directly; instead they attach a tag to the device carrying the meat.

Regarding tracking technology in the meat industry, the following approaches have been suggested. Mousavi, Sarhadi, Fawcett, Bowles, and York (2005) present a conveyor belt system capable of tracking

* Corresponding author at: Richard Petersens Plads, bldg. 324, 2800 Lyngby, Denmnark. Tel.: +45 45255228.
E-mail addresses: abll@dtu.dk (A.B.L. Larsen), mahd@teknologisk.dk (M.S. Hviid), mej@teknologisk.dk (M.E. Jørgensen), rlar@dtu.dk (R. Larsen), abda@dtu.dk (A.L. Dahl).

0309-1740/$ – see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.meatsci.2013.07.023

meat cuts in a boning hall. To facilitate the tracking, RFID chips are embedded in carrier hooks for the meat cuts. Fröschle, Gonzales-Barron, McDonnell, and Ward (2009) examine the usability of barcodes printed on the beak and legs of chickens. This approach does not generalize well to other meat tracking scenarios because it requires the meat product to have non-edible parts suitable for barcode printing. Arana, Soret, Lasa, and Alfonso (2002); Suekawa et al. (2010) perform breed identification of beefs based on DNA analysis, and Tate (2001) investigates the possibility of using DNA analysis for tracing individual meat cuts back to the original carcass. Our vision-based approach is reminiscent of DNA identification in the way identification is derived from the object rather than from a tag attached to the object. However, DNA identification is still a cumbersome process for a slaughterhouse environment.

Of the three tracking technologies mentioned above, the conveyor belt system is most representative of current slaughterhouse practice. Typically, meat cuts are tracked individually or in a batch by attaching a tag to the container or carrier device. A drawback of this method is that it is prone to accidents where pieces are lost or exchanged between carrier devices. Such accidents may happen since the meat cuts cannot be directly connected to the carrier device at all times. With a vision-based approach, this scenario will not be a problem since the meat cut carries identification in its appearance.

For both the food and the non-food industry we have not been able to find examples of visual recognition methods similar to ours applied in a tracking/identification setup. Weichert et al. (2010) propose combining RFID tracking with a vision system that can recognize and decode 2D barcodes. Using cheap cameras they can offer a more continuous identification and localization of the products and thereby improve fault detection. Again, this approach is not viable for meat cuts as the goal is to avoid foreign objects (both barcodes and RFID tags) that can end up in the product. Therefore, to the best of our knowledge, tracking from visual recognition of the products directly has not been attempted before.

1.2. Contributions

In this work we investigate a new technology for enabling traceability of individual meat cuts in a slaughterhouse environment. The investigation extends the work presented by Hviid, Jørgensen, and Dahl (2011) by scaling up the experiment to 211 pork loins and introducing *nuisance factors* to simulate a slaughterhouse environment. We show that the pork loins can be recognized and identified correctly between the two photo sessions. These results indicate that current computer vision methods for object recognition are mature for integration in production lines.

2. Experiment setup

The dataset for our experiment is constructed using 211 pig loins. The pig loins are photographed in two sessions separated by 1 day. Overnight, the loins are hanging on *Christmas trees* stored in a chill room, see Fig. 1.

The photographing setup (see Fig. 2) is the same for both photo sessions. We use the popular and inexpensive Microsoft Kinect camera that captures a depth map along with a standard RGB image of the loin. Examples of both images are shown in Fig. 3. Next to the camera a fluorescent tube is mounted spreading light at a wide angle. The loins are photographed separately by placing them one by one on a table and capturing a photo.

A selection of the loins undergoes different perturbation scenarios in an attempt to simulate slaughterhouse treatment. All perturbations occur after the first and before the second photo session. The perturbations are:

Rough treatment 19 loins are knocked hard onto a table before the second photo session.

Fig. 1. Pork loins are stored overnight on *Christmas trees* between the two photo sessions.

Incorrect trimming Pieces of meat and bones are cut off from 18 loins before the second photo session.

Incorrect hanging 19 loins are stored overnight by hanging them sideways on Christmas trees which causes bends.

