

DR. ALVIN'S PUBLICATIONS

# CONVOLUTIONAL NEURAL NETWORKS (CNN)

---

HOW IT WORKS  
DR. ALVIN ANG



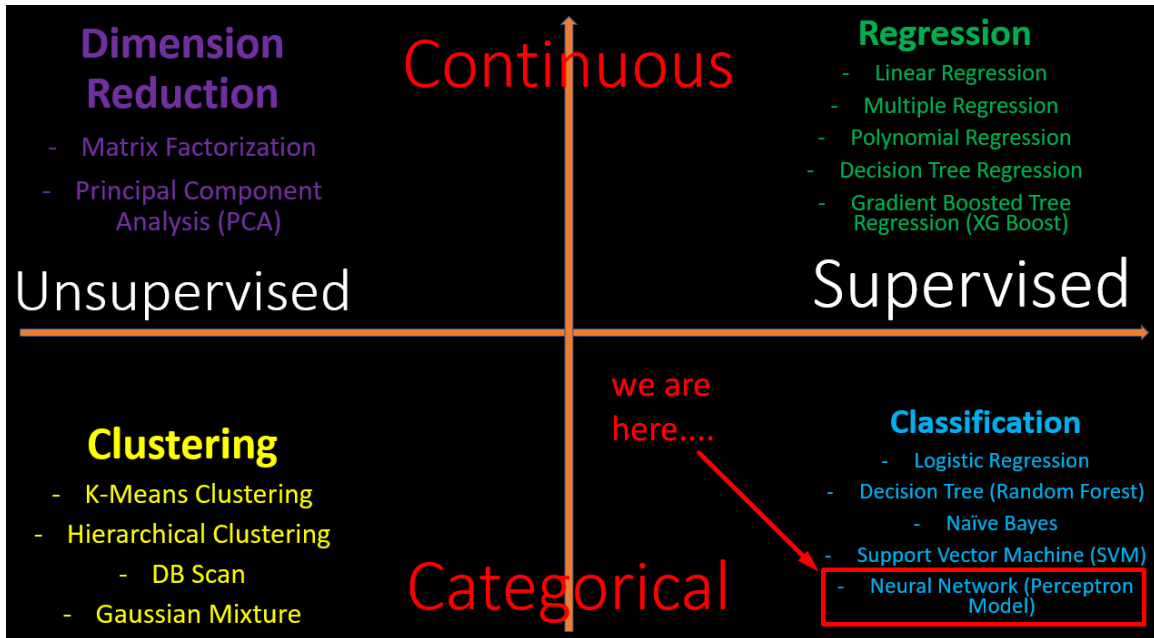
---

1 | PAGE

COPYRIGHTED BY DR ALVIN ANG  
WWW.ALVINANG.SG

# CONTENTS

<b>I. Deep Learning = Neural Network (NN)</b> .....	<b>3</b>
<b>II. Why CNN?</b> .....	<b>4</b>
<b>A. Brief Recap of Artificial Neural Network (ANN)</b> .....	<b>4</b>
<b>B. Flattening is VERY BAD</b> .....	<b>5</b>
<b>C. Basic Idea Behind CNN</b> .....	<b>7</b>
<b>III. CNN Architecture</b> .....	<b>8</b>
<b>IV. Understanding CNN Using Keras</b> .....	<b>10</b>
<b>A. Step 1: Import All Libraries</b> .....	<b>10</b>
<b>B. Step 2: Load the photo</b> .....	<b>11</b>
<b>C. Step 3: What is a Convolution?</b> .....	<b>12</b>
1. Convolution Does This .....	13
2. Sequential Modeling .....	14
3. What is Padding? Why do we Need It? .....	16
4. Convolution Output / Effect of Convolution .....	17
<b>D. Step 4: What is the Effect of using Relu Activation?</b> .....	<b>18</b>
1. Sequential Modeling .....	19
2. Relu Output / Effect of using Relu Activation .....	19
<b>E. What is Pooling?</b> .....	<b>20</b>
1. Sequential Modeling .....	22
2. Pooling Output / Effect of Pooling .....	22
<b>References</b> .....	<b>24</b>
<b>About Dr. Alvin Ang</b> .....	<b>25</b>



