Computer Vision

Neurobio 230

Bill Lotter

Exciting time: Neuroscience ⇔ computer vision

-Traditionally: computer vision relied on hand crafted features

-Today: “Deep Learning”

- -loosely based on how the brain does computations

- -most of components learned from data

- -a lot of commonalities between computer vision models and the visual ventral stream in the brain

Overview of Computer Vision Problems

Object Recognition

Object Detection

Image Segmentation

Face Identification

Action Recognition

Video Prediction

...

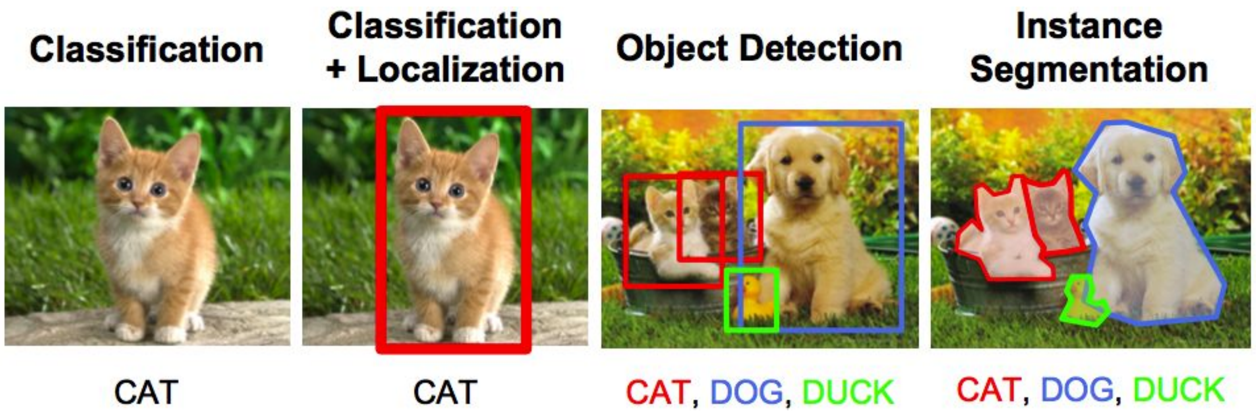

Stanford CS 224d

Common Testbeds for Computer Vision

MNIST

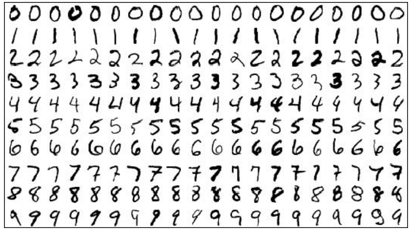

MS COCO

Imagenet

General Problem Formulation

