Supplemental Material for
VCSeg: Virtual Camera Adaptation for Road Segmentation

1. Qualitative Results

In the paper, we show qualitative results of all segmentation algorithms for the images from each of the five datasets on which VCSeg obtains the *highest* IoU (successes) and also on the images on which VCSeg obtains the *lowest* IoU (failures). Fig. 1 below complements these demonstrations by showing qualitative results for the three images on which VCSeg obtains its *median* (typical) performance. This figure demonstrates that VCSeg typically yields superior qualitative results to SegNet, GeoSeg and AdaptSegNet.

2. Projective Completion

VCSeg depends upon a novel projective completion method to complete gaps created when reprojecting images to the virtual camera domain. Fig. 2 shows more examples of the projectively completed images..

To demonstrate the impact on performance, Fig. 3 shows examples of VCSeg segmentations on Cityscapes images computed with and without projective completion. We see that segmentations without projective completion are consistently more fragmented and less accurate.

3. Visualizing Shift in Geometric Camera Parameters

In the standard approach to road segmentation (e.g., SegNet, GeoSeg, AdaptSegNet), images are fed into a fixed network without regard to the intrinsic or extrinsic parameters of the camera. VCSeg attempts to improve on this by correcting for geometric shift prior to input. To visualize the effect of this intervention on the input, we can average the imagery both within and between datasets, with and without the VCSeg correction (Fig. 4). Note that the scene structure is much more crisply apparent with geometric correction than without, particularly for the Qian dataset, where variations in camera parameters are most extreme.

