Deep neural networks

On-line Resources

- http://neuralnetworksanddeeplearning.com/index.html Online book by Michael Nielsen
- http://matlabtricks.com/post-5/3x3-convolution-kernels-with-online-demo - of convolutions
- https://cs.stanford.edu/people/karpathy/convnetjs/demo/mnist.html - demo of CNN
- http://scs.ryerson.ca/~aharley/vis/conv/ - 3D visualization
- http://cs231n.github.io/ Stanford CS class CS231n: Convolutional Neural Networks for Visual Recognition.

- http://www.deeplearningbook.org/ MIT Press book in prep from Bengio

A history of neural networks

- 1940s-60's:
  - McCulloch & Pitts; Hebb: modeling real neurons
  - Rosenblatt, Widrow-Hoff: : perceptrons
  - 1969: Minskey & Papert, *Perceptrons* book showed formal limitations of one-layer linear network
- 1970's-mid-1980's: …
- mid-1980's – mid-1990's:
  - backprop and multi-layer networks
  - Rumelhart and McClelland *PDP* book set
  - Sejnowski's NETTalk, BP-based text-to-speech
  - Neural Info Processing Systems (NIPS) conference starts
- Mid 1990's-early 2000's: …
- Mid-2000's to current:
  - More and more interest and experimental success

TECH INSIDER INNOVATION

'Mach
the in

Drake B
Apr. 1,

FACEBO

Don't worry,

We're in the

It used to b

SECTIONS

Facebook
and Intrus
Paying Off

TECHNOLOGY

The Economist

World politics | Business & finance | Economics | Science & technology | Culture

Artificial intelligence

Million-dollar babies

As Silicon Valley fights for talent, universities struggle to hold on to their stars

Apr 2nd 2016 | SAN FRANCISCO | From the print edition

Timekeeper | Like 6.8K | Tweet

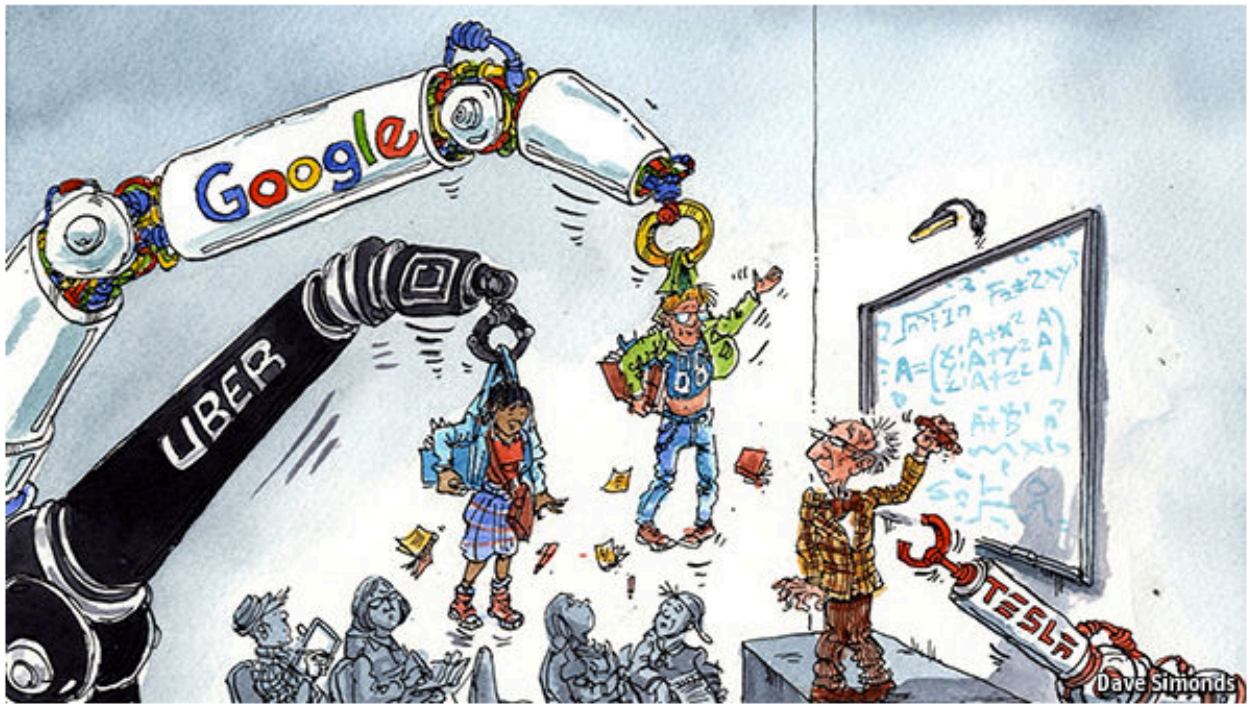

CIENCE LIFE T-LOUNGE TE

odfather Of
aid Of AI

SUBSCRIBE

program AlphaGo
world Go champion Lee

ence experts, who
er program would need at
ough to be able to beat a

uter program is that
said that AlphaGo could
h with Sedol was able to

st skilled humans in
d Garry Kasparov two
lerts

Talks

Silicon Valley Looks to Artificial Intelligence for the Next Big Thing

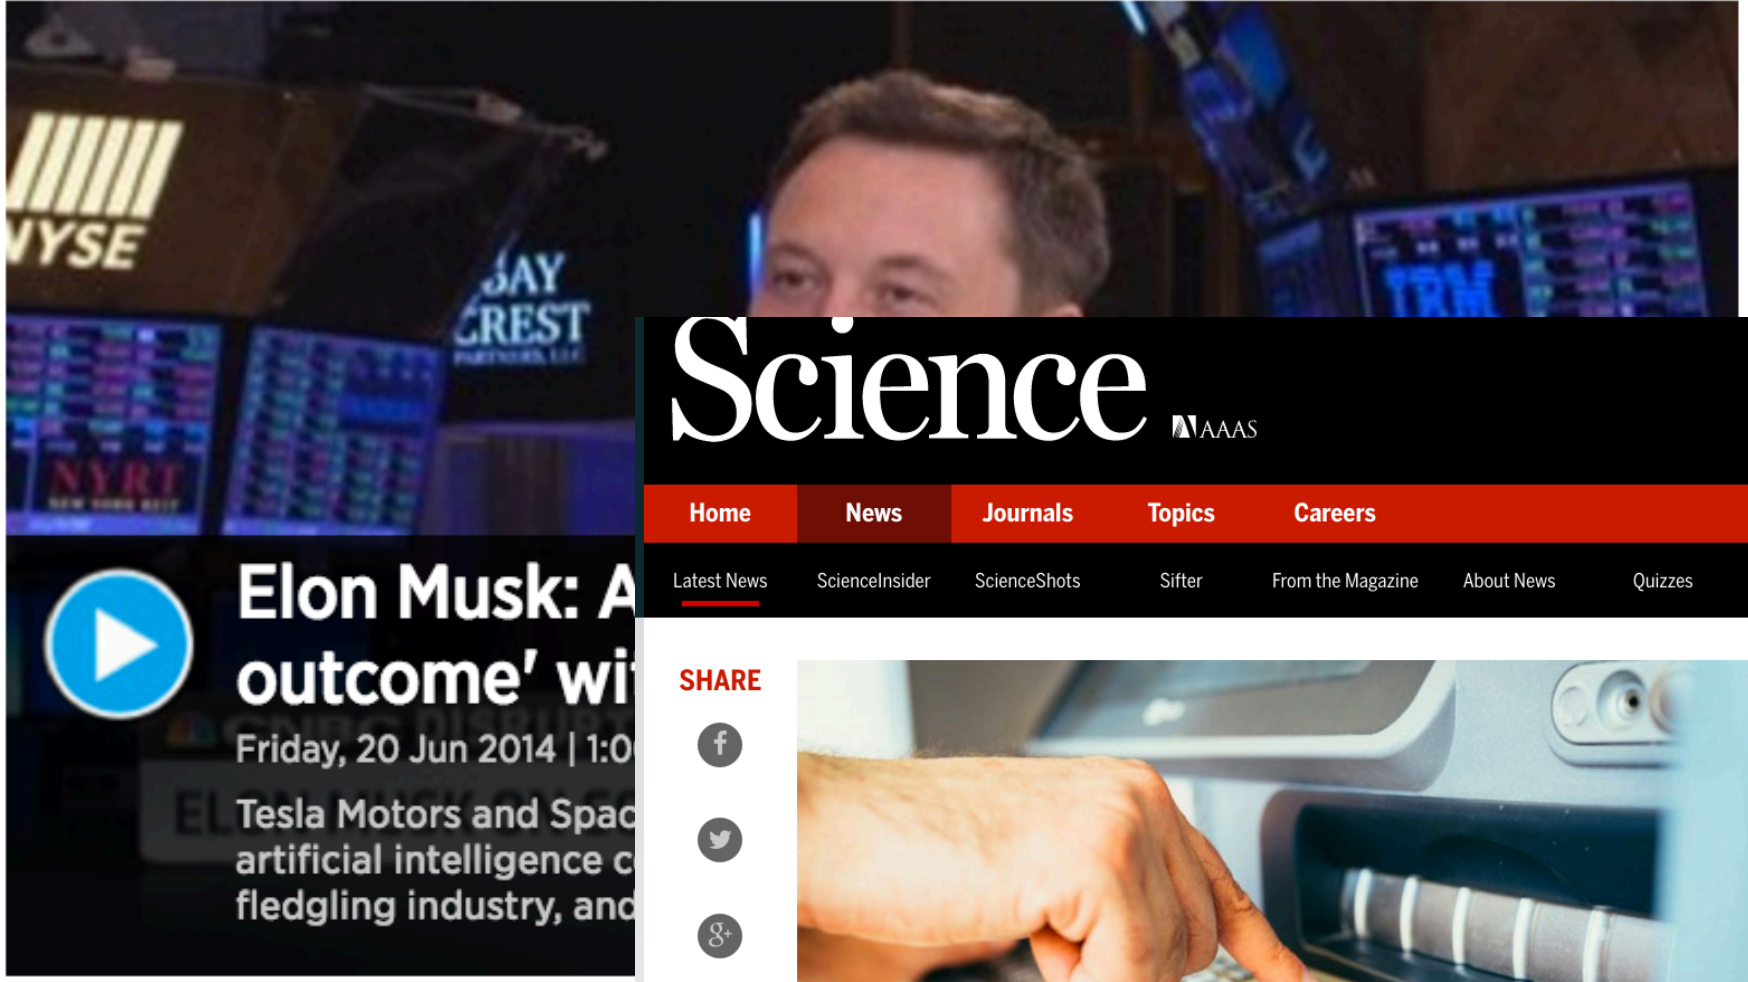

Science AAAS

SHARE

© FilippoBacci/iStockphoto

Artificial intelligence steals money from banking customers

By Adrian Cho | Apr. 1, 2016 , 3:00 AM

ANNs in the 90's

- Mostly 2-layer networks or else carefully constructed "deep" networks
- Worked well but training was slow and finicky

Nov 1998 – Yann LeCunn, Bottou, Bengio, Haffner

PROC. OF THE IEEE, NOVEMBER 1998 7

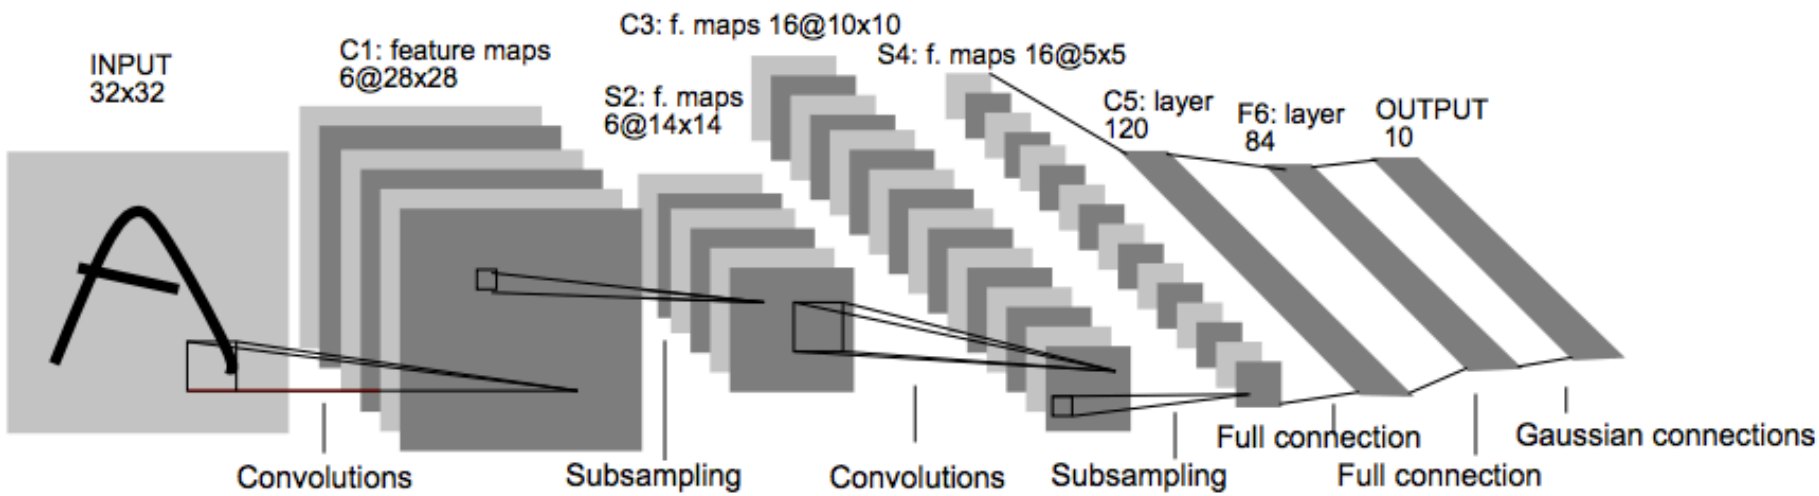

Fig. 2. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

ANNs in the 90's

- Mostly 2-layer networks or else carefully constructed "deep" networks
- Worked well but training typically took weeks when guided by an expert

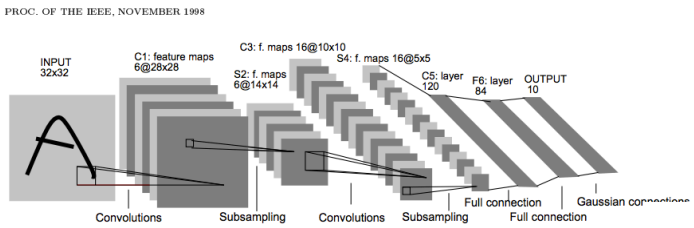

Fig. 2. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a s whose weights are constrained to be identical.

SVM: 98.9-99.2% accurate

CNNs: 98.3-99.3% accurate

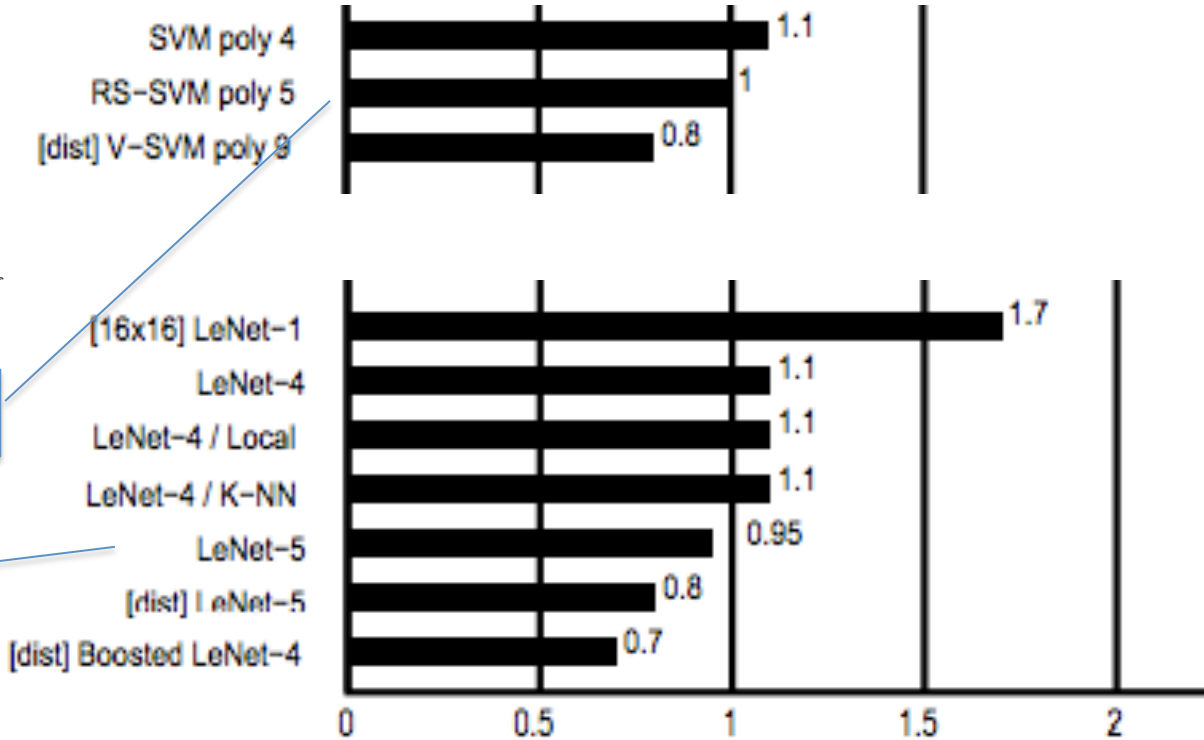

A history of neural networks

- 1940s-60's:
  - McCulloch & Pitts; Hebb: modeling real neurons
  - Rosenblatt, Widrow-Hoff: : perceptrons
  - 1969: Minskey & Papert, *Perceptrons* book showed formal limitations of one-layer linear network
- 1970's-mid-1980's: …
- mid-1980's – mid-1990's:
  - backprop and multi-layer networks
  - Rumelhart and McClelland *PDP* book set
  - Sejnowski's NETTalk, BP-based text-to-speech
  - Neural Info Processing Systems (NIPS) conference starts
- Mid 1990's-early 2000's: …
- **Mid-2000's to current: The next two lectures**

Outline of lectures

- What's new in ANNs in the last 5-10 years?
- What types of ANNs are most successful and why?
- What are the hot research topics for deep learning?

Outline

- What's new in ANNs in the last 5-10 years?
  - **Deeper networks**, more data, and faster training
    - **Scalability and use of GPUs**
    - Symbolic differentiation
    - Some subtle changes to cost function, architectures, optimization methods
- What types of ANNs are most successful and why?
  - Convolutional networks (CNNs)
  - Long term/short term memory networks (LSTM)
  - Word2vec and embeddings
- What are the hot research topics for deep learning?

PARALLEL TRAINING FOR ANNS

Recap: logistic regression with SGD

$$P(Y=1 \mid X=\mathbf{x}) = p = \frac{1}{1+e^{-\mathbf{x}\cdot\mathbf{w}}}$$

$$\mathbf{w}^{(t+1)} = \mathbf{w}^{(t)} + \lambda(y-p)\mathbf{x}$$

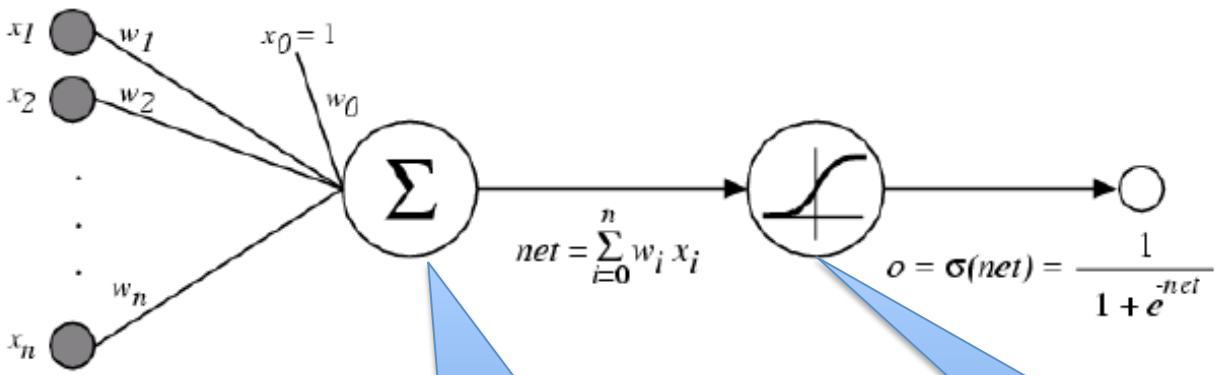

This part computes inner product **<x,w>**

This part logistic of **<x,w>**

Recap: logistic regression with SGD

$$P(Y=1 \mid X=\mathbf{x}) = p = \frac{1}{1+e^{-\mathbf{x}\cdot\mathbf{w}}}$$

$$\mathbf{w}^{(t+1)} = \mathbf{w}^{(t)} + \lambda(y-p)\mathbf{x}$$

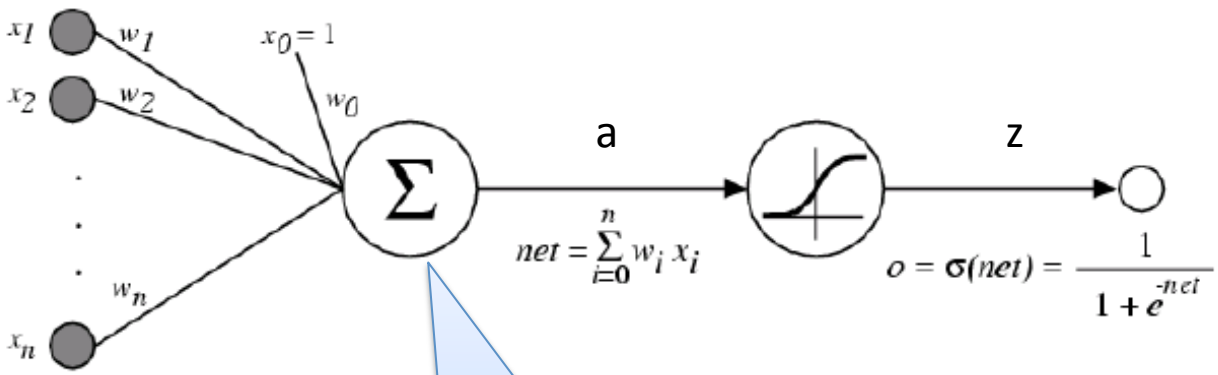

On one example: computes inner product **<x,w>**

$$\sum_j x^j w^j$$

There's some chance to compute this in parallel…can we do more?

