
Intro to Keras

Hello world!

Overview

- How to build your own Neural Network?
- Keras Libraries.
- First example.
- Exercise: Recognition of handwritten digits.

How to build your own NN?

Input X & Output Y

nature (fixed, sequential, ..),
type, shape

Architecture

Type of layers, # of layers,
kernels size, ...

The task

regression, classification, ...,
and therefore, your loss
function.

Tune hyperparameters

Learning rate, batch size, # of
epochs, the loss function...

Specify Input (X) & Output (Y)

- Input:
 - Vector,
 - n-D matrix,
 - sequential data,
 - Multimodal input, ...
- Output:
 - discrete scalar,
 - vector,
 - n-D matrix,
 - sequential output, ...

Define the task

Classification predictive modeling is the task of approximating a mapping function (f) from input variables (X) to discrete output variables (y). The output variables are often called labels or categories. The mapping function predicts the class or category for a given observation.

- A classification problem requires that examples be classified into one of two or more classes.
- A classification can have real-valued or discrete input variables.
- A problem with two classes is often called a two-class or **binary classification problem**.
- A problem with more than two classes is often called a **multi-class classification problem**.
- A problem where an example is assigned multiple classes is called a **multi-label classification problem**.


There are many ways to evaluate a classification predictive model, but perhaps the most common is to calculate the classification accuracy.

Define the task

Regression predictive modeling is the task of approximating a mapping function (f) from input variables (X) to a continuous output variable (y). A continuous output variable is a real-value, such as an integer or floating point value. These are often quantities, such as amounts and sizes.

- A regression problem requires the prediction of a quantity.
- A regression can have real valued or discrete input variables.
- A problem with multiple input variables is often called a **multivariate regression problem**.
- A regression problem where input variables are ordered by time is called a **time series forecasting problem**.

A regression predictive model predicts a quantity, therefore to evaluate the model we report an error in those predictions.



Keras

Keras is a high-level neural networks API, written in Python and capable of running on top of TensorFlow, CNTK, or Theano. It was developed with a focus on enabling fast experimentation.

- Allows for easy and fast prototyping (through user friendliness, modularity, and extensibility).
- Supports both **convolutional networks** and **recurrent networks**, as well as combinations of the two.

[keras.io]

Sequential Model

The **Sequential** model is a linear stack of layers.

You can create a **Sequential** model by passing a list of layer instances to the constructor:

```
from keras.models import Sequential
from keras.layers import Dense


model = Sequential([
    Dense(32, input_shape=(784,), activation='relu'),
    Dense(10, activation='softmax')
])
```

```
from keras.models import Sequential
model = Sequential()
```


from keras import layers

- Core layers:
 - Dense layer: fully connected layer

Fully-connected layer



```
model.add(Dense(4, activation='softmax'))
```

from keras import layers

- Convolutional layers:
 - Conv1D, Conv2D, Conv3D
 - UpSampling1D, UpSampling2D...

1 _{x1}	1 _{x0}	1 _{x1}	0	0
0 _{x0}	1 _{x1}	1 _{x0}	1	0
0 _{x1}	0 _{x0}	1 _{x1}	1	1
0	0	1	1	0
0	1	1	0	0

Image


4		

Convolved
Feature

```
model.add(Conv2D(filters=5, kernel_size=(3,3), activation='sigmoid'))
```

from keras import layers

- Pooling Layers:
 - MaxPooling1D, MaxPooling2D, ...
 - AveragePooling1D, AveragePooling2D, ...



Convolved
feature



Pooled
feature


```
model.add(MaxPooling2D(pool_size=(8,8)))
```

from keras.layers import activations

Activation layers in neural networks, takes a value that is passed through a function which *squashes* the value into a range.


Sigmoid

$$\sigma(x) = \frac{1}{1+e^{-x}}$$




tanh

$$\tanh(x)$$




ReLU

$$\max(0, x)$$



Leaky ReLU

$$\max(0.1x, x)$$




Maxout

$$\max(w_1^T x + b_1, w_2^T x + b_2)$$

ELU

$$\begin{cases} x & x \geq 0 \\ \alpha(e^x - 1) & x < 0 \end{cases}$$




from keras.layers import activations

Softmax

- calculates the probabilities of each target class over all possible target classes.
- is often used in the final layer of a neural network-based classifier.

$$\sigma(z)_j = \frac{e^{z_j}}{\sum_{k=1}^K e^{z_k}}$$




```
model.add(Activation('softmax'))
```

Sigmoid

- returns a real-valued output.
- is often used as the activation function for artificial neurons.

$$\text{sigmoid}(x) = \frac{1}{1 + e^{-x}}$$



```
model.add(Activation('sigmoid'))
```

from keras import losses

A **loss function** or **cost function** is a function that maps values of one or more variables onto a real number intuitively representing some associated "cost". An optimization problem seeks to minimize a loss function.

The loss function lets us quantify the quality of any particular set of parameters (weights **W** and biases **B**). Some available loss functions:

- `mean_squared_error`
- `mean_absolute_error`
- ...
- `categorical_crossentropy`
- `binary_crossentropy`
- ...

from keras import optimizers

The goal of optimization is to find the set of parameters (**W** and **B**) that minimizes the loss function.

Some available optimizer:

- **SGD** #Stochastic gradient descent optimizer.
- **Adagrad** #Adaptive gradient descent optimizer
- **Adadelta** #Adaptive learning rate optimizer
- **Adam**
- ...

from keras import applications

Keras Applications are deep learning models that are made available alongside pre-trained weights. These models can be used for prediction, feature extraction, and fine-tuning.


- Xception
- VGG16
- VGG19
- ResNet50
- InceptionV3
- InceptionResNetV2
- MobileNet
- DenseNet
- NASNet

```
from keras.applications.vgg16 import VGG16

model = VGG16(weights='imagenet', include_top=True)
```


from `keras import data augmentation`

Image Augmentation is the process of taking images that are already in a training dataset and manipulating them to create many altered versions of the same image.



from keras import data augmentation

Generate batches of tensor image data with real-time data augmentation. The data will be looped over (in batches) indefinitely.

```
datagen = ImageDataGenerator(  
    featurewise_std_normalization=True,  
    rotation_range=20,  
    width_shift_range=0.2,  
    height_shift_range=0.2,  
    horizontal_flip=True)  
  
datagen.fit(x_train)  
  
for x, y in datagen.flow(x_train, y_train, batch_size=32):  
    ...
```

from keras import wrap-up!

```
model = Sequential()
model.add(Dense(32, input_dim=784))
model.add(Activation(sigmoid))

opt = Adam(lr=0.01)
model.compile(loss='categorical_crossentropy',
              optimizer=opt,
              metrics=['accuracy'])

model.fit(x_train, y_train,
          epochs=20,
          batch_size=32)
score = model.evaluate(x_test, y_test, batch_size=32)
```

Exercise: handwritten digits

label = 5



label = 0



label = 4



label = 1



label = 9



label = 2



label = 1



label = 3



label = 1



label = 4



label = 3



label = 5



label = 3



label = 6



label = 1



label = 7



label = 2



label = 8




label = 6




label = 9



Exercise: handwritten digits



Exercise: handwritten digits



Exercise: handwritten digits