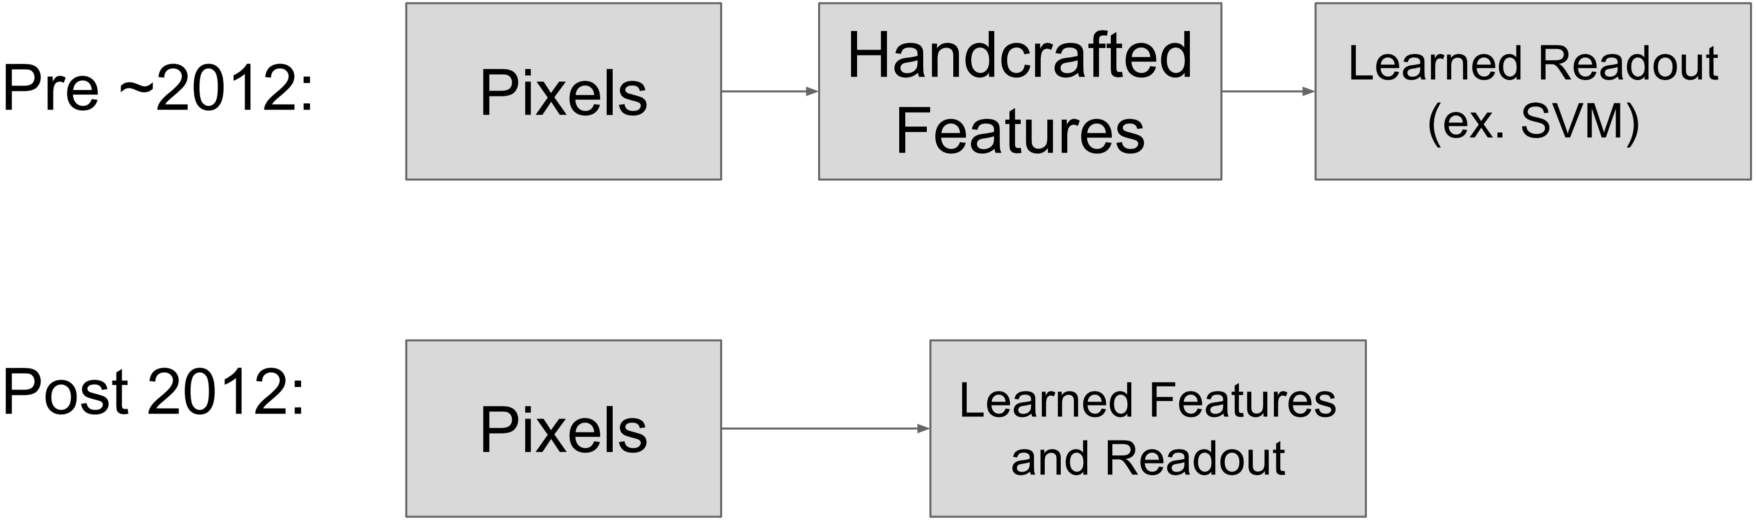

Focusing on Object Recognition: Convolutional Neural Networks (CNNs)

Background:

- Hubel and Wiesel Simple and Complex Cells (1959, 1960s)
- Neocognitron (Fukushima, 1980)
- HMAX (Riesenhuber & Poggio 1999, Serre, Kreiman et al. 2007)
- Yann LeCun’s work on MNIST with CNNs (1998)

What is an Artificial Neural Network?

a lot of variations, hard to generalize, but a simple ANN looks something like this..

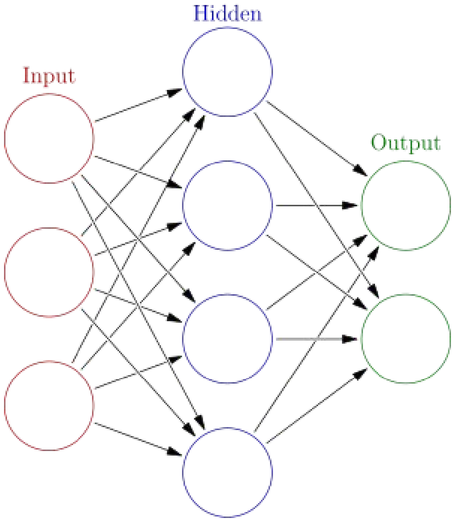

Training the Network: Backprop

Backpropagation (Rumelhart, Hinton, Williams 1986): way to calculate gradient of error in terms of network parameters

Today: gradient descent with some bells and whistles

Formulating for object recognition...