In ANNs we have many many logistic regression nodes

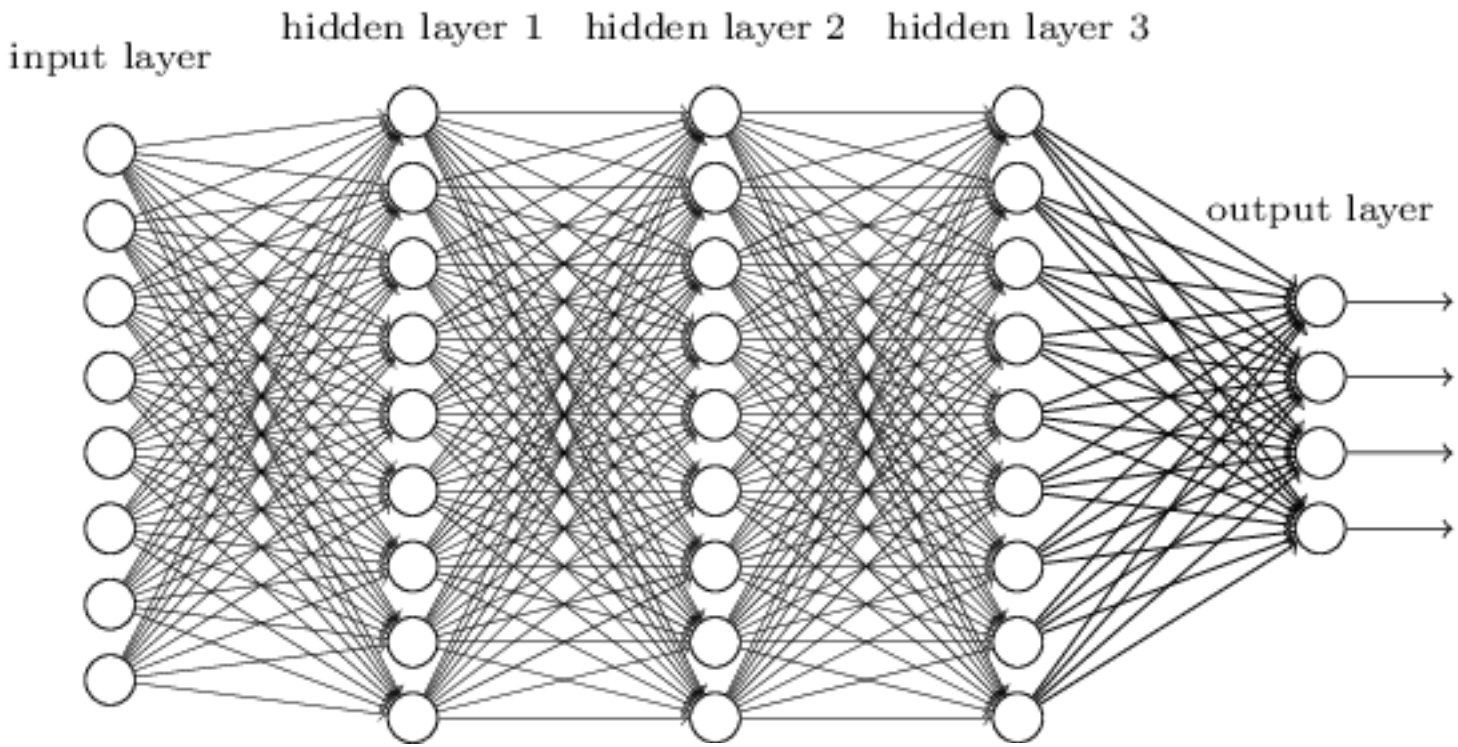

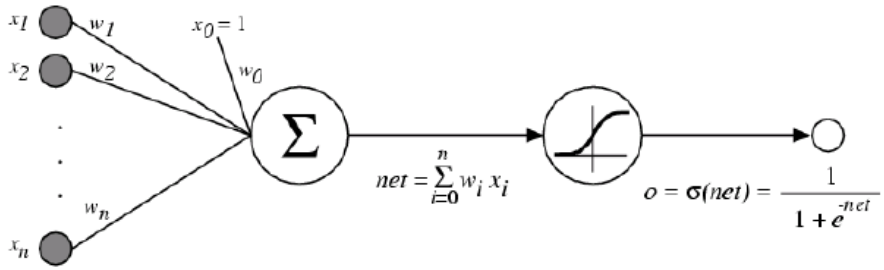

Recap: logistic regression with SGD

Let **x** be an example
Let $\mathbf{w}_i$ be the input weights for the i-th hidden unit
Then output $\mathbf{a}_i = \mathbf{x} \cdot \mathbf{w}_i$

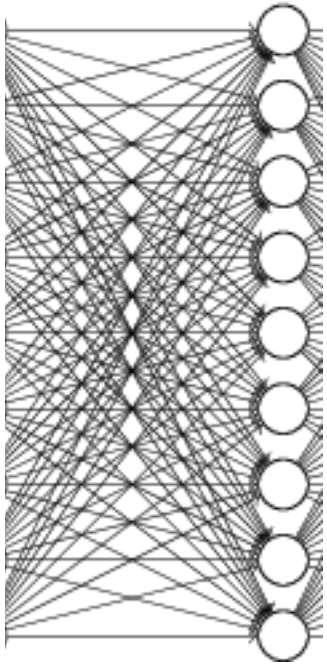

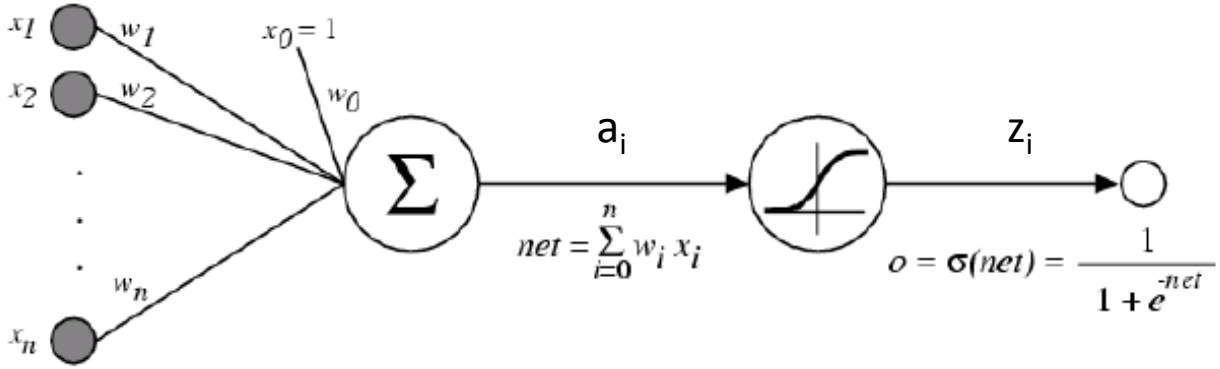

Recap: logistic regression with SGD

Let **x** be an example

Let $\mathbf{w}_i$ be the input weights for the i-th hidden unit

Then **a** = **x** W

is output for all *m* units

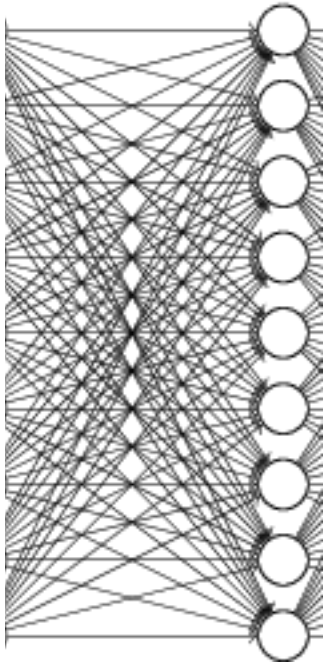

W =

w_1	w_2	w_3	…	w_m
0.1	-0.3	…		
-1.7	…			
..				
…				

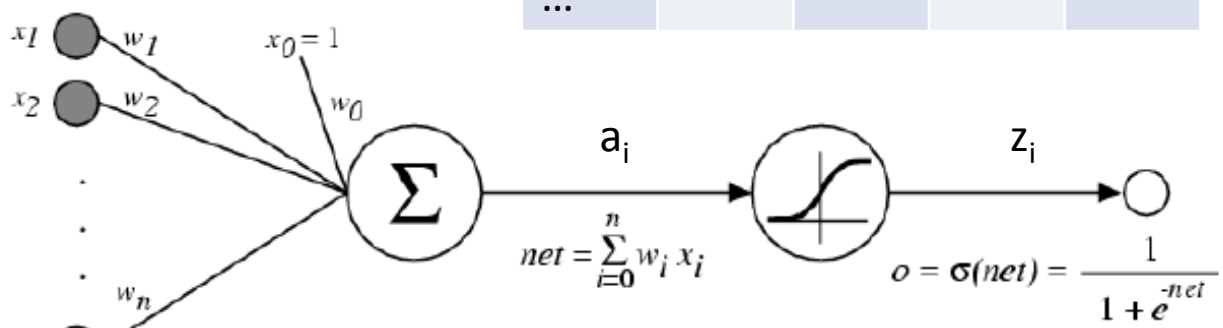

Recap: logistic regression with SGD

Let X be a matrix with *k* examples
Let $\mathbf{w}_i$ be the input weights for the i-th hidden unit
Then A = X W is output for all *m* units
for all *k* examples

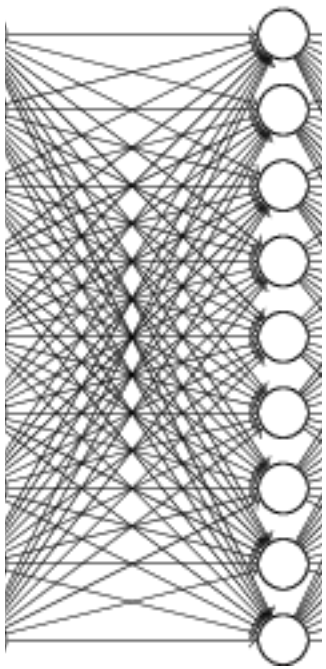

x_1	1	0	1	1
x_2	…			
…				
x_k				

w_1	w_2	w_3	…	w_m
0.1	-0.3	…		
-1.7	…			
0.3	…			
1.2				

There's a lot of chances to do this in parallel

XW =

$x_1.w_1$	$x_1.w_2$	…	$x_1.w_m$
$x_k.w_1$	…	…	$x_k.w_m$

ANNs and multicore CPUs

- Modern libraries (Matlab, numpy, …) do matrix operations fast, in parallel
- Many ANN implementations exploit this parallelism automatically
- Key implementation issue is working with matrices comfortably

ANNs and GPUs

- GPUs do matrix operations very fast, in parallel
  - For dense matrixes, not sparse ones!
- Training ANNs on GPUs is common
  - SGD and minibatch sizes of 128
- Modern ANN implementations can exploit this
- GPUs are not super-expensive
  - $500 for high-end one
  - large models with $O(10^7)$ parameters can fit in a large-memory GPU (12Gb)
- Speedups of 20x-50x have been reported

ANNs and multi-GPU systems

- There are ways to set up ANN computations so that they are spread across multiple GPUs
  - Sometimes needed for very large networks
  - Not especially easy to implement and do with most current tools

Outline

- What's new in ANNs in the last 5-10 years?
  - **Deeper networks**, more data, and faster training
    - Scalability and use of GPUs
    - Symbolic differentiation
    - Some subtle changes to cost function, architectures, optimization methods
      - But first – why?
- What types of ANNs are most successful and why?
  - Convolutional networks (CNNs)
  - Long term/short term memory networks (LSTM)
  - Word2vec and embeddings
- What are the hot research topics for deep learning?

WHY ARE DEEPER NETWORKS USEFUL?

Recap: multilayer networks

- Classifier is a multilayer *network* of *logistic units*
- Each *unit* takes some inputs and produces one output using a logistic classifier
- Output of one unit can be the input of another

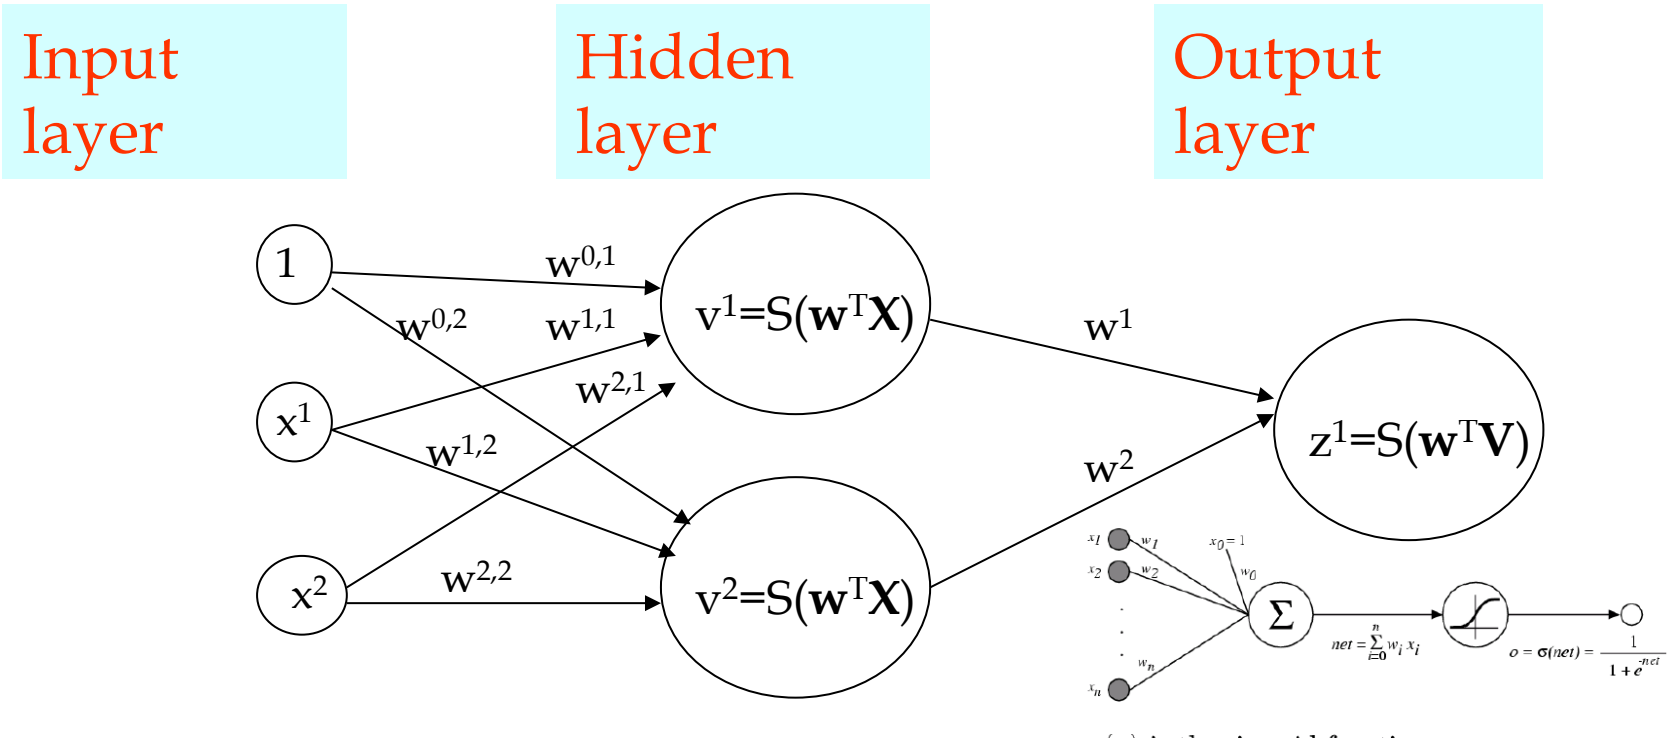

Recap: Learning a multilayer network

- Define a loss which is squared error
  - But over a network of logistic units
- Minimize loss with gradient descent

$$J_{\mathbf{X},\mathbf{y}}(\mathbf{w}) = \sum_i \left(y^i - \hat{y}^i \right)^2$$

  - But output is network output

Recap: weight updates for multilayer ANN

For nodes k in output layer:

$$\delta_k \equiv \left(t_k - a_k\right)\ a_k\left(1 - a_k\right)$$

For nodes j in hidden layer:

$$\delta_j \equiv \sum_k \left(\delta_k w_{kj}\right)\ a_j\left(1 - a_j\right)$$

For all weights:

$$w_{kj} = w_{kj} - \varepsilon\ \delta_k a_j$$

$$w_{ji} = w_{ji} - \varepsilon\ \delta_j a_i$$

"Propagate errors backward"
BACKPROP

Can carry this recursion out further if you have multiple hidden layers

ANNs are expressive

- One logistic unit can implement and AND or an OR of a subset of inputs
  - e.g., (x_3 AND x_5 AND … AND x_{19})
- Every boolean function can be expressed as an OR of ANDs
  - e.g., (x_3 AND x_5) OR (x_7 AND x_{19}) OR …
- So one hidden layer can express any BF

(But it might need lots and lots of hidden units)

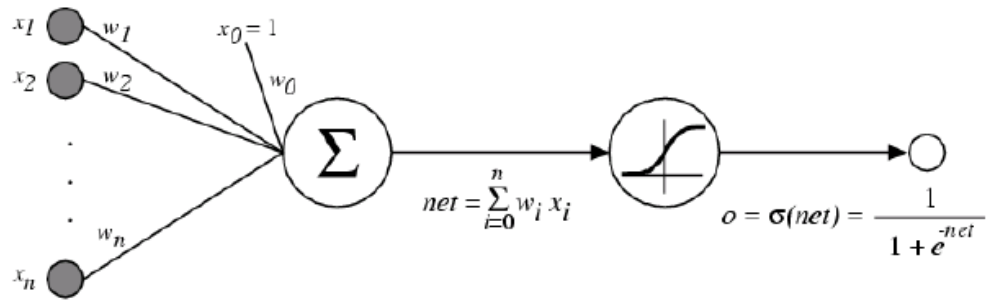

ANNs are expressive

- One logistic unit can implement and AND or an OR of a subset of inputs
  - e.g., (x_3 AND x_5 AND … AND x_{19})
- Every boolean function can be expressed as an OR of ANDs
  - e.g., (x_3 AND x_5) OR (x_7 AND x_{19}) OR …
- So one hidden layer can express any BF

- Example: parity(x1,…,xN)=1 iff off number of xi's are set to one

Parity(a,b,c,d) =
(a & -b & -c & -d) OR (-a & b & -c & -d) OR … #list all the "1s"
(a & b & c & -d) OR (a & b & -c & d) OR … #list all the "3s"

Size in general is $O(2^N)$

Deeper ANNs are more expressive

- A two-layer network needs $O(2^N)$ units
- A two-layer network can express binary XOR
- A 2*logN layer network can express the parity of N inputs (even/odd number of 1's)
  - With O(logN) units in a binary tree
- Deep network + parameter tying ~= subroutines

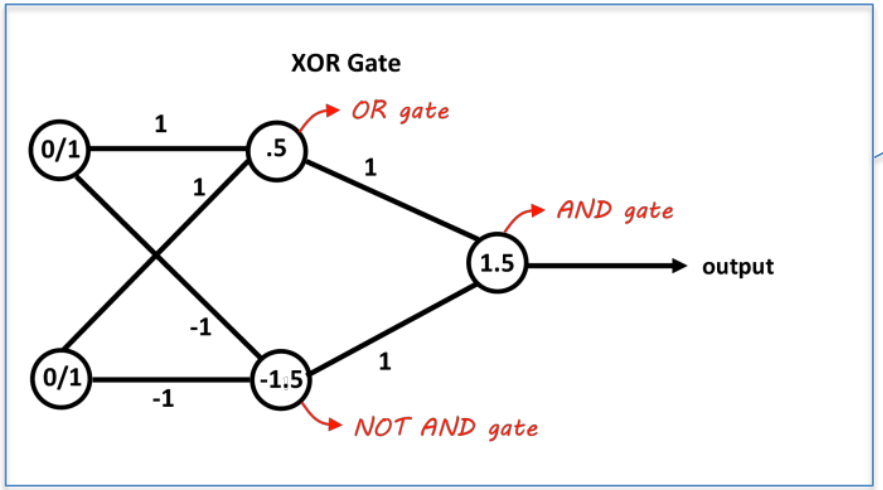

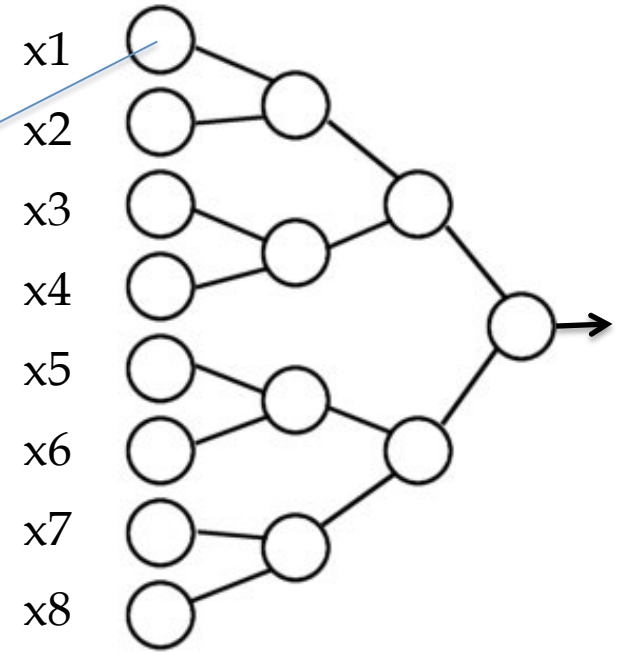

http://neuralnetworksanddeeplearning.com/chap1.html

Hypothetical code for face recognition

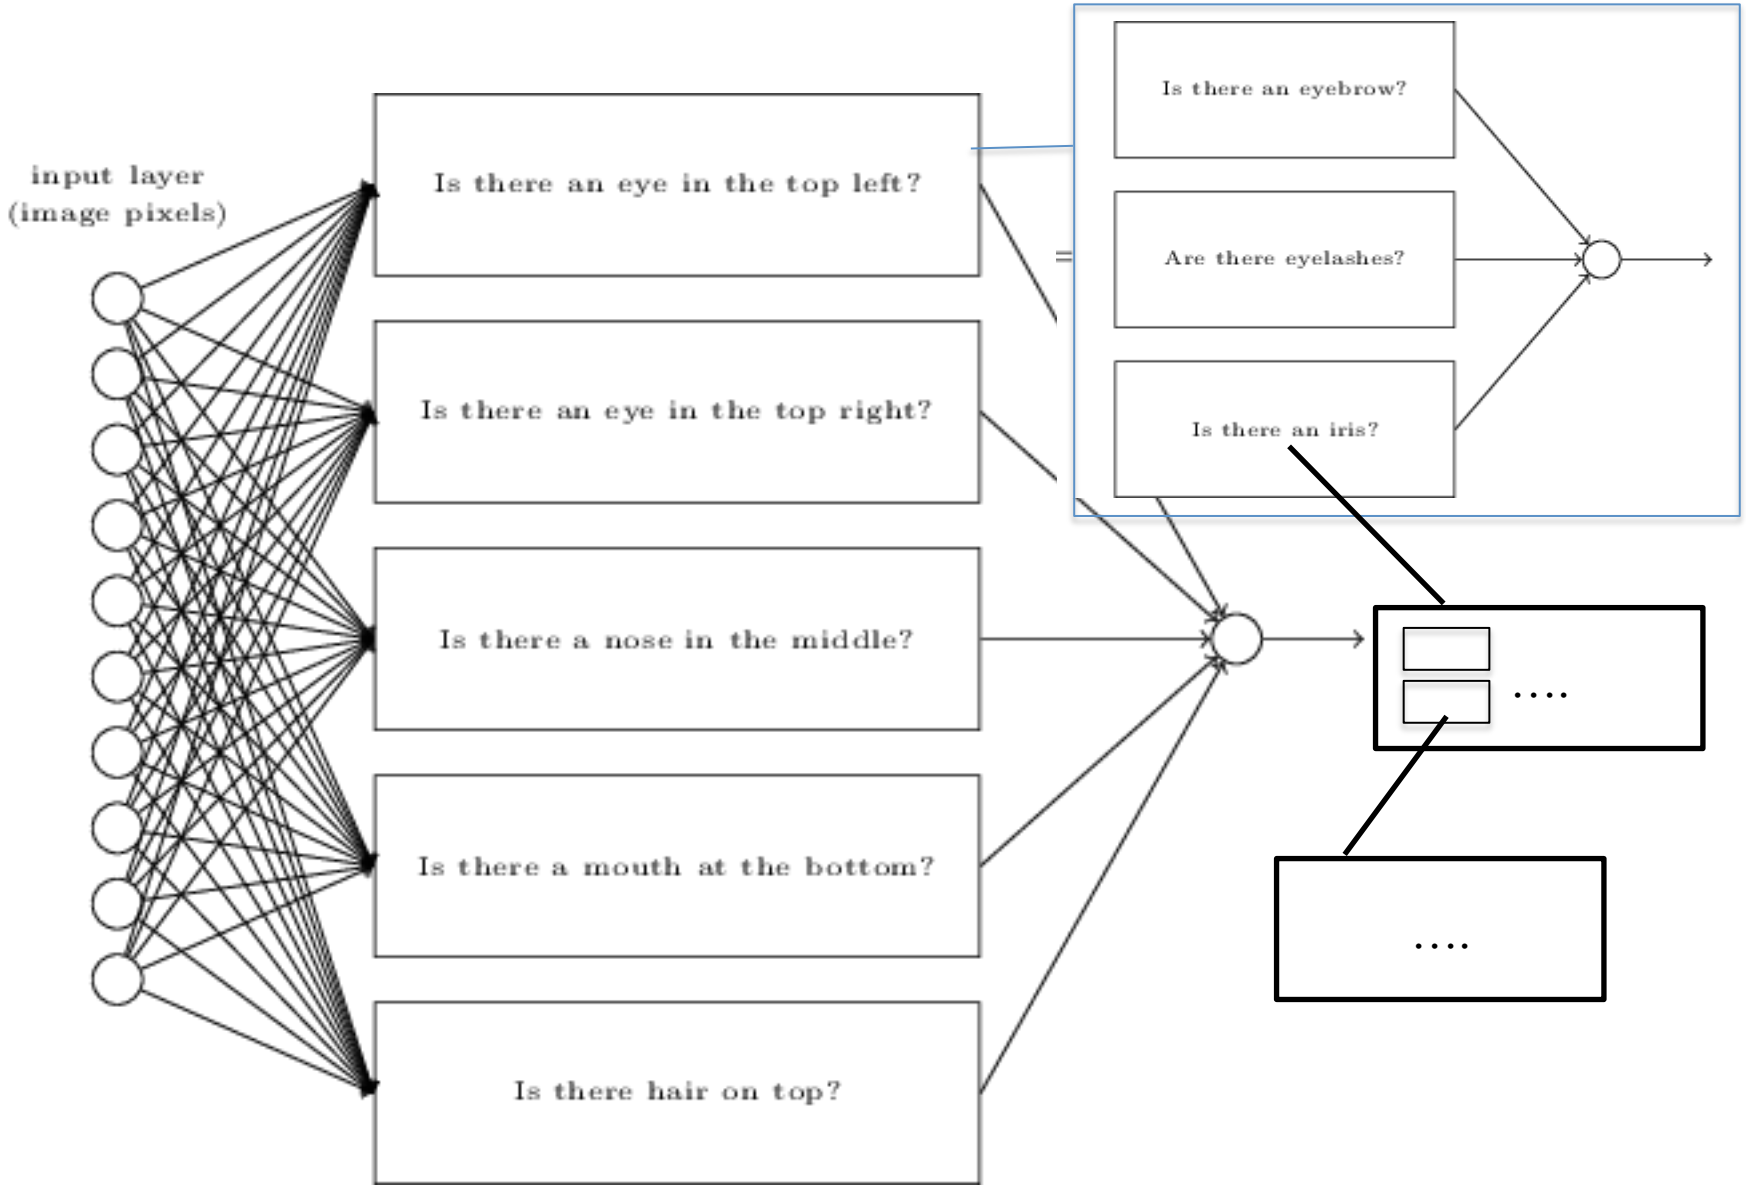

WHY ARE DEEP NETWORKS HARD TO TRAIN?

