A Comparative Study of Stereo-Matching Algorithms for Road-Modeling in the Presence of Windscreen Wipers

Konstantin Schauwecker, Sandino Morales, Simon Hermann, and Reinhard Klette

Computer Science Department, The University of Auckland
Private Bag 92019, Auckland 1142, New Zealand

***Abstract*—In this study we examine three road-modeling methods, which we evaluate on seven stereo matching algorithms. The road-modeling methods we consider are a B-spline modeling technique based on region-growing and two versions of the popular v-disparity approach. The used stereo algorithms are variations or different parameterizations of belief propagation, graph cut and semi-global matching.**

We examine the performance of the possible combinations of modeling methods and stereo algorithms by comparing the deviation towards a reference road profile. Two evaluation sequences are used, of which one features switched on windscreen wipers that are visible in the the recorded imagery.

Our findings are that the examined B-spline modeling method provides the best results in most cases. In terms of the modeled distance, belief propagation is the most suitable stereo matching method, which also performs best with the wiper sequence. Semi-global matching achieves smaller model deviations, but with drastically reduced modeled distances.

I. INTRODUCTION

Computer vision is gaining importance in the automotive environment. Some modern series-production vehicles are already equipped with computer vision based *driver assistance systems* (DAS). Examples are lane departure warnings, traffic sign recognition and blind spot monitoring, which all employ *monocular vision*. It is also possible to use *binocular vision* for driver assistance purposes, which in addition to normal video imagery can provide depth information for the recorded scene. While this additional information could open new application areas, gathering and analyzing it is challenging.

For understanding a traffic situation from binocular imagery, it is particularly important that we recognize the geometry of the road. Previous research has been conducted on modeling the road profile from the *disparity maps* received through binocular vision and stereo matching [1]–[5]. Knowing the road geometry facilitates segmentation and recognition of obstacles on the road, such as other traffic participants. If we want to use this information in novel computer vision based DAS, it is important that the road-modeling process works robustly under all weather and illumination conditions.

In this research we evaluate three road-modeling techniques using seven stereo matching algorithms. The road modeling algorithms we consider are the region-growing B-spline approach we presented in [5] and two variations of the popular v-disparity method, as introduced in [1]. The stereo matching algorithms we examine are adaptations or different parameterizations of belief propagation (BP), graph cut (GC) and semi-global matching (SGM).

Further, we take a close look on how windscreen wipers disturb the stereo matching algorithms and road-modeling methods. As the common setup for vehicle mounted cameras is behind the windscreen, moving windscreen wipers appear in the recorded imagery. A computer vision based method is required to cope with wipers occluding part of the visible scene in one camera, if it should be applicable on rainy weather conditions. We hence test all methods on how they perform in this situation.

II. RELATED WORK

The common approach for judging the performance of a stereo method is to process stereo pairs for which known ground truths are available. The generated disparity maps can then be compared to the ground truths and the deviation can be quantified. This approach was applied in [6], [7] to evaluate the performance of a large selection of stereo matching algorithms. The used stereo pairs were, however, either synthetic or photographed in a controlled lighting environment. The performance of stereo algorithms on real-world imagery is generally much lower. Hence, the information we gain from an evaluation on such stereo pairs is only of limited use.

We are interested in the stereo matching performance on real-world image data. In particular, we are focusing on stereo sequences recorded from vehicle mounted cameras. Performing an evaluation on such driving sequences is more challenging as ground truths are usually not available. In [8] an evaluation of stereo algorithms on driving sequences is performed, by using a third camera for judging the found stereo correspondences. In the evaluation presented in [9], the stereo matching results are compared to the data from a vehicle mounted LIDAR. To our knowledge, no evaluation has yet been performed on how windscreen wipers in driving sequences influence the stereo matching process.

In this study, we focus on using stereo-vision for modeling the ground manifold. In particular, we are interested in the *vertical road profile* that for a given road describes the elevation change over distance. To model the vertical profile of non-flat roads, previous attempts in literature used envelopes of lines [1], quadratic curves [2], clothoids [3], and recently B-spline curves [4]. In [5] we published a new method for modeling the vertical road profile with a B-spline curve, which relies on a region growing process. Similarly, a region-growing technique is used in [2] for modeling the road with a quadratic curve. We will use our previous region-growing B-spline method in

this research as one road-modeling technique for evaluating the performance of the used stereo algorithms.