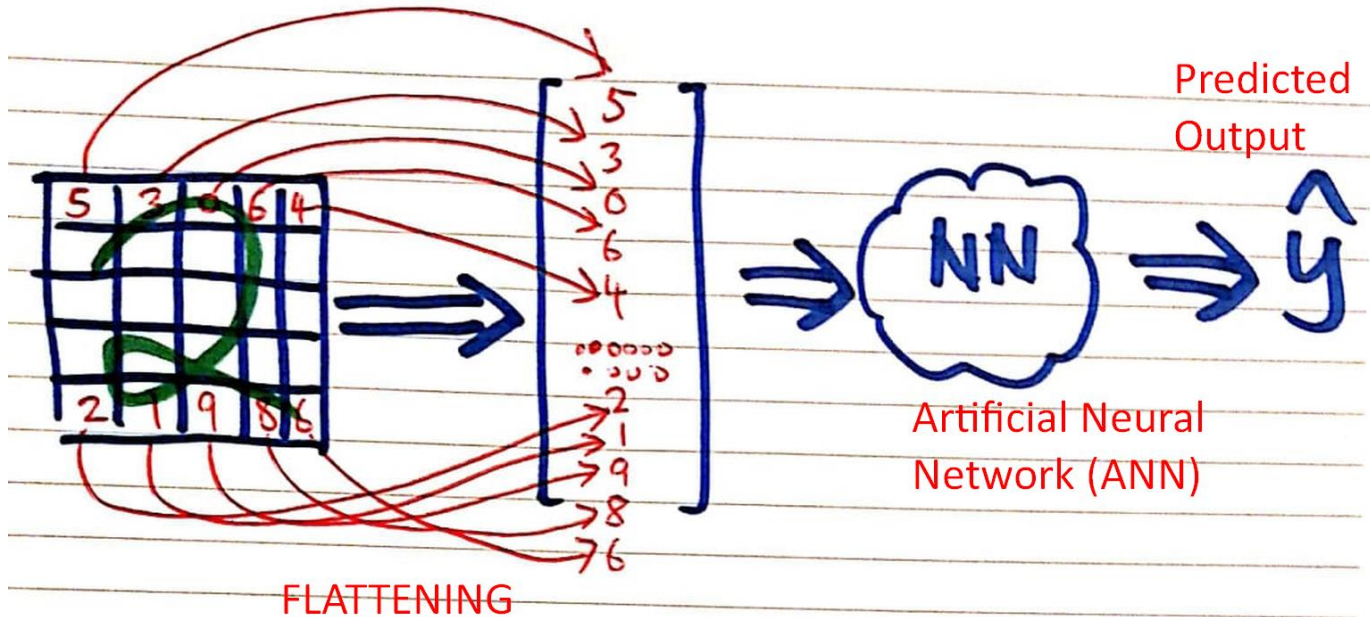
- Above is a table categorizing the different Machine Learning algorithms.
- Objective of Neural Network is to predict a CATEGORY.

---

## II. WHY CNN?

---


### A. BRIEF RECAP OF ARTIFICIAL NEURAL NETWORK (ANN)



- In my previous article, I talked about how ANN worked (with images).
- An image “2” needs to be first FLATTENED, then fed into the NN for training / learning.
- Subsequently, once the NN is trained, it can predict new jpgs and output the label / classification.


## B. FLATTENING IS VERY BAD


- However, FLATTENING is a very DESTRUCTIVE process.
- It breaks down the jpg into an array (which obviously you can't see a 2 anymore in the vector).



- Flattening converts a 3 x 3 array  $\rightarrow$  9 x 1 array


FLATTENING DOES THIS!!!






- Now if I write a 2 (tilted) like the above, the FLATTENED array obviously will be different from the first “2”’s array (the straight up one in the previous page).
- However, we all know a “2” is still just a “2” and they are both the same! But their arrays are completely different from one another!
- Thus, FLATTENING KILLS A LOT OF IMPORTANT INFORMATION.
- We need a way to PREVENT INFORMATION LOSS.

