Supplementary Material

1. Running-time Analysis

In this supplementary material, we provide more running-time analysis for the models that we used in the ablation studies. We use the same setup as in Section 4.4. We measure the running time for models with and without auto-registration or box merging and scoring. Table 1 shows the results. Optimizing the implementation is not the focus of this work. However, the running time results indicate that both auto-registration and box merging and score do not largely increase the running time.

2. More Qualitative Results on Auto-Registration

In Figure1, we show more qualitative results on the *test* set of the KITTI dataset. In Figure2, we show more visualization of the auto-registration on the KITTI *validation* split. The settings follow the same as in Section 4.4. The same observation can be made that the vertex positions with added offsets move towards the center of vehicles. The offset is computed based on vertex states, and it moves the vertex position closer to the center of the car regardless of the original vertex position. The offset reduces the original translation of the vertex position, thus reduces the translation variance.

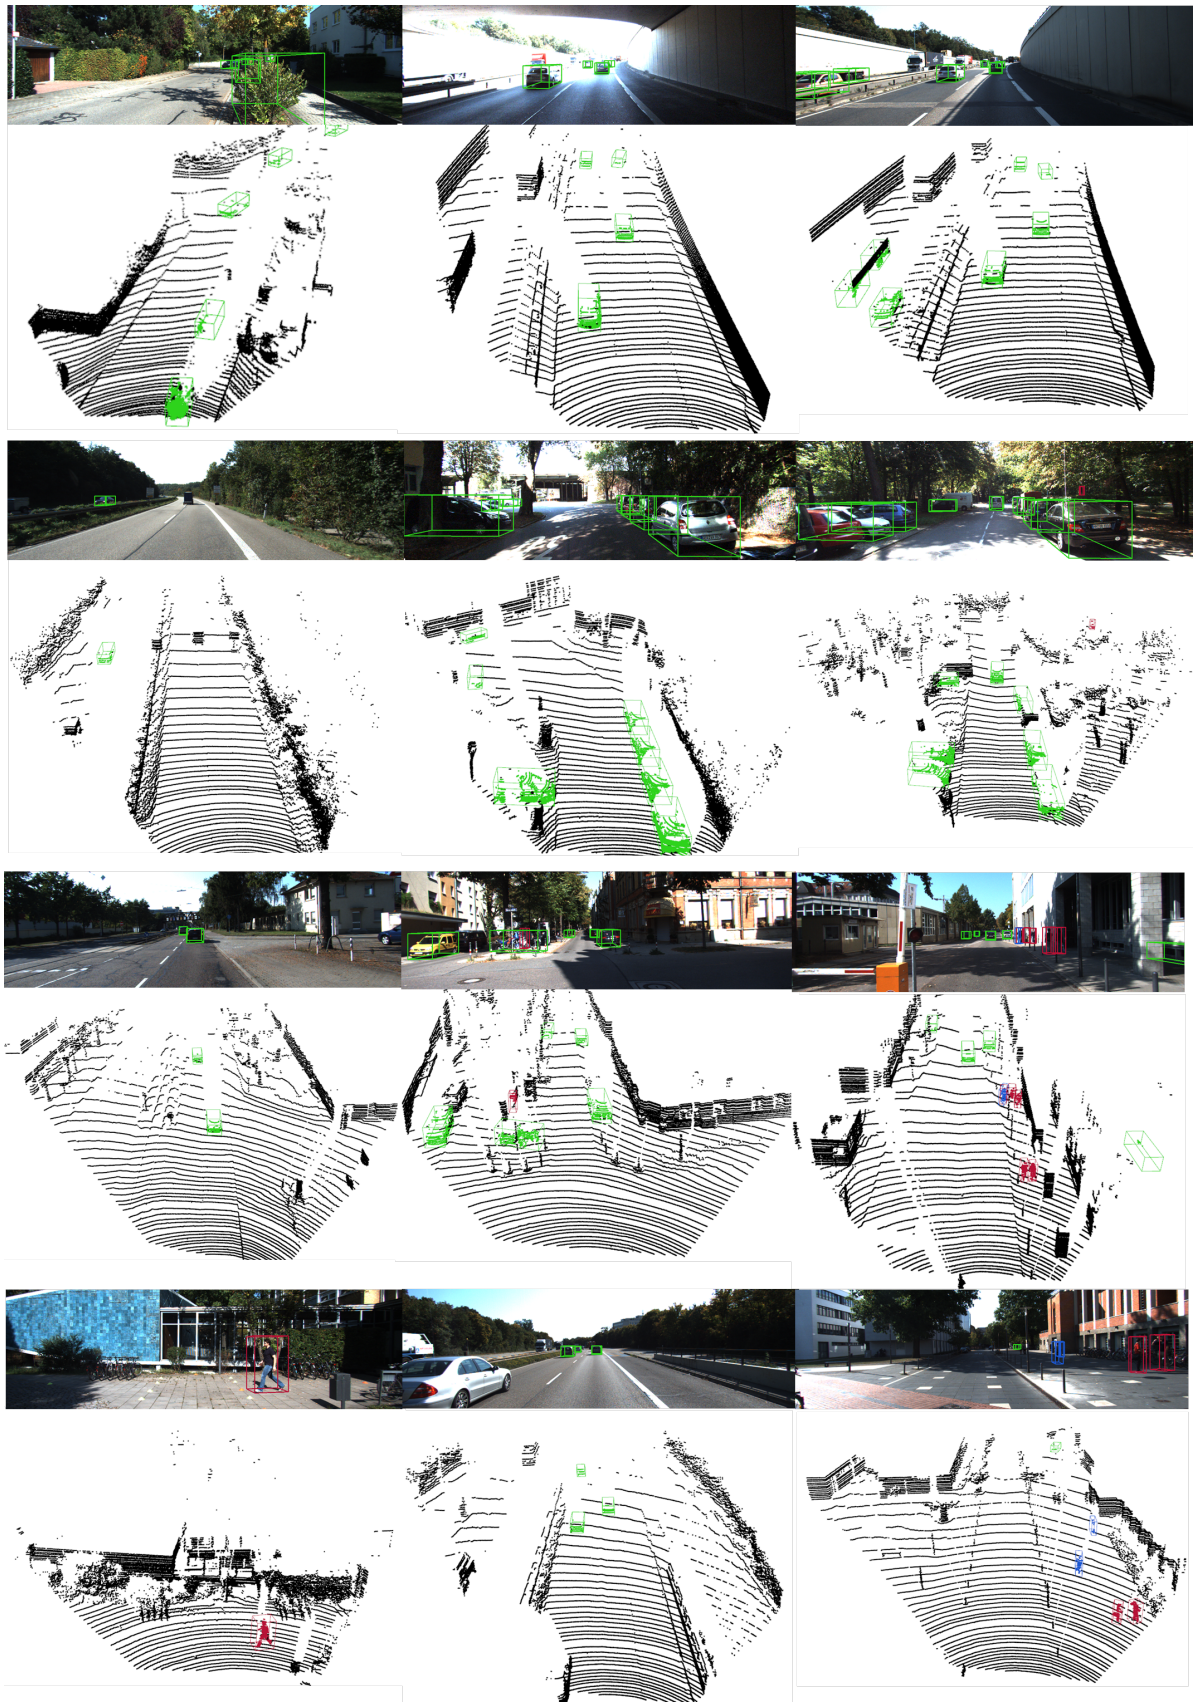

Figure 1. More qualitative results on the KITTI *test* dataset using Point-GNN. We show the predicted 3D bounding box of Cars (green), Pedestrians (red) and Cyclists (blue) on both the image and the point cloud. Best viewed in color.

Box Merging and Scoring	Auto Registration	Running Time					
		Read Input	Graph Construction	GNN	Box Decoding	Box Merging and Score	Total
-	-	48 ms	123 ms	361 ms	4 ms	48 ms	584 ms
-	✓	48 ms	123 ms	362 ms	4 ms	50 ms	587 ms
✓	-	71 ms	125 ms	362 ms	4 ms	83 ms	645 ms
✓	✓	70 ms	121 ms	363 ms	5 ms	84 ms	643 ms

Table 1. The average running time for one sample in the KITTI *validation* split.

