THURSDAY LEARNING HOUR

Computer Vision Using Deep Neural Networks

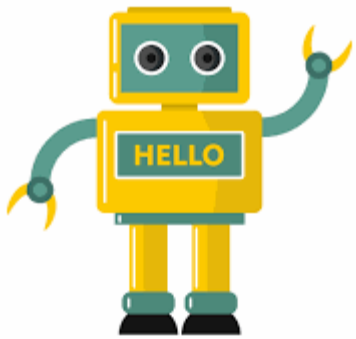

Prabakaran Chandran

Agenda

Computer Vision using Deep Neural Networks

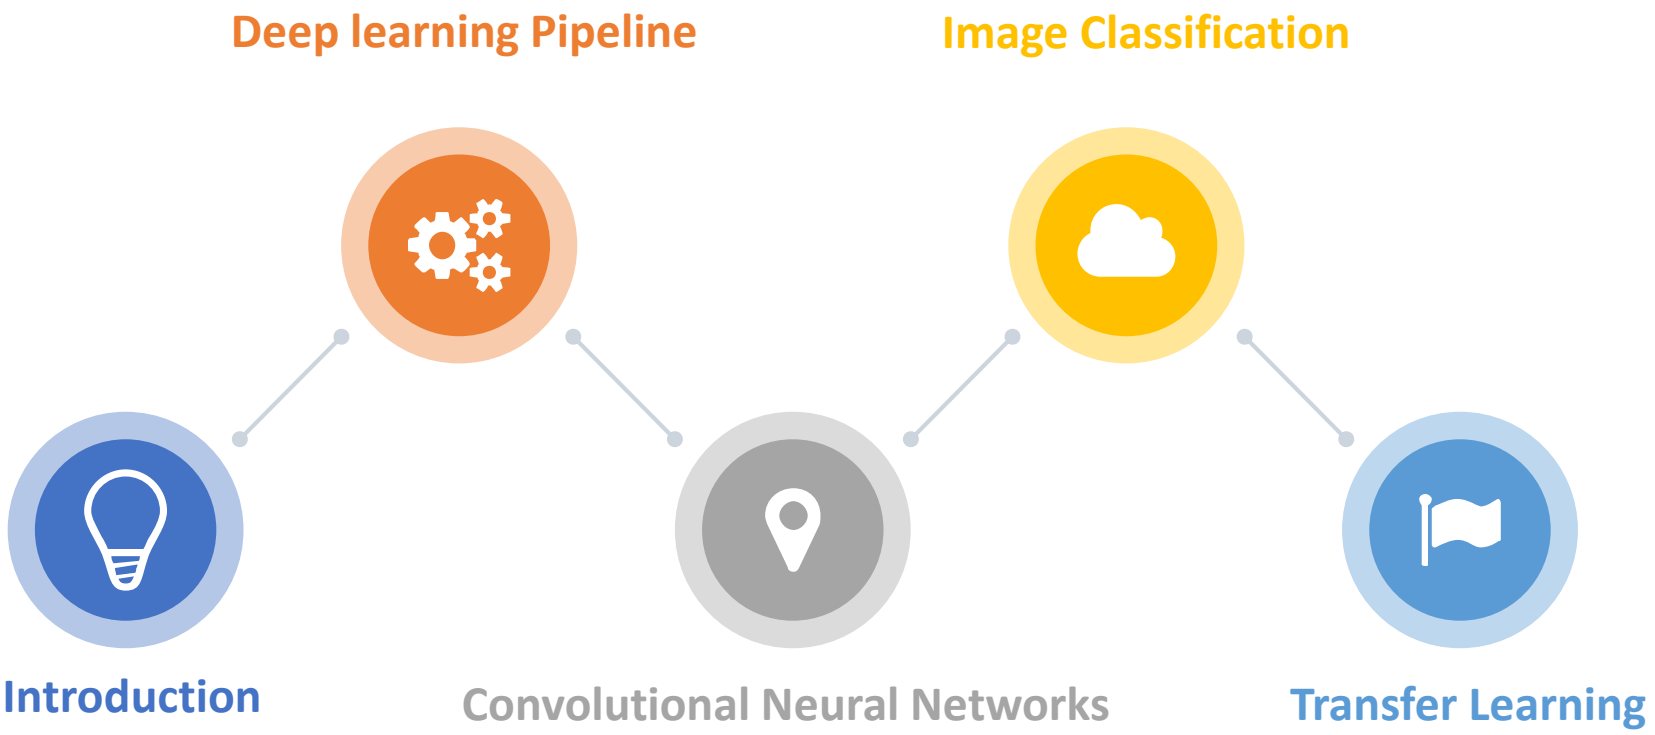

What is Computer Vision? :

- Computer vision is the field of computer science that focuses on replicating parts of the complexity of the human vision system and enabling computers to identify and process objects in images and videos in the same way that humans do

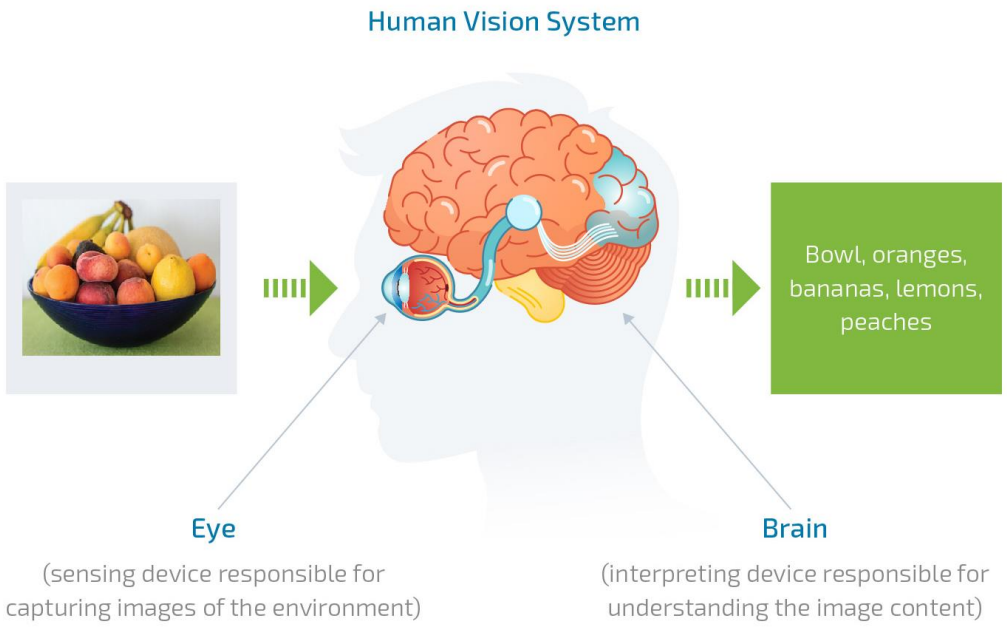

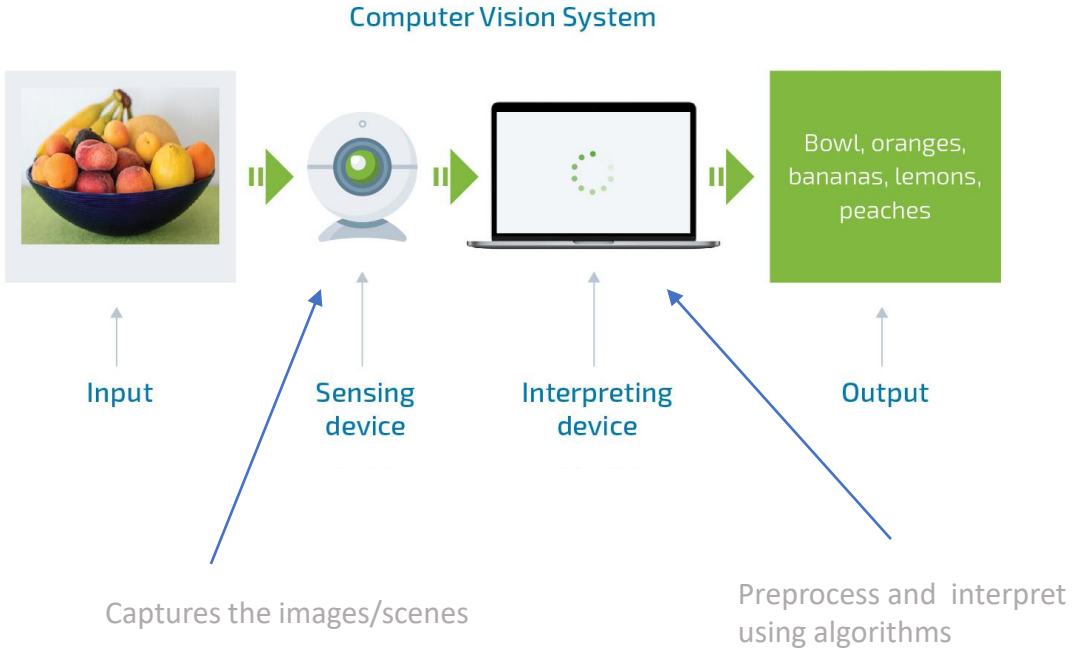

Major Computer Vision tasks:

Classification

CAT

Classification + Localization

CAT

Object Detection

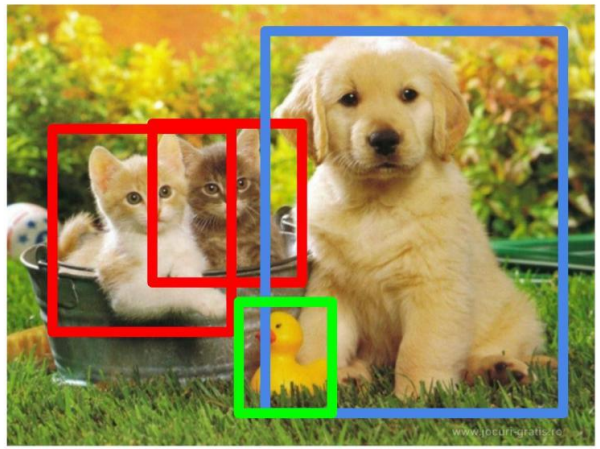

CAT, DOG, DUCK

Instance Segmentation

CAT, DOG, DUCK

Single object

Multiple objects

Computer Vision Tasks:

Classification

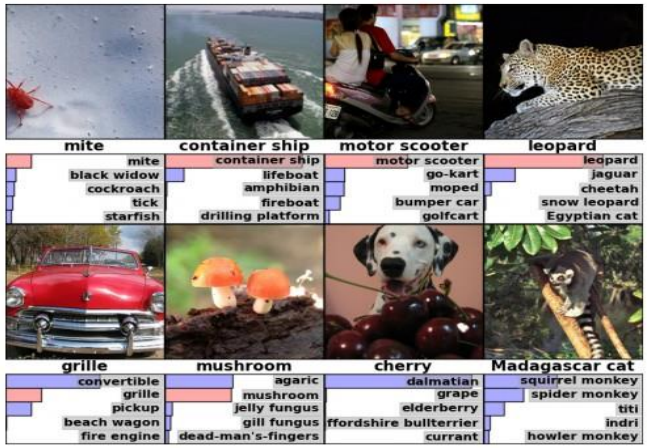

Detection

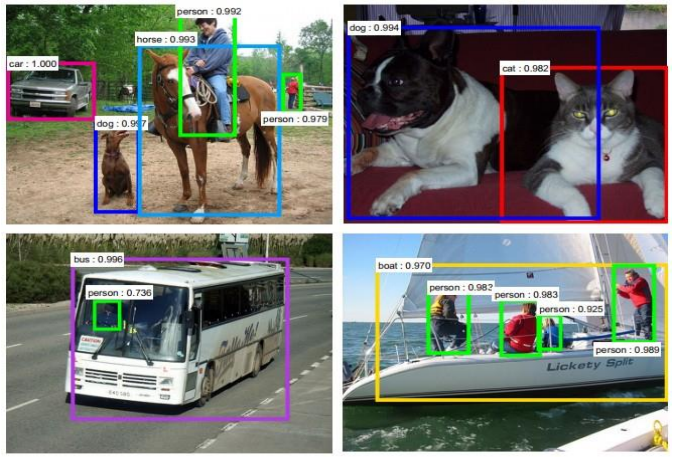

Segmentation

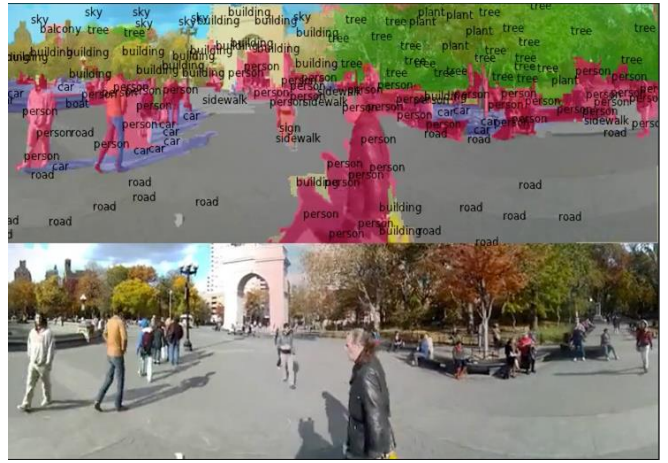

Image Captioning

Image Cartooning

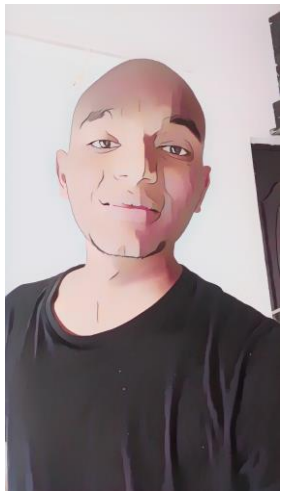

How can a Machine do Brain's work ? – Artificial Neural Network

- Artificial Neural Networks or ANN is an information processing paradigm that is inspired by the way the biological nervous system such as brain process information
- It is composed of large number of highly interconnected processing elements(neurons) working in unison to solve a specific problem

