



DiSTRO

Deep Learning Course

Introduction to Convolutional Neural Networks

Alejandro Sztrajman

Vienna, Austria, November 2017



Funded by
the European Union

Index

- Linear Image Classifier
- Neural Network Representation
- Deep Neural Networks
- Convolutional Neural Networks
- CNN Example with Keras

Linear Image Classifier

Image Classification

$$f: \text{image} \rightarrow \text{category}$$

Image Classification


$$f\left(\text{img}\right) \rightarrow \text{category}$$


Image Classification


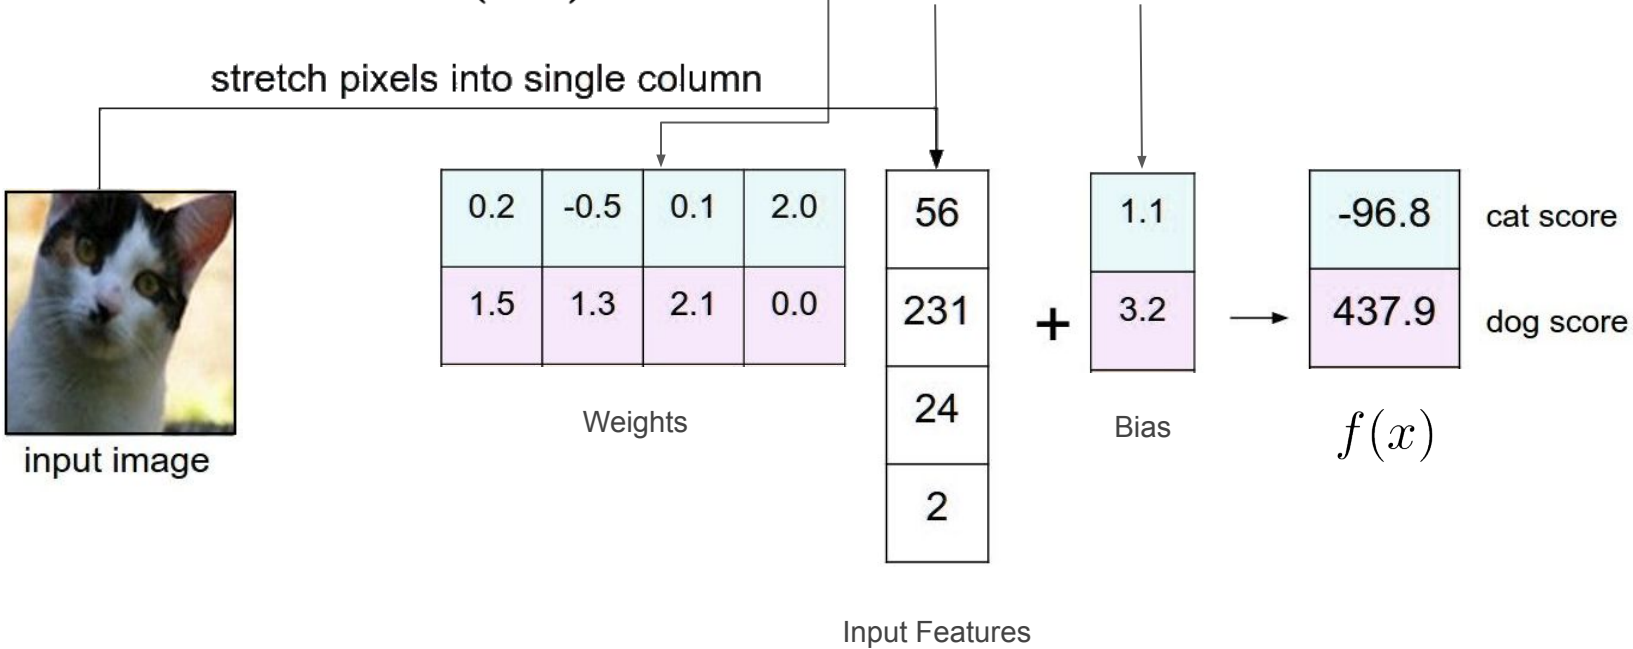
$$f\left(\text{img}\right) \rightarrow \text{cat}$$
A small square image of a black and white cat's face, positioned within the parentheses of the function notation in the equation above.