III. ROAD MODELING METHODS

The region-growing B-spline method (RGBS) that we use for road-modeling, first selects a small group of near-range points from the set of 3D points gained through stereo triangulation. A B-spline curve is then fitted to those points and serves as the first hypothesis for the vertical road profile. In subsequent iterations, more points are selected that are close to the previous hypothetical road model. A new and potentially more accurate curve is then fitted to the increased set of points, which is repeated for every iteration step.

To increase accuracy, a set of constraints is introduced that limit the number of points, which can be selected in one iteration. For example, points are required to be connected in the xy-plane, and their distance in z-direction must not exceed a threshold. Further constraints penalize high gradients or curvatures, and enforce a flat gradient at the camera position (as we assume that the car is standing flat on the road).

The v-disparity method, which is another road-modeling method we consider in this study, was initially introduced in [1]. In this approach, the disparity map is first transformed into a virtual image, in which the intensity of a pixel with the coordinates (u, v) is proportional to the frequency of the disparity u in image row v. This virtual image is called *v-disparity image*, and in it the disparity values corresponding to the visible road are likely to form a curve. This curve is detected with a Hough transform that delivers a set of best matching straight lines, which are then combined by selecting either the upper or lover envelope. In addition to this envelope-based approach (VENV), which matches the method initially published in [1], we also examine a simplified version that only selects the best matching straight line (VPLAN).

IV. EVALUATION METHOD

The methodology we use for evaluating the quality of a generated road model matches the method we previously applied in [5]. We manually create a set of mask images that identify the visible road in the left image (the reference image) of an evaluation stereo pair. Only frames are selected in which no wiper is visible in the right image. If a wiper is visible in the left image, the region covered by it is excluded in the corresponding mask. In total, 40 masks were created for each of the two evaluation sequences.

We use one such mask-image and the disparity map obtained through the examined stereo algorithm, to extract a reference road profile. First, the mask is used to extract all road-pixels from the disparity map. We then select one best matching disparity value for every image row in which the road is visible. For this task, we use the median of the extracted disparity row, rather than the average, to make the method more robust towards outliers. A generated road model is then compared to this set of road profile points, by calculating the distance in y-direction between a road profile point and the road model curve. We take the sum of those distances, which is the *sum of absolute differences* (SAD) between the y-coordinates, as measure for the goodness of a model.

The number of compared road profile points depends on the length of the modeled and visible road section. This means that a different number of points might be used each time we calculate an SAD value, which makes it impossible to compare the results for two different stereo pairs. To make our measure independent of the number of compared profile points, we divide the SAD value by the total number of points we evaluated, and receive a new *relative SAD*.

V. STUDY OUTLINE

For our evaluation we use two stereo sequences that were both recorded on the same road. The considered road has a wavy profile that should be ideal for evaluating road-modeling techniques. The first sequence, to which we refer as `Midday`, was recorded when the sun was close to its zenith and features a good illumination without reflections. The second sequence, which we call `Wiper`, was recorded shortly afterwards on the same road, with switched on windscreen wipers (but no rain). The illumination conditions have not changed noticeably between both recordings. The used sequences have been published online as Set 9 of the EISATS website [10].

An example image from each sequence is shown in Fig. 1a and 2a. By recording both sequences on the same road, we allow for an easy comparison of the gained results. This enables us to draw conclusions on how the presence of wipers influences the performance of the tested algorithms.

We evaluate the considered road-modeling methods on a selection of well-performing stereo algorithms. We chose to use BP [11] and GC [12], as they are among the most popular global energy minimization methods. Our implementations are based on the code published along with [13] and [14].

Unfortunately, both algorithms are too slow to enable real-time stereo-matching, which would be required for driver assistance purposes. We hence chose to further evaluate SGM, which was initially published in [15]. Unlike BP and GC, the SGM algorithm does not perform a global optimization but only optimizes along a set of scan lines. Nevertheless, SGM provides excellent results that may well compete with the ones gained though global methods such as BP or GC.

