

# Lab 11

# Tensorflow Basics

Song Chaoyang

Assistant Professor


Department of Mechanical and Energy Engineering

[songcy@sustech.edu.cn](mailto:songcy@sustech.edu.cn)

# Let's Try

<https://github.com/GoogleCloudPlatform/tensorflow-without-a-phd>

>TensorFlow and deep learning\_  
without a PhD



#Tensorflow

 Google Cloud Platform


@martin\_gorner

# A 1-layer NN

○○○


```
1 $ git clone https://github.com/GoogleCloudPlatform/tensorflow-without-a-phd.git
2 $ cd tensorflow-without-a-phd/tensorflow-mnist-tutorial
```

## 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



AncoraSIR.com

DEVONX

@martin\_gorner


Google Cloud Platform

## Training digits

Training digits

- =>
- updates to weights and biases
- =>
- better recognition (loop)

Handwritten digits in the MNIST dataset are 28x28 pixel greyscale images.



# Gradient Descent

- Training digits and labels =>
- Loss function =>
- Gradient (partial derivatives) =>
- Steepest descent =>
- Update weights and biases =>
- Repeat with next mini-batch of training images and labels

## Softmax, on a batch of images

## Now in TensorFlow (Python)

Predictions  $Y[100, 10]$

Images  $X[100, 784]$    Weights  $W[784, 10]$    Biases  $b[10]$

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


applied line by line   matrix multiply   broadcast on all lines

tensor shapes in [ ]

tensor shapes:  $X[100, 784]$     $W[784, 10]$     $b[10]$

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

matrix multiply   broadcast on all lines



weighted sum of all pixels + bias

$$\text{softmax}(L_n) = \frac{e^{L_n}}{\|e^L\|}$$

neuron outputs

Partial derivative of cross-entropy gives the "gradient", computed for a given image, label and present value of weights and biases

## Success ?

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

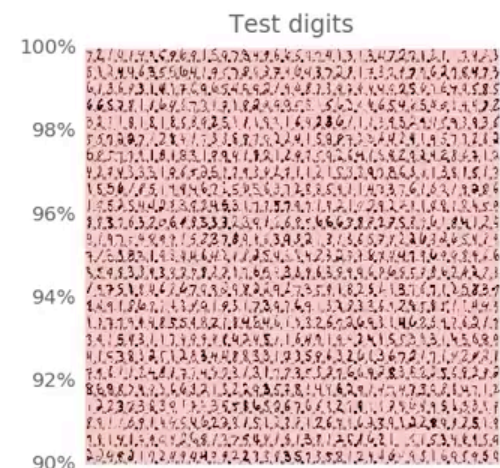
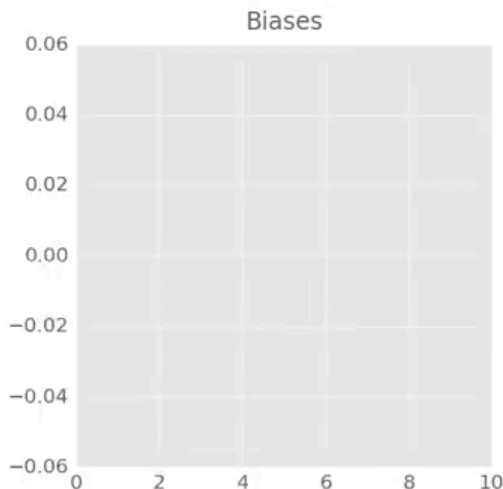



actual probabilities, "one-hot" encoded

$$\text{Cross entropy: } -\sum Y_i' \cdot \log(Y_i)$$

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

computed probabilities   this is a "6"

# python mnist\_1.0\_softmax.py => 92%



# Lab: let's jump into the code

## 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)
```

*initialisation*  
*this will become the batch size, 100*  
*28 x 28 grayscale images*  
*model*  
*flattening images*  
*Training = computing variables W and b*

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

*"one-hot" encoded*  
*success metrics*

