
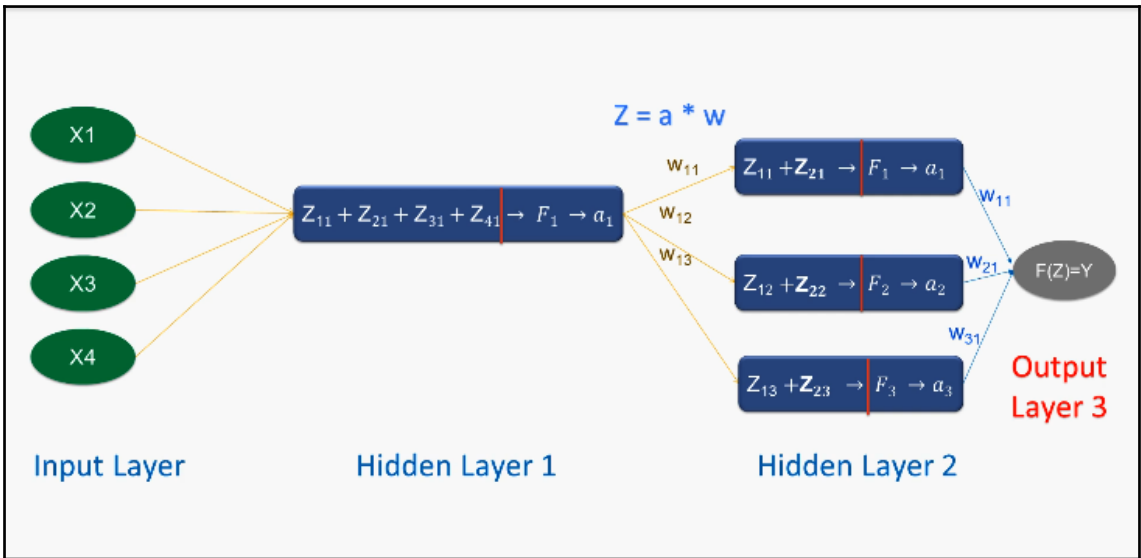


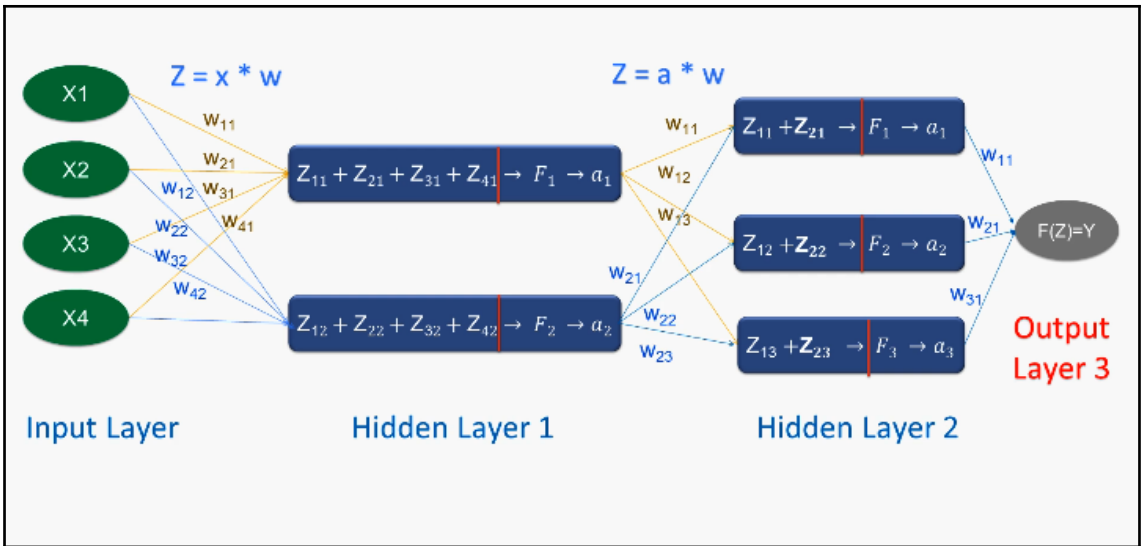
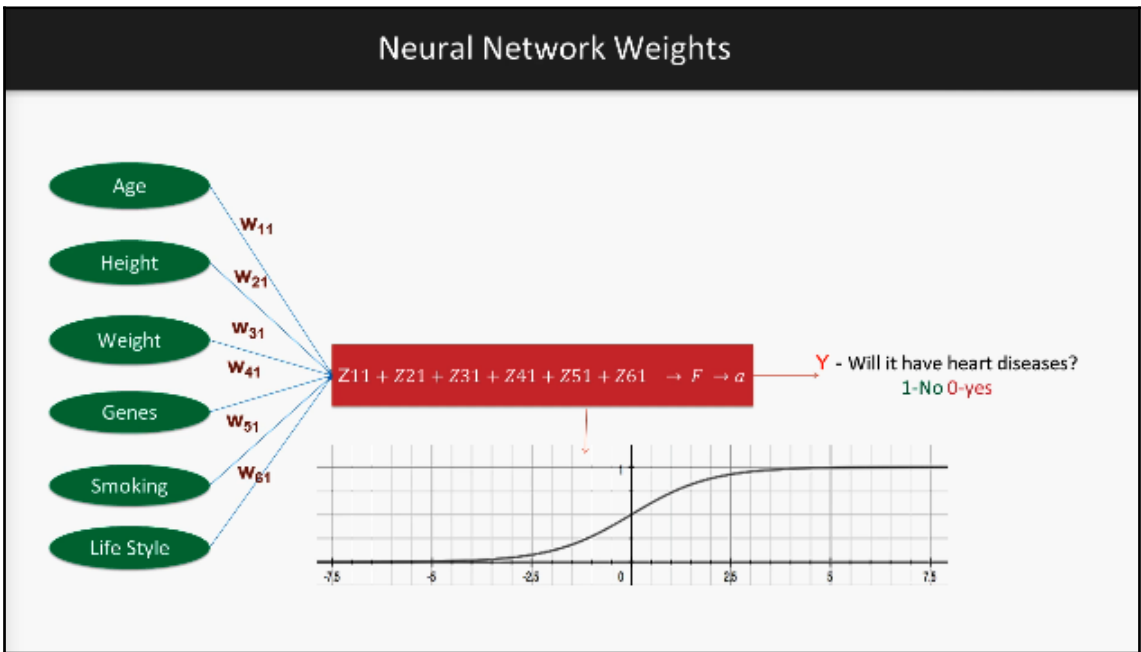



# Chapter 1: Introduction to Computer Vision and Training Neural Networks

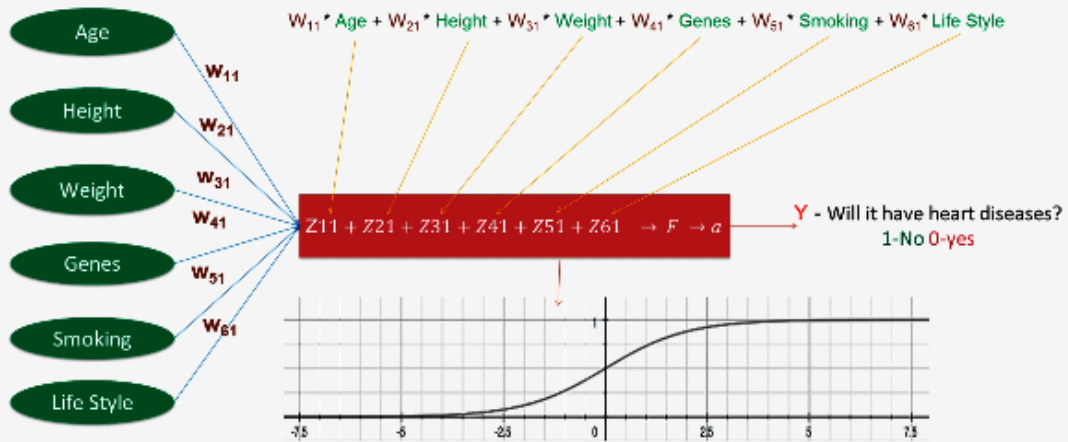




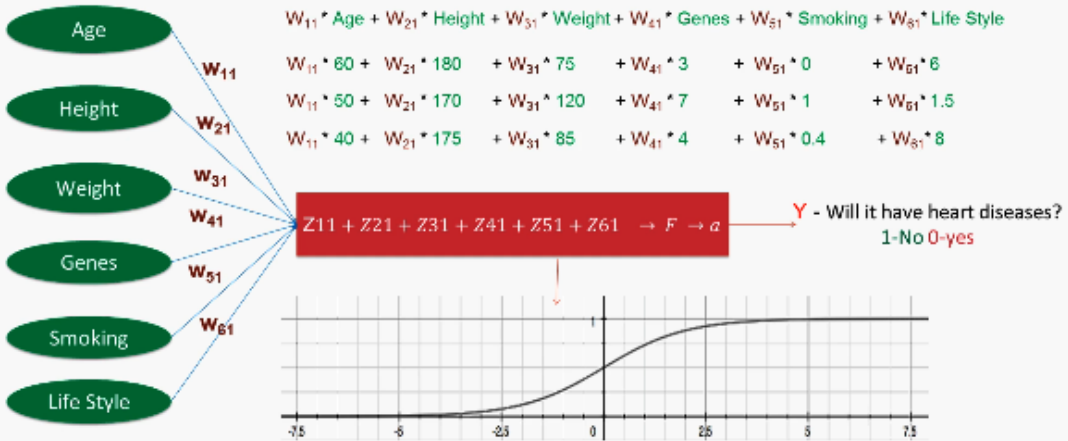





## Neural Network Weights



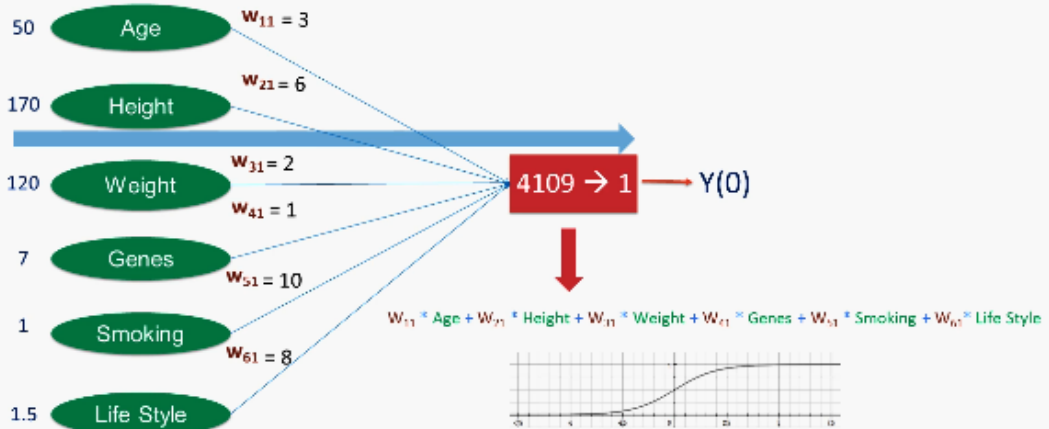
## Neural Network Weights



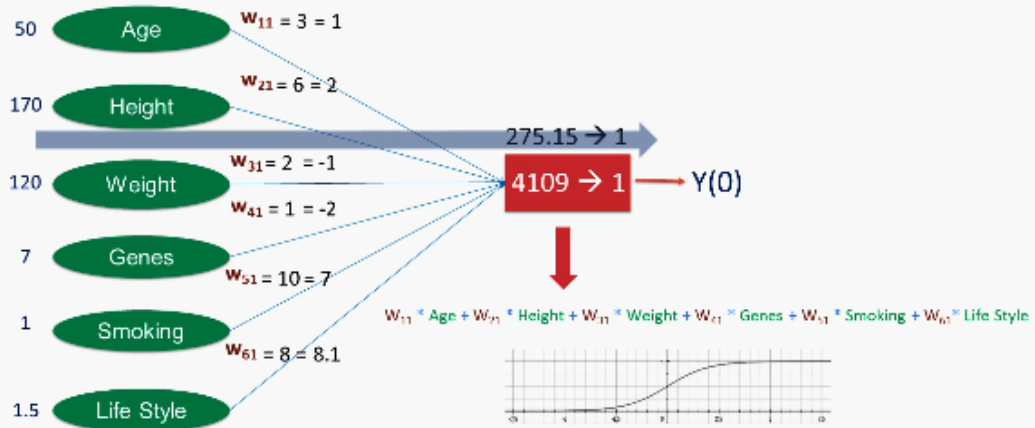
## Neural Network Weights



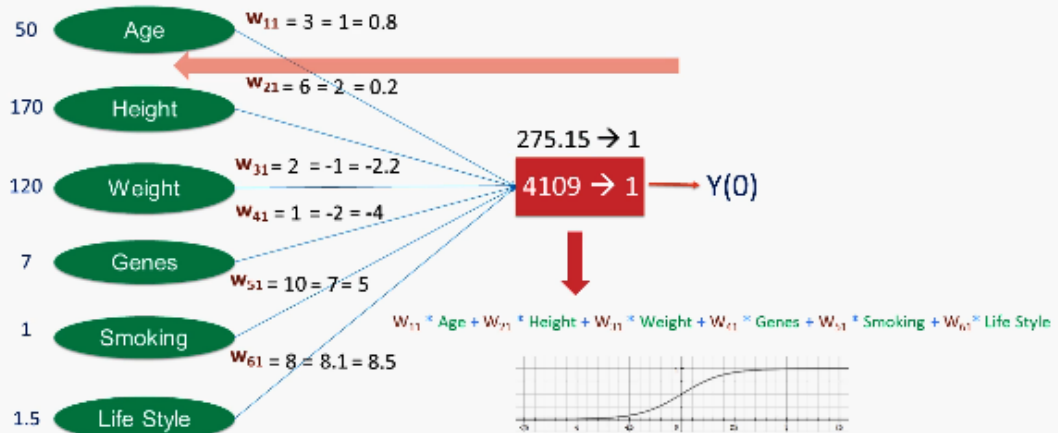
## Learning Neural Network Weights




## Learning Neural Network Weights



## Learning Neural Network Weights








---

# High Variance

Low Error Rate on **Training Data Set**

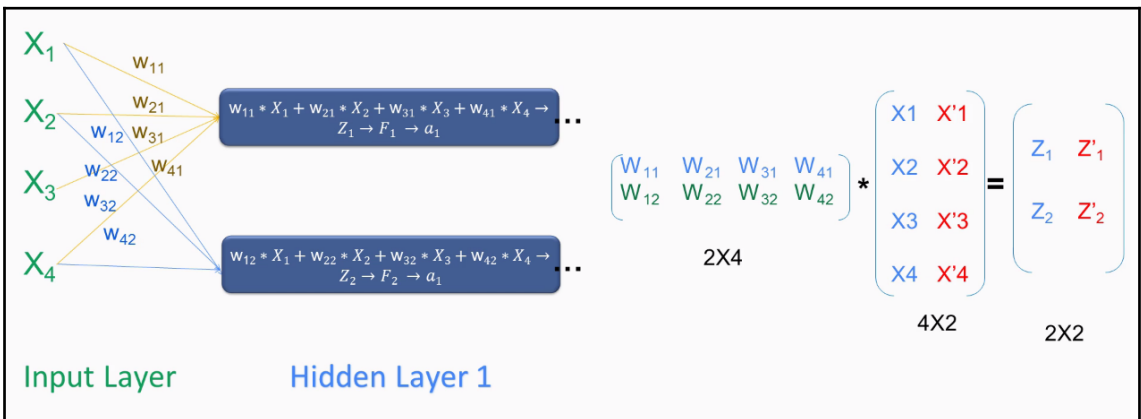

High Error Rate on **Dev Data Set**

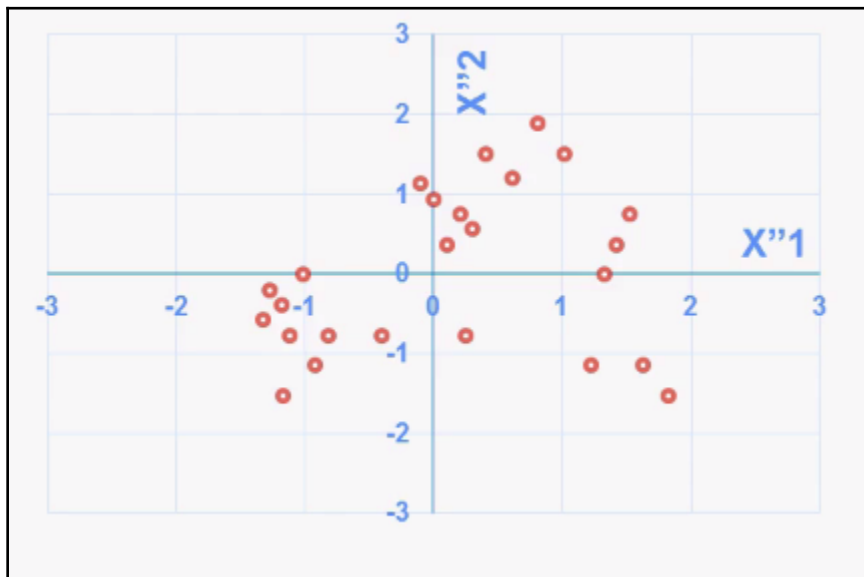
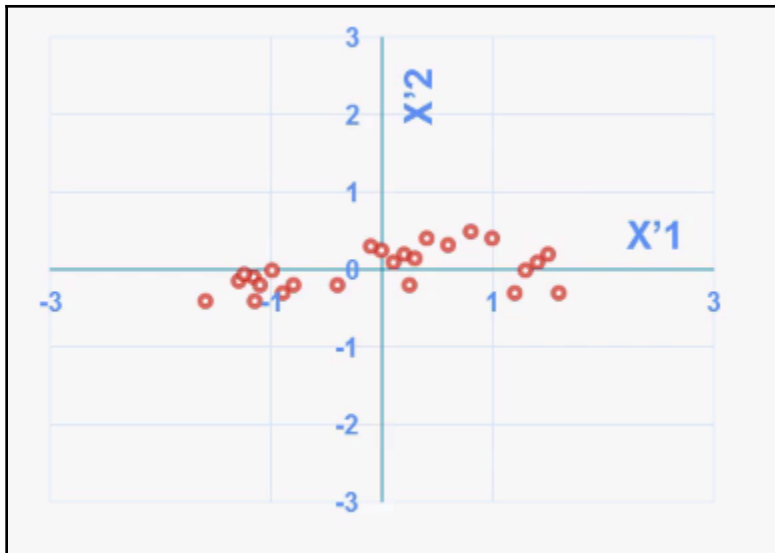




