Some Agricultural applications of Machine Vision

John Billingsley, *Senior Member, IEEE*

***Abstract*— Although these applications have been published before, the aim here is to provide tutorial information in a workshop environment. Examples include simple vision guidance, shape analysis by simple means and by s-psi encoding of the boundary, the use of DirectX filters for image acquisition and processing and standalone camera hardware including the use of embedded processing power.**

I. Introduction

THE National centre for Engineering in Agriculture was founded in Toowoomba some fifteen years ago. It was a three-way partnership between the University of Southern Queensland, the Department of Primary Industries and a firm of consultants. While one arm has been firmly based on agricultural concerns such as water efficiency, the Centre has provided fascinating opportunities for the application of mechatronic techniques to tasks such as automatic vehicle guidance, visual grading of produce, weed detection and eradication and identification and segregation of animal species.

Over the years, the enabling technology has changed considerably. Cameras are now commonplace and processor power is escalating. Nevertheless many of the underlying principles remain the same.

It is likely that as much can be learned from the Centre's mistakes as from its successes. The transition from 'interesting research output' to 'successful commercial product' is not an easy one.

II. Automatic guidance of tractors

The vision guidance project really commenced before the Centre was founded. A Masters student, Murray Schoenfisch, gave a presentation that investigated the use of buried cable for sensing. It was clear from his slides that rows of crop could readily be discerned, so the emphasis turned to machine vision.

The first step was to capture some video footage from a car driving somewhat erratically along the rows. This was captured to a PC using a primitive frame-grabber, only capable of capturing a binary image at relatively low resolution.

In a way this was fortunate, since the PCs in those days had very limited power to process any large quantity of image data. The image was captured via direct-memory-access (DMA) and access to it was kept sparse.

The capture board allowed the threshold to be varied by an output command. The simple expedient was used of entering a 'farmer guess' at the proportion of the image that should be green, then automatically varying the threshold level until the proportion of 'plant' pixels in the window of interest reached that level.

For row tracking, 'keyholes' were selected that should contain only one row. The 'plant' pixels were treated as blobs and a regression line was computed for them. This allowed the keyholes to track the rows from frame to frame, while the hardware could take its time to steer to bring the keyhole back to its proper position.

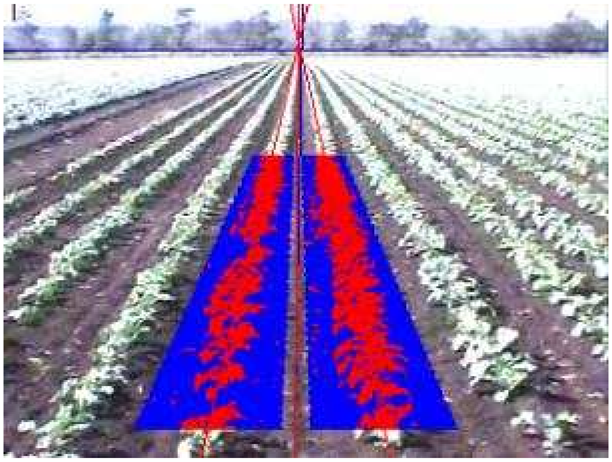

Fig. 1. A more recent example of 'keyholes' fitted to the rows.

Almost at once, a succession of Video Blaster cards came onto the market, each with a different organization of data in the machine memory. System level calls were needed to retrieve the image, now in colour, while the software ran under the DOS operating system. With commercialisation in view, parts of the software had to be rewritten each time the capture card dropped from the market and another was substituted.

Meanwhile, the tasks of implementing the signals to achieve steering and the need for an ergonomic interface to the driver occupied the ingenuity of Murray and of Jason

Manuscript received February 1, 2008.

John Billingsley is Professor of Mechatronic Engineering at the University of Southern Queensland, 4350 Australia, and directs technology research in the National Centre for Engineering in Agriculture. (Telephone +61 7 4631 2513, fax +61 7 4631 2526, e-mail: billings@usq.edu.au).

Stone. A module with an embedded processor took commands from the master and signals from a novel steering-angle sensor and operated a pair of hydraulic valves that had been cut into the steering circuitry.

Marketing was to be performed by CASE IH, Australia. Farmers who had sworn in trials that they could not live without the system now kept their money firmly in their pockets.

