

CS-C3240 - Machine Learning D, Lecture, 10.1.2022-8.4.2022

Report Outline

Introduction:

- Explain the background (real-life scenario) of your ML application.
- Briefly outline the structure of this report

Problem Formulation:

- Formalise the application as an ML problem.
- Clearly explain the data points, features and labels of this ML problem
- Explain the source of the dataset

Stage 1
10 Feb, 20:00

Methods:

- State the number of datapoints, briefly describe the dataset and/or any data preprocessing needed.
- Explain your feature selection process (no theoretical justification needed).
- Describe and explain (why?) your choice of ML model(s)/hypothesis space(s)*, e.g., linear predictors, etc.
- Describe and explain (why?) your choice of loss function(s)*, e.g., logistic loss
- Explain the process of model validation - how did you split the data into training, validation and test sets. What are the sizes of each set and why did you make such design choice.

Stage 2
3 Mar, 20:00
You are required to discuss only 1 method.

Results:

- Compare and discuss the training and validation errors obtained for all ML methods considered.
- Which is the final chosen method and why?
- What is the test error of the final chosen method?

Conclusion:

- Summarise the report and your findings.
- Are the results suggesting that the problem is solved satisfactorily, or might there be room for improvement?
- Explain the limitation of the methods and how it can be further improved.

Bibliography/References

Appendices

- Your code with which you obtained the results discussed in the report

Stage 3
31 Mar, 20:00
You are required to discuss at least 2 methods

*) Choose from the ones covered in the course.