

Tracking Lane and Pavement Edges Using Deformable Templates

K. C. Kluge^a, C. M. Kreucher^b, S. Lakshmanan^b

^aUniversity of Michigan Artificial Intelligence Lab
1101 Beal Avenue, Ann Arbor, MI 48109-2110 USA

^bUniversity of Michigan-Dearborn Vehicle Electronics Lab
4901 Evergreen Road, Dearborn, MI 48128-1491 USA

ABSTRACT

Experiments with the LOIS (*Likelihood Of Image Shape*) Lane detector have demonstrated that the use of a deformable template approach allows robust detection of lane boundaries in visual images. The same algorithm has been applied to detect pavement edges in millimeter wave radar images. In addition to ground vehicle applications involving lane sensing, the algorithm is applicable to airplane applications for tracking runways in either visual or radar data. Previous work on LOIS has focused on the problem of detecting lane edges in individual frames. This paper describes extensions to the LOIS algorithm which allow it to smoothly track lane edges through maneuvers such as lane changes.

Keywords: deformable templates, object detection, Bayesian analysis, image sequences

1. INTRODUCTION

Over the last decade a great deal of research has been performed in the area of vision-based detection of lane boundaries. This technology has a number of potential applications. One high-value potential application is drowsy driver warning. There are over three million traffic accidents each year in the U.S. in which a vehicle leaves the roadway without colliding with another vehicle. These accidents result in 13,000 deaths annually, and are responsible for 16.5% of all traffic delays. It is estimated that up to 53% of these accidents could be avoided if vehicles had lane departure warning systems [1]. Another potential application is to enhance the accuracy of tracking the leading vehicle for intelligent cruise control. Loss of correct tracking when the leading vehicle enters a curve is a significant cause of problems for ICC systems, and methods for detecting curves based on estimates of the motion of either the leading vehicle or one's own vehicle have limitations [2]. A longer-term application is to provide autonomous lateral vehicle control [3]. Vision-based techniques for autonomous lateral control have the advantage of using existing visual cues in the road environment, compared to infrastructure-based methods that require modification of the road.

A distinction can be made between the problems of lane detection and lane tracking. Lane detection involves determining the location of the lane boundaries in a single image without strong prior knowledge regarding the lane position. Lane tracking involves determining the location of the lane boundaries in a sequence of consecutive images, using information about the lane location in previous images in the sequence to constrain the probable lane location in the current image. Some systems use different algorithms for lane detection and tracking. The VaMoRs system, for instance, uses the algorithm described in [4] to perform the initial detection of the road, then switches to the algorithm described in [5] to perform frame-to-frame tracking of the lane location. The approach taken in the work described in this paper is to use the same basic image processing for lane detection and tracking, the LOIS Lane Detector.

LOIS (for *Likelihood Of Image Shape*) uses a deformable template approach. A parametric family of shapes describes the set of all possible ways that the lane edges could appear in the image. A function is defined whose value is proportional to how well a particular set of lane shape parameters matches the pixel data in a specified image. Lane detection is performed by finding the lane shape parameters that maximizes the function for the current image. LOIS uses a weaker prior model of lane location when performing initial lane detection, then uses information from the previous frame to constrain the probable lane location when performing lane tracking. Previous articles on LOIS [6][7][8] have focused on locating the lane boundaries in single images in situations where the vehicle remained near the center of the lane. The more recent work reported here deals

with tracking the lane boundaries through sequences of images, including sequences where the vehicle performs maneuvers such as lane changes and excursions which take it far away from the lane center.

The next section reviews related work in the area of lane detection. This is followed by a detailed description of the LOIS algorithm. Modifications made to the basic algorithm to improve performance when tracking lane edges through image sequences are described, and results are shown. Finally, directions for further extensions to the research are described.