It is likely that the commercial error was in setting the price too low.

Agents were unwilling to enter a commitment to maintain a system that they did not understand, that might require servicing in the back of beyond, for a mere $5,000, an amount that doubled the price at which the units were supplied to them.

A further reason could have been the extravagant claims being made for the GPS systems entering the market. Such a system gave a guidance display to the driver at a cost nearer to $100,000 than the price of our system. And these systems certainly did sell. There was some consolation in the fact that our “steering ready kit” sold well at $20,000, consisting merely of the control module that had been such a small part of the system. But it did not take the GPS company very long to back-engineer and duplicate it.

Meanwhile the spin-off company founded by Murray Schoenfisch and Jason Stone remained afloat with sales of products such as a mechanical ‘feeler’ for guiding a cotton harvester along rows of stems and another to guide a tractor to follow a furrow.

III. Grading of Broccoli

This was another early project. Broccoli heads were manually cut in the field, dumped into bins and brought to a shed for grading. The clients wished to see a fully automated process of retrieving the heads from a saline bath into which they had been dumped, inspection of size, shape stem length and colour and sorting into bins for shipping.

Once again the vision system was propped up on pragmatism, minimising the accesses to vision memory. In a top view, a central vertical trail of inspection points, quite widely spaced, allowed the head to be found, whereupon a vertical chord was explored. From its centre, a horizontal chord measured the width. A second vertical chord through the centre of this then gave a more accurate measure of height. Now that the centre is known, an elliptical ring of points that should be outside the circular head can be tested, together with a ring of points at a somewhat smaller radius that should lie inside the head.

This gave a count of ‘lumps’ and ‘dents’, points that were green but should not be and points that should be green but were not. A scatter of points within the head were also tested for colour, to detect any yellowing.

The data items, size, width-to-height ratio, lumps, dents, yellow and stem length (measured by means of its reflection in a 45 degree mirror) served to classify the head into a grade and a destination, a genuine use of fuzzy logic.

The vision system did all that was expected of it, but the associated mechanical system gave problems. The heads could not be retrieved and loaded automatically. Even though the vision system would have given much better consistency of grading and perfect bookkeeping, the clients would not consider manual loading of the heads.

Many years later, I believe that the clients are at last considering grading within the harvesting vehicle, where the heads can have been loaded into the inspection holders by the human pickers.

IV. Counting Macadamia Nuts

When selecting which trees to propagate, the growers of macadamias wish to assess the yield of each individual tree. Since the method of harvesting is based on picking up nuts that have fallen to the ground, the task becomes one of localisation of each individual fallen nut.

Vision has been used to track the kernels as they are picked up by a 'bristle roller' which is coloured blue as a means of simplifying the discrimination problem. The simple algorithm used in the broccoli application can be used to locate the centres of circular kernels, while the 'inside and outside circle' method can check for circularity. To discriminate between leaf litter and a pair of touching kernels, however, calls for more sophisticated methods.

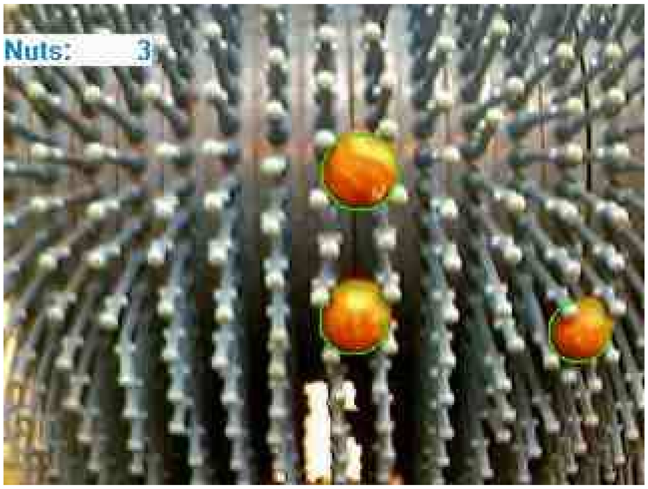

Fig. 2. Nuts have been located and circles marked around them

In the history of the centre's research, the Blaster card frame grabbers were now a thing of the past. Microsoft had perceived the huge market for the PC as an entertainment machine and much effort had been put into adding video capabilities.