## Very Important Feature



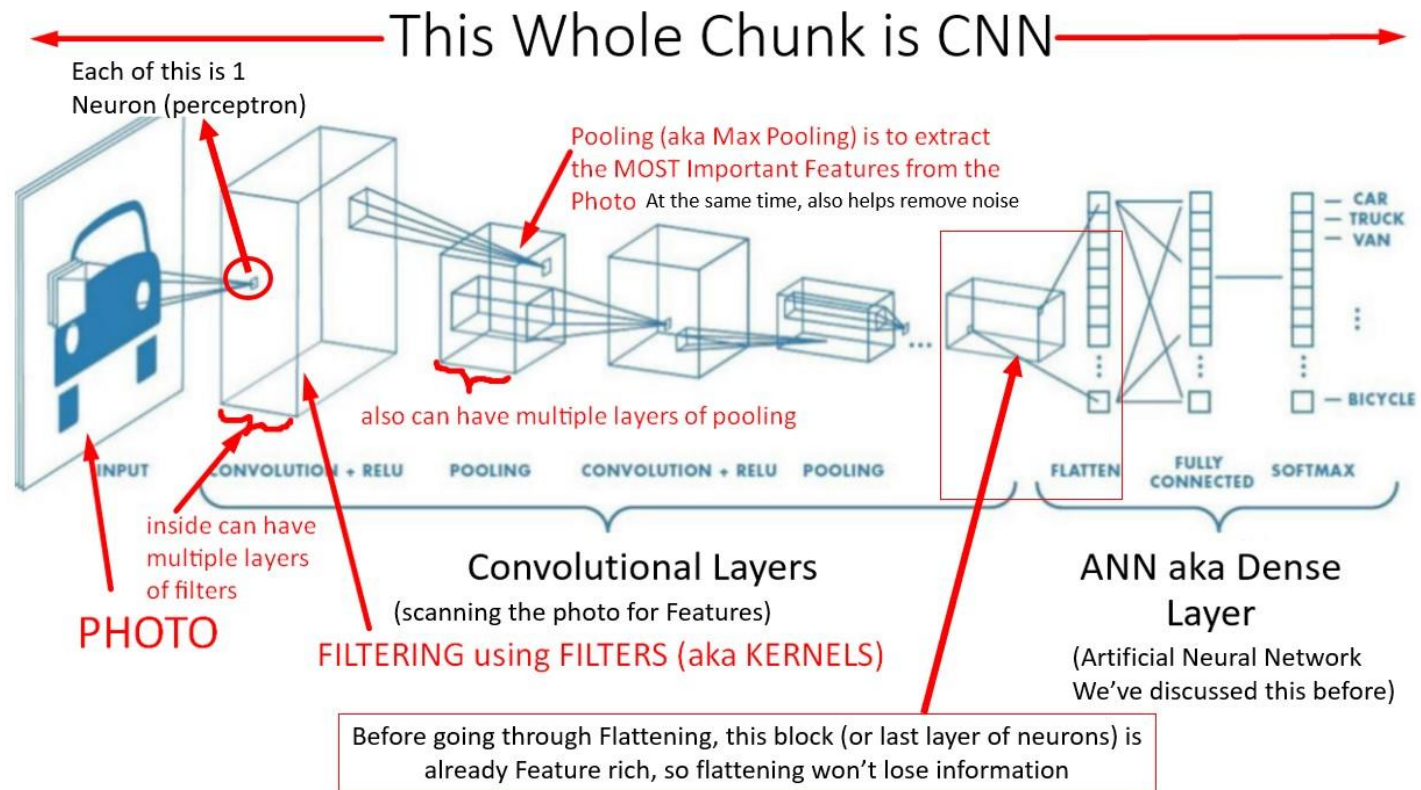
- For example, a “4” important feature is the cross → without it, its hard for us to see it’s a “4”.
- We need a way to feed the IMPORTANT FEATURES into the NN for it to recognize the number.