Both ☹️



High Error Rate on Training Data Set


Even Higher Error Rate on Dev Data Set











### RMSProp

$$S_0 = 0$$

$$S_t = \beta * S_{t-1} + (1 - \beta) * \left( \frac{dJ}{dW_{xy}} \right)^2$$


$$\cancel{w_{xy} = w_{xy} - \alpha \frac{dJ}{dW_{xy}}} \quad w_{xy} = w_{xy} - \alpha * \frac{\frac{dJ}{dW_{xy}}}{\sqrt{S_t}}$$

$$V_0 = 0 \quad S_0 = 0$$

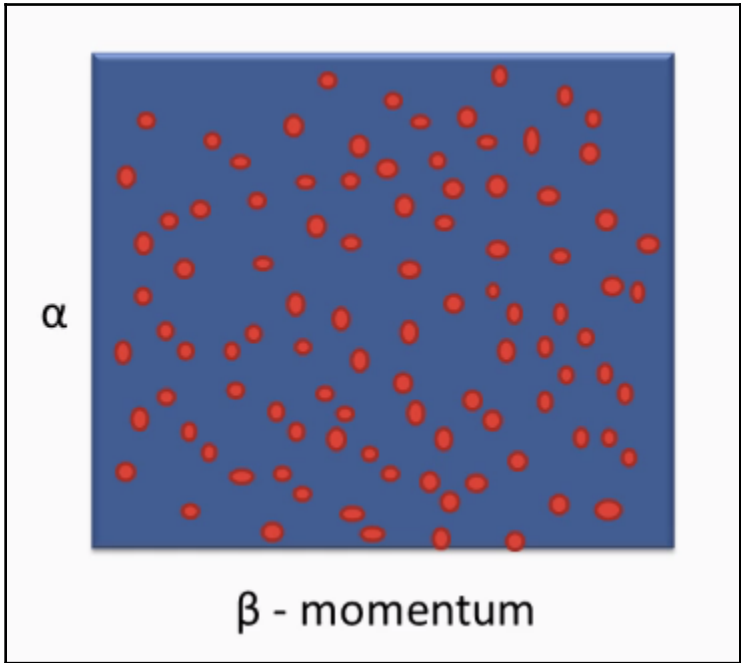
$$V_t = \beta * V_{t-1} + (1 - \beta) * \frac{dJ}{dW_{xy}} \quad S_t = \beta_2 * S_{t-1} + (1 - \beta_2) * \left( \frac{dJ}{dW_{xy}} \right)^2$$


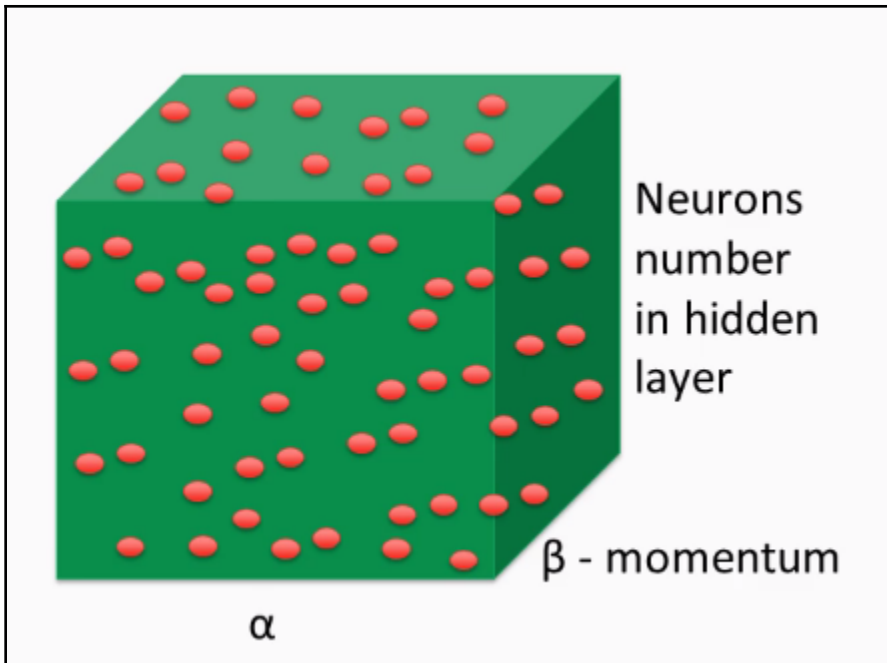
---

$$w_{xy} = w_{xy} - \alpha \frac{df}{dW_{xy}} \quad w_{xy} = w_{xy} - \alpha * \frac{V_t}{\sqrt{S_t}}$$

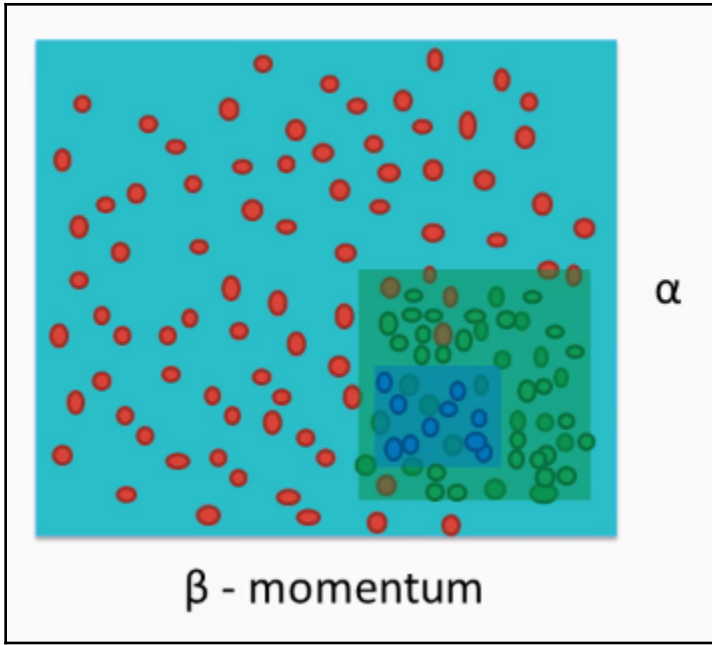


```
ThreadLocalRandom.current().nextDouble(-4, 0);
```









Gray Scale 0(white) → 255(black)

14,211,0,255,23,45,11

05,255,1,255,10,17,23

77,167,9,112,56,16,90

45,245,0,145,22,55,48


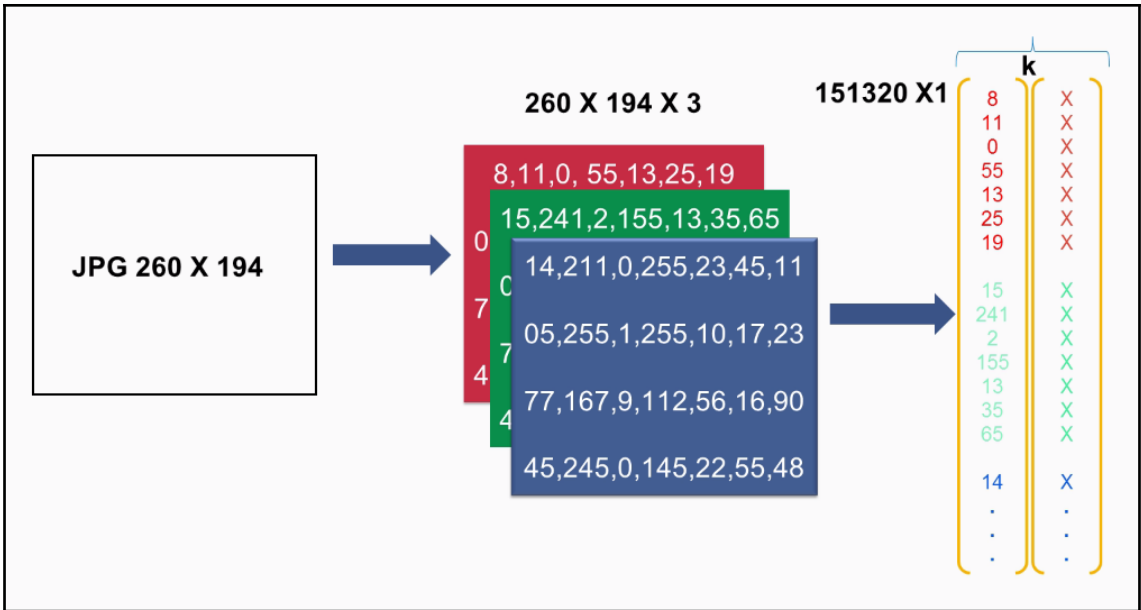
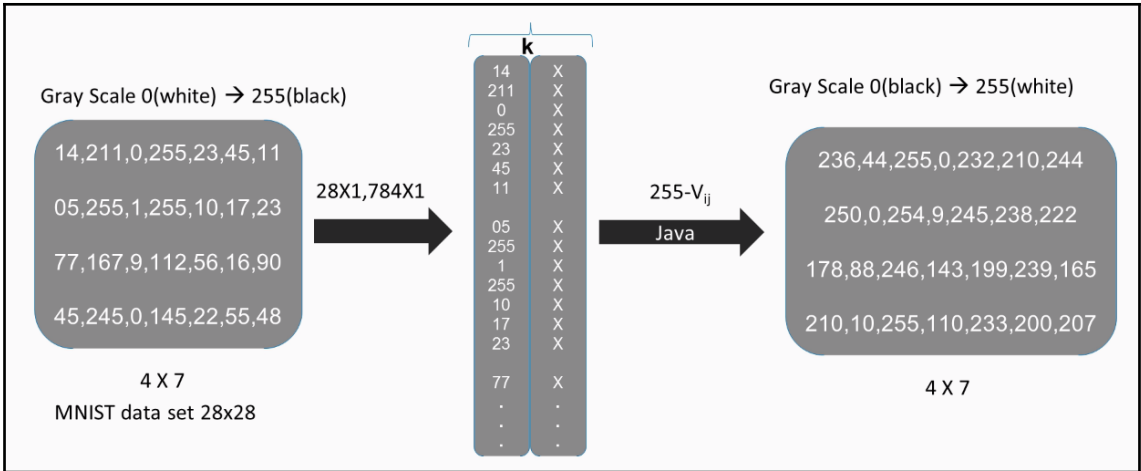
28X1,784X1

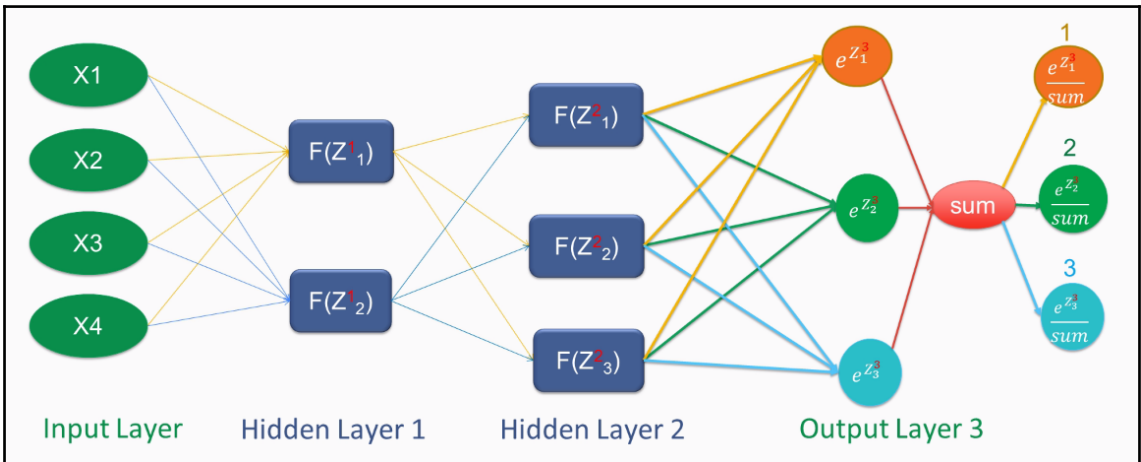
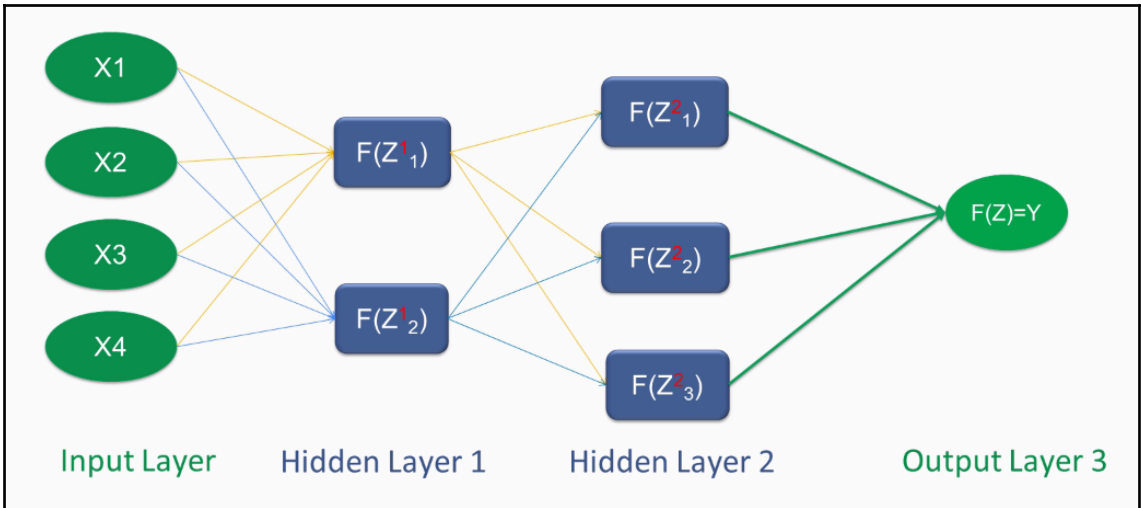



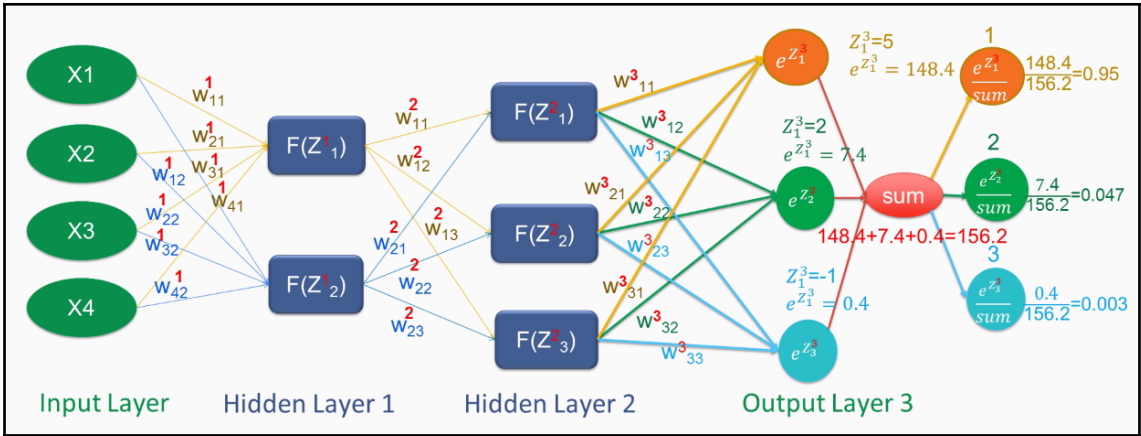
14  
211  
0  
255  
23  
45  
11  
  
05  
255  
1  
255  
10  
17  
23  
  
77  
.  
.  
.

4 X 7

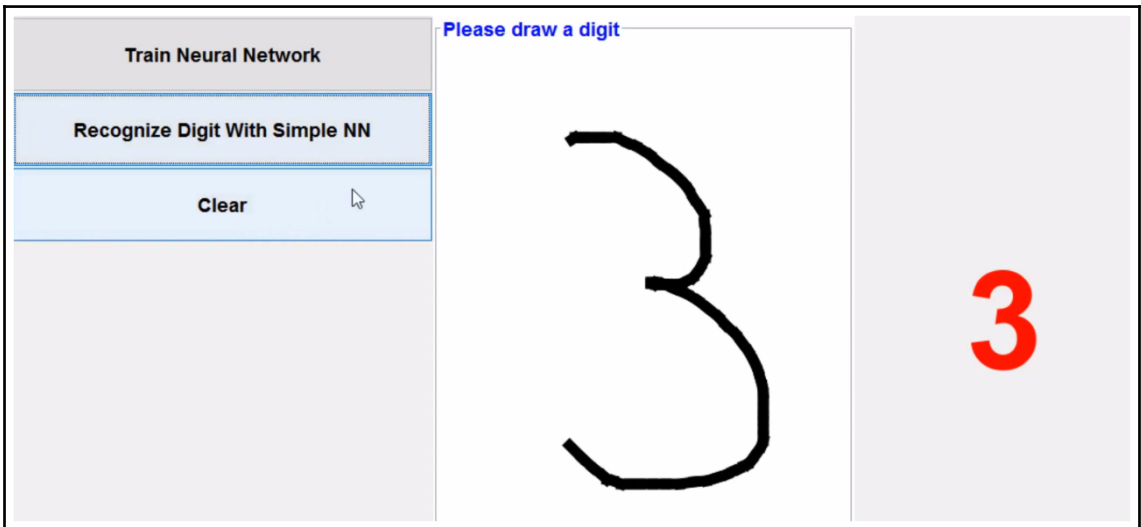
MNIST data set 28x28








```
# of classes:    10
Accuracy:       0.9731
Precision:      0.9731
Recall:         0.9728
F1 Score:       0.9729
Precision, recall & F1: macro-averaged (equally weighted avg. of 10 classes)
=====
[2018-04-07 01:55:04]Saving model at C:\devSources\ComputerVision\HandWrittenDigitReco
[2018-04-07 01:55:04]Congratulations,the desired score found,!
[2018-04-07 01:55:04]*****Example finished*****
Process finished with exit code 0
```



---

# Chapter 2: Convolutional Neural Network Architectures





255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10

255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10

 $*$ 

1	0	-1
1	0	-1
1	0	-1

Vertical Filter

255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10

 $*$ 

1	0	-1
1	0	-1
1	0	-1

Vertical Filter

255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10

 $*$ 

1	0	-1
1	0	-1
1	0	-1

Vertical Filter



255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10

\*

1	0	-1
1	0	-1
1	0	-1

Vertical Filter

$1 * 255 + 0 * 255 + -1 * 255 +$   
 $1 * 255 + 0 * 255 + -1 * 255 +$   
 $1 * 255 + 0 * 255 + -1 * 255 = 0$