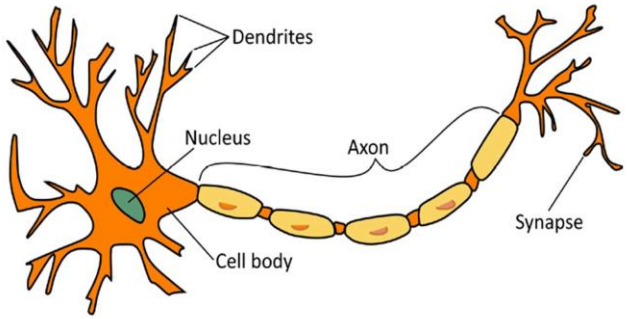

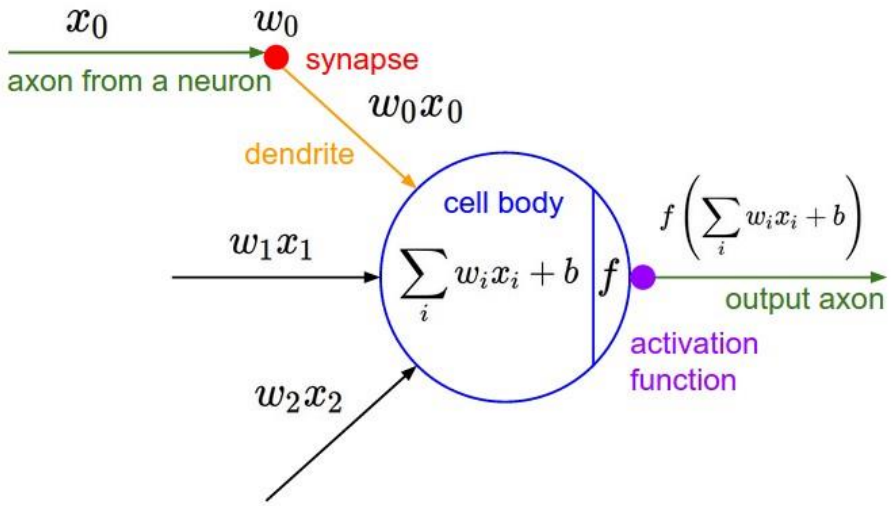

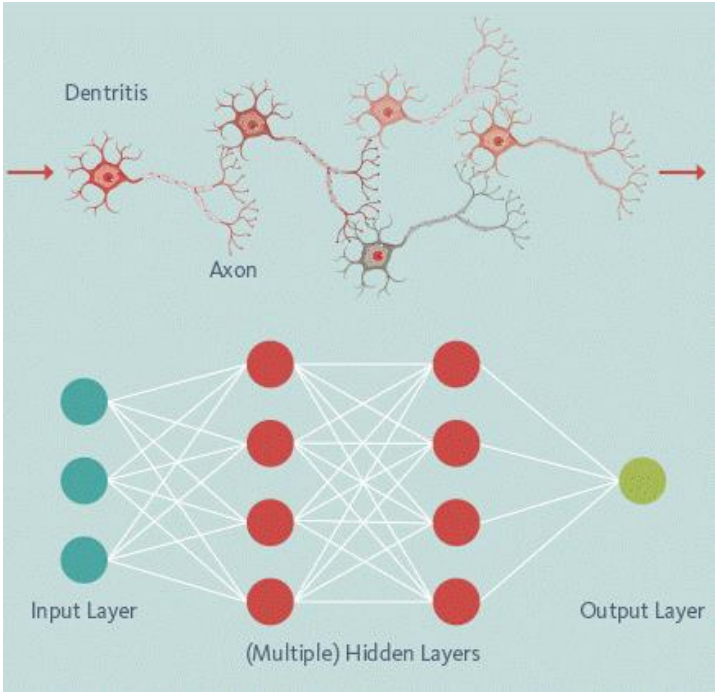

Neuron

Neural Network

How does a DNN learn things? – Fundamental Work flow

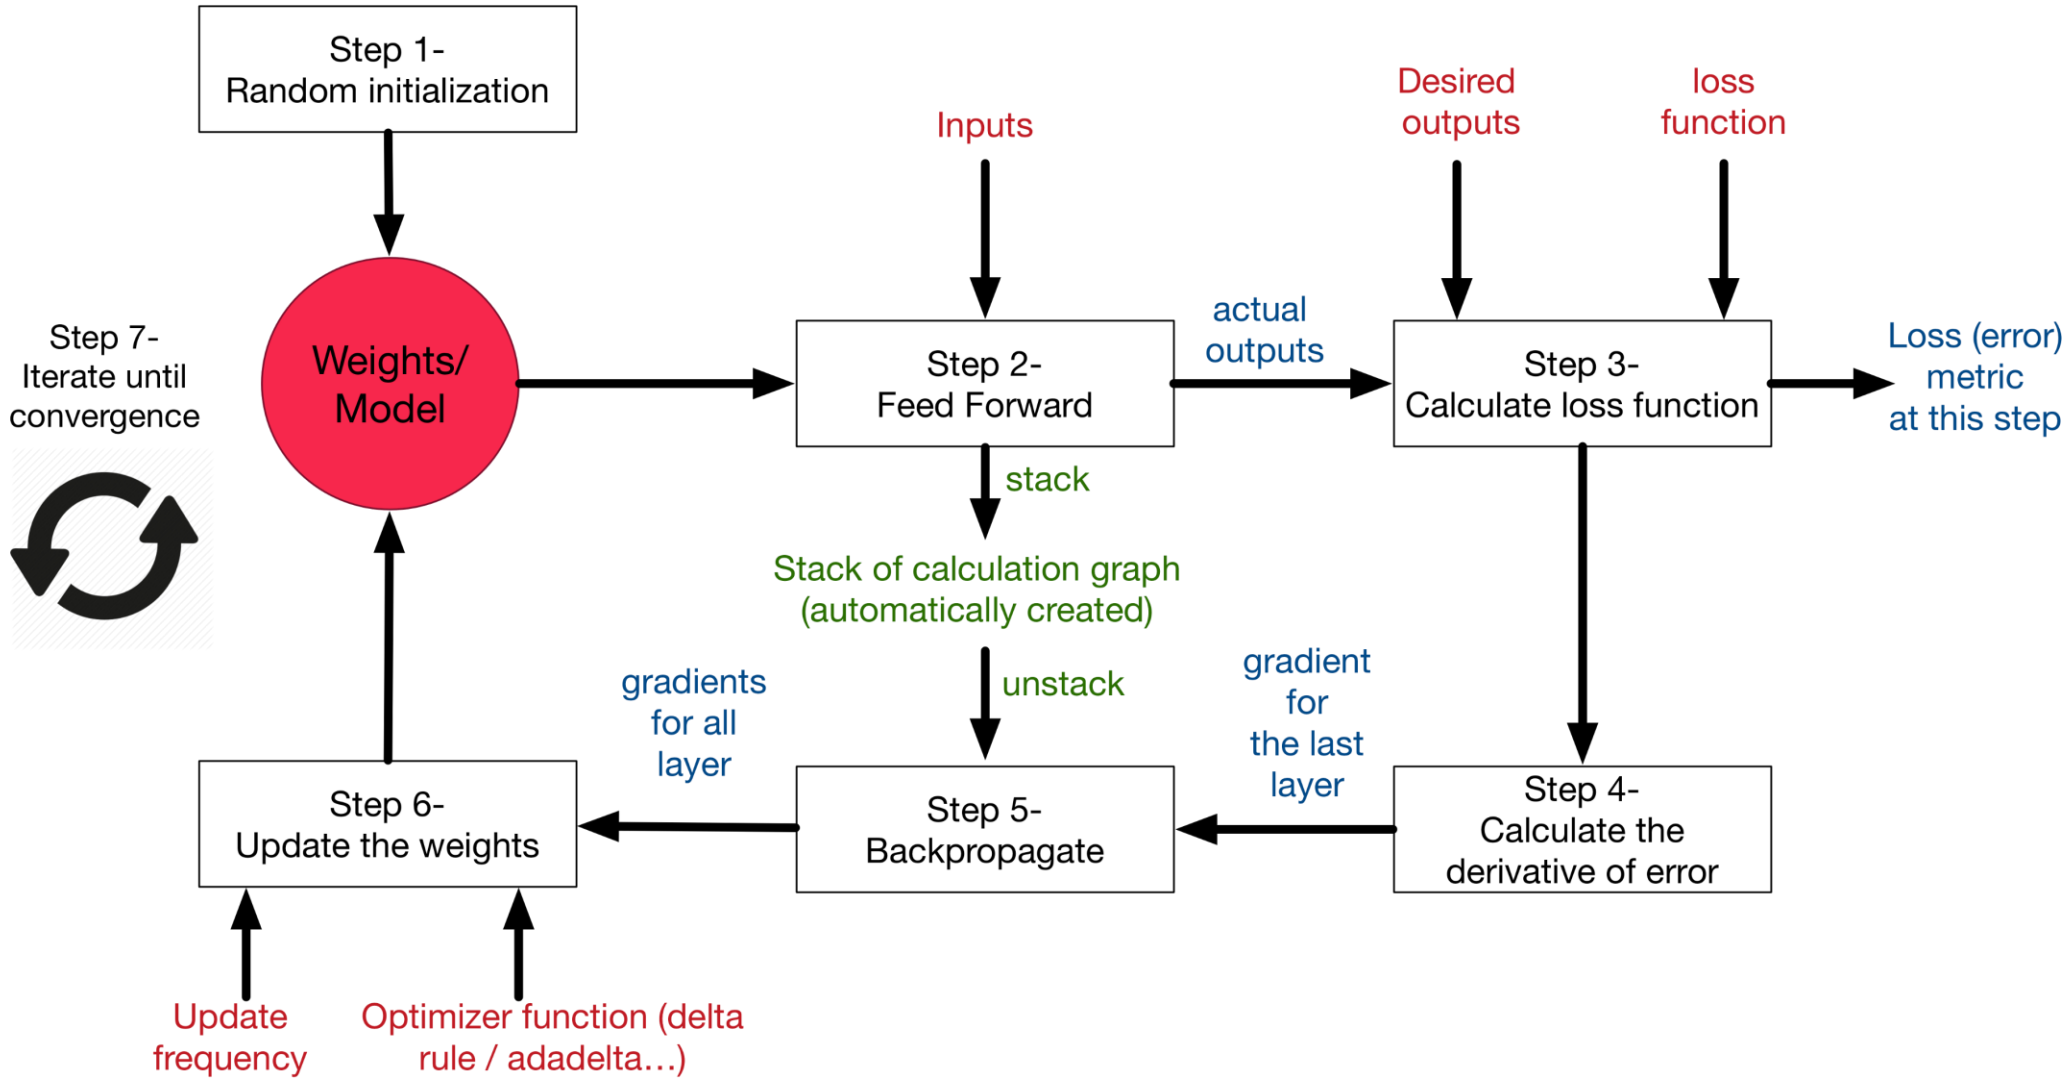

Fundamental Work flow and Key terms

I. Forward-propagate Input Signal

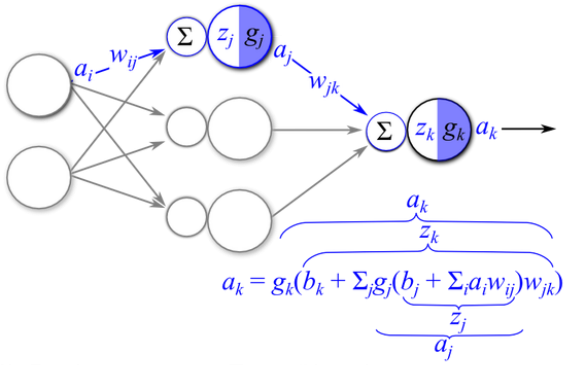

II. Back-propagate Error Signals

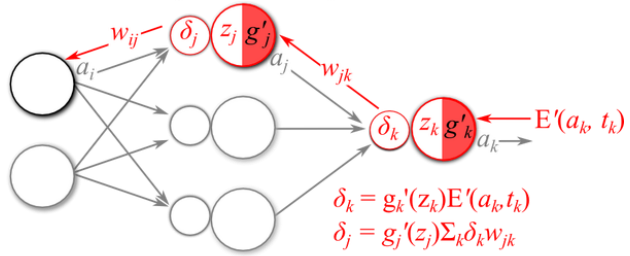

III. Calculate Parameter Gradients

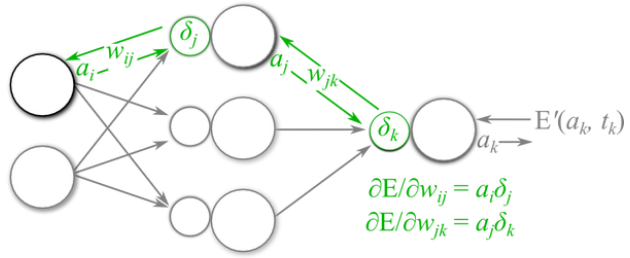

IV. Update Parameters

$w_{ij} = w_{ij} - \eta(\partial E/\partial\, w_{ij})$

$w_{jk} = w_{jk} - \eta(\partial E/\partial w_{jk})$

for learning rate η

a_i $-$ *inputs* , w_i $-$ *weights are the actual learning paramerters*

Σ $-$ *function that sums up input* $*$ *weights*

g_i, g_k $-$ *activation function* (*Relu* , *softmax* , tanh, *and sigmoid*)

b_i $-$ *bias term*

Forward pass calculates the output - a_k t_k $-$ *target E or J* $-$ *loss*

The aim of the back propagation is to calculate gradients to update the learning parameters which reduce the loss J

Loss : Negative log likelihood , Categorical Cross entropy , MSE

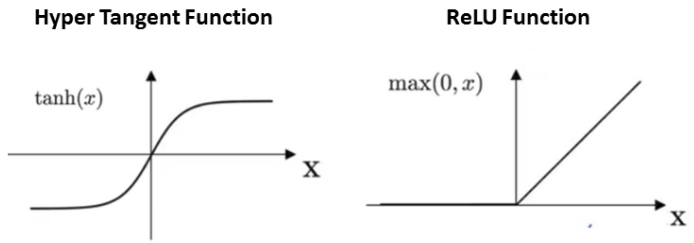

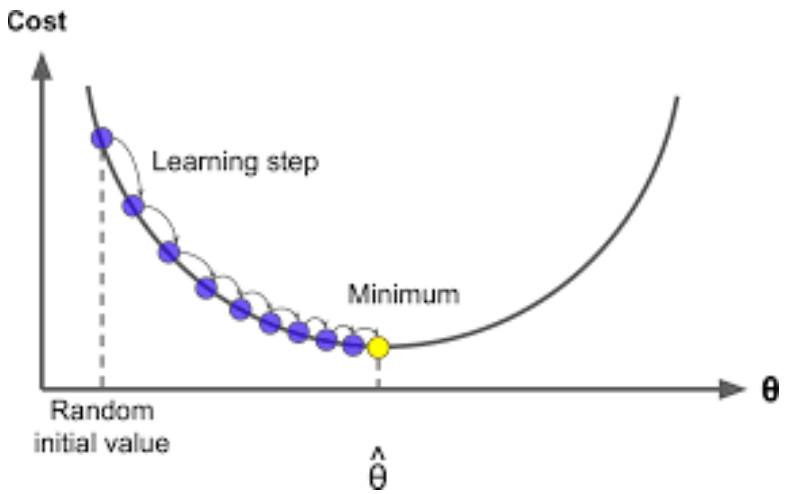

Fundamental Work flow – Loss function and Optimizers

Cross-entropy loss

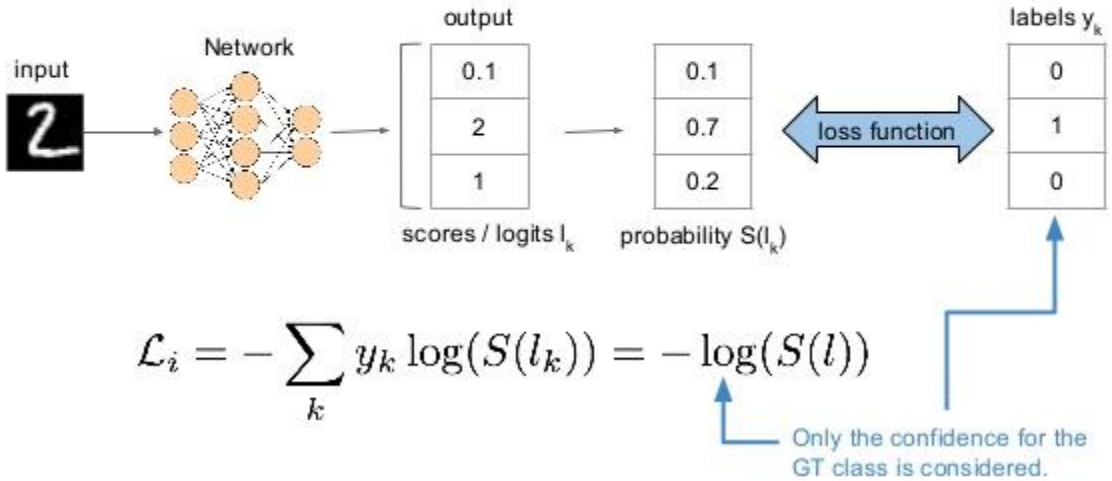

33

Other loss functions:
1.Negative log likelihood loss
2.Binary cross entropy
3.Mean square error
4.Mean absolute error
5.Hinge loss

