Deep Residual Learning for Image Recognition

Authors: Kaiming He, Xiangyu Zhang, Shaoqing Ren, Jian Sun

Presenter: Venkata Ravi Teja Ailavarapu

Very Deep Networks are useful

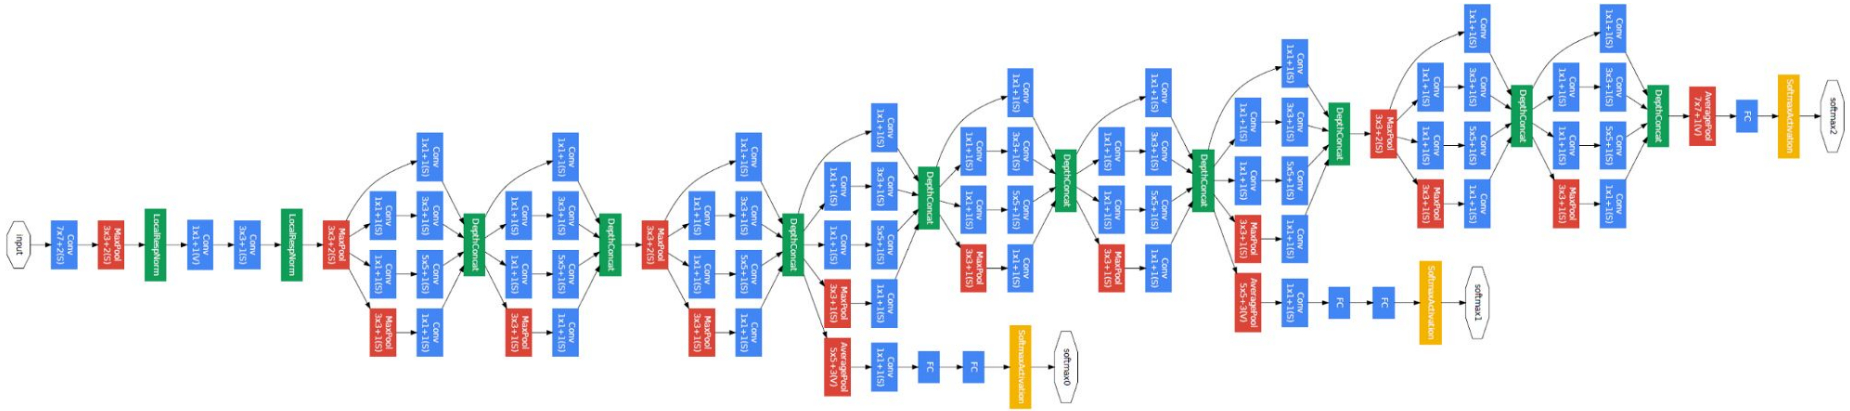

Image Courtesy: Szegedy, Christian, et al. "Going deeper with convolutions." *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2015.