1 * 255	0 * 255	-1 * 255	10	10	10
1 * 255	0 * 255	-1 * 255	10	10	10
1 * 255	0 * 255	-1 * 255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10

\*

1	0	-1
1	0	-1
1	0	-1

0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0

$1 * 255 + 0 * 255 + -1 * 10 +$   
 $1 * 255 + 0 * 255 + -1 * 10 +$   
 $1 * 255 + 0 * 255 + -1 * 10 = 735$

255	1 * 255	0 * 255	-1 * 10	10	10
255	1 * 255	0 * 255	-1 * 10	10	10
255	1 * 255	0 * 255	-1 * 10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10

\*

1	0	-1
1	0	-1
1	0	-1

0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0

$1 * 255 + 0 * 255 + -1 * 10 +$   
 $1 * 255 + 0 * 255 + -1 * 10 +$   
 $1 * 255 + 0 * 255 + -1 * 10 = 735$

255	255	1 * 255	0 * 10	-1 * 10	10
255	255	1 * 255	0 * 10	-1 * 10	10
255	255	1 * 255	0 * 10	-1 * 10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10

\*

1	0	-1
1	0	-1
1	0	-1

0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0

$1 * 10 + 0 * 10 + -1 * 10 +$   
 $1 * 10 + 0 * 10 + -1 * 10 +$   
 $1 * 10 + 0 * 10 + -1 * 10 = 0$

255	255	255	1 * 10	0 * 10	-1 * 10
255	255	255	1 * 10	0 * 10	-1 * 10
255	255	255	1 * 10	0 * 10	-1 * 10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10

\*

1	0	-1
1	0	-1
1	0	-1

0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0

$1 * 255 + 0 * 255 + -1 * 255 +$   
 $1 * 255 + 0 * 255 + -1 * 255 +$   
 $1 * 255 + 0 * 255 + -1 * 255 = 0$

255	255	255	10	10	10
1 * 255	0 * 255	-1 * 255	10	10	10
1 * 255	0 * 255	-1 * 255	10	10	10
1 * 255	0 * 255	-1 * 255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10

\*

1	0	-1
1	0	-1
1	0	-1

0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0

$1 * 255 + 0 * 255 + -1 * 10 +$   
 $1 * 255 + 0 * 255 + -1 * 10 +$   
 $1 * 255 + 0 * 255 + -1 * 10 = 735$

255	255	255	10	10	10
255	1 * 255	0 * 255	-1 * 10	10	10
255	1 * 255	0 * 255	-1 * 10	10	10
255	1 * 255	0 * 255	-1 * 10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10

\*

1	0	-1
1	0	-1
1	0	-1

0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0

255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10
255	255	255	10	10	10

8 X 6

\*

1	0	-1
1	0	-1
1	0	-1

3 X 3

=

0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0
0	735	735	0

6 X 4

255	255	255	255	255	255
255	255	255	255	255	255
255	255	255	255	255	255
255	255	255	255	255	255
10	10	10	10	10	10
10	10	10	10	10	10
10	10	10	10	10	10
10	10	10	10	10	10

**8 X 6**

**\***

1	1	1
0	0	0
-1	-1	-1

**3 X 3**

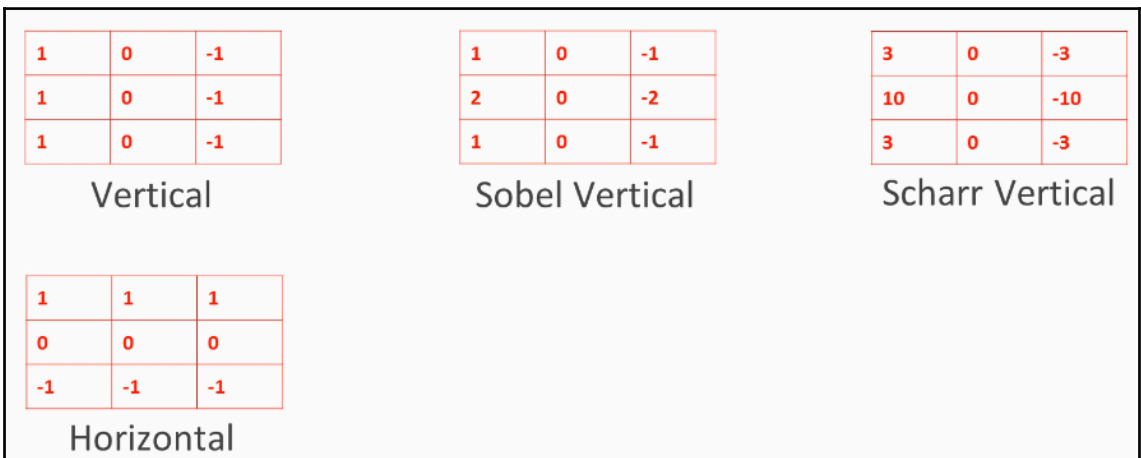


**=**

0	0	0	0
0	0	0	0
735	735	735	735
735	735	735	735
0	0	0	0
0	0	0	0

**6 X 4**

1	255	0	255	-1	255	255	255	255
1	255	0	255	-1	255	255	255	255
1	255	0	255	-1	255	255	255	255
1	255	1	255	1	255	255	255	255
0	10	0	10	0	10	10	10	10
-1	10	-1	10	-1	10	10	10	10
	10		10		10	10	10	10
	10		10		10	10	10	10

1	255	0	255	-1	255	255	255	255
1	255	0	255	-1	255	255	255	255
1	255	0	255	-1	255	255	255	255
1	255	1	255	1	255	255	255	255
0	10	0	10	0	10	10	10	10
-1	10	-1	10	-1	10	10	10	10
	10		10		10	10	10	10
	10		10		10	10	10	10



