

First and second days: Homework 1

Data Analysis and Machine Learning

May 19, 2020

Day one and two exercises

Exercise 1. The first exercise here is of a mere technical art. We want you to have

- git as a version control software and to establish a user account on a provider like GitHub. Other providers like GitLab etc are equally fine.
- Install various Python packages

We will make extensive use of Python as programming language and its myriad of available libraries. You will find IPython/Jupyter notebooks invaluable in your work. You can run **R** codes in the Jupyter/IPython notebooks, with the immediate benefit of visualizing your data. You can also use compiled languages like C++, Rust, Fortran etc if you prefer. The focus in these lectures will be on Python.

If you have Python installed (we recommend Python3) and you feel pretty familiar with installing different packages, we recommend that you install the following Python packages via **pip** as

1. pip install numpy scipy matplotlib ipython scikit-learn sympy pandas pillow

For **Tensorflow**, we recommend following the instructions in the text of Aurelien Geron, [Hands-On Machine Learning with Scikit-Learn and TensorFlow](#), O'Reilly

We will come back to **tensorflow** later.

For Python3, replace **pip** with **pip3**.

For OSX users we recommend, after having installed Xcode, to install **brew**. Brew allows for a seamless installation of additional software via for example

1. brew install python3

For Linux users, with its variety of distributions like for example the widely popular Ubuntu distribution, you can use **pip** as well and simply install Python as

1. `sudo apt-get install python3` (or `python` for Python2.7)

If you don't want to perform these operations separately and venture into the hassle of exploring how to set up dependencies and paths, we recommend two widely used distributions which set up all relevant dependencies for Python, namely

- [Anaconda](#),

which is an open source distribution of the Python and R programming languages for large-scale data processing, predictive analytics, and scientific computing, that aims to simplify package management and deployment. Package versions are managed by the package management system **conda**.

- [Enthought canopy](#)

is a Python distribution for scientific and analytic computing distribution and analysis environment, available for free and under a commercial license.

We recommend using **Anaconda**.

Exercise 2: Python getting started. This exercise has as its aim to write a small program which reads in data from a **csv** file on the equation of state for dense nuclear matter. The file is localized at <https://github.com/mhjensen/MachineLearningMSU-FRIB2020/blob/master/doc/pub/Regression/ipybn/datafiles/EoS.csv>. Thereafter you will have to set up the design matrix \mathbf{X} for the n datapoints and a polynomial of degree 3. The steps are:

- Write a Python code which reads the in the above mentioned file.
- Use for example **pandas** to order your data and find out how many data points there are.
- Set thereafter up the design matrix with dimensionality $n \times p$ where $p = 4$ and where you have defined a polynomial of degree $p - 1 = 3$. Print the matrix and check that the numbers are correct.

We recommend looking at the examples in the [regression slides](#).

Exercise 3. We will generate our own dataset for a function $y(x)$ where $x \in [0, 1]$ and defined by random numbers computed with the uniform distribution. The function y is a quadratic polynomial in x with added stochastic noise according to the normal distribution $\mathcal{N}(t, \infty)$. The following simple Python instructions define our x and y values (with 100 data points). `x = np.random.rand(100,1)` `y = 5*x*x+0.1*np.random.randn(100,1)`

1. Write your own code (following the examples under the [regression slides](#)) for computing the parametrization of the data set fitting a second-order polynomial.

2. Use thereafter **scikit-learn** (see again the examples in the regression slides) and compare with your own code.
3. Using scikit-learn, compute also the mean square error, a risk metric corresponding to the expected value of the squared (quadratic) error defined as

$$MSE(\hat{y}, \tilde{y}) = \frac{1}{n} \sum_{i=0}^{n-1} (y_i - \tilde{y}_i)^2,$$

and the R^2 score function. If \tilde{y}_i is the predicted value of the i -th sample and y_i is the corresponding true value, then the score R^2 is defined as

$$R^2(\hat{y}, \tilde{y}) = 1 - \frac{\sum_{i=0}^{n-1} (y_i - \tilde{y}_i)^2}{\sum_{i=0}^{n-1} (y_i - \bar{y})^2},$$

where we have defined the mean value of \hat{y} as

$$\bar{y} = \frac{1}{n} \sum_{i=0}^{n-1} y_i.$$

You can use the functionality included in scikit-learn. If you feel for it, you can use your own program and define functions which compute the above two functions. Discuss the meaning of these results. Try also to vary the coefficient in front of the added stochastic noise term and discuss the quality of the fits.

Exercise 4, mean values and variances in linear regression. This exercise deals with various mean values and variances in linear regression method (here it may be useful to look up chapter 3, equation (3.8) of [Trevor Hastie, Robert Tibshirani, Jerome H. Friedman, The Elements of Statistical Learning, Springer](#)).

The assumption we have made is that there exists a function $f(\mathbf{x})$ and a normal distributed error $\varepsilon \sim \mathcal{N}(0, \sigma^2)$ which describes our data

$$\mathbf{y} = f(\mathbf{x}) + \varepsilon$$

We then approximate this function with our model from the solution of the linear regression equations (ordinary least squares OLS), that is our function f is approximated by $\tilde{\mathbf{y}}$ where we minimized $(\mathbf{y} - \tilde{\mathbf{y}})^2$, with

$$\tilde{\mathbf{y}} = \mathbf{X}\boldsymbol{\beta}.$$

The matrix \mathbf{X} is the so-called design matrix.

Show that the expectation value of \mathbf{y} for a given element i

$$\mathbb{E}(y_i) = \mathbf{X}_{i,*} \boldsymbol{\beta},$$

and that its variance is

$$\text{Var}(y_i) = \sigma^2.$$

Hence, $y_i \sim \mathcal{N}(\mathbf{X}_{i,*}\boldsymbol{\beta}, \sigma^2)$, that is \mathbf{y} follows a normal distribution with mean value $\mathbf{X}\boldsymbol{\beta}$ and variance σ^2 .

With the OLS expressions for the parameters $\boldsymbol{\beta}$ show that

$$\mathbb{E}(\boldsymbol{\beta}) = \boldsymbol{\beta}.$$

This means that the estimator of the regression parameters is unbiased.

Show finally that the variance of $\boldsymbol{\beta}$ is

$$\text{Var}(\boldsymbol{\beta}) = \sigma^2 (\mathbf{X}^T \mathbf{X})^{-1}.$$

Exercise 5. Finally, try now to write your own code (you can use the example the nuclear masses in the lecture slides on Regression and Getting started, see <https://compphysics.github.io/MLErasmus/doc/pub/Regression/html/Regression-bs.html> and <https://compphysics.github.io/MLErasmus/doc/pub/How2ReadData/html/How2ReadData-bs.html>) that reads in the nuclear masses and compute the proton separation energies, the two-neutron and two-proton separation energies and finally the shell gaps for selected nuclei.

Finally, try to compute the Q -values for β^- decay for selected nuclei.

We will use this code later as a starting point for our discussions on linear regression.