






# Introduction to Artificial Intelligence

ACCA6011- Sem 1  
Week 1- Lecture 1




# ACCA6011

- 
- Textbook: S. Russell and P. Norvig *Artificial Intelligence: A Modern Approach* Prentice Hall, 2003, **Second Edition**
- Assessment: Coursework and Exam (50% each)



# Course overview

- Introduction and Agents (chapters 1,2)
- Search (chapters 3,4)
- Games (chapter 5)
- Constraints processing (chapter 6)
- Representation and Reasoning with Logic (chapters 7,8,9)
- Learning (chapters 18,20)
- Planning (chapter 11)
- Uncertainty (chapters 13,14)
- Natural Language Processing (chapter 22,23)




# Course Outline

## Resources on the Internet

- [AI on the Web](#): A very comprehensive list of Web resources about AI from the Russell and Norvig textbook.

## Essays and Papers

- [What is AI](#), John McCarthy
- [Computing Machinery and Intelligence](#), A.M. Turing
- [Rethinking Artificial Intelligence](#), Patrick H. Winston




# Today's class


- What is Artificial Intelligence?
- A brief History
- Intelligent agents
- State of the art

# What is Artificial Intelligence

## (John McCarthy, Basic Questions)



- **What is artificial intelligence?**
- It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.
- **Yes, but what is intelligence?**
- Intelligence is the computational part of the ability to achieve goals in the world. Varying kinds and degrees of intelligence occur in people, many animals and some machines.
- **Isn't there a solid definition of intelligence that doesn't depend on relating it to human intelligence?**
- Not yet. The problem is that we cannot yet characterize in general what kinds of computational procedures we want to call intelligent. We understand some of the mechanisms of intelligence and not others.
- **More in:** <http://www-formal.stanford.edu/jmc/whatisai/node1.html>




# What is AI?

Views of AI fall into four categories:

|                |                   |
|----------------|-------------------|
| Think Humanly  | Think rationally  |
| Acting Humanly | Acting rationally |


The textbook advocates "acting rationally"

[List of AI-topics](#)



# What is Artificial Intelligence?

- **Human-like** (“How to simulate humans intellect and behavior on by a machine.”)
  - Mathematical problems (puzzles, games, theorems)
  - Common-sense reasoning (*if there is parking-space, probably illegal to park*)
  - Expert knowledge: lawyers, medicine, diagnosis
  - Social behavior
- **Rational-like:**
  - achieve goals, have performance measure



# What is Artificial Intelligence

- **Thought processes**


- “The exciting new effort to make computers think .. Machines with minds, in the full and literal sense” (Haugeland, 1985)

- **Behavior**

- “The study of how to make computers do things at which, at the moment, people are better.” (Rich, and Knight, 1991)


# The Turing Test

(Can Machine think? A. M. Turing, 1950)




## ○ Requires

- Natural language
- Knowledge representation
- Automated reasoning
- Machine learning
- (vision, robotics) for full test



# What is AI?

- Turing test (1950)
- Requires:
  - Natural language
  - Knowledge representation
  - automated reasoning
  - machine learning
  - (vision, robotics.) for full test
- Thinking humanly:
  - Introspection, the general problem solver (Newell and Simon 1961)
  - Cognitive sciences
- Thinking rationally:
  - Logic
  - Problems: how to represent and reason in a domain
- Acting rationally:
  - Agents: Perceive and act



# AI examples

## Common sense reasoning

- Tweety
- Yale Shooting problem

## Update vs revise knowledge

- The OR gate example:  $A \text{ or } B \rightarrow C$
- Observe  $C=0$ , vs Do  $C=0$

## Chaining theories of actions

Looks-like(P)  $\rightarrow$  is(P)

Make-looks-like(P)  $\rightarrow$  Looks-like(P)


-----

Makes-looks-like(P)  $\rightarrow$  is(P) ???

Garage-door example: garage door not included.


- Planning benchmarks
- 8-puzzle, 8-queen, block world, grid-space world

Abduction: cambridge parking example




# History of AI

- McCulloch and Pitts (1943)
  - Neural networks that learn
- Minsky (1951)
  - Built a neural net computer
- Darmouth conference (1956):
  - McCarthy, Minsky, Newell, Simon met,
  - Logic theorist (LT)- proves a theorem in Principia Mathematica-Russel.
  - The name “Artificial Intelligence” was coined.
- 1952-1969
  - GPS- Newell and Simon
  - Geometry theorem prover - Gelernter (1959)
  - Samuel Checkers that learns (1952)
  - McCarthy - Lisp (1958), Advice Taker, Robinson’s resolution
  - Microworlds: Integration, block-worlds.
  - 1962- the perceptron convergence (Rosenblatt)




# The Birthplace of “Artificial Intelligence”, 1956

- **Darmouth workshop, 1956:** historical meeting of the precieved founders of AI met: John McCarthy, Marvin Minsky, Alan Newell, and Herbert Simon.
- **A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence. J. McCarthy, M. L. Minsky, N. Rochester, and C.E. Shannon. August 31, 1955. "We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it." *And this marks the debut of the term "artificial intelligence."***
- **50 anniversery of Darmouth workshop**




