

Learning

Neural network (very simplified 1/2)

• NN = h, a large function with numerous parameters (weights θ)

Think about a polynomial function ...

$$f(x) = ax + bx^2 + cx^3 + dx^4..$$



Neural network (very simplified 1/2)

- NN = h, a large function with numerous parameters (weights θ)
- Input: x, as a vector (matrix)
 - Columns: dimension
 - Example: tumors properties (size, genes, etc)
 - Example: 128x128 HU values (a CT slice)
 - Row: samples
- Output: y as a vector (matrix)
 - Example: one value, probability to be maligne/benigne
 - Example: 128x128 values, a segmented image
 - One row per input sample



Neural network (very simplified 2/2)

- Training a NN = find the values of the parameters θ
 - Input: a large database with **both x** and **y**
 - Loss function: a function to estimate the predicting error bw real samples from the db and estimated ones from h
 - Optimisation: an iterative algorithm to find theta
- Long process, need GPU
- To be performed only once
- The results is h, a set of parameters values
- Can be used with any x, to guess y

$$y = h(x, \theta)$$

$$Loss(y, \hat{y})$$

Known y Predicted y





Cell body

Information is segregated into "useful" and "not useful" categories.

Dendrite Input to a neuron

Axon

Output to the neuron





Find the parameters values

Optimisation = found the min

- Parameters = research space dim
- Loss function = diff x and y
- Iterative, number of *epoch*
- Direction = loss gradient
- Learning rate = convergence step





Back-propagation Learning process



Neural Network Activation Functions: a small subset!



10 Most Common Loss Functions		
in Machine Learning		blog.DailyDoseofD5.com
Loss Function Name	Description	Function
Regression Loss Functions		
Mean Bias Error	Captures average bias in prediction. But is rarely used for training.	$\mathcal{L}_{\mathcal{MBE}} = \frac{1}{N} \sum_{i=1}^{N} (y_i - f(x_i))$
Mean Absolute Error	Measures absolute average bias in prediction. Also called L1 Loss.	$\mathcal{L}_{\mathcal{MAE}} = \frac{1}{N} \sum_{i=1}^{N} y_i - f(x_i) $
Mean Squared Error	Average squared distance between actual and predicted. Also called L2 Loss.	$\mathcal{L}_{MSE} = \frac{1}{N} \sum_{i=1}^{N} (y_i - f(x_i))^2$
Root Mean Squared Error	Square root of MSE. Loss and dependent variable have same units.	$\mathcal{L}_{\mathcal{RMSE}} = \sqrt{\frac{1}{N} \sum_{i=1}^{N} (y_i - f(x_i))^2}$
Huber Loss	A combination of MSE and MAE. It is parametric loss function.	$\mathcal{L}_{\text{Huberloss}} = \begin{cases} \frac{1}{2}(y_i - f(x_i))^2 & : y_i - f(x_i) \le \delta\\ \delta(y_i - f(x_i) - \frac{1}{2}\delta) & : otherwise \end{cases}$
Log Cosh Loss	Similar to Huber Loss + non-parametric. But computationally expensive.	$\mathcal{L}_{\text{LogCosh}} = \frac{1}{N} \sum_{i=1}^{N} log(cosh(f(x_i) - y_i))$
Classification Loss Functions (Binary + Multi-class)		
Binary Cross Entropy (BCE)	Loss function for binary classification tasks.	$\mathcal{L}_{BCE} = \frac{1}{N} \sum_{i=1}^{N} y_i \cdot log(p(x_i)) + (1 - y_i) \cdot log(1 - p(x_i))$
Hinge Loss	Penalizes wrong and right (but less confident) predictions. Commonly used in SVMs.	$\mathcal{L}_{\text{Hinge}} = max(0, 1 - (f(x) \cdot y))$
Cross Entropy Loss	Extension of BCE loss to multi- class classification.	$\mathcal{L}_{C\mathcal{E}} = -\frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{M} y_{ij}.log(f(x_{ij}))$ $N: samples; M: classes$
KL Divergence	Minimizes the divergence between predicted and true probability distribution	$\mathcal{L}_{\mathcal{KL}} = \sum_{i=1}^{N} y_i \cdot log(rac{y_i}{f(x_i)})$

https://playground.tensorflow.org









• Neural network : large non-linear function, with "neurons"

Summary

- Layers: input / parameters / output
- Training process: optimization, loss, backpropagation
- Image specific: **convolutional** neural network



