



Τεχνητή Νοημοσύνη και Μηχανική Μάθηση

Κωνσταντίνος Καραμανής

The University of Texas at Austin & Archimedes/Athena RC

constantine@utexas.edu

<https://caramanis.github.io/>

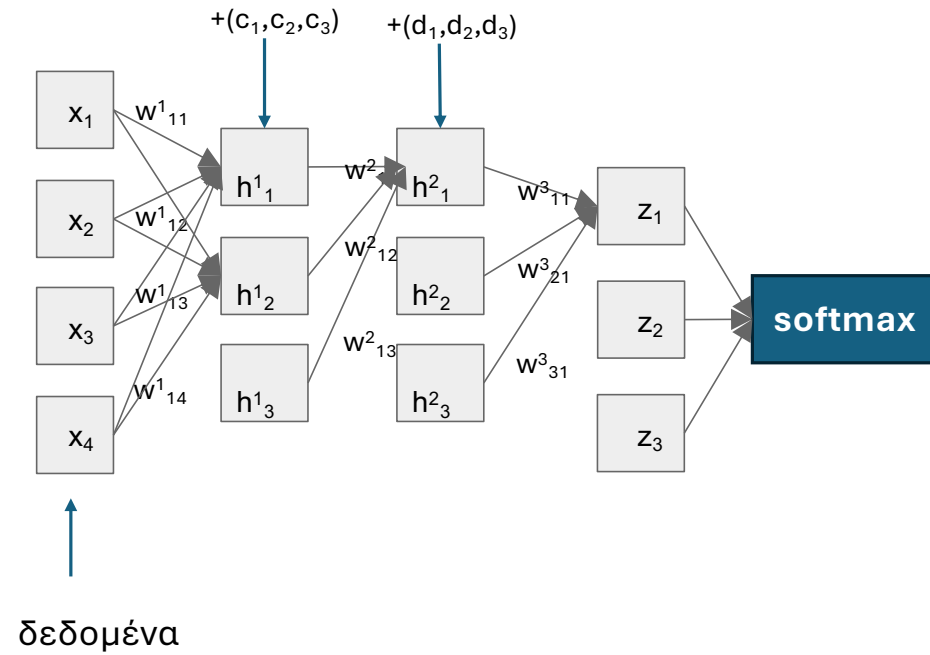
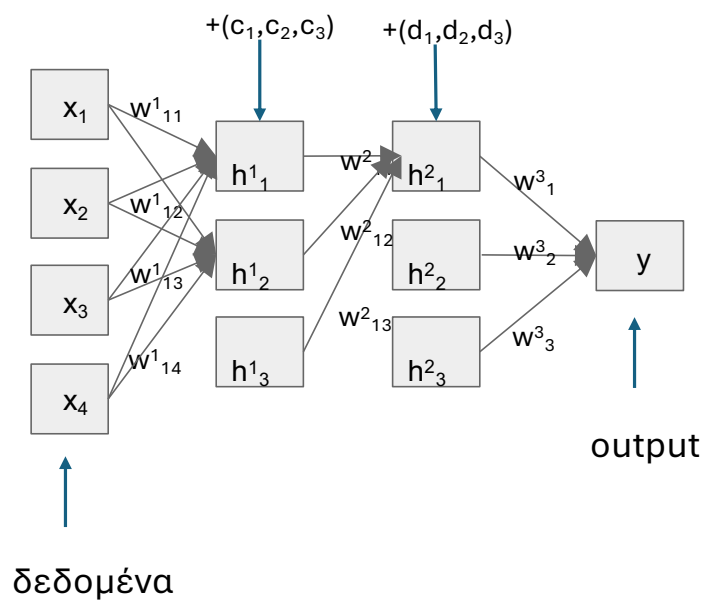





Ας θυμηθούμε τα
προηγούμενα...




Fully Connected + Softmax





Fully Connected + ReLU




Fully Connected + ReLU




Αρκούν τα Γραμμικά Επίπεδα;




Αρκούν τα Γραμμικά Επίπεδα;




$$y = w_1 x_1 + w_2 x_2 + c$$



Αρκούν τα Γραμμικά Επίπεδα;




$$y = w_1 x_1 + w_2 x_2 + c$$




$$y = \alpha h_1 + \beta h_2 + c_3$$

Αρκούν τα Γραμμικά Επίπεδα;




$$y = w_1x_1 + w_2x_2 + c$$




$$y = \alpha h_1 + \beta h_2 + c_3 = \alpha(w_{11}x_1 + w_{12}x_2 + c_1) + \beta(w_{21}x_1 + w_{22}x_2 + c_2) + c_3$$

