

Reinforcement Learning (RL)

Response to student feedback in 2017/18 course survey

I have taken note of the feedback from students in this course survey, and have the following comments in response:

- On Course Organisation: It is true that the course this year suffered from some organisational issues because of the late assignment of the course, which in turn led to limited resources in terms of time and teaching support. I had to prioritise where these resources would be spent. There were regrettable consequences, specifically as refers to the slides (page numbering, notational consistency, and occasional late release) and the feedback on coursework (late, and not personalised). Such a situation will not be repeated, since the lecturer for 2018/19 has already been appointed.
- On Course Content: One of the tasks given to me at the beginning of the semester was to reduce the amount of material covered by the course, which was close to that of a 20 credit course (Reinforcement Learning (RL) is worth 10 credits). This was done by not covering Bandit Algorithms, and focussing on a few key algorithms on later topics. The fundamentals were covered in significant depth, and a number of advanced topics received their own lectures.
- On Teaching Style: As a lecturer, I like to discourage using the slides as a primary source of study material. Instead, I use the slides as an aide to help focus on fostering understanding through examples and discussion. The students are expected to spend a good amount of time studying the material specified, and the lecture is intended to give them a way to approach that study successfully. I complement this with post-lecture discussions, extended office hours, and very active forum engagement (Piazza). The coursework's main aim was to push for this interaction, and bring students face to face with decisions that have to be made before deploying RL algorithms.
- On Examinable Material: The examinable material is clearly specified in the course website. I complemented this with a detailed list of what knowledge one should be leaving the course with, which also functioned as a study guide. In addition to this, a mock exam in a similar vein to that of the official examination was held. This was discussed in class, and solutions were released the week thereafter.
- On Deep RL: When considering more advanced topics, I shared a survey with the students to gauge their interests. One of the options was Deep RL. Though this was the topic in which students showed the most interest, their interest spread across topics. This survey was intended, and announced, as advisory. I decided against teaching Deep RL, given that the course did not have Deep Learning as a prerequisite and the necessary material could not be adequately covered within the 10 credit course.
- On the Programming Environment: I would tend to agree that Python is seeing more widespread use for Data Analytics and in Machine Learning research. However, the programme description for the course specified the use of Matlab for this year. I considered allowing both, but this would have taxed my already limited resources. Ultimately, the teaching of Python is not a goal of the course. If you have some knowledge of programming, and have built a good understanding of the topic from the course, then you should have no problem transferring that to any programming language. The same applies to the use of the OpenAI framework. More importantly, such a framework would abstract away from some of the implementations that I want the students to have to go through. Even if the course had been based on Python, I would

not have made use of OpenAI. An additional point to consider is that Matlab is a simpler language to get into and, being a postgraduate course, RL brings people in from diverse backgrounds. Lastly, it is worth noting that the restriction in the DRPS on using Matlab has now been lifted.

Pavlos Andreadis, September 2018