

HW4: Classic papers in artificial intelligence

This homework requires you to read and comment on three famous papers of artificial intelligence.

Alan Turing: Computing Machinery and Intelligence (1950)

Many people consider this paper as the initiator of the field that later became known as artificial intelligence. It introduces the “imitation game”, the evaluation technique later known as the Turing test. Overall, Turing argues that in about 50 years, computers will be able to pass this test, and presents arguments against a number of possible objections.

The version of the paper for the homework is from a 2009 edited book. In addition to the original paper, it has extensive footnotes from three different groups of researchers.

- You only need to read the original paper, not the footnotes. (But feel free to parse them for context).
- You can safely skip Sections 3.4 and 3.5. These sections present digital computers to a reader who, in 1950, needed explanations about what a computer is.

Questions

Write a 400 word essay about the paper. Make sure to:

- Describe things that impress you from this article.
- Describe things where you think that Turing was wrong.
- Describe things that you were surprised by in this article.

Rodney A. Brooks: Intelligence without representation (1987)

This paper is part of a series of papers in which Rodney Brooks challenged what it was the common wisdom of AI at that time. Instead of the symbolic, logic-based techniques that aimed to create high level reasoning, he aimed to replicate efficient behaviors as found animals such as insects. Understandably, this triggered a major dispute within the field of AI in the 1990s.

As part of the later work, Brooks co-founded iRobot - thus current robotic vacuums are a direct result of this paper.

Questions:

Write a 400 word essay about the paper. Make sure to:

- Describe whether you think that Brooks was right about the approach he proposed.
- What do you think about the relationship between what Brooks discusses in this paper and the current generation of AI products, such as ChatGPT? Is ChatGPT a validation or a refutation of the ideas in this paper?

Volodymyr Mnih et al.: Human level control through deep reinforcement learning (2015)

This paper introduced the world to the recent results of deep reinforcement learning (in particular, a variant of called Deep Q-network DQN). In a number of ways, this paper raised the awareness of the world to the work that had been done at DeepMind, a British startup just acquired by Google. The paper shows that DQN is able to learn control strategies that perform better than humans for a large number of Atari games.

Note that the paper is composed of the 4 page description, and an appendix containing the methods. In order to answer the questions, you will also need to at least skim through the appendix.

Questions:

Write a 400 word essay about the paper. In it:

- As you learned about Q-learning in class, describe what are the main differences between the method proposed in this paper and the original Q-learning we covered.
- Find a way to play or at least look at games of Breakout (for which DQN was very successful) and Montezuma's Revenge (for which it failed completely). Explain why do you think that this massive difference exists.