Αρκούν τα Γραμμικά Επίπεδα;



$$y = w_1 x_1 + w_2 x_2 + c$$




$$\begin{aligned} y &= \alpha h_1 + \beta h_2 + c_3 = \alpha(w_{11}x_1 + w_{12}x_2 + c_1) + \beta(w_{21}x_1 + w_{22}x_2 + c_2) + c_3 \\ &= (\alpha w_{11} + \beta w_{21})x_1 + (\alpha w_{12} + \beta w_{22})x_2 + (\alpha c_1 + \beta c_2) + c_3 \\ &= \widetilde{w}_1 x_1 + \widetilde{w}_2 x_2 + \widetilde{c} \end{aligned}$$

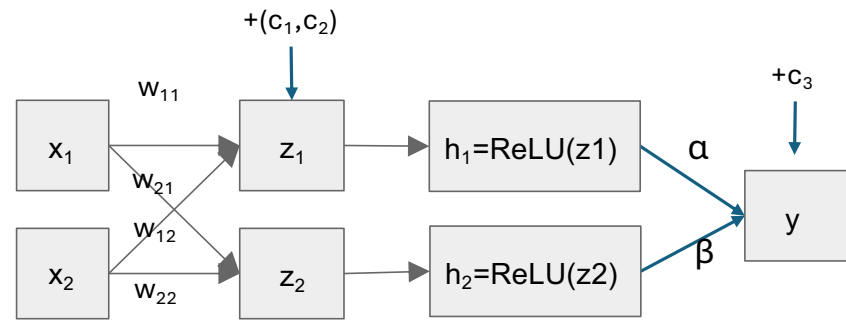
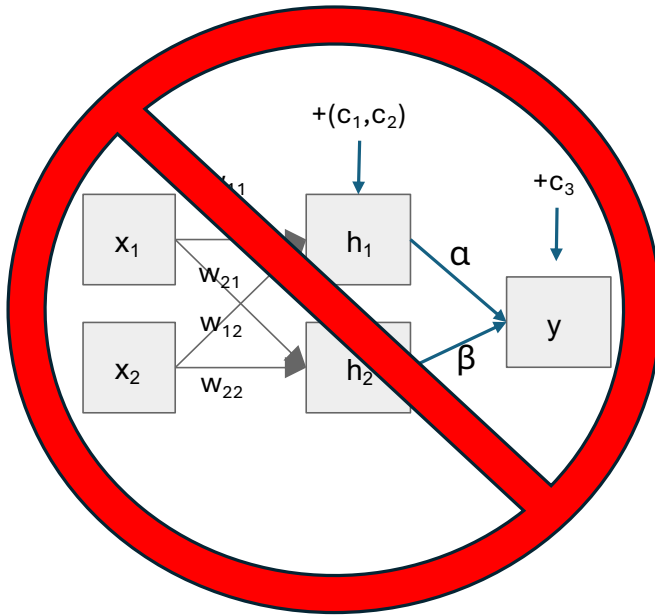
Μη-Γραμμικά Επίπεδα: ReLU

$$\text{ReLU}(\alpha) = \max(0, \alpha)$$

α	$\text{ReLU}(\alpha)$
2	2
-1	0
4	4




Μη-Γραμμικά Επίπεδα: ReLU




Γιατί ReLU και όχι κάτι άλλο;

