Application of the Hough Transform to Lane Detection and Following on High Speed Roads

John B. Mc Donald*
Signals & Systems Group, Department of Computer Science,
National University of Ireland, Maynooth, Ireland.

Abstract

A new Hough transform based technique for the automated detection and tracking of road contours on high speed (*i.e.* low curvature) roads is presented. Given a scene acquired from a dashboard mounted camera system, the road contours of the current lane in the near field of view are automatically detected and tracked. Parameter space search optimisations and tracking strategies are outlined. Results are presented from the application of the technique to real data acquired from a motorway driving scenario.

Keywords: Automated vehicle guidance, lane detection, Hough transform.

1 Introduction

The last two decades have seen a proliferation in the number of advanced automated vehicle subsystems in use, which has in turn resulted in a driving environment which is both safer and more efficient [Bertozzi and Broggi, 1997]. If a focal point of these systems was to be defined it would be total automation of vehicle guidance.

In this paper we are primarily concerned with lane detection and following. Given a scene acquired from a dashboard mounted camera system, we wish to automatically detect the road contours of the current lane in the near field of view and, to track these contours in subsequent frames. It is also desirable, (but not required), to locate other road contours in the scene. Solutions to this problem are applicable in semi-autonomous cruise control systems, lane departure warning systems, and play an integral role in automated vehicle guidance systems.

A preliminary computational model for road detection and following based around the Hough transform was presented in [McDonald et al., 2001a]. In this paper that computational model is extended through the development of geometric constraints in Hough space, and by incorporating extra temporal processing for the reduction of false road contour detections. Other extensions detailed in this paper include the detection of all road contours, as well as the identification of occluding objects such as cars in the neighbouring lane. We also provide details and results from more comprehensive experiments.

The problem domain constitutes a well developed area with the main work being done by groups at Carnegie Mellon University [Taylor et al., 1999, Thorpe et al., 1998], Universität der Bundeswehr [Dickmanns and Mysliwetz, 1992], and Universitá di Parma [Bertozzi et al., 2000].

A review of the state of the art in automated vehicle guidance systems is presented in Section 2. This is followed, in Sections 3 and 4, by a description of the Hough based approach developed at NUIM. Qualitative results from the technique are presented in Section 5, with an analysis and concluding remarks presented in Section 6.

2 Automated Lane Following

2.1 State of the Art

The field of automated vehicle guidance is a well developed area, with many successful approaches previously presented in the relevant literature. In this section a number of the more prominent approaches are described, which although not complete should illustrate the range of schools of thought which have been applied to the problem.

Major research in the field started during the mid-eighties with various projects being initiated in Europe, the US, and Japan. In Europe the most prominent project was known as PROMETHEUS

*Corresponding author.*Email address: johnmcd@cs.may.ie*

(PROGraM for a European Traffic with Highest Efficiency and Unprecedented Safety), and involved many vehicle manufacturers working in cooperation with research institutes and universities.

The US and Japan on the other hand initiated many projects which recently resulted in national efforts being formed in both countries (known as the AHSRA and NAHSC respectively). For more information see [Bertozzi et al., 2000, Koller et al., 1999] and the references therein.

In the past, researchers have applied many different types of sensor to this problem but of all the modalities investigated vision has proven to be the most successful. The main reasons for this are due to its high information content and the fact that it is a passive technology. Passive sensor technologies are advantageous since active sensors tend to have lower resolution and slower scanning speeds [Bertozzi et al., 2000]. Another major drawback of active sensors is their tendency to interfere with each other when in close proximity. With this said active sensors do have their own particular advantages in the area and, hence many approaches have opted for a sensor fusion scheme where redundancy provides enhanced robustness [Hofmann et al., 2000].

As a result of the projects described above, and other initiatives many techniques have been developed. One of the most successful approaches reported to date has been that of work done by Dickmanns et al. [Dickmanns and Mysliwetz, 1992] at the Universität der Bundeswher. This work centers around what has been termed the *4D-approach* to machine vision. Here dynamical models are used to recursively estimate vehicle and road state parameters in a predefined look-ahead range [Dickmanns, 1988].

