

# Building Blocks of Deep Learning – Classification with (Dense) Neural Networks

*Cris Ababei*



**BE THE DIFFERENCE.**

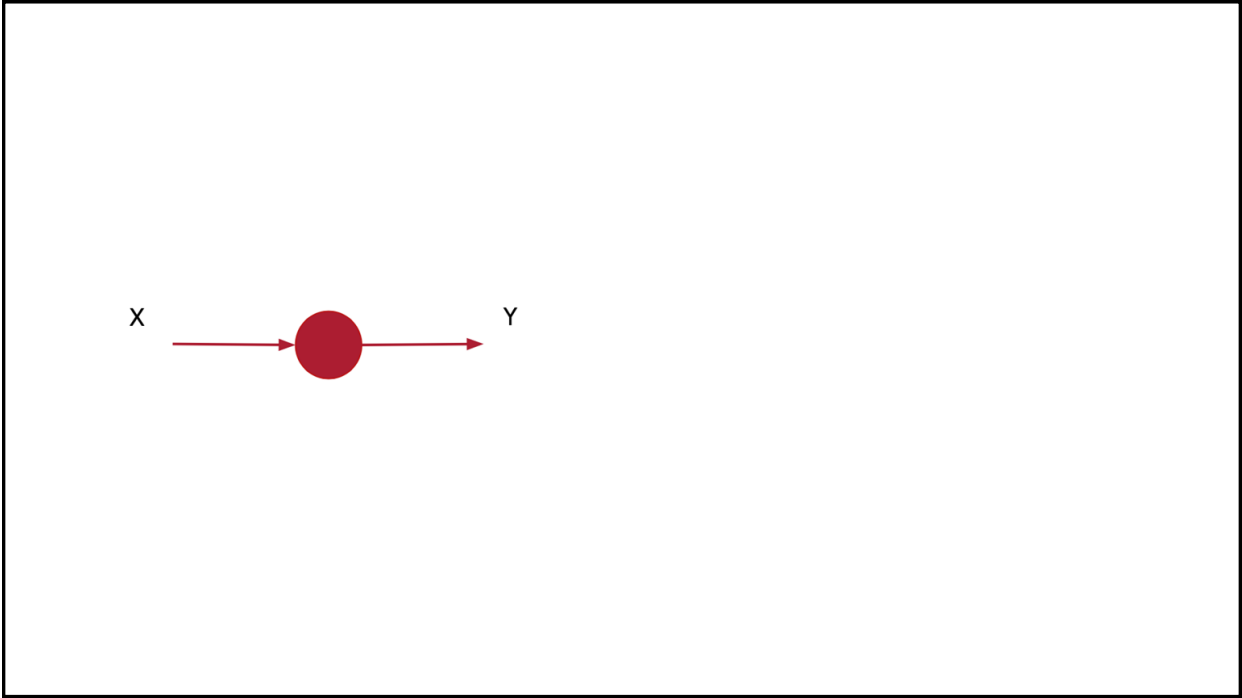
1

1

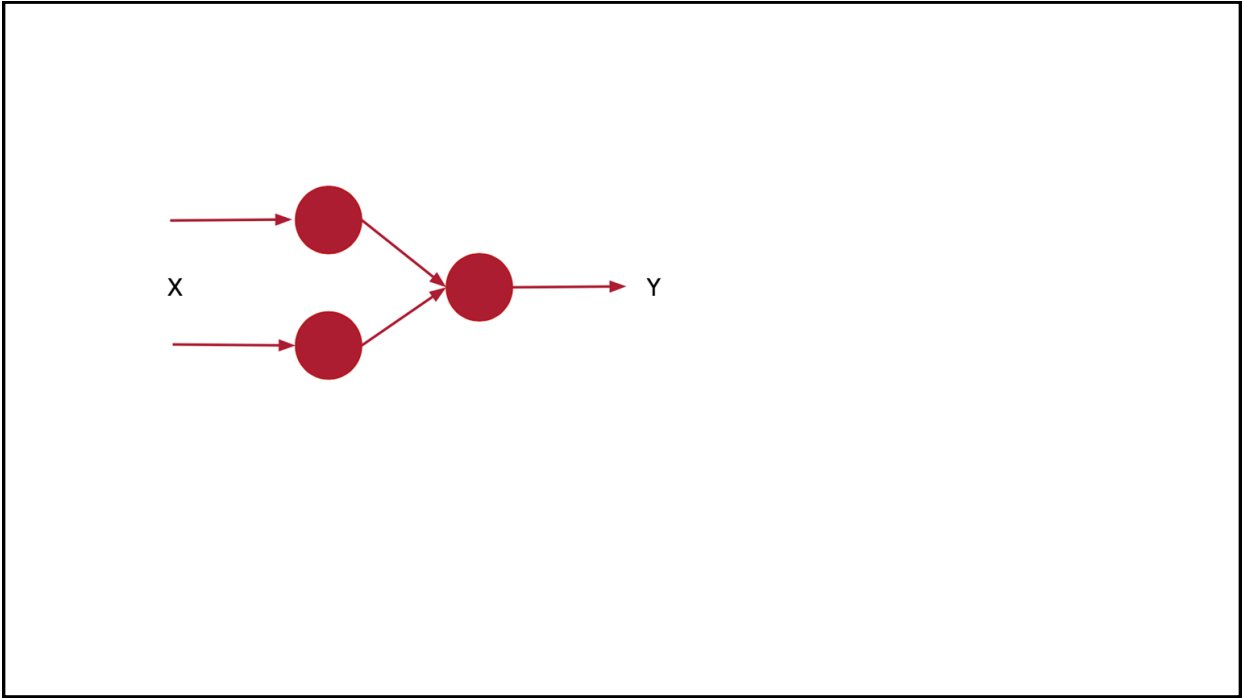
## From Regression to Classification

2

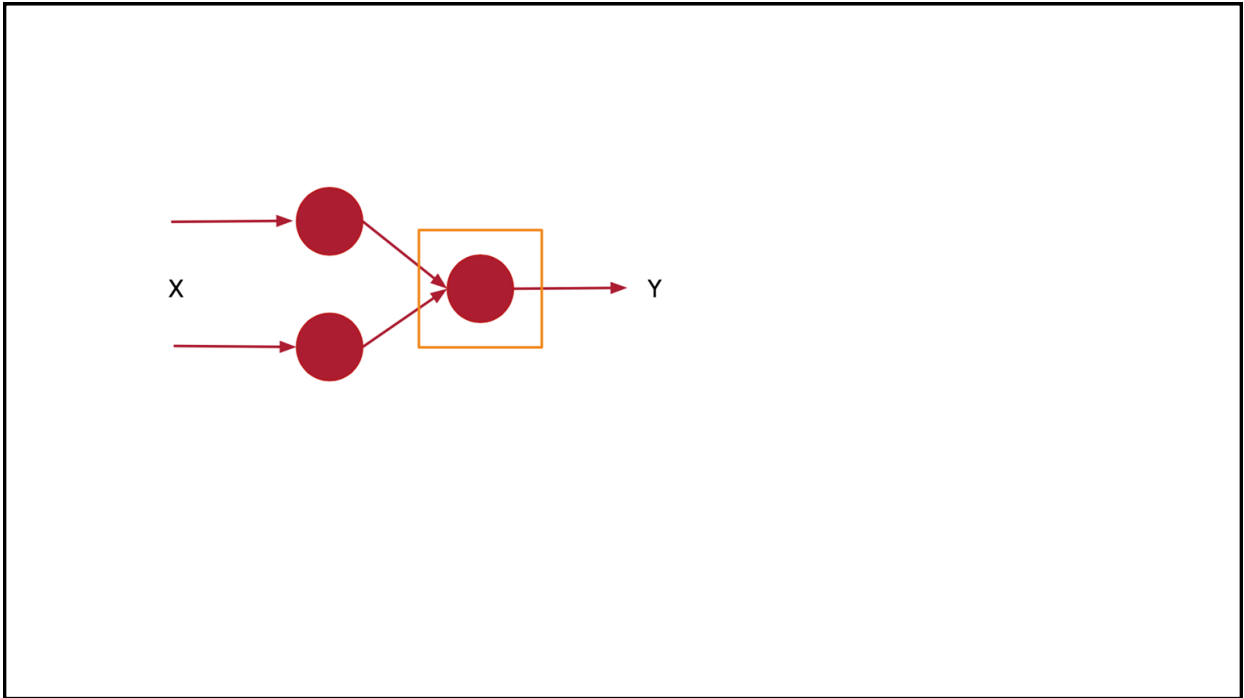
2



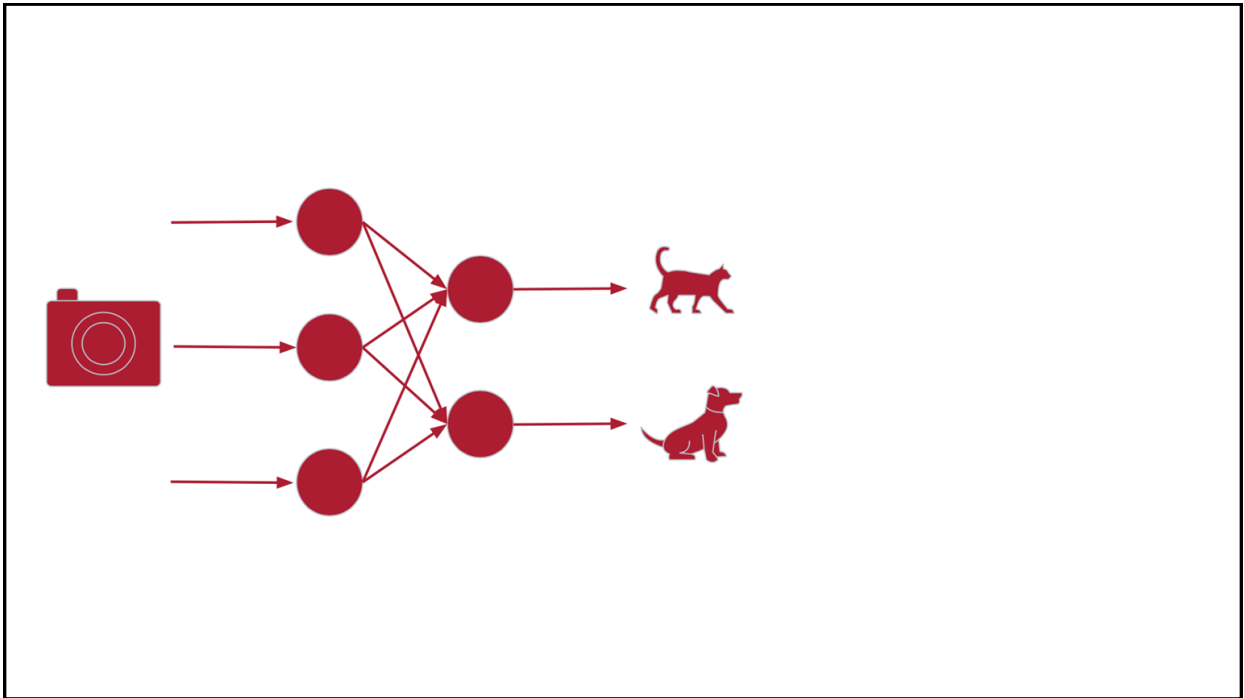
3



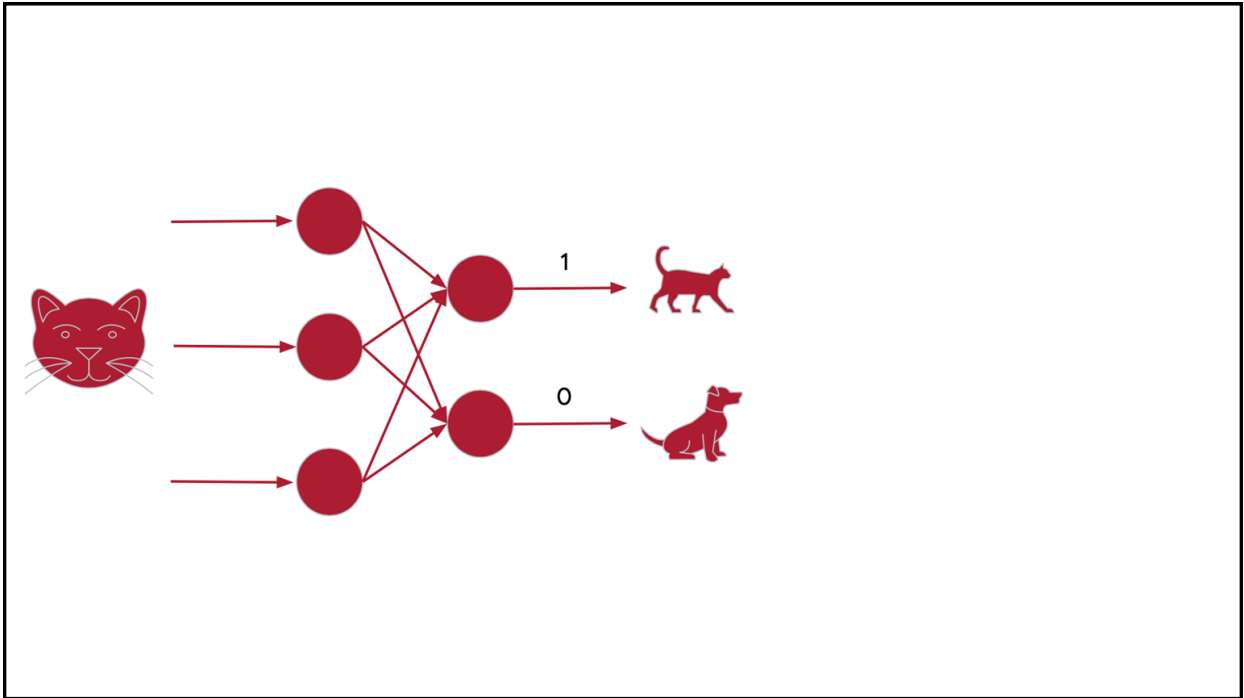