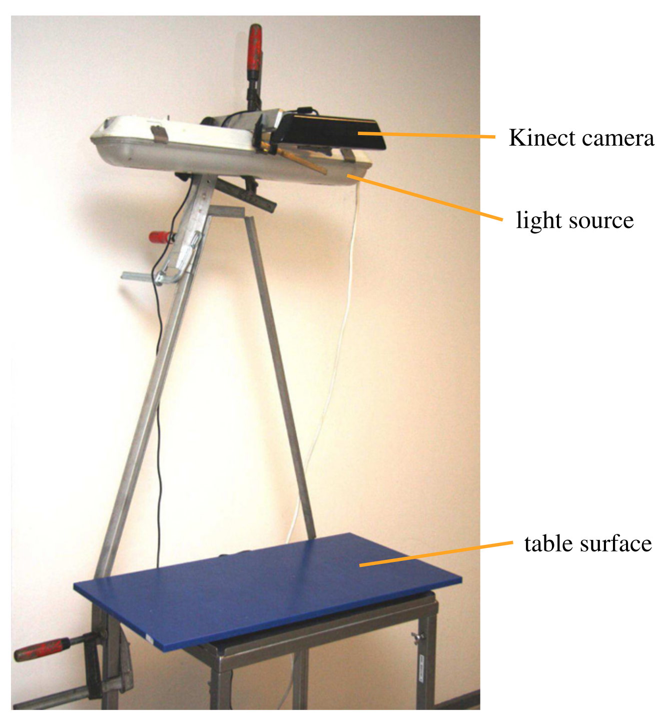

Fig. 2. Camera setup. Pork loins are placed on the table and are photographed from above.

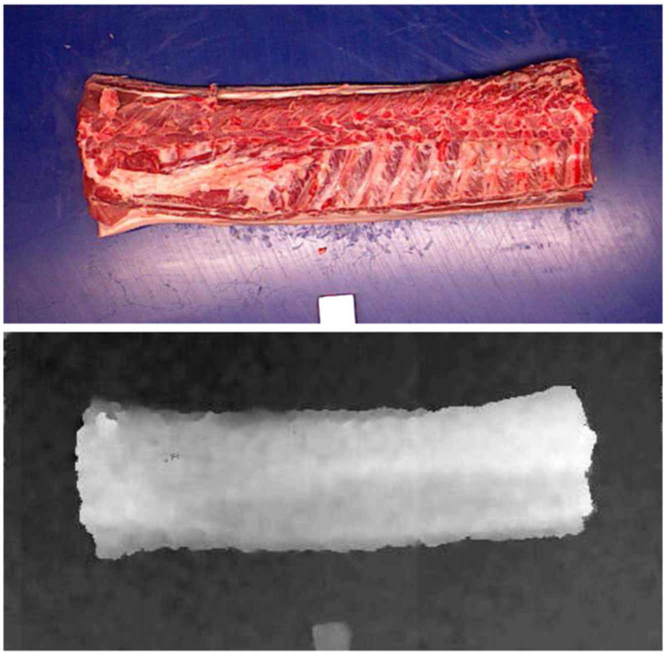

Fig. 3. RGB and depth images of a pork loin as captured by the Kinect camera.

Fig. 4. Pork loin segmentation. Top image shows the segmentation mask derived from the depth image (see Fig. 3). Bottom image shows the pork loin cut out using the segmentation mask.

Illumination and orientation changes 37 loins are rotated between 45° and 180° around the optical axis before being photographed. This creates variations in lighting because the light falls differently on a rotated object. Moreover, the rotated loins serve as a check to see if our algorithm is invariant to different object orientations.

All loins except those subject to incorrect hanging are photographed normally on day 2 before any perturbations occur. Because some loins are reused in multiple perturbation scenarios (e.g. a loin is photographed at a different orientation and again after a trimming), we cannot perform a matching on all 211 loins by combining all perturbations. Instead, we combine each perturbation scenario with the remaining unperturbed images from day 2 in 4 separate experiments. This means that for the incorrect hanging scenario we want to match all 211 loins whereas for the other scenarios we want to match 192 loins.

3. Visual recognition method

The purpose of the visual recognition method is to match the pork loin images correctly between the two photo sessions. Our recognition method is divided into 4 steps listed here as an overview.

1. *Segmentation*. First, we perform a segmentation of the pork loin. That is, we cut the pork loin out from the background image pixels.
2. *Canonization*. The segmented pork loin images are then brought to a canonized form that minimizes variability from external sources, e.g. illumination.
3. *Description*. From the canonized images we generate a description of the image structure.
4. *Matching*. Finally, we perform the pork loin matching by comparing the descriptors from the previous step.

