

CSEN901

Introduction to Artificial Intelligence

Course Material

- **Text:** Stuart Russell and Peter Norvig (2003). *Artificial Intelligence: A Modern Approach*. Prentice Hall.
- **Internet:** met.guc.edu.eg/courses/Winter2017_CSEN901.aspx

Course Assessment

- Your grade in this course will be based on your scores in one midterm exam, one final exam, three programming projects, and your best four out of five quiz scores.
- Weights

Projects	20%
Quizzes	15%
Midterm Exam	25%
Final Exam	40%

Origins of Artificial Intelligence

Lecture 1

Outline

- 1 What is AI?
- 2 What Does it Take to be an (Intelligent) Agent?
- 3 The Turing Test
- 4 Structure of Intelligent Agents
- 5 Problem Solving by Search
- 6 Conclusion

Outline

- 1 What is AI?
- 2 What Does it Take to be an (Intelligent) Agent?
- 3 The Turing Test
- 4 Structure of Intelligent Agents
- 5 Problem Solving by Search
- 6 Conclusion

Answers

Charniak and McDermott

Artificial intelligence is the study of mental faculties through the use of computational models

Shapiro

Artificial intelligence is a field of science and engineering concerned with the computational understanding of what is commonly called intelligent behavior, and with the creation of artifacts that exhibit such behavior

Rich and Knight

The study of how to make computers do things at which, at the moment, people are better

Answers

Charniak and McDermott

Artificial intelligence is the study of mental faculties through the use of computational models

Shapiro

Artificial intelligence is a field of science and engineering concerned with the computational understanding of what is commonly called intelligent behavior, and with the creation of artifacts that exhibit such behavior

Rich and Knight

The study of how to make computers do things at which, at the moment, people are better

Answers

Charniak and McDermott

Artificial intelligence is the study of mental faculties through the use of computational models

Shapiro

Artificial intelligence is a field of science and engineering concerned with the computational understanding of what is commonly called intelligent behavior, and with the creation of artifacts that exhibit such behavior

Rich and Knight

The study of how to make computers do things at which, at the moment, people are better

Answers (Cont'd)

Russell and Norvig

AI is viewed as the study and construction of rational agents

- **Agent:** “anything that can be seen as perceiving its environment through sensors and acting upon that environment through effectors.”
- **Rational:** “A rational agent is one that does the right thing.”
 - Whatever that means?

Answers (Cont'd)

Russell and Norvig

AI is viewed as the study and construction of rational agents

- **Agent:** “anything that can be seen as perceiving its environment through sensors and acting upon that environment through effectors.”
- **Rational:** “A rational agent is one that does the right thing.”
 - Whatever that means?

Outline

- 1 What is AI?
- 2 What Does it Take to be an (Intelligent) Agent?**
- 3 The Turing Test
- 4 Structure of Intelligent Agents
- 5 Problem Solving by Search
- 6 Conclusion

Agents

- According to Russell and Norvig, agents are entities that can be viewed as **perceiving** and **acting upon** their environment.
- Agents use **sensors** to perceive, and **effectors** to act.

Example

- Animals.
- Robots.
- Calculators.
- Thermostats.
- Clocks.

Rational Agents

- Rational agents do “the right thing”.
- “Right”: Whatever causes the agent to be “most successful”.
- Need a measure of success.

Example: Coffee-Delivery Agent

- Operate in a suite of offices.
- Receives coffee from kitchenette and delivers to offices.
- What is a suitable measure of success?

Rationality vs. Success

- Note that rational \neq successful.
- Rationality depends not only on the performance measure, but also on:
 - 1 Agent's sensory capabilities (what it can sense).
 - Coffee-agents that detect obstacles visually vs. those that use collision detection.
 - 2 Agent's actuator capabilities (what it can do).
 - Coffee-agents that can heat coffee with a built-in coil vs. those that cannot.
 - 3 Agent's knowledge.
 - Omniscient agents, uninformed agents, misinformed agents.