Recap: weight updates for multilayer ANN

For nodes k in output layer L:

$$\delta_k^L \equiv \left(t_k - a_k\right)\ a_k\left(1 - a_k\right)$$

For nodes j in hidden layer h:

$$\delta_j^h \equiv \sum_k \left(\delta_j^{h+1}\ w_{kj}\right)\ a_j\left(1 - a_j\right)$$

What happens as the layers get further and further from the output layer? E.g., what's gradient for the bias term with several layers after it?

$$\frac{\partial C}{\partial b_1} = \sigma'(z_1) \times w_2 \times \sigma'(z_2) \times w_3 \times \sigma'(z_3) \times w_4 \times \sigma'(z_4) \times \frac{\partial C}{\partial a_4}$$

Gradients are unstable

Max at 1/4

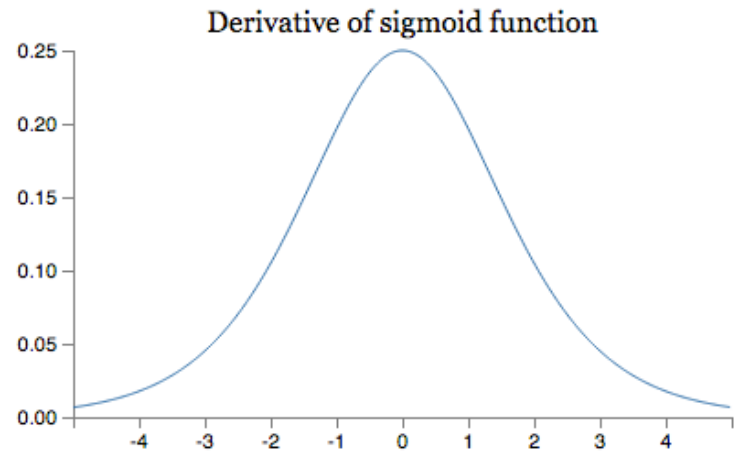

If weights are usually < 1 then we are multiplying by many numbers < 1 so the gradients get very small.

The **vanishing gradient** problem

What happens as the layers get further and further from the output layer? E.g., what's gradient for the bias term with several layers after it in a trivial net?

$$\frac{\partial C}{\partial b_1} = \sigma'(z_1) \times w_2 \times \sigma'(z_2) \times w_3 \times \sigma'(z_3) \times w_4 \times \sigma'(z_4) \times \frac{\partial C}{\partial a_4}$$

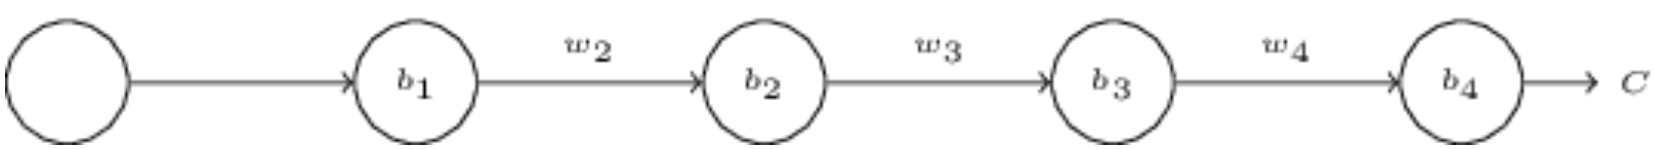

Gradients are unstable

Max at 1/4

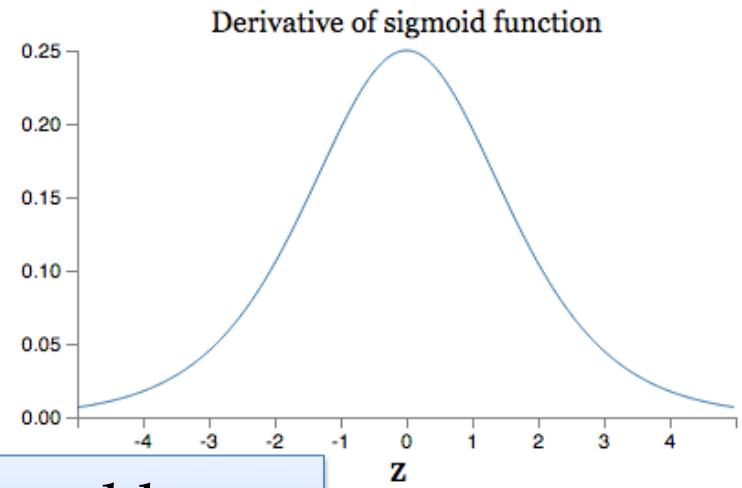

If weights are usually > 1 then we are multiplying by many numbers > 1 so the gradients get very big.

The **exploding gradient** problem (less common but possible)

What h

further from the output layer? E.g., what's gradient for the bias term with several layers after it in a trivial net?

$$\frac{\partial C}{\partial b_1} = \sigma'(z_1) \times w_2 \times \sigma'(z_2) \times w_3 \times \sigma'(z_3) \times w_4 \times \sigma'(z_4) \times \frac{\partial C}{\partial a_4}$$

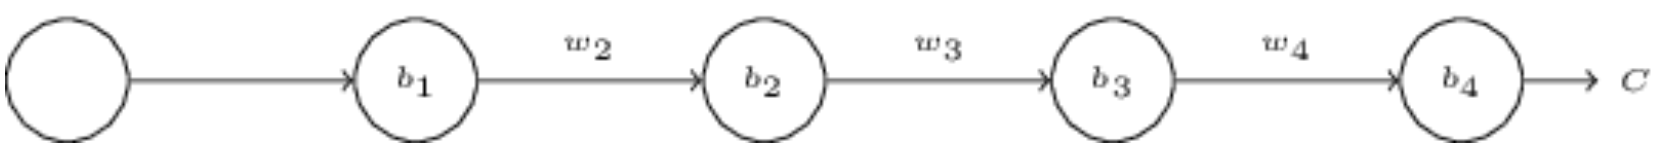

Understanding the difficulty of training deep feedforward neural networks

AI Stats 2010

Xavier Glorot **Yoshua Bengio**
DIRO, Université de Montréal, Montréal, Québec, Canada

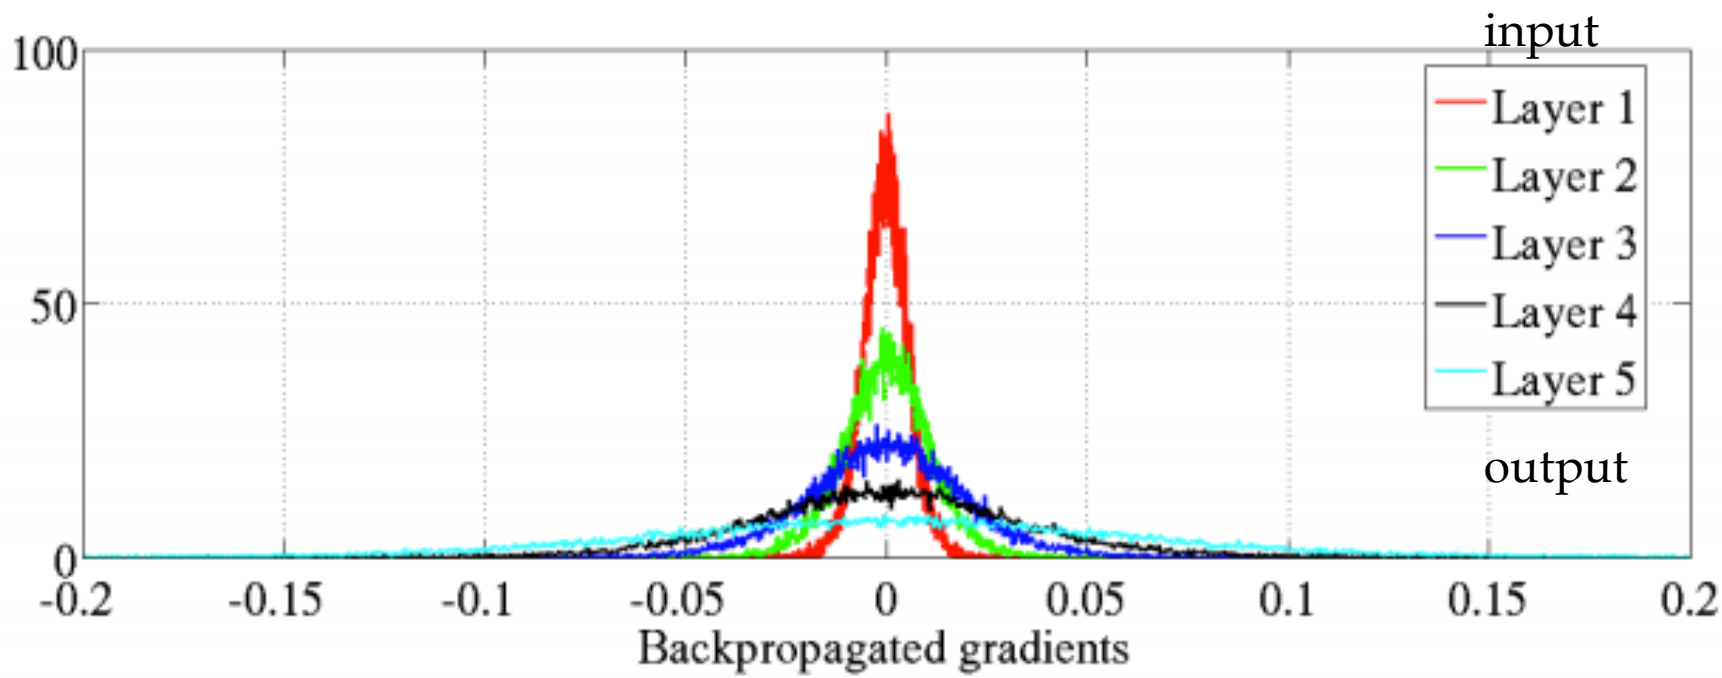

Histogram of gradients in a 5-layer network for an artificial image recognition task

Understanding the difficulty of training deep feedforward neural networks

AI Stats 2010

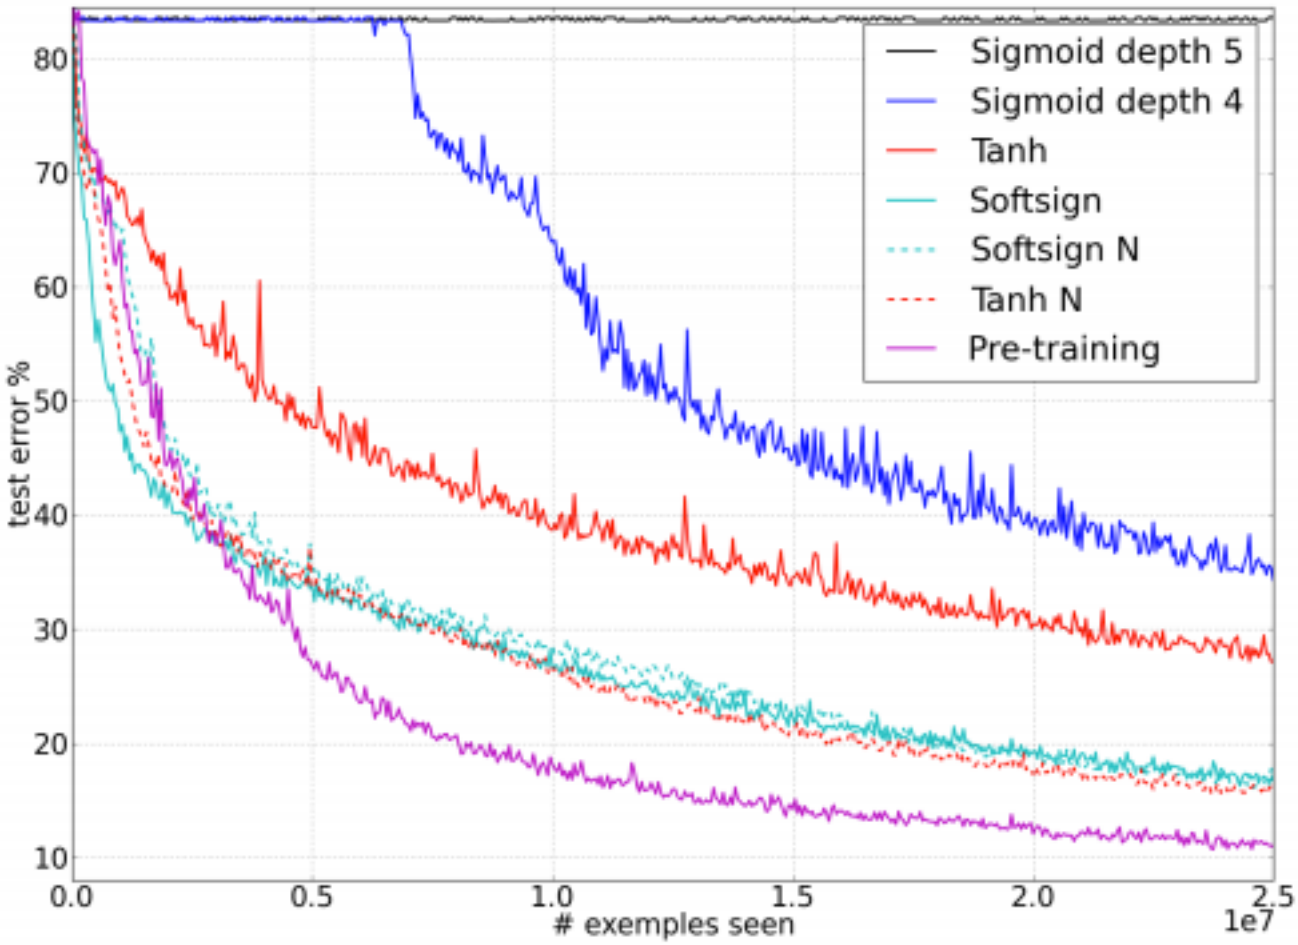

We will get to these tricks eventually….

It's easy for sigmoid units to saturate

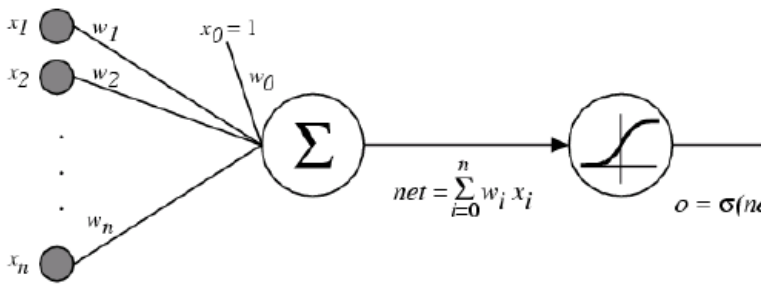

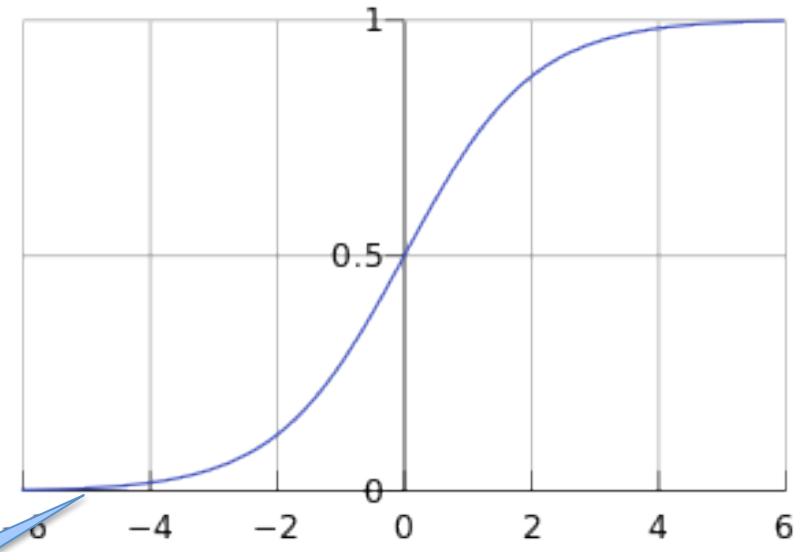

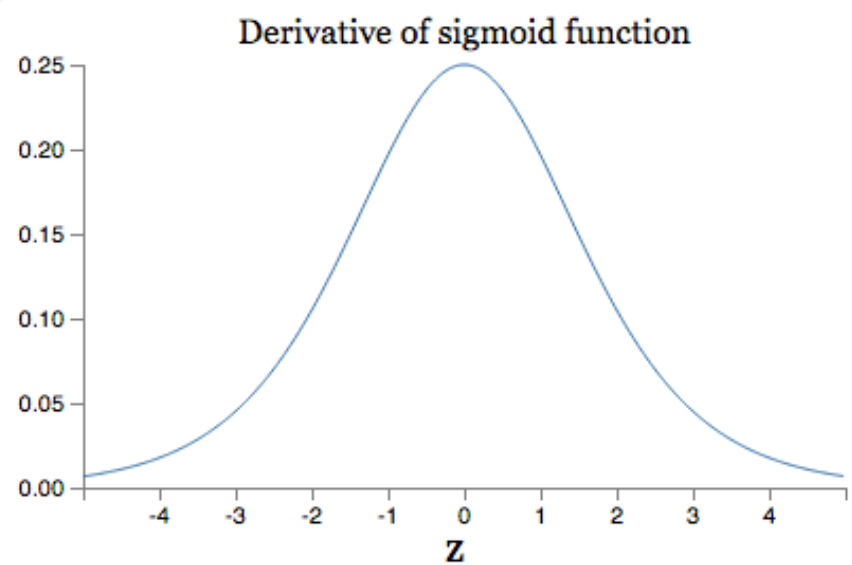

Learning rate approaches zero and unit is "stuck"

It's easy for sigmoid units to saturate

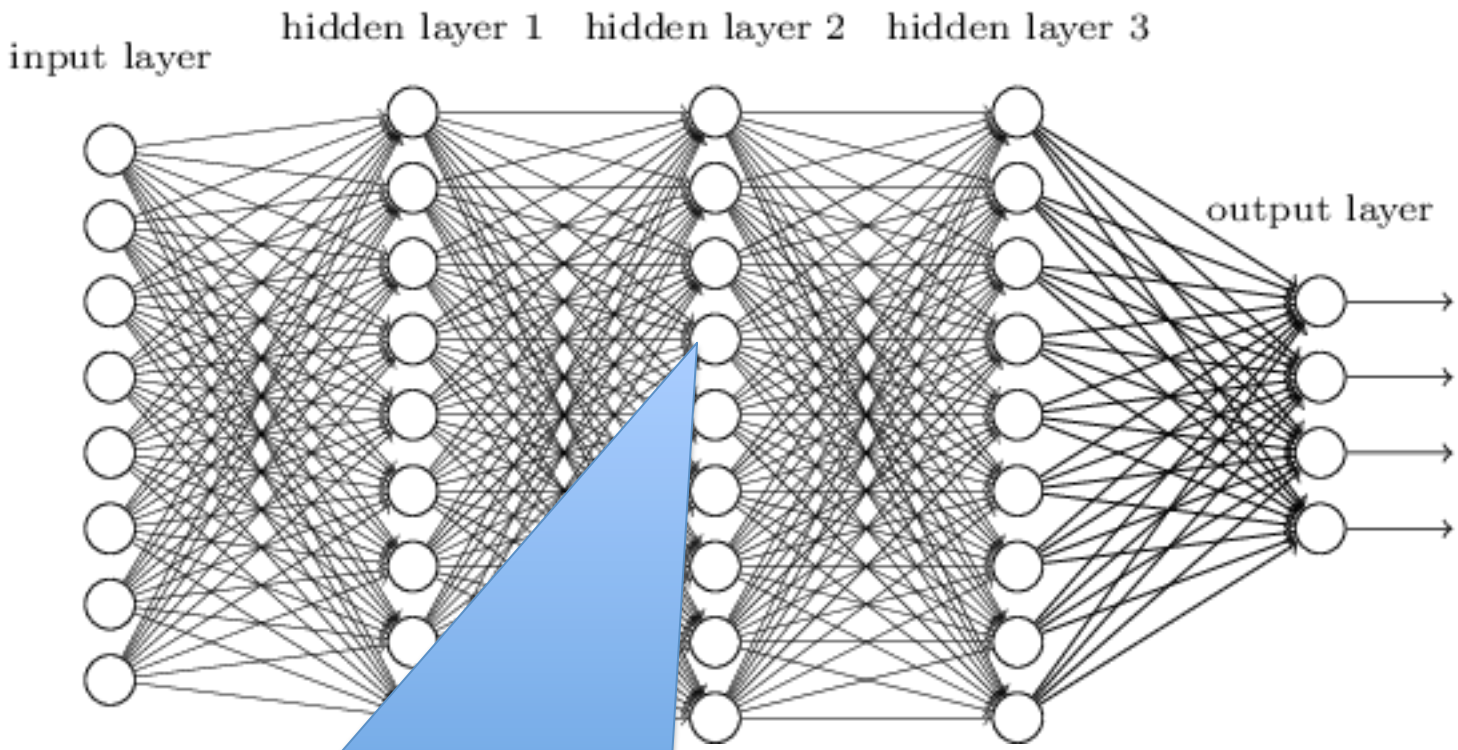

For a big network there are lots of weighted inputs to each neuron. If any of them are too large then the neuron will saturate. So neurons get stuck with a few large inputs OR many small ones.