'Video for Windows' tools appeared, but these were soon replaced with 'DirectX' and the concept of ‘graphs’ and ‘filters’, a technology seized eagerly by Mark Dunn to pursue many of the following projects.

Each software element of the vision system is termed a ‘filter’. One such filter can be a webcam, appearing as a ‘video capture source’. When this is inserted into the ‘graph’, a diagram of boxes and interconnections, it is seen to have one or more output ‘pins’. These define data-streams that can be connected to other filters.

For the webcam, an option is to ‘render’ the pin. Two boxes appear in the graph, one to transform the colour space and one to perform the actual display. Clicking a ‘run’ icon causes a window to appear with the moving image in it.

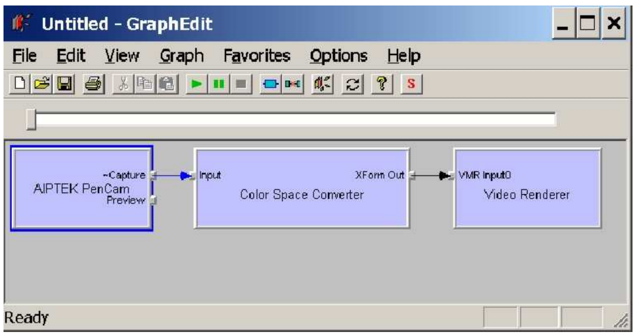

Fig. 3. A simple DirectX graph

Other filters can be inserted in the graph, software that has been written to perform a specific purpose such as identifying the coordinates of nut kernels.

When a new image frame has been received and processed by any preceding filters, the code of the filter starts to run. Pointers locate the new data and a frame for the processed data, an image that can be passed on for processing and perhaps display. Other filters allow data to be written to disk.

Of course, while it is running, the filter code can write data to a ‘blackboard’ shared with other modules or can directly execute routines for outputting data.

From the simple location of a kernel in an image, the macadamia software must incorporate the vehicle’s location with the aid of odometry, GPS where possible and radio transponders attached to tree trunks.

V. Identification of Animal Species

Australia has long been in drought and many of the bores in the great Artesian basin have been capped. With the number of waterholes limited it has become possible to consider control of feral species, the animals that are wild but not native. These include pigs and goats, though in some regions wild horses, camels and water buffalo are also a nuisance.

A vision system captures images from a laneway on the approach to the waterhole and their species is recognised. A gate moves to divert them, either into an enclosure that they can leave through a one-way gate or into a closed pen from which the feral animals can be ‘harvested’. There is a lucrative overseas market for wild pork.

First the outline of the animal is traced. When viewed against a blue background this is simplified, but work has also been performed on detecting the change in a general background to isolate the silhouette of the animal.

Fig. 3. The outline of a goat.

The outline trace is in the form of an s-psi plot, in which the angle of the tangent is plotted against the distance moved around the circumference. There are many properties of the s-psi plot that make it attractive for matching the outline against a template.

It is possible to normalise the circumference to a value corresponding to unity, so that there are, say, 256 values of the angle to represent the closed shape. Now the matching of shape is independent of size.

If the outline is tilted, a constant will be added to each value of the plot, corresponding to the angle of tilt.

After one complete circuit, the angle will have increased or decreased by 2 pi.

If the plot commences at a different starting point, the values will be shifted cyclically.

When the object is an animal, there are attendant problems, but solutions are at hand. Much of the circumference comprises the legs, some of which can be hidden behind others. The data will tend to be dominated by the way the animal is standing. The solution is to encode only the top of the animal, as defined by the path between the leftmost and rightmost extremity. This also solves the ‘starting point’ problem.

It might be necessary to ‘flip’ the data horizontally, or to hold two templates for the animal facing left and right. Usually, however, the orientation can be determined from the direction of travel.

Several prototypes were built, the software being embodied in a laptop computer. To be left unattended, this required a special housing, partly to protect against the climate and partly to avoid theft. For reasons including battery endurance and weatherproofing, it was seen to be desirable to develop a standalone capture, analysis and recording platform.

Mark’s ROC, or Rugged Outdoor Camera, has as its principal components a camera module, a powerful embedded processor and a memory card, of the sort used in cameras. With a modest battery, this can give continuous operation of several days.

