Supplementary Material for
Beat the MTurkers: Automatic Image Labeling from Weak 3D Supervision

Liang-Chieh Chen[1] Sanja Fidler[2] Alan L. Yuille[1] Raquel Urtasun[2]
[1]University of California- Los Angeles [2]University of Toronto
{lcchen@cs, yuille@stat}.ucla.edu {urtasun, fidler}@cs.toronto.edu

In the supplementary material for paper [1], we show

1. one more experimental result about our ad-hoc to remove some point cloud outliers for the appearance model,
2. more segmentation results by our model, similar to Fig. 7 in the main paper.

Effect of removing point cloud outliers for appearance model: Recall in our appearance model, we remove the background seeds that both fall inside the car convex hull and have larger depth values than the mean. Using this outlier removal for the appearance model slightly improves performance. For the Discrete Depth model in Table 2 of the main paper, which directly uses the seeds as a potential, it improves from 70.1% to 74.2%. However, with the full model, it improves only 0.1% (1.8% if using average shapes instead of CAD). This shows that our full model is robust to outliers. Note that even with this ad-hoc method, the dataset is still very noisy as shown in Fig.1 (bottom) of the main paper.

More segmentation results We show more segmentation results in the following pages.

References

[1] L.-C. Chen, S. Fidler, A. Yuille, and R. Urtasun. Beat the mturkers: Automatic image labeling from weak 3d supervision. In *CVPR*, 2014. 1

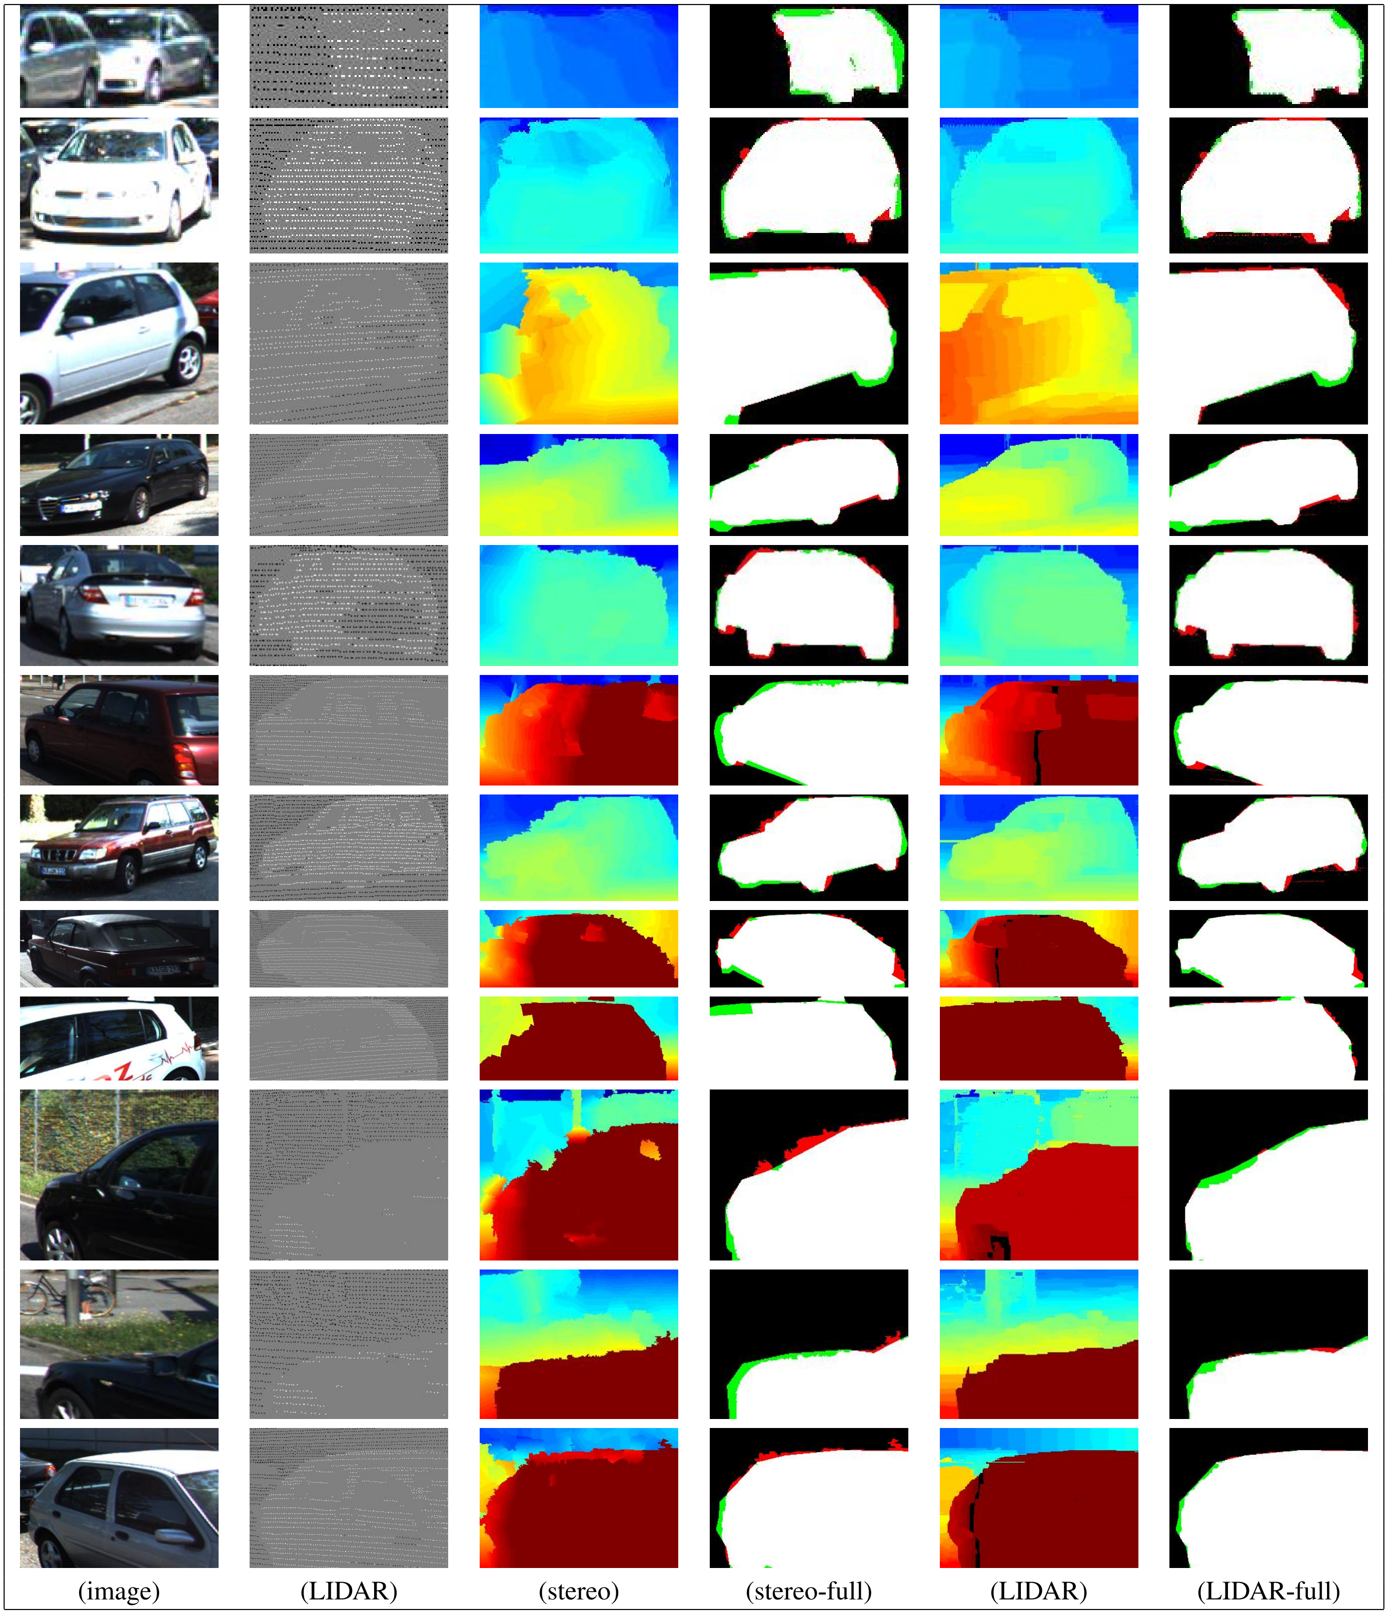

Figure 1: **Segmentation results:** For each row, we show original image, projected point clouds (White: car, Black: background), stereo depth, segmentation results by our full model with stereo (White: True Positive, Black: True Negative, Red: False Positive, Green: False Negative), depth images reconstructed by Laplacian MRF, and segmentation results by our full model with LIDAR. Note that all images are rescaled for display.

