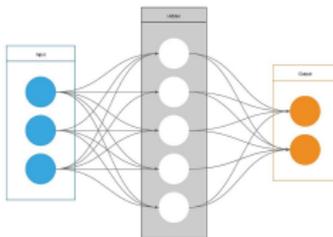


# Data Analysis, Neural Networks and the use of Keras

Cecilia Jarne

cecilia.jarne@unq.edu.ar



Twitter: @ceciliajarne



# What is this talk about?

- What is Machine Learning, Neural Networks, and Deep Learning?
- Neural Networks concepts and topologies.
- A brief review of the use of python ML libraries.

# What can we do with data?



We are interested in:

- Data Visualization.
- Data Analysis (Identify features from the data).
- Data Classification.
- Implementation of different algorithms as intelligent as we can get.

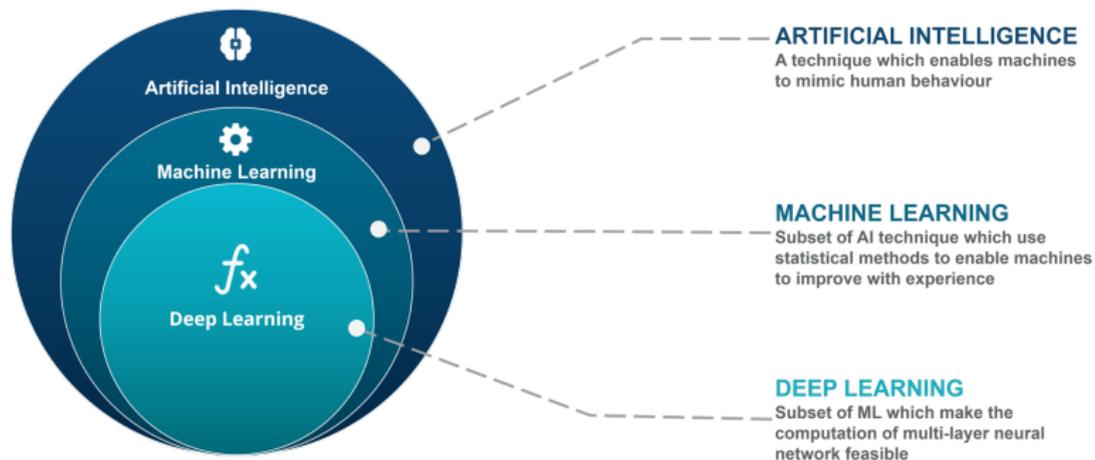
# Which are the available algorithms?

	<u>TYPE</u>	<u>NAME</u>	<u>DESCRIPTION</u>	<u>ADVANTAGES</u>	<u>DISADVANTAGES</u>
Linear		 <p>Linear regression</p>	The "best fit" line through all data points. Predictions are numerical.	Easy to understand – you clearly see what the biggest drivers of the model are.	<ul style="list-style-type: none"> <li>X Sometimes too simple to capture complex relationships between variables.</li> <li>X Tendency for the model to "overfit".</li> </ul>
		 <p>Logistic regression</p>	The adaptation of linear regression to problems of classification (e.g., yes/no questions, groups, etc.)	Also easy to understand.	<ul style="list-style-type: none"> <li>X Sometimes too simple to capture complex relationships between variables.</li> <li>X Tendency for the model to "overfit".</li> </ul>
Tree-based		 <p>Decision tree</p>	A graph that uses a branching method to match all possible outcomes of a decision.	Easy to understand and implement.	<ul style="list-style-type: none"> <li>X Not often used on its own for prediction because it's also often too simple and not powerful enough for complex data.</li> </ul>
		 <p>Random Forest</p>	Takes the average of many decision trees, each of which is made with a sample of the data. Each tree is weaker than a full decision tree, but by combining them we get better overall performance.	A sort of "wisdom of the crowd". Tends to result in very high quality models. Fast to train.	<ul style="list-style-type: none"> <li>X Can be slow to output predictions relative to other algorithms.</li> <li>X Not easy to understand predictions.</li> </ul>
		 <p>Gradient Boosting</p>	Uses even weaker decision trees, that are increasingly focused on "hard" examples.	High-performing.	<ul style="list-style-type: none"> <li>X A small change in the feature set or training set can create radical changes in the model.</li> <li>X Not easy to understand predictions.</li> </ul>
Neural networks		 <p>Neural networks</p>	Mimics the behavior of the brain. Neural networks are interconnected neurons that pass messages to each other. Deep learning uses several layers of neural networks put one after the other.	Can handle extremely complex tasks - no other algorithm comes close in image recognition.	<ul style="list-style-type: none"> <li>X Very, very slow to train, because they have so many layers. Require a lot of power.</li> <li>X Almost impossible to understand predictions.</li> </ul>

# Which are the available algorithms?

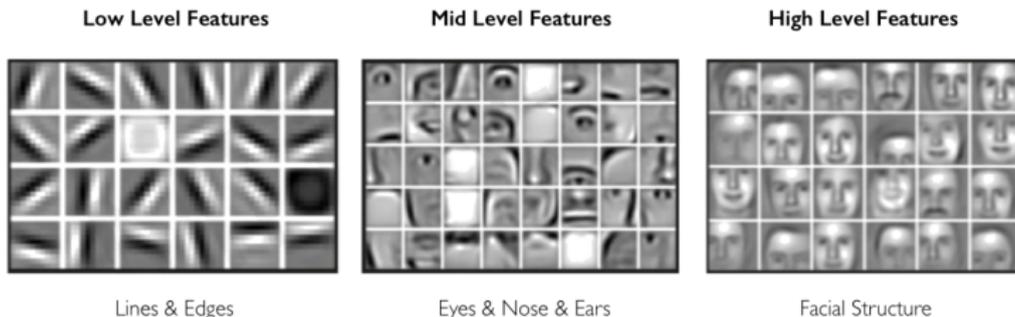
	<u>TYPE</u>	<u>NAME</u>	<u>DESCRIPTION</u>	<u>ADVANTAGES</u>	<u>DISADVANTAGES</u>
Linear		Linear regression	The "best fit" line through all data points. Predictions are numerical.	Easy to understand – you clearly see what the biggest drivers of the model are.	<ul style="list-style-type: none"> <li>X Sometimes too simple to capture complex relationships between variables.</li> <li>X Tendency for the model to "overfit".</li> </ul>
		Logistic regression	The adaptation of linear regression to problems of classification (e.g., yes/no questions, groups, etc.)	Also easy to understand.	<ul style="list-style-type: none"> <li>X Sometimes too simple to capture complex relationships between variables.</li> <li>X Tendency for the model to "overfit".</li> </ul>
Tree-based		Decision tree	A graph that uses a branching method to match all possible outcomes of a decision.	Easy to understand and implement.	<ul style="list-style-type: none"> <li>X Not often used on its own for prediction because it's also often too simple and not powerful enough for complex data.</li> </ul>
		Random Forest	Takes the average of many decision trees, each of which is made with a sample of the data. Each tree is weaker than a full decision tree, but by combining them we get better overall performance.	A sort of "wisdom of the crowd". Tends to result in very high quality models. Fast to train.	<ul style="list-style-type: none"> <li>X Can be slow to output predictions relative to other algorithms.</li> <li>X Not easy to understand predictions.</li> </ul>
		Gradient Boosting	Uses even weaker decision trees, that are increasingly focused on "hard" examples.	High-performing.	<ul style="list-style-type: none"> <li>X A small change in the feature set or training set can create radical changes in the model.</li> <li>X Not easy to understand predictions.</li> </ul>
Neural networks		Neural networks	Mimics the behavior of the brain. Neural networks are interconnected neurons that pass messages to each other. Deep learning uses several layers of neural networks put one after the other.	Can handle extremely complex tasks - no other algorithm comes close in image recognition.	<ul style="list-style-type: none"> <li>X Very, very slow to train, because they have so many layers. Require a lot of power</li> <li>X Almost impossible to understand predictions.</li> </ul>

# What about machine learning?

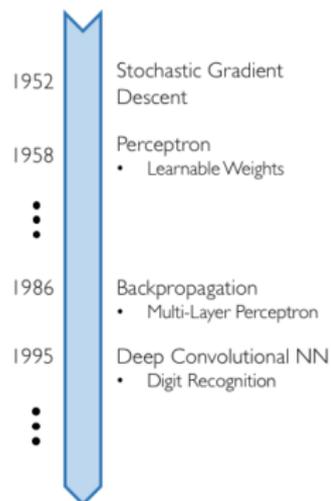


# Why Deep Learning?

- Data driven approach.
- Can we learn the underlying features directly from data?



# Why Now?



Neural Networks date back decades, so why the resurgence?

