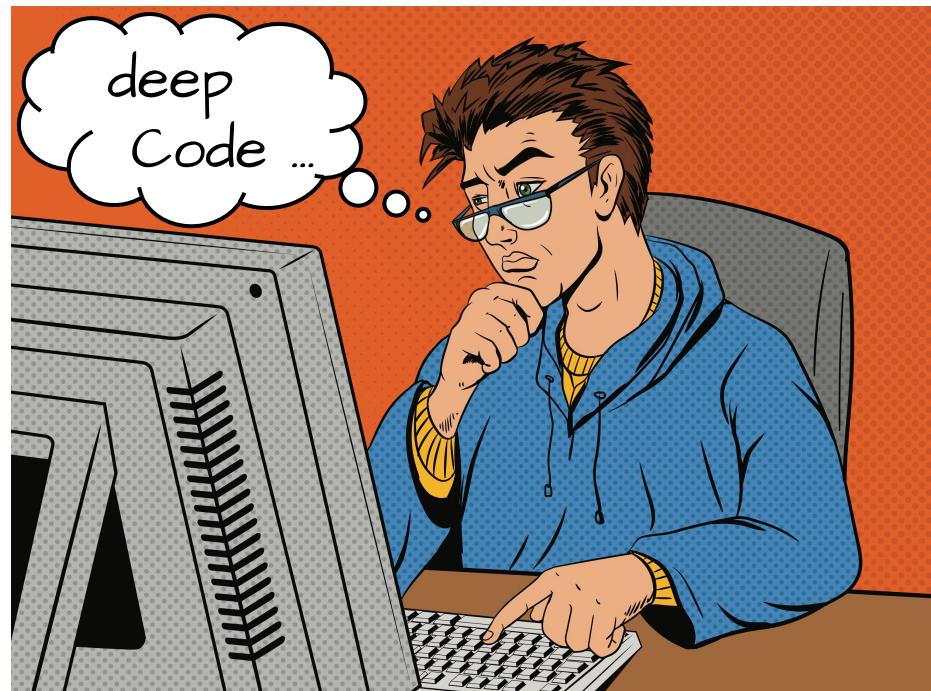
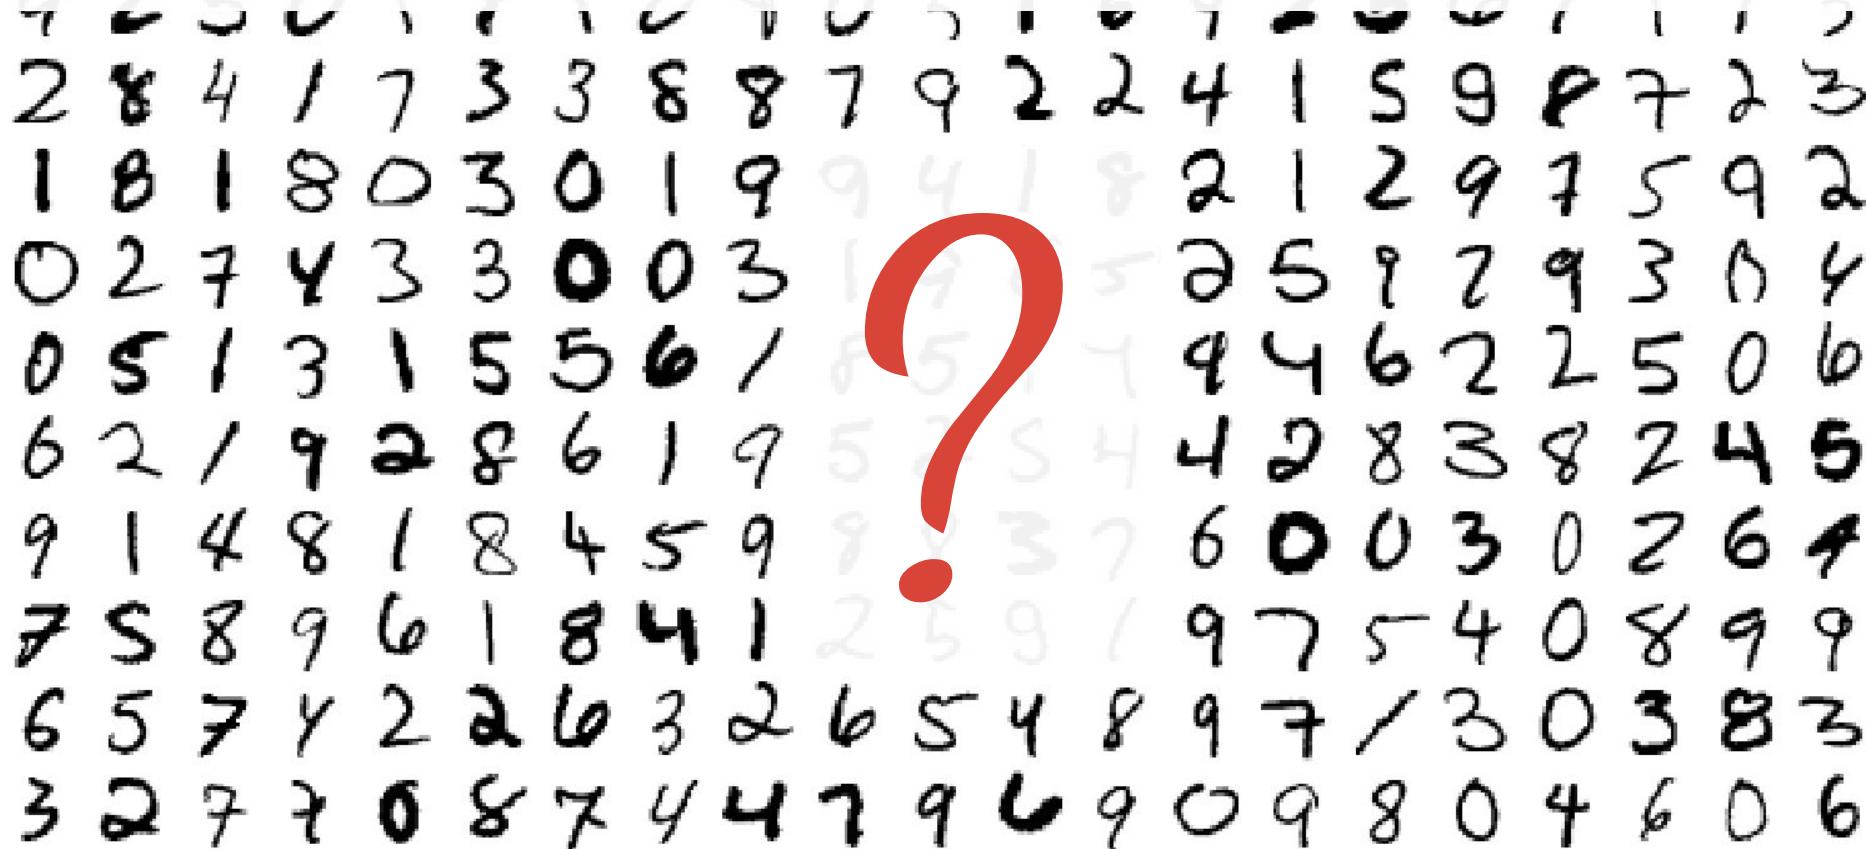


