Honda Research Institute EU
Honda Research Institute Europe GmbH
https://www.honda-ri.de/

Vision-based Lateral Position Improvement of RADAR Detections

Morimichi Nishigaki, Sven Rebhan, Nils Einecke

2012

Preprint:

This is an accepted article published in ITSC. The final authenticated version is available online at: https://doi.org/[DOI not available]

2012 15th International IEEE Conference on Intelligent Transportation Systems
Anchorage, Alaska, USA, September 16-19, 2012

Vision-based Lateral Position Improvement of RADAR Detections

Morimichi Nishigaki, Sven Rebhan and Nils Einecke

***Abstract*— Detecting vehicles ahead by sensors mounted on an ego-vehicle is an essential element of Advanced Driver Assistance Systems (ADAS). Although millimeter-wave radar sensors are very robust at detecting vehicles, the lateral position resolution is very low. However, a more precise lateral position would not only improve the performance of existing ADAS, but it would also enable a wide range of additional applications. In this paper, we propose a method for improving the lateral position of radar sensor detections by fusing them with a vision sensor. Our method is based on a robust symmetry axis detection of vehicles' rear sides. In contrast to the typical approach of detecting symmetry by means of edges, we use a patch-matching-based approach together with a robust cost function. The initial lateral position of the radar detection is improved by conducting a local symmetry search in the image region corresponding to the objects detected by the radar. The experimental results show that the proposed method is able to improve the lateral position accuracy by a factor of seven. Furthermore, additional consistency checks discover miss-detections of the symmetry axis, thus ensuring that the lateral position never drops below the accuracy of the radar detection.**

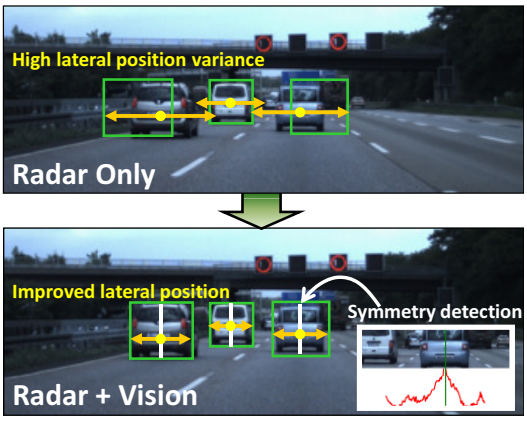

Fig. 1. Lateral position improvement by means of vision: Radar detections have a notoriously bad lateral position. A symmetry analysis in image regions corresponding to the radar detections improves the accuracy of the lateral position by aligning the radar detections to the center of the detected vehicles.

I. Introduction

Detecting surrounding traffic participants is a crucial prerequisite for all Advanced Driver Assistant Systems (ADAS). Apart from pure detections, most systems also benefit from an accurate lateral and longitudinal positioning of vehicles in front of the ego-vehicle. Though a radar sensor, employed in most modern ADAS, is able to detect those vehicles and measure their distance reliably, it suffers from a poor lateral position accuracy. Often a strong temporal filtering is used to compensate this handicap. However, this strong filtering leads to an increased latency in the reported lateral position, which is disadvantageous for functions depending on precise lateral positioning or movements of vehicles. Contrary to the aforementioned approach, the fusion of radar results with visual information of a camera promises both increased lateral precision and a low latency.

Applying radar-vision fusion in the automotive domain is well-known. The authors of [1] describe a system that uses visual information to validate radar detections. However, the position improvement is not investigated and an implicit vehicle model is used due to the classifier. In [2] symmetry is employed to validate radar detections and to improve the position accuracy. The authors report that "... *the position and size [of vehicles] are detected with an extremely high precision (with an error lower than 40 centimeters) ...*"[2]. Even though these are convincing results for the precision of vehicle positions, it is reported that there are problems with traffic and complex situations. Those two issues are likely related to unhandled partial or total occlusion destroying the symmetry property of the occluded vehicle. Another issue reported by Bombini *et al.* are illumination problems originating from the symmetry algorithm itself, which is described in [3], [4] and [5]. The algorithm proposed there uses contour symmetry based on binarized edges of the vehicle. By doing so, the symmetry calculation is prone to failing edge detection caused by strong reflections or low contrast. Furthermore, the threshold used for edge binarization is an additional parameter strongly influencing the resulting symmetry. Similar problems exist for the approaches in [6] and [7] which are also proposing edge-based contour symmetry. In [8] the author describes a symmetry detection that is based on contour symmetry, gray-level symmetry and the maximum horizontal line length. Again, the approach suffers from failing edge detections due to illumination common in real-world images. Furthermore, the contour computation heavily depends on the thresholds regarding a pixel as a contour pixel or not. Contrary to the contour-based symmetry, Zielke *et al.* [9] operate on the intensity values of an image. Here, illumination effects like reflections also affect the symmetry measure because gray values in the image differ strongly, especially on rear windows of vehicles. Except [2], all the aforementioned approaches search for symmetry values on the overall image, neglecting the high

M. Nishigaki is with Honda R&D Europe (Deutschland) GmbH, Carl-Legien-Strasse 30, D-63073 Offenbach/Main, Germany morimichi_nishigaki@de.hrdeu.com

S. Rebhan and N. Einecke are with Honda Research Institute Europe GmbH, Carl-Legien-Strasse 30, D-63073 Offenbach/Main, Germany {sven.rebhan,nils.einecke}@honda-ri.de

978-1-4673-3063-3/12/$31.00 ©2012 IEEE

potential for faster processing by selectively processing only the regions around a radar detection. Furthermore, none of the approaches explicitly handles occlusions in the image, destroying the symmetry property of vehicles, nor stabilizes the detected symmetry axis over time.