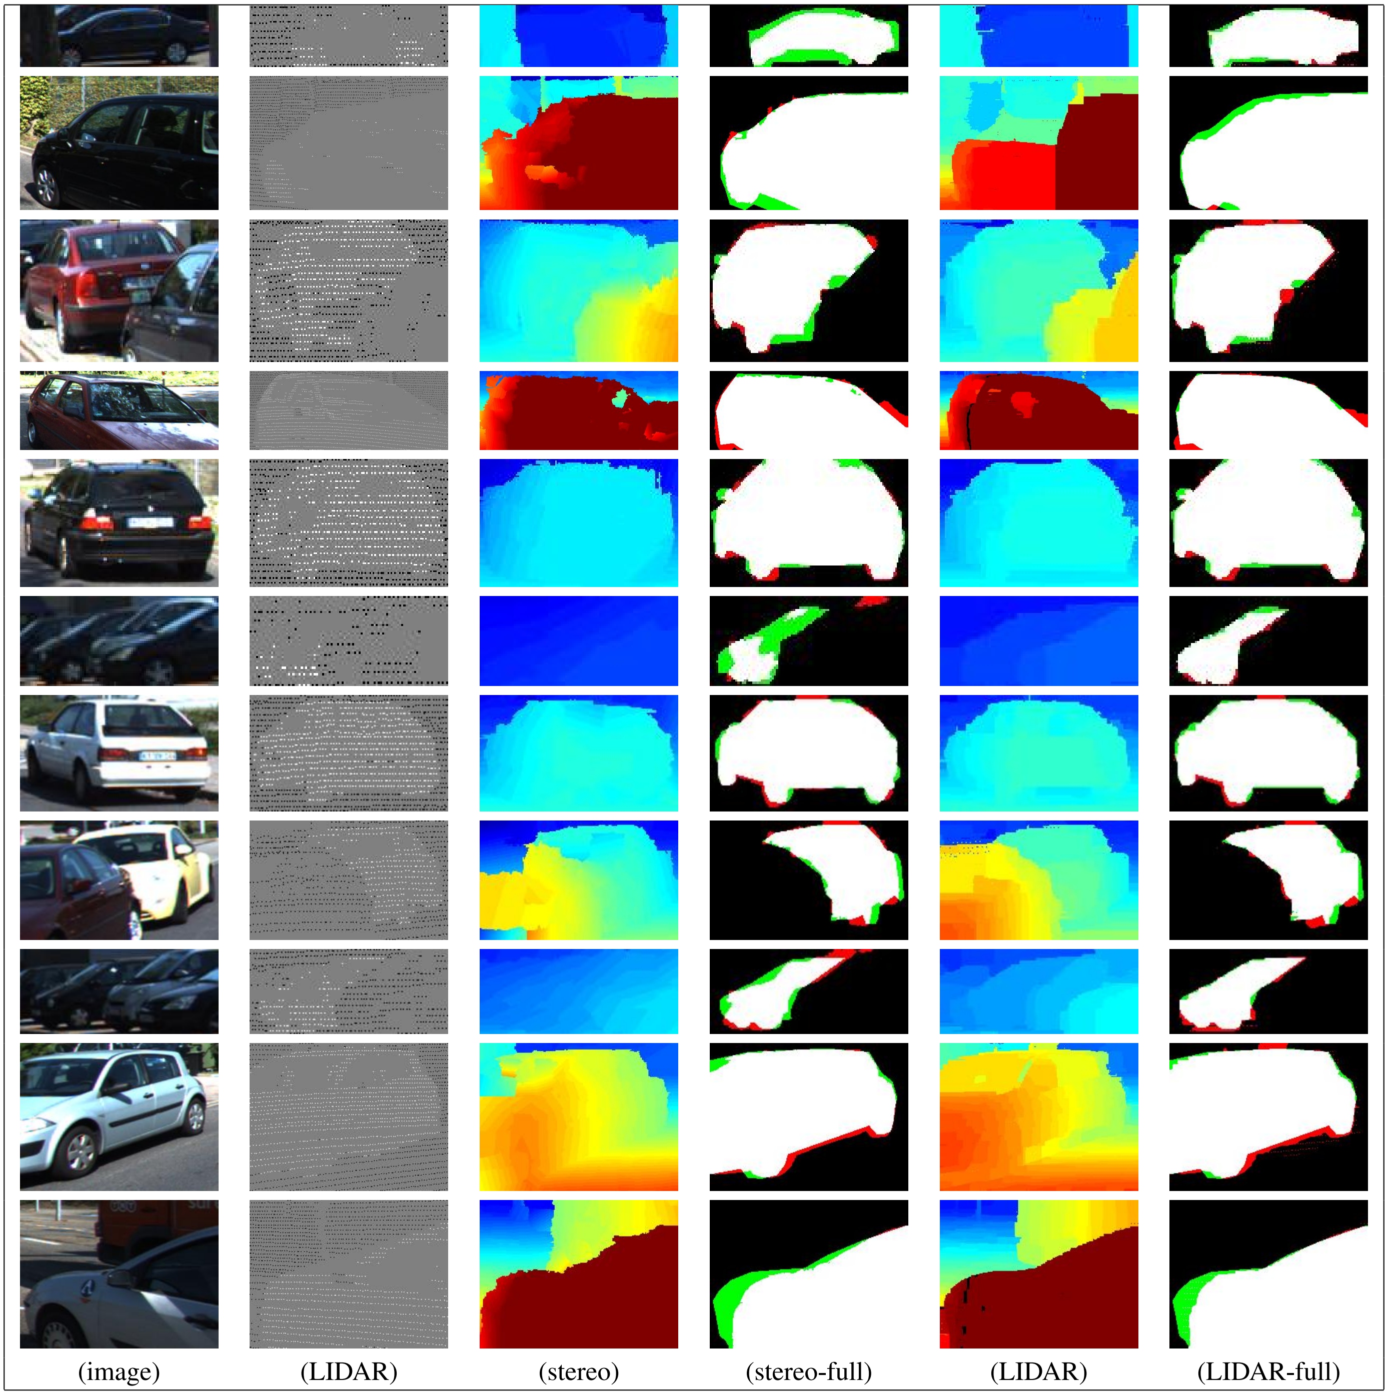

Figure 2: **Segmentation results:** For each row, we show original image, projected point clouds (White: car, Black: background), stereo depth, segmentation results by our full model with stereo (White: True Positive, Black: True Negative, Red: False Positive, Green: False Negative), depth images reconstructed by Laplacian MRF, and segmentation results by our full model with LIDAR. Note that all images are rescaled for display.