# >TensorFlow and deep learning\_

without a PhD

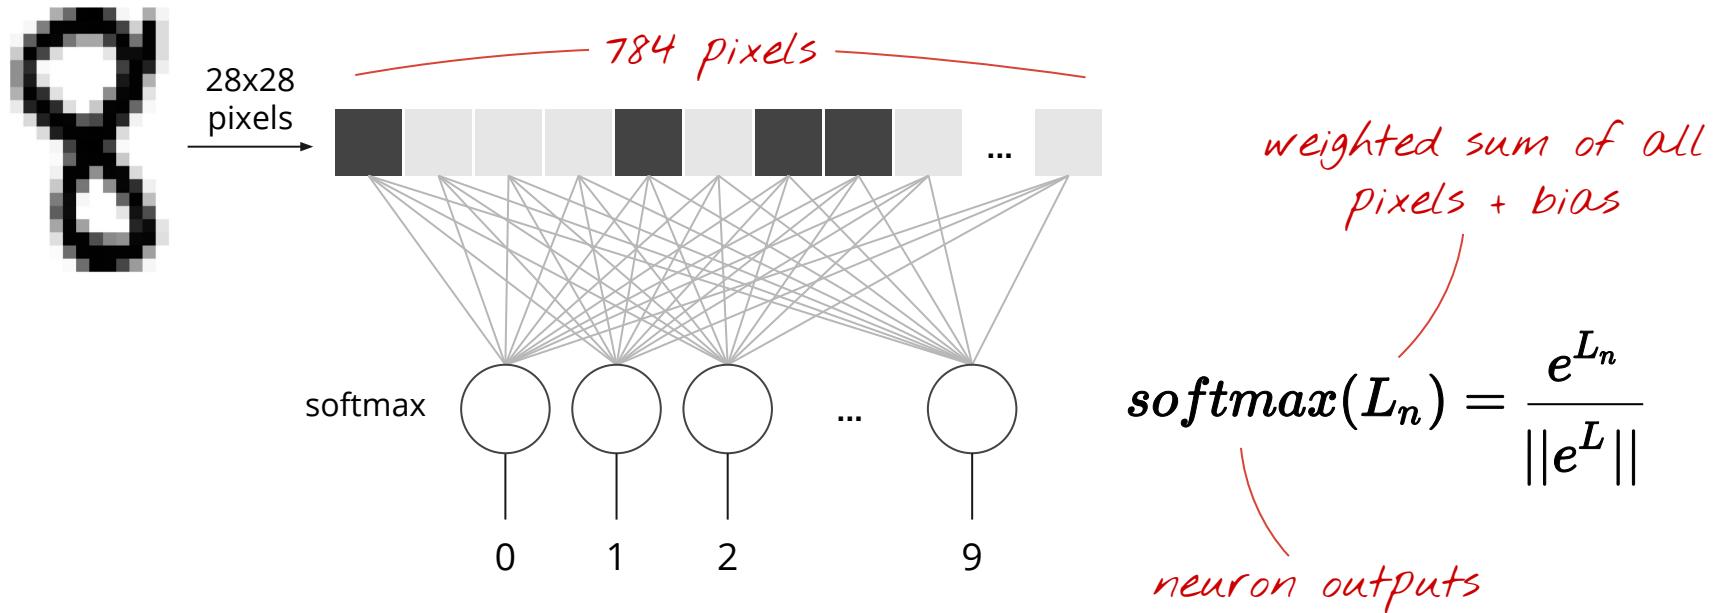


# Hello World: handwritten digits classification - MNIST

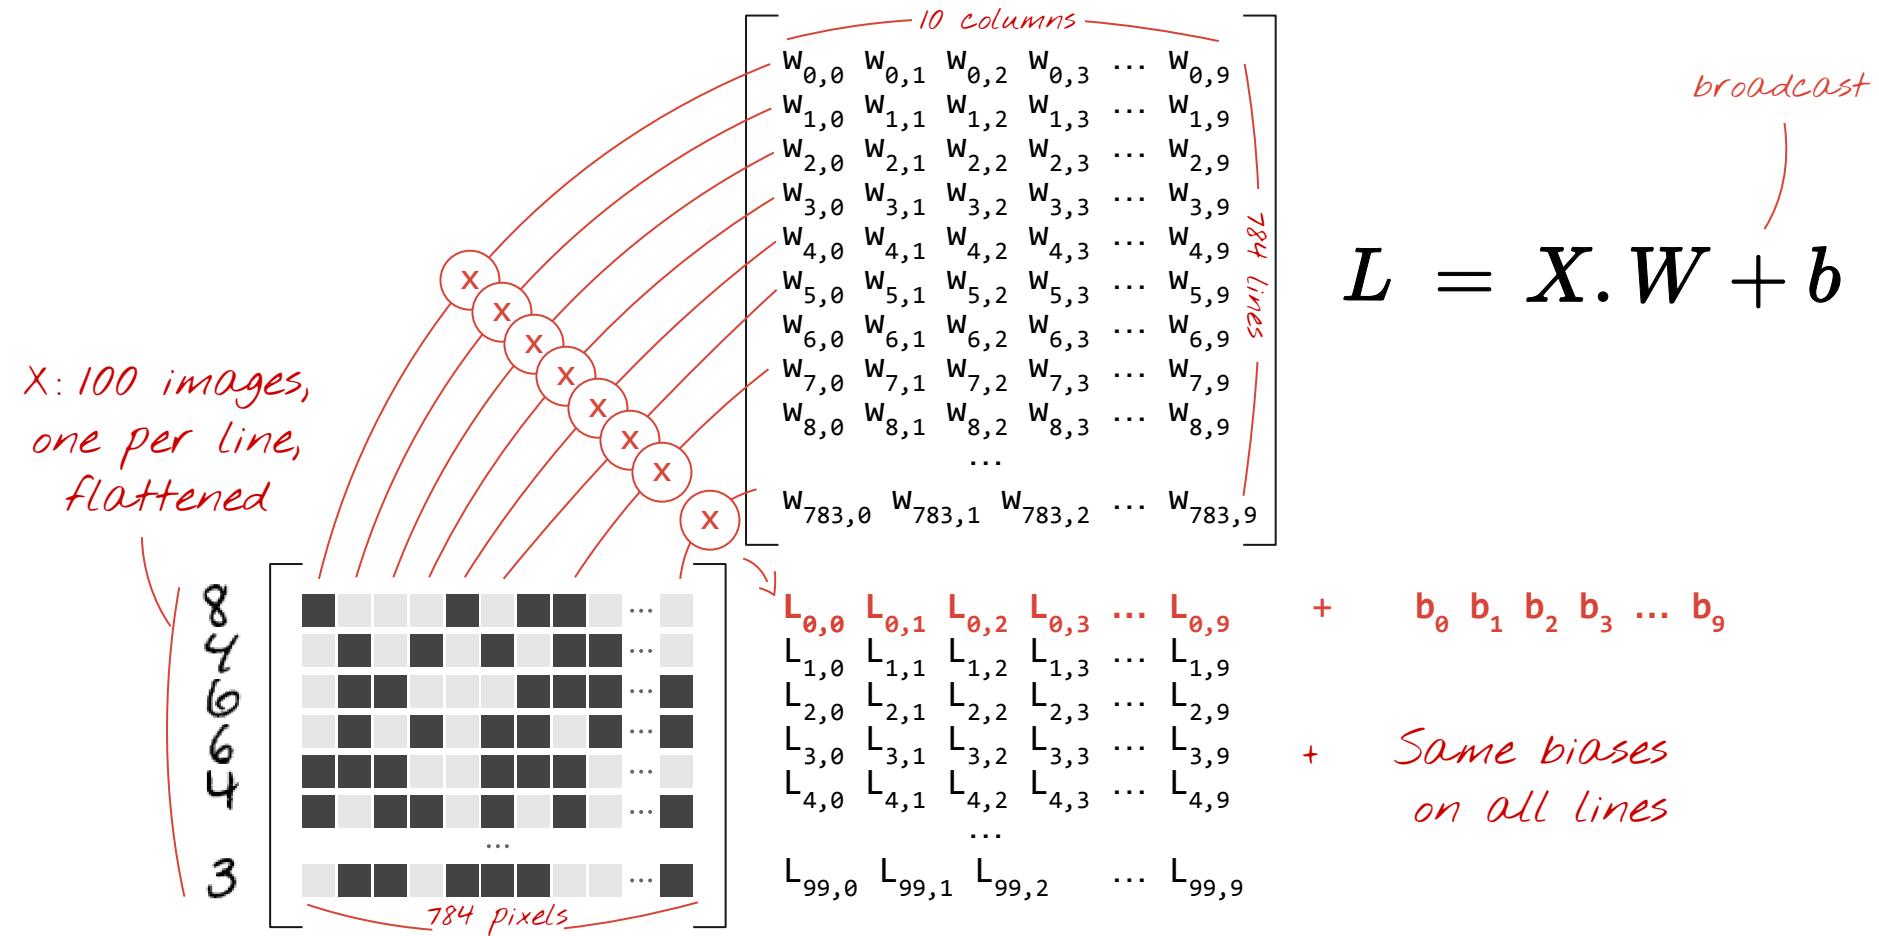


MNIST = Mixed National Institute of Standards and Technology - Download the dataset at <http://yann.lecun.com/exdb/mnist/>

# Very simple model: softmax classification



# In matrix notation, 100 images at a time



# Softmax, on a batch of images

Predictions

$Y[100, 10]$

$$Y = \text{softmax}(X \cdot W + b)$$

applied line  
by line

tensor shapes in [ ]

Images

$X[100, 748]$

Weights

$W[748, 10]$

Biases

$b[10]$

matrix multiply

broadcast  
on all lines

# Now in TensorFlow (Python)

tensor shapes:  $X[100, 748]$     $W[748, 10]$     $b[10]$

$Y = \text{tf.nn.softmax}(\text{tf.matmul}(X, W) + b)$

*matrix multiply*      *broadcast  
on all lines*

# Success ?

|   |   |   |   |   |   |   |   |   |   |
|---|---|---|---|---|---|---|---|---|---|
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 |

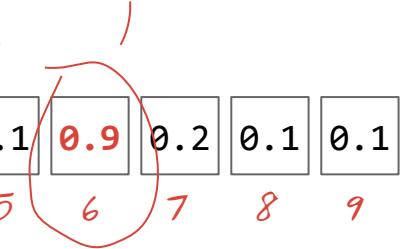
actual probabilities, "one-hot" encoded



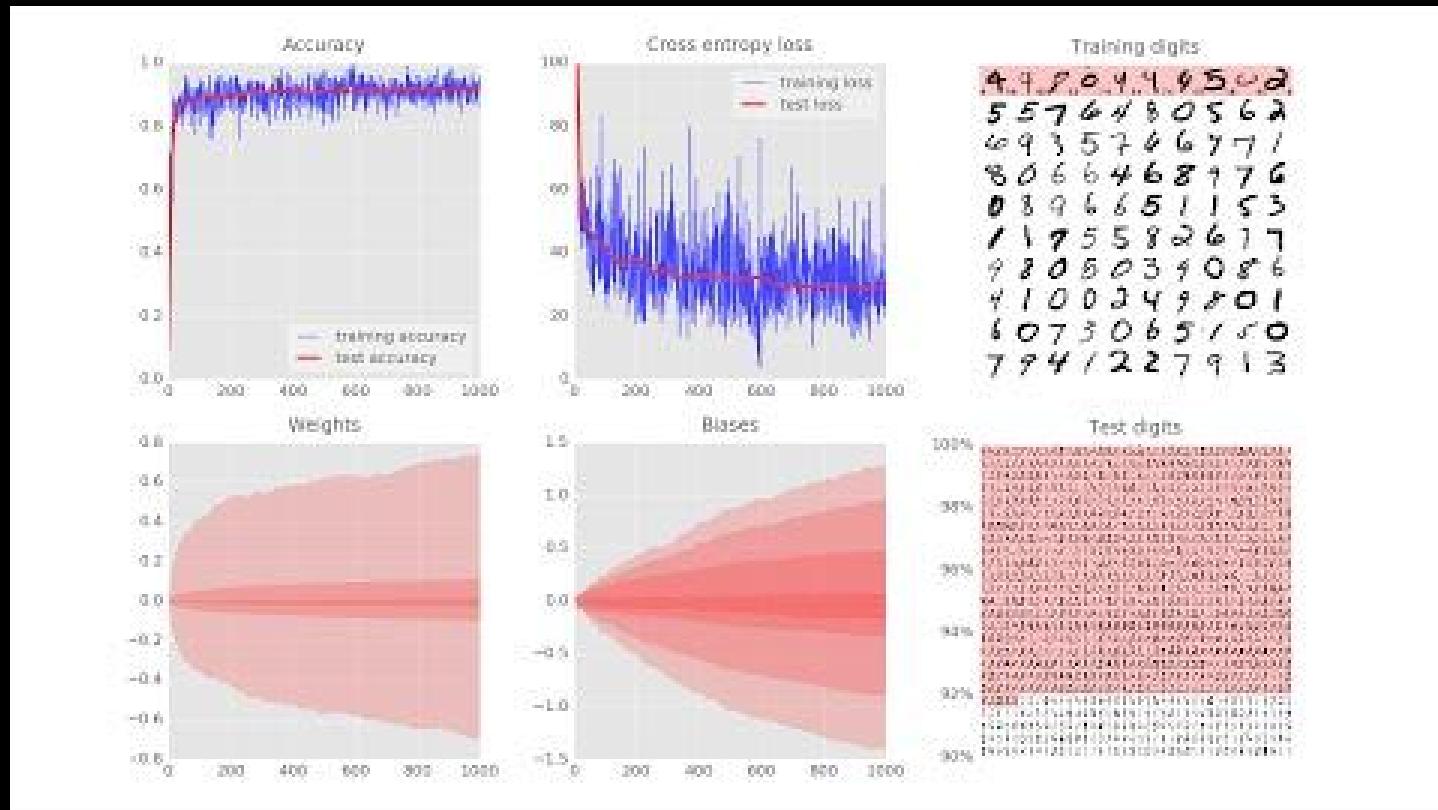
Cross entropy:  $-\sum Y'_i \cdot \log(Y_i)$

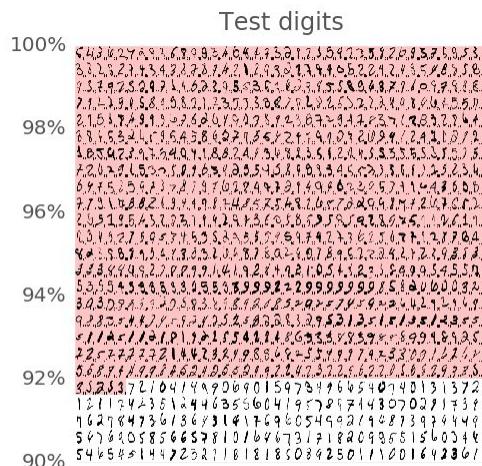
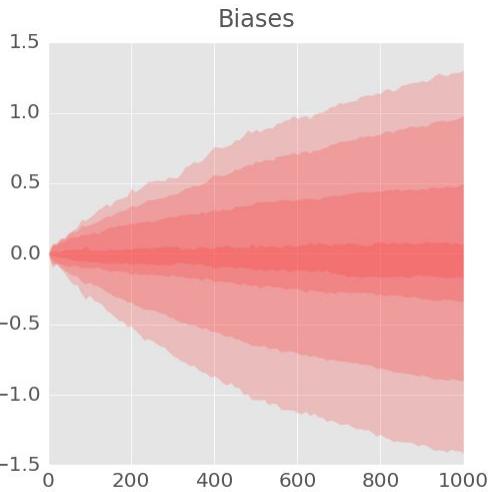
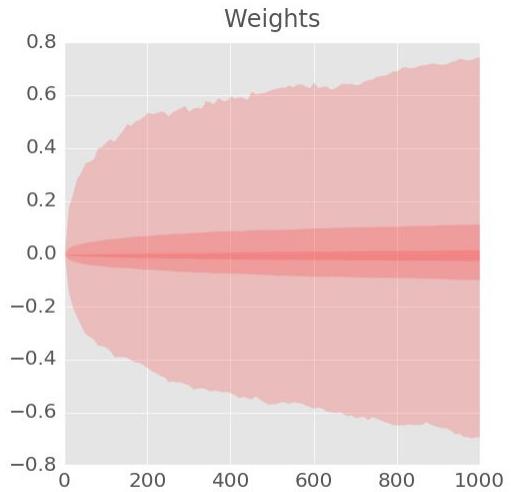
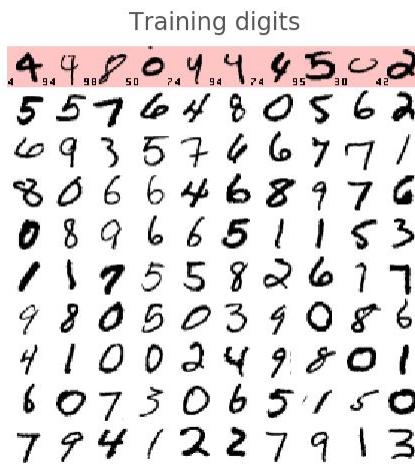
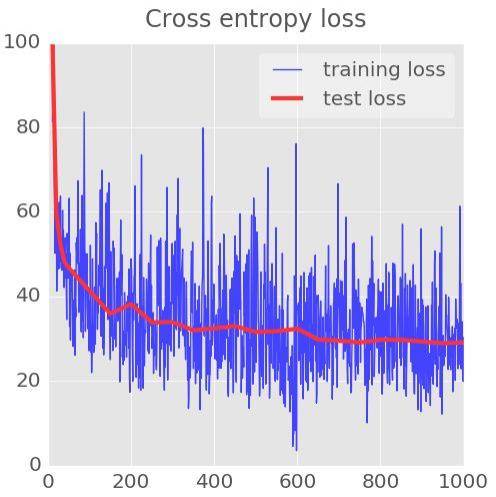
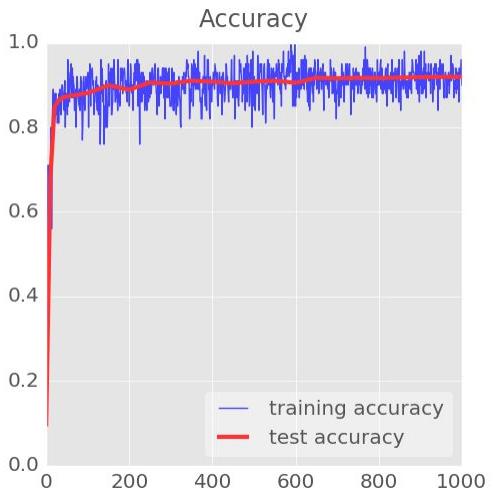
)  
computed probabilities  
this is a "6"

|     |     |     |     |     |     |     |     |     |     |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| 0.1 | 0.2 | 0.1 | 0.3 | 0.2 | 0.1 | 0.9 | 0.2 | 0.1 | 0.1 |
| 0   | 1   | 2   | 3   | 4   | 5   | 6   | 7   | 8   | 9   |



# Demo





92%

# TensorFlow - initialisation

```
import tensorflow as tf
```

this will become the batch size, 100  
|

```
x = tf.placeholder(tf.float32, [None, 28, 28, 1])
```

```
w = tf.Variable(tf.zeros([784, 10]))
```

```
b = tf.Variable(tf.zeros([10]))
```

|  
28 x 28 grayscale images

```
init = tf.initialize_all_variables()
```

Training = computing variables W and b

# TensorFlow - success metrics

```
# model  
Y = tf.nn.softmax(tf.matmul(tf.reshape(X, [-1, 784]), W) + b)  
# placeholder for correct answers  
Y_ = tf.placeholder(tf.float32, [None, 10])  
  
# loss function  
cross_entropy = -tf.reduce_sum(Y_ * tf.log(Y))  
  
# % of correct answers found in batch  
is_correct = tf.equal(tf.argmax(Y, 1), tf.argmax(Y_, 1))  
accuracy = tf.reduce_mean(tf.cast(is_correct, tf.float32))
```

*flattening images*

/

*"one-hot" encoded*

*"one-hot" decoding*

# TensorFlow - training

```
optimizer = tf.train.GradientDescentOptimizer(0.003)  
train_step = optimizer.minimize(cross_entropy)
```

