

USQ4C 1995

Design & Commissioning of a vision based
guidance system.

A Vision Guided Agricultural Tractor

J Billingsley, M Schoenfisch
University of Southern Queensland

Abstract

A vision guidance system has been designed, built and commissioned which steers a tractor relative to the rows of a crop such as cotton. It was required to be insensitive to additional visual "noise" from weeds, while tolerating the fading out of one or more rows in a barren patch of the field. The system integrates data from several crop rows, testing for image quality. At the same time, the data processing requirements have been limited by the use of frame-sequential strategies to reduce the image space which must be processed. The current prototype employs an embedded 386 PC notebook computer and shows great promise of cost effective commercial exploitation.

Experimental results are reported and further sensing systems are explored to enhance performance in difficult environments.

Introduction

There is a need for automated guidance of agricultural vehicles, not to remove the presence of a driver but to allow greater attention to be given by the driver to the cultivation operation. Automatic steering also promises to improve the effectiveness of "controlled traffic", a technique where by repeatedly using the same "footprint", vehicles minimise compaction damage to the soil. Under manual control, this increases the pressure on the driver to maintain precise control of the track of the vehicle. The experimental vehicles are already capable of much more accurate sustained control.

For spraying operations, high speeds are desirable to enable a ground vehicle to challenge the role of a crop-spraying aircraft. Once again, the driver's task is made more demanding and an "autopilot" becomes highly desirable.

Many guidance methods can be considered, ranging from buried leader cables to beacons, surveying instruments or even satellite navigation. All have their drawbacks. The most appealing method is to follow human practice and take guidance from the crop itself, steering the vehicle by means of the view of the rows ahead.

There are, however, many complications as the condition of the crop changes through the growing cycle. Initially the plants appear as rows of small dots among scattered random dots, the weeds. Later they fuse to form a clear solid line. Before long, however, the lines have thickened and threaten to block the laneways. Many variations of the vision algorithm are thus required to fulfil all the seasonal requirements.

Vision systems can acquire data at a very large rate. A full-colour high-resolution image can require 1.4 megabytes of memory to hold it, and twenty-five such images are received from a conventional camera each second. Many vision projects have become congested by such data rates, requiring massive computing power to extract the simplest of features. The policy in this project has been to limit the acquired visual data to a modest level and to use frame-sequential analysis methods to select only that fraction of the data for analysis

that will yield the necessary steering data.

Within a single 386 computer, software was implemented for both image analysis and on-line control of the tractor. Initially a stepper motor was used as the steering actuator, but the very restricted steering slew which this could achieve meant that limit-cycle instability was prone to occur at all but the lowest velocities. The use of a variable-structure control algorithm gave substantial improvement, but the system was still sensitive to the accuracy of initial calibration.

The stepper was replaced with a geared DC motor as used in a cordless drill. The computer interface employed two relays for the final output, so that the bang-bang operation would emulate the control of hydraulic steering valves in a future implementation. The performance was dramatically improved and the steering mechanism was no longer seen as a limitation.

Although the machine performed well, its appearance was most inelegant. A 250 volt inverter and a computer tower-case were taped to the roof of a tractor, while the monitor screen was held on the bonnet in front of the driver by even more adhesive tape. A camcorder was similarly taped to the front of the tractor.

Funding of A\$150,000 was at that stage granted by the Cotton Research and Development Council and J I Case donated the long-term use of a Maxxum 100 horsepower tractor. The research which had up to then proceeded with minimal resources was put onto a sound footing.

The cordless drill motor was replaced by a specially designed hydraulic valve system for direct actuation of the steering. A notebook computer with expansion rack was substituted for the tower case. The vision sensor is still a camcorder, now neatly fixed to the bonnet with velcro, but the overall appearance is much more impressive. The vehicle is already capable of travelling through a crop at over 25 kilometres per hour with only a centimetre or two of waver.

Image acquisition

The first experiments were based on a binary "frame-grabber" which yielded a black-and-white image - no grey levels - with a resolution of 768 horizontal points by 96 rows vertically. The board allowed multiple images to be interlaced to build up a resolution of 576 rows, but that feature was ignored. Similarly the discrimination level could be varied. By sweeping this level between frames, an image with multiple grey levels could be constructed. Instead the guidance software drove the discrimination threshold to achieve a chosen "spot density" in the area of greatest interest.

The actual transfer was performed by DMA - direct memory access - and allowed processing of the data to go on undisturbed. However each transfer had to pause until the incoming television image frame started so that some amount of waiting for synchronisation was inevitable. This time-loss was reduced by the use of a double-buffering technique. As soon as a frame of data was seen to have arrived, the "grabber" was primed to load the next image into the second buffer. Processing of the first image then started, and when all necessary actions had been taken the program checked for the filling of the second buffer and waited if necessary.

