

Compulsory Exercise 2

TMA4268 Statistical Learning V2026

The submission deadline is: **Sunday the 12th of April, 23:59h**

Introduction

In this project, you will be working on analyzing a data set by developing *prediction/classification* models using statistical learning techniques.

The goal of this relatively open project is to give you hands-on experience with the methods learned in the course, and to help you develop your skills in data pre-processing, model selection, and evaluation.

To get started, here is a broad guideline for your project:

1. **Choose a data set for your project.** We list two possible data sets below, or you can find one on your own.
 - **Heart Failure Prediction Dataset:** [The data set](#) is provided by user *Fedesoriano* on Kaggle. More information can be found in the provided link.
 - **AirBnB Prices in European Cities:** [The data set](#) is provided by user *The Devastator* on Kaggle. More information can be found in the provided link. Note that the description says that the dataset is designed for inference, but we will use it for prediction or classification. You can choose to focus on data from one city, or combine data from different cities.
 - **Choosing your own data set:** If you choose to use another dataset, ensure that that it is diverse enough and contains enough data points to train a good model. There are many example data sets provided by R packages (for example `carData`), and you can find a variety of data sets on [Kaggle](#) or [TidyTuesday](#). Try searching for a topic that you find interesting and would like to work with. To find a “well-organized” data on Kaggle set we recommend choosing a data set with sufficiently-large amount of upvotes. Otherwise, a data set could be unorganized, poorly structured, or contain a lot of missing values. We do not want you to spend much time on data cleaning!
2. **Decide on a prediction/classification task.** What are you trying to predict/classify with the data set? Decide on this *before* you try out any of the models - it’s fine if it turns out in the end that you are not able to get good predictions, as long as you have done a sound, meaningful analysis. Some ideas are provided here, but in principle you can decide on anything that makes sense.
 - For the heart failure data set you could try to classify whether a patient gets heart disease or not, predict their age, or predict their cholesterol.
 - For the AirBnB price data set, you could try to predict the price or the rating of of accommodation, or classify the room type or “superhost” status of the host.
 - If you chose another data set you should come up with a prediction or classification task that makes sense from the data.
3. **Data pre-processing.** Before building your model, you might need to pre-process the data (data wrangling), depending on what format your methods require.
4. **Choose appropriate models and methods.** You should use methods that you have learned about in the course. Make sure to justify your choice of algorithm, check model assumptions and, if relevant, tune the hyperparameters to improve model performance. **Use methods from at least two different**

modules of the course. You can also consider transformations of the variables in your data set, and interaction effects. Make sure that you do the model selection in a *valid* way, as you have learned in the course, and remember the bias-variance tradeoff.

- If you are solving a prediction task, you could try methods such as multiple linear regression, GAMs, ridge/lasso regression, trees, random forests, boosting *etc.*
 - If you are solving a classification task, you could try methods such as logistic regression, LDA, QDA, KNN, regression trees, random forests and boosting *etc.*
5. **Model assessment.** Evaluate your model's performance using appropriate metrics or evaluation tools, such as accuracy, MSE, sensitivity, specificity, *etc.* Again, make sure that you do this in a *valid* way, as you have learned in the course.
 6. **Reporting.** Present your results and state your findings and interpretation of the results.

We hope you find this project both challenging and rewarding. Best of luck!

Supervision

We will use the times where we would have lectures and exercises for supervision.

Supervision hours (in the usual lecture room):

- Monday the 23rd of March, 08:15-10:00, S6
- Tuesday the 23rd of March, 10:15-10:00, KJL22
- Wednesday the 24th of March, 10:15-12:00, KJL5

If you have questions outside of these hours, feel free to send me an email (simen.k.furset@ntnu.u.no), or come knock on my door (SB2, 1004).

Practical issues (Please read carefully)

