ENCS3340 - Artificial Intelligence

Introduction

What is Artificial Artificial Intelligence?

- Let’s define intelligence first.

- Intelligence, based on Wikipedia, is the capacity for:
  - Abstraction
  - Logic
  - Understanding
  - Self-awareness
  - Learning
  - Emotional knowledge
  - Reasoning
  - Planning
  - Creativity
  - Critical thinking
  - Problem-solving.

I.e. the ability to perceive or infer information, and to retain it as knowledge to be applied towards adaptive behaviors within an environment or context.

What is Artificial Artificial Intelligence?

- No standard definition of AI among those working in the field.

- AI has even been defined as:

"... the collection of problems and methodologies studied by artificial intelligence researchers." - Luger and StubbleField

- The textbook discusses four different schools of thought:
  - Machines that **think like humans**
  - Machines that **act like humans**
  - Machines that **think rationally**
  - Machines that **act rationally**

2 Dimensions to Describe AI

- 4 points of views: Systems that can ….

	Cognitive Approach	Engineering / Rational Approach
Behavior	act like humans	act rationally
Reasoning	think like humans	think rationally

- Engineering Approach:
  - Tries to find optimal solution.
  - No matter how (not necessarily what human do)

- Cognitive Approach:
  - Tries to understand the process
  - Tries to reproduce human behavior (even if wrong results)

A Test for Intelligence

The Turing Test

- Proposed by Alan Turing in 1950
- If a human interrogator cannot tell the computer and human then the computer is intelligent
- Measures the intelligence of a computer vs. a human
- Turing predicted that by 2000, a machine might have a 30% chance of fooling a person for 5 minutes

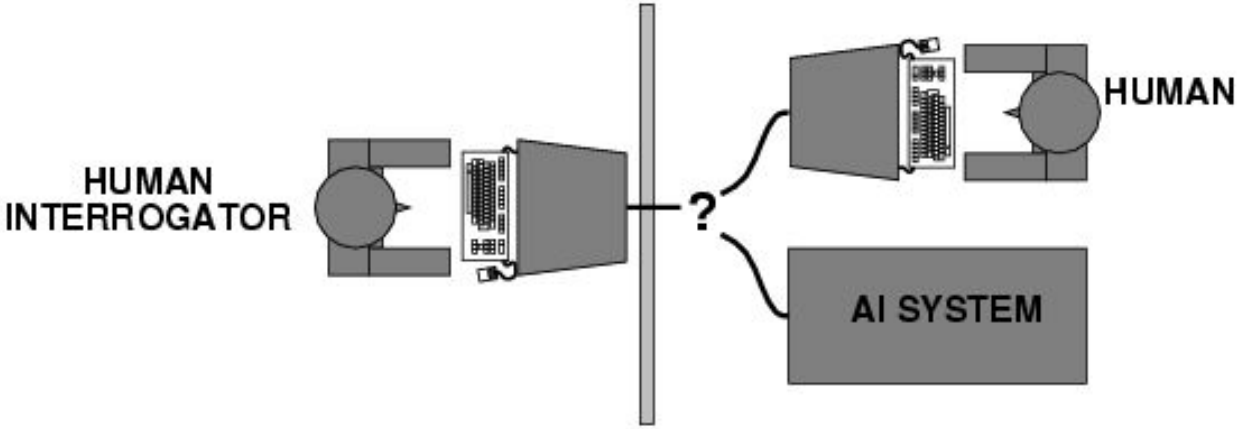

The Turing Test

What capabilities would a computer need to have to pass the Turing Test?

- Natural language processing
- Knowledge representation
- Automated reasoning
- Machine learning
- ….

