Introduction to Artificial Intelligence

Prof. Dechter

ICS 271

Fall 2012

Our trip to Namibia and AI

Examples of thinking/acting
The flat tire scenario

- We drove on unpaved bumpy rocky road…. then we heard a bump …
- Will we have a flat tire?... We hear unpleasant noise (evidence)… we have a flat tire.
- facts… we have two flat tires…luckily we have two spare tires. Query: how could it be?
- Explanation: it probably happened when we went over the rock. Can it be that the same rock impact both tiers? Not likely
- Evidence: people are coming: is it good? Bad?... Good.
- Chana: be careful… (expects the worse)
- Rina: they want to help…… they helped (optimistic)
- Question: can it be that they put rocks on the road so that people will have flat tiers? So they can get some money from helping out?

On the way to Kowalib Lodge

- The miles calculation show that we should have been there… so how come we are in the middle of nowehere? What happened?
- Chana: we just missed a turn
- Amikam: maybe our kilometrage was wrong?,
- perhaps the information is not exact?
- Chana: we made a mistake: I now remember that there were turns… we are not on road 35. I am sure of it.
- Rina: Lets go a little further, and if we see nothing we will ask… (we found it a little further)

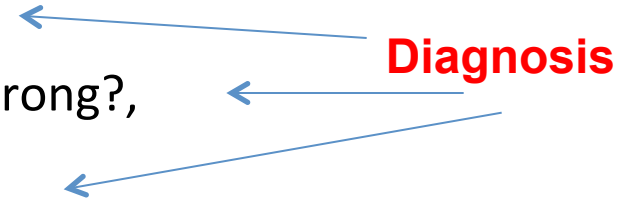

Who does the diagnosis?
Rina, Chana?

Robot navigation and object recognition examples

- **Reasoning about navigation:** next day from Kowalib Lodge to the main road:
  - We wanted to follow the reverse path… then discovered something that did not look familiar. We turned back (but we asked first)
- **Object recognition:**
  - we saw from far… an elephant… a tied elephant… actually a simple tent

What's AI?
Examples from our trip

- (Knowledge + evidence)→ answer query
- Knowledge can be deterministic (we could have been either is either in whiteok or Soseflei but not both),
- A constraint: we must get to the lodge while there is light.
- Or probabilistic: it will take us around an hour to to get to the lodge.
- Type of questions: commonsense, expert systems, playing games
- Predictions (depends by whom)? E.g., approaching a restaurant after 3… they will not give us Café now. If we will jump from a balcony we will fall on the ground

Course Outline

http://www.ics.uci.edu/~dechter/courses/ics-271/fall-12/

Course requirement

Assignments:

- There will be weekly homework-assignments, a project, a midterm or a final.

Course-Grade:

- Homeworks plus project will account for 50% of the grade, midterm or final 50% of the grade.

.

Course overview

- Introduction and Agents (chapters 1,2)
- Search (chapters 3,4,5,6)
- Logic (chapters 7,8,9)
- Planning (chapters 10,11)

new

Plan of the course

Part I Artificial Intelligence

1 Introduction

2 Intelligent Agents

Part II Problem Solving

3 Solving Problems by Searching

4 Beyond Classical Search

5 Adversarial Search

6 Constraint Satisfaction Problems

Part III Knowledge and Reasoning

7 Logical Agents

8 First-Order Logic

9 Inference in First-Order Logic

10 Classical Planning

11 Planning and Acting in the Real World

12 Knowledge Representation

Resources on the internet

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

- AI Topics: http://aitopics.net/index.php

Today's class

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

Today's class

- What is Artificial Intelligence?

What is Artificial Intelligence

(John McCarthy , Basic Questions

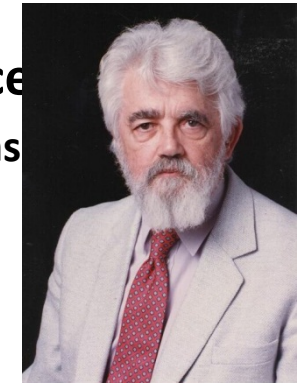

- **What is artificial intelligence?**
- It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically 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 Artificial Intelligence?

- Thought processes vs behavior
- Human-like vs rational-like
- How to simulate humans intellect and behavior by a machine.
  - Mathematical problems (puzzles, games, theorems)
  - Common-sense reasoning
  - Expert knowledge: lawyers, medicine, diagnosis
  - Social behavior

What is AI?