The system developed at UBM was implemented in two separate test-beds. The first was a 5 tonne van known as VaMoRs, which was tested on a construction site for a new strip of Autobahn. This vehicle was driven autonomously at a speed of approximately 100km/h over a total length of 24km [Dickmanns, 1988]. VaMoRs was subsequently upgraded to a Mercedes saloon car known as VaMP. This latter vehicle was proven robust in autonomously driving for 95% of a 1600km journey from Munich to Odense, in Germany.

One of the main problems with the approaches similar to UBMs is their difficulty in coping with variation in the driving environment i.e. less structured environments. This has lead researchers to investigate more adaptive approaches based on machine learning. Here it is hoped that since a rigid model does not exist in the system it should be more capable of adapting to unseen environments.

One such system known as ALVINN (Autonomous Land Vehicle in a Neural Network) has arisen from the long running CMU based NAVLAB project. Here low resolution images (to allow real-time capability) are taken from a vehicle mounted camera and passed as input to a 3-layer feed-forward neural network. During the learning phase control inputs are also passed to the ANN. This allows the network to learn by simply observing and coupling particular driving scenarios with human driver responses.

As with many systems based purely on learning strategies, it was seen that increased performance could be achieved by incorporating high level knowledge with a learning based approach. The resulting hybrid strategy should capitalise on world models whilst maintaining the ability to generalise to new environments. In reaction to this CMU (Pomerlau et al. [Pomerlau, 1995]) developed a second system, known as RALPH (Rapidly Adapting Lateral Position Handler), which although it does not use a neural network model, avoids explicit representation of road parameters.

The GOLD (General Obstacle and Lane Detection) system developed at the Universitá di Parma uses an inverse perspective mapping to remove perspective distortions introduced by the imaging system [Bertozzi et al., 2000, Bertozzi and Broggi, 1997]. It then enhances and binarizes the birds-eye view which results in an image with enhanced road contour pixels. Next short line segments are extracted and concatenated into larger segments, which are then selected or rejected based on an estimated road geometry.

The GOLD system was installed and tested on a Lancia Therma passenger vehicle (named Argo) and exercised extensively over a 2000km journey during which the vehicle drove itself for about 94% of the total trip. [Bertozzi et al., 2000]

From the systems described in this section it should be clear that the state of the art in this area is quite advanced, with autonomous highway cruise control predicted to be brought to market within the next 5-10 years. Fully autonomous vehicles are however some distance away. Prior to such systems being developed problems such as dealing with unstructured environments, user interfaces, formal safety verification, along with many others need to be solved.

3 The Hough Transform

The *Hough transform* (HT) [Duda and Hart, 1972, Hough, 1962] is a technique which, using some curve representation, transforms a set of points defined over the image space to a set of points defined over some parameter space (known as *Hough space*) [McDonald and Vernon, 1998, McDonald et al., 2001b]. Points in Hough space represent particular instances of the curve in the image. Therefore, the strategy used by the HT is to map *sets* of points from a particular instance of the considered curve, *i.e.* the parameterised curve, to a *single* point representing the curve in Hough space and, in effect, cause a peak to occur at that point.

It has long been known that the HT can be considered as a special case of the more general *Radon transform* (RT) [Deans, 1981]. So much so, that both have been presented using an identical computational model(Eq. 1) [Barducci and Pippi, 1999, Deans, 1981].

$$H(\rho,\theta) = \int\int I(x,y)\delta(\rho - x\cos\theta - y\sin\theta)dxdy. \tag{1}$$

Equation 1 can be interpreted as evaluating the total pixel intensity along a line parameterised by ρ and θ. Other interpretations of this transform view the process as that of accumulating votes for each possible line in the image. Lines for which a high number of votes accumulate result in occurence of peaks in *Hough space*.

One of the main advantages of the HT is its robustness to noise and occlusion. This is due to the fact that each image point is considered independently of the others. Hence, within reason, the removal of valid edge points, or the introduction of spurious noise points only alters peak intensities in Hough space. A direct consequence of this is the techniques capability to detect broken line patterns used in central road markings, or road markings partially occluded by obstacles (*e.g.* vehicles).