Activation Functions

Basic non-linear functions




Softmax



! Vanishing Gradient


Advanced non-linear functions

ReLU




! "dying" ReLU

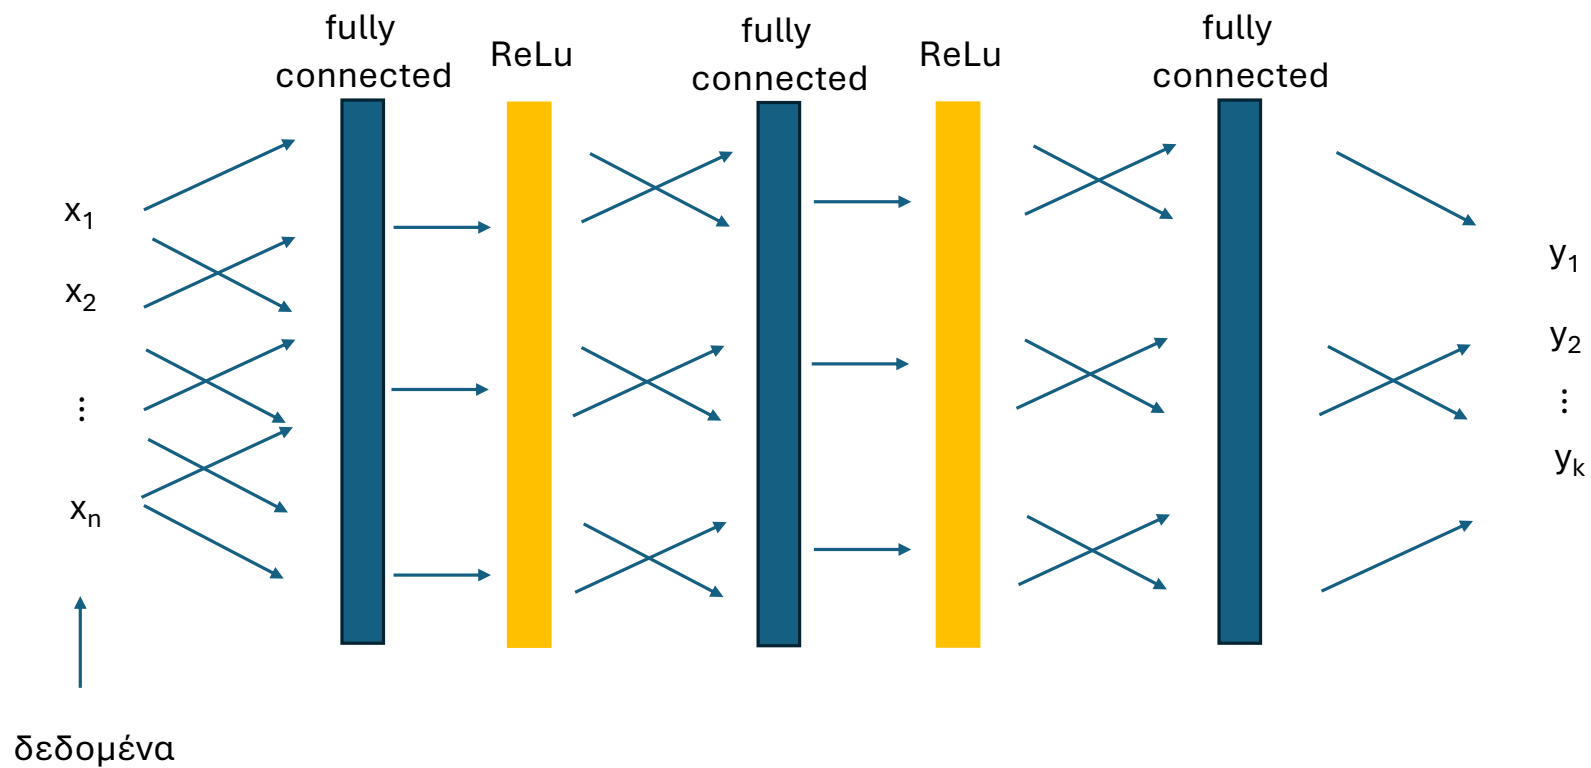
Leaky ReLU



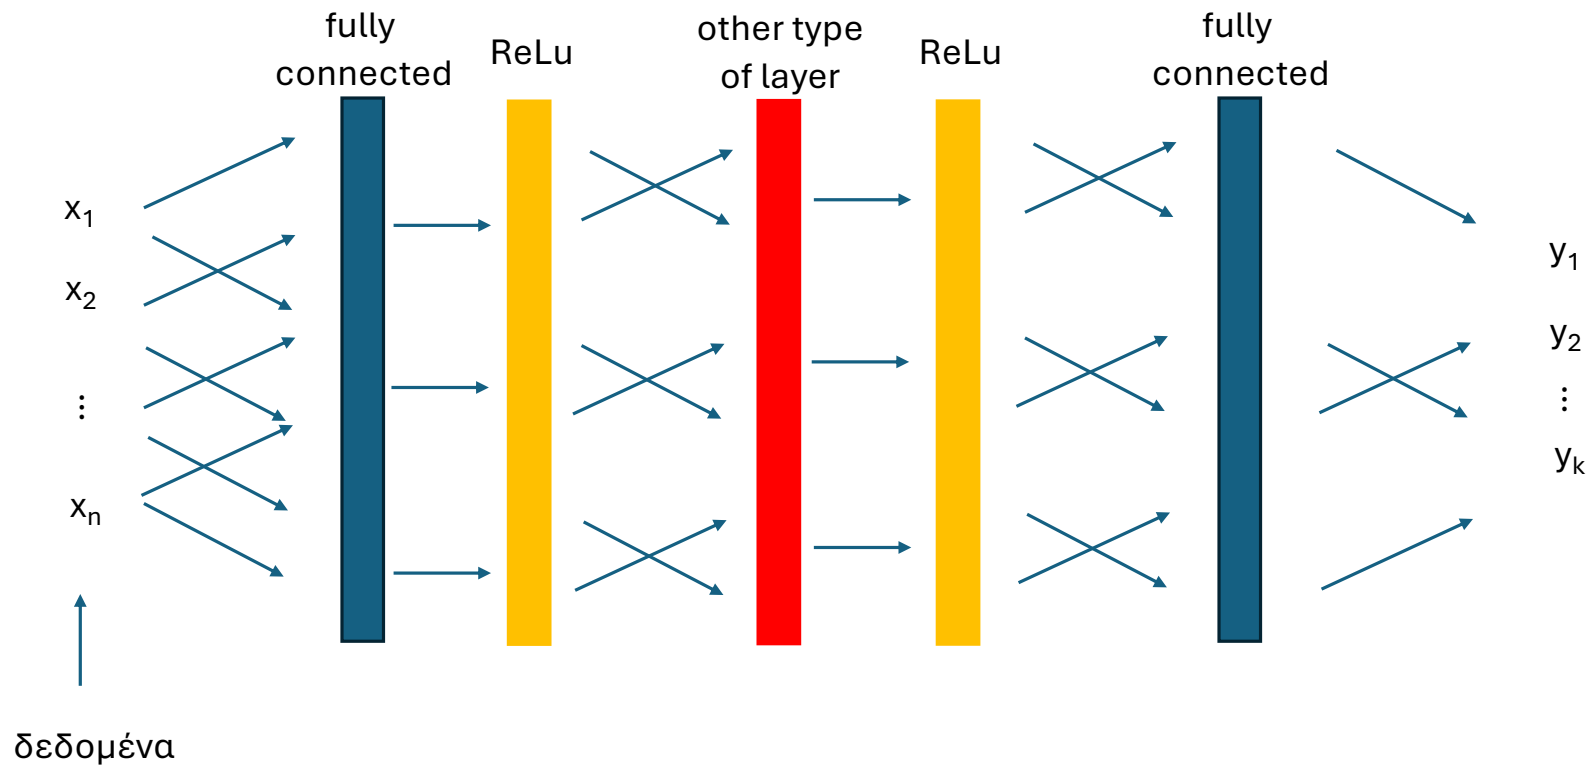
ELU



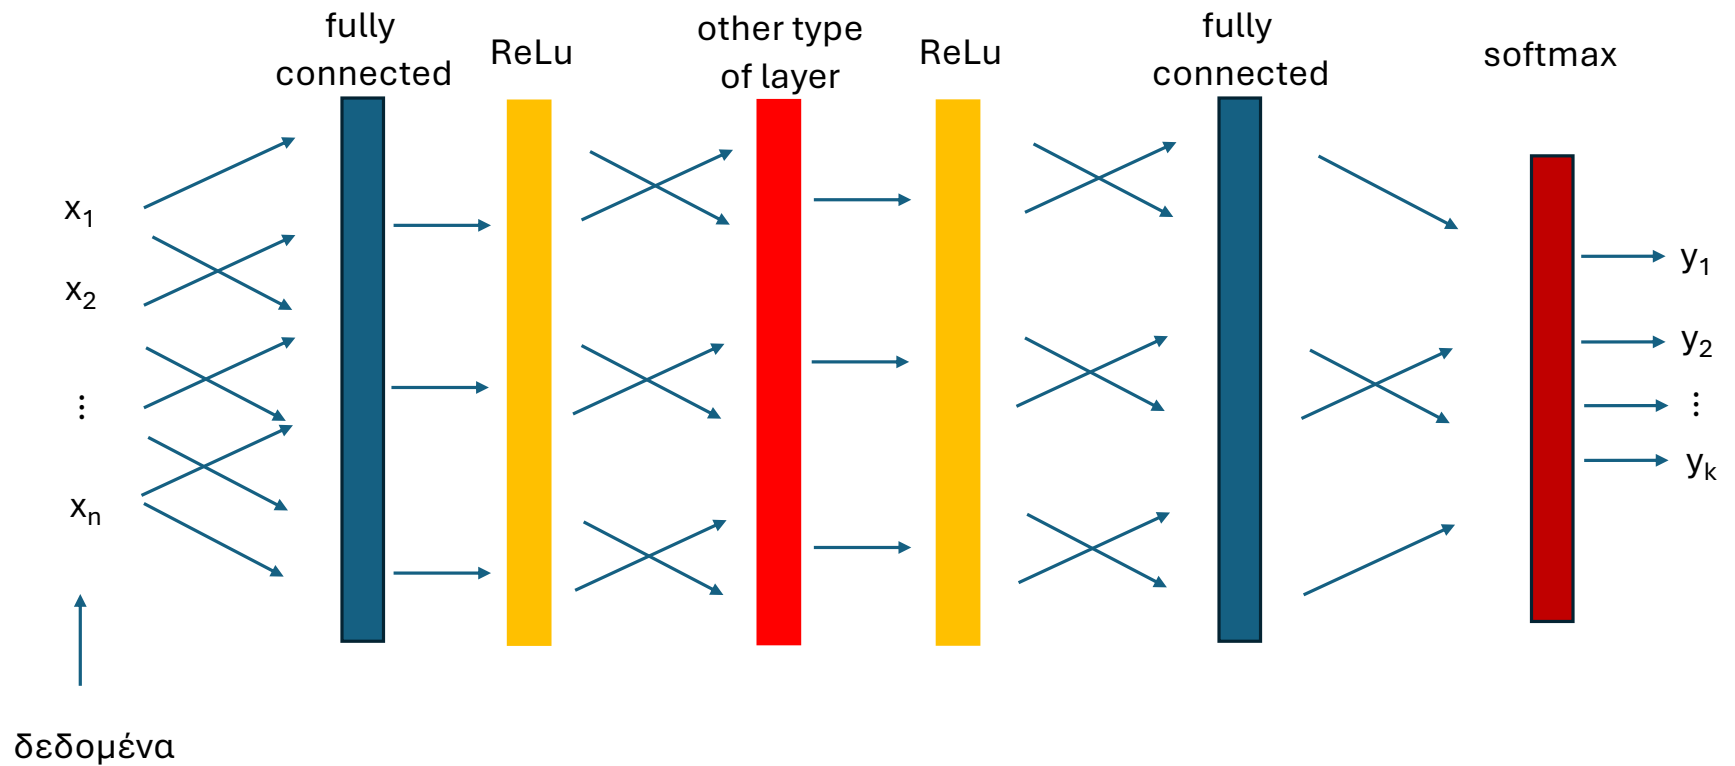
Fully Connected + ReLU



Fully Connected + ReLU + other



Fully Connected + ReLU + SoftMax + other



Πως τα
χρησιμοποιούμε

Δέντρα απόφασης

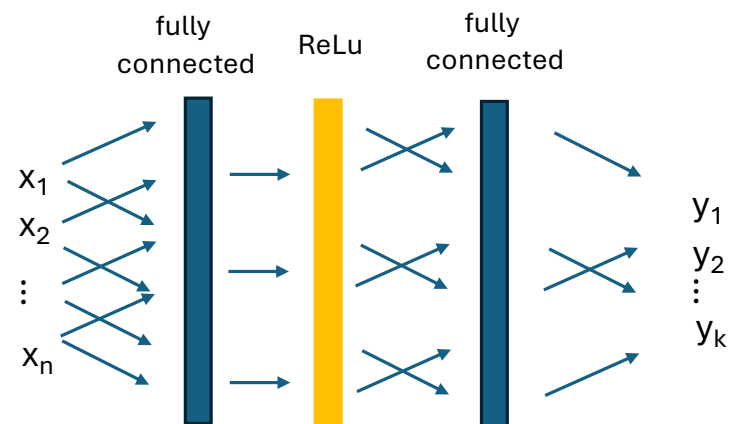
1. `mymodel = DecisionTrees(max_depth=3)` – ορίζουμε την οικογένεια