## 1. Big Data

- Larger Datasets
- Easier Collection & Storage

IMAGENET



## 2. Hardware

- Graphics Processing Units (GPUs)
- Massively Parallelizable

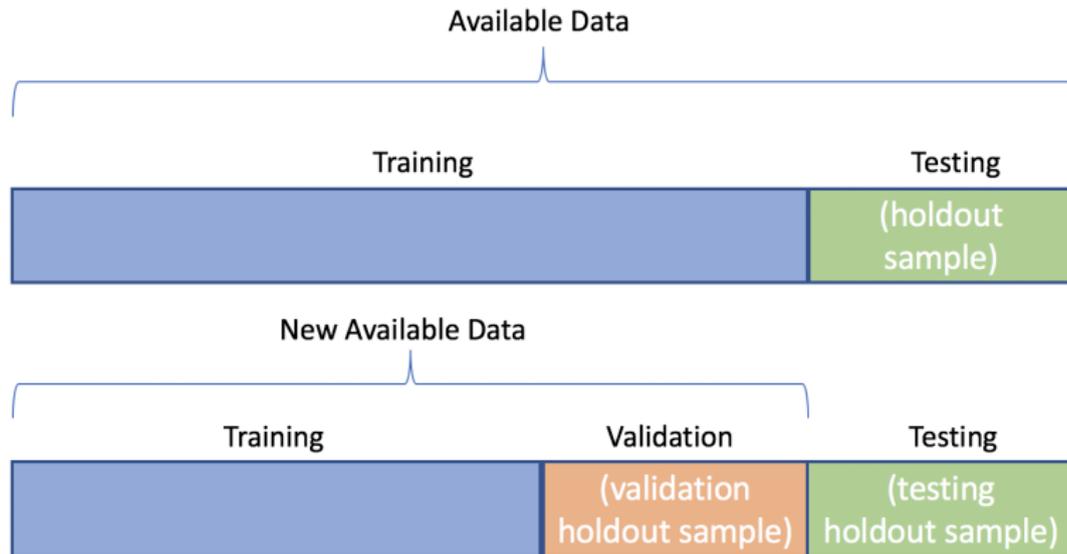


## 3. Software

- Improved Techniques
- New Models
- Toolboxes



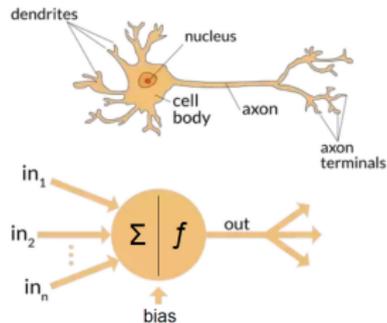
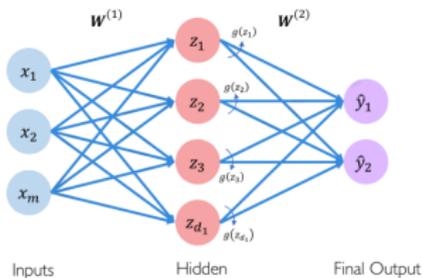
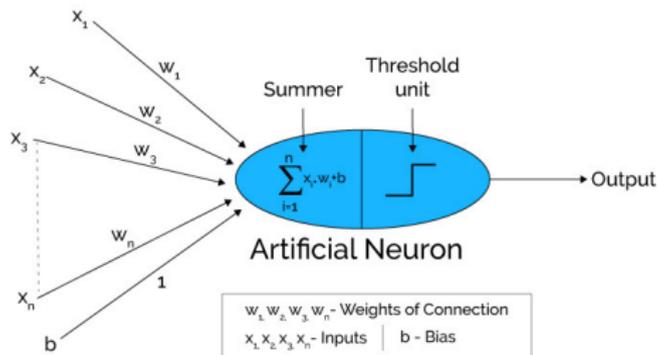
# How we split our data to train a model?



# How we use (split) our data to train a model?

- **Training set:** The data where the model is trained on. We train our model, by pairing the input with the expected output.
- **Validation set:** Data the model has not been trained on and is used to tune hyperparameters. Here we estimate how well your model has been trained.
- **Test set:** Same like the validation set.. just used at the final end. This is an Application phase: we apply our developed model to the real-world data and get the results. This fase is split into two parts:
  - First you look at your models and select the best performing approach using the validation data (=validation).
  - Then you estimate the accuracy of the selected approach (=test).

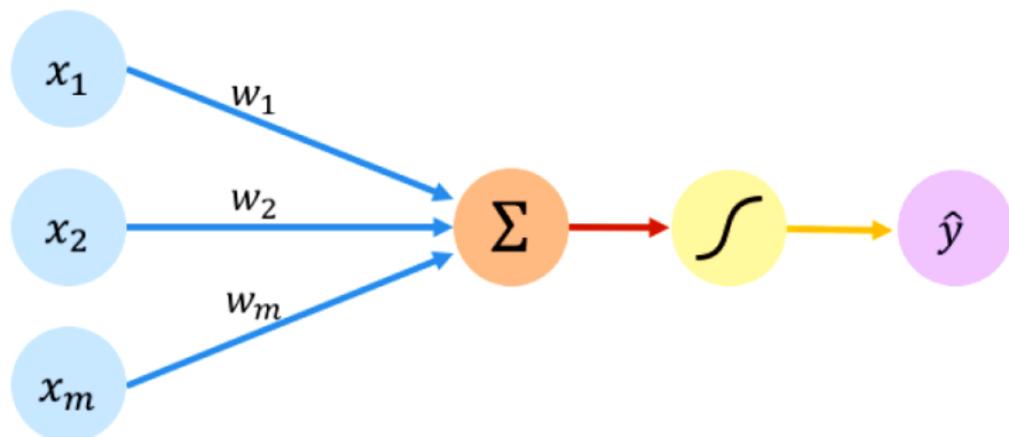
# Neural Networks:



# The Perceptron: Forward Propagation

The structural building block of Deep Learning

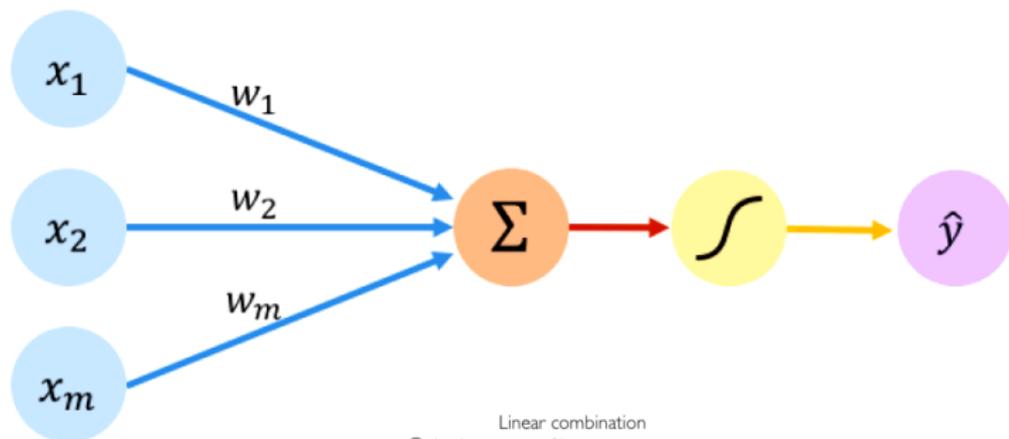
Input  $\rightarrow$  Weights  $\rightarrow$  Sum  $\rightarrow$  Non linearity  $\rightarrow$  Output



# The Perceptron: Forward Propagation

The structural building block of Deep Learning:

Input  $\rightarrow$  Weights  $\rightarrow$  Sum  $\rightarrow$  Non linearity  $\rightarrow$  Output



Linear combination of inputs

$$\hat{y} = g \left( \sum_{i=1}^m x_i w_i \right)$$

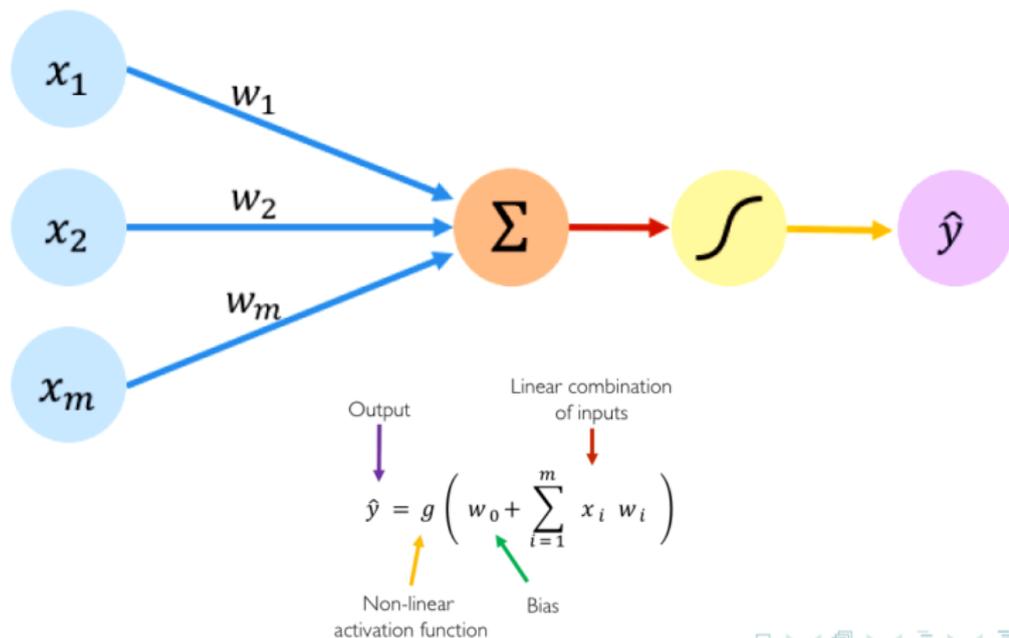
Non-linear activation function

Output

# The Perceptron: Forward Propagation

The structural building block of Deep Learning:

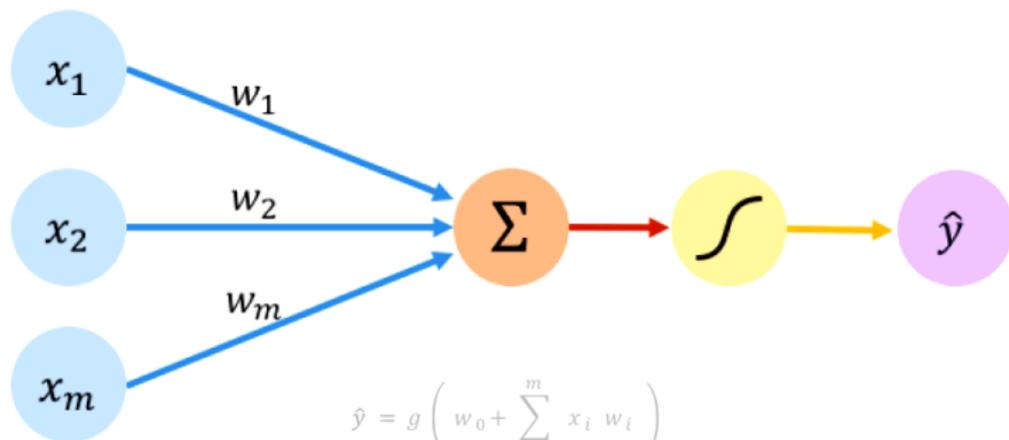
Input  $\rightarrow$  Weights  $\rightarrow$  Sum  $\rightarrow$  Non linearity  $\rightarrow$  Output



# The Perceptron: Forward Propagation

The structural building block of Deep Learning:

Input  $\rightarrow$  Weights  $\rightarrow$  Sum  $\rightarrow$  Non linearity  $\rightarrow$  Output

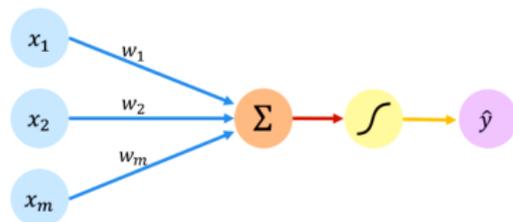


$$\hat{y} = g \left( w_0 + \sum_{i=1}^m x_i w_i \right)$$

$$\hat{y} = g ( w_0 + \mathbf{X}^T \mathbf{W} )$$

$$\text{where: } \mathbf{X} = \begin{bmatrix} x_1 \\ \vdots \\ x_m \end{bmatrix} \text{ and } \mathbf{W} = \begin{bmatrix} w_1 \\ \vdots \\ w_m \end{bmatrix}$$

# The Perceptron: Forward Propagation

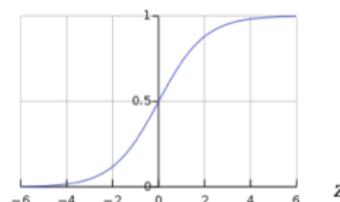


## Activation Functions

$$\hat{y} = g(w_0 + \mathbf{X}^T \mathbf{W})$$

- Example: sigmoid function

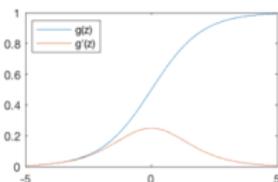
$$g(z) = \sigma(z) = \frac{1}{1 + e^{-z}}$$



# Neural Networks: Activation Functions

The purpose of Activation functions is to introduce non linearities into the network

Sigmoid Function

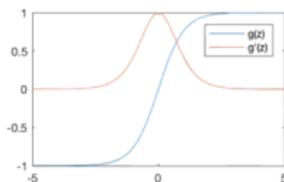


$$g(z) = \frac{1}{1 + e^{-z}}$$

$$g'(z) = g(z)(1 - g(z))$$

 `tf.math.sigmoid(z)`

Hyperbolic Tangent

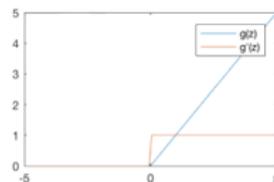


$$g(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}}$$

$$g'(z) = 1 - g(z)^2$$

 `tf.math.tanh(z)`

Rectified Linear Unit (ReLU)



$$g(z) = \max(0, z)$$

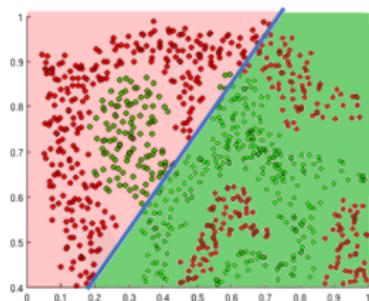
$$g'(z) = \begin{cases} 1, & z > 0 \\ 0, & \text{otherwise} \end{cases}$$

 `tf.nn.relu(z)`

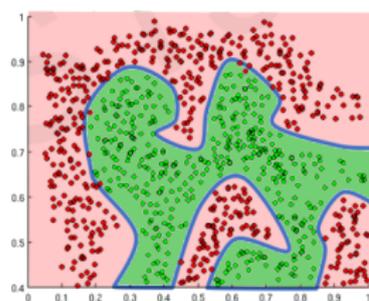
# Neural Networks: Activation Functions

Name	Plot	Equation	Derivative
Identity		$f(x) = x$	$f'(x) = 1$
Binary step		$f(x) = \begin{cases} 0 & \text{for } x < 0 \\ 1 & \text{for } x \geq 0 \end{cases}$	$f'(x) = \begin{cases} 0 & \text{for } x \neq 0 \\ ? & \text{for } x = 0 \end{cases}$
Logistic (a.k.a. Soft step)		$f(x) = \frac{1}{1 + e^{-x}}$	$f'(x) = f(x)(1 - f(x))$
Tanh		$f(x) = \tanh(x) = \frac{2}{1 + e^{-2x}} - 1$	$f'(x) = 1 - f(x)^2$
ArcTan		$f(x) = \tan^{-1}(x)$	$f'(x) = \frac{1}{x^2 + 1}$
Rectified Linear Unit (ReLU)		$f(x) = \begin{cases} 0 & \text{for } x < 0 \\ x & \text{for } x \geq 0 \end{cases}$	$f'(x) = \begin{cases} 0 & \text{for } x < 0 \\ 1 & \text{for } x \geq 0 \end{cases}$
Parametric Rectified Linear Unit (PReLU) <sup>[2]</sup>		$f(x) = \begin{cases} \alpha x & \text{for } x < 0 \\ x & \text{for } x \geq 0 \end{cases}$	$f'(x) = \begin{cases} \alpha & \text{for } x < 0 \\ 1 & \text{for } x \geq 0 \end{cases}$
Exponential Linear Unit (ELU) <sup>[3]</sup>		$f(x) = \begin{cases} \alpha(e^x - 1) & \text{for } x < 0 \\ x & \text{for } x \geq 0 \end{cases}$	$f'(x) = \begin{cases} f(x) + \alpha & \text{for } x < 0 \\ 1 & \text{for } x \geq 0 \end{cases}$
SoftPlus		$f(x) = \log_e(1 + e^x)$	$f'(x) = \frac{1}{1 + e^{-x}}$

# Neural Networks: Non linear decision



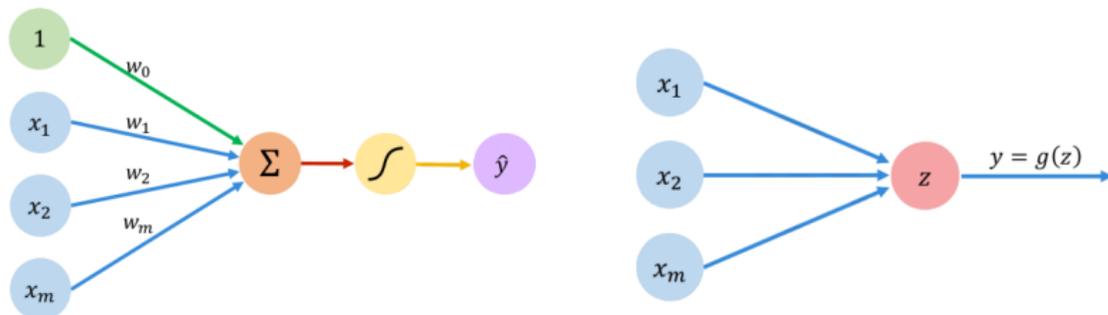
Linear activation functions produce linear decisions no matter the network size



