

Want to get started with deep learning?







Prototyping a deep learning image classifier

Thomas Ellebæk, Ferring Pharmaceuticals,
ML02 PhUSE EU Connect 2018,
November 5th 2018

Pattern detection

FERRING

PHARMACEUTICALS



Classification



{Alligator, Beaver, Cat, Dog, ..., Zebra}

Classification




{Alligator, Beaver, Cat, Dog, ..., Zebra}
0.01 0.00 0.87 0.07 ... 0.00

Deep learning applications

Teradata report based on survey conducted July 2017:

”80% report that some form of AI is already in production in their organization”

(EB9867_State_of_Artificial_Intelligence_for_the_Enterprises.pdf)




Deep learning applications

Teradata report based on survey conducted July 2017:


”80% report that some form of AI is already in production in their organization”

(EB9867_State_of_Artificial_Intelligence_for_the_Enterprises.pdf)



Growing Use of Deep Learning at Google

of directories containing model description files



Across many products/areas:

Android
Apps
drug discovery
Gmail
Image understanding
Maps
Natural language understanding
Photos
Robotics research
Speech
Translation
YouTube
... many others ...



Deep learning breakthrough

ImageNet Classification with Deep Convolutional Neural Networks

Alex Krizhevsky
University of Toronto
kriz@cs.utoronto.ca

Ilya Sutskever
University of Toronto
ilya@cs.utoronto.ca


Geoffrey E. Hinton
University of Toronto
hinton@cs.utoronto.ca

Abstract

We trained a large, deep convolutional neural network to classify the 1.2 million high-resolution images in the ImageNet LSVRC-2010 contest into the 1000 different classes. On the test data, we achieved top-1 and top-5 error rates of 37.5% and 17.0% which is considerably better than the previous state-of-the-art. The neural network, which has 60 million parameters and 650,000 neurons, consists of five convolutional layers, some of which are followed by max-pooling layers, and three fully-connected layers with a final 1000-way softmax. To make training faster, we used non-saturating neurons and a very efficient GPU implementation of the convolution operation. To reduce overfitting in the fully-connected layers we employed a recently-developed regularization method called “dropout” that proved to be very effective. We also entered a variant of this model in the ILSVRC-2012 competition and achieved a winning top-5 test error rate of 15.3%, compared to 26.2% achieved by the second-best entry.

that proved to be very effective. We also entered a variant of this model in the ILSVRC-2012 competition and achieved a winning top-5 test error rate of 15.3%, compared to 26.2% achieved by the second-best entry.

ImageNet results (top-5)



Feedforward Neural Network

$$\begin{matrix} z_1^{[0]} \\ = x_1 \end{matrix}$$

$$\begin{matrix} z_2^{[0]} \\ = x_2 \end{matrix}$$

Input layer

$$\begin{matrix} z_1^{[3]} \\ = \hat{y}_1 \end{matrix}$$

$$\begin{matrix} z_2^{[3]} \\ = \hat{y}_2 \end{matrix}$$

Output layer

Feedforward Neural Network

$z_1^{[0]} = x_1$ Shoe size

$z_2^{[0]} = x_2$ Hair length


Input layer

Male $z_1^{[3]} = \hat{y}_1$

Female $z_2^{[3]} = \hat{y}_2$

Output layer

Feedforward Neural Network



Input layer

Hidden layer 1

Output layer

Shoe size and
hair length

Gender


Neuron output:

$$z_j^{[l]} = g \left(\sum_{k=0}^{d^{[l-1]}} w_{kj}^{[l-1]} z_k^{[l-1]} \right)$$

Activation function:
(rectified linear unit – ReLU)

$$g(a) = \max(0, a)$$

Feedforward Neural Network



Input layer

Hidden layer 1

Hidden layer 2

Output layer

Shoe size and
hair length

Gender


Neuron output:

$$z_j^{[l]} = g \left(\sum_{k=0}^{d^{[l-1]}} w_{kj}^{[l-1]} z_k^{[l-1]} \right)$$

Activation function:
(rectified linear unit – ReLU)

$$g(a) = \max(0, a)$$

Feedforward Neural Network



Input layer

Hidden layer 1

Hidden layer 2

Output layer

Shoe size and
hair length

Gender


Neuron output:

$$z_j^{[l]} = g \left(\sum_{k=0}^{d^{[l-1]}} w_{kj}^{[l-1]} z_k^{[l-1]} \right)$$

Activation function:
(rectified linear unit – ReLU)

$$g(a) = \max(0, a)$$

Feedforward Neural Network



Input layer

Hidden layer 1

Hidden layer 2

Output layer

Shoe size and
hair length

Gender

Neuron output:

$$z_j^{[l]} = g \left(\sum_{k=0}^{d^{[l-1]}} w_{kj}^{[l-1]} z_k^{[l-1]} \right)$$


Activation function:
(rectified linear unit – ReLU)

$$g(a) = \max(0, a)$$


Loss function:
(categorical cross entropy)

$$L = - \sum_{i=1}^N \sum_{k=1}^{d^{[out]}} y_k^{(i)} \log \left(f(x^{(i)})_k \right)$$

Convolutional Neural Network



Convolution



Kernel

w_1	w_2
w_3	w_4


Input

a	b	c	d
e	f	g	h
i	j	k	l
m	n	o	p

2D convolution

$aw_4 + bw_3 + ew_2 + fw_1$	...	...
...	...	...
...	...	...
...	...	...