# History, continued

- 1966-1974 a dose of reality
  - Problems with computation
- 1969-1979 Knowledge-based systems
  - Weak vs. strong methods
  - Expert systems:
    - Dendral: Inferring molecular structures
    - Mycin: diagnosing blood infections
    - Prospector: recommending exploratory drilling (Duda).
  - Roger Shank: no syntax only semantics
- 1980-1988: AI becomes an industry
  - R1: Mcdermott, 1982, order configurations of computer systems
  - 1981: Fifth generation
- 1986-present: return to neural networks
- Recent event:
  - AI becomes a science: HMMs, planning, belief network




# Abridged history of AI

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1956 Dartmouth meeting: "Artificial Intelligence" adopted
- 1952—69 Look, Ma, no hands!
- 1950s Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1965 Robinson's complete algorithm for logical reasoning
- 1966—73 AI discovers computational complexity  
Neural network research almost disappears
- 1969—79 Early development of knowledge-based systems
- 1980-- AI becomes an industry
- 1986-- Neural networks return to popularity
- 1987-- AI becomes a science
- 1995-- The emergence of intelligent agents




# State of the art

- Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997
- Proved a mathematical conjecture (Robbins conjecture) unsolved for decades
- No hands across America (driving autonomously 98% of the time from Pittsburgh to San Diego)
- During the 1991 Gulf War, US forces deployed an AI logistics planning and scheduling program that involved up to 50,000 vehicles, cargo, and people
- NASA's on-board autonomous planning program controlled the scheduling of operations for a spacecraft
- `Proverb` solves crossword puzzles better than most humans
- DARPA grand challenge 2003-2005, Robocup



# Agents (chapter 2)


- Agents and environments
- Rationality
- PEAS (Performance measure, Environment, Actuators, Sensors)
- Environment types
- Agent types



# Agents

- An **agent** is anything that can be viewed as **perceiving** its **environment** through **sensors** and **acting** upon that environment through **actuators**
- Human agent: eyes, ears, and other organs for sensors; hands,
- legs, mouth, and other body parts for actuators
- Robotic agent: cameras and infrared range finders for sensors;
- various motors for actuators

# Agents and environments




- The **agent function** maps from percept histories to actions:


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

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

# 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.



# What's involved in Intelligence?

## Intelligent agents


- Ability to interact with the real world
  - to perceive, understand, and act
  - e.g., speech recognition and understanding and synthesis
  - e.g., image understanding
  - e.g., ability to take actions, have an effect
- Knowledge Representation, Reasoning and Planning
  - modeling the external world, given input
  - solving new problems, planning and making decisions
  - ability to deal with unexpected problems, uncertainties
- Learning and Adaptation
  - we are continuously learning and adapting
  - our internal models are always being “updated”
    - e.g. a baby learning to categorize and recognize animals



# Implementing agents

- Table look-ups
- Autonomy
  - All actions are completely specified
  - no need in sensing, no autonomy
  - example: Monkey and the banana
- Structure of an agent
  - agent = architecture + program
  - Agent types
    - medical diagnosis
    - Satellite image analysis system
    - part-picking robot
    - Interactive English tutor
    - cooking agent
    - taxi driver

## Vacuum-cleaner world



Percepts: location and contents, e.g.,  $[A, \textit{Dirty}]$

Actions: *Left*, *Right*, *Suck*, *NoOp*

## A vacuum-cleaner agent

| Percept sequence              | Action       |
|-------------------------------|--------------|
| <i>[A, Clean]</i>             | <i>Right</i> |
| <i>[A, Dirty]</i>             | <i>Suck</i>  |
| <i>[B, Clean]</i>             | <i>Left</i>  |
| <i>[B, Dirty]</i>             | <i>Suck</i>  |
| <i>[A, Clean], [A, Clean]</i> | <i>Right</i> |
| <i>[A, Clean], [A, Dirty]</i> | <i>Suck</i>  |
| <i>⋮</i>                      | <i>⋮</i>     |

**function** REFLEX-VACUUM-AGENT( *[location, status]*) **returns** an action

**if** *status = Dirty* **then return** *Suck*  
**else if** *location = A* **then return** *Right*  
**else if** *location = B* **then return** *Left*

What is the **right** function?

Can it be implemented in a small agent program?

## PEAS

To design a rational agent, we must specify the **task environment**

Consider, e.g., the task of designing an automated taxi:

Performance measure??

Environment??

Actuators??

Sensors??

## PEAS

To design a rational agent, we must specify the **task environment**

Consider, e.g., the task of designing an automated taxi:

Performance measure?? safety, destination, profits, legality, comfort, ...

Environment?? US streets/freeways, traffic, pedestrians, weather, ...

Actuators?? steering, accelerator, brake, horn, speaker/display, ...

Sensors?? video, accelerometers, gauges, engine sensors, keyboard, GPS, ...