For every algorithm, we examine several different variations or different parameter settings. In total, we employ three BP, two GC and two SGM variants. This gives us a total number of seven distinct stereo-matching algorithms that we examine in this study. Those algorithms are:

BP-EPE BP with a data-term based on the gradient end point error (equivalent to the data-term evaluated in [16]).

BP-SEPE BP with a gradient end point error data-term and a low data-term weighting factor λ. This leads to smooth disparity maps, but blurred depth discontinuities.

BP-CEN BP with a data-term based on the *census transform* [17]. The census transform is an image transformation that has previously been used for robust stereo-matching. Recent studies have shown that the census transform is very robust against illumination differences [16], [18].

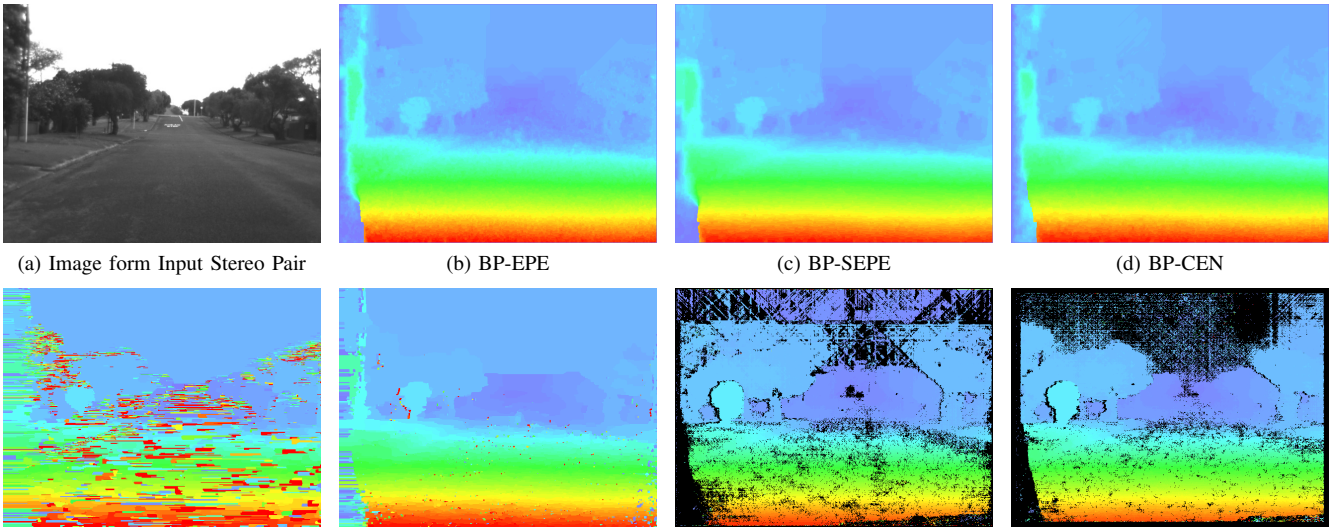

(a) Image form Input Stereo Pair (b) BP-EPE (c) BP-SEPE (d) BP-CEN

(e) GC-EPE (f) GC-CEN (g) SGM-EPE (h) SGM-CEN

Fig. 1: (a) Image of a stereo pair from `Midday` and (b-h) corresponding disparity maps.