4 Lane Detection and Following using the Hough Transform

In this section we apply the Hough transform to automated lane detection and following. Specifically we are interested in the automatic extraction of road position during motorway driving scenarios. Although the assumption that the road is linear is not true, it is compensated for by the fact that the Hough transform may be *tuned* to treat a road segment of low curvature as a linear segment. This is achieved by decreasing the resolution of the Hough space.

4.1 Geometric constraints

Initial lane contour locations are unknown but are assumed to lie within a certain predefined range. This is possible due to the fact that the range of lateral vehicle positions on the road, and hence camera viewpoints, is bounded. Since the road contours in the near field of view closely approximate a set of parallel lines, the set of possible projections is easily determined. It should be noted however that changes in the camera parameters will effect these ranges.

A simple algorithm is used in detecting the initial contour locations by first setting *θ-windows* for the contours at both the extreme left and extreme right of the road. The algorithm procedes by first searching for the extreme left lane in the scence by computing the Hough transform over the set of θ values defined by the left-lane θ-window. Next, the same procedure is performed for the extreme right lane using its corresponding θ-window. The θ values for both of these contours are then used to dynamically set the θ-window for the central road marking. The bounds for this window are computed such that they never overlap with the previously detected contours *i.e.* an offset is added (or subtracted) from a detected side marking to arrive at a window boundary. This results in a system within which ambiguity between the various road contours is minimised.

No explicit constraints were placed on the permittable values of ρ for the various road contours. However implicit constraints were placed on this parameter through the use of a geometrically inspired Hough space optimisation. If we again consider the fact that the road contours are parallel and lie in a common plane, then the projection of these contours into the image plane results in a set of lines all of which intersect in a common point on the horizon line. From the Hough transform it is known that parameter pairs for any set of lines which intersect on a common point in image space must lie on the hough transform of that point of intersection. That is, since all the road contours intersect at a point on the horizon, their parameters must therefore lie on the hough transform of that horizon point.

Figure 1 illustrates this point. Here images (a) and (b) show a frame taken from both extremities of the lateral road positions for the car. As can be seen from these images, in both cases the points on the horizon where the road contours intersect lies within a small sub-region of the image. Hence we can assume that, in any given frame, all road contours will be colinear with some pixels inside this sub-region. To translate this into a contraint in Hough space we compute the Hough tranform of the pixels in this sub-region (i.e. compute the hough transform of image (c)). Image (d) shows a thresolded version of this hough space where the white region delimits the permittable (ρ,θ) pairs i.e. all the points in this region represent lines in image space which pass close to the hoirzon point.

This results in an implicit constraint on the ρ values in that for any particular θ value only the ρ values which lies in the white region for that column may be considered.