2. RELATED WORK

Many algorithms for lane detection are described in the literature. A number of systems make the assumption that the lane boundaries are straight lines, and detect them by extracting edge points from the image and using the Hough transform. An example of such a system is a lane detector developed at Honda [9]. The system developed at Matsushita [10] and later versions of the LANELOK system developed at General Motors [11] enforce a global scene constraint between the left and right lane edges (that they should meet at a vanishing point on the horizon row in the image plane). A system by Polk and Jain [12] and the SHIVA system [13] further extend these methods to better handle curved roads by dividing the image into a small number of horizontal sections and finding linear approximations to the lane edges within each section. None of these systems provides an explicit estimate of the reliability of the results it produces, and the use of a linear or piecewise-linear model of the lane boundaries limits the accuracy of the feature location in the image.

Two systems explicitly use curved models of the lane boundaries. The system developed at the University of Bristol [14] performs an intensity-based segmentation to extract regions corresponding to solid painted lines or sections of dashed painted lines. The system looks for the set of concentric arcs that define the lane structure. While the system handles curved lane boundaries and applies a global scene constraint to filter the results of the image segmentation, it assumes that there are painted stripes marking both edges of all the lanes, and provides no reliability estimate. ARCADE [15] uses a simple edge detection scheme in conjunction with the Least Median Squares robust estimator to find the road curvature and orientation, then does a simple 1-D segmentation of the intensity profile of the road to locate the offsets of the lane edges. One unique feature of ARCADE is that the road curvature and orientation estimation does not require any perceptual grouping of the extracted edge points into individual lane edges.

The GOLD algorithm [16] can detect curved lane boundaries, but places no restrictions on the shape of the lane other than assuming a constant lane width on a flat ground plane. The image of the scene is backprojected onto the ground plane. Simple 1-D image processing is performed on each row of the transformed image to identify possible lane markings, modeled as narrow bright features against a darker background. The lane width is identified by constructing a histogram of the horizontal separation between all pairs of potential lane edge points and finding the peak value. The lane is found by identifying the pairs of edge points in each row forming the longest continuous lane.

3. THE LOIS ALGORITHM

The deformable template approach to object detection has three components: a parametric family of shapes which describes the possible ways that the object can appear in the image, a likelihood function which measures how well a particular hypothesized object shape matches a particular image, and a method for finding the shape parameters which maximizes the likelihood function for the image being examined. Each of these components is described in detail below.

3.1 Defining a class of parametric lane shapes

Assume that the lane edges are circular arcs on a flat ground plane. For small to moderate curvatures, a circular arc with curvature k can be closely approximated by a parabola of the form

$$x = 0.5 * k * y^2 + m * y + b \quad (1)$$

The derivation of the class of corresponding curves in the image plane is given for the case of an untilted camera, but it can be shown that the same family of curves results when the camera is tilted. Assuming perspective projection, a pixel (r, c) in the image plane projects onto the point (x, y) on the ground plane according to the equations

$$x = c * cf * y \quad (2)$$

and

$$y = \frac{H}{r * rf} \quad (3)$$

where H is the camera height, rf is the height of a pixel on the focal plane divided by the focal length, and cf is the width of a pixel on the focal plane divided by the focal length. Substituting (2) and (3) into (1) and performing some simple algebraic manipulation results in the image plane curve

$$c = \frac{0.5 * k * H}{rf * cf * r} + \frac{b * rf * r}{H * cf} + \frac{m}{cf} \quad (4)$$

or, combining the ground plane and camera calibration parameters together,

$$c = k' / r + b' * r + vp \quad (5)$$

In the case of a tilted camera, the same family of curves results if the image coordinate system is defined so that row 0 is the horizon row. For left and right lane edges defined by concentric arcs the approximation is made that the arcs have equal curvature and equal tangential orientation where they intersect the X axis, so k' and vp will be equal for the left and right lane edges. As a result, the lane shape in an image can be defined by the four parameters k' , b'_{LEFT} , b'_{RIGHT} , and vp .