In this paper, we propose a comprehensive system refining the lateral position of vehicles detected by the radar using a novel symmetry detection approach (see Fig. 1). Contrary to the contour-based approaches, we employ a correlation-based measure that efficiently makes use of even weak texture information on vehicles. An interesting additional property of correlation-based methods is that they are independent from illumination effects to a large extend. Apart from the novel symmetry measure, we also propose means to care about potentially difficult situations such as vehicle-vehicle occlusion and incorporate temporal integration. We show that our approach yields a high precision in the lateral position making a strong temporal filtering superfluous.

In the next section we explain the overall system structure and its building blocks in more detail. Afterwards, we evaluate our system using about one hour of driving on a highway. Eventually, we conclude this paper and give an outlook on future work.

II. Lateral Position Improvement

The complete system we use for improving the lateral position of radar targets consists of several modules. Figure 2 gives an overview of the different modules and their interaction.

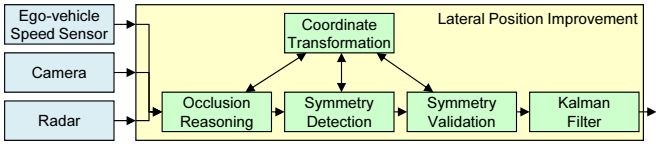

Fig. 2. Block diagram of the system: Radar data, camera image, and ego-vehicle speed are the input to the system. First, an *occlusion reasoning* selects only non-occluded radar targets. The lateral position of these radar targets are improved by means of *symmetry detection*. In order to make the whole system robust, an additional *symmetry validation* analysis is performed on the results of the *symmetry detection* which rejects implausible results. Eventually, a *Kalman filter* tracks position, velocity and acceleration over time.

As can be seen in the block diagram, prior to a symmetry calculation the radar targets are checked for occlusion in the *Occlusion Reasoning* module. As input, the *Symmetry Detection* module receives all non-occluded radar targets as well as the images from the front camera and the ego-vehicle speed sensor data. The input from the radar data consists of lateral and longitudinal positions of the radar targets. The camera images and the radar are synchronized, i.e. all radar targets from one radar time frame are visible (if not occluded by other objects) in the corresponding camera image frame. By combining the information of the radar with the camera images the lateral position of the detected objects is improved. The main idea here is to locate the symmetry axis in the local image area that corresponds to a radar detection. While radar detections might arise from any lateral point on a car, a visual symmetry axis is always found at the center of a car's rear side. Hence, the symmetry axis provides a much more accurate position signal. In order to prevent miss-alignments the *Symmetry Validation* module checks the found symmetry axis for plausibility and consistency.

For converting coordinates between the camera image and the radar a *Coordinate Transformation* module is used. This module is calibrated in advance. Eventually, the refined position information is fed into a *Kalman Filter* module to further stabilizing the detected objects over time and to estimate object velocity and acceleration.

A. Occlusion Reasoning

Before the actual symmetry axis search takes place, some pre-processing is performed in order to remove unsuitable radar detections. First, static objects are removed. These are identified by checking if the relative speed of an object is close to ego-vehicle speed in the opposite direction. Second, for each radar target a rectangular region of interest (ROI) in the camera image is computed. This ROI is supposed to approximate the visual area of a vehicle and is computed by means of a pre-defined vehicle width and height as well as the known radar mount position and the camera calibration parameters.

In a third step, the complete set of ROIs is used to identify occluded objects. This is necessary because calculating the symmetry axis of a laterally occluded object is prone to gross errors. In these cases it is more efficient to rely on the radar information and the temporal evidence from the Kalman filter. For checking if an object is occluded by another one, a margin is added to the left and right of each ROI. The margin is necessary to cope with the low lateral position accuracy of the radar. After the margin has been added, it is checked if any other region that is closer to the ego-vehicle is horizontally overlapping with the margin-extended ROI. In this case, the object corresponding to this ROI is considered as occluded. The distances for checking the "closer to the ego-vehicle" relation are available by the radar measurement. Figure 3 displays an example. The ROI of object 3 with the margins on the left and right is overlapping with the ROI of object 2. Therefore object 3 is considered as occluded.