GoogLeNet

Very Deep Networks are hard to train

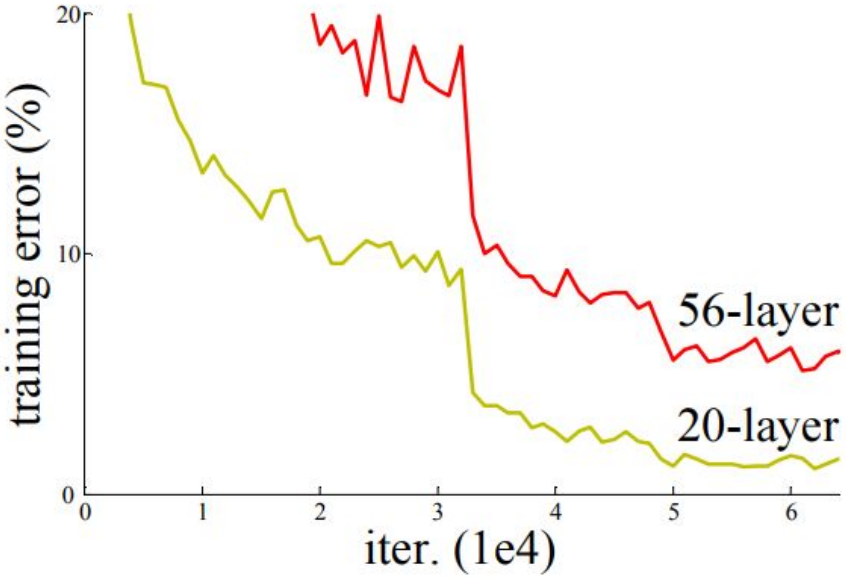

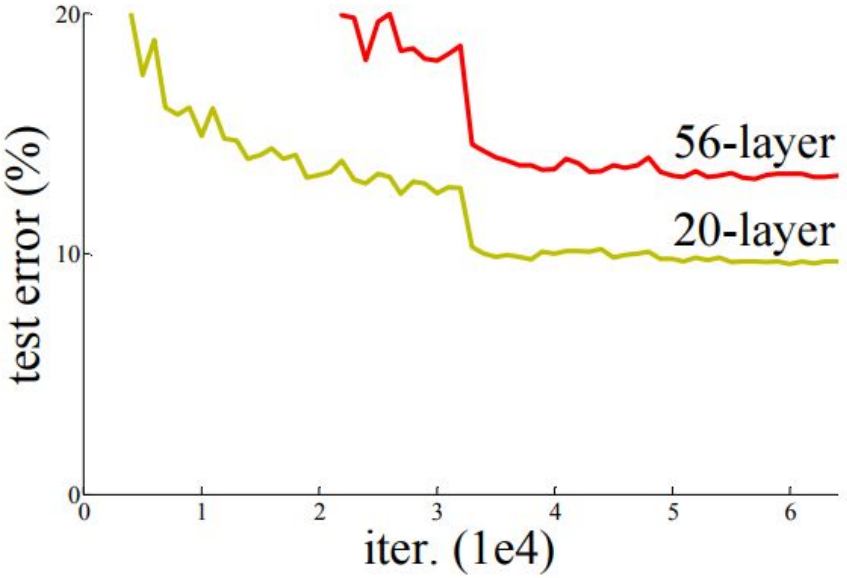

Degradation Problem

Solution by Construction

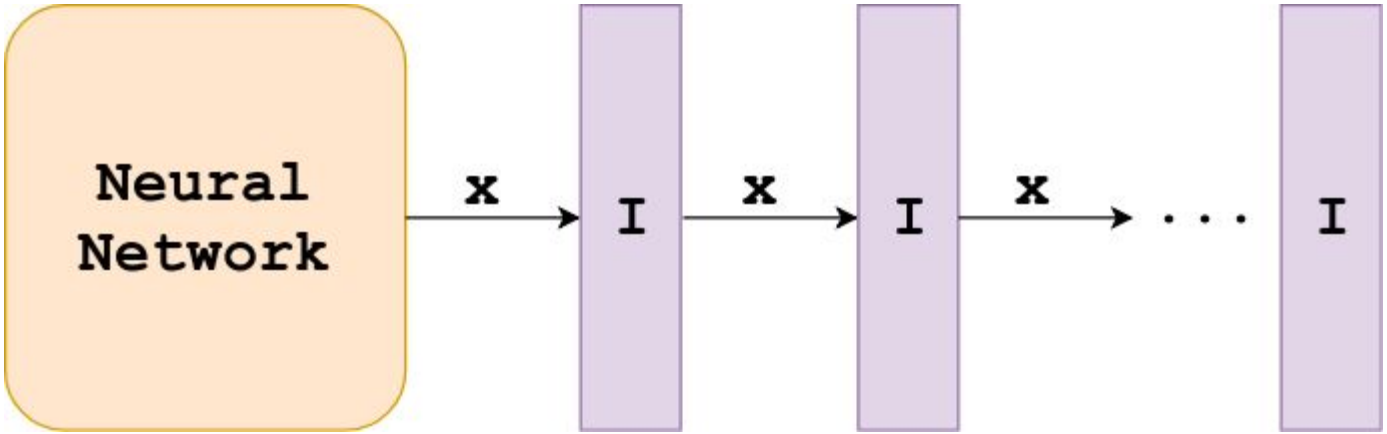

Residual Block

