Machine Learning Practical (MLP; INFR11132) – mid-semester student feedback
21 October 2016 (Survey: 19 October 2016)
Steve Renals

This is a concise overview that summarises the feedback collected during week 5 of the Machine Learning Practical. 69 students (32% of those registered for the course) completed the form.

1 Level of challenge

The feedback form asked students to choose one of the following option best describes how they feel about the level of challenge in MLP this course

(a) I would prefer a more challenging course (moving faster and/or more depth).

(b) I am happy with the level of challenge in this course.

(c) I can probably manage, but would prefer a slightly less challenging course.

(d) So far I have found this course much too challenging for me.

The results are summarised in figure 1.

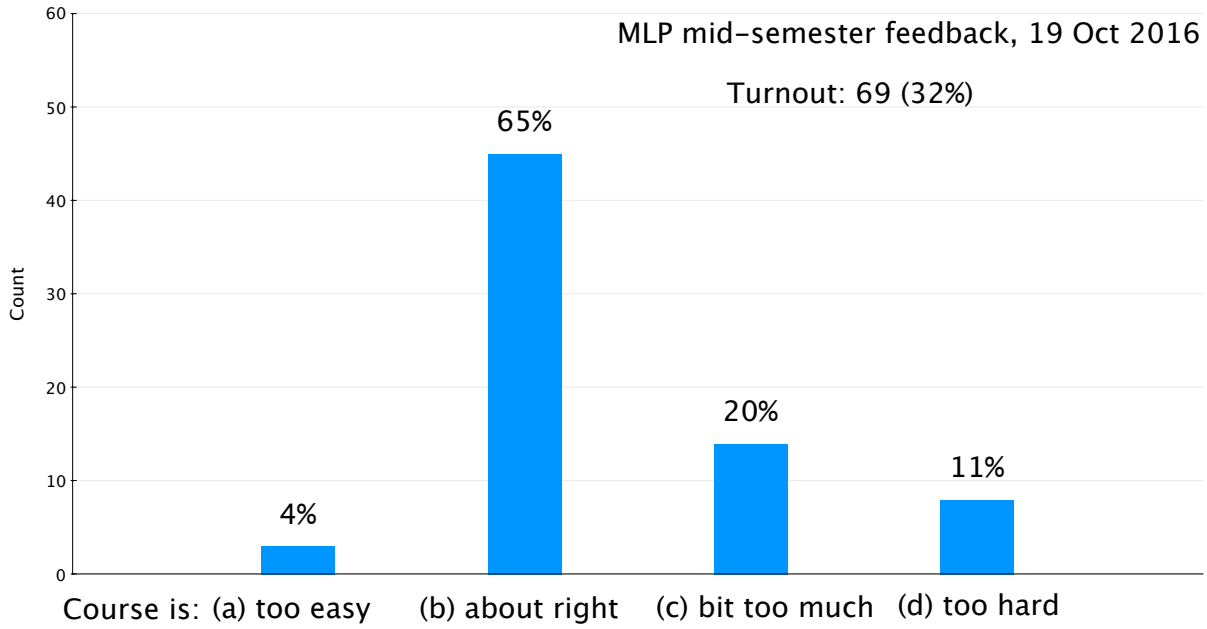

Figure 1: Summary of responses to level of challenge in the MLP course. The a/b/c/d options have been paraphrased to make them more concise as labels.