### C. BASIC IDEA BEHIND CNN

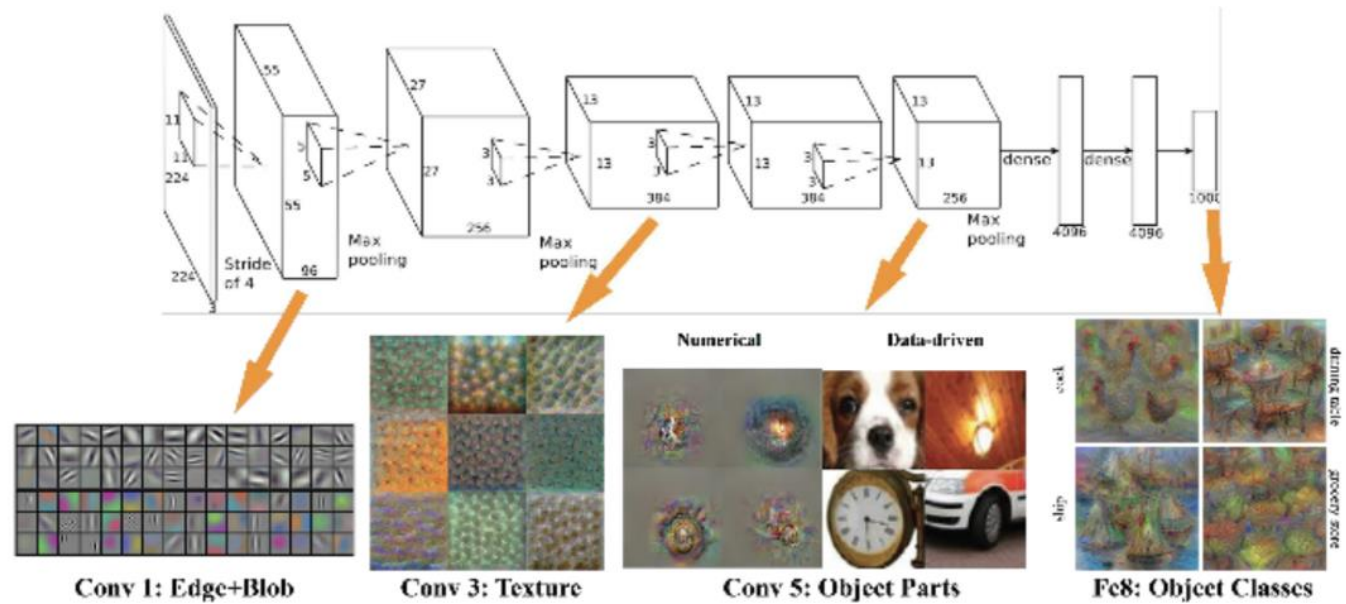


- We need a way for the NN to pick up important features from the image (in order to recognize it).
- Mario's Ear... + Moustache... + Shoulder... each part will piece together to help the NN recognize that its Mario.

### III. CNN ARCHITECTURE







- At every Convolution layer, Features are extracted.
- 1<sup>st</sup> Convolution normally only extracts very simple Features like Vertical and Horizontal Lines.
- Subsequent layers search for more complex Features like Texture and Parts of an Object.

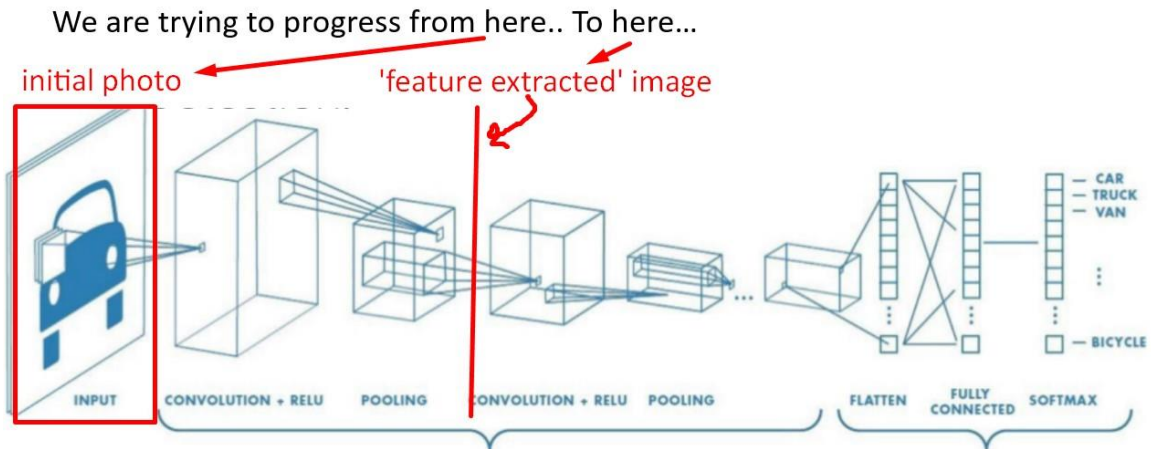
---

## IV. UNDERSTANDING CNN USING KERAS

---

IPYNB:

[https://www.alvinang.sg/s/Understanding\\_CNN\\_by\\_Dr\\_Alvin\\_Ang.ipynb](https://www.alvinang.sg/s/Understanding_CNN_by_Dr_Alvin_Ang.ipynb)



### A. STEP 1: IMPORT ALL LIBRARIES

#### Step 1: Import All Libraries

```
1 from tensorflow.keras.models import Sequential
2 from tensorflow.keras.layers import Conv2D, MaxPooling2D
3
4 import numpy as np
5 import matplotlib.pyplot as plt
6 import cv2 # only used for loading the image
7
8 %matplotlib inline
```


## B. STEP 2: LOAD THE PHOTO

### Step 2: Show the Photo

```
[ ] alvin = cv2.imread('/content/ALVIN.jpg')
```

```
▶ plt.imshow(alvin)
```

```
↳ <matplotlib.image.AxesImage at 0x7f23e250e8d0>
```



```
▶ # what does the image look like?  
alvin.shape
```

