Tools and Resources for Cow Behaviour Monitoring

Green, J.[1], French, A.P.[2], Brignell, C.[3], Huxley, J.[4] and Pridmore, T.P.[1]
{[1]*School of Computer Science,* [2]*School of Biosciences,* [3]*School of Mathematical Sciences,* [4]*School of Veterinary Medicine and Science*}, *University of Nottingham, University Park, Nottingham, NG7 2RD, UK*
andrew.p.french@nottingham.ac.uk

Abstract

This paper describes a feasibility study which focuses on developing tools and resources for monitoring cow behaviour. Monitoring equipment was installed into a working dairy, and an algorithm developed to track and log the details of the appearance of regions on the backs of the observed cows. A replay tool was developed to play back this rich data set. The system was evaluated during a 24-hour monitoring session. The resulting data is well suited to the shape analysis and statistical modelling needed to recognize specific animals and detect motion events, which is the subject of current work.

1. Introduction

We have carried out a feasibility study into an automated visual surveillance system capable of identifying, tracking and recording the movements of multiple cows within a dairy centre (Figure 1a). Such data can potentially be automatically analysed and used as an early warning system in order to detect disease or the onset of lameness, by recognising deviations from learnt models of behaviour and/or movement; previous work has demonstrated this using individual animal behavioural models (e.g.[1]). Initial work has concentrated on developing a prototype system which has been used to monitor the movement of cows over an extended period (24 hours), with the aim of discovering the challenges and issues that development of a more robust system would need to overcome. This pilot system has involved developing and installing a method of capturing images from inside a working dairy centre, developing a method to extract the markings on cows' backs, and storing such information in a novel data structure. Additionally, a replay tool has been developed to visualize and navigate the resulting large and detailed data sets.

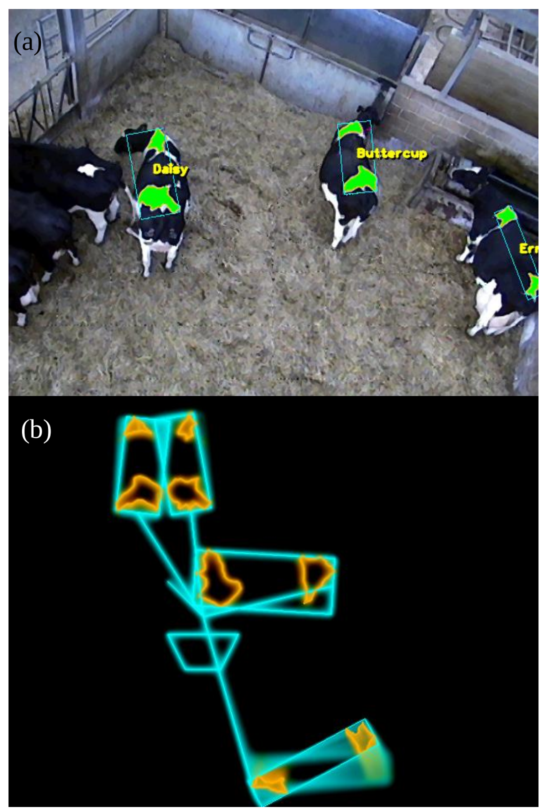

Figure 1. (a) Overlaid output on a video feed. (b) Replay of four cows being tracked

In this work we will present the approaches we have taken so far, the algorithms developed, some initial results, and a discussion of future work.

1.1. Motivation

The ability to closely monitor animals for welfare and production efficiency purposes is increasingly important in today's climate of high animal welfare standards and given the need to ensure food security whilst maintaining production efficiency. Monitoring animals is highly labour intensive for the stockman; an automated system to draw attention to particular animals and watch for particular behavioural events is clearly valuable for the industry.

The automated analysis of motion data holds much promise. Giving researchers access to the raw video feeds during the development of this system has already allowed previously impractical research work to be carried out. Throughout the pilot stage of this project, the setup has been actively used by vets and students to record footage of, and remotely observe, cows in different areas of the dairy centre. One study carried out using this equipment, on the effect of lameness on the feeding habits of cows, has already shown promising results (data demonstrates that lame cows spend significantly less time eating compared to the sound controls (196 mins/day vs 307 mins/day, P=0.03. [2])). The ability to automatically produce such results is a clear motivation for this work. Additionally, the unobtrusive nature of a remote, automated surveillance setup as compared to manual observation has been noted, and allows the observation of behaviours unaffected by human presence.