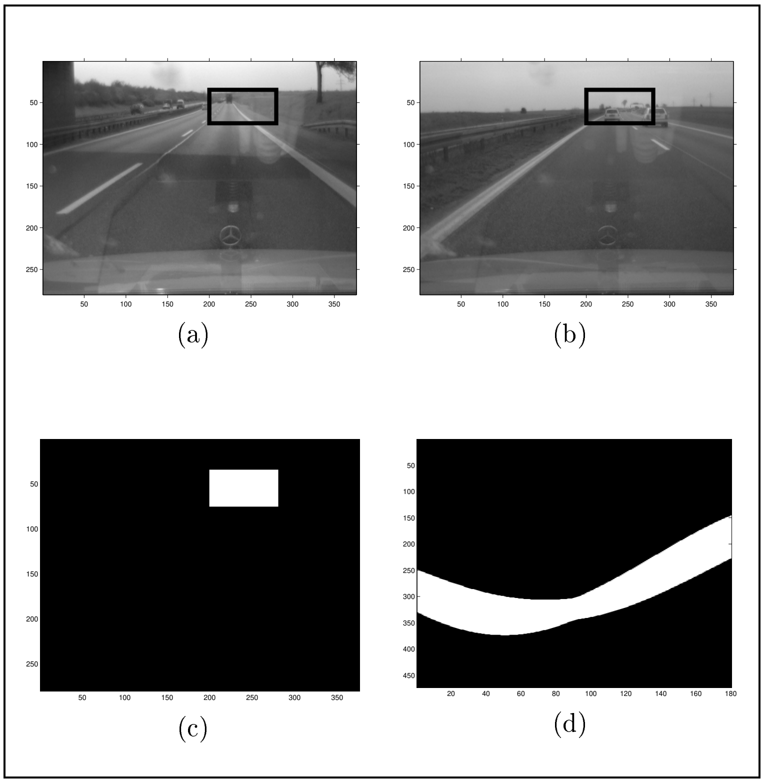

Figure 1: Hough space optimisations

4.2 Temporal filtering

Temporal variation in contour parameters is catered for by two complimentary techniques. Firstly as a new frame is processed the range of parameter values searched is recentered around the parmeters from the previous frame. Since road width is bounded these windows are also bounded. This provides a safe guard in instances where a particular contour becomes obstructed and the tracking strategy breaks down. Rather that allowing the window to continue tracking invalid contours beyond the possible range of θ values, instead the window will sit at the boundary. When the valid contour reappears the occlusion strategy described below reactivates the tracking and proceeds as normal. Since the occlusion strategy opens up the θ-window for the obstructed curve, even if the lateral position of the vehicle, (and hence the θ value for the contour in the image), has changed the system will re-identify the correct contour.

Secondly in order to reduce the occurence of outliers in the data, temporal filtering is employed in the form of an exponential averaging process(Eq. 2)

$$\tau_i = \alpha * x_i + (1 - \alpha) * \tau_{(i-1)}, \tag{2}$$

where τ_i is the predicted value, x_i is the measured value, and τ_{i-1} is the predicted value from the previous iteration. Here the α parameter is used to weight the influence of the current measurement and past predictions on the new prediction. High values of α result in a measurement process which is dynamic and prone to outliers, whereas low values of α result in a measurement process where past predictions outweigh outliers but which is overall more “lethargic”.

Hence a key issue here is how to set the α parameter, which is obviously related to the current driving situation. For example, the optimal α parameter used for when driving straight will obviously be different to that used when involved in overtaking maneuvers.

Current work at NUIM is addressing this issue by coupling the α parameter with ego-motion information from vehicle sensors. From initial tests this approach has shown promising results and we believe lends itself well to integration in closed-loop control strategies using more complex temporal processing algorithms. In the future we hope to replace this exponential averaging with a more optimal approach such as that of [Dickmanns and Mysliwetz, 1992] which uses the extended Kalman filter.

One problem with exponential averaging is that since the prediction is always incorporated in the final value, slight deviations from the real value generally ensue. To overcome this problem if the new computed value is within a specified threshold the system accepts it as being the true value. Hence the predicted value is only considered if it is radically different from the measured line parameters.

4.3 Dealing with occlusion

In any road detection system some mechanism is required for dealing with occlusion of road contours due to other vehicles, lack of well defined road markings, *etc.* In this system such problems are dealt with directly in Hough space by monitoring variations in peak intensities. Occlusion due to other vehicles results in an edge image where some, or even none, of the road contour's edges exist. In such instances the behaviour of the peak due to the obstructed road contours is characterised by a gradual degradation in peak height. To identify these scenarios as the contour peaks are tracked the various peak heights are monitored. As the peak height for any contour drops below a certain threshold value the system assumes that the tracked peak is no longer valid and resets the the tracking window for that contour.

This reseting process consists of using the position of the other road contours in the image to make a heuristic estimation of the the position of the occluded road contour. The tracking window is set such that inaccuracies in this estimation do not affect the techniques ability to re-identify the contour at a later stage. As the occluding object moves out of the way of the occluded contour the peak moves above the threshold and reappears within the tracking window.

5 Experiments and Results

5.1 System configuration and parameters

Experiments were carried out on a 650 frame (25fps) frame image sequence of normal motorway driving. The system was implemented in Matlab using a standard Pentium II 350MHz with 64MBytes of RAM. Images were processed at a resolution of 376×280 where only the second quarter of the image rows were processed. With the origin defined to be the top left-hand corner of the image, this image region consists of rows 94 to 188. This section of the image contains the majority of the road in the near field of view.

The window of permittable horizon points in which the road contours could intersect was defined to be a region of size 50×40 pixels centered horizontally and overlayed directly above the processed image region. Both this horizon region, and the region within which the road was detected are illustrated in all figures by a dashed black lines.