Stochastic Gradient descent

Randomly shuffle (reorder) training examples

Repeat {
  for $i := 1, \ldots, m$ {
    $\theta_j := \theta_j - \alpha(h_\theta(x^{(i)}) - y^{(i)})x_j^{(i)}$
    (for every $j = 0, \ldots, n$)
  }
}

Other optimizers:
1.SGD with momentum
2.Adagrad
3.AdaDelta
4.Adam
5.RMSprop

Computer Vision + Deep learning : Convolutional Neural Network.

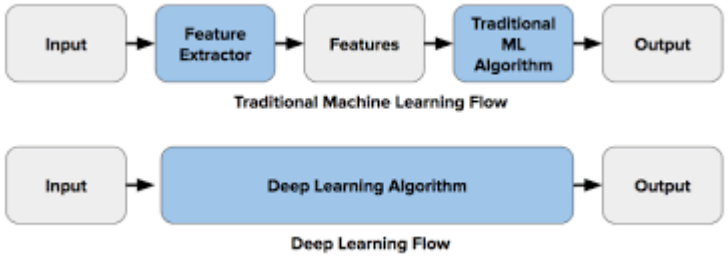

Traditional way of CV extracts features using some explicit tasks like wavelet transformation , image processing , but in Deep learning everything is being handled by the network itself.

