Jana Tumova, LH216V VT19

Assignment 2

Context: This assignment focuses on the course DD2380 Artificial Intelligence. The course is mandatory to several Master programs in the first year and during past years attracted increasing numbers of students from other programs, ending up with 490+ students in HT18.

Part I

Intended learning outcomes:

After completing this course the student should be able to

1. interpret various concepts in Artificial Intelligence (AI)
2. develop effective and efficient solutions to address sub-problems from the AI domain
3. integrate tools to design a computer program exhibiting multiple traits expected from an intelligent system
4. explain own solution to a problem from the AI domain
5. demonstrate an insight into societal and ethical aspects of AI

so that the student can

- make use of methods from artificial intelligence in the design, implementation, and analysis of computer programs
- contribute to the design of an intelligent system in academic as well as industrial applications.

Grading criteria:

ILO		*Grade*	*Criteria*
1	interpret various concepts in Artificial Intelligence (AI)	E	can interpret basic concepts including modeling techniques, algorithms, and other tools in all topics covered in the course
		C	can interpret more advanced concepts including modeling techniques, algorithms, and other tools in all topics covered in the course
2	develop effective and efficient solutions to address sub-problems from the AI domain	E	can solve selected sub-problems from the AI domain at a basic level (in terms of effectiveness, i.e. the output quality and efficiency, i.e. the code optimality) by implementing and debugging given algorithms
		C	can propose and implement simple modifications to the given algorithms resulting in a performance improvement, based on an evaluation of weaknesses of own implementation of the given algorithms in terms of effectiveness, i.e. the output quality and efficiency, i.e. the code optimality

		A	can solve given sub-problems from the AI domain at a level significantly outperforming basic solutions in terms of effectiveness, i.e. the output quality and efficiency, i.e. the code optimality by selecting, implementing, and fine-tuning a combination of several algorithms, and can do so with limited resources in terms of time
3	integrate tools to design a computer program exhibiting multiple traits expected from an intelligent system	E	can design a computer program that is able to significantly improve the outcome of problem solving, planning, or decision making over time through the use of prediction or machine learning in an uncertain environment
4	explain own solution to a problem from the AI domain	E	can clearly explain details of own implementation, principles of the implemented algorithms, and fitness of the algorithm to solving the problem
5	demonstrate an insight into societal and ethical aspects of AI	E	can independently reflect on ethical and societal aspects of AI, discuss these, and support arguments with references to relevant scientific studies

Alignment of ILOs, grading criteria and assignments

The planned assignments are:

- A quiz on each topic covered in the course in two variants: basic QB and advanced QA. This assignment is individual, in Canvas. Both QB and QA yield grade Pass/Fail. The quiz questions will regard interpretation of various concepts in AI discussed in the lectures, either in general, or in small concrete instances. QB will focus on concepts marked as basic, while QA as concepts marked as advanced in the lecture materials. Some examples of the quiz question types are:
    - Which of the following is the correct description of a particular modeling formalism? (multiple choice).
    - Consider a particular type of problem. Which of the following algorithms can be used to solve it? (multiple answers)
    - Consider a small concrete instance of a problem and a concrete algorithm designed to address it. What is the output of the algorithm on the instance? (open text)
- A sequence of programming exercises A1, A2, and A3 to address a simple sub-problems in 1) adversarial search, 2) reinforcement learning, and 3) probabilistic inference. They are submitted to Kattis and presented in person. This assignment is in teams of up to two students. Each of these assignments yield grades A-F. For each of the programming exercises, a basic algorithm will be given as a starting point for the implementation, as well as a set of test cases for debugging.
- A project, whose goal is to implement an intelligent system successfully playing an arcade game, such as Pacman. The project is presented in person. This assignment is in teams of up to four students. The project is graded Pass/Fail. The students here build on their implementations from assignments A1 – A3.

- An essay on ethical and societal aspects of AI. This assignment involves an individual part including reading and own reflection, and a reflection based on discussion with team of four students. The essay is graded Pass/Fail.

There is no exam in this course.

ILO		*Q1 - Q9*	*A1*	*A2*	*A3*	*Project*	*Essay*
1	interpret various concepts in Artificial Intelligence (AI)	X					
2	develop effective and efficient solutions to address sub-problems from the AI domain		X	X	X		
3	integrate tools to design a computer program exhibiting multiple traits expected from an intelligent system					X	
4	explain own solution to a problem from the AI domain		X	X	X	X	
5	demonstrate an insight into societal and ethical aspects of AI						X