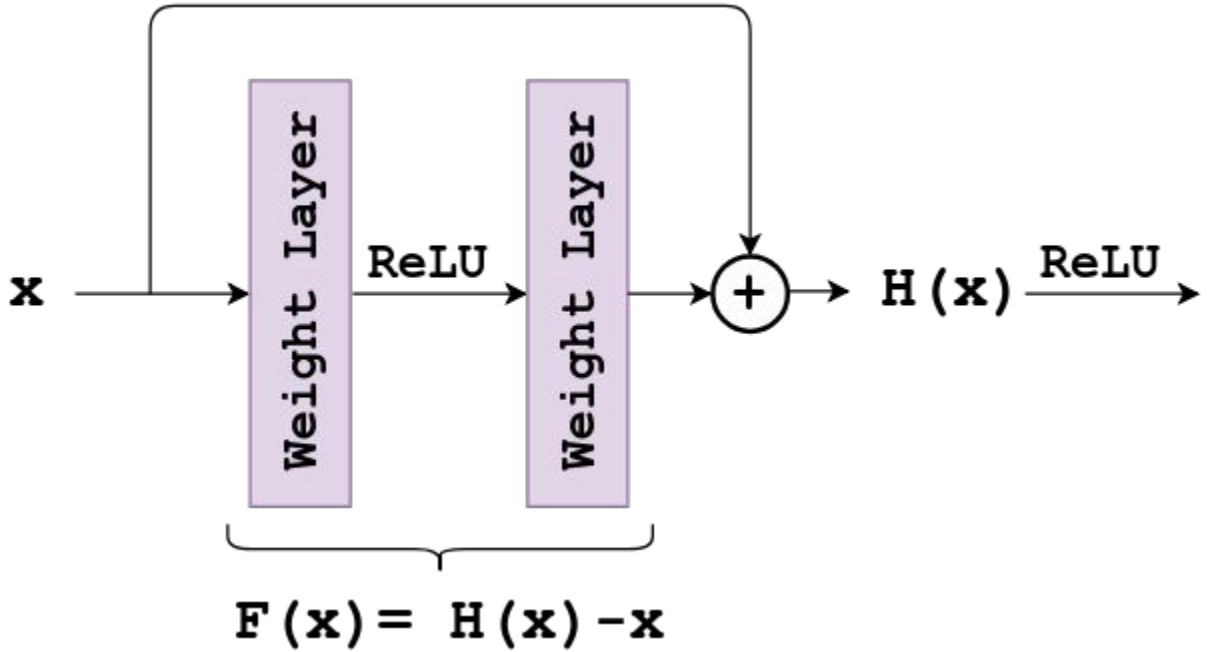

ResNet Architecture

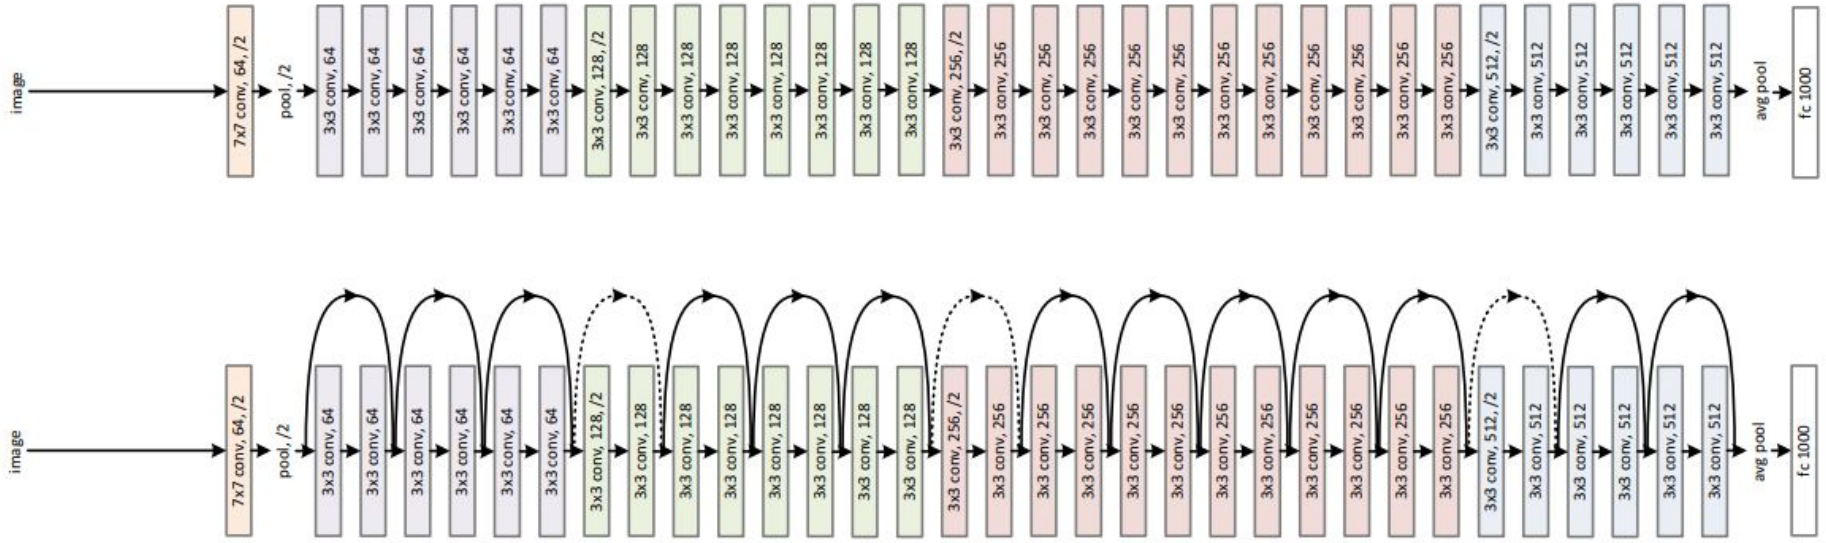

Deep Bottleneck Architecture

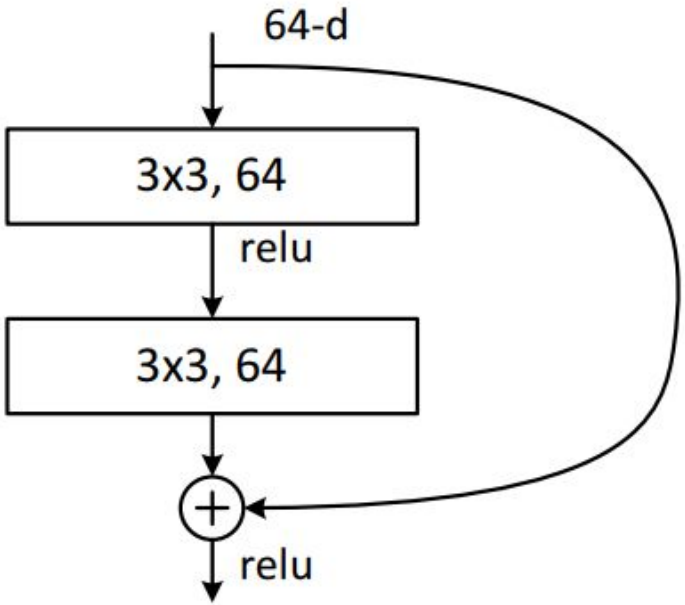

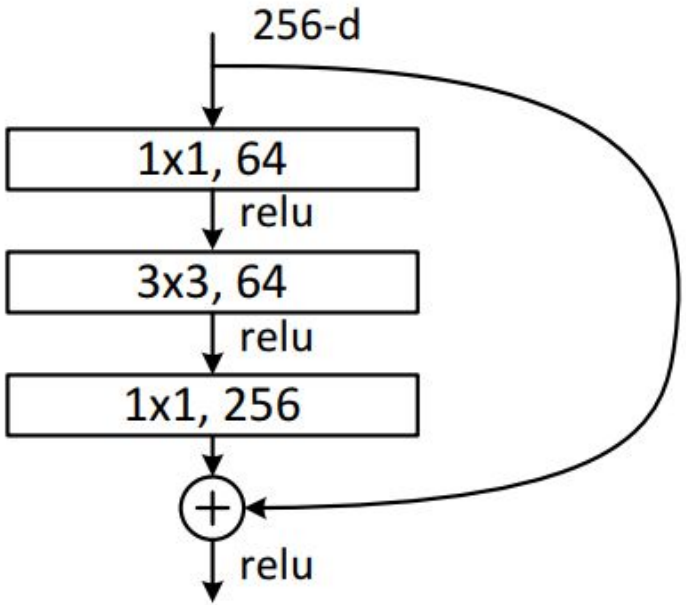

Training Error in Deep Networks