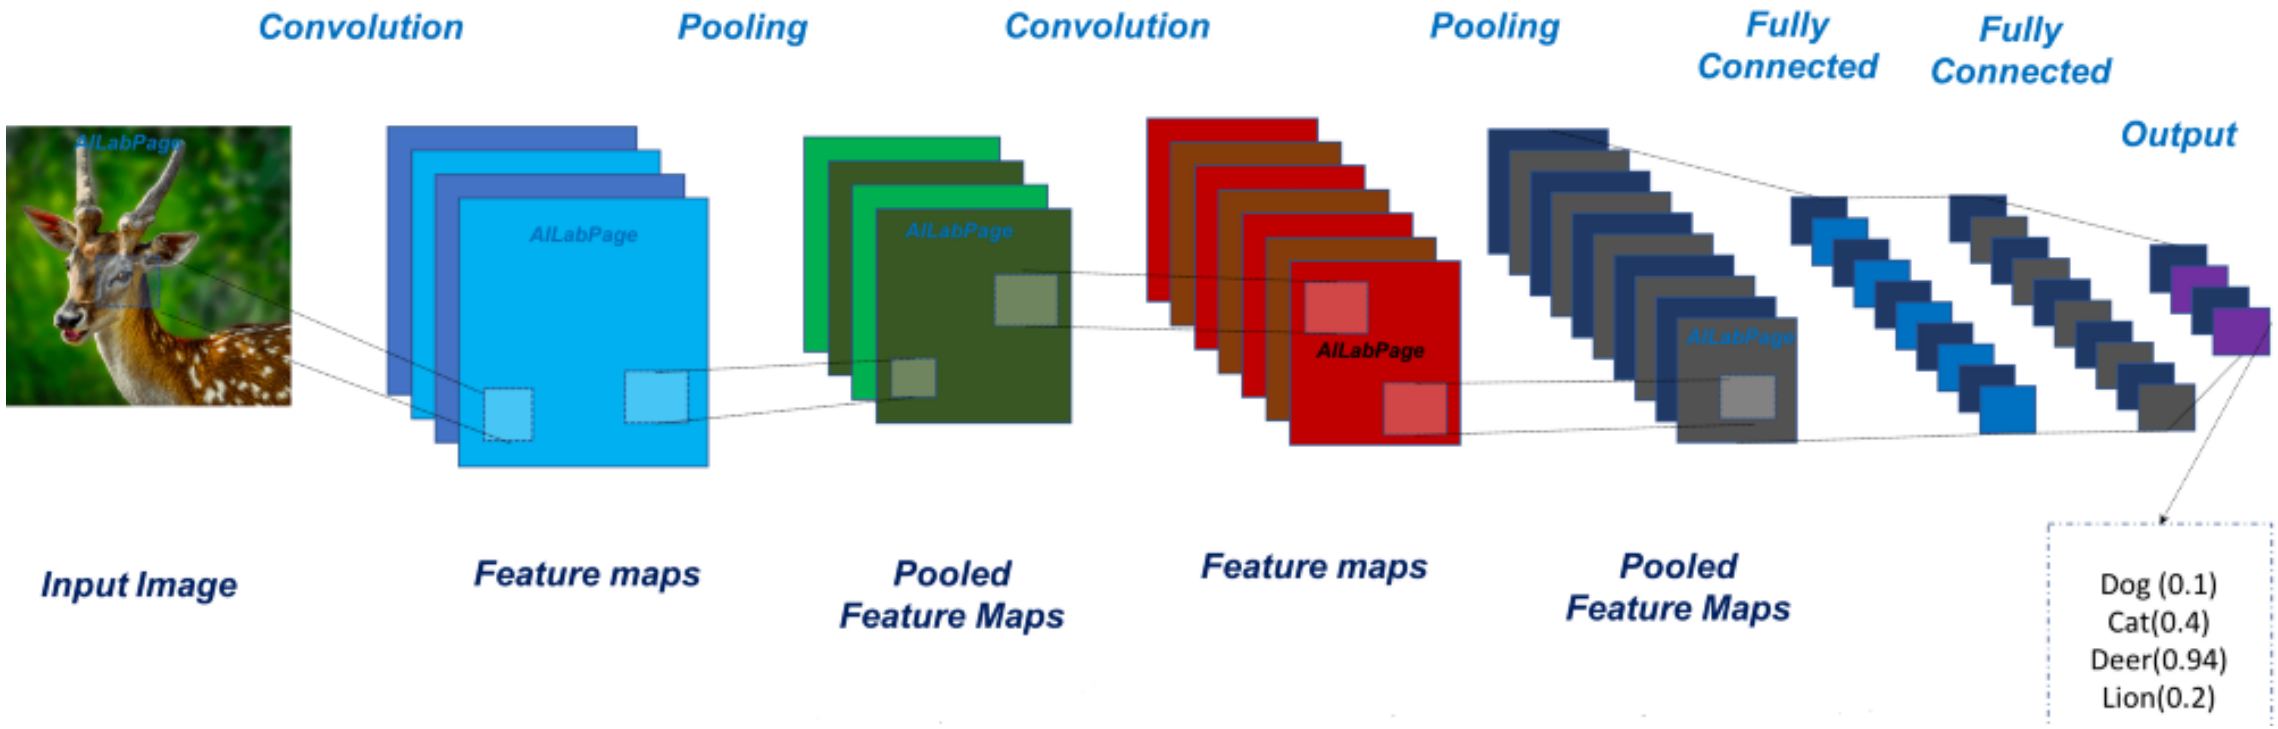

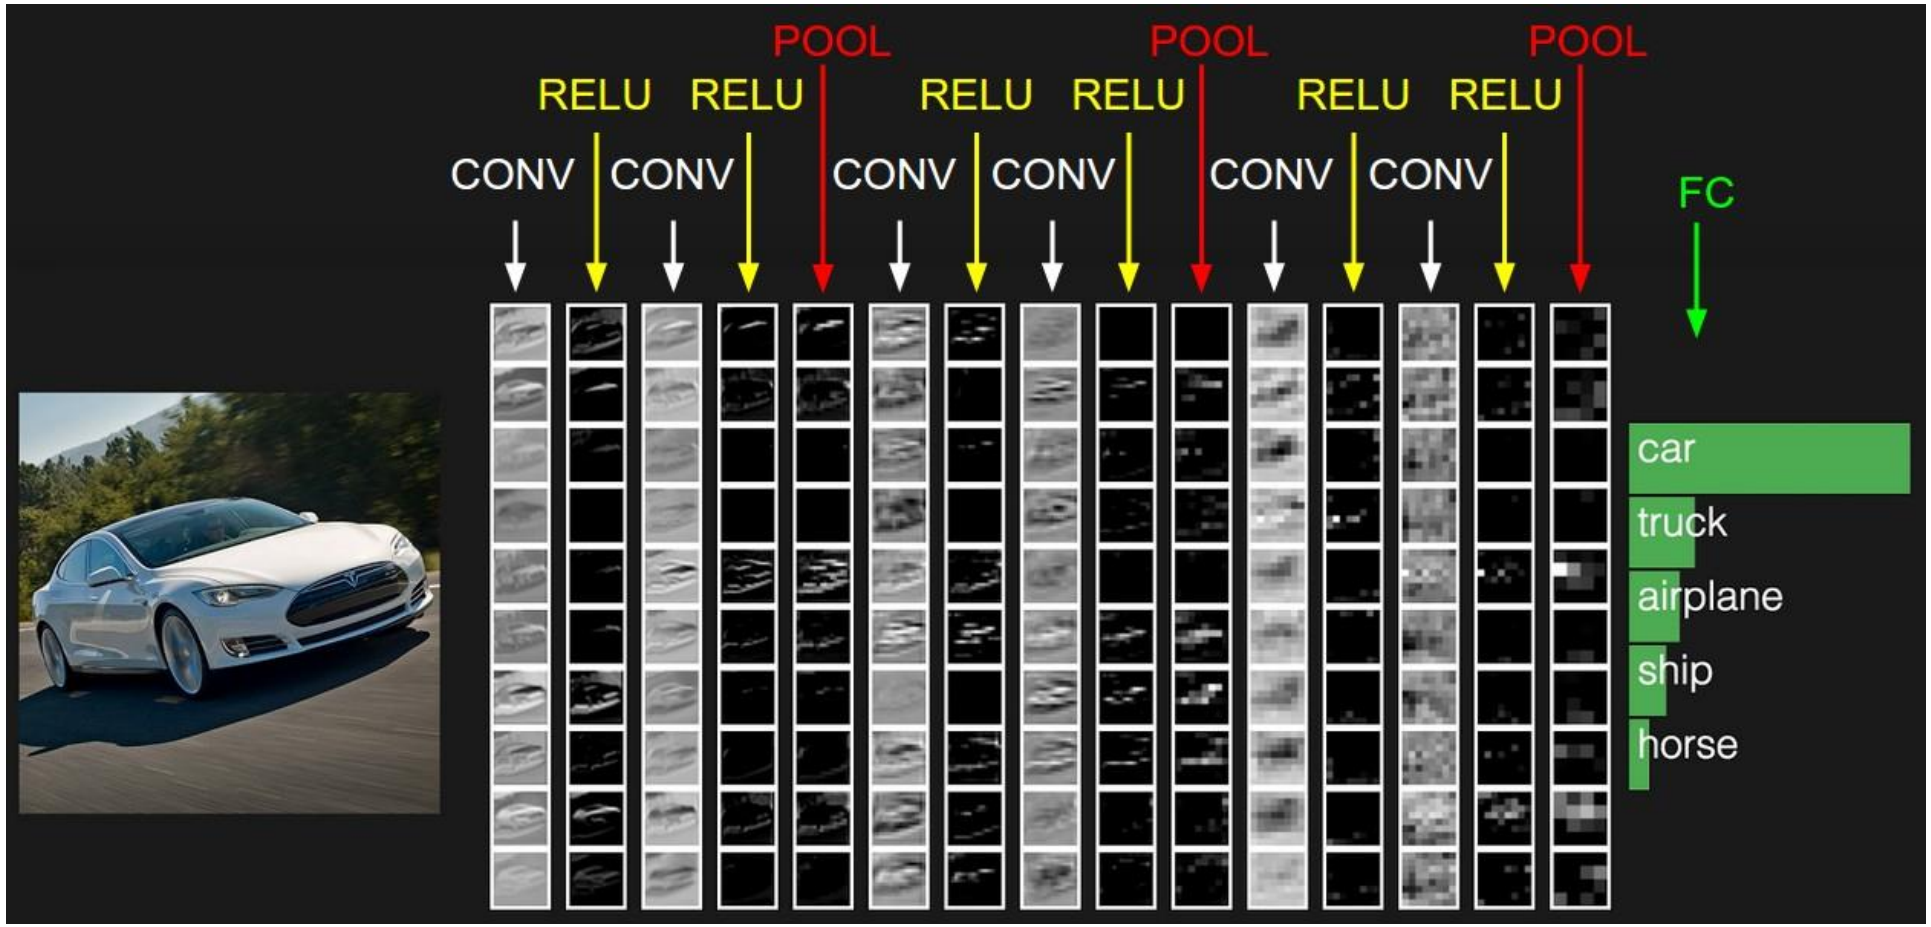
POOL
POOL
POOL
RELU
RELU
RELU
RELU
RELU
RELU
CONV
CONV
CONV
CONV
CONV
CONV
FC
car
truck
airplane
ship
horse

Convolutional Neural Networks – different layers of CNN

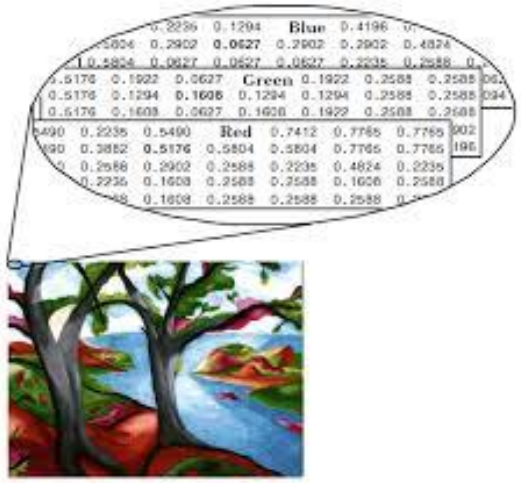

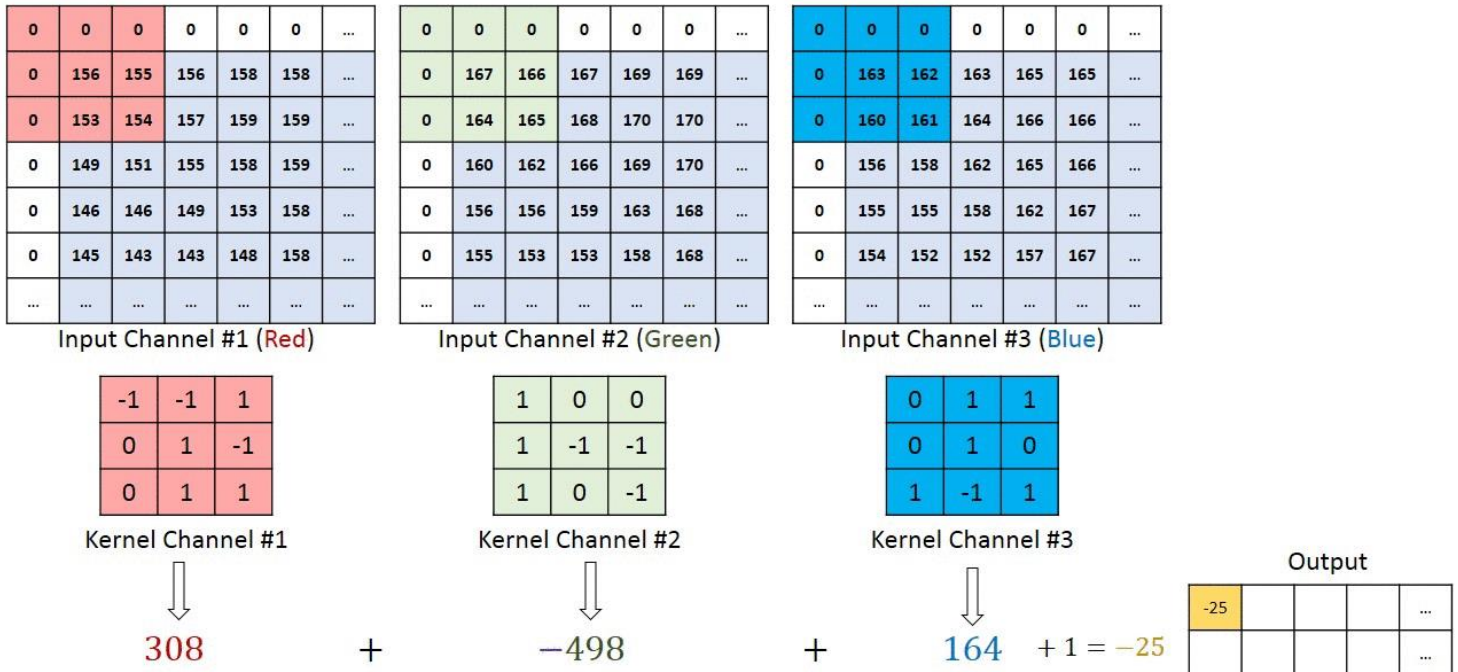

Action :

- Apply filters to extract features
- Filters are composed of small kernels, learned
- One bias per filter
- Apply activation function on every value of feature map

