



MTH786P Machine Learning with Python, Semester A, 2023/24

Final project late summer resit

1st examiner Dr. N. Perra, 2nd examiner Dr. O. Bobrowski

For the final project you will have to:

1. Select **one** of the five datasets provided on the QM+ module page and described below;
2. Analyse the dataset by studying the statistical characteristics of features and target variables. Plot and discuss their frequency/distributions and correlations. You are allowed to use Pandas, seaborn, and other libraries for plotting purposes;
3. Apply the relevant algorithms discussed during the module to solve a prediction task;
4. Write a detailed report (written preferably in $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$) of no more than 8 pages length structured as described below;
5. Submit a zipped folder with: i) the notebook you wrote with all the codes and results, ii) the data needed to run the codes. The name of the folder should be your student ID.

While you are allowed to use different libraries for visualisation purposes (e.g., seaborn and pandas), you can only use NumPy to implement your algorithms/models; tools from libraries such as SciPy/Scikit-learn, etc. are not allowed.

Datasets description (select one)

You will need to select **one** of the following datasets:

Recognition of handwritten digits: MNIST dataset

If you select this dataset, the task is to implement, describe and present (multi-class) classification models of your choice for the recognition of handwritten digits. More precisely, the goal is to implement (multi-class) classifiers that take an image of size 28×28 of a handwritten digit in the range from zero to nine as their input and automatically determine which class the input image most likely belongs to. The dataset to be used is the so-called MNIST dataset and will be provided on the QM+ module page.

Fires prediction: Forest fires dataset

If you select this dataset, the task is to implement, describe and present regression models of your choice to predict the size of the burnt area of forest fires. The goal is to develop regression models using any number of the variables provided, which describe fires' features, to predict the size of the burnt area. The dataset will be provided on the QM+ module page.

Wine quality prediction: Wine dataset

If you select this dataset, the task is to implement, describe and present regression and/or classification models of your choice to predict the quality of red wines given a range of their features. The goal is to develop regression and/or classification models using any number of the variables provided, which describe wines' features, to predict their quality (measured as a score from 0 to 10 based on sensory data from three experts). The dataset will be provided on the QM+ module page.

Booking cancellation prediction: Hotel dataset

If you select this dataset, the task is to implement, describe and present (binary) regression/classification models of your choice to predict which customers will cancel an hotel reservation. The goal is to develop (binary) regression/classification models using any number of the variables provided, which describe reservations' features, to predict whether they will be cancelled. The dataset will be provided on the QM+ module page.

Music taste prediction: Spotify dataset

If you select this dataset, the task is to implement, describe and present (binary) regression/classification models of your choice to predict the musical taste of a given Spotify user. The goal is to develop (binary) regression/classification models using any number of the variables provided, which describe songs' features, to predict whether the Spotify user will like them or not. The dataset will be provided on the QM+ module page.

Structure of the written report

You conclude the assessment by writing a report. The report should be no longer than eight pages and be written preferably in L^AT_EX with your favourite editor. If no editor is at hand, please feel free to use online editors such as Overleaf. The structure of the written report should be the following:

1. Title;
2. Your name and ID;
3. Introduction and problem statement;

4. Analysis of the dataset. Dive into the features and targets. Study their key statistical characteristics by plotting and discussing their frequency/distributions and correlations;
5. Methods. Describe in some details the algorithms/models you will be using. Use the notation and definitions discussed in the lectures and lecture notes;
6. Results of the prediction task. Present your results, compare the performance of the models implemented, and discuss the best you found;
7. Conclusions. Provide a summary of the problem you tackled and reflect on the results you found.
8. References. Reference i) the data you use, ii) the external resources like websites, articles, you might find and use, iii) the methods and approaches used. These do not count against the 8 page limits.

You need to submit also a zipped folder, via a dedicated submission link, with: the notebook you wrote with all the codes and results (name the folder as your student ID). This is key to show the work done. Failure to submit the code will affect the final mark.

We refer to the following table for marking guidance on the final project. Note how the choice of dataset by does not impact your final grade.

Approximate weight	Description
30 %	Overall quality of the project work (as described in the report), taking into account its difficulty.
30 %	Quality/relevance of the methodology and choice of resources, the application of these resources, explanation/exposition and correctness of mathematical notation and computer programming and analysis of results.
30 %	Quality of presentation: Logical structure of the document, clarity and coherence of exposition, correct use of English (punctuation, spelling and grammar) and precise mathematical writing, layout and style, sensible use of sections and subsections, lack of typing mistakes, choice of sensible (and standard) notation and other conventions, sensible use of equation numbering, appropriate use of figures, tables, charts, diagrams, etc.
10 %	Initiative and ambition (as described in the report): How far did you go, or attempt to go? How many models did you try/implement?