Image Classification

$$f(x) = Wx + b$$



Training

$$f(x_i) = Wx_i + b = y_i$$

Dataset of images and annotated categories

Parameters to optimise

Evaluation

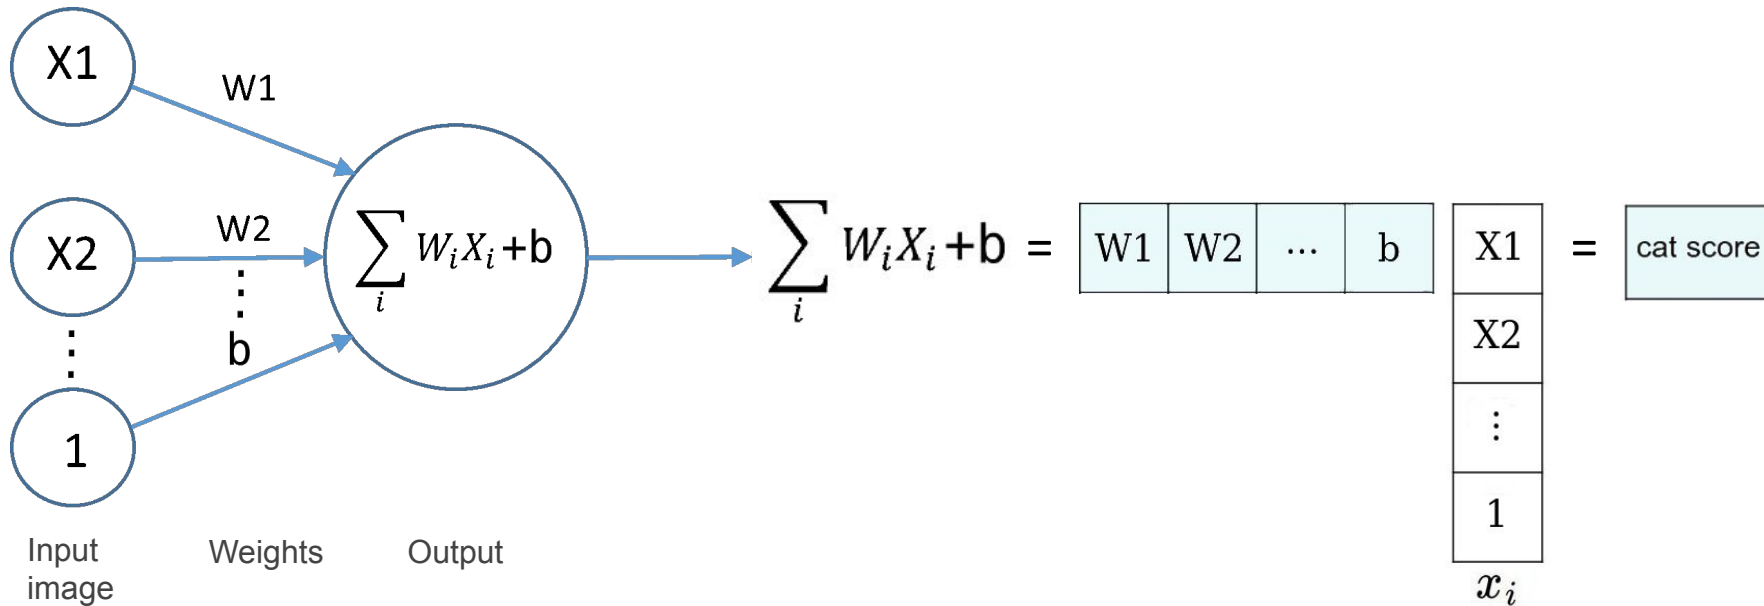
$$f(x_i) = Wx_i + b$$

Known


- All the weight of the computation is carried by the training. Evaluation is a simple matrix multiplication.
- Once we compute the parameters W, b we don't need the training data anymore.