4



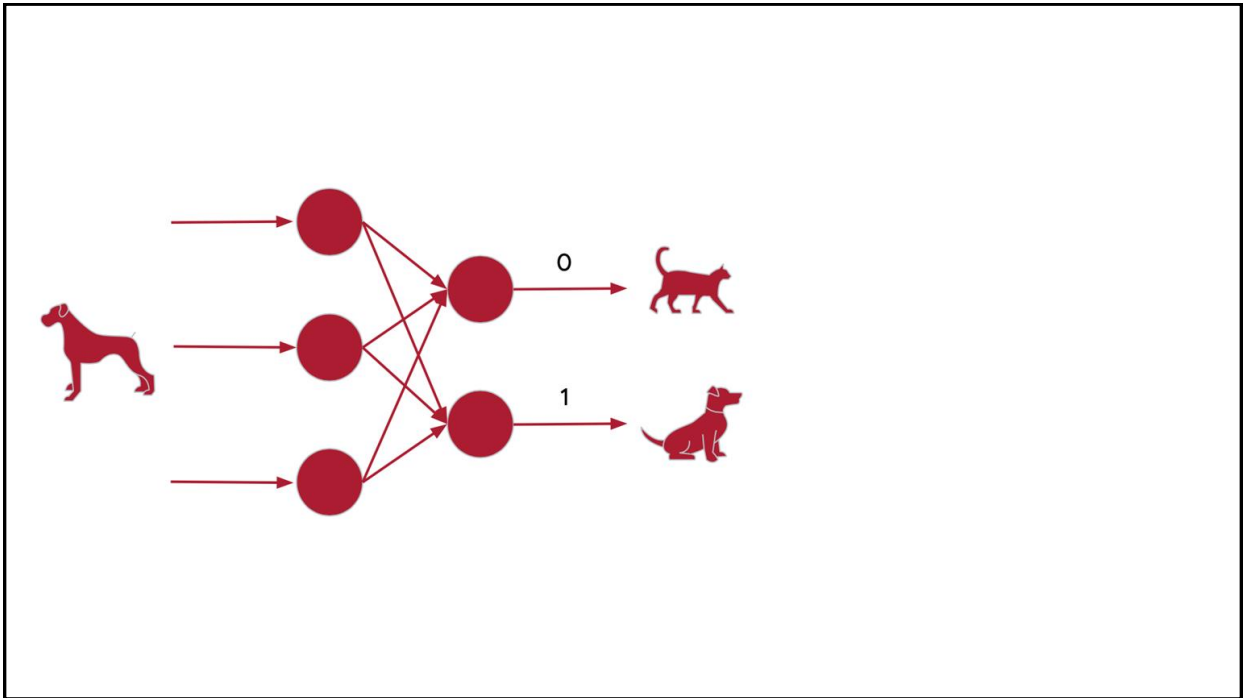
5





6



7



8

Data	Label
	[1, 0]
	[0, 1]

9

Data	Label
0	[1, 0, 0, 0, 0, 0, 0, 0, 0, 0]
1	[0, 1, 0, 0, 0, 0, 0, 0, 0, 0]
2	[0, 0, 1, 0, 0, 0, 0, 0, 0, 0]
3	[0, 0, 0, 1, 0, 0, 0, 0, 0, 0]
4	[0, 0, 0, 0, 1, 0, 0, 0, 0, 0]
5	[0, 0, 0, 0, 0, 1, 0, 0, 0, 0]
6	[0, 0, 0, 0, 0, 0, 1, 0, 0, 0]
7	[0, 0, 0, 0, 0, 0, 0, 1, 0, 0]
