Simultaneous Detection of Lane and Pavement Boundaries Using Model-Based Multisensor Fusion

Bing Ma, Sridhar Lakshmanan, Alfred O. Hero

Abstract—**This paper treats a problem arising in the design of intelligent vehicles: automated detection of lane and pavement boundaries using forward-looking optical and radar imaging sensors mounted on an automobile. In previous work, lane and pavement boundaries have always been located separately. This separate detection strategy is problematic in situations when either the optical or the radar image is too noisy.**

In this paper, we propose a Bayesian multisensor image fusion method to solve our boundary detection problem. This method makes use of a deformable template model to globally describe the boundaries of interest. The optical and radar imaging processes are described with random field likelihoods. The multisensor fusion boundary detection problem is reformulated as a joint MAP estimation problem. However, the joint MAP estimate is intractable, as it involves the computation of a notoriously difficult normalization constant, also known as the partition function. Therefore we settle for the so-called empirical MAP estimate, as an approximation to the true MAP estimate. Several experimental results are provided to demonstrate the efficacy of the empirical MAP estimation method in simultaneously detecting lane and pavement boundaries. Fusion of multimodal images is not only of interest to the intelligent vehicles community, but to others as well, such as biomedicine, remote sensing, target recognition. The method presented in this paper is also applicable to image fusion problems in these other areas.

I. Introduction

This paper treats an important problem concerning navigation of intelligent vehicles: automated detection of lane and pavement boundaries using forward-looking optical and millimeter-wave radar imaging sensors mounted on an automobile. This boundary detection problem is particularly difficult when no prior knowledge of the road geometry is available (such as from previous time instants – see [1–3]), and when the detection algorithms have to locate the boundaries even in situations where there may be a great deal of clutter in the images. The ability to detect pavement and lane boundaries in the radar and optical images of a road scenario is an enabling or enhancing technique for a number of driver assisting applications, such as road departure or lane excursion warning, intelligent cruise control and autonomous driving. All of these applications have potential use in both military and civilian contexts.

Many gradient-based detection algorithms, which are applicable for structured edges including lane and pavement boundaries, apply a threshold to the image gradient magnitude to detect edges – see [4–7] and the references therein. When a visual image has clearly visible lanes, and when the radar image has uniform regions with good separation between the regions, good performance can be obtained with these algorithms. However, real road scenes seldom give rise to such clean images — there may be a great deal of clutter in the images, radar images have extremely low signal-to-noise ratio, or there are spurious boundaries in the images.

Needless to say, it is difficult to select thresholds which eliminate noise edges without also eliminating many of the edge points of interest, and so the conventional edge detection algorithms described in [4–7] are not suitable for our boundary detection problem under the above mitigating conditions.

A class of successful methods that overcome the thresholding problem are studied in [7–11]. These methods work directly with the image intensity array, as opposed to separately detected edge points, and use a global model of lane and pavement boundary shape. Two examples from this class are particularly relevant to this paper:

1. Reference [9] presents a vision-based real-time algorithm called LOIS for locating lane and pavement boundaries using a deformable template global shape model. The global shape model adaptively adjusts and aligns a template so that it best matches the underlying features of the lane and pavement-boundary over the entire image. At the heart of LOIS is a matching function that encodes the knowledge that the edges of the lane should be near intensity gradients whose orientation are perpendicular to the lane edge. This allows strong magnitude gradients to be discounted if they are improperly oriented and weak magnitude gradients to be boosted if they are properly oriented. LOIS is shown to work well under a wide variety of conditions, including cases with strong mottled shadows and broken or interrupted lane markings, which pose a challenge for gradient-based lane detection schemes.
2. Reference [11] presents a method for detecting pavement boundaries in radar images. Like in LOIS, here too a deformable template model is used. The biggest difference though is in the matching function: [11] uses the log-normal probability model of the radar imaging process, which was proposed in [8]. This function encodes the knowledge that boundaries of the pavement should divide the image into three "relatively" homogeneous regions.

In both references [9] and [11] the boundary detection problem on hand is reformulated as a Bayesian estimation problem, where the deformable template model plays the role of a prior probability density function(pdf) and the matching

Bing Ma and Alfred O. Hero are with the Department of Electrical Engineering and Computer Science, the University of Michigan, Ann Arbor, MI 48109. E-mail: bingm, hero@eecs.umich.edu .

Sridhar Lakshmanan is with the Department of Electrical and Computer Engineering, the University of Michigan, Dearborn, MI 48128. E-mail: lakshman@umich.edu .

function plays the role of a likelihood pdf, respectively.

Note that previously lane boundary detection in optical images [5–7, 9, 10], and pavement boundary detection in radar images [8, 11–13] have always been studied separately. However, a single sensor, either optical or radar sensor, limits itself in the ability to sense and identify the relevant features in varying environments. For example, the optical sensor is not able to operate in a poorly illuminated environment, while the radar sensor can not distinguish the lane markers on the road. To take advantage of the strengths (and overcome the weaknesses) of both the optical and radar sensors, it is natural to think of combining the two different types of sensed data together since multiple sensors will provide more information and hence a better and more precise interpretation of the sensed environment.

In this paper, we present a method for fusing the two multi-modal images — vision and radar. The problem of interest is simultaneous detection of lane and pavement boundaries using images from both the modalities. Consideration of both the modalities is meaningful because lane and pavement boundaries for the same road scene are highly correlated. The lane and pavement boundary detection problem is re-formulated in a Bayesian framework in this paper. The lane and pavement boundary shapes are represented by a pair of deformable template models, one in the ground plane and the other in the image plane, respectively. And each model has an associated set of deformation parameters. However, the model parameters are correlated. A statistical distribution that expresses the variability of and correlation between these parameters constitutes the prior pdf for the Bayesian reformulation. A pair of likelihood pdfs (matching functions) provide a relative measure of how well the deformed templates of the lane and pavement boundaries match the observed vision and radar images of a road scene. Fusion of the multi-modal information is realized by obtaining the joint maximum *a posteriori* (MAP) estimate – see Figs. 7–10. The same framework can be extended to fuse data from other modalities such as from infra-red, ultrasound or single-beam radar sensors. We refer the reader to [14, 15] for examples.

Computing the joint MAP estimate involves the relative weighting of the matching functions, corresponding to the two likelihood pdfs. Without a proper choice of this weighting either the vision or the radar matching function will always dominate the other. A theoretically correct choice of weighting involves a normalizing constant, whose determination is computationally prohibitive. Instead we adopt the so called empirical MAP (Bayes') estimate [16,17]. The two matching functions are linearly combined with a fixed weighting. This fixed weight is experimentally derived by examining the individual matching functions for a (training) set of vision and radar image pairs. The weight is chosen so that the scales of variation for the individual matching function contributions are approximately the same over the entire training set.

Several experimental results are provided to demonstrate the efficacy of this empirical MAP estimation method in simultaneously detecting lane and pavement boundaries. The information gained by fusing data from both modalities is evident when these results are compared to those obtained by processing each data modality separately. The fused boundary detection results are also compared to those obtained when no likelihood weighting is used, to illustrate the advantage of weighting.

The rest of this paper is organized as follows: Section II contains a detailed discussion on multisensor image fusion, especially as it applies to the lane and pavement boundary detection problem. In Section III, the deformable template global shape models for both lane and pavement boundaries are described, along with the prior pdf over the template shape parameters. Sections IV and V treat the radar and optical imaging likelihoods, respectively. Section VI presents the Bayesian reformulation of the problem, along with a description of the joint MAP method used to implement the multisensor image fusion. Section VII presents the empirical MAP estimate. Sections VIII presents the computational issue of the empirical MAP estimate and Section IX gives experimental results that illustrate the performance of the proposed data fusion method. The paper concludes with a detailed discussion of several relevant issues, including possible future work, in Section X.

II. Multisensor Image Fusion

There has been a growing interest over the past few years in the use of multiple sensors to increase the capabilities of intelligent machines and systems. Multisensor fusion deals with integration of information from several different sources, aiming at an improved quality of results, e.g., better decisions, robust behavior of algorithms, or improved classification accuracies. In other words, fused images can significantly enhance system interpretation capabilities and reliabilities. The major application areas for image fusion include biomedicine [18, 19], remote sensing [20–22], automatic target recognition [23, 24], and autonomous landing guidance [25].

The intelligent vehicles community has been very active in the image fusion area as well [14, 26, 27]. The problem of co-registering vehicle navigation data (proximity of the vehicle ahead, map of the scene ahead, host vehicle position, speed and heading direction, etc.) with concurrently obtained visual images is of much interest to this community. The use of this co-registered information for affecting various driver warning and vehicle control tasks is also of interest. The problem at hand, namely, the simultaneous detection of lane and pavement boundaries from vision and radar images of the same road scene has not been considered before. A related problem of simultaneous detection of lanes and obstacles from vision and radar images of the same road scene has indeed received much attention recently [27]. However, the type of radar used, the objects of interest, and the estimation methods employed are all distinctly different from the ones in this paper.

