Road Segmentation for Classification of Road Weather Conditions

Emilio J. Almazan, Yiming Qian(✉), and James H. Elder

Centre for Vision Research, York University, Toronto, Canada
ealmazanm@gmail.com, {yimingq,jelder}@yorku.ca

Abstract. Using vehicle cameras to automatically assess road weather conditions requires that the road surface first be identified and segmented from the imagery. This is a challenging problem for uncalibrated cameras such as removable dash cams or cell phone cameras, where the location of the road in the image may vary considerably from image to image. Here we show that combining a spatial prior with vanishing point and horizon estimators can generate improved road surface segmentation and consequently better road weather classification performance. The resulting system attains an accuracy of 86 % for binary classification (bare vs. snow/ice-covered) and 80 % for 3 classes (dry vs. wet vs. snow/ice-covered) on a challenging dataset.

Keywords: Linear perspective · Vanishing point · Horizon · Road segmentation · Weather conditions

1 Introduction

Automatic assessment of road weather conditions using vehicle camera data can be used to inform the human driver, driver-assist controls and autonomous control systems. Moreover, the information can be shared across connected vehicles, alerting following vehicles to conditions ahead. Another application is automatic dispatch and verification of snow ploughs and service vehicles. Given their typically wide geographic distribution, these service vehicles can provide real-time data on road conditions to central management, which can then use the data to verify maintenance and optimize dispatch.

While future generations of service vehicles may be manufactured with appropriate built-in cameras, in the meantime there is interest in retrofitting existing vehicles with removable dash cams that can be used for multiple purposes. This poses a challenge for video analytics, as the pose of the camera relative to the road surface may vary considerably. Since the cameras are mounted inside the vehicle, imagery may be partially occluded by the hood of the vehicle, For snow ploughs, the road surface may also be occluded further by the plough, depending on its position (Fig. 1).

To address these challenges a reliable algorithm for segmenting the road surface from the imagery is required. This method must be able to handle variations

G. Hua and H. Jégou (Eds.): ECCV 2016 Workshops, Part I, LNCS 9913, pp. 96–108, 2016.
DOI: 10.1007/978-3-319-46604-0_7

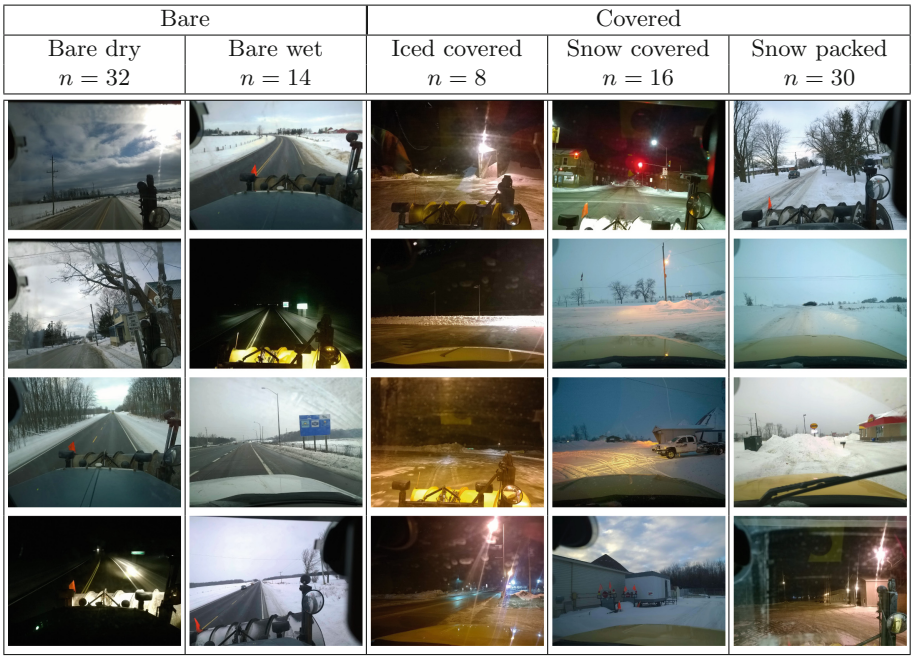

Fig. 1. Example images from training dataset, and the number n in each class.

in the position and pose of the camera as well as geometry of the road surface. Using appearance features of the road surface (e.g., texture) for segmentation is unlikely to be reliable across diverse weather conditions, since the road appearance will vary considerably and sometimes may strongly resemble other surfaces in the scene. For these reasons, we focus here on geometric methods for identifying the road surface and show that by fusing a combination of these methods we can significantly improve road weather classification performance.