Non-linearities allow us to approximate arbitrarily complex functions

# Building Neural Networks with Perceptrons

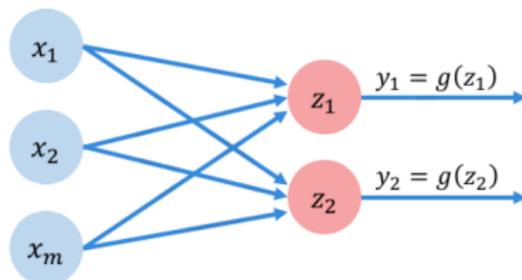
A simplified perceptron:



$$z = w_0 + \sum_{j=1}^m x_j w_j$$

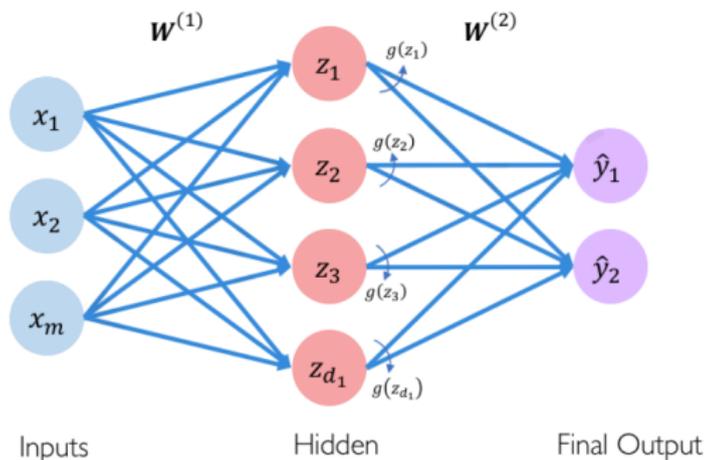
# Building Neural Networks: Multi Outputs

Because all inputs are connected to all outputs these are called Dense layers:



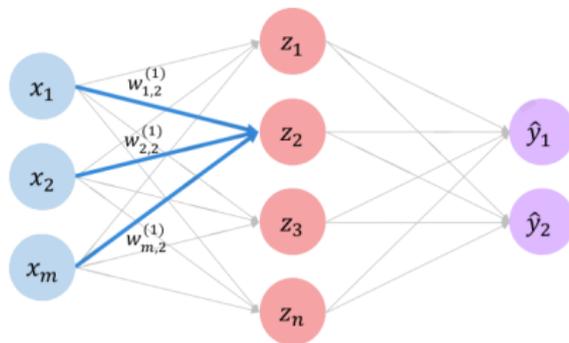
$$z_i = w_{0,i} + \sum_{j=1}^m x_j w_{j,i}$$

# Single Layer Neural Network



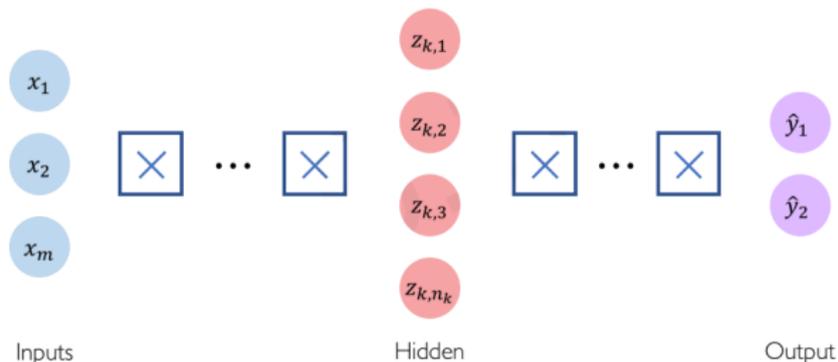
$$z_i = w_{0,i}^{(1)} + \sum_{j=1}^m x_j w_{j,i}^{(1)} \quad \hat{y}_i = g \left( w_{0,i}^{(2)} + \sum_{j=1}^{d_1} g(z_j) w_{j,i}^{(2)} \right)$$

# Single Layer Neural Network



$$\begin{aligned} z_2 &= w_{0,2}^{(1)} + \sum_{j=1}^m x_j w_{j,2}^{(1)} \\ &= w_{0,2}^{(1)} + x_1 w_{1,2}^{(1)} + x_2 w_{2,2}^{(1)} + x_m w_{m,2}^{(1)} \end{aligned}$$

# A deep neural network structure



$$z_{k,i} = w_{0,i}^{(k)} + \sum_{j=1}^{n_{k-1}} g(z_{k-1,j}) w_{j,i}^{(k)}$$

(... or Cost Function or Objective Function)

Depends on the kind of problem.

- Regression → Mean square error
- Classification → Cross entropy, binary cross entropy.

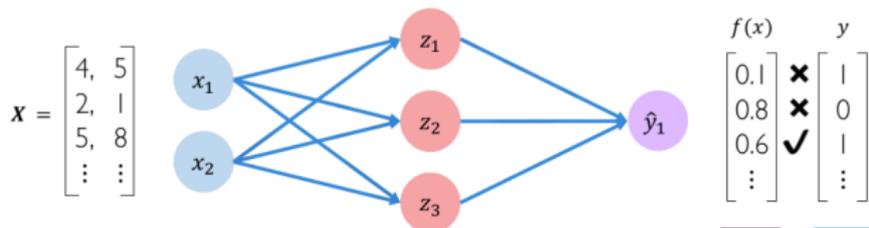
# Neural Networks: Loss Functions

symbol	name	equation
$\mathcal{L}_1$	L <sub>1</sub> loss	$\ \mathbf{y} - \mathbf{o}\ _1$
$\mathcal{L}_2$	L <sub>2</sub> loss	$\ \mathbf{y} - \mathbf{o}\ _2^2$
$\mathcal{L}_1 \circ \sigma$	expectation loss	$\ \mathbf{y} - \sigma(\mathbf{o})\ _1$
$\mathcal{L}_2 \circ \sigma$	regularised expectation loss <sup>1</sup>	$\ \mathbf{y} - \sigma(\mathbf{o})\ _2^2$
$\mathcal{L}_\infty \circ \sigma$	Chebyshev loss	$\max_j  \sigma(\mathbf{o})^{(j)} - \mathbf{y}^{(j)} $
hinge	hinge [13] (margin) loss	$\sum_j \max(0, \frac{1}{2} - \hat{\mathbf{y}}^{(j)} \mathbf{o}^{(j)})$
hinge <sup>2</sup>	squared hinge (margin) loss	$\sum_j \max(0, \frac{1}{2} - \hat{\mathbf{y}}^{(j)} \mathbf{o}^{(j)})^2$
hinge <sup>3</sup>	cubed hinge (margin) loss	$\sum_j \max(0, \frac{1}{2} - \hat{\mathbf{y}}^{(j)} \mathbf{o}^{(j)})^3$
log	log (cross entropy) loss	$-\sum_j \mathbf{y}^{(j)} \log \sigma(\mathbf{o})^{(j)}$
log <sup>2</sup>	squared log loss	$-\sum_j [\mathbf{y}^{(j)} \log \sigma(\mathbf{o})^{(j)}]^2$
tan	Tanimoto loss	$\frac{-\sum_j \sigma(\mathbf{o})^{(j)} \mathbf{y}^{(j)}}{\ \sigma(\mathbf{o})\ _2^2 + \ \mathbf{y}\ _2^2 - \sum_j \sigma(\mathbf{o})^{(j)} \mathbf{y}^{(j)}}$
D <sub>CS</sub>	Cauchy-Schwarz Divergence [3]	$-\log \frac{\sum_j \sigma(\mathbf{o})^{(j)} \mathbf{y}^{(j)}}{\ \sigma(\mathbf{o})\ _2 \ \mathbf{y}\ _2}$