The radar data and visual images used in this work were all obtained from an imaging platform mounted atop WOLVERINE I (*Wheeled Onroad Lab Vehicle Enabling Re-*

search Into New Environments), a self-contained test-bed vehicle (Fig. 1). The imaging platform uses a 77GHz frequency modulated continuous wave (FMCW) radar sensor, with a maximum range of 128 meters (resolution 0.5 meters) and angular field of view of 64° (resolution 1°). The imaging platform is also equipped with a video camera that is bore-sighted with the radar and captures a visual image of the same road scene – see Figs. 1 and 2 for an illustration and the setup. The radar and optical sensors sense the same road scenario simultaneously in order to acquire pair of co-registered images of that scene.

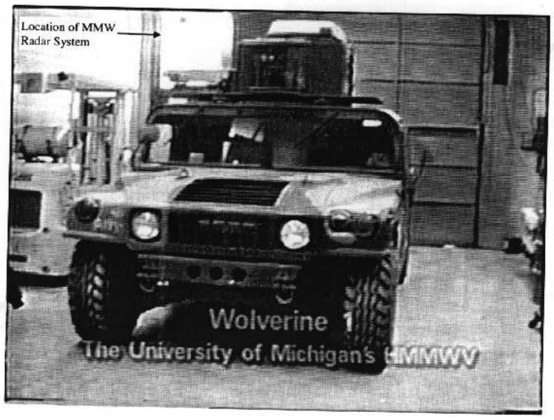

(a) WOLVERINE I

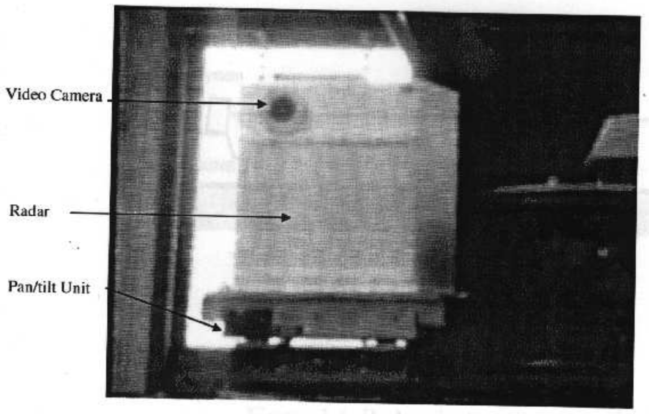

(b) Imaging platform

Fig. 1. Configuration of the test-bed vehicle

The information contained in the optical image depends on the reflectivity of the road illuminated by natural visible light. The lane boundaries, i.e., the white or yellow lane markers, constitute one of the two boundaries of interest to us and are clearly visible in the optical image. The radar image is obtained by illuminating the road scene with electromagnetic radiation in the millimeter-wave spectrum – see [8] and [13] for a detailed discussion of this image acquisition process. The relatively smooth road surface forward scatters much of this incident electromagnetic power and hence returns very little power back to the radar; the side of the road, because it is made up of a coarser structure than the road, returns a slightly higher amount of power. Thus, in the radar image, the road region is a little darker than

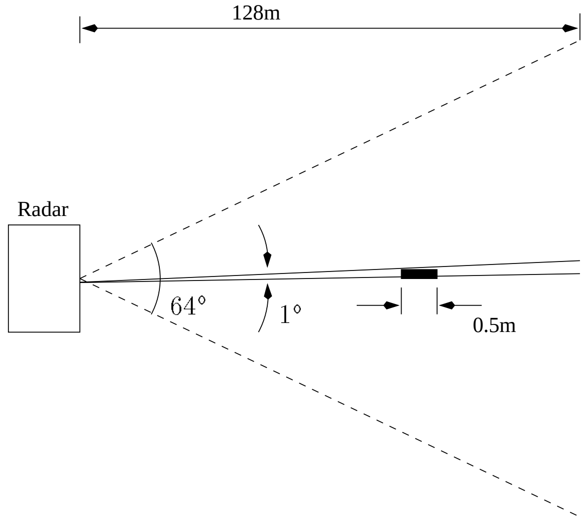

(a) Top view

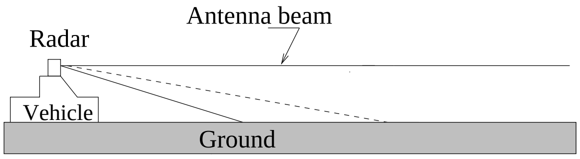

(b) Side view

Fig. 2. Radar setup

the road sides, and the boundaries of the road delineate these three regions of different contrast.

Evidently, either optical or radar sensor used alone has limited capabilities for resolving ambiguities and providing consistent descriptions of the road scenario due to the operating range and limitation which characterize the sensor. The optical sensor can provide high signal-to-noise ratio images in a well-illuminated environment such as a sunny day. Such images, which clearly reflect the lane information, are sufficient for the lane boundary detection task. However, since it is a passive sensor and works at visible light wavelengths, in an ill-illuminated environment, e.g., at night or in foggy weather, the optical sensor will fail to provide sufficient information about the lane boundary. The radar sensor, on the contrary, being an active sensor and operating at millimeter wavelengths, has the ability to penetrate through rain, snow, fog, darkness, etc., i.e., it can operate under all weather conditions and provides an “alternate” image of the scenario in front of the vehicle. Thus the radar image, regardless the illumination situations, can give us the pavement boundary information, and the precise geometry of the pavement boundaries can be subsequently used in a number of driver warning and vehicle control tasks. The downside of the radar image though is its notoriously poor signal-to-noise ratio and low spatial resolution when compared to a visual image of the same road scene.

Since the optical and radar sensors provide different but complementary information about the road scene ahead of the vehicle, if the two types of information are combined appropriately and efficiently, the accuracy of the detected lane and pavement boundaries can be improved. The optical and radar fusion system shown in Fig. 3 exploits this redundancy, diversity and complementarity between the two modalities.

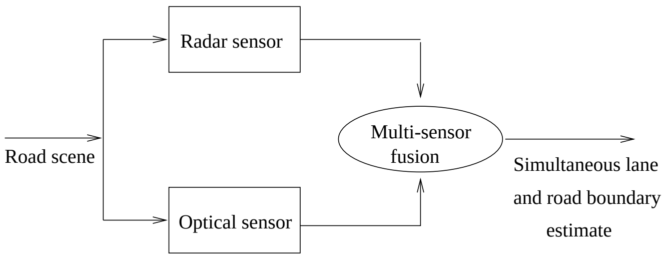

Fig. 3. Diagram of the optical and radar multisensor fusion

Redundancy is caused by the use of two sensors to measure the same entity. It is well known that redundancy reduces uncertainty. This can be appreciated from the fact that for two sensors, the signal related to the measured quantity is often correlated, whereas the uncertainty associated with each individual sensor tends to be uncorrelated. To be specific, the lane and pavement boundaries in the visual and radar images, respectively, are highly correlated since they are concentric arcs on the ground-plane and the lane boundaries are inside the road region, while the degradations of this boundary information introduced by the visual and radar imaging processes have completely different characteristics.

Since the vision and radar sensors measure the same scene with different laws of physics, we obtain physical sensor diversity. Another diversity, spatial diversity, which offers different viewpoints of the sensed environment simply by having sensors in different locations, also plays a very important role in multisensor fusion. In our case, since both sensors are placed at almost the same location, spatial diversity is not an issue in our work. However, it is because of the absence of spatial diversity that we can greatly simplify the registration process, which is usually very complicated.

Since the imaging system is made up of sensors each of which observes a subset of the environment space (the optical sensor offers information about the lane boundaries, while the radar sensor offers information about the pavement boundaries), and the union of these subsets makes up the whole road scenario, we achieve data complementarity.

In order to realize the advantage of multisensor systems, information from multiple sources needs to be effectively combined in a coherent and efficient manner to compensate for the individual limitations and deficiencies of sensors. According to the correlation level among data from different sensors, the fusion methods can be classified into two major classes — weakly coupled and strongly coupled data fusion [28]. In weak coupling, the outputs of two or more algorithms that produce information independently are combined. The basic approach of weakly coupled sensory fusion has been to pool the information using what is essentially "weighted averaging" techniques of varying degrees of complexity. In strong coupling, on the other hand, the operation of one sensory module is affected by the output of a second sensory module, so that the outputs of the two modules are no longer independent. Since the lane and pavement boundaries are concentric and the lanes are restricted inside the pavement region, their geometric shapes are highly related. Specifically, the curvature and orientation of the lane and pavement boundaries are nearly identical in the ground plane. Also, the left and right lane markers have to necessarily fall in the pavement region. Thus the problem of simultaneous estimation of the lane and pavement boundaries falls into the strongly coupled fusion category.