In total, the recognition method takes under 2 s per image in processing time on a 2.67 GHz CPU. It should be possible to speed this up significantly since our method has not been implemented with speed efficiency as a priority. Note that we have made all implementation details available online.[1]

3.1. Segmentation

To separate the loin from the background we use the depth image provided by the Kinect camera. We know the depth of the table surface which makes it easy to differentiate between the surface and the meat. To account for noisy depth data, we employ the *max-flow/min-cut* graph cut algorithm to perform segmentation of the depth image (Boykov & Kolmogorov, 2004). This yields a binary *mask* specifying which pixels belong to the loin and which pixels belong to the background. The result of the segmentation algorithm is shown in Fig. 4.

3.2. Canonization

The goal of the canonization step is to bring the pork loin images to a common form making them invariant to changes in *illumination*, *rotation* and *size*.

Since the pork loin primarily consists of red color nuances we can discard the colors by converting the image to greyscale without losing significant information. Moreover, we perform a *histogram equalization* to increase the contrast and compensate for differences in lighting.

To ensure the same orientation for all pork loins, we use the segmentation mask and calculate the *image moments* of the mask region. The second-order moments can be used to derive the covariance matrix of the image region. The dominant orientation of the region is then calculated from the angles of the covariance matrix eigenvectors. We rotate the loin such that the dominant orientation is parallel to the x axis, that is, the loin is orientated horizontally along its broad side. Notice that this rotation does not consider if the loin is placed upside-down. We handle this situation by performing a pixel-wise comparison of a loin image with the average of 20 upright loin images. If a loin is pixel-wise closer to the upside-down version rather than upright version, it should be rotated 180°.

[1] Source code available at http://compute.dtu.dk/~abll/meat_recognition.

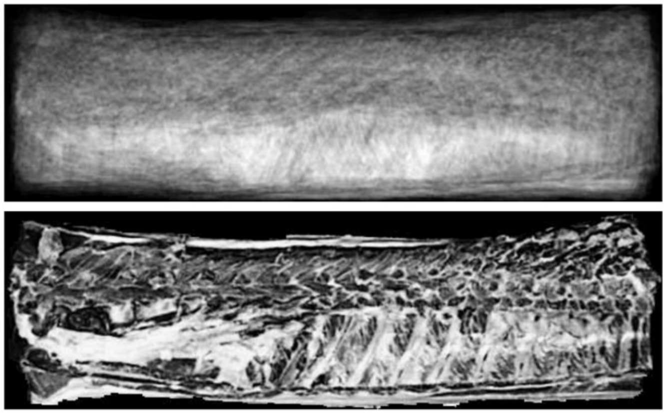

Fig. 5. The result of the canonization step. Top image shows an average of 20 pork loin images with the same orientation. The average image is used to check if loin images are placed upside down. Bottom image shows a canonized loin image.

Finally, the pork loin images are trimmed to remove the background border followed by a scaling to 600 × 180 pixels giving all loins the same dimensions. An example of the canonization is shown in Fig. 5.

3.3. Description

In the image description step we seek to achieve an image representation that captures the image structure in a manner suitable for comparative purposes. E.g., the standard pixel representation is not suitable because it is sensitive to object translations.

We employ the popular *bag-of-words* approach (Prince, 2012) and perform *K-means clustering* on image patches extracted from 30 out of the 211 loin images. The cluster centers yield a finite set of different image patches (the *visual vocabulary*). An image can now be described by extracting numerous image patches and mapping each patch to its nearest entry (aka. *visual word*) in the vocabulary and counting the number of occurrences of each visual word. The bag-of-words image characterization thus constitutes a histogram over visual words. An overview of bag-of-words description is shown in Fig. 6.

3.3.1. Feature description

Instead of using raw image patches as mentioned above, we perform *feature description* of these patches. Feature description yields a low-dimensional representation of an image patch that attempts to capture the image structure while being invariant to various image perturbation factors. A wide selection of feature descriptors exists in the literature but their performance varies only little for general applications (Dahl, Aanæs, & Pedersen, 2011; Kaneva, Torralba, & Freeman, 2011). We use the DAISY descriptor (Tola, Lepetit, & Fua, 2010) as it is formulated for dense extraction. However, we use our own variation of DAISY since the original is based on unsigned gradient orientations spanning 180° (whereas signed orientations span 360°). Unsigned orientations offer invariance towards more complicated illumination situations where light areas become dark and vice versa due to surface reflectance properties. Our scenario is sufficiently constrained allowing us to benefit from signed orientations.