In particular we develop a novel method for estimating the road vanishing point, which yields a triangular road segmentation hypothesis. This vanishing point method also delivers a measure of reliability, which can be used to identify when the vanishing point is ill-defined (in a parking lot, for example). Under these conditions we revert to a weaker segmentation based upon detection of the horizon line. Combining these with a spatial prior then delivers an estimated road segmentation tailored to each image.

This paper is organized as follows: Sect. 2 reviews prior work, Sect. 3 details our road segmentation algorithm, Sect. 4 describes our classification process, Sect. 5 reports results and finally Sect. 6 presents our conclusions and plans for future work.

2 Prior Work

2.1 Road Segmentation

One might initially imagine using surface appearance properties to distinguish the road pavement surface from other surfaces in the scene [2,3]. However, this approach is problematic here for at least three reasons: (1) pavement appearance varies depending upon the exact road materials employed and age of the road; (2) other nearby surfaces (e.g., sidewalks, driveways, buildings) may be constructed from similar materials; (3) the road surface may be partially or completed covered by snow and/or ice.

For these reasons, we use geometric methods to estimate the road segmentation. This is still non-trivial due to diversity in camera pose and road geometry. Roads vary in shape, are sometimes relatively unstructured, non-homogeneous and vary in appearance under varying weather and illumination conditions.

Previous approaches have estimated the vanishing point [12,18,19], horizon [1] and/or border lines of the road [13]. The vanishing point is typically detected using texture information from Gabor wavelet filters [1,12,18], using a Hough transform [24] or with a line segment voting scheme [20]. In this work we adapt recent work on Hough-based vanishing point detection [21] that has proven effective for Manhattan frame estimation.

The horizon line is often estimated as the line that partitions the image into two regions that differs maximally in appearance [6,7], however more elaborate approaches based on gist descriptors [1,15] have also been employed. Here we identify the horizon as the vertical location that maximizes the RMS first derivative in the vertical direction across all horizontal locations and colour channels.

2.2 Road Weather Classification

Since the focus of this paper is on road segmentation, we will review the literature on road weather classification only briefly.

Much of the prior work on road condition classification has focused on the use of polarization and infrared cameras [5,8–10,14,25], which can be expensive and installation can be complex. However, there are also a number of efforts employing standard RGB cameras. Omer and Fu [16] used an SVM with RGB and gradient histogram features to classify conditions as bare, covered or covered with bare tire tracks. However their approach required manual cropping of each image to extract the image region projecting from the ground surface, which is impractical for a real system. Kauai et al. [11] used colour cues to add some degree of illumination invariance, however their approach depends on detecting white line markings to identify the road area, which will fail under snowy conditions or for roads that are poorly marked. In a very recent paper, Amthor et al. [4] proposed a spatiotemporal approach that integrates over many frames to detect specular reflections indicative of wet conditions. While their method improves over prior approaches, it requires integration over many frames, increasing computational load and delay.

For the present paper we employ the classifier reported recently by Qian et al. [17], which uses a naive Bayes classifier over texton and luminance features. Please see Sect. 4 for more details.

2.3 Dataset

To train and evaluate our algorithm, we employ the challenging dataset of 100 2048×1536 pixel images (Fig. 1) introduced by Qian et al. [17], obtained directly from the authors. The dataset contains roads under different weather conditions, from bare dry to snow packed. Each of the five classes was randomly and evenly split between training and test datasets, each consisting of 50 images.

The pictures were taken at different times of the day, thus covering a wide range of illumination conditions and the camera pose varied considerably. The road condition class was identified manually by our industrial partner. For training and evaluation purposes, we manually segmented the road surface from the background. Running our classifier on the manually segmented imagery allows us to estimate the potential for increasing classification performance through further improvements in segmentation.

3 Road Segmentation

The main challenge for segmenting the road surface is the variability of the road appearance under different weather conditions, which limits the utility of appearance features such as luminance, colour, texture and detailed road markings. We therefore propose a method that relies on contextual information to define the vanishing point of the road and horizon line. The process consists of four stages: (1) estimation of the vanishing point of the road; (2) assessment of the reliability of this vanishing point estimate; (3) direct estimation of the horizon line, for vanishing points assessed to be unreliable; (4) fusion with a spatial prior to identify the region of the image corresponding to the road surface.