Rationality vs. Success

- Note that rational \neq successful.
- Rationality depends not only on the performance measure, but also on:
 - 1 Agent's sensory capabilities (what it can sense).
 - Coffee-agents that detect obstacles visually vs. those that use collision detection.
 - 2 Agent's actuator capabilities (what it can do).
 - Coffee-agents that can heat coffee with a built-in coil vs. those that cannot.
 - 3 Agent's knowledge.
 - Omniscient agents, uninformed agents, misinformed agents.

Rationality vs. Success

- Note that rational \neq successful.
- Rationality depends not only on the performance measure, but also on:
 - ① Agent's sensory capabilities (what it can sense).
 - Coffee-agents that detect obstacles visually vs. those that use collision detection.
 - ② Agent's actuator capabilities (what it can do).
 - Coffee-agents that can heat coffee with a built-in coil vs. those that cannot.
 - ③ Agent's knowledge.
 - Omniscient agents, uninformed agents, misinformed agents.

Rationality vs. Success

- Note that rational \neq successful.
- Rationality depends not only on the performance measure, but also on:
 - 1 Agent's sensory capabilities (what it can sense).
 - Coffee-agents that detect obstacles visually vs. those that use collision detection.
 - 2 Agent's actuator capabilities (what it can do).
 - Coffee-agents that can heat coffee with a built-in coil vs. those that cannot.
 - 3 Agent's knowledge.
 - Omniscient agents, uninformed agents, misinformed agents.

Rationality vs. Success

- Note that rational \neq successful.
- Rationality depends not only on the performance measure, but also on:
 - 1 Agent's sensory capabilities (what it can sense).
 - Coffee-agents that detect obstacles visually vs. those that use collision detection.
 - 2 Agent's actuator capabilities (what it can do).
 - Coffee-agents that can heat coffee with a built-in coil vs. those that cannot.
 - 3 Agent's knowledge.
 - **Omniscient** agents, uninformed agents, misinformed agents.

Rationality vs. Success

- Note that rational \neq successful.
- Rationality depends not only on the performance measure, but also on:
 - 1 Agent's sensory capabilities (what it can sense).
 - Coffee-agents that detect obstacles visually vs. those that use collision detection.
 - 2 Agent's actuator capabilities (what it can do).
 - Coffee-agents that can heat coffee with a built-in coil vs. those that cannot.
 - 3 Agent's knowledge.
 - **Omniscient** agents, uninformed agents, misinformed agents.

Ideal Rational Agents

Russell and Norvig

For each possible percept sequence, an ideal rational agent should do whatever action is expected to maximize its performance measure, on the basis of the evidence provided by the percept sequence and whatever built-in knowledge the agent has.

- Choosing to perform a **knowledge-producing action** is part of rationality.
 - Ostrich-attitude is not very rational.

Autonomy

- Intelligent agents should also be **autonomous**:
 - Agent's behavior not completely based on built-in knowledge, but also on its own experience.
- Provide agent with enough built-in knowledge to get started, and a learning mechanism to allow it to derive knowledge from percepts (and other knowledge).

Outline

- 1 What is AI?
- 2 What Does it Take to be an (Intelligent) Agent?
- 3 The Turing Test**
- 4 Structure of Intelligent Agents
- 5 Problem Solving by Search
- 6 Conclusion

What is “Intelligence”?

- This is actually the main problem of AI.
 - If we know a precise-enough definition, AI would not be a problem.
 - At least not as intriguing as it is.
- How would we know that we’ve succeeded?
- Turing proposed a way.

Alan Turing

- In the first paragraph of his article *Computing Machinery and Intelligence*,¹ Turing considered the question:

“Can machines think?”



- On pain of having to struggle with the exact meaning of “think”, Turing proposed to replace the question with a related, albeit relatively unambiguous, one.
- This new question is based on what Turing calls **the imitation game**.

¹*Mind*, vol. 59, 1950.