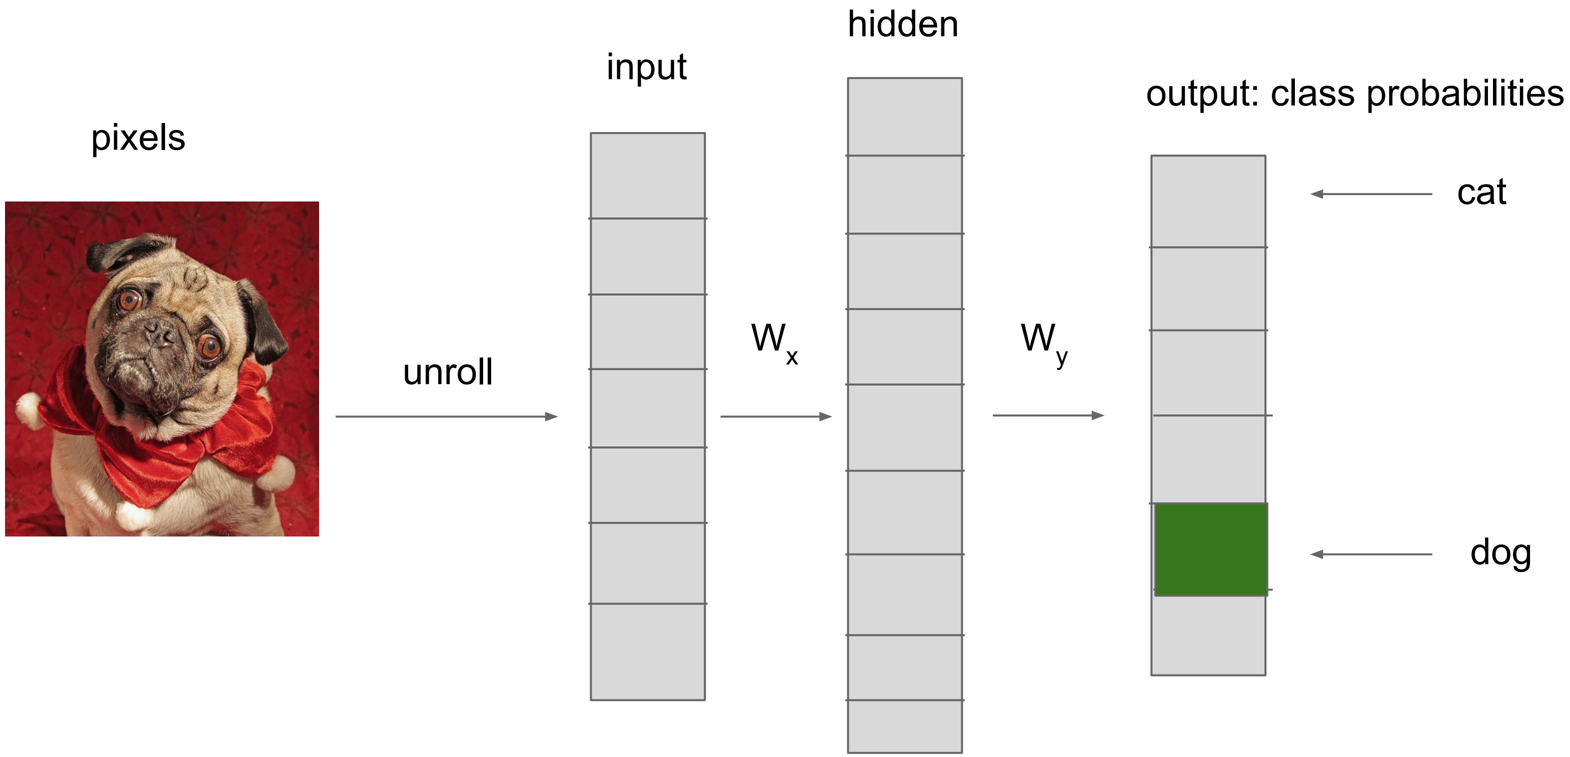

Taking a look at parameters..

image: 256x256x3 = 196,608 inputs

outputs: 1000 categories

even if just go directly from image to outputs:

    1000 x 196,608 = 196 million params!!

even if you have 1 million training images, you would severely overfit the network

Using Convolutions

Natural images aren't just random arrays, they have structure

Two things to exploit while designing networks: locality and ~spatial invariance

Relating to neuroscience: weights for a given unit can be thought of as receptive field