3.1 Vanishing Point Estimation

A vanishing point is defined as a point in the image plane where parallel lines converge. We use the line detector algorithm of Tal and Elder, 2012 [21] (code obtained directly from authors). The detector returns betwen 122 and 746 lines for each of the images in the training dataset. It also returns the estimated total length l of the line segments along each detected line. As revealed in Fig. 2, the geometry (position ρ, orientation θ, length l) of each line provides some information about the likelihood that it is generated by the vanishing point (ON) versus a background process (OFF). We therefore rerank the lines using a naive Bayes model to approximate the likelihood ratio L_i for each line i:

$$L_i = \frac{p(\rho_i, \theta_i, l_i|\text{ON})}{p(\rho_i, \theta_i, l_i|\text{OFF})} \approx \frac{p(\rho_i|\text{ON})p(\theta_i|\text{ON})p(l_i|\text{ON})}{p(\rho_i|\text{OFF})p(\theta_i|\text{OFF})p(l_i|\text{OFF})}. \quad (1)$$

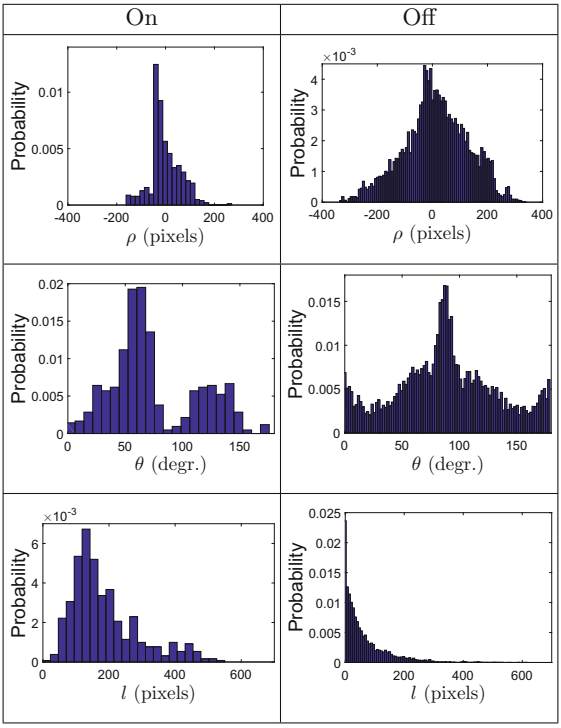

Fig. 2. Likelihood distributions for the three line features (ρ, θ and l).

We estimate the vanishing point as the point in the image that minimizes the distance to the top-ranked n lines. Specifically, we adopt a naive Bayes approach and choose the point v that maximizes

$$p(v|\mathrm{L}) \propto \prod_{i}^{n} p(d_i|v)p(v) \tag{2}$$

where $\mathrm{D} = (d_1, d_2, ..., d_n)$ are the distances to the detected lines. The spatial prior $p(v)$ is modelled as a Gaussian distribution, and the likelihoods $p(d_i|v)$ are determined from the training data as shown in Fig. 3.

Equation 2 is not convex in general; to maximize we select the optimal solution from 50 gradient descent solutions (MATLAB fminsearch), initialized by randomly sampling from the prior $p(v)$.

Figure 4 compares performance of the vanishing point algorithm on the training dataset with and without the re-ranking step. We find that generally the re-ranking improves results and that error is minimizec by using the top-ranked 20 lines. Note also that using the lines to estimate the vanishing point yields much lower error than using the centroid of vanishing points from the training set (Prior model). Figure 5 shows examples of automatically estimated vanishing points.

