

Report - EL2805 – Fall 2021

Please note that there is only one respondent to this form: the person that performs the course analysis.

Course analysis carried out by (name, e-mail):

Alexandre Proutiere, alepro@kth.se

COURSE DESIGN

Briefly describe the course design (learning activities, examinations) and any changes that have been implemented since the last course offering.

The course is based on 14 standard lectures based on slides and blackboards, 6 exercise sessions and 3 computer exercises. The assessments are based on 1 homework, 2 laboratory exercises and a final exam. Since the first course (2018-2019), we have made efforts each year to polish and improve the content of the course. Most efforts have focused on introducing the rapidly evolving tools used in deep RL – we are always trying to integrate novel algorithms in the lectures but also in the labs.

This year, we have completely reformulated the exercise sessions and the labs. For the exercise sessions we decided to have a better mix of theory and practice.

We have also introduced lectures on TRPO and PPO as promised last year.

THE STUDENT'S WORKLOAD

Does the students' workload correspond to the expected level (40 hours/1.5 credits)? If there is a significant deviation from the expected, what can be the reason?

It is always hard to measure the exact workload that the students spent on the course. Based on the students' comments in the course evaluation, the students did spend around less than 18 hours per week on the course. The students come with very heterogenous backgrounds, which induces a large variance in the efforts put in the course. As in previous years, some students found the course too easy, some others too demanding. We have created extra labs for students willing to gather bonus points – these labs were very time consuming but they were not mandatory. This year, we made the bonus points easier to get.

THE STUDENTS' RESULTS

How well have the students succeeded on the course? If there are significant differences compared to previous course offerings, what can be the reason?

This year, the pass rate for the exam was around 70% overall. The pass rate is similar to those in previous years.

OVERALL IMPRESSION OF THE LEARNING ENVIRONMENT

What is your overall impression of the learning environment in the polar diagrams, for example in terms of the students' experience of meaningfulness, comprehensibility and manageability? If there are significant differences between different groups of students, what can be the reason?

From the course evaluation, the students are in general positive about the learning environment. This year, we have focused on trying to provide as much feedback to the students as possible (this was asked in previous course evaluations). This remains a challenging task due to the large number of students. Again, as in previous years, the group is extremely heterogenous, and this constitutes the main difficulty for the course design. To accommodate for this heterogeneity, we have tried to adapt the content and labs so as to provide a meaningful experience to every student.

ANALYSIS OF THE LEARNING ENVIRONMENT

Can you identify some stronger or weaker areas of the learning environment in the polar diagram - or in the response to each statement - respectively? Do they have an explanation?

Strong aspects of the courses: the lectures, the new exercise sessions, the labs.

Weaker aspects: last year the exercise sessions and their connection to the course were the weak point. From the evaluation, we are in a good direction to improve this aspect.

ANSWERS TO OPEN QUESTIONS

What emerges in the students' answers to the open questions? Is there any good advice to future course participants that you want to pass on?

The course blends theory and applications. Students should primarily understand the basics before experimenting in the labs. By mastering the basics, the students will understand the applicability and the limits of RL techniques.

PRIORITY COURSE DEVELOPMENT

What aspects of the course should primarily be developed? How could these aspects be developed in the short or long term?

We are working on developing a lecture on RL with human feedback, an essential component of the fine-tuning procedure for LLMs.

OTHER INFORMATION

Is there anything else you would like to add? Nope.