See: <https://arxiv.org/pdf/1702.05659.pdf>s

# Neural Networks: Crossentropy

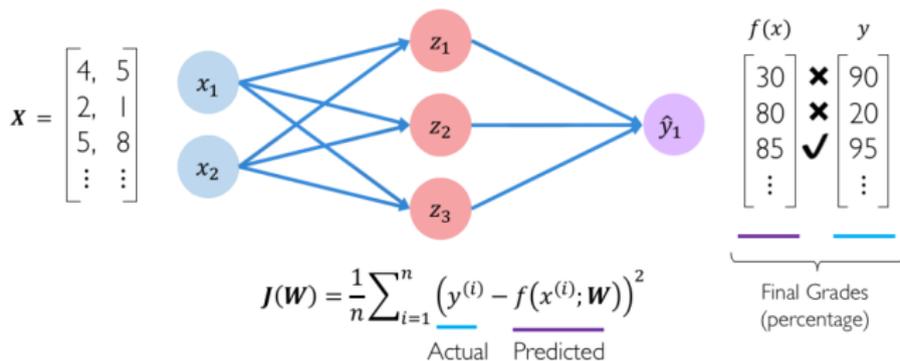
*Cross entropy loss* can be used with models that output a probability between 0 and 1



$$J(\mathbf{W}) = \frac{1}{n} \sum_{i=1}^n \underbrace{y^{(i)}}_{\text{Actual}} \log \left( \underbrace{f(x^{(i)}; \mathbf{W})}_{\text{Predicted}} \right) + (1 - \underbrace{y^{(i)}}_{\text{Actual}}) \log \left( 1 - \underbrace{f(x^{(i)}; \mathbf{W})}_{\text{Predicted}} \right)$$

# Neural Networks: Mean square error

*Mean squared error loss* can be used with regression models that output continuous real numbers



# Training Neural Networks: Loss Optimization

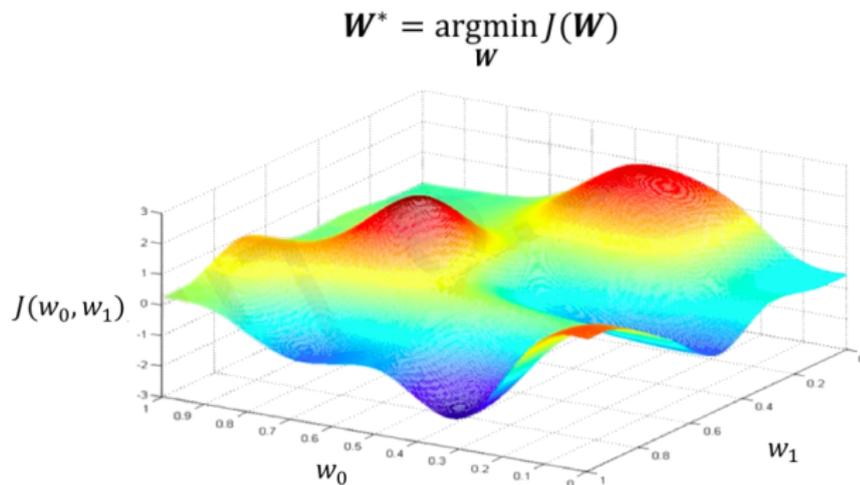
We want to find the network weights that achieve the lowest loss

$$\mathbf{W}^* = \operatorname{argmin}_{\mathbf{W}} \frac{1}{n} \sum_{i=1}^n \mathcal{L}(f(x^{(i)}; \mathbf{W}), y^{(i)})$$

$$\mathbf{W}^* = \operatorname{argmin}_{\mathbf{W}} J(\mathbf{W})$$

# Training Neural Networks: Loss Optimization

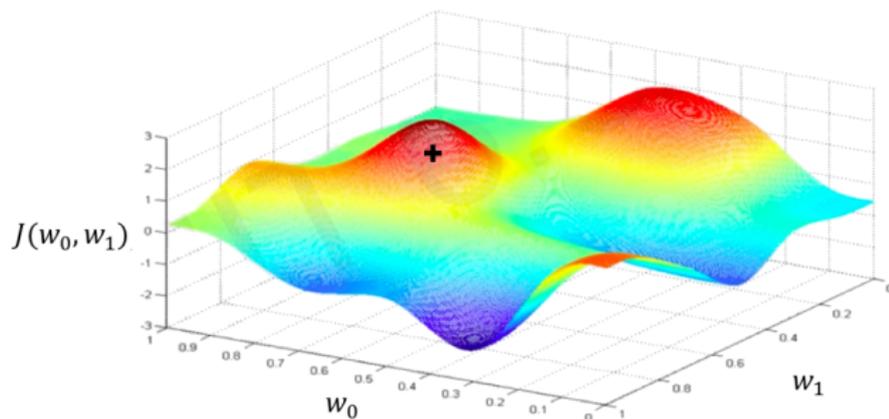
We want to find the network weights that achieve the lowest loss



# Training Neural Networks: Loss Optimization

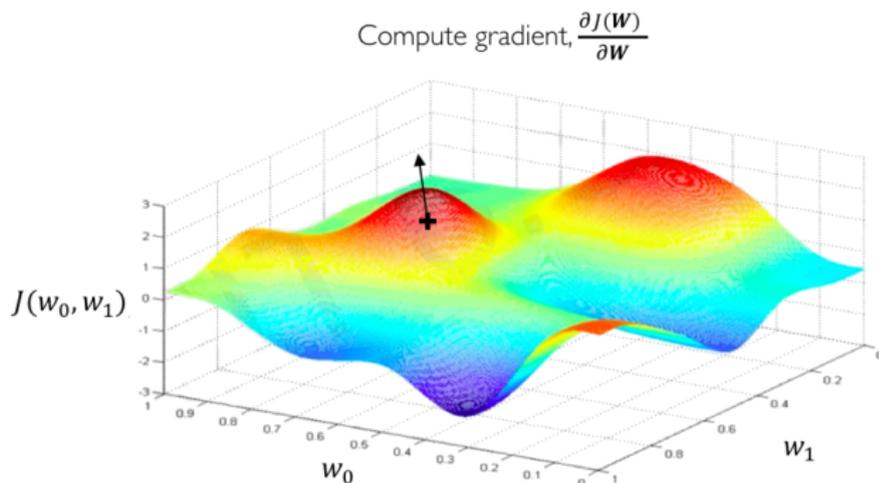
We want to find the network weights that achieve the lowest loss

Randomly pick an initial  $(w_0, w_1)$



# Training Neural Networks: Loss Optimization

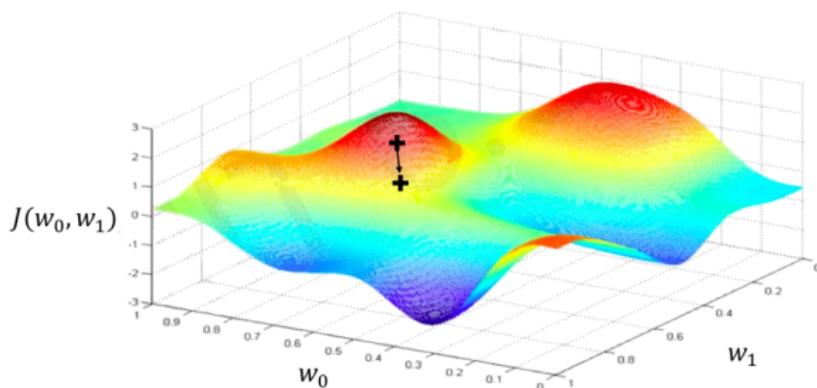
We want to find the network weights that achieve the lowest loss



# Training Neural Networks: Loss Optimization

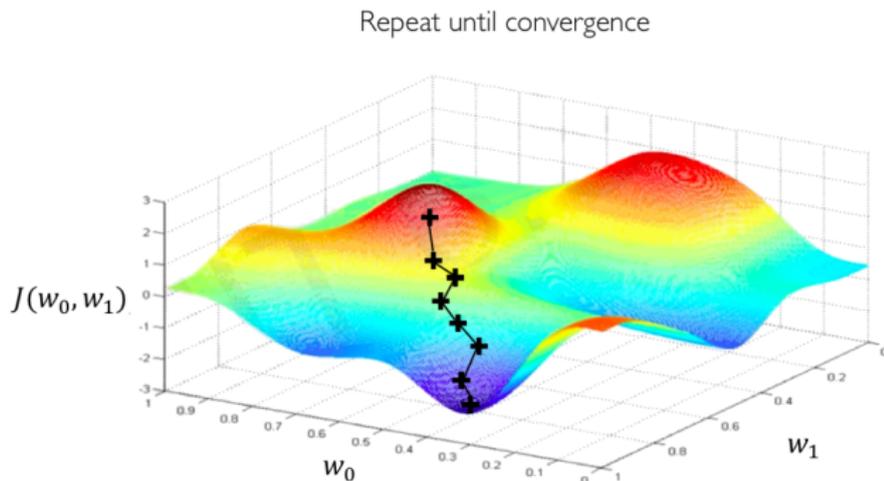
We want to find the network weights that achieve the lowest loss

Take small step in opposite direction of gradient



# Training Neural Networks: Loss Optimization

We want to find the network weights that achieve the lowest loss



## Algorithm

1. Initialize weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence:
3.     Compute gradient,  $\frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
4.     Update weights,  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
5. Return weights

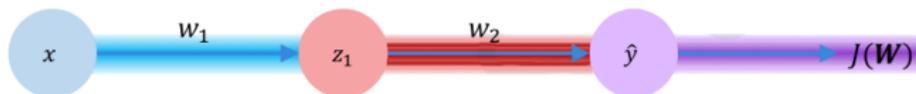
# Training Neural Networks: Back propagation

How a small change in one weight affect the loss



# Training Neural Networks:

We apply chain rule



$$\frac{\partial J(\mathbf{W})}{\partial w_1} = \frac{\partial J(\mathbf{W})}{\partial \hat{y}} * \frac{\partial \hat{y}}{\partial z_1} * \frac{\partial z_1}{\partial w_1}$$

Repeat this for **every weight in the network** using gradients from later layers

Optimization through gradient descent

$$\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$$

Optimization through gradient descent

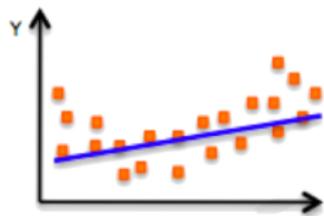
$$\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$$

How can we set the  
learning rate?