It's easy for sigmoid units to saturate

- If there are 500 non-zero inputs initialized with a Gaussian ~N(0,1) then the SD is $\sqrt{500} \approx 22.4$

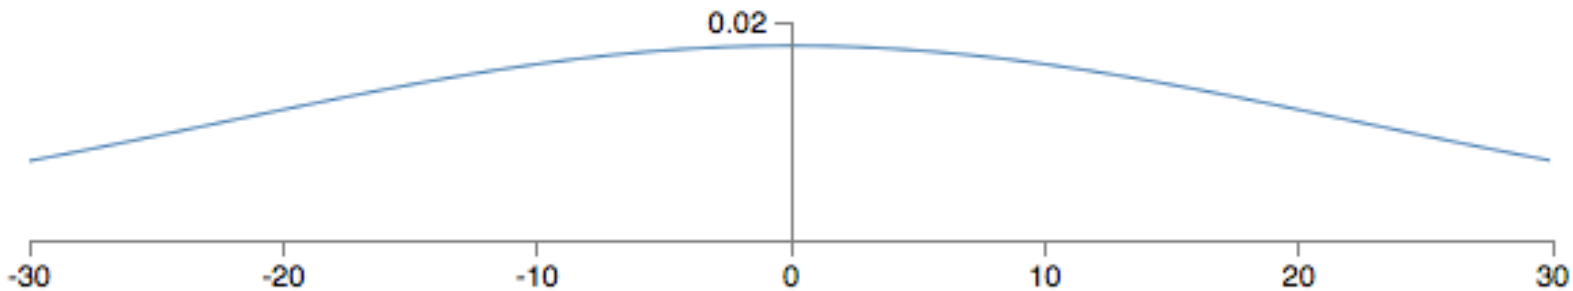

It's easy for sigmoid units to saturate

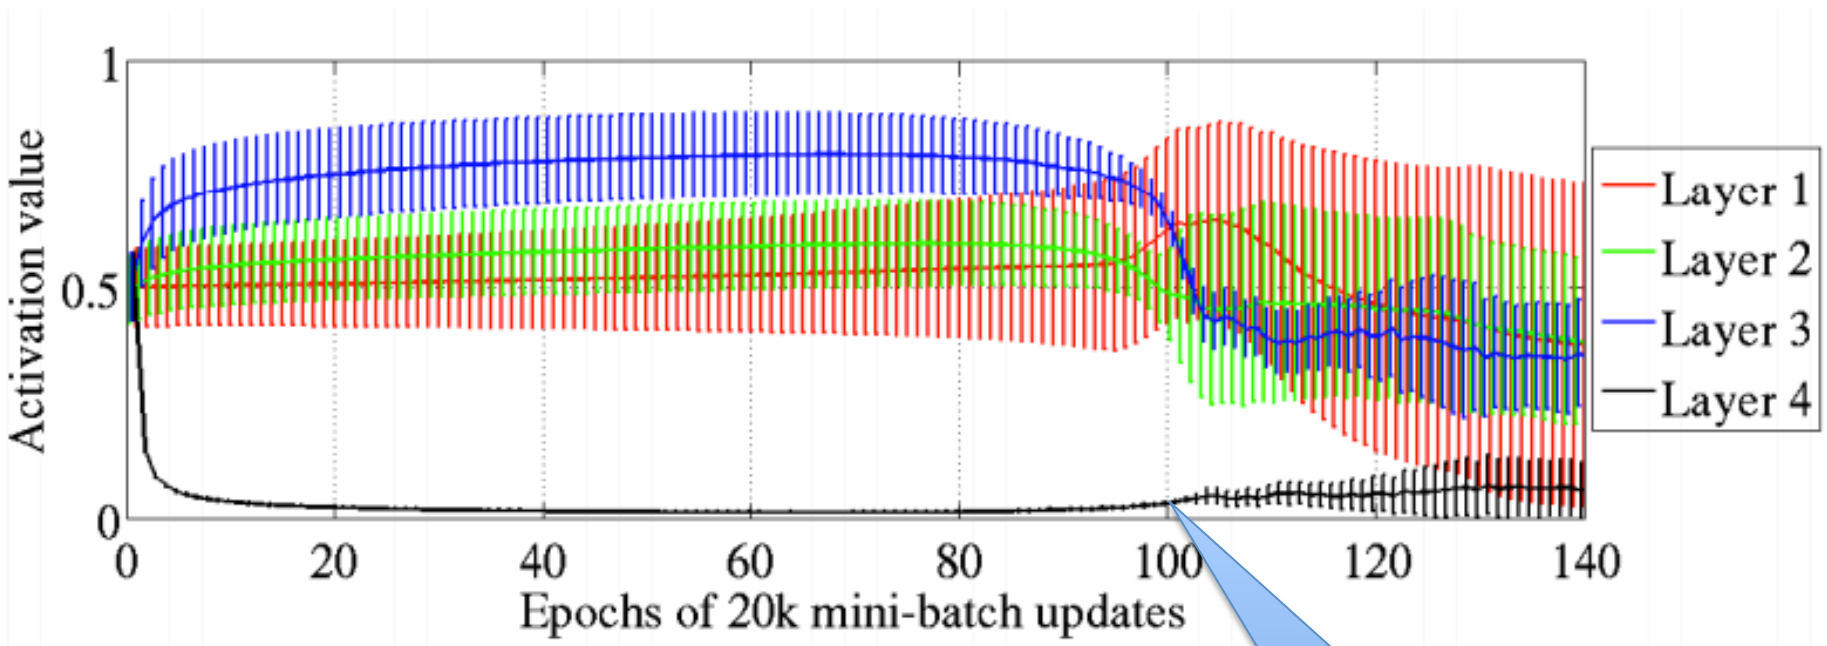

- Saturation visualization from Glorot & Bengio 2010 -- using a smarter initialization scheme

Bottom layer still stuck for first 100 epochs

Outline

- What's new in ANNs in the last 5-10 years?
  - Deeper networks, more data, and faster training
    - Scalability and use of GPUs
    - Symbolic differentiation
    - **Some subtle changes to cost function, architectures, optimization methods**
- What types of ANNs are most successful and why?
  - Convolutional networks (CNNs)
  - Long term/short term memory networks (LSTM)
  - Word2vec and embeddings
- What are the hot research topics for deep learning?

WHAT’S DIFFERENT ABOUT MODERN ANNS?

Some key differences

- Use of softmax and entropic loss instead of quadratic loss.
- Use of alternate non-linearities
  - reLU and hyperbolic tangent
- Better understanding of weight initialization
- Data augmentation
  - Especially for image data

Cross-entropy loss

$$C = -\frac{1}{n}\sum_x \left[y \ln a + (1-y)\ln(1-a)\right],$$

$$\frac{\partial C}{\partial w_j} = \frac{1}{n}\sum_x x_j(\sigma(z) - y).$$

Compare to gradient for square loss when a~=1 y=0 and x=1

$$\frac{\partial C}{\partial w} = \sigma(z) - y$$

$$\frac{\partial C}{\partial w} = (a-y)\sigma'(z)x = a\sigma'(z)$$

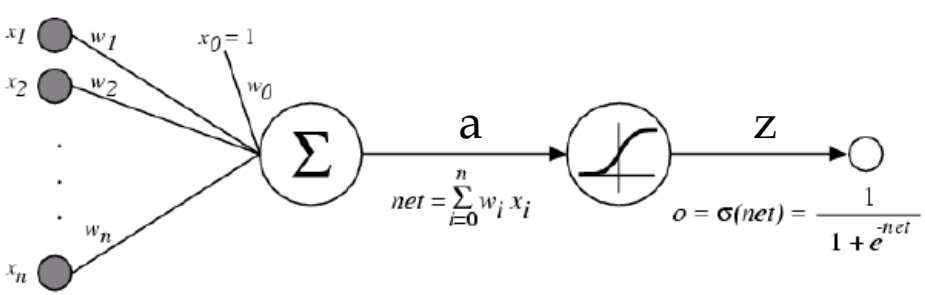

Cross-entropy loss

$$C = -\frac{1}{n} \sum_x \left[y \ln a + (1-y) \ln(1-a) \right],$$

$$\frac{\partial C}{\partial w_j} = \frac{1}{n} \sum_x x_j (\sigma(z) - y).$$

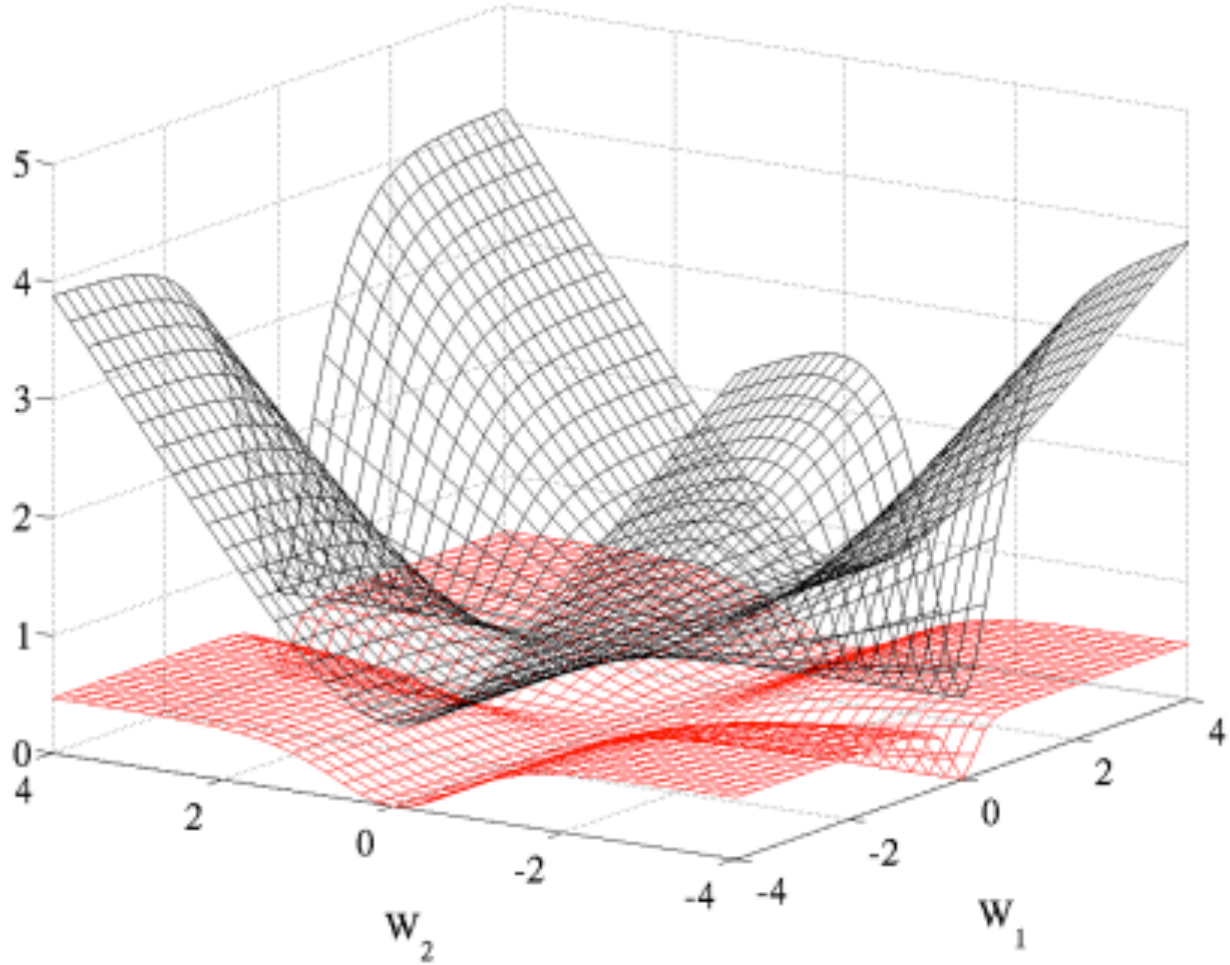

Figure 5: *Cross entropy (black, surface on top) and quadratic (red, bottom surface) cost as a function of two weights (one at each layer) of a network with two layers, W_1 respectively on the first layer and W_2 on the second, output layer.*

Softmax output layer

$$a_j^L = \frac{e^{z_j^L}}{\sum_k e^{z_k^L}},$$

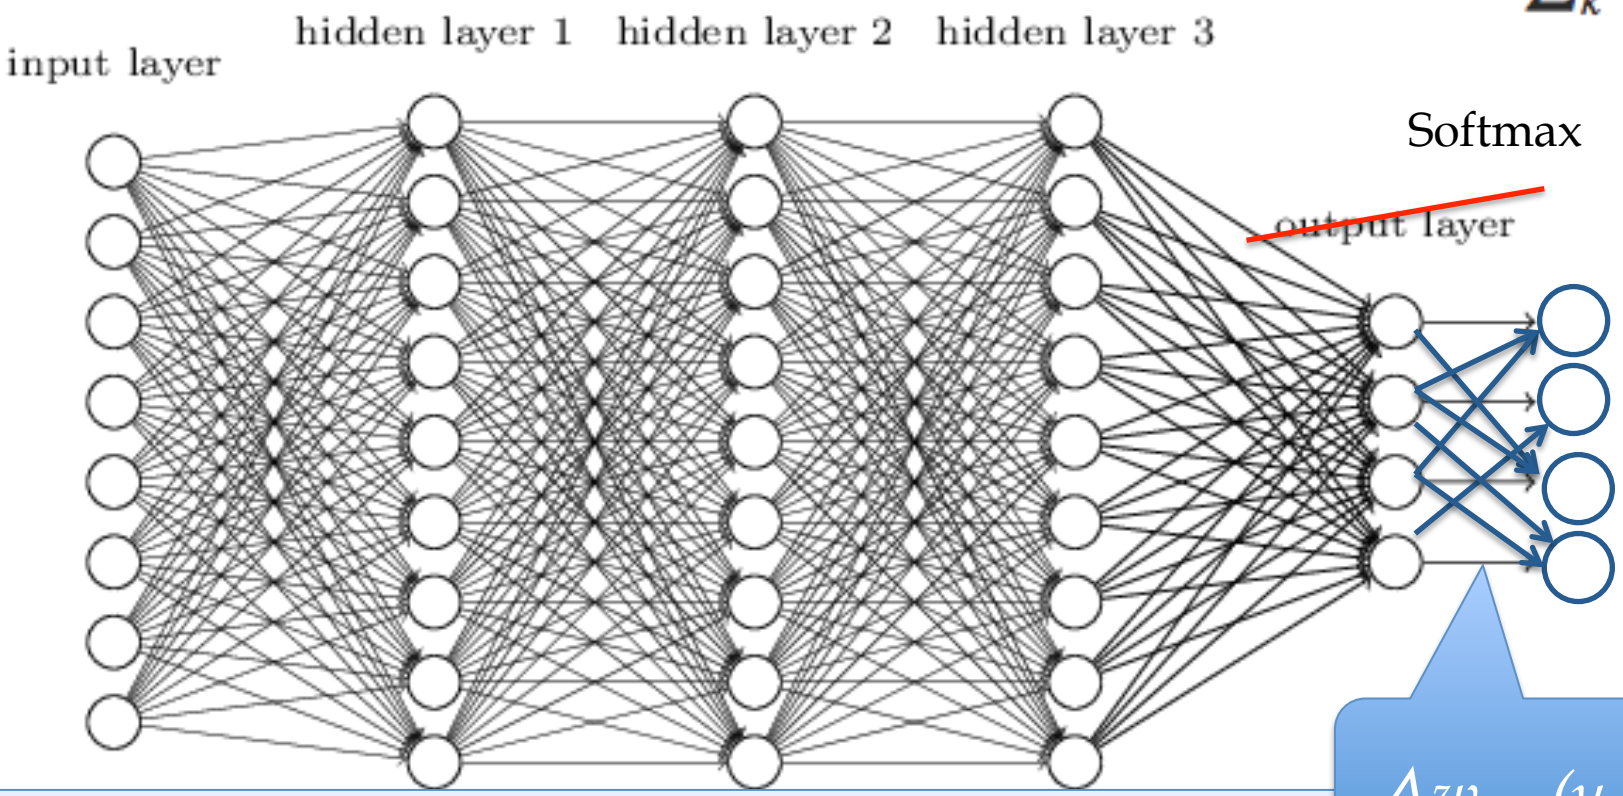

$$\Delta w_{ij} = (y_i - z_i) y_j$$

Network outputs a probability distribution!

Cross-entropy loss after a softmax layer gives a very simple, numerically stable gradient

Some key differences

- Use of softmax and entropic loss instead of quadratic loss.
  - Often learning is faster and more stable as well as getting better accuracies in the limit
- Use of alternate non-linearities
- Better understanding of weight initialization
- Data augmentation
  - Especially for image data

Some key differences

- Use of softmax and entropic loss instead of quadratic loss.
  - Often learning is faster and more stable as well as getting better accuracies in the limit
- Use of alternate non-linearities
  - reLU and hyperbolic tangent
- Better understanding of weight initialization
- Data augmentation
  - Especially for image data

Alternative non-linearities

- Changes so far
  - Changed the **loss** from square error to cross-entropy (no effect at test time)
  - Proposed adding another output layer (softmax)
- A new change: modifying the nonlinearity
  - The logistic is not widely used in modern ANNs

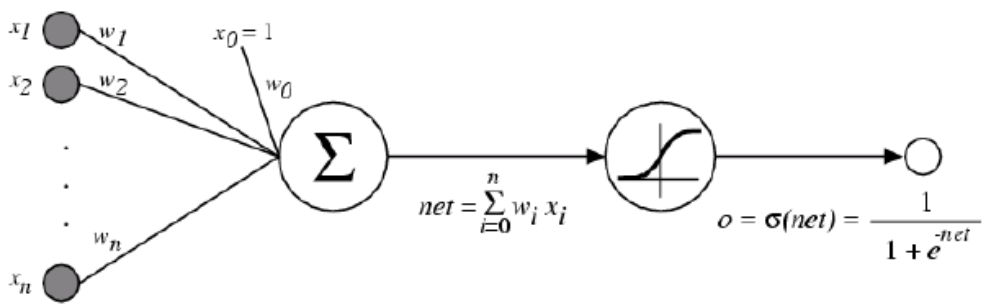

Alternative non-linearities

- A new change: modifying the nonlinearity
  - The logistic is not widely used in modern ANNs

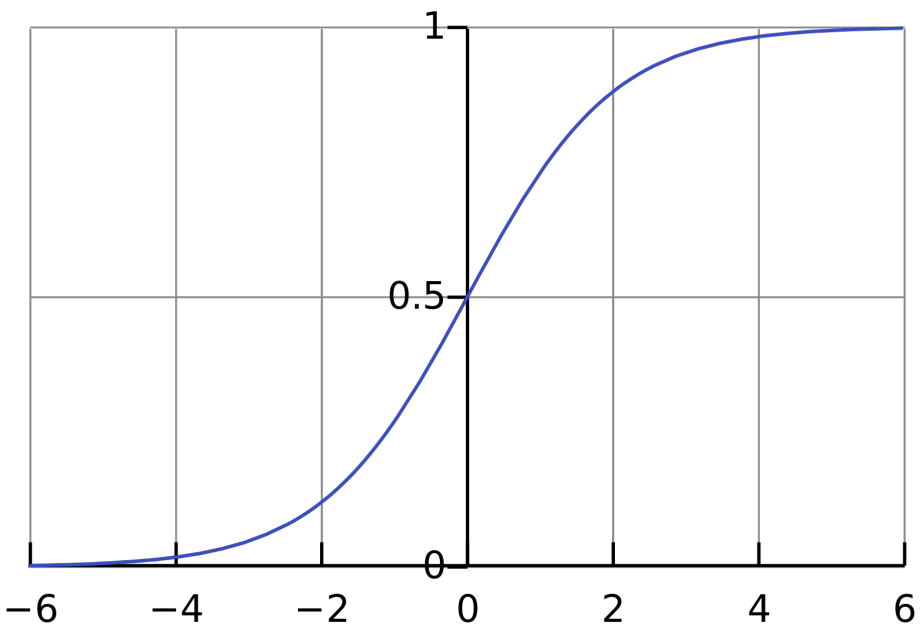

Alternate 1: tanh

Like logistic function but shifted to range [-1, +1]

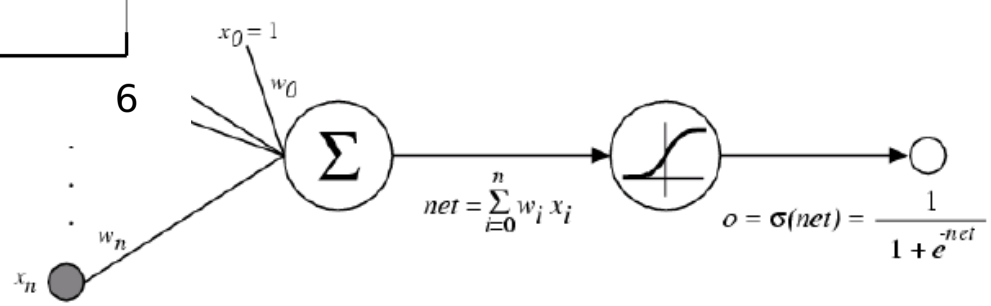

Understanding the difficulty of training deep feedforward neural networks

AI Stats 2010

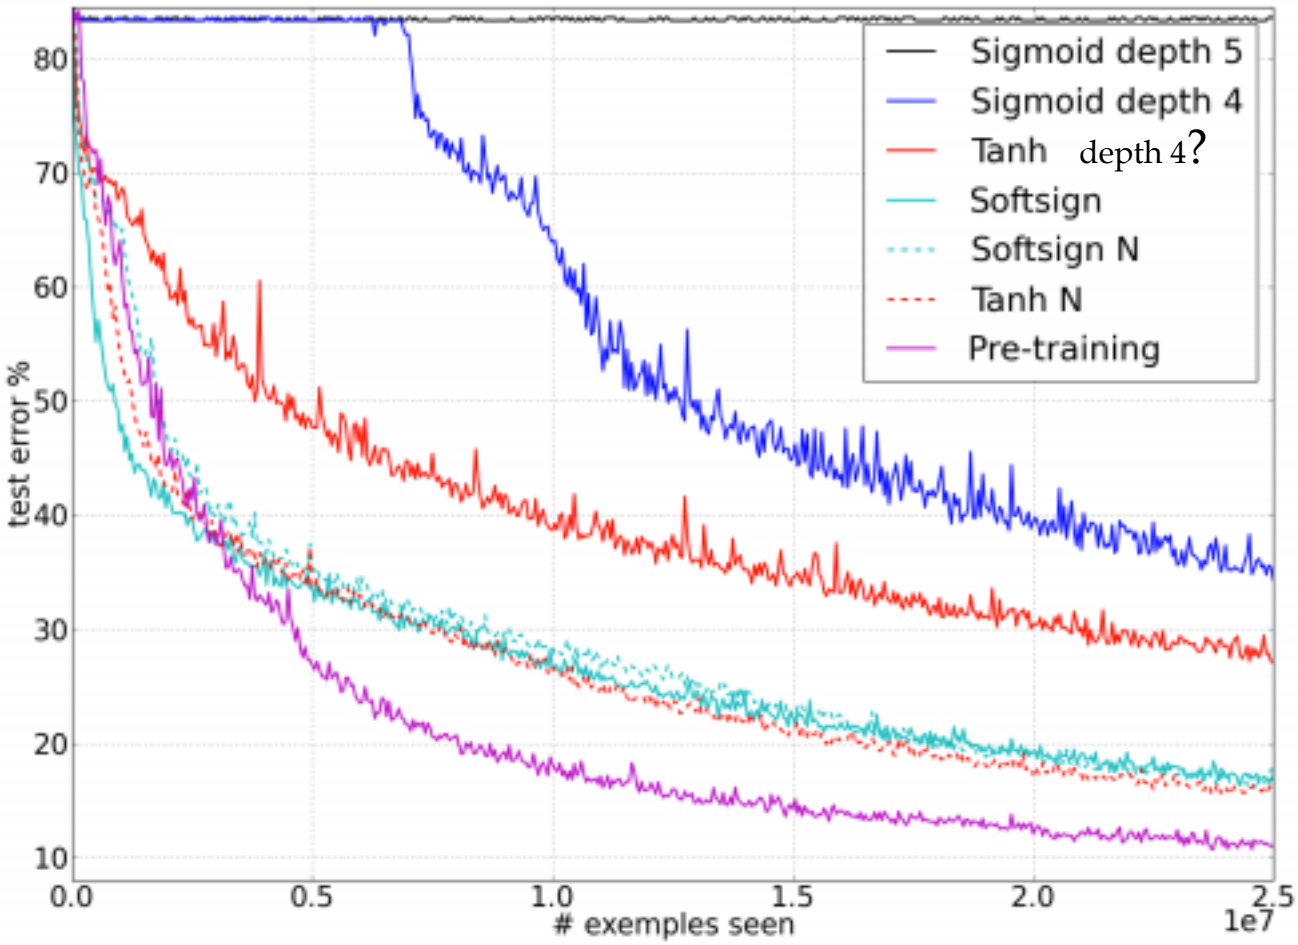

We will get to these tricks eventually….