DAISY follows a popular approach to feature description based on image gradient orientations summarized in histograms. From an image I we can extract the x and y directional derivatives,

$$\mathbf{L}_x = \frac{\partial \mathbf{G}_\sigma}{\partial x} * \mathbf{I} \quad , \quad \mathbf{L}_y = \frac{\partial \mathbf{G}_\sigma}{\partial y} * \mathbf{I} \quad , \tag{1}$$

by convolving the image with a Gaussian window $\mathbf{G}$ differentiated along the x and y axis respectively. * denotes convolution and σ adjust the width of the Gaussian kernel, i.e., the scale at which we extract the derivatives. We can then calculate the image gradient orientations θ and their magnitude $\mathbf{m}$ from

$$\theta = \arctan 2\left(\mathbf{L}_x, \mathbf{L}_y\right) \quad , \quad \mathbf{m} = \sqrt{\mathbf{L}_x^2 + \mathbf{L}_y^2}. \tag{2}$$

To describe the gradient orientation statistics we select a number of bins N in the angular range. For each angle a_i,$i = \{1,\ldots,N\}$ we calculate the bin contribution $\mathbf{b}_i$ using the circular normal distribution to smooth out contributions among neighbor bins to be invariant towards small rotations. Moreover, we weigh the bin contributions by the gradient magnitude such that small gradients have less influence than large gradients.

$$\mathbf{b}_i = exp\ (\kappa * cos\ (\theta - a_i)) \circ \mathbf{m} \tag{3}$$

$\circ$ denotes the element-wise product. κ adjusts the scale of the bin smoothing in the angular range. To gather bin contributions spatially, we convolve with a Gaussian window of scale γ.

$$\mathbf{b}_{\gamma,i} = \mathbf{G}_\gamma * \mathbf{b}_i \tag{4}$$

Finally, we assemble a histogram at the spatial location (u,v) from

$$\mathbf{h}_\gamma(u,v) = \left[\mathbf{b}_{\gamma,1}(u,v), \ldots, \mathbf{b}_{\gamma,N}(u,v)\right]. \tag{5}$$

To perform DAISY description of an image patch, we sample $\mathbf{h}_\gamma$ in a log-polar grid similar to the original formulation. The histogram vectors are then concatenated and the entire descriptor vector is L1-normalized. This normalization makes the descriptor invariant to affine illumination variations.

Feature descriptors like DAISY have become very popular for visual recognition. They are effective at capturing both texture information and local image structure while being robust towards various image perturbations.

3.3.2. Pork loin image representation

The bag-of-words representation disregards all spatial layout of the extracted image patches. This is good for achieving invariance to object translations, but not so good for providing a distinctive object description. We reestablish some of the spatial layout information by sampling multiple bag-of-words histograms at different positions in the image. More specifically, we generate 8 histograms from the pork loin image by weighing the different histogram contributions using Gaussian windows placed in a 2 × 4 grid to reflect the oblong shape of a pork loin. See Fig. 7. The reason for using Gaussian windows to gather bin contributions is because the smoothed weighting handles object translations more gracefully leading to a more robust description.

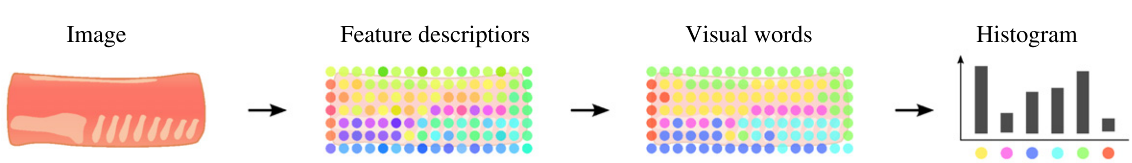

Fig. 6. Overview of the description pipeline. From the raw image we perform a dense extraction of local feature descriptors. The feature descriptors are then quantized into visual words. Finally the occurrences of visual words are summarized in a histogram that becomes the final image description.

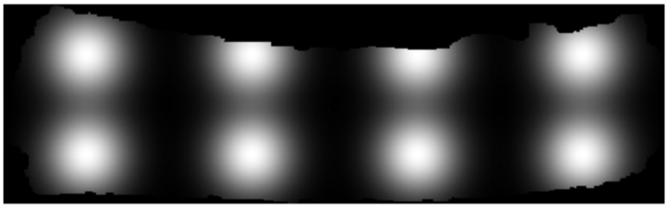

Fig. 7. The visual word contribution to each bag-of-words histogram is weighted using a Gaussian window. Visual words that lie outside the segmentation have 0 weight.

