

Real-time radial symmetry for speed sign detection

Nick Barnes
National ICT Australia
Canberra
ACT 2601, AUSTRALIA
Nick.Barnes@nicta.com.au

Alex Zelinsky
Department of Systems Engineering
Australian National University
ACT 0200, AUSTRALIA
Alex.Zelinsky@anu.edu.au

Abstract—Algorithms for classifying road signs have a high computational cost per pixel processed. A promising approach to real-time sign detection is to reduce the number of pixels to be classified as being a particular sign to a minimum by some form of sign detection on the image using less time expensive algorithms. In this paper, we adapt the fast radial symmetry detector to the image stream from a camera mounted in a car eliminate almost all non-sign pixels from the image stream. We then are able to apply normalised cross-correlation to classify the signs. This method is suitable for circular signs only; we apply it to Australian speed signs in this paper. Our results show that it is robust to a broad range of lighting conditions. Also, as the method is fast, there is no need to make unrealistically strict assumptions about image structure.

I. INTRODUCTION

A key technological goal in road vehicles today is to improve safety. One way this can be achieved is by creating systems within the vehicle that support the driver in reacting to changing road conditions. In our research we are particularly concerned with driver support systems. Systems that support the driver in controlling the car, but keep the driver in the loop. Within this area of driver support systems, it is important to consider that roads are highly structured environments, designed to simplify the driving task where possible.

Sign recognition is an important task for a driver support system. Signs giving information that is relevant to the local conditions appear clearly in the environment, however, a driver may not notice a particular sign due to distractions or a lack of concentration. In this case it may be helpful to make them aware of the information that they have missed. An obvious question is how the system within the car can perceive whether the driver is already aware of a particular event. This is difficult to judge for many types of sign. However, in this study we choose to address speed signs in particular because, if the driver is aware of the sign they should have adjusted their speed to a safe level. Thus, if a system can monitor speed signs, and inform the driver if they are currently driving too fast (or too slow), then the driver will only gain information that they don't know. (Alternatively, the driver may have decided not to react to the information, in which case it may not be a bad thing to inform them anyway).

Road sign recognition research has been around since the mid 1980's. A direct approach is to apply normalised cross correlation to the raw traffic scene image. This brute force method is computationally prohibitive, but this can be

eased somewhat by approaches such as simulated annealing [1]. Another method for controlling computation is to apply templates to an edge image of the road scene [2]. The method applies a distance to nearest feature transform to smooth the matching space for coarse-to-fine matching, and makes a hierarchy of templates to reduce the number to be compared. However, this is still computationally intensive and so unsuitable for an in the loop system.

Many approaches have introduced separate stages for sign detection and classification of sign type (e.g., [3], [4], [5], [6]). This is particularly used when a large number of sign types are to be classified. We argue that this can be an effective means of managing computation for even a small number of sign types if the detection algorithm has low computational cost, facilitating real-time operation. In this manner, computationally intensive classification can be performed on only a small fraction of the input image stream, without requiring assumptions about where signs may appear.

Colour segmentation is the most common method for the detection stage. Typically, this is based on the assumption that the wavelength that arrives at the camera from a traffic sign is invariant to the intensity of incident light. This assumption usually manifests in the statement that HSV (or HSI) space is invariant to lighting conditions [6]. A great deal of the research in this area exploits a detection stage based on this assumption (e.g., [6], [7], [3], [5], [8], [9], [10]), either finding the signs, or eliminating much of the image from further processing. However, the camera image is not invariant to changes in the chromaticity of the incident light. Such variance may occur in changes between conditions, such as direct sunlight, heavy cloud, smog, and under headlights at night. Further, once we consider that an imaging device must convert this wavelength to signals, the invariance of HSV space to even the intensity of incident light becomes problematic as the response to varying wavelength and intensity of standard imaging devices is non-linear and interdependent [11]. Further, as signs fade over time the colour of the signs cannot be invariant. It may be considered that HSV colour invariance is still sufficient in typical conditions. However, as opposed to ideogram based signs, speed signs only have narrow bands of colour, only a few pixels wide at most if a wide field of view is to be maintained. CCD sensors return an integration over a small field of view. If more than one colour falls on a CCD element, the colour is not easily predicted, see

Figure 1. Dealing with this problem could easily remove the fast segmentation advantage of the two step approach and requires further research.