8	[0, 0, 0, 0, 0, 0, 0, 0, 1, 0]
9	[0, 0, 0, 0, 0, 0, 0, 0, 0, 1]

10

```
import tensorflow as tf
```

```
data = tf.keras.datasets.mnist  
(training_images, training_labels), (val_images, val_labels) = data.load_data()
```

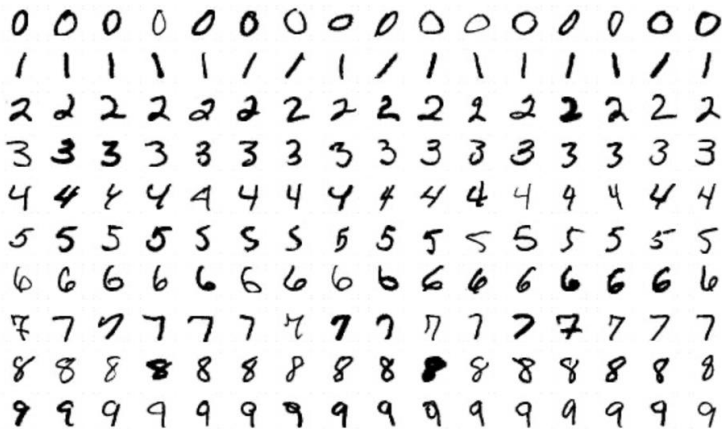
```
training_images = training_images / 255.0
```

```
val_images = val_images / 255.0
```

```
model = tf.keras.models.Sequential(  
    [tf.keras.layers.Flatten(input_shape=(28,28)),  
     tf.keras.layers.Dense(20, activation=tf.nn.relu),  
     tf.keras.layers.Dense(10, activation=tf.nn.softmax)])
```

Collect  
Data

11



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

60,000 Labelled Training Examples  
10,000 Labelled Validation Examples

12

```
import tensorflow as tf

data = tf.keras.datasets.mnist
(training_images, training_labels), (val_images, val_labels) = data.load_data()

training_images = training_images / 255.0
val_images = val_images / 255.0

model = tf.keras.models.Sequential(
    [tf.keras.layers.Flatten(input_shape=(28,28)),
     tf.keras.layers.Dense(20, activation=tf.nn.relu),
     tf.keras.layers.Dense(10, activation=tf.nn.softmax)])
```