Grading criterion		*Assignment*
1	can interpret basic concepts including modeling techniques, algorithms, and other tools in all topics covered in the course	passes all basic quizzes in Canvas
	can interpret more advanced concepts including modeling techniques, algorithms, and other tools in all topics covered in the course	passes all advanced quizzes in Canvas
2	can solve selected sub-problems from the AI domain at a basic level (in terms of effectiveness, i.e. the output quality and efficiency, i.e. the code optimality) by implementing and debugging given algorithms	solves A1/A2/A3 by implementing the given simple algorithms and passes tests in Kattis with a required threshold of x points and computation budget of y seconds.
	can propose and implement simple modifications to the given algorithms	solves additional questions in A1/A2/A3 regarding the sources of ineffectiveness and

	resulting in a performance improvement, based on an evaluation of weaknesses of own implementation of the given algorithms in terms of effectiveness, i.e. the output quality and efficiency, i.e. the code optimality	inefficiency of the given simple algorithms, and optimizes the implementation to pass tests in Kattis with a required threshold of x' points and computation budget of y' seconds.
	can solve given sub-problems from the AI domain at a level significantly outperforming basic solutions in terms of effectiveness, i.e. the output quality and efficiency, i.e. the code optimality by selecting, implementing, and fine-tuning a combination of several algorithms, and can do so with limited resources in terms of time	solves A1/A2/A3 by implementing an own choice of algorithms and passes tests in Kattis with a required threshold of x'' points and computation budget of y'' seconds by the deadline
3	can design a computer program that is able to significantly improve the outcome of problem solving, planning, or decision making over time through the use of prediction or machine learning in an uncertain environment	solves project and gains required score in the arcade games
4	can clearly explain details of own implementation, principles of the implemented algorithms, and fitness of the algorithm to solving the problem	presents A1, A2, A3 orally, writes a comprehensible project report, and answers complementary questions regarding the presented solution and its principles asked at the presentation of A1, A2, A3 and the project
5	can independently reflect on ethical and societal aspects of AI, discuss these, and support arguments with references to relevant scientific studies	writes a comprehensible essay supported references to literature and writes a comprehensible summary of a discussion with colleagues

Combining grades

For passing grade, the student has to achieve a passing grade in ILOs 1-4, i.e. at worst P from QB, E from A1,A2, and A3, P from the Project, and P from Essay.

QA fail	WORSTOF(AVG(A1,A2,A3), "D")
QA pass	AVG(A1,A2,A3)

Part II: Reflection

I am the course responsible since HT17, and "inherited" the course after Patric Jensfelt, who, I think, designed excellent ILOs, grading criteria, and assignments. However, as the field of artificial intelligence rapidly evolves, and the number of students grows, we need to adjust the ILOs, design new assignments, and also grading criteria.

The new ILOs, criteria, and assignments were suggested while keeping in mind that the course has had received plenty of positive feedback in the past, but also that several weaknesses were repeatedly mentioned in the course evaluation:

- Previously, the course consisted of two large programming assignments. They were of very different difficulty and the difficulty between the grades varied a lot, too. For future, we aim to introduce three smaller and less time-consuming programming assignments, with less difference between their difficulty and demands for different grades. The grading criteria for ILO2 keep gradually growing demands for better grades.
- One of the frequent feedbacks we receive is that the quizzes are a great way to reflect on material, but some of the questions are too easy. Other students thought that the quizzes were too hard. Motivated by this, and by the fact that quizzes are essentially the only individual assignment in the course, we would like to give them more importance, and develop a set of more advanced questions. For that reason, we introduced two-stage grading criteria for quizzes here.
- Another feedback that we often hear from our students is that in fact, they do not get to develop an "intelligent" system within the course. For that reason, I propose ILO3 and suggest to replace the previous project that was aimed again at a particular sub-area of AI with an integration project.

These three major changes in the assignment structure naturally triggered the need for the update of grading criteria and ILOs, to maintain a meaningful alignment, and motivated my suggestions.

Following the feedback I received on the previous version of this report, the most notable change with respect to the first draft is that now all ILOs are aligned with an E-level grading criterion and a corresponding assignment. Previously, I suggested to have ILO3 to be graded at A-level only; however, it became apparent that this ILO is indeed desirable to meet for all levels.

In comparison how the grading criteria are used today (posted on Canvas with the hope that students will go through them), I plan to use the new grading criteria more actively. As the assignments are introduced gradually in this course, I plan to introduce them coupled with the course ILOs, grading criteria, and precise details on how that grading criteria project onto the grading of the assignment.

With the system of combining grades, I believe that the students will be motivated to invest energy into selected parts of the course that they are interested in and "compensate" for parts that they are not. This fits well the large breadth of topics covered in the course and the fact that the assignments will be of similar difficulty. I believe that this scheme will have a better effect on students' learning than having the "worst-of" or "best-of" grade combination, potentially leading to early giving up once an extremely bad/good grades are achieved. On the other hand, the fact that the students cannot get a grade higher than D if they do not pass the advanced quizzes aims to encourage them to attain deeper understanding of theory and principles individually next to being able to apply the techniques in practice. This suggestion also aims to mitigate the missing exam.