



Your First MLP Code

Recitation 1, part 1
Spring 2022



Overview

- Neural Networks
- Perceptrons
- Multilayer perceptrons
 - Forward Pass
 - Backpropagation
 - Update Weights


Neural Networks

- The brain, made up of connected neurons, are the inspirations for artificial neural networks.



Neural Networks

- A neuron is a node with many inputs and one output.
- A neural network consists of many interconnected neurons -- a “simple” device that receives data at the input and provides a response.
- Information are transmitted from one neuron to another by electrical impulses and chemical signals.



Perceptrons

- Perceptron is a single layer neural network.

Perceptrons

- Perceptron is a single layer neural network.
- The perceptron consists of 4 parts.

Perceptrons

- Perceptron is a single layer neural network.
- The perceptron consists of 4 parts.
 - Input values



x_1

The diagram illustrates the input layer of a perceptron. It consists of a vertical column of four light blue circles. The top circle contains the label x_1 . The second circle contains a vertical ellipsis (three dots). The third circle contains the label x_{n-1} . The bottom circle contains the label x_n .


•
•
•

x_{n-1}

x_n


Perceptrons

- Perceptron is a single layer neural network.
- The perceptron consists of 4 parts.
 - Input values
 - Weights




Perceptrons

- Perceptron is a single layer neural network.
- The perceptron consists of 4 parts.
 - Input values
 - Weights
 - Weighted sums




Perceptrons

- Perceptron is a single layer neural network.
- The perceptron consists of 4 parts.
 - Input values
 - Weights
 - Weighted sums
 - Threshold / Activation functions




Perceptrons

- Perceptron is a single layer neural network.
- The perceptron consists of 4 parts.
- The perceptron works on the following steps:
 - Multiply all inputs with their weights







Perceptrons

- Perceptron is a single layer neural network.
- The perceptron consists of 4 parts.
- The perceptron works on the following steps:
 - Multiply all inputs with their weights
 - Add all multiplies values \rightarrow weighted sum





Perceptrons

- Perceptron is a single layer neural network.
- The perceptron consists of 4 parts.
- The perceptron works on the following steps:
 - Multiply all inputs with their weights
 - Add all multiplies values → weighted sum
 - Apply the weighted sum to activation function




Perceptrons

- Perceptron is usually used to classify the data into two parts -- **Linear Binary Classifier.**



Perceptrons

- Perceptron is usually used to classify the data into two parts --
Linear Binary Classifier.
 - **Weights** shows the **strength** of the particular node.
 - **Activation functions** are used to map the input between the required values





Multilayer Perceptrons

What if we want to be able to distinguish between more classes?


Multilayer Perceptrons

What if we want to be able to distinguish between more classes?

- Introduce more perceptrons !




Multilayer Perceptrons



In order to correctly classify things, the network must be **learned**.

But first, **what** do we need to learn?


The parameters (or the weights)




How do we learn?

→ Actual Function that we are trying to model:

◆ Note: We don't know the actual function.



→ We only have several sample data points on this function.



→ Our goal:

◆ **Estimate the function with the given samples.**

How do we learn?

→ A measurement of **error**

- ◆ How much off is the **network output** with respect to the **desired output**

The diagram shows the loss function $Loss(W) = \frac{1}{N} \sum_i div(f(X_i, W), d_i)$ with several annotations. An arrow points from 'Number of samples' to $\frac{1}{N}$. An arrow points from 'For each sample' to the summation index i . An arrow points from 'Divergence function' to div . A box encloses the function $f(X_i, W)$, with an arrow pointing to it from 'MLP'. An arrow points from 'Sample value' to X_i . An arrow points from 'Current weights of estimated function' to W . An arrow points from 'Network Output' to the output of the MLP box. An arrow points from 'Desired Output' to d_i .

$$Loss(W) = \frac{1}{N} \sum_i div(f(X_i, W), d_i)$$


→ Our goal (more specifically):

- ◆ Minimize the loss

$$\hat{W} = \arg \min_W Loss(W)$$


How do we learn?