Neural Network Representation

Neuron




Two neurons



$$Wx + b = \begin{array}{|c|c|c|c|} \hline W_{11} & W_{12} & \dots & b_1 \\ \hline W_{21} & W_{22} & \dots & b_2 \\ \hline \end{array} \begin{array}{|c|} \hline x_1 \\ \hline x_2 \\ \hline \vdots \\ \hline 1 \\ \hline \end{array} = \begin{array}{|c|} \hline \text{cat score} \\ \hline \text{dog score} \\ \hline \end{array}$$

x_i

Multiple neurons and layers



Input
layer

Hidden
Layer 1

Hidden
Layer 2

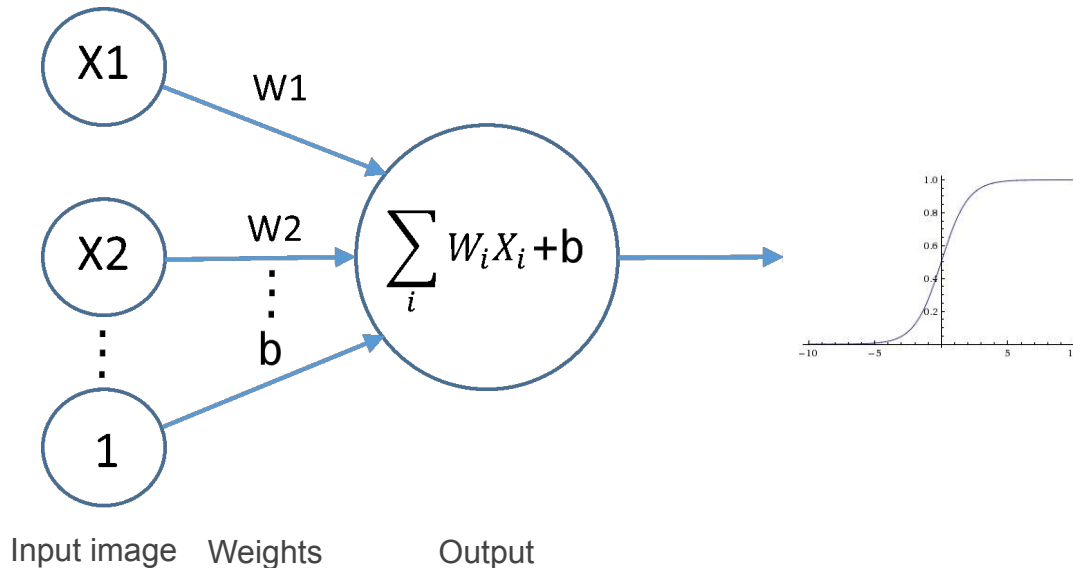
Output
Layer

- We can keep adding neurons in each layer and this means adding more classifiers.
- However adding more layers of neurons does not increase the complexity of the classifier, the model is still linear.

Activation Functions*

* Based on Slides by A. Peñate Sanchez.

Neuron with activation function




- We add a nonlinear operation applied to the output of the neuron.
- This increases the complexity of the classifier and allows the optimisation to favor certain features over others (automatic feature selection).

Activation functions

- Sigmoid


- Inspired by probability theory
- Single neuron corresponds exactly to the input-output mapping defined by logistic regression.
-

$$f(x) = \frac{1}{1 + \exp(-x)}$$




Activation functions

- Hyperbolic tangent
 - rescaled version of the sigmoid



Activation functions


- ReLU (Rectified Linear Unit)
 - $f(x) = \max(0, x)$
 - Out = 0 if (Input of i) < 0
 - Out = linear otherwise
 - **The most popular activation function for deep neural networks**



ReLU typically learns much faster than tanh, [Glorot et al. ICAIS 2011]


Connection Weights

- Effect of different weights in a neuron connection (sigmoid activation)




BIAS

- The weights can change the shape of the activation function, the BIAS shifts the function activation response.



