

Teaching Compositionality to CNNs (supplement material)

Austin Stone
Yi Liu

Huayan Wang
D. Scott Phoenix

Michael Stark
Dileep George

Vicarious FPC, San Francisco, CA, USA

{austin, huayan, michael, yi, scott, dileep}@vicarious.com

Additional visualizations

In Figure 6 of the paper we showed guided-backprop [39] visualization on a number of images that better illustrate the effect of adding our compositionality constraint. To confirm that this observation is consistent over a large portion of images, we sampled 243 images completely at random from MS-COCO validation set and applied that same visualization (see Figure 1-9).




Figure 1: Same as Figure 6 in the main paper.




Figure 2: Same as Figure 6 in the main paper.




Figure 3: Same as Figure 6 in the main paper.




Figure 4: Same as Figure 6 in the main paper.




Figure 5: Same as Figure 6 in the main paper.




Figure 6: Same as Figure 6 in the main paper.




Figure 7: Same as Figure 6 in the main paper.




Figure 8: Same as Figure 6 in the main paper.




Figure 9: Same as Figure 6 in the main paper.