


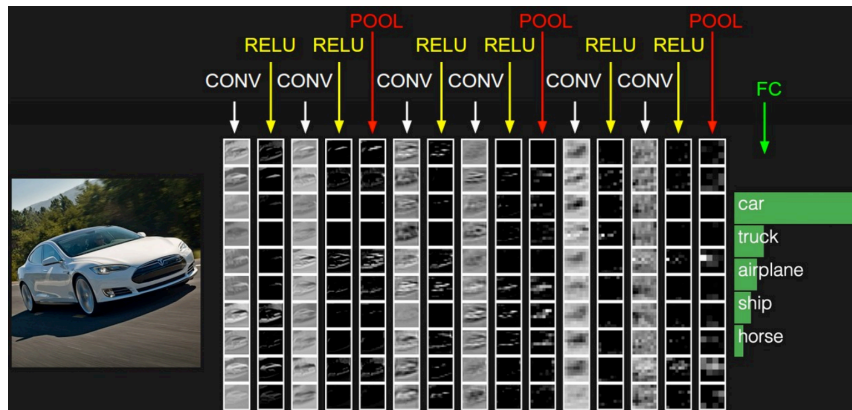
Convolutional Neural Networks (CNNs)

Fei-Fei Li, Justin Johnson, Serena Yeung (<http://cs231n.stanford.edu/>)

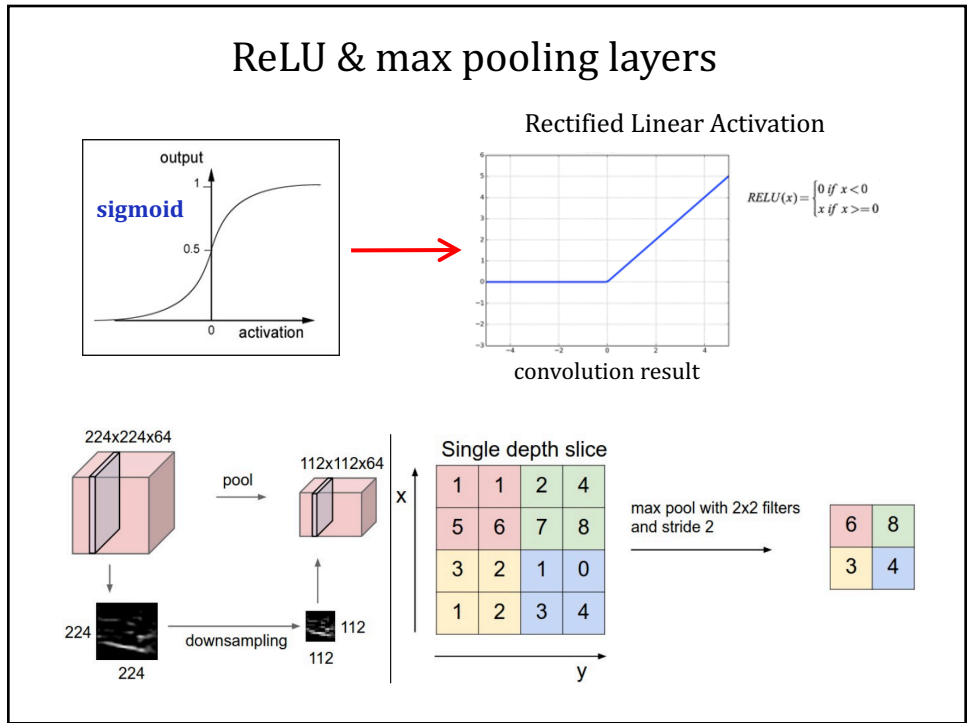
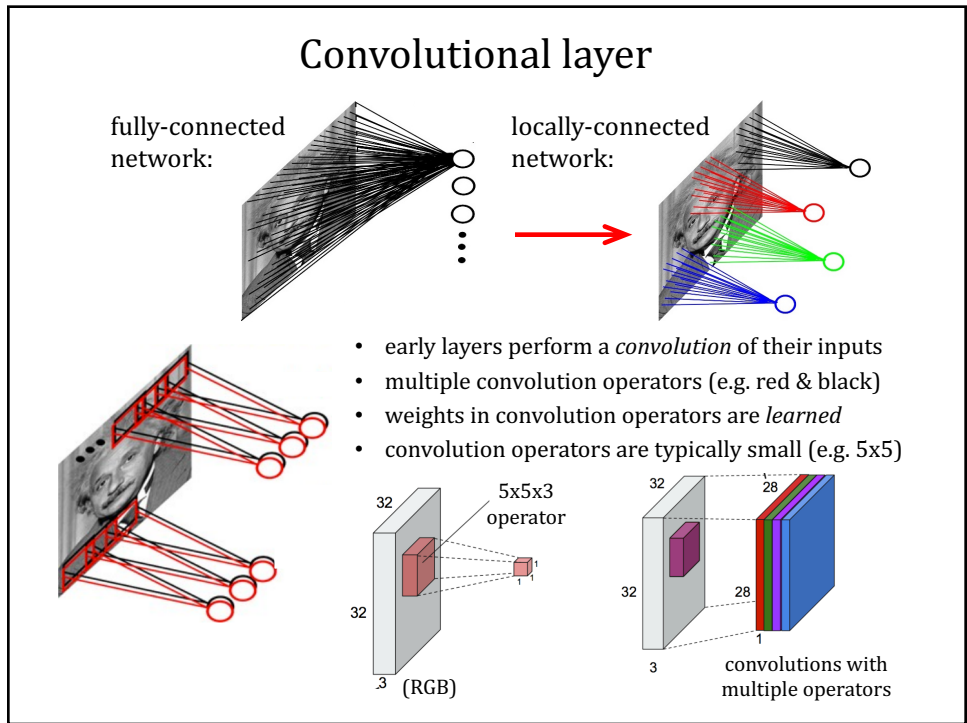


