


Convolutional Neural Nets

October 24, 2013


Convolution




- Blur operation on 512×512 image of the moon's surface
- Matrix multiplication: $512^2 \times 512^2$ matrix
... 256GBytes of memory & 128GFlops :/
- Convolution (conv2): 9×9 "kernel" (aka "filter") & 20.25 MFlops

Convolution

The Convolution Operation Sequence




Orientation Maps



Miikkulainen, 2005

Edge Filters



Sobel

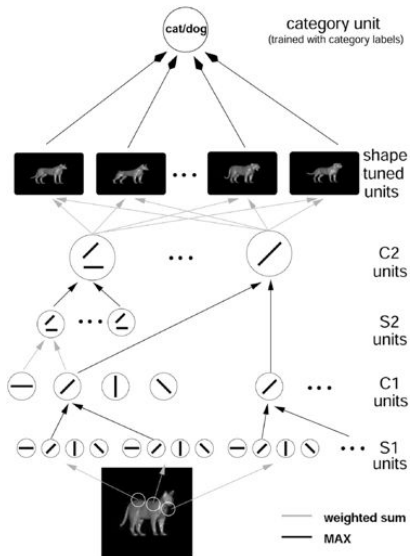
$$G_x = \begin{bmatrix} +1 & 0 & -1 \\ +2 & 0 & -2 \\ +1 & 0 & -2 \end{bmatrix} * A, \quad G_y = \begin{bmatrix} +1 & +2 & +1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix} * A$$
$$G = \sqrt{G_x^2 + G_y^2}$$

Filterbanks



Histogram of Oriented Gradients (HoG)


Biologically-inspired filters: HMAX




HoG + SVM




Person detection



Opponent Colors



Part-Based models



Different scenarios




Using

and

In

Dalal 2005


Multi-Layer Convolutional Neural Nets (MLCNNs)



“The basic design principle is to reduce the number of free parameters in the network as much as possible without overly reducing its computational power. Application of this principle increases the probability of correct generalization because it results in a specialized network architecture that has a reduced entropy”

“LeNet” (LeCun 1989)

Winning the competition



Krizhevsky & Hinton 2012 LSVRC Challenge (1K classes, 1M images):

convolution + ReLU + max-pooling + dropout + transforms

*“Trained with stochastic gradient descent on two NVIDIA **GPUs** for about a week ... 650,000 neurons, 60,000,000 parameters, 630,000,000 connections”*

Object Recognition



mite

container ship

motor scooter

leopard

	mite	container ship	motor scooter	leopard
	black widow	lifeboat	go-kart	jaguar
	cockroach	amphibian	moped	cheetah
	tick	fireboat	bumper car	snow leopard
	starfish	drilling platform	golfcart	Egyptian cat



grille

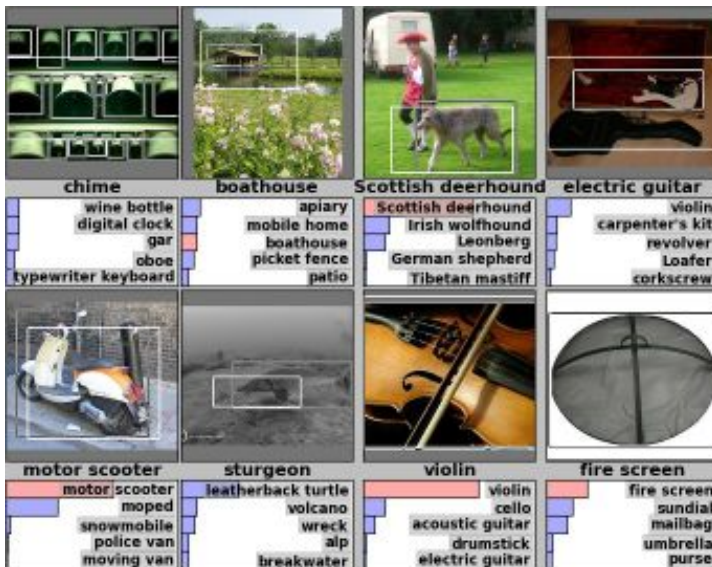
mushroom

cherry



Madagascar cat

convertible	agaric	dalmatian	squirrel monkey
grille	mushroom	grape	spider monkey
pickup	jelly fungus	elderberry	titi
beach wagon	gill fungus	ffordshire bullterrier	indri
fire engine	dead-man's-fingers	currant	howler monkey

+ Localization



Cancer Detection



Schmidhuber, MICCAI 2013 Grand Challenge on Mitosis Detection

Code

Matlab: <https://github.com/rasmusbergpalm/DeepLearnToolbox>

Python + CUDA: <http://deeplearning.net/software/theano/>

C++ CUDA: <http://code.google.com/p/cuda-convnet/>

Lua: <http://www.torch.ch>