Turing Test: Criticism

- Not reproducible
- Not constructive
  - Not a good way to solve practical problems
- Not amenable to mathematical analysis
- Machine intelligence designed w.r.t. humans
  - Some human behavior is not intelligent
  - Some intelligent behavior may not be human

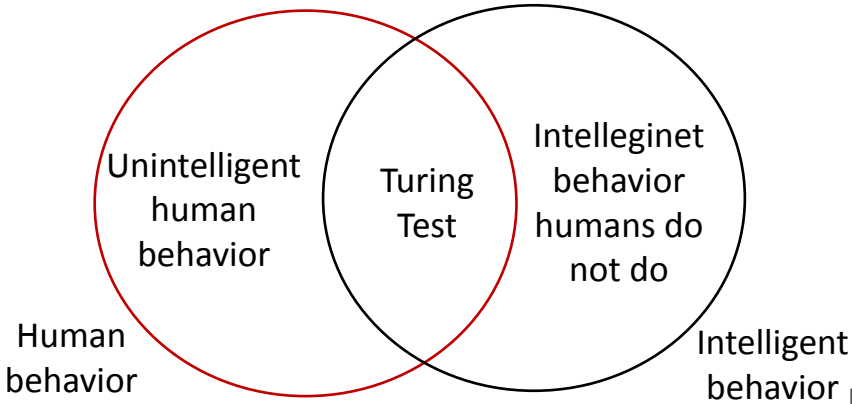

What do we do in AI?

- Search / Heuristic Search
- Automatic Game Playing
- Planning Systems
- Knowledge Representation
- Logic and Inference
- Natural Language Processing
- Dealing with uncertainty – probability and decision theory
- Machine Learning
- Vision & Robotics
- Philosophical Issues
- …

Disciplines involved in AI

- Computer science:
  - Algorithms and data structures to implement theories
- Philosophy:
  - logic, methods of reasoning, foundation of learning, rationality
- Mathematics:
  - formal representation and proof, algorithms, computations, (un)decidability, (in)tractability, probability
- Psychology:
  - Experimental techniques (psycholinguistics, ...)
- Linguistics:
  - Knowledge representation, grammars
- mechanical engineering.
- ….

Intelligent Agents

Recap: What is AI?

- 4 points of views: Systems that can ….

	Cognitive Approach	Engineering / Rational Approach
Behavior	act like humans	act rationally
Reasoning	think like humans	think rationally

- The textbook advocates "acting rationally”

- Acting rationally: Rational Agent
  - Rational behavior: doing the right thing
  - The right thing: that which is expected to maximize goal achievement, given the available information
  - Doesn't necessarily involve thinking – e.g., blinking reflex – but thinking should be in the service of rational action

Agents

- An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators

- Abstractly, an agent is a function from percept histories to actions

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

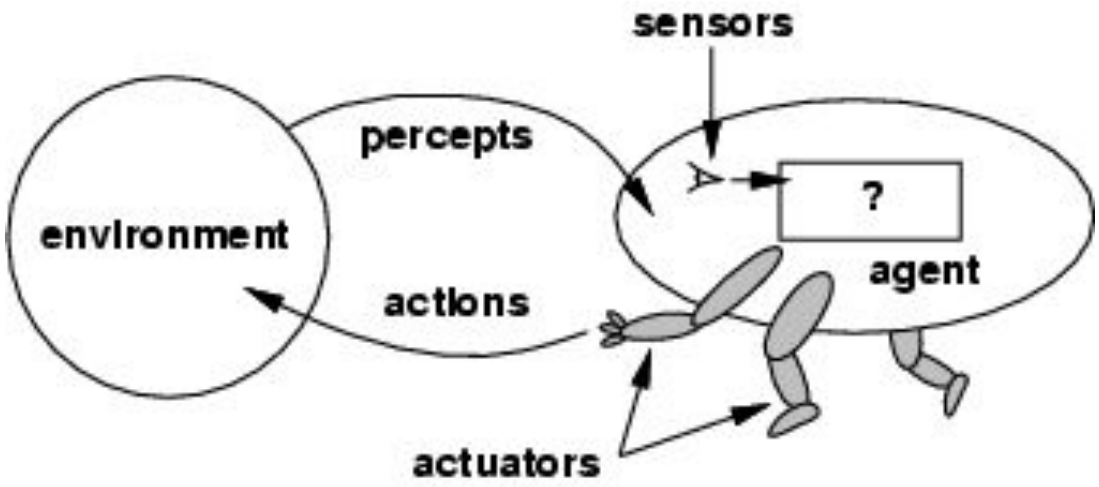

- The agent program runs on the physical architecture to produce f

Example of Agents

- 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

- software agent:
  - functions as sensors: information provided as input to functions in the form of encoded bit strings or symbols
  - functions as actuators: results/outputs of the function