- By adjusting the weights and the BIAS a single neuron can represent many different activation functions

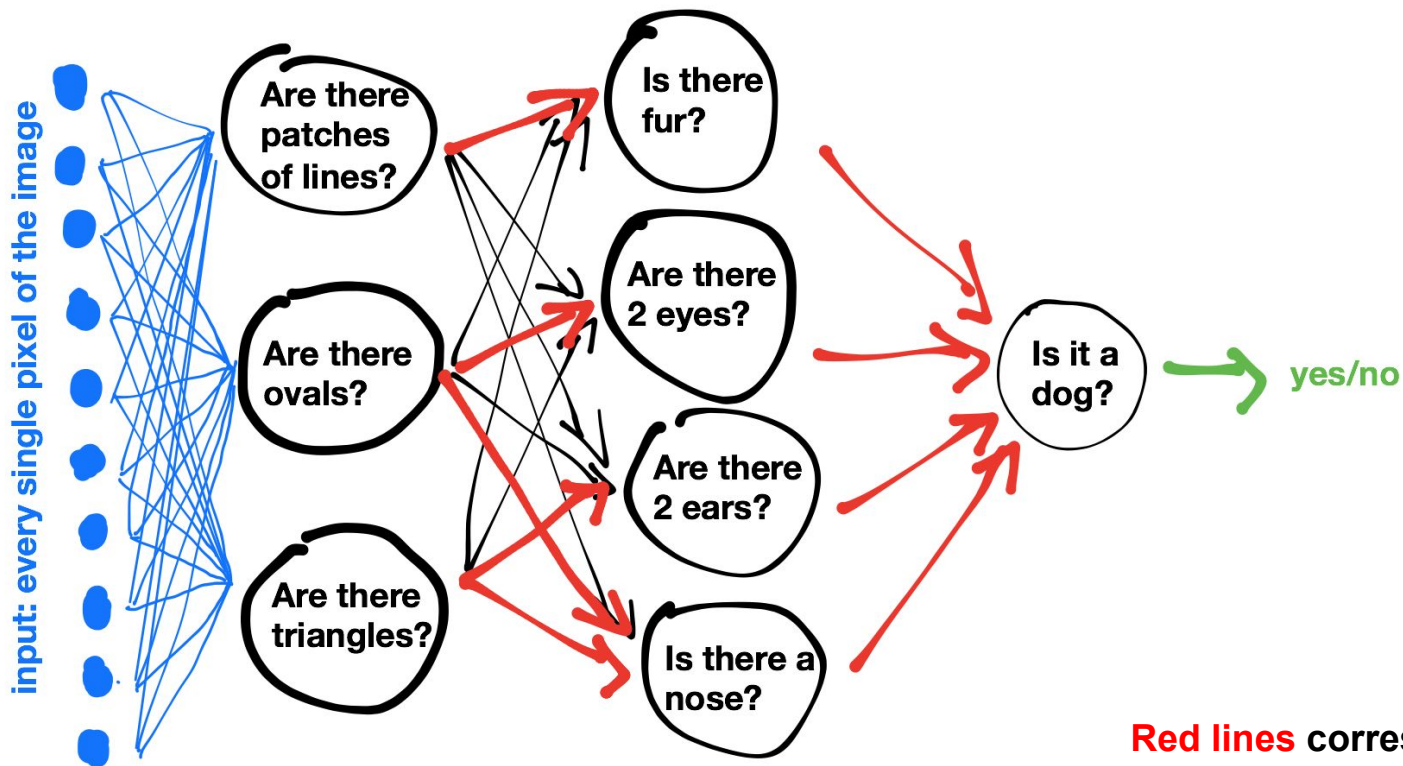
Fully Connected Deep Neural Networks



3 Hidden layers

- The final result is a composition of non linear functions applied to the features (pixels).
- The intermediate values that emerge as outputs of one layer of neurons are used as new input features (of increasing complexity) for the next layer.