The processing rate for each frame was approximately 1Hz. It should be noted that this processing rate is not representative of real-world system performance since at present it is implemented in a prototyping language. As such it is simply mentioned here for completeness.

5.2 Preprocessing

Since the Hough tranform works on binary valued images it was first necessary to preprocess the input frames. This preprocessing consisted of a combination of thresholding and edge detection through the use of a sobel operator. The process is illustrated in Figure 2. Here, image (a) shows an original image frame from which we wish to extract lane contour edge points. All other edge points are considered as noise.

In our initial experiments we assumed that thresholding the image (Figure 2(b)) would yield the necessary road contour pixels. However, as can be seen, this approach suffers due to the glare from the road in the far field of view. Increasing threhold levels here fails to yield improved performance. To overcome these difficulties, it was decided to instead use edge detection which led to improved perfomance but resulted in many false indifications of central reservation barriers as road contours. Since such barriers are parrallel to the road contours they satifies the horizon points constraints detailed in the previous section, and hence result in difficult ambiguities.

To overcome these difficulties we used the fact that the invalid edge pixels, due to road barriers, were uncorrelated with the invalid pixels due to road glare in the thresholded image. However the valid road contour edge pixels existed in both images. Hence by logically ANDing images (b) and (c) together invalid edge pixels are supressed, as demonstrated in image (d). Obviously this is approach is not *fool-proof* but has proven to dramitically increase the systems performance.

5.3 Results

Quantitaive analysis of the system caused some dificulty mainly due to the small amount of test data, the lack of ground truth measurements, and the subjective nature of the notion of accuracy (especially without ground truth data). However, in order to validate the system, tests on the real-world data were carried out with the results evaluated by a human operator (*i.e.* for each frame a human operator determined whether the positions of the detected contours were correct).

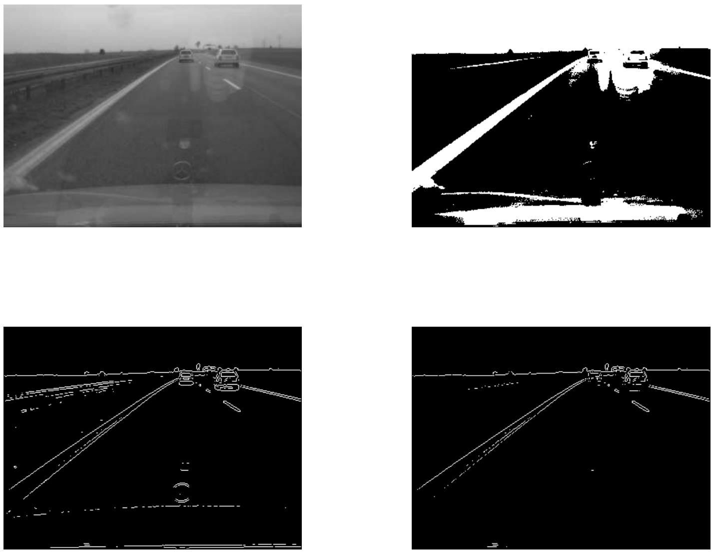

Figure 2: Frame pre-processing

To perform the tests two sub-sequences were extract from the main image sequence, consisting of driving whilst mantaining road position, and driving during a lane change maneuver, respectively. The former consisted of 500 frames whilst the latter consisted of 150 frames. Both of these sequences were then processed with α set to 0.6.

In the first sequence the outer left contour was tracked almost perfectly with misclassifcation only occuring in two frames. This misidentification occurred due to failure of the preprocessing step to suppress edge pixels due to the central reservation barrier. Hence ambiguity between the central reservation barrier and the road contour occurred. The other two contours had a much higher misidentification rate, with errors occurring in 66 out of the 500 frames. These misclassifations occurred due to the window for the outer right contour overlapping with the central contour and hence the central contour was misidentified as the right lane contour. When this happens the window for the central contour is displaced and hence correct classification is impossible.