The final description of a pork loin image is the concatenation of the 8 bag-of-words histograms. Our bag-of-words vocabulary consists of 1500 visual words meaning that each histogram can be represented by a 1500-dimensional vector. Thus, the concatenation of the 8 bag-of-words histograms yields a 12,000-dimensional image description vector.

3.4. Matching

We assess the similarity of two pork loin images by calculating the histogram distance between their two description vectors generated in the previous step. For every pork loin from day 1 a match is established to the pork loin from day 2 with the smallest χ^2 distance defined as

$$\chi^2(\mathbf{x},\mathbf{y}) = \sum_{n=1}^{D} \frac{(\mathbf{x}(n)-\mathbf{y}(n))^2}{\mathbf{x}(n)+\mathbf{y}(n)} \quad , \tag{6}$$

where D is the dimensionality of the vectors $\mathbf{x}$ and $\mathbf{y}$ and $\mathbf{x}(n)$ is the nth element of $\mathbf{x}$.

4. Results

We run our recognition method on the 4 different experiments listed in Section 2. In all experiments we are able to match all pork loins between the two photo sessions correctly.

To demonstrate the visual impact of the perturbation scenarios, we show examples of pork loins from both days in Fig. 8. We show the canonized versions rather than the original camera images as the canonization makes visual comparison easier.

Fig. 8 shows a loin without perturbations. i.e. proper hanging overnight. We observe both local pixel translations due to minor object deformations and global pixel translations due to improper alignment in the canonization step. In Fig. 8b and c, we observe local deformations caused by rough handling of the meat and incorrect trimming. Fig. 8d shows perturbations due to incorrect hanging overnight. The twist causes translation, local deformation in the right end of the loin, and minor local rotation. Finally in Fig. 8e and f, the illumination changes caused by object rotation are shown. These perturbations are significantly diminished by the canonization step, however, we still see that specularities and shadows change indicating that the experiment setup could be improved with a more diffuse illumination.

To investigate the robustness of the recognition method we inspect loins that have been poorly matched in our experiments. We measure the quality of a match by its distinctiveness d computed by subtracting the descriptor distance of the nearest incorrect match from the descriptor distance of the correct match.

A large difference means that the matching pork loin image pair from day 1 and 2 stands out from the rest of the loins. A small difference means that there exists a mismatching loin from day 2 with an image description similar to the pork loin from day 1. In Fig. 9, we show 3 examples of poorly matched pork loin image pairs along with the second-closest match from day 2. In Fig. 9a and b we see two examples where the appearances of the second-closest matches are similar to the loins from day 1. If a human were to tell the loins apart, he/she would most likely rely on smaller details in their appearances. In Fig. 9c a significant bend affects the aspect ratio of the loin image leading to a poor canonization caused by improper alignment. Thus, it is the canonization rather than the image description that fails.

Finally in Fig. 10, we illustrate the distinctiveness statistics for each experiment. We see that our recognition method is very close to yielding a few mismatches as the distinctiveness of the lowest outliers come close to 0 (a negative value means an incorrect match). However, the main part of the remaining loins (around 200) are matched with a comfortable margin to the nearest incorrect match.

5. Discussion

In Fig. 8, we have seen examples of different perturbation scenarios in a slaughterhouse environment. Our image description algorithm is

