

National University of Singapore
School of Computing
CS3243: Foundations of Artificial Intelligence
Solutions for Tutorial 1

Readings: AIMA Chapters 1 & 2

1. Rich and Knight defined AI as “the study of how to make computers do things at which, at the moment, people are better”. Can you name a computer system that does things at which people are not doing better and hence the system cannot be called an AI system? Similarly, name a system that is not doing better than people and hence can qualify as an AI system. Do you think this definition is reasonable?

A database management system can keep track of huge amounts of data, involving query search and updating. Such a task is beyond the capability of any human. A DBMS is thus normally not considered an AI system.

Visual perception at the moment is still a very difficult task for computers, whereas human beings generally have no problem in understanding any scene captured by the eyes. A machine vision system thus qualifies to be called an AI system.

Rich and Knight’s definition of AI is quite reasonable, as they view human intelligence as far superior to the machines today. Hence they define AI based on “doing things at which, at the moment, people are better”. However, one may also argue that there are certain AI systems that are performing far better than people today. For instance, data mining makes use of an AI technique, namely, machine learning, but humans are not good at data mining tasks because of the huge amounts of data involved.

2. Is the Turing test reproducible? Is it amenable to mathematical analysis? Justify your answer.

The Turing test is very subjective. The interrogator’s assessment is subjective, and different interrogators may come to a different conclusion, or the same interrogator may come to a different conclusion at different times. Hence the test is not reliably reproducible. As it is not an objective and reproducible test, it is not amenable to mathematical analysis.

3. Consider a system that provides on-line translation of telephone conversation between English and Japanese speakers. Discuss its performance measure, environment, actuators, and sensors. Is its environment fully observable? deterministic? episodic? static? discrete? single agent? Justify your answer.

Performance measure: Accurate translation output (meaning preserving, grammatical, audible, unnoticeable delay)

Environment: human users and physical surroundings

Actuators: speaker

Sensors: microphone

Fully observable: No (the background noise may interfere at times)

Deterministic: No (the noise level may change in the environment)

Episodic: No (sentence translation may require knowledge of earlier translation of previous sentences), or Yes (if you are able to process the speech in segments that you know for a fact are independent – this will be difficult to prove)

Static: No (the human users may get impatient and walk away if the speech translator takes too long to produce a translation or is too inaccurate)

Discrete: No (speech signal)

Single agent: No (if you consider the human users as agents), or Yes (otherwise)

4. In the framework of Chapter 2, what is the difference between a performance measure and a utility function?

A performance measure is used by an outside observer to evaluate how successful an agent is. A utility function is used by an agent itself to evaluate how desirable states are. In our framework, an agent may have no explicit utility function (like a simple reflex agent), but there is always a performance measure.

5. Consider an agent which functions as a medical diagnosis system. Determine what type of agent design is the most appropriate for such an agent (simple reflex, model-based reflex, goal-based, or utility-based). Justify your answer.

The most appropriate design is a utility-based agent, since there are uncertainties and trade-offs involved (benefits of treatment vs. cost).

6. Weizenbaum's ELIZA program simulates the behaviour of a psychotherapist carrying out a conversation with a patient. It basically works by finding keywords in the user's input so as to fire certain rules based on the keywords. Which AI definition does ELIZA fit in? (Thinking humanly? Acting humanly? Thinking rationally? Acting rationally?) Discuss how an ELIZA-like system will behave, if it is modelled according to each of the four agent types, namely, "simple reflex agent", "model-based reflex agent", "goal-based agent", and "utility-based agent".

ELIZA fits in the "acting humanly" definition because it tries to mimic human behavior but it does not make use of any human thinking or reasoning to give an answer.

A "simple reflex agent" will simply reply by matching the first successful rule or any randomly selected successful rule, whereas a "model-based reflex agent" will try to keep track of the state to monitor, for instance, what replies had been used before so that it would not keep replying with the same answer to the same or similar inputs. A "goal-based agent" will try to meet the goal of giving appropriate psychiatric help to the user. A "utility-based agent" will consider various trade-offs among conflicting goals, like giving the quickest response or the least-cost treatment plan or the safest treatment plan, etc.

Key Points:

- Given a specified problem, you should be able to perform a PEAS analysis and also determine the type of agent (model) that is most suitable for solving the problem.
- Start considering the simplest models first, then subsequently more complicated models.
- For the goal-based agent, the agent must be able to determine that it has reached its goal (if not, try utility-based agent) and the knowledge of the goal should be somehow helpful in solving the problem.
- Consider the model-based agent, if knowledge of the goal is not necessary for solving the problem.
- If the environment is fully-observable, the reflex agent might work. If the environment is not fully-observable, you might need a model-based agent because you need to construct an estimate of the environment from the available inputs and then use your model to make your decisions.
- The models are there to help you start somewhere. They are not prescriptive, i.e. you can choose to organize your agent slightly differently if you are convinced that your way is better.