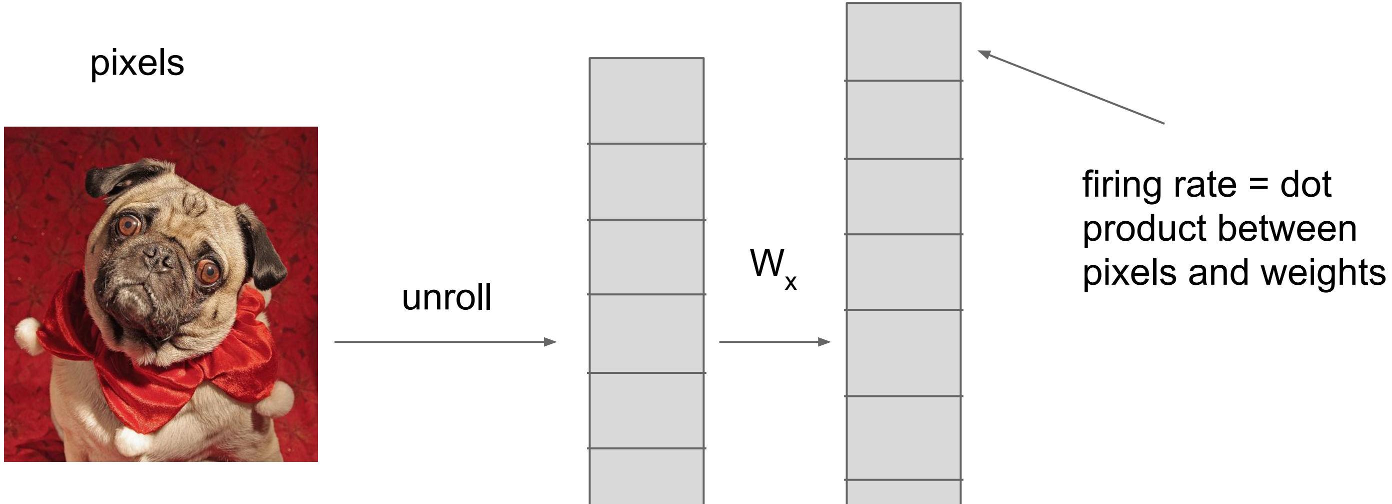

Using Convolutions

Weights as receptive fields: localized and can replicate over visual field

=> It makes sense to use convolutions

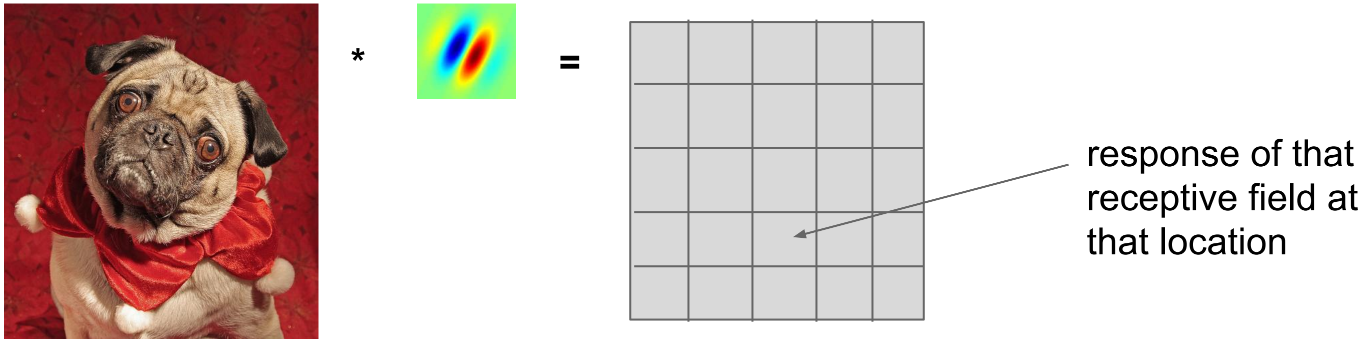

Using Convolutions

Full formulation: layers have “depth” as well

(x, y) pixel position and 3 color channels

We want a bunch of different filters to convolve the image with

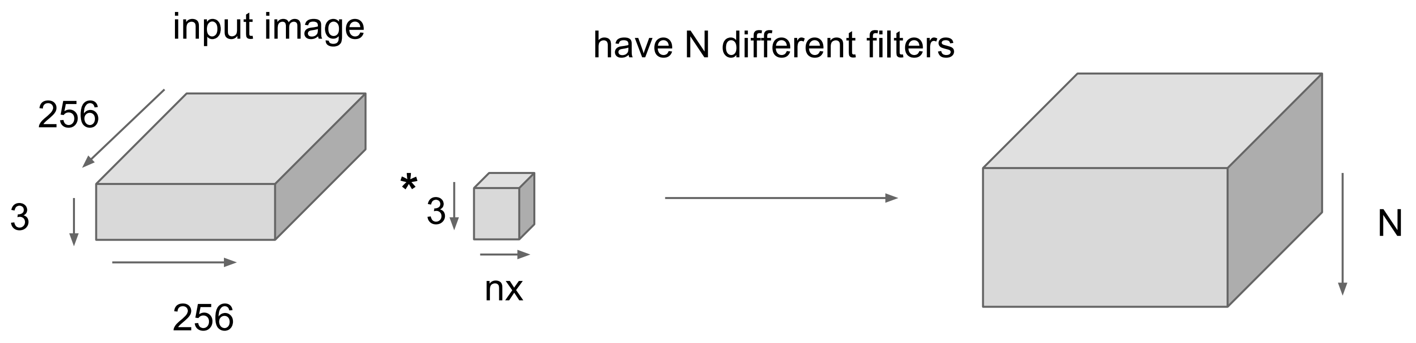

Incorporating other stuff we know is important in biology

Hierarchy: ventral stream has several layers (V1, V2,...)

Neurons are nonlinear: common non-linearity used today is rectified linear units (don’t allow neurons to have negative firing rate)

“Complex”-type cells: incorporating pooling

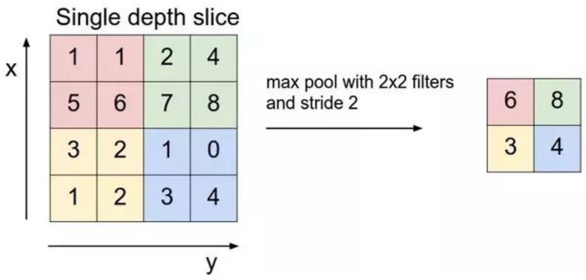

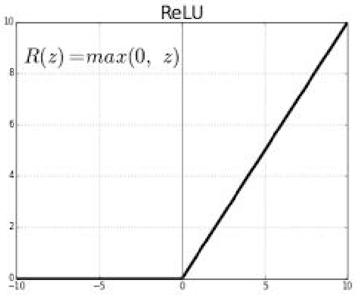

Putting it all together…

Krizhevsky et al. 2012 (Alexnet)