2. Method

2.1. Tracking

In this work, use is made of the black and white patches on the back of a cow. We track the different patches on a cow's back independently using a flood fill algorithm from a seed point, which fills pixels outwards by comparing candidate pixels with neighbouring pixels already in the connected component, and filling into them if they are sufficiently similar. The seed is initially manually set by clicking on the image inside the region to be tracked, and then propagated to future frames by locating the point furthest from the region boundary and using this as the seed point in the next frame (Figure 2). This does of course place a restriction upon the speed at which cows can move while being reliably tracked; however at a 25 frames per second capture rate, this has yet to be a problem. Under favourable conditions, fixed threshold flood filling has proven extremely effective at re-growing regions and hence updating their location and appearance in the tracker. The system is implemented using the OpenCV vision library [3] as a foundation.

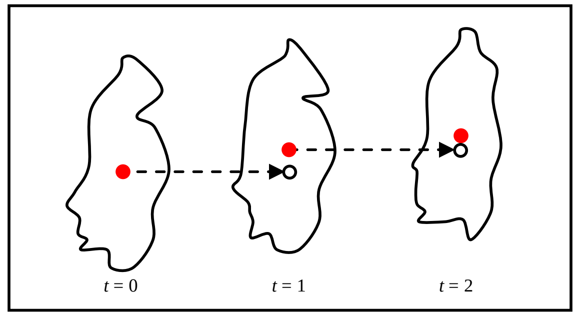

Figure 2. Illustration of flood-filling tracking algorithm over three time points for a moving and deforming region. Filled circle = point furthest from boundary, open circle = fill seed point.

The tracking algorithm contains some error checking to alert the user to severe tracking problems. These include detecting when a region becomes abnormally small, or when a region grows abnormally fast.

Future modelling work lies in statistical analysis of a region's description before each tracking update in order to determine if its newly found position is valid. If this is not the case, region growing parameters could then be adjusted in order to obtain a more likely match. Additionally, the shape of the boundaries could be constrained using shape analysis techniques. Development of such a technique is currently work in progress. The output of the tracker consists of a set of contours outlining patches on a cow's back, and the relationship between them (Figure 1b).

2.2. Target specification

Following Tsutsumi and Kita [4], we have developed a comprehensive data structure for defining information that can be used to locate and track a cow as it moves about the dairy centre. This information consists of two types: reference data and feature data (regions represented as connected data points), the latter being sub divided further into features found in three areas which correspond to a cow's neck, head and torso. It is further possible to define such features as being coloured either black or white in order to accommodate the tracking of cows which are more predominantly one colour. Any subset or all of these features together can be used for tracking; however for this pilot study, experiments have concentrated on the use of white torso regions to define the appearance and location of a predominantly black cow. The aforementioned reference data consists of a human readable name, unique identification number (for future association of position data with other database records e.g. milking statistics) and annotations, for example regarding known health issues of an individual cow.

To define individual cows for tracking in our prototype system, a frame is captured and a graphical interface allows the user to select the types of region to be defined (e.g. white torso regions). The user then clicks on areas of the displayed frame, once for each region to be recorded. Once all regions have been defined, reference data (name, etc.) can be entered and the information is saved as XML. This process is repeated for each subsequent cow to be tracked.

2.3. Data Logging

The system currently incorporates three user-selectable detail levels for logging, producing one (.csv) or two (.csv and .xml) files for each cow that is being tracked. The basic level simply logs location data, in effect a single location coordinate for each cow at each frame (calculated by fitting a bounding box around all tracked regions and taking the centre point), with an additional value representing orientation. A secondary level also includes data to enable the reproduction from logs of a crude representation of the cow's posture showing the approximate area taken up by its torso and head position. Level three logging allows the complete reproduction of tracker state and produces an .xml file detailing the exact positions of each tracked region for every frame as well as flagging if any update errors have been detected (Figure 1b). This log can be loaded into a separate replay application for analysis.

2.4. Replay and analysis

In addition to the tracker, a replay program has also been produced primarily for observational analysis of tracking results but also to allow automatic offline analysis of a cow's movement and behaviour over a period of time. The tool allows the user to load in and replay simultaneously one or several logs pertaining to particular cows and select all, or a subset of, these for display. The exact data shown can be adjusted allowing users to view centre coordinates, a wire frame representation or a selection of region types (e.g. white torso), as well as a combinations of these. It is also possible to toggle between a motion history and single frame display, with the former of these representing for any frame, *n* prior frames blurred and faded in accordance with age. This has the advantage of allowing observation of speed of movement in a single frame as well as giving the illusion of averaged position data (Figure 1b).