2. `mymodel.fit(X,y)` – βρίσκουμε παραμέτρους που συμφωνούν με (X,y)
3. `mymodel.predict(x)` – υπολογίζουμε προβλέψεις για τα x

Για νευρωνικά δίκτυα, αυτές οι τρεις ιδέες παραμένουν, αλλά οι εντολές αλλάζουν

model = MyNNNet()


Ορίζουμε κάθε επίπεδο:

MyNNNet = $x \rightarrow fc1 \rightarrow ReLU \rightarrow fc2 \rightarrow y$



model = MyNNNet()

`model = MyNNNet()`



Ορίζουμε κάθε επίπεδο:


$\text{MyNNNet} = x \rightarrow \text{fc1} \rightarrow \text{ReLU} \rightarrow \text{fc2} \rightarrow y$

`model = MyNNNet()`

```
# Define the neural network
class MyNNNet(nn.Module):
    def __init__(self):
        super(MyNNNet, self).__init__()
        self.fc1 = nn.Linear(3, 2) # 3 inputs to 2 outputs
        self.fc2 = nn.Linear(2, 1) # 2 inputs to 1 output
        self.relu = nn.ReLU()

    def forward(self, x):
        x = self.fc1(x)
        x = self.relu(x)
        x = self.fc2(x)
        return x
```

model = MyNNNet()



```
def __init__(self):  
    super(MyNNNet, self).__init__()  
    self.fc1 = nn.Linear(3, 2)  
    self.fc2 = nn.Linear(2, 1)  
    self.relu = nn.ReLU()
```