## Gradient Descent Algorithms

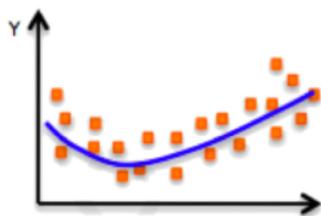
Algorithm	TF Implementation	Reference
• SGD	 <code>tf.keras.optimizers.SGD</code>	Kiefer & Wolfowitz. "Stochastic Estimation of the Maximum of a Regression Function." 1952.
• Adam	 <code>tf.keras.optimizers.Adam</code>	Kingma et al. "Adam: A Method for Stochastic Optimization." 2014.
• Adadelata	 <code>tf.keras.optimizers.Adadelta</code>	Zeiler et al. "ADADELTA: An Adaptive Learning Rate Method." 2012.
• Adagrad	 <code>tf.keras.optimizers.Adagrad</code>	Duchi et al. "Adaptive Subgradient Methods for Online Learning and Stochastic Optimization." 2011.
• RMSProp	 <code>tf.keras.optimizers.RMSProp</code>	

# The problem of Overfitting

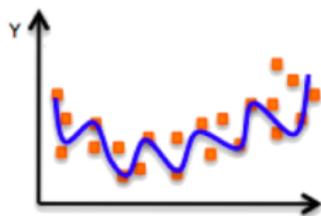


## Underfitting

Model does not have capacity to fully learn the data



## Ideal fit



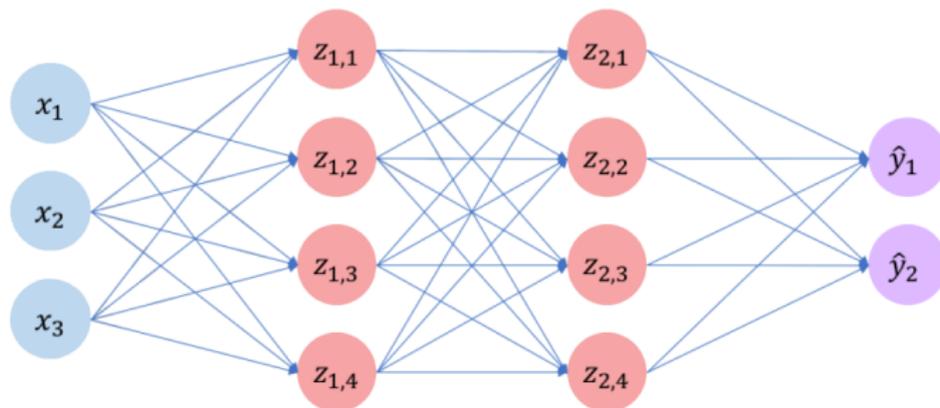
## Overfitting

Too complex, extra parameters, does not generalize well

- A technique that constraint the optimization problem to avoid complex models.
- We use it to improve the generalization on our model to unseen data.
- There are different kind of methods.

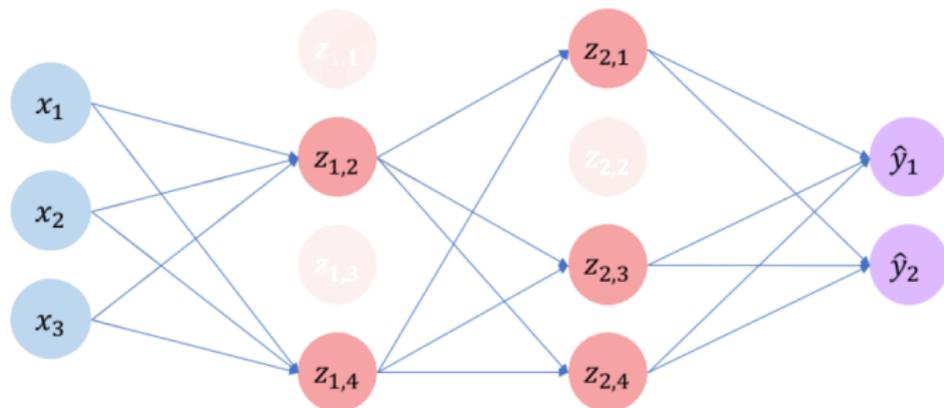
# Regularization I: Dropout

During training randomly set some activations to 0.



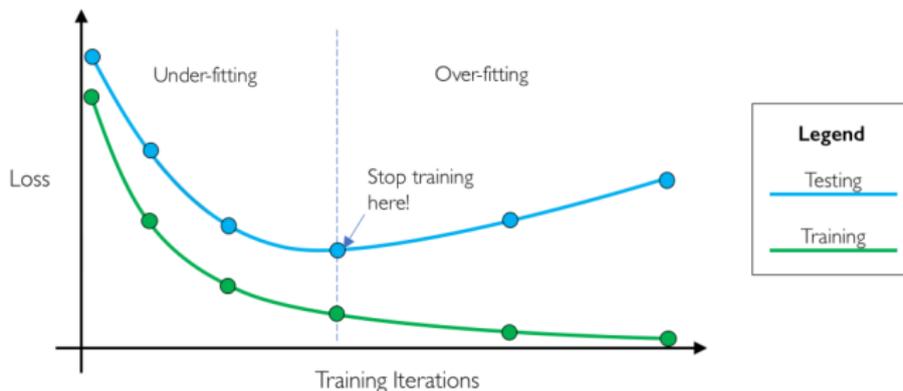
# Regularization I: Dropout

During training randomly set some activations to 0.



# Regularization II: early stopping

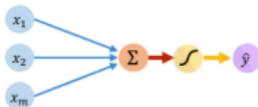
To stop before having the opportunity to overfit by monitoring testing and training data.



# Summary

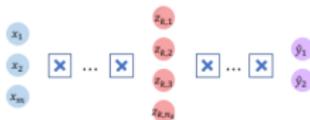
## The Perceptron

- Structural building blocks
- Nonlinear activation functions



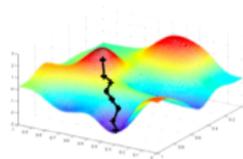
## Neural Networks

- Stacking Perceptrons to form neural networks
- Optimization through backpropagation



## Training in Practice

- Adaptive learning
- Batching
- Regularization

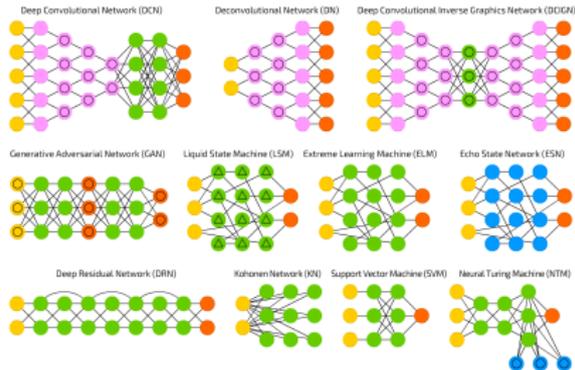
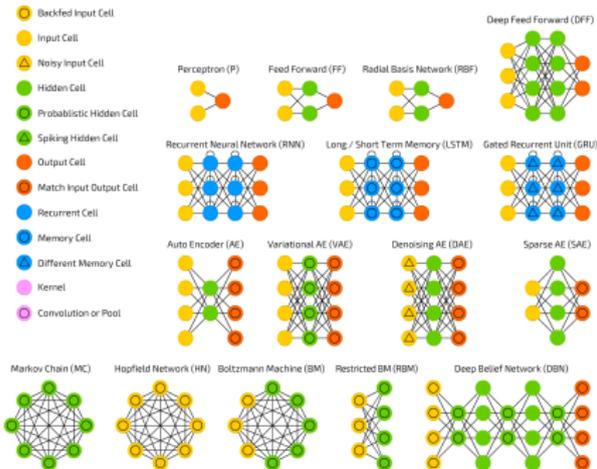


