WORKING NOTE AC112

Alan Creak
11 July 1997

A VELOCITY SENSOR FOR SMALL MACHINES

This idea is merely a translation into digital of a proposal[1] for an analogue sensor embodying the same principle, so it isn't new. The method depends on measuring the rate of motion of an image over a photosensitive surface, and requires no feature recognition. Using a two-dimensional sensor array – or perhaps a light-sensitive RAM chip, which is addressable – gives the information required very quickly, and also allows the velocity calculation to be performed in two perpendicular directions more or less simultaneously, and rather cheaply.

THE PRINCIPLE.

Blake, Hamid, and Tarassenko[1] propose, and demonstrate, a simple method for measuring velocity by – in effect – watching an image of the passing scene as it is seen through a grating. The total illumination perceived through the grating is measured, and the velocity is related to the frequency of any oscillations observed.

If the view is quite featureless, the input will be constant and no oscillation will be observed, and the device will not work. In practice, this is unlikely to happen with real views, so the sensor will work most of the time. Its one defect is that with the obvious construction it is impossible to determine the direction of motion.

The clever bit is to use a virtual grating, simulated by electronics. By moving this simulated grating at a steady speed, greater than the maximum plausible speed of an image across the grating, the detected signal for zero velocity is not a zero frequency, but whatever corresponds to the rate of movement of the simulated grating. With a virtual grating, it is also possible to use a differential measure, which is the difference between the intensity of illumination "passing through" the grating and that "stopped" by the grating; by this means, a signal oscillating around zero is guaranteed.

The authors describe the phenomenon by means of an extraordinarily complicated mathematical treatment which I don't understand, so I begin with a simpler treatment which is doubtless less rigorous, but which I do understand.

MATHEMATICS.

I analyse (not very formally) a very simple example. To begin with an illustration of the principle, suppose that the grating is composed of alternating transparent and opaque strips of equal width, and that readings are taken at intervals during which the image moves exactly one strip across the grating.

Let the intensities in subsequent strips of the image be t_i. Then, if the grating comprises n opaque strips and n transparent strips (peculiar, perhaps, but easy enough if you're simulating it) we can write down the total intensities transmitted and blocked by the grating when the lowest numbered transmitted strip is strip k as :

$$T_{t,k} = \Sigma_{i=0}^{n-1} t_{k+2i}$$

and

$$T_{b,k} = \Sigma_{i=0}^{n-1} t_{k+2i+1}$$

At the next observation, the image has moved one strip further along, so the new values of the intensities become :

$$T_{t,k+1} = \Sigma_{i=0}^{n-1} t_{k+2i+1} = T_{b,k}$$

and

$$T_{b,k+1} = \Sigma_{i=1}^{n} t_{k+2i} = T_{t,k} - t_k + t_{k+2n}$$

The blocked view is now just the previous transmitted view, while the transmitted view is the previous blocked view, without the first strip but with an additional new strip at the leading end.

One step further still :

$$T_{t,k+2} = \Sigma_{i=1}^{n}\, t_{k+2i} = T_{b,k+1} = T_{t,k} - t_k + t_{k+2n}$$

and

$$T_{b,k+2} = \Sigma_{i=1}^{n}\, t_{k+2i+1} = T_{t,k+1} - t_{k+1} + t_{k+2n+1} = T_{b,k} - t_{k+1} + t_{k+2n+1}$$

Now the whole scene has moved two strip widths beyond its original position, so both blocked and transmitted views are almost the same as the corresponding originals, but in each case one strip has been lost from the trailing edge and one new strip is included at the leading edge.

Provided that n is not too small, $T_{b,k+2}$ is likely to be rather similar to $T_{b,k}$ in most cases, as each is the sum of n values of t_i, and they have n-1 values in common. Further, at each step the image strips which were obscured become visible, and vice versa, so both $T_{t,k}$ and $T_{b,k}$ are likely to oscillate in value.