Views of AI fall into four categories:

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

The textbook advocates "acting rationally“

How to simulate humans intellect and behavior by a machine.
- Mathematical problems (puzzles, games, theorems)
- Common-sense reasoning
- Expert knowledge: lawyers, medicine, diagnosis
- Social behavior

The Turing Test

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

http://aitopics.net/index.php

http://amturing.acm.org/acm_tcc_webcasts.cfm

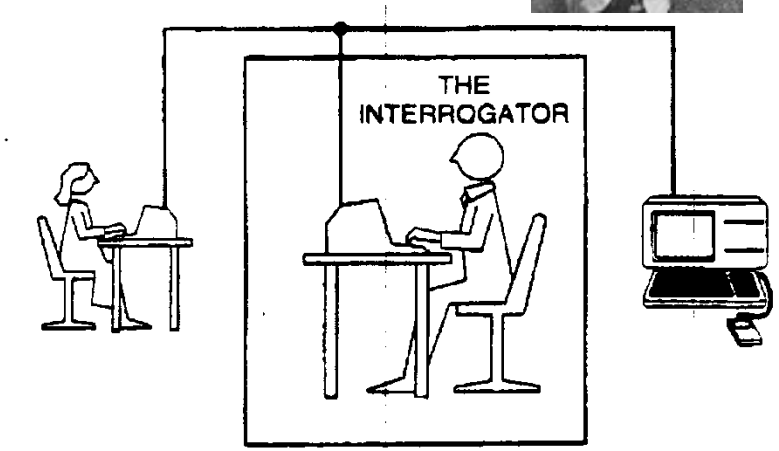

Figure 1.1 The Turing test.

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

Acting/Thinking Humanly/Rationally

- Turing test (1950)
- Requires:
  - Natural language
  - Knowledge representation
  - automated reasoning
  - machine learning
  - (vision, robotics.) for full test
- Methods for 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

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)

More AI examples

Common sense reasoning (1980-1990)

- Tweety
- Yale Shooting problem

Update vs revise knowledge

The OR gate example: A or B → C

- Observe C=0, vs Do C=0

Chaining theories of actions

Looks-like(P) → is(P)
Make-looks-like(P) → Looks-like(P)

Makes-looks-like(P) ---is(P) ???

Garage-door example: garage door not included.

- Planning benchmarks
- 8-puzzle, 8-queen, block world, grid-space world
- Cambridge parking example

Smoked fish example… what is this?

The foundation of AI

Philosophy, Mathematics, Economics,Neuroscience, Psychology, Computer Engineering,

Today's class

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

Histroy of AI

- McCulloch and Pitts (1943)
  - Neural networks that learn
- Minsky and Edmonds (1951)
  - Built a neural net computer
- Darmouth conference (1956):
  - McCarthy, Minsky, Newell, Simon met,
  - Logic theorist (LT)- Of Newell and Simon proves a theorem in Principia Mathematica-Russel.
  - The name "Artficial Intelligence" was coined.
- 1952-1969 (early enthusiasm, great expectations)
  - 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**
- List of AI-topics

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(Buchanan et. Al. 1969)
    - Mycin: diagnosing blood infections (Shortliffe et. Al, certainty factors)
    - Prospector: recomending 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
- 1987-present :
  - **AI becomes a science**: HMMs, planning, belief network
- 1995-present: The emergence of intelligent agents
  - Ai agents (SOAR, Newell, Laired, 1987) on the internet, technology in web-based **applications** , recommender systems. Some researchers (Nilsson, McCarthy, Minsky, Winston) express discontent with the progress of the field. AI should return to human-level AI (they say).
- 2001-present: The availability of data;
  - The knowledge bottleneck may be solved for many applications: learn the information rather than hand code it .

State of the art

- **Game Playing:** Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997
- **Robotics vehicles**: (Staneley (Thrun 2006). No hands across America (driving autonomously 98% of the time from Pittsburgh to San Diego)
- **Autonomous planning and scheduling**:
  - 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
- **Speech recognition**
- DARPA grand challenge 2003-2005, Robocup
- **Machine translation** (From English to arabic, 2007)
- **Natural language processing**: Watson won Jeopardy (Natural language processing), IBM 2011.

