IESTI01 – TinyML

Embedded Machine Learning

11. Image Classification using Convolutions (CNN)

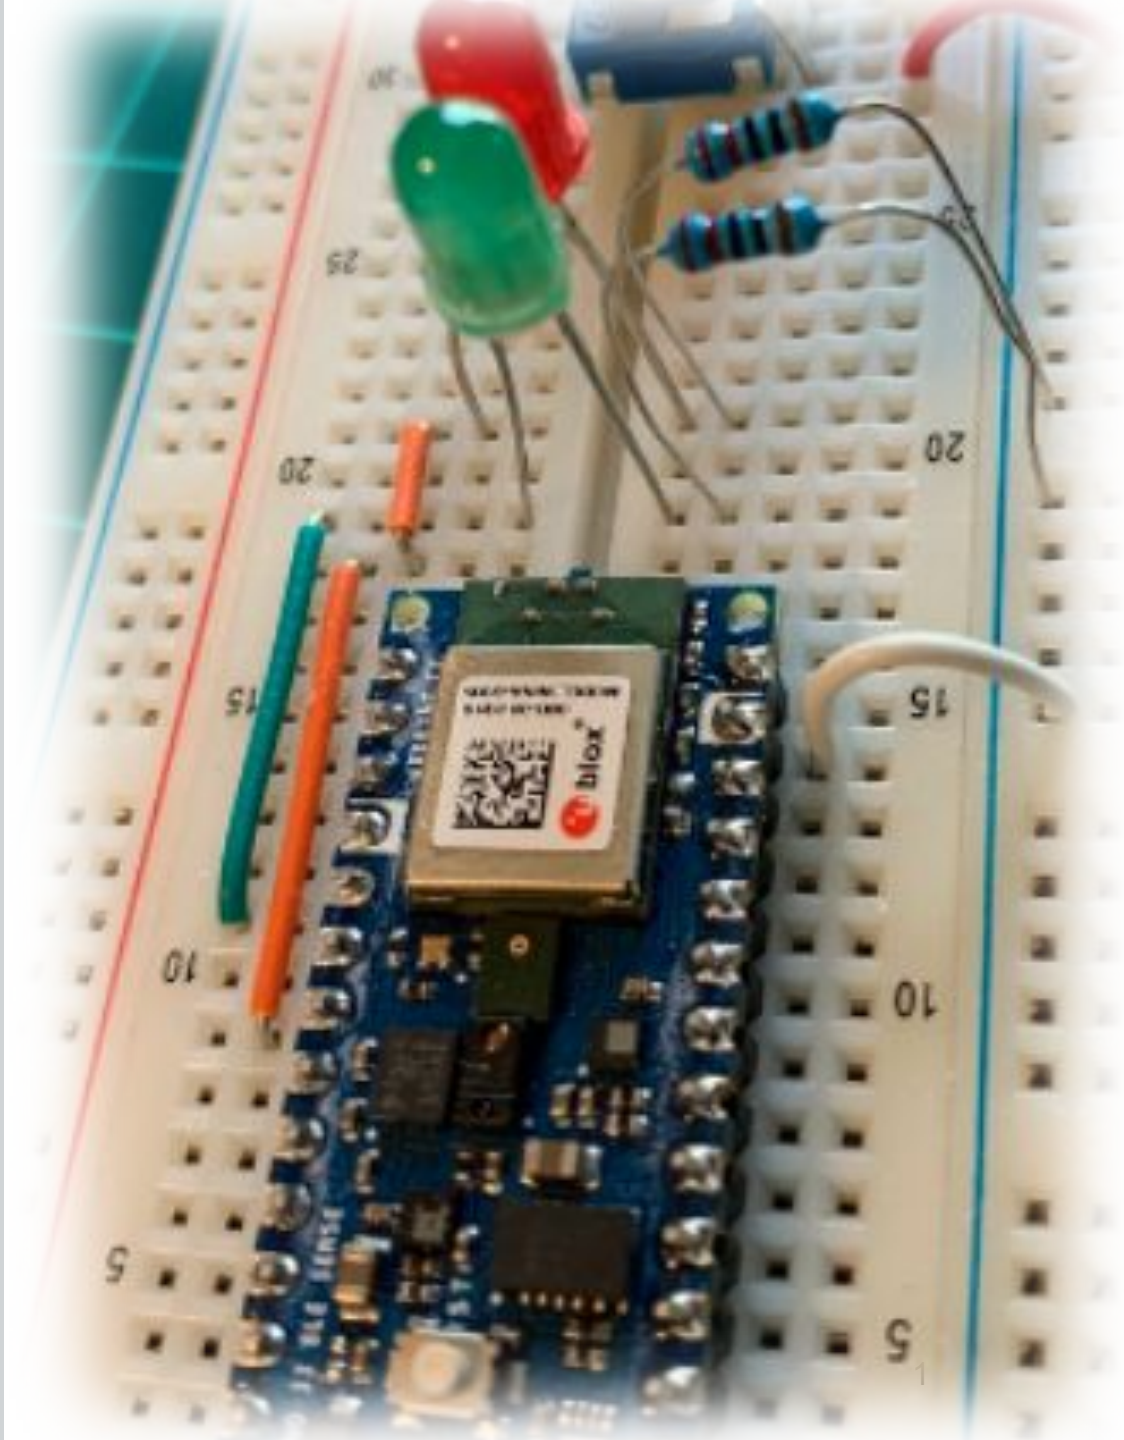

Prof. Marcelo Rovai
UNIFEI

Convolutions

CNN EXPLAINER:

Learning Convolutional Neural Networks with Interactive Visualization

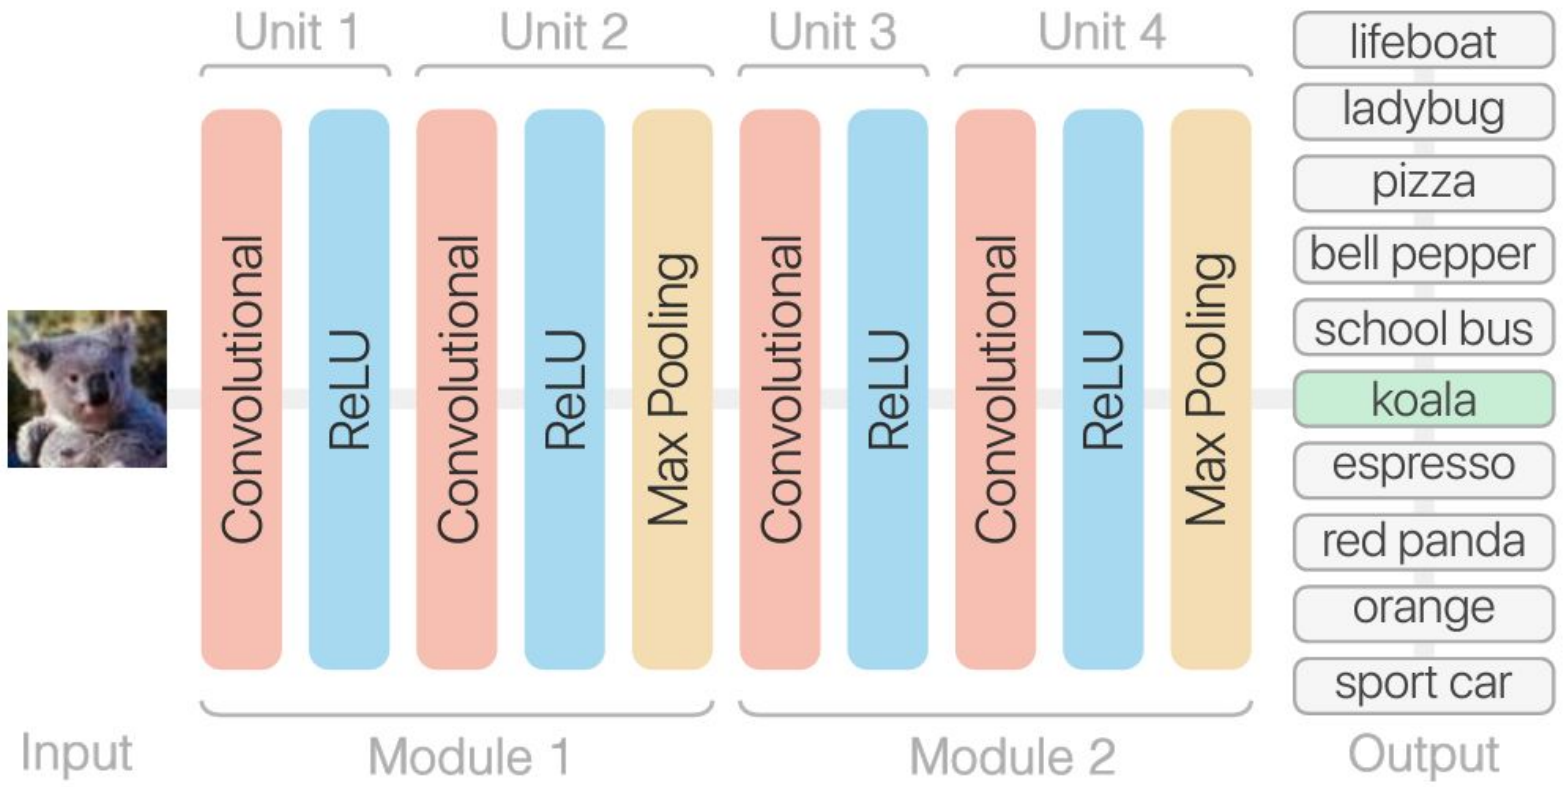

https://poloclub.github.io/cnn-explainer/

The convolution operation

The fundamental difference between a densely connected layer and a convolution layer is this: **Dense layers learn global patterns** in their input feature space (for example, for an MNIST digit, patterns involving all pixels), whereas **convolution layers learn local patterns**—in the case of images, patterns found in small 2D windows of the inputs In the previous example, these windows were all 3 × 3.