- Remember to write your names and group number on top of your submission file.
- The exercise should be handed in as two files: **one pdf file with your main report and a supplementary file with your code (.R or .py)** We will read the pdf-file and use the code file in case we need to check details in your submission.
- Do not include the text from the file that you are reading now. We want your (relevant) R code, plots and written solutions - if you are using R-markdown, you can use the attached template `Compulsory2_template.Rmd`.
- Any math equations and/or computations should be formatted using TeX or similar. No scans or images of handwritten computations!
- Please **no more than 14 pages** in your pdf-file. **We will stop reading your report after page 14.** Keep this in mind when choosing what R-code/output to include, and when sizing your figures.
- Do not submit a word-file or a zip-file.
- **Use of AI tools:** make sure you have read and understood NTNU's guidelines on the use of such tools, found [here](#). Note especially the part about submitting an AI declaration, which can be submitted as an additional document when you hand in your report. As a rule of thumb, we expect you to show that you understand what you are doing when you analyze the data for your project. Points obviously cannot be given for automatically generated reports.
- **Bonus hint: Neat reports are easier to understand and may result in a better grade - simply because we cannot give full points if things are unclear, ambiguous or messy.** Pretend that the task you decided on was given to you by your boss at a company, and the report is what you will deliver to the boss. When writing the report, keep in mind that boss has limited time and attention, and has not spent as much time as you have on getting familiar with the problem.

If you follow the template below you should be able to produce a reasonable (i.e. passable) project. Feel free to deviate from the template if it makes sense for your project; if you do make sure your project has the same ‘scope’ as outlined in the template.

Abstract (max. 350 words)

The purpose of the abstract is to give a short and concise summary of your project. It is a stand-alone text that is given before the actual report starts. It includes the following components:

1. Begin your abstract by clearly stating the purpose of your project. What problem are you trying to solve? What question do you want to answer? It is important to be concise and to the point.
2. In the next few sentences, describe the data and methods you used to conduct your study. What kind of data did you use? How did you analyze it? What tools, techniques, or methods did you use? Be specific, but avoid going into too much detail.
3. Summarize your key findings: In the main part of your abstract, summarize the most important results of your project and interpret them briefly (i.e., what do your results mean?). Highlight the most significant findings, and provide enough detail to give the reader a sense of what you discovered.
4. (optional) Emphasize the significance of your results: Explain why and/or how your findings are important (or not important). Highlight any novel or unexpected findings, and explain how they add to our understanding of the topic.

Introduction: Scope and purpose of your project

- Briefly introduce the broad idea of the problem or task that you chose and the respective data set that you use. This could be a classification task (e.g., predicting whether a patient gets heart disease or not) or a prediction task (e.g., predicting the price of an AirBnB). Clearly define the scope of your project. What specific problem are you trying to solve?
- Describe the source and give a reference to where the data set is coming from.
- Describe the purpose of your project in more detail. What are the specific question that you want to answer in your project? Are you trying to find the best performing method or a good performing and light method that is easy to use? Who is your audience? Are you trying to discover the relations between different variables? Are you trying to find important predictors for your classification? Are you trying to draw some insightful understanding in a particular topic/domain?

Descriptive data analysis/statistics

Conduct descriptive data analysis to get an overview over your data (see [this example](#) for inspiration). Try to **focus on what will be relevant for your modelling** and use common sense. For example, too much detail, or figures without any explanation or axis labels, are not useful to the reader. Be selective, and focus on what is most important.

For example, you could:

- Report measures such as mean, median, range, standard deviation, and variance to describe the central tendency, variability, and distribution of a data set.
- Make scatter plots and correlation matrices across different variables and histograms of variables (see [this example](#)).
- Make box plots of variables.

Methods

- Describe the methods that you are using in your project and explain in detail how you applied them. You should use at least 2 methods for your problem so that you can compare their performance.
- Explain briefly how each method works, what its strengths and weaknesses are, both in general but also in the light of your project (how suitable is the method *in your case?*).

- If relevant, describe which hyperparameters are optimized for the methods (e.g., the shrinkage factor is a hyperparameter in Lasso regression).
- Describe clearly how you evaluate the performance of the different models and methods (accuracy, MSE, misclassification error, CV error, ...). Explain how each performance metric is calculated, and on what data, and why it is a useful measure of model performance.
- (optional) Consider and describe potential limitations of the methods and the chosen evaluation metrics.

Results and interpretation

1. Present your results in a *clear and organized* manner. This could include tables, graphs, or other visualizations that help to convey your findings. Report also all the hyperparameters, the performance (e.g., test error) etc. that you introduced in the Methods section.
2. Interpret the results. **You can compare the different methods in terms of aspects such as performance, computational cost, flexibility, bias-variance trade-off, etc.** The interpretation should depend on the prediction/classification task you decided on.
3. Discuss any limitations or caveats that are important to keep in mind when interpreting your results.
4. (Optional) Give an outlook on potential alternative/better ways to analyze your data in the future.

Summary

Summarize the main findings of your project. What did you discover, and what were the key insights that you gained from your analysis?