

Machine Learning Practical

Part 2: Group Projects

MLP Semester 2: Group Projects

Steve Renals, Hakan Bilen, Pavlos Andreadis

Machine Learning Practical — MLP Lecture 11
27 November 2018

<http://www.inf.ed.ac.uk/teaching/courses/mlp/>

You're not expected to start work on your project until next semester
– but it might be helpful to know more about what is coming up

You're not expected to start work on your project until next semester
– **but it might be helpful to know more about what is coming up**

- Semester two will be based on **group projects**
- 2–3 students per group
 - You can discuss any aspects of the assignment with your group
 - Divide up the tasks any way you like
 - Best if the team collaborates on each part
- Can use Piazza “Search for Teammates” to help form a group - but make sure you use the MLP Piazza for Semester 2 (signup at piazza.com/ed.ac.uk/spring2019/infr11132)

What should a project look like?

- Feasible for a group of 2–3 to do in 7–8 weeks, not forgetting your other courses!
- Needs to have a significant amount of experimentation
- Should link to the main themes of MLP so far, but you can extend things
 - Conv nets, recurrent networks, feed-forward networks, ...
 - Classification, density estimation, reinforcement learning, ...

Choosing a project

- How to choose a project?
 - Begin with an interesting data set or task, and focus on engineering fairly standard approaches to work well
 - Begin with a more challenging approach and work on a dataset you already understand and for which you have good baselines
 - Take an interesting seed paper and extend it in one or more ways

All these types of project are valid, and you can get excellent marks on for projects of all these types

- Over the next few weeks we'll construct a list on the course website designed to stimulate ideas for projects
 - Recent interesting papers in areas related to the MLP course
 - Kaggle competitions which might be feasible for MLP
 - Other possible datasets

You are also free to develop your own project entirely separate to this!

Some example projects from 2017/18

- Investigating the properties of capsule networks
- Image captioning with neural networks
- Generating facial images from sketches
- LSTMs for financial data time series
- Tackling the one-shot problem
- Predicting cryptocurrency prices
- Nerve segmentation in ultrasound images
- Music genre recognition using LSTMs and CNNs
- Using CNNs for painter identification
- Sequence-to-sequence models for text summarisation
- Transfer learning on CIFAR-10
- Weight sharing in deep reinforcement learning for continuous control
- Painting generation using conditional GANs
- Word2Vec and RNNs in automatic playlist construction
- Using GANs to generate audio
- Improved optimisation for neural machine translation

Interactions with instructors in semester 2

- Weekly lectures on advanced topics in deep learning
- MLP Helpdesk – best place for technical queries
- Tutorials – discuss the progress of your project
 - Each project group is assigned to a tutor, who will discuss and review progress
 - A project group will set up a Google Doc for the group (shared with instructors) that is used to report progress and experimental results, give plans, raise questions
 - Weekly tutorial sessions to meet with tutor – tutorial sessions will typically involve 5 groups
- Piazza – ask and answer questions, search for teammates, ...
- No scheduled labs although we will provide Jupyter Notebooks to help to get you started in TensorFlow and/or PyTorch

Computing... “you’re gonna need a bigger boat”

- Deep learning uses up a lot of compute cycles...
- The MLP GPU system available for use in semester 2 – 25 servers each with 8 NVidia 1060 Ti GPU cards
- Why 1060Ti? Need to make a balance between power consumption*, compute performance, and cost ...
(*) When running 200 GPUs, the issue of power consumption becomes really important!
- Also possible to use cloud services (e.g. Google Cloud, Microsoft Azure) – it’s not yet clear if we will have additional coupons specifically for the MLP course for semester 2

Coursework 3 – Interim Report

- Motivation and introduction to the project
- Aims and objectives – be precise
- Data set and task
- Research questions
- First phase of experiments
- Any interim conclusions
- Plan for the remainder of the project, including discussion of risks, backup plans
- *Feedback only, no numerical mark*
- Submission deadline: Thursday 14 February 2019, 16:00

- Brief introduction, including a reprise of the aims and objectives, the data and the task
- Experiments
 - Methodology
 - Results
 - Discussion and interpretation
- Conclusions and discussion
 - Conclusions with respect to aims and objectives, research questions
 - Any changes with respect to the original plans
 - Discussion of what was achieved and learned in the project
 - Potential further work
- Submission deadline: Friday 22 March 2019, 16:00

- *Can I do the project alone?*

We won't stop you, but it is not recommended. We are expecting projects to be have the amount of work from a 2-3 person group; interacting with your team is an important experience.

- *Do we have to use PyTorch or TensorFlow?*

No: Keras, MXNet, OpenAI Gym, . . . would all be OK.

- *Can this be part of my dissertation project?*

No, it should be completely separate.

One more thing...

We have prizes!

IBM UK have again kindly donated a prize
for the best project in MLP.

More details next semester...

We have prizes!

IBM UK have again kindly donated a prize
for the best project in MLP.

More details next semester...

See the coursework web page for
the short-listed projects for last years prize

One more thing...



Thank you very much,
and see you in 2019