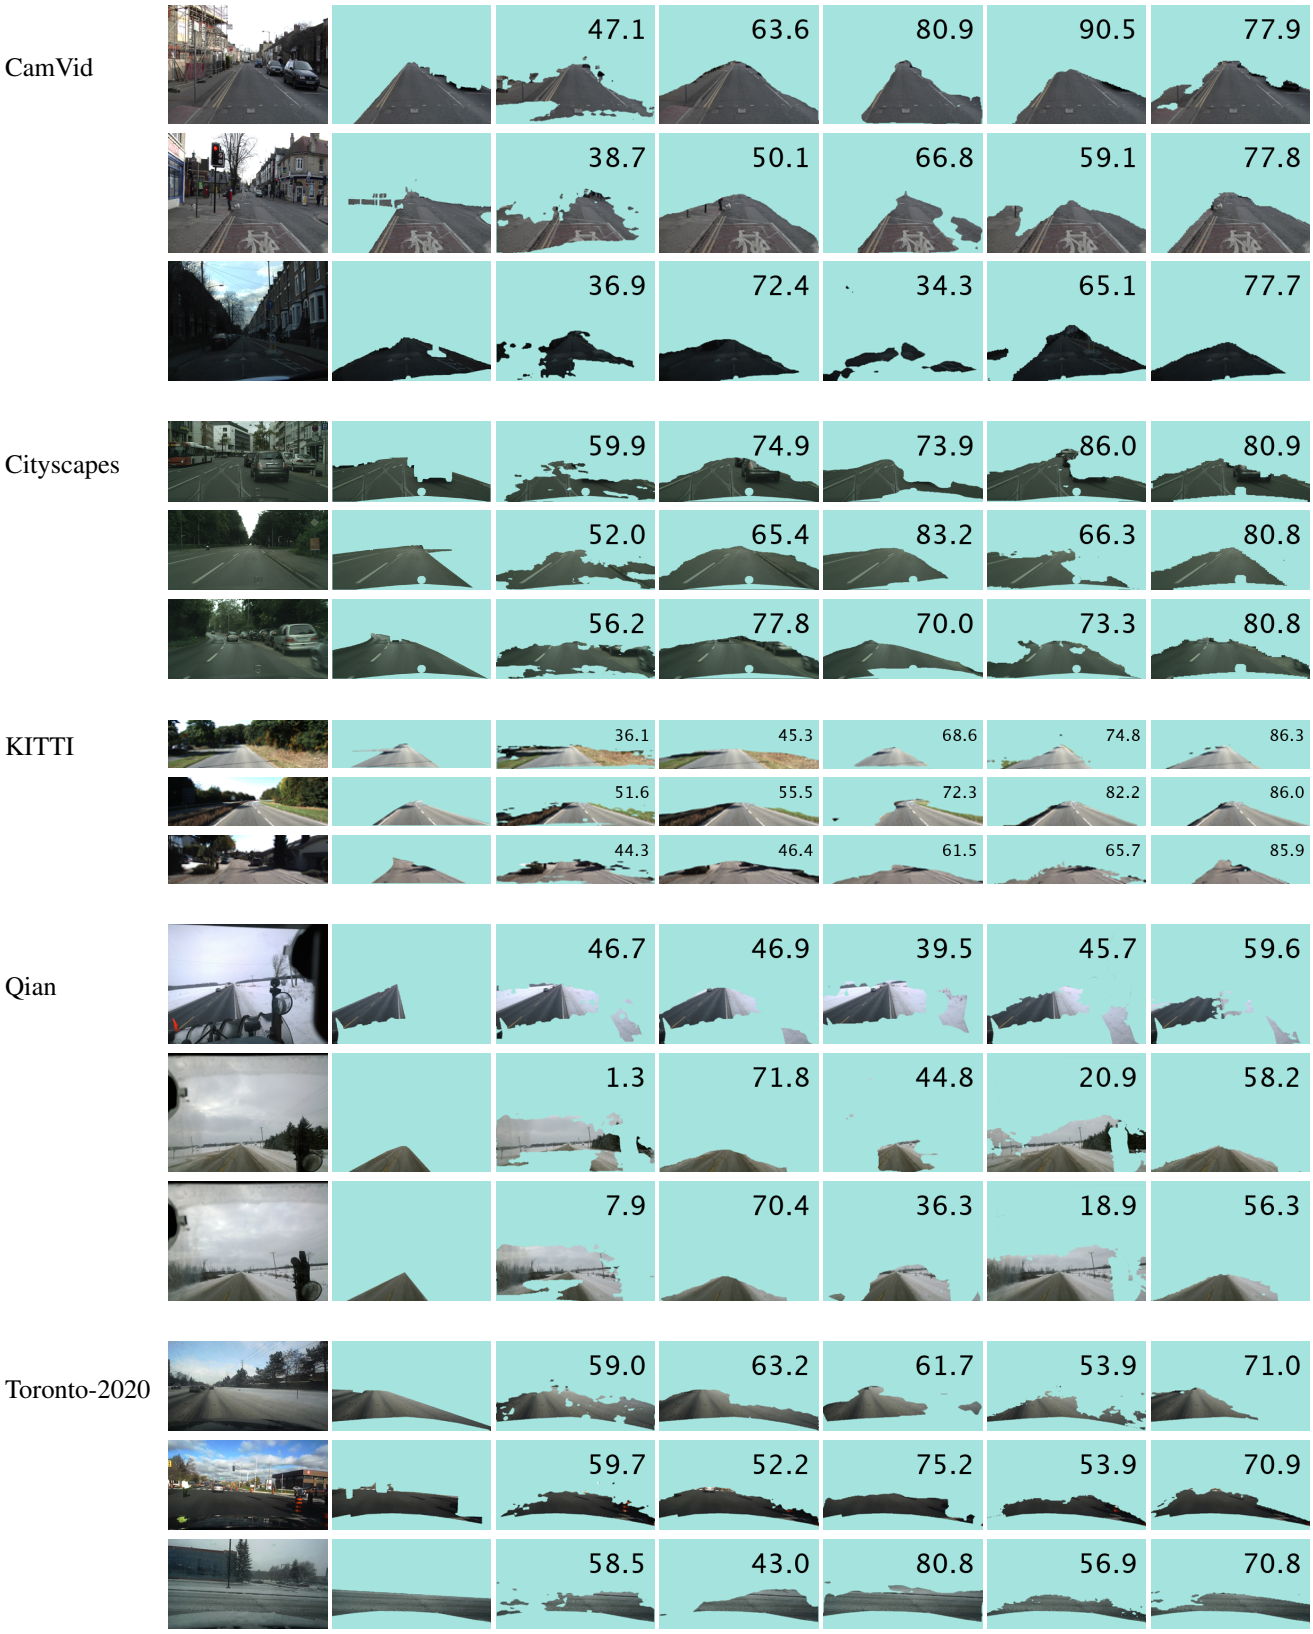

Figure 1. Medians: Three images from each dataset for which our full VCSeg method obtains roughly median IoU. The IoU obtained by each algorithm is indicated in the top right corner of its segmentation.

