

SSD++ for multiple Object Detection from clutter and indoor shots

Computer Vision (CS6350)
TPA - 9

1 Problem Statement

Localization of different types of objects in indoor scenes with the help of bounding-boxes using Deep Learning techniques. The images may include challenges like occlusion, background clutter, camera shake, varied object size and illumination condition. The aim of the project is to detect and recognize objects in real-time, taking live feed from web-camera (for live demo). Models based on recent state-of-the-art methods like YOLO v3 [6], RefineDet [7], Relation-Net [8], RFBNet [9], CornerNet [10] etc. or anything new showing substantial performance over the existing techniques are expected.

2 Input

- Images containing objects along with class-labels and bounding box information.

3 Assumptions

- Images may contain multiple objects of same or different classes.

4 Output

- *Detect, identify and visualize* the location of each of the objects in the input image using *bounding boxes*.
- Obtain the *class label* of the detected objects.
- Calculate *Rank-1 recognition rate* of each of the objects detected.
- *Online demo* taking live feed from web-camera.
- *Off-line performance* on a set of test images (to be provided during evaluation) using *mean Average Precision* (mAP@IoU=0.5) as evaluation metric.

5 Datasets

- *ImageNet*, *MS COCO* and *Pascal VOC (2011)* to be used for training followed by fine-tuning in *VPLab-Indoor* Dataset (to be provided).

6 References

1. Jifeng Dai, Yi Li, Kaiming He, and Jian Sun. “R-FCN: Object Detection via Region-based Fully Convolutional Networks.”, NIPS 2016.
2. Joseph Redmon and Ali Farhadi. “YOLO9000: Better, Faster, Stronger.” CVPR 2017.
3. Tsung-Yi Lin, Piotr Dollr, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. “Feature pyramid networks for object detection.” CVPR 2017.
4. Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollr. “Focal Loss for Dense Object Detection.” ICCV 2017.
5. Kaiming He, Georgia Gkioxari, Piotr Dollr, and Ross Girshick. “Mask r-cnn.” ICCV 2017.
6. Joseph Redmon and Ali Farhadi. “YOLOv3: An Incremental Improvement.” arXiv 2018.
7. Shifeng Zhang, Longyin Wen, Xiao Bian, Zhen Lei, and Stan Z. Li. “Single-Shot Refinement Neural Network for Object Detection.” CVPR 2018.
8. Han Hu, Jiayuan Gu, Zheng Zhang, Jifeng Dai, and Yichen Wei. “Relation networks for object detection.” CVPR 2018.
9. Songtao Liu, and Di Huang. “Receptive field block net for accurate and fast object detection.” ECCV 2018.
10. Hei Law, and Jia Deng. “Cornersnet: Detecting objects as paired keypoints.” ECCV 2018.