Another approach to detection is *a priori* assumptions about image formation. At its simplest, one can assume that the road is approximately straight, so large portions of the image can be ignored as signs will not appear in them. Combined with colour segmentation, Hsu and Huang [10] look for signs in only a restricted part of the image. However, such assumptions can break down on curved roads, or with bumps such as speed humps. A more sophisticated approach is to use some form of detection to facilitate scene understanding, and thus eliminate a large region of the image. For example, Piccoli *et. al* [7] suggest large uniform regions of the image that correspond to the road and sky, and thus only looking in the region along side the road and below the sky where signs are likely to appear. However, this will not be adequate in more difficult road scenes, such as shown in Figure 4 (b). They also suggest ignoring one side of the image as signs will only come up on one side. This is inadequate for multi-lane highways, and throwing away information in scenes such as Figure 4 (a).




Fig. 1. A typical detected candidate sign (a) at input image size, and (b) close-up. The outer circle and numbers are narrow. Despite its consistent appearance as a small image to our eyes, it contains few pixels that could be said to be red, black, or even white.

We propose a new efficient method for sign detection: the fast radial symmetry detector [12]. It is applicable to signs with a circular feature, a significant subset of signs. Many shape detectors are non-robust because they require closed shapes. Robust techniques such as Hough circle detection [13] are slow to compute over large images. The fast radial symmetry detector can be run as a detector at frame rate. All Australian speed signs have a red circle on a white

background with black numbers. We are able to eliminate the vast majority of false positives by considering only radially symmetric regions, that are stable across several images, have a high count of pixels in ratio to the radius.

This has now reduced the locations that are possible for signs to occur to very few. The typical next stage is cross-correlation. However, now there are very few pixels where templates must be applied. Further, for cross-correlation, scale is a problem, typically requiring multiple templates at different resolutions. However, from the radius returned from the fast radial symmetry detector we know the approximate scale of the template.

This paper presents an application of the radial symmetry operator to visual speed sign recognition. It identifies possible candidate speed signs in an image stream from a video camera mounted within the car looking along the road. This information can then be compared against the vehicle's speedometer, and passed onto the driver if it appears that they have not reacted to the change of conditions. Our approach exploits the structured nature of the road to facilitate fast processing. We also demonstrate how this operator can be combined with cross-correlation on candidates to yield classification.

II. SPEED SIGN RECOGNITION

There is much possible variation in the appearance of a speed sign in an image. Throughout the day, and at night time, lighting conditions can vary enormously. A sign may be well lit by direct sunlight, or headlights, or on the other hand it may be completely in shadow on a bright day. Further, heavy rain may blur the image of the sign. Ideally, signs have clear colour contrast, but over time they can become quite faded, but still be quite clear to drivers. Although signs appear by the road edge, this may be far from the car on a multi-lane highway - to the left or right, or very close on a single lane exit ramp. Further, although signs are generally a standard distance above the ground, they can also appear on temporary roadworks signs at ground level. With this nature, it is not simple to restrict the possible viewing positions of a sign within the image. By modelling the road [14], it may be possible to dictate parts of the image where a sign cannot appear, but road modelling has its own computational expense, and, as discussed previously, colour-based methods are not robust.

However, the roadway is well structured. Under Australian law speed sign appearance is highly restricted: a particular size, and must be a white sign with black numbers surrounded by a red circle. Unless the sign has been tampered with, signs will appear approximately orthogonally to the road direction. Finally, signs are always placed to be easily visible, so the driver can easily see them without having to look away from the road. In this paper we assume that the signs are vertical along side the road, however, with minor accidents, signs can occasionally appear tilted.

This algorithm searches for near circular features. As a legal speed sign must have a red circle around it, and

signs almost always appear in the orthogonal direction to the road, provided our camera points in the direction of vehicle motion, the surface of signs will be approximately parallel to the image plane. On a rapidly curving road it may be that the sign only appears parallel to the image plane briefly, but this will be when the vehicle is close to the sign, so it will appear large in the image. If we are processing images at around the standard frame rate, and we are able to recognise a sign reliably from only a small number of frames then generally we are safe to assume that the sign is approximately parallel to the image plane.

