

Visualizing and Understanding Convolutional Networks

Large Convolutional Network models have recently demonstrated impressive classification performance on the ImageNet Benchmark (Krizhevsky et al., 2012). However there is no clear understanding of why they perform so well, or how they might be improved. In this paper we address both issues. We introduce a novel visualization technique that gives insight into the function of intermediate feature layers and the operation of the classifier. Used in a diagnostic role, these visualizations allow us to find model architectures that outperform Krizhevsky et al. on the ImageNet classification benchmark. We also perform an ablation study to discover the performance contribution from different model layers. We show our ImageNet model generalizes well to other datasets: when the softmax classifier is retrained, it convincingly beats the current state-of-the-art results on Caltech-101 and Caltech-256 datasets.

Author: Mr MATTHEW, D. Zeiler (New York University)

Presenter: Mr MATTHEW, D. Zeiler (New York University)