*learning rate*



*loss function*



# TensorFlow - run !

```
sess = tf.Session()  
sess.run(init)  
  
for i in range(1000):  
    # Load batch of images and correct answers  
    batch_X, batch_Y = mnist.train.next_batch(100)  
    train_data={X: batch_X, Y_: batch_Y}  
  
    # train  
    sess.run(train_step, feed_dict=train_data)
```

running a Tensorflow computation, feeding placeholders

```
# success ?  
a,c = sess.run([accuracy, cross_entropy], feed_dict=train_data)  
  
# success on test data ?  
test_data={X: mnist.test.images, Y_: mnist.test.labels}  
a,c = sess.run([accuracy, cross_entropy, It], feed=test_data)
```

Tip:

do this  
every 100  
iterations

# TensorFlow - full python code

```
import tensorflow as tf

X = tf.placeholder(tf.float32, [None, 28, 28, 1])
W = tf.Variable(tf.zeros([784, 10]))
b = tf.Variable(tf.zeros([10]))
init = tf.initialize_all_variables()

# model
Y = tf.nn.softmax(tf.matmul(tf.reshape(X, [-1, 784]), W) + b)

# placeholder for correct answers
Y_ = tf.placeholder(tf.float32, [None, 10])

# loss function
cross_entropy = -tf.reduce_sum(Y_ * tf.log(Y))

# % of correct answers found in batch
is_correct = tf.equal(tf.argmax(Y, 1), tf.argmax(Y_, 1))
accuracy = tf.reduce_mean(tf.cast(is_correct, tf.float32))
```

initialisation

model

success metrics

```
optimizer = tf.train.GradientDescentOptimizer(0.003)
train_step = optimizer.minimize(cross_entropy)
```

training step

```
sess = tf.Session()
sess.run(init)

for i in range(10000):
    # load batch of images and correct answers
    batch_X, batch_Y = mnist.train.next_batch(100)
    train_data = {X: batch_X, Y_: batch_Y}

    # train
    sess.run(train_step, feed_dict=train_data)

    # success ? add code to print it
    a, c = sess.run([accuracy, cross_entropy], feed=train_data)

    # success on test data ?
    test_data = {X: mnist.test.images, Y_: mnist.test.labels}
    a, c = sess.run([accuracy, cross_entropy], feed=test_data)
```

Run

# Cookbook

Softmax  
Cross-entropy  
Mini-batch

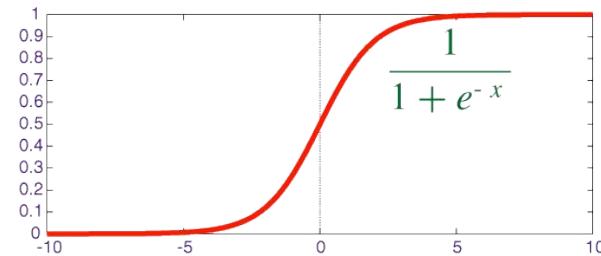
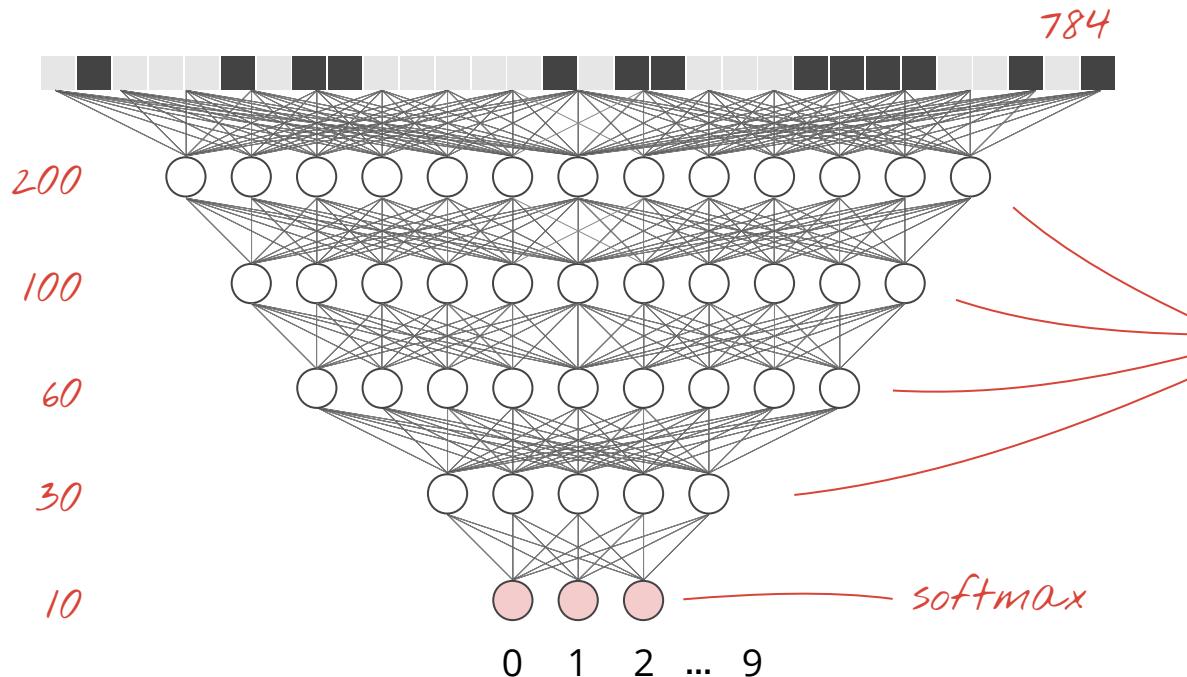




Go deep !

# Let's try 5 fully-connected layers !

overKill ↗



sigmoid function

softmax

# TensorFlow - initialisation

```
K = 200  
L = 100  
M = 60  
N = 30
```

*weights initialised  
with random values*

```
W1 = tf.Variable(tf.truncated_normal([28*28, K] ,stddev=0.1))  
B1 = tf.Variable(tf.zeros([K]))
```

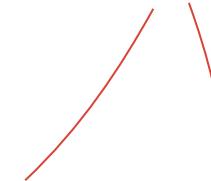
```
W2 = tf.Variable(tf.truncated_normal([K, L], stddev=0.1))  
B2 = tf.Variable(tf.zeros([L]))
```

```
W3 = tf.Variable(tf.truncated_normal([L, M], stddev=0.1))  
B3 = tf.Variable(tf.zeros([M]))  
W4 = tf.Variable(tf.truncated_normal([M, N], stddev=0.1))  
B4 = tf.Variable(tf.zeros([N]))  
W5 = tf.Variable(tf.truncated_normal([N, 10], stddev=0.1))  
B5 = tf.Variable(tf.zeros([10]))
```

# TensorFlow - the model

```
X = tf.reshape(X, [-1, 28*28])
```

*weights and biases*



```
Y1 = tf.nn.sigmoid(tf.matmul(X, W1) + B1)
```

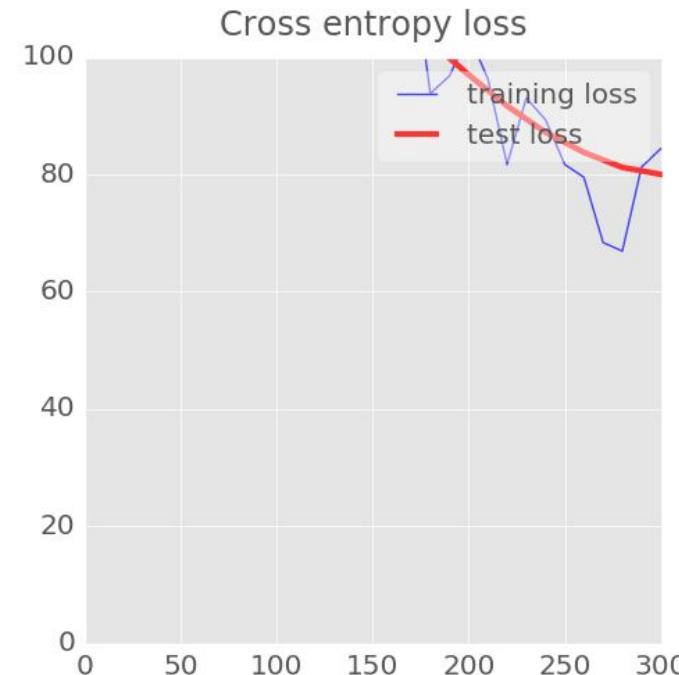
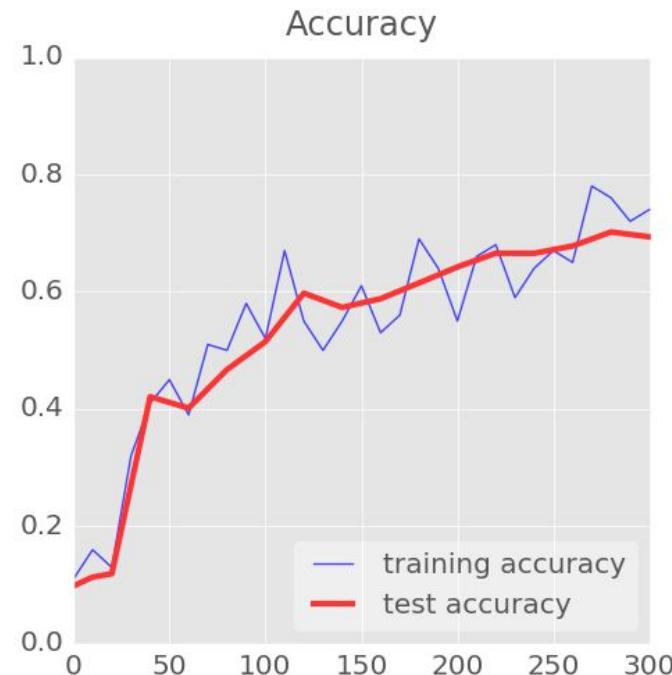
```
Y2 = tf.nn.sigmoid(tf.matmul(Y1, W2) + B2)
```

```
Y3 = tf.nn.sigmoid(tf.matmul(Y2, W3) + B3)
```

```
Y4 = tf.nn.sigmoid(tf.matmul(Y3, W4) + B4)
```

```
Y = tf.nn.softmax(tf.matmul(Y4, W5) + B5)
```

# Demo - slow start ?



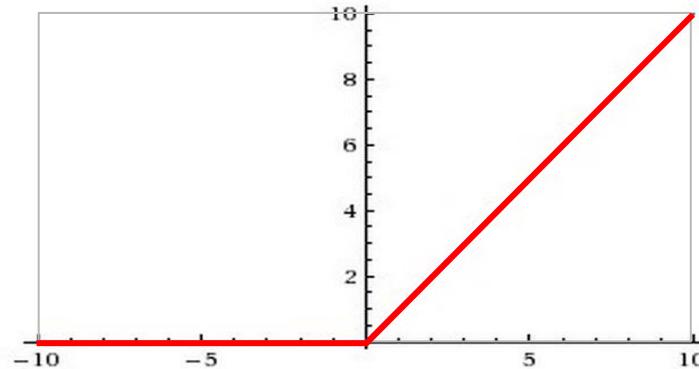
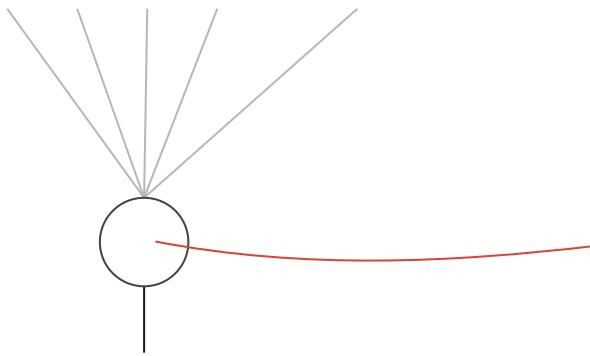


Relu !

