Introduction to Artificial Intelligence
COMP 3501 / COMP 4704-4
Lecture 1: Course Intro

Prof. Nathan Sturtevant

Lecture Overview

- What is AI?
- AI History
- Views/goals of AI
- Course Overview

Artificial Intelligence

- As humans we have intelligence
  - But what is intelligence?
  - What does it mean to build artificial intelligence?

A history of AI

- Darthmouth conference (1956)
  - 10 attendees spend two months discussing AI
  - John McCarthy, Marvin Minsky, Claude Shannon, Arthur Samuel, Allen Newell, Herbert Simon
  - From MIT, CMU, Stanford and IBM
- Newell and Simon had already developed a logical reasoning program

A history of AI (1952-1969)

- It seems like AI can do anything:
  - General Problem Solver; imitates human thinking
    - Newell and Simon
  - Checkers program by Samuel learns to play
  - LISP invented by John McCarthy
  - Minsky & students work on small problems requiring intelligence
  - Early work on learning and perceptrons

Reality (1966-1973)

- Program previously only run on very small problems
- Complexity theory develops proving “hard” problems
- Perceptrons shown to have learning limitations
- AI research nearly killed in the UK

Knowledge-based systems (1969-1979)

- Knowledge from experts distilled into rules
  - First systems enhanced by human expertise
  - MYCIN system diagnosed blood infections at the level of experts (450 rules)
- Many specialized representations and reasoning languages

Commercial Success (1980 - present)

- Many commercial companies started using AI techniques internally
  - A DEC program helped configure new orders
  - Saved ~$40 million a year

The scientific method (1986-present)

- Neural networks come back into favor
- Add-hoc methods start to drop away
- New work borrows ideas from mathematics and statistics providing a stronger foundation

Intelligent agents & large data sets

- Architectures such as SOAR use many agents for simulating behavior
- The internet provided a rich application domain
- Internet also provides very large data sets
  - Plurality of examples allows simple learning algorithms

State of Art

- Robotic cars: DARPA grand challenge
- Speech recognition: commonly used for phone systems
- Autonomous planning: performed on spacecraft
- Game playing: Deep Blue & Watson
- Spam Fighting: Learning algorithms classify spam
- Logistics: 1991 Persian Gulf planned automatically
  - DARPA states this paid off all investments in AI
- Machine Translation: Reasonable automatic translation

"Your automated surgery should go just fine, but if you feel them taking out anything but your gallbladder, pull that cord immediately."

Homework:
Russell and Norvig
Chapter 1, questions 1.11; 1.12; 1.13
Due before lecture Wednesday

Intelligent Agents

- Definitions:
  - *Environment*
  - *Sensors*
  - *Actuators*
- An *agent* perceives the environment via sensors and acts on environment through actuators
  - A *percept* describes an agents inputs

Agents

- Agent → f(*percepts*)
  - If percepts are finite, we can measure the agent behavior exactly and write it into a table
  - Internally an agent will have some program to implement its own behavior
    - This could just be a table
    - But, could be something more complex

Example: Vacuum World

- Percepts:
- Actions:

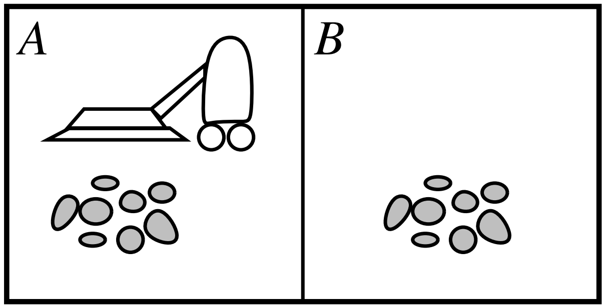

Example: Vacuum World

- Percepts: Location, contents
- Actions: Left, right, suck, no-op