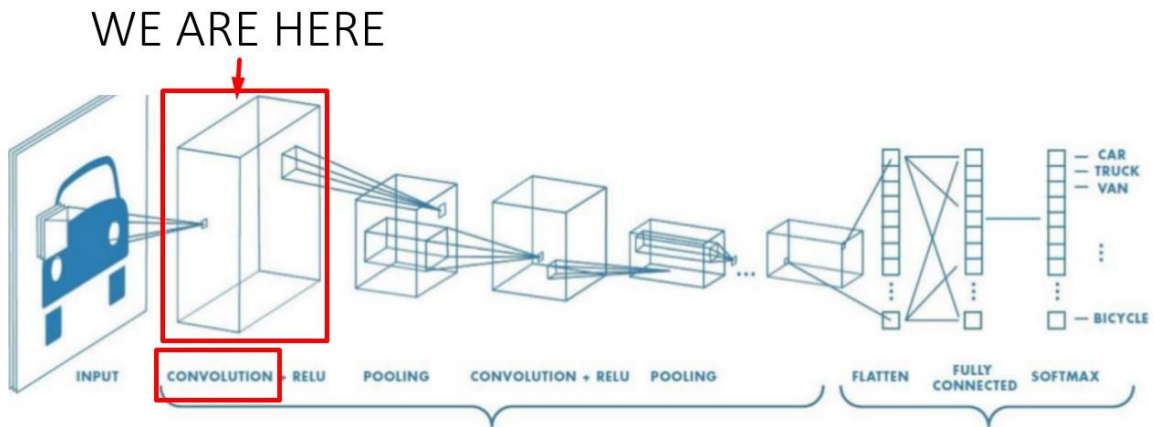
```
#1408 pixels by 793 pixels with 3 channels RGB
```

```
↳ (836, 789, 3)
```

C. STEP 3: WHAT IS A CONVOLUTION?

# Step 3: Demonstrating effect of 1 Convo Layer

## What Does a Convolution Layer Do? aka Filtering




# 1. CONVOLUTION DOES THIS

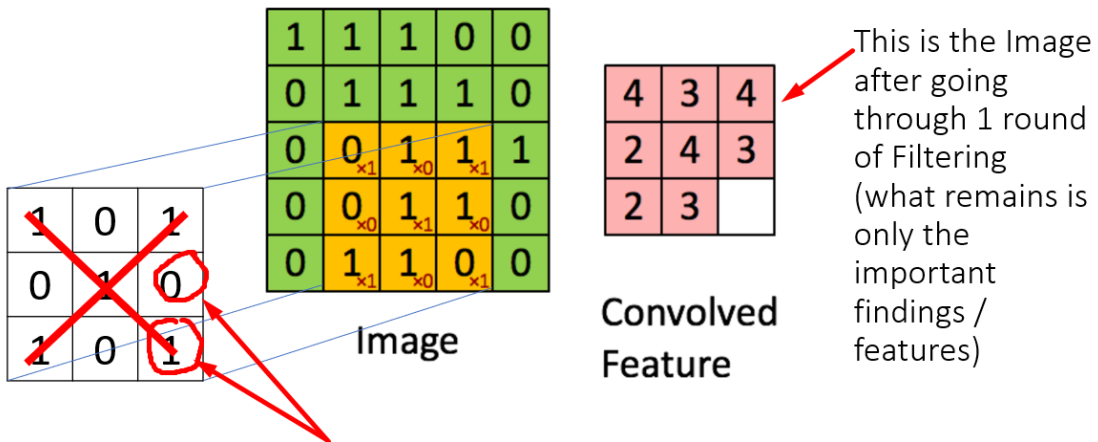
To see animated .gif, go here to understand how the Filter slides across the image:

<https://towardsdatascience.com/a-comprehensive-guide-to-convolutional-neural-networks-the-eli5-way-3bd2b1164a53>

This Filter  
(or Kernel) is 3x3



- Filter (or Kernel) is trying to seek out an 'X' Feature
- Imagine that particular yellow portion of the image is all 1
- Then  $(1 \times 1) + (1 \times 1) + (1 \times 1) + \dots$  Will get you MAXIMUM number 9  $\rightarrow$  means that 'X' feature is found!
- Imagine that is all 0... then  $(0 \times 1) + (0 \times 1) + \dots$  You will get 0  $\rightarrow$  No 'X' feature there!



- Note that these are Weights which will be automatically found by the CNN (updated with each Epoch)
- In other words, CNN will find out the "best" features to look out for automatically (in each Filter).

