Quiz 2:

Introduction to Supervised Learning

Email address *

mikesymmonds@gmail.com

Please enter your name: *

Michael Symmonds

Linear Regression/Logistic Regression

Which model aims to fit the best line based on the following data ? 1 point

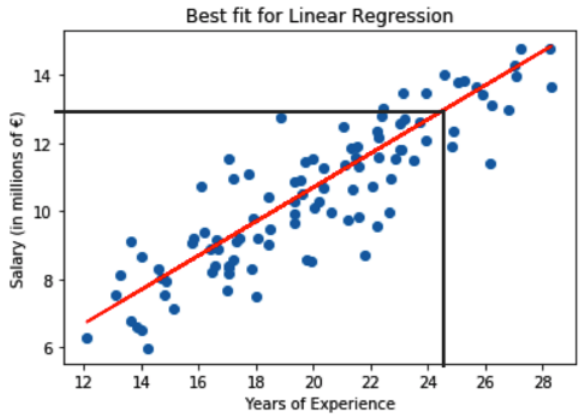

- ○ Logistic Regression
- ◉ Linear Regression
- ○ Hidden Markov Model

What model is summarized as follows? 1 point

$$\forall i \in \{1, \ldots, N\} \quad Y_i | X_i = x_i \sim \mathcal{B}(\sigma(w^T x_i))$$

- ○ Bernoulli model
- ◉ Logistic Regression
- ○ Linear Regression

What algorithm should be used to learn the parameters of Linear Regression and Logistic Regression? 1 point

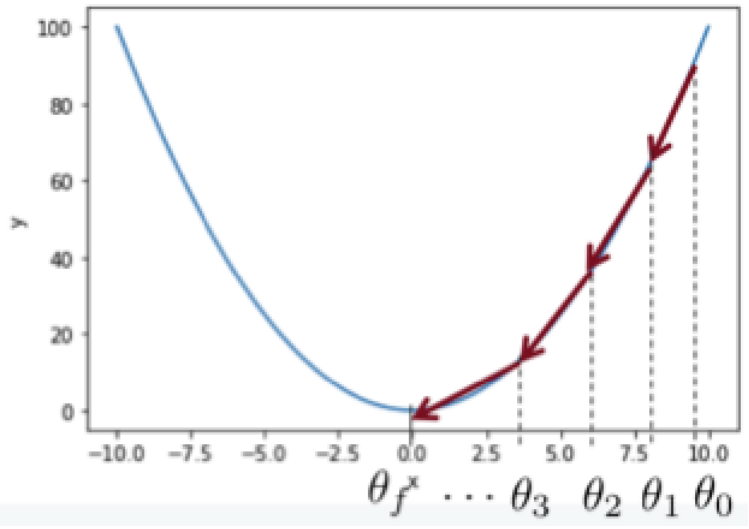

Gradient Descent

What are the two hyperparameters that should be chosen before applying the Gradient Descent algorithm? 2 points

The learning rate (eta)
The number of iterations of gradient descent until convergence

Decision Trees Algorithm

We want to predict whether someone is going to succeed or fail in a Machine Learning Interview based on the following features: "years of experience", "coding skills" (with discrete values in [0, 5]), "familiar with the models" (with discrete values in [0, 5]), and "like chocolate"(with binary output 0/1). We obtain the following graph of decision

Graph of decision

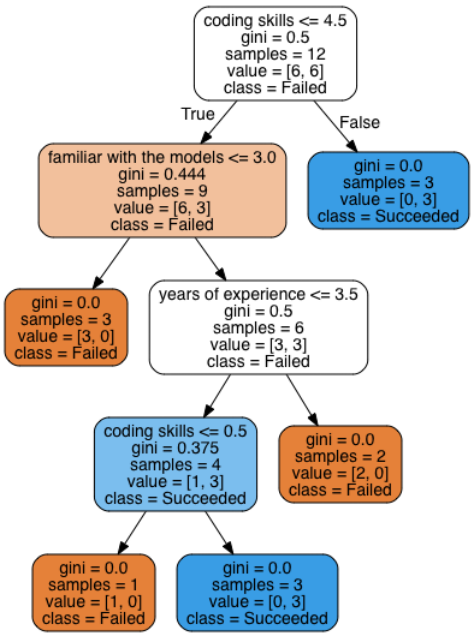

Quiz 2:

How many candidates have succeeded? 1 point

- [ ] 4
- [ ] 5
- [x] 6

How many candidates have failed? 1 point

- [ ] 2
- [ ] 4
- [x] 6

If a candidate has the following characteristics: 2 years of experience, 0 for coding skills, and 4 for the familiarity with the models. What would the algorithm predict? 1 point

- [ ] Successed
- [x] Failed

What is the minimum value of "coding skills" that can change the prediction value in the previous example? 1 point

- [x] 1
- [ ] 2
- [ ] 3

Give one of the hyperparameters for the Decision Trees Algorithm and one for the Random Forest Algorithm? 1 point

Decision Tree Algorithm - Depth of the tree (and the impurity measure)
Random Forest Algorithm - The number of trees K to include in the forest (and the sample size n of the bootstrap sampling & the number of attributes d at each split)

Programming Session

Did you understand the problem?

- [x] Yes
- [ ] No

Google