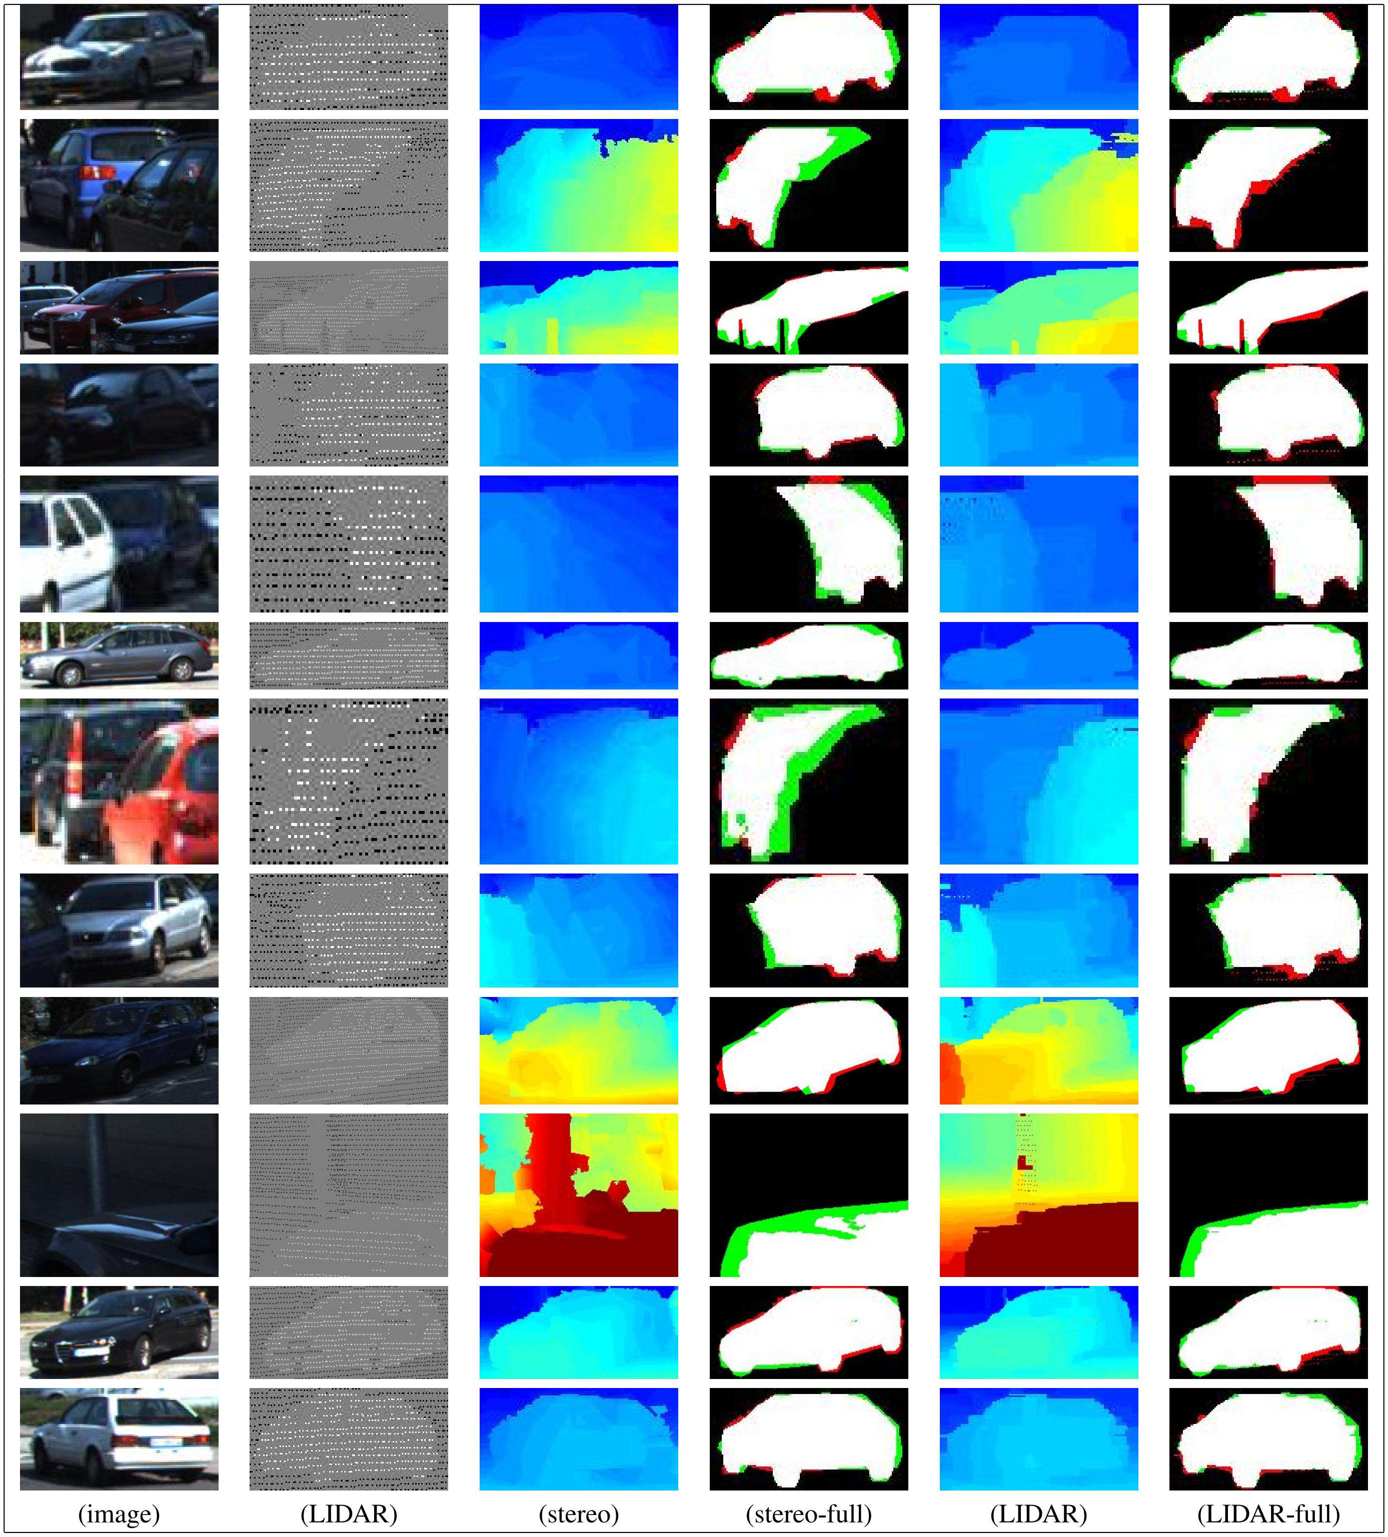

Figure 3: **Segmentation results:** For each row, we show original image, projected point clouds (White: car, Black: background), stereo depth, segmentation results by our full model with stereo (White: True Positive, Black: True Negative, Red: False Positive, Green: False Negative), depth images reconstructed by Laplacian MRF, and segmentation results by our full model with LIDAR. Note that all images are rescaled for display.