III. Deformable Template Models of Lane and Pavement Boundaries

In most cases, we can assume that *a priori* knowledge regarding the shape of the lane and pavement boundaries in the optical and radar images is available. A commonly used shape model for lane and pavement boundaries assumes that they can be approximated by concentric circular arcs on a flat ground plane. Such arcs, at least within a reasonable field of view, for small-to-moderate curvatures, are well approximated by parabolic curves,

$$x = \frac{1}{2}ky^2 + my + b \tag{1}$$

where the parameter k is the curvature of the arc, m is the tangential orientation, and b is the offset. While the radius of curvature and tangential orientation of left and right lane/pavement boundaries will differ slightly, constraining the left and right lane/pavement boundaries to have the same parameters k and m closely approximates the actual edge shapes for all but very small radii of curvature. So, it is assumed that concentric lane and pavement boundaries share the same parameters k and m.

The radar image is composed of reflections from the ground, and its domain is indeed the ground plane. So, (1) can be directly applied to model the shape of pavement boundaries in the radar image. The domain of the visual image, however, is a perspective projection of the ground plane, and hence (1) needs to be rewritten in order to model the shape of lane boundaries in the image plane. Assuming a tilted pinhole camera perspective projection model, parabolic curves in the ground plane (1) transform into hyperbolic curves in the image plane:[1]

$$c = \frac{k'}{r - hz} + b'(r - hz) + vp' \tag{2}$$

where

$$\begin{aligned} k' &= \eta_k \; k, \\ vp' &= \eta_m \; m + \eta_{m,k} \; k + \eta, \text{ and} \\ b' &= \eta_b \; b + \eta_{b,m} \; m + \eta_{b,k} \; k. \end{aligned} \tag{3}$$

[1] See [6] for a derivation

In other words, the k' parameter is linearly proportional to the curvature of the arc on the ground plane. The vp parameter is a function of the tangential orientation of the arc on the ground plane, with some coupling to the arc curvature as well. The b' parameter is a function of the offset of the arc from the camera on the ground plane, with couplings to arc curvature and tangential orientation. The constants $\eta_k, \eta_m, \eta_{m,k}, \eta, \eta_b, \eta_{b,m}$, and $\eta_{b,k}$ depend on the camera geometry (resolution, focal length, height of the camera from the ground plane, and camera tilt).

Let $\underline{\theta}^o = \{k', vp', b'_L, b'_R\}$ and $\underline{\theta}^r = \{k, m, b_L, b_R\}$ denote the unknown lane and pavement boundaries' parameters, respectively. Let $\underline{\theta} = \{\underline{\theta}^r, \underline{\theta}^o\}$ denote their adjoinment. By changing the values of $\underline{\theta}$, various lane and pavement boundary shapes can be realized – see Fig. 4.

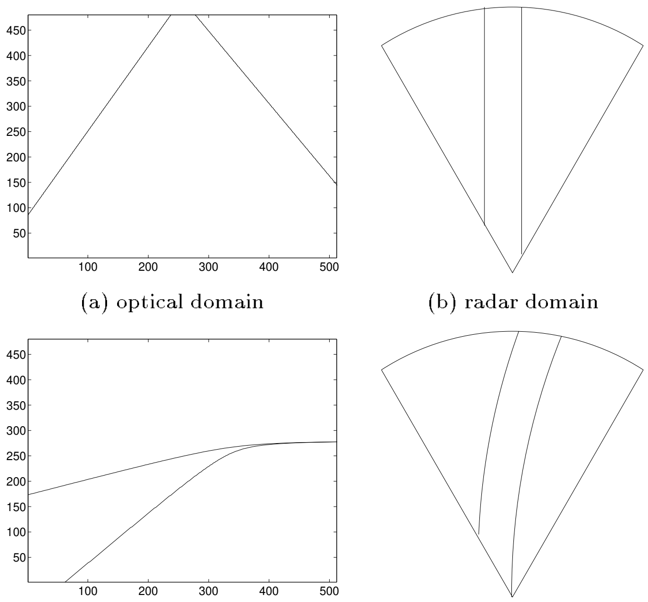

(a) optical domain (b) radar domain

(c) optical domain (d) radar domain

Fig. 4. Boundary templates for the optical and radar images

The templates of the upper row in Fig. 4 illustrate a straight road scenario with the deformation parameters $k' = 0, vp' = 256, b'_L = -0.6, b'_R = 0.7, hz = -30, b_L = -12.0, b_R = 2.4$. The templates of the lower row demonstrate a curved road scene with $k' = 600, vp' = 340, b'_L = -3.2, b'_R = -1.0, hz = 200, b_L = -14.2, b_R = 0.4$.

The problem of simultaneous detection of the lane and pavement boundaries is now equivalent to the problem of estimating $\underline{\theta}$. The elements of $\underline{\theta}$ have to satisfy some constraints, and for some elements of $\underline{\theta}$ the range of physically meaningful values they can possibly assume is known *a priori*. Given a hypothetical $\underline{\theta}$, its fidelity to the observed optical and radar images can also be assessed. In this paper, we choose a probabilistic Bayesian framework to express the constraints, the *a priori* beliefs, and the assessment of fidelity to data.

We present the so-called prior pdf here (the likelihood pdfs are in the next section):

$$
\begin{aligned}
& p(\underline{\theta}) = p(\underline{\theta}^o, \underline{\theta}^r) \\
&= I_{b'_L > b_L}(b'_L, b_L) \times I_{b'_R < b_R}(b'_R, b_R) \times \delta(k' - \eta_k\ k) \\
&\quad \times \delta(vp' - [\eta_m\, m + \eta_{m,k}\, k + \eta]) \\
&\quad \times \frac{2}{\pi}\,\mathrm{atan}[\beta_{o,1} \times (b'_R - b'_L)] \times \frac{2}{\pi}\,\mathrm{atan}\left[\frac{\beta_{o,2}}{b'_R - b'_L}\right] \\
&\quad \times \frac{2}{\pi}\,\mathrm{atan}[\beta_{r,1} \times (b_R - b_L)] \times \frac{2}{\pi}\,\mathrm{atan}\left[\frac{\beta_{r,2}}{b_R - b_L}\right] \quad (4)
\end{aligned}
$$

where $I_A(x, y)$ is an indicator function,

$$
I_A(\{x, y\}) = \begin{cases} 1, & \text{if } (x, y) \text{ satisfies relation } A \\ 0, & \text{otherwise} \end{cases}
$$

and $\delta(x)$ is the Kronecker delta function,

$$
\delta(x) = \begin{cases} 1, & \text{if } x = 0 \\ 0, & \text{otherwise} \end{cases}
$$

The terms on the first two lines of (4)'s RHS, correspond to the constraints that the elements of $\underline{\theta}$ have to satisfy. The first two terms impose the constraint that the lane markers be contained within the pavement region, the last two terms impose the constraint that the lane boundaries' curvature and orientation be precisely related to the pavement boundaries' curvature and offset via (3). The terms on the last two lines of (4)'s RHS expresses the *a priori* beliefs that lanes and pavements can be neither be too narrow nor too wide.

IV. Radar Imaging Likelihood

The true boundaries of the pavement separate the observed radar image into three regions associated with the road surface, the left side of the road, and the right side of the road. Given a $\underline{\theta}^r$, its fidelity to the observed radar image is assessed by how homogeneous the corresponding three (road, left, and right) regions are. A log normal probability law [29] is used to derive the homogeneity criteria.

We caution the reader that the radar returns over the the left, right and road regions are not truly homogeneous. Sources for non-homogeneity including point-like scatterers (cars, sign posts, retro-reflectors, etc.), changes in the off-road surface (grass, gravel, etc.) and presence of periodic off-road structures (bridges, fences, trees, etc.). Modeling all such variability is impossible. The log-normal pdf is meant to be a reasonable but low-complexity (two parameters per region) approximation to the actual variations in the data.

The rationale for using the log-normal law, as opposed to normal, exponential, or Rayleigh laws, is due to previous studies [30–32]. To appreciate the appropriateness of the log-normal pdf to describe radar returns, we refer the reader to Fig. 5, which demonstrates that log-normal is an excellent approximation to the radar return.