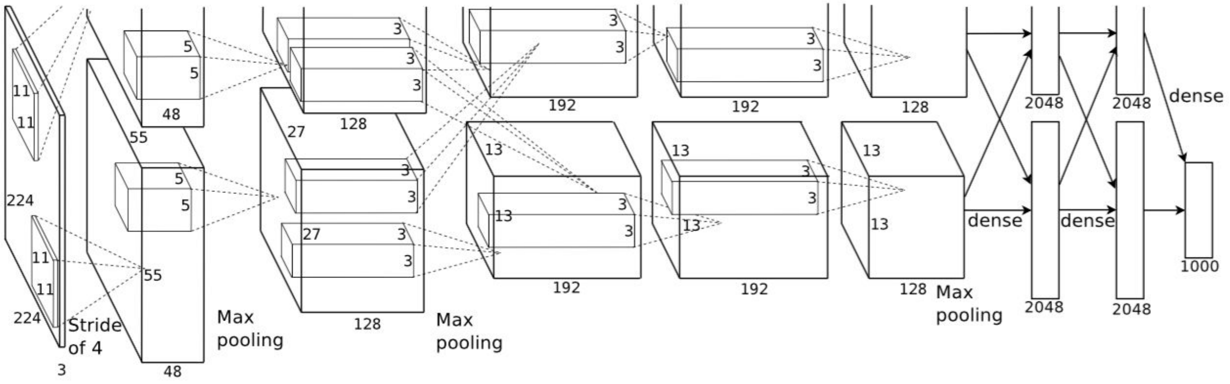

Comparing with Biology

Similarities

hierarchical

receptive fields get bigger as go higher

first layer trained weights look like V1 receptive fields

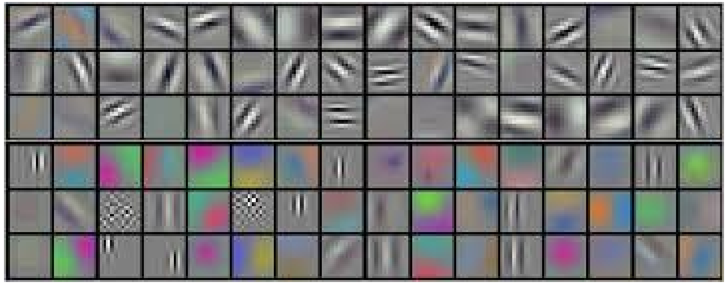

Differences

backprop?

supervised vs. unsupervised learning

Models often purely feedforward

Other Cool Stuff

Learned features from ImageNet are useful for other tasks

Can do other tasks like object localization (Oquab et al. 2015)

Often used as initialization when training on a different dataset

Other Cool Stuff

Models that are better at classification tend to be better at predicting neural responses (Yamins 2014)

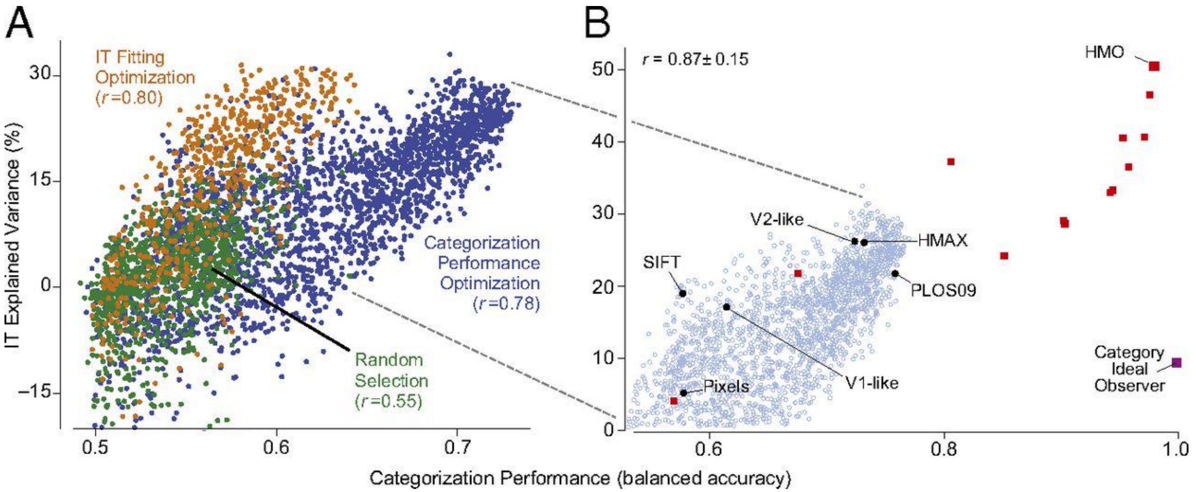

Other Cool Stuff

Nonetheless, it is easy to fool convnets (Szegedy 2013): "Adversarial Examples"