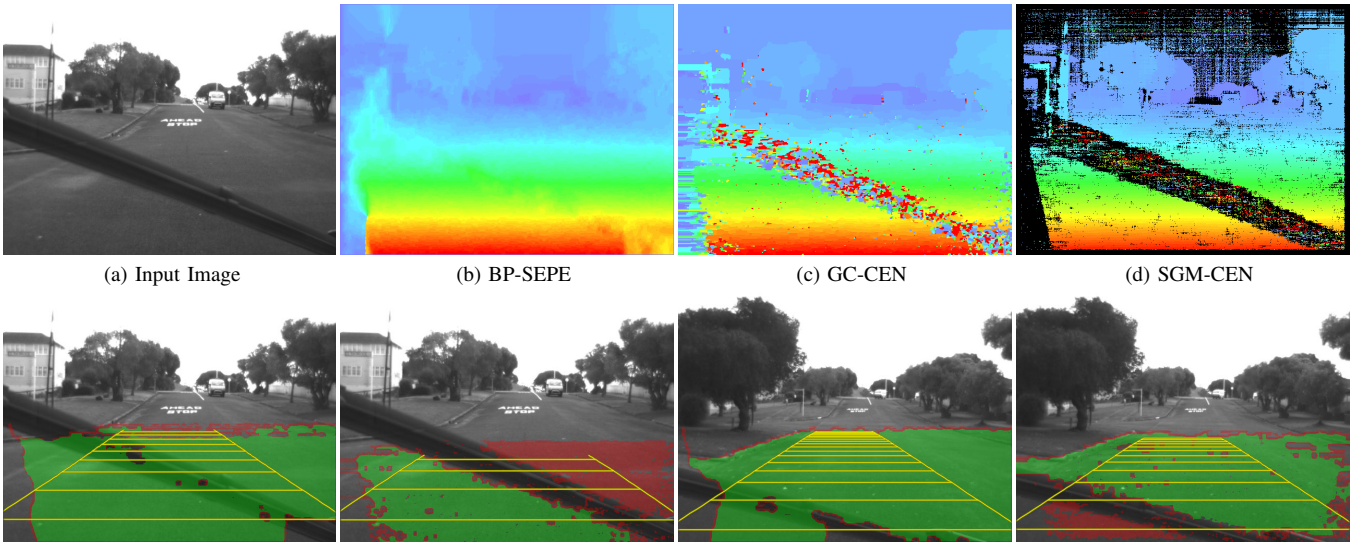

(a) Input Image (b) BP-SEPE (c) GC-CEN (d) SGM-CEN

(e) BP-SEPE (f) GC-CEN (g) BP-SEPE (h) GC-CEN

Fig. 2: (a) Image of a stereo pair from `Wiper` and (b-d) corresponding disparity maps; (e-h) regions selected by RGBS.

GC-EPE GC algorithm based on α-expansion and with a data-term based on the gradient end point error.

GC-CEN GC based on α-expansion and with a data cost term based on the census transform.

SGM-EPE SGM algorithm using a data-term based on the gradient end point error.

SGM-CEN SGM using a data-term based on the census transform.

VI. EVALUATION RESULTS

We applied the seven selected stereo-matching algorithms to the two evaluation sequences. This resulted in a total of 14 disparity sequences, which we each processed with the three road-modeling methods we consider in this study. Hence, we performed a total of 42 test runs. The results we gained are summarized in the following subsections.

A. *Sequence* `Midday`

Sequence `Midday` provides excellent conditions for stereo-matching and we thus received accurate results with this sequence for all examined stereo-matching algorithms. An example for the disparity maps generated with each algorithm for this sequence is shown in Fig. 1.

The disparity maps obtained with all examined BP variants appear very similar and are all of a very high quality. It is difficult to spot differences between the end point error and

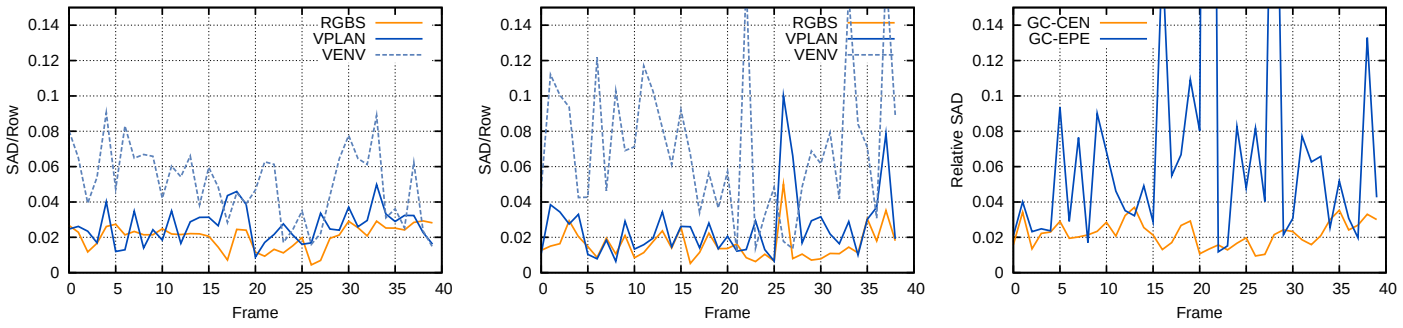

(a) (b) (c)

Fig. 3: Results on Sequence `Midday` for (a) BP-CEN, (b) SGM-CEN, and (c) GC-CEN / GC-EPE with B-splines.

the census transform based BP versions. The differences between BP-EPE and BP-SEPE are slightly smoother transitions between the disparity levels.