Define the neural network

```
MyNNNet(nn.Module):
```

```
def __init__(self):
```

```
    super(SimpleNet, self).__init__()
```

```
    self.fc1 = nn.Linear(3, 2) # 3 inputs to 2 outputs
```

```
    self.fc2 = nn.Linear(2, 1) # 2 inputs to 1 output
```

```
    self.relu = nn.ReLU()
```

```
def forward(self, x):
```

```
    x = self.fc1(x)
```


```
    x = self.relu(x)
```

```
    x = self.fc2(x)
```

```
    return x
```

model = MyNNNet()

model = MyNNNet()




```
def __init__(self):  
    super(MyNNNet, self).__init__()  
    self.fc1 = nn.Linear(3, 2)  
    self.fc2 = nn.Linear(2, 1)  
    self.relu = nn.ReLU()
```

```
# Define the neural network  
MyNNNet(nn.Module):  
    def __init__(self):  
        super(MyNNNet, self).__init__() # Call the superclass's __init__ method  
        self.fc1 = nn.Linear(3, 2) # 3 inputs to 2 outputs  
        self.fc2 = nn.Linear(2, 1) # 2 inputs to 1 output  
        self.relu = nn.ReLU()  
  
    def forward(self, x):  
        x = self.fc1(x)  
        x = self.relu(x)  
        x = self.fc2(x)  
        return x
```

model = MyNNNet()

model = MyNNNet()



```
def __init__(self):  
    super(MyNNNet, self).__init__()  
    self.fc1 = nn.Linear(3, 2)  
    self.fc2 = nn.Linear(2, 1)  
    self.relu = nn.ReLU()
```

```
def forward(self, x):  
    x = self.fc1(x)  
    x = self.relu(x)  
    x = self.fc2(x)  
    return x
```

Define the neural network

```
MyNNNet(nn.Module):
```


```
def __init__(self):  
    super(MyNNNet, self).__init__()  
    self.fc1 = nn.Linear(3, 2) # 3 inputs to 2 outputs  
    self.fc2 = nn.Linear(2, 1) # 2 inputs to 1 output  
    self.relu = nn.ReLU()
```

```
def forward(self, x):  
    x = self.fc1(x)  
    x = self.relu(x)  
    x = self.fc2(x)  
    return x
```

Op
My

model = 1


model = MyNNNet()



```
def __init__(self):  
    super(MyNNNet, self).__init__()  
    self.fc1 = nn.Linear(3, 2)  
    self.fc2 = nn.Linear(2, 1)  
    self.relu = nn.ReLU()
```

```
def forward(self, x):  
    x = self.fc1(x)  
    x = self.relu(x)  
    x = self.fc2(x)  
    return x
```


```
# Define the neural network  
MyNNNet(nn.Module):  
    def __init__(self):  
        super(MyNNNet, self).__init__()  
        self.fc1 = nn.Linear(3, 2) # 3 inputs to 2 outputs  
        self.fc2 = nn.Linear(2, 1) # 2 inputs to 1 output  
        self.relu = nn.ReLU()  
  
    def forward(self, x):  
        x = self.fc1(x)  
        x = self.relu(x)  
        x = self.fc2(x)  
        return x
```



Op
My

model = 1

model = BigNet()



Ορίζουμε κάθε επίπεδο:

BigNet = $x \rightarrow fc1 \rightarrow ReLU \rightarrow fc2$


$ReLU \rightarrow fc3 \rightarrow y$

model = BigNet()

```
# Define the neural network
class BigNet(nn.Module):
    def __init__(self):
        super(BigNet, self).__init__()
        self.fc1 = nn.Linear(3, 2) # 3 inputs to 2 outputs
        self.fc2 = nn.Linear(2, 2) # 2 inputs to 2 outputs
        self.fc3 = nn.Linear(2, 1) # 2 inputs to 1 output
        self.relu = nn.ReLU()

    def forward(self, x):
        x = self.fc1(x)
        x = self.relu(x)
        x = self.fc2(x)
        x = self.relu(x)
        x = self.fc3(x)
        return x
```


model = BigNetNet()



```
def __init__(self):  
    super(BigNetNet, self).__init__()  
    self.fc1 = nn.Linear(3, 2)  
    self.fc2 = nn.Linear(2, 2)  
    self.fc3 = nn.Linear(2, 1)  
    self.relu = nn.ReLU()
```


```
# Define the neural network  
class BigNetNet(nn.Module):  
    def __init__(self):  
        super(BigNetNet, self).__init__()  
        self.fc1 = nn.Linear(3, 2) # 3 inputs to 2 outputs  
        self.fc2 = nn.Linear(2, 2) # 2 inputs to 2 outputs  
        self.fc3 = nn.Linear(2, 1) # 2 inputs to 1 output  
        self.relu = nn.ReLU()  
  
    def forward(self, x):  
        x = self.fc1(x)  
        x = self.relu(x)  
        x = self.fc2(x)  
        x = self.relu(x)  
        x = self.fc3(x)  
        return x
```

