Machine Learning Practical: Coursework 3 – Project Interim Report

Due date: 16:00 Monday 19 February 2018

Report length: 6 pages + references

MLP in semester 2 will be based on projects done in groups of 2–3 students. The projects, which should be done using one of the deep learning toolkits (TensorFlow is recommended) can be chosen from a variety of topics and data sets. By working in a small group you can discuss ideas and work things out together.

1 Project Groups

You should do your project in a group of 2–3 students. Form your own groups. You can use the Piazza 'Search for Teammates' to help you form a group, if you like.

You may discuss any aspects of the assignment with your group and divide up the tasks however you wish; but we encourage you to collaborate on each part rather than doing a strict division of tasks, as this will enable better learning for all of you.

If you have not done so already, register your group on the Project Groups Sheet. Then sign-up your group for a weekly tutorial session on the Tutorial Signup Sheet.

2 Projects

Once you have formed a group, your first task is to scope out your project. Although we can give some pointers (via your tutor, via Piazza, via the daily helpdesk), you must scope your own project. These are the most important things to bear in mind:

- Your project should be feasible for the group to achieve in seven weeks (given that you have other courses!)
- It should have a significant amount of experimentation
- It should link to the main themes of MLP so far – deep learning – but you may also extend things beyond what has been covered. Thus projects may focus on convolutional networks, recurrent networks, or deep feed-forward networks, and the task under consideration may be classification, density estimation, sequence learning, unsupervised learning, or reinforcement learning.

There different ways to choose a project: a group might begin with an interesting data set or task, and focus on engineering fairly standard approaches to work well on that task; or a group might begin with a more challenging approach, and work on a dataset they already understand and for which it is straightforward to produce good baselines. Both types of project are valid, and you can get excellent marks for projects of either type.

To scope your project make a plan which addresses the following issues:

- what data you will be using?
- what approaches you will investigate?
- what are the research questions?
- what are the project's objectives?

There are many possible data sets that you could explore including CIFAR-10/100 object recognition, the Million Song Database (or a subset of it) for tasks like music genre recognition, Painter-by-numbers to predict if images of two paintings are by the same artist, and the large movie review dataset for sentiment analysis. Also, if you have a suitable data set and task then you are free to Bring Your Own Data (BYOD).

There are also many possible technical approaches you could investigate such as multitask learning, curriculum learning, one-shot learning, generative models, Bayesian deep learning, and meta-learning.

It is important to note that the aim of a project is not necessarily to obtain the most accurate possible results on a task. You might be exploring if a particular technique is useful. If, after careful experiments, it results in a less accurate system this is OK. Especially if you can provide some analyses of why it failed to improve.

3 Weekly updates

Each group should set up a Google Doc which you should use to report progress (including results), give plans, and raise any questions the group has, updated each week. In addition to be being shared between the group members, this Doc should be shared with your tutor, and with the course lecturer (Steve Renals).

Tutors will review Docs of the groups that they tutor each week, before the tutorial. You should make sure your group's Doc is updated 24 hours before the tutorial.

4 Software and Hardware

You are free to use any of the deep learning toolkits including TensorFlow, Keras, and PyTorch. Please report this clearly (but concisely). You are also free to use other externally produced software; in this case it is very important to be clear and precise about your own contribution, especially if you modify such code.

We are in the process of making available a GPU cluster (based on Nvidia GTX 1060 Ti GPUs) for the MLP course. Please see the project website, emails, and Piazza for more information about this.

We shall also create a branch `mlp2017-8/semester2` on the MLP github (https://github.com/CSTR-Edinburgh/mlpractical) which will contain useful scripts and other files.

5 Backing up your work

It is **strongly recommended** you use some method for backing up your work. Those working in their AFS homespace on DICE will have their work automatically backed up as part of the routine backup of all user homespaces. If you are working on a personal computer you should have your own backup method in place (e.g. saving additional copies to an external drive, syncing to a cloud service or pushing commits to your local Git repository to a private repository on Github). **Loss of work through failure to back up does not constitute a good reason for late submission**.

It is recommended that project groups keep their work under version control using git. If you make regular commits of your work, then this will allow you to keep track of the changes you have made and if necessary revert to previous versions of files and/or restore accidentally deleted work. Using git will also support effective collaboration within a group.

6 The Interim Report

Each group's Interim Report should cover the following issues:

- Motivation and introduction to the project
- Research questions
- Objectives – be precise
- Data set and task
- Methodology
- Baseline experiments (and any further experiments that have been done so far)
- Interim conclusions
- Plan for the remainder of the project, including discussion of risks, backup plans

Motivation: You should give a broad introduction to the project, including citations to related work. Your aim here is to explain why the project is addressing an interesting topic, and how it relates to things that have been done in the area.

Research questions: The research questions then be more specific, addressing precisely what the aims of the project are. In this section you should make clear what the project's contribution is: how is it different to what is already done. Please note that projects do not have to make novel research contributions; however they should make a contribution in the sense that they extend what is already available. Perhaps it will be a more detailed investigation of some particular aspects of a model/algorithm on a particular dataset; perhaps you will be looking at an approach not previously reported on a particular dataset. Your project should involve more than using existing code and running it without modification.

Objectives: The interim report should state the objectives of the project, which are related to the research questions. What experiments do you plan to carry out? You can differentiate between core objectives, and optional objectives you hope to achieve if things go well. The conclusions in your final report should relate to these objectives.

