272SM: Introduction to Artificial Intelligence
Homework Assignment 1: Introduction/Agents

Thursday 5th October, 2023 - Sunday 22nd October, 2023

Instructions

Report your answers to the following exercises and submit them. Submitting homework solutions is highly recommended; Correct solutions will be counted as bonus points towards the final grade. You can team up with your colleagues into groups of max. 5 people.

Exercise 1

Read Turing's original paper on AI (Turing, 1950). In the paper, he discusses several objections to his proposed enterprise and his test for intelligence. Which objections still carry weight? Are his refutations valid? Can you think of new objections arising from developments since he wrote the paper? In the paper, he predicts that, by the year 2000, a computer will have a 30% chance of passing a five-minute Turing Test with an unskilled interrogator. What chance do you think a computer would have today? In another 25 years?

Link: https://redirect.cs.umbc.edu/courses/471/papers/turing.pdf

Exercise 2

Choose one of the following tasks:

- Study the 2023 EU "Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)" (or its final version, if available). Which provisions appear to have the most direct and tangible impact on what kinds of AI systems can be deployed? (link: https://www.europarl.europa.eu/doceo/document/
9-2023-0188$_E N.html$)
- Summarize the pros and cons of allowing the development, deployment, and use of lethal autonomous weapons.
- Many researchers have pointed to the possibility that machine learning algorithms will produce classifiers that display racial, gender, or other forms of bias. How does this bias arise? Is it possible to constrain machine learning algorithms to produce rigorously fair predictions?

Prepare to present your answers/discussion in about 5 minutes in class (using blackboard or slides is allowed).

Exercise 3

Implement a simple reflex agent for the vacuum environment. Run the environment with this agent for all possible initial dirt configurations and agent locations. Record the performance score for each configuration and the overall average score. Extend the environment so to get differently shaped rooms, obstacles, dirt generation process, or other features you consider important.

(Hint: Use resources from https://github.com/aimacode/aima-python to get inspired.)