


Deep learning without topology


Pingyao

Feb, 29, 2024


Image data



Size: (369, 496, 3)



voiced



voiceless

CNN architecture

```
model = Sequential()
model.add(Conv2D(64, (3,3), input_shape = x.shape[1:]))
model.add(Activation("relu"))
model.add(MaxPooling2D(pool_size=(2,2)))
model.add(Conv2D(128, (3,3)))
model.add(Activation("relu"))
model.add(MaxPooling2D(pool_size=(2,2)))
model.add(Conv2D(256, (3,3)))
model.add(Activation("relu"))
model.add(MaxPooling2D(pool_size=(2,2)))
model.add(Flatten())
model.add(Dense(64))
model.add(Dense(1))
model.add(Activation('sigmoid'))
model.compile(loss='binary_crossentropy',
              optimizer="adam",
              metrics=['accuracy'])
model.fit(x_train, y_train, batch_size=64, epochs=25)
```

```
Epoch 1/25
18/18 [=====] - 128s 7s/step - loss: 2.7118 - accuracy: 0.5201
Epoch 2/25
18/18 [=====] - 121s 7s/step - loss: 0.6794 - accuracy: 0.5852
Epoch 3/25
18/18 [=====] - 120s 7s/step - loss: 0.6197 - accuracy: 0.6637
Epoch 4/25
18/18 [=====] - 120s 7s/step - loss: 0.5740 - accuracy: 0.6896
Epoch 5/25
18/18 [=====] - 120s 7s/step - loss: 0.5356 - accuracy: 0.7217
Epoch 6/25
18/18 [=====] - 119s 7s/step - loss: 0.5065 - accuracy: 0.7538
Epoch 7/25
18/18 [=====] - 119s 7s/step - loss: 0.4606 - accuracy: 0.7690
Epoch 8/25
18/18 [=====] - 119s 7s/step - loss: 0.4055 - accuracy: 0.7984
Epoch 9/25
18/18 [=====] - 118s 7s/step - loss: 0.3618 - accuracy: 0.8323
Epoch 10/25
18/18 [=====] - 119s 7s/step - loss: 0.3305 - accuracy: 0.8412
Epoch 11/25
18/18 [=====] - 118s 7s/step - loss: 0.2854 - accuracy: 0.8814
Epoch 12/25
18/18 [=====] - 119s 7s/step - loss: 0.2770 - accuracy: 0.8742
Epoch 13/25
...
Epoch 24/25
18/18 [=====] - 124s 7s/step - loss: 0.1181 - accuracy: 0.9500
Epoch 25/25
18/18 [=====] - 124s 7s/step - loss: 0.1037 - accuracy: 0.9581
```

Result

```
from sklearn.metrics import auc
auc_keras = auc(fpr_keras, tpr_keras)
auc_keras
```

```
0.8956711414103158
```


```
y_pred = (y_pred >= 0.5).astype(int)
print(accuracy_score(y, y_pred))
```

```
0.8508114856429463
```


Model: "sequential"

Layer (type)	Output Shape	Param #
conv2d (Conv2D)	(None, 367, 494, 64)	1792
activation (Activation)	(None, 367, 494, 64)	0
max_pooling2d (MaxPooling2D)	(None, 183, 247, 64)	0
conv2d_1 (Conv2D)	(None, 181, 245, 128)	73856
activation_1 (Activation)	(None, 181, 245, 128)	0
max_pooling2d_1 (MaxPooling2D)	(None, 90, 122, 128)	0
conv2d_2 (Conv2D)	(None, 88, 120, 256)	295168
activation_2 (Activation)	(None, 88, 120, 256)	0
max_pooling2d_2 (MaxPooling2D)	(None, 44, 60, 256)	0
...		
Total params: 43,624,705		
Trainable params: 43,624,705		
Non-trainable params: 0		

Kernels in CNN



PCA



Persistence diagram