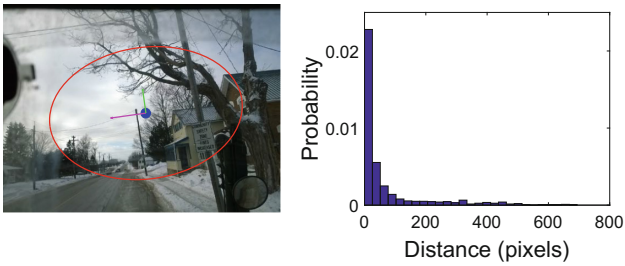

(a) Prior (b) Likelihood

Fig. 3. (a) Vanishing point prior distribution plotted on a sample image from the dataset. The red ellipse indicates the 95 % confidence interval for the vanishing point. (b) Likelihood distribution for the distance of the top-ranked 20 detected lines from the vanishing point. (Color figure online)

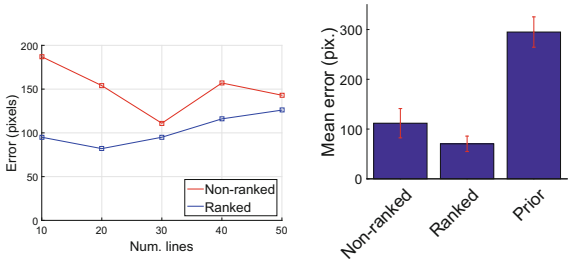

(a) Number of lines (b) Mean error

Fig. 4. (a) Average Euclidean error of the estimated vanishing point as a function of the number n of lines employed. (b) Mean error and standard error of the mean, collapsing over n.

There are some situations where the vanishing point of the road is not readily apparent, such as in parking lots or intersections (Fig. 10(c–f)). Our vanishing point algorithm will tend to produce large errors for these cases, which could in turn lead to large errors in the road segmentation. To prevent this, we assess the reliability of vanishing point estimates as the average distance $\bar{D}_k$ of the top k lines closest to the estimated vanishing point; If $\bar{D}_k$ exceeds a threshold t, we reject the vanishing point estimate and use an alternate method to determine the horizon line (see below). There are two free parameters for this reliability measure: the number k of lines and the threshold t. We optimize using the training data based on the ultimate error in estimating the horizon line, assessed as the average absolute error at left and right image boundaries (Fig. 6): values of $k = 12$ and $t = 21$ were found to be optimal.

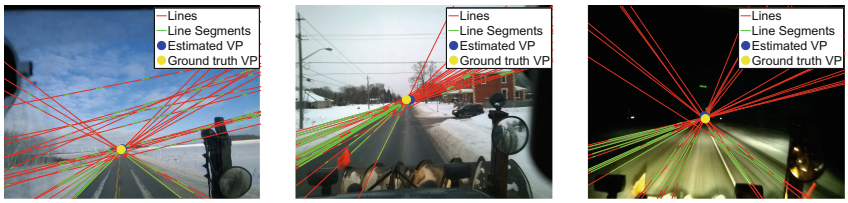

Fig. 5. Examples of automatically estimated vanishing points.

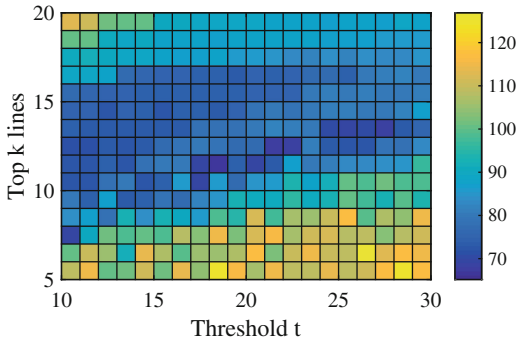

Fig. 6. The error graph of different combinations of threshold t and top k lines.

3.2 Horizon Estimation

When vanishing point estimation fails it may still be possible to constrain the road location using the horizon line. The horizon line can be defined by two parameters, for example, its vertical location and orientation. For our dataset, however, we found that orientation estimation could be quite unreliable. We therefore fixed the orientation to horizontal.

To estimate the vertical location of the horizon, we used the training images for which our vanishing point estimates were judged (automatically) to be unreliable as a horizon training dataset ($k = 23$ images in all). We then (1) normalized each image to have zero mean and unit variance, (2) registered them vertically so that their horizons aligned, (3) extracted the luminance channel from each image, (4) averaged over horizontal position (see Fig. 7(a)) and (5) cropped to extract a luminance vector $\mathbf{l}_i$ of length n centred at the horizon. Finally, we computed the first m principal components $\mathbf{u}_i$ (see Fig. 7(b)) of length n over these k vectors.