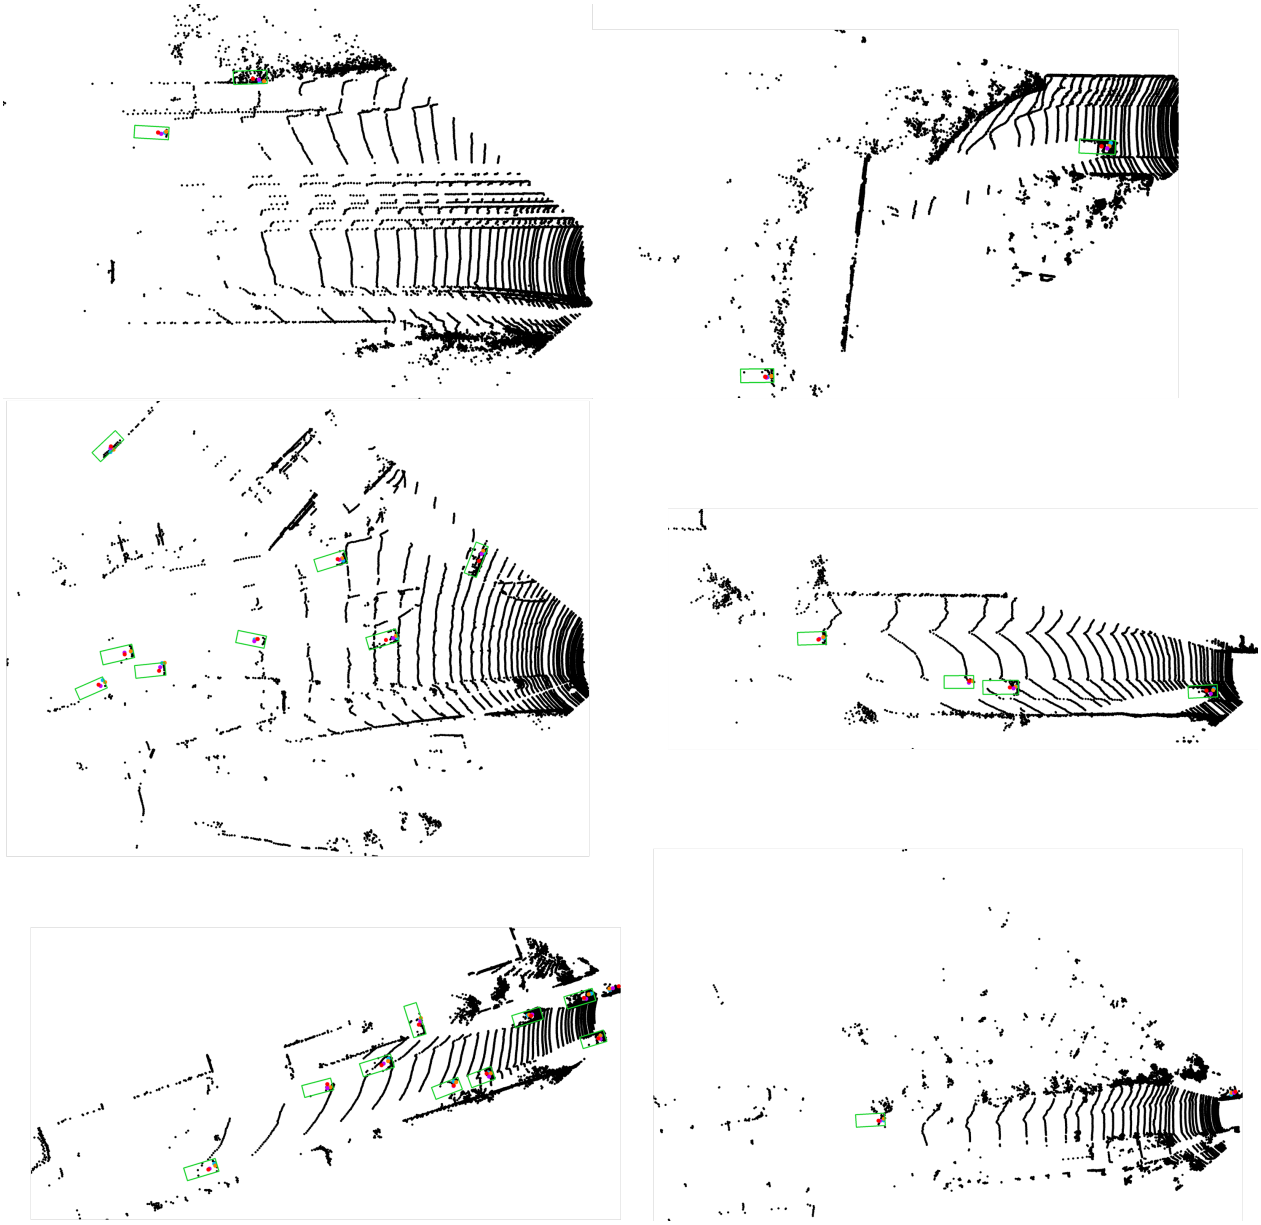

Figure 2. More examples from the *validation* split showing the vertex locations with added offsets. The blue dot indicates the original position of the vertices. The orange, purple, and red dots indicate the original position with added offsets from the first, the second, and the third graph neural network iterations. Best viewed in color.