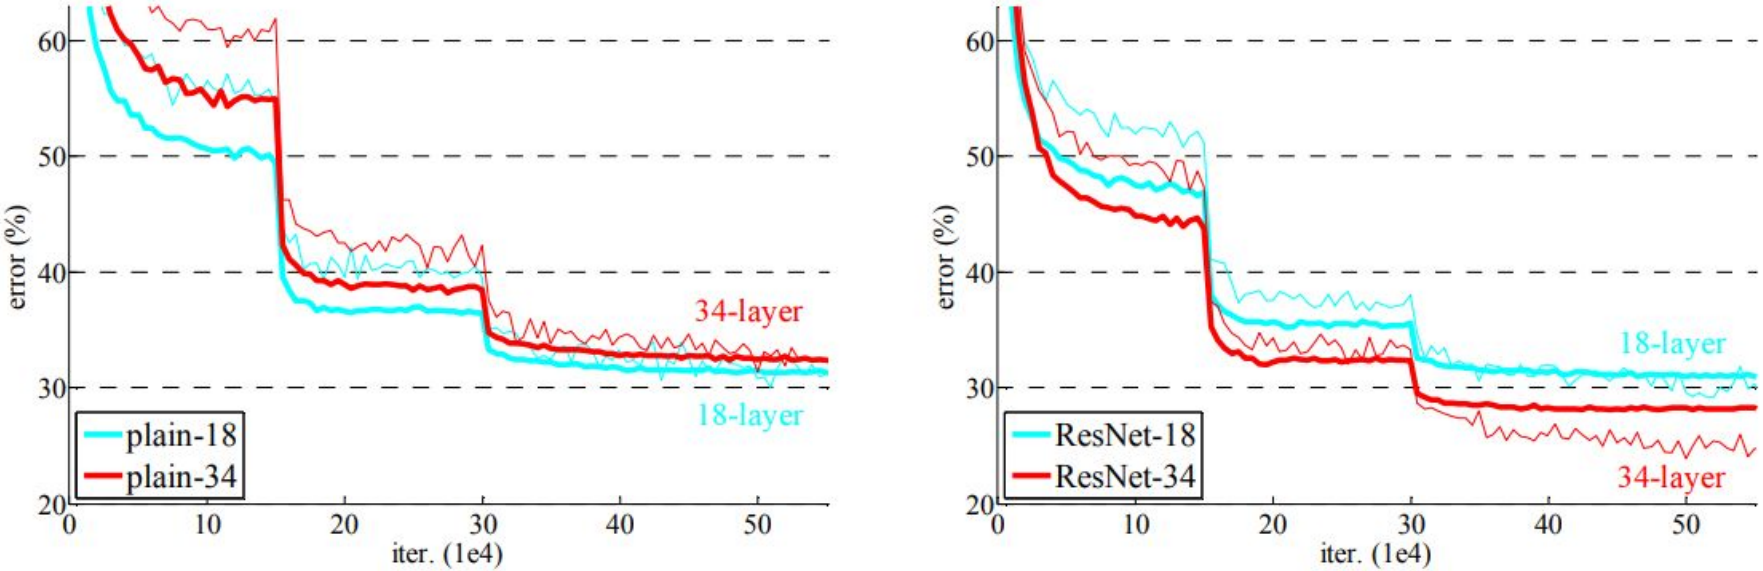

Training Results on ImageNet

Results on Deep Networks

method	top-1 err.	top-5 err.
VGG [40] (ILSVRC'14)	-	8.43†
GoogLeNet [43] (ILSVRC'14)	-	7.89
VGG [40] (v5)	24.4	7.1
PReLU-net [12]	21.59	5.71
BN-inception [16]	21.99	5.81
ResNet-34 B	21.84	5.71
ResNet-34 C	21.53	5.60
ResNet-50	20.74	5.25
ResNet-101	19.87	4.60
ResNet-152	**19.38**	**4.49**

Table 4. Error rates (%) of **single-model** results on the ImageNet validation set (except † reported on the test set).