## Environment types

|                        | Solitaire | Backgammon | Internet shopping | Taxi |
|------------------------|-----------|------------|-------------------|------|
| <u>Observable??</u>    |           |            |                   |      |
| <u>Deterministic??</u> |           |            |                   |      |
| <u>Episodic??</u>      |           |            |                   |      |
| <u>Static??</u>        |           |            |                   |      |
| <u>Discrete??</u>      |           |            |                   |      |
| <u>Single-agent??</u>  |           |            |                   |      |

## Environment types

|                        | Solitaire | Backgammon | Internet shopping | Taxi |
|------------------------|-----------|------------|-------------------|------|
| <u>Observable??</u>    | Yes       | Yes        | No                | No   |
| <u>Deterministic??</u> |           |            |                   |      |
| <u>Episodic??</u>      |           |            |                   |      |
| <u>Static??</u>        |           |            |                   |      |
| <u>Discrete??</u>      |           |            |                   |      |
| <u>Single-agent??</u>  |           |            |                   |      |

## Environment types

|                        | Solitaire | Backgammon | Internet shopping | Taxi |
|------------------------|-----------|------------|-------------------|------|
| <u>Observable??</u>    | Yes       | Yes        | No                | No   |
| <u>Deterministic??</u> | Yes       | No         | Partly            | No   |
| <u>Episodic??</u>      |           |            |                   |      |
| <u>Static??</u>        |           |            |                   |      |
| <u>Discrete??</u>      |           |            |                   |      |
| <u>Single-agent??</u>  |           |            |                   |      |

## Environment types

|                        | Solitaire | Backgammon | Internet shopping | Taxi |
|------------------------|-----------|------------|-------------------|------|
| <u>Observable??</u>    | Yes       | Yes        | No                | No   |
| <u>Deterministic??</u> | Yes       | No         | Partly            | No   |
| <u>Episodic??</u>      | No        | No         | No                | No   |
| <u>Static??</u>        |           |            |                   |      |
| <u>Discrete??</u>      |           |            |                   |      |
| <u>Single-agent??</u>  |           |            |                   |      |

## Environment types

|                        | Solitaire | Backgammon | Internet shopping | Taxi |
|------------------------|-----------|------------|-------------------|------|
| <u>Observable??</u>    | Yes       | Yes        | No                | No   |
| <u>Deterministic??</u> | Yes       | No         | Partly            | No   |
| <u>Episodic??</u>      | No        | No         | No                | No   |
| <u>Static??</u>        | Yes       | Semi       | Semi              | No   |
| <u>Discrete??</u>      |           |            |                   |      |
| <u>Single-agent??</u>  |           |            |                   |      |

## Environment types

|                        | Solitaire | Backgammon | Internet shopping | Taxi |
|------------------------|-----------|------------|-------------------|------|
| <u>Observable??</u>    | Yes       | Yes        | No                | No   |
| <u>Deterministic??</u> | Yes       | No         | Partly            | No   |
| <u>Episodic??</u>      | No        | No         | No                | No   |
| <u>Static??</u>        | Yes       | Semi       | Semi              | No   |
| <u>Discrete??</u>      | Yes       | Yes        | Yes               | No   |
| <u>Single-agent??</u>  |           |            |                   |      |

## Environment types

|                        | Solitaire | Backgammon | Internet shopping    | Taxi |
|------------------------|-----------|------------|----------------------|------|
| <u>Observable??</u>    | Yes       | Yes        | No                   | No   |
| <u>Deterministic??</u> | Yes       | No         | Partly               | No   |
| <u>Episodic??</u>      | No        | No         | No                   | No   |
| <u>Static??</u>        | Yes       | Semi       | Semi                 | No   |
| <u>Discrete??</u>      | Yes       | Yes        | Yes                  | No   |
| <u>Single-agent??</u>  | Yes       | No         | No (except auctions) | No   |

The environment type largely determines the agent design


The real world is (of course) partially observable, stochastic, sequential, dynamic, continuous, multi-agent

## Agent types

Four basic types in order of increasing generality:


- simple reflex agents
- reflex agents with state
- goal-based agents
- utility-based agents

All these can be turned into learning agents



# Agent types

- Example: Taxi driver
- Simple reflex
  - **If** car-in-front-is-breaking **then** initiate-breaking
- Agents that keep track of the world
  - If car-in-front-is-breaking **and** on fwy **then** initiate-breaking
  - needs internal state
- goal-based
  - If car-in-front-is-breaking **and** needs to get to hospital **then** go to adjacent lane and plan
  - search and planning
- utility-based
  - If car-in-front-is-breaking **and** on fwy **and** needs to get to hospital alive **then** search of a way to get to the hospital that will make your passengers happy.
  - Needs utility function that map a state to a real function (am I happy?)



# Summary

- What is Artificial Intelligence?
  - modeling humans thinking, acting, should think, should act.
- History of AI
- Intelligent agents
  - We want to build agents that act rationally
- Real-World Applications of AI
  - AI is alive and well in various “every day” applications
    - many products, systems, have AI components
- Assigned Reading
  - Chapters 1 and 2 in the text R&N