To estimate the vertical position y_h of the horizon for a target image, we (1) normalize the image to have zero mean and unit variance, (2) extract the luminance channel, (3) average over horizontal position and (4) convolve the resulting vector $\mathbf{l}$ with each of the m principal component vectors to generate m

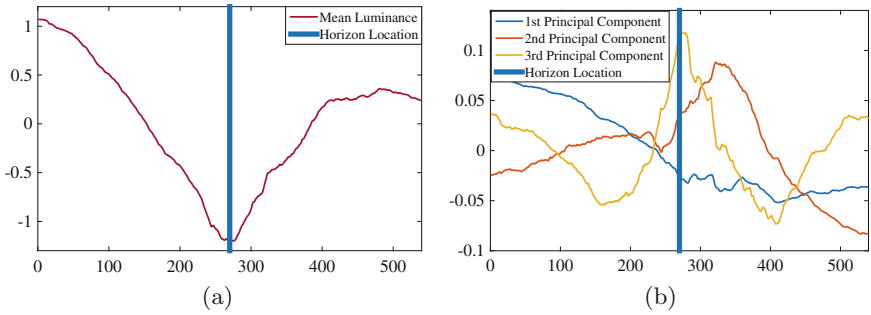

Fig. 7. (a) Mean normalized luminance of horizon training images as a function of vertical displacement from horizon. (b) First three principal components of vertical luminance distribution around horizon location.

projection vectors $\tilde{\mathbf{l}}_i$. Approximating the training data as multivariate normal, the log probability that the horizon lies at vertical location y can be estimated as:

$$\log p\left(y_h = y\right) \propto -\sum_{i=1}^{m} \left(\tilde{l}_i(y) - \bar{\mathbf{l}}^T \mathbf{u}_i\right)^2 / \lambda_i, \tag{3}$$

where $\bar{\mathbf{l}}$ is the mean over the training vectors $\mathbf{l}_i$ and λ_i is the ith eigenvalue. The horizon is then estimated by maximizing this log probability over y.

The two free parameters of this method are the length n of the principal component filters and the number m of these filters to employ. We optimized these parameters by grid search over the training data (Fig. 8), finding $n = 300$ pixels and $m = 6$ to be optimal. Examples of horizon lines estimated on the test dataset are shown in Fig. 10(c–f).

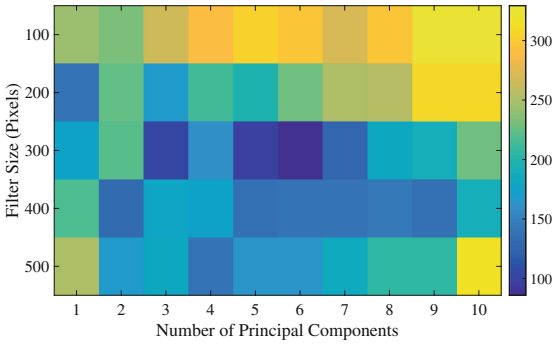

Fig. 8. Mean vertical location error in horizon estimate over training data, as a function of the length n and number m of principal components filters.

3.3 Region of Interest

The road vanishing point and horizon line provide crucial geometric constraints on the location of the road surface. To turn these constraints into an approximate segmentation of the road, we fuse them with a spatial prior that has been conditioned on the estimated vanishing point or horizon. Figure 9(a) shows the spatial prior, learned over the training dataset, without conditioning. The prior is quite diffuse, due not only to the variability in road geometry but also to the variation in camera placement. Figures 9(b–c) show the same prior, computed relative to the vanishing point location for images where a vanishing point can be identified (a) and relative to the horizon line for images where it cannot (c). Note that conditioning on the geometry leads to a more focused and accurate indicator of the road surface location.

In order to segment the road for a novel image, we first estimate the vanishing point. If the estimate is judged reliable, we register the associated prior (Fig. 9(b)) with the estimated vanishing point and label all pixels above a threshold probability p_0 to be road pixels. If the vanishing point is judged to be unreliable we follow the same procedure for the estimated horizon line and prior. If our vanishing point and horizon line estimates were perfectly accurate, a threshold of $p_0 = 0.5$ would maximize proportion correct pixel labelling (road/non-road). Figure 9(d) shows that our road segmentation algorithm performs substantially better than the method of Qian et al. [17], which involved simply thresholding the spatial prior, with no estimation of road geometry. This figure also confirms that a threshold of 0.5 works well in practice and is the value we adopt for road weather classification.

Figure 10 shows representative road segmentation results on the test dataset. There remain some failure modes for both road vanishing point and horizon detection, but generally speaking the results are good, as indicated by the median examples (b).