1	0	-1
1	0	-1
1	0	-1

Vertical

1	0	-1
2	0	-2
1	0	-1

Sobel Vertical

3	0	-3
10	0	-10
3	0	-3

Scharr Vertical

1	1	1
0	0	0
-1	-1	-1

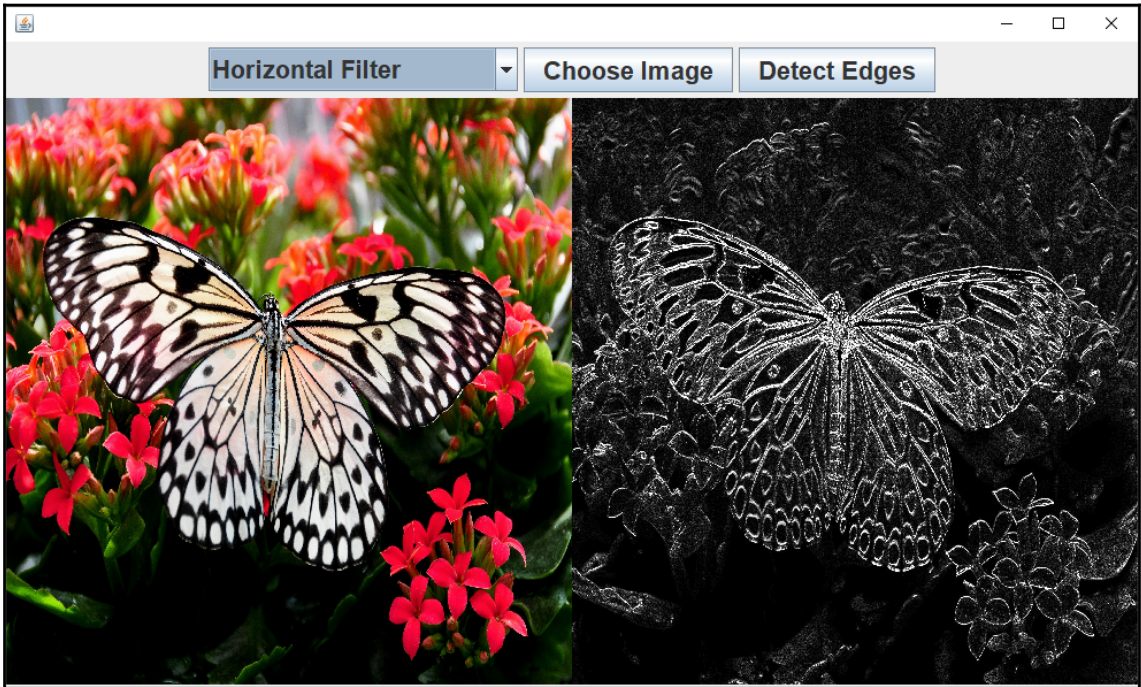
Horizontal


1	2	1
0	0	0
-1	-2	-1

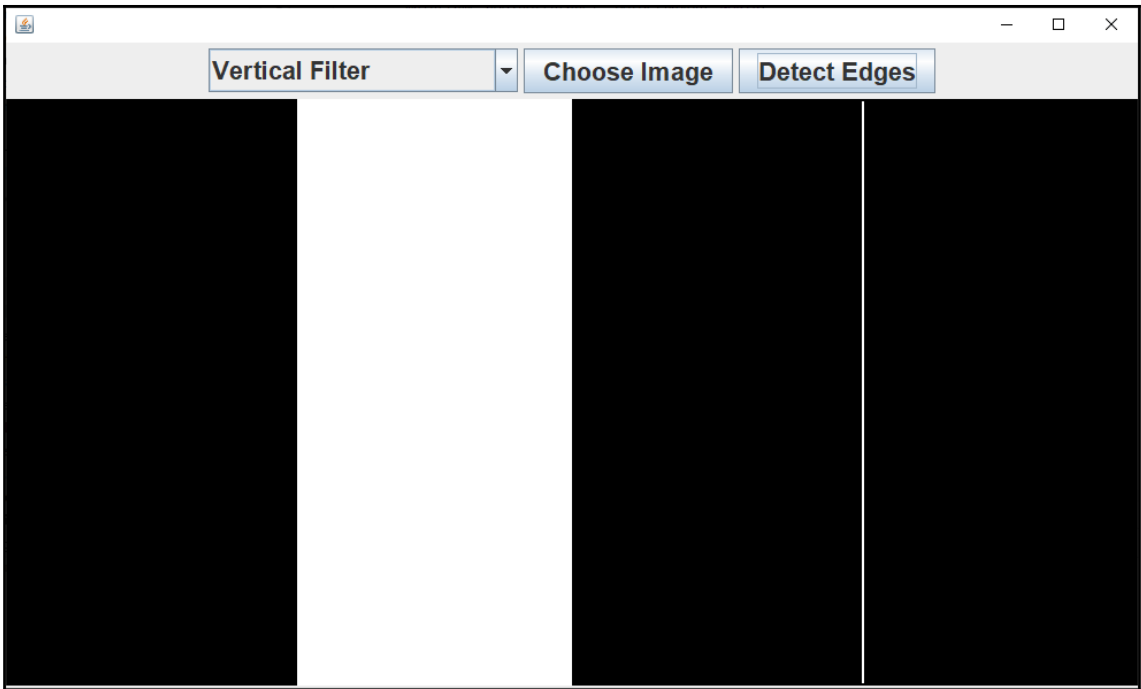
Sobel Horizontal

3	10	3
0	0	0
-3	-10	-3

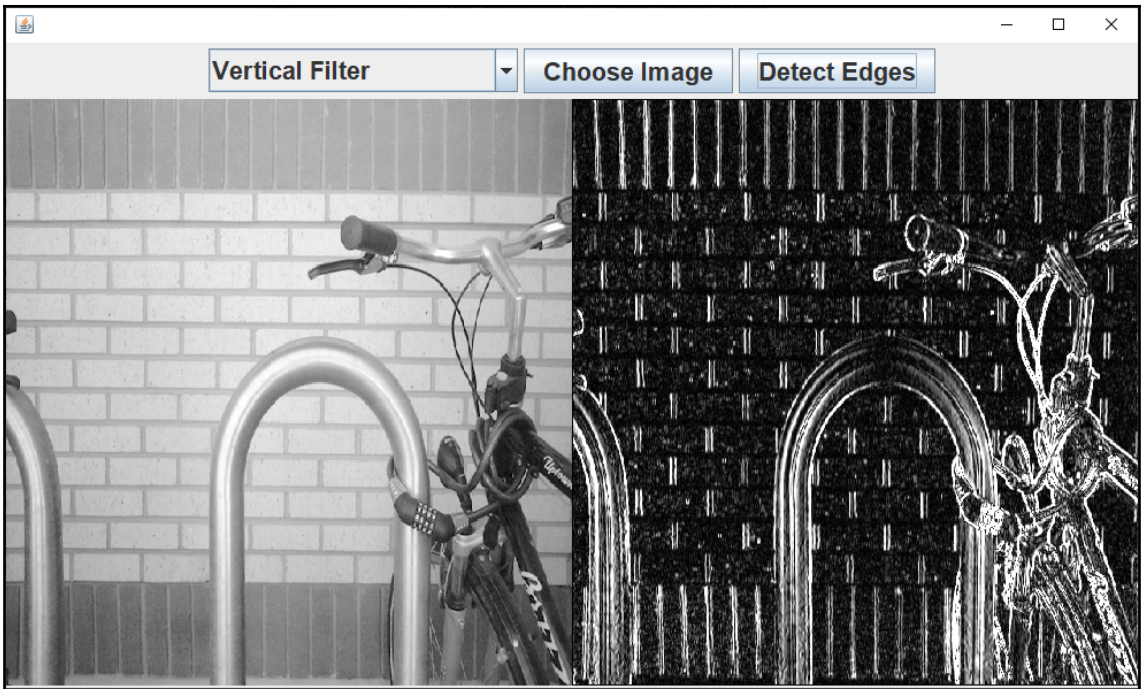
Scharr Horizontal

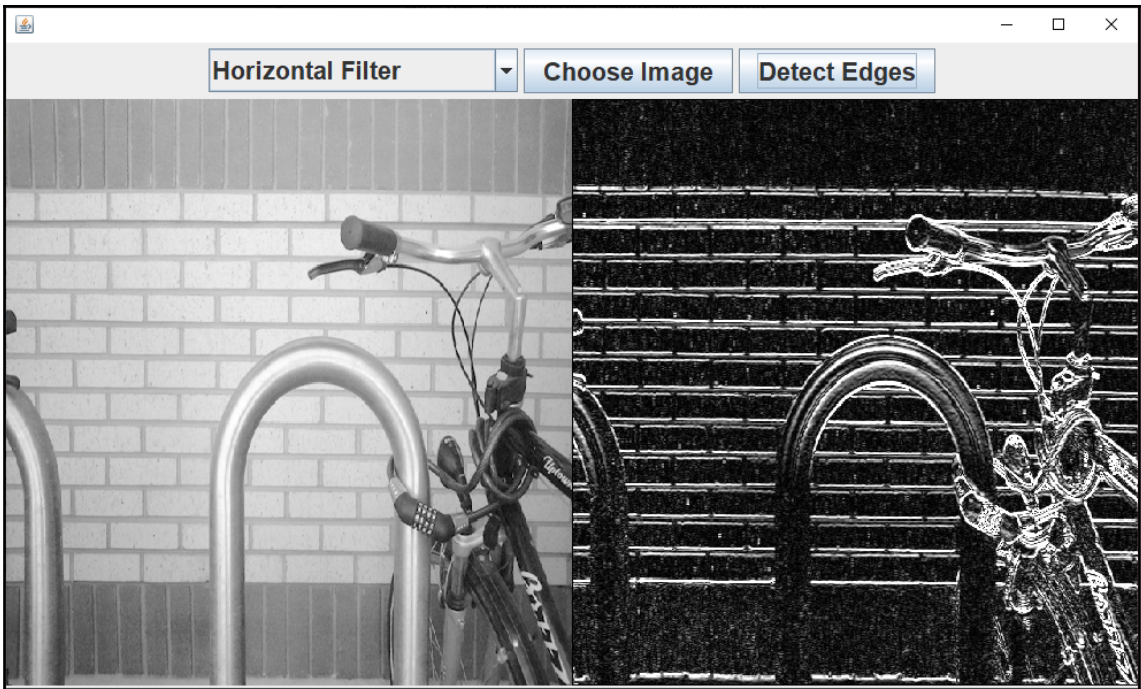


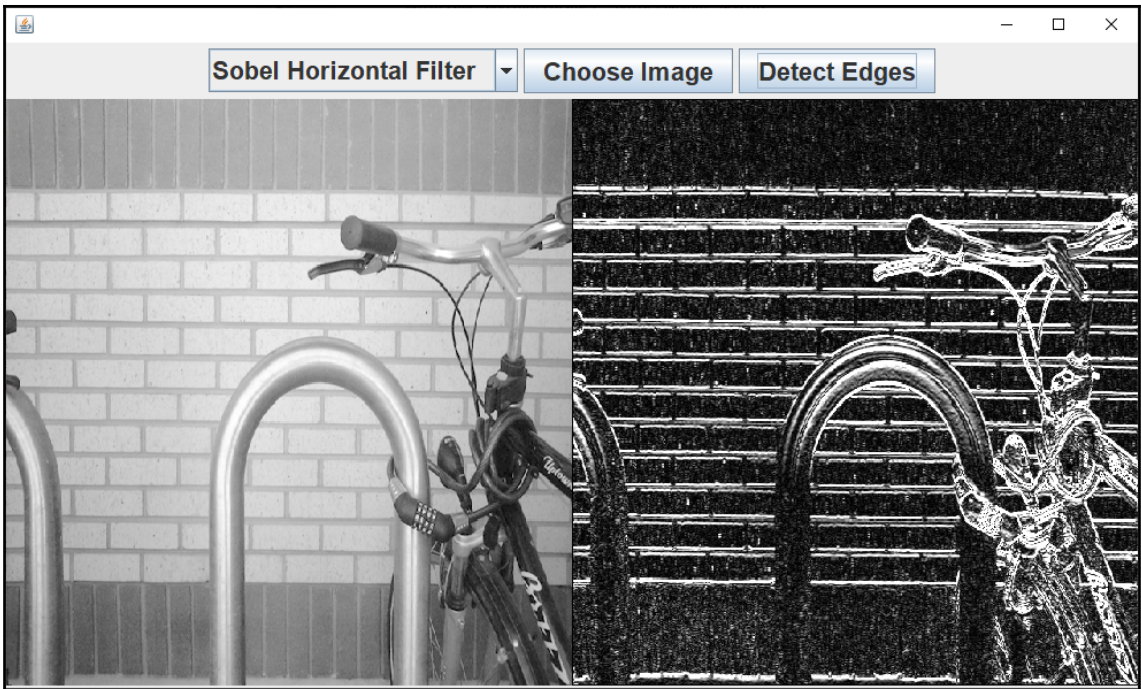





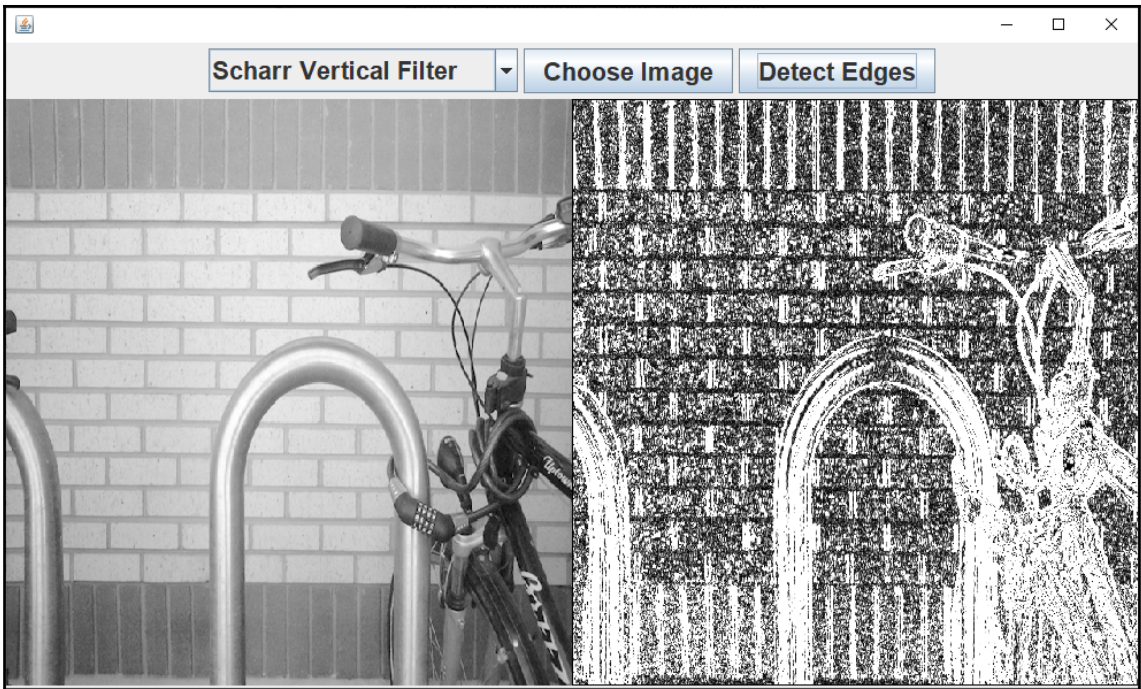


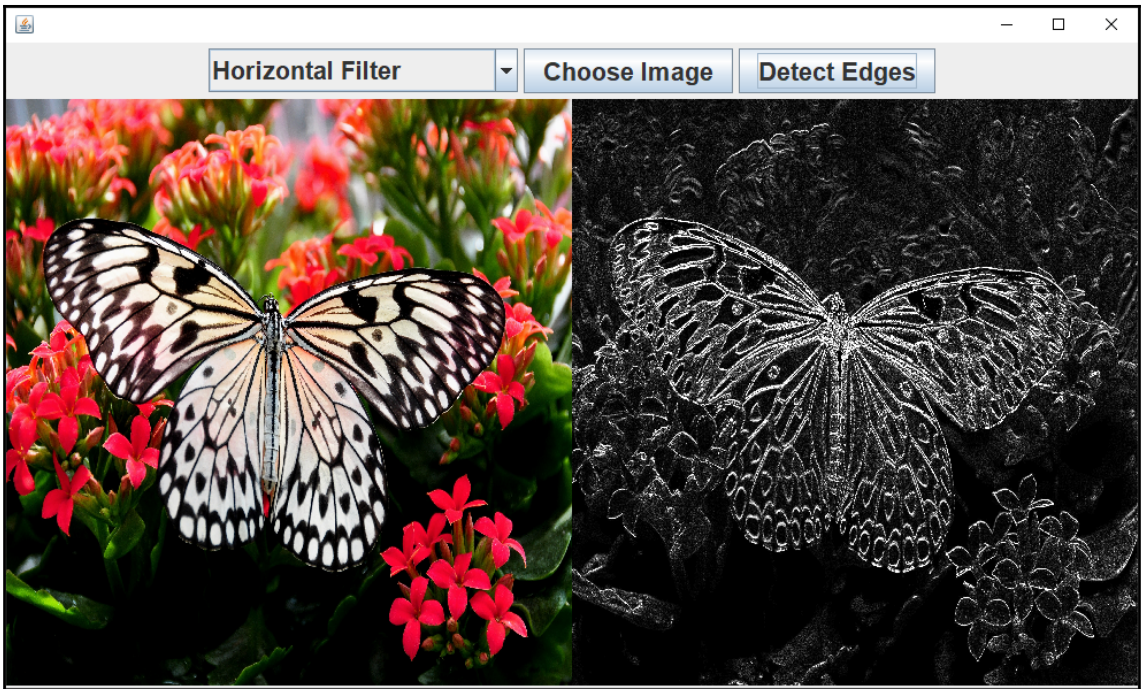


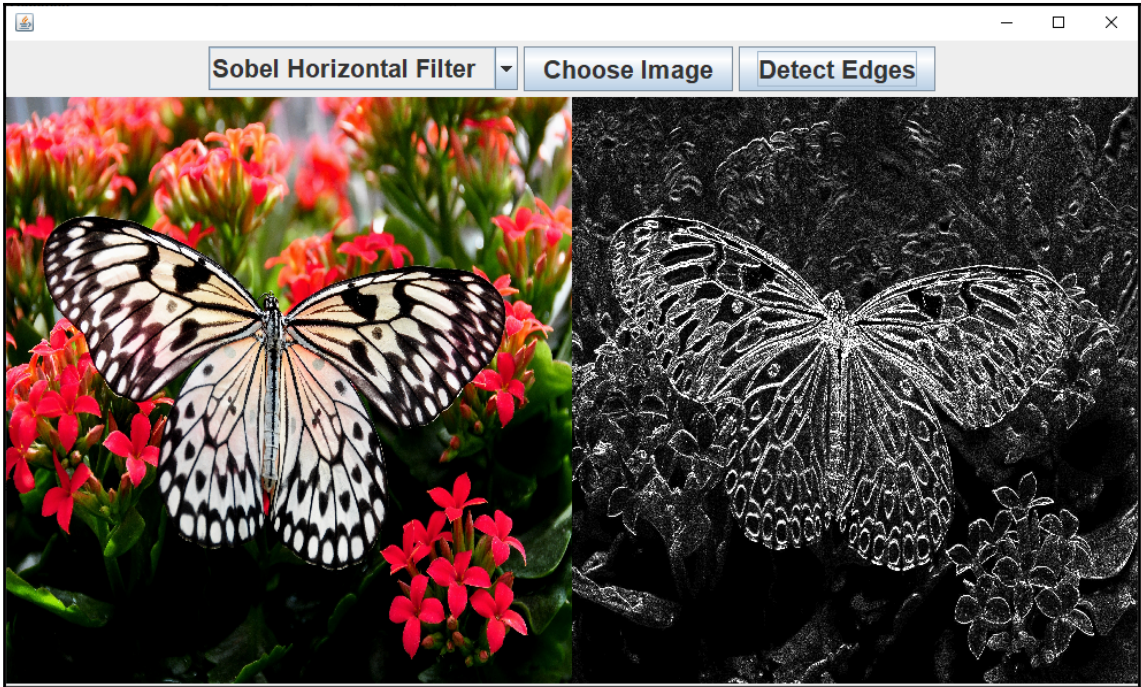


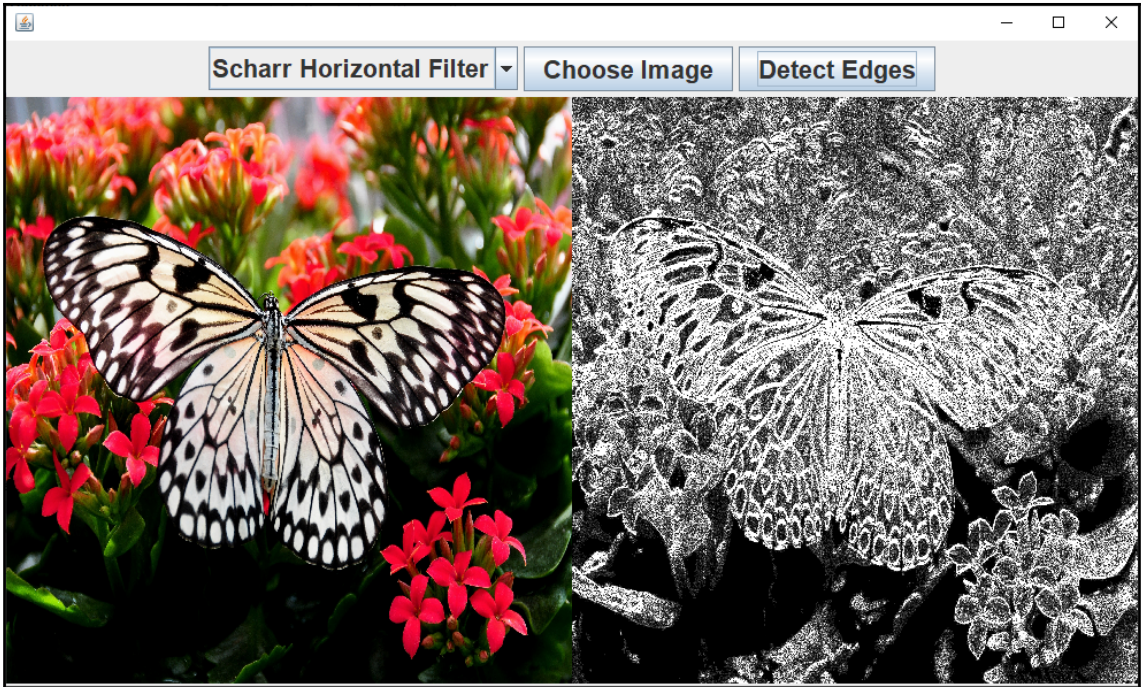




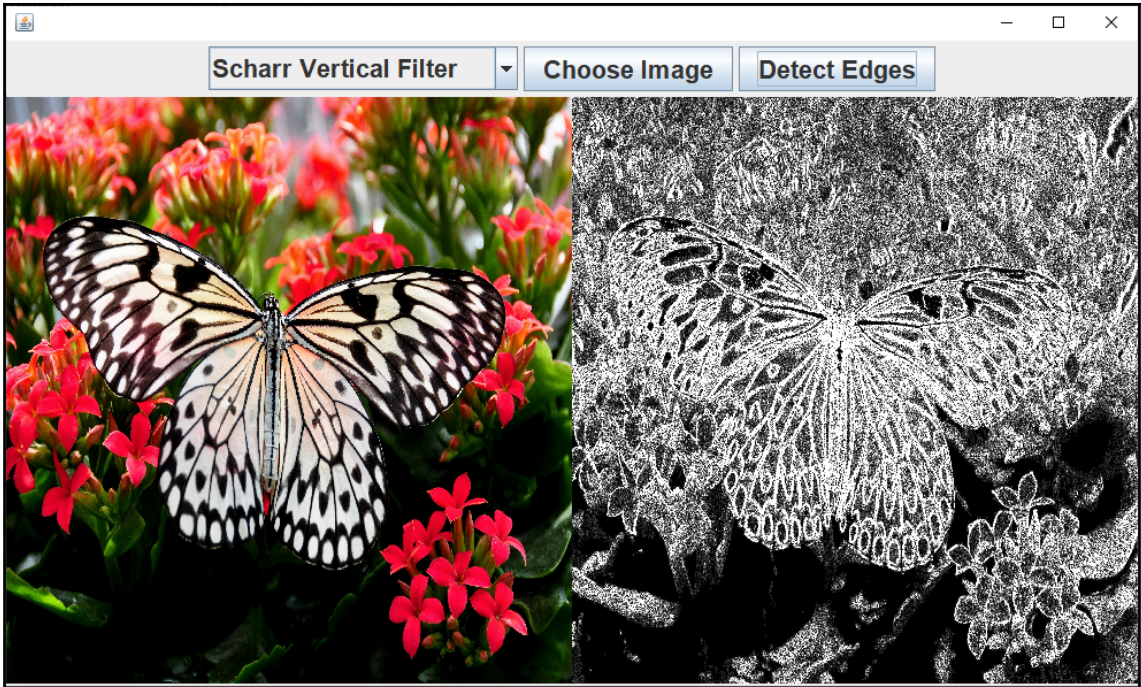












1	1	1
0	0	0
-1	-1	-1

Horizontal

1	0	-1
1	0	-1
1	0	-1

Vertical

1	0	-1
2	0	-2
1	0	-1

Sobel Vertical

3	0	-3
10	0	-10
3	0	-3

Scharr Vertical

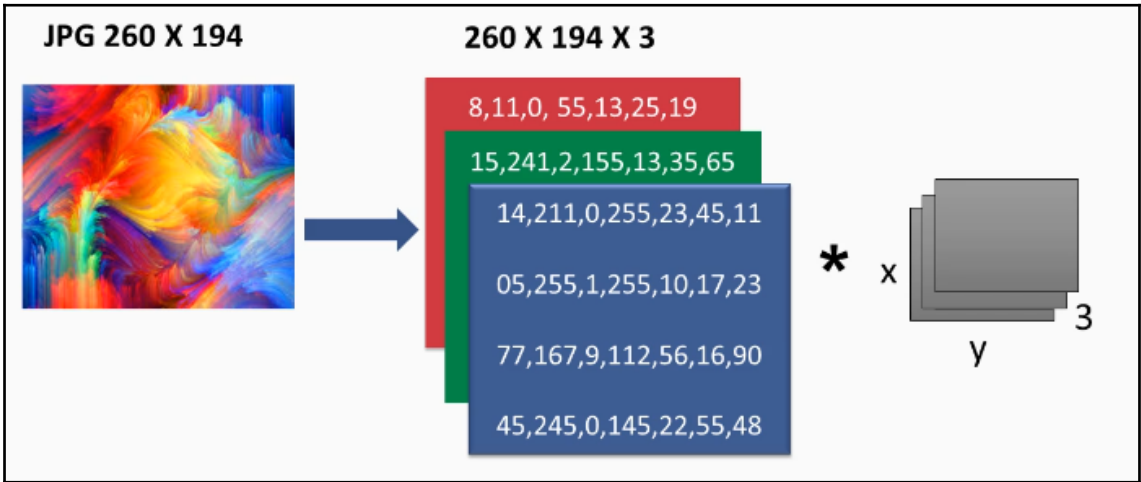

W11	W12	W13
W21	W22	W23
W31	W32	W33

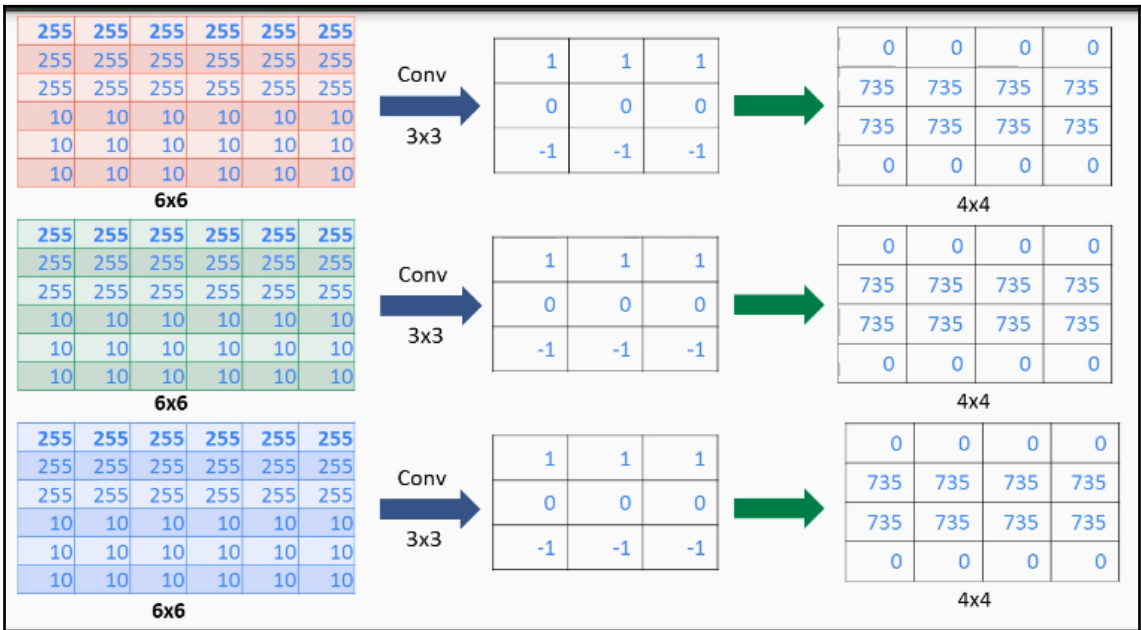

1	2	1
0	0	-
-1	-2	-1

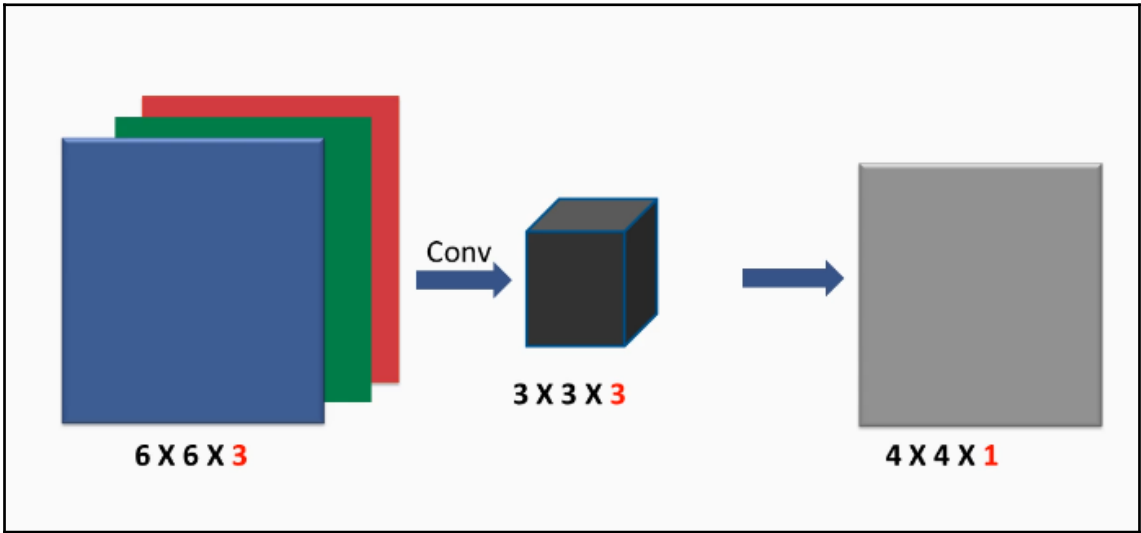
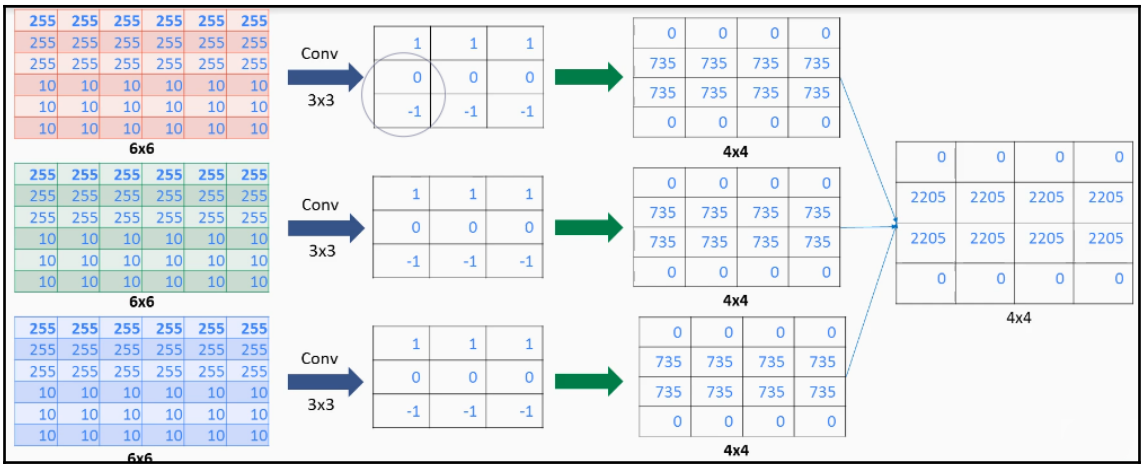
Sobel Horizontal