Again because each is the sum of n values of t_i, the oscillations are likely to be small as compared with the magnitude of the quantities, but the differential measure brings out the differences effectively. Consider the value at each step of the difference between the two quantities :

$$T_{d,k} = T_{t,k} - T_{b,k}.$$

$$T_{d,k+1} = T_{b,k} - (T_{t,k} - t_k + t_{k+2n}) = - T_{d,k} + t_k - t_{k+2n}$$

$$T_{d,k+2} = (T_{t,k} - t_k + t_{k+2n}) - (T_{b,k} - t_{k+1} + t_{k+2n+1}) = T_{d,k} - t_k + t_{k+1} + t_{k+2n} - t_{k+2n+1}$$

The major term (again making plausible assumptions about n and the comparative uniformity of the t_i values) is $T_{d,k}$, which appears with alternating sign. Unless n is small, or there are some very unusual t_i values, this will lead to slowly modified oscillatory behaviour. We will expect breaks in the oscillatory nature of the differential signal at points where the sample received by the grating changes suddenly in intensity of illumination.

Clearly enough, this will work best if the width of the strips is comparable with the dimensions of image features, but except with a very bland image one might expect to end up with something plausibly like an oscillatory signal. Anyway, it works in the paper.

PICTURES.

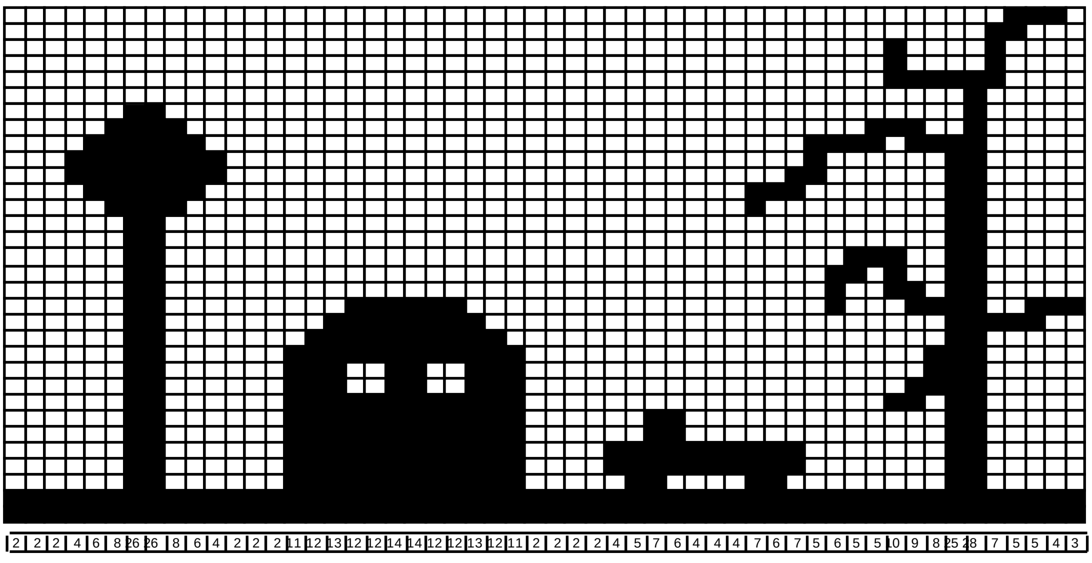

The numbers measure the dark intensity in the corresponding vertical strips. They are used in calculating the graphs below.

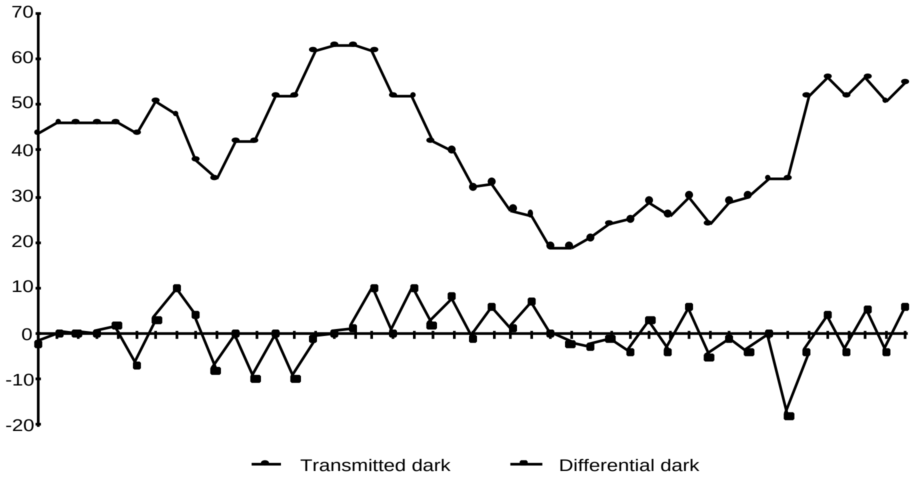

The upper line shows the intensity of darkness passed through a grid of five slots. The lower line is the differential plot of the darkness. Though there are occasional faults, the essential periodicity of the differential darkness is very clear. The faults are, as predicted, at points where the sample received by the grating enters or leaves areas of very different illumination – the house, and the right-hand tree trunk. The horizontal position of a point on the graph corresponds fairly closely to the position on the picture.