Let $\mathcal{L} = \{(r, \phi),\ 1 \le r \le r_{max},\ \phi_{min} \le \phi \le \phi_{max}\}$ denote the range and azimuth coordinates of the pixels in the millimeter-wave radar image Z^r. Z^r is a random field

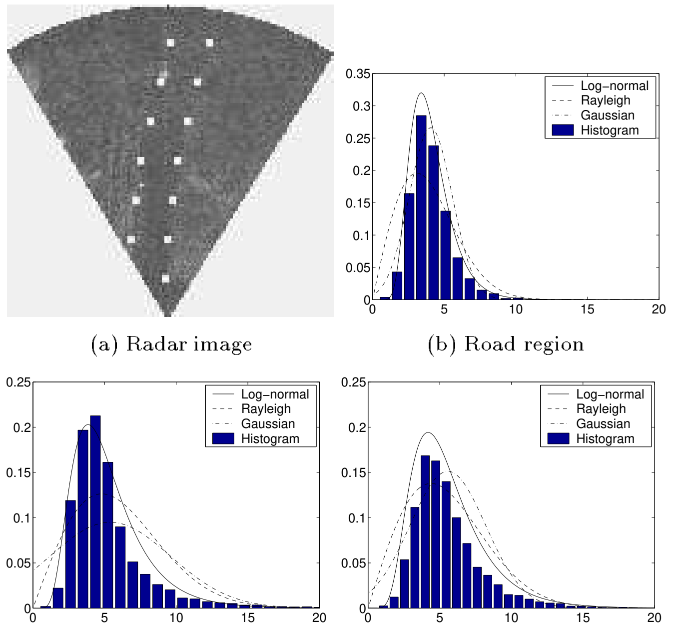

(a) Radar image (b) Road region

(c) Left-side region (d) Right-side region

Fig. 5. Rationale for using log-normal pdf. (a) shows a radar image overlayed with the correct positions for the left and right pavement boundaries. (b), (c), and (d) show histograms of the actual radar returns for the three regions – the road, the left-side and right-side of the road. Also shown in (b), (c), and (d) are the maximum likelihood fits of the log-normal, Rayleigh, and Gaussian pdf's to the radar return histograms.

representing the power of radar returns over the domain $\mathcal{L}$. Given the pavement boundaries, it is assumed that the conditional probability of Z^r taking on a realization z^r (corresponding to a single observation) is given by the log-normal pdf

$$p(z^r \mid \underline{\theta}^r) = \prod_{(r,\phi)\in\mathcal{L}} \frac{1}{z^r_{r\phi}\sqrt{2\pi\sigma^2_{r\phi}(\underline{\theta}^r)}} \times \exp\left\{-\frac{1}{2\sigma^2_{r\phi}(\underline{\theta}^r)}\left[\log z^r_{r\phi} - \mu_{r\phi}(\underline{\theta}^r)\right]^2\right\} \quad (5)$$

where $\mu_{r\phi}(\underline{\theta}^r)$, $\sigma^2_{r\phi}(\underline{\theta}^r)$ denote the mean and variance of the region to which the pixel (r,ϕ) belongs. In (5) $\underline{\theta}^r$ is used to emphasize the dependencies of means $\mu_{r\phi}$ and variances $\sigma^2_{r\phi}$ on the unknown parameters $\underline{\theta}^r$. However, henceforth, in order to make the representations concise, we will omit the explicit references to $\underline{\theta}^r$.

In its present form, (5) is not very useful for assessing the data fidelity of $\underline{\theta}^r$, due to the presence of nuisance parameters $\mu_{r\phi}$ and $\sigma^2_{r\phi}$. In the sequel, (5) is rewritten as a function of these nuisance parameters, and it reveals an intuitively appealing homogeneity criterion. Some additional notations are necessary:

- Let N denote the number of pixels in $\mathcal{L}$
- Also let $\mathcal{L}^{rd}$, $\mathcal{L}^{lt}$, $\mathcal{L}^{rt}$ denote the Cartesian coordinates of the pixels in the road, left-side and right-side regions, respectively, and $(\mu^{rd}, [\sigma^{rd}]^2, N^{rd})$, $(\mu^{lt}, [\sigma^{lt}]^2, N^{lt})$ and $(\mu^{rt}, [\sigma^{rt}]^2, N^{rt})$ denote the means, variances and the numbers of pixels of the corresponding regions.
- Define

$$\begin{aligned}
\underline{\mu} &= \left[\mu^{rd} \quad \mu^{lt} \quad \mu^{rt}\right]^T, \\
\underline{\sigma}^2 &= \left[[\sigma^{rd}]^2 \quad [\sigma^{lt}]^2 \quad [\sigma^{rt}]^2\right]^T, \\
\log \underline{Z}^r &= \left[\log z^r_{1\phi_{min}}, \cdots, \log z^r_{1\phi_{max}}, \log z^r_{2\phi_{min}}, \cdots,\right. \\
&\qquad \left.\log z^r_{2\phi_{max}}, \cdots, \log z^r_{N\phi_{min}}, \cdots, \log z^r_{N\phi_{max}}\right]^T,
\end{aligned}$$

and

$$\bar{Z}^r = \prod_{(r,\phi)\in\mathcal{L}} z^r_{r\phi}$$

- Finally, let $\mathcal{I}$ denote an indicator matrix defined as follows

$$\mathcal{I} = \begin{bmatrix} \mathcal{I}(1,1) & \mathcal{I}(1,2) & \mathcal{I}(1,3) \\ \mathcal{I}(2,1) & \mathcal{I}(2,2) & \mathcal{I}(2,3) \\ & \vdots & \\ \mathcal{I}(N,1) & \mathcal{I}(N,2) & \mathcal{I}(N,3) \end{bmatrix}$$

where for every p such that $1 \le p \le N$, if p belongs to j-th region,

$$\mathcal{I}(p,k) = \begin{cases} 1, & \text{if } k = j, \quad k \in \{1,2,3\} \\ 0, & \text{if } k \ne j, \quad k \in \{1,2,3\} \end{cases}$$

These notations are used to modify (5) can be expressed in terms of the above notations:

$$\begin{aligned}
& p(z^r \mid \underline{\theta}^r) \\
&= \left[\prod_{(r,\phi)\in\mathcal{L}} \frac{1}{z^r_{r\phi}}\right] \prod_{(r,\phi)\in\mathcal{L}} \frac{1}{\sqrt{2\pi\sigma^2_{r\phi}}} \exp\left\{-\frac{1}{2\sigma^2_{r\phi}}[\log z^r_{r\phi} - \mu_{r\phi}]^2\right\} \\
&= \frac{1}{\bar{Z}^r} \prod_{(r,\phi)\in\mathcal{L}_{rd}} \frac{1}{\sqrt{2\pi[\sigma^{rd}]^2}} \exp\left\{-\frac{1}{2[\sigma^{rd}]^2}[\log z^r_{r\phi} - \mu^{rd}]^2\right\} \\
&\quad \prod_{(r,\phi)\in\mathcal{L}_{lt}} \frac{1}{\sqrt{2\pi[\sigma^{lt}]^2}} \exp\left\{-\frac{1}{2[\sigma^{lt}]^2}[\log z^r_{r\phi} - \mu^{lt}]^2\right\} \\
&\quad \prod_{(r,\phi)\in\mathcal{L}_{rt}} \frac{1}{\sqrt{2\pi[\sigma^{rt}]^2}} \exp\left\{-\frac{1}{2[\sigma^{rt}]^2}[\log z^r_{r\phi} - \mu^{rt}]^2\right\} \\
&= \frac{1}{\bar{Z}^r} \frac{1}{(2\pi)^{N/2}\|\mathrm{diag}[\mathcal{I}\,\underline{\sigma}^2]\|^{1/2}} \\
&\quad \exp\left\{-\frac{1}{2}\left[\log \underline{Z}^r - \mathcal{I}\,\underline{\mu}\right]^T \left[\mathrm{diag}(\mathcal{I}\,\underline{\sigma}^2)\right]^{-1} \left[\log \underline{Z}^r - \mathcal{I}\,\underline{\mu}\right]\right\}
\end{aligned} \quad (6)$$

Given a hypothetical shape of the pavement boundaries, $\underline{\theta}^r$, the nuisance parameters $\underline{\mu}$ and $\underline{\sigma}^2$, corresponding to the means and variances of the three regions, can be empirically estimated from the observed radar image z^r by a maximum likelihood method. Note that the likelihood is in a normal form, and so the maximum likelihood and least squares estimate of the nuisance parameters are equivalent,

$$\begin{aligned}
\hat{\underline{\mu}} &= \left[\mathcal{I}^T\mathcal{I}\right]^{-1} \mathcal{I}^T \log \underline{Z}^r \\
\hat{\underline{\sigma}}^2 &= \left[\mathcal{I}^T\mathcal{I}\right]^{-1} \mathcal{I}^T \left(\log \underline{Z}^r - \mathcal{I}\,\hat{\underline{\mu}}\right)^T \left(\log \underline{Z}^r - \mathcal{I}\,\hat{\underline{\mu}}\right)
\end{aligned} \quad (7)$$

Substituting these estimates of $\underline{\mu}$ and $\underline{\sigma}^2$ back into (6), and taking the logarithm results in

$$\log p(z^r|\underline{\theta}^r) = -N^{rd}\log\hat{\sigma}^{rd} - N^{lt}\log\hat{\sigma}^{lt} \\ -N^{rt}\log\hat{\sigma}^{rt} - \log\bar{Z}^r - \frac{1}{2}N(1+\log 2\pi) \quad (8)$$

This result leads to the following interpretation:

> Given a hypothetical pavement boundary shape $\underline{\theta}^r$, the fidelity of this shape to the observed radar image z^r is assessed by the sum of the variances of the observed image over the corresponding three (road, left-side, and right-side) regions. The smaller the sum of these variances, the better the fidelity.