The Convolutional Neural Network in this example is classifying images live in your browser using Javascript, at about 10 milliseconds per image. It takes an input image and transforms it through a series of functions into class probabilities at the end. The transformed representations in this visualization can be loosely thought of as the activations of the neurons along the way. The parameters of this function are learned with backpropagation on a dataset of (image, label) pairs. This particular network is classifying CIFAR-10 images into one of 10 classes and was trained with ConvNetJS. Its exact architecture is [conv-relu-conv-relu-pool]x3-fc-softmax, for a total of 17 layers and 7000 parameters. It uses 3x3 convolutions and 2x2 pooling regions. By the end of the class, you will know exactly what all these numbers mean.

Sample stages of a CNN

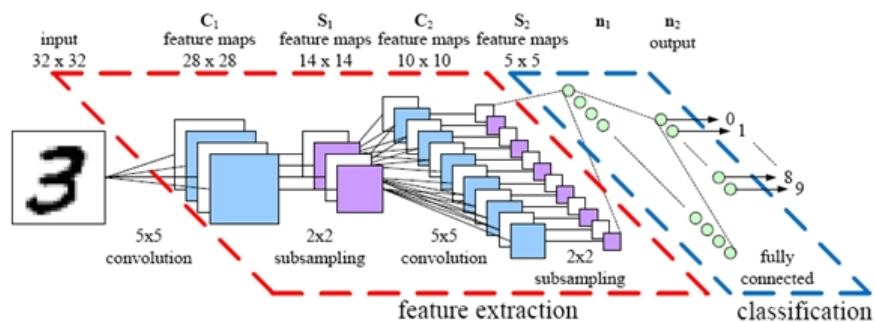


- CONV: "convolution" layer with weights that are learned
- RELU: "rectified linear unit" applies an activation function
- POOL: "pooling" selects maximum value in small neighborhoods
- FC: "fully-connected" neural network



Adding a fully-connected neural net layer


Recognizing digits from the MNIST database with a CNN:



LeNet

LeCun, Bottou, Bengio, Haffner (1998)

AlexNet, ZF Net, GoogLeNet, VGGNet, ResNet, ...




