

Full-Velocity Radar Returns by Radar-Camera Fusion

– Supplementary Material –

Abstract

In the supplementary material, we illustrate predicted radar-camera association. Second, we evaluate the influences of two factors, depth and the angle between actual velocity and radial direction, on the performance of point-wise full velocity estimation. Moreover, we report the computational time of three components of the estimation pipeline. Finally, we present a video showing point-wise velocity estimation in real driving scenes.

1. Visualization of Predicted Radar-Camera Association

Fig. 1 shows the mean of predicted association A for the test set. It appears the radar point is more likely associated with pixels above the raw projection, as the measured radar height is always on the radar plane which is typically lower than vehicle height.

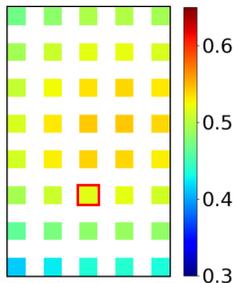


Figure 1: Mean of predicted association over the test set in the neighborhood of raw projection (marked by a red square). The neighborhood region has a size of 9×15 pixels and association estimation skips every other pixel to improve the computational efficiency.

2. Velocity Estimation Error for Different Depths and α

This experiment extends the evaluation of point-wise velocity estimation discussed in Section 4.1 of the main paper. In Fig. 2, each heat map shows point-wise velocity error under different depth ranges, i.e. $[0, 25)$, $[25, 50)$ and $[50, \infty)$

meters as well as various α ranges, i.e. $[0, 30)$, $[30, 60)$ and $[60, 90]$ degrees, where α is the angle between actual moving direction and radial direction of a radar point and ranges from 0 to 90 degrees. Results of the proposed method and baseline are show in the first row and second row, respectively. The baseline (second row), with only radial measurement, suffers from large α since the the actual moving direction is very different from radial direction under large α . The proposed method outperforms the baseline in all depth and α ranges for full velocity estimation.