With such a device, there is renewed virtue in the use of compact software and many of the earlier algorithms once again become relevant. Soon the experimental prototypes will have been replaced with this dedicated hardware.

VI. Irrigation of Cotton

Water is of great concern to Queensland farmers, particularly during a lengthy period of drought. Crops such as cotton receive irrigation and there is concern to optimise the amount of water that is applied. While instrumentation can assess soil moisture, the most direct measurement of effectiveness must be made on the plants themselves, as in a project conducted by Cheryl McCarthy.

The growth rate of the cotton plant is assessed in terms of the inter-node stem length. A vision system captures numerous snapshots of the plants from which the stems are detected and the distance measured between nodes, the thickening in the stem where branching or leaves occur.

Assessment of growth rate has to be based on a statistical measure. It is of course possible to return to the same location for the next pass, but at this stage it would be difficult to ensure that we had returned to the same individual stems of a particular plant.

Fig. 4. Vehicle for scanning internode length

Fig. 5. Stem image and analysis

VII. Assessment of Fodder Quality

When animal feedstuffs are grown for sale abroad, it is important to agree on some standard on which the price can be fixed. At present judgment is excessively subjective. It relies on the extraction by hand of ten stems of which the thickness is measured and an average is taken. An assessment of colour is made by eye.

What was wanted was an objective measurement that would show consistency between a machine used by the vendor and one used by the recipient. Once more, machine vision was selected as the method.

A sample handful is placed in a box with controlled lighting, thereby limiting lighting variation. From the captured image, the width is measured not just of ten stems but of all that can be discerned. Instead of taking an average, a histogram is drawn. Boundary analysis allows an analysis to be made of stem-leaf ratio.

As well as the more obvious means of analysis, microspectral methods were considered. Detailed spectra were taken, extending as far into the infra-red as a monochrome camera could be expected to detect. Any spectral differences could lead to a means of discrimination, even if the colours appeared the same to the eye.

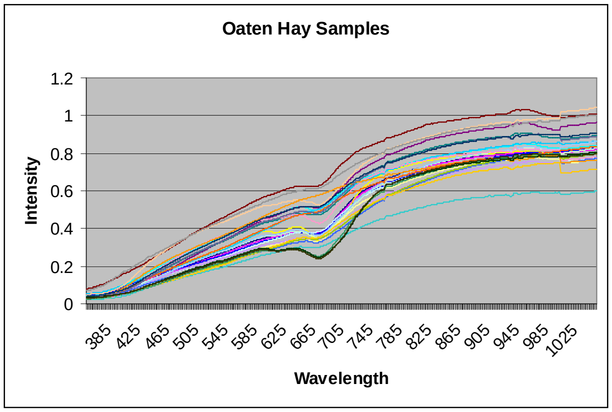

Fig. 6. Comparison of spectra.

This aspect can be pursued in a number of ways. With suitable illumination a diffraction grating can be used to give spectra at an array of points on a line across the image. Alternatively full two dimensional images can be captured at various wavelengths, either by selective illumination or by filtering the returning light.

VIII. Weed Detection

When spraying to kill weeds, two strategies are possible. A toxin can be sprayed overall that will selectively kill the weeds while leaving the crop unharmed. (Some crops have been genetically modified to withstand the 'Roundup' weedkiller.) Alternatively, if the weeds can be identified and located, a small directed burst of a general purpose herbicide can be used to kill them. This second approach obviously has great value in reducing the amount of chemical used.

Sugarcane plantations can be beset by weeds such as 'green panic', 'guinea grass' and wild sorghum. Discrimination can depend on a variety of strategies. In some cases there is a colour difference that can be perceived from the three channels of a conventional colour camera. In others, leaf shape and the orientation and clumping of the growing stems are the best indicators.

Steven Rees is pursuing this project with funding from the Sugar Research and Development corporation.

Many of the fundamental directions taken in the fodder project can thus be adapted for use in the field. If the scan is to be conducted with controlled lighting conditions in the daytime, a mobile canopy is needed to exclude the bright Queensland sun.

Fig. 7. Field testing the weed detector and canopy

The prototype equipment sprays brightly coloured foam so that its effectiveness can be analysed.