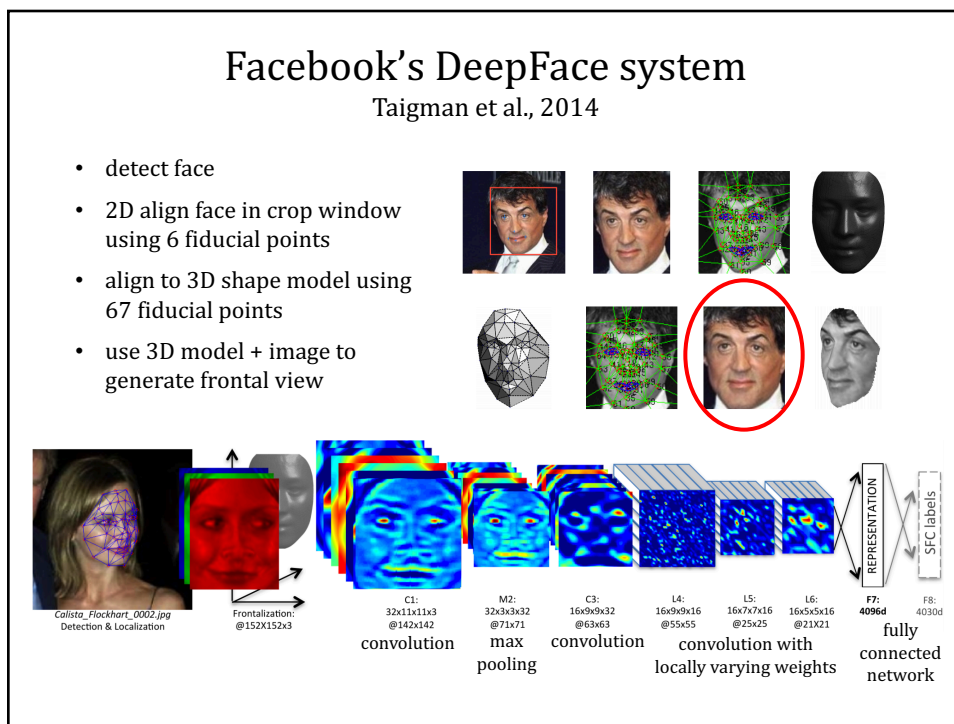
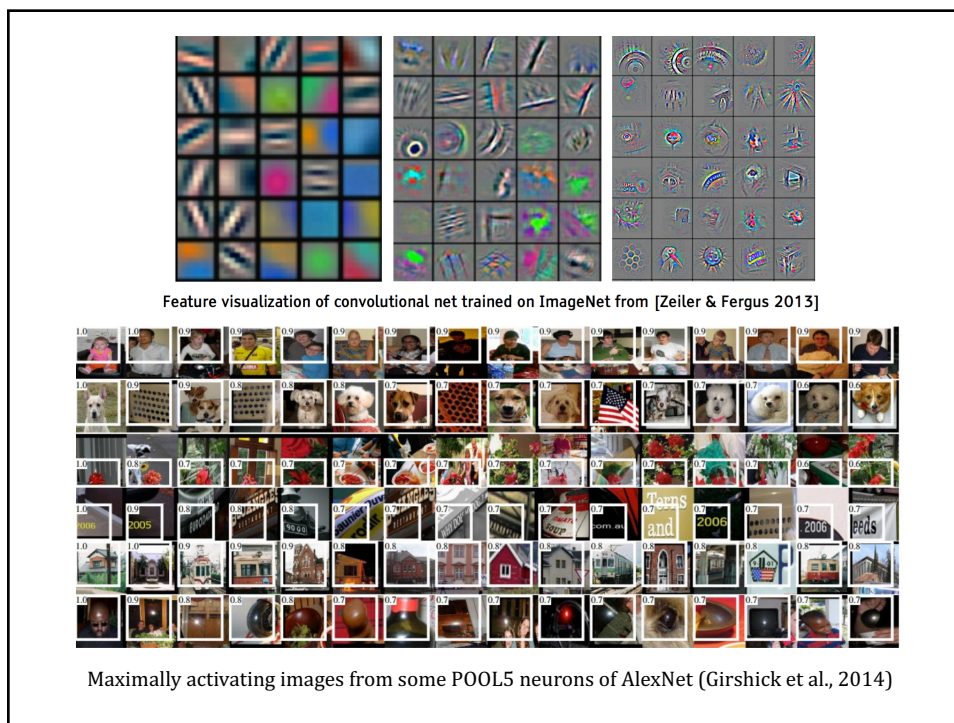
AlexNet: Krizhevsky, Sutskever, Hinton (2012)

ImageNet Large Scale Visual Recognition Challenge (ILSVRC)

IMAGENET


Annually since 2010






Google's FaceNet system Schroff et al., 2015

FaceNet also uses a deep convolutional network



threshold = 1.1 classifies pairs correctly
(smaller value means more similar)

- learns mapping from images to a space where distance between images captures similarity
- training data: triplets of face thumbnails
 - two same ID, one different ID
- learning process: minimize distance between anchor & positive images (same ID), maximize distance between anchor & negative images



OpenFace implementation of FaceNet system Amos et al., 2015

<https://cmusatyalab.github.io/openface/>

OpenFace

Home Demos - User Guide - DNN Models - Release Notes

OpenFace 5,417 Stars

Free and open source face recognition with deep neural networks.

News

- 2016-09-15: We presented OpenFace in the *Data (after)Lives* art exhibit at the University of Pittsburgh and have released the code as [Demo 4: Real-time Face Embedding Visualization](#).
- 2016-08-09: [New blog post: \(Face\) Image Completion with Deep Learning in TensorFlow](#). ([OpenFace group discussion on it](#))
- 2016-06-01: [OpenFace tech report released](#)
- 2016-01-19: OpenFace 0.2.0 released! See [this blog post](#) for more details.

OpenFace is a Python and [Torch](#) implementation of face recognition with deep neural networks and is based on the CVPR 2015 paper [FaceNet: A Unified Embedding for Face Recognition and Clustering](#) by Florian Schroff, Dmitry Kalenichenko, and James Philbin at Google. Torch allows the network to be executed on a CPU or with CUDA.

Crafted by [Brandon Amos](#), [Bartosz Ludwiczuk](#), and [Mahadev Satyanarayanan](#).

- The code is available on GitHub at [cmusatyalab/openface](#).
- [API Documentation](#)
- Join the [cmu-openface group](#) or the [gitter chat](#) for discussions and installation issues.
- Development discussions and bugs reports are on the [issue tracker](#).