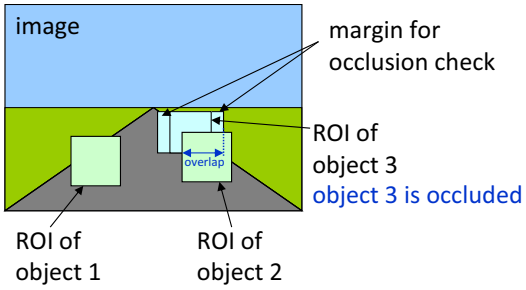

Fig. 3. Identification of occluded objects: In this example the extended ROI of object 3 is partially overlapping with the ROI of object 2. Since object 2 is closer to the ego-vehicle than object 3, the overlap actually means that object 3 is partially occluded by object 2. Hence, a symmetry calculation for object 3 is prone to errors and thus skipped.

B. Symmetry Detection

In order to refine the lateral position of a valid radar target by means of vision we define a search area for the symmetry axis detection as depicted in Fig. 4. The reason for limiting

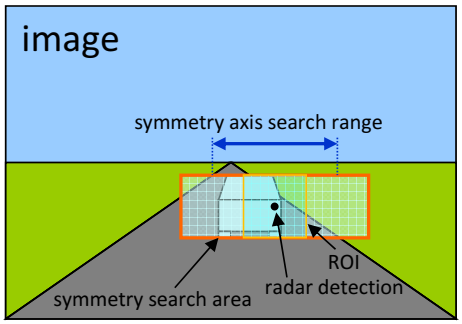

Fig. 4. Symmetry search area: The ROI of a radar target is determined based on the projection of the radar target into the camera image and the corresponding object distance. The symmetry search area is centered at the ROI and has the same height like the ROI but three times the width of the ROI.

the search area is mainly for computational efficiency but also for robustness. As the radar detection beam can be reflected by any point of the vehicle, the radar position can differ up to half a vehicle width from the symmetry axis which we assume to lie in the center of a vehicle rear or vehicle front. Together with the width of the ROI this leads to a minimal image area of twice the width of a vehicle. However, as we use a mean vehicle width the search area might be too small for some vehicles like trucks. Hence, we use a search range that is three times wider than the ROI width and centered at the ROI.

For finding the symmetry axis, a window of the same size as the ROI is swept over the symmetry search area and for each position a symmetry correlation value is calculated (Fig. 5). In contrast to other approaches we do not base

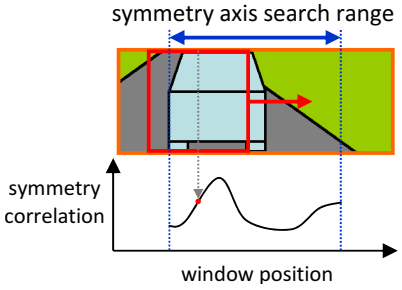

Fig. 5. Symmetry correlation curve: For finding the symmetry axis of a radar target, a filter window of the same size as the object ROI is shifted over the symmetry search area (see Fig. 4) and a symmetry correlation value is computed at each shift position. This results in a correlation curve whose peaks denote the potential position of the vehicle's symmetry axis.

our symmetry value calculation on visual edges. Instead we calculate the symmetry value in a patch-match fashion as described in the following and depicted in Fig. 6.

For the symmetry value calculation, first, the image patch within the filter window is split at the hypothetical symmetry axis in the center. Second, the right part of the image patch



Fig. 6. Computation of the symmetry value of an image patch: First, the image patch is divided in a left and a right part. Then the right part is mirrored and finally compared to the left part using a cost function.

is mirrored and compared to the left part. This comparison leads to a symmetry value that describes the amount of symmetry in the image patch according to the hypothetical symmetry axis. For the comparison of the two parts of the image patch different cost function are conceivable. Here, we use the *summed normalized cross-correlation* (SNCC) [10] because it has several favorable properties concerning the symmetry calculation as discussed in the next section.

C. Summed Normalized Cross-Correlation (SNCC)

The SNCC cost function is a variation of the *normalized cross-correlation* (NCC). For comparing two half-patches I^L and I^R the NCC cost is defined as:

$$\rho = \frac{\frac{1}{|X|}\sum_{x\in X}(I_x^L - \mu^L)(I_x^R - \mu^R)}{\sigma^L \sigma^R} \;, \tag{1}$$

where

$$\mu = \frac{1}{|X|}\sum_{x\in X} I_x \;, \quad \sigma = \sqrt{\frac{1}{|X|}\sum_{x\in X}(I_x - \mu)^2} \;. \tag{2}$$

In the above equations X is the set of all pixel positions in the compared patches. In the way NCC works it makes a statement about the statistical correlation between two image patches. The main problem of NCC is that already a small number of outlier-pixels in the two patches can lead to a low correlation value, although the remaining major part of the pixels are a perfect match. Most critical are outlier-pixels that exhibit strong contrast changes, i.e. white pixels in one patch are matched to black pixels in the other patch. This phenomenon leads to the well-known fattening effect in block-matching stereo calculations. Thus, the authors of [10] have proposed a slight change to the correlation calculation. Instead of calculating a holistic correlation value for whole patches, they subdivide the patches into multiple overlapping sub-patches. For all sub-patches the correlation coefficients are calculated according to (1). The actual patch correlation is then defined as the mean of all sub-patch correlations:

$$\bar{\rho} = \frac{1}{|X'|}\sum_{x'\in X'} \rho_{x'} \;, \tag{3}$$

where $\rho_{x'}$ is the correlation coefficient of the sub-patch centered at pixel x'. Since the correlation is calculated for each sub-patch independently, outlier-pixels can only influence the correlation values of sub-patches they are located within. Because of this, SNCC degrades very smoothly with the number of outliers or non-symmetric pixels. Furthermore, SNCC can be computed very efficiently by calculating the

NCC for the overlapping sub-patches by means of box-filters [11]. In doing so, the computational complexity of the symmetry value computation for a single patch is linear with the number of pixels within a patch.