2.5. Modelling of data

The white regions defined by the flood fill algorithm form three dimensional shapes viewed in a two dimensional image. Statistical models for the shape of objects will in future work be applied to perform automatic cow identification even though the shape of the regions may change when viewed at different distances or orientations. This can be exploited to predict the location, shape and orientation of individual regions based on previous frames to validate and improve the flood fill algorithm. A statistical model will also be developed to perform automatic event detection tests. Given the location of the cow, the shape of individual regions and the relationship between regions we aim to be able to assess the probability of a specific event, such as feeding, taking place.

3. Results

Being a tools and resources development project, the main outcomes of this work can be divided into three main sections: infrastructure, algorithm development, and data.

3.1. Infrastructure

Two CCTV dome cameras ('Pan Tilt Zoom', SystemQ, PTZ505K), were permanently installed in the roofing of the centre to provide overhead coverage of the area of interest, and these were cabled to a remote computer in an adjoining room. Both cameras were fully movable by remote control software. The camera streams can be recorded for future use, or analyzed live.

3.2. Algorithm

The algorithm described in Section 2.1 was used to track through 24 hours of recorded image sequences. The tracker had to be restarted only 33 times due to tracking failures, including failures caused by serious unforeseeable occlusions which changed the layout of the area under observation. The longest consecutive sequence of successful tracking lasted 4.1 hours, with the average successful consecutive sequence being 44 minutes.

The major cause of tracking failure was due to fast illumination changes, or high dynamic illumination range within the scene. This was in large part due to sunlight streaming into the shed, causing very abrupt lighting changes which caused regions to be mis-tracked. Often these errors could be picked up by the algorithms self-checking code which ensured sensible region size changes, and the user could then manually restart the tracker. It should be noted that the system was able to track regions both during daylight and night time light levels.

3.3. Data

From the 24-hour tracked sequence, a large and rich data set was produced, comprising contour data of regions on the cow's back, from which cow orientation can be estimated, as can the cow position on the ground.

Figure 3 illustrate the kind of rich contour data stored by the system, in this case for one region on a cow as it moves and changes posture over time. Typically a cow will have at least two of these patches being tracked, and so the potential for learning to recognize interesting deformations using statistical shape analysis techniques is high; this is the subject of our future work.

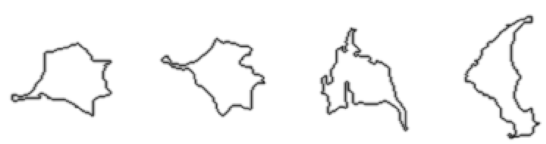

Figure 3. Shape of one region sampled at four points within a 5 minute period.

4. Conclusion

The algorithm was demonstrated to be suitable for tracking over long time sequences. Although manual restarts were required, running for an average of 44 minutes before these were necessary makes the tracking method much less labour intensive than any comparable manual approach. Additionally, the tracking method outputs a rich data set, which lends itself to shape analysis techniques; such approaches are actively being researched by the authors, with the hypothesis that using such shape-related information can be used to detect behavioural events and potentially aid tracking by providing a meaningful description of the boundaries of interest. Data from this experiment is currently being analyzed by mathematical modellers with a view to improving tracking, recognising individual animals and modelling behaviour. An approach of modelling 3D shapes from 2D images is being investigated

5. Acknowledgments

This project has been funded by an EPSRC Bridging the Gaps award to the University of Nottingham, which aims to bring together researchers from different schools for interdisciplinary research.

References

[1] Magee, D.R., Boyle, R. D. Spatio-Temporal Modeling in the Farmyard Domain. *Proceedings of the First International Workshop on Articulated Motion and Deformable Objects*. pp. 83-95, 2000.

[2] Voss, 2009, *University of Nottingham, BVMedSci Research Project*

[3] OpenCV computer vision library. http://opencv.willowgarage.com/wiki/. Accessed: 1-06-2010

[4] Tsutsumi, D., Kita, Y.. Motion Tracking of Cattle with a Constrained Deformable Model. *16th International Conference on Pattern Recognition (ICPR'02) , vol. 1, p.10372,* 2002.