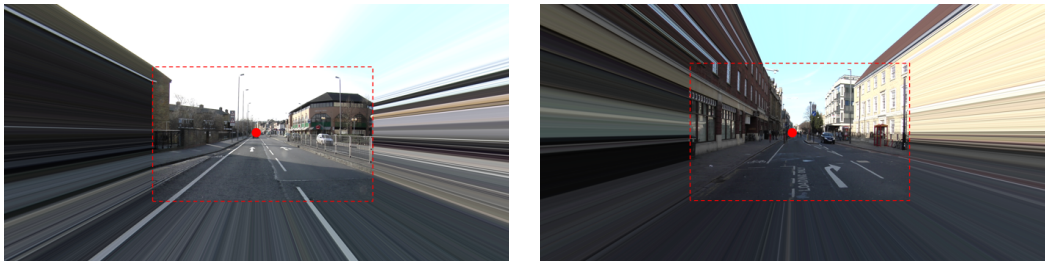

CamVid

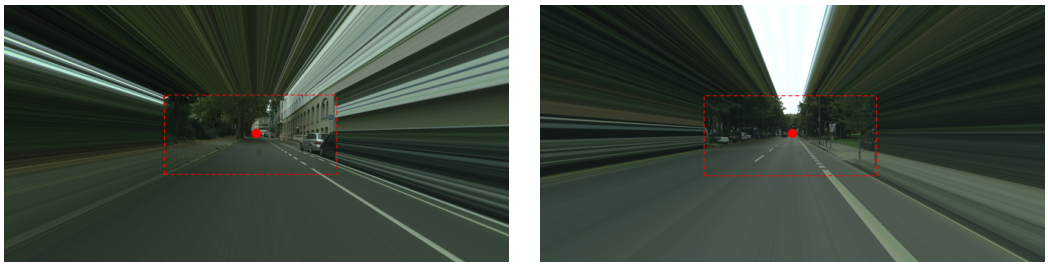

Cityscapes

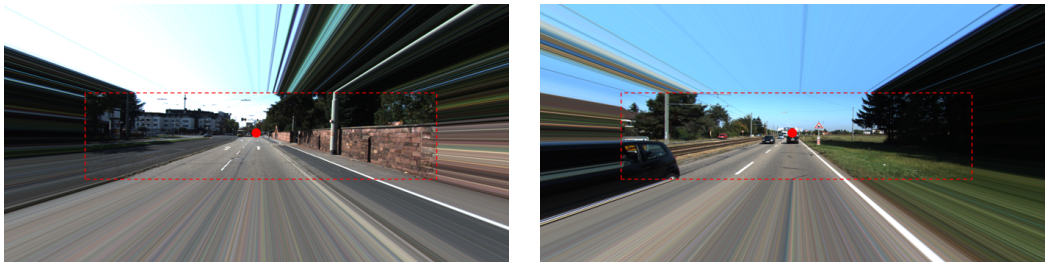

KITTI

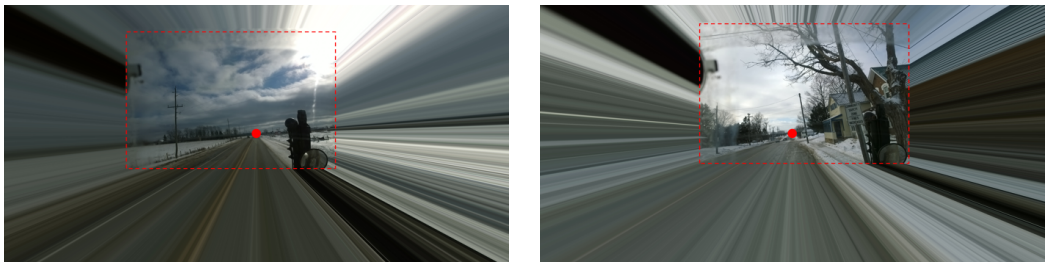

Qian

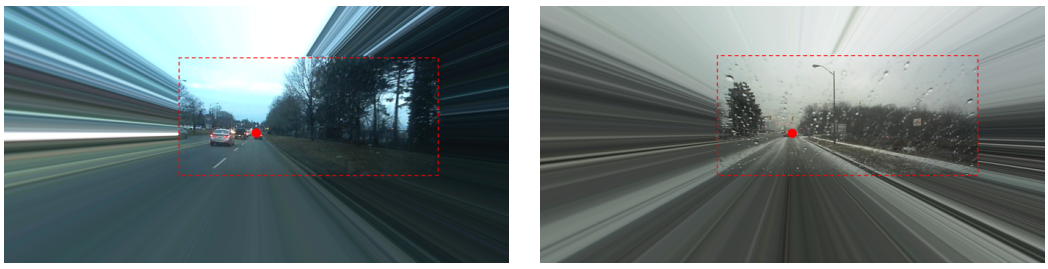

Toronto 2020

Figure 2. Example projective completions. Red dot indicates ground truth vanishing point and red lines indicate original image boundaries.