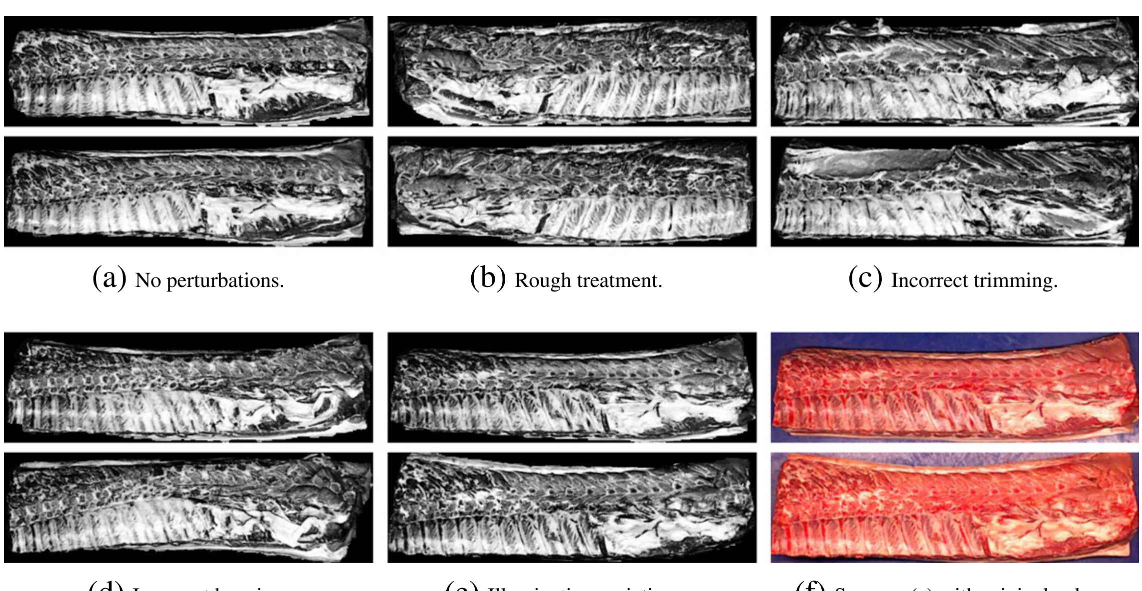

(a) No perturbations. (b) Rough treatment. (c) Incorrect trimming.

(d) Incorrect hanging. (e) Illumination variations. (f) Same as (e) with original colors.

Fig. 8. Examples of perturbation scenarios between day 1 (upper image) and day 2 (lower image). Canonized images are shown for better visual comparison (except for (f)).

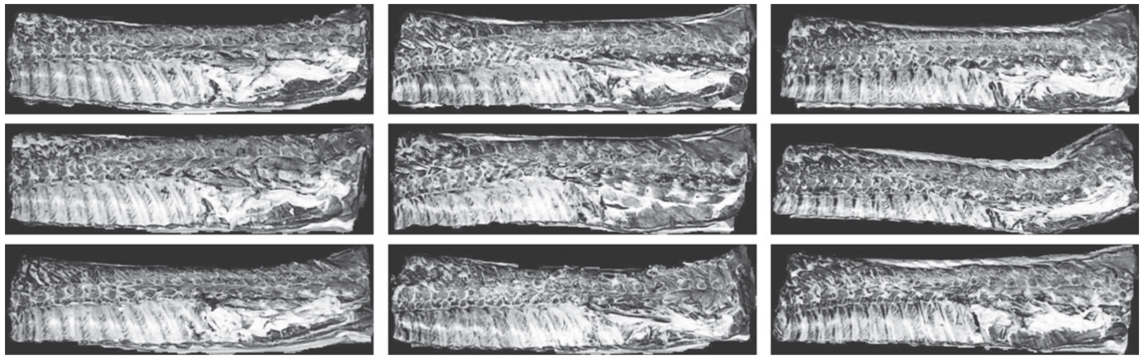

(a) Rough treatment. (b) Incorrect trimming. (c) Incorrect hanging.

Fig. 9. Examples of pork loins for which our recognition method yield image descriptions with little distinctiveness compared to the other image descriptions. Top row shows the pork loin on day 1. Middle row shows the same loin on day 2. Bottom row shows the closest candidate among the other loins.

constructed to be robust towards such perturbations and our experiments have confirmed this. While our results seem promising, we should note that we do not have sufficient image data to create proper *training*, *validation* and *test* sets. Therefore, our method parameters are likely to be overfitted because of improper training on the test set. However, as we consider this work a proof of concept, we still believe that the results show that our approach is feasible. In this connection we add that pork loins in slaughterhouses are typically processed in batches of significantly fewer pieces than in our experiment.

From the results in Fig. 9c, we have identified an important shortcoming of our canonization method. The image alignment is not suitable for non-rigid deformations because it leads to improper scaling and placement of the object. We consider this an important bottleneck of our current recognition method because bad alignment directly influences the image description.

It is possible to improve the recognition task even further by disallowing a loin from day 1 to be matched to multiple loins from day 2 and vice versa (by considering the problem an instance of *bipartite matching*). However, since the experiment images do not challenge our recognition method sufficiently, it will be difficult to draw conclusions from improvements to the method. It should be noted that by introducing bipartite matching, we lose the ability to perform any matching before all loins have been photographed for the second time. We have not investigated whether this will be a critical point in practice along a production line.