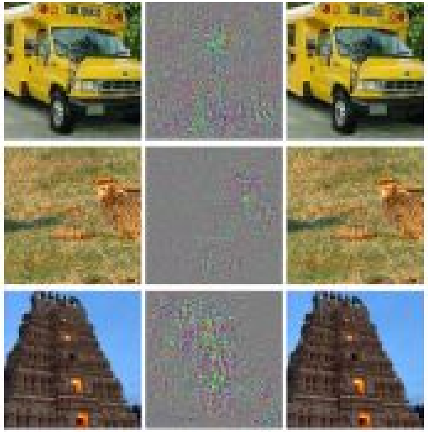

(a)

classified as ostrich

Other Cool Stuff

Computing optimal stimuli

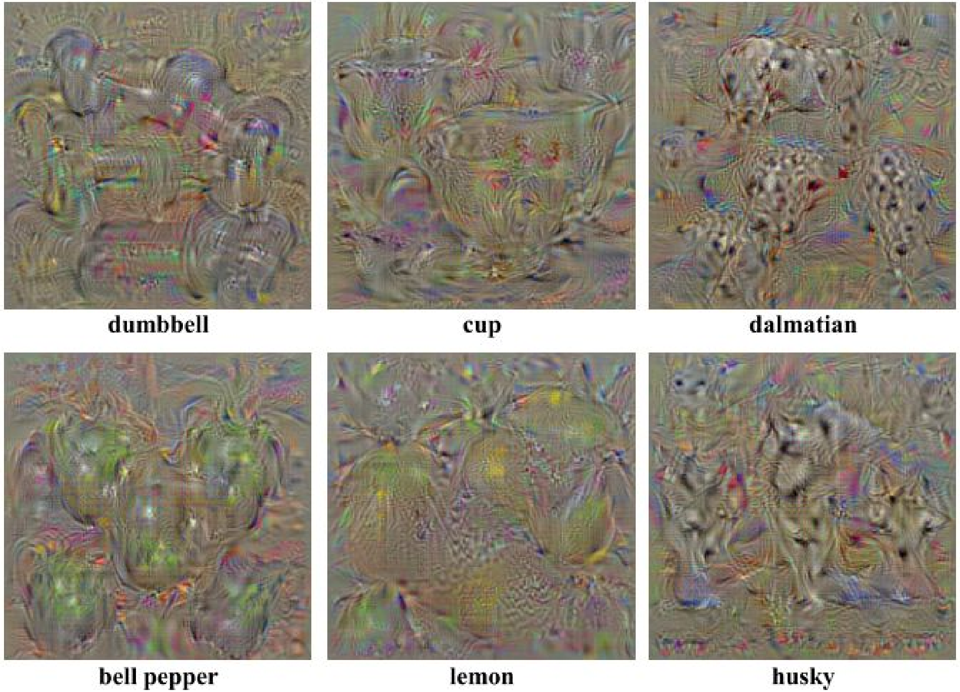

Simonyan et al. 2014

Other Cool Stuff

Generative Models, ex. Generative Adversarial Networks (GANs) (Goodfellow 2014)

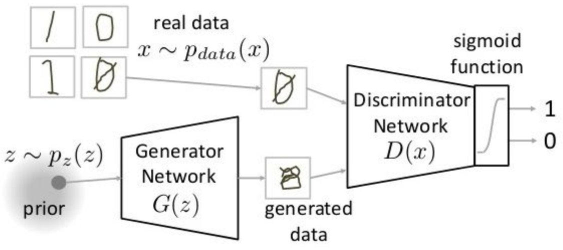

https://www.analyticsvidhya.com/blog/2017/06/introductory-generative-adversarial-networks-gans/

Radford 2015

Other Cool Stuff

Style Transfer

Gatys 2015

Final Thoughts

Still far away from making machines that can perform as well as humans, but making steady progress by designing models that share many features with brain

Neuroscience has informed computer vision, but computer vision models also allow for testing of neuroscience theories

Much easier to do “neuroscience” on models than real brains