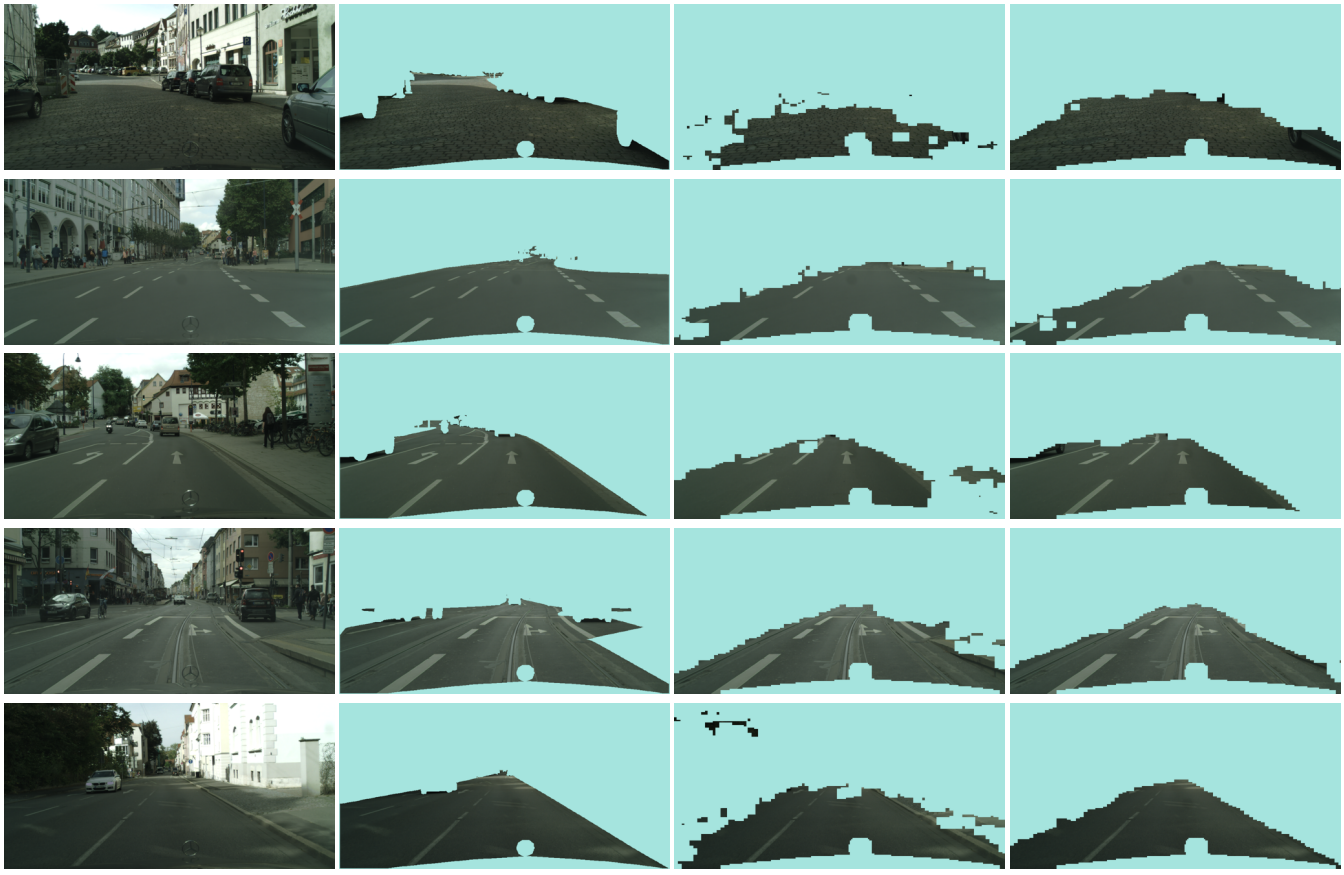

Image Ground truth VCSeg (no completion) VCSeg (with completion)

Figure 3. Demonstration of the importance of projective completion on images from the Cityscapes dataset.