4 Road Condition Classification

To assess the utility of the improved road segmentation, we use it to define the region of interest (ROI) for road weather classification. In particular, we use an adaptation of the method of Qian et al. [17], which employs scale- and orientation-invariant MR8 filters [23]. These filters produce an 8-dimensional feature vector at each pixel in the ROI. Each ground truth ROI in the training set is divided into 8×8 pixel non-overlapping patches, each of which is then represented by an $8 \times 8 \times 8 = 512$-dimensional feature vector. K-means is then used to cluster the feature vectors from the training data into $k = 74$ textons.

While MR8 is roughly luminance-invariant, luminance can carry information about weather (snow is bright, wet roads tend to be darker). Qian et al. therefore augmented this texton descriptor with a 20-bin histogram of grey level deviations from the mean image luminance.

Qian et al. based their classifier on a naive Bayes model of exponential-χ^2 distances of input vectors from their class-conditional means. Here we take a

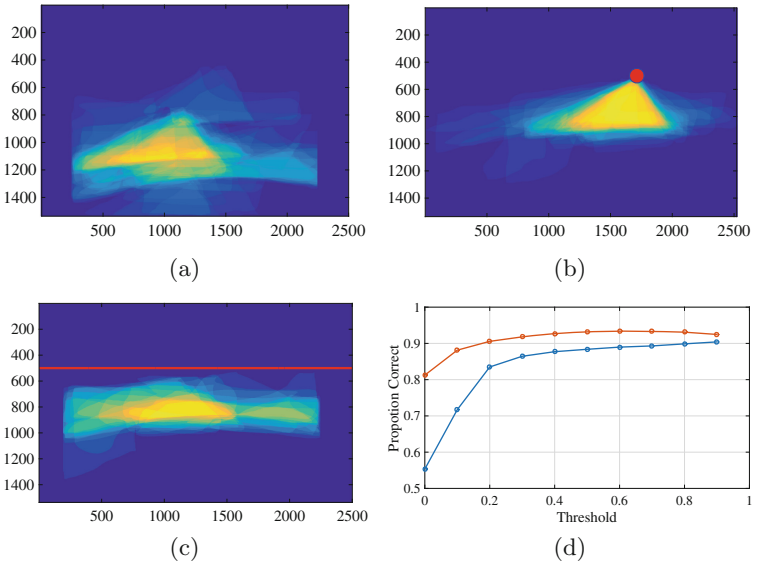

Fig. 9. (a) Prior spatial distribution of road pixels, in absolute image coordinates. (b) Prior horizontally and vertically registered to vanishing point. (c) Prior vertically registered to horizon. (d) Proportion of correctly labelled image pixels for the test set, as a function of the probability threshold p_0.

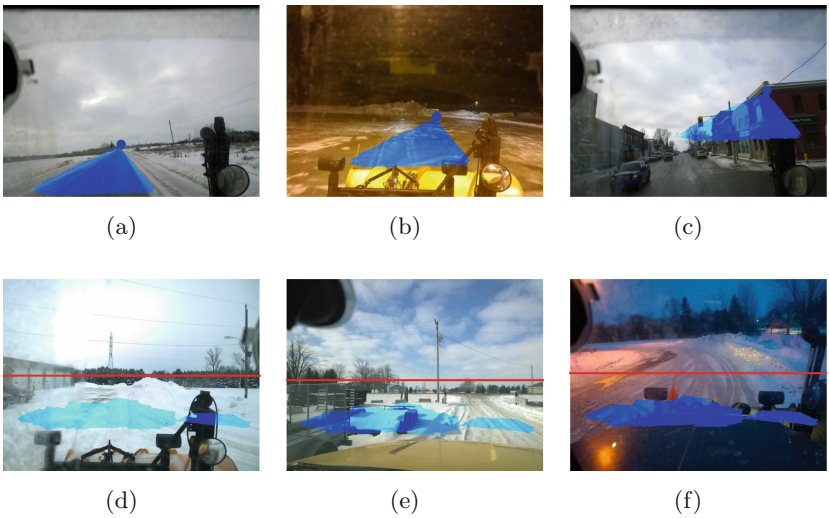

Fig. 10. Example road segmentations on the test dataset. Results are evaluated based on proportion of correctly labelled pixels (road/non-road). (a)–(c) show best, median and worst-case (failure mode) examples for cases where a vanishing point could be estimated. (d)–(f) show best, median and worst-case examples for cases where the vanishing point was deemed unreliable and a horizon estimator was used.

simpler approach, using an SVM with RBF kernel in a one-versus all design (the SVM implementation provided in the MATLAB Statistics and Machine Learning Toolbox). We found this to yield very similar results.