## 2. SEQUENTIAL MODELING

```
▶ model = Sequential()

model.add(Conv2D(filters = 3,
                 #we need to use 3 filters because its RGB

                 kernel_size = (3, 3),
                 #we use a commonly used kernel_size of 3x3

                 #strides = 1,
                 #default stride = 1, even if you don't declare

                 #padding = 'same',
                 #default padding = 'valid', even if you don't declare
                 #we use 'same' to prevent loss of information

                 #activation = None,
                 #default activation = 'none', even if you don't declare
                 #but for now, we purposely force our activation to 'none'
                 #for purpose of demonstrating the effect of 1 Conv layer
                 #without activation


                 input_shape = alvin.shape))
```

```
model = Sequential()

model.add(Conv2D(filters = 3,
                 #we need to use 3 filters because its RGB

                 kernel_size = (3, 3),
                 #we use a commonly used kernel_size of 3x3
```

- We need 3 filters because our ALVIN.jpg is made up of 3 primary colours = RED + GREEN + BLUE




## 3 Filters need for 3 Layers of Colours RED + GREEN + BLUE

To see animated .gif, go here to understand how the 3 Filters slide across the image:

<https://towardsdatascience.com/a-comprehensive-guide-to-convolutional-neural-networks-the-eli5-way-3bd2b1164a53>

So basically, this happens.....



Movement of the Kernel

### 3. WHAT IS PADDING? WHY DO WE NEED IT?


```
model.add(Conv2D(filters = 3,  
                 #we need to use 3 filters because its RGB  
                 kernel_size = (3, 3),  
                 #we use a commonly used kernel_size of 3x3  
                 #strides = 1,  
                 #default stride = 1, even if you don't declare  
                 #padding = 'same',  
                 #default padding = 'valid', even if you don't declare  
                 #we use 'same' to prevent loss of information
```

To see animated .gif, go here to understand how padding is done:

[https://github.com/vdumoulin/conv\\_arithmetic](https://github.com/vdumoulin/conv_arithmetic)

<https://towardsdatascience.com/a-comprehensive-guide-to-convolutional-neural-networks-the-eli5-way-3bd2b1164a53>

Stride = 1 means it moves 1 pixel by 1 pixel  
Stride = 2 means it moves 2 pixels one shot



- **Valid** Padding = No Padding
  - No Padding leads to smaller **output**
  - Means information loss
- **Same** Padding = 1 Layer of Padding
  - 1 Padding leads to Same **output size**
  - Means **NO** information loss




#### 4. CONVOLUTION OUTPUT / EFFECT OF CONVOLUTION

### 3c) Convo Output

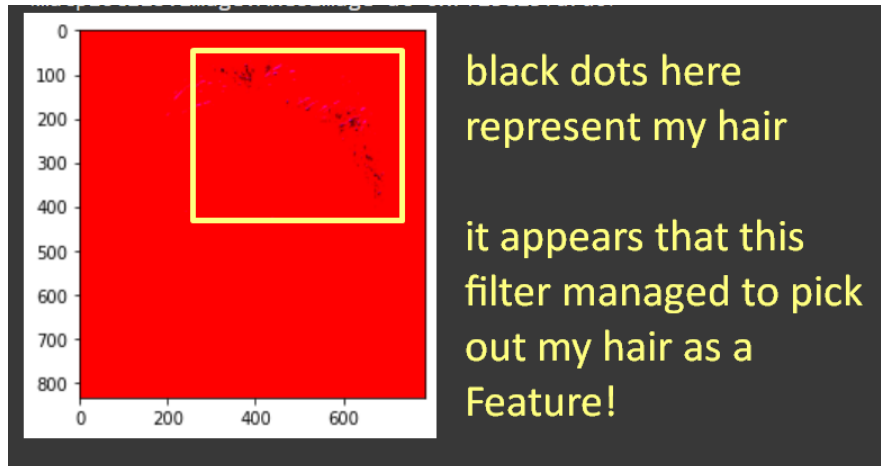
```
▶ plt.imshow(conv_alvin)
```

```
#sometimes the photo is yellow, sometimes red, sometimes blue  
#it will change upon different runs
```

```
↳ Clipping input data to the valid range for imshow with RGB data ([0..1] for floats or [0..255]  
<matplotlib.image.AxesImage at 0x7f23e23fa7d0>
```




- What do we see there?
- After 1 x Convolution with 3 filters, the entire image turned RED! (but you can try running it several times... it will change colours)



D. STEP 4: WHAT IS THE EFFECT OF USING RELU ACTIVATION?

## Step 4: Demonstrating effect of adding Relu activation

What happens when you Add a Relu Activation?



- Relu is an activation function and has already been explained here: <https://www.alvinang.sg/s/Artificial-Neural-Network-ANN-How-It-Works-by-Dr-Alvin-Ang.pdf>

## 1. SEQUENTIAL MODELING

### 4a) Sequential Modeling

```
▶ model2 = Sequential()

model2.add(Conv2D(filters = 3,
                  kernel_size = (3, 3),
                  #strides = 1,
                  #padding = 'same',

                  activation = 'relu',
                  #now we purposely change to Relu Activation
                  #to observe its effect


                  input_shape = alvin.shape))
```

## 2. RELU OUTPUT / EFFECT OF USING RELU ACTIVATION

### 4c) Relu Output

```
▶ plt.imshow(conv_alvin_2)
```


```
↳ Clipping input data to the valid range for imshow with RGB data ([0..1] for float)
   <matplotlib.image.AxesImage at 0x7f23e2340e90>
```



by using Relu, it appears that more dots are found...