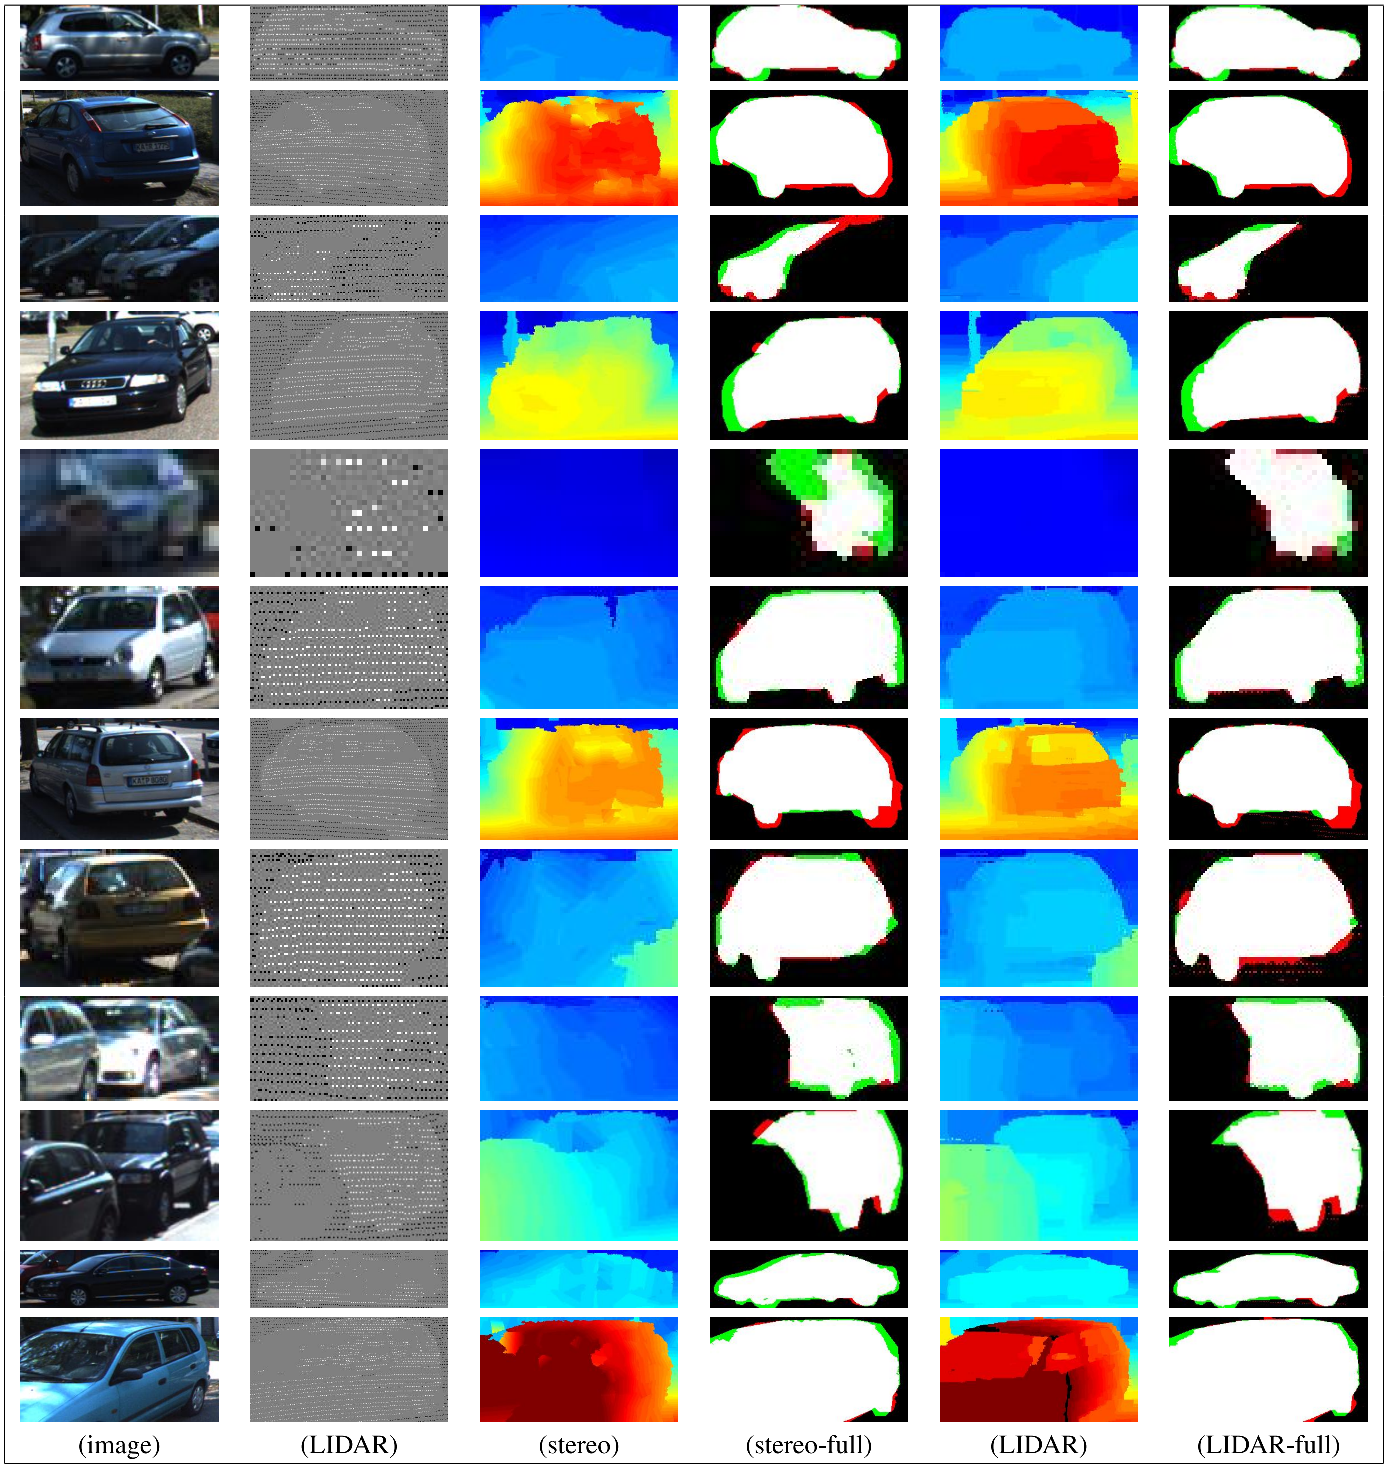

Figure 4: **Segmentation results:** For each row, we show original image, projected point clouds (White: car, Black: background), stereo depth, segmentation results by our full model with stereo (White: True Positive, Black: True Negative, Red: False Positive, Green: False Negative), depth images reconstructed by Laplacian MRF, and segmentation results by our full model with LIDAR. Note that all images are rescaled for display.

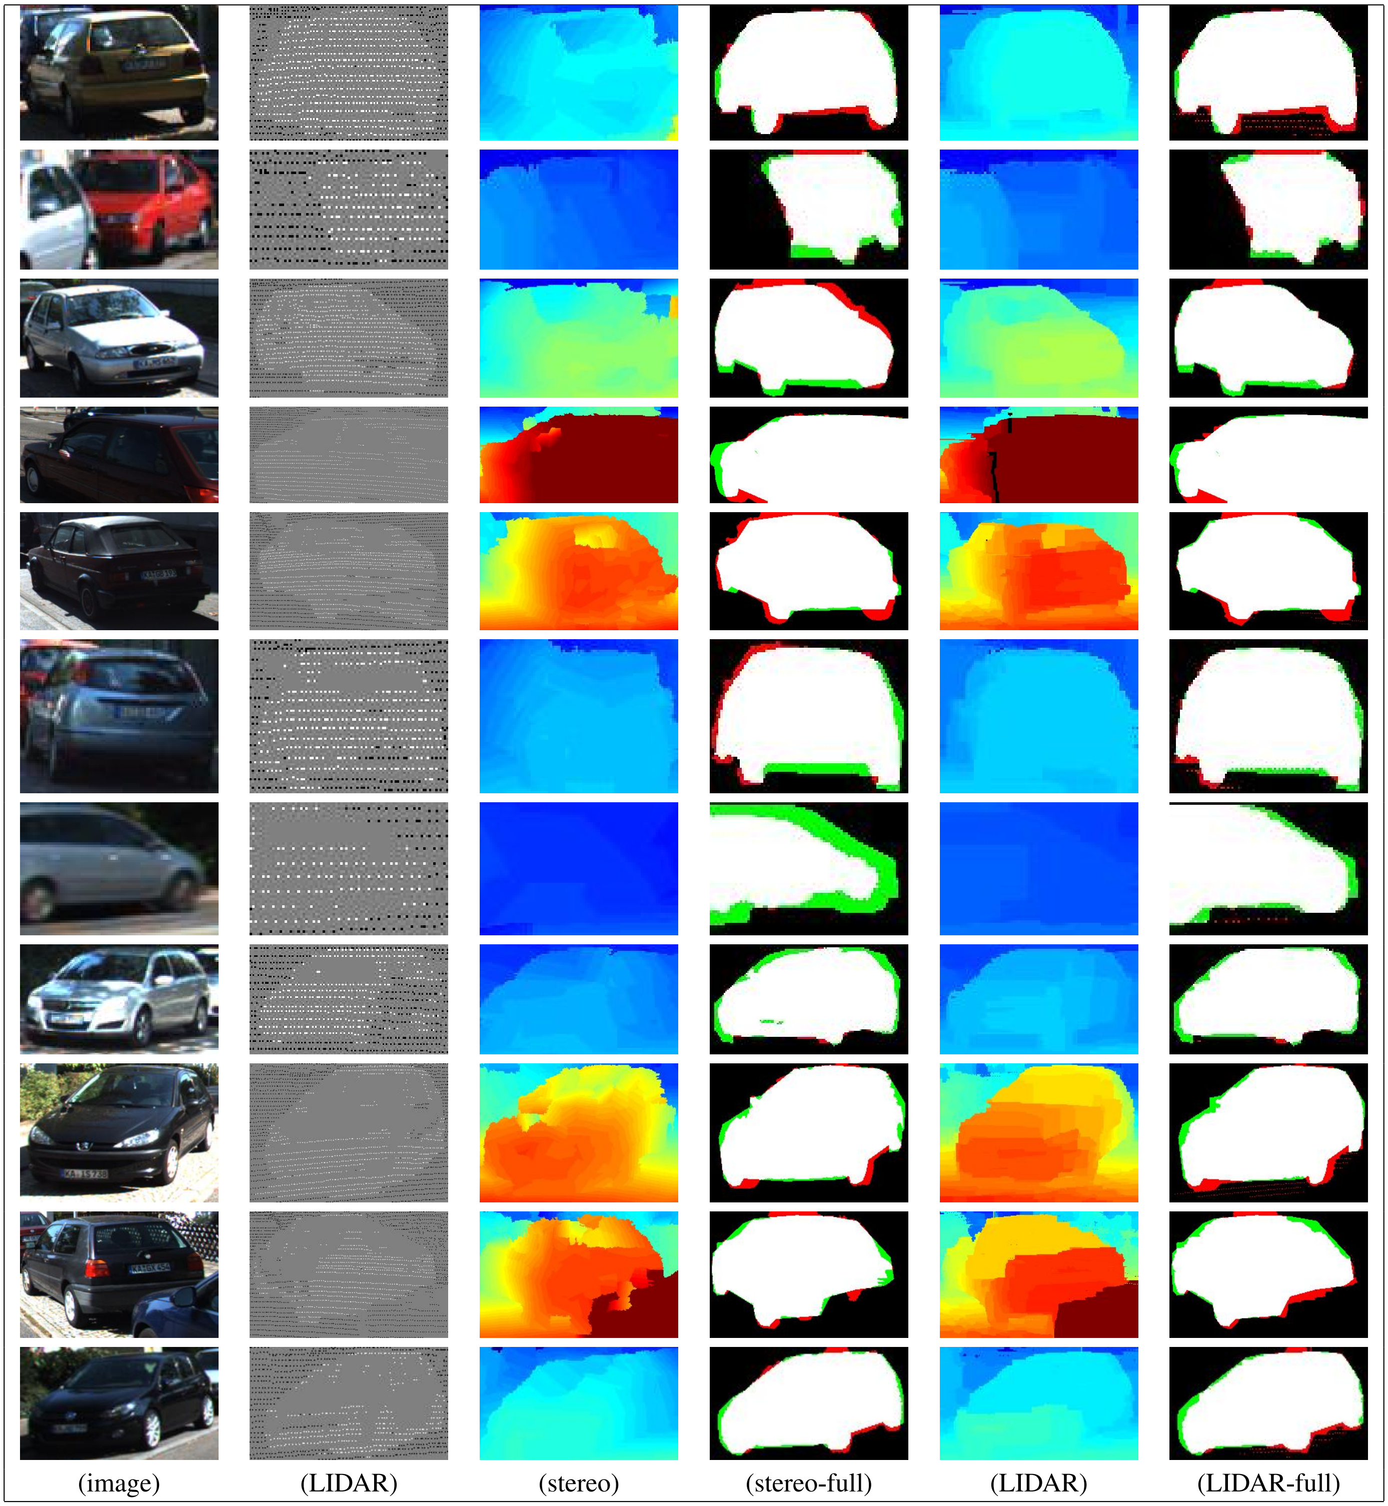

Figure 5: **Segmentation results:** For each row, we show original image, projected point clouds (White: car, Black: background), stereo depth, segmentation results by our full model with stereo (White: True Positive, Black: True Negative, Red: False Positive, Green: False Negative), depth images reconstructed by Laplacian MRF, and segmentation results by our full model with LIDAR. Note that all images are rescaled for display.