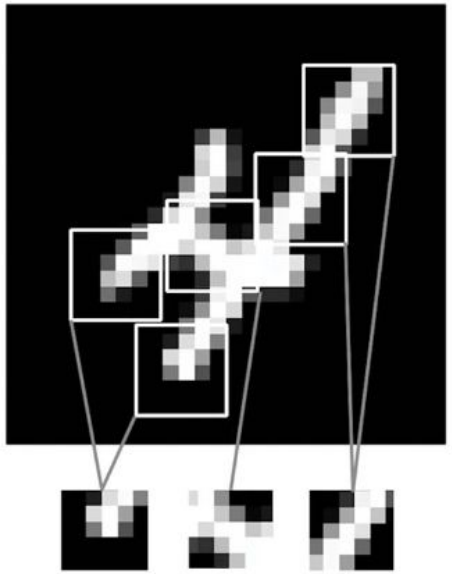

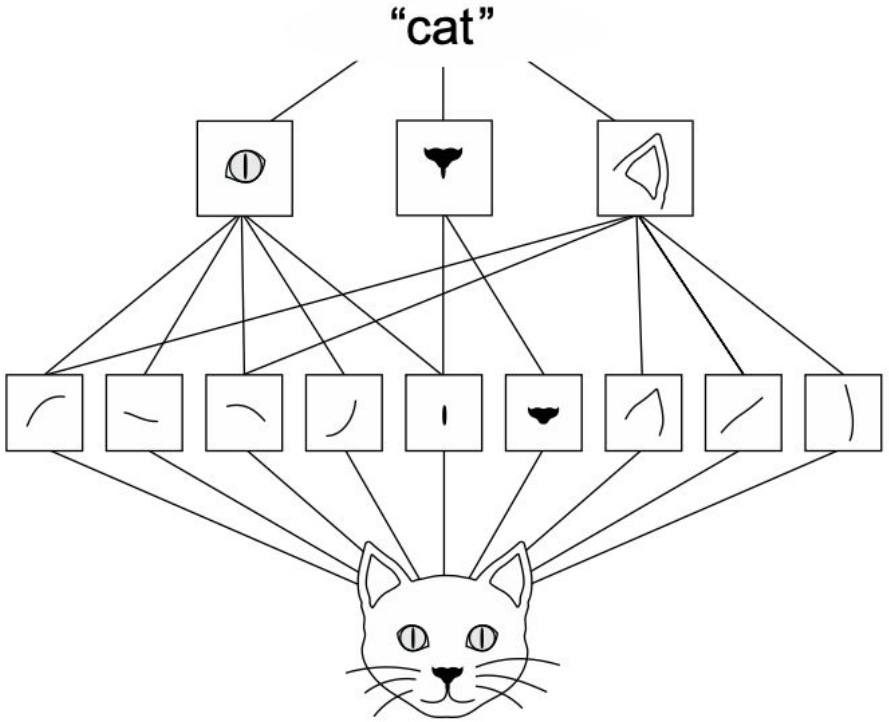

They can learn spatial hierarchies of patterns. A first convolution layer will learn small local patterns such as edges, a second convolution layer will learn larger patterns made of the features of the first layers, and so on.