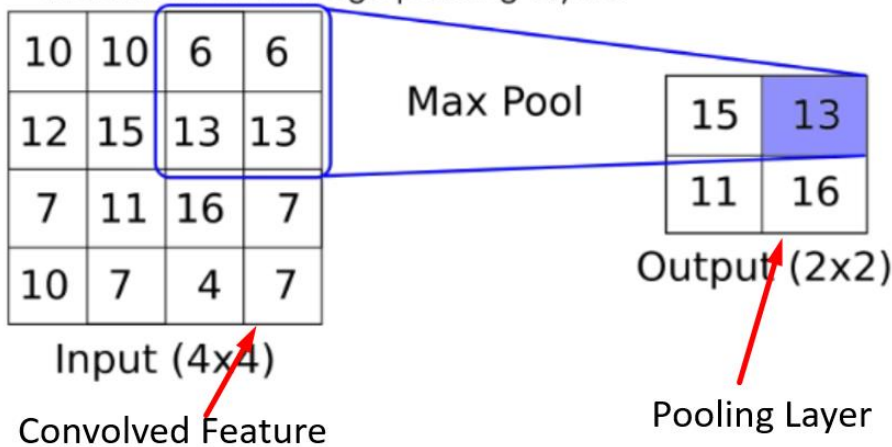
meaning, Relu helps the Filter to identify even more Features now...

Step 5: Demonstrating effect of Pooling  
 What does Pooling do?



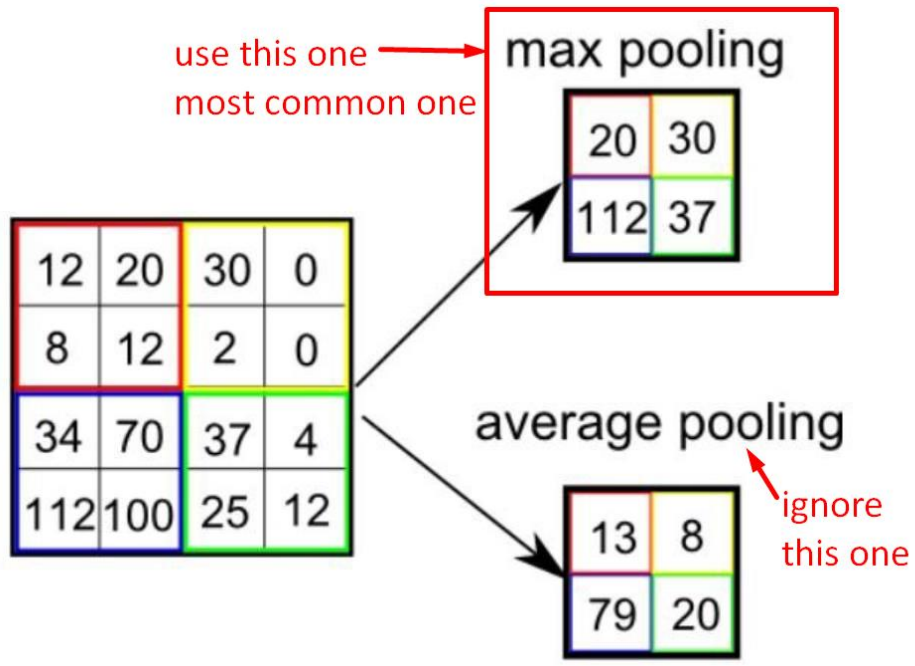
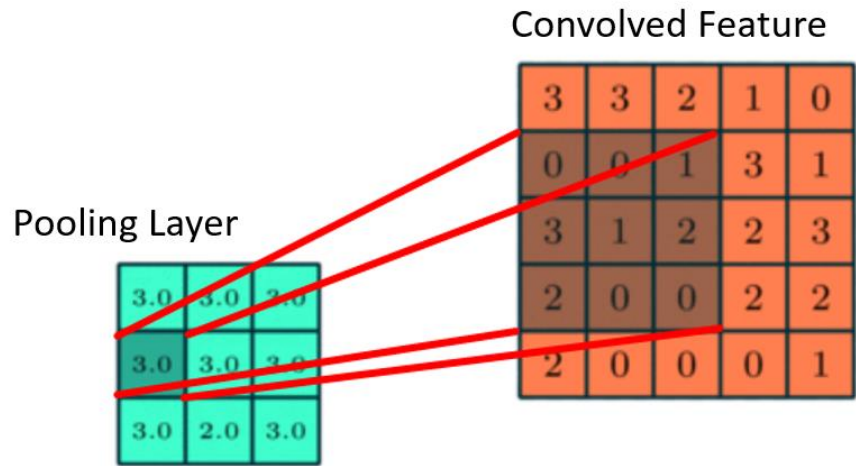
### Pooling

