

Lecture Overview

- What is Artificial Intelligence?
- Agents acting in an environment
- Representations


What is Artificial Intelligence?

- Artificial Intelligence is the synthesis and analysis of computational agents that act intelligently.
- An agent is something that acts in an environment.
- An agent that acts intelligently if:
 - ▶ its actions are appropriate for its goals and circumstances
 - ▶ it is flexible to changing environments and goals
 - ▶ it learns from experience
 - ▶ it makes appropriate choices given perceptual and computational limitations

Goals of Artificial Intelligence?

- **Scientific goal:** to understand the principles that make intelligent behavior possible, in natural or artificial systems.
 - ▶ analyze natural and artificial agents;
 - ▶ formulate and test hypotheses about what it takes to construct intelligent agents;
 - ▶ design, build, and experiment with computational systems that perform tasks that require intelligence.
- **Engineering goal:** to specify methods for the design of useful, intelligent artifacts.
- Analogy between studying flying machines and thinking machines.

Agents acting in an environment



Example agent: robot

- **abilities:** movement, grippers, speech, facial expressions, . . .
- **observations:** vision, sonar, sound, speech recognition, gesture recognition, . . .
- **goals:** deliver food, rescue people, score goals, explore, . . .
- **past experiences:** effect of steering, slipperiness, how people move, . . .
- **prior knowledge:** what is important feature, categories of objects, what a sensor tell us, . . .

Example agent: teacher

- **abilities:** present new concept, drill, give test, explain concept, . . .
- **observations:** test results, facial expressions, errors, focus, . . .
- **goals:** particular knowledge, skills, inquisitiveness, social skills, . . .
- **past experiences:** prior test results, effects of teaching strategies, . . .
- **prior knowledge:** subject material, teaching strategies, . . .


Example agent: medical doctor

- **abilities:** operate, test, prescribe drugs, explain instructions, . . .
- **observations:** verbal symptoms, test results, visual appearance. . .
- **goals:** remove disease, relieve pain, increase life expectancy, reduce costs, . . .
- **past experiences:** treatment outcomes, effects of drugs, test results given symptoms. . .
- **prior knowledge:** possible diseases, symptoms, possible causal relationships. . .

Example agent: user interface

- **abilities:** present information, ask user, find another information source, filter information, interrupt, . . .
- **observations:** users request, information retrieved, user feedback, facial expressions. . .
- **goals:** present information, maximize useful information, minimize irrelevant information, privacy, . . .
- **past experiences:** effect of presentation modes, reliability of information sources, . . .
- **prior knowledge:** information sources, presentation modalities. . .

Representations



Example representations: machine language, C, Java, Prolog, natural language

Defining a Solution

Given a problem, what is a solution?

- An **optimal solution** to a problem is one that's the best solution according some measure of solution quality.
- A **satisficing solution** is one that's good enough.
- An **approximately optimal solution** is one whose measure of quality is close to the best that could theoretically be obtained.
- A **probable solution** is one that is likely to be a solution.

What do we want in a representation?

We want a representation to be

- rich enough to express the knowledge needed to solve the problem;
- as close to the problem as possible: compact, natural and maintainable;
- amenable to efficient computation;
able to express features of the problem that can be exploit for computational gain and able to trade off accuracy and computation time/space
- able to be acquired from people, data and past experiences.

Physical symbol system hypotheses

- A **symbol** is a meaningful physical pattern that can be manipulated.
- A **symbol system** creates, copies, modifies and destroys symbols.

Physical symbol system hypothesis:

- A physical symbol system has the necessary and sufficient means for general intelligent action.

Levels of abstraction

- The **knowledge level** is in terms of what an agent knows and what its goals are.
- The **symbol level** is a level of description of an agent in terms of what reasoning it is doing.

Reasoning is the computation required to determine what an agent should do.

- **Design time reasoning and computation** is carried out by the designer the agent.
- **Offline computation** is the computation done by the agent before it has to act.
Background knowledge and data \rightsquigarrow **knowledge base**.
- **Online computation** is the computation that's done by an agent between receiving information and acting.