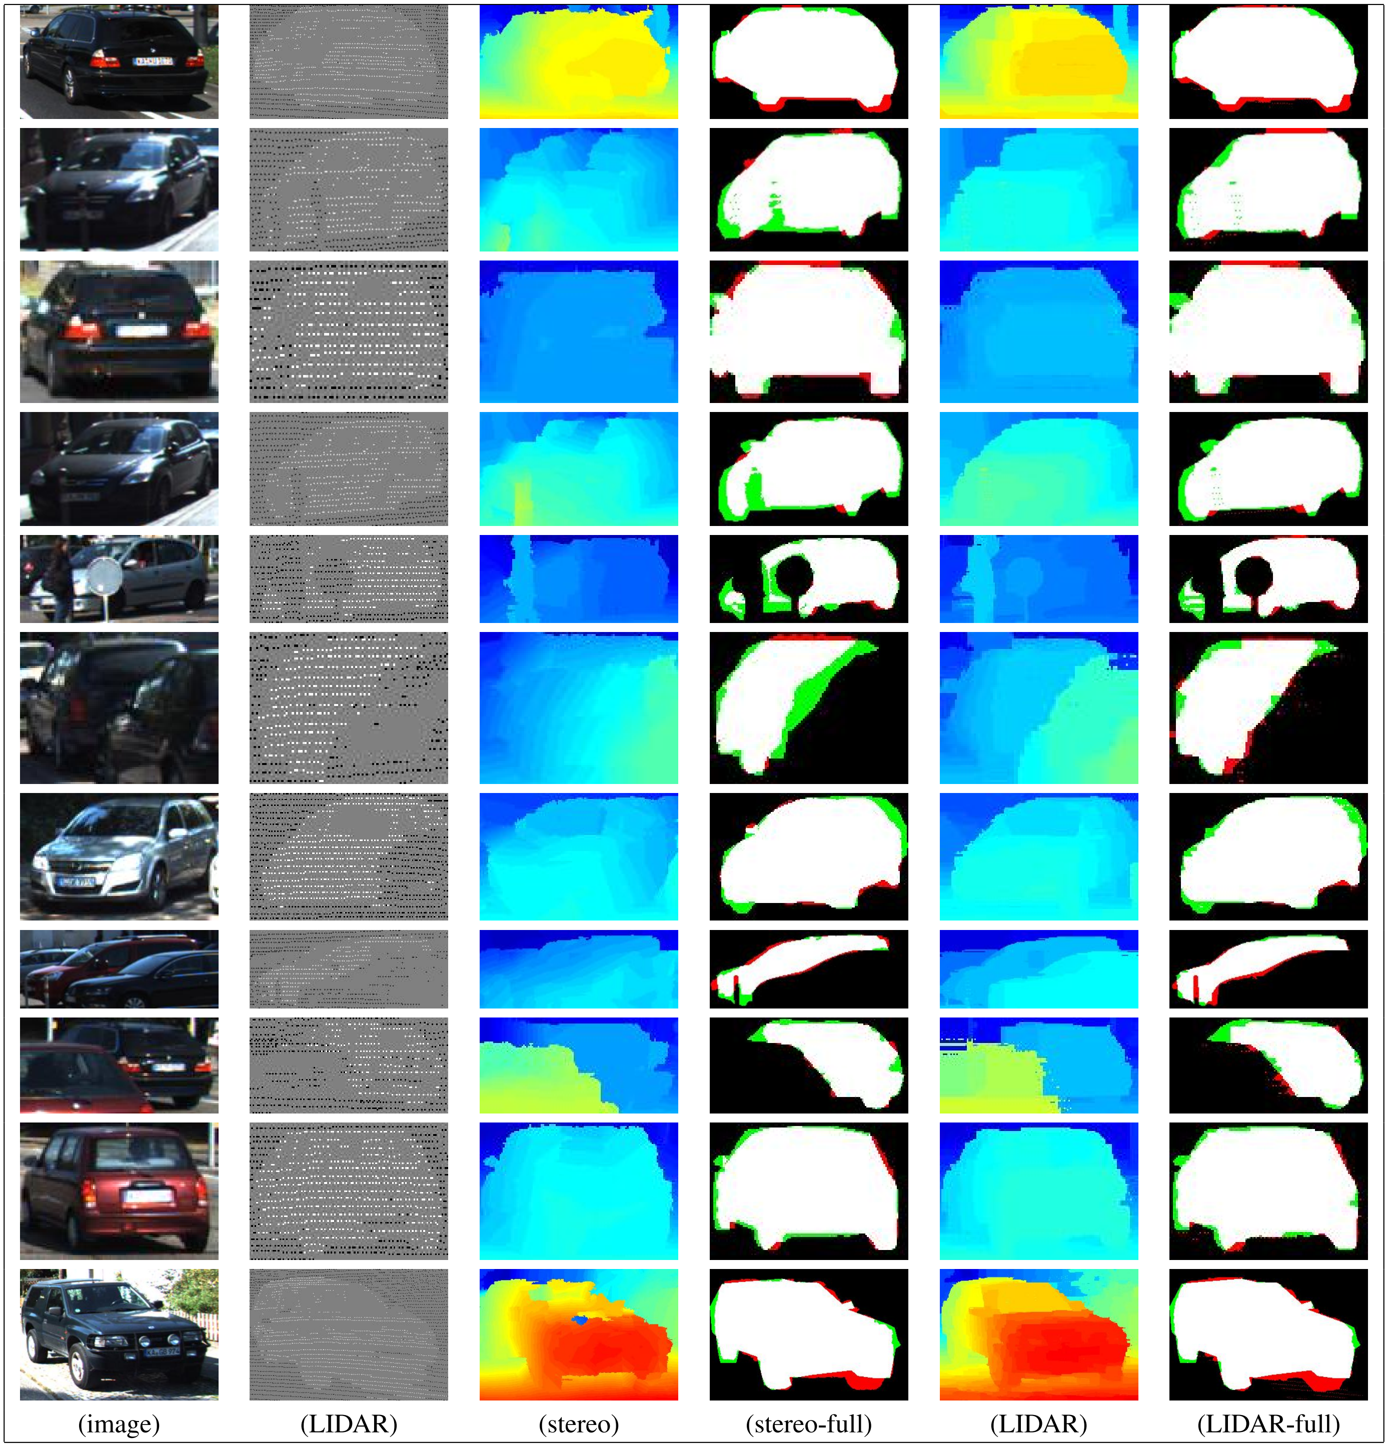

Figure 6: **Segmentation results:** For each row, we show original image, projected point clouds (White: car, Black: background), stereo depth, segmentation results by our full model with stereo (White: True Positive, Black: True Negative, Red: False Positive, Green: False Negative), depth images reconstructed by Laplacian MRF, and segmentation results by our full model with LIDAR. Note that all images are rescaled for display.