Alternative non-linearities

- A new change: modifying the nonlinearity
  - reLU often used in vision tasks

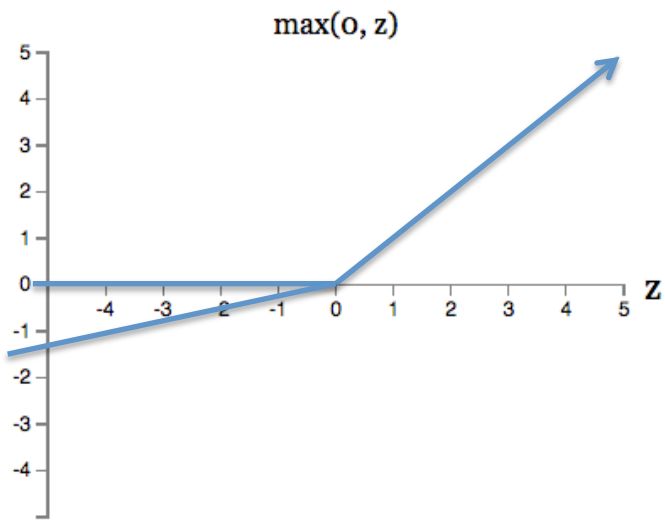

Alternate 2: rectified linear unit

Linear with a cutoff at zero

(Implement: clip the gradient when you pass zero)

$$\max(0, w \cdot x + b).$$

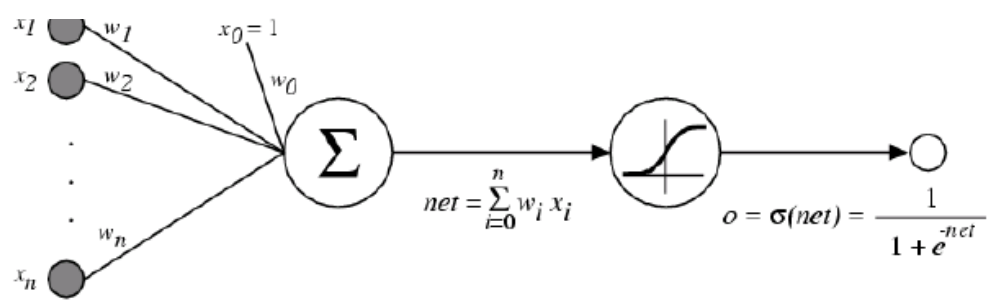

Alternative non-linearities

- A new change: modifying the nonlinearity
  - reLU often used in vision tasks

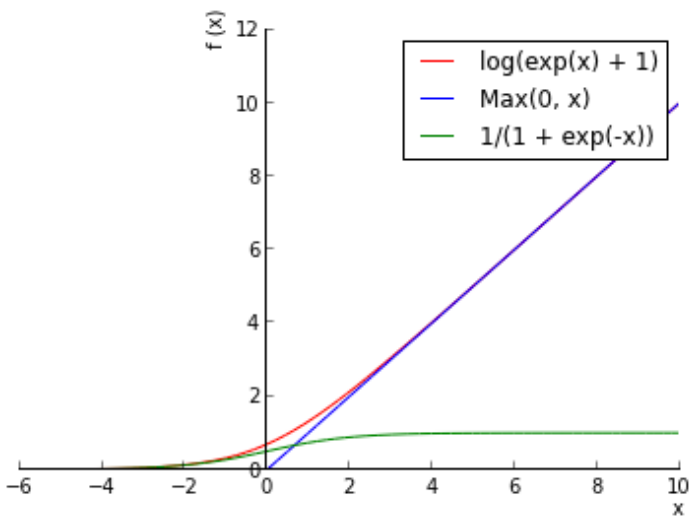

Alternate 2: rectified linear unit

Soft version: log(exp(x)+1)

Doesn't saturate (at one end)
Sparsifies outputs
Helps with vanishing gradient

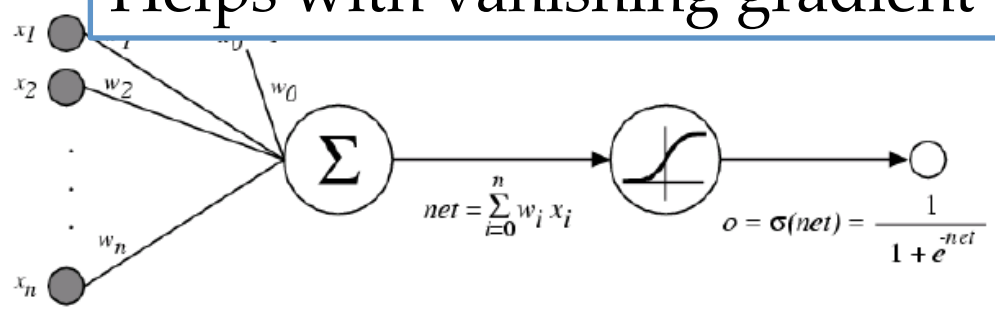

Some key differences

- Use of softmax and entropic loss instead of quadratic loss.
  - Often learning is faster and more stable as well as getting better accuracies in the limit
- Use of alternate non-linearities
  - reLU and hyperbolic tangent
- Better understanding of weight initialization
- Data augmentation
  - Especially for image data

It's easy for sigmoid units to saturate

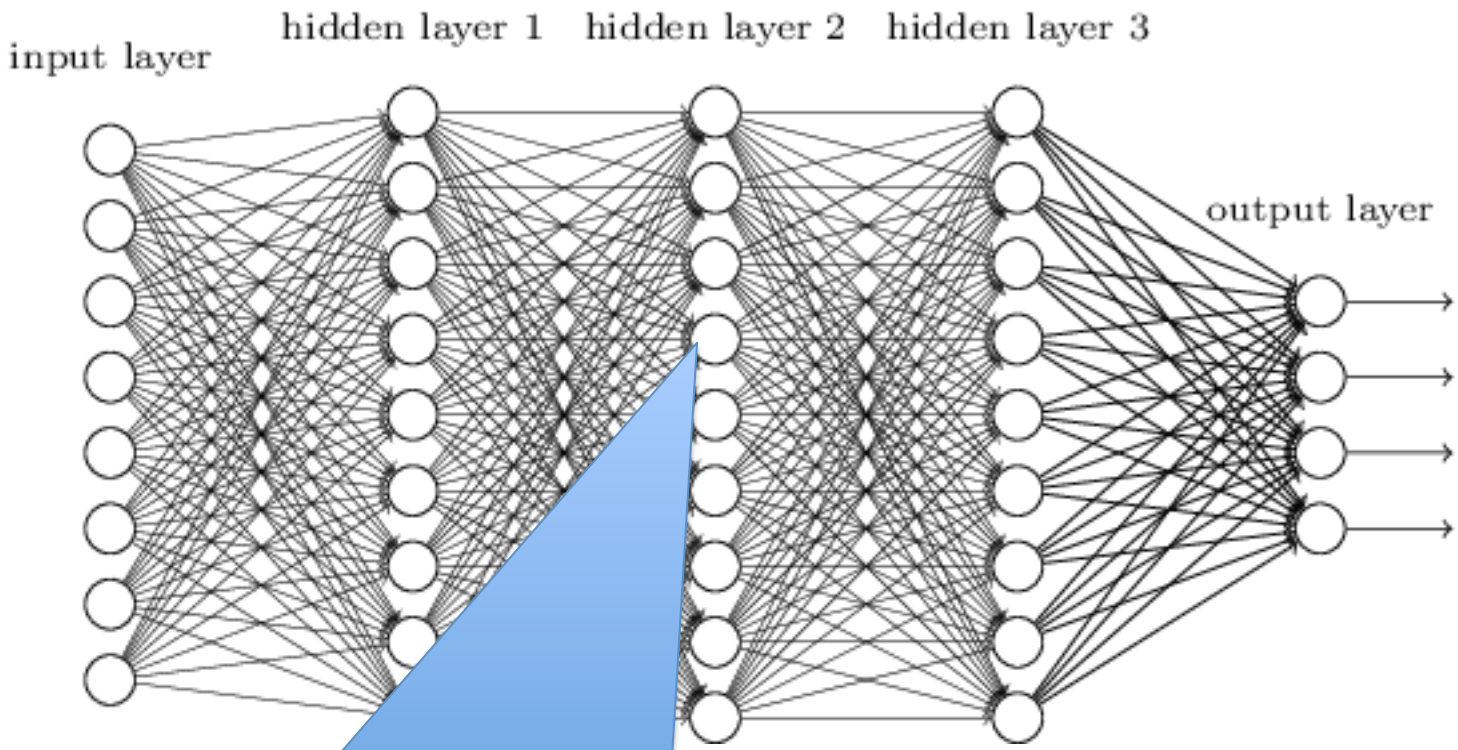

For a big network there are lots of weighted inputs to each neuron. If any of them are too large then the neuron will saturate. So neurons get stuck with a few large inputs OR many small ones.

It's easy for sigmoid units to saturate

- If there are 500 non-zero inputs initialized with a Gaussian ~N(0,1) then the SD is $\sqrt{500} \approx 22.4$

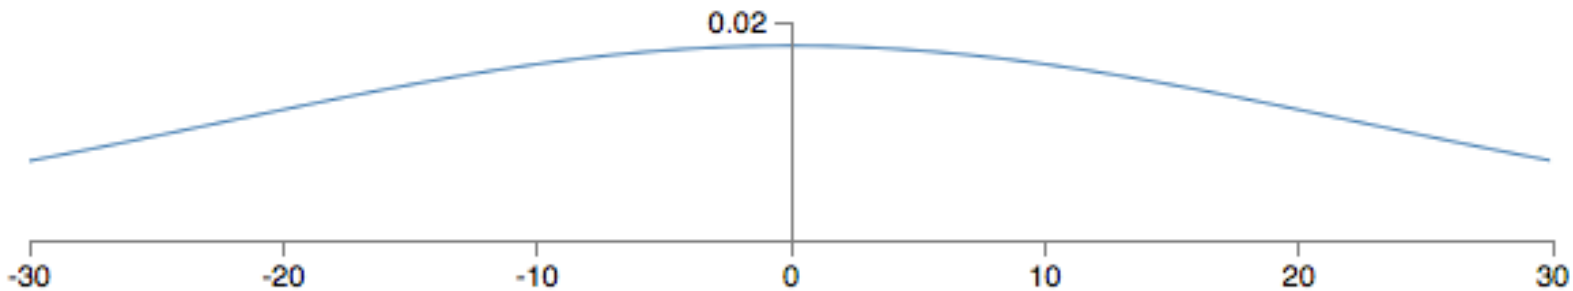

- Common heuristics for initializing weights:

$$N\left(0, \frac{1}{\sqrt{\#\text{inputs}}}\right) \qquad U\left(\frac{-1}{\sqrt{\#\text{inputs}}}, \frac{-1}{\sqrt{\#\text{inputs}}}\right)$$

It's easy for sigmoid units to saturate

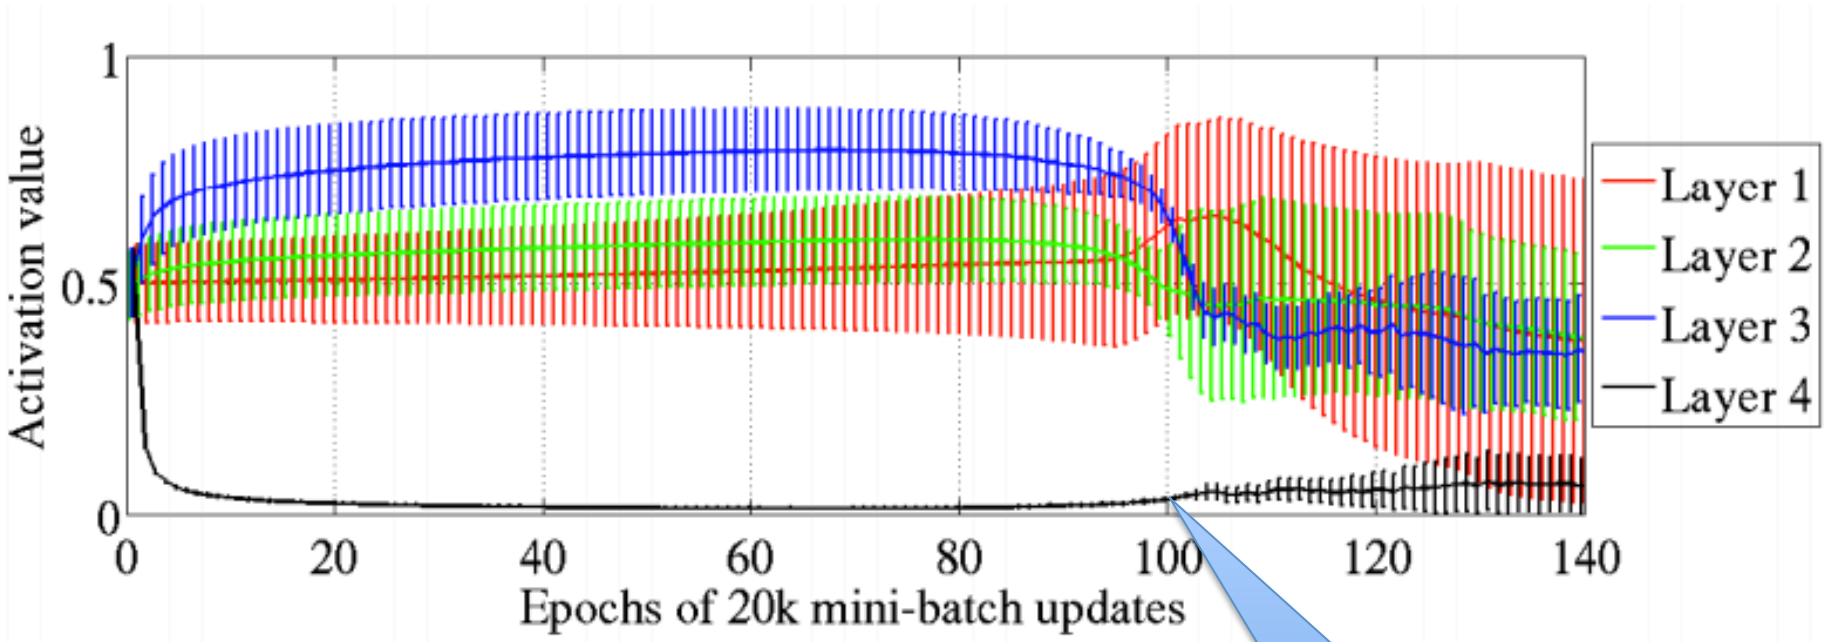

- Saturation visualization from Glorot Bengio 2010 using

$$U\left(\frac{-1}{\sqrt{\#\text{inputs}}}, \frac{1}{\sqrt{\#\text{inputs}}}\right)$$

Bottom layer still stuck for first 100 epochs

Initializing to avoid saturation

- In Glorot and Bengio they suggest weights if leve

First breakthrough deep learning results were based on clever pre-training initialization schemes, where deep networks were seeded with weights learned from unsupervised strategies

TYPE	Shapeset	MNIST	CIFAR-10	ImageNet
Tanh	27.15	**1.76**	55.9	70.58
Tanh N	**15.60**	**1.64**	**52.92**	**68.57**
Sigmoid	82.61	2.21	57.28	70.66

This is not always the solution – but good initialization is very important for deep nets!

Outline

- What's new in ANNs in the last 5-10 years?
  - **Deeper networks**, more data, and faster training
    - Scalability and use of GPUs
    - Symbolic differentiation
    - Some subtle changes to cost function, architectures, optimization methods
- What types of ANNs are most successful and why?
  - **Convolutional networks (CNNs)**
  - Long term/short term memory networks (LSTM)
  - Word2vec and embeddings
- What are the hot research topics for deep learning?

WHAT’S A CONVOLUTIONAL NEURAL NETWORK?

Model of vision in animals

- **[Hubel & Wiesel 1962]:**
  - simple cells detect local features
  - complex cells "pool" the outputs of simple cells within a retinotopic neighborhood.

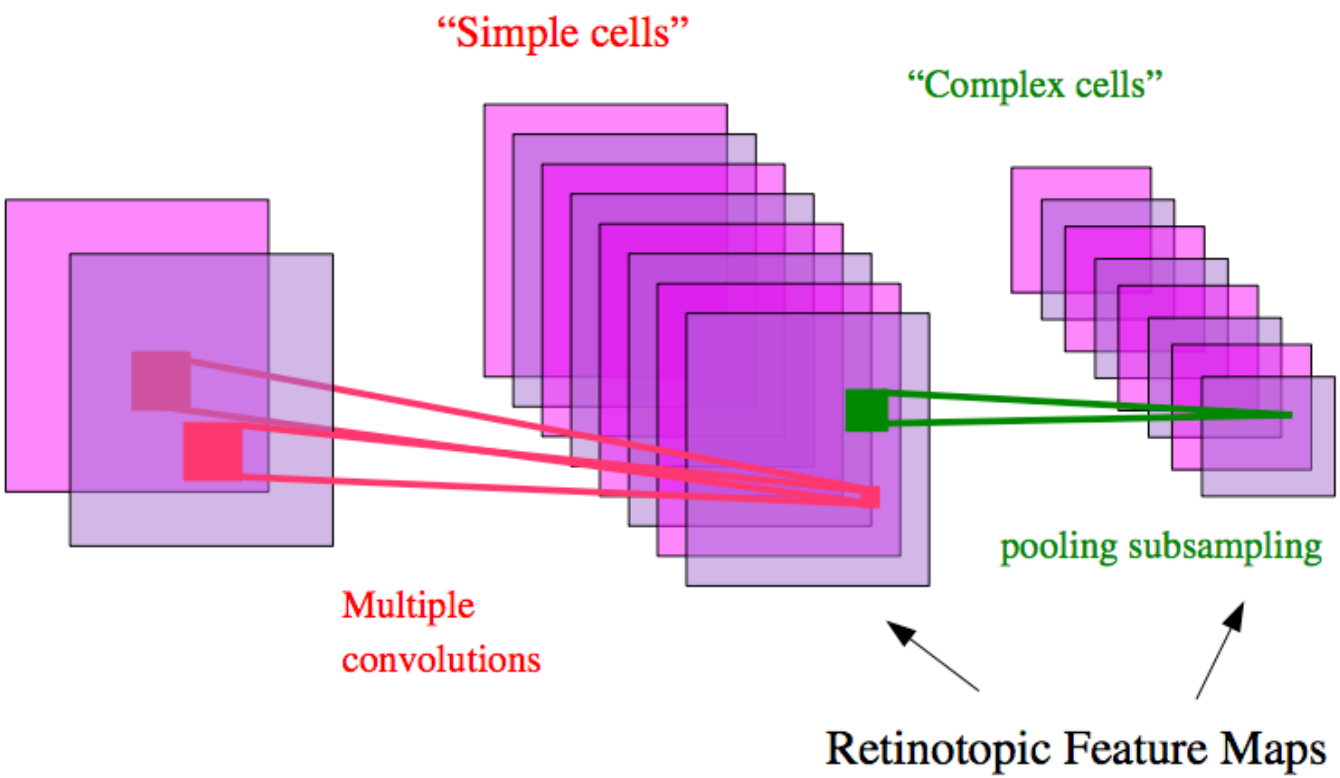

Vision with ANNs

(LeCun et al., 1989)

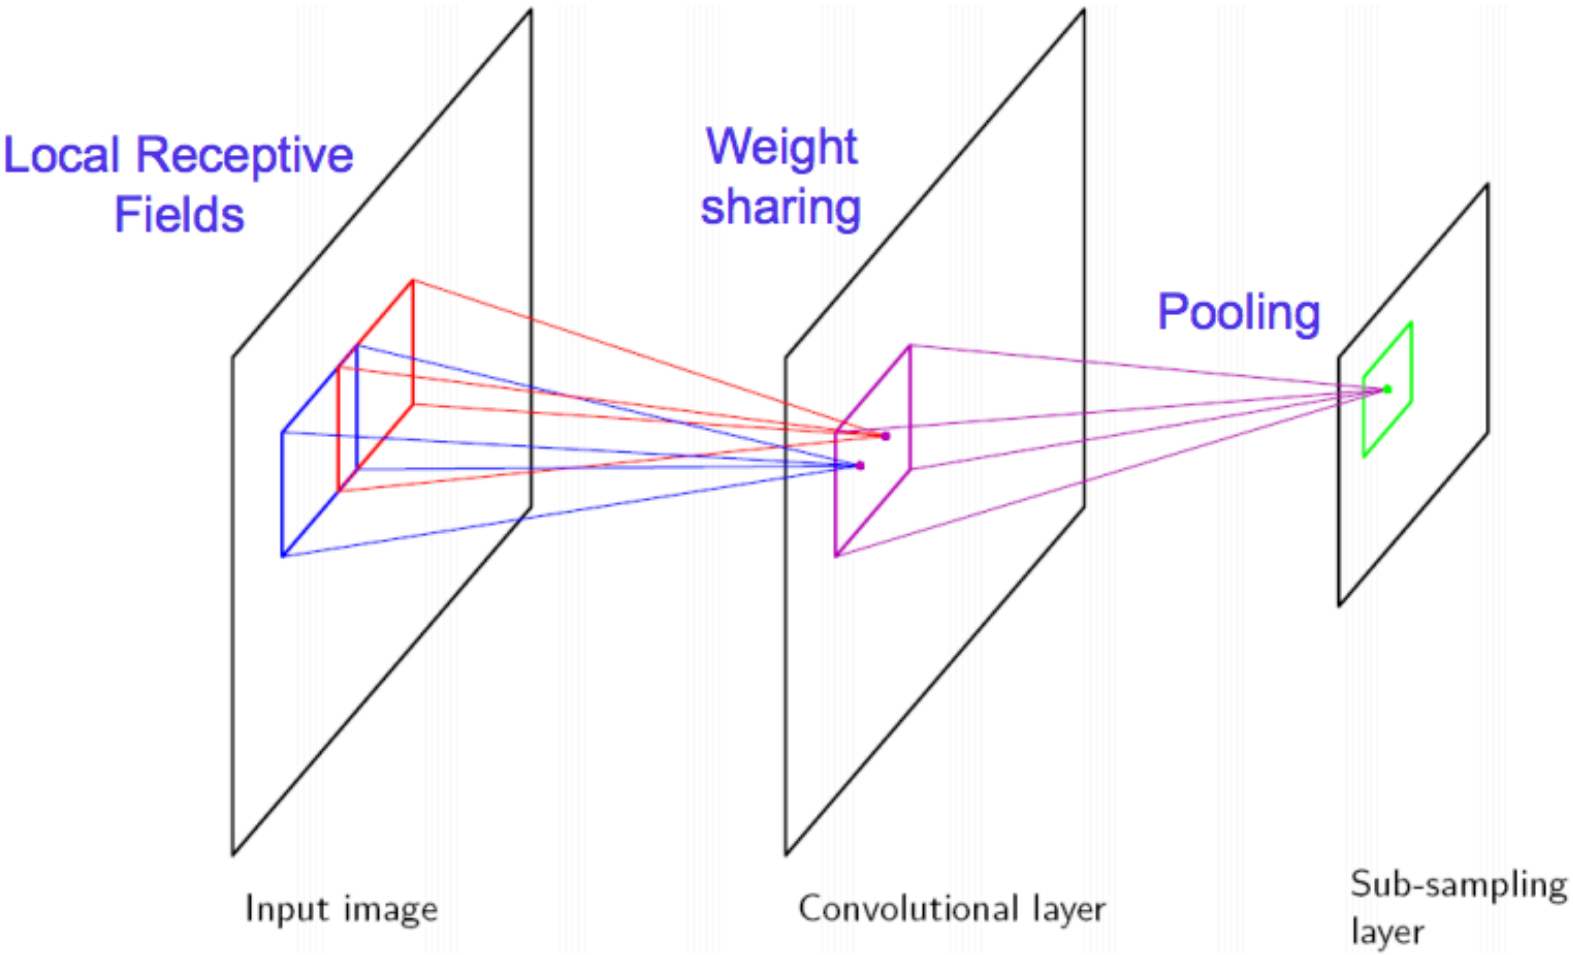

What's a convolution?

https://en.wikipedia.org/wiki/Convolution

1-D

$$(f * g)(t) \stackrel{\text{def}}{=} \int_{-\infty}^{\infty} f(\tau)\, g(t-\tau)\, d\tau$$
$$= \int_{-\infty}^{\infty} f(t-\tau)\, g(\tau)\, d\tau.$$

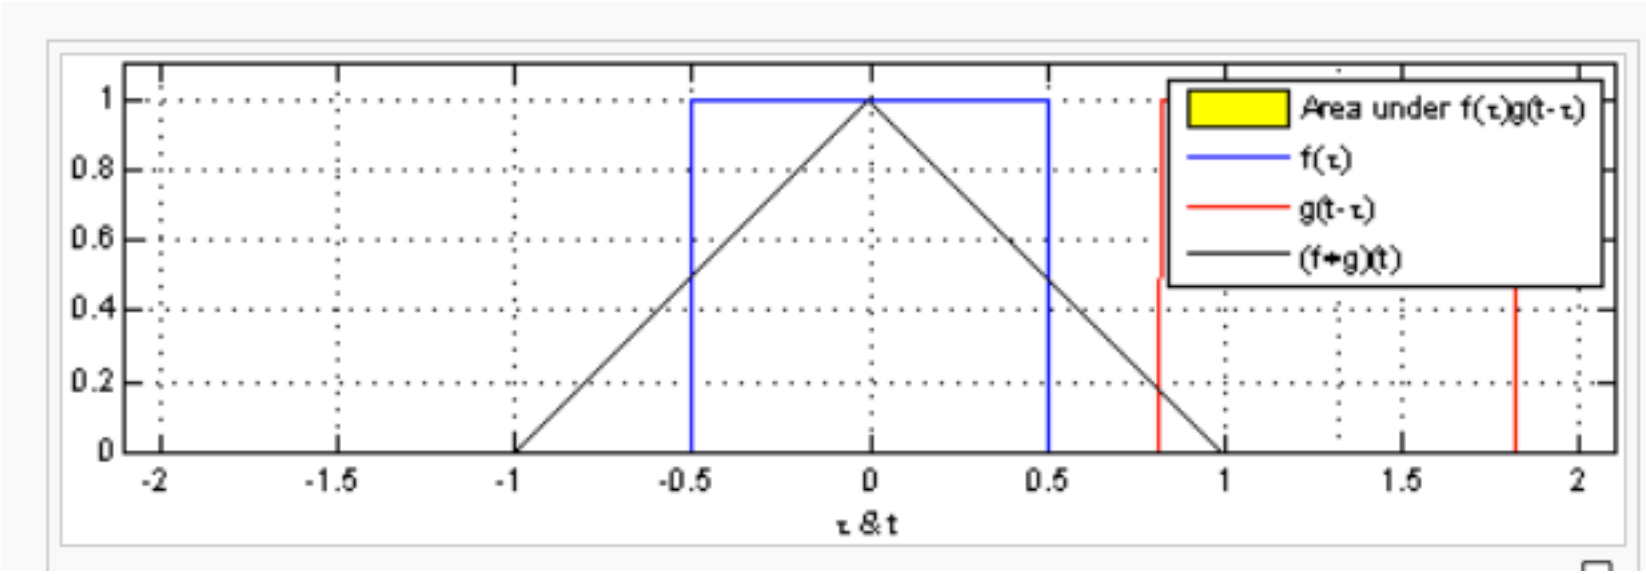

What's a convolution?