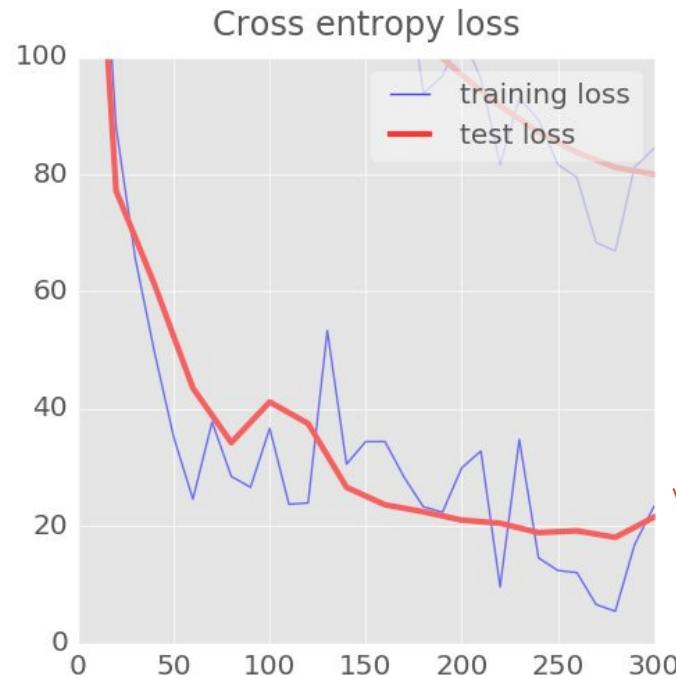


# RELU

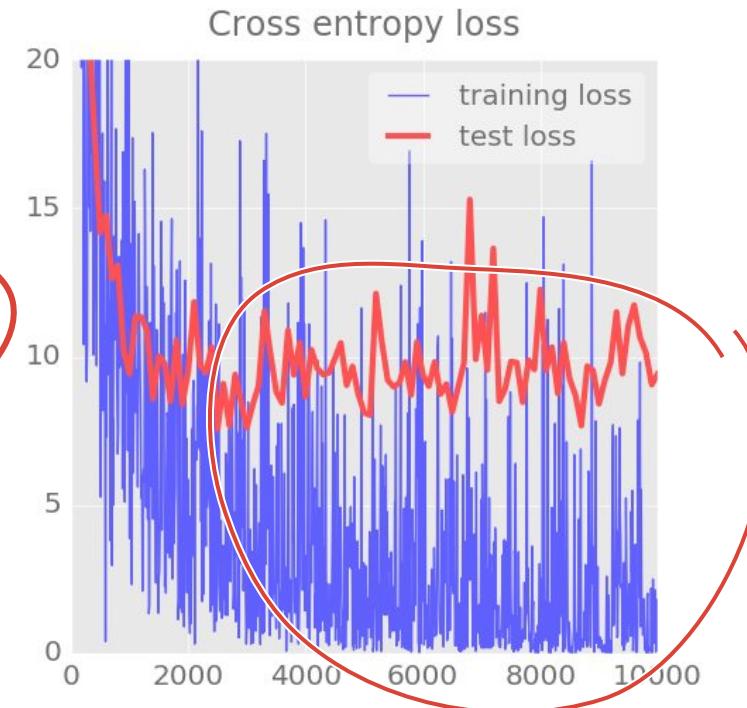
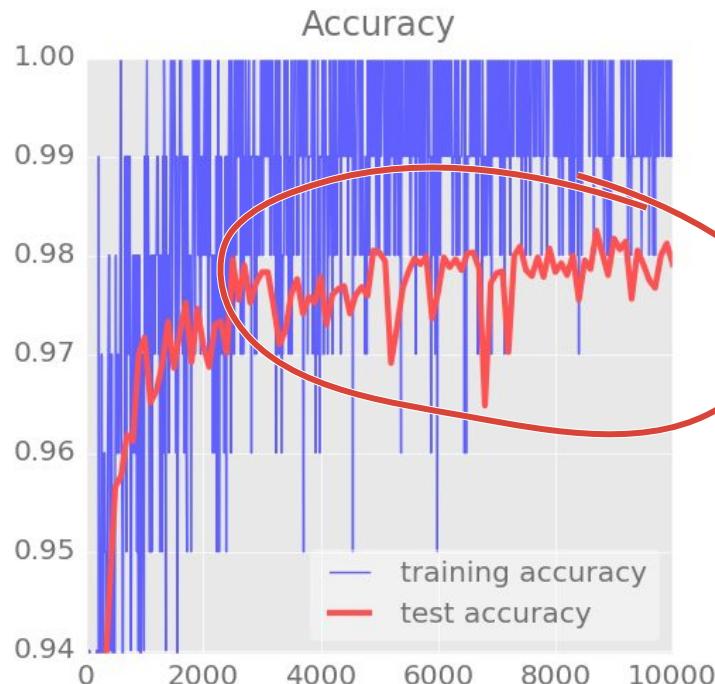
RELU = Rectified Linear Unit


$$Y = \text{tf.nn.relu}(\text{tf.matmul}(X, W) + b)$$

# RELU



# Demo - noisy accuracy curve ?



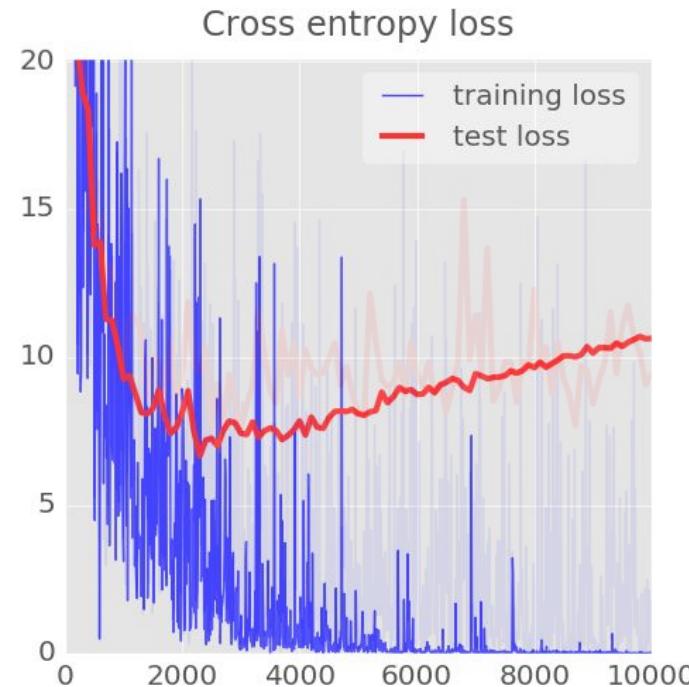
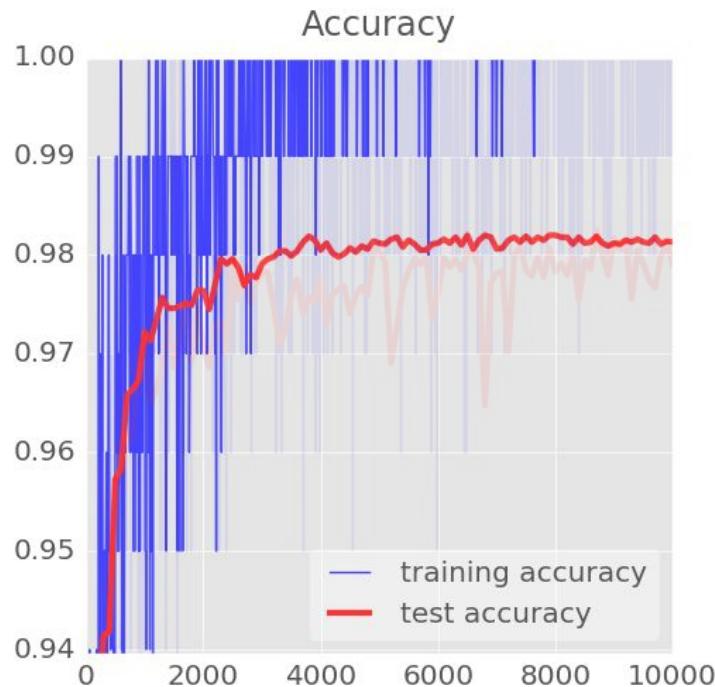
yuck!

Slow down...

Learning  
rate decay

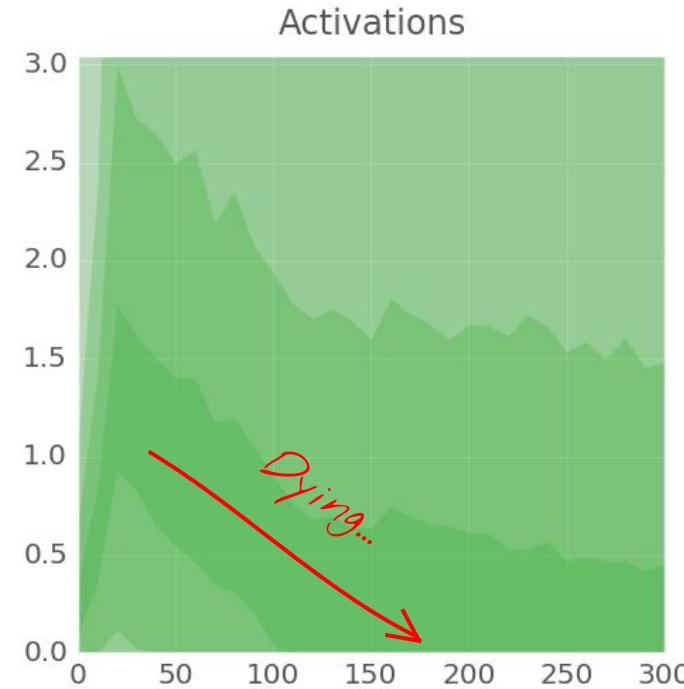
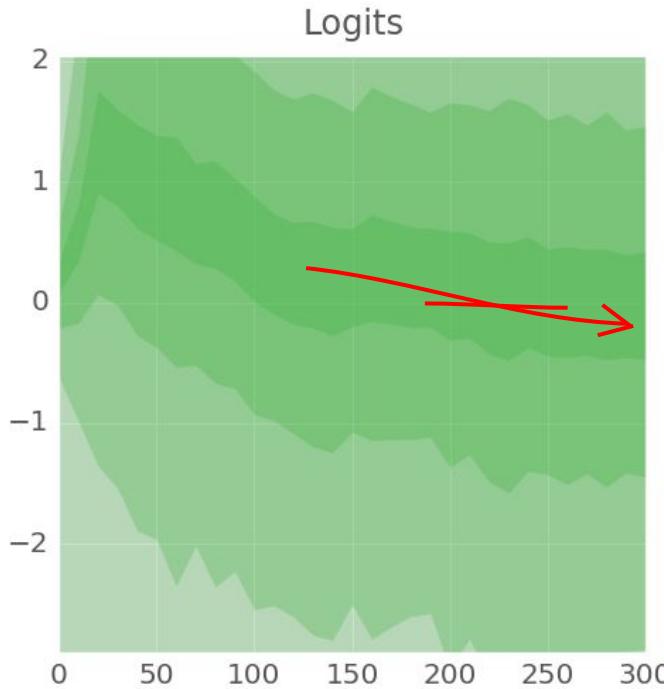


# Learning rate decay



Learning rate 0.003 at start then dropping exponentially to 0.0001

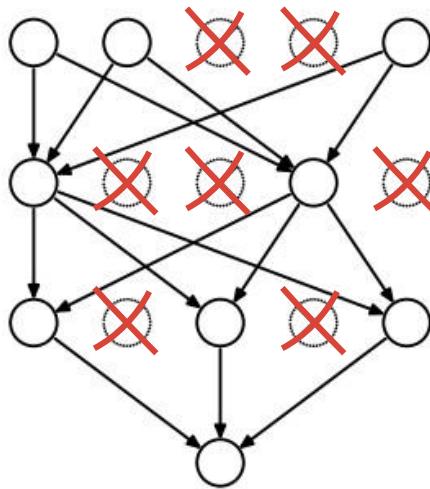
# Demo - dying neurons



Dropout



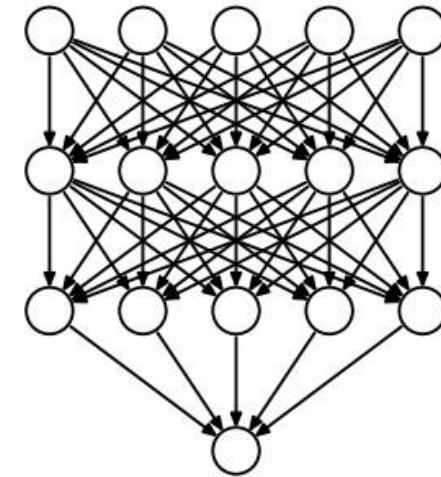
# Dropout



*TRAINING*  
 $pKeep=0.75$

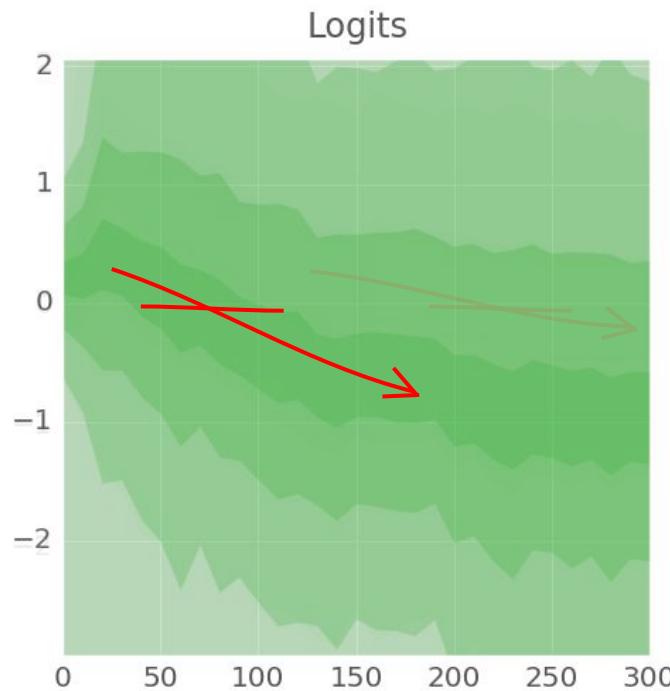
```
pkeep =  
tf.placeholder(tf.float32)
```

```
Yf = tf.nn.relu(tf.matmul(X, W) + B)  
Y = tf.nn.dropout(Yf, pkeep)
```



