Contribution ID: **14**   Type: **not specified**

ImageNet Classification with Deep Convolutional

We trained a large, deep convolutional neural network to classify the 1.2 million high-resolution images in the ImageNet LSVRC-2010 contest into the 1000 different classes. On the test data, we achieved top-1 and top-5 error rates of 37.5% and 17.0% which is considerably better than the previous state-of-the-art. The neural network, which has 60 million parameters and 650,000 neurons, consists f five convolutional layers, some of which are followed by max-pooling layers, nd three fully-connected layers with a final 1000-way softmax. To make training aster, we used non-saturating neurons and a very efficient GPU implementation of the convolution operation. To reduce overfitting in the fully-connected layers we employed a recently-developed regularization method called "dropout"that proved to be very effective. We also entered a variant of this model in the ILSVRC-2012 competition and achieved a winning top-5 test error rate of 15.3%, compared to 26.2% achieved by the second-best entry.

Author: Mr KRIZHEVSKY, Alex (University of Toronto)

Presenter: Mr KRIZHEVSKY, Alex (University of Toronto)