


ذكاء اصطناعي

Artificial Intelligence (AI)

Comp 408

Book:



Artificial Intelligence :

A Modern Approach

Second Edition by Stuart J. Russell & Peter Norvig

(2003, Prentice Hall)

- Lectures, Hws, ...etc will be available on
- <http://scholar.cu.edu.eg/?q=areegsaid/node/25645>

AI

- G. F. Luger & W. F. Stubblefield (1993),
G. F. Luger (2005)
- AI is the branch of computer science concerned with the automation of intelligent behavior.

The sources of studying AI

- Philosophy: Logic, methods of reasoning
- Mathematics: Formal representation and proof algorithms, (un)decidability, probability
- Economics: Decision theory
- Neuroscience: physical state for mental activity
- Psychology: perception, experimental techniques
- Linguistics, psycholinguistics
- Control System: design systems that maximize an objective function

Tentative syllabus for an AI course

Introduction and Agents (chapters 1,2)

Search (chapters 3,4,5,6)

Logic (chapters 7,8,9)

Natural Language Processing (chapter 22,23)

Remark: Comp 408 will not necessarily finish the above

Grades

60 % Final


10 % Midterm

10 % Quiz

20 % Lab Assignments/ Quiz/ project

Hw will be assigned frequently to help you study.
But no grades on it (you know why !)

Disciplines which form the core of AI inside
Fields which draw from these discipline outside.



Human vs Machine Intelligence

- Intelligence is not a unique and unshared capability of human. It is more an open collection of attributes than it is a single well-defined entity
- Humans embody many aspects of intelligence while animals typically embody a smaller number of intelligent characteristics, and usually at a much lower level
- The advent of digital computers made possible credible attempts to fulfill the AI dreams
- Computer based intelligence must be specialized to very *restricted* domains to be comparable to human performance

7

Human vs Machine Intelligence

- The two hemispheres of the human brain deal with problems in two distinct paradigms:
 - sequential (or logical) approach that considers only a small portion of the available data at a time
 - parallel processing looks at data on a global basis

Many tasks which we might reasonably think require **intelligence** are performed by computers **without even thinking**

Complex Arithmetic

Other tasks that people do without thinking are extremely difficult to automate

Recognizing a Face

8

What is AI

For Defining AI, Books go in four directions:

- Think Like a human
- Act Like a human
- Think rationally
- Act rationally

The textbook adopts "acting rationally"

Thinking humanly

- 1960s "cognitive revolution"
 - How the computer performs tasks does matter
 - The reasoning steps are important
 - Requires scientific theories of the activities of the brain
 - Need ability to manipulate symbols
 - How to validate? Requires
 - 1) Predicting and testing behavior of human subjects (top-down)--- Cognitive Science or
 - 2) Direct identification from neurological data (bottom-up)
cognitive neuroscience
- Both approaches are now distinct from AI Approaches

Thinking rationally

Aristotle: what are correct arguments processes?

- Logic and reasoning
- Requires reasoning structures
- Not easy to represent informal information
- Approaches: Machine Learning, NN, Reasoning
- Problems:
 - ✗ Not all things can be formally represented in logic notation, particularly if there is any uncertainty
 - ✗ It is usually the case that even small scale problems can exhaust the computational power of any computer unless heuristics are used

Acting humanly

- Turing (1950) "Computing machinery and intelligence"
- "Can machines think?" → "Can machines behave intelligently?"
- Take a task that people normally do : playing chess, diagnosis a disease, navigating a building ... etc
then build a computer system that does it automatically

Acting humanly

To pass the Turing Test, the computer would need:

1. Natural Language Processing → Communication
2. Knowledge Representation → store info before and during interrogation
3. Automated Reasoning → answer questions and draw new conclusions
4. Machine learning → adapt to new circumstances

Acting Rationally: rational agent

- Rational behavior: doing the right / best thing
- The right thing: maximize goal achievement, given the available information
- Does not necessarily involve thinking
- Act rationally = reason logically to the conclusion and act on that conclusion

The Rational Agent Approach

- An agent is something that perceives and acts
- A rational agent is one that acts so as to achieve the best outcome

AI Languages

- Programming languages best suited to AI tasks are **Lisp** (1960) and **Prolog** (1972).
 - **PRO**gramming in **LOG**ic
 - Emphasis on *what* rather than *how*
- There also have been specialized knowledge representation systems and languages, used to develop knowledge bases and knowledge-based systems. This includes expert systems, in which probability and beliefs play an important role.