https://en.wikipedia.org/wiki/Convolution

1-D

$$(f * g)(t) \stackrel{\text{def}}{=} \int_{-\infty}^{\infty} f(\tau)\, g(t-\tau)\, d\tau$$
$$= \int_{-\infty}^{\infty} f(t-\tau)\, g(\tau)\, d\tau.$$

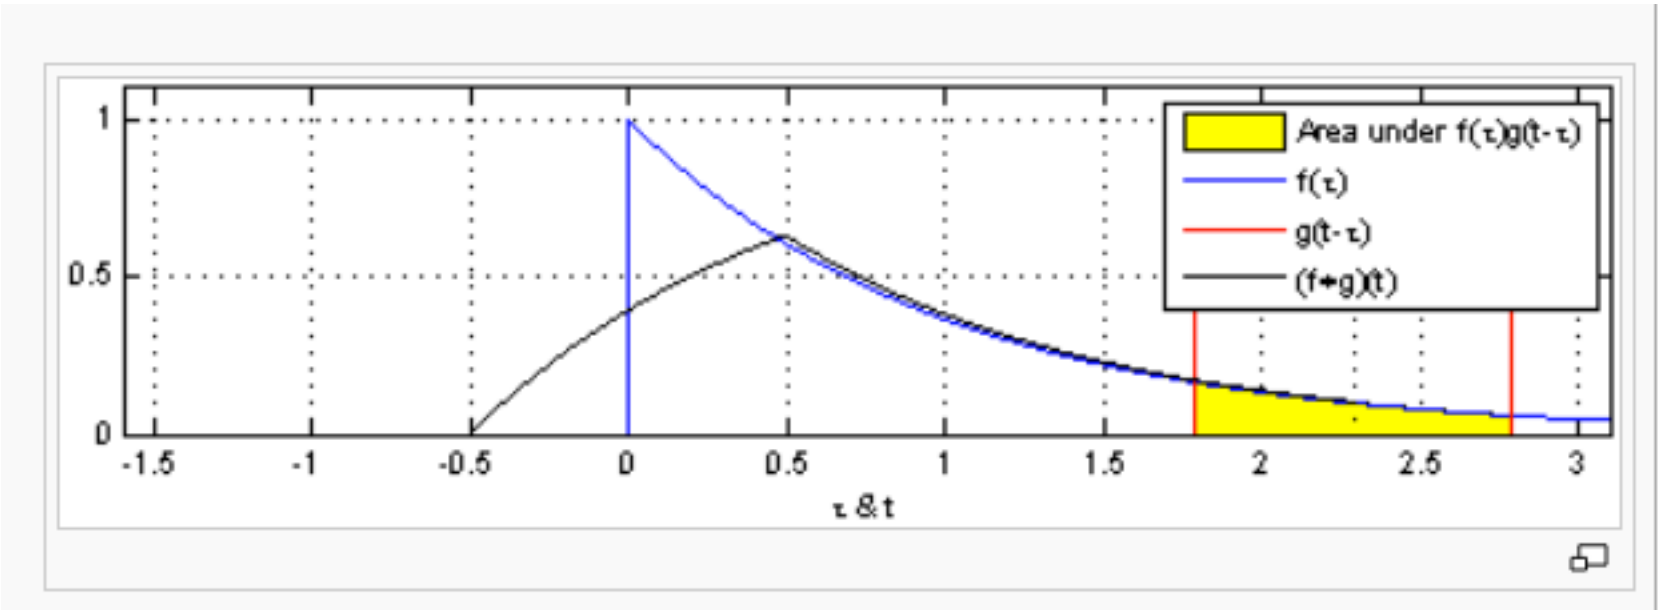

What's a convolution?

http://matlabtricks.com/post-5/3x3-convolution-kernels-with-online-demo

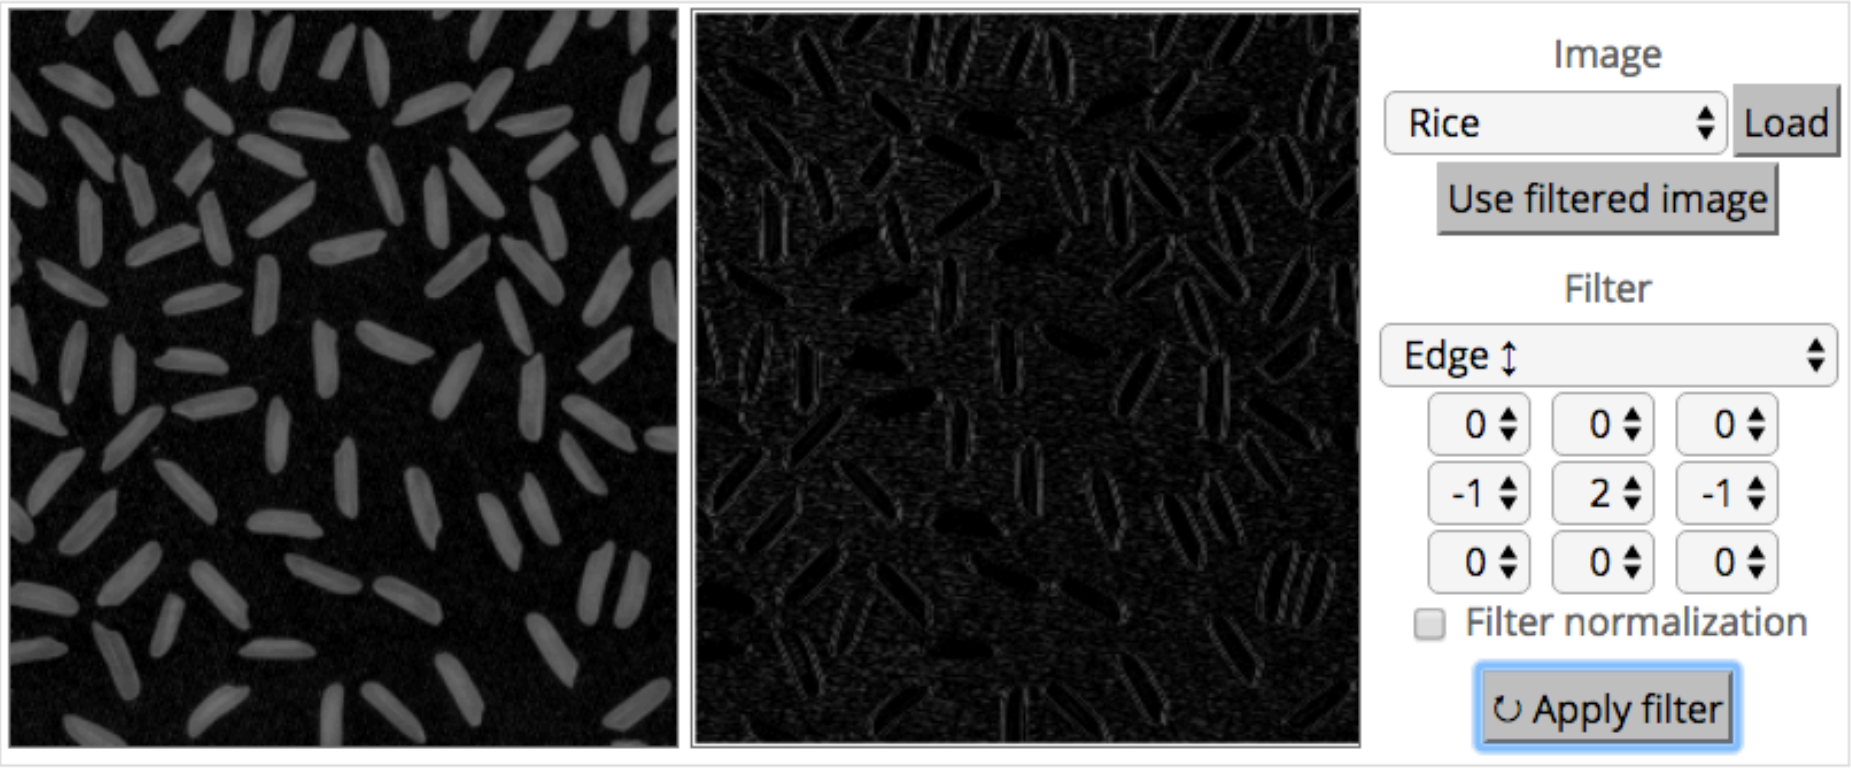

What's a convolution?

http://matlabtricks.com/post-5/3x3-convolution-kernels-with-online-demo

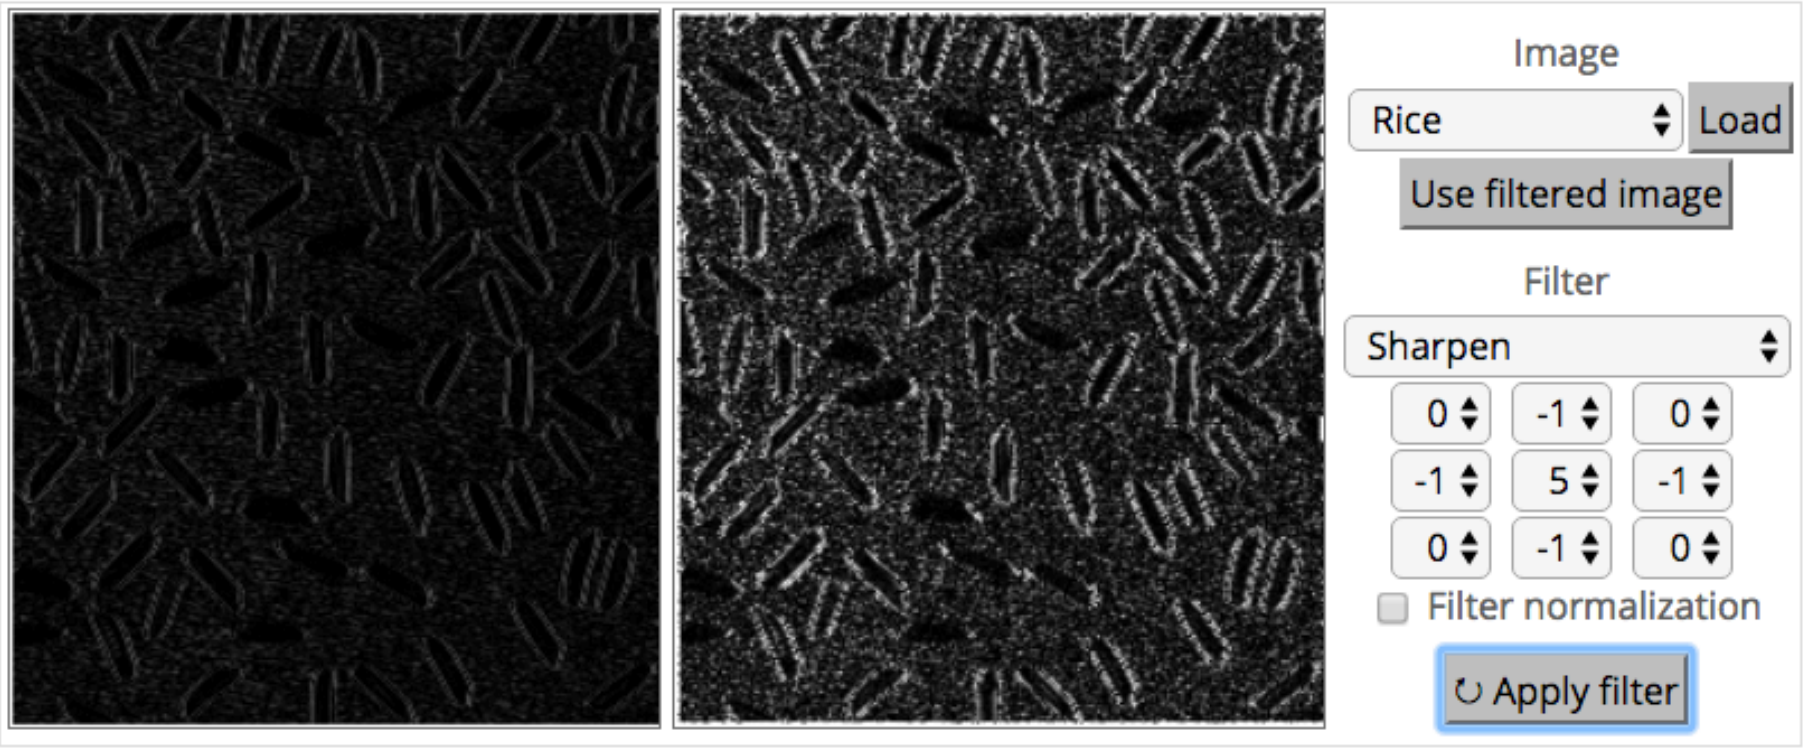

What's a convolution?

http://matlabtricks.com/post-5/3x3-convolution-kernels-with-online-demo

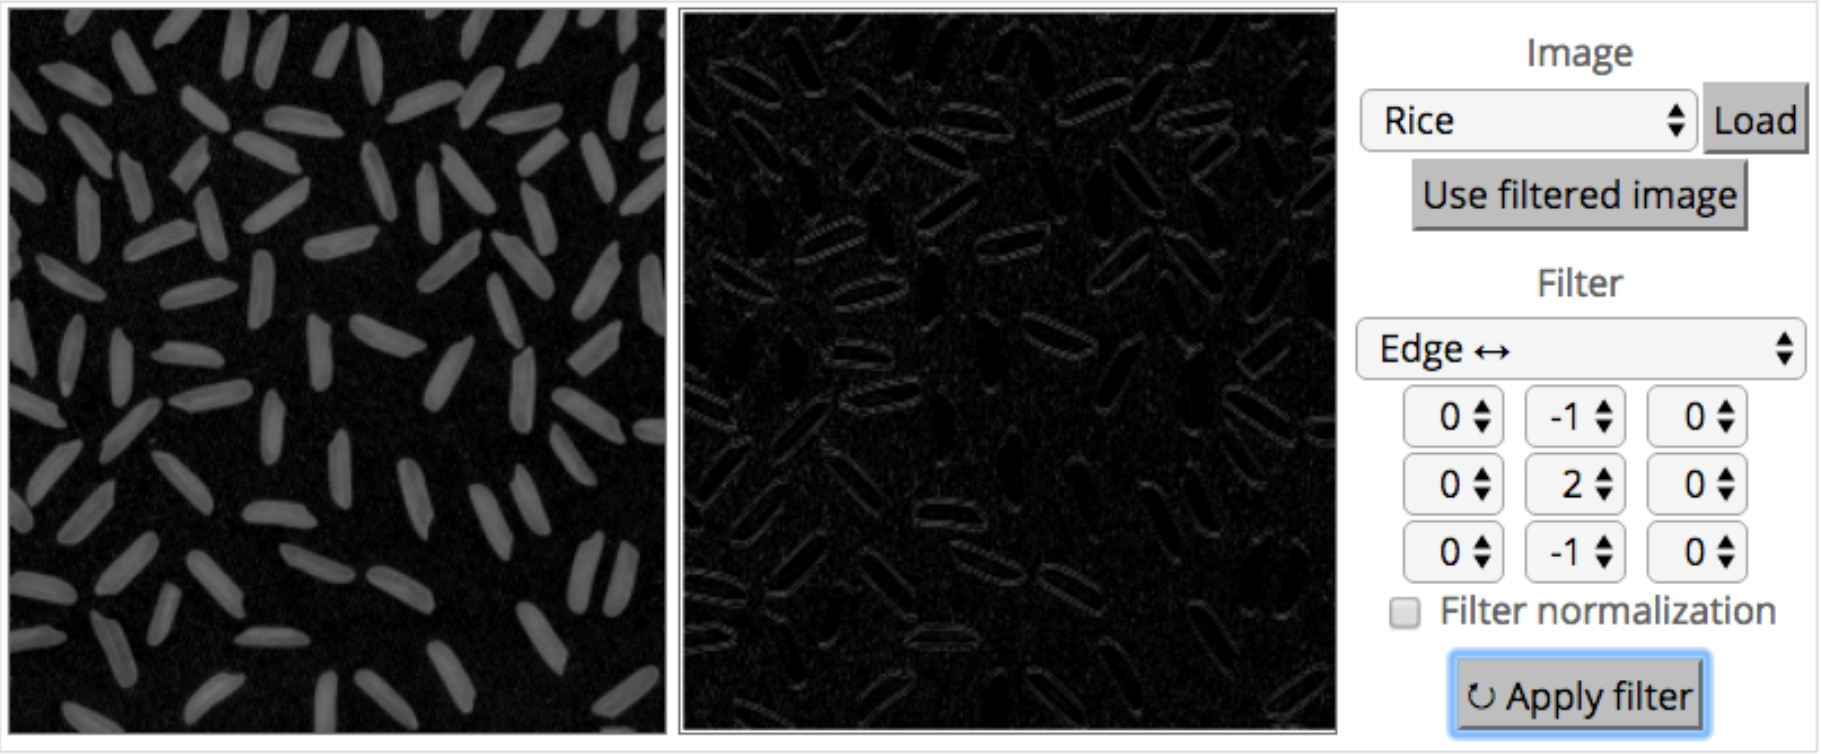

What's a convolution?

http://matlabtricks.com/post-5/3x3-convolution-kernels-with-online-demo

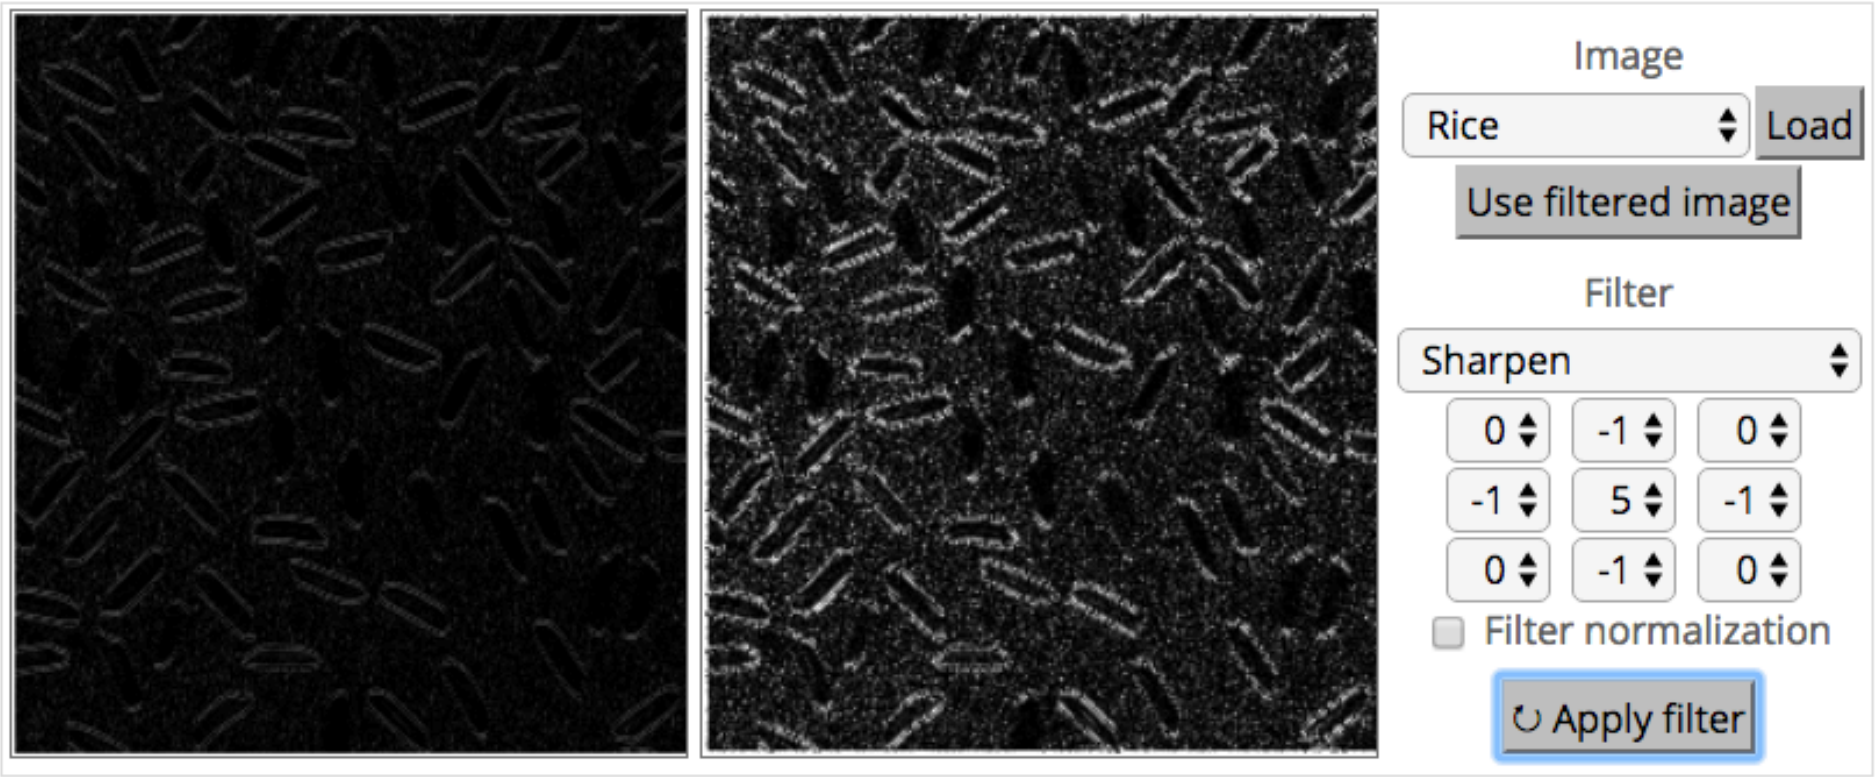

What's a convolution?

