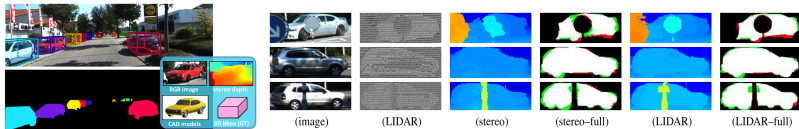


Beat the MTurkers: Automatic Image Labeling from Weak 3D Supervision

L.-C. Chen, S. Fidler, A. L. Yuille, and R. Urtasun (CVPR 2014)



- ▶ Automatically segmentation of objects given annotated 3D bounding boxes
- ▶ Inference in a binary MRF using appearance models, stereo and/or noisy point clouds, 3D CAD models, and topological constraints
- ▶ 10 to 20 labeled objects to train the system
- ▶ Evaluated using 3D boxes available on KITTI
- ▶ 86% IOU score on segmenting cars (performance of MTurkers)
- ▶ It can be used to de-noise MTurk annotations.
- ▶ Segmenting big cars is easier than smaller ones.
- ▶ Each potential increases performance (CAD model most).
- ▶ Same performance with stereo or LIDAR (highest using both)
- ▶ Fast: 2 min for training and 44 seconds for full test set
- ▶ Robust to low-resolution, saturation, noise, sparse point clouds, depth estimation errors and occlusions