A. Candidate detection

The fast radial symmetry detector [12] is a variant on the circular Hough transform that executes in order kp , where p is the number of pixels, and k is the number of discrete radii that are searched. This is as opposed to the traditional circular Hough transform that executes in order kbp . For the traditional circular Hough transform, each edge pixel votes on all circles over a discrete set of radii k that could pass through that edge pixel. The factor b comes from the a discretisation into a number of bins on the gradient of circular tangents that could pass through this point. The fast radial symmetry detector eliminates the factor b by taking the gradient of the edge point directly from the output of the Sobel edge detector. In this way, the computation of the radial symmetry detector is reduced by a factor of b , but also the resulting circle map is simplified by a dimension. This makes it suitable for real time use, 13.2 ms for a 240 x 320 image [12].

To better explain our adaptation of the fast radial symmetry detector, we include a description, largely taken from [12]. For a given pixel, p , the gradient, g , is calculated using an edge operator that yields orientation, such as Sobel. If this pixel lay on the arc of a circle, then its centre would be in the direction of the gradient, at the size of the radius. To achieve lighting invariance, we apply the discrete form of the detector. The location of a pixel that will gain a vote as a potential centre is defined:

$$\mathbf{p}_{+ve} = \mathbf{p} + \text{round} \left(\frac{g(\mathbf{p})}{\|g(\mathbf{p})\|} n \right), \quad (1)$$

where $n \in N$ is the radius, and N is the set of possible radii. In our application to sign detection, this is defined by our expectations about the apparent sign size. (A negative image is defined similarly, facilitating constraining the operator to find only dark light circles, or light dark only.) A vote image is defined based on these orientation votes as:

$$O_n(\mathbf{p}_{+ve}) = O_n(\mathbf{p}_{+ve}) + 1 \quad (2)$$

The vote image is defined as:

$$\hat{F}_n(\mathbf{p}) = \text{sgn}(\tilde{O}_n(\mathbf{p})) \left(\frac{|\tilde{O}_n(\mathbf{p})|}{k_n} \right)^\alpha, \quad (3)$$

where α is the radial strictness parameter, and k_n is a scaling factor that normalises O_n across different radii. Also,

$$\tilde{O}_n(\mathbf{p}) = \begin{cases} O_n(\mathbf{p}), & \text{if } O_n(\mathbf{p}) < k_n, \\ k_n, & \text{otherwise.} \end{cases} \quad (4)$$

To obtain the radial symmetry image, \hat{F}_n is convolved with a Gaussian. There are several images produced by the transform. Each radii of N votes into a separate image. There is also a full transform image that averages all the symmetry contributions over all radii considered:

$$S = \frac{1}{|N|} \sum_{n \in N} S_n. \quad (5)$$

See [12] for full details.

In order to adapt the transform for the sign detection case, we apply it to only to radii that are practical for detecting speed signs in traffic images, this is between five and nine pixels on our images. A circle with a small number of pixels as its radius may well constitute a speed sign, however, there will not be enough pixels present to discern what the sign says, and so there is no point in further processing. We should wait until the sign is close enough to be recognised. Further, in normal driving conditions a sign will never appear closer to the camera than several metres. Given a camera of approximately known focal length, we can impose an upper limit on the possible radius of circles that we are interested in. In our system these limits were empirically derived based on sample road images. As the majority of circle pixels will be detectable even on heavily degraded signs, we require that a large percentage of any candidate circle is covered by edge pixels with gradients that indicate the same centre. The fading of signs tends to be fairly even, so it is likely either most pixels will be above the detection threshold, or almost none will be. This restriction eliminates partially circular false positives.

To detect the signs, we search the full transform image, Equation (5), for values above a threshold defined as the required number of pixels for the smallest radius. We then check all the individual radii images, Equation (3), to find if any are above the threshold of edge pixels for that radius. It could be that a high vote in Equation (5) results only from a collection of accidental votes from several contributing radii. If a particular radius is above the threshold, then this is a possible circle candidate.

Further consistency checking can be performed over time. A circle must appear for at least two concurrent frames, it cannot have changed radius by more than two pixels during that time, and it must not move by more than 20 pixels in the image. It may be possible to model car motion in predicting the new location of the image where the sign should appear. However, one would have to assume a smooth straight road, which is quite restrictive, or include feedback of steering direction, which is complex, but may contribute further.

We track up to three candidates in any single image, and given temporal consistency, may find up to three candidates.

The output is the region of the image in which the candidate appears, a radius, representing the scale of the candidate, and a centre point of the candidate.