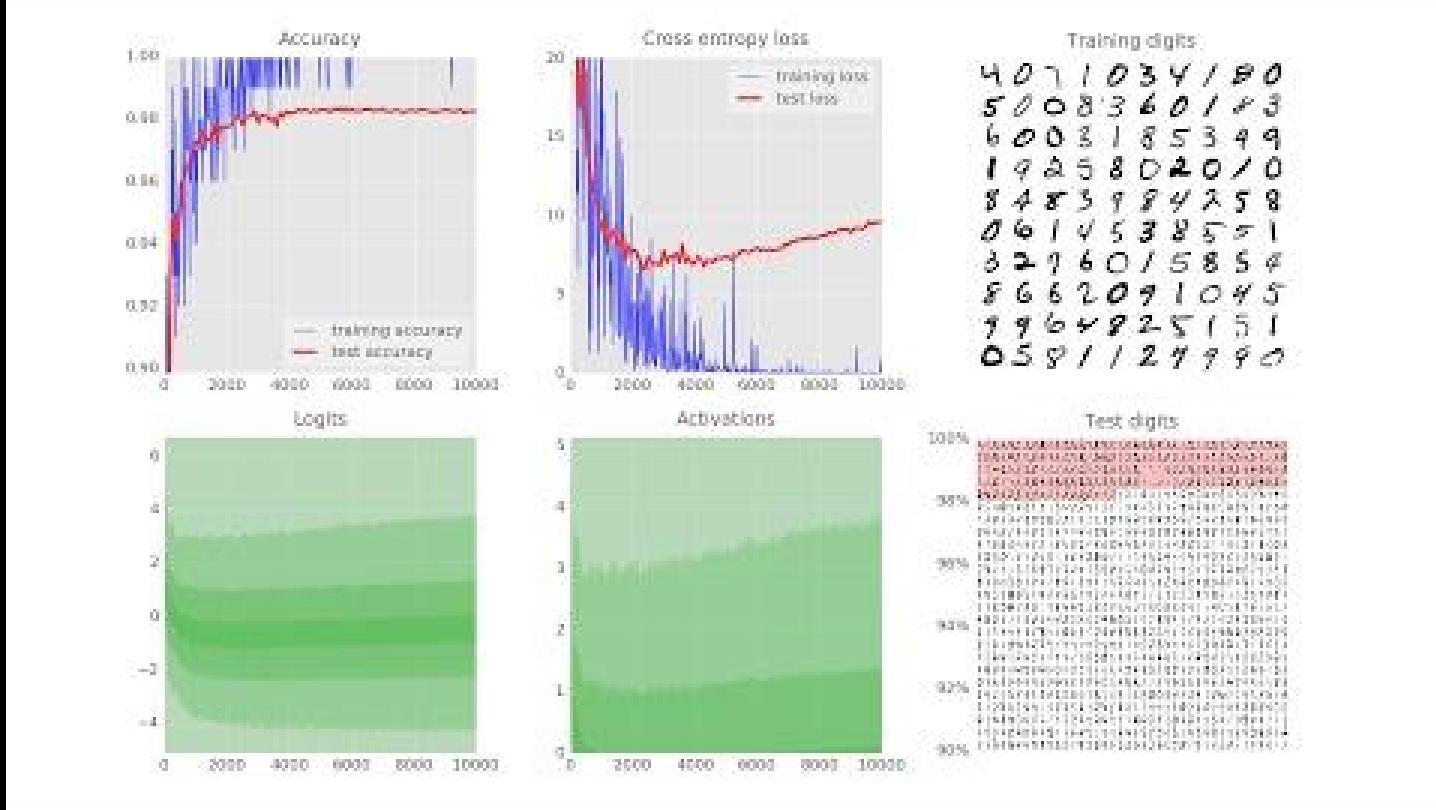
*EVALUATION*  
 $pKeep=1$

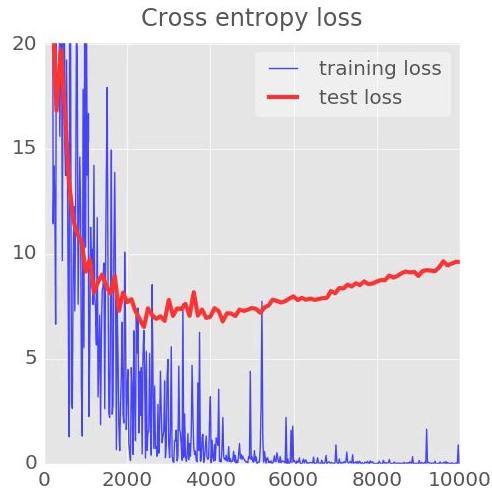
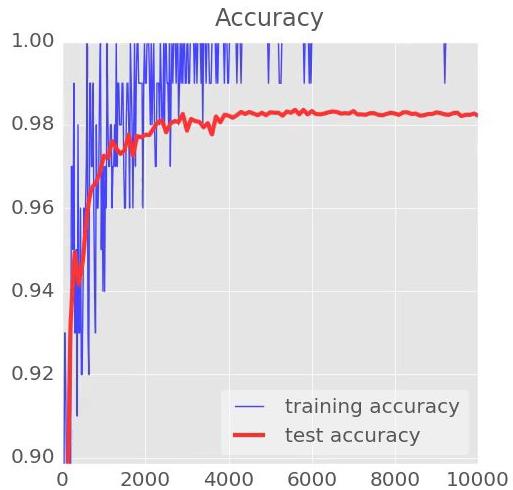
# Dropout



with dropout

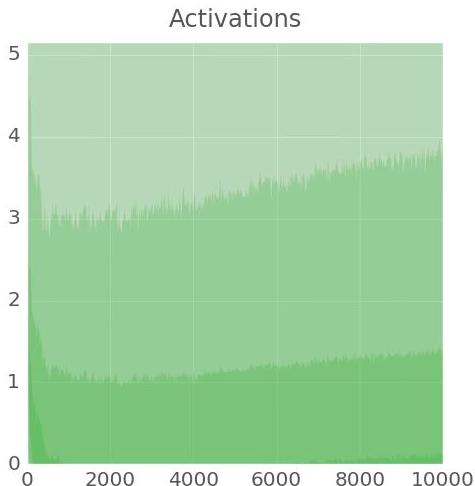
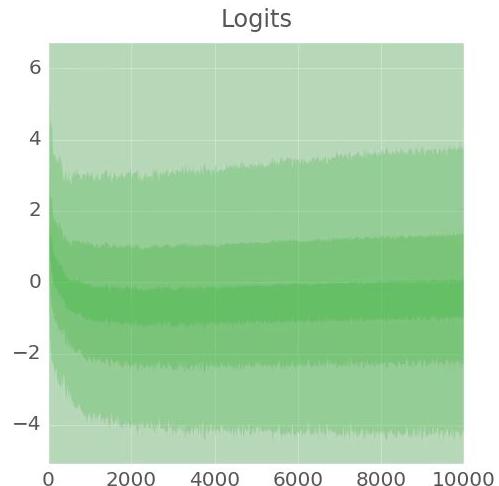
# Demo





Training digits

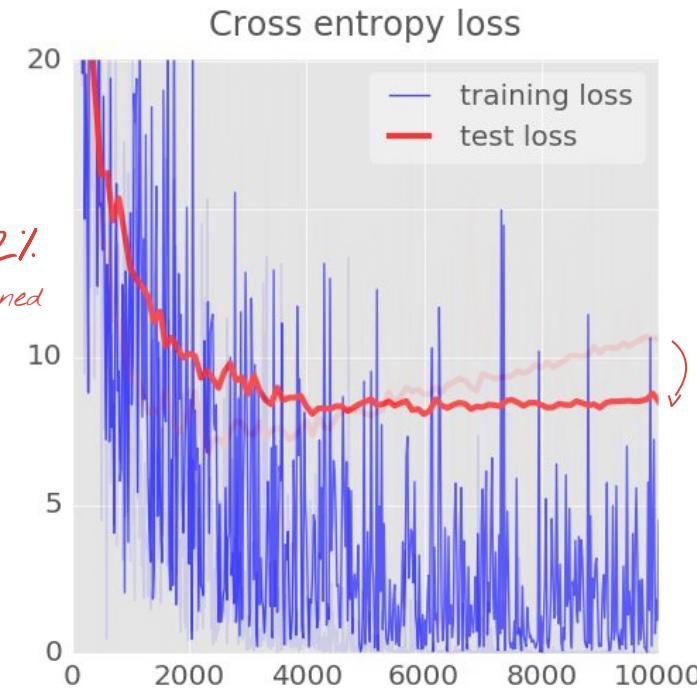
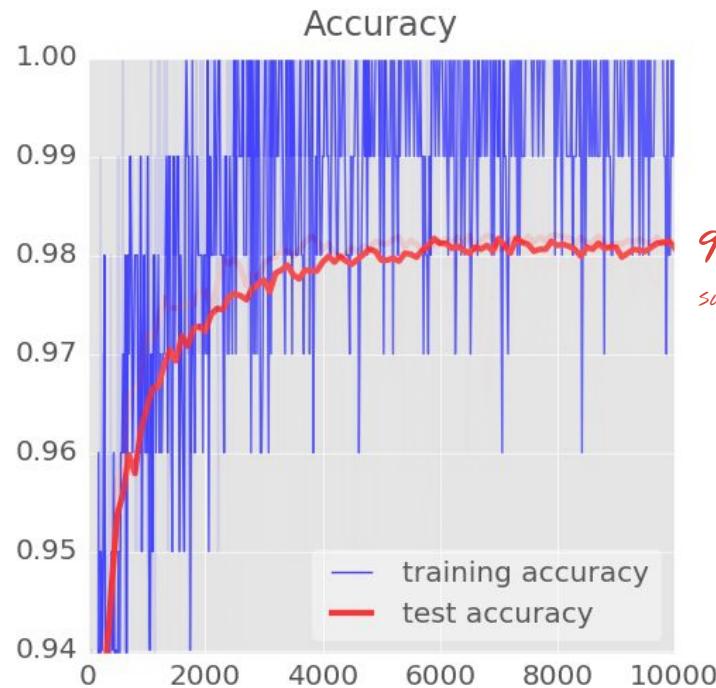
4 0 7 1 0 3 4 1 8 0  
5 0 0 8 3 6 0 1 8 3  
6 0 0 3 1 8 5 3 4 9  
1 9 2 5 8 0 2 0 1 0  
8 4 8 3 9 8 4 2 5 8  
0 6 1 4 5 3 8 5 5 1  
3 2 7 6 0 1 5 8 5 4  
8 6 6 2 0 9 1 0 4 5  
9 9 6 4 8 2 5 1 5 1  
0 5 8 1 1 2 7 9 9 0



Test digits

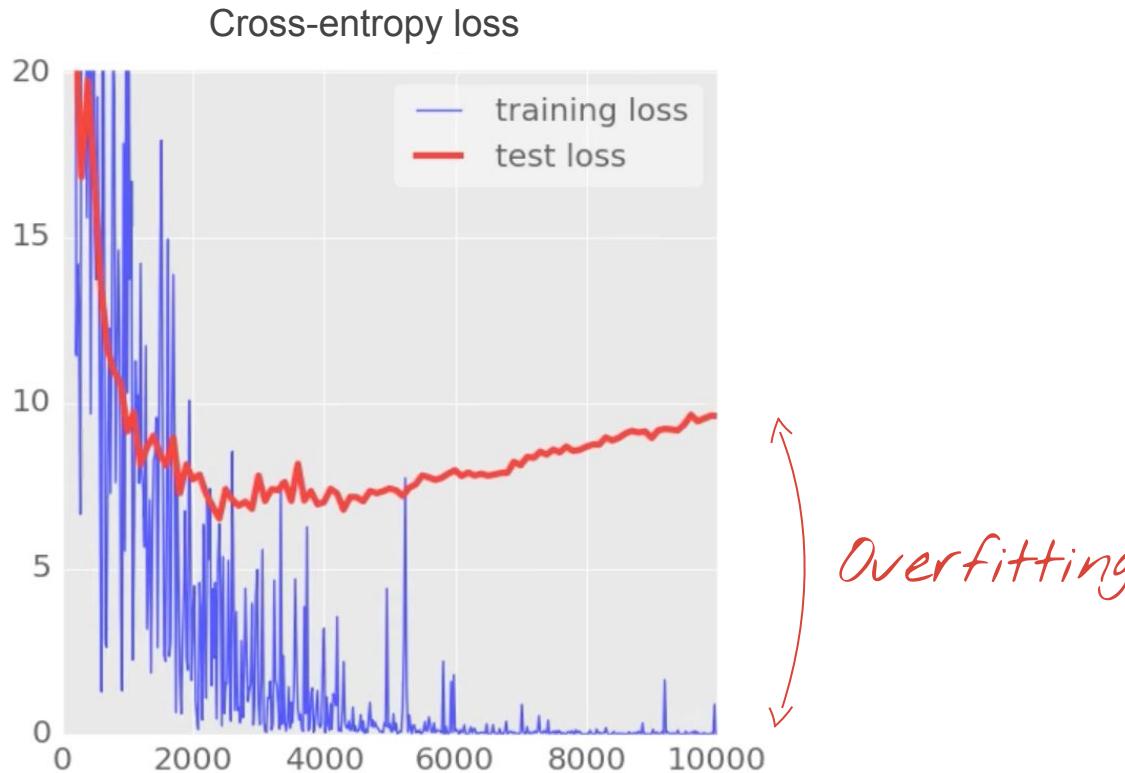
| Accuracy (%) | Digit |
|--------------|-------|
| 100%         | 9     |
| 98%          | 8     |
| 96%          | 7     |
| 94%          | 6     |
| 92%          | 5     |
| 90%          | 4     |