Example: Vacuum-cleaner world

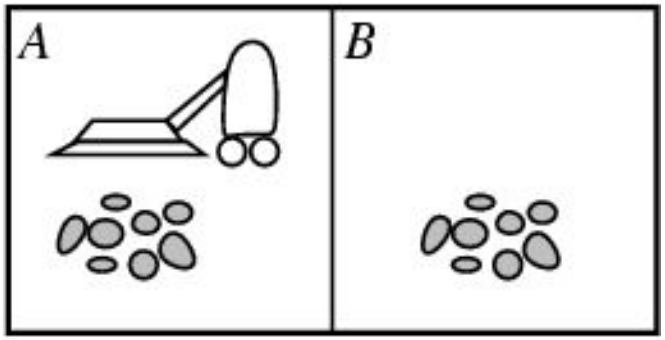

- Percepts: location and contents, e.g.,

  [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
- Possible performance measures of a vacuum-cleaner agent:
  - amount of dirt cleaned up,
  - amount of time taken,
  - amount of electricity consumed,
  - amount of noise generated, etc
- **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.

PEAS

- PEAS: Performance measure, Environment, Actuators, Sensors
- Performance measure: used to evaluate how well an agent solves the task at hand
- Environment: surroundings beyond the control of the agent
- Actuators: determine the actions the agent can perform
- Sensors: provide information about the current state of the environment

PEAS examples

Agent: automated taxi driver

- **Performance measure**: Safe, fast, legal, comfortable trip, maximize profits
- **Environment**: Roads, other traffic, pedestrians, customers
- **Actuators**: Steering wheel, accelerator, brake, signal, horn
- **Sensors**: Cameras, sonar, speedometer, GPS, odometer, engine sensors, keyboard

PEAS examples

Agent: Medical diagnosis system

- **Performance measure**: Healthy patient, minimize costs, lawsuits
- **Environment**: Patient, hospital, staff
- **Actuators**: Screen display (questions, tests, diagnoses, treatments, referrals)
- **Sensors**: Keyboard (entry of symptoms, findings, patient's answers)

PEAS examples

Agent: Part-picking robot in a factory

- **Performance measure**: Percentage of parts in correct bins
- **Environment**: Conveyor belt with parts, bins
- **Actuators**: Jointed arm and hand
- **Sensors**: Camera, joint angle sensors

Environment Types

- **Fully observable (vs. partially observable):**
  - An agent's sensors give it access to the complete state of the environment at each point in time.

- **Deterministic (vs. stochastic):**
  - The next state of the environment is completely determined by the current state and the action executed by the agent.

- **Episodic (vs. sequential):**
  - The agent's experience is divided into atomic "episodes" (each episode consists of the agent perceiving and then performing a single action), and the choice of action in each episode depends only on the episode itself.

Environment Types

- **Static (vs. dynamic):**
  - The environment is unchanged while an agent is deliberating. (The environment is **semidynamic** if the environment itself does not change with the passage of time but the agent's performance score does)

- **Discrete (vs. continuous):**
  - A limited number of distinct, clearly defined percepts and actions.

- **Single agent (vs. multiagent):**
  - An agent operating by itself in an environment.

Environment Types

	Chess with a clock	Chess without a clock	Taxi driving
Fully observable	Yes	Yes	No
Deterministic	Strategic	Strategic	No
Episodic	No	No	No
Static	Semi	Yes	No
Discrete	Yes	Yes	No
Single agent	No	No	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

Five basic types in order of increasing generality:

- Simple reflex agents
  - select actions on the basis of the current percept, ignoring the rest of the percept history.
- Model-based reflex agents
  - maintain some sort of internal state that depends on the percept history and thereby reflects at least some of the unobserved aspects of the current state.
- Goal-based agents
  - needs some sort of goal information that describes situations that are desirable
- Utility-based agents
  - chooses actions to maximize its utility
- Learning agents