# Neural Networks: The Zoo of topologies



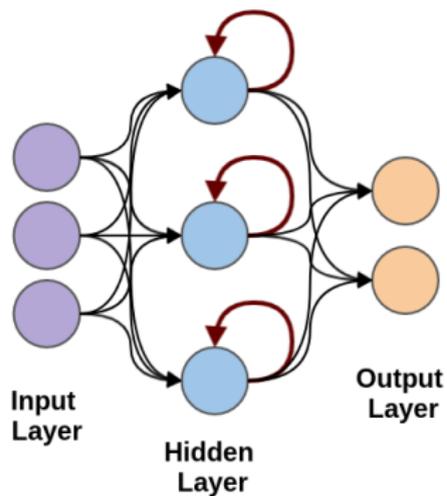
# Neural Networks: The Zoo of topologies

©2016 Fjodor van Veen - asimovinstitute.org

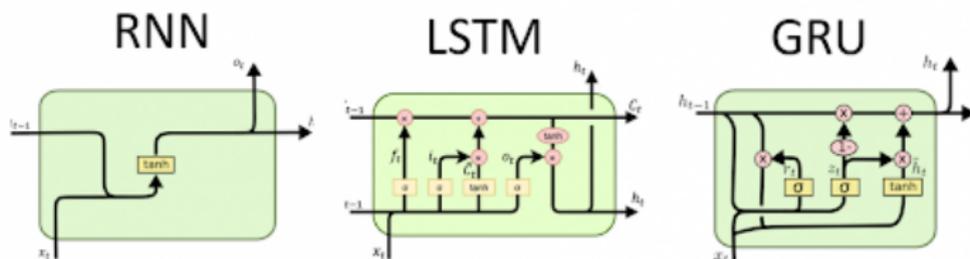


- Feed Forward networks (already seen).
- Recurrent neural networks — simple, LSTM, and GRU.
- Convolutional neural networks

# Topologies: Recurrent neural networks



# Topologies: Recurrent neural networks

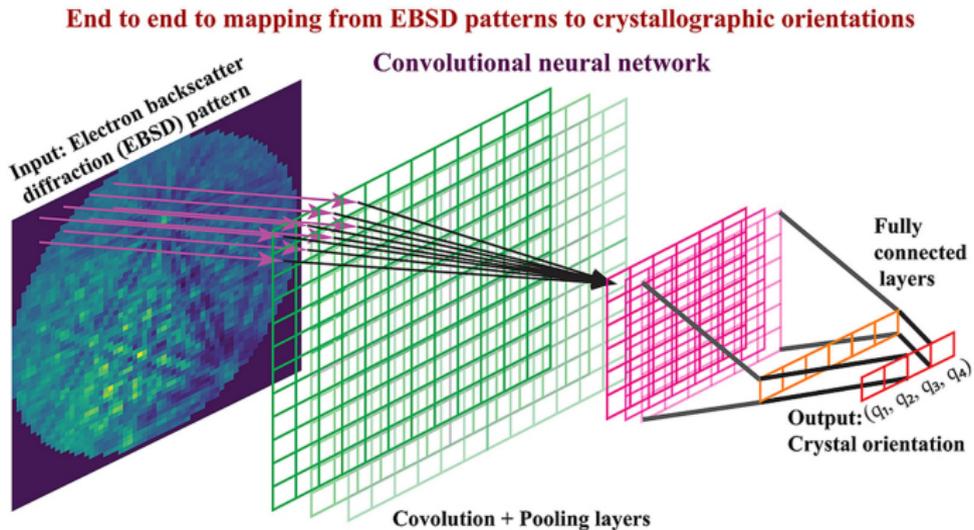


- Recurrent neural networks are a class of neural networks that exploit the sequential nature of their input.
- Inputs could be: a text, a speech, time series, and anything else where the occurrence of an element in the sequence is dependent on the elements that appeared before it.

# Convolutional and pooling layers

- ConvNets are a class of neural networks using convolutional and pooling operations for progressively learning rather sophisticated models based on progressive levels of abstraction.
- This learning via progressive abstraction resembles vision models that have evolved over millions of years inside the human brain.
- People called it deep with 3-5 layers a few years ago, and now it has gone up to 100-200.

# Convolutional Neural Networks.

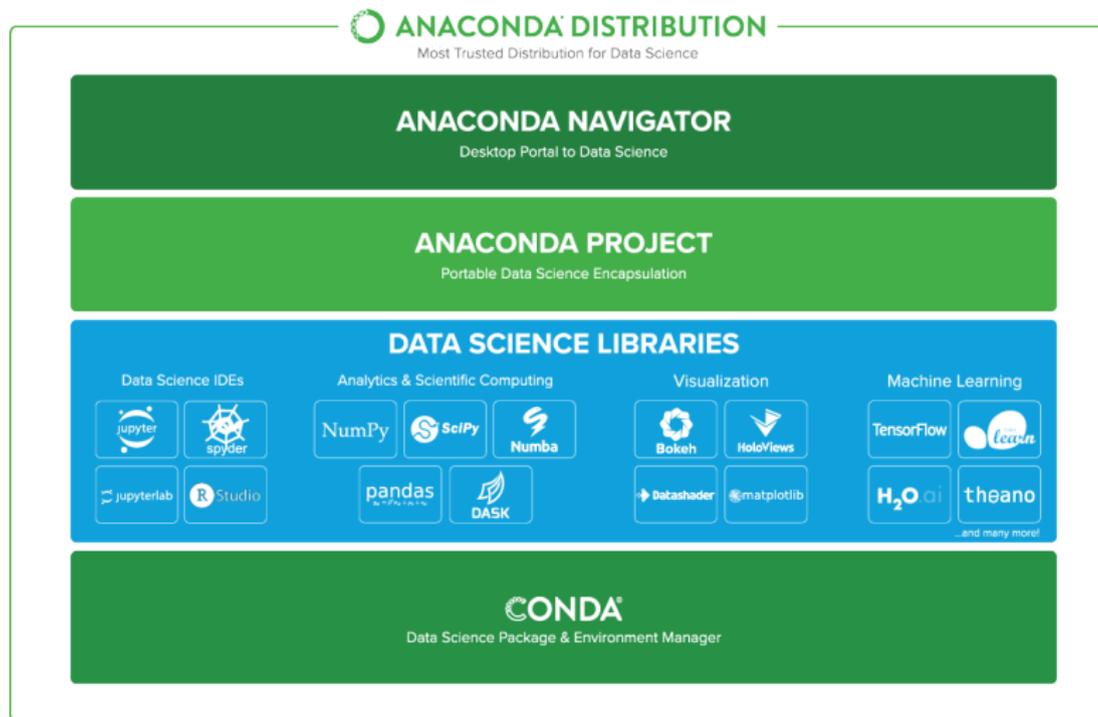




Where can we work?

- Locally in Virtual env:  
Main purpose of Python virtual env is to create an isolated environment for Python projects. Each project can have its own dependencies, regardless of what dependencies every other project has.  
<https://realpython.com/python-virtual-environments-a-primer/>
- Google Cloud ML.  
<https://cloud.google.com/ai-platform/docs/getting-started-keras>
- other Services

# About installation Local: What is Anaconda for?





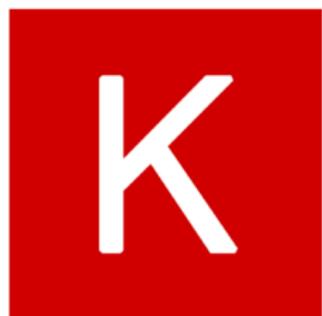
# How do we implement this algorithms?

- From zero with math libraries and python.
- Using dedicated open source frameworks:
  - Tensorflow.
  - Keras.



## TensorFlow

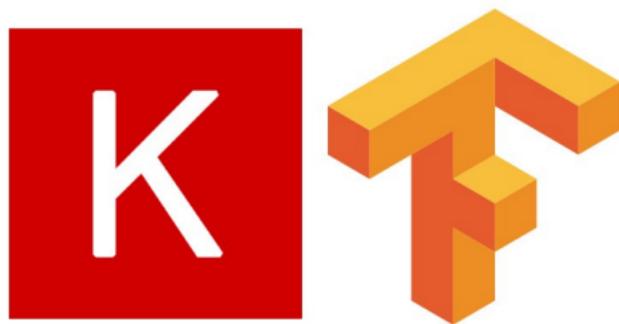
Tensorflow is an end-to-end open source platform for machine learning. It has a comprehensive, flexible ecosystem of tools, libraries and community resources that lets researchers push the state-of-the-art in ML and developers easily build and deploy ML powered applications.



# Keras

Keras: A high-level neural networks API, written in Python and capable of running on top of TensorFlow, CNTK, or Theano. It was developed with a focus on enabling fast experimentation. Being able to go from idea to result with the least possible delay is key to doing good research.

# Keras is now embedded in Tensorflow



<https://www.tensorflow.org/guide/keras/functional>

- **Modularity:** A model is either a sequence or a graph of standalone modules that can be combined together like LEGO blocks for building neural networks
- **The libraries** predefines a large number of modules implementing different types of neural layers, cost functions, optimizers, initialization schemes, activation functions, and regularization schemes.
- **Minimalism:** The library is implemented in Python and each module is kept short and self-describing.
- **Easy extensibility:** The library can be extended with new functionalities.