D. Symmetry Validation

As explained in section II-A, a symmetry axis search is only conducted for radar targets that are neither static nor occluded. This skips unnecessary error prone calculations. However, the symmetry axis search for the remaining targets might still lead to wrong results. In order to mitigate this problem we employ some security checks on the correlation curve. In pseudo-code the overall process flow reads as follows:

Algorithm 1 Update Lateral Position

```
for each object detected by radar do
  if the object is not static, not occluded, and in the pre-
  defined position range then
    Symmetry axis is searched
    if the correlation curve meets the security criteria, and
    the position of symmetry is consistent with previous
    positions then
      Lateral position is updated with the computed
      symmetry position
    end if
  end if
end for
```

As we use the SNCC cost function for calculating the symmetry value, high correlation values are expected when the image patch inside the filter window is left-right symmetric. This means that all peaks in the correlation curve resulting from the symmetry sweep search (Fig. 5) are considered as candidates for the position of the symmetry axis of an object. Less robust candidate peaks are sieved out by means of the following two shape criteria.

First, the correlation value of a peak has to be greater than a threshold. This suppresses all peaks created by image noise. Second, the segment of the correlation curve above the threshold that contains the peak has to be narrower than a certain ratio of the ROI width, e.g. the segment must not be wider than 1.5 times of the ROI width. This ratio has been determined empirically. Figure 7 shows some typical correlation curves for visualizing the two selection criteria. While the first criterion removes noisy peaks, the second criterion removes peaks in homogeneous areas. If no peak passes both criteria, the lateral position of the object is not refined, but the position measured by radar is used.

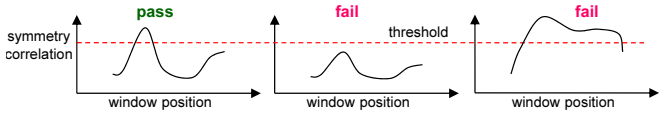

Fig. 7. Examples of peaks that pass and fail the correlation curve criteria: The left peak is a favorable peak. In contrast, the middle peak fails for being confident enough to be above noise level. Although the right peak surpasses the noise threshold, it fails for having a nicely pointed shape because the corresponding line segment is much too wide.

In addition to the shape criteria on correlation curves explained above we also check the remaining peaks for consistency with the positions from previous processing cycles. To this end, the lateral image positions of all valid peaks are converted to lateral positions in the radar coordinate system. Furthermore, a list of the past positions is maintained for each object. From this list the median is calculated and the peak closest to the median is considered the potential current position. If the difference between the median and the closest peak is less than a threshold, then the peak passes into the next step. Otherwise, the lateral position of the object is not refined by vision but the radar's lateral position estimation is left untouched.

In case a symmetry peak has passed all security checks it is further stabilized to counteract position jitter. Jitter occurs when multiple peaks with similar height appear close by and the maximum frequently switches among them. Although this might also be handled by the subsequent Kalman filter, it turned out to be better to directly threat this problem. For doing so, we compute the weighted mean of the lateral position in a certain range around the peak (Fig. 8):

$$x^* = \sum_{x \in \mathcal{N}(x^p)} \frac{c(x) \cdot x}{Z} \tag{4}$$

$$Z = \sum_{x \in \mathcal{N}(x^p)} c(x), \tag{5}$$

where x^* is the stabilized lateral position, $c(x)$ is the correlation value at the lateral position x, x^p is the lateral position of the peak and $\mathcal{N}(x^p)$ is the accumulation range centered at x^p. In our current implementation the range is fixed to a quarter of the ROI width. The stabilized lateral position in image pixel coordinates is eventually converted to radar coordinates and the lateral position measured by the radar is updated with the computed lateral position based on visual symmetry.

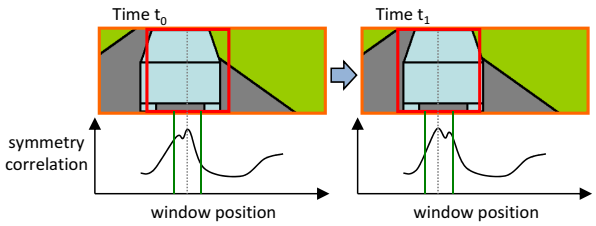

Fig. 8. Symmetry peak stabilization: Sometimes the correlation curve does not give a clear single peak but multiple peaks. In this example, the highest peak appears a little bit offset to the right (gray dotted line) of the actual center of the vehicle at time step t_0, and then switches to the left peak at the next time step t_1. In order to counteract this problem, the weighted mean of an area around the maximal peak (green vertical lines) is calculated.