→ Gradient Descent




Forward Pass


- For each single perceptron




Forward Pass




Forward Pass




Forward Pass



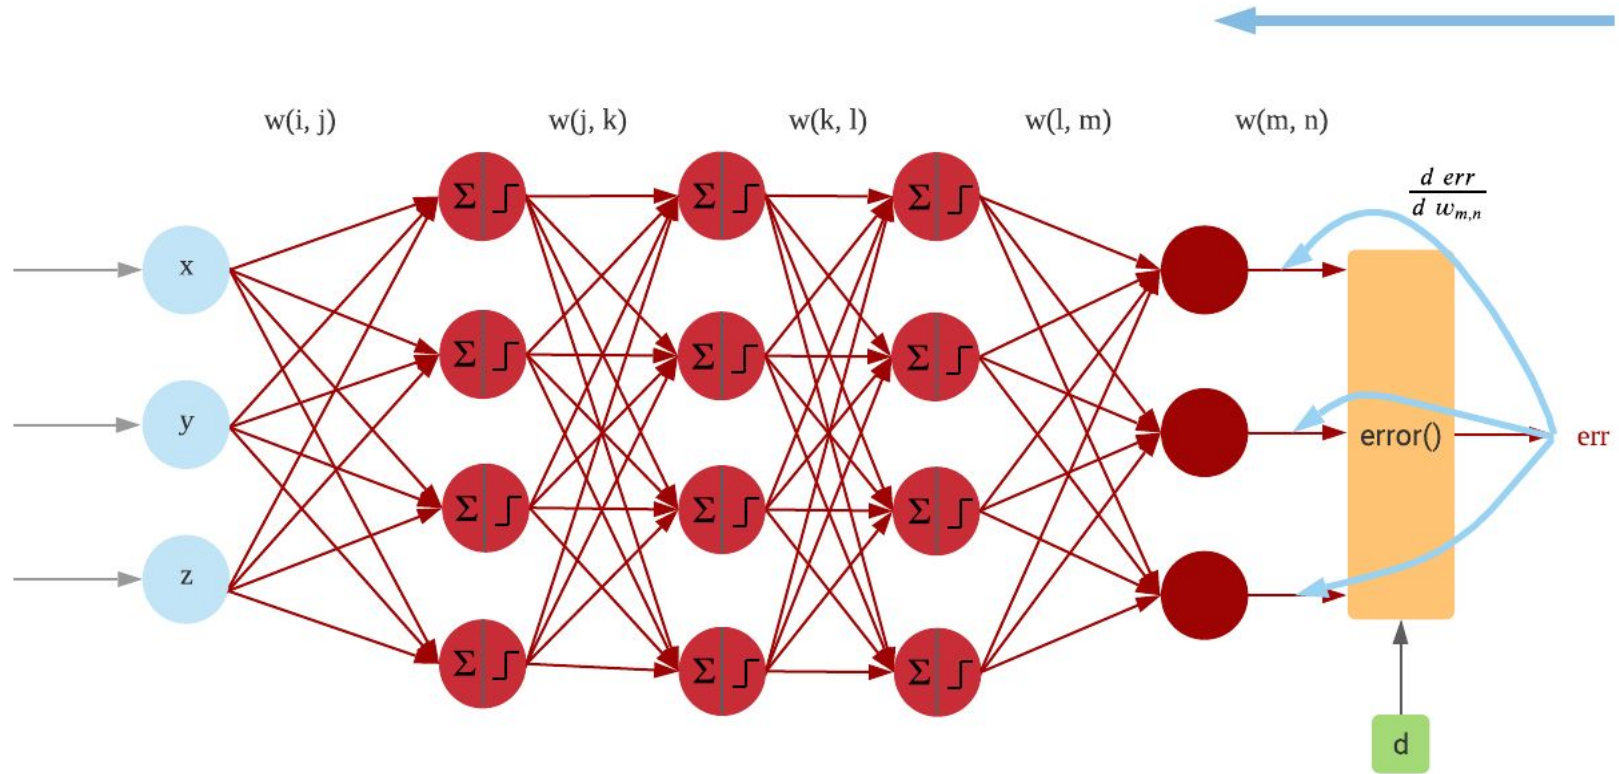
Forward Pass




Error




Backpropagation




Backpropagation




Backpropagation




Backpropagation



Backpropagation



Backpropagation



Backpropagation

...

Backpropagation

All gradients of weights w.r.t error are calculated!



Update Weights

$$W \leftarrow W - \eta \cdot \nabla_W \text{Loss}(W)$$

learning
rate

gradient

What should be the learning rate?



Optimizers

Gradient Descent:

$$\theta_{t+1} = \theta_t - \alpha \cdot \nabla_{\theta} J(\theta)$$

Momentum (<http://proceedings.mlr.press/v28/sutskever13.pdf>):

$$\begin{aligned} m_{t+1} &= \mu \cdot m_t + \alpha \cdot \nabla_{\theta} J(\theta) \\ \theta_{t+1} &= \theta_t - m_{t+1} \end{aligned}$$

Adagrad (<https://jmlr.org/papers/volume12/duchi11a/duchi11a.pdf>):

$$\begin{aligned} g &\leftarrow \nabla_{\theta} J(\theta) \\ r &\leftarrow r + g^2 \\ \Delta\theta &\leftarrow \frac{\delta}{\sqrt{r + \epsilon}} \cdot g \\ \theta &\leftarrow \theta - \Delta\theta \end{aligned}$$

Optimizers (Cont')

Adam (<https://arxiv.org/pdf/1412.6980.pdf>):

$$m_t = \beta_1 \cdot m_{t-1} + (1 - \beta_1) \cdot g_t$$


$$v_t = \beta_2 \cdot v_{t-1} + (1 - \beta_2) \cdot g_t^2$$

$$\hat{m}_t = m_t / (1 - \beta_1^t)$$

$$\hat{v}_t = v_t / (1 - \beta_2^t)$$

$$\theta_t = \theta_{t-1} - \alpha \cdot \hat{m}_t / (\sqrt{\hat{v}_t} + \epsilon)$$

Visualization



<https://github.com/Jaewan-Yun/optimizer-visualization>

Some fun with TF Playground

Tinker With a **Neural Network** Right Here in Your Browser.
Don't Worry, You Can't Break It. We Promise.