Robotic links

- Deep Blue: http://en.wikipedia.org/wiki/Deep_Blue_(chess_computer)
- **Robocup Video**
  - Soccer Robocupf
- Darpa Challenge
  - Darpa's-challenge-video
- Watson
- http://www.youtube.com/watch?v=seNkjYyG3gI

Today's class

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

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

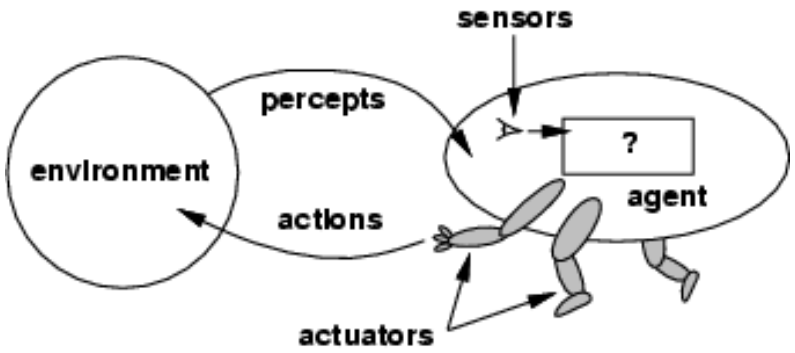

- The agent function maps from percept histories to actions:

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

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

- agent = architecture + program

What's involved in Intelligence?

- **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
    - **Graduate student**

Vacuum-cleaner world

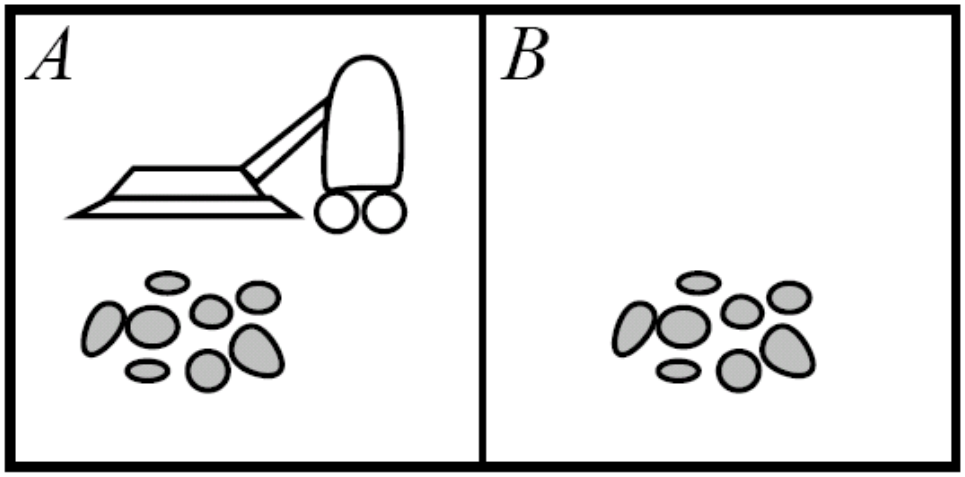

Percepts: location and contents, e.g., $[A, Dirty]$

Actions: $Left$, $Right$, $Suck$, $NoOp$

A vacuum-cleaner agent

Percept sequence	Action
$[A, Clean]$	$Right$
$[A, Dirty]$	$Suck$
$[B, Clean]$	$Left$
$[B, Dirty]$	$Suck$
$[A, Clean]$, $[A, Clean]$	$Right$
$[A, Clean]$, $[A, Dirty]$	$Suck$
⋮	⋮

```
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?

Rationality

Fixed performance measure evaluates the environment sequence
- one point per square cleaned up in time T?
- one point per clean square per time step, minus one per move?
- penalize for $> k$ dirty squares?

A rational agent chooses whichever action maximizes the expected value of the performance measure given the percept sequence to date

Rational $\neq$ omniscient
Rational $\neq$ clairvoyant
Rational $\neq$ successful

Rational $\Rightarrow$ exploration, learning, autonomy

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
Observable??				
Deterministic??				
Episodic??				
Static??				
Discrete??				
Single-agent??				

Environment types

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	Yes	Yes	No	No
Deterministic??				
Episodic??				
Static??				
Discrete??				
Single-agent??				

Environment types

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	Yes	Yes	No	No
Deterministic??	Yes	No	Partly	No
Episodic??				
Static??				
Discrete??				
Single-agent??				

Environment types

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	Yes	Yes	No	No
Deterministic??	Yes	No	Partly	No
Episodic??	No	No	No	No
Static??				
Discrete??				
Single-agent??				

Environment types

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	Yes	Yes	No	No
Deterministic??	Yes	No	Partly	No
Episodic??	No	No	No	No
Static??	Yes	Semi	Semi	No
Discrete??				
Single-agent??				

Environment types

	Solitaire	Backgammon	Internet shopping	Taxi
Observable??	Yes	Yes	No	No
Deterministic??	Yes	No	Partly	No
Episodic??	No	No	No	No
Static??	Yes	Semi	Semi	No
Discrete??	Yes	Yes	Yes	No
Single-agent??				

Environment types

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

Grad student

The environment type largely determines the agent design

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

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