3.2 The likelihood function

The intuition underlying the likelihood function used by LOIS is that there should be a brightness gradient near every point along the lane edges. The larger the magnitude of that gradient, the more likely it is to correspond to a lane edge. Also, the closer the orientation of that gradient is to perpendicular to the lane edge, the more likely it is to correspond to a lane edge. This likelihood function operates on raw image gradient information without the need for explicit thresholding to select edge points. As a result, weak edges with consistent gradient orientations can support the correct lane shape hypothesis, while strong edges with inconsistent orientations (such as those resulting from shadows) do not distract LOIS from finding the correct lane shape.

More formally, define the penalty function

$$f(\alpha, x) = 1 / (1 + \alpha * x^2) \quad (6)$$

where α determines how fast $f(\alpha, x)$ decreases as x increases. Then the contribution of each pixel to the likelihood value is equal to

$$GradientMagnitude_{pixel} * f(\alpha_1, Dist_{pixel}) * f(\alpha_2, \cos(GradientOrientation_{pixel} - LaneEdgeTangent)) \quad (7)$$

where $Dist_{pixel}$ is the distance in columns from the closest lane edge (left or right), and $LaneEdgeTangent$ is the tangential orientation of the closest lane edge calculated for the pixel's row. The calculation of the likelihood function clips the portions of the image more than a specified distance from the hypothesized lane edges in order to increase the speed of the function. Also, lookup tables are used for $\cos()$ and the penalty function $f()$ in order to further increase the speed of the likelihood function calculation.

3.3 MAP Estimation and the Metropolis Algorithm

The prior and likelihood models are combined in a Bayesian framework, resulting in the lane detection problem being posed as finding the *Maximum A Posteriori* estimate of the lane shape parameters,

$$\begin{aligned}
 (k^*, b'_{LEFT}, b'_{RIGHT}, \nu p^*) &= \underset{k', b'_{LEFT}, b'_{RIGHT}, \nu p}{\operatorname{argmax}} P(k', b'_{LEFT}, b'_{RIGHT}, \nu p | \text{image intensity gradient field}) \\
 &= \underset{k', b'_{LEFT}, b'_{RIGHT}, \nu p}{\operatorname{argmax}} (\operatorname{atan}((b'_{RIGHT} - b'_{LEFT}) - 1) - \operatorname{atan}((b'_{RIGHT} - b'_{LEFT}) - 3)) * L(k', b'_{LEFT}, b'_{RIGHT}, \nu p)
 \end{aligned} \tag{8}$$

In general $(\operatorname{atan}((b'_{RIGHT} - b'_{LEFT}) - 1) - \operatorname{atan}((b'_{RIGHT} - b'_{LEFT}) - 3)) * L(k', b'_{LEFT}, b'_{RIGHT}, \nu p)$ is not a concave function, and often contains several local maxima. As a result, local optimization techniques such as gradient descent are not appropriate. The current implementation of LOIS uses the Metropolis algorithm with a geometric annealing schedule [17] to perform this maximization (see [6][7][8] for a more detailed description).

All the results shown in this paper were generated by running the Metropolis algorithm for 40 iterations. In each iteration a small step away from the current value is tested for each of the lane shape parameters. The initial temperature is 10.0, and the final temperature is 0.01.

3.4 Experimental Results Using LOIS

Figure 1 shows examples LOIS' lane detection ability under a variety of road and environmental conditions.




Figure 1. Examples of LOIS detecting the lanes correctly

We have tested the LOIS lane detection algorithm on a very large number of images. Shown in Figure 2 are the cumulative results of the center of the lane as determined by LOIS, on a data set of approximately 1,400 images acquired in sequence as a vehicle (an U.S. Army HMMWV) was being driven for over 45 miles of regular Michigan highways (the route is also shown in Figure 2, starting from Dearborn, MI and back and forth to Rockwood, MI). We determined that the standard deviation in offset with respect to the center of the lane (a combination of errors due to both LOIS and the driver) is close to 13cms.