The Imitation Game (I)

“It is played with three people, a man (A), a woman (B), and an interrogator (C) who may be of either sex. The interrogator stays in a room apart from the other two. The object of the game for the interrogator is to determine which of the other two is the man and which is the woman. He knows them by labels X and Y, and at the end of the game he says either ‘X is A and Y is B’ or ‘X is B and Y is A’. The interrogator is allowed to put questions to A and B thus:

C: Will X please tell me the length of his or her hair?

Now suppose X is actually A, then A must answer. It is A’s object in the game to try and cause C to make the wrong identification. His answer might therefore be

‘My hair is shingled, and the longest strands are about nine inches long.’ ”

(Turing, 1950, pp. 433–434)

The Imitation Game (II)

“In order that tones of voice may not help the interrogator the answers should be written, or better still, typewritten. The ideal arrangement is to have a teleprinter communicating between the two rooms. Alternatively the questions and answers can be repeated by an intermediary. The object of the game for the third player (B) is to help the interrogator. The best strategy for her is probably to give truthful answers. She can add such things as ‘I am the woman, don’t listen to him!’ to her answers, but it will avail nothing as the man can make similar remarks.” (Turing, 1950, pp. 434)

The Imitation Game (III)

“We now ask the question, ‘What will happen when a machine takes the part of A in this game?’ Will the interrogator decide wrongly as often when the game is played like this as he does when the game is played between a man and a woman? These questions replace our original, ‘Can machines think?’”
(Turing, 1950, pp. 434)

The Turing Test

- A machine/program is said to have passed the **Turing test** if it can fool the interrogator as often as the man does.
- By defining the test thus, Turing has set a busy agenda for AI research.
 - Passing the Turing test requires solving most of the major AI problems (natural language competence, reasoning, planning—the so-called “AI-complete problems”).
 - Extensions of the Turing test propose changing the setting of the game to account for the AI-complete problems of vision and robotics.
- Though AI researchers do not explicitly target the Turing test, it remains the ultimate measure of success in AI.

Outline

- 1 What is AI?
- 2 What Does it Take to be an (Intelligent) Agent?
- 3 The Turing Test
- 4 Structure of Intelligent Agents**
- 5 Problem Solving by Search
- 6 Conclusion

Agent Design

- To design an agent, we have to, first, specify four things:
 - 1 Performance measure.
 - 2 Environment.
 - 3 Actuators.
 - 4 Sensors.
- We, then, need to consider different agent structures.

Programs and Architecture

- Agent = Agent program + Architecture.
- Agent program maps percept sequences to actions.
- Architecture is whatever the agent program runs on.
- The architecture has three functions:
 - ① Maps sensory input to data structures that are made available to the program.
 - ② Runs the program.
 - ③ Maps program outputs onto signals to the effectors.

Skeleton Agent

function SKELETON-AGENT(*percept*) **returns** action
static *memory*

memory ← UPDATE-MEMORY(*memory*, *percept*)

action ← CHOOSE-BEST-ACTION(*memory*)

memory ← UPDATE-MEMORY(*memory*, *action*)

return *action*

Coffee-Delivery Agent

- Performance measure: Delivering coffee to correct room.
- Environment: Suite of offices.
- Actions: translate, rotate.
- Sensations: location (where am I now?), orientation (35° North-East?).

Sample Environment

- Assume the suite of offices is a grid.
- A room occupies a number of contiguous grid cells.
- Each wall has a door to the neighboring room, if there is one.

R1	R2	R2	R3
R4	R2	R5	R3
R6	R2	R7	R8
R9	R10	R8	R8

Goal Formulation

- Suppose agent is in R1.
- It needs to deliver coffee to R8.
- What does it mean to deliver coffee to the right room?
- Well, to be in R8.
- Thus, the goal is to put the world in a state in which the agent is in R8.
- The goal, thus, may be considered to be a **set of states**.