13

```
import tensorflow as tf

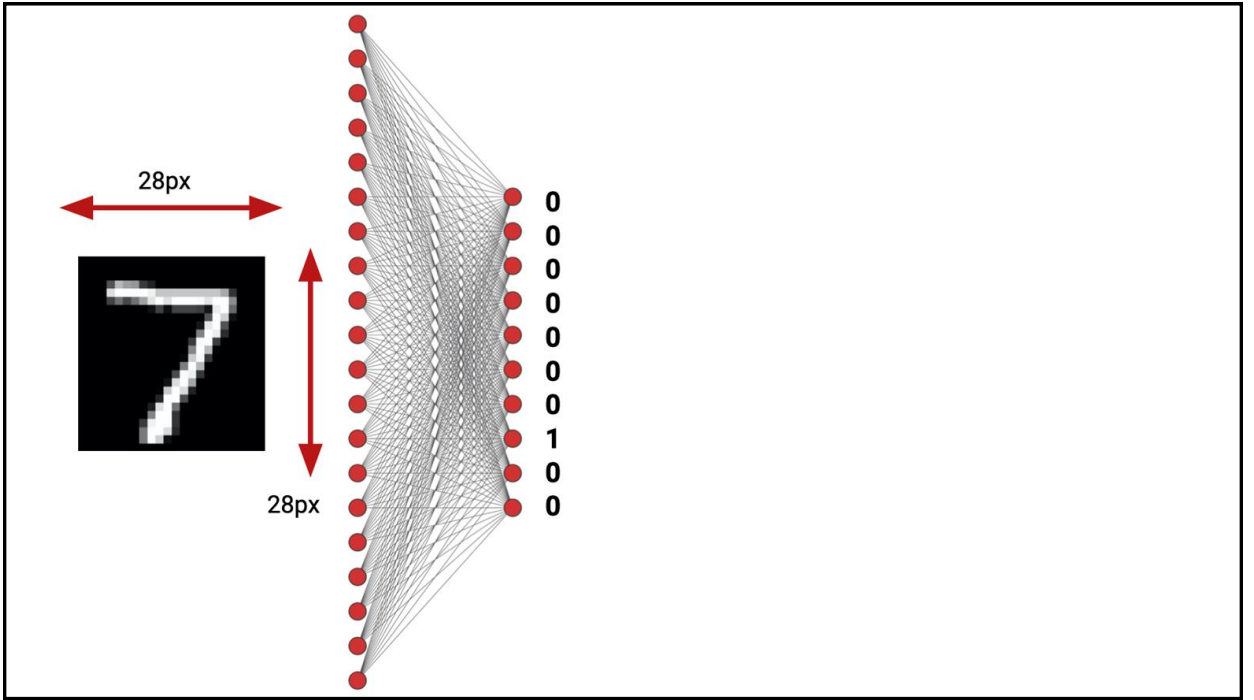
data = tf.keras.datasets.mnist
(training_images, training_labels), (val_images, val_labels) = data.load_data()

training_images = training_images / 255.0
val_images = val_images / 255.0

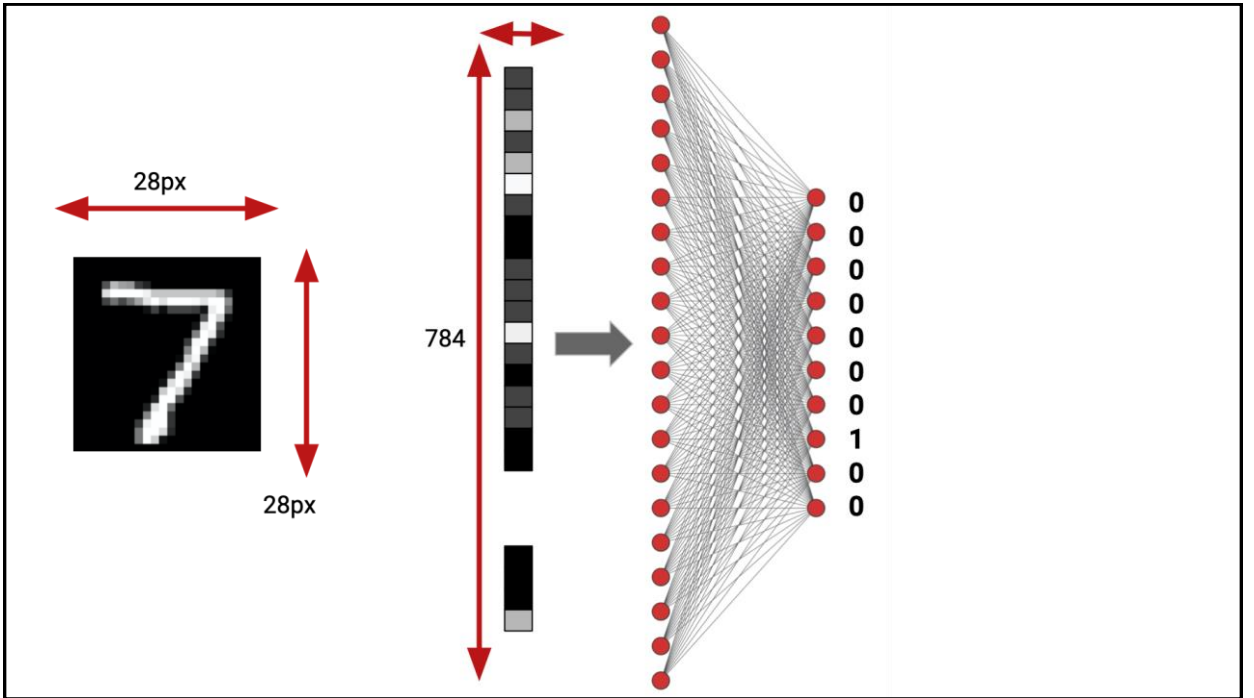
model = tf.keras.models.Sequential(
    [tf.keras.layers.Flatten(input_shape=(28,28)),
     tf.keras.layers.Dense(20, activation=tf.nn.relu),
     tf.keras.layers.Dense(10, activation=tf.nn.softmax)])
```