A small sample of results from the system are shown in Fig. 3 and Fig. 4. In Fig. 3(a) the system is actively tracking and correctly indentifying the three road contours. Figures 3(b) to (e) show a passenger car in the right adjacent lane moving into the analysed window and occluding the outside right contour. The system indentifies this problem by detecting a sharp drop in the height of the corresponding peak in Hough space. It illustrates this fact by overlaying a broken black line on the last known position of the identified contour. As explained above, when this occurs the θ window for that contour is reset and the system repeatedly searches for the reappearance of the peak. Finally, Fig. 3(f) shows the first frame in the sequence where the system reidentifies the contour of interest.

Test performed on the lane change sequence were resulted in 100% accuracy for each contour. One point to note however is that no occlusion occurred during this sequence and hence performance was expected to improve.

In Figure 4 tracking of the contours over a lane change maneuver is demonstrated. Here it can be seen that each of the contour is track robustly with the θ-windows constantly recentering themselves about the previously detect peak. Figure 5 shows a plot of the the detected θ-parameters with their corresponding θ-windows. As can be seen from this plot the central window marking is particularly wide yet during this sequence the contour was detected robustly.

6 Concluding Remarks and Future Work

In this paper a new technique, based on the Hough transform, for the detection and tracking of road contours in images was presented. Geometric constraints due to the road contours being parallel and lying in a ground plane were outlined. These contraints dramatically reduced the region of Hough space searched for peaks. Temporal contraints were developed in the form of a dynamic *focus-of-attention* windows in Hough space. In conjunction with this an exponential averaging filter was applied to reduce the occurence of outliers in the data affecting the detection process. Finally a simple method for indentifying

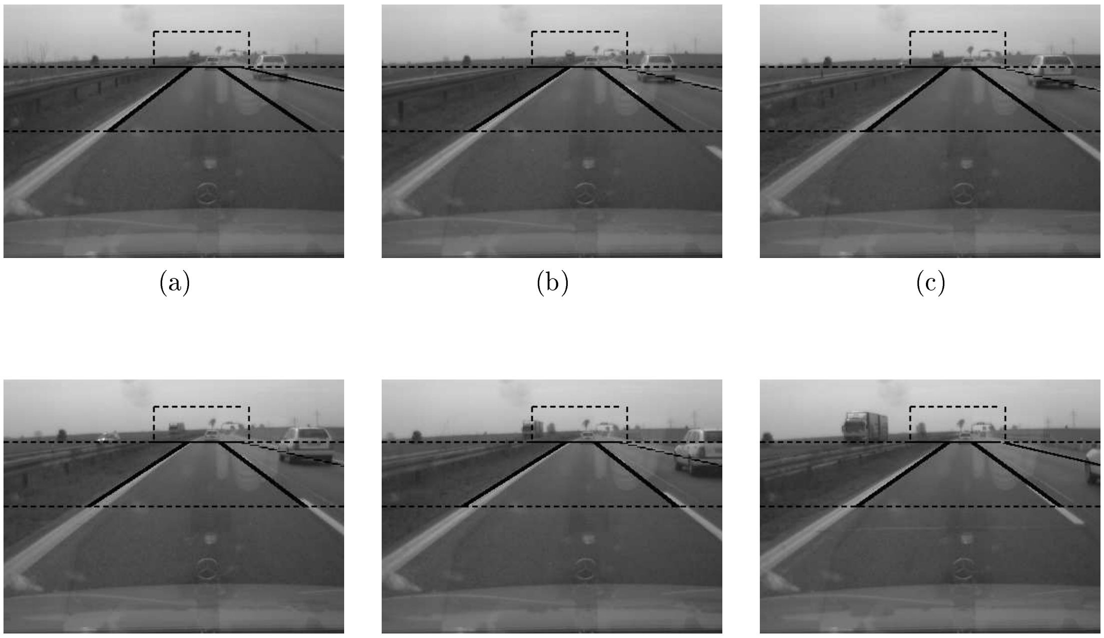

(a) (b) (c)

(d) (e) (f)

Figure 3: Lane following from motorway image sequence: (a) oncoming vehicle, (b)-(e) obstalce detected (indicated by dashed line), (f) contour tracking reactivated.

occlusion of road contours based on peak heights in Hough space was outlined.