3. Inference Time

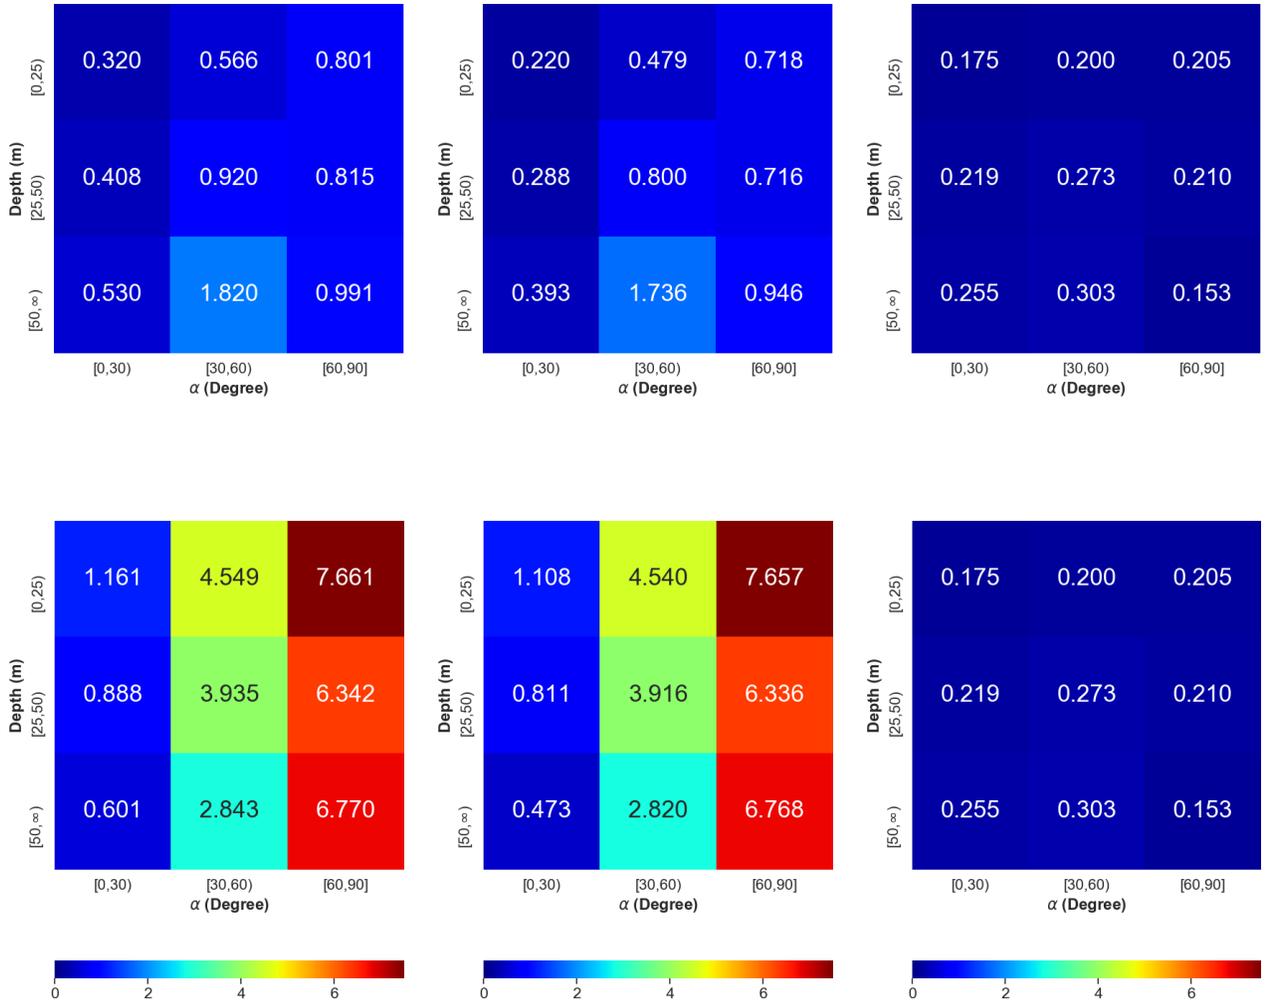
The pipeline of our full velocity estimation includes three major components, optical flow computation, radar-camera association estimation and closed-form solution of full velocity. The time used by each component per frame is listed in Table 1. Our computational platform includes Intel Core i7-8700 CPUs and a NVIDIA GeForce RTX 2080 Ti GPU. The proposed closed-form solution achieves highly efficient computation. Note the computational cost of optical flow can be improved by limiting the region of flow computation to areas with radar projections.

| Components | Time per Frame (s) |
|----------------------------------|-----------------------|
| Optical Flow [1] | 4.47×10^{-1} |
| Radar-camera Association | 2.58×10^{-3} |
| Closed-form Velocity Computation | 2.49×10^{-4} |

Table 1: Computational time of three components in the method pipeline.

4. Video File

In the video, we show point-wise velocity estimation (black arrow) of dynamic radar points in bird-eye view of radar coordinates. Moving radar points are also plotted with radial velocity (red arrow) and static points are shown in orange. The true velocity of vehicles are plotted as green arrow. The GT moving and static vehicles are plotted as solid and dashed bounding boxes, respectively. Images with radar projections are shown at top-left corner for reference.



(a) Full Velocity Error

(b) Tangential Component Error

(c) Radial Component Error

Figure 2: Comparison of average error (in meters) of point-wise velocity estimates by the proposed method (first row) and baseline (second row). Columns 1, 2 and 3 are error of full velocity, tangential component and radial component, respectively. Each heat map shows the error for radar points in different depth and α ranges, where α is the angle between full velocity and radial direction.

References

- [1] Zachary Teed and Jia Deng. RAFT: Recurrent all-pairs field transforms for optical flow. In *European Conference on Computer Vision*, pages 402–419, 2020. [1](#)