Considering the similar disparity maps, all road-modeling methods delivered similar results for each BP algorithm. Fig. 3a shows the road model deviation for each evaluated frame, when using BP-CEN. In this diagram, we see that RGBS predominantly delivers the best results throughout the sequence. Apart from a few large deviations, VPLAN is not far behind, while VENV clearly delivers the worst results for most of the stereo pairs in the sequence.

Despite the very similar performance of the end point error and census transform based data-terms for BP, there appears to be a large difference when using GC. While GC-CEN delivers relatively good results with few noisy mismatches, GC-EPE contains a large amount of wrong disparity values. This lower performance is clearly visible with our evaluation method. Figure 3c shows the results we received with RGBS for GC-CEN and GC-EPE. However, this bad performance might be exaggerated because the reference road profile extraction is not as accurate for the noisy GC-EPE disparity maps.

If we compare the disparity maps from Fig. 1g and 1h that we received with the tested SGM variants, we notice a difference in the density of disparity values. Compared to SGM-EPE, the disparity maps from the SGM-CEN algorithm have a higher density for the road region. For RGBS, the disparity map density has a strong effect on the modeled distance due to the region growing process.

Figure 3b shows the evaluation results we received for SGM-EPE, which shall serve as an example for the performance of SGM on Sequence `Midday`. In this figure we see that VENV performs much worse than for the other stereo algorithms. The same observation was also made for SGM-CEN. The performance of VPLAN is better but suffers from several high outliers. For this stereo algorithm, RGBS again produces the best results.

A summary of the results gained with each stereo algorithm and road-modeling method on this sequence is shown in the upper half of Table I. This table contains the mean and median relative SAD values and maximum modeled distances that we received for the entire sequence. For each stereo algorithm, the mean and median of the best performing road-modeling method has been printed in bold.

Except for GC-EPE (and in case of the median also SGM-EPE), RGBS performed best for all stereo algorithms, while VENV always performed worst. The bad performance of RGBS with GC-EPE is not surprising, as the large amount of noise in the disparity maps impedes the region-growing process, while the Hough transform used by VENV and VPLAN is still be able to find a robust fit.

In terms of the average model deviation, BP-CEN and BP-EPE deliver equally good results with RGBS; the lowest deviation is, however, caused by BP-SEPE. For GC, the version performing best with RGBS is obviously GC-CEN, whose average deviation is slightly higher than for BP-EPE and BP-CEN, but the median deviation is less. However, the modeled distance with GC-CEN is approximately 5 m less than for the three BP algorithms.

The SGM algorithm performing best with RGBS is SGM-CEN, which has the lowest average and median deviation of all tested algorithms. However, the average modeled distance is approximately 9 m less than for the BP algorithms, and 6 m less than for SGM-EPE.

In conclusion, BP-SEPE provided the highest maximum modeled distance, while SGM-CEN caused the lowest average and median model deviation. However, the modeled distance for SGM-CEN was much lower than for BP-SEPE. It is likely that if we limit the modeling distance for BP-SEPE, we will too receive a lower deviation as the close-range road profile can be modeled with higher accuracy. It is thus wise to pay more attention to the maximum modeled distance than to the model deviation.

B. Sequence `Wiper`

Figures 2b to 2d show the disparity maps we received with BP-SEPE, GC-CEN and SGM-CEN for one stereo pair from Sequence `Wiper`, in which a wiper is visible in the left image. The wiper can only be seen by one camera, and it is thus impossible to find any matching pixel pairs for the image region it covers. In the shown disparity maps, we can see that each algorithm responds differently to this situation.