As mentioned above, at present parameter ranges are set empirically. Current work is dealing with investigating a projective framework within which these values may be computed more robustly. Here it is assumed that the road contours are parallel, equidistant from each other, and contained with the ground plane. Projection of points on this ground plane into image space is governed by a homography. The implication of this is that assuming it is possible to track any single one of these contours, which we have shown to be the case, the entire projection of the road is directly computable.

That is to say, strong detection of a particular road contour (through the use of the peak size in the Hough space), could be used in conjunction with the projective model of the ground plane and salient road parameters, to infer possible positions of the other contours in the image. Using this type of approach would make it possible adapt more easily to the current lateral position of the vehcile on the road, and to variations in road geometry, such as road width, *etc.*

The other area of possible improvement, as mentioned above, is in the dynamical models used. From our experiments so far although the exponential averaging technique does aid in the removal of outliers in the data, it does so in a unbiased manner *i.e.* it does so without taking into account the current state of the vehcile, image acquisition system, or road geometry. Since much of this information is available to us, some way of intergrating it into the system is required. The work of Dickmanns et al. [Dickmanns and Mysliwetz, 1992, Dickmanns, 1988], and Taylor et al. [Taylor et al., 1999] have addressed this problem through the use of various control strategies.

7 Acknowledgements

The author would like to express his gratitude to Dr. Charles Markham, Mr. Simon McLoughlin, Dr. Robert Shorten and the other members of Signals & Systems Group for their help in the many fruitful discussions leading up to this work. He would also like to acknowledged the contribution of Automative Distance Control GmBH., Germany, and specifically Mr. Jochen Franz for aquisition and provision of the test data used throughout this paper.

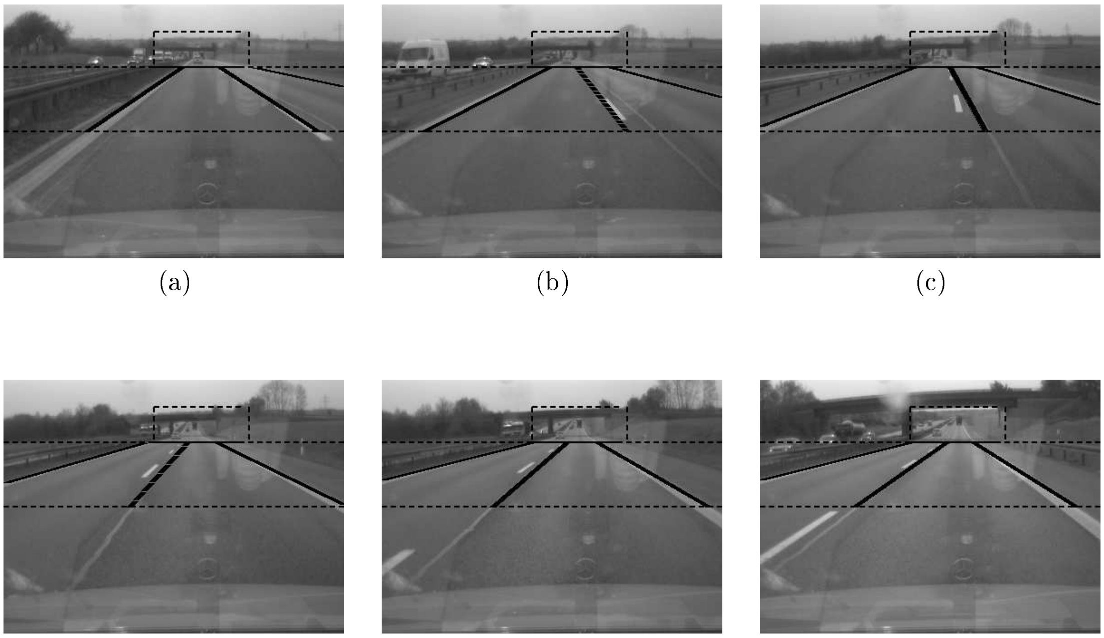

(a) (b) (c)

(d) (e) (f)

Figure 4: Lane change maneuver

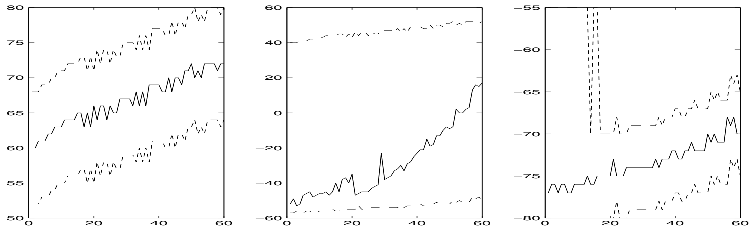

Figure 5: Theta values during lane change maneuver. The left, center, and right plots show the values for the left, center and right contours respectively. Also shown on each plot, is a dashed plot of the θ-windows.