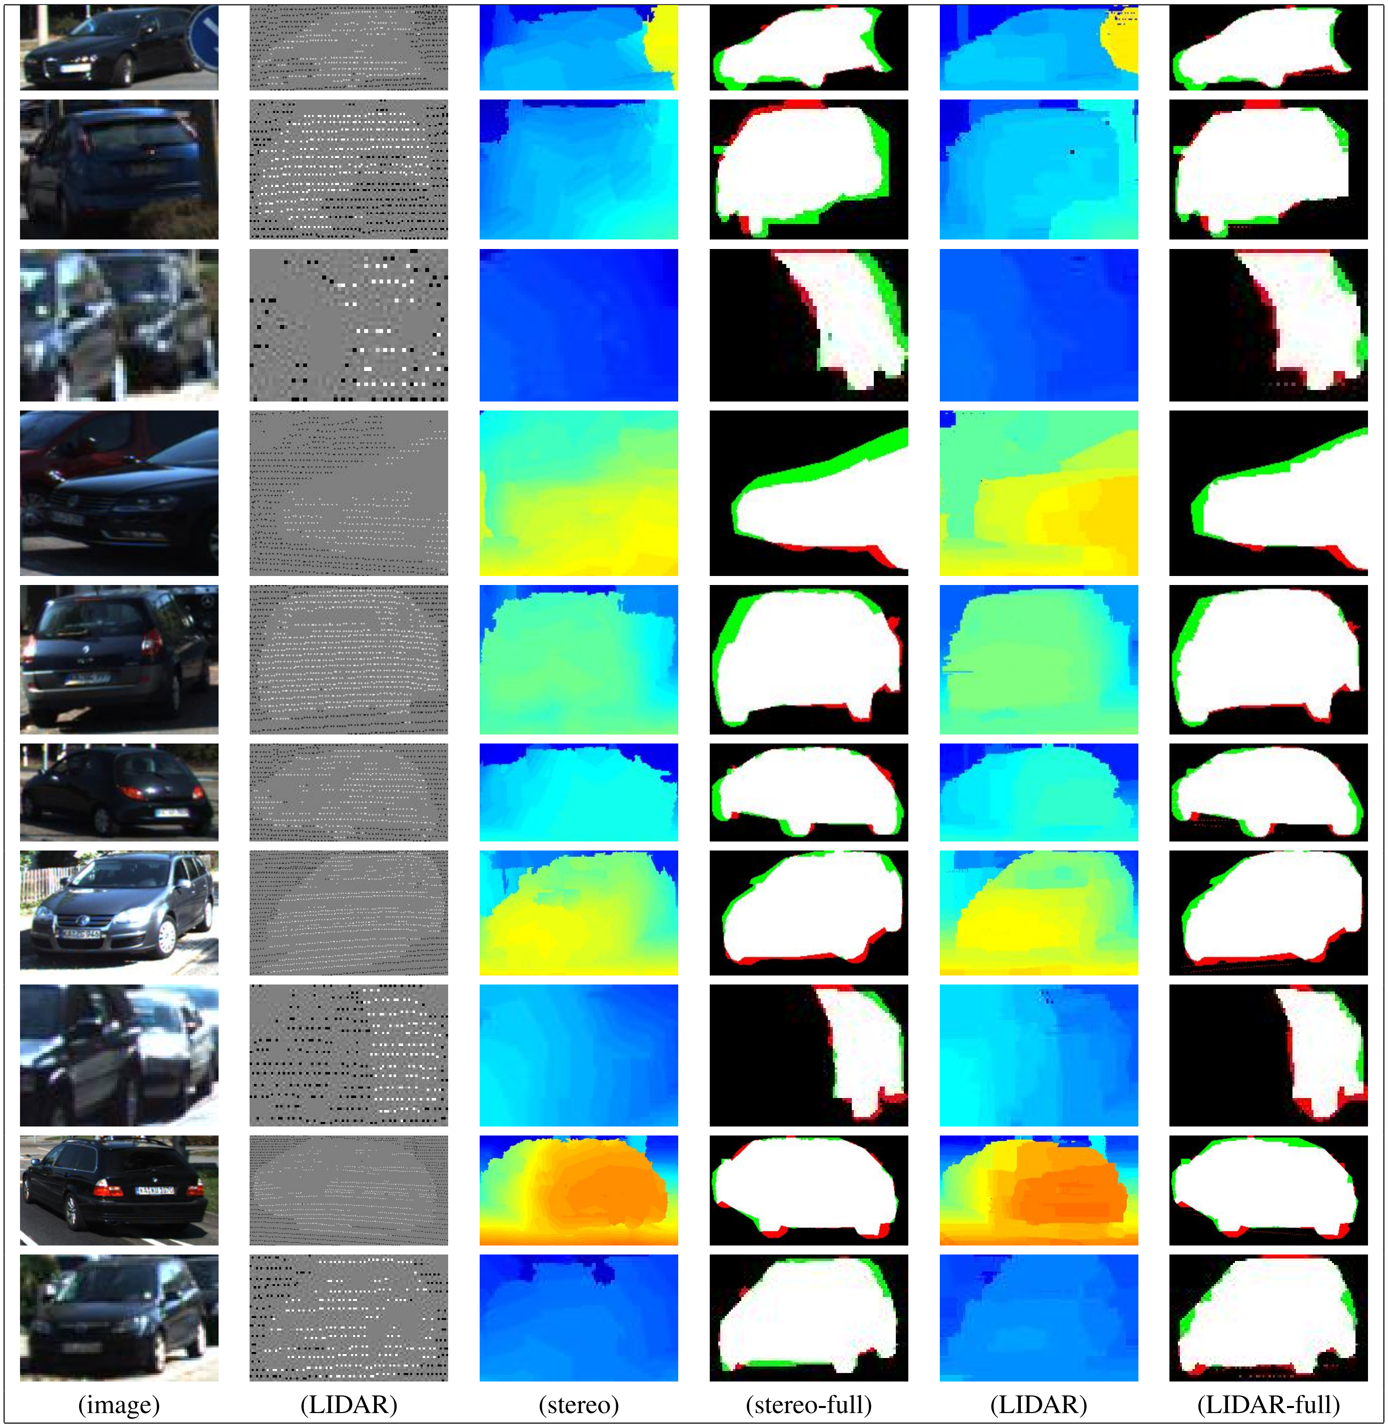

Figure 7: **Segmentation results:** For each row, we show original image, projected point clouds (White: car, Black: background), stereo depth, segmentation results by our full model with stereo (White: True Positive, Black: True Negative, Red: False Positive, Green: False Negative), depth images reconstructed by Laplacian MRF, and segmentation results by our full model with LIDAR. Note that all images are rescaled for display.