Weighted Log-normal Likelihood

It is indicated in (8) that part of the objective is trying to maximize $\log p(z^r|\underline{\theta}^r)$, i.e., to minimize $N^{rd}\log\hat{\sigma}^{rd} + N^{lt}\log\hat{\sigma}^{lt} + N^{rt}\log\hat{\sigma}^{rt}$, i.e., trying to minimize the standard deviations, and hence, variances of the three regions. However, in the radar images, it is clear that the backscatter distribution of the pavement is virtually homogeneous while the backscatter distributions of the roadsides are much more complicated. Therefore, we might prefer having an even smaller variance in the road region at the price of having slightly larger variances in the roadside regions.

In the above method, the variances are weighted proportionately to the number of pixels in their respective regions. In order for the standard deviation of the road, $\hat{\sigma}^{rd}$, to weigh more heavily, in terms of its contribution to the likelihood, N^{rd} would have to be proportionately large. The same is true for the other two regions as well.

In order to re-enforce our *a priori* belief that road pixels tend to be homogeneous (at least compared to the pixels belonging to either side of the road), and to overcome the undue influence of bright point scatterers in the roadside regions, we propose a function that gives the region of the road a different weight w^{rd} $(0 < w^{rd} < 1)$ from those given to the roadside regions.

$$L^r(z^r,\underline{\theta}^r) = w^{rd}N^{rd}\log\hat{\sigma}^{rd} + N^{lt}\log\hat{\sigma}^{lt} + N^{rt}\log\hat{\sigma}^{rt} \quad (9)$$

In this paper, we will utilize (9) to describe the log-likelihood of radar imaging process in a Bayesian estimation scheme:

$$\log p(z^r|\underline{\theta}^r) = -L^r(z^r,\underline{\theta}^r) - c^r \quad (10)$$

where

$$c^r \;=\; \log\bar{Z}^r + \frac{1}{2}N(1+\log 2\pi) \quad (11)$$

$L^r(z^r,\underline{\theta}^r)$ is the effective matching part of the log-likelihood while c^r is a constant for a specified image and thus can be neglected in the estimation process.

V. Optical Imaging Likelihood

Recall that the radar likelihood evaluates the merit of a pavement boundary hypothesis based on how homogeneous the corresponding radar returns over the corresponding left, right and road regions are. In the case of the optical image, given a hypothetical lane boundary shape, the expectation is that pixels in the observed optical image on and around the hypothetical lane boundary will possess "lane-like" brightness features. The brightness features of pixels that are further away from the hypothetical lane boundary are unpredictable, and hence trimmed. The optical likelihood evaluates the merit of each hypothesis in terms how lane-like the coresponding (trimmed) optical image pixels are. In this paper we use the likelihood pdf in [33], without any modification, for evaluating the merit of a hypothetical lane shape.

The likelihood described in [33] encodes the knowledge that the edges of the lane should be near intensity gradients whose orientation should be perpendicular to the lane edge. More specifically, given a hypothetical parameter set of underlying edges $\underline{\theta}^o = \{k', vp', b'_L, b'_R\}$, we assume that the likelihood of observing the optical image Z^o is given by

$$p(z^o|\underline{\theta}^o) = \gamma(\underline{\theta}^o)\; e^{-L^o(z^o,\underline{\theta}^o)} \quad (12)$$

where $L^o(z^o,\underline{\theta}^o)$ denotes an energy function and $\gamma(\underline{\theta}^o)$ is a normalizing constant for a given optical deformation parameter set $\underline{\theta}^o$,

$$\gamma(\underline{\theta}^o) = \frac{1}{\int e^{-L^o(z^o,\underline{\theta}^o)}dz^o} \quad (13)$$

The energy function $L^o(z^o,\underline{\theta}^o)$ can be described in the following steps:

- Define the Cauchy density function

$$f(\alpha,x) \triangleq \frac{\alpha}{\pi}\frac{1}{1+\alpha^2x^2}, \quad (14)$$

where α is a per-determined constant to control the width of the Cauchy density function. It takes different values when evaluating the contribution of the magnitude and direction of the gradient at a pixel to the likelihood. In (16) α_m controls the portion of the contribution of magnitude to the energy function, which is 0.01 in our experiment; while α_d controls the portion of the contribution of direction to the energy function, which is 1.13 in our experiment.

- Let $gm(r,c)$ be the gradient magnitude at pixel (r,c), with $gd(r,c)$ being the corresponding gradient direction. Instead of the observed optical image, the gradient features are used in the representation of the energy function.
- Define the edges of the lane in the image by the curves:

$$E_L(r,c,\underline{\theta}^o) \triangleq \frac{k'}{r-hz} + b'_L(r-hz) + vp' \\ E_R(r,c,\underline{\theta}^o) \triangleq \frac{k'}{r-hz} + b'_R(r-hz) + vp' \quad (15)$$

Given these definitions, the energy function that is related to the probability of observing an image gradient field given a set of lane shape parameters $\underline{\theta}^o$ is

$$\begin{aligned} L^o(z^o, \underline{\theta}^o) &= L^o(gm, gd, \underline{\theta}^o) \\ &= -\sum_{(r,c)} gm(r,c) \times f(\alpha_m, c - E_L(r,c,\underline{\theta}^o)) \\ &\qquad \times f(\alpha_d, \cos(gd(r,c) - \mathrm{atan}(\frac{d}{dr} E_L(r,c,\underline{\theta}^o)))) \\ &\quad - \sum_{(r,c)} gm(r,c) \times f(\alpha_m, c - E_R(r,c,\underline{\theta}^o)) \\ &\qquad \times f(\alpha_d, \cos(gd(r,c) - \mathrm{atan}(\frac{d}{dr} E_R(r,c,\underline{\theta}^o)))) \end{aligned} \tag{16}$$

Then $e^{-L^o(z^o, \underline{\theta}^o)}$ is an unnormalized likelihood function. In other words, the contribution made by a pixel to this function is the gradient magnitude (gm) at that pixel, multiplied by a function whose value decreases as the pixel column gets further from the lane edge (effectively the data is trimmed) and a function whose value decreases as the gradient direction (gd) at the pixel becomes less perpendicular to the tangent to the lane edge.

The logarithm of the optical imaging likelihood is

$$\log p(z^o|\underline{\theta}^o) = -L^o(z^o, \underline{\theta}^o) + \log \gamma(\underline{\theta}^o) \tag{17}$$

VI. Joint MAP Estimate for Lane and Pavement Boundaries

Since the prior distributions of the deformation parameters and the imaging likelihood pdfs are available, we shall pose the lane and pavement edge detection problem in a Bayesian framework. Let z^r be a realization of the radar random field Z^r and z^o be a realization of the optical random field Z^o. The optical and radar fusion detection problem can be solved by the joint MAP estimate

$$\begin{aligned} \hat{\underline{\theta}} \;=\; \{\hat{\underline{\theta}}^r, \hat{\underline{\theta}}^o\} &= \arg\max_{\underline{\theta}} p(\underline{\theta}|z^r, z^o) \\ &= \arg\max_{\{\underline{\theta}^r, \underline{\theta}^o\}} p(\underline{\theta}^r, \underline{\theta}^o | z^r, z^o) \end{aligned} \tag{18}$$

According to the Bayes' rule, we have

$$\hat{\underline{\theta}} = \arg\max_{\{\underline{\theta}^r, \underline{\theta}^o\}} \frac{p(z^r, z^o, \underline{\theta}^r, \underline{\theta}^o)}{p(z^r, z^o)} \tag{19}$$