The image was captured in an array in main memory, where the software was able to access it for processing. At some cost in overall speed, part of the image was intermittently copied to the display memory so that it could be seen on the computer screen and the effectiveness of the algorithm could be assessed.

The second system used a camera interface targeted at the consumer market, the "Video Blaster". It is available at a relatively low price and has some very impressive features. A full colour image is captured in the on-board memory, and can be merged "live" as a window forming part of the VGA display. The image can be scaled horizontally and vertically with no use of the processor time of the host computer. Lines and other graphics can be superimposed on the screen image, so that the performance of the analysis system becomes very clear to see.

The system does have attendant disadvantages. First is the massive amount of data held in each frame. The image memory is mapped at a high address in extended memory, chosen to be at seven megabytes, and occupies three-quarters of a megabyte of addressing space. By scaling the image to be about one third of the height and width of the screen, the useful data is reduced to some eighty kilobytes, still a lot to process between scans.

The second disadvantage is the difficulty of accessing this high memory except in "protected mode". The technique used at present is to call the routine at DOS interrupt 15 to transfer lines of data into a local buffer, but this is inefficient in processor time.

A great advantage of the new interface is the provision of colour. A field with a newly shooting crop may be littered with light-coloured detritus which makes discerning the crop rows difficult if brightness alone is used. Even the use of a green filter over the lens makes little improvement. It is possible to use the chrominance signal rather than luminance to capture an image based on the "greenness" of each point. The spatial resolution of chrominance is nowhere near as sharp as that of luminance, but resolution is not of the greatest importance.

Excellent performance has been achieved, although some processing speed is lost in decoding the colour information. Commonality between the various hardware versions has been achieved by the use of a function, *picbit(x,y)* which presents the image in a standard form to the analyser whether acquired from the binary grabber, the luminance signal or the chrominance signal of the Video Blaster or from any future system.

Vision analysis

The task is to identify a row of crops and locate its displacement from some datum position. There will certainly not be a well-defined object with shape which could be analysed by outline methods, even if time permitted. In the early stages of growth, the crop takes the form of a spotty row of variously-sized blobs. At its best, it is a linearly-connected domain with a highly irregular outline. If a window can be established within which members of only a single crop row will be present, however, then a relatively straightforward averaging technique can be used.

The analysis method makes heavy use of information learned from previous frames. With knowledge of the location of a row, a window can be set for the next frame where movement of the vehicle should not have carried it as far as an adjoining row. If all goes well, the new frame will yield a new window for searching the following frame and so on.

Now the task becomes one of making the best estimate of a line through a row of blobs within the frame - and for this sort of problem the technique of regression analysis can be used.

Regression is conventionally used to fit the best straight line to a collection of points, usually measurement samples or readings from which statistics are to be drawn. The regression line minimises a quadratic cost function, the sum of the weights of the points times the squares of their distances from the line. This cost function can be thought of as analogous to the moment of inertia of the data points, represented as masses, when spun about the regression line.

In this case, the "points" are the image pixels (picture elements) and their weights are related to the effective brightness measured at each location within the window. (In the original binary frame grabber, image samples were considered in groups of eight and the effective brightness was counted as the number of pixels within the group which were white.) The cost function becomes:

$$C = \sum_{x=mx_0}^{nx_0} \sum_{y=my_0}^{ny_0} m(x,y) (a - b \cdot y)^2$$

The best-fit line is defined by the parameters a and b which minimise C ; a and b can be solved from the equations given by

$$\frac{\partial C}{\partial a} = 0 \quad \text{and} \quad \frac{\partial C}{\partial b} = 0$$

A routine within the program adds the weights $m(x,y)$ given by the function $picbit(x,y)$ to form a variable m , adds their weights times distance from the window centre-line to give variable mx , adds the weights times the vertical displacement to give my - and so on. Appendix A gives more details of the software.

The result of the computation is a variable $xfit$ which indicates the displacement of the "best line" from the centre of the window and another, $sfit$, which describes the slope of the regression line. The window is moved to a new position, provided certain conditions (described below) are satisfied.

In addition to moving the window from side to side to track the row, the window is distorted by shearing it horizontally so that its centre line lies parallel to the row. Since the function is evaluated over a lozenge defined by the previous best estimates a' and b' , the variable x in the equation for C should really be replaced by $(x - a' - b' \cdot y)$. The fit is thus computed over the sheared window and the result is used to update the window shear for the next frame.

Now the moment of inertia about the regression line can be computed - most of the calculations have already been performed. If the fit is good, the result should be small. If the crop is scattered, however, the moment of inertia will be larger. As a test, this moment of inertia is compared against myy , the moment of inertia about a horizontal axis. The ratio gives a measure of the quality of fit and the information is only acted upon if the quality is sufficiently high.

Often a row may fade out half way down the field. For this reason, the computation is performed not just for one row but for two or for three.