http://matlabtricks.com/post-5/3x3-convolution-kernels-with-online-demo

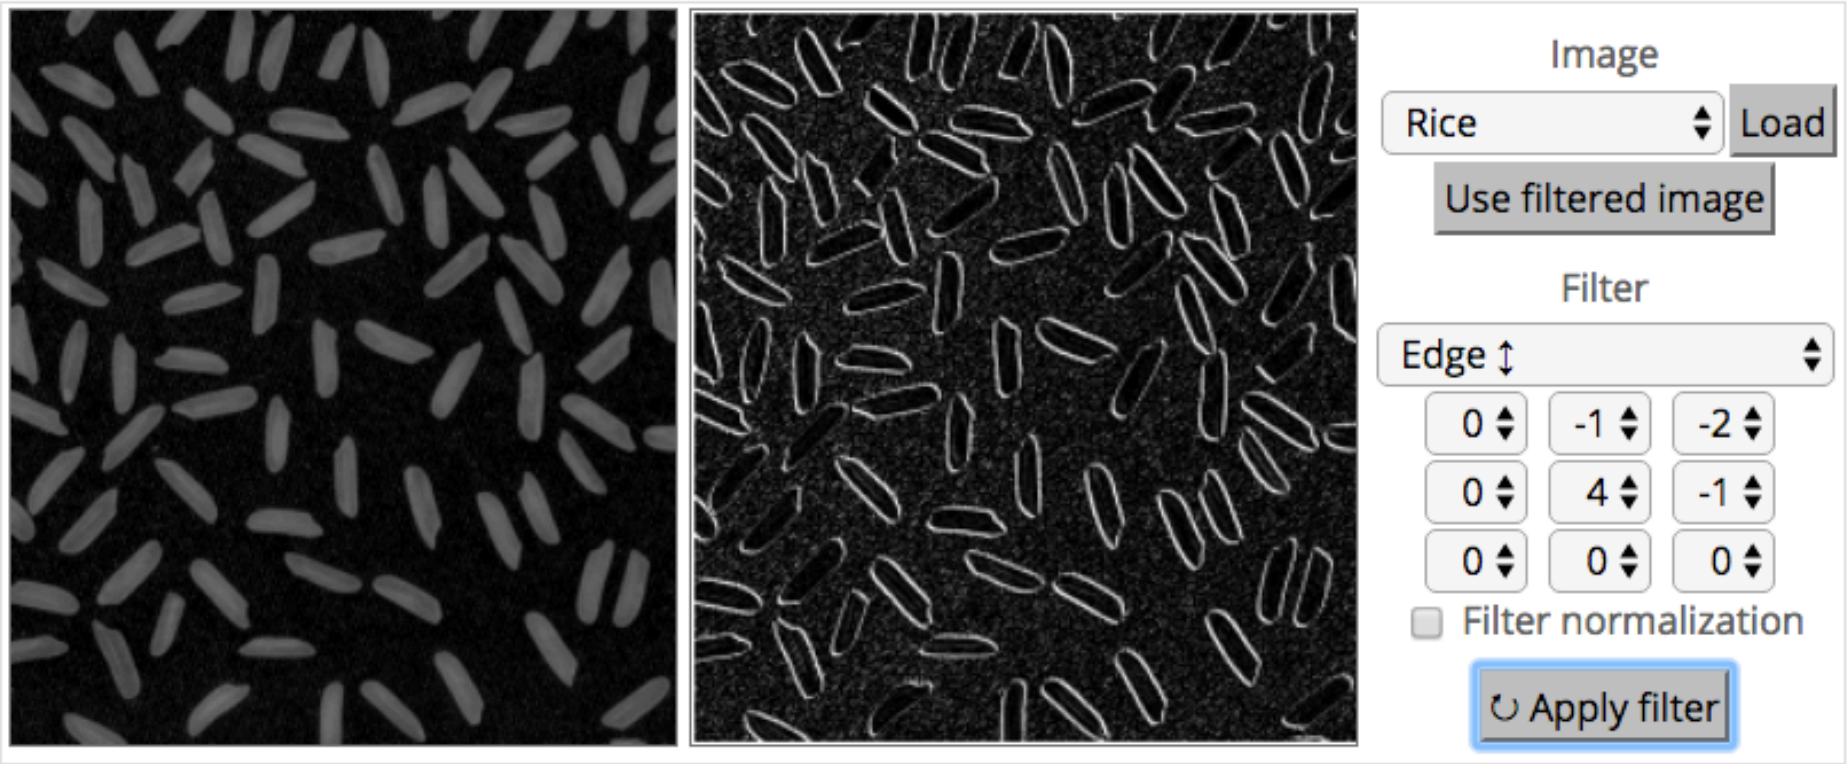

What's a convolution?

http://matlabtricks.com/post-5/3x3-convolution-kernels-with-online-demo

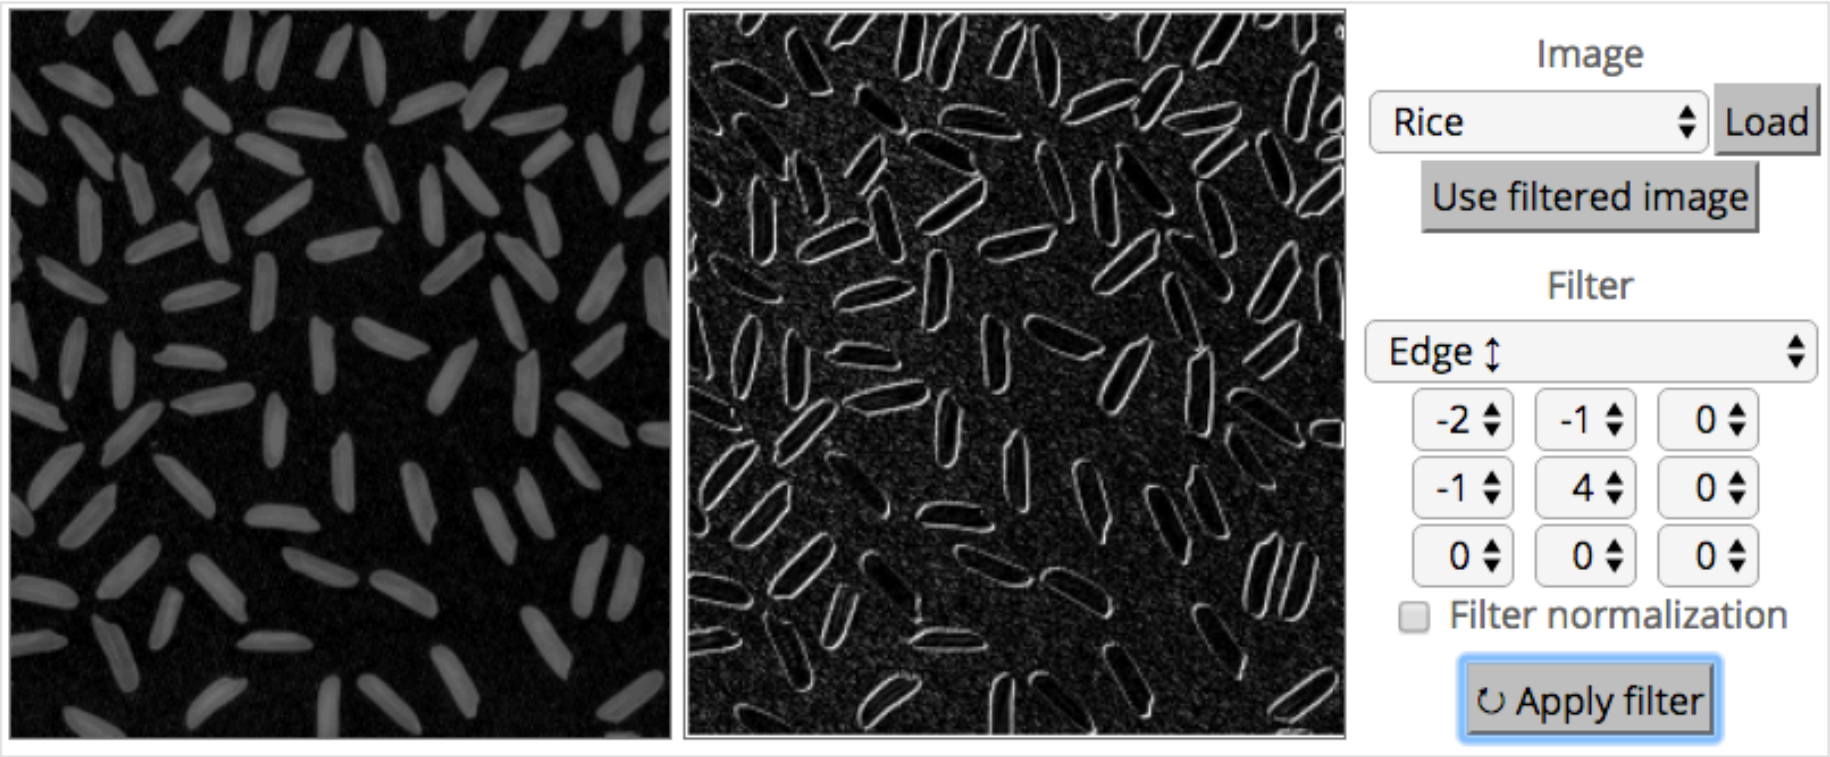

What's a convolution?

- Basic idea:
  - Pick a 3-3 matrix F of weights
  - Slide this over an image and compute the "inner product" (similarity) of F and the corresponding field of the image, and replace the pixel in the center of the field with the output of the inner product operation
- Key point:
  - Different convolutions extract different types of low-level "features" from an image
  - All that we need to vary to generate these different features is the weights of F

How do we convolve an image with an ANN?

Note that the parameters in the matrix defining the convolution are **tied** across all places that it is used

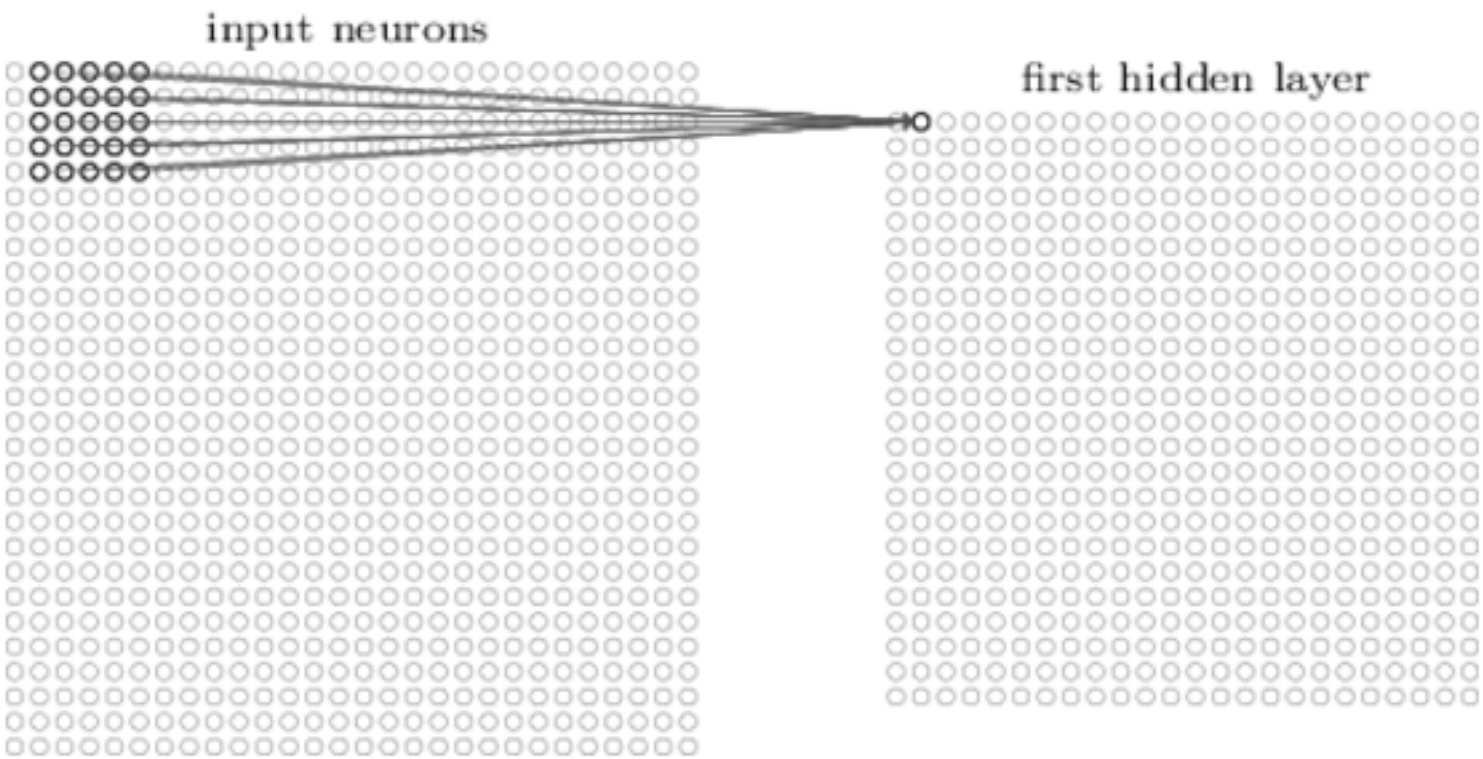

How do we do many convolutions of an image with an ANN?

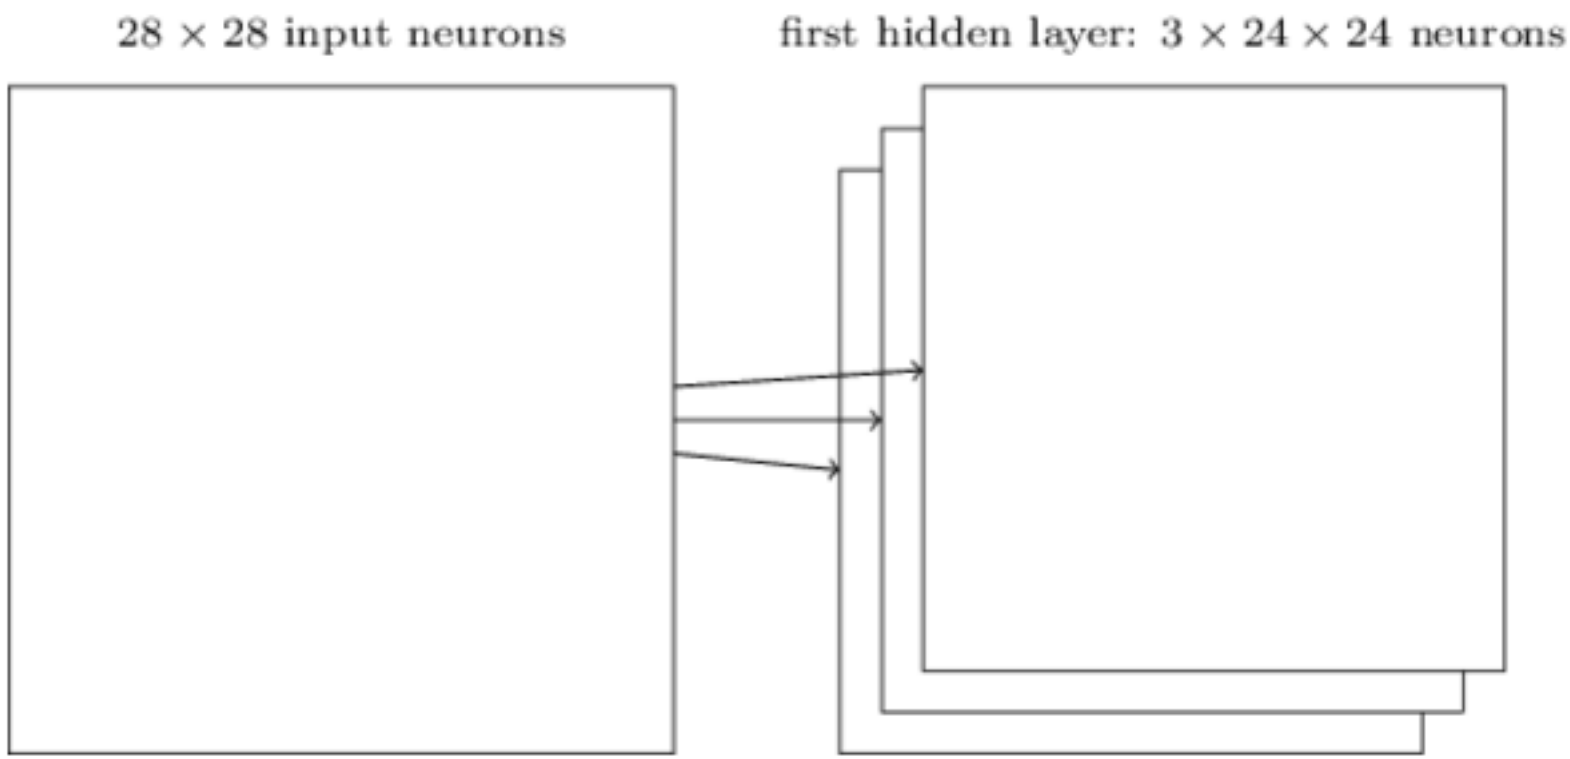

Example: 6 convolutions of a digit

http://scs.ryerson.ca/~aharley/vis/conv/

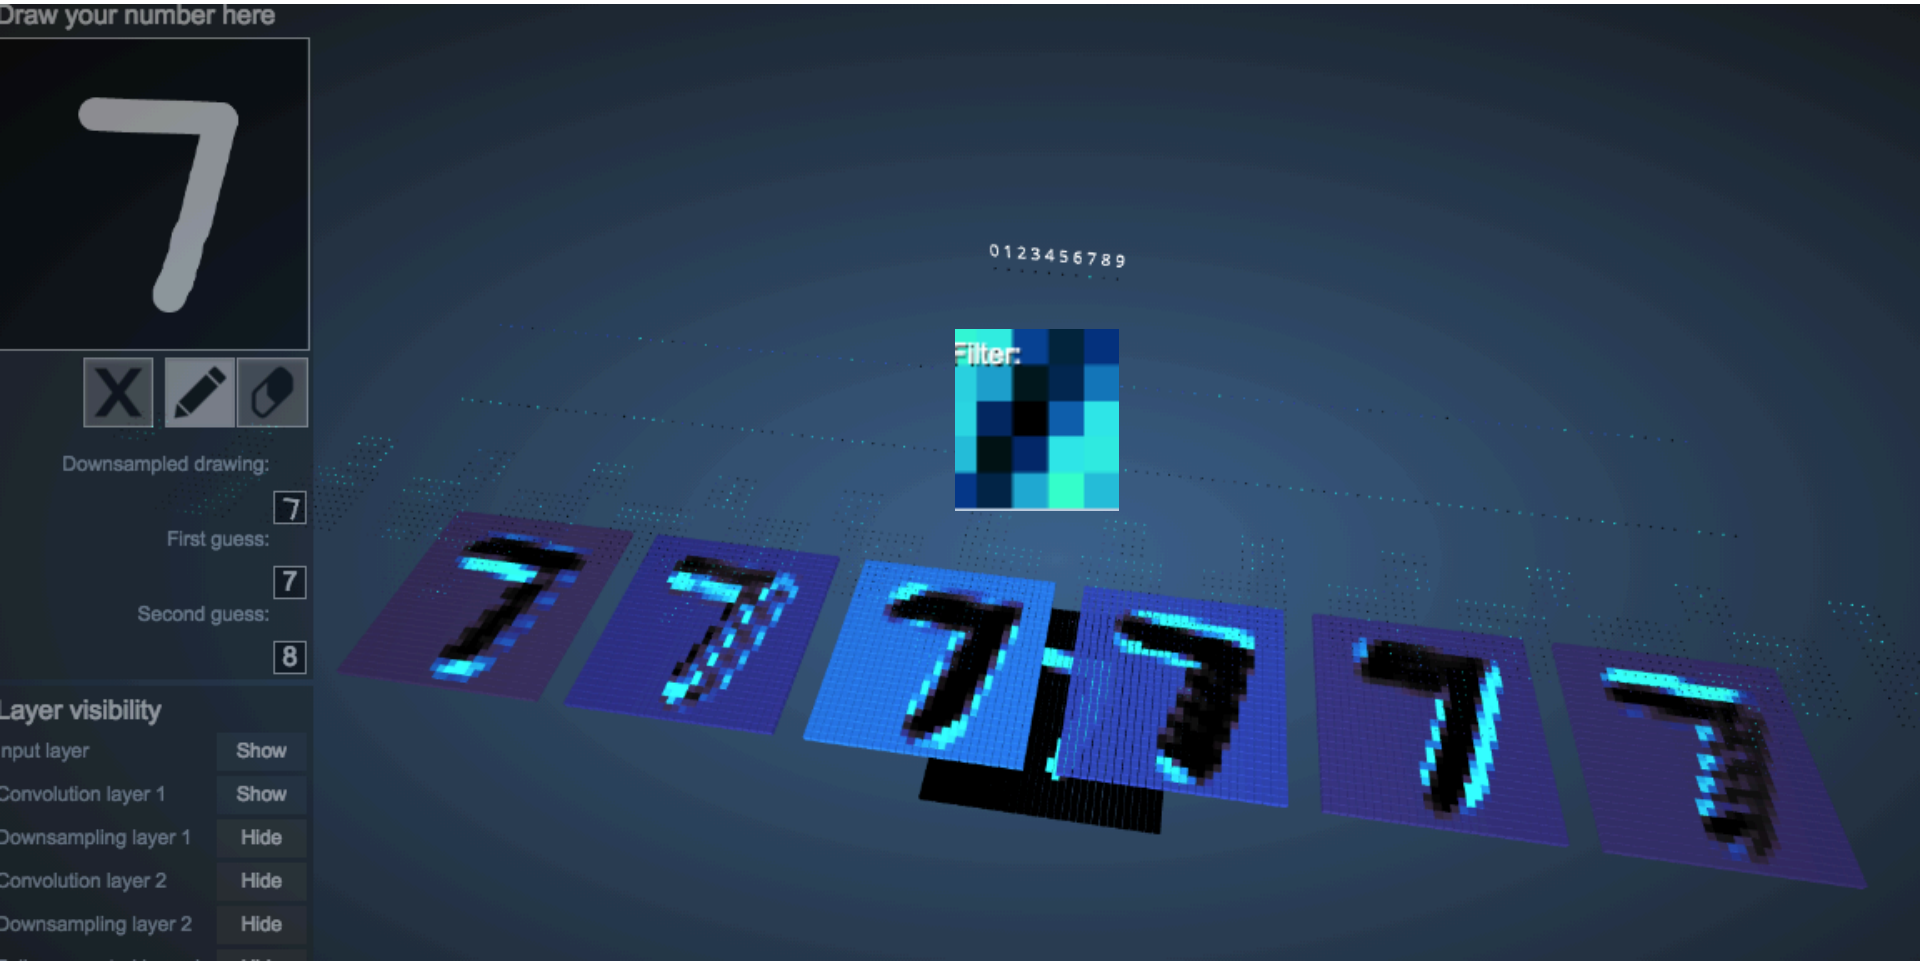

CNNs typically alternate convolutions, non-linearity, and then downsampling

Downsampling is usually averaging or (more common in recent CNNs) max-pooling

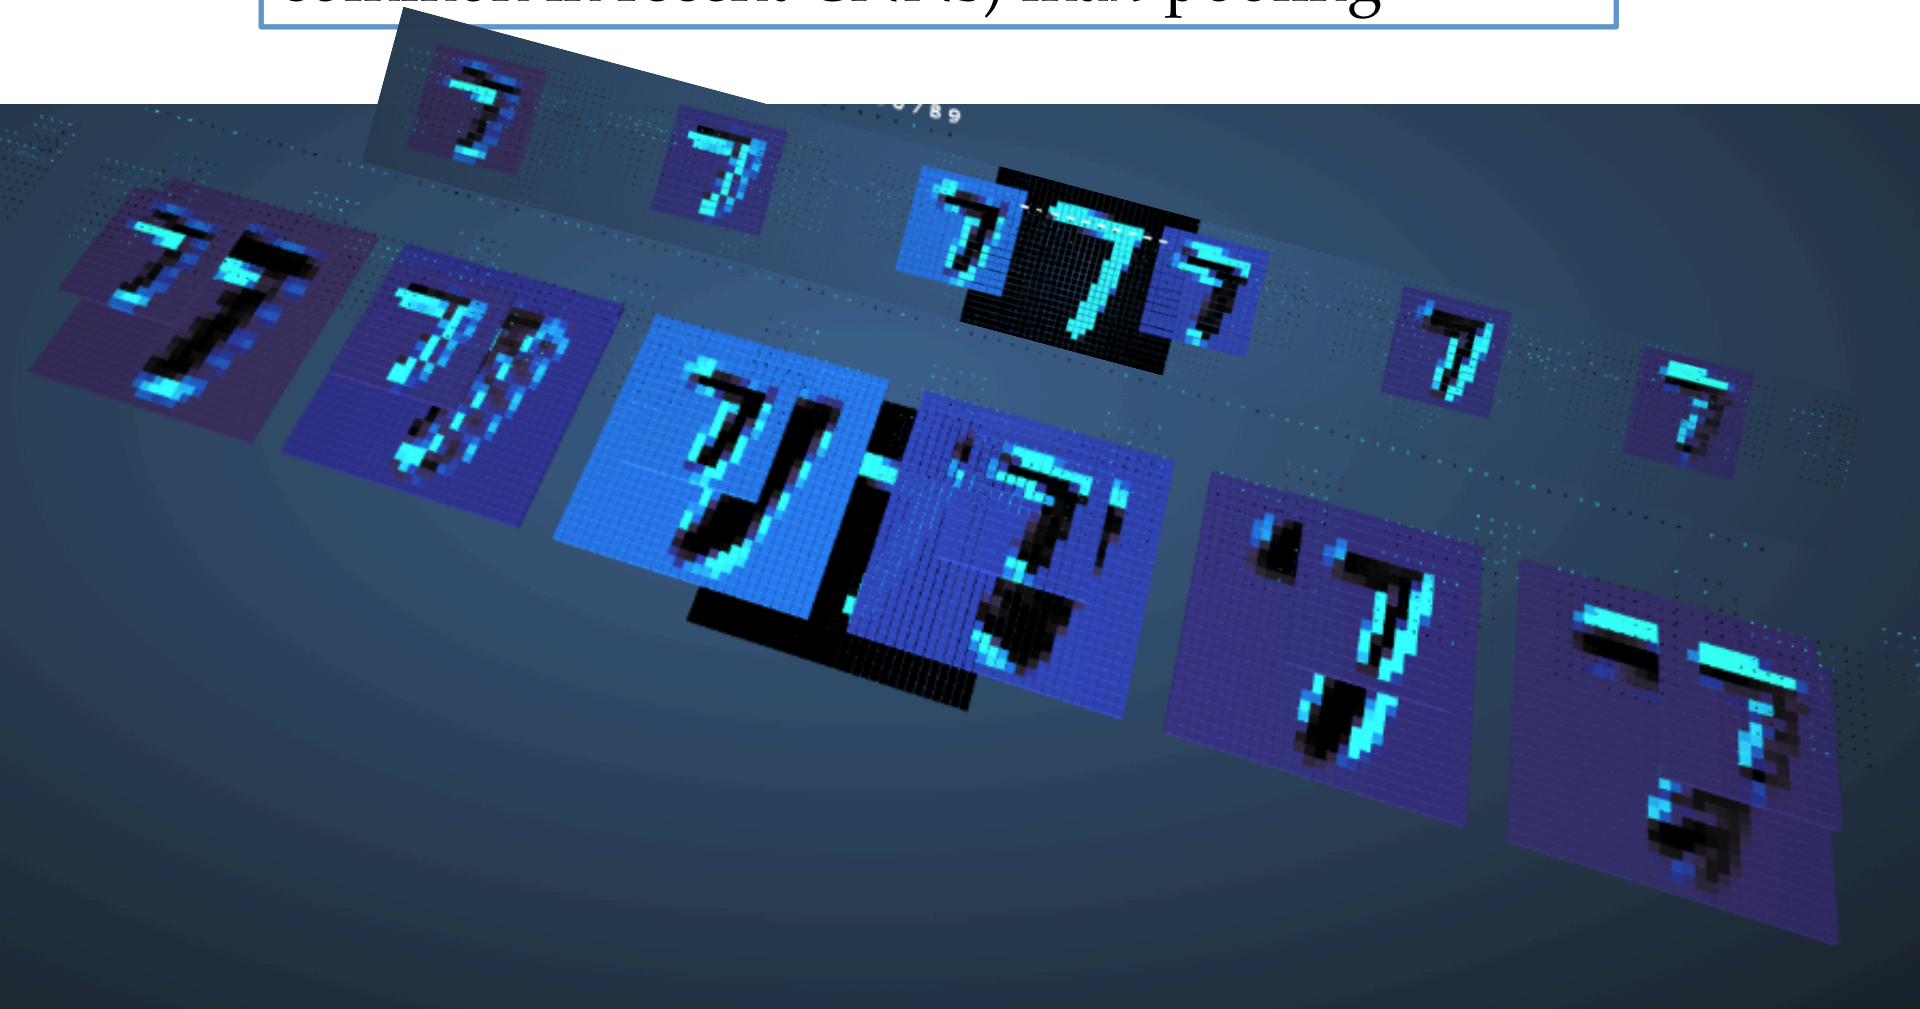

Why do max-pooling?