```
# 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))
```

*"one-hot" decoding*  
*Tip: do this every 100 iterations*

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

*training step*  
*learning rate*  
*loss function*


```
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)
```

*running a TensorFlow computation, feeding placeholders*  
*Run*

# Lab: adding layers

Let's try 5 fully-connected layers !



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]))
```

DEVONX

@martin\_gorner

Google Cloud Platform

DEVONX

@martin\_gorner

Google Cloud Platform

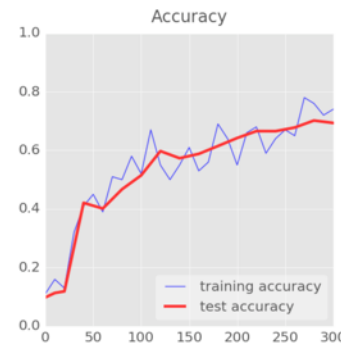
TensorFlow - the model

Demo - slow start ?

```
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)
```



DEVONX

@martin\_gorner

Google Cloud Platform

DEVONX

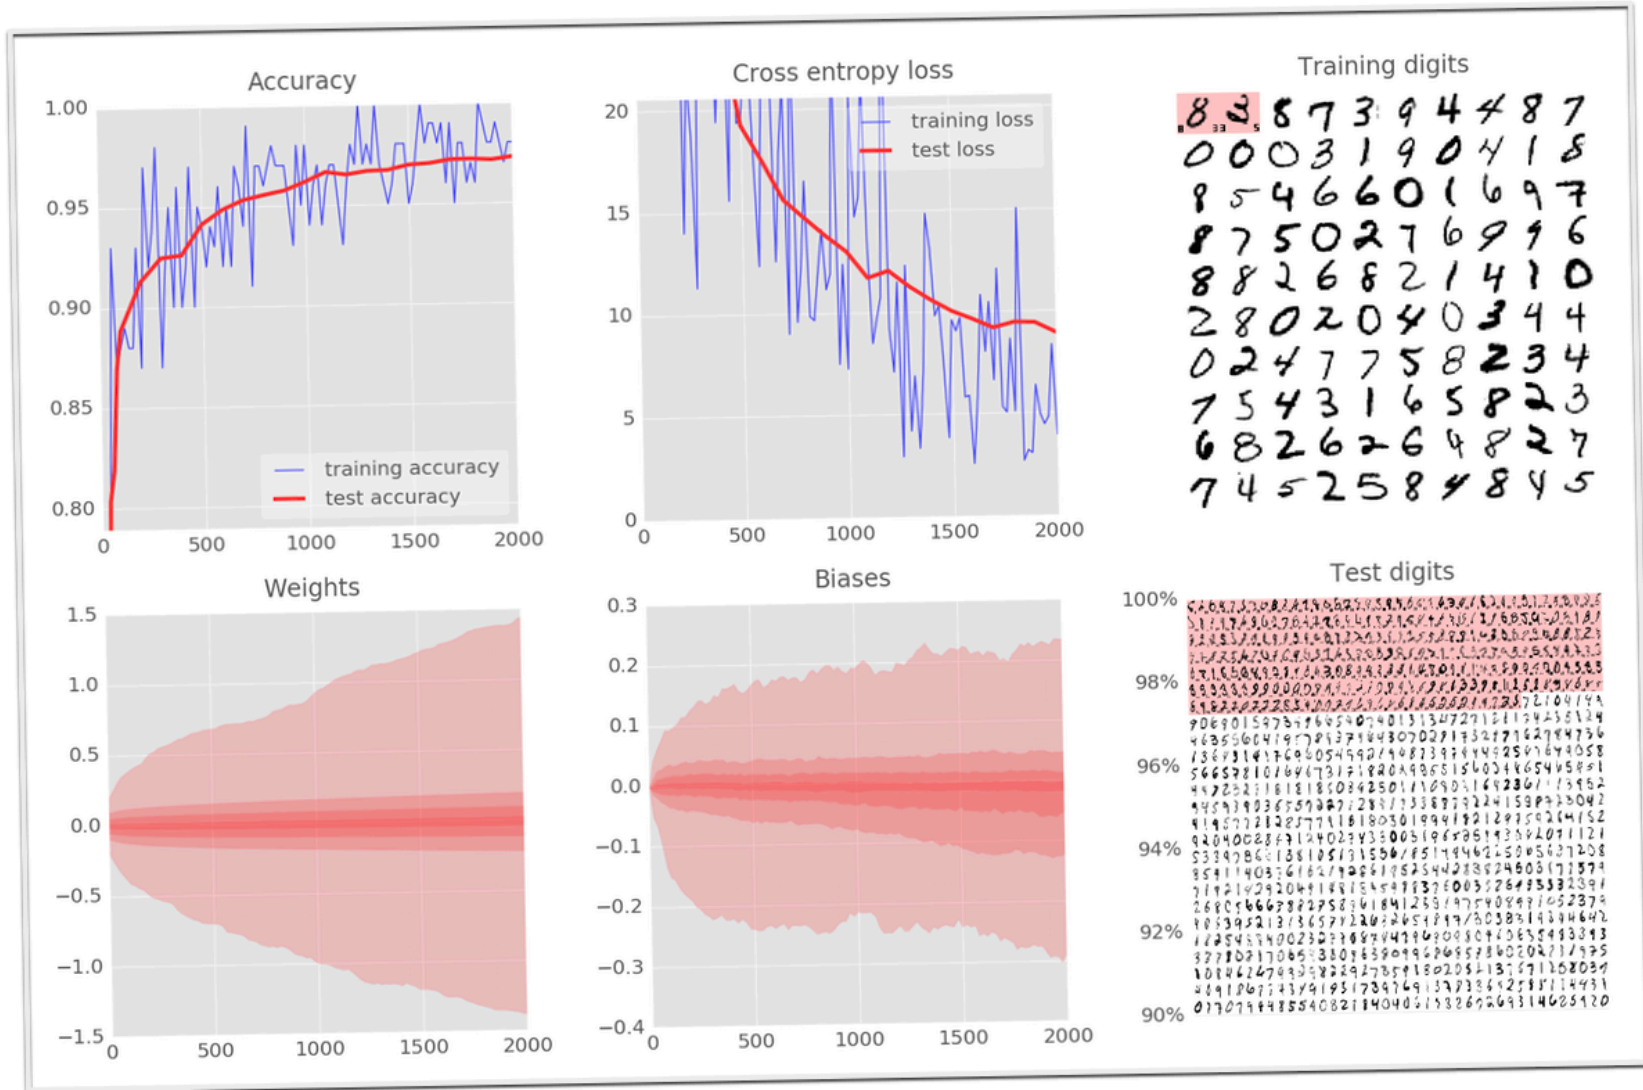
@martin\_gorner

Google Cloud Platform

AncoraSIR.com

Southern University of Science and Technology

# python mnist\_2.0\_five\_layers\_sigmoid.py => 97%






# Lab: special care for deep networks


## RELU

RELU = Rectified Linear Unit



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

## RELU



DEVONX

@martin\_gorner

Google Cloud Platform


DEVONX

@martin\_gorner

Google Cloud Platform

## Demo - noisy accuracy curve ?

## RELU



yuck! 🤔



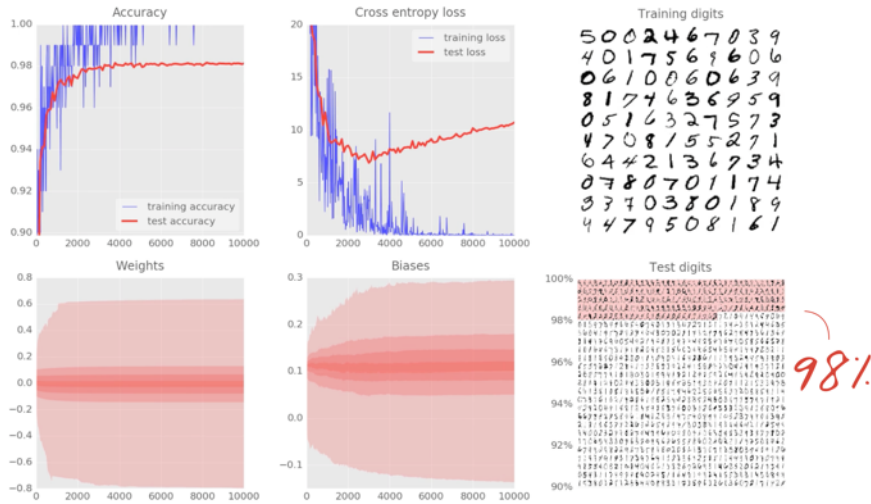
DEVONX

@martin\_gorner


Google Cloud Platform

Southern University of Science and Technology


# python mnist\_2.1\_five\_layers\_relu\_lrdecay.py => 98%



## Learning rate decay



## Dropout



pkeep =  
tf.placeholder(tf.float32)