E. Kalman Filter

We apply a Kalman filter [12] to estimate the position, speed and acceleration of the detected vehicles. Since the noise in the lateral position is much lower when it is refined

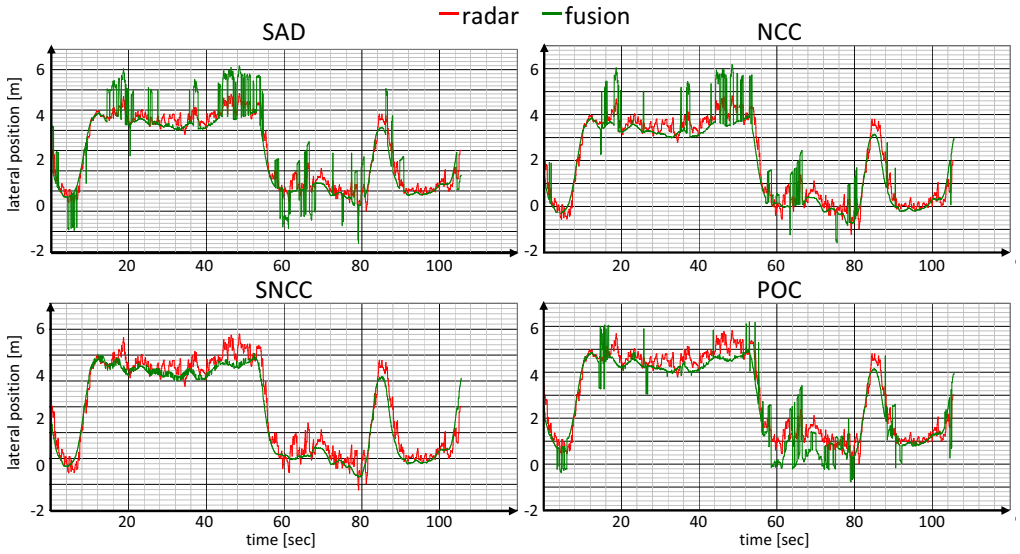

Fig. 9. Comparison of different correlation measures for the patch-matching-based symmetry detection: The lateral position estimation of the symmetry is plotted together with the lateral position of the radar. Note that in this experiment the criteria explained in section II-D were not applied in order to see the plain performance of the different correlation measures. It can be clearly observed that SNCC outperforms the other correlation measures by showing no signs of outliers.

by the *Lateral Position Improvement* as compared to the lateral position by the radar only, the noise parameters of the Kalman filter dynamically adapt to the source of lateral position either being from radar or visual symmetry. In particular, the noise parameter of the lateral position in the Kalman filter is set to nearly zero for the visual symmetry. This leads to the advantageous side-effect that the latency of the lateral movement estimation becomes very small, which is vital for a swift reaction time in several ADAS.

III. Experiments

A. Experimental Setup

For testing our proposed approach we have installed a recording system into a vehicle in order to acquire data for experimental analyses. The system records data from a radar, a camera and an inertial sensor for ego-vehicle speed. The radar and the camera are mounted on a vehicle so that both of their sensing areas in front of the vehicle are overlapping (Fig. 10). Furthermore, the camera and the radar are calibrated in order to perform coordinate transformations between the image domain and the 3-D world.

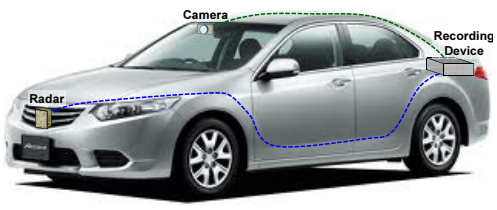

Fig. 10. Configuration of the experimental setup: The test vehicle is equipped with a radar and a camera. Both are scanning the front side of the vehicle and they are mounted in a way so that their sensing areas are overlapping. For storing the sensor data, a recording device is installed in the trunk.

B. Comparison of Different Correlation Measures

As explained in section II-B we calculate a symmetry value for an image patch by splitting it in half and comparing the left side to the mirrored right side by means of matching these half-patches. In general, this patch matching can be performed by a variety of different cost functions. In section II-B we argued that the *summed normalized cross-correlation* (SNCC) is best-suited for such a task. To verify this hypothesis we compare the performance of SNCC in a first experiment with the correlation measures *sum of absolute difference* (SAD), *normalized cross correlation* (NCC) and *phase only correlation* (POC) [13]. For a 100 second test stream, the lateral position corresponding to the highest symmetry correlation in the image is converted to the radar coordinate system and plot over time in comparison to the radar signal as shown in Fig. 9. In this figure, the red curve shows the lateral position measured by the radar, and the green curve shows the lateral position estimated by the symmetry detection. Please note that in this experiment the criteria explained in section II-D were not applied in order to see the plain performance of the different correlation measures. Apparently, SNCC qualitatively outperforms the other correlation measures for robustness in symmetry detection. While the lateral position curve of SNCC is smooth and in correlation with the radar curve, the curves of the other measures show many outliers that even surpass the noise of the radar signal. Empirically we found that sudden jumps of the lateral position of more than 0.5m can be regarded as a wrong detection of the symmetry axis. In order to compare performances quantitatively, Table I shows the numbers of such jumps in lateral position for each cost function. The results in the table support the observation that SNCC clearly outperforms the other cost functions.