- Saves space
- Reduces overfitting?
- Because I'm going to add *more* convolutions after it!
  - Allows the short-range convolutions to extend over larger subfields of the images
    - So we can spot larger objects
    - Eg, a long horizontal line, or a corner, or …

PROC. OF THE IEEE, NOVEMBER 1998 7

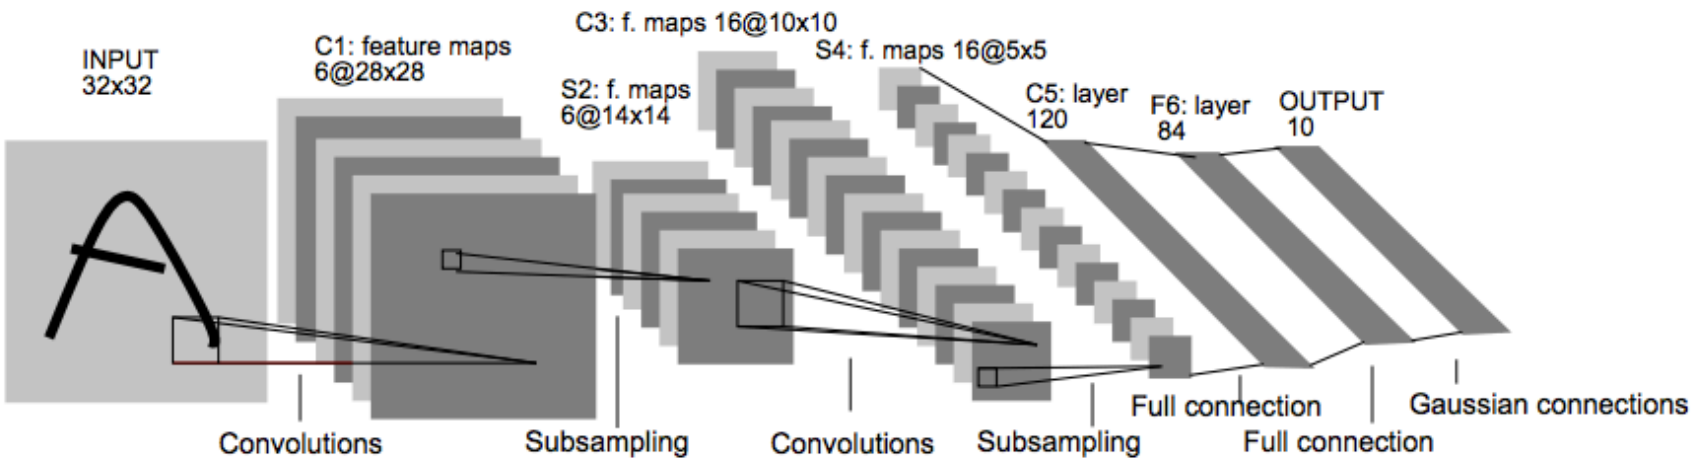

Fig. 2. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

Another CNN visualization

https://cs.stanford.edu/people/karpathy/convnetjs/demo/mnist.html

Network Visualization

input (24x24x1)
max activation: 0.99607, min: 0

Activations:

filter size 5x5x1, stride 1
max activation: 2.96187, min: -5.43735
max gradient: 0.00068, min: -0.00102
parameters: 8x5x5x1+8 = 208

Activation Gradients:

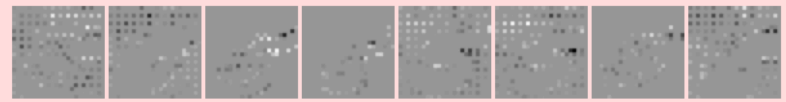

Weights:

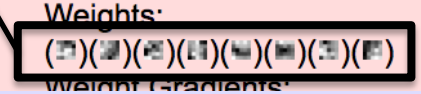

Activations:

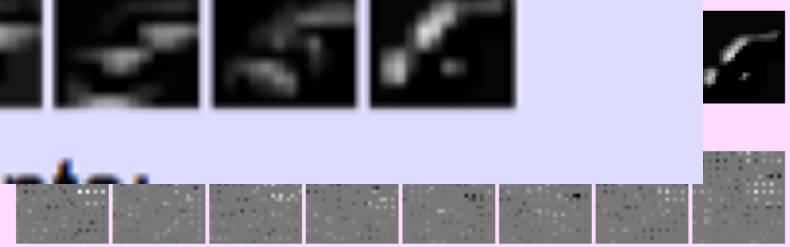

pool (12x12x8)
pooling size 2x2, stride 2
max activation: 2.96187, min: 0
max gradient: 0.00106, min: -0.00102

Activations:

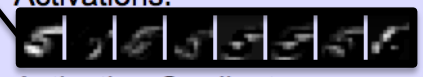

Activation Gradients:

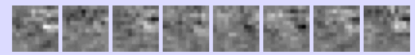

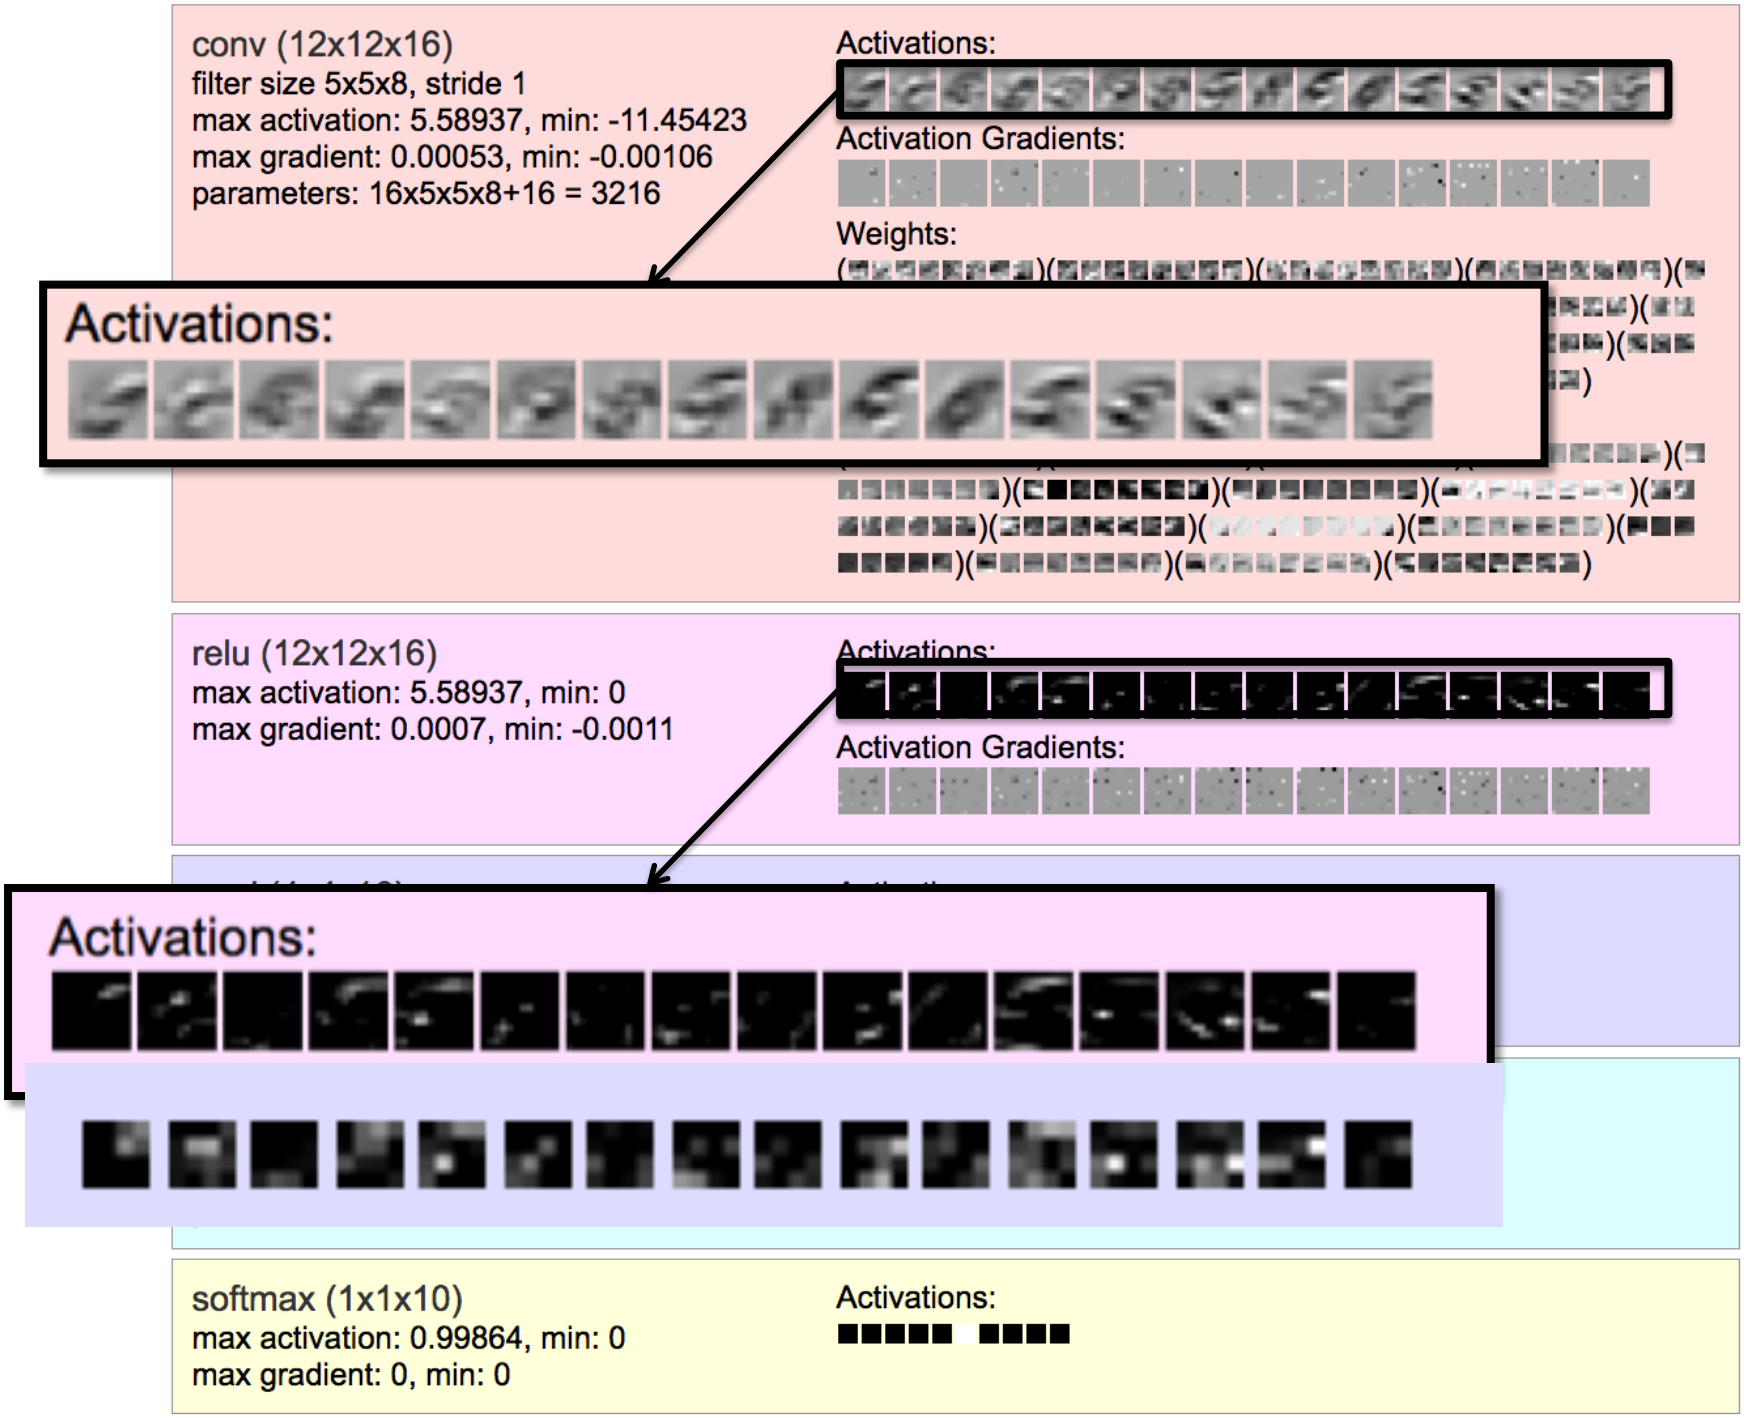
conv (12x12x16)
filter size 5x5x8, stride 1
max activation: 5.58937, min: -11.45423
max gradient: 0.00053, min: -0.00106
parameters: 16x5x5x8+16 = 3216
Activations:
Activation Gradients:
Weights:
Activations:
relu (12x12x16)
max activation: 5.58937, min: 0
max gradient: 0.0007, min: -0.0011
Activations:
Activation Gradients:
Activations:
softmax (1x1x10)
max activation: 0.99864, min: 0
max gradient: 0, min: 0
Activations:

Why do max-pooling?

- Saves space
- Reduces overfitting?
- Because I'm going to add *more* convolutions after it!
  - Allows the short-range convolutions to extend over larger subfields of the images
    - So we can spot larger objects
    - Eg, a long horizontal line, or a corner, or …
- At some point the feature maps start to get very sparse and blobby – they are indicators of some semantic property, not a recognizable transformation of the image
- Then just use them as features in a "normal" ANN

Why do max-pooling?

- Saves space
- Reduces overfitting?
- Because I'm going to add *more* convolutions after it!
  - Allows the short-range convolutions to extend over larger subfields of the images
    - So we can spot larger objects
    - Eg, a long horizontal line, or a corner, or …

PROC. OF THE IEEE, NOVEMBER 1998 7

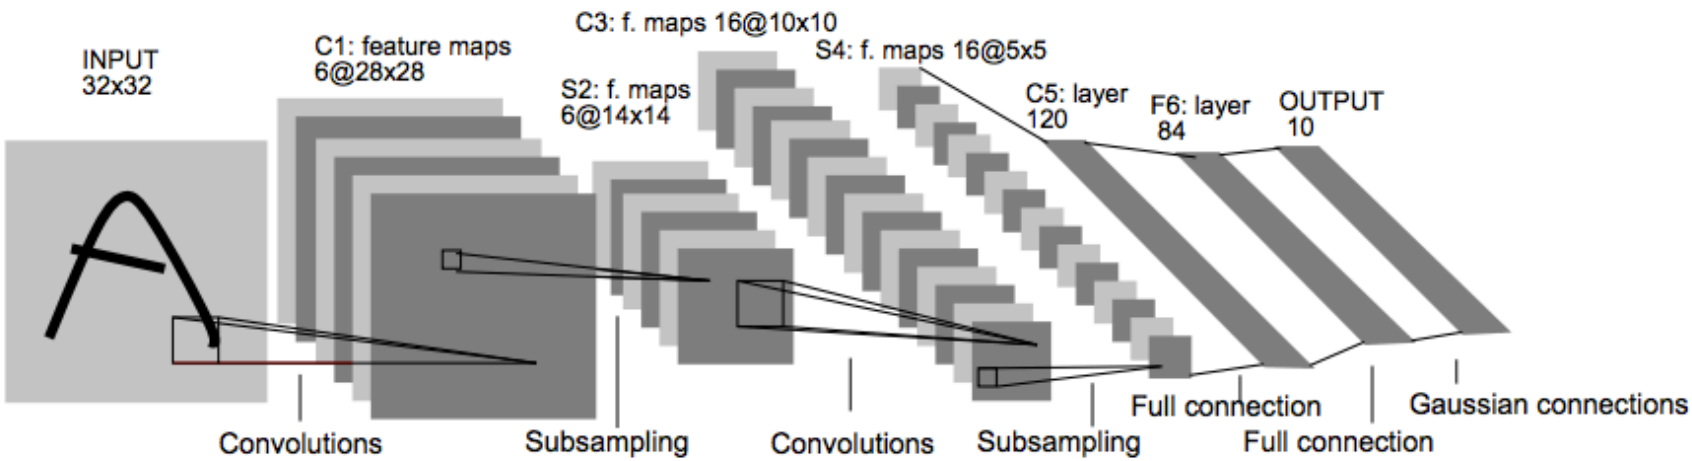

Fig. 2. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

Alternating convolution and downsampling

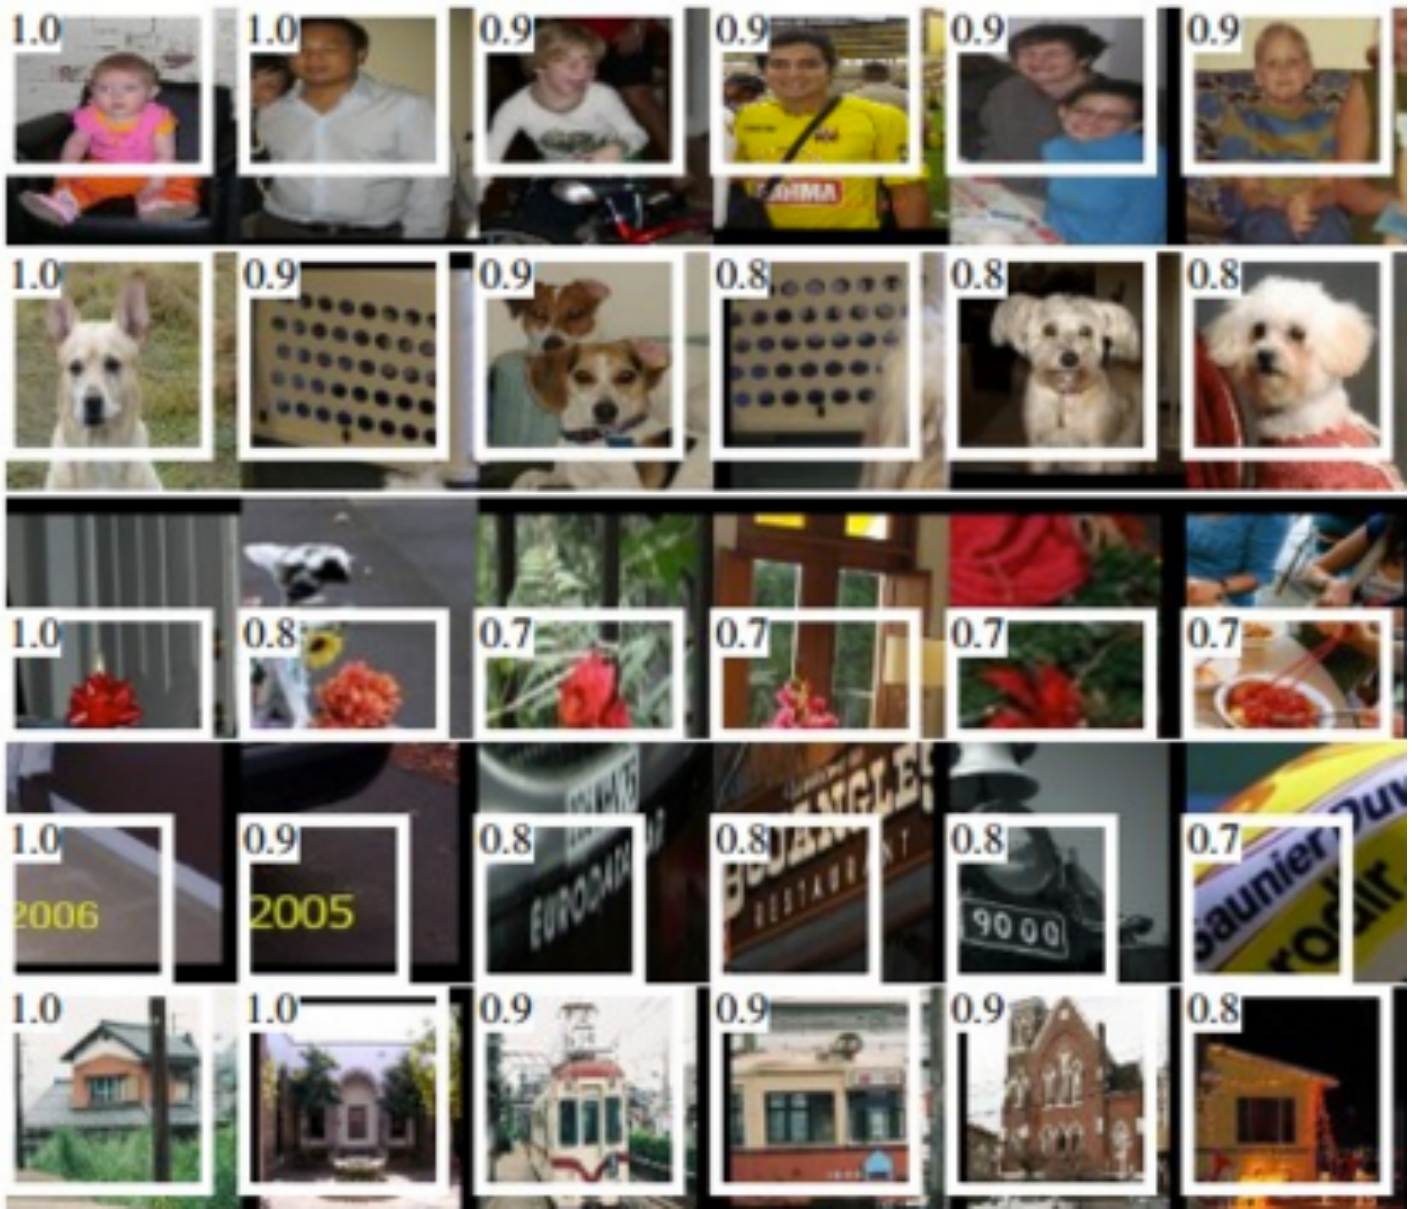

5 layers up

The subfield in a large dataset that gives the strongest output for a neuron