References

[Barducci and Pippi, 1999] Barducci, A. and Pippi, I. (1999). Object recognition by edge analysis: a case study. *Optical Engineering*, 38(2):284–294.

[Bertozzi and Broggi, 1997] Bertozzi, M. and Broggi, A. (1997). Vision-based vehicle guidance. *IEEE Computer*, 30(7):49–55.

[Bertozzi et al., 2000] Bertozzi, M., Broggi, A., and Fascioli, A. (2000). Vision-based intelligent vehicles: State of the art and perspectives. *Robotics and Autonomous Systems*, 32:1–16.

[Deans, 1981] Deans, S. R. (1981). Hough transform from the Radon transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 3(2):185–188.

[Dickmanns, 1988] Dickmanns, E. D. (1988). Dynamic computer vision for mobile robot control. In *Proc. of 19th International Sympos. Expos. Robots*, pages 314–327, Sydney, Australia.

[Dickmanns and Mysliwetz, 1992] Dickmanns, E. D. and Mysliwetz, B. D. (1992). Recursive 3-d road and relative ego-state recognition. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 14(2):199–213.

[Duda and Hart, 1972] Duda, R. and Hart, P. (1972). Use of the Hough transform to detect lines and curves in pictures. *Communications of the ACM*, 15:11–15.

[Hofmann et al., 2000] Hofmann, U., Rieder, A., and Dickmanns, E. D. (2000). EMS-Vision: Application to Hybrid Adaptive Cruise Control. In *Procs. IEEE Intelligent Vehicles Symposium 2000*, pages 468–473, Detroit, USA.

[Hough, 1962] Hough, P. (1962). Method and means of recognizing complex patterns. U.S. Patent 3069654.

[Koller et al., 1999] Koller, D., Luong, Q.-T., Weber, J., and Malik, J. (1999). *Handbook of Pattern Recognition & Computer Vision*, chapter 4.6 Vision-Based Automatic Road Vehicle Guidance, pages 817–857. World Scientific Publishing Co. Ltd., 2 edition. eds. Chen C.H., Pau L.F., and Wang P.S.P.

[McDonald et al., 2001a] McDonald, J., Franz, J., and Shorten, R. (2001a). Application of the hough transform to lane following in motorway driving scenarios. In Shorten, R., Ward, T., and Lysaght, T., editors, to appear *Proc. of the Irish Signals and Systems Conference.*

[McDonald et al., 2001b] McDonald, J., Markham, C., and McLoughlin, S. (2001b). Selected problems in automated vehicle guidance. Technical Report NUIM/SS/01/05, Signals & Systems Group, Department of Computer Science, National University of Ireland, Maynooth, Ireland.

[McDonald and Vernon, 1998] McDonald, J. and Vernon, D. (1998). A new Hough transform for the detection of arbitrary 3-dimensional objects. In Vernon, D., editor, *Proc. of Optical Engineering Society of Ireland and Irish Machine Vision and Image Processing Joing Conference*, pages 243–255, NUI Maynooth, Ireland.

[Pomerlau, 1995] Pomerlau, D. (1995). Ralph: Rapidly adapting lateral position handler. In *Proc. of Intelligent Vehicles*, pages 54–59.

[Taylor et al., 1999] Taylor, C. J., Košecká, J., Blasi, R., and Malik, J. (1999). A comparitive study of vision-based lateral control strategies for autonomous highway driving. *International Journal of Robotics Research*, 18(5):442–453.

[Thorpe et al., 1998] Thorpe, C., Hebert, M. H., Kanade, T., and Shafer, S. A. (1998). Vision and navigation for the Carnegie-Mellon Navlab. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(3):362–373.