Fully Connected Deep Neural Networks



Red lines correspond to high activation values.

Convolutional Neural Networks (CNNs)

Neural Networks on Images - Problem

So far, the design choice has been to **fully connect** all the hidden units to all the input units.


With images of size 96x96 this start to become *unfeasible*:
Learning features that span the entire image is very computationally expensive

Solution:

Locally Connected Networks

Restrict the connections between the hidden units and the input units

Convolutional Layer



- We leave the input images as 3D arrays (links channels of the same pixel).
- The convolutional layer is composed of a set of kernels which are convolved with the input image.
- Each kernel filters localised features in the input image. The size of the kernels determines the size of the features that can be learned.

Convolutional Layer

Operation: convolution

1	0	1
0	1	0
1	0	1

Convolutional Kernel

1 _{x1}	1 _{x0}	1 _{x1}	0	0
0 _{x0}	1 _{x1}	1 _{x0}	1	0
0 _{x1}	0 _{x0}	1 _{x1}	1	1
0	0	1	1	0
0	1	1	0	0

Image

Activation map


4		

Convolved
Feature

- The weights of the kernels are learned as part of the optimisation (feature selection).
- The output of the conv layer is a stack of 2D activation maps.

Convolutional Layer

Operation: convolution




- The weights of the kernels are learned as part of the optimisation (feature selection).
- The output of the conv layer is a stack of 2D activation maps.

Pooling Layer

Goal: progressively reduce the spatial size of the representation to reduce the amount of parameters and computation in the network, and hence to also control overfitting.


- Max pooling
- Min pooling
- Average pooling



Pooling Layer

Goal: progressively reduce the spatial size of the representation to reduce the amount of parameters and computation in the network, and hence to also control overfitting.