TABLE I
NUMBER OF WRONG SYMMETRY DETECTIONS.

SNCC	SAD	NCC	POC
0	142	89	107

C. Comparison between Color and Gray-level Image

So far we have been using color images for the symmetry detection. However, in seek of a fast processing we checked the difference in the performance of the symmetry axis detection between color images and gray-level images. It has been shown in stereo image processing [14] that for normalizing measures the performance with gray-level images is the same or even better as compared to color images. Hence, it is reasonable to assume that the same holds for the symmetry detection based on patch matching. Figure 11 shows a comparative plot. The yellow curve is

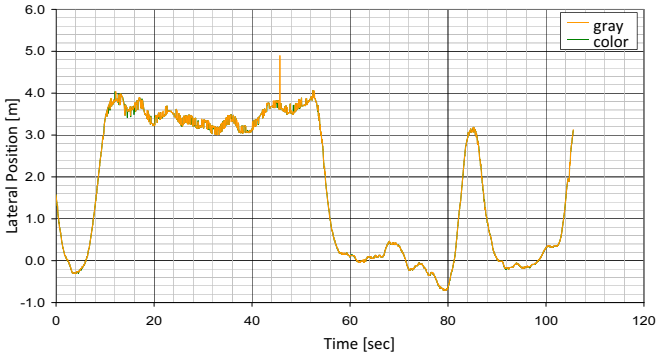

Fig. 11. Comparison of the lateral position estimation with symmetry using color images and gray-level images: Except for one outlier the performance in both cases is the same. Hence, the green curve for the color images is hard to see due to the large overlap with the yellow curve for the gray-level images.

the lateral position estimation using gray-level images while the green curve shows the results for RGB color images. Even though the processing time for gray-level images is about one third of the color images, the quality of the lateral position estimation is about the same. In this experiment only one error occurred for the sequence of gray-level images. However, it is easily sorted out by the consistency checks explained in Section II-D. In conclusion, this experiment is in line with the findings by Bleyers *et al.* [14] for stereo processing that gray-level images yield the same performance as color images.

D. Improvement of Lateral Position Accuracy

In this second experiment, we demonstrate that the accuracy of the lateral position measured by a radar is strongly improved by the fusion with symmetry detection. For this purpose we have manually labeled the center of one car in about 2400 image frames of a 120 second stream. We use this labeled center position to extract the ground truth lateral position of the car.

The top plot in Fig. 12 shows a temporal plot of the ground truth lateral position (cyan curve) as well as the lateral position measured by the radar (red curve), the lateral position estimated by a Kalman filter applied to the radar signals (blue curve) and the position estimated by fusing the radar signals with the proposed visual symmetry detection (green curve). This plot is done in metric coordinates to

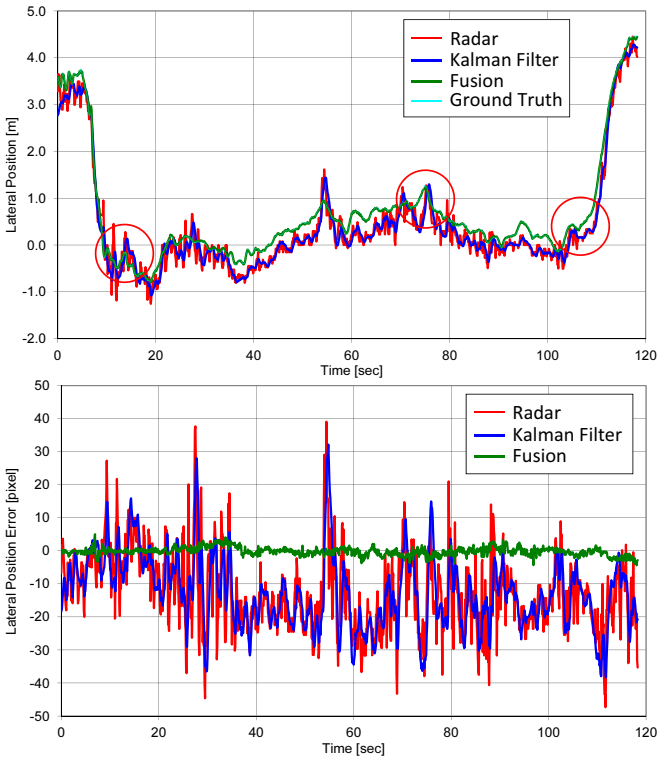

Fig. 12. Comparison of the lateral position measured by the radar, estimated by a Kalman filter applied to the radar signals and estimated by fusing radar signals with the proposed symmetry axis detection. It can be observed that the fusion of radar and visual symmetry leads to much more accurate estimates without latency compared to the pure radar or Kalman-filtered radar signals. The latency of the radar signals is particularly visible at the spots marked by the red circles in the top plot.

visualize the impact of the fusion on the accuracy of the 3-D position and movement estimation, which is important for subsequent ADAS. Please note that the ground truth (cyan curve) is hardly visible because it is overlapped by the green curve for the fusion of radar with visual symmetry. Thus, the bottom plot in Fig. 12 depicts the lateral position error of the estimated lateral position, i.e. the difference to the ground truth. This second plot is done in image coordinates in order to have an evaluation that is independent from the distance to the vehicle.