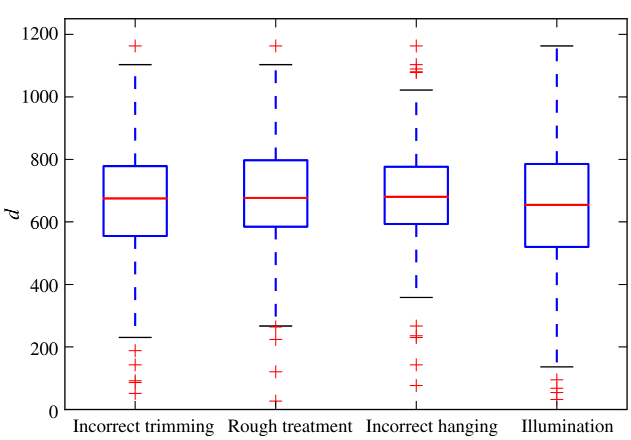

Fig. 10. A box plot describing the statistics of the match distinctiveness d for each experiment. Rectangles represent the *interquartile range* IQR = Q3 − Q1. The whiskers are placed at Q1 − 1.5 IQR and Q3 + 1.5 IQR. The plusses denote outliers.

5.1. Perspectives in a slaughterhouse environment

Based on our results, we believe that the proposed method is a competitive alternative to current technology relying on RFID tags of carrier devices. Vision-based tracking is less intrusive as it does not require physical contact with the tracked objects. Moreover, our relatively simple camera setup should be easy to integrate in a production line. As our experiments shows, our method does not enforce strict requirements to the camera stations wrt. lighting or light shielding. Though, one should still strive for a good diffuse illumination of the objects as it improves the matching distinctiveness.

Regarding the IT infrastructure needed to implement this system, we believe that the requirements of vision-based tracking are similar to what is currently required by RFID tracking. For both tracking methods we need IT systems for bookkeeping to keep track of which products have been seen where and when. A consequence of image-based identification is that the amount of identification data is many orders of magnitude bigger than with physical tags (the entire image description versus a single number per tag). With current computer networking speeds, however, we do not believe that this will impose any problems.

We imagine that the visual recognition should supplement the RFID tracking of carrier devices and ameliorate the tracking granularity from batches to individual meat cuts. Thus, from the RFID tag we can identify which batch is currently being processed and perform visual recognition within this batch. This is a subject for further investigation when our approach is to be tested on a real production line.

So far, we have only experimented with pork loins that exhibit a very characteristic image structure. It is likely that other meat cuts are more difficult to represent distinctively using our method. More experiments are needed to assess the robustness of our recognition method in more challenging situations.

Finally, as a more speculative perspective, we imagine that the image data gathered can be used for further analysis as a part of a quality assurance and process control stage. E.g. the fat percentage or the quality of the cutting process could be quantified by an image-analysis program using images from camera stations along the production line.

6. Conclusion

Tracking of individual meat cuts is an important part of facilitating meat traceability from farmer to consumer. In this work we have demonstrated a vision-based system that enables meat traceability in a slaughterhouse environment. By combining off-the-shelf vision and image processing technology we are able to track around 200 pig loins between two points without errors. This approach is meant as an alternative to current more intrusive tracking methods and our investigation shows that it is feasible. Further experiments are needed to determine the limitations of our method.

References

Arana, A., Soret, B., Lasa, I., & Alfonso, L. (2002). Meat traceability using DNA markers: Application to the beef industry. *Meat Science*, *61*, 367–373. http://dx.doi.org/10.1016/S0309-1740(01)00206-6.

Boykov, Y., & Kolmogorov, V. (2004). An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, *26*, 1124–1137. http://dx.doi.org/10.1109/TPAMI.2004.60.

Buhr, B. (2003). Traceability and information technology in the meat supply chain: Implications for firm organization and market structure. *Journal of Food Distribution Research*, *34*, 13–26.

Carriquiry, M., & Babcock, B.A. (2007). Reputations, market structure, and the choice of quality assurance systems in the food industry. *American Journal of Agricultural Economics*, *89*, 12–23. http://dx.doi.org/10.1111/j.1467-8276.2007.00959.x (arXiv:http://ajae.oxfordjournals.org/content/89/1/12.full.pdf+html).

Cimino, M., & Marcelloni, F. (2012). Enabling traceability in the wine supply chain. In G. Anastasi, E. Bellini, E. Nitto, C. Ghezzi, L. Tanca, & E. Zimeo (Eds.), *Methodologies and Technologies for Networked Enterprises. Lecture Notes in Computer Science*, *Vol. 7200*. (pp. 397–412). Springer Berlin Heidelberg. http://dx.doi.org/10.1007/978-3-642-31739-2_20.