BP deals with this problem by smoothly interpolating the

Sequence	Algorithm	RGBS deviation		VPAN deviation		VENV deviation		Max. modeled dist. / m	
		Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Med.
Midday	**BP-EPE**	**0.0205**	**0.0225**	0.0264	0.0251	0.0504	0.0497	29.3	29.4
	BP-SEPE	**0.0191**	**0.0209**	0.0266	0.0263	0.0483	0.0456	30.1	30.8
	BP-CEN	**0.0205**	**0.0217**	0.0266	0.0261	0.0504	0.0512	29.4	29.4
	GC-EPE	0.0798	0.0470	**0.0776**	**0.0399**	0.0916	0.0527	13.7	15.4
	GC-CEN	**0.0222**	**0.0214**	0.0252	0.0252	0.0478	0.0484	25.2	25.7
	SGM-EPE	**0.0252**	0.0237	0.0267	**0.0236**	0.0413	0.0395	14.0	18.1
	SGM-CEN	**0.0154**	**0.0137**	0.0251	0.0200	0.0690	0.0643	20.6	20.6
Wiper	**BP-EPE**	**0.0231**	**0.0236**	0.0247	0.0231	0.0494	0.0501	27.9	25.7
	BP-SEPE	**0.0223**	**0.0216**	0.0259	0.0254	0.0484	0.0468	31.1	30.8
	BP-CEN	**0.0230**	**0.0230**	0.0259	0.0238	0.0510	0.0499	30.5	28.0
	GC-EPE	**0.7372**	0.2632	0.7742	**0.1290**	2.2340	0.1640	14.7	15.4
	GC-CEN	0.0571	**0.0244**	**0.0279**	**0.0244**	0.0506	0.0506	25.7	25.7
	SGM-EPE	**0.0218**	**0.0192**	0.0245	0.0212	0.0428	0.0429	15.5	19.3
	SGM-CEN	**0.0192**	**0.0209**	0.0266	0.0257	0.0542	0.0475	19.8	20.6

TABLE I: Average and median deviation for examined algorithms and road-modeling methods on both evaluation sequences.

disparity values that surround the wiper. The result is a disparity map in which the occluded part of the road still has fairly accurate disparity values, due to the fact that they were interpolated form the closest non-occluded road sections. This smooth interpolation is caused by the smoothness-term used in BP that penalizes non-gradual disparity changes.

The disparity map generated with GC for the same stereo pair shows a different behavior. For the region covered by the wiper, GC finds mostly random disparity values. SGM on the other hand produces only few disparity values for this region, and the values that it selects appear random as well. The reason why the used SGM implementation performed differently than GC for this stereo pair is that it includes a *left-right consistency check* [19] and removes inconsistent pixels. A left-right consistency check is, however, not specific to SGM and can just as well be included in a GC algorithm.

For RGBS, the behavior of BP is the most favorable. Because BP produces a smooth interpolation of the disparity values surrounding the wiper, the region-growing method will be able to select pixels from the wiper region and thus pass through to the road section cut off by the wiper. For GC and SGM, the erroneous or missing disparity values will prevent any further expansion of the selected road region. The road that is depicted beyond the wiper will hence never have the chance of being selected. This behavior is shown in Fig. 2e and 2f for BP-SEPE and GC-CEN.

In case where an initial region above the wiper is selected, the wiper does not have an impeding effect for GC or SGM, as shown in Fig. 2h. Rather, in this case the region-growing process does not select the region below the wiper. This region is usually redundant, as close road sections to the opposite side of the wiper can still be selected.

In terms of the modeling error, there seems to be only a small difference between BP and GC or SGM, as can be seen in Fig. 4a and 4b for BP-SEPE and GC-CEN (if we neglect the high error observed for the last two frames). The modeled distance on the other hand is significantly reduced for several frames when using GC instead of BP, as is shown in Fig. 4c.The fact that we didn't observe a significantly higher model deviation reveals that our road models do not degrade in accuracy, but only in modeled distance.

The summary for all stereo algorithms and road-modeling methods is shown in the bottom half of Table I. If we compare those results to the ones that we received for the previous sequence, we realize that RGBS performs worse for most stereo algorithms (SGM-EPE is an exception). This is, however, not true for VENV and VPLAN that in many cases performed slightly better, which can be credited to the robustness of the used Hough transform. Nevertheless, RGBS still performs best for most stereo algorithms.