"Deep Learning with Python" by François Chollet

Image Classification with CNN

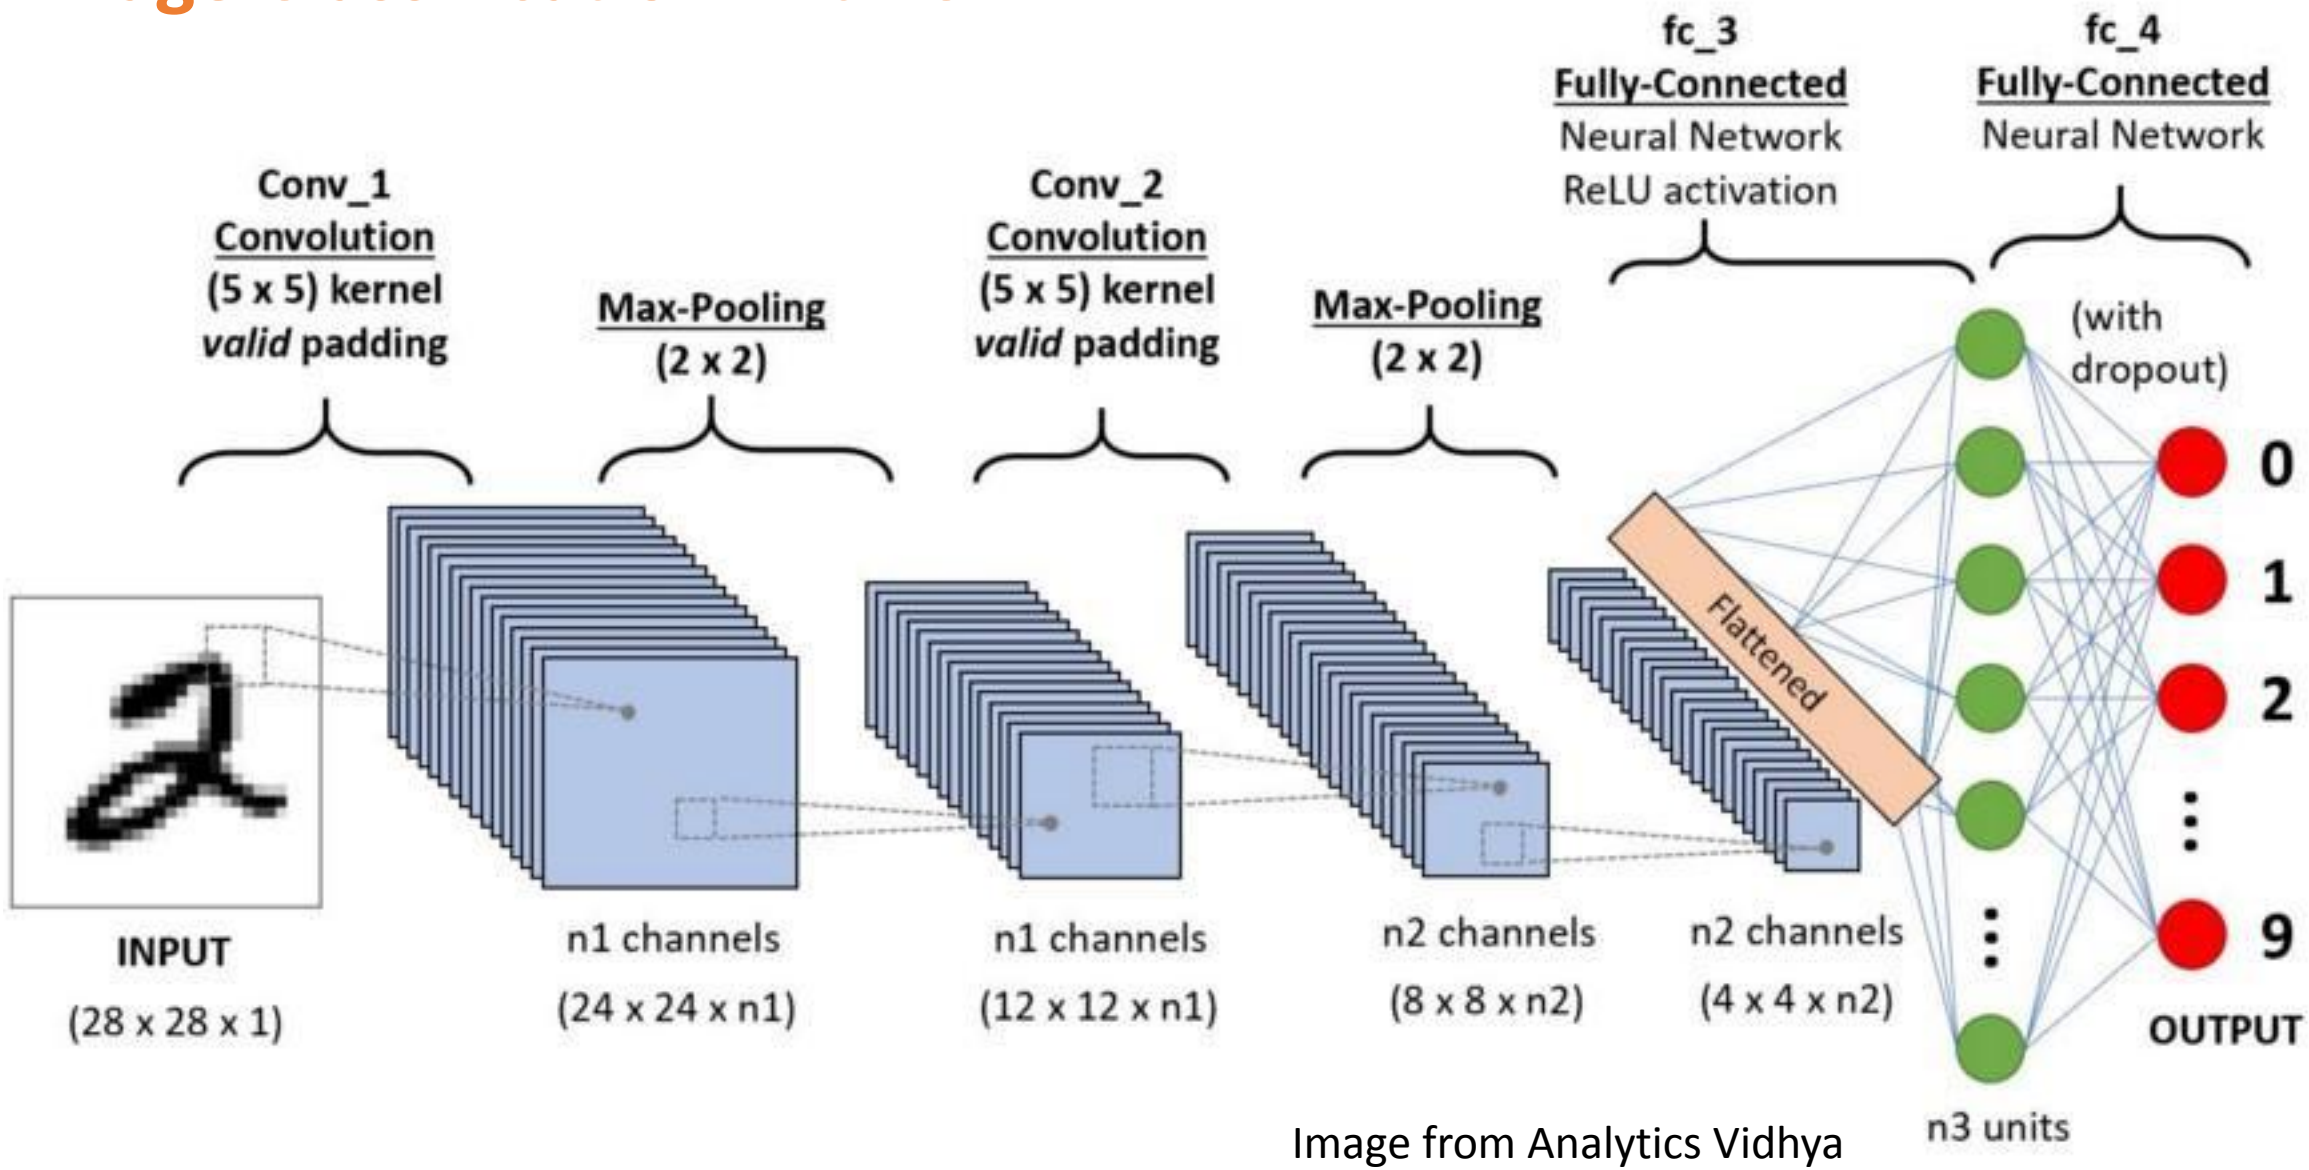

Image from Analytics Vidhya

Cifar-10