5 Performance Evaluation

Classification results on the test set for two classes (bare vs snow/ice covered), three classes (dry, wet, snow/ice covered) and five classes (dry, wet, snow, ice, packed) are shown in Fig. 11. We evaluate results using three different methods to define the ROI: (1) manual segmentation; (2) the automatic segmentation method proposed here; (3) automatic segmentation using the fixed prior of Fig. 9(a) [17]. As a baseline we also show performance for a classifier that simply selects the highest *a priori* probability. Note that the manual segmentation provides an upper bound on the possible payoff from further improvements to segmentation.

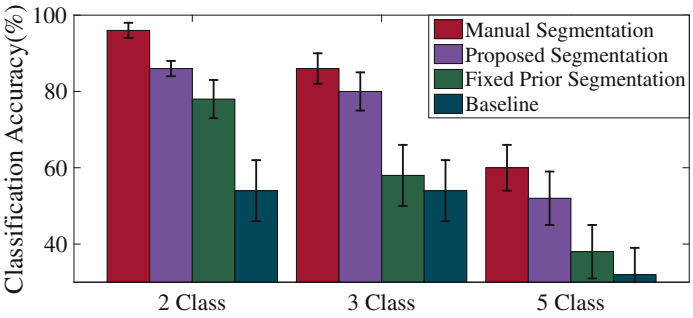

Fig. 11. Results comparison for 2, 3 and 5 classification using manual, Proposed Segmentation, Fixed prior segmentation against random guesses

Our proposed segmentation method achieved an accuracy of 86 % for two-classes, 80 % for three-classes, and 52 % for 5 classes. Given the challenging nature of the dataset, these are promising results.

The proposed segmentation method improves classification accuracy over the fixed segmentation method used by Qian et al. [17] in all cases. Statistical significance of this improvement can be assessed by computing the posterior probability that the underlying probability of correct classification is greater for our proposed method, assuming a flat prior over the performance for the two methods. This yields a posterior probability of 0.81 for the 2-class case, 0.99 for the 3-class case and 0.90 for the 5-class case, corresponding to p-values of .19, .01 and .10 in the language of null-hypothesis testing.

6 Summary

In this paper we have proposed a novel algorithm for road segmentation from uncalibrated dash cameras, to support road weather analysis. The algorithm was designed to operate robustly over a a diverse range of camera poses on both structured (highway/local road) and unstructured roads (parking lots). The approach consist of finding the vanish point, or horizon and fusing with a registered spatial prior. Classification performance on a challenging dataset was 86 % and 80 % for two- and three-class problems, respectively, representing a significant improvement over prior work [17]. Our analysis reveals that further improvements in performance will likely depend on improvements in both the segmentation and classification stages.

In recent years concern has been raised regarding the lack of generalization of traditional supervised road segmentation techniques, which tend to fail when facing situations not covered in the training set. To overcome this issue, attempts have been made to use online learning methods to adapt the parameters of the algorithm to the image data [1,22]. These methods typically assume that the bottom of the image projects from the road surface [1,2,7,22], which unfortunately is not the case for our dataset (Fig. 1). Nevertheless, this is an important issue, and we hope to increase the adaptiveness of our method in the near future.