If we compare the maximum modeled distances between both sequences, we realize that they do not differ much. Given our previous observation on the limited region expansion, this appears counter intuitive. However, most stereo pairs do not depict a wiper, and for those stereo pairs that do, only few provide a scenario where the modeled distance is significantly reduced. Hence, the the maximum modeled distance is only slightly smaller for most algorithms.

For BP-EPE and BP-CEN, the modeled distance has in fact increased. As the wiper occlusion does not form an obstacle for BP, this is not surprising. The smooth interpolation caused by BP for the road section occluded by the wiper might even facilitate the region-growing process for RGBS.

The relative performance of the individual algorithms did not change significantly from Sequence `Wiper`. BP-SEPE still performs best among the BP algorithms and achieves the highest modeled distance. Likewise, SGM-CEN still causes the lowest model deviation, but the advance is much less than it was for Sequence `Wiper`, while the difference in modeled distance remains approximately equal.

VII. CONCLUSIONS

We have performed an evaluation of the three considered road-modeling methods using seven stereo algorithms and two evaluation sequences. In most of our experiments, RGBS provided the most accurate models. If we compare VENV and VPLAN, we realize that VENV performed worse than VPLAN for all sequences and stereo algorithms. There is hence no

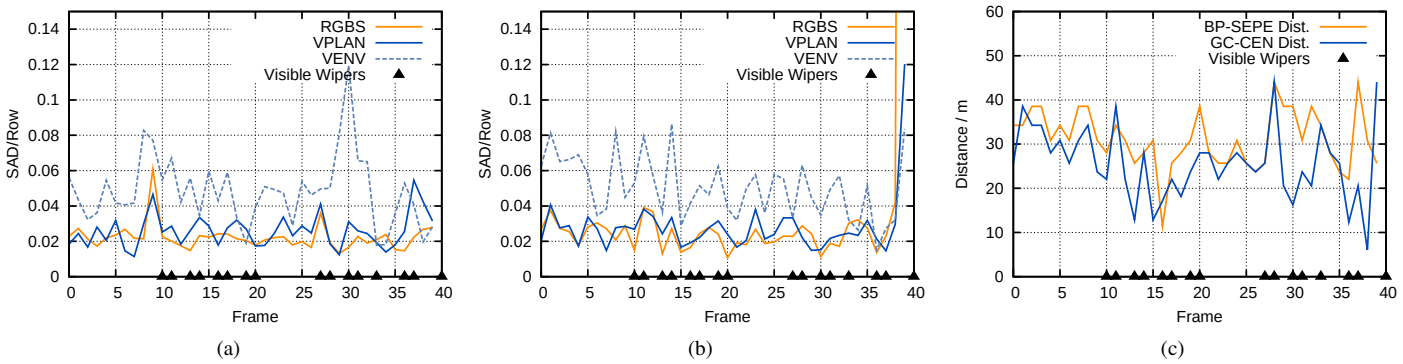

(a) (b) (c)

Fig. 4: Deviation for `Wiper` with (a) BP-SEPE and (b) GC-CEN, and (c) maximum distance with BP-SEPE and GC-CEN.

reason for ever using VENV in favor of VPLAN. However, this is likely to be caused by the way the envelope function is extracted from the v-disparity image, which matches the method used in [1]. In principle, an envelope function can model more complex road geometries and should be able to provide more accurate results.

It is difficult to choose a best performing stereo algorithm from the experiments we conducted in this study. Clear winners are the BP algorithms in the case when a windscreen wiper is visible. Losers are the GC variations that failed to convince on any sequence, and especially GC-EPE, which always performed worst. In terms of the maximum modeled distance, BP-SEPE is the best performing algorithm. This algorithm produces smoother disparity maps than the other examined BP versions. It is, however, possible to also create a smoother BP algorithm that is based on the census transform. What our experiments show is that this increased smoothness is beneficial for the region-growing process.

In terms of the model deviation, SGM-CEN produced the best results for RGBS with both evaluation sequences. However, in both cases the maximum modeled distance was far lower than for BP-SEPE. Because the close range road profile can be modeled with higher accuracy, we cannot conclude that SGM-CEN provides the better road models. As BP-SEPE caused a fairly low model deviation and on average the farthest modeled distance, this algorithm is likely to be the most suitable stereo algorithm for RGBS. Considering the similar performance of BP-EPE and BP-CEN, it is likely that a smooth census-based BP algorithm would produce similar results to BP-SEPE. Such an algorithm should be included in future evaluations.