Parameters

- Number of Kernels , size of kernels
- Activation function ,striding , padding

Convolutional Neural Networks – different layers of CNN

Before and after Convolution

52	34	14	5
45	12	17	11
29	20	19	27
99	85	60	55
120	112	88	29

131	122	97	17
114	104	78	88
74	144	99	112
156	213	113	98
146	177	120	130

Edge detection

Kernel

 * $\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$ = 

Sharpen

 * $\begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix}$ = 

Convolutional Neural Networks – different layers of CNN

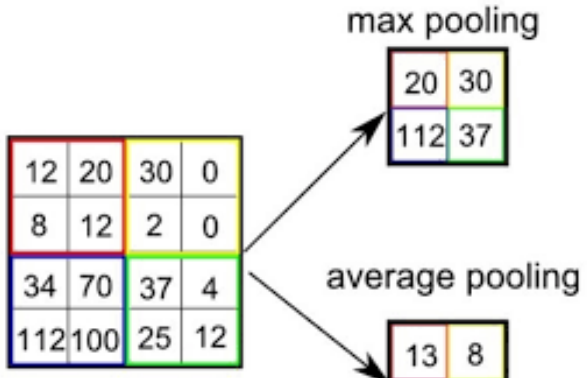

Max Pooling

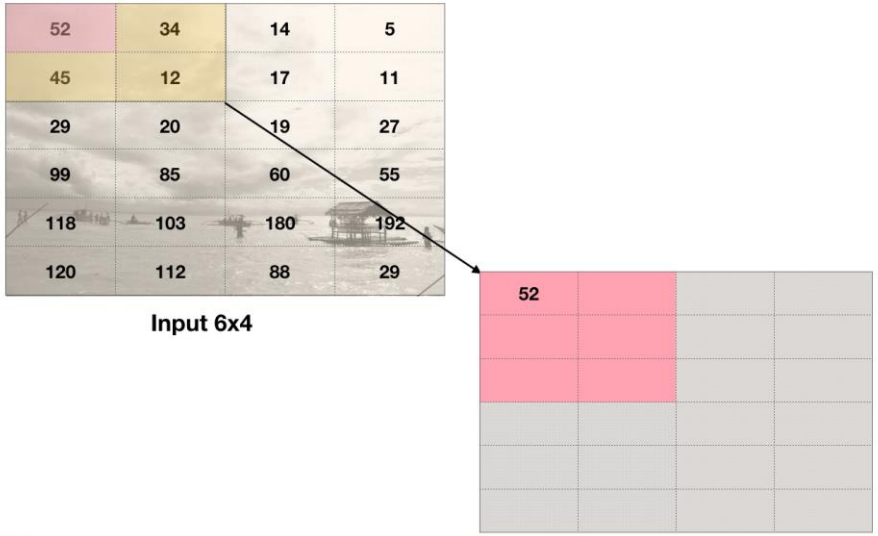

Input 6x4

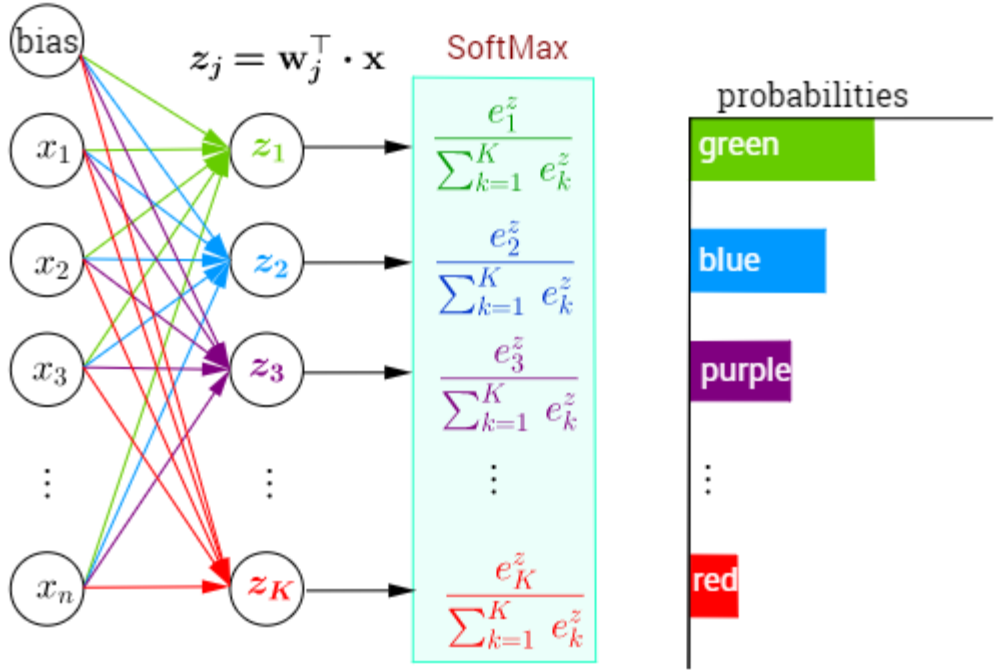

Pooling layers:

- Reduce Dimensionality
- Extract maximum of / average of a region
- Follow Sliding window approach

Fully connected layers:

- Aggregate information from final feature maps
- Flatten the feature maps for final classification
- Generate final classification with the use of Sigmoid/SoftMax

CNN for Image classification – ImageNet Competition

From ~2012 we see that Convolutional Neural Network (CNN) have become an important tool for object recognition

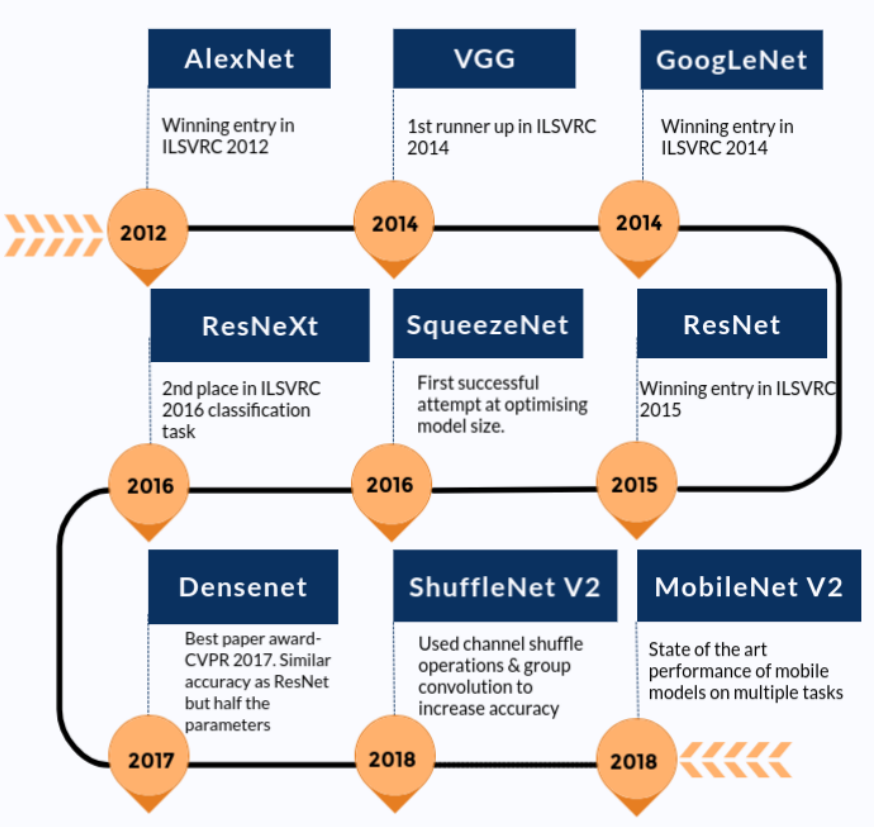

IMAGENET **Large Scale Visual Recognition Challenge**

The Image Classification Challenge:
1,000 object classes
1,431,167 images

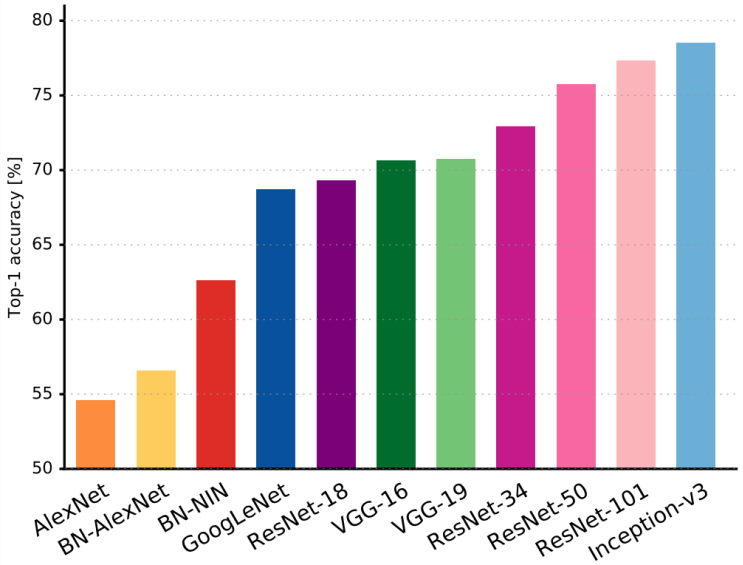

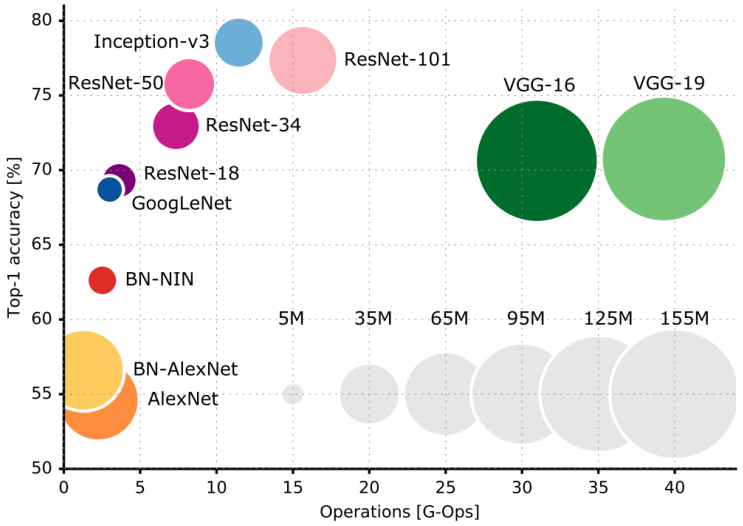

CNN for Image Classification : Various architectures for Image Classification

LeNet (1998)

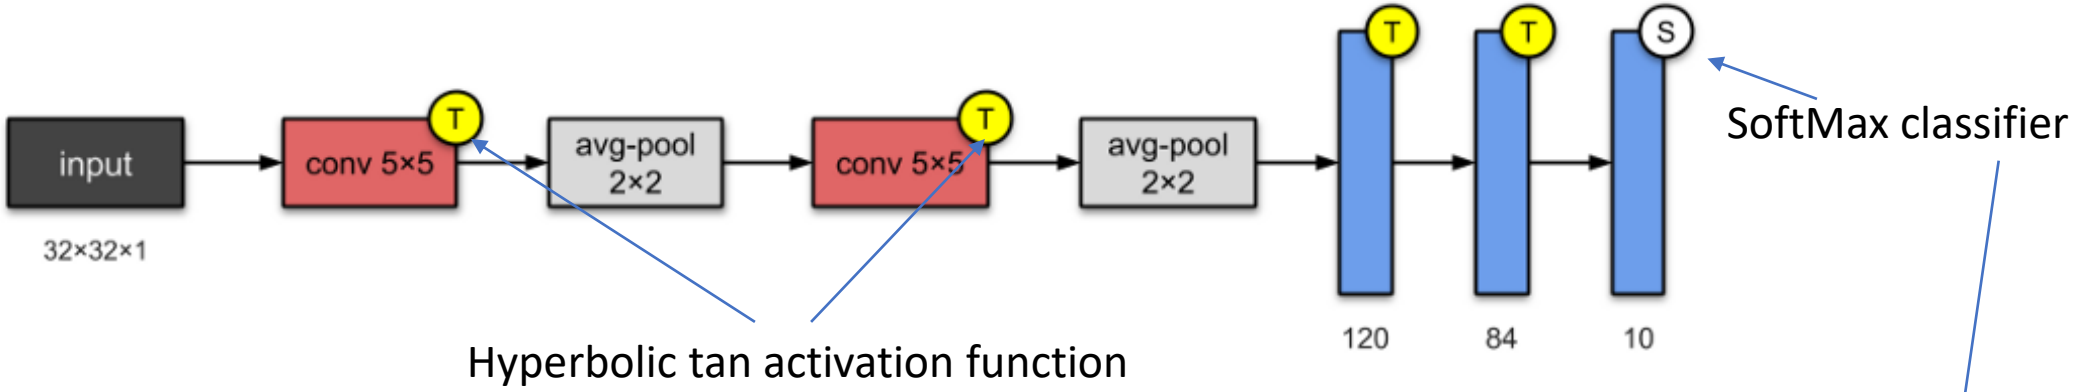

Alexnet (2012)