# All the party tricks



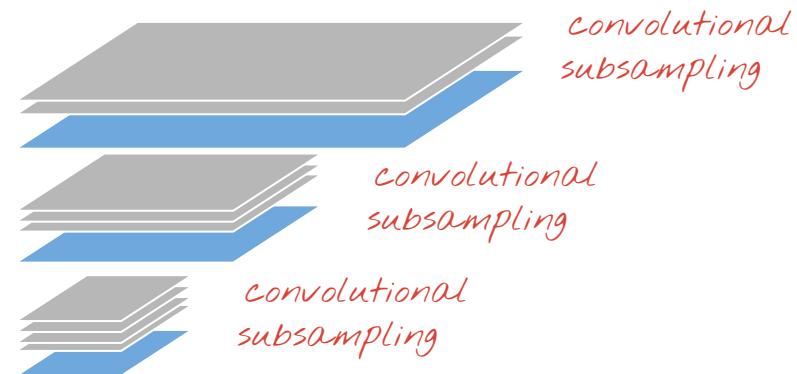
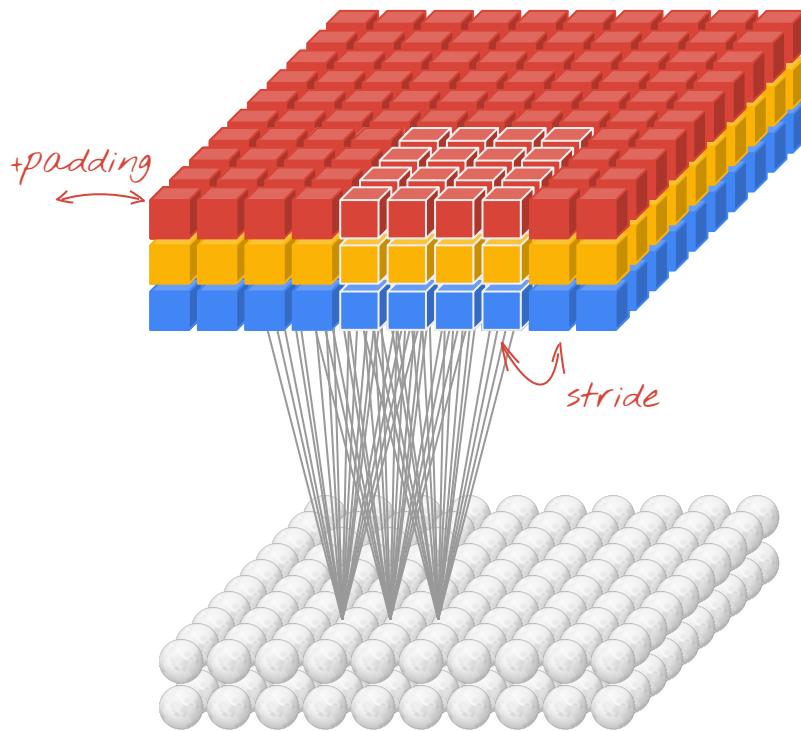
*RELU, decaying learning rate 0.003 → 0.0001 and dropout 0.75*

# Overfitting





# Convolutional layer



$W[4, 4, 3]$   
 $W_2[4, 4, 3]$  |  $W[4, 4, 3, 2]$

filter size      input channels      output channels

Hacker's tip

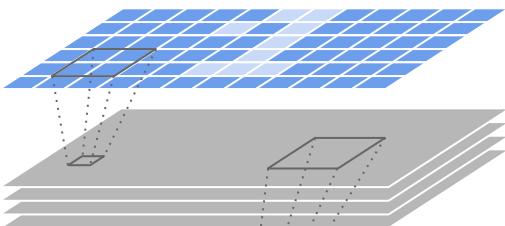


ALL  
Convolu-  
tional

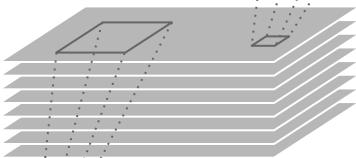
# Convolutional neural network

+ biases on  
all layers

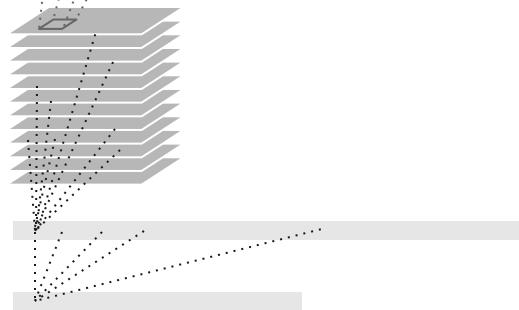
$28 \times 28 \times 1$



$28 \times 28 \times 4$



$14 \times 14 \times 8$



$7 \times 7 \times 12$

200

10

convolutional layer, 4 channels

$W1[5, 5, 1, 4]$  stride 1

convolutional layer, 8 channels

$W2[4, 4, 4, 8]$  stride 2

convolutional layer, 12 channels

$W3[4, 4, 8, 12]$  stride 2

fully connected layer

$W4[7 \times 7 \times 12, 200]$

softmax readout layer

$W5[200, 10]$

# Tensorflow - initialisation

K=4

L=8

M=12

*filter  
size*      *input  
channels*      *output  
channels*

W1 = tf.Variable(tf.truncated\_normal([5, 5, 1, K], stddev=0.1))

B1 = tf.Variable(tf.ones([K])/10)

W2 = tf.Variable(tf.truncated\_normal([5, 5, K, L], stddev=0.1))

B2 = tf.Variable(tf.ones([L])/10)

W3 = tf.Variable(tf.truncated\_normal([4, 4, L, M], stddev=0.1))

B3 = tf.Variable(tf.ones([M])/10)

N=200

*weights initialised  
with random values*

W4 = tf.Variable(tf.truncated\_normal([7\*7\*M, N], stddev=0.1))

B4 = tf.Variable(tf.ones([N])/10)

W5 = tf.Variable(tf.truncated\_normal([N, 10], stddev=0.1))

B5 = tf.Variable(tf.zeros([10])/10)

# Tensorflow - the model

input image batch  
 $X[100, 28, 28, 1]$

weights

stride

biases

```
Y1 = tf.nn.relu(tf.nn.conv2d(X, W1, strides=[1, 1, 1, 1], padding='SAME') + B1)
Y2 = tf.nn.relu(tf.nn.conv2d(Y1, W2, strides=[1, 2, 2, 1], padding='SAME') + B2)
Y3 = tf.nn.relu(tf.nn.conv2d(Y2, W3, strides=[1, 2, 2, 1], padding='SAME') + B3)
```

```
YY = tf.reshape(Y3, shape=[-1, 7 * 7 * M])
```

```
Y4 = tf.nn.relu(tf.matmul(YY, W4) + B4)
```

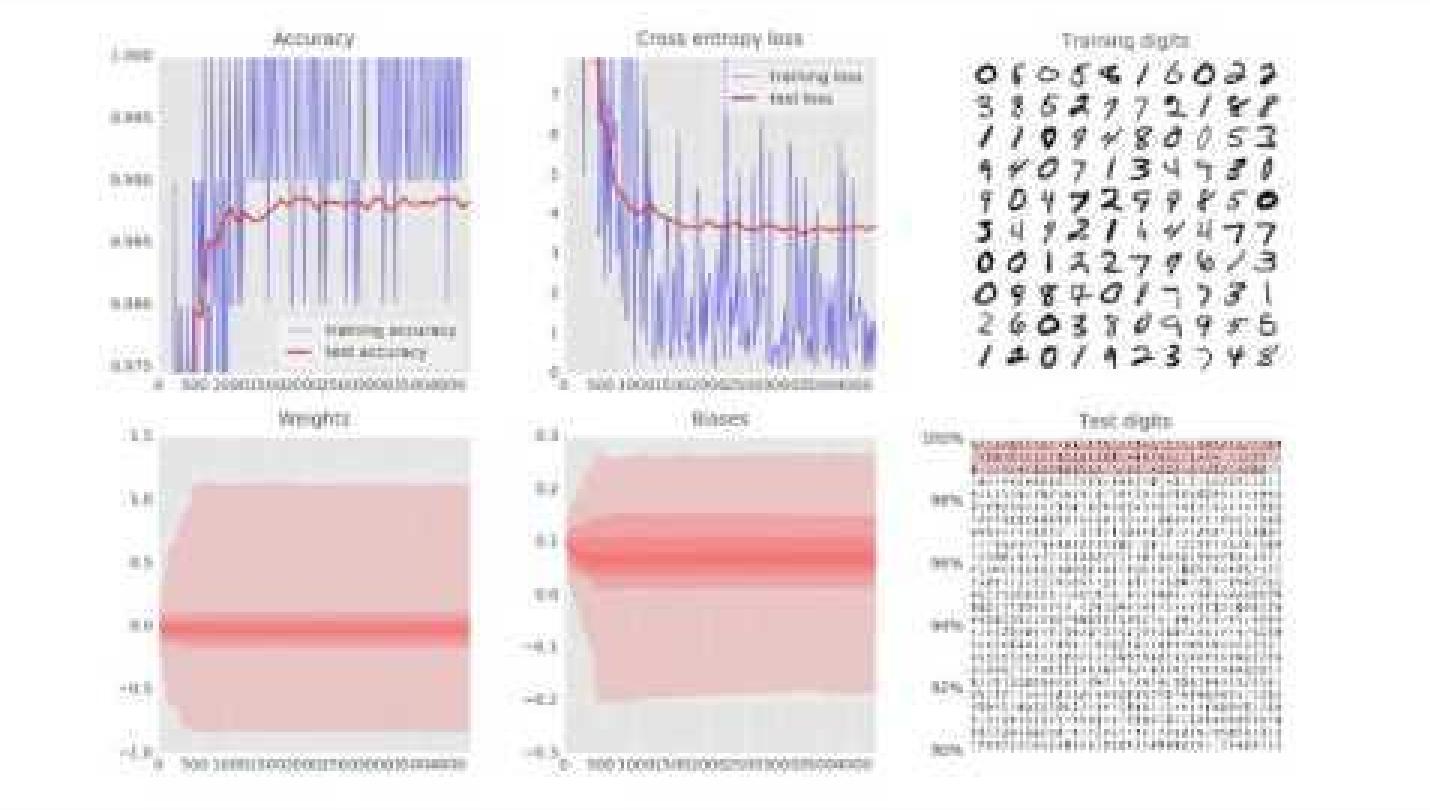
```
Y  = tf.nn.softmax(tf.matmul(Y4, W5) + B5)
```

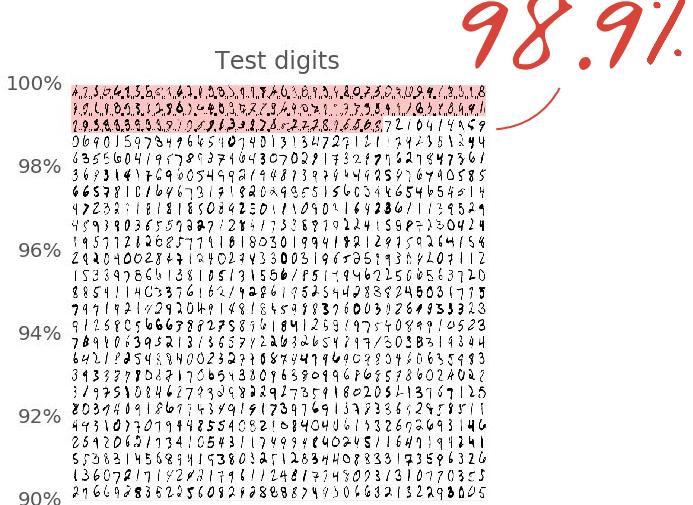
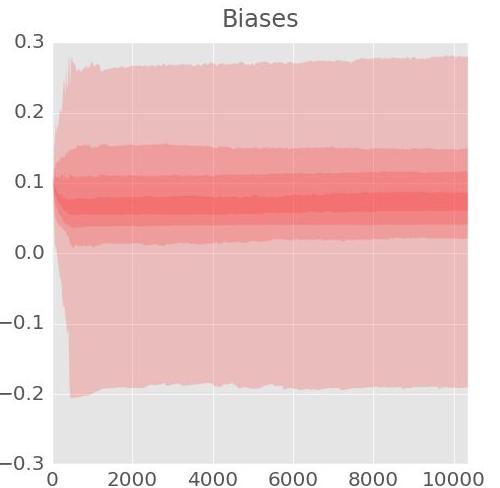
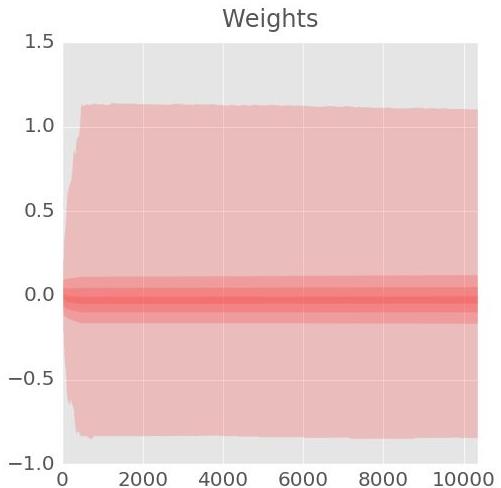
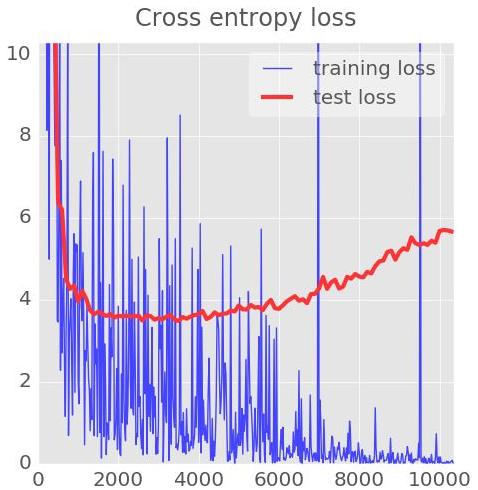
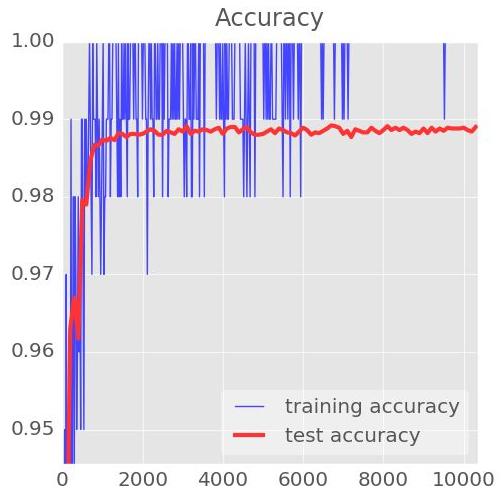
flatten all values for  
fully connected layer