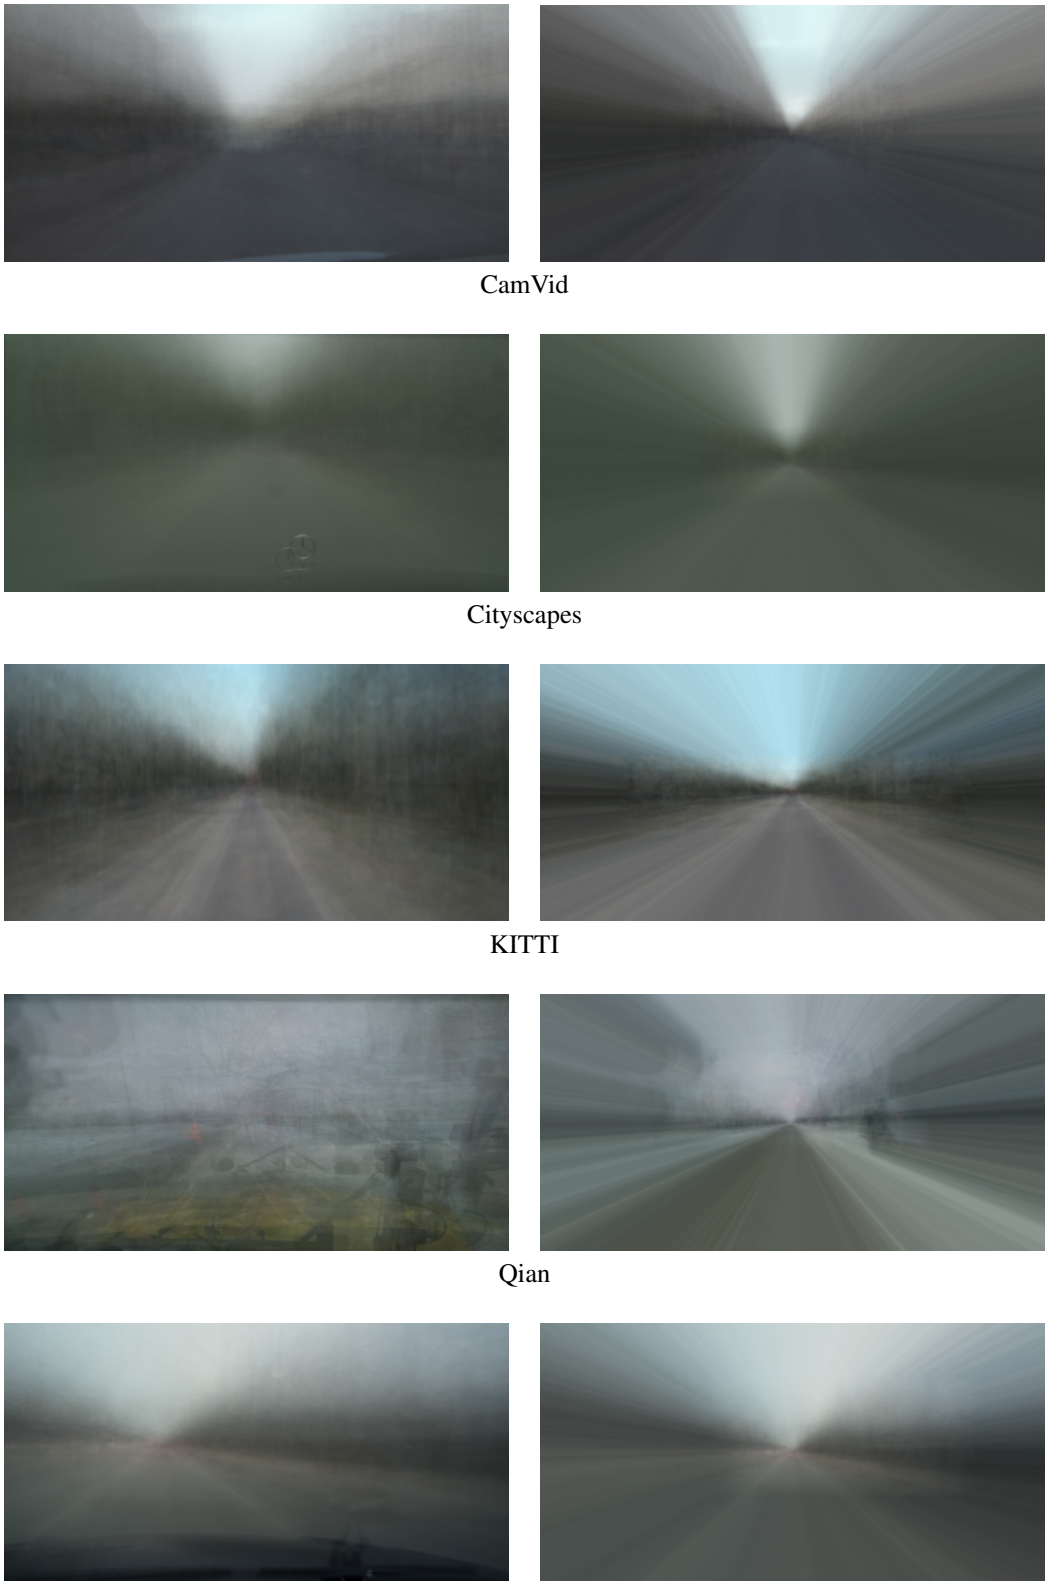

Figure 4. Averages across training datasets, computed in original source domains and in the virtual camera domain, after reprojection, correcting for both between-camera and within-camera shift. Source domain images have been reshaped to match virtual camera dimensions to ease comparison.