14



15



16

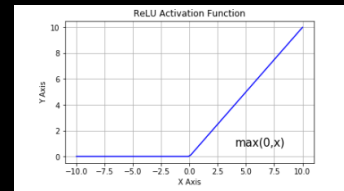


```
import tensorflow as tf

data = tf.keras.datasets.mnist
(training_images, training_labels), (val_images, val_labels) = data.load_data()

training_images = training_images / 255.0
val_images = val_images / 255.0

model = tf.keras.models.Sequential(
    [tf.keras.layers.Flatten(input_shape=(28, 28)),
     tf.keras.layers.Dense(20, activation=tf.nn.relu),
     tf.keras.layers.Dense(10, activation=tf.nn.softmax)]])
```



“ReLU applies much-needed non-linearity into the model. Non-linearity is necessary to produce non-linear decision boundaries, so that the output cannot be written as a linear combination of the inputs.”

[https://en.wikipedia.org/wiki/Activation\\_function](https://en.wikipedia.org/wiki/Activation_function)

Design a Model

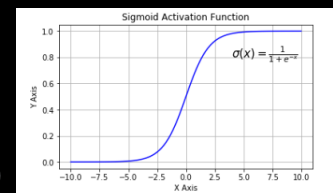
17

```
import tensorflow as tf

data = tf.keras.datasets.mnist
(training_images, training_labels), (val_images, val_labels) = data.load_data()

training_images = training_images / 255.0
val_images = val_images / 255.0

model = tf.keras.models.Sequential(
    [tf.keras.layers.Flatten(input_shape=(28, 28)),
     tf.keras.layers.Dense(20, activation=tf.nn.relu),
     tf.keras.layers.Dense(10, activation=tf.nn.softmax)]])
```



**SOFTMAX:** Generalization of the logistic function (or Sigmoid) to multiple dimensions. A softmax operation serves a key purpose: making sure the Neural Network (in this case, a Dense NN) outputs sum to 1. Because of this, softmax operations are useful to scale model outputs into probabilities.