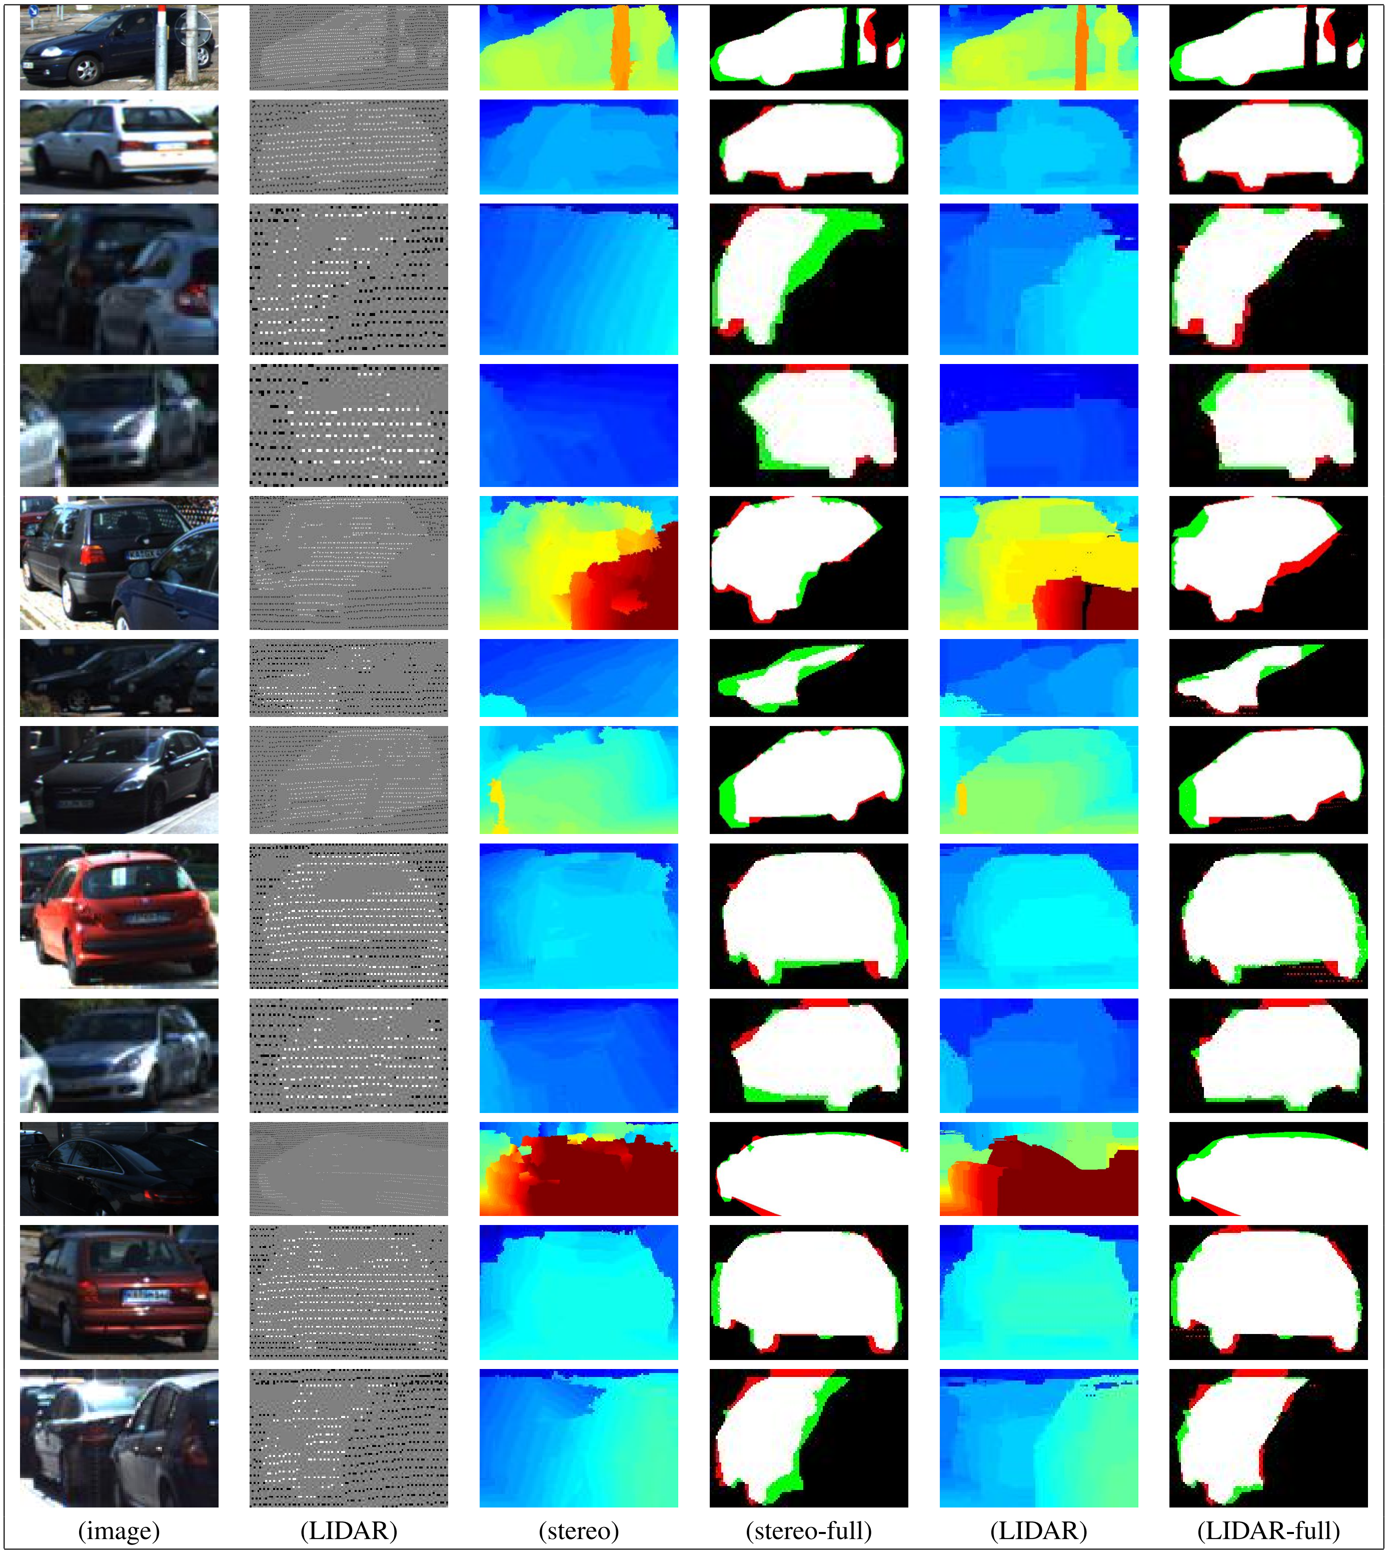

Figure 8: **Segmentation results:** For each row, we show original image, projected point clouds (White: car, Black: background), stereo depth, segmentation results by our full model with stereo (White: True Positive, Black: True Negative, Red: False Positive, Green: False Negative), depth images reconstructed by Laplacian MRF, and segmentation results by our full model with LIDAR. Note that all images are rescaled for display.