Op
Big

ReLU → fc3 → y

model = BigNetNet()

model = BigNet()



```
def __init__(self):  
    super(BigNet, self).__init__()  
    self.fc1 = nn.Linear(3, 2)  
    self.fc2 = nn.Linear(2, 2)  
    self.fc3 = nn.Linear(2, 1)  
    self.relu = nn.ReLU()
```


```
# Define the neural network  
class BigNet(nn.Module):  
    def __init__(self):  
        self.fc1 = nn.Linear(3, 2) # 3 inputs to 2 outputs  
        self.relu = nn.ReLU() # 2 inputs to 2 outputs  
        self.fc2 = nn.Linear(2, 2) # 2 inputs to 2 outputs  
        self.fc3 = nn.Linear(2, 1) # 2 inputs to 1 output  
        self.relu = nn.ReLU()  
  
    def forward(self, x):  
        x = self.fc1(x)  
        x = self.relu(x)  
        x = self.fc2(x)  
        x = self.relu(x)  
        x = self.fc3(x)  
        return x
```

Op
Big

ReLU → fc3 → y

model = BigNet()

model = BigNNet()



```
def __init__(self):  
    super(BigNNet, self).__init__()  
    self.fc1 = nn.Linear(3, 2)  
    self.fc2 = nn.Linear(2, 2)  
    self.fc3 = nn.Linear(2, 1)  
    self.relu = nn.ReLU()
```


```
def forward(self, x):  
    x = self.fc1(x)  
    x = self.relu(x)  
    x = self.fc2(x)  
    x = self.relu(x)  
    x = self.fc3(x)  
    return x
```

```
# Define the neural network  
class BigNNet(nn.Module):  
    def __init__(self):  
        self.fc1 = nn.Linear(3, 2) # 3 inputs to 2 outputs  
        self.relu = nn.ReLU() # 2 inputs to 2 outputs  
        self.fc2 = nn.Linear(2, 2) # 2 inputs to 2 outputs  
        self.fc3 = nn.Linear(2, 1) # 2 inputs to 1 output  
        self.relu = nn.ReLU()  
  
    def forward(self, x):  
        x = self.fc1(x)  
        x = self.relu(x)  
        x = self.fc2(x)  
        x = self.relu(x)  
        x = self.fc3(x)  
        return x
```

Op
Big

model :


model = BigNNet()



```
def __init__(self):  
    super(BigNNet, self).__init__()  
    self.fc1 = nn.Linear(3, 2)  
    self.fc2 = nn.Linear(2, 2)  
    self.fc3 = nn.Linear(2, 1)  
    self.relu = nn.ReLU()
```

```
def forward(self, x):  
    x = self.fc1(x)  
    x = self.relu(x)  
    x = self.fc2(x)  
    x = self.relu(x)  
    x = self.fc3(x)  
    return x
```

```
# Define the neural network  
class BigNNet(nn.Module):  
    def __init__(self):  
        super(BigNNet, self).__init__()  
        self.fc1 = nn.Linear(3, 2) # 3 inputs to 2 outputs  
        self.fc2 = nn.Linear(2, 2) # 2 inputs to 2 outputs  
        self.fc3 = nn.Linear(2, 1) # 2 inputs to 1 output  
        self.relu = nn.ReLU()
```




Op
Big

model :

model.fit(X,y)


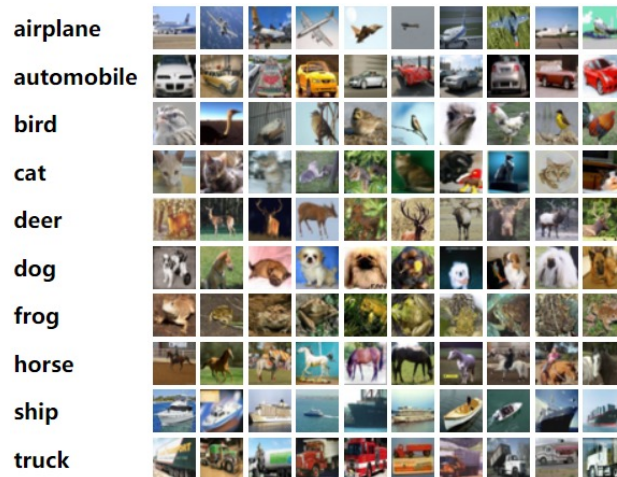
Η εντολή `model.fit(X,y)` ψάχνει για παραμέτρους που ελαχιστοποιούν την απώλεια / μεγιστοποιούν την ακρίβεια, στα δεδομένα εκπαίδευσης.



model.fit(X,y)

Η εντολή `model.fit(X,y)` ψάχνει για παραμέτρους που ελαχιστοποιούν την απώλεια / μεγιστοποιούν την ακρίβεια, στα δεδομένα εκπαίδευσης.


CIFAR-10



model.fit(X,y)

Η εντολή `model.fit(X,y)` ψάχνει για παραμέτρους που ελαχιστοποιούν την απώλεια / μεγιστοποιούν την ακρίβεια, στα δεδομένα εκπαίδευσης.

