

BOVIFOCR Project

Biometria Ocular, Vivacidade de
Imagens Faciais e Reconhecimento de
Texto (OCR) em Documentos Oficiais

Aula - 06/04/2024

Bernardo Biesseck

Ph.D. student at the PPGInf (UFPR)


Advisor: David Menotti Gomes

TÓPICOS


1. Redes Neurais Convolucionais (CNN)
2. Convolução
3. Multilayer perceptron (MLP)
4. Tutorial em Python

Rede Neural Convolucional (CNN)


Convolution Neural Network (CNN)




Rede Neural Convolucional (CNN)



Rede Neural Convolucional (CNN)





Rede Neural Convolucional (CNN)




Convolução

- Convolution



- https://youtu.be/CXOGvCMLrkA?si=a97_j8xxxknldASU
- <https://youtu.be/f0t-OCG79-U?si=WYMbuxs57PZ6R7FB>
- <https://youtu.be/JboZfxUjLSk?si=HMQ3v22WS8TJLQCL>


Background concepts - MLP



Cross-Entropy/Softmax Loss Function:

$$L = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{w_{y_i}^T x_i}}{\sum_{j=1}^C e^{w_j^T x_i}}$$


Background concepts - MLP



Cross-Entropy/Softmax Loss Function:

$$L = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{w_{y_i}^T x_i}}{\sum_{j=1}^C e^{w_j^T x_i}}$$


Background concepts - MLP



Cross-Entropy/Softmax Loss Function:

$$L = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{w_{y_i}^T x_i}}{\sum_{j=1}^C e^{w_j^T x_i}}$$


Background concepts - MLP



Cross-Entropy/Softmax Loss Function:

$$L = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{w_{y_i}^T x_i}}{\sum_{j=1}^C e^{w_j^T x_i}}$$


Background concepts - MLP



Cross-Entropy/Softmax Loss Function:

$$L = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{w_{y_i}^T x_i}}{\sum_{j=1}^C e^{w_j^T x_i}}$$

Background concepts - MLP




$$\sigma(X, i) = \frac{e^{w_i^T X}}{\sum_{j=1}^C e^{w_j^T X}}$$

Cross-Entropy/Softmax Loss Function:

$$L = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{w_{y_i}^T x_i}}{\sum_{j=1}^C e^{w_j^T x_i}}$$

Background concepts - MLP




$$\sigma(X, i) = \frac{e^{w_i^T X}}{\sum_{j=1}^C e^{w_j^T X}}$$

Cross-Entropy/Softmax Loss Function:

$$L = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{w_{y_i}^T x_i}}{\sum_{j=1}^C e^{w_j^T x_i}}$$

Background concepts - MLP




$$\sigma(X, i) = \frac{e^{w_i^T X}}{\sum_{j=1}^C e^{w_j^T X}}$$

Cross-Entropy/Softmax Loss Function:

$$L = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{w_{y_i}^T x_i}}{\sum_{j=1}^C e^{w_j^T x_i}}$$

Four samples
(mini-batch)

Background concepts - MLP



$$\sigma(X, i) = \frac{e^{w_i^T X}}{\sum_{j=1}^C e^{w_j^T X}}$$


Cross-Entropy/Softmax Loss Function:

$$L = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{w_{y_i}^T x_i}}{\sum_{j=1}^C e^{w_j^T x_i}}$$

$$\frac{\partial L}{\partial W_O} \cdot \frac{\partial L}{\partial W_H}$$

Four samples
(mini-batch)

Background concepts - MLP




- <https://youtu.be/sLsCN9ZL9RI?si=WAI5TuvvptfRjSdZ>

Bibliotecas Python:

- Pytorch - <https://pytorch.org/>
- TensorFlow - <https://www.tensorflow.org/>
- MXNet - <https://mxnet.apache.org/>
- Keras - <https://keras.io/>


Toy example:

- Dataset MNIST (1994)
 - 28x28 pixels (1x784)
 - 60,000 training images
 - 10,000 testing images




Dataset MNIST

<https://www.kaggle.com/datasets/oddrational/mnist-in-csv>



- archive.zip
 - mnist_train.csv
 - mnist_test.csv




	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y
1	label	1x1	1x2	1x3	1x4	1x5	1x6	1x7	1x8	1x9	1x10	1x11	1x12	1x13	1x14	1x15	1x16	1x17	1x18	1x19	1x20	1x21	1x22	1x23	1x24
2	5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
4	4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
5	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
6	9	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
7	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
8	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
9	3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
10	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
11	4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
12	3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
13	5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0


```
import pandas as pd
df = pd.read_csv('../input/mnist_train.csv')
torch_X_train = df.view(-1, 1,28,28).float()
```

Dataset MNIST

<https://www.kaggle.com/datasets/oddrational/mnist-in-csv>




- archive.zip
 - mnist_train.csv
 - mnist_test.csv



	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y
1	label	1x1	1x2	1x3	1x4	1x5	1x6	1x7	1x8	1x9	1x10	1x11	1x12	1x13	1x14	1x15	1x16	1x17	1x18	1x19	1x20	1x21	1x22	1x23	1x24
2	5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
4	4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
5	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
6	9	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
7	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
8	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
9	3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
10	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
11	4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
12	3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
13	5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

```
import pandas as pd
df = pd.read_csv('../input/mnist_train.csv')
torch_X_train = df.view(-1, 1,28,28).float()
```






Dataset MNIST

Pytorch tutorial

- <https://www.kaggle.com/code/sdelecourt/cnn-with-pytorch-for-mnist>

Install dependencies

pip3 install numpy pandas scikit-learn torch



```
import torch.nn as nn
```

```
import torch.nn.functional as F
```

```
class CNN(nn.Module):
```

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

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

```
        self.conv1 = nn.Conv2d(1, 32, kernel_size=5)
```

```
        self.conv2 = nn.Conv2d(32, 32, kernel_size=5)
```

```
        self.conv3 = nn.Conv2d(32, 64, kernel_size=5)
```

```
        self.fc1 = nn.Linear(3*3*64, 256)
```

```
        self.fc2 = nn.Linear(256, 10)
```

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

```
        x = F.relu(self.conv1(x))
```

```
        x = F.relu(F.max_pool2d(self.conv2(x), 2))
```

```
        x = F.relu(F.max_pool2d(self.conv3(x), 2))
```

```
        x = F.dropout(x, p=0.5, training=self.training)
```

```
        x = x.view(-1, 3*3*64)
```

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

```
        x = F.dropout(x, training=self.training)
```

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

```
        return F.log_softmax(x, dim=1)
```

Dataset MNIST

Keras tutorial

- <https://www.kaggle.com/code/tobikaggle/keras-mnist-cnn-learning-curve>

Install dependencies

pip3 install numpy pandas scikit-learn keras matplotlib