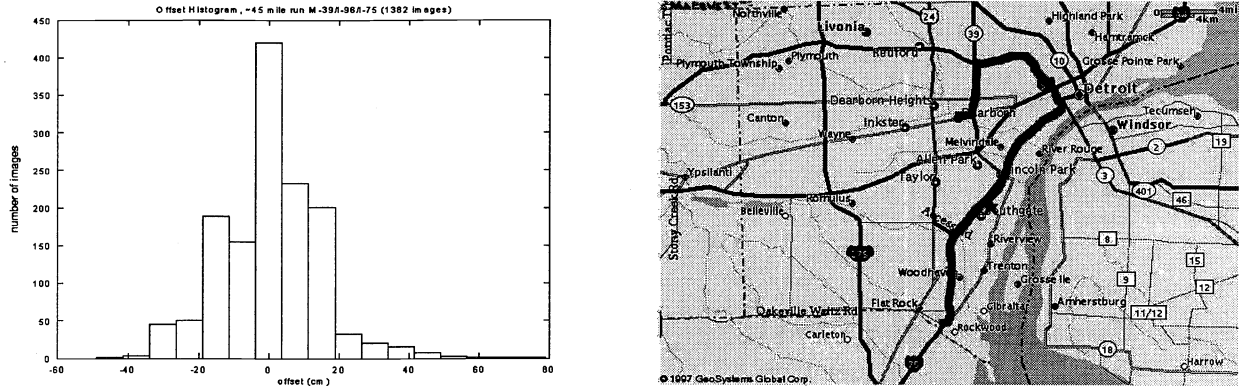


Figure 2. Results of lane detection using LOIS on a very large data set

Occasionally, as seen in Figure 2, the lanes detected by LOIS corresponds to a local maxima of the likelihood function, and the resulting detected lanes would have large errors compared to the true ones. This is an artifact of the Metropolis algorithm, which guarantees a convergence to the global maximum only if iterated infinite number of time [20] and not if iterated only a fixed number of times, as in LOIS. Figure 3 shows some typical examples of LOIS' failure to find the correct lane markers.




Figure 3. Examples of incorrect lane detection by LOIS

Shown subsequently in Figure 4, are lane detection results on the same set of images as Figure 3, but with the Metropolis algorithm replaced with an exhaustive search over the lane shape parameters.




Figure 4. The lanes detected by exhaustive search

Evidently, the global maximum (as found by exhaustive search) of the likelihood that LOIS uses is not the same as that obtained by using the Metropolis algorithm, and in each case the global maximum is closer to the "true" lane markers. One way to alleviate this difficulty is by starting the Metropolis algorithm with a "smarter" initial lane position.

4. LOIS-BASED LANE TRACKER

Clearly, roadway images taken in succession and spaced closely together (at a 10Hz rate, as for example in [16]) have very similar lane locations from frame to frame. This section describes how we exploit this temporal similarity when looking for lanes in individual images, especially as a means of increasing estimation accuracy of the standard single frame LOIS algorithm.

Specifically, the lane shape parameters k' , b'_{LEFT} , b'_{RIGHT} , and vp of the previous frame are used as a starting point of the Metropolis algorithm for the succeeding image. Our specific implementation of the Metropolis algorithm, remembers the likelihood values corresponding to points that are in the “neighborhood” of the current lane shape parameters. So, if the previous lane shape parameters correspond to a point “near” the peak of the current likelihood, then using the previous values of k' , b'_{LEFT} , b'_{RIGHT} , and vp as a starting point also results in a speed-up in the Metropolis algorithm’s rate of convergence to the peak of the current likelihood. The sequence of images shown in Figure 5, represent a vehicle keeping within its lane. In each individual still image, the location of the lane in the previous frame has been used as a starting point for lane searches in the current frame. Such a strategy seems to achieve the intended purpose.