$Y3 [100, 7, 7, 12]$

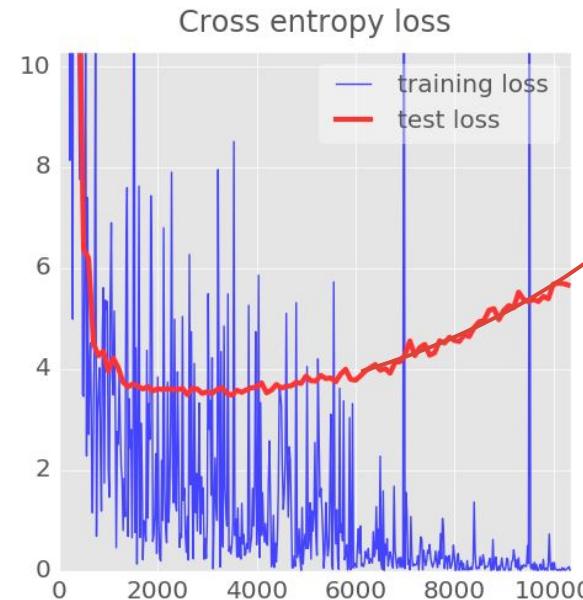
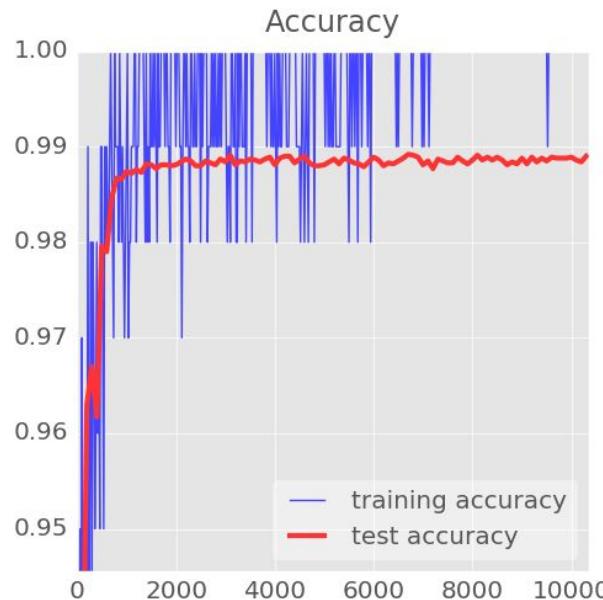
$YY [100, 7x7x12]$

# Demo

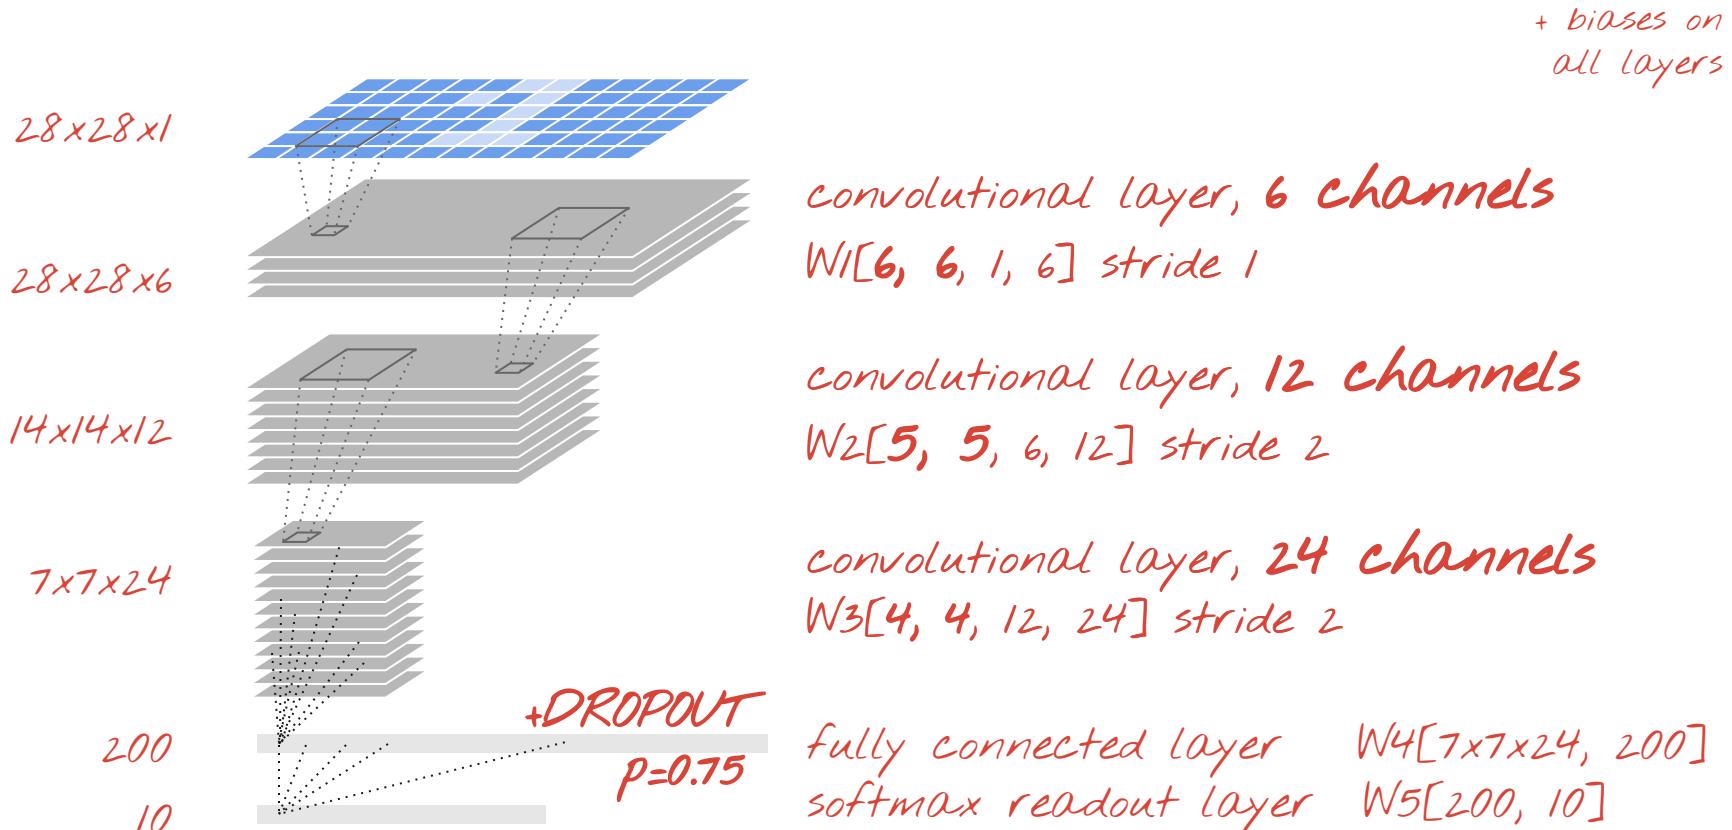




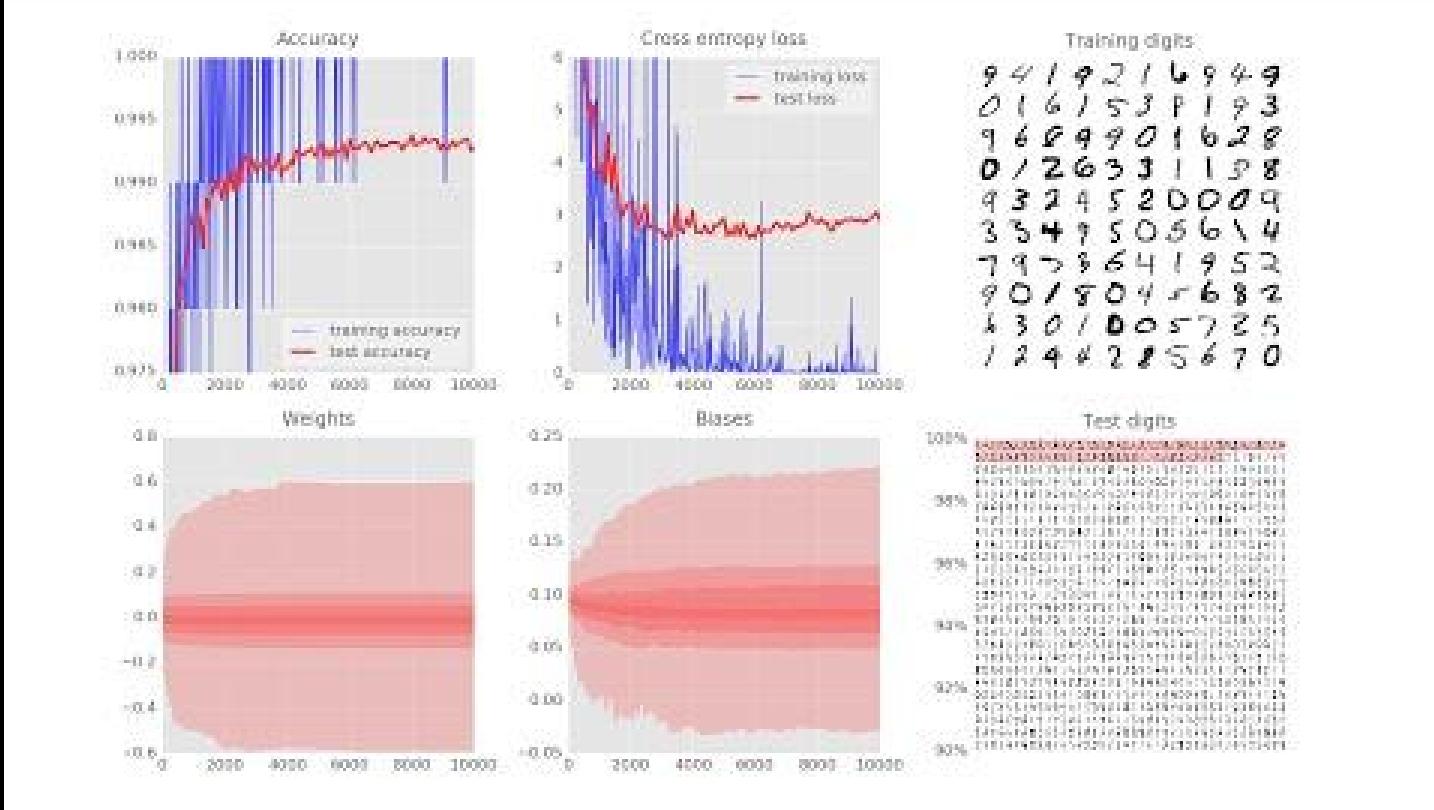
# WTXH ???

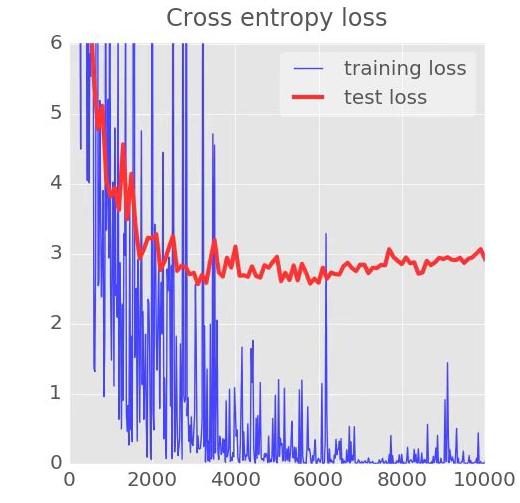
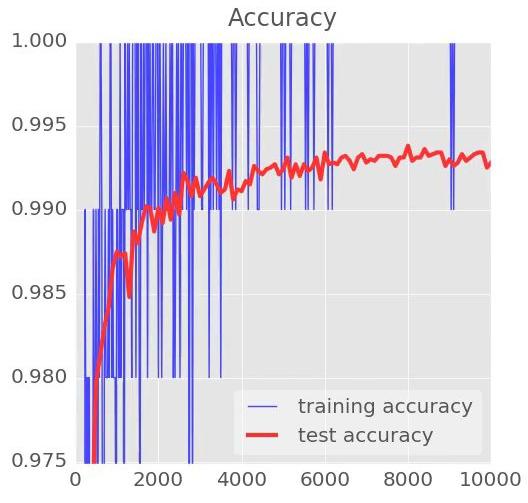


# Bigger convolutional network + dropout

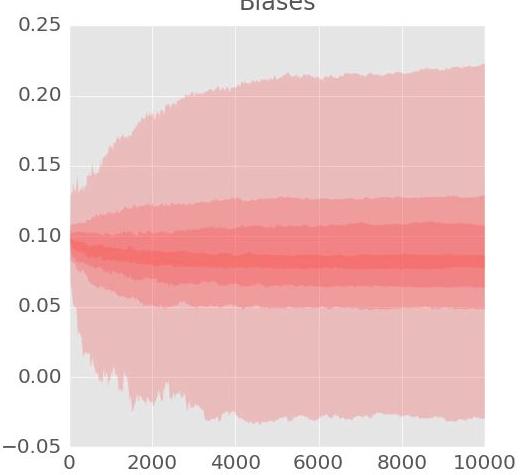
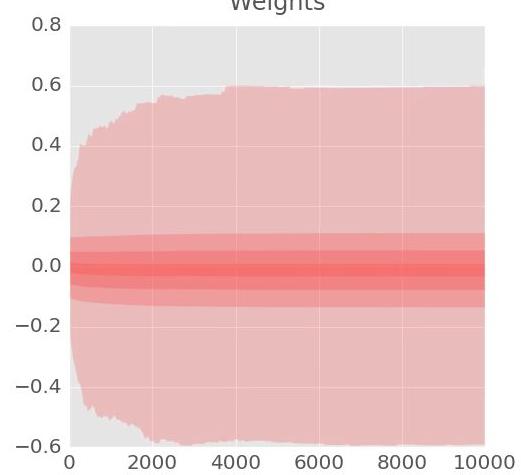