Data: Clearly describe the data set and task you will be exploring. If the data requires any preprocessing, then explain this. The description should be in enough detail such that your work would be reproducible by another group. Describe how you will evaluate the task (for example, classification accuracy). Use citations where appropriate.

Methodology: Explain clearly the technical methodology, the models and algorithms that are used. Approaches that were covered in MLP in semester 1 can be described briefly. Approaches beyond this should be explained in more detail. Again use citations to the literature.

Experiments: The interim report should include some experimental results. In most cases these will be baseline experiments. Baseline experiments refer to experiments conducted using well-understood approaches against which you can compare later results. For example if you were exploring a new data set, the baselines might include linear networks and deep neural networks with different numbers of hidden layers; if you were exploring a different approach to regularisation, then the baselines would include no regularisation, and conventional techniques such as L1, L2,and dropout. Do include the results of any further experiments in your interim report.

Interim conclusions: What have you learned so far? Do the experiments indicate that the project is feasible? Do the experiments indicate that you should consider changes to your original plan? Can you compare your results so far to what has been reported in the literature?

Plan: Based on what you have done so far, present a plan for the rest of the project. Are there any changes to the objectives? What are the risks? Do you need a backup plan?

7 Report Details

A single interim report should be submitted for each group. The report should show the project group number (e.g. G123) and the student matriculation numbers of the team members.

Format and length: Use the same document style for the interim report as for courseworks 1 and 2. The interim report should be a maximum of **6 pages** long, not including references. We will not read or assess any parts of the report beyond the allowed 6 pages + references.

Citations: If you make use of any any books, articles, web pages or other resources you should appropriately cite these in your report. You do not need to cite material from the course lecture slides or lab notebooks.

Marks: This assignment will be assessed out of 100 marks and forms 25% of your final grade for the course.

Academic conduct: Assessed work is subject to University regulations on academic conduct: http://web.inf.ed.ac.uk/infweb/admin/policies/academic-misconduct

Submission: You can submit more than once up until the submission deadline. All submissions are timestamped automatically. Identically named files will overwrite earlier submitted versions, so we will mark the latest submission that is submitted up to 7 days after the deadline.

Late penalty: Reports submitted after the deadline will be recorded as late and will be penalised as follows:

> *Following the University guidelines, late coursework submitted without an authorised extension will be recorded as late and the following penalties will apply: 5 percentage points will be deducted for every calendar day or part thereof it is late, up to a maximum of 7 calendar days. After this time a mark of zero will be recorded.*

Warning: The `submit` command will allow you to submit after the deadline. The late penalty will be based on the timestamp of the latest submitted report. Thus if you submit before the deadline, then make a submission after the deadline you will be subject to a late penalty.

Extension requests: For information about extension requests (and additional information about late penalties), see http://web.inf.ed.ac.uk/infweb/student-services/ito/admin/coursework-projects/late-coursework-extension-requests
Course lecturers cannot give extensions. Therefore please do not contact any course staff directly about extension requests; you must follow the instructions on the web page.

8 Submission

Your coursework submission should be done electronically using the `submit` command available on DICE machines.

Each group should nominate one team member to make the submission for the group. Multiple submissions of the report are not required. Please make sure that your report includes the student ID numbers of the team members, as well as your project group ID.

Your submission should include

- your completed interim report as a PDF file, using the provided template
- a directory files containing code, scripts and configurations you used in the project

Please do not submit anything else (e.g. log files, parameter files, outputs, ...).

You should copy all of the files to a single directory, `coursework3`, e.g.

```
mkdir coursework3
cp -a interimReport-<groupID>.pdf coursework3
cp -a <groupID>-codeDirectory coursework3
```

(where <groupID> corresponds to your project group ID e.g. G123)

and then submit this directory using

```
submit mlp cw3 coursework3
```

Please submit the directory, not a zip file, not a tar file.

The `submit` command will prompt you with the details of the submission including the name of the files / directories you are submitting and the name of the course and exercise you are submitting for and ask you to check if these details are correct. You should check these carefully and reply `y` to submit if you are sure the files are correct and `n` otherwise.

You can amend an existing submission by rerunning the `submit` command any time up to the deadline. It is therefore a good idea (particularly if this is your first time using the DICE submit mechanism) to do an initial run of the `submit` command early on and then rerun the command if you make any further updates to your submisison rather than leaving submission to the last minute.

9 Marking Guideline

- Abstract – how clear is it? does it cover what is reported in the document
- Introduction – do you clearly outline and motivate the paper, and describe the research questions investigated, with appropriate citation to the literature?
- Objectives – are the objectives clearly and precisely stated?
- Data set and task – is the data set clearly presented? is it clear what preprocessing has been done, and how the data is split into training/test/validation? is the task explained? how is the task evaluated?
- Methodology – are the technical approaches adopted well-explained, with reference to the literature?
- Experiments – have appropriate baseline experiments been carried out? Are the results presented in a clear way? Is it clear what the hypothesis underlying each experiment is?
- Interim conclusions – how does the work done so far match to the objectives? are changes necessary? are the conclusions clearly expressed?
- Plan – is there a clear plan for the remainder of the project? have any risks been discussed? is there a need for a backup plan, and if so what is it?

In general your report should be clearly written and presented, well-structured, make good use of citations to the literature, and enable your experiments to be reproduced.