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

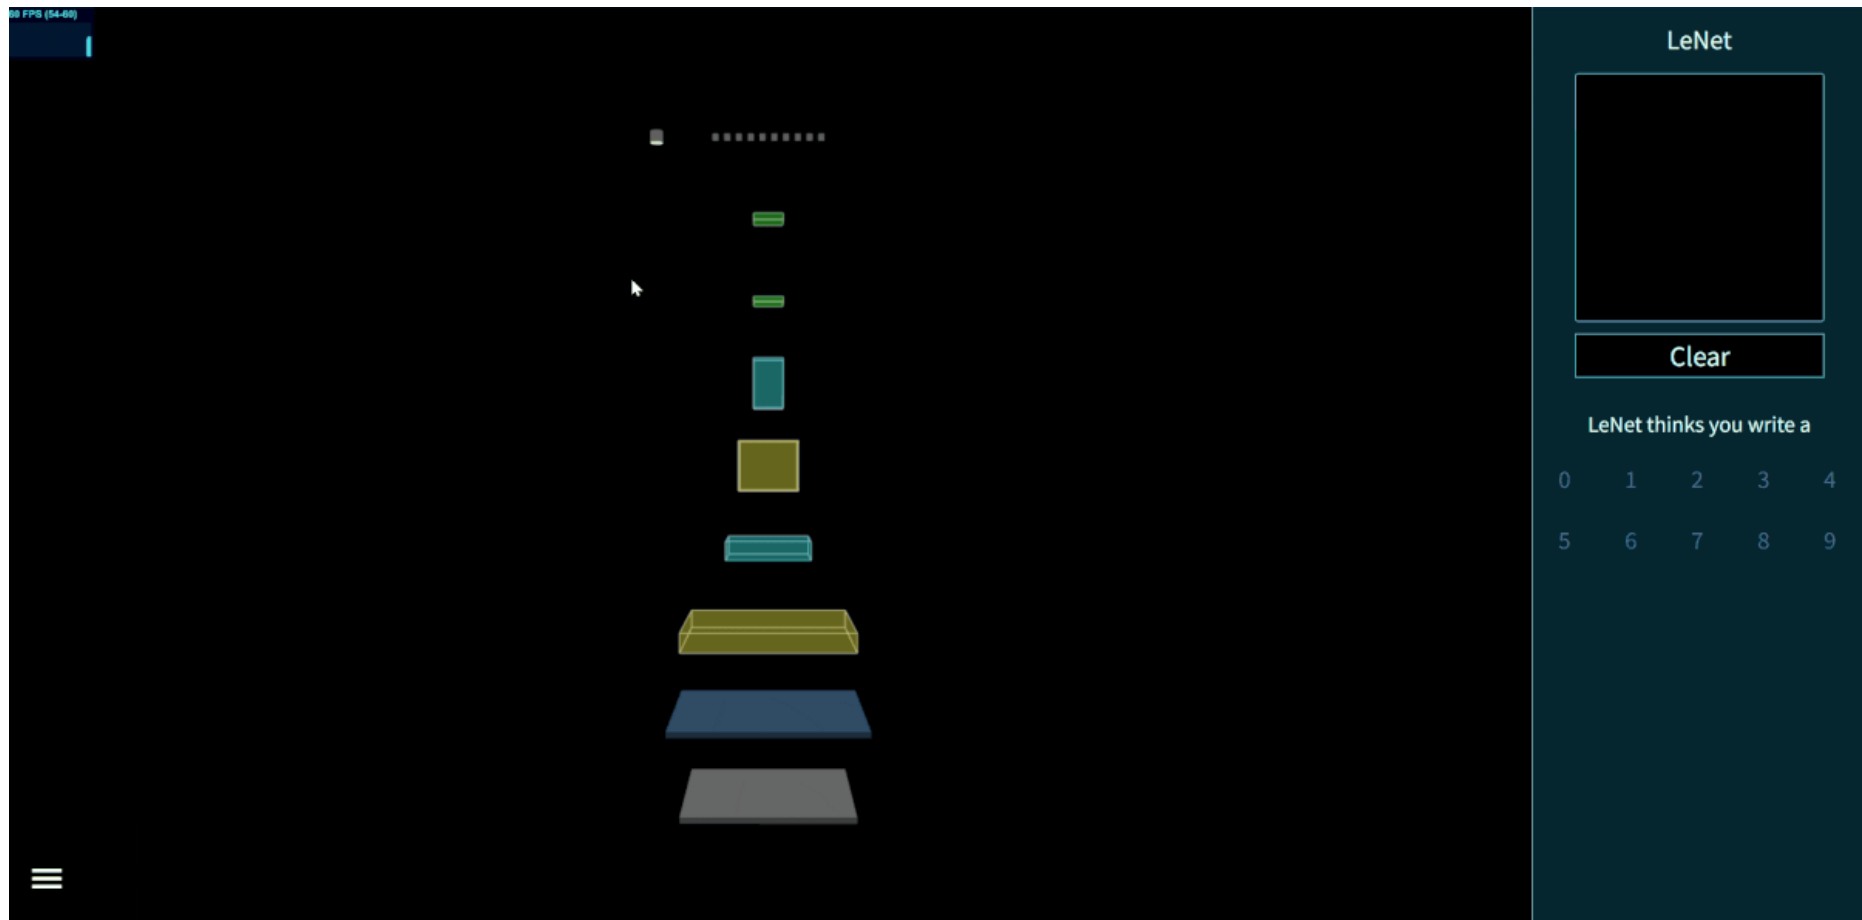
## All the party tricks



mnist\_2.2\_five\_layers\_relu\_lrdecay\_dropout.py

# TensorSpace.js

*A neural network 3D visualization framework*



YOLOv2-tiny

Select an image

MobileNetv1

Select an image

MobileNet thinks its a  
fox squirrel

AlexNet

Select an image

AlexNet thinks its a  
beagle

ResNet-50

Select an image

ResNet thinks its a  
Old English sheepdog,  
bobtail

VGG-16

Select an image

VGG thinks its a  
coffeepot

InceptionV3  
(Model Size: 96MB)

Select    reset

InceptionV3 thinks its a  
Persian cat

[Live demo](#)



**SUSTech**  
Southern University  
of Science and Technology

# Thank you!

Prof. Song Chaoyang

- Dr. Wan Fang ([sophie.fwan@hotmail.com](mailto:sophie.fwan@hotmail.com))