States: Abstraction

- What are the possible states of the world?
 - Uncountably many.
- Use **abstraction** to focus only on relevant aspects.
- How about percepts as states?
 - That is, a state is a pair: (location, orientation)

States: More Abstraction

- Well, neither the exact location, nor the exact orientation are relevant.
- What is relevant is
 - 1 the grid cell where the agent is, and
 - 2 whether it is facing north, south, east, or west.
- Thus, consider a state to be a pair: (cell, dir), where
 - cell $\in \{1, 2, \dots, 16\}$ (assuming, for example, row-major indexing), and
 - dir $\in \{N, S, E, W\}$

Actions

- We can also use abstraction in formalizing possible actions.
- Instead of continuous translation and rotation, consider
 - ① Forward: “translate” till the next facing cell in the grid.
 - What if there is no such cell?
 - ② Right: “rotate” 90° clock-wise.
 - It is 12 O’clock when the agent is facing north.
 - ③ Left: “rotate” 90° counter clock-wise.
- This way, if the agent starts out facing one of the principal directions, it will also be facing one of them after performing an action.

The Goal

- The goal, now, corresponds to the set of states:
 $\{ (12, N) (12, S) (12, E)(12, W)$
 $(15, N) (15, S) (15, E)(15, W)$
 $(16, N) (16, S) (16, E)(16, W) \}$
- Why didn't we choose location to be room-, rather than cell-, number?

Problem Types

- Depending on the agent's knowledge and the environment, search problems may be divided into four types:
 - ① Single-state.
 - ② Multiple-State.
 - ③ Contingency.
 - ④ Exploration.
- Contingency and exploration problems require **interleaving** search and execution. We discuss these under the topics of **planning** and **learning**.

Single-State Problems

A single next-state for every state-action pair

- The agent's sensors provide enough information for the agent to tell which state it is in.
- The agent knows the effects of all its actions.

Multiple-State Problem

At any time, the agent is not certain about which state in a given set of states it is in.

- Perfect knowledge of effects of actions, but insufficient state-sensing capabilities.
 - Coffee-agent with only cell information or only direction information (or neither).
- Alternatively, the agent may have perfect sensors, but uncertainty about the effects of its actions.
 - Sometimes translation moves the agent backwards.

The Search Tree

- The search process may be thought of as building a tree through the state space.
- The root of the **search tree** is the initial state of the problem.
- A leaf of the search tree is a node which has either
 - 1 not yet been expanded, or
 - 2 been expanded, but has no successors.

The State Space and the Search Tree

- The state space is *not* the search tree.
 - ① Topologically, the state space is a directed graph; the search tree is a tree.
 - ② Nodes of the search tree are not mere states; they typically carry more information.
 - ③ The state space is finite; the search tree may be infinite.

Search Tree Nodes

- Following Russell and Norvig, we shall consider nodes to be 5-tuples:
 - 1 The **state** of the state space that this node corresponds to.
 - 2 The **parent node**.
 - 3 The **operator** applied to generate this node.
 - 4 The **depth** of the node in the tree.
 - 5 The **path cost** from the root.

General Search

```
function GENERAL-SEARCH(problem, QING-FUN)
  returns a solution, or failure
  nodes  $\leftarrow$  MAKE-Q(MAKE-NODE(INIT-STATE(problem)))
  loop do
    If nodes is empty then return failure
    node  $\leftarrow$  REMOVE-FRONT(nodes)
    If GOAL-TEST(problem)(STATE(node)) then return node
    nodes  $\leftarrow$  QING-FUN(nodes, EXPAND(node, OPER(problem)))
  end
```

Outline

- 1 What is AI?
- 2 What Does it Take to be an (Intelligent) Agent?
- 3 The Turing Test
- 4 Structure of Intelligent Agents
- 5 Problem Solving by Search
- 6 Conclusion**

Turing's Vision

“I believe that in about fifty years' time it will be possible to programme computers, with a storage capacity of about 10^9 , to make them play the imitation game so well that an average interrogator will not have more than 70 per cent. chance of making the right identification after five minutes of questioning. [. . .] I believe that at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted.”
(Turing, 1950, p. 442)