Since the denominator $p(z^r, z^o)$ is fixed by the observation, the above formula can be rewritten as

$$\hat{\underline{\theta}} = \arg\max_{\{\underline{\theta}^r, \underline{\theta}^o\}} p(z^r, z^o, \underline{\theta}^r, \underline{\theta}^o) \tag{20}$$

By the chain rule of conditional probability, we have

$$p(z^r, z^o, \underline{\theta}^r, \underline{\theta}^o) = p(\underline{\theta}^r) p(z^r|\underline{\theta}^r) p(\underline{\theta}^o|\underline{\theta}^r, z^r) p(z^o|\underline{\theta}^o, z^r, \underline{\theta}^r) \tag{21}$$

Given the road shape parameter $\underline{\theta}^r$, the lane shape parameter $\underline{\theta}^o$ is independent of the radar observation z^r, i.e.,

$$p(\underline{\theta}^o|\underline{\theta}^r, z^r) = p(\underline{\theta}^o|\underline{\theta}^r), \tag{22}$$

and also given the lane shape parameter $\underline{\theta}^o$, the optical observation z^o is independent of the radar observation z^r and road shape parameter $\underline{\theta}^r$, i.e.,

$$p(z^o|\underline{\theta}^o, z^r, \underline{\theta}^r) = p(z^o|\underline{\theta}^o) \tag{23}$$

Substituting (22) and (23) into (21), and then substituting (21) into (20), we have

$$\begin{aligned} \hat{\underline{\theta}} &= \arg\max_{\{\underline{\theta}^r, \underline{\theta}^o\}} p(\underline{\theta}^r)\; p(z^r|\underline{\theta}^r)\; p(\underline{\theta}^o|\underline{\theta}^r)\; p(z^o|\underline{\theta}^o) \\ &= \arg\max_{\{\underline{\theta}^r, \underline{\theta}^o\}} p(\underline{\theta}^r, \underline{\theta}^o)\; p(z^r|\underline{\theta}^r)\; p(z^o|\underline{\theta}^o) \end{aligned} \tag{24}$$

Making use of the logarithm of the density functions, the joint MAP estimate turns to

$$\begin{aligned} \hat{\underline{\theta}} &= \arg\max_{\underline{\theta}} \; \log p(z^r, z^o, \underline{\theta}) \\ &= \arg\max_{\{\underline{\theta}^r, \underline{\theta}^o\}} \{\log p(\underline{\theta}^r, \underline{\theta}^o) + \log p(z^r|\underline{\theta}^r) + \log p(z^o|\underline{\theta}^o)\} \\ &= \arg\max_{\{\underline{\theta}^r, \underline{\theta}^o\}} \{\log p(\underline{\theta}^r, \underline{\theta}^o) - L^r(z^r, \underline{\theta}^r) - L^o(z^o, \underline{\theta}^o) + \log \gamma(\underline{\theta}^o)\} \end{aligned} \tag{25}$$

Calculating $\gamma(\underline{\theta}^o)$ is intractable, as it involves an integration over all the realizations of Z^o (see (13)). We will use a relative weighting to compensate for neglecting $\gamma(\underline{\theta}^o)$.

VII. Empirical MAP

Since $\gamma(\underline{\theta}^o)$ in (25) is intractable, it is impossible to obtain a theoretical estimate for deformation parameters $\underline{\theta}$ based on (25). The primary difficulty we met in this edge detection problem is that for the two imaging likelihood pdfs, the radar imaging likelihood is normalized, while the optical imaging likelihood is not. Instead of computing the normalizing factor $\gamma(\underline{\theta}^o)$, we turn to the empirical MAP estimate,

$$\hat{\underline{\theta}} = \arg\max_{\{\underline{\theta}^r, \underline{\theta}^o\}} \{\log p(\underline{\theta}^r, \underline{\theta}^o) - L^r(z^r, \underline{\theta}^r) - \beta \cdot L^o(z^o, \underline{\theta}^o)\} \tag{26}$$

We expect the weighting factor β to play the same role as the normalizing constant $\gamma(\underline{\theta}^o)$.

The matching functions $L^r(z^r, \underline{\theta}^r)$ and $L^o(z^o, \underline{\theta}^o)$ are relative rather than absolute measure. This is the root cause of the problem — the matching functions are of different dynamic ranges within the parameter space of interest. Fig. 6 shows an example of the different dynamic ranges. The dynamic range of the radar matching function is 1.4×10^3, while the dynamic range of the optical matching function is 4.5×10^4. The instant result of the difference in dynamic ranges is that in most cases, the optical image dominates the joint estimate(Fig. 7).

In order to overcome this inherent deficiency, we scale (using the weight β) one of the matching functions so that the ranges of variation for the weighted matching functions are approximately the same. This fixed weight β is empirically derived by examining the individual matching functions for a (training) set of vision and radar image pairs. Our extensive experiments seem to indicate that $\beta = 0.01$ gives good performance for most image pairs(Fig.8).

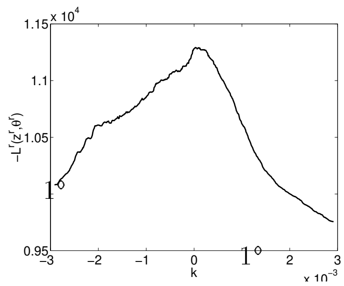

(a) radar matching function

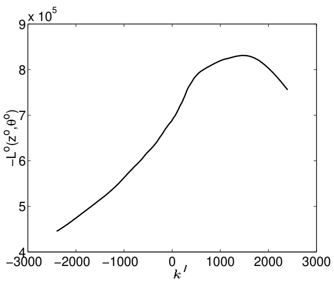

(b) vision matching function

Fig. 6. Different dynamic range of the matching functions for optical and radar images

(a) Lane edge detection

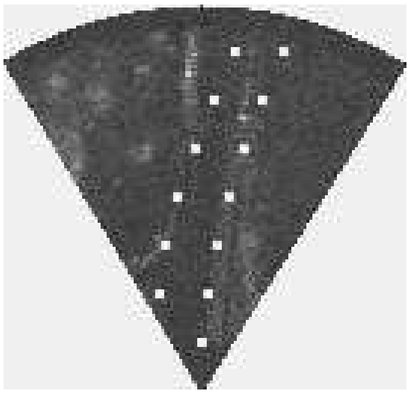

(b) Pavement edge detection

Fig. 7. Wrong edge detection by no-weight scheme

VIII. Computation of the Empirical MAP Estimate

The problem is one of obtaining the maximum in (26), which is equivalent to finding the mode of a six-dimensional density surface. The surface is non-concave with many local maxima, hence we can not just apply the greedy search algorithms such as conjugate gradient methods. In this paper, we suggest two techniques to find the global maximum:

A. Multi-resolution Pseudo-Exhaustive Search

Exhaustive search can find the optimal solutions at the cost of unacceptable computation resources in some optimization problems. For the problem addressed in this paper, exhaustive search is not feasible due to the

(a) Lane edge detection

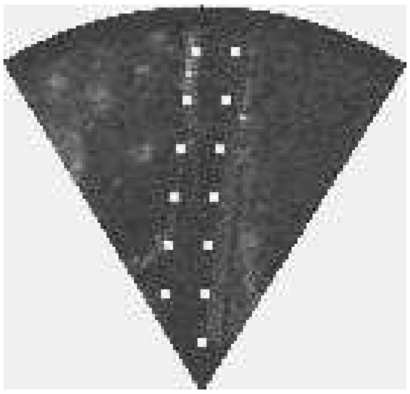

(b) Pavement edge detection

Fig. 8. Edge detection by fixed weight scheme

large searching space. Instead, a multi-resolution pseudo-exhaustive search method is studied. First, we constrain the parameters in appropriate ranges. Then we select a set of coarse step sizes (coarse grid) for the parameters and do the pseudo-exhaustive search to find the maximum of the joint MAP objective function (26). Once this coarse maximum is found, the corresponding estimated parameters are taken as the center of a finer search procedure with finer step sizes (finer grid) and smaller ranges of the parameters. Repeat the above step until the desired parameter grid size is reached.