0 airplane
1 automobile
2 bird
3 cat
4 deer
5 dog
6 frog
7 horse
8 ship
9 truck

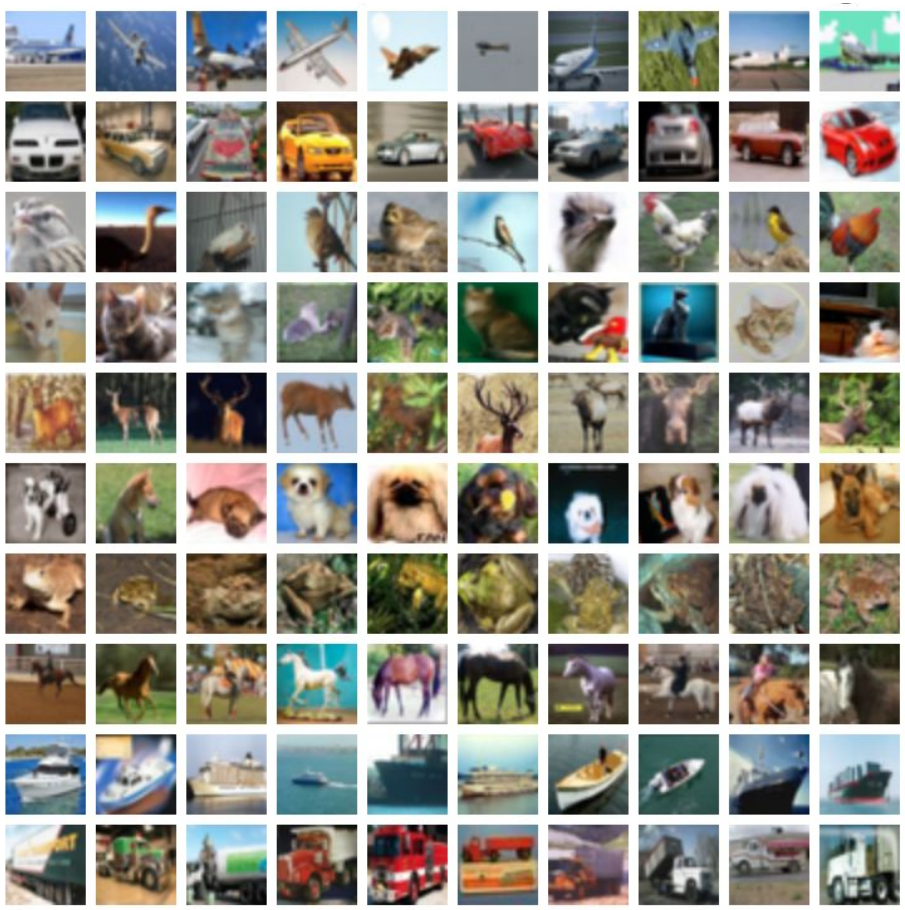

https://www.tensorflow.org/datasets/catalog/cifar10

So far …

We saw how to build Neural Networks (DNN and CNN) that classify images of digits (**MNIST**) or even fashion images (**Fashion-MNIST**).

