

# ImageNet Classification with Deep Convolutional Neural Networks

Alex Krizhevsky  
Ilya Sutskever  
Geoffrey Hinton

University of Toronto  
Canada

# ImageNet

- 15M images in 22K categories
- For this contest: 1.2M images in 1K categories
- Classification: make 5 guesses about label




# Results

Team	Score
SuperVision	<b>0.15315</b>
ISI	0.26172
OXFORD_VGG	0.26979
XRCE/INRIA	0.27058
University of Amsterdam	0.29576
LEAR-XRCE	0.34464

# Overview of our model

- **Deep:** 7 hidden “weight” layers
- **Learned:** all feature extractors initialized at white Gaussian noise and learned from the data
- Entirely supervised
- More data = good



**Convolutional layer:** convolves its input with a bank of 3D filters, then applies point-wise non-linearity



**Fully-connected layer:** applies linear filters to its input, then applies point-wise non-linearity

# Overview of our model


- Trained with stochastic gradient descent on two NVIDIA GPUs for about a week
- 650,000 neurons
- 60,000,000 parameters
- 630,000,000 connections
- **Final feature layer: 4096-dimensional**



**Convolutional layer:** convolves its input with a bank of 3D filters, then applies point-wise non-linearity



**Fully-connected layer:** applies linear filters to its input, then applies point-wise non-linearity



Image