3	0	-3
10	0	-10
3	0	-3




Scharr Horizontal

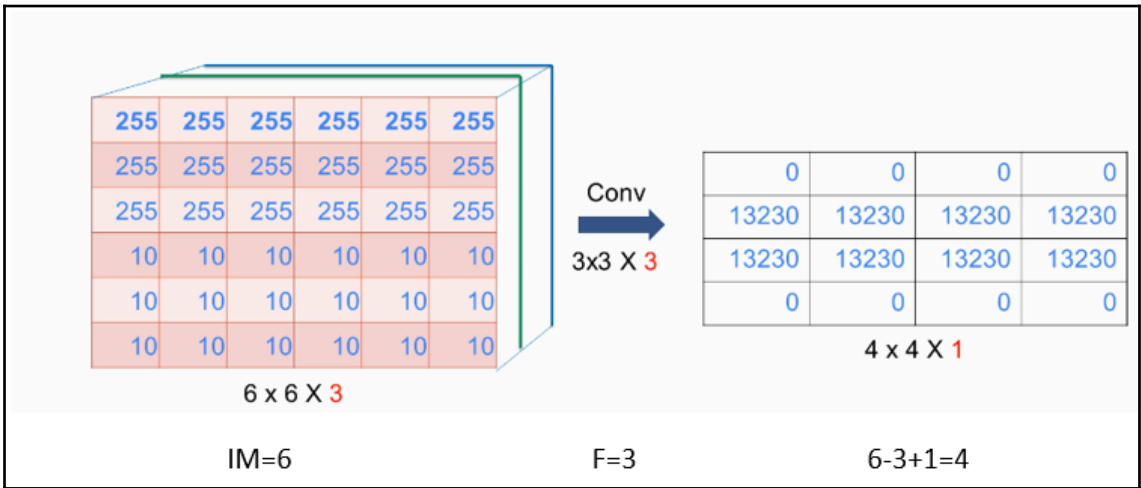
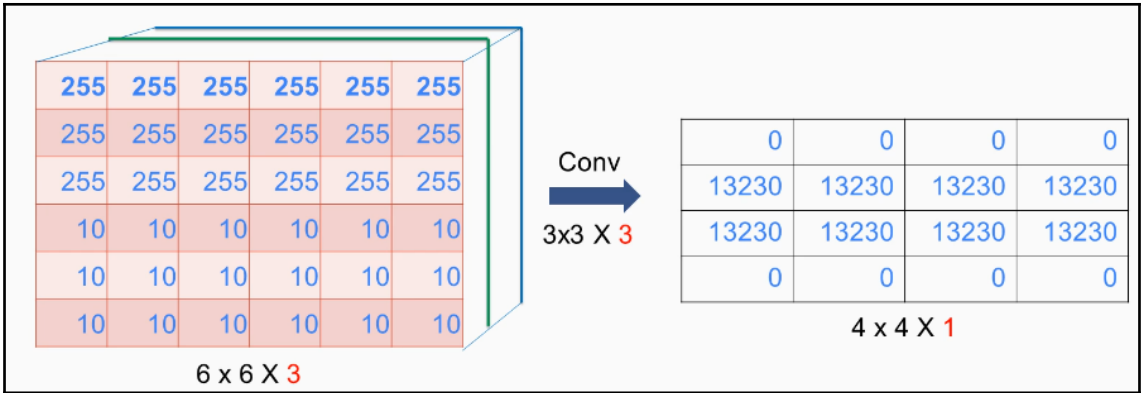


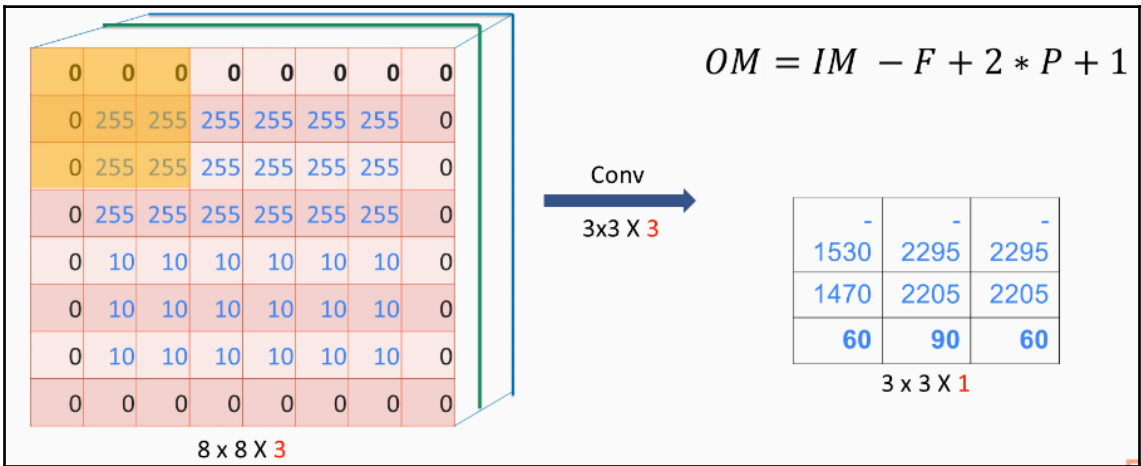
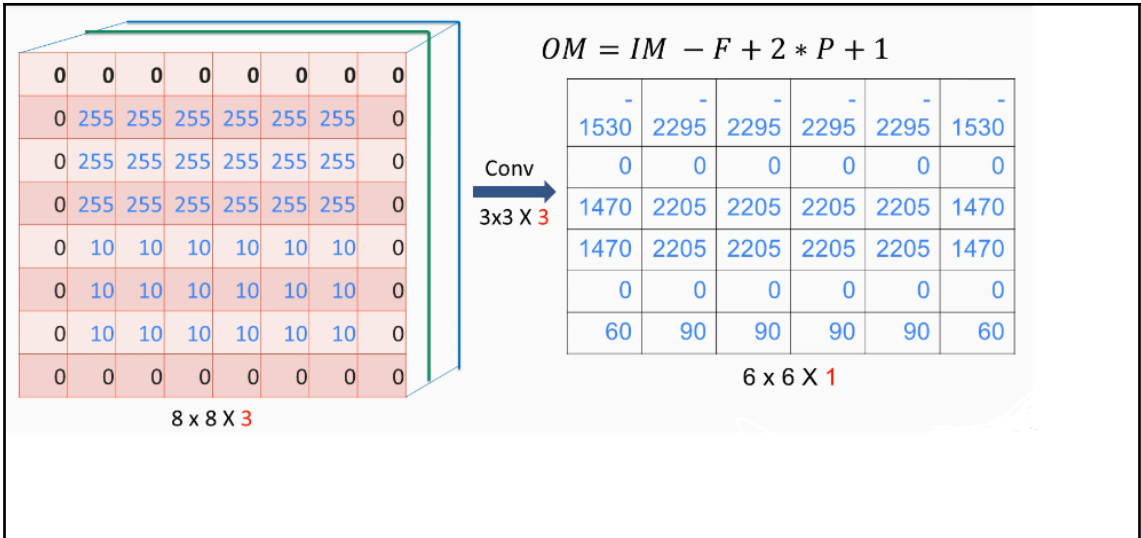





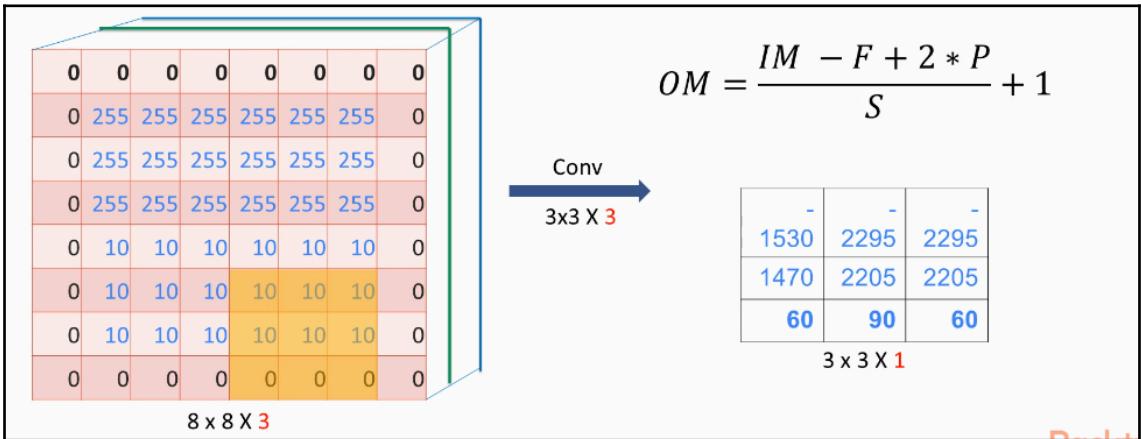



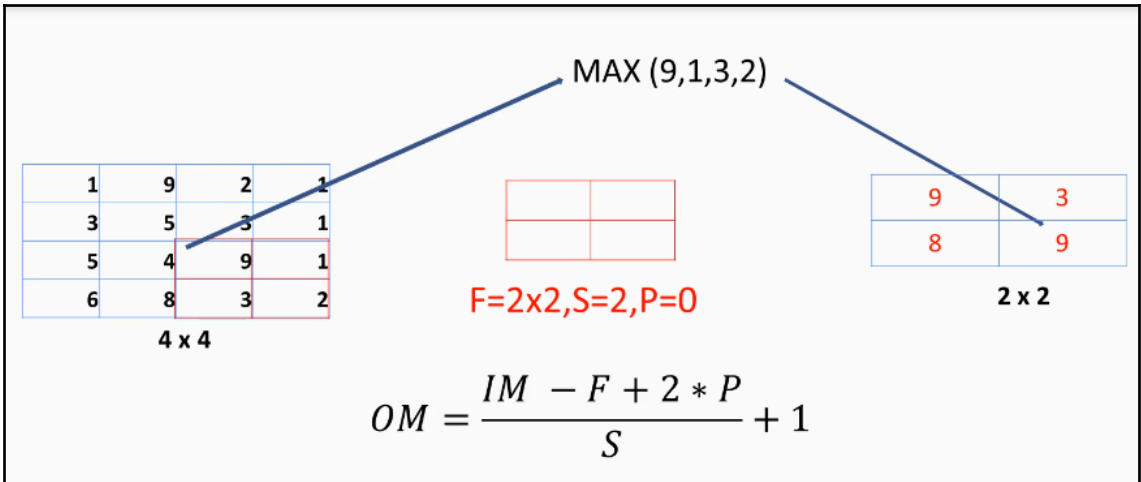





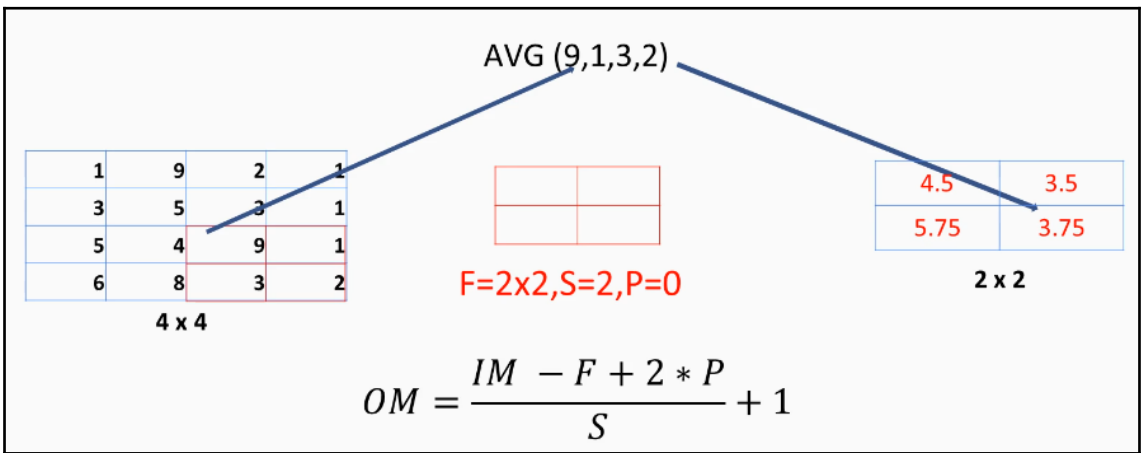


















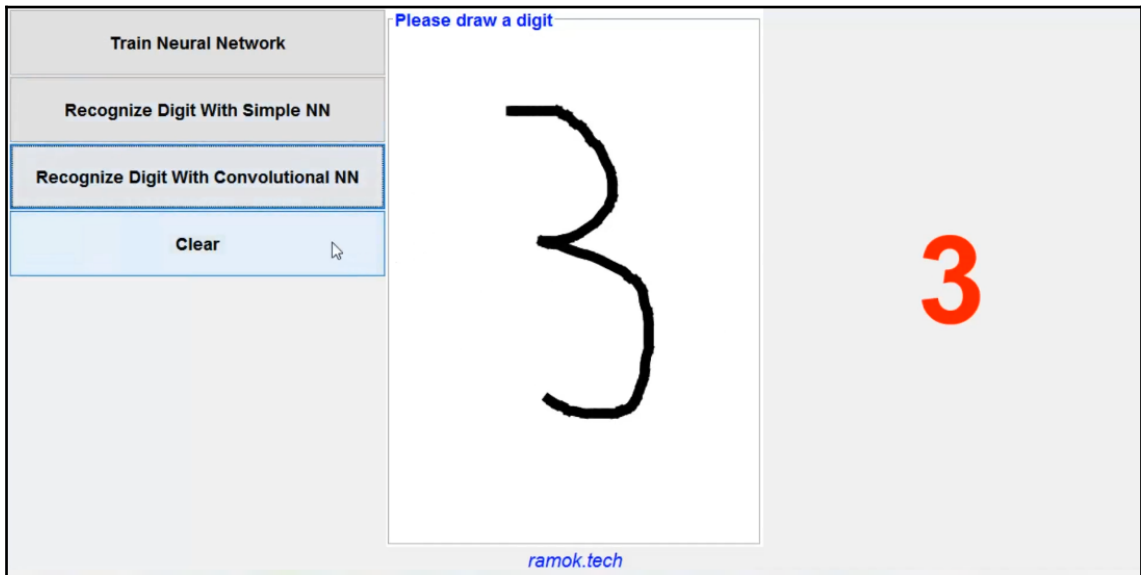
	Activation Matrix (Input)	Activation Size	Parameters Size
Image Matrix	28 X 28 X 1	784	0
Conv 5X5X20,S=1	24 X 24 X 20	11520	500
Max 2 X 2 S=2	12 X 12 X 20	2880	0
Conv 5X5X50,S=1	8 X 8 X 50	3200	1250
Max 2 X 2 S=2	4 X 4 X 50	800	0
FC1	128 X 1	128	102400
FC2	64 X 1	64	8192
OUT	10 X 1	10	640

```



[2018-04-26 22:52:41]Cores: [4]; Memory: [13.3GB];
[2018-04-26 22:52:41]Blas vendor: [OPENBLAS]
[2018-04-26 22:52:43]Reflections took 1269 ms to scan 116 urls, producing 2333 keys and 9715 values
[2018-04-26 22:52:43]Starting MultiLayerNetwork with WorkspaceModes set to [training: NONE; inference:
[2018-04-26 22:52:43]Reflections took 115 ms to scan 13 urls, producing 387 keys and 1538 values
[2018-04-26 22:52:44]Starting early stopping training
[2018-04-26 22:54:36]Completed training epoch 0
[2018-04-26 22:54:40]Accuracy at iteration0 0.0657
[2018-04-26 22:54:40]Score at epoch 0: 0.0343
[2018-04-26 22:59:33]Completed training epoch 1
[2018-04-26 22:59:52]Accuracy at iteration1 0.983
[2018-04-26 22:59:52]New best model: score = 0.017000000000000015, epoch = 1 (previous: score = 0.0343
[2018-04-26 23:05:16]Completed training epoch 2
[2018-04-26 23:05:39]Accuracy at iteration2 0.98
[2018-04-26 23:14:13]Completed training epoch 3
[2018-04-26 23:14:39]Accuracy at iteration3 0.9841
[2018-04-26 23:14:39]New best model: score = 0.0159000000000000025, epoch = 3 (previous: score = 0.0170
[2018-04-26 23:22:02]Completed training epoch 4
[2018-04-26 23:22:38]Accuracy at iteration4 0.9865
[2018-04-26 23:22:38]New best model: score = 0.0134999999999999956, epoch = 4 (previous: score = 0.0159
[2018-04-26 23:29:07]Completed training epoch 5
[2018-04-26 23:29:39]Accuracy at iteration5 0.9867
[2018-04-26 23:29:39]New best model: score = 0.0132999999999999979, epoch = 5 (previous: score = 0.0134
[2018-04-26 23:35:27]Completed training epoch 6
[2018-04-26 23:36:01]Accuracy at iteration6 0.988



