

Report - EL2805 – 2020

Respondents: 1

Answer Count: 1

Answer Frequency: 100.00 %

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), the course has put more emphasis on RL with function approximation. The computer exercises have changed.

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 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 37 hours /1.5 credits. However, the students come with very heterogenous backgrounds, which induces a large variance in the efforts put in the course. From the students' comments, some students found the course too easy, some others too demanding.

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 time the pass rate for the exam was 83% overall, which is a positive result. The pass rate is similar to last year's.

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, it seems that the students are in general positive about the learning environment. However, the group is extremely heterogenous, and this constitutes the main difficulty for the course design. The content seems obvious to some students, and very challenging to others. We have decided in the course to start from the basics, and hence to make the course accessible to all. This choice may have made some students frustrated, but overall, I think this was the right choice. I do not see very strong differences between student perceptions of the course (remember that we have 39 students who answered the questionnaire – so very few per group).

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?

In general, the comments of the students about the learning environment are positive. As last year, the less strong aspect of the course is the feedback given to students. The course has attracted lots of students, and it has become even more difficult to give proper feedback.

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 is on RL, a discipline that has experienced a lot of empirical developments in the past ten years. Every year, hundreds of algorithms are proposed – essentially based on deep learning techniques. As one student commented in her response, DQN is old, and the same student would like to see the most recent algorithms. Actually, DQN was published in 2015. Since then, many algorithms were developed; the latter are heuristics, have no foundation at all, and from my perspective, are often wrongly designed.

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 new labs to smoothly go from tabular RL to deep RL. Since the students are willing to know more about deep RL, I will improve this part of the course, and present more of the recently proposed algorithms.

OTHER INFORMATION

Is there anything else you would like to add?

Nope.