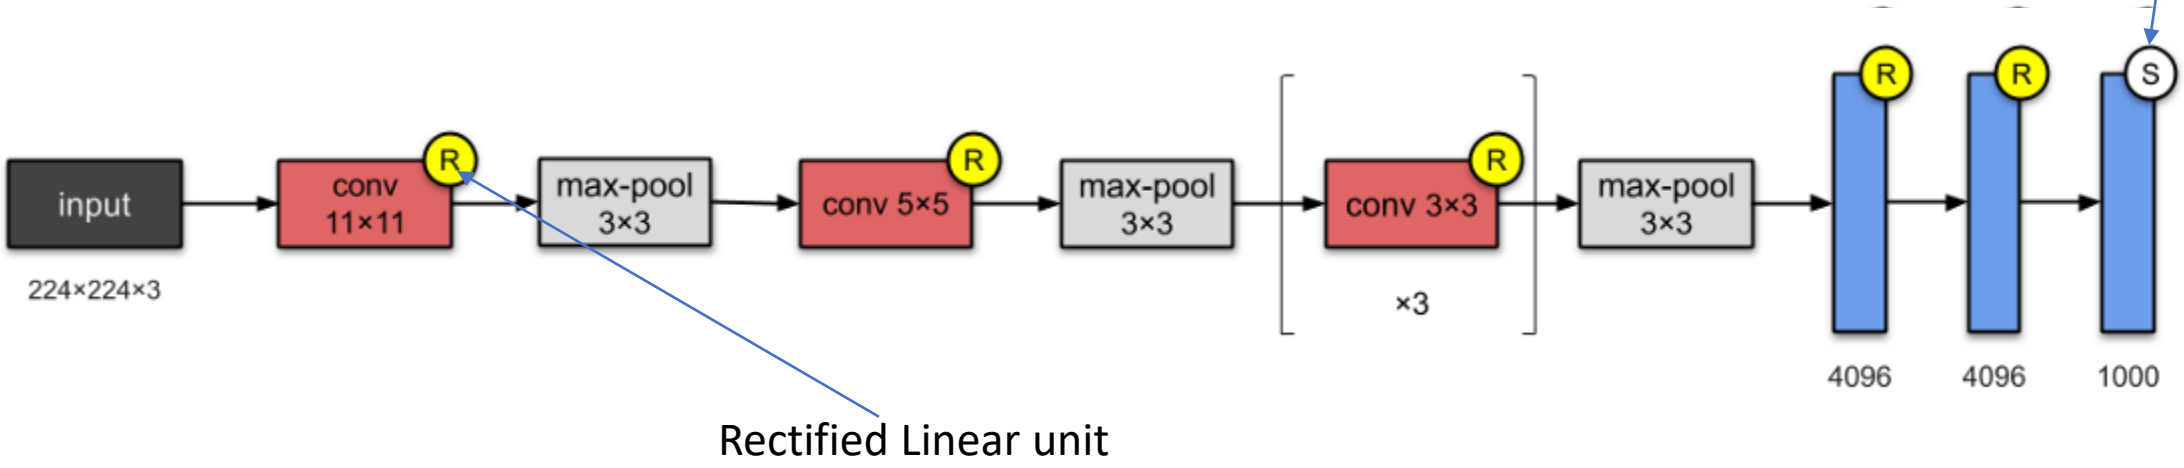

CNN for Image Classification : Various architectures for Image Classification

VGG-16 Visual Geometry Group (2014)

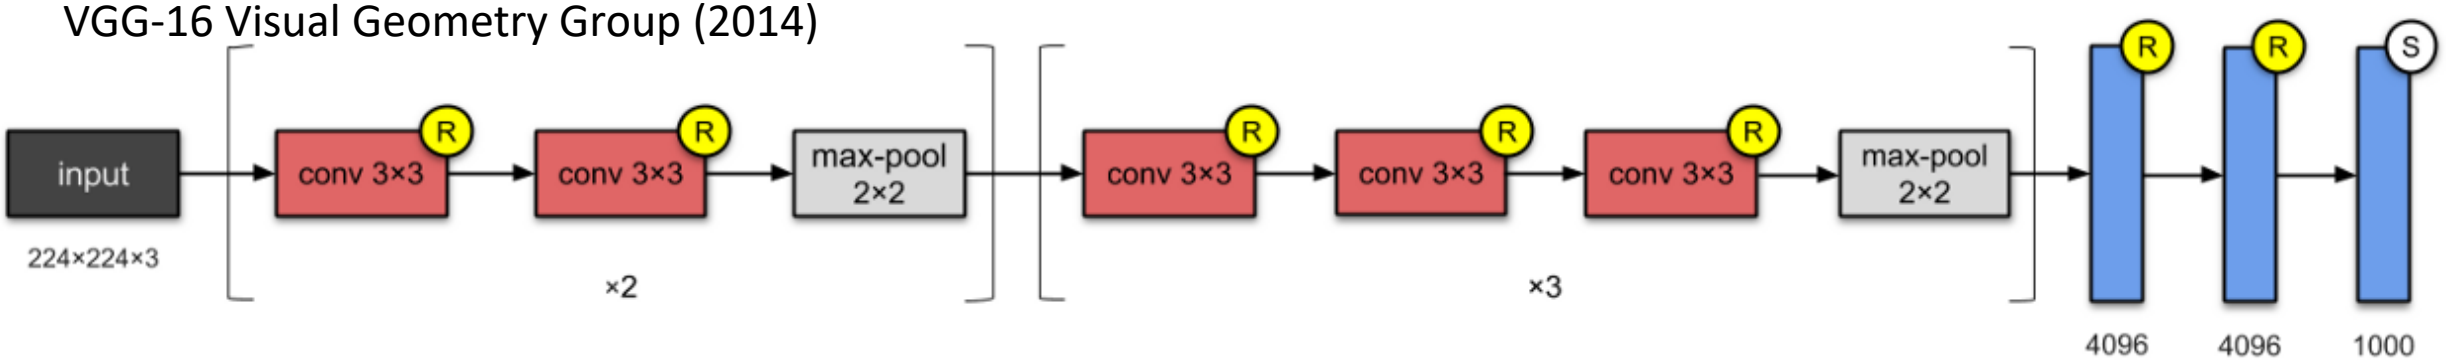

Residual Net – 50 (2015)