```

```
ComputerVision > HandWrittenDigitRecognizer > src > main > resources > cnnTrainedModels > prevLog.txt
DigitRecognizerConvolutionalNeuralNetwork.java x AccuracyCalculator.java x prevLog.txt x
13 [2018-04-26 22:54:36]Completed training epoch 0
14 [2018-04-26 22:54:40]Accuracy at iteration0 0.9657
15 [2018-04-26 22:54:40]Score at epoch 0: 0.0343
16 [2018-04-26 22:59:33]Completed training epoch 1
17 [2018-04-26 22:59:52]Accuracy at iteration1 0.983
18 [2018-04-26 22:59:52]New best model: score = 0.017000000000000015, epoch = 1 (previous: score = 0.0343)
19 [2018-04-26 23:05:16]Completed training epoch 2
20 [2018-04-26 23:05:39]Accuracy at iteration2 0.98
21 [2018-04-26 23:14:13]Completed training epoch 3
22 [2018-04-26 23:14:39]Accuracy at iteration3 0.9841
23 [2018-04-26 23:14:39]New best model: score = 0.015900000000000025, epoch = 3 (previous: score = 0.0170)
24 [2018-04-26 23:22:02]Completed training epoch 4
25 [2018-04-26 23:22:38]Accuracy at iteration4 0.9865
26 [2018-04-26 23:22:38]New best model: score = 0.013499999999999956, epoch = 4 (previous: score = 0.0159)
27 [2018-04-26 23:29:07]Completed training epoch 5
28 [2018-04-26 23:29:39]Accuracy at iteration5 0.9867
29 [2018-04-26 23:29:39]New best model: score = 0.013299999999999979, epoch = 5 (previous: score = 0.0134)
30 [2018-04-26 23:35:27]Completed training epoch 6
31 [2018-04-26 23:36:01]Accuracy at iteration6 0.9898 |
32 [2018-04-26 23:36:01]New best model: score = 0.010199999999999987, epoch = 6 (previous: score = 0.0132)
```





---

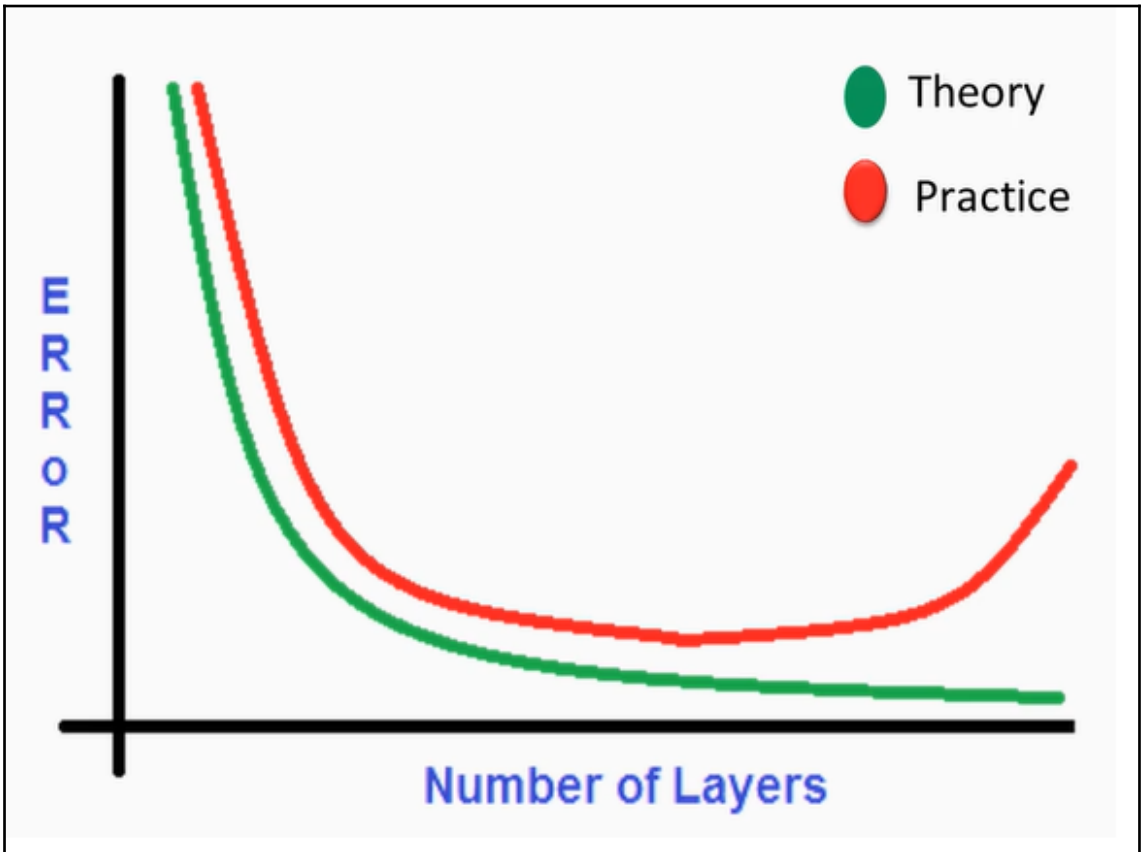

<b>Train Neural Network</b>	<p>Please draw a digit</p>  <p><i>ramok.tech</i></p>	
<b>Recognize Digit With Simple NN</b>		
<b>Recognize Digit With Convolutional NN</b>		
<b>Clear</b>		




<b>Train Neural Network</b>	<p>Please draw a digit</p>  <p><i>ramok.tech</i></p>	
<b>Recognize Digit With Simple NN</b>		
<b>Recognize Digit With Convolutional NN</b>		
<b>Clear</b>		

# Chapter 3: Transfer Learning and Deep CNN Architectures









16	22	18	44
111	34	41	78
1	21	17	76
45	52	78	91

**4x4x1**

\*

4

**1x1**

64	88	72	176
444	136	164	312
4	84	86	304
180	208	312	364

**4x4x1**

4 * 16	22	18	44
111	34	41	78
1	21	17	76
45	52	78	91

**4x4x1**

\*

4
---

**1x1**

64	88	72	176
444	136	164	312
4	84	86	304
180	208	312	364

**4x4x1**

16	4 * 22	18	44
111	34	41	78
1	21	17	76
45	52	78	91

**4x4x1**

\*

4
---

**1x1**

64	88	72	176
444	136	164	312
4	84	86	304
180	208	312	364

**4x4x1**

16	22	4 * 18	44
111	34	41	78
1	21	17	76
45	52	78	91

**4x4x1**

\*

4
---

**1x1**

64	88	72	176
444	136	164	312
4	84	86	304
180	208	312	364

**4x4x1**

16	22	18	4 * 44
111	34	41	78
1	21	17	76
45	52	78	91

**4x4x1**

\*


4
---



**1x1**



64	88	72	176
444	136	164	312
4	84	86	304
180	208	312	364


**4x4x1**


16	22	18	44	*	<div style="border: 1px solid black; display: inline-block; padding: 2px;">4</div> <b>1x1</b>	64	88	72	176
<b>4 *</b> 111	34	41	78			444	136	164	312
1	21	17	76			4	84	86	304
45	52	78	91			180	208	312	364
<b>4x4x1</b>				<b>4x4x1</b>					





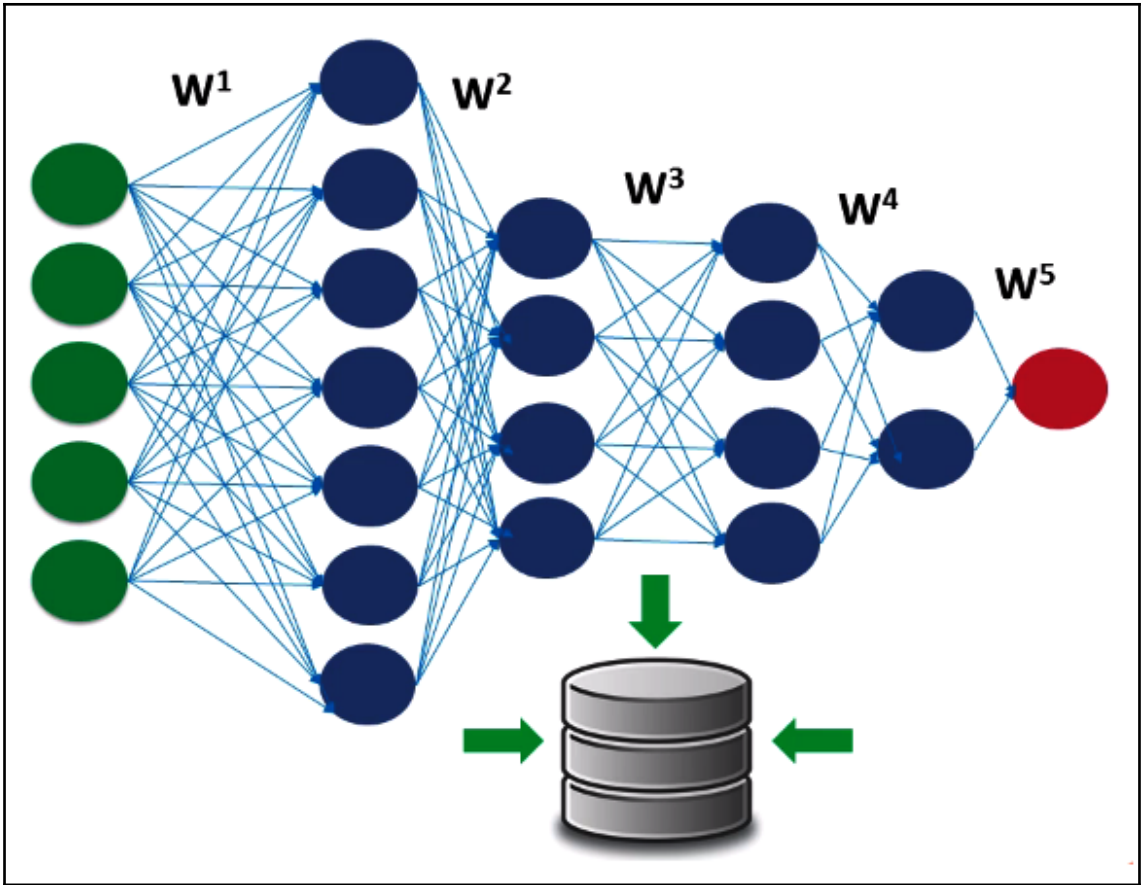


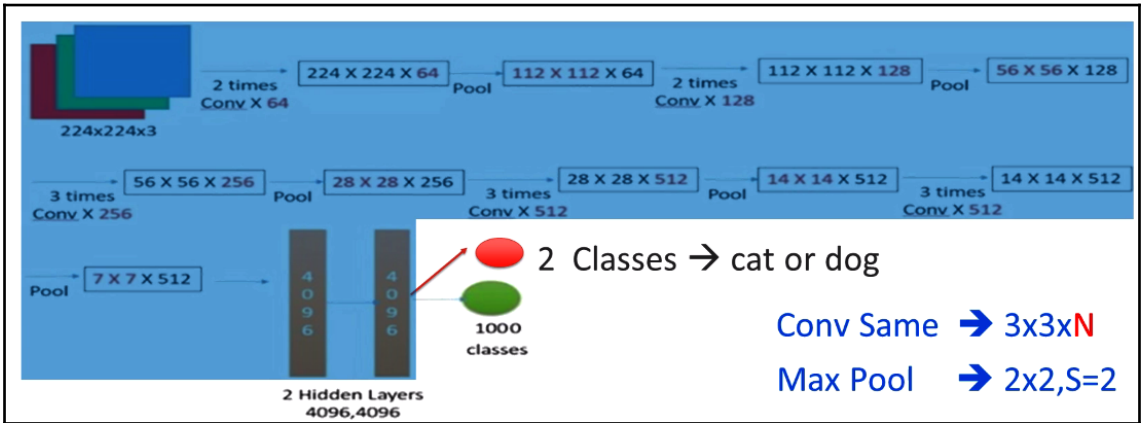













VertexName (VertexType)	nIn, nOut	TotalParams	ParamsShape	Vertex Inputs
input_1 (InputVertex)	-,-	-	-	-
block1_conv1 (ConvolutionLayer)	3, 64	1792	b: {1, 64}, W: {64, 3, 3, 3}	[input_1]
block1_conv2 (ConvolutionLayer)	64, 64	36928	b: {1, 64}, W: {64, 64, 3, 3}	[block1_conv1]
block1_pool (SubsamplingLayer)	-,-	0	-	[block1_conv2]
block2_conv1 (ConvolutionLayer)	64, 128	73856	b: {1, 128}, W: {128, 64, 3, 3}	[block1_pool]
block2_conv2 (ConvolutionLayer)	128, 128	147584	b: {1, 128}, W: {128, 128, 3, 3}	[block2_conv1]
block2_pool (SubsamplingLayer)	-,-	0	-	[block2_conv2]
block3_conv1 (ConvolutionLayer)	128, 256	295168	b: {1, 256}, W: {256, 128, 3, 3}	[block2_pool]
block3_conv2 (ConvolutionLayer)	256, 256	590080	b: {1, 256}, W: {256, 256, 3, 3}	[block3_conv1]
block3_conv3 (ConvolutionLayer)	256, 256	590080	b: {1, 256}, W: {256, 256, 3, 3}	[block3_conv2]
block3_pool (SubsamplingLayer)	-,-	0	-	[block3_conv3]
block4_conv1 (ConvolutionLayer)	256, 512	1180160	b: {1, 512}, W: {512, 256, 3, 3}	[block3_pool]

block4_conv2 (ConvolutionLayer)	512, 512	2359808	b: {1, 512}, W: {512, 512, 3, 3}	[block4_conv1]
block4_conv3 (ConvolutionLayer)	512, 512	2359808	b: {1, 512}, W: {512, 512, 3, 3}	[block4_conv2]
block4_pool (SubsamplingLayer)	-,-	0	-	[block4_conv3]
block5_conv1 (ConvolutionLayer)	512, 512	2359808	b: {1, 512}, W: {512, 512, 3, 3}	[block4_pool]
block5_conv2 (ConvolutionLayer)	512, 512	2359808	b: {1, 512}, W: {512, 512, 3, 3}	[block5_conv1]
block5_conv3 (ConvolutionLayer)	512, 512	2359808	b: {1, 512}, W: {512, 512, 3, 3}	[block5_conv2]
block5_pool (SubsamplingLayer)	-,-	0	-	[block5_conv3]
flatten (PreprocessorVertex)	-,-	-	-	[block5_pool]
fc1 (DenseLayer)	25088, 4096	102764544	b: {1, 4096}, W: {25088, 4096}	[flatten]
fc2 (DenseLayer)	4096, 4096	16781312	b: {1, 4096}, W: {4096, 4096}	[fc1]
predictions (DenseLayer)	4096, 1000	4097000	b: {1, 1000}, W: {4096, 1000}	[fc2]

VertexName (VertexType)	nIn,nOut	TotalParams	ParamsShape	Vertex Inputs
input_1 (InputVertex)	-, -	-	-	-
block1_conv1 (Frozen ConvolutionLayer)	3, 64	1792	b:{1, 64}, W:{64, 3, 3, 3}	[input_1]
block1_conv2 (Frozen ConvolutionLayer)	64, 64	36928	b:{1, 64}, W:{64, 64, 3, 3}	[block1_conv1]
block1_pool (Frozen SubsamplingLayer)	-, -	0	-	[block1_conv2]
block2_conv1 (Frozen ConvolutionLayer)	64, 128	73856	b:{1, 128}, W:{128, 64, 3, 3}	[block1_pool]
block2_conv2 (Frozen ConvolutionLayer)	128, 128	147584	b:{1, 128}, W:{128, 128, 3, 3}	[block2_conv1]
block2_pool (Frozen SubsamplingLayer)	-, -	0	-	[block2_conv2]
block3_conv1 (Frozen ConvolutionLayer)	128, 256	295168	b:{1, 256}, W:{256, 128, 3, 3}	[block2_pool]
block3_conv2 (Frozen ConvolutionLayer)	256, 256	590080	b:{1, 256}, W:{256, 256, 3, 3}	[block3_conv1]
block3_conv3 (Frozen ConvolutionLayer)	256, 256	590080	b:{1, 256}, W:{256, 256, 3, 3}	[block3_conv2]
block3_pool (Frozen SubsamplingLayer)	-, -	0	-	[block3_conv3]

**Total Parameters: 138357544**  
**Trainable Parameters: 138357544** I  
**Frozen Parameters: 0**

---

## Chapter 4: Real-Time Object Detection

**Image Classification**



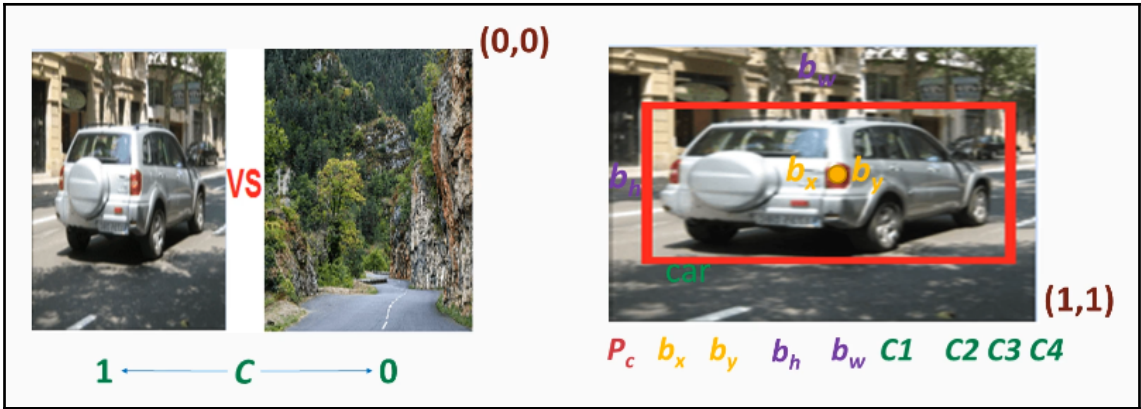
*Is it a car?*

**Object Localization**





*Where is the car?*

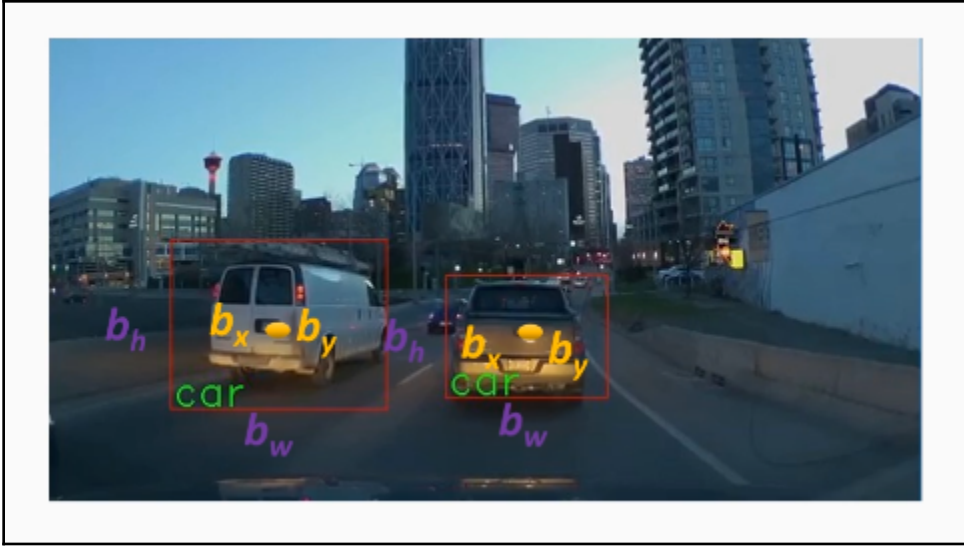




$P_c$	$b_x$	$b_y$	$b_h$	$b_w$	$C1$	$C2$	$C3$	$C4$
1	0.5	0.5	0.4	0.8	1	0	0	0
0	?	?	?	?	?	?	?	?