B. Metropolis Algorithm with Geometric Annealing

Although the multi-resolution pseudo-exhaustive search gives us relatively accurate solutions, since we have six parameters to estimate, the search procedure is very time consuming (almost minutes per image pair). To accelerate the maximization procedure, we employ a sub-optimal approach, the Metropolis algorithm [9] with a geometric annealing schedule [34], to perform this maximization,

1. Set $i = 0$, and initialize $\underline{\theta}^{(0)}$.
2. Calculate $\log P(z^r, z^o, \underline{\theta}^{(i)})$.
3. Pick $\tilde{\underline{\theta}}$ at random among all the possible parameter values in the neighborhood of $\underline{\theta}^{(i)}$.
4. Calculate $\log P(z^r, z^o, \tilde{\underline{\theta}})$.
5. Calculate $\rho^{(i)} = \exp\left(\frac{\log P(z^r, z^o, \tilde{\underline{\theta}}) - \log P(z^r, z^o, \underline{\theta}^{(i)})}{T^{(i)}}\right)$, where $T^{(i)} = T_{init}\left(\frac{T_{final}}{T_{init}}\right)^{\frac{i+1}{max\_iter}}$.

6. Update the curve deformation parameters

$$\underline{\theta}^{(i+1)} = \begin{cases} \underline{\tilde{\theta}} & \text{if } \rho^{(i)} \geq 1 \\ \underline{\tilde{\theta}} & \text{w.p. } \rho^{(i)} \text{ if } \rho^{(i)} < 1 \\ \underline{\theta}^{(i)} & \text{otherwise} \end{cases}$$

7. Set i = i+1 and go to step 2.

The running time is now about 20 seconds per image pair. We are confident that by optimizing our implementation, we can run this algorithm in real-time.

IX. Performance Comparison

We have applied the proposed multisensor fusion method to jointly detect the lane and pavement boundaries in registered radar and optical images. We have implemented the scheme described in this paper on a data set containing 25 vision and radar image pairs. These image pairs were acquired under a variety of imaging conditions – see http://www.eecs.umich.edu/~bingm/database.pdf – including nighttime, fog, rain, snow, etc. In some of the optical images the lane boundaries are not so easily discerned and extracted, conversely in some radar images the pavement boundaries are not so easily discerned and extracted. Figs. 9(a),(b) and 10(a),(d) show examples of optical and radar image pairs where such problems occur.

Since multiple (radar and optical) sensors provide more information and hence a more precise interpretation of the sensed environment, the performance of lane and pavement boundary detection is robust and accurate. For representative road scenes, the results obtained via independent optical and radar edge detection algorithms are illustrated in Figs. 9(a) and (b), Figs. 10(a) and (b), and the results with the fusion method are shown in Figs. 9(c) and (d), Figs. 10(c) and (d). Fig. 9 demonstrates that the radar image improves the lane detection in the optical image in some situations, while Fig. 10 demonstrates that the optical image improves the pavement edge detection in the radar image in other situations. Both examples indicate that the fusion method outperforms edge detection independently extracted from single sensors. We also noticed that fusion does not degrade the performance of the individual detection results when they are good by themselves.

To compare the advantage of the fusion method over the single sensor detection methods and to appreciate the role of likelihood weighting, we undertake a large experiment. For the database of 25 vision and radar image pairs referred to earlier, we hand-picked ground truth and plot the average detection errors compared to ground truth in Figs. 11 and 12. Fig. 11 shows the reduction in detection errors when the two data modalities are fused, compared to processing them individually. Fig. 12 shows the reduction in detection errors when the likelihoods are relatively weighted using $\beta = 0.01$, compared to when no weighting is used.

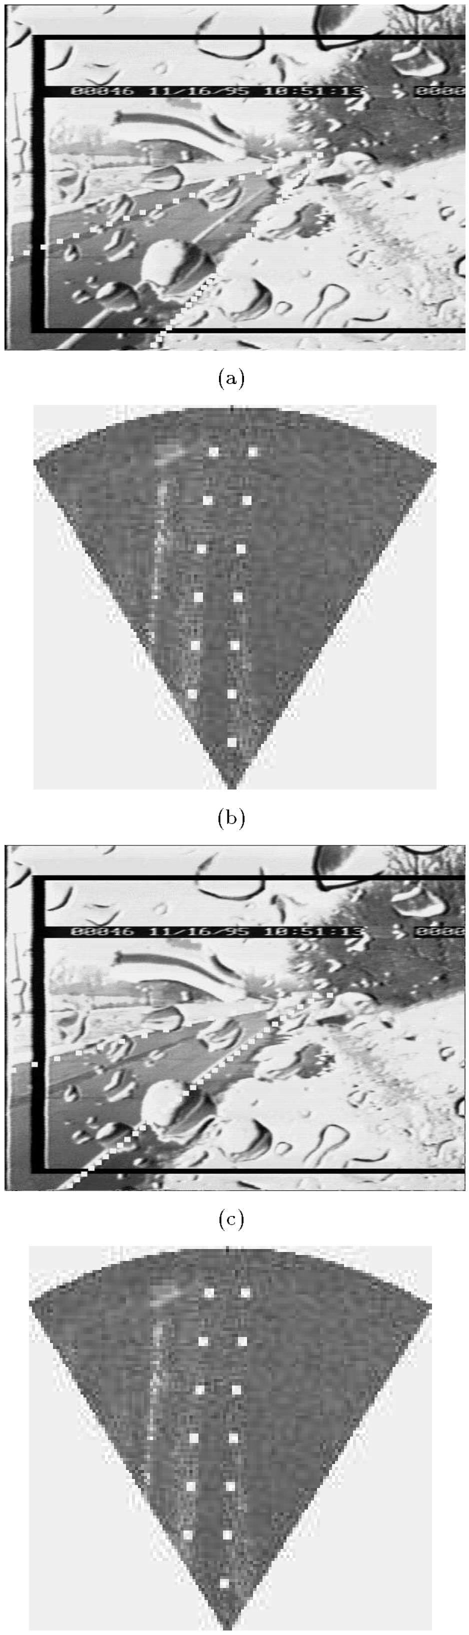

(a) (b) (c) (d)

Fig. 9. Comparison of the fusion and single sensor based methods

X. Discussion and Conclusion

One of the emerging problems in intelligent vehicle and highway systems, the simultaneous detection of lane and pavement boundaries, is investigated in this paper. We employ a novel Bayesian multisensor fusion technique to

(a)

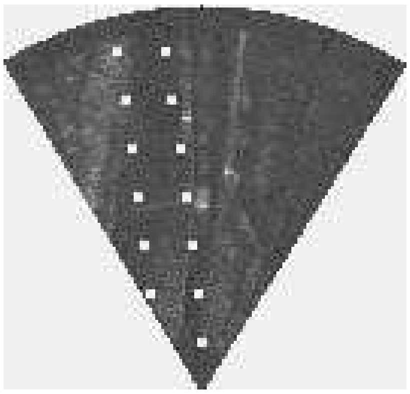

(b)

(c)

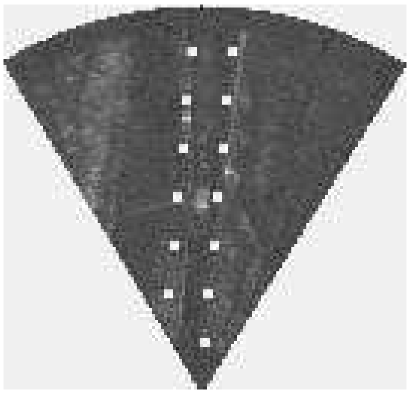

(d)

Fig. 10. Comparison of the fusion and single sensor based methods

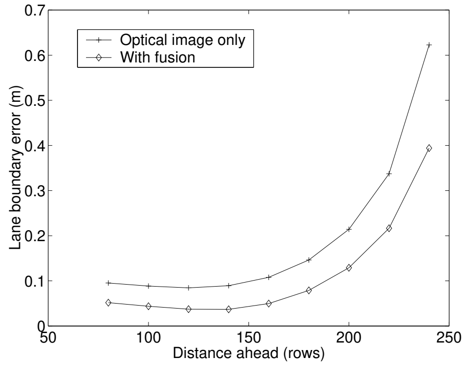

(a) Errors for lane boundary detection

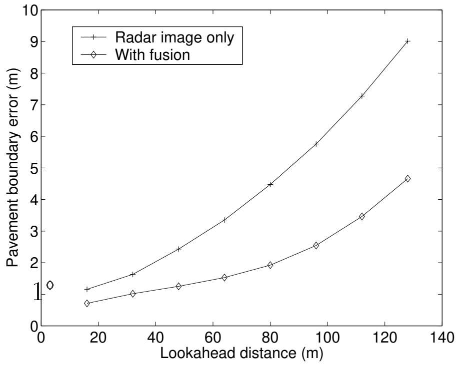

(b) Errors for pavement boundary detection

Fig. 11. Advantage of the fusion method over single sensor based method

locate the lane and pavement edges in the optical and radar images. We make use of the deformable template model to globally describe the lane and pavement boundary shapes. This global shape model enables the method to successfully disregard extraneous edges. Likelihood models of the radar and optical imaging processes are introduced, and the simultaneous boundary detection problem is posed in a Bayesian framework. It is shown that the resulting joint MAP estimate is intractable. This difficulty is overcome by a likelihood weighting scheme, resulting in an empirical Bayes' estimate, which is entirely tractable. Experimental results are used to demonstrate that this fusion method operates robustly in detecting the lane and pavement boundaries. A comparison is made between the results of this fusion method and those obtained by processing the individual modalities separately, and also those obtained when no likelihood weighting is used.