Overall about 2/3 of those who responded found the level of challenge to be about right; and total of 85% found it about right or something that they can manage. A couple of people would like it to be a lot more challenging, and 11% of respondants found it much too challenging,

2 Free text comments

Students were invited to write comments relating to the following prompts about the course:

- **Keep doing this!** It's working well
- **Start this!** I think it's worth a try
- **Stop this!** I don't find it helpful

They were also invited to comment on individual steps that they could take to improve their own learning in MLP. 87% of the responses included some written comments. These comments have been (manually) clustered together revealing the following trends:

30–40% of respondents...
- positive about labs

20–30% of respondents...
- positive about lectures
- thought they should do more reading

10–20% of respondents...
- positive about lecture recordings
- would like more interactive/graphical lecture slides
- would like more examples in lectures
- would like more lab time
- thought they should spend more time on self-study

5–10% of respondents...
- positive about the software environment
- negative about the software environment
- positive about the demonstrators
- positive about the link between lectures and labs
- would prefer to use Piazza rather than NB
- thought they should work on their own experiments

3 Response to feedback

This feedback has been very helpful, and it is good that the majority of stuidents are happy with the level of challenge offered by MLP. Some responses to the main comments raised are included below:

1. *Why aren't there more examples in lectures?*
   This is a good question. Basically the course is designed so that the lectures focus on the theory and methodology, and there are examples to explore this in the labs. We have carefully designed the course such that the material covered in lectures is explored in the labs and the coursework. So you are creating the examples yourself. I believe that together the lectures and the lab notebooks provide what is required.

2. *Please make the lecture slides more interactive. Please use less text and more pictures. Please provide notes.*
   I am doing the best job I can with the lecture slides and do try to provide diagrams and illustrations where appropriate. The slides are basically the concise notes for this course together with references to reading material (which is included in every lecture). So far most (but not all) the concepts are covered in the online book `http://neuralnetworksanddeeplearning.com` ; as

we move to more recently developed topics, so there will be more links to the literature. I take the point about providing notes. Unfortunately it will not happen this year. But I am putting effort into making the slides as informative as possible (although I understand that this can make them less appealing to some).

3. *Why not have a lecture (or part of a lecture) about the software framework and its architecture?*
   To be honest, I did not really consider it, as I thought it was very well introduced in the labs. But I can see that having at least some videolecture material about to introduce the software framework could be helpful. This is probably something to do for next year's course.

4. *Lectures are too fast-paced.*
   There are two issues here. First apologies if I speak too fast sometimes, I do appreciate that this is difficult especially if English is not your first language. I will try harder. The second issue is that each lecture contains too much new material. I don't think they do (and please note that the course explicitly recommends previous experience of machine learning – it is not an introductory course). Part of the aim of the course is to encourage self-study, in particular via the labs. Basically every concept discussed in the lectures (with the possible exception of recurrent networks which may be the focus of the final lecture) is linked to either the labs or the coursework, so you get the opportunity to explore it yourself and in detail.

5. *Why are we using this codebase? Why not TensorFlow?*
   We have a clear vision for how MLP should operate: in the first semester one of the key objectives is to learn in detail how to implement different neural network architectures and to understand how gradient descent training works. We have provided a codebase which is attempts to be efficient and clear. I have many years experience of different neural network and machine learning software frameworks, and I believe that Pawel and Matt (the authors of the MLP framework) have done an excellent job producing software which is efficient, extensible, and as clear as it might be. The codebase is also well documented. Maybe we should have made the following point more strongly: spend some time studying the MLP codebase, so you understand how everything works.

   Regarding toolkits, it is our plan to use TensorFlow (or possibly another toolkit) in semester 2, learning to use such a toolkit is another goal of the course. Please note that it probably takes longer to get up to speed with TensorFlow compared with the current framework.

6. *Why are we using NB and not Piazza?*
   This is a very fair question. What I like about NB is the ability to highlight part of a pdf thus making it easy to ask specific questions. But the obvious downside of NB is that it is cumbersome, at best, as an interactive forum. While Piazza is much better forum, software, I would miss the core NB functionality. I'm on the lookout for better solutions, and open to suggestions.

7. *Please can we have study groups?*
   I am very happy to provide any support necessary (e.g. help with booking rooms to meet in) - if you would like study groups then please go ahead and self-organise.

8. *More links to current research would add value.*
   We are planning on four guest lectutes in semester two which will indeed address current research topics related to MLP.