Pros

Experiments

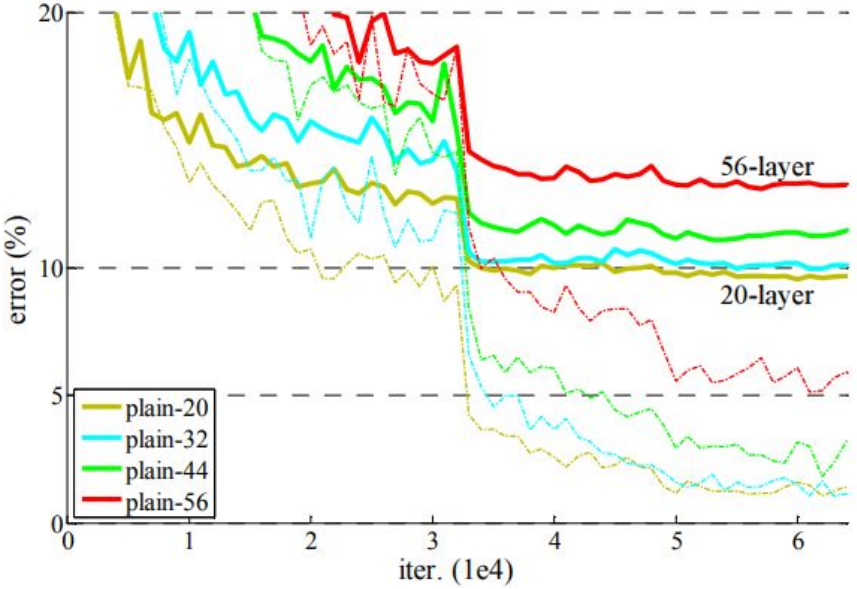

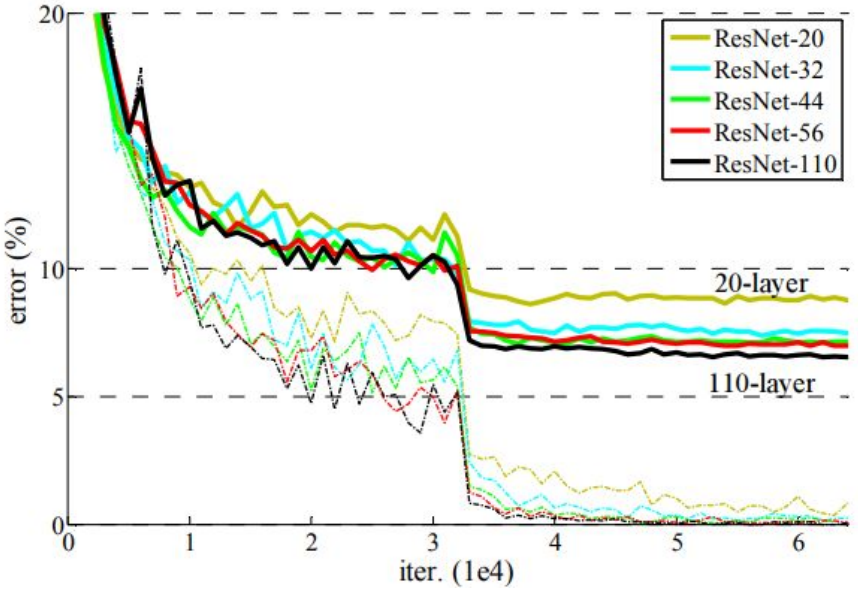

Analysis of Deep(er) Networks

Fair Comparison

	plain	ResNet
18 layers	27.94	27.88
34 layers	28.54	**25.03**

Table 2. Top-1 error (%, 10-crop testing) on ImageNet validation. Here the ResNets have no extra parameter compared to their plain counterparts. Fig. 4 shows the training procedures.

Applicability

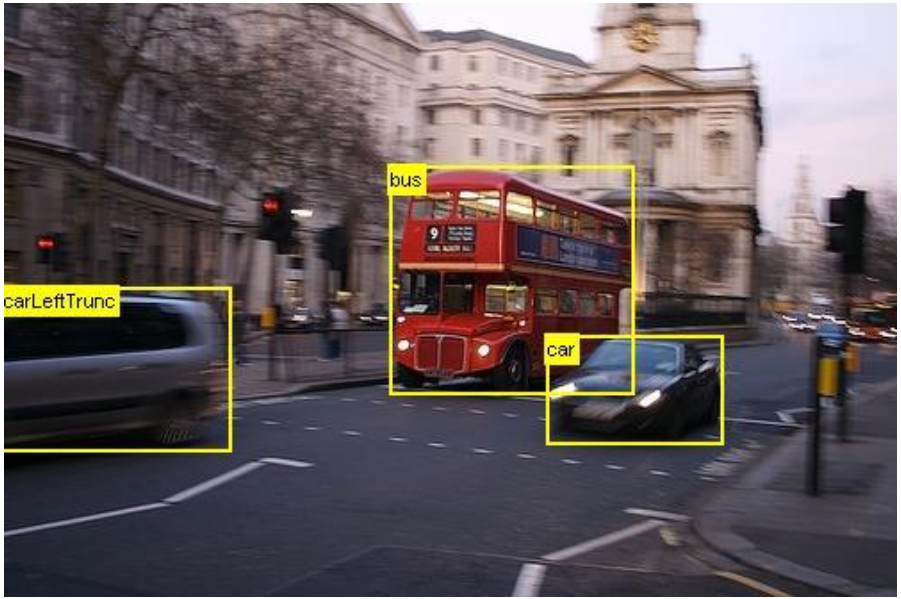

Image Courtesy: http://host.robots.ox.ac.uk/pascal/VOC/voc2007/

Image Courtesy: Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems* (pp. 1097-1105).

Paper is well-structured