THE PROPOSAL.

The proposal is simply to simulate the method described in the paper using a digital computer and some sort of addressable photosensitive array. It works; they show some examples as preliminaries to their "real" measurements, but don't consider it satisfactory for use in a real sensor.

Their objections to such systems are of several sorts. They comment particularly on the expense of CCD camera, frame stores, etc., and also on the limited sampling speed (typically around 50Hz) imposed by the CCD architecture.

They don't point out that, given the two-dimensional nature of the picture, there's enough information to compute two perpendicular components of velocity, which yields the true vector velocity. Neither do they discuss the possibility of using other sensors with which arbitrary image points can be directly addressed, thereby avoiding the sampling speed problem.

Experiment 1 :

The first experiment is to use a conventional CCD camera with framestore to experiment with the determination of two-dimensional velocities. Problems might be expected from diminished correlations between successive readings in directions almost perpendicular to the true velocity, as the image will "slide across" the sensor, and the successive image strips in the equation above will not be the same at different time steps. Calculations of the velocity component in many directions should therefore be carried out; once a direction of motion has been established, calculations in that direction and others at small angles on each side of the direction should be sufficient to keep track of changes. This experiment seems likely to work, but is still constrained by the slow frame sampling rate.

Experiment 2 :

It would be interesting to attempt to minimise the number of intensity readings taken from the camera. In the treatments above, it is assumed that the intensity values used are the integrals across the area of the image strips defined by the grating. That's what you get from a physical grating used with a single photocell, but to get the same effect with a digitised image a lot of addition is necessary. An effective alternative would be to identify small samples of the surface as representatives of their strips; that would work just as well, and would materially reduce the number of intensity readings which had to be taken from the image. Unfortunately, it would be necessary to track the samples so that the same ones were always used, which isn't very practicable.

Another way to reduce the number of readings might be to use optical means to spread the image so that any point on the image received light from an area of picture amounting to one strip in dimension; a measure of the total illumination of the strip could then be obtained with just one intensity reading. This could be achieved by judiciously defocussing the image, when the intensity of illumination at any image point would measure some sort of average illumination over a circle of some diameter, which would correspond to the slit width. A reasonable estimate of the illumination over a slit could then be obtained by adding the intensity values at a series of points along the slit axis spaced at successive distances of one slit width.

Experiment 3 :

The third experiment is to explore the use of a directly-addressable sensor based on a dynamic RAM chip. The rate of discharge of an exposed RAM cell is affected by illumination, being faster in brighter light. A type of sensor has been designed based on this effect; it has the advantage that any bit can be directly addressed at RAM speeds, so access is fast. The sensors were specially designed so that a simple addressing sequence corresponded to the surface geometry in a comprehensible way, but the sensor chips were not available when we looked for them some years ago (around 1987). Ordinary dynamic RAM chips are also photosensitive[2], but do not have their cells arranged in a geometry which produces a nice picture with a simple addressing sequence; in fact, that doesn't matter much provided that we know how to address different parts of the surface.

One complication is that the sensors, being composed of single RAM cells which store one bit, are binary in nature. The value returned is determined by the rate of decay of the charge on the cell, which is determined by the light intensity, and the time since the cell was refreshed. It is possible to synthesise a somewhat quantised grey-scale picture by repeated scans at with different delays, but that depends on the object standing more or less still for some time. Experiments will be necessary to determine whether enough information can be retrieved for the velocity measurement to be effective. If only binary values are available, the sensor might only work with specialised images; chessboards come to mind.

REFERENCE.

1 : A. Blake, G. Hamid, L. Tarassenko : "A design for a visual motion transducer", *IEEE Trans. Robotics Automation* **11**, 625-633 (1995).

2 : J. Johns : *Camera* (Auckland University Computer Science Department, 07.473 Assignment 2, 1987).