# Demo



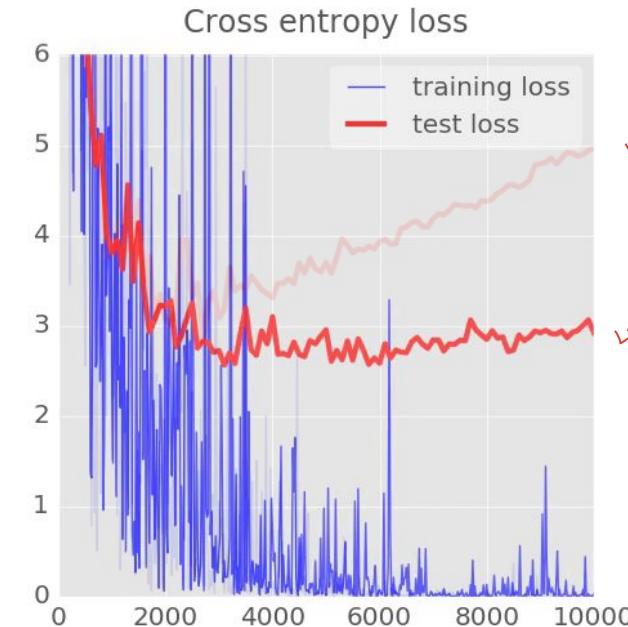
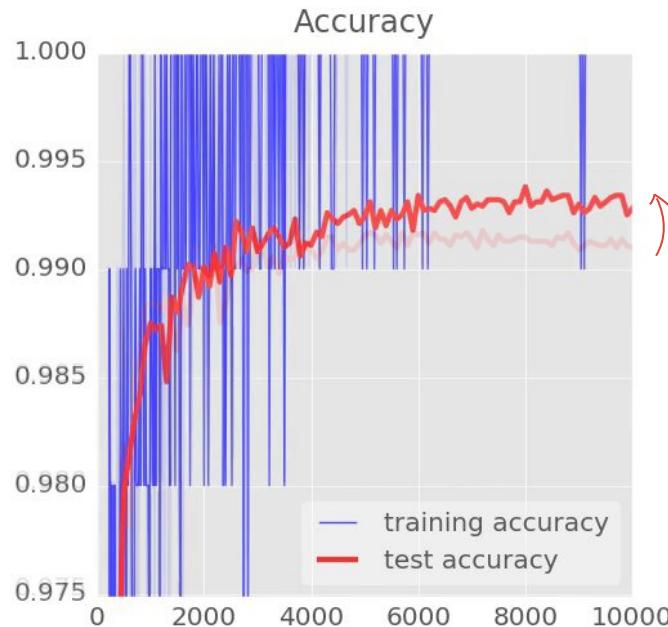


99.3%



|      |  |
|------|--|
| 100% | 64753232153243131037401231143102131003110310<br>62132135321301121431333521101121321227104919<br>67069015973949164901401313472712111743313<br>446353560419737164307027173271712178673<br>6436293116169605499014813371442501674958<br>5665331016967317162034955154034465495821<br>44725271818186084250111093164236711352<br>44839130365598872841133897922411594735042<br>41957712826877118180501994192291924115<br>629200002347134027133003192581131120711<br>215339286311351051215567F511946225065637<br>200854114033761621286119534428362460317<br>737911927292019148181599137500312649332<br>36116905663829758361691259197403991052<br>37891912392131365712483126581871305821330<br>464212054134002325168744139906804163348<br>33933770170654310963509976557860240223<br>179511844247938982927359118020511311125<br>8039460130714579171730769133723361291811<br>443107301944285405216450061326206314<br>62512062113410543117493446402451164919124<br>151363124568941538038128344688331935632<br>613607217117241111248179490931301707035<br>3416662832235509216806874930603215229560<br>5131416029117473938347121232339391190358 |
| 98%  | 64753232153243131037401231143102131003110310<br>62132135321301121431333521101121321227104919<br>67069015973949164901401313472712111743313<br>446353560419737164307027173271712178673<br>6436293116169605499014813371442501674958<br>5665331016967317162034955154034465495821<br>44725271818186084250111093164236711352<br>44839130365598872841133897922411594735042<br>41957712826877118180501994192291924115<br>629200002347134027133003192581131120711<br>215339286311351051215567F511946225065637<br>200854114033761621286119534428362460317<br>737911927292019148181599137500312649332<br>36116905663829758361691259197403991052<br>37891912392131365712483126581871305821330<br>464212054134002325168744139906804163348<br>33933770170654310963509976557860240223<br>179511844247938982927359118020511311125<br>8039460130714579171730769133723361291811<br>443107301944285405216450061326206314<br>62512062113410543117493446402451164919124<br>151363124568941538038128344688331935632<br>613607217117241111248179490931301707035<br>3416662832235509216806874930603215229560<br>5131416029117473938347121232339391190358 |
| 96%  | 64753232153243131037401231143102131003110310<br>62132135321301121431333521101121321227104919<br>67069015973949164901401313472712111743313<br>446353560419737164307027173271712178673<br>6436293116169605499014813371442501674958<br>5665331016967317162034955154034465495821<br>44725271818186084250111093164236711352<br>44839130365598872841133897922411594735042<br>41957712826877118180501994192291924115<br>629200002347134027133003192581131120711<br>215339286311351051215567F511946225065637<br>200854114033761621286119534428362460317<br>737911927292019148181599137500312649332<br>36116905663829758361691259197403991052<br>37891912392131365712483126581871305821330<br>464212054134002325168744139906804163348<br>33933770170654310963509976557860240223<br>179511844247938982927359118020511311125<br>8039460130714579171730769133723361291811<br>443107301944285405216450061326206314<br>62512062113410543117493446402451164919124<br>151363124568941538038128344688331935632<br>613607217117241111248179490931301707035<br>3416662832235509216806874930603215229560<br>5131416029117473938347121232339391190358 |
| 94%  | 64753232153243131037401231143102131003110310<br>62132135321301121431333521101121321227104919<br>67069015973949164901401313472712111743313<br>446353560419737164307027173271712178673<br>6436293116169605499014813371442501674958<br>5665331016967317162034955154034465495821<br>44725271818186084250111093164236711352<br>44839130365598872841133897922411594735042<br>41957712826877118180501994192291924115<br>629200002347134027133003192581131120711<br>215339286311351051215567F511946225065637<br>200854114033761621286119534428362460317<br>737911927292019148181599137500312649332<br>36116905663829758361691259197403991052<br>37891912392131365712483126581871305821330<br>464212054134002325168744139906804163348<br>33933770170654310963509976557860240223<br>179511844247938982927359118020511311125<br>8039460130714579171730769133723361291811<br>443107301944285405216450061326206314<br>62512062113410543117493446402451164919124<br>151363124568941538038128344688331935632<br>613607217117241111248179490931301707035<br>3416662832235509216806874930603215229560<br>5131416029117473938347121232339391190358 |
| 92%  | 64753232153243131037401231143102131003110310<br>62132135321301121431333521101121321227104919<br>67069015973949164901401313472712111743313<br>446353560419737164307027173271712178673<br>6436293116169605499014813371442501674958<br>5665331016967317162034955154034465495821<br>44725271818186084250111093164236711352<br>44839130365598872841133897922411594735042<br>41957712826877118180501994192291924115<br>629200002347134027133003192581131120711<br>215339286311351051215567F511946225065637<br>200854114033761621286119534428362460317<br>737911927292019148181599137500312649332<br>36116905663829758361691259197403991052<br>37891912392131365712483126581871305821330<br>464212054134002325168744139906804163348<br>33933770170654310963509976557860240223<br>179511844247938982927359118020511311125<br>8039460130714579171730769133723361291811<br>443107301944285405216450061326206314<br>62512062113410543117493446402451164919124<br>151363124568941538038128344688331935632<br>613607217117241111248179490931301707035<br>3416662832235509216806874930603215229560<br>5131416029117473938347121232339391190358 |
| 90%  | 64753232153243131037401231143102131003110310<br>62132135321301121431333521101121321227104919<br>67069015973949164901401313472712111743313<br>446353560419737164307027173271712178673<br>6436293116169605499014813371442501674958<br>5665331016967317162034955154034465495821<br>44725271818186084250111093164236711352<br>44839130365598872841133897922411594735042<br>41957712826877118180501994192291924115<br>629200002347134027133003192581131120711<br>215339286311351051215567F511946225065637<br>200854114033761621286119534428362460317<br>737911927292019148181599137500312649332<br>36116905663829758361691259197403991052<br>37891912392131365712483126581871305821330<br>464212054134002325168744139906804163348<br>33933770170654310963509976557860240223<br>179511844247938982927359118020511311125<br>8039460130714579171730769133723361291811<br>443107301944285405216450061326206314<br>62512062113410543117493446402451164919124<br>151363124568941538038128344688331935632<br>613607217117241111248179490931301707035<br>3416662832235509216806874930603215229560<br>5131416029117473938347121232339391190358 |

# YEAH !



*with dropout*



# Have fun !



# TensorFlow

[tensorflow.org](http://tensorflow.org)



Martin Görner

Google Developer relations

[@martin\\_gorner](https://twitter.com/martin_gorner)

[plus.google.com/+MartinGorner](https://plus.google.com/+MartinGorner)



Google Cloud Platform - [cloud.google.com](http://cloud.google.com)



**Cloud ML** ALPHA

your TensorFlow models  
trained in Google's cloud,  
fast.



Pre-trained models:



Cloud Vision API



Cloud Speech API ALPHA



Google Translate API

All code snippets are on  
GitHub:

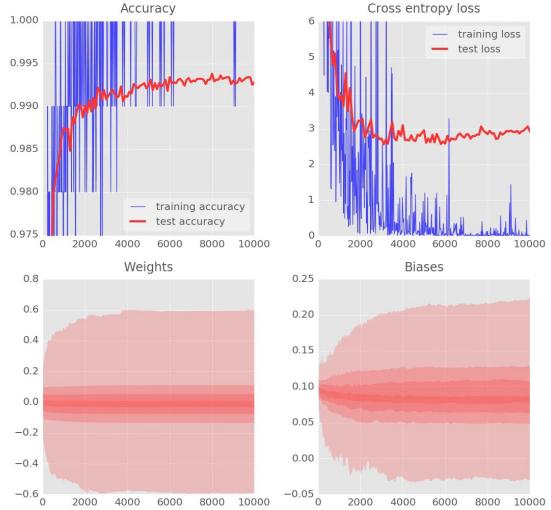
[github.com/martin-gorner/tensorflow-mnist-tutorial](https://github.com/martin-gorner/tensorflow-mnist-tutorial)

This presentation:

[goo.gl/pHeXe7](https://goo.gl/pHeXe7)



# Workshop



Keyboard shortcuts for the visualisation GUI:

|          |       |                                    |
|----------|-------|------------------------------------|
| 1        | ..... | display 1 <sup>st</sup> graph only |
| 2        | ..... | display 2 <sup>nd</sup> graph only |
| 3        | ..... | display 3 <sup>rd</sup> graph only |
| 4        | ..... | display 4 <sup>th</sup> graph only |
| 5        | ..... | display 5 <sup>th</sup> graph only |
| 6        | ..... | display 6 <sup>th</sup> graph only |
| 7        | ..... | display graphs 1 and 2             |
| 8        | ..... | display graphs 4 and 5             |
| 9        | ..... | display graphs 3 and 6             |
| ESC or 0 | ..    | back to displaying all graphs      |
| SPACE    | ..... | pause/resume                       |
| 0        | ..... | box zoom mode (then use mouse)     |
| H        | ..... | reset all zooms                    |
| Ctrl-S   | ....  | save current image                 |

# Workshop

Starter code and solutions:  
[github.com/martin-gorner/tensorflow-mnist-tutorial](https://github.com/martin-gorner/tensorflow-mnist-tutorial)

## 1. Theory (sit back and listen)

Softmax classifier, mini-batch, cross-entropy and how to implement them in Tensorflow (slides 1-14)

## 2. Practice

Open file: `mnist_1.0_softmax.py`  
Run it, play with the visualisations (see instructions on previous slide), read and understand the code as well as the basic structure of a Tensorflow program.

## 3. Theory (sit back and listen)

Hidden layers, sigmoid activation function (slides 16-19)

## 4. Practice

Start from the file you have and add one or two hidden layers. Use [cross\\_entropy\\_with\\_logits](#) to avoid numerical instabilities with  $\log(0)$ .

Solution in: `mnist_2.0_five_layers_sigmoid.py`

## 5. Theory (sit back and listen)

The neural network toolbox: RELUs, learning rate decay, dropout, overfitting (slides 20-35)

## 6. Practice

Replace all your sigmoids with RELUs. Test. Then add learning rate decay from 0.003 to 0.0001 using the formula  $lr = lr_{min} + (lr_{max} - lr_{min}) * \exp(-i/2000)$ .

Solution in: `mnist_2.1_five_layers_relu_lrdecay.py`

## 7. Practice (if time allows)

Add dropout on all layers using a value between 0.5 and 0.8 for `pkeep`.

Solution in: `mnist_2.2_five_layers_relu_lrdecay_dropout.py`

## 8. Theory (sit back and listen)

Convolutional networks (slides 36-42)

## 9. Practice

Replace your model with a convolutional network, without dropout.

Solution in: `mnist_3.0_convolutional.py`

## 10. Practice (if time allows)

Try a bigger neural network (good hyperparameters on slide 44) and add dropout on the last layer.

Solution in: `mnist_3.0_convolutional_bigger_dropout.py`