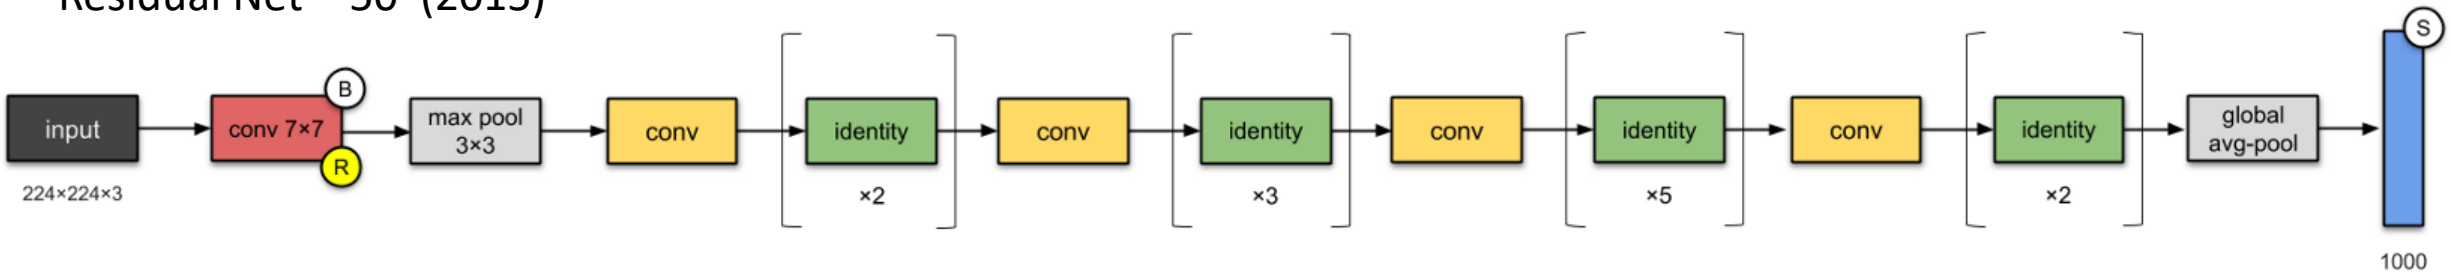

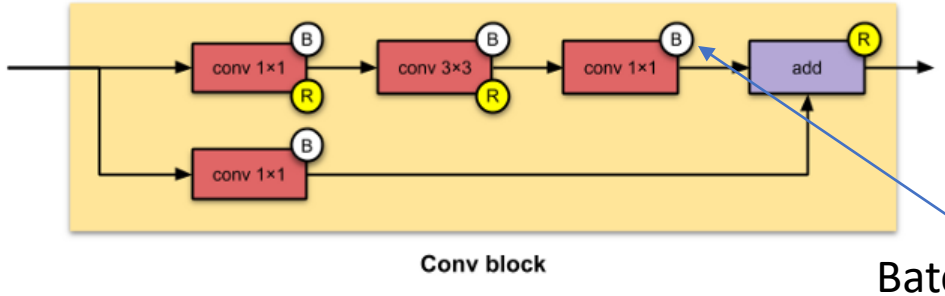

Conv block

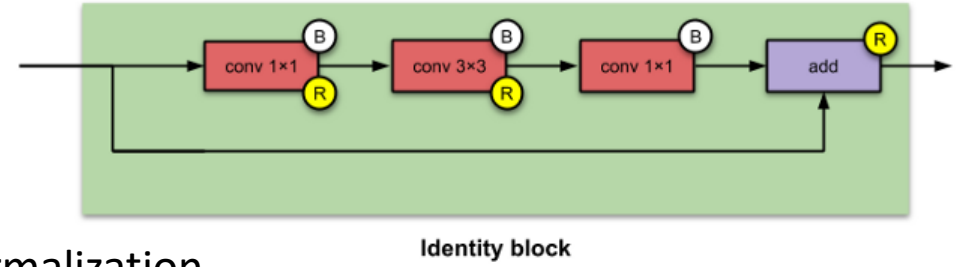

Identity block

Batch Normalization

CNN for Image Classification : Transfer Learning

Transfer learning or fine tuning refers to training a network on a huge standard data set (e.g. ImageNet) and then re-tuning just the last few layers of the network for the required specific task

Pretrained models can be downloaded from the model zoos

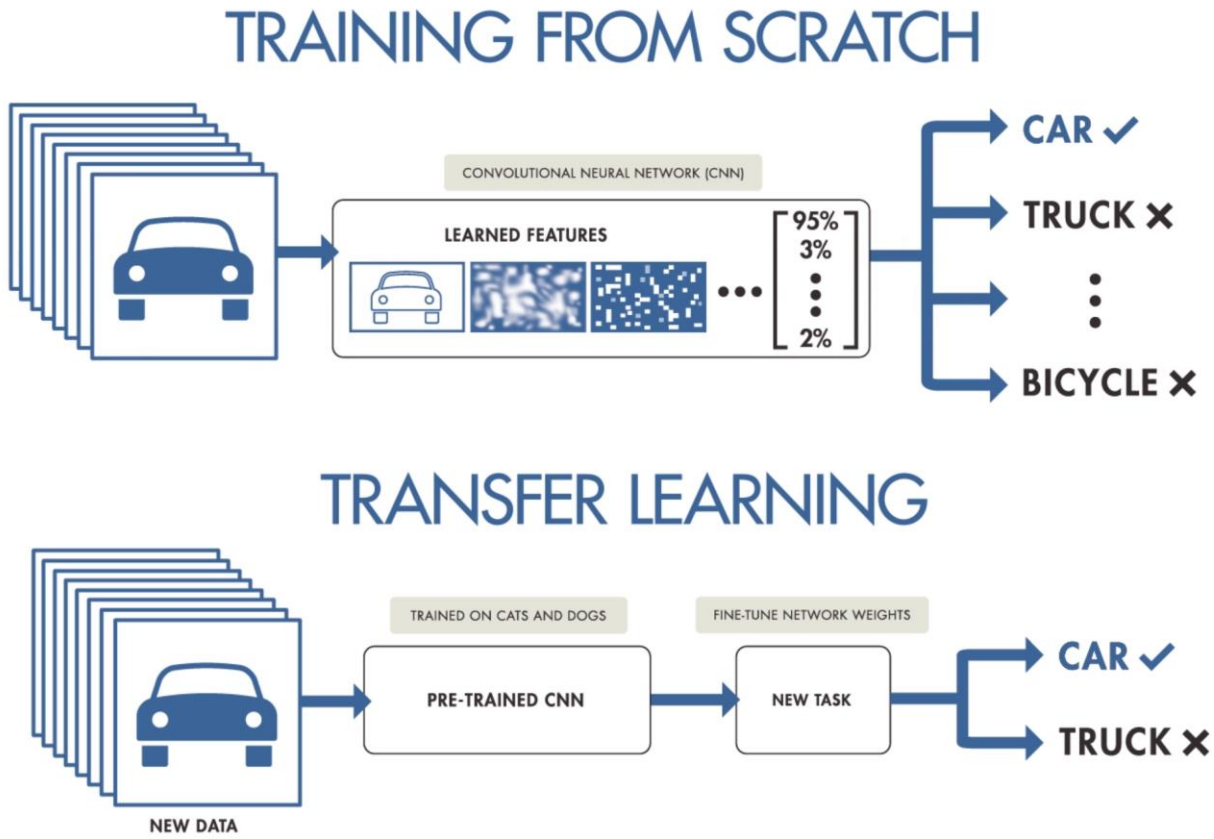

Caffe2 · Microsoft CNTK · TensorFlow · theano · Chainer · PYTORCH · dmlc mxnet · K

A Deep learning and CNN based Image Classification pipeline

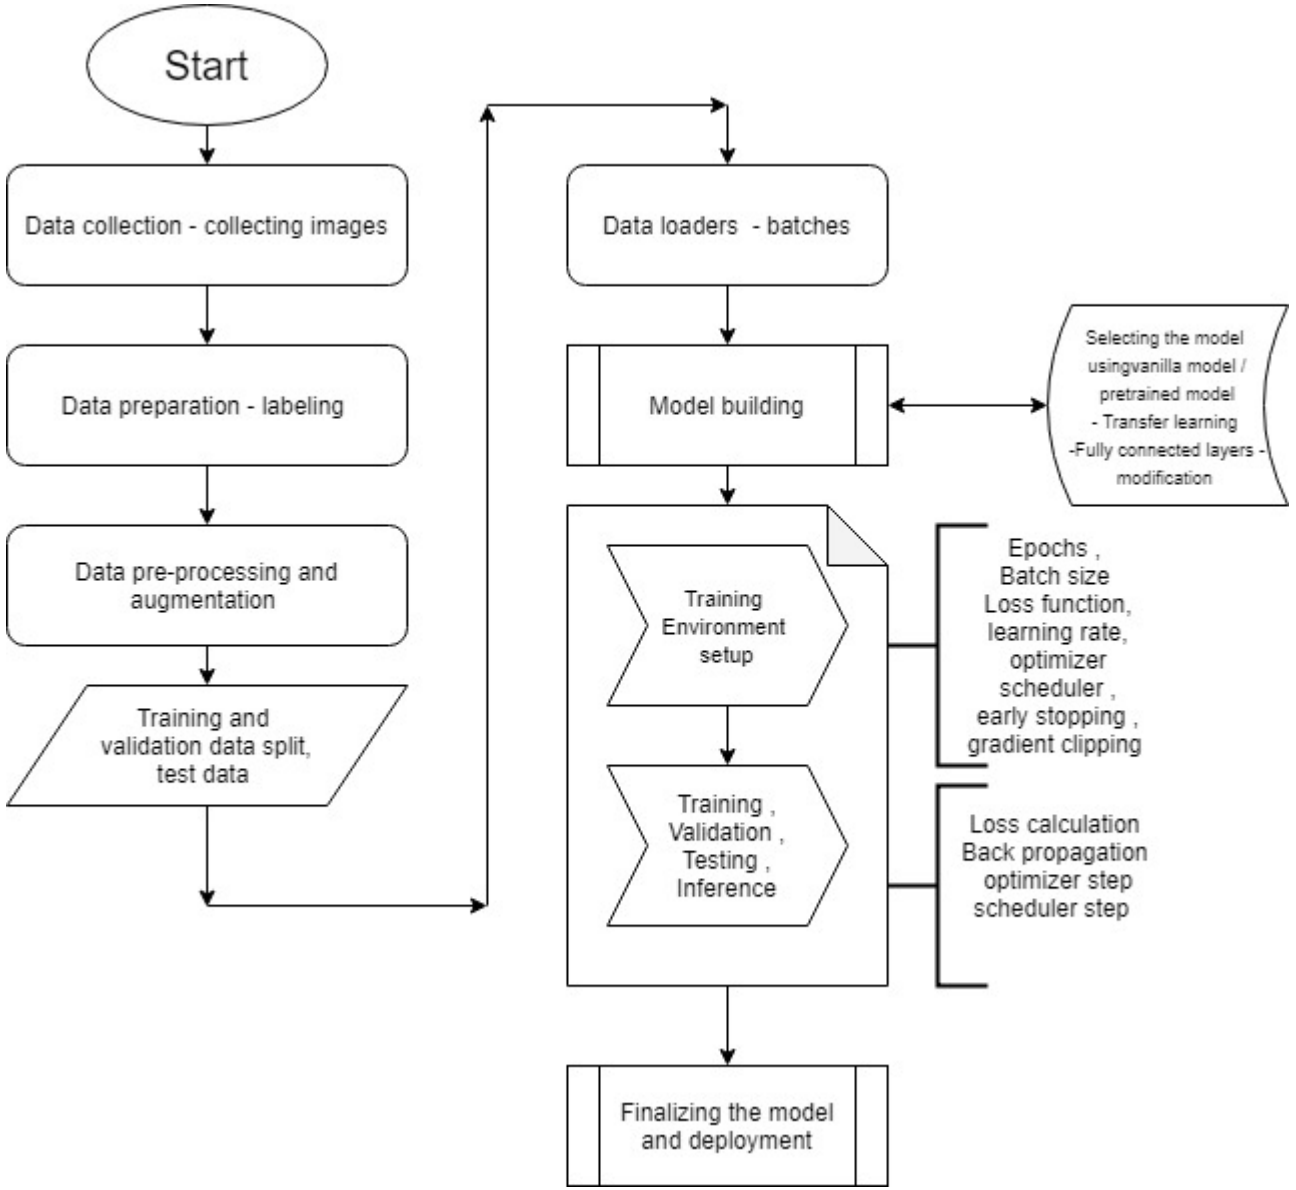

Appendix

Questions!