Although we have found the performance of the empirically determined weight to be satisfactory, we are still interested in a better weight selection scheme. We are currently studying the role of likelihood curvatures in determining this weight. Since likelihood curvature is intimately related to estimator variance, we hope to obtain a more mathematically appealing approximation to the true joint MAP estimate in future.

References

[1] W. Kasprzak and H. Niemann, "Adaptive road recognition and ego-state tracking in the presence of obstacles," *International Journal of Computer Vision*, vol. 28, no. 1, pp. 5–26, June 1998.

[2] E. Mizutani, T. Kozek, and L. O. Chua, "Road lane marker

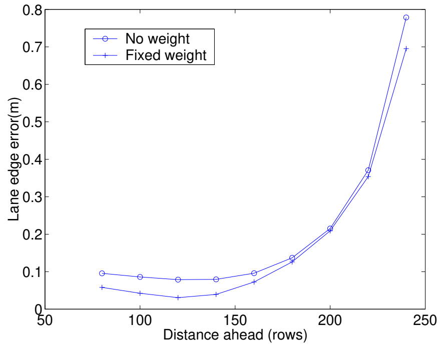

(a) Errors for lane boundary detection

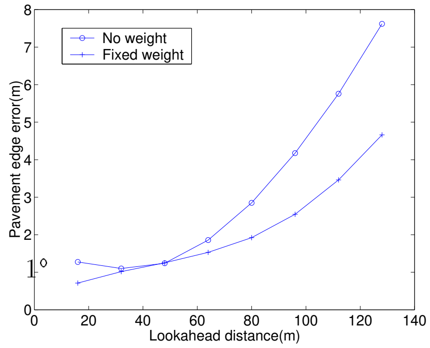

(b) Errors for pavement boundary detection

Fig. 12. Advantage of the fixed weighting scheme over no weight scheme

extraction by motion-detector cnns," in *Proc. IEEE Intl Joint Conf. on Neural Networks*, 1998, pp. 503–508.

[3] E. D. Dickmanns and B. D. Mysliwetz, "Recursive 3-d road and relative ego-state recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 2, pp. 199–213, Feb. 1992.

[4] S. K. Kenue, "LANELOK: Detection of lane boundaries and vehicle tracking using image-processing techniques, parts I and II," in *SPIE Mobile Robots IV*, 1989.

[5] K. C. Kluge, *YARF: An open-ended framework for robot road following*, Ph.D. thesis, Carnegie Mellon University, 1993.

[6] K. C. Kluge, "Extracting road curvature and orientation from image edge points without perceptual grouping into features," in *Proc. Intell. Vehic. '94 Symp.*, Sept. 1994, pp. 109–114.

[7] K. C. Kluge, "Performance evaluation of vision-based lane sensing: some preliminary tools, metrics and results," in *IEEE Conference on Intelligent Transportation Systems*, 1997.

[8] S. Lakshmanan and D. Grimmer, "Deformable template approach to detecting straight edges in radar images," *IEEE Trans. on PAMI*, vol. 18, no. 4, pp. 438–443, Apr. 1996.

[9] S. Lakshmanan and K. C. Kluge, "LOIS: A real-time lane detection algorithm," in *Proc. Conf. information Sciences and Systems*, Princeton, 1996.

[10] D. Pomerleau and T. Jochem, "Rapidly adapting machine vision for automated vehicle steering," *IEEE Expert*, vol. 11, no. 2, pp. 19–27, 1996.

[11] B. Ma, S. Lakshmanan, and A. Hero, "Detection of curved road edges in radar images via deformable templates," in *Proc. IEEE. Intl. Conf. Image Proc.*, 1997.

[12] S. Lakshmanan, A. K. Jain, and Y. Zhong, "Detecting straight edges in millimeter-wave images," in *IEEE Inter. Conf. Image Processing*, 1996, vol. 1.

[13] K. Kaliyaperumal, S. Lakshmanan, and K. Kluge, "Detecting roads and obstacles in radar images," *to appear in IEEE Transactions on Vehicular Technology*, 2000.

[14] M. Beauvais, C. Kreucher, and S. Lakshmanan, "Building world models for mobile platforms using heterogeneous sensors fusion and temporal analysis," in *Proc. IEEE Conference on Intelligent Transportation Systems*, Boston, MA, USA, 1997.

[15] M. Beauvais and S. Lakshmanan, "CLARK: A heterogeneous sensor fusion method for finding lanes and obstacles," *Image and Vision Computing*, vol. 18, no. 5, pp. 397–413, 2000.

[16] N. M. Laird, "Empirical bayes estimates using the nonparametric maximum likelihood estimate for the prior," *Journal of Statistical Computation and Simulation*, pp. 211–220, 1982.

[17] K. Alam and A. Mitra, "An empirical bayes estimate of multinomial probabilities," *Communications in Statistics. A. Theory and Methods*, pp. 3103–3127, 1986.

[18] A. O. Hero, R. Piramuthu, J. A. Fessler, and S. R. Titus, "Minimax emission computed tomography using high resolution anatomical side information and B-spline models," *IEEE Trans. on Information Theory*, vol. 45, no. 3, pp. 920–938, 1999.

[19] A. Hernandez, O. Basset, I. Magnin, A. Bremond, and G. Gimenez, "Fusion of ultrasonic and radiographic images of the breast," in *Proc. IEEE Ultrasonics Symposium*, San Antonio, TX, USA, 1996, pp. 1437–1440.

[20] L. Bruzzone and S. B. Serpico, "Fusion of multisensor and multitemporal data in remote-sensing image analysis," in *Proc. IEEE Intl Geoscience and Remote Sensing Symposium*, Seattle, WA, USA, 1998.

[21] C. Pohl and J. L. Van Genderen, "Multisensor image fusion in remote sensing: concepts, methods and applications," *International Journal of Remote Sensing*, vol. 19, no. 5, pp. 823–854, 1998.

[22] B. Solaiman, R. K. Koffi, M. Mouchot, and A. Hillion, "Information fusion method for multispectral image classification postprocessing," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 36, no. 2, pp. 395–406, 1998.

[23] R. C. Harney, "Sensor fusion for target recognition: a review of fundamentals and a potential approach to multisensor requirement allocation," in *Proc. of SPIE*, 1994, vol. 2269, pp. 316–335.

[24] D. C. Borghys, P. S. Verlinde, C. Perneel, and M. Acheroy, "Long-range target detection in a cluttered environment using multisensor image sequences," in *Proc. of SPIE*, 1997, vol. 3068, pp. 569–578.

[25] B. T. Sweet and C. L. Tiana, "Image processing and fusion for landing guidance," in *Proc. of SPIE*, 1996, vol. 2736, pp. 84–95.

[26] D. Langer and T. Jochem, "Fusing radar and vision for detecting, classifying and avoiding roadway obstacles," in *IEEE Intelligent Vehicles Symposium*, 1996, pp. 333–338.

[27] C. Thorpe, T. Jochem, and D. Pomerleau, "1997 automated highway free agent demonstration," in *IEEE Conference on Intelligent Transportation Systems*, 1997, pp. 496–501.

[28] R. R. Brooks, *Multi-sensor fusion : fundamentals and applications with software*, McGraw-Hill, 1992.

[29] J. Aitchison and J. A. C. Brown, *The Log-normal Distribution*, Cambridge University Press, 1957.

[30] A. Farina, A. Russo, and F. A. Studer, "Coherent radar detection in log-normal clutter," *IEE Proc., Part F: Communications, Radar and Signal Processing*, vol. 133, no. 1, pp. 39–54, Feb 1986.

[31] J. S. Lee, "Speckle analysis and smoothing of synthetic aperture radar images.," *Comput Graphics Image Process*, vol. 17, no. 1, pp. 24–32, Sept 1981.

[32] Slolnik M. L., *Introduction to radar systems*, McGraw-Hill, 1980.

[33] S. Lakshmanan and K. C. Kluge, "Lane detection for automotive sensors," in *Proceedings of the 20th International Conference on Acoustics, Speech, and Signal Processing*, 1995, pp. 2955–2958.

[34] D. N. Strenki and S. Kirkpatrick, "Analysis of finite-length annealing schedules," *Algorithmica*, pp. 346–366, 1991.