B. Classification

To demonstrate that the symmetry operator can combine effectively with classification, we applied normalised cross-correlation from the Intel Integrated Performance Primitives library. We made a template of the sign text from a close by sign, and scaled it down by linear interpolation across all colour channels to all expected radii N , to form a total of eight templates for each sign. At run-time, the appropriate subset of templates was chosen based on the radius found from the operator. Other authors have scaled the actual input image [7], but scaling the templates can be performed *a priori* giving better performance, provided the scale of the input candidate is known.

As the sign is assumed to be vertical and parallel to the image plane, no scaling or rotation of the template is necessary. Also, as the centroid is isolated, it was only necessary to apply cross-correlation over a 5×5 window. Thus, for each candidate, we only apply cross-correlation to $5 \times 5 \times 3 \times s$ pixels, where 3 is the number of radii we allow, and s is the number of possible speed signs. For many signs s could be reduced by the use of hierarchical templates. The base templates used for 40 and 60 signs are shown in Figure 2.




Fig. 2. The base templates used, these were scaled down in size to match to multiple sizes. (Note that these are enlarged in this figure for better visibility.)

III. RESULTS



Fig. 3. The ANU/NICTA Intelligent Vehicle. The cameras are mounted where the rear view mirror would be.



Fig. 4. Some sample images with speed signs present. The quality of sign and the lighting varied within our sequences, along with the scale of the sign that appeared. Also, more than one sign may appear in a single image.

The system was run over several raw image sequences taken from the ANU/NICTA intelligent vehicle, see Figure 3. The sequences come from cameras mounted in a binocular head mounted approximately in the position of the rear-view mirror. All images used in the experiments were taken of signs on public roads around Canberra, including on campus. Some of the sequences were taken at frame rate while driving at around the speed limit, while others were taken while stationary on the road in front of the signs. Several sample signs from those used in our experiments are shown in Figure 4 and 5. As can be seen, there is great variation in the appearance of the signs, including their scale in the image, lighting, and the deterioration of the sign viewed due to weather conditions.

The radial symmetry detector was run over a total of 1107 frames from the camera. From this sequence 152 sign candidates were detected. Of these 90% were correctly detected. This is quite a good number from a fast early classification stage. It is not simple to say how many false negatives were reported, certainly there were a number of cases where a speed sign was not found, but some of these cases the sign is poor enough that no sign should have been reported, or small. For most scenes, as the sign approached, it was detected many times before it passed out of view, this would give classification a good chance to correctly classify the sign. For the clear scenes, it was detected for most frames where it appeared. An exception occurred in a sequence when it was raining, here no sign was detected. Some of the input image stream was taken under quite difficult conditions including shadowing and rain. Some difficult successful examples are shown in Figure 5. Note that (c) would be a difficult image for any detectors that rely on uniform colour regions to reduce computation by eliminating the sky or road. Figure 6 shows some images where a sign was present, but not detected.

The outputs of the classifier were not very helpful in discriminating candidates from non-candidates. Many non-candidates still produced reasonably high correlation values with the templates. However, the class they were classified into was random. Our results indicate that requiring a classification to be consistent for several frames would

eliminate many of these false positives.




Fig. 5. Some images where candidate detection was difficult, but a candidate was found.

To demonstrate that the radial symmetry operator was capable of being used efficiently in concert with the standard cross-correlation, we ran the two together as a system. Each candidate that was returned was individually evaluated for being either a 40 sign or a 60 sign. From 126 valid candidates 96% were correctly classified. Of these, unfortunately, only 75% of the 25 40 signs returned as candidates were correctly classified. Note that classification is not the emphasis of this paper, so the performance was not tuned. It does clearly demonstrate, however, that the radial symmetry operator can be combined effectively with cross-correlation to classify roadsigns in real time.

A. Real-time results

The full system was implemented in c++ and set up to run directly from a camera, or captured sequence to test the real-time performance. It was found that for a 320x240 image the full detection and classification was able to run at 20Hz. This was only with 40 and 60 signs, however, it would scale well to more signs as the time taken to run a single template against a candidate was ≤ 1 ms.

IV. CONCLUSION

Detection using the fast radial symmetry detector has been demonstrated to be highly successful, with only 10% false positives over many varied scenes, and a total of 1107 input images. It would be desirable to throw away these false positives before proceeding to classification. Additional detector phases could be introduced after radial symmetry detection. These methods can have some greater computational expense at this stage, as they would only operate on a tiny fraction of the input data.




