

Project Details

ACTL3143 & ACTL5111 Deep Learning for Actuaries
Patrick Laub





Warning

This page is out of date for 2024, and will be updated shortly.

A complete deep learning project

Individual project over the term. You will:

1. specify a *supervised learning* problem,
2. collect and clean the data,
3. perform an exploratory data analysis (EDA),
4. create a simple (non-deep learning) benchmark model,
5. fit two different deep learning architectures,
6. perform hyperparameter tuning,
7. write a discussion of the results.



Project components

The deliverables for the project will include:

1. draft due at noon on Friday in Week 5 (10%),
2. recorded presentation due at noon on Friday in Week 8 (15%),
3. final report due at noon on Monday of Week 10 (15%).



Project draft (10%)

Draft should show that you have:

1. specified your supervised learning problem,
2. collected and cleaned the data,
3. performed a basic exploratory data analysis,
4. create a simple (non-deep learning) benchmark model.

Upload to Moodle by **noon on Friday in Week 5**, no late submissions.



Presentation

Create a 3–5 minute recording covering:

1. the problem you are investigating,
2. the source of the data,
3. the deep learning approaches you are using, and
4. preliminary results you have (table of metrics).

Deliverable: YouTube link (public or unlisted) to a special StoryWall page. Presentations will be “public” to the class.

Suggestions: aim to be fully public and give peer feedback.



Presentation marking scheme

- **Content** (6%): did you cover the four points on previous slide?
- **Style** (6%): are your slides/figures professional and do they enhance the presentation?
- **Delivery** (3%): is the presentation interesting and within the time limit?



Tip

It is a critical skill to be able to condense a complicated project into a short pitch. The project report is where you will give us all the details.



Presentation tips

- Each project is different, you decide which parts to focus on.
- Not necessary to film yourself.
- Nice to *briefly* show the data (look at my lecture slides for example).
- Don't go overboard on EDA. Mention the *most important* 1–2 facts (if any!) about the data. E.g. imbalanced classes for classification.
- You can avoid adding 'UNSW' & the course code.



Report requirements

You are asked to cover the four requirements in the draft, and also:

- fit two different deep learning architectures,
- perform hyperparameter tuning,
- write a discussion of the results and any potential ethical concerns.

Deliverable: Report (PDF file), Jupyter Notebook, and dataset (e.g. CSV or ZIP file). Submission not public, probably to Moodle.



Report marking criteria

- **Content (8%)**: did you cover the seven points in the ML workflow?
- **Style (5%)**: does your report look professional, are your plots/tables useful and unpixelated, do you have spelling or grammar errors, are you within the page limit, and is the text easy to read?
- **Code (2%)**: is your code clean and well-commented, have useless cells been pruned, does it give errors when the “Run All” button is pressed?

Avoid screenshots & code in the report.



Some comments on the report

- **Focus on deep learning:** I'm most interested in seeing your ability to use and explain your neural networks. For example, your mastery of the Lee–Carter model is less important to demonstrate.
- **Hyperparameter tuning:** The tuning is one significant change from the weekly StoryWall tasks. Add a table (for each neural network) showing (at least) two hyperparameters that you tuned.
- **Use appendices:** If you run out of space, use appendices which are not counted in the page limit. E.g., the less urgent parts of your EDA can go in here.