Dahl, A., Aanæs, H., & Pedersen, K. (2011). Finding the best feature detector–descriptor combination. *3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT), 2011 International Conference on* (pp. 318–325). http://dx.doi.org/10.1109/3DIMPVT.2011.47.

Fröschle, H. K., Gonzales-Barron, U., McDonnell, K., & Ward, S. (2009). Investigation of the potential use of e-tracking and tracing of poultry using linear and 2d barcodes. *Computers and Electronics in Agriculture*, *66*, 126–132. http://dx.doi.org/10.1016/j.compag.2009.01.002.

Hviid, M., Jørgensen, M., & Dahl, A. (2011). Vision based meat tracking. *58th International Congress of Meat Science and Technology*.

Kaneva, B., Torralba, A., & Freeman, W. (2011). Evaluation of image features using a photorealistic virtual world. *Computer Vision (ICCV), 2011 IEEE International Conference on* (pp. 2282–2289). http://dx.doi.org/10.1109/ICCV.2011.6126508.

Karlsen, K. M., Dreyer, B., Olsen, P., & Elvevoll, E. O. (2013). Literature review: Does a common theoretical framework to implement food traceability exist? *Food Control*, *32*, 409–417. http://dx.doi.org/10.1016/j.foodcont.2012.12.011.

Lefebvre, L. A., Castro, L., & Lefebvre, É. (2011). *RFID: An emerging paradigm for the agrifood supply chain.* Wiley-Blackwell, 109–129 (chapter 7).

Mousavi, A., Sarhadi, M., Fawcett, S., Bowles, S., & York, M. (2005). Tracking and traceability solution using a novel material handling system. *Innovative Food Science & Emerging Technologies*, *6*, 91–105. http://dx.doi.org/10.1016/j.ifset.2004.10.006.

Pouliot, S., & Sumner, D. A. (2008). Traceability, liability, and incentives for food safety and quality. *American Journal of Agricultural Economics*, *90*, 15–27. http://dx.doi.org/10.1111/j.1467-8276.2007.01061.x (arXiv:http://ajae.oxfordjournals.org/content/90/1/15.full.pdf+html).

Prince, S. (2012). *Computer vision: Models learning and inference.* Cambridge University Press.

Regattieri, A., Gamberi, M., & Manzini, R. (2007). Traceability of food products: General framework and experimental evidence. *Journal of Food Engineering*, *81*, 347–356. http://dx.doi.org/10.1016/j.jfoodeng.2006.10.032.

Suekawa, Y., Aihara, H., Araki, M., Hosokawa, D., Mannen, H., & Sasazaki, S. (2010). Development of breed identification markers based on a bovine 50 k snp array. *Meat Science*, *85*, 285–288. http://dx.doi.org/10.1016/j.meatsci.2010.01.015.

Tate, M. (2001). Traceability of meat products—application of DNA technology. *Proceedings of the New Zealand Grassland Association* (pp. 255–258).

Tola, E., Lepetit, V., & Fua, P. (2010). Daisy: An efficient dense descriptor applied to wide-baseline stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *32*, 815–830.

Trienekens, J., & Zuurbier, P. (2008). Quality and safety standards in the food industry, developments and challenges. *International Journal of Production Economics*, *113*, 107–122. http://dx.doi.org/10.1016/j.ijpe.2007.02.050 (research and Applications in E-Commerce and Third-Party Logistics Management — Special Section on Meta-standards in Operations Management: Cross-disciplinary perspectives.).

Vanhonacker, F., Verbeke, W., Poucke, E. V., & Tuyttens, F. A. (2008). Do citizens and farmers interpret the concept of farm animal welfare differently? *Livestock Science*, *116*, 126–136. http://dx.doi.org/10.1016/j.livsci.2007.09.017.

Wang, X., & Li, D. (2006). Value added on food traceability: A supply chain management approach. *Service Operations and Logistics, and Informatics, 2006. SOLI'06. IEEE International Conference on* (pp. 493–498). http://dx.doi.org/10.1109/SOLI.2006.329074.

Weichert, F., Fiedler, D., Hegenberg, J., Mller, H., Prasse, C., Roidl, M., & Hompel, M. (2010). Marker-based tracking in support of rfid controlled material flow systems. *Logistics Research*, *2*, 13–21. http://dx.doi.org/10.1007/s12159-010-0025-6 (URL: http://dx.doi.org/10.1007/s12159-010-0025-6).