References

[1] R. Labayrade, D. Aubert, and J.-P. Tarel, "Real time obstacle detection in stereovsion on non flat road geometry through "v-disparity" representation," in *IEEE Intelligent Vehicle Symposium (IV)*, vol. 2, 2002, pp. 646–651.

[2] F. Oniga, S. Nedevschi, M. Marc, and B. Thanh, "Road surface and obstacle detection based on elevation maps from dense stereo," in *IEEE Conference on Intelligent Transportation Systems (ITSC)*, 2007, pp. 859–865.

[3] S. Nedevschi, R. Danescu, and D. Frentiu, "High accuracy stereovision approach for obstacle detection on non-planar roads," in *IEEE Intelligent Engineering Systems (INES)*, 2004, pp. 292–297.

[4] A. Wedel, H. Badino, C. Rabe, H. Loose, W. Franke, and D. Cremers, "B-spline modeling of road surfaces with an application to free-space estimation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 10, no. 4, pp. 572–583, 2009.

[5] K. Schauwecker and R. Klette, "A Comparative Study of Two Vertical Road Modelling Techniques," in *Computer Vision in Vehicle Technology: From Earth to Mars (CVVT:E2M)*, 2010.

[6] D. Scharstein and R. Szeliski, "A taxonomy and evaluation of dense two-frame stereo correspondence algorithms," *International Journal of Computer Vision*, vol. 47, no. 1, pp. 7–42, 2002.

[7] R. Szeliski, R. Zabih, D. Scharstein, O. Veksler, V. Kolmogorov, A. Agarwala, M. Tappen, and C. Rother, "A comparative study of energy minimization methods for markov random fields with smoothness-based priors," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 30, no. 6, pp. 1068–1080, 2007.

[8] S. Morales and R. Klette, "A third eye for performance evaluation in stereo sequence analysis," in *Computer Analysis of Images and Patterns*, vol. 5702, 2009, pp. 1078–1086.

[9] ——, "Ground Truth Evaluation of Stereo Algorithms for Real World Applications," in *Computer Vision in Vehicle Technology: From Earth to Mars (CVVT:E2M)*, 2010.

[10] The University of Auckland, "Multimedia Imaging Technology Portal – EISATS," http://www.mi.auckland.ac.nz/EISATS, 2011.

[11] J. Sun, N. Zheng, and H. Shum, "Stereo matching using belief propagation," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, pp. 787–800, 2003.

[12] Y. Boykov, O. Veksler, and R. Zabih, "Fast approximate energy minimization via graph cuts," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 23, no. 11, pp. 1222–1239, 2002.

[13] P. Felzenszwalb and D. Huttenlocher, "Efficient Belief Propagation for Early Vision," *International Journal of Computer Vision*, vol. 70, no. 1, pp. 41–54, 2006.

[14] V. Kolmogorov and R. Zabih, "What energy functions can be minimized via graph cuts?" *European Conference on Computer Vision (ECCV)*, pp. 185–208, 2002.

[15] H. Hirschmüller, "Accurate and Efficient Stereo Processing by Semi-Global Matching and Mutual Information," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 2, 2005, pp. 807–814.

[16] S. Hermann, S. Morales, T. Vaudrey, and R. Klette, "Illumination Invariant Cost Functions in Semi-Global Matching," in *Computer Vision in Vehicle Technology: From Earth to Mars (CVVT:E2M)*, 2010.

[17] R. Zabih and J. Woodfill, "Non-parametric local transforms for computing visual correspondence," *Computer Vision – ECCV'94*, vol. 801, pp. 151–158, 1994.

[18] H. Hirschmüller and D. Scharstein, "Evaluation of stereo matching costs on images with radiometric differences," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 31, no. 9, pp. 1582–1599, 2008.

[19] P. Fua, "Combining stereo and monocular information to compute dense depth maps that preserve depth discontinuities," in *International Joint Conferences on Artificial Intelligence (IJCAI)*, 1991, pp. 1292–1298.