---

$$\begin{Bmatrix} P_c \\ L_{1x} \\ L_{1y} \\ L_{2x} \\ L_{2y} \\ L_{3x} \\ L_{3y} \\ \cdot \\ \cdot \\ \cdot \\ L_{25x} \\ L_{25y} \end{Bmatrix}$$






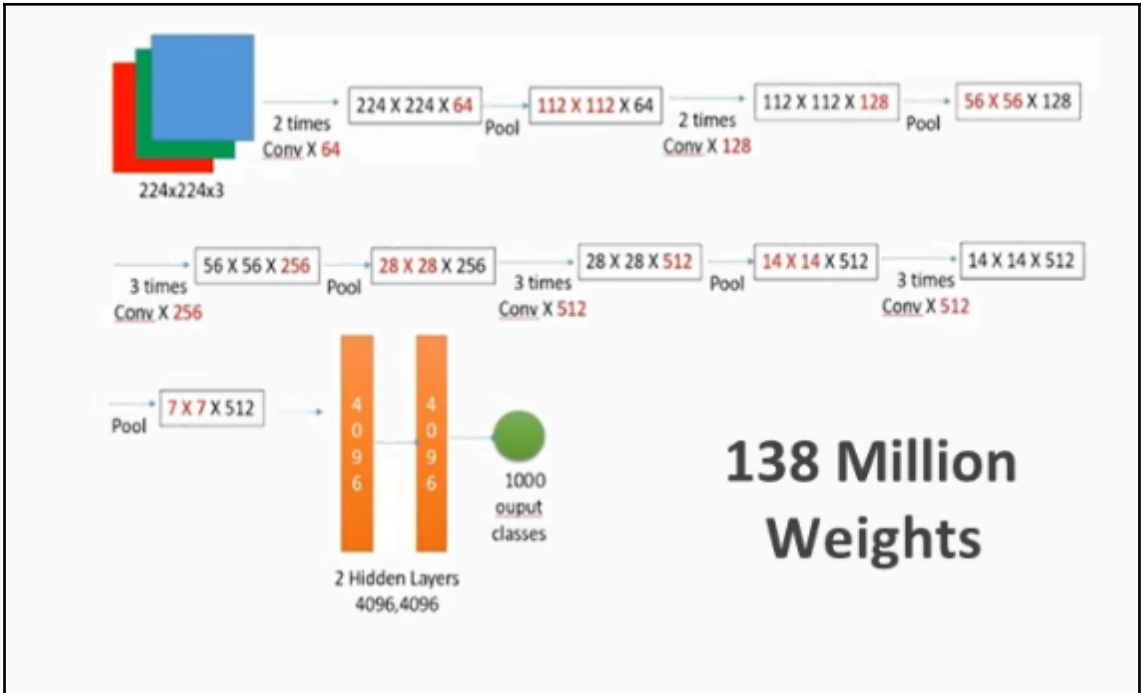
Chosen Windows Sizes













Chosen Windows Sizes











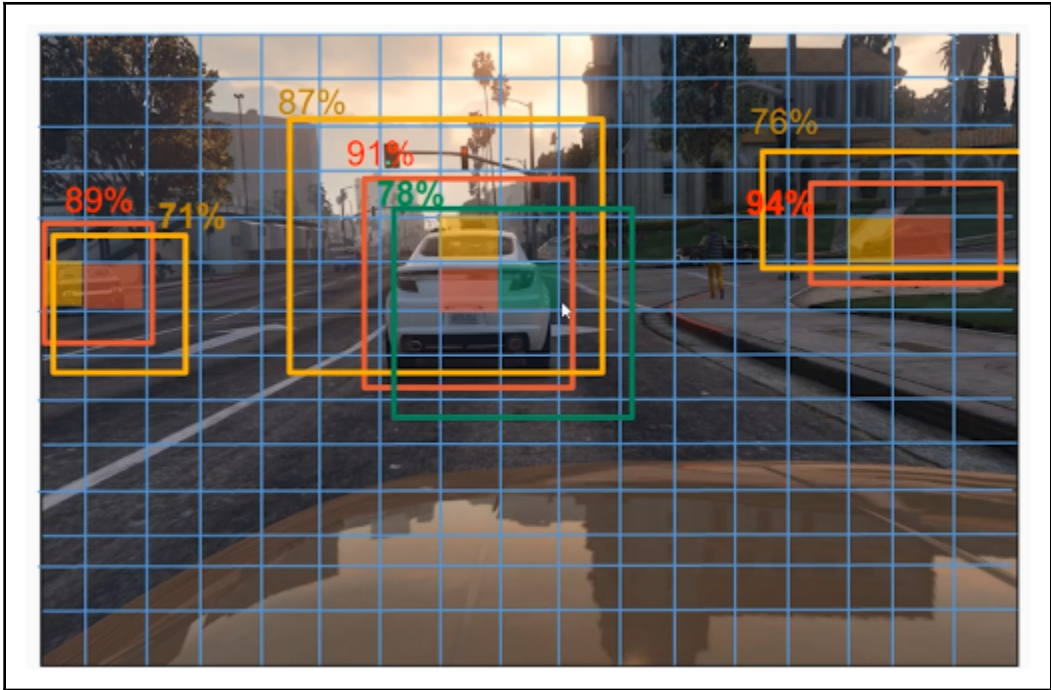






$$\begin{Bmatrix} 0.6 \\ b_x \\ b_y \\ b_h \\ b_w \\ 1 \\ 0 \\ 0 \end{Bmatrix}$$





---

*list = all bounding boxes*

*repeat until no boxes left in list*

*maxBox=choose box with max  $P_c$*


*list.remove(maxBox)*



*for box in list*

*if(loU(box, maxBox)>0.5)*

*list.remove(box)*

*showBox(maxBox);*





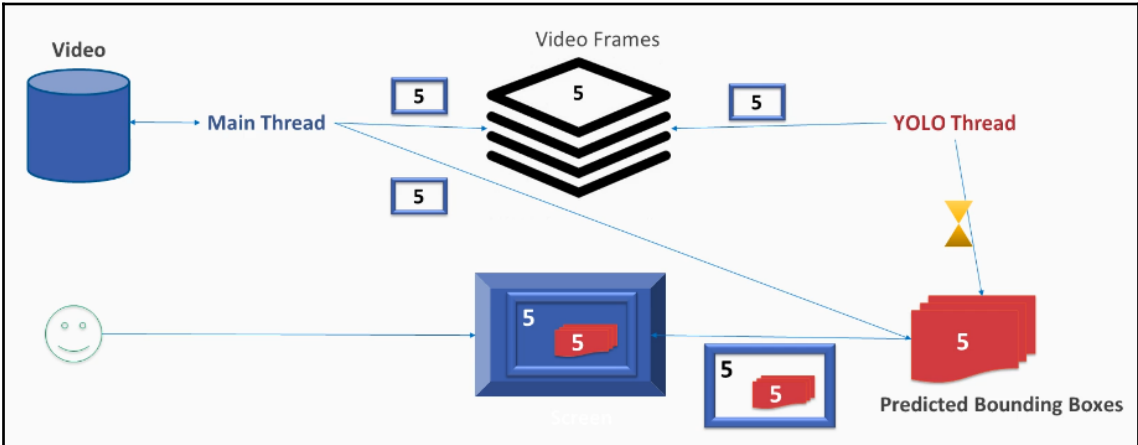
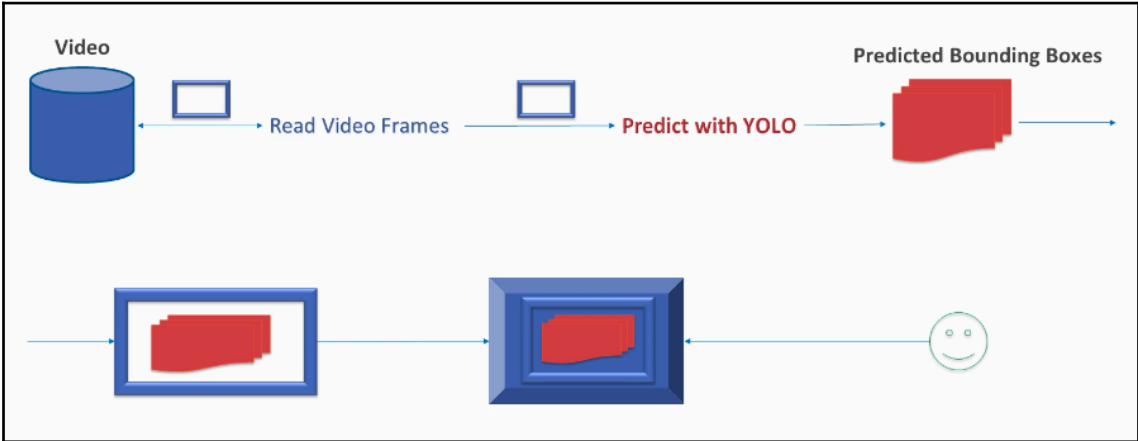
$$\begin{Bmatrix} P_c \\ b_x \\ b_y \\ b_h \\ b_w \\ 1 \\ 0 \\ 0 \end{Bmatrix}$$


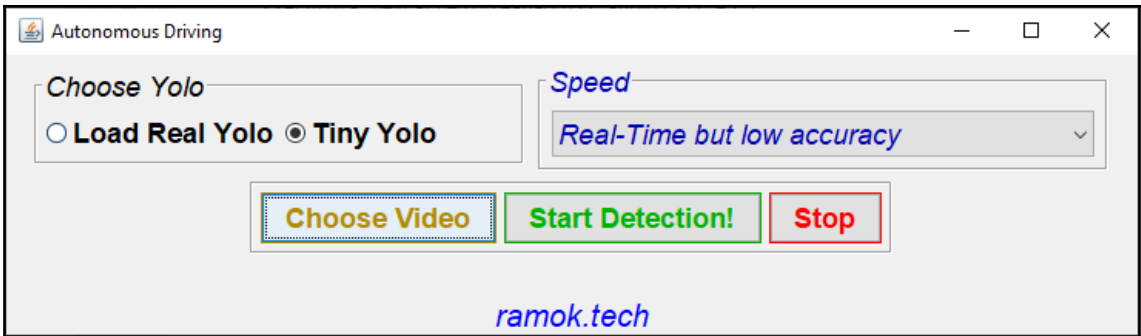
$$\begin{matrix} 8 \times 9 \times (5+C) \\ 8 \times 9 \times 8 \end{matrix}$$


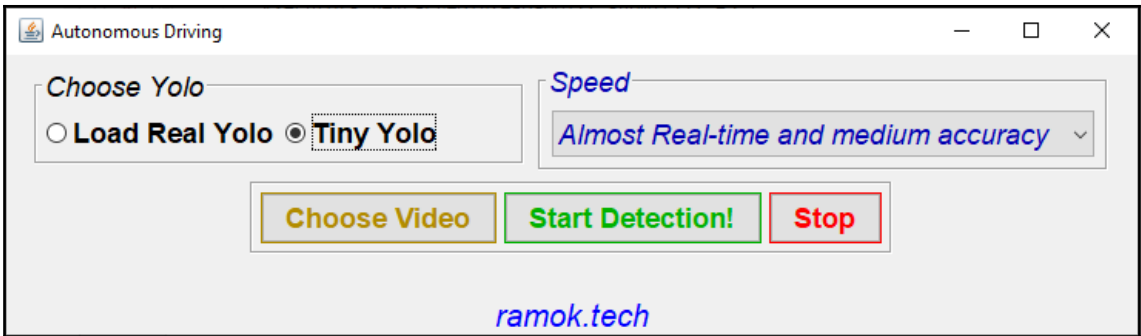
$$\begin{matrix} \boxed{\phantom{000000000}} \\ \boxed{\phantom{000000000}} \end{matrix} \begin{Bmatrix} P_c & b_x & b_y & b_h & b_w & 1 & 0 & 0 \\ P_c & b_x & b_y & b_h & b_w & 0 & 1 & 0 \end{Bmatrix}$$