Prolog's strong and weak points

- Assists thinking in terms of *objects* and *entities*
- *Straight forward translation of logic statements*
- ⊗ **Not good for *number manipulation***
- Useful applications of Prolog in
 - *Expert Systems* (Knowledge Representation and Inferencing)
 - *Natural Language Processing*
 - *Relational Databases*


Back to AI

Agent

Agents

- An **agent** is anything that can be viewed as **perceiving** its **environment** through **sensors** and **acting** upon that environment through **actuator**
- The complete set of inputs at a given time is called a **percept**
- The current percept, or a sequence of percepts can influence the actions of an agent.
- The agent can change the environment through actuators.
- An operation involving an actuator is called an action. Actions can be grouped into action sequences.
- The agent can have goals which it tries to achieve.

Agents and environments



- An agent can be looked upon as a system that implements a mapping from percept sequences to actions

$$[f: P^* \rightarrow \mathcal{A}]$$

- The **agent program** runs on the physical **architecture** to produce f
- agent = architecture + program

Examples of Agents


Humans: senses, body parts move

Programs: keyboard, mouse, speakers

Robots: cameras, motors

- An Intelligent Agent must **sense**, must **act**, must be **autonomous** (to some extent), must be **rational**
 - A **rational agent** is one that does the right thing
1. What are the functionalities (goals)?
 2. What are the components?
 3. How do we build them?

Vacuum-cleaner world



- Percepts: location and contents, e.g., [A,Dirty]
- Actions: *Left*, *Right*, *Suck*, *NoOp*

Rational agents

- An agent should strive to "do the right thing", based on what it can perceive and the actions it can perform. The right action is the one that will cause the agent to be most successful
- Performance measure: An objective criterion for success of an agent's behavior
- E.g., performance measure of a vacuum-cleaner agent could be amount of dirt cleaned up, amount of time taken, amount of electricity consumed, amount of noise generated, etc.

Rational agents

- **Rational Agent:** For each possible percept sequence, a rational agent should select an action that is expected to maximize its performance measure, given the evidence provided by the percept sequence and whatever built-in knowledge the agent has.
- **Rational Action** is the action that maximizes the expected value of the performance measure given the percept sequence .

Bounded Rationality

- Perfect Rationality assumes that the rational agent will take the action that maximizes its utility.
- Human beings do not satisfy this definition of rationality 😊
- Herbert Simon, 1972 "Because of the limitations of the human mind, humans must use approximate methods to handle many tasks."
- The agent has to select the best action to the best of its knowledge depending on its percept sequence, its background knowledge and its feasible actions.
- Therefore rationality must take into account the limitations of the agent.

Rational vs. Omniscient Agents

- An omniscient agent knows the *actual outcome of its actions and can act accordingly*;
- *omniscience is impossible in reality*
- Our definition of rationality does not require omniscience because the rational choice depends only on the percept sequence to date
- Rationality is distinct from omniscience (all-knowing with infinite knowledge)

learning and autonomy

- Agents can perform actions in order to modify future percepts so as to obtain useful information (information gathering, exploration)
- An agent is **autonomous** if its behavior is determined by its own experience (with ability to learn and adapt)

Summary of agents kinds

- A **rational** agent is one that does the right thing
- The **omniscience** of an agent knows the *actual outcome assuming infinite knowledge*
- The **autonomy** of an agent is the extent to which its behaviour is determined by its own experience