References

1. Alvarez, J.M., Gevers, T., Lopez, A.M.: 3D scene priors for road detection. In: 2010 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 57–64. IEEE (2010)
2. Álvarez, J.M., Ĺopez, A.M.: Road detection based on illuminant invariance. IEEE Trans. Intell. Transp. Syst. **12**(1), 184–193 (2011)
3. Álvarez, J.M., López, A.M., Baldrich, R.: Shadow resistant road segmentation from a mobile monocular system. In: Martí, J., Benedí, J.M., Mendonça, A.M., Serrat, J. (eds.) IbPRIA 2007. LNCS, vol. 4478, pp. 9–16. Springer, Heidelberg (2007). doi:10.1007/978-3-540-72849-8_2
4. Amthor, M., Hartmann, B., Denzler, J.: Road condition estimation based on spatio-temporal reflection models. In: Gall, J., Gehler, P., Leibe, B. (eds.) GCPR 2015. LNCS, vol. 9358, pp. 3–15. Springer, Heidelberg (2015). doi:10.1007/978-3-319-24947-6_1
5. Casselgren, J.: Road surface classification using near infrared spectroscopy. Ph.D. thesis
6. Dahlkamp, H., Kaehler, A., Stavens, D., Thrun, S., Bradski, G.R.: Self-supervised monocular road detection in desert terrain. In: Proceedings of Robotics: Science and Systems, Philadelphia (2006)
7. De Cristóforis, P., Nitsche, M.A., Krajník, T., Mejail, M.: Real-time monocular image-based path detection. J. Real Time Image Process. **11**, 1–14 (2013)
8. Jokela, M., Kutila, M., Le, L.: Road condition monitoring system based on a stereo camera. In: IEEE 5th International Conference on Intelligent Computer Communication and Processing, ICCP 2009, pp. 423–428. IEEE (2009)
9. Jonsson, P.: Remote sensor for winter road surface status detection. In: Proceedings of 2011 IEEE Sensors, pp. 1285–1288. IEEE (2011)

10. Jonsson, P., Casselgren, J., Thornberg, B.: Road surface status classification using spectral analysis of NIR camera images. IEEE Sens. J. **15**(3), 1641–1656 (2015)
11. Kawai, S., Takeuchi, K., Shibata, K., Horita, Y.: A method to distinguish road surface conditions for car-mounted camera images at night-time. In: 2012 12th International Conference on ITS Telecommunications (ITST), pp. 668–672. IEEE (2012)
12. Kong, H., Audibert, J.Y., Ponce, J.: Vanishing point detection for road detection. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2009, pp. 96–103. IEEE (2009)
13. Kong, H., Audibert, J.Y., Ponce, J.: General road detection from a single image. IEEE Trans. Image Process. **19**(8), 2211–2220 (2010)
14. Lim, S.H., Ryu, S.K., Yoon, Y.H.: Image recognition of road surface conditions using polarization and wavelet transform. J. Korean Soc. Civil Eng. **27**(4D), 471–477 (2007)
15. Oliva, A., Torralba, A.: Modeling the shape of the scene: a holistic representation of the spatial envelope. Int. J. Comput. Vis. **42**(3), 145–175 (2001)
16. Omer, R., Fu, L.: An automatic image recognition system for winter road surface condition classification. In: 2010 13th International IEEE Conference on Intelligent Transportation Systems (ITSC), pp. 1375–1379. IEEE (2010)
17. Qian, Y., Almazan, E.J., Elder, J.H.: Evaluating features and classifiers for road weather condition analysis. In: 2016 IEEE International Conference on Image Processing (ICIP). IEEE (2016)
18. Rasmussen, C.: Grouping dominant orientations for ill-structured road following. In: Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR 2004, vol. 1, pp. I-470. IEEE (2004)
19. Rasmussen, C.: Texture-based vanishing point voting for road shape estimation. In: BMVC, pp. 1–10. Citeseer (2004)
20. Suttorp, T., Bucher, T.: Robust vanishing point estimation for driver assistance. In: IEEE Intelligent Transportation Systems Conference, ITSC 2006, pp. 1550–1555. IEEE (2006)
21. Tal, R., Elder, J.H.: An accurate method for line detection and Manhattan frame estimation. In: Park, J.-I., Kim, J. (eds.) ACCV 2012. LNCS, vol. 7729, pp. 580–593. Springer, Heidelberg (2013). doi:10.1007/978-3-642-37484-5_47
22. Tan, C., Hong, T., Chang, T., Shneier, M.: Color model-based real-time learning for road following. In: IEEE Intelligent Transportation Systems Conference, ITSC 2006, pp. 939–944. IEEE (2006)
23. Varma, M., Zisserman, A.: Classifying images of materials: achieving viewpoint and illumination independence. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) ECCV 2002. LNCS, vol. 2352, pp. 255–271. Springer, Heidelberg (2002). doi:10.1007/3-540-47977-5_17. http://www.robots.ox.ac.uk/ vgg
24. Wang, Y., Teoh, E.K., Shen, D.: Lane detection and tracking using B-snake. Image Vis. Comput. **22**(4), 269–280 (2004)
25. Yang, H.J., Jang, H., Kang, J.W., Jeong, D.S.: Classification algorithm for road surface condition. IJCSNS **14**(1), 1 (2014)