Fig. 6. Images where a candidate was present, but not detected: in (a) and (c) the sign is self-shadowed on a bright day, and so is too dark for adequate contrast for edge detection, (b) is too small, the sign was detected in a subsequent image, slightly closer to the vehicle, and, (d) the faded sign produces insufficient contrast for edge detection.

The performance of the classifier on candidates produced from by the detection phase demonstrated that the detector can be effectively included in a classification system. The classification results could be improved through tuning, and requiring temporal consistency of classification.

V. FURTHER WORK

Classifying 40 and 60 signs, we have just proved the concept. The system needs to be able to recognise the full range of speed signs that can occur on the Australian road. It would be interesting to further develop other detectors that are capable of finding signs with other shapes, such as stop signs and give-way signs.

VI. ACKNOWLEDGMENTS

National ICT Australia is funded by the Australian Department of Communications, Information Technology and the Arts and the Australian Research Council through *Backing Australia's ability* and the ICT Centre of Excellence Program.

REFERENCES

- [1] M. Betke and N. Makris, "Fast object recognition in noisy images using simulated annealing," A.I. Lab, M.I.T., Cambridge, Mass, USA, Tech. Rep. AIM-1510, 1994. [Online]. Available: citeseer.nj.nec.com/betke94fast.html
- [2] D. M. Gavrilu, "A road sign recognition system based on dynamic visual model," in *Proc 14th Int. Conf. on Pattern Recognition*, vol. 1, Aug 1998, pp. 16–20.
- [3] P. Paclik, J. Novovicova, P. Somol, and P. Pudil, "Road sign classification using the laplace kernel classifier," *Pattern Recognition Letters*, vol. 21, pp. 1165–1173, 2000.

- [4] J. Miura, T. Kanda, and Y. Shirai, "An active vision system for real-time traffic sign recognition," in *Proc 2000 IEEE Int Vehicles Symposium*, Oct 2002, pp. 52–57.
- [5] B. Johansson, "Road sign recognition from a moving vehicle," Master's thesis, Centre for Image Analysis, Swedish University of Agricultural Sciences, 2002.
- [6] L. Priese, J. Klieber, R. Lakmann, V. Rehrmann, and R. Schian, "New results on traffic sign recognition," in *Proceedings of the Intelligent Vehicles Symposium*. Paris: IEEE Press, Aug. 1994, pp. 249–254. [Online]. Available: citeseer.nj.nec.com/priese94new.html
- [7] G. Piccioli, E. D. Micheli, P. Parodi, and M. Campani, "Robust method for road sign detection and recognition," *Image and Vision Computing*, vol. 14, no. 3, pp. 209–223, 1996.
- [8] C. Y. Fang, C. S. Fuh, S. W. Chen, and P. S. Yen, "A road sign recognition system based on dynamic visual model," in *Proc IEEE Conf. on Computer Vision and Pattern Recognition*, vol. 1, 2003, pp. 750–755.
- [9] D. G. Shaposhnikov, L. N. Podladchikova, A. V. Golovan, and N. A. Shevtsova, "A road sign recognition system based on dynamic visual model," in *Proc 15th Int Conf on Vision Interface*, Calgary, Canada, 2002.
- [10] S.-H. Hsu and C.-L. Huang, "Road sign detection and recognition using matching pursuit method," *Image and Vision Computing*, vol. 19, pp. 119–129, 2001.
- [11] M. Thomson and S. Westland, "Colour-imager characterization by parametric fitting of sensor responses," *Color Research and Application*, vol. 26, no. 6, pp. 442–449, Dec. 2001.
- [12] G. Loy and A. Zelinsky, "Fast radial symmetry for detecting points of interest," *IEEE Trans Pattern Analysis and Machine Intelligence*, vol. 25, no. 8, pp. 959–973, Aug. 2003.
- [13] L. G. Minor and J. Sklansky, "Detection and segmentation of blobs in infrared images," *IEEE Trans. on Systems, Man, and Cybernetics*, vol. 11, no. 3, pp. 194–201, Mar. 1981.
- [14] R. Labayrade, D. Aubert, and J.-P. Tarel, "Real time obstacle detection in stereovision on non flat road geometry through v-disparity representation," in *Proc IEEE Int Vehicles Symposium*, June 2002.