

Course report

Knowledge Representation & Reasoning (INFO282) 2023h

What are the teaching methods and forms of assessment used in the course?

The course consisted of 15 weekly lectures (1h45m each) and 14 weekly seminars (1h45m each). The final two lectures were used for summarising the contents and for the students to ask questions. For assessment, the final grade was given by a single individual (digital) exam. Through the course, the students were asked to submit six mandatory assignments (a single 'resubmission' was possible for each assignment). They needed to answer correctly at least 50% of each mandatory assignment to have the right of presenting the final exam.

Follow up from previous evaluations

Two of the main points of previous evaluations were the heavy workload and the seemingly disconnection between the course's topics and others that are currently discussed more frequently in the AI community (e.g., neural networks and machine learning). For the first, we reduced the number of mandatory assignments (from 7 to 6) and reduce the number of exercises in each one of them. For the second, we spend additional time in the introductory lecture to place the course's contents within the current AI situation, highlighting the relevance of the ideas, tools and methods the course introduces. A third point from previous evaluations was about the mandatory assignments not counting towards the final grade. This was not addressed as, when it became clear I will be teaching the course again, the deadline for making this kind of 'drastic' modification has already passed.

Form of evaluation

Student evaluation

Summarize the results from the student evaluation

Only 9 students filled the student evaluation. One of them pointed some issues, including the clashing of the last lecture and mandatory assignment with other course's exams (the lecturer has no control on other courses' schedules), lectures whose content was not present in the main syllabus (note: sources for those lectures were provided), and a better indication about how the final exam would be. Another mentioned that there were too many mandatory assignments, and most of them agreed that the workload of the course was high. Still, the students (who submitted the evaluation) were in overall happy with the course and with the teaching.

The course's coordinator evaluation

In my opinion, the course improved from the previous edition in the following aspects. The initial discussion helped to frame the contents of the course within the large AI picture. There were also less 'on the fly' changes, as I knew better what to

expect from the students. Finally, my impression (from personal comments) was that the students were happy with the lectures.

Grade distribution

The exam consisted of 24 questions for a total of 231pts. One of the questions (for 15pts) was wrongly formulated, so it was excluded from the grading. Then exam was then graded as a set of 23 questions for a total of 216pts. The following threshold was used for grading: F:[0%,50%); E:[50%,55%); D:[55%,65%); C:[65%,80%); B:[80%,90%); A: [90%,100%]. The resulting grade distribution among the 87 students was as follows: F: 19 students; E: 5 students; D: 12 students; C: 32 students; B: 14 students; A: 5 students.

Comments on the grade distribution

The grade distribution is what one might expect from similar theoretical courses, with few getting the highest grade (5 with grade A), most of them falling in the middle (58 with grades among B, C, D), few getting a barely passing grade (5 with grade E) and some (19) failing.

Goals for the next evaluation - what can be improved?

One aspect to address is making the mandatory assignments count towards the final grade (and maybe make them more accessible). This change will be proposed (even though I might not be the lecturer of this course next year, and thus someone else will need to implement it). Additionally, it would be also good to include material related with recent developments in the area.