## How transferable are features in deep neural

Many deep neural networks trained on natural images exhibit a curious phenomenon

in common: on the first layer they learn features similar to Gabor filters

and color blobs. Such first-layer features appear not to be specific to a particular dataset or task, but general in that they are applicable to many datasets and tasks. Features must eventually transition from general to specific by the last layer of the network, but this transition has not been studied extensively. In this paper we experimentally quantify the generality versus specificity of neurons in each layer of a deep convolutional neural network and report a few surprising results.

Author: Mr JASON, Yosinski (Cornell University)

**Presenter:** Mr JASON, Yosinski (Cornell University)

Track Classification: Machine Learning