Now,

We will instead, recognize the 10 classes of **CIFAR** ('airplane', 'automobile', 'bird', 'cat', 'deer', 'dog', 'frog', 'horse', 'ship' and 'truck').

There are some key differences between these image datasets that we need to take into account:

- While MNIST has 28x28 monochrome images (1 color channel), CIFAR is 32x32 color images (3 color channels).
- Besides, MNIST images are simple, containing just the object centered in the image, with no background. Conversely, CIFAR ones are not centered and can have the object with a background, such as airplanes that might have a cloudy sky behind them!

Those differences are the main reason to use a CNN instead of a DNN.

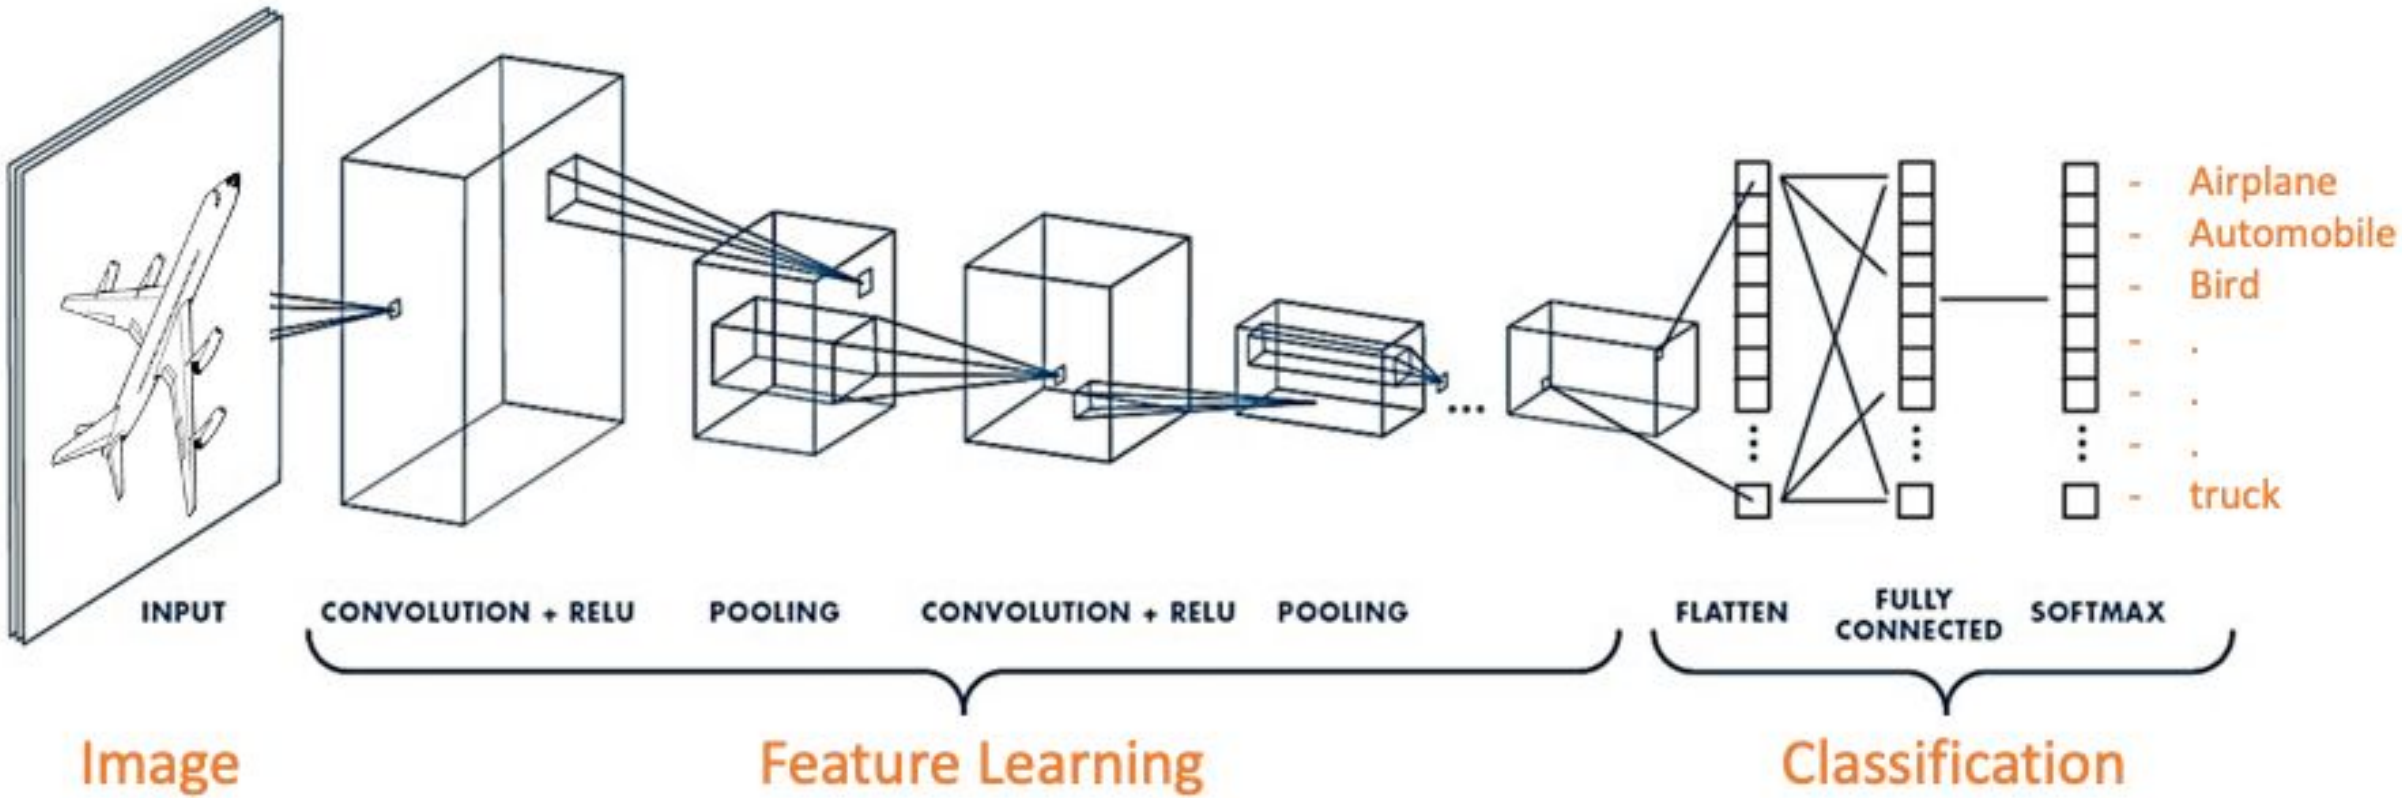
- Airplane
- Automobile
- Bird
- .
- .
- .
- truck
INPUT
CONVOLUTION + RELU
POOLING
CONVOLUTION + RELU
POOLING
FLATTEN
FULLY CONNECTED
SOFTMAX
Image
Feature Learning
Classification

Image Classification using CNN

Code Time!

CNN_Cifar-10.ipynb

Reading Material

Main references

- Harvard School of Engineering and Applied Sciences - CS249r: Tiny Machine Learning
- Professional Certificate in Tiny Machine Learning (TinyML) – edX/Harvard
- Introduction to Embedded Machine Learning - Coursera/Edge Impulse
- Computer Vision with Embedded Machine Learning - Coursera/Edge Impulse
- Fundamentals textbook: “Deep Learning with Python” by François Chollet
- Applications & Deploy textbook: “TinyML” by Pete Warden, Daniel Situnayake
- Deploy textbook “TinyML Cookbook” by Gian Marco Iodice

I want to thank Shawn Hymel and Edge Impulse, Pete Warden and Laurence Moroney from Google, Professor Vijay Janapa Reddi and Brian Plancher from Harvard, and the rest of the TinyMLedu team for preparing the excellent material on TinyML that is the basis of this course at UNIFEI.

The IESTI01 course is part of the **TinyML4D**, an initiative to make TinyML education available to everyone globally.

Thanks

UNIFEI