A common malady among most lane tracking systems is finding lanes during a sequence in which a vehicle is making a lane change, as it represents a situation where temporal correlation is detrimental. This problem is further compounded by the fact that the *a priori* constraints on lane shapes that are commonly used become invalid during this maneuver - e.g., that the right lane is to the right of the vehicle or that the left lane is within a certain distance of the vehicle.

The lane tracking system described here overcomes this problem by a simple scheme. We describe the scheme via an example: Shown in Figure 6 is an image sequence as a vehicle performs a lane change maneuver. Notice that as the right lane marker is being crossed, the previous frame’s right lane becomes the left lane of the current frame. The LOIS-based lane tracker uses the previous estimate of the right lane location as a starting point for estimating the current left lane. Furthermore, the location of the new right lane is now unknown, and there is no temporal information of its location. By explicitly recognizing this scenario, the LOIS-based lane tracking system acts like a standard LOIS lane detection system as far as the right lane is concerned. The net effect is an accurate tracking of the lane makers by our LOIS-based algorithm even when the vehicle makes a lane change.

One danger of using past lane estimates to influence the current lane finding procedure is that bad estimates in the past may confine the current search to a poor solution space. Without proper consideration, the lane tracker could lose the lanes in one frame and never be able to recapture them in subsequent frames. This, however, is not a concern in the LOIS-based lane tracker, because the past lane shape estimate is used only as an initial guess of the current lane locations, without any additional constraints. We have found that this allows the lane tracker to accurately follow the lanes when it has a good estimate while simultaneously allowing it to discard erroneous starting estimates in the search for lanes in the current image when the initial estimate is very poor. Shown in Figure 7 is a brief three image sequence along with the corresponding LOIS-lane tracker outputs. Notice how the lane tracking “survives” a poor lane shape estimate (as seen in the middle image of Figure 6) caused by a severe glare saturating the imaging sensor. Using a Kalman filter to predict the lane shape parameters in the current frame, based on the past lane shape parameter estimates, and penalizing the likelihood for deviations of the current lane shape parameters from the predicted one could have also conceivably overcome this problem.




Figure 5. Lane tracking when there is no lane change




Figure 6. Lane tracking through a lane change maneuver




Figure 7. The LOIS-based lane tracker recovering from a “bad” lane shape estimate

5. DISCUSSION

The LOIS-based lane tracker described in this paper can be used to provide data (lane shape parameters as a function of time) for a subsequent driver warning or vehicle steering system. Indeed, a forthcoming paper [19] describes a lane departure warning system that is built using the lane tracking method described herein in conjunction with a Kalman filter. We have found the lane shape parameters determined by the LOIS-based tracker to be reliable/robust enough for such a driver warning system. Also, the same method employed in this paper can be used to track pavement or runway edges in radar image sequences acquired via an alternative (all-weather) millimeter-wave modality – see [18][21][22] for related works.