# Keras Basics:

## Python For Data Science Cheat Sheet

### Keras

Learn Python for data science interactively at [www.DataCamp.com](http://www.DataCamp.com)



#### Keras

Keras is a powerful and easy-to-use deep learning library for Theano and TensorFlow that provides a high-level neural networks API to develop and evaluate deep learning models.

#### A Basic Example

```
>>> import numpy as np
>>> from keras.models import Sequential
>>> from keras.layers import Dense
>>> data = np.random.random((1000,100))
>>> labels = np.random.randint(2,size=(1000,1))
>>> model = Sequential()
>>> model.add(Dense(100,
                    activation='relu',
                    input_dim=100))
>>> model.add(Dense(1, activation='sigmoid'))
>>> model.compile(optimizer='rmsprop',
                loss='binary_crossentropy',
                metrics=['accuracy'])
>>> model.fit(data, labels, epochs=10, batch_size=32)
>>> predictions = model.predict(data)
```

#### Data

Also see NumPy, Pandas & Scikit-Learn

Your data needs to be stored as NumPy arrays or as a list of NumPy arrays. Ideally, you split the data in training and test sets, for which you can also resort to the `train_test_split` module of `sklearn.cross_validation`.

#### Keras Data Sets

```
>>> from keras.datasets import mnist, fashion_mnist, cifar10, cifar100
>>> (x_train, y_train), (x_test, y_test) = mnist.load_data()
>>> (x_train, y_train2), (x_test, y_test2) = fashion_mnist.load_data()
>>> (x_train, y_train3), (x_test, y_test3) = cifar10.load_data()
>>> (x_train, y_train4), (x_test, y_test4) = mnist.load_data(num_words=20000)
>>> num_classes = 10
```

#### Other

```
>>> from urllib.request import urlopen
>>> data = np.loadtxt(urlopen("http://archive.ics.uci.edu/ml/machine-learning-databases/pima-indians-diabetes/pima-indians-diabetes.data"), delimiter=",")
>>> X = data[:,0:8]
>>> y = data[:,8]
```

#### Preprocessing

##### Sequence Padding

```
>>> from keras.preprocessing import sequence
>>> train = sequence.pad_sequences(x_train, maxlen=80)
>>> x_test = sequence.pad_sequences(x_test, maxlen=80)
```

##### One-Hot Encoding

```
>>> from keras.utils import to_categorical
>>> y_train = to_categorical(y_train, num_classes)
>>> y_test = to_categorical(y_test, num_classes)
>>> Y_train = to_categorical(x_train, num_classes)
>>> Y_test = to_categorical(x_test, num_classes)
```

#### Model Architecture

##### Sequential Model

```
>>> from keras.models import Sequential
>>> model = Sequential()
>>> model2 = Sequential()
>>> model3 = Sequential()
```

##### Multilayer Perceptron (MLP)

###### Binary Classification

```
>>> from keras.layers import Dense
>>> model.add(Dense(12,
                  input_dim=8,
                  kernel_initializer='uniform',
                  activation='relu'))
>>> model.add(Dense(8, kernel_initializer='uniform', activation='relu'))
>>> model.add(Dense(1, kernel_initializer='uniform', activation='sigmoid'))
```

###### Multi-Class Classification

```
>>> from keras.layers import Dropout
>>> model.add(Dense(512, activation='relu', input_shape=(784,)))
>>> model.add(Dropout(0.2))
>>> model.add(Dense(512, activation='relu'))
>>> model.add(Dropout(0.2))
>>> model.add(Dense(10, activation='softmax'))
```

###### Regression

```
>>> model.add(Dense(64, activation='relu', input_dim=train_data.shape[1]))
>>> model.add(Dense(1))
```

##### Convolutional Neural Network (CNN)

```
>>> from keras.layers import Activation, Conv2D, MaxPooling2D, Flatten
>>> model2.add(Conv2D(32, (3, 3), padding='same', input_shape=x_train_shape[1]))
>>> model2.add(Activation('relu'))
>>> model2.add(Conv2D(32, (3, 3)))
>>> model2.add(Activation('relu'))
>>> model2.add(MaxPooling2D(pool_size=(2, 2)))
>>> model2.add(Dropout(0.25))
>>> model2.add(Conv2D(64, (3, 3), padding='same'))
>>> model2.add(Activation('relu'))
>>> model2.add(Conv2D(64, (3, 3)))
>>> model2.add(Activation('relu'))
>>> model2.add(MaxPooling2D(pool_size=(2, 2)))
>>> model2.add(Dropout(0.25))
>>> model2.add(Flatten())
>>> model2.add(Dense(512))
>>> model2.add(Activation('relu'))
>>> model2.add(Dropout(0.5))
>>> model2.add(Dense(num_classes))
>>> model2.add(Activation('softmax'))
```

##### Recurrent Neural Network (RNN)

```
>>> from keras.layers import Embedding, LSTM
>>> model3.add(Embedding(10000, 128))
>>> model3.add(LSTM(128, dropout=0.2, recurrent_dropout=0.2))
>>> model3.add(Dense(1, activation='sigmoid'))
```

Also see NumPy & Scikit-Learn

##### Train and Test Sets

```
>>> from sklearn.model_selection import train_test_split
>>> X_train, X_test, y_train, y_test = train_test_split(X,
                                                    y,
                                                    test_size=0.33,
                                                    random_state=43)
```

##### Standardization/Normalization

```
>>> from sklearn.preprocessing import StandardScaler
>>> scaler = StandardScaler().fit(x_train2)
>>> x_train_scaled = scaler.transform(x_train2)
>>> standardized_x_test = scaler.transform(x_test2)
```

#### Inspect Model

```
>>> model.output_shape
>>> model.summary()
>>> model.get_config()
>>> model.get_weights()
```

Model output shape  
Model summary representation  
Model configuration  
List all weight tensors in the model

#### Compile Model

##### MLP: Binary Classification

```
>>> model.compile(optimizer='adam',
                loss='binary_crossentropy',
                metrics=['accuracy'])
```

##### MLP: Multi-Class Classification

```
>>> model.compile(optimizer='rmsprop',
                loss='categorical_crossentropy',
                metrics=['accuracy'])
```

##### MLP: Regression

```
>>> model.compile(optimizer='rmsprop',
                loss='mse',
                metrics=['mse'])
```

##### Recurrent Neural Network

```
>>> model3.compile(loss='binary_crossentropy',
                 optimizer='adam',
                 metrics=['accuracy'])
```

#### Model Training

```
>>> model3.fit(x_train,
             y_train,
             batch_size=32,
             epochs=15,
             verbose=1,
             validation_data=(x_test, y_test))
```

#### Evaluate Your Model's Performance

```
>>> score = model3.evaluate(x_test,
                          y_test,
                          batch_size=32)
```

#### Prediction

```
>>> model3.predict(x_test, batch_size=32)
>>> model3.predict_classes(x_test, batch_size=32)
```

#### Save/Reload Models

```
>>> from keras.models import load_model
>>> model3.save('model.h5')
>>> my_model = load_model('my_model.h5')
```

#### Model Fine-tuning

##### Optimization Parameters

```
>>> from keras.optimizers import RMSProp
>>> opt = RMSProp(lr=0.001, decay=1e-6)
>>> model2.compile(loss='categorical_crossentropy',
                 optimizer=opt,
                 metrics=['accuracy'])
```

##### Early Stopping

```
>>> from keras.callbacks import EarlyStopping
>>> early_stopping_monitor = EarlyStopping(patience=2)
>>> model3.fit(x_train,
            y_train,
            batch_size=32,
            epochs=15,
            validation_data=(x_test, y_test),
            callbacks=[early_stopping_monitor])
```

DataCamp

Learn Python For Data Science Interactively



# Keras Basics: code example

You can create a Sequential model by passing a list of layer instances to the constructor:

```
1
2 from keras.models import Sequential
3 from keras.layers import Dense,
  Activation
4
5 model = Sequential([
6     Dense(32, input_shape=(784,)),
7     Activation('relu'),
8     Dense(10),
9     Activation('softmax'),
10 ])
11
```

You can also simply add layers via the `.add()` method:

```
1 model = Sequential()
2 model.add(Dense(32, input_dim=784))
3 model.add(Activation('relu'))
```

# More:

[https://github.com/katejarne/Keras\\_tensorflow\\_course](https://github.com/katejarne/Keras_tensorflow_course)

Classes and material will be at:

<http://ceciliajarne.web.unq.edu.ar/cns-2020-tutorial/>

## References:

- Deep Learning (The MIT Press Essential Knowledge series)
- <http://introtodeeplearning.com/> MIT course.
- <https://www.tensorflow.org/>
- <https://keras.io/> Francois Chollet et al. Keras. 2015.
- Martín Abadi, et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015.