$8 \times 9 \times 2 \times (5+C)$   
 $8 \times 9 \times 16$

$8 \times 9 \times 2 \times (5+C)$   
 $8 \times 9 \times 16$

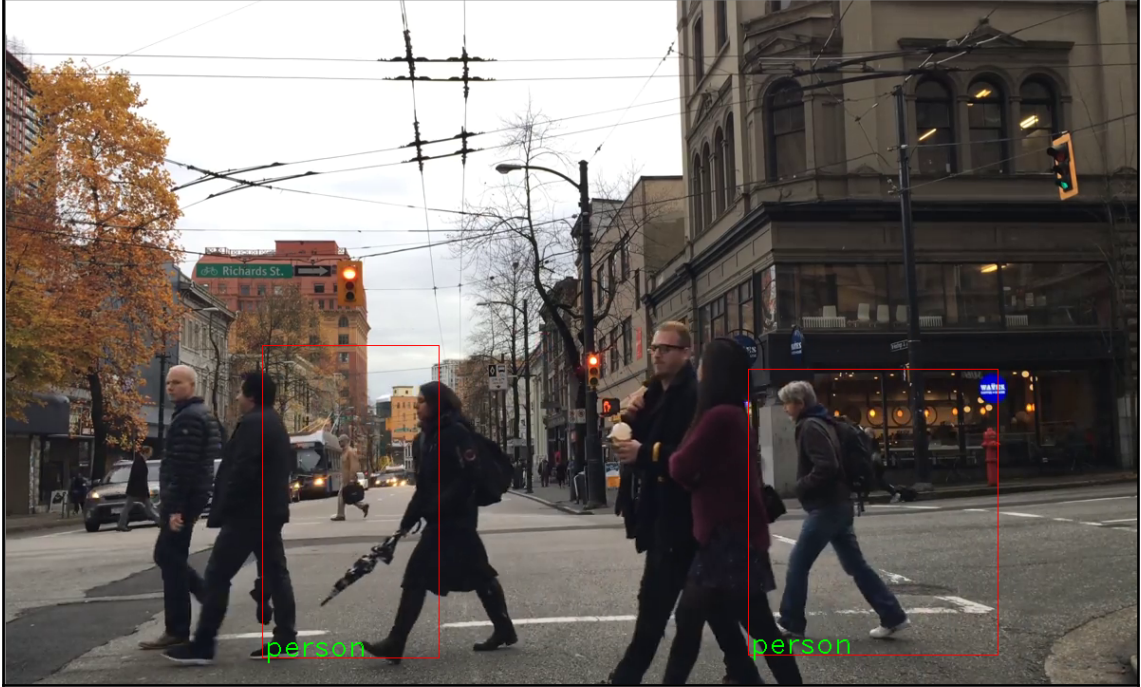







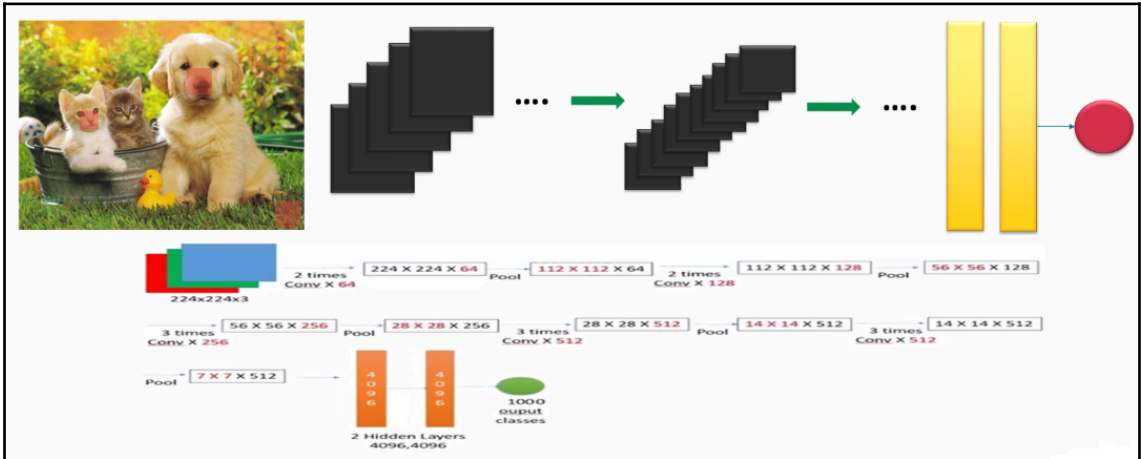



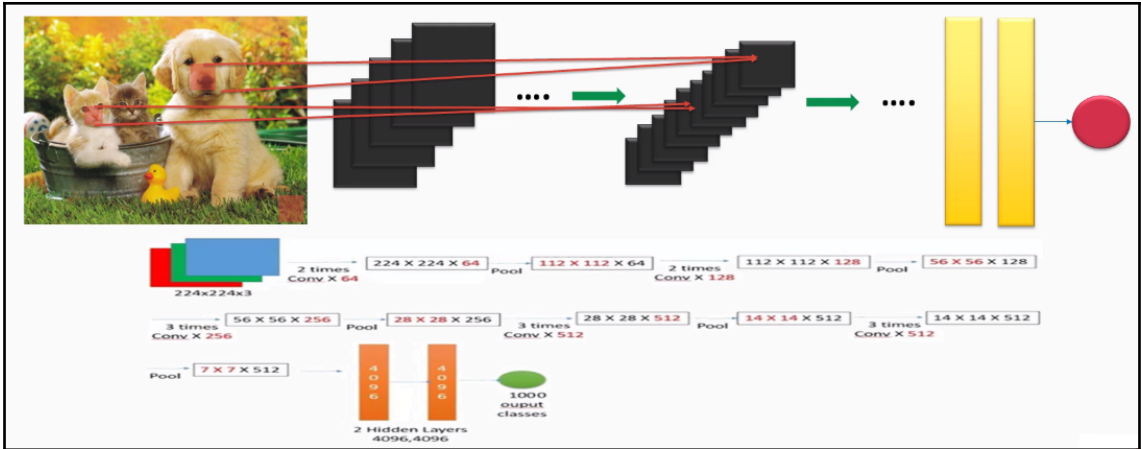








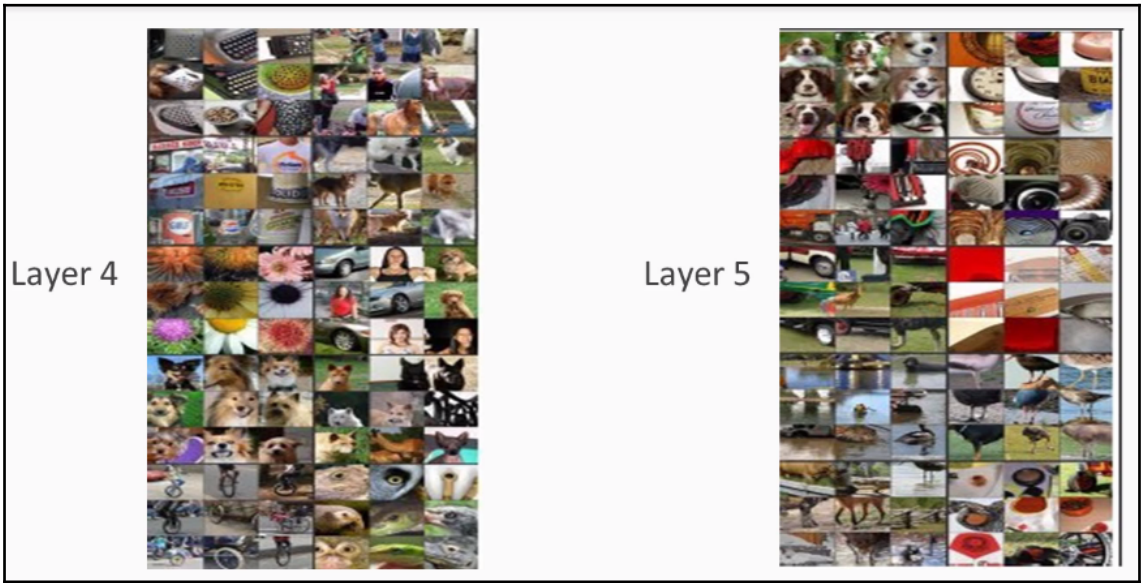

# Chapter 5: Creating Art with Neural Style Transfer











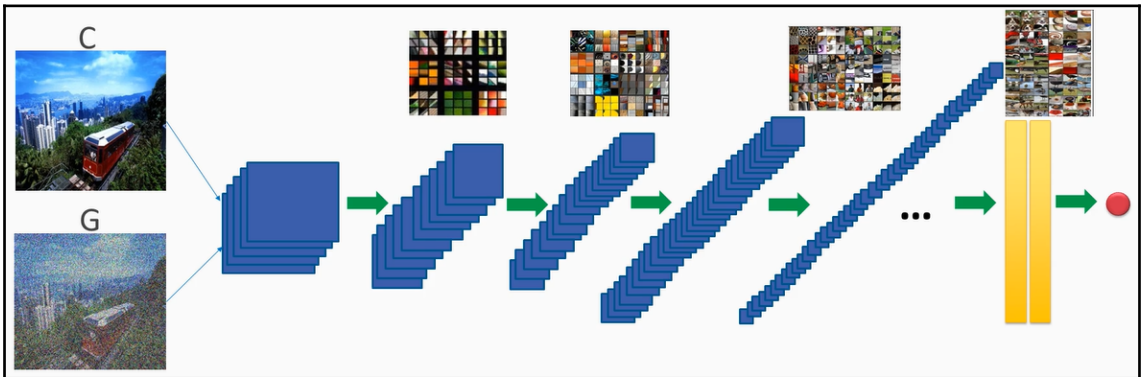
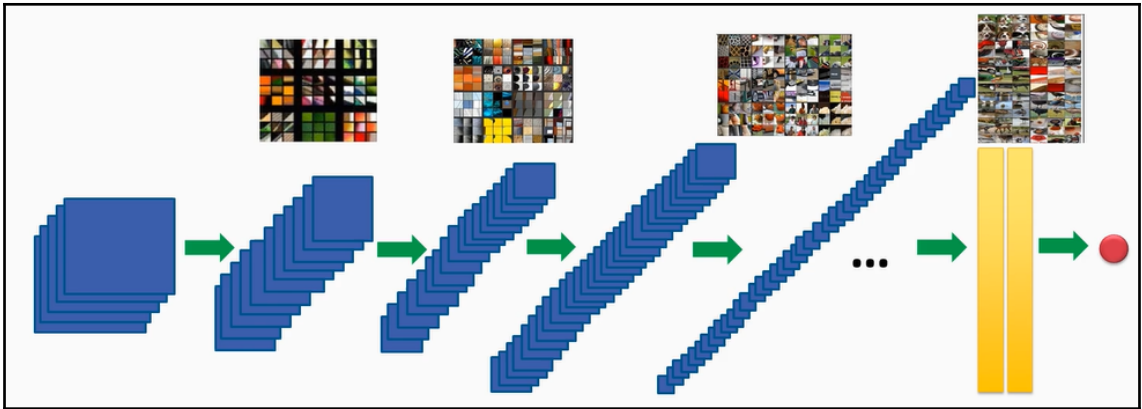




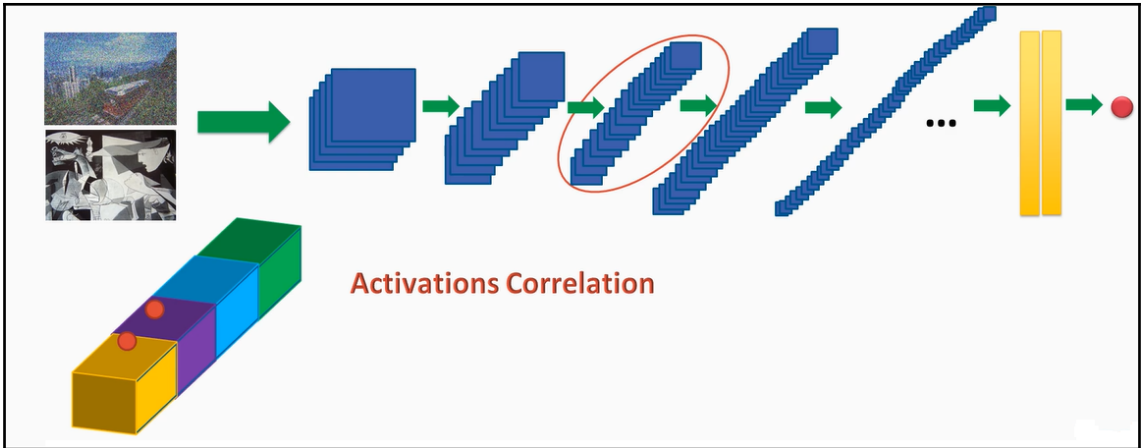















Diagram illustrating the comparison of  $g_i^S$  with  $g_i^G$ .

The equation is:  $\mathbf{g} = \mathbf{C} * \mathbf{H}^*W$ , where  $\mathbf{C}$  is a vertical column of colored dots and  $\mathbf{H}^*W$  is a horizontal row of colored dots. The result is a 4x4 grid of elements  $g_{ij}$ .

The resulting matrix  $\mathbf{g}$  is shown as a 4x4 grid of elements  $g_{ij}$ :



$g_{11}$	$g_{12}$	$g_{13}$	$g_{14}$
$g_{21}$	$g_{22}$	$g_{23}$	$g_{24}$
$g_{31}$	$g_{32}$	$g_{33}$	$g_{34}$
$g_{41}$	$g_{42}$	$g_{43}$	$g_{44}$


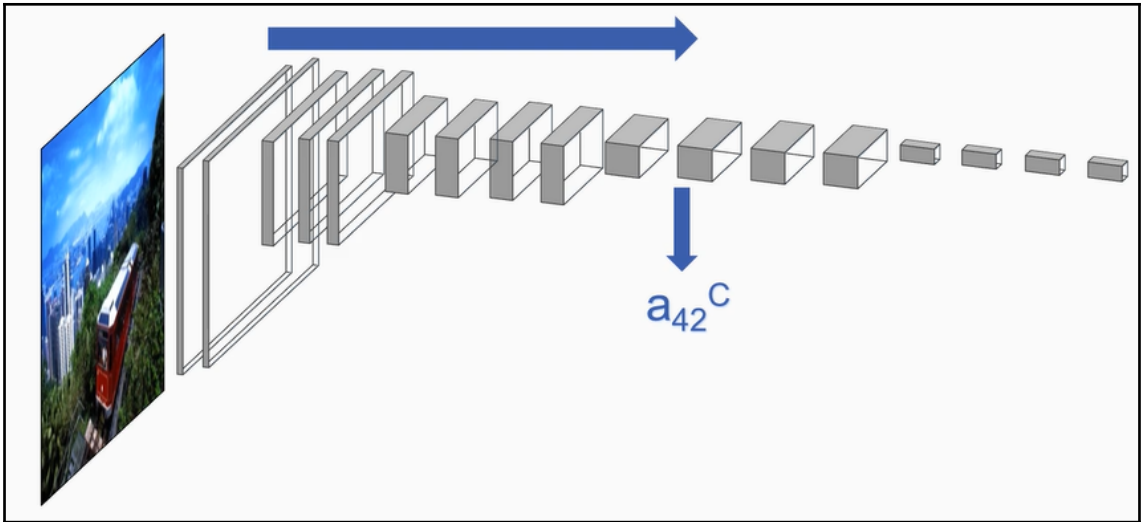
$$J(G) = \alpha * J(C, G) + \beta * J(S, G)$$

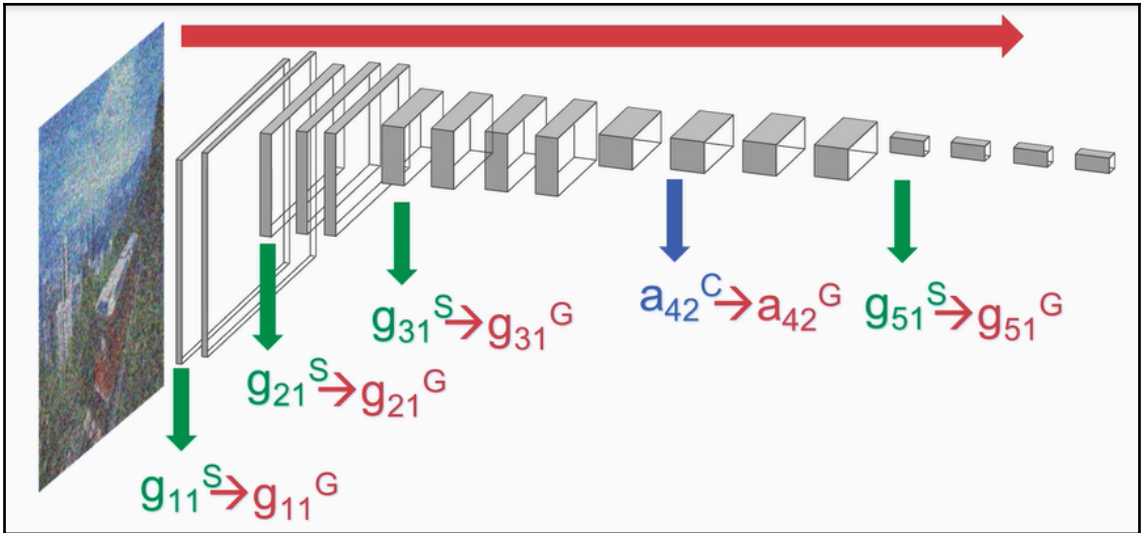

$$J(S, G) = \frac{1}{4 * (W * H * C)^2} * \|g_i^S - g_i^G\|^2$$

$$\mathbf{g} = \mathbf{a}_l * \mathbf{a}_l^T; d\mathbf{g} = 2 * \mathbf{a}_l^T$$

$$\frac{dJ(S, G)}{dG} = \frac{1}{(W * H * C)^2} * \|g_i^S - g_i^G\| * (\mathbf{a}_l^G)^T$$







$$J(G) = \alpha * J(C, G) + \beta * (k1 * J(S, G)^{l1} + k2 * J(S, G)^{l2} + k3 * J(S, G)^{l3} + k4 * J(S, G)^{l4})$$



---

$$J(C, G) = \frac{1}{4 * W * H * C} * \|a_i^C - a_i^G\|^2$$

$$J(S, G) = \frac{1}{4 * (W * H * C)^2} * \|g_i^S - g_i^G\|^2$$

```
"C:\Program Files\Java\jdk-11.0.2\bin\java.exe" ...  
[2019-02-06 17:51:30]Loaded [CpuBackend] backend  
[2019-02-06 17:51:30]given scan urls are empty. set urls in the configuration  
[2019-02-06 17:51:30]Number of threads used for NativeOps: 4  
[2019-02-06 17:51:30]Reflections took 47 ms to scan 13 urls, producing 31 keys and 227 values  
[2019-02-06 17:51:30]Number of threads used for BLAS: 4  
[2019-02-06 17:51:30]Backend used: [CPU]; OS: [Windows 10]  
[2019-02-06 17:51:30]Cores: [8]; Memory: [4.0GB];  
[2019-02-06 17:51:30]Blas vendor: [OPENBLAS]  
[2019-02-06 17:51:31]Downloading model to C:\Users\Admin\.deeplearning4j\vgg16_d14j_inference.zip
```





---

## Chapter 6: Face Recognition

### Image Classification



*Is it a car?*

---

## Face Verification



*Is Person with ID X?*

---

## Face Recognition



*Is any of known Persons in group?*

98%

### Face Verification




*Is Person with ID X?*



16\*2%=32%

### Face Recognition









*Is any of known Persons in group?*

















Positive

$$d(A,P)=(F(A) - F(P))^2 \approx 0$$



Negative

$$d(A,N)=(F(A) - F(N))^2 > 0$$

$$(F(A) - F(P))^2 \approx 0 ; (F(A) - F(N))^2 > 0$$

$$(F(A) - F(P))^2 \leq (F(A) - F(N))^2$$

$$(F(A) - F(P))^2 - (F(A) - F(N))^2 \leq 0$$

---

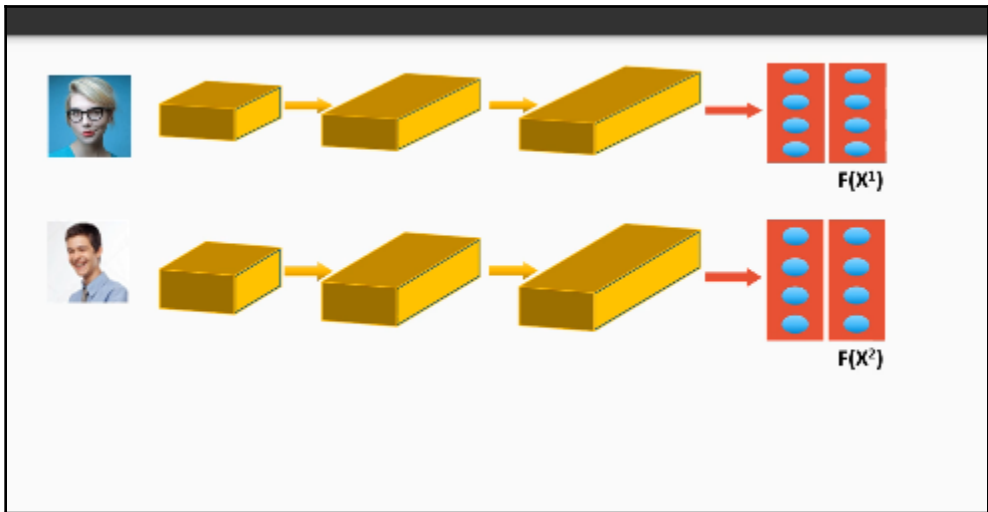
$$(F(A) - F(P))^2 - (F(A) - F(N))^2 \leq -\epsilon$$


$$(F(A) - F(P))^2 - (F(A) - F(N))^2 + \epsilon \leq 0$$

$$0.5 - (F(A) - F(N))^2 + 0.2 \leq 0$$

When negative examples N are chosen randomly then the condition is easily satisfied:  $d(A,P) + \epsilon \leq d(A,N)$ ; since  $d(A,N)$  will be always big enough

Choose triplets that are hard to train  $d(A,P) + \epsilon \leq d(A,N)$ ; so network has to push up hard  $d(A,N)$  and down  $d(A,P)$  to satisfy the margin " $\epsilon$ "





$$Z = \text{sum}(W_k * (f(X_i^1) - f(X_i^2)) + b)$$

$$Y = \text{sig}(Z); 0|1$$

$$Z = \text{sum}(W_k * \frac{(f(X_i^1) - f(X_i^2))^2}{f(X_i^1) + f(X_i^2)} + b)$$