Convolution



FERRING
PHARMACEUTICALS

Kernel

w_1	w_2
w_3	w_4

Input

a	b	c	d
e	f	g	h
i	j	k	l
m	n	o	p


2D convolution

$aw_4 + bw_3 + ew_2 + fw_1$	...	...
...	...	...
...	...	...
...	...	...

2D cross-correlation

$aw_1 + bw_2 + ew_3 + fw_4$	...	...
...	...	...
...	...	...
...	...	...

Max pooling



Input


1	4	9	16
2	3	8	15
5	6	7	14
10	11	12	13

(2,2) max pooling


4	16
11	14

CASE: Cell Counter


Dataset	Split	Examples
Training	70%	2246
Validation	15%	440
Test	15%	438



true label: 1




true label: 2




true label: 3



true label: 4



true label: 5





Computing environment



Machine	Standard epoch	Datagen epoch
Intel® Core™ i7-6500U CPU @ 2.50GHz, 4 logical processors, 16 GB DDR3 RAM	285 s	840 s
Azure Standard_NC6, 6 CPU, 1 GPU (one-half NVIDIA Tesla K80 card), 56 GiB RAM	7 s	8 s


Fitting AlexNet-type model



Validation accuracy: 35.7%

		Predicted labels				
		1	2	3	4	>=5
Actual labels	1	0	0	0	79	0
	2	0	0	0	96	0
	3	0	0	0	56	0
	4	0	0	0	157	0
	>=5	0	0	0	52	0

Final model




3 convolutional layers and
2 dense fully connected
layers



Final model

FERRING

PHARMACEUTICALS



3 convolutional layers and
2 dense fully connected
layers



Error analysis


Test accuracy: 55.9% >> 35.7%

		Predicted labels				
		1	2	3	4	>=5
Actual labels	1	79	18	2	2	0
	2	8	76	17	19	0
	3	1	20	12	20	5
	4	0	18	19	67	18
	>=5	0	3	5	18	11


		Predicted (%)				
		1	2	3	4	>=5
Actuals (%)	1	0,78	0,18	0,02	0,02	0,00
	2	0,07	0,63	0,14	0,16	0,00
	3	0,02	0,34	0,21	0,34	0,09
	4	0,00	0,15	0,16	0,55	0,15
	>=5	0,00	0,08	0,14	0,49	0,30

2 correct predicted examples

pred: 2 (78.4%), true label: 2




pred: 4 (97.6%), true label: 4



5 wrongly predicted examples


pred: 3 (51.1%), true label: 2



pred: 4 (77.0%), true label: 5



pred: 1 (87.8%), true label: 2



pred: 3 (93.0%), true label: 4





pred: 2 (65.2%), true label: 4



Live demo!



Takeaways and suggested learning



Andrew Ng

Co-founder, Coursera; Adjunct Professor, Stanford University;
formerly head of Baidu AI Group/Google Brain

[Stanford University](https://stanford.edu/)
deeplearning.ai





<https://machinelearningmastery.com/>

Welcome to Machine Learning Mastery!

Hi, I'm Jason Brownlee Ph.D. and I help developers like you skip years ahead.

Takeaways and suggested learning

- ▶ Technology is ready
- ▶ Data is the most important asset



Andrew Ng

Co-founder, Coursera; Adjunct Professor, Stanford University;
formerly head of Baidu AI Group/Google Brain

[Stanford University](https://stanford.edu/)
deeplearning.ai





<https://machinelearningmastery.com/>

Welcome to Machine Learning Mastery!

Hi, I'm Jason Brownlee Ph.D. and I help developers like you skip years ahead.

Takeaways and suggested learning

- ▶ Technology is ready
- ▶ Data is the most important asset



Questions?



Andrew Ng

Co-founder, Coursera; Adjunct Professor, Stanford University;
formerly head of Baidu AI Group/Google Brain

Stanford University
deeplearning.ai



<https://machinelearningmastery.com/>

Welcome to Machine Learning Mastery!

Hi, I'm Jason Brownlee Ph.D. and I help developers like you skip years ahead.