The two plots in Fig. 12 highlight three things. First, they demonstrate that the proposed fusion of the radar signal with the symmetry detection leads to a strong improvement of the lateral position accuracy. The large error for the pure radar signals arises from the fact that a radar has difficulties to calculate a precise lateral position within a radar beam. As the number of beams is limited, so is the lateral precision of the radar. Because of this, typically a Kalman filter is used to smooth the radar signal, but as the two plots show this alleviates the problem only little.

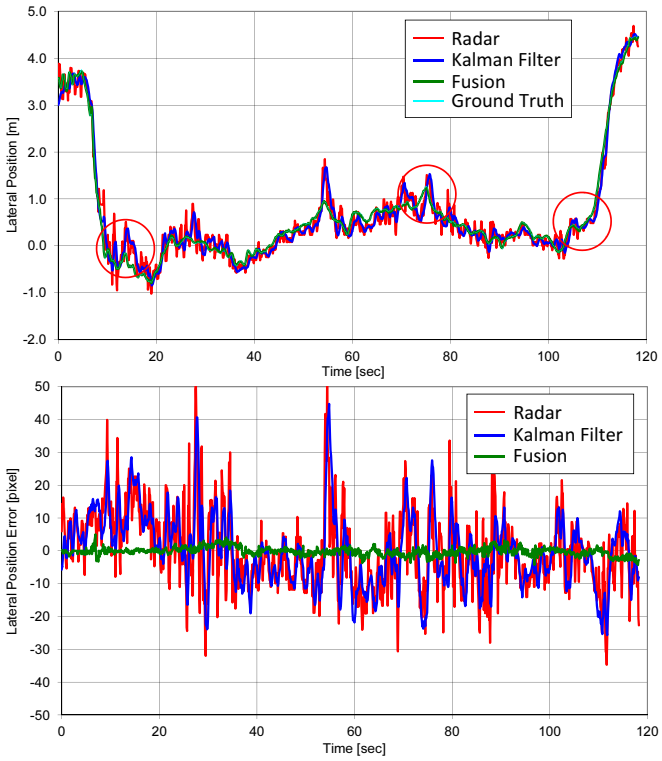

Fig. 13. This figure shows the same plots as in Fig. 12, but the bias for the radar and the radar filter with Kalman is removed. They show that even with removed bias, the lateral position estimated by fusing the radar with symmetry is much better than applying a Kalman filter or using the pure radar signals.

TABLE II

ERRORS IN LATERAL POSITION ESTIMATION.

	Error in metric [m]			Error in image [pixel]		
	Radar	Filter	Fusion	Radar	Filter	Fusion
Mean	-0.23	-0.24	0.00	-12.6	-12.7	-0.2
Std. dev.	0.23	0.21	0.03	11.7	11.0	1.3

Second, the symmetry-fused lateral position in Fig.12 shows no hints of latency. In contrast, the radar signals show a strong increase in latency due to the applied Kalman filter. This means that the onset of a change of movement direction is delayed in the lateral position estimation. At the spots marked by the red circles in the top plot of Fig. 12 it is observable that the onset of the lateral movement of the vehicle in the symmetry fusion is detected several seconds before the radar or the Kalman-filtered radar signals. Hence, the lateral position improvement by symmetry allows for ADAS with strongly reduced latency which increases the safety.

The third observation is that in our setup the radar has a bias towards the left side (negative values in lateral position). This is visible in both plots in Fig. 12. One reason for this might be a calibration error between the camera and the radar. In order to quantify this bias, Table II shows the mean and standard deviation of the lateral position error of the radar, the radar with Kalman and the radar fused with symmetry for metric 3-D world coordinates and 2-D pixel image coordinates. As can be seen, the mean error of the radar or the radar filtered by Kalman is roughly 23cm in 3-D world coordinates. In contrast, the mean error of the lateral position estimated by means of the fusion symmetry is negligible. In order to give a better image, Fig. 13 shows the same plots as Fig. 12 but with removed bias.

These plots demonstrate that even with removed bias the lateral position estimation of the symmetry fusion is much better than the radar with Kalman. This is further supported by the standard deviation of the lateral position error shown in Table II. While the lateral position estimated by the radar without symmetry varies about 22cm from the ground truth in world coordinates or 11 pixels in image coordinates, the lateral position estimated with symmetry varies only 3cm from the ground truth in world coordinates or 1.3 pixel in image coordinates. This confirms that the lateral position accuracy is tremendously improved by fusing the radar signals with the proposed visual symmetry.

E. Lateral Movement Events

For this final evaluation, we manually annotated events in images of one hour of driving where vehicles start a lateral movement, e.g. vehicles moving into a curve or starting to overtake another car. Altogether, we counted 102 lateral movement events. For these events an accurate lateral position is very important and we wanted to analyse how often the symmetry-improved results can be used.