A second analysis method is also interspersed with the first. Samples are totalled over a set of lines parallel to the sheared window axis, forming "slices" of the window. The largest slice-total will indicate the row position. Its ratio with respect to the row average will indicate the quality. No slope information is given, however, so the regression method must be interspersed frequently despite involving lengthier computation.

Finally, the mean value of all the samples in the window is used to adjust the brightness setting for the next frame. In the original binary grabber, the threshold was moved to seek a proportion of white cells of, say, ten percent of total. In the later system, the pixel brightness or the chrominance value is rescaled to give a figure from zero to fifteen for the function $pcbtt(x,y)$. The datum brightness of this calculation is again moved to standardise the average.

Control implementation

Within the program, two variables are maintained which represent the state of the vehicle. These are the lateral movement of the vanishing point and the change in the aggregate slope of the rows. Alternatively, the lateral displacement of the rows in the centre of the view may be regarded as defining the second variable. These variables must be used to steer the vehicle.

The primary error signal is the mid-view lateral displacement, the apparent lateral shift of the rows half-way down the picture. This will have a value which is a sum of vehicle displacement and heading terms, since it is measured ahead of the vehicle.

In a simple strategy, the demanded steering angle could be made to be proportional to this error, causing the tractor to turn to bring the row laterally to the centre of view. If this can be achieved, the resulting track will follow an exponential decay towards the row. In lieu of a time-constant, there will be a "distance constant" set by the distance of the centre of view ahead of the vehicle.

When the change of steering angle is rate-limited, however, an abrupt onset of limit-cycle instability will occur for larger errors. This effect can be countered by modifying the strategy to limit the commanded steering angle and also limit the heading angle at which the vehicle is required to approach the row.

In the algorithm, the lateral displacement of the rows is multiplied by an appropriate gain to yield a desired heading for correction of the error. A magnitude constraint is applied to this demanded heading change. The heading is expressed in terms of the lateral displacement of the vanishing point.

In the original prototype the error between demanded and perceived heading was converted to a demanded steering angle, again with a magnitude constraint imposed. The task of commanding the stepper motor to follow the demand was left to an interrupt routine which dictated the stepping rate of 32 steps per second.

The stepper motor was replaced first with a geared DC motor for which the computer output used two relays, then by a pair of hydraulic valves driving a steering ram. A potentiometer and analogue-to-digital convertor provided the computer with information of the steering angle and a timed interrupt examined the discrepancy between actual steering angle and computed demand. The actuator was pulsed with a mark-space drive to close the

steering loop. This performed exactly as designed.

Continuing work

With a series of very successful demonstrations the project is by no means at an end. The final objective is a system which can be used literally "in the field" on a commercial basis.

There are many decisions to be made concerning additional sensors for operating in other modes. When the field is first to be marked out, the task of ruling straight furrows perhaps two kilometres long is a taxing one. A flux-gate compass unit has been interfaced to the system to address this problem.

At planting time, the furrows exist but no crop is present. It is possible that under suitable daylight conditions or at night with suitable headlights the furrows could be made to stand out with sufficient contrast. The attractive alternative is to use tactile sensing of the furrows with an electromechanical transducer.

When the canopy has closed in, no gaps can be perceived between the rows. The addition of tactile stalk-sensors can be used for more accurate guidance of harvesters and post-harvest stalk-pullers.

These and many others are aspects under investigation for the generation of a wide-capability system.

Conclusions

A program of research combining theory and experimentation has resulted in the verification of a practical guidance system, despite early limitations of very meagre resources. Now that adequate funding has been allocated, the system can be prototyped to professional standards and its performance enhanced to achieve industry acceptance.

References

Smith A.L, Schafer R.L, Bailey A.C (1987); Verification of Tractor Guidance Algorithms. Transactions of the ASAE, Vol. 30(2):305-310.

Smith A.L, Schafer R.L, Young R.E (1985); Control Algorithms for Tractor Implement Guidance. Transactions of the ASAE, Vol. 29(2):415-419.


Campbell N.W, Thomas B.T (1993); Navigation of an Autonomous Road Vehicle Using Lane Boundary Markings. Paper Preprint 1st IFAC International Workshop on Intelligent Autonomous Vehicles, Pgs. 169-180.

A Vision-guided Agricultural Tractor, J Billingsley, M Schoenfisch, Australian Robot Association Conference, Robots for Competitive Industries, Brisbane, July 14-16 1993.



Newly sprouted plants appear in relatively neat rows

1.




Steering information can be derived from lines fitted to the row images

2.




Part of the image is selected in a window

3.




A regression line is fitted to the points

4.




The moment about the regression line gives a measure to guard against errors

5.




A group of three windows is updated by corrections to all valid regression lines

6.



Movement of the vanishing point indicates a change in heading

7.



Movement of the pattern centre indicates lateral displacement of the vehicle

8.

Fig. 1. Slides illustrating the image analysis algorithm.