Report - DD2380 - 2019-05-09

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):

Iolanda Leite, iolanda@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 arranged as a series of lectures and three tutorial sessions to deepen the understanding of selected areas. Thanks to the breadth of AI, several lectures were given by guest lecturers, who are experts in the field (Johan Boye and Gabriel Skantze on Natural Language Processing, Danica Kragic on Robotics, Mårten Björkman on Computer Vision, and Josephine Sullivan on Machine Learning). The rest of the lectures were given by Jana Tumova, Patric Jensfelt and Iolanda Leite and focused on topics from three areas: taming uncertainty, problem-solving, knowledge representation and planning.

A criteria-based grading scheme is used with TEN1 2hp consisting of a series of 9 online quizzes released after lectures, and an essay on ethics and risks of AI (with individual reflection and a team discussion part) , and LAB1 4ph with 2 programming assignments, and a choice between an individual pen and paper assignment or an open-ended project on planning in teams of 4 students for higher grades. The programming assignments are conducted typically in pairs, in C++ or Java, evaluated in Kattis, and also presented in person to teaching assistants. There is no written exam, the final grade A-F is determined from the grades of LAB1.

This year, Iolanda Leite was the course responsible for the first time. The lecture topics were updated (especially the Reinforcement Learning lecture) and the quizzes were revised. Some minor improvements were introduced also to the programming assignments, for example, the list of questions asked by the TAs for evaluation was revised. The essay assignment also followed the same format as in previous course rounds, with individual reflection and a team discussion.

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?

The reported students' workload varied quite a lot, but there is a "normal distribution" trend around the expected 20 hours/week workload, with 20% of the students selecting that option.
The comments from the students show that one of the reasons could be a big difference between requirements for A and requirements for E. One student commented that the requirements for A-B level grades demands a higher workload than many other courses.
Another possible explanation for the reported workload variability is the fact that in this course round the background of the students was particularly diverse, with a significant part of the students coming from international programs or KTH programs where this course is not mandatory.

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?

The students who passed did approximately as well as the year before. However, in this course round there is a higher percentage of the students who did not fully complete the course yet (aprox. 40% compared to 20% in the previous round), but this analysis is being conducted before the re-examination period. There are many students who are missing only one assignment to complete the course and we expected them to take part in the re-examination. Therefore, we predict that the course completion rates will get closer to previous years after the re-examination phase.

Nevertheless, one reason for the higher non-completion rate this round can again be explained by the diversity of the students. Because many of the students were taking this as an elective course, they might want to learn about AI as a very timely topic but not expect the content (a wide range of AI topics and algorithmic principles), and requirements (especially on programming and understanding of probabilistic inference) to be so demanding.

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?

The course seems to be well-received and the students found it interesting. The course received good scores in meaningfulness, comprehensibility, and manageability; meaningfulness being the strongest one of the three. There seemed to be no gender differences or differences among the different groups of students, but given the low response rate of the questionnaire (20 participants), it is difficult to draw strong conclusions about different groups.

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?

Most students reported that the course covered interesting and challenging issues (highlighting the lectures and the lab assignments in their comments) and that they were offered opportunities to collaborate and discuss with others. They were also quite positive on the ability to get support. We offered consultations through a discussion forum on Canvas, quick in-person consultations during lecture breaks, as well as scheduled consultations, but we made an effort not to help with coding and debugging itself. Project advisors followed each project group from the stage that they proposed the work up to the evaluation, as it happened in the previous course round.

There were some mixed responses on the fairness of the assessments, which based on the open-ended questions seem related to the grading of the lab assignments. Some students were not happy with the fact that their evaluation is in part based on their Kattis submission results (something that is difficult to change in such a large course), while others believe that there were differences between TAs grading criteria. We will keep working to improve the preparation of the TAs for the presentation sessions and continue refining the evaluation procedure of the assignments by analyzing the question banks again and making sure the questions for the different grade level have the same level of difficulty.

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 future course participants were advised to start on the assignments early and to do the quizzes right after lectures. These are the most recurring comments in all course evaluations and we highlight them in the first lecture (and remind the students about the quizzes after each lecture).

PRIORITY COURSE DEVELOPMENT

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

In the short term: Further improvements in A3 and in making sure the evaluation sessions of the lab assignments are as coherent as possible between each other (regardless of the TA).

In long term: Together with the other course teachers and a team of TAs that we selected only for this particular task, we are working on a larger set of smaller programming assignments (more balanced in terms of difficulty) and introduce some integration in an intelligent agent. We hope these new assignments will solve many of the weaknesses pointed out on the current assignments, namely the fact that A1 requires a much higher workload than A2, and making sure that students don't spend a lot of time "tweaking" parameters to get a higher score in their submissions.