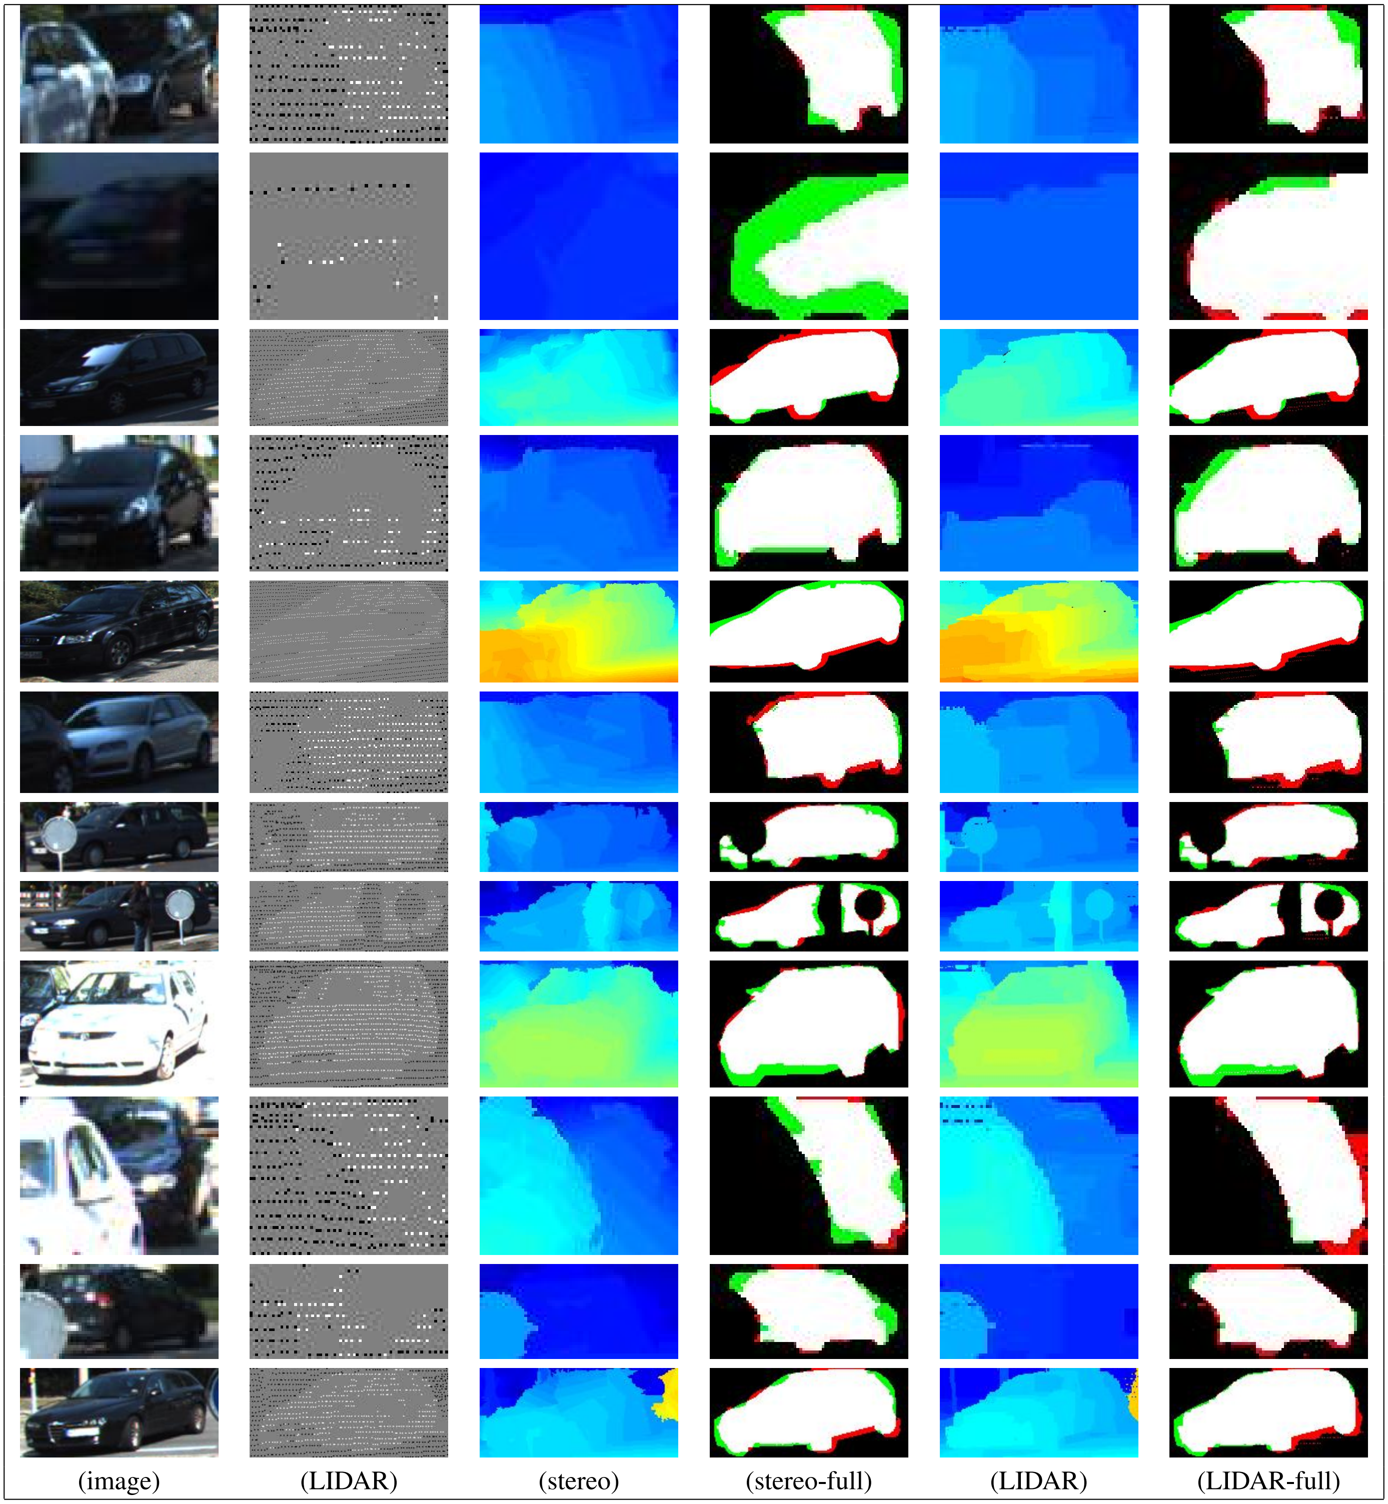

Figure 9: **Segmentation results:** For each row, we show original image, projected point clouds (White: car, Black: background), stereo depth, segmentation results by our full model with stereo (White: True Positive, Black: True Negative, Red: False Positive, Green: False Negative), depth images reconstructed by Laplacian MRF, and segmentation results by our full model with LIDAR. Note that all images are rescaled for display.

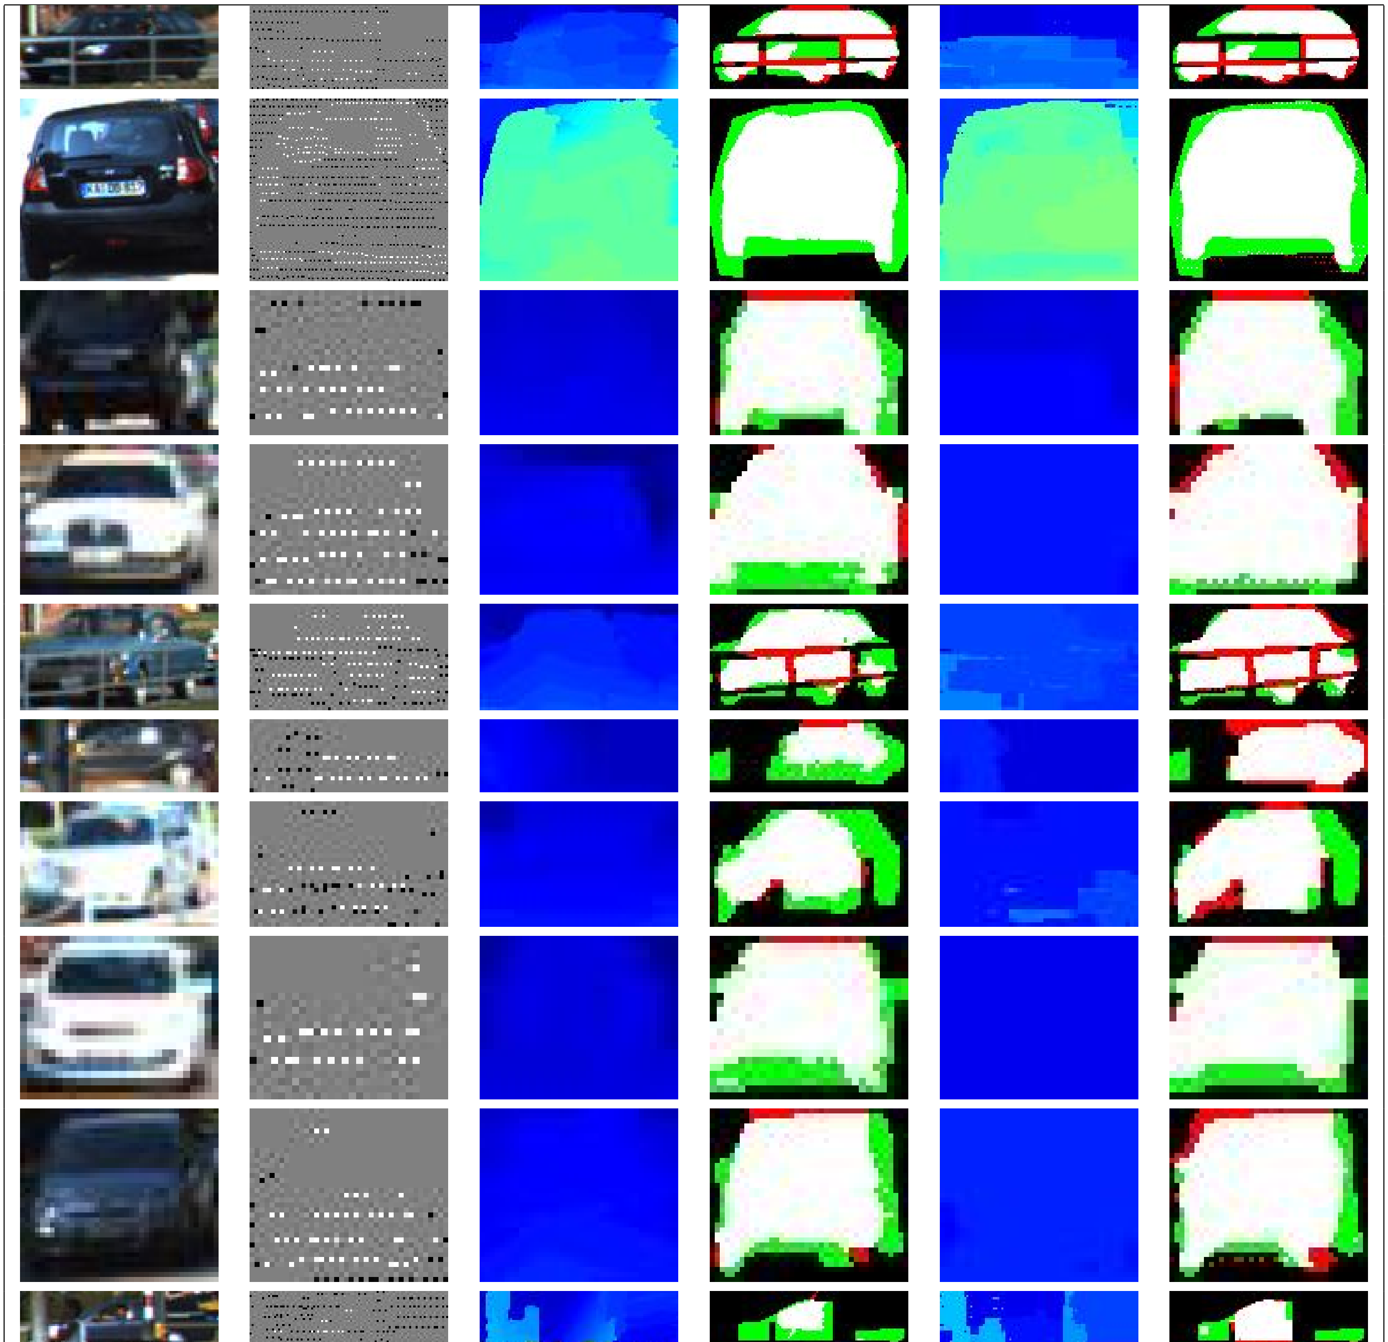

Figure 10: **Failure modes:** For each row, we show original image, projected point clouds (White: car, Black: background), stereo depth, segmentation results by our full model with stereo (White: True Positive, Black: True Negative, Red: False Positive, Green: False Negative), depth images reconstructed by Laplacian MRF, and segmentation results by our full model with LIDAR. Note that all images are rescaled for display.