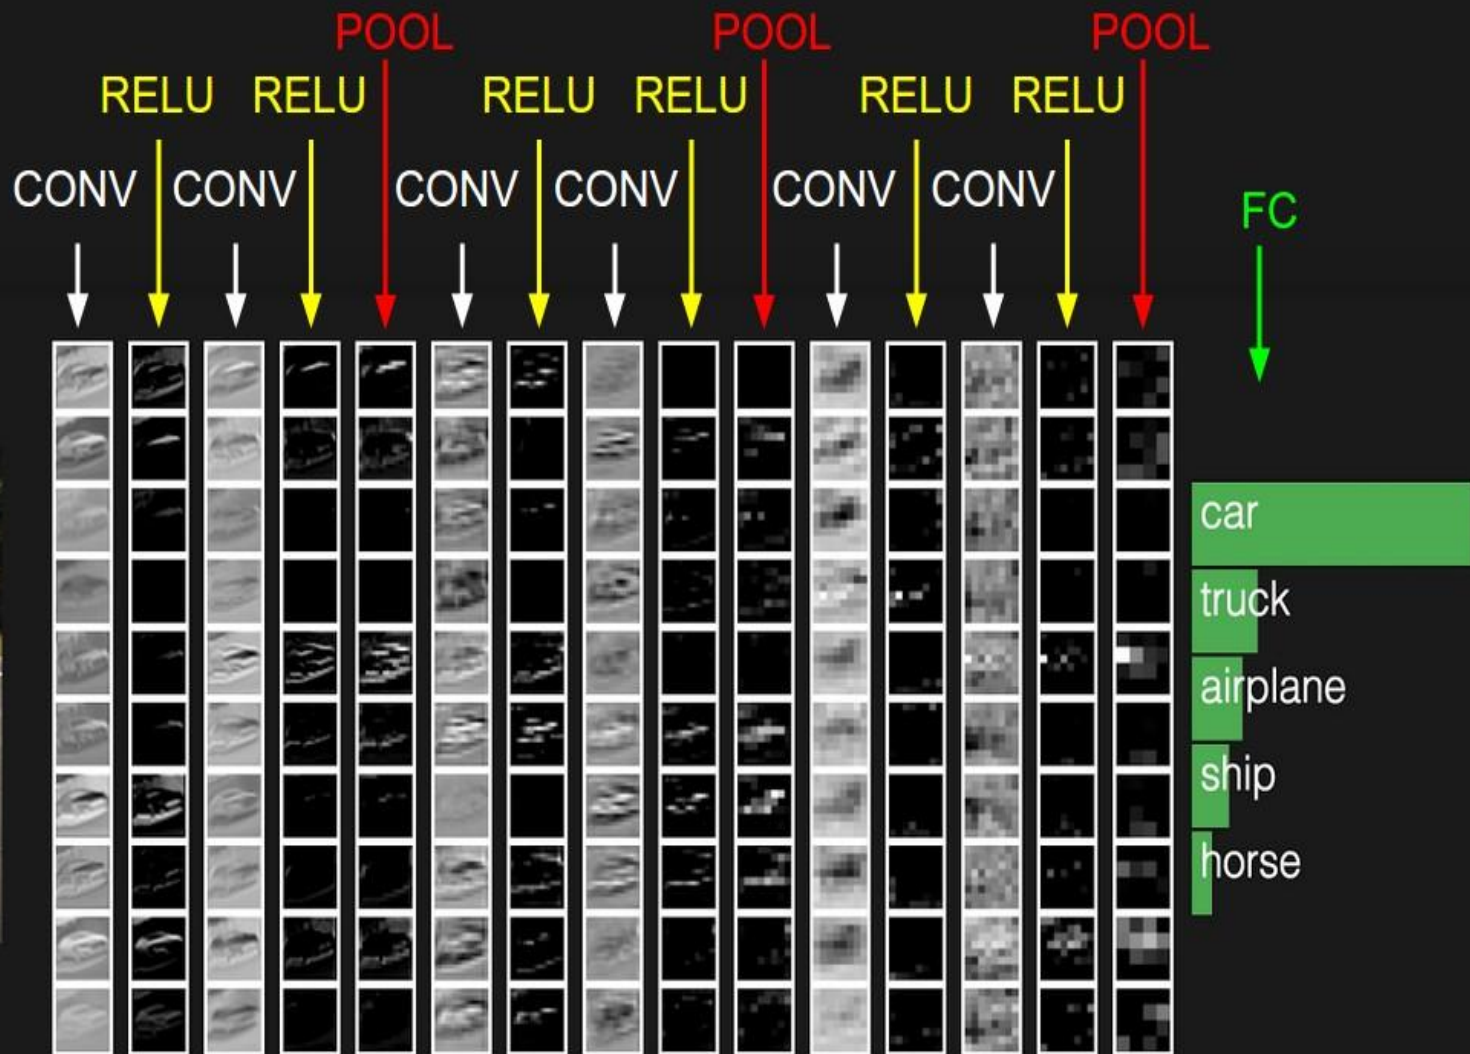
- Max pooling
- Min pooling
- Average pooling



Pooling Layer

If the pooled regions are contiguous areas in the input image, then the pooling units will be **translation invariant**, something necessary:

we want the classifier to still accurately classify objects regardless of their relative position in the image frame.



Dropout Layer

Technique used only in the training phase, addressing the **overfitting problem**.

Idea:


- Randomly drop units during training

Drop out means **temporarily** removing the unit from the network along with all its incoming and outgoing connections


Dropout Layer

Example on fully connected layers

Each unit is retained with a fixed probability p **independent** of other units



(a) Standard Neural Net



(b) After applying dropout.

In addition to prevent overfitting, it also provides a way of approximately **combining** exponentially many **different neural network architectures** efficiently.

A bit of code:
simple example with Keras