IX. Guidance Revisited

When farmers had realised the limitations of GPS guidance, it seemed time to revive the concept of guiding the tractor by means of the crop itself. By now, a number of technologies had matured enough to be cost effective. Inertial sensors such as rate gyroscopes and accelerometers were available in an 'on-chip' form, so that an inertial package could smooth over any breaks in vision or GPS guidance. At the end of a row, a 'headland turn' was necessary that could not rely on vision.

All the vision processing power of DirectX was now available and computer speeds were many times greater. Nevertheless the early 'keyhole' approach produced excellent results when a crop could be clearly seen.

However the commercial director of the project was keen that little should be left to the setting-up abilities of the driver, and that the system should work with green crop on dark soil, dark wheel-tracks on grey soil, rows of pale stubble or anything with a discernible array of lines.

The most promising technique involved capturing an image to use as a template, then correlating an array of horizontal stripes against the incoming vision stream. The correlation will show a strong peak at a value corresponding to the displacement of the 'slice', from which position and heading errors can be derived. If the tractor pitches, the spatial frequency in the slice will change, causing a decrease in the peak value of the correlation, but the tracking will be undisturbed. In mid field, this event is unlikely.

In the overall project, steering is a relatively minor part of the software burden. High in importance is the driver's interface, with choice of suitable icons and menu trees, not to mention the actual size and resolution of the display. With networks so fashionable, it is tempting to distribute the processing power rather than concentrate it in one location. However a more appealing approach is to use the ROC not merely to capture the image but also to extract all the steering commands.

X. Conclusions

In the consumer field, the concept of computer as entertainer has spun off a dazzling array of products and services, including the embedding of cameras in cellular telephones, internet download of media such as iTunes with attendant MP3 players capable of displaying video, even digital 'picture frames' to display the vast amount of data in shots captured by digital cameras.

In this flood of products and technologies are many that can be exploited for solving 'real' engineering problems, such as those that have so long confronted the farmer. As these new solutions are seen to emerge, new problems are brought to light.

There have been very many other projects for which there is no time or space here. These include that assessment of the roughness of oranges, detection of foreign objects among new-laid eggs, assessment of the growth of pigs and cattle, observation of the behaviour of small marsupials and even the use of machine vision to measure the density of dingo teeth.

On the one hand, the rate of progress is most exciting. On the other, today's groundbreaking innovation will tomorrow be seen as commonplace.

Acknowledgments

The research described here has been supported by grants from many organisations.

CASE IH supplied funding for the commercialisation of the early guidance work, while the later programme has been performed in collaboration with AgGuide with support from Aus Industry.

The animal identification work received funding from the National Heritage Trust, collaborators are the University of Queensland and RPM Rural Products.

The macadamia project is funded by Horticulture Australia Ltd., collaborators are Hidden Valley Plantations and NSW DPI.

The work on spot spraying is supported by the Sugar Research Development Corporation and the collaborators are Bundaberg Sugar Ltd.

Work on fodder quality assessment is supported by the Rural Industry Research and Development Corporation.

The work on cotton growth rate is supported by the Cotton Research and Development Corporation.

References

[1] Information sheets on many of the projects mentioned can be found at http://www.ncea.org.au/Page_SheetsList.html (visited February 1, 2008.

[2] Dunn, M., Billingsley, J. and Bell, D. (2006). "Vision based macadamia yield assessment" *Sensor Review*, 26(4), 312-317.

[3] McCarthy, C L, Hancock, N H and Raine, S R (2006) "On-the-go machine vision sensing of cotton plant geometric parameters – first results," *Proc Conf. Mechatronics and Machine Vision in Practice (M2VIP) 2006*, Toowoomba, Australia, December 2006.

[4] J Billingsley, "The Counting of Macadamia Nuts", Mechatronics and Machine Vision 2002: Current Practice, Research Studies Press Ltd. ISBN 0-86380-278-8, pp 221-7.

[5] J Billingsley and M Schoenfisch, "Vision and Mechatronics Applications at the NCEA", Fourth IARP workshop on Robotics in Agriculture and the Food Industry, Toulouse, October 1995.

[6] M Dunn, J Billingsley and N Finch, "Machine Vision Classification of Animals", Mechatronics and Machine Vision 2003:Future trends, pub Research Studies Press Ltd, Baldock, UK, ISBN 0 86380 290 7, pp 157-163.