6. REFERENCES

1. M. Chen, T. Jochem, and D. Pomerleau, “AURORA: A Vision-Based Roadway Departure System,” in *Proceedings, IEEE Conference on Intelligent Robots and Systems*, vol. 1, pp. 243-248, August 1995.
2. M. Baret, S. Baillarin, C. Calesse, and L. Martin, “Sensor Fusion: Lane Marking Detection and Autonomous Cruise Control System,” *Collision Avoidance and Automated Traffic Management (Proc. SPIE vol. 2592)*, pp. 150-162, 1995.
3. D. Pomerleau and Todd Jochem, “Rapidly Adapting Machine Vision for Automated Vehicle Steering,” *IEEE Expert*, 11, (2), pp. 19-27, April 1996.
4. R. Behringer, “Road Recognition From Multifocal Vision,” in *Proceedings of the Intelligent Vehicles '94 Symposium*, pp. 302-7, October 1994.
5. E. D. Dickmanns and V. Graefe, “Applications of Dynamic Monocular Machine Vision,” *Machine Vision and Applications*, 1, pp. 241-261, 1988.
6. K. Kluge and S. Lakshmanan, “A Deformable Template Approach to Lane Detection,” in *Proceedings of the Intelligent Vehicles '95 Symposium*, pp. 54-59, 1995.
7. S. Lakshmanan and K. Kluge, “LOIS: A Real-Time Lane Detection Algorithm,” in *Proceedings of the 30th Annual Conference on Information Sciences and Systems*, Princeton, NJ, March 1996.
8. K. Kluge and S. Lakshmanan, “Lane Boundary Detection Using Deformable Templates: Effects of Image Subsampling on Detected Lane Edges,” in *Recent Developments in Computer Vision*, S. Z. Li, D. P. Mital, E. K. Teoh, and H. Wang, eds., Springer-Verlag, 1996.
9. K. Hashimoto, S. Nakayama, T. Saito, N. Oono, S. Ishida, K. Unoura, J. Ishii, and Y. Okada, “An Image-Processing Architecture and a Motion Control Method for an Autonomous Vehicle,” in *Proceedings of the Intelligent Vehicles '92 Symposium*, pp. 213-18, June 1992.
10. A. Suzuki, N. Yasui, N. Nakano, and M. Kaneko, “Lane Recognition System for Guiding of Autonomous Vehicle,” in *Proceedings of the Intelligent Vehicles '92 Symposium*, pp. 196-201, June 1992.
11. S. K. Kenue, “LANELOCK: Detection of Lane boundaries and Vehicle Tracking Using Image- Processing Techniques -- Parts I and II,” in *Proceedings, SPIE Mobile Robots IV*, pp. 221-244, 1989.

12. A. Polk, and R. Jain, "A Parallel Architecture for Curvature-Based Road Scene Classification," in *Roundtable Discussion on Vision- Based Vehicle Guidance '90 (in conjunction with IROS)*, July 1990.
13. K. C. Kluge, *YARF: An Open-Ended Framework for Robot Road Following*, PhD thesis, Carnegie Mellon University, 1993.
14. L. T. Schaaser, and B. T. Thomas, "Finding Road Lane Boundaries for Vision Guided Vehicle Navigation," in *Vision-Based Vehicle Guidance*, Ichiro Masaki ,ed, Springer-Verlag, 1992, Chapter 11.
15. K. C. Kluge, "Extracting Road Curvature and Orientation From Image Edge Points Without Perceptual Grouping Into Features," in *Proceedings of the Intelligent Vehicles '94 Symposium*, pp. 109-114, October 1994.
16. M. Bertozzi and A. Broggi, "GOLD: A parallel real-time stereo vision system for generic obstacle and lane detection." *IEEE Transactions on Image Processing*, vol. 7, pp. 62-80, Jan. 1998.
17. P.N. Strenski and S. Kirkpatrick, "Analysis of Finite Length Annealing Schedules." *Algorithmica*, vol. 6, pp. 346-366, 1991.
18. S. Lakshmanan, K. Kaliyaperumal, and K. C. Kluge, "LEXLUTHER: An algorithm for detecting roads and obstacles in radar images," in *Proceedings of the 1st IEEE Intelligent Transportation Systems Conference*, Nov. 1997.
19. C. Kreucher, S. Lakshmanan, and K. C. Kluge, "A Driver Warning System Based on the LOIS Lane Detection Algorithm." In preparation Mar. 1998.
20. B. Gidas, "Nonstationary Markov Chains and Convergence of the Annealing Algorithm." *J. Statistical Physics*, Vol. 39, pp. 73-131, 1985.
21. B. Ma, S. Lakshmanan, and A. O. Hero, "Detecting Curved Roads in Radar Images Using Deformable Templates," in *Proceedings of the 4th IEEE International Conference on Image Processing*, Oct. 1997.
22. S. Lakshmanan, A. K. Jain, and Y. Zhong, "Detecting Straight Edges in Millimeter Wave Images," in *Proceedings of the 4th IEEE International Conference on Image Processing*, Oct. 1997.