As explained in section II-D, the symmetry axis search sometimes can not be applied or trusted. Thus, we counted the number of events that actually included the refinement of radar with visual symmetry. In our recorded streams, the symmetry detection was used in 54% (55 of 102) of the events of laterally moving vehicles. One might wonder about the discrepancy between the strong improvement of accuracy demonstrated in the last section and the seemingly small number of cases were symmetry was actually trustful in this experiment. One reason for this is that the evaluation in this section also includes the difficult cases of overlapping vehicles while the experiment in the last section was only done with non-overlapping vehicles. On the other hand it may also indicate that the rejection is too strict and could be improved.

IV. CONCLUSION AND FUTURE WORK

A. Conclusion

In this paper, we proposed a method for improving the lateral position accuracy of radar detections by performing a symmetry axis search in camera images. The effectiveness of our method has been shown in several experiments. One key element of our approach is that we use the *summed normalized cross-correlation* (SNCC) for finding the symmetry axis of vehicle rear sides in a patch-matching fashion. The merit of SNCC is its high robustness against local pixel outliers thus allowing for stable results even if a vehicle has some non-symmetric parts. Our experiments have shown that due to this robustness SNCC leads to a much better symmetry axis than other typical correlation measures.

A second important element is a set of checks that detects situations where the symmetry calculation is prone to fail. This ensures that the fusion of radar and symmetry does never lead to a lateral position estimation that is worse than by radar alone. The experiments we conducted demonstrate that the accuracy (standard deviation from ground truth) of the lateral position improves from 21cm for radar with Kalman filter to 3cm for radar fused with symmetry. This amounts to an error reduction by a factor of seven. Furthermore, the symmetry-fused lateral position shows no sights of latency while the Kalman filtered radar lateral position can be delayed by some seconds. Thus, the proposed approach leads to a significant increase in safety for subsequent ADAS because these can react more swiftly and robustly.

B. Future Work

In our current implementation, it takes about 70ms to process one vehicle with a gray-level image on Intel(R) Xeon(R) CPU X5550(2.67GHz). In future work, we want to investigate on how the correlation computation can be speeded up. Another target is to further improve the robustness of the symmetry axis search so that there are fewer rejections and the actual lateral improvement is applied more often. Furthermore in this paper, we have done quantitative analysis only for one short image stream of 120 seconds. Hence, one of our next targets is to manually label larger parts of our recorded streams for the lateral position of the vehicles in order to allow for a more detailed quantitative analysis which also covers the critical cases of occlusion.

REFERENCES

[1] A. Haselhoff, A. Kummert, and G. Schneider, "Radar-vision fusion for vehicle detection by means of improved haar-like feature and adaboost approach," in *15th European Signal Processing Conference*, September 2007, pp. 2070–2074.

[2] L. Bombini, P. Cerri, P. Medici, and G. Alessandretti, "Radar-vision fusion for vehicle detection," in *Proceedings of the International Workshop on Intelligent Transportation*, March 2006, pp. 65–70.

[3] A. Broggi, M. Bertozzi, A. Fascioli, C. G. Lo Bianco, and A. Piazzi, "Visual perception of obstacles and vehicles for platooning," in *IEEE Transactions on Intelligent Transportation Systems*, vol. 1, no. 3, September 2000, pp. 164–176.

[4] M. Bertozzi, A. Broggi, A. Fascioli, and A. Nichele, "Stereo vision-based vehicle detection," in *IEEE Intelligent Vehicles Symposium*, October 2000, pp. 39–44.

[5] A. Broggi, P. Cerri, and P. C. Antonello, "Multi-resolution vehicle detection using artificial vision," in *IEEE Intelligent Vehicles Symposium*, 2004, pp. 310–314.

[6] Y. Du and N. P. Papanikolopoulos, "Real-time vehicle following through a novel symmetry-based approach," in *IEEE International Conference on Robotics and Automation*, vol. 4, April 1997, pp. 3160–3165.

[7] C. Hoffmann, T. Dang, and C. Stiller, "Vehicle detection fusing 2d visual features," in *IEEE Intelligent Vehicles Symposium*, June 2004, pp. 280–285.

[8] A. Kuehnle, "Symmetry-based recognition of vehicle rears," *Pattern Recognition Letters*, vol. 12, no. 4, pp. 249–258, April 1991.

[9] T. Zielke, M. Brauckmann, and W. von Seelen, "Intensity and edge-based symmetry detection applied to car-following," in *Computer Vision - ECCV'92*, ser. LNCS, vol. 588, 1992, pp. 865–873.

[10] N. Einecke and J. Eggert, "A two-stage correlation method for stereoscopic depth estimation," in *DICTA*, 2010, pp. 227–234.

[11] M. J. McDonnell, "Box-filtering techniques," *Computer Graphics and Image Processing*, vol. 17, no. 1, pp. 65–70, 1981.

[12] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal Of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.

[13] C. D. Kuglin and D. C. Hines, "The phase correlation image alignment method," *IEEE Conference on Cybernetics and Society*, pp. 163–165, 1975.

[14] M. Bleyer and S. Chambon, "Does color really help in dense stereo matching?" in *International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT)*, 2010, pp. 1–8.