

Analyzing the Efficiency of Subsequent Convolutional Layers with Small-Scale Images

Pagdanganan, Anjo (School: Salinas High School)

Convolutional neural networks (ConvNets) are currently the state-of-the-art model for image classification; however, training them is computationally expensive. In addition, there are no general guidelines on what makes an accurate ConvNet. In this experiment, I attempted to find the optimal depth of convolutional layers (conv. layers) between pooling layers in order to maximize training efficiency on small images. To do this, four ConvNets were trained on the CIFAR-10 dataset (Krizhevsky et al.), each model n with n conv. layers between each pooling layers. The accuracy improvements between each model were then compared using the Kolmogorov-Smirnov test. I found that there was no significant improvement in accuracy between models #3 and #4. Thus, when training on small images, up to three conv. layers should be used between pooling layers. These findings could be used in situations requiring rapid prototyping of image classifiers, such as disease detection with low-resolution images.