

Amodal Completion and Size Constancy in Natural Scenes: Supplementary Material

Abhishek Kar, Shubham Tulsiani, João Carreira and Jitendra Malik
University of California, Berkeley - Berkeley, CA 94720
{akar, shubhtuls, carreira, malik}@eecs.berkeley.edu

1. Additional Results

We present more qualitative results of our amodal completion and size inference algorithm in this document. All the results are obtained by using our class agnostic amodal completion network on ground truth boxes in the PASCAL VOC *val* set and estimating sizes using our size inference algorithm on these amodal boxes. Each row shows the original image, the ground truth box annotations in PASCAL and finally the amodal boxes along with inferred sizes and the horizon (shown by the horizontal red line).

We also show the distribution of camera heights inferred using our algorithm in Figure 1 and the focal lengths used for training our CNN focal length predictor in Figure 2.




Figure 1: Distribution of camera heights as inferred on PASCAL VOC. It can be seen that the distribution is peaked around the height at which humans normally take pictures (1.4m) with a long tail.

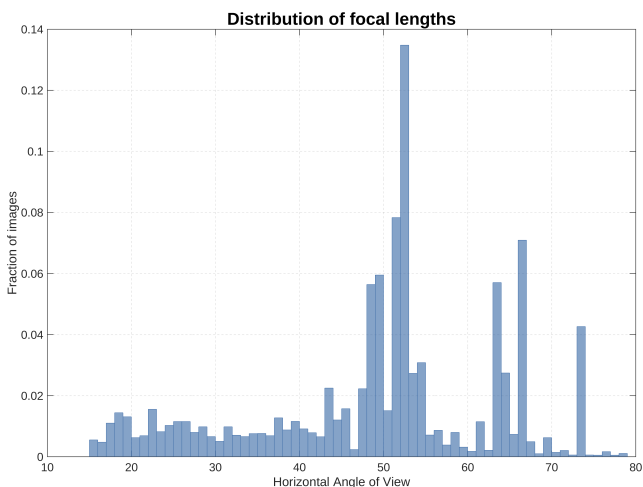


Figure 2: Distribution of camera focal lengths in the Places dataset shown as horizontal angle of view. The peaks in the distribution correspond to canonical focal lengths in popular wide angle and telephoto lenses. It can be seen that the images cover the full spectrum from wide angle views of scenes (corresponding to the right end of the plot) to close-ups (left end of the histogram). Interestingly, the largest peak in the distribution is close to the field of view of iPhones.

