

CS158 - Assignment 1

Data

Due: Friday, August 29 at 5pm



<https://xkcd.com/1429/>

0. Slack

If you are not a member of the `cs158-fall12025` slack channel, message me and I'll add you.

1. You know the drill

Read through the administrative handout on the course web page.

- (a) *What is the late policy for the course?*
- (b) *What is the late policy on using ChatGPT (and similar tools) for the course?*

2. Data, data, data

There are now lots of really interesting data sets publicly available to play with. They range in size, quality and the type of features and have resulted in many new machine learning techniques being developed.

Find a public, free, supervised (i.e. it must have features *and* labels), machine learning dataset. You may NOT list a data set from 1) The UCI Machine Learning Repository or 2) from Kaggle.com. Once you've found the data set, provide the following information:

- (a) The name of the data set.
- (b) Where the data can be obtained.

- (c) A brief (i.e. 1-2 sentences) description of the data set including what the features are and what is being predicted.
- (d) The number of examples in the data set.
- (e) The number of features for each example. If this isn't concrete (i.e. it's text), then a short description of the features.

Extra credit will be given for particularly interesting data sets, e.g. the most unique, the data set with the largest number of examples and the data set with the largest number of features.

3. Data analysis

One of the first things to do before trying any formal machine learning technique is to dive into the data. This can include looking for funny values in the data, looking for outliers, looking at the range of feature values, what features seem important, etc. At:

`http://www.cs.pomona.edu/classes/cs158/assignments/assign1/titanic-train.csv`

I have provided a modified version of passenger survival data for the Titanic¹.

This data set has six binary features:

- **First_class** (whether the passenger was in first class or not)
- **Sex** (0 = Male, 1 = Female)²
- **Age** (0 = <25, 1 = 25+)
- **SibSp** (had siblings/spouses aboard?)
- **ParCh** (had parents/children aboard?)
- **Embarked** (Left from Southhampton?)

Based on these features, the Titanic task is to learn to predict the last column, whether or not the passenger survived (1 = survived).

- (a) For each of the features calculate (and write down) the *training error* if you used **only** that feature to classify the data. To do this you will need to do the following for each feature:
 - Split the data based on that feature. Call bin_0 all examples that have 0 for that feature and bin_1 all examples that have 1 for that feature.
 - Calculate the majority count for the label in each bin, i.e. for bin_0 ,

$$majority(bin_0) = \max(count(bin_0 = survive), count(bin_0 = notsurvive))$$

This value is how many examples you would get right in bin_0 if you split on that feature. Make sure you understand why!

¹The original data can be found at: <http://www.kaggle.com/c/titanic-gettingStarted>

²I recognize that there are diverse gender identities. In this case, we are limited by the data and will focus on sex as a binary feature.

- Calculate the training error for that feature. The accuracy on the training set (i.e. percentage correct) can be calculate as:

$$accuracy = \frac{majority(bin_0) + majority(bin_1)}{totalNumberOfTrainingExamples}$$

and then

$$error = 1 - accuracy$$

You can either write a program to do this in any language you'd like (you don't need to submit the code) or you could also do this in a spreadsheet program like excel. Your answer to this problem should be 6 error training error rates, one for each feature.

- (b) Which feature would be the best to use? Put another way, if we were building a 1-level decision tree using Algorithm 1 from the book, which feature would it pick?
- (c) Do you agree that this is the best choice to make? Just 1-2 sentences explaining yes/no is sufficient.

4. Ethics and ML

As machine learning becomes more commonplace, it is increasingly important to think about the ethical implications, including when and how we apply the models, data collection practices, and biases that can be introduced into the models from the data. Find one article that discusses machine learning ethics:

- (a) What is the article (i.e., title and url)?
- (b) Give a couple of sentence sentence summary of the article?
- (c) Are there other ethical concerns that you can think of related to the article topic that weren't mentioned?

When you're done

Put your answers to the questions above in either a `.txt` or `.pdf` file *with your name at the top of the file* and named *lastname.1* (with the corresponding file extension).

Submit your assignment via gradescope.