Άσκηση: η CIFAR-10 έχει 3.072 features. Εάν το 1^ο επίπεδο είναι fully-connected και έχει 1.000 νευρώνες, πόσες παραμέτρους έχει το νευρωνικό δίκτυο μόνο σε αυτό το πρώτο επίπεδο;

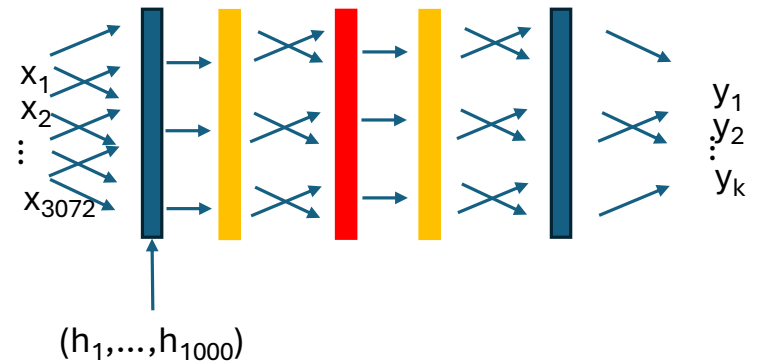


model.fit(X,y)

Η εντολή `model.fit(X,y)` ψάχνει για παραμέτρους που ελαχιστοποιούν την απώλεια / μεγιστοποιούν την ακρίβεια, στα δεδομένα εκπαίδευσης.

Άσκηση: η CIFAR-10 έχει 3.072 features. Εάν το 1^ο επίπεδο είναι fully-connected και έχει 1.000 νευρώνες, πόσες παραμέτρους έχει το νευρωνικό δίκτυο μόνο σε αυτό το πρώτο επίπεδο;

Απάντηση: $3.072 * 1.000 + 1.000 = 3.073.000$





model.fit(X,y)

Η εντολή `model.fit(X,y)` ψάχνει για παραμέτρους που ελαχιστοποιούν την απώλεια / μεγιστοποιούν την ακρίβεια, στα δεδομένα εκπαίδευσης.

Τα νευρωνικά δίκτυα έχουν **πολλές παραμέτρους**, και η εύρεση των καλύτερων παραμέτρων είναι δύσκολο/αδύνατον.


Αντί αυτού, ψάχνουμε **καλύτερες αρκετά καλές** παραμέτρους.

`model.fit(X,y)` → `train(model, data, optimizer, epochs)`



model.fit(X,y)

Η εντολή `model.fit(X,y)` ψάχνει για παραμέτρους που ελαχιστοποιούν την απώλεια / μεγιστοποιούν την ακρίβεια, στα δεδομένα εκπαίδευσης.



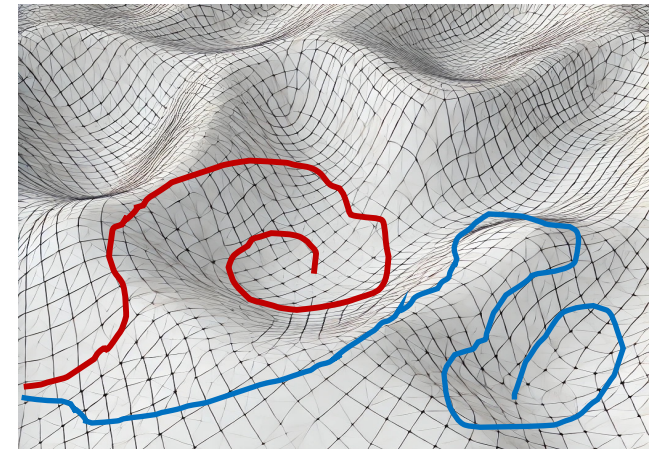
Τα νευρωνικά δίκτυα έχουν **πολλές παραμέτρους**, και η εύρεση των καλύτερων παραμέτρων είναι δύσκολο/αδύνατον.

Αντί αυτού, ψάχνουμε **καλύτερες αρκετά καλές** παραμέτρους.

`model.fit(X,y)` → `train(model, data, optimizer, epochs)`

Πώς ψάχνουμε: `optim.Adam`

πόσο ψάχνουμε



model.predict(x)

model(x)



Πως τα
χρησιμοποιούμε

Νευρωνικά Δίκτυα

1. Ορίζουμε τα επίπεδα: `MyNeuralNetwork`
2. `mymodel = MyNeuralNetwork()` – ορίζουμε την οικογένεια
3. `train(mymodel, data, optimizer, epochs)` – βρίσκουμε παραμέτρους που συμφωνούν με (X, y)
4. `mymodel(x)` – υπολογίζουμε προβλέψεις για τα x