- Similar to stride, the purpose of a pooling layer is to downsample the output size. and extract the salient features.
- Maximum or average pooling layers



To see animated .gif, go here to understand how the Pooling Layer slides across the Convolved Feature:

<https://towardsdatascience.com/a-comprehensive-guide-to-convolutional-neural-networks-the-eli5-way-3bd2b1164a53>



## 1. SEQUENTIAL MODELING

### 5a) Sequential Modeling

```
[ ] model3 = Sequential()

model3.add(Conv2D(filters = 3,
                 kernel_size = (3, 3),
                 #strides = 1,
                 #padding = 'same',

                 activation = 'relu',
                 #activation is set to None by default
                 #but over here, we want to progressively see the effects
                 #of Convolution--> Relu --> Pooling
                 #so we use Relu

                 input_shape = alvin.shape))

model3.add(MaxPooling2D((2, 2)))
#2x2 receptive field is default and commonly used
```

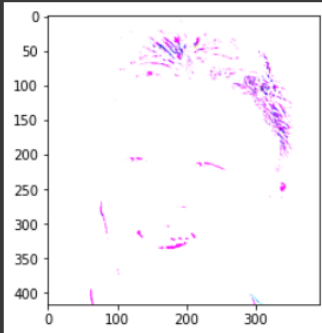
pooling layer  
is added

## 2. POOLING OUTPUT / EFFECT OF POOLING

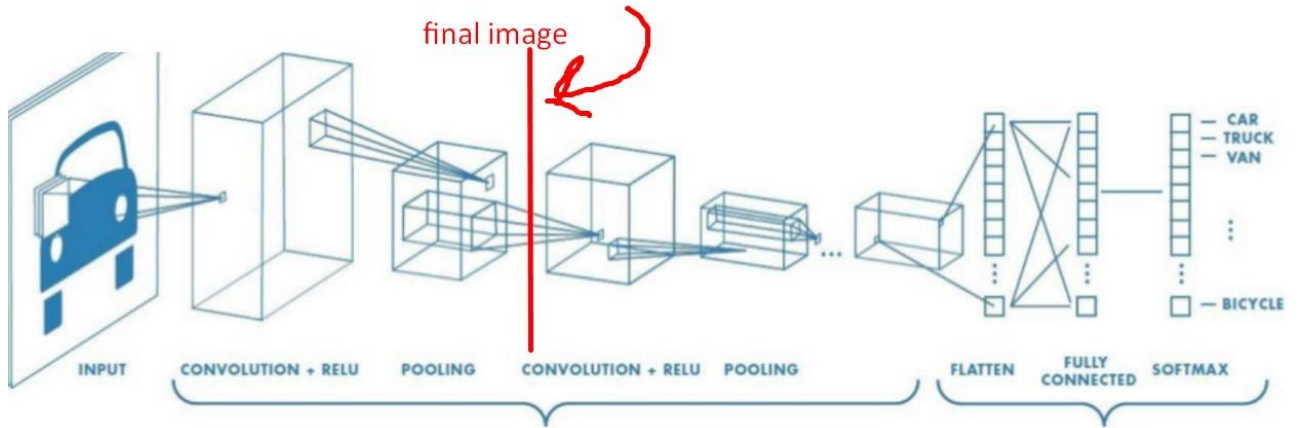
```
plt.imshow(conv_alvin_3)
```

```
#you can see that the Convo + Relu + Pooling extracted  
#my hair + eyes + teeth
```

Clipping input data to the valid range for imshow with RGB data ([0..1] for floats or [0..255] for integers).  
<matplotlib.image.AxesImage at 0x7f23e23cd590>



WE ARE HERE



---

## REFERENCES

---

- <https://towardsdatascience.com/a-comprehensive-guide-to-convolutional-neural-networks-the-eli5-way-3bd2b1164a53>



---

## ABOUT DR. ALVIN ANG

---



Dr. Alvin Ang earned his Ph.D., Masters and Bachelor degrees from NTU, Singapore. He is a scientist, entrepreneur, as well as a personal/business advisor. More about him at [www.AlvinAng.sg](http://www.AlvinAng.sg).