Design a Model

18

```
model.compile(optimizer='adam',  
              loss='sparse_categorical_crossentropy',  
              metrics=['accuracy'])
```



19

```
model.compile(optimizer='adam',  
              loss='sparse_categorical_crossentropy',  
              metrics=['accuracy'])
```

### Mean Squared Error

$$MSE = \frac{1}{N} \sum (t_i - s_i)^2$$

Prediction  $s_i$   
Ground Truth  $t_i$

### Cross Entropy Loss

$$CE = - \sum_i^C t_i \log(s_i)$$

Classes  $C$   
Prediction  $s_i$   
Ground Truth  $\{0,1\}$   $t_i$



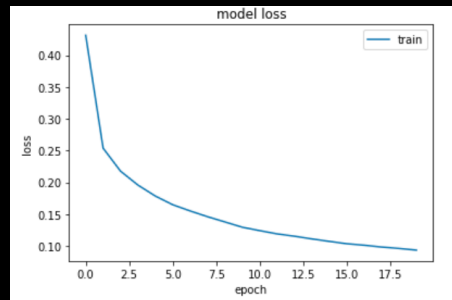
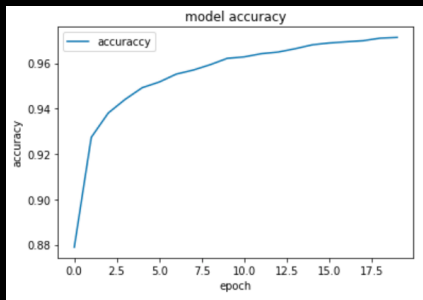
20

```
model.compile(optimizer='adam',  
              loss='sparse_categorical_crossentropy',  
              metrics=['accuracy'])  
  
model.fit(training_images, training_labels, epochs=20)
```



21

```
model.compile(optimizer='adam',  
              loss='sparse_categorical_crossentropy',  
              metrics=['accuracy'])  
  
model.fit(training_images, training_labels, epochs=20)
```



22

```
classifications = model.predict(val_images)
print(classifications[0])
print(test_labels[0])
```

```
[2.4921512e-09 1.3765138e-10 8.8281205e-08
1.0477231e-03 2.8455029e-12 4.0820678e-06
2.0070659e-16 9.9894780e-01 1.0296049e-07
2.9972372e-07]
```

7

Make  
Inferences

23

## Digit Classification using Dense NN with TF2

Code Time!

TF\_MNIST\_Classification.ipynb



24

24

# Credits

- A previous edition of this course was developed in collaboration with Dr. Susan C. Schneider of Marquette University.
- We are very grateful and thank all the following professors, researchers, and practitioners for jump-starting courses on TinyML and for sharing their teaching materials:
  - Prof. Marcelo Rovai - TinyML - Machine Learning for Embedding Devices, UNIFEI
    - <https://github.com/Mjrovai/UNIFEI-IESTI01-TinyML-2022.1>
  - Prof. Vijay Janapa Reddi - CS249r: Tiny Machine Learning, Applied Machine Learning on Embedded IoT Devices, Harvard
    - <https://sites.google.com/g.harvard.edu/tinyml/home>
  - Prof. Rahul Mangharam – ESE3600: Tiny Machine Learning, Univ. of Pennsylvania
    - <https://tinyml.seas.upenn.edu/#>
  - Prof. Brian Plancher - Harvard CS249r: Tiny Machine Learning (TinyML), Barnard College, Columbia University
    - [https://a2r-lab.org/courses/cs249r\\_tinyml/](https://a2r-lab.org/courses/cs249r_tinyml/)

15

25

# References

- Additional references from where information and other teaching materials were gathered include:
  - Applications & Deploy textbook: “TinyML” by Pete Warden, Daniel Situnayake
    - <https://www.oreilly.com/library/view/tinyml/9781492052036/>
  - Deploy textbook “TinyML Cookbook” by Gian Marco Iodice
    - <https://github.com/PacktPublishing/TinyML-Cookbook>
  - Jason Brownlee
    - <https://machinelearningmastery.com/>
  - TinyMLedu
    - <https://tinyml.seas.harvard.edu/>
  - Professional Certificate in Tiny Machine Learning (TinyML) – edX/Harvard
    - <https://www.edx.org/professional-certificate/harvardx-tiny-machine-learning>
  - Introduction to Embedded Machine Learning - Coursera/Edge Impulse
    - <https://www.coursera.org/learn/introduction-to-embedded-machine-learning>
  - Computer Vision with Embedded Machine Learning - Coursera/Edge Impulse
    - <https://www.coursera.org/learn/computer-vision-with-embedded-machine-learning>

16

26