AI programs as *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, arms, and other body parts for actuators
- Robotic agent: cameras and infrared range finders for sensors;
- various motors for actuators

Russell & Norvig's book gives a good background on agents
The first 9 slides are based on instructors material at aima.cs.berkeley.edu

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

- Rationality is distinct from omniscience (all-knowing with infinite knowledge)
- 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)

Wumpus World

- Simple "game" by Genesereth, discussed in Russell-Norvig
- In the square containing the wumpus and in the directly (not diagonally) adjacent square the agent will perceive a stench
- In the squares directly adjacent to a pit, the agent will perceive a breeze.
- In the square where the gold is, the agent will perceive a glitter.
- When an agent walks into a wall, it will perceive a bump.
- When the wumpus is killed, it gives out a woeful scream that can be perceived anywhere in the cave.
- The locations of the gold and the wumpus are chosen randomly.
- Each square other than the start can be a pit, with a probability of 0.2

Logic can solve this

		G	
	W		P
P		P	
A			

W = wumpus (eats agents)
P = pit, G = gold
A = agent - go right, left, up, and down, grab (pick up the gold, if it is in the same square), shoot (fire an arrow in a straight line in the direction the agent is facing)

PEAS for Agent Design

- PEAS: Performance measure, Environment, Actuators, Sensors
- Specify the setting for intelligent agent design
- 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
- Agent: Interactive English tutor
  - Performance measure: Maximize student's score on test
  - Environment: Set of students
  - Actuators: Screen display (exercises, suggestions, corrections)
  - Sensors: Keyboard

Aspects of the environment

- 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.
- 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. (If the environment is deterministic except for the actions of other agents, then the environment is strategic)

Agent types

- Four basic types in order of increasing generality:
  - Simple reflex agents
  - Model-based reflex agents
  - Goal-based agents
  - Utility-based agents

Simple reflex agents

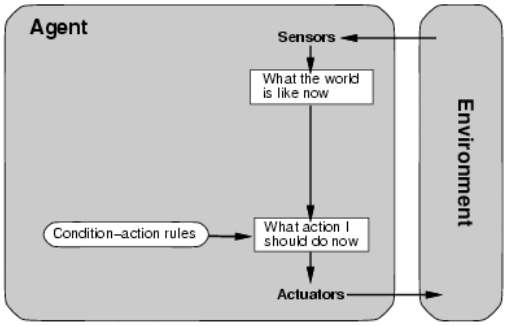

Learning agents

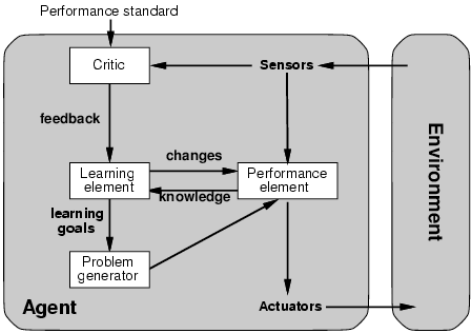

n-player iterated prisoner's dilemma

- A simple testbed for agent systems
- Game-theoretic model of co-operation in a world of egoists without central authority
  - applications - commerce, nation-states, trench warfare, …

Prisoner's dilemma : two players

each can either co-operate or defect

choices are made simultaneously

rewards according to the table (right)

non-zero sum game - win-win situation

	co-operate	defect
co-operate	3 / 3	0 / 5
defect	5 / 0	1 / 1

Iterated version - repeat *m* times, accumulate reward
n-player version - interact with *n* players

n-IPD tournaments

- include some simple strategies
  - always co-operate
  - always defect
  - tit-for-tat - echo whatever your opponent did last time
  - grim - co-operate until opponent defects, then always defect
  - tit-for-two-tats - co-operate unless opponents last two moves = defect
  - random (with different probabilities of co-operate/defect)
- and some complex strategies
  - try to find exploitable opponents
- There is no overall best strategy
  - tournaments show high scoring strategies are
    - responsive i.e. they punish defection
    - generally nice i.e. don't defect without provocation
    - forgiving i.e. don't over-punish defection

Example : robocup soccer agents

www.robocup.org

Objective : "By the year 2050, develop a team of fully autonomous humanoid robots that can win against the human world soccer champion team".

Intermediate level

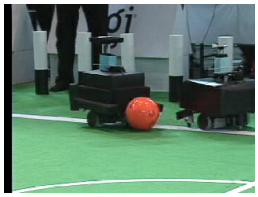

Reduce hardware problems by using wheeled robots, also 4-legged.
Main problems - how to program individual behaviour, team tactics, learning, reaction to opponents, …

Simulations (RSS)

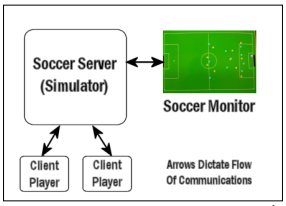

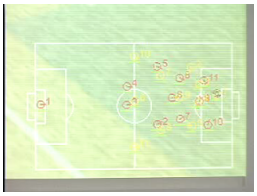

- artificial players
  - can move as freely as humans
  - sense their surroundings as well as a human
  - receive regular information from server
  - reply with their next action
- individual agents
  - limited information
  - limited time
  - learning vs pre-programming
- work towards emergent behaviour
  individual actions → team performance
- easy benchmarking

Projects

- n-player iterated prisoners dilemma
  - investigate multi-agent interactions
- eliza
  - e.g. elicit information for customer service
  - grammar or keyword approach
- wumpus world
  - representation, search, solution
  - possibly also a "world generator"
- sudoku solver
  - representation, search, difficulty level
  - need to be familiar with solving the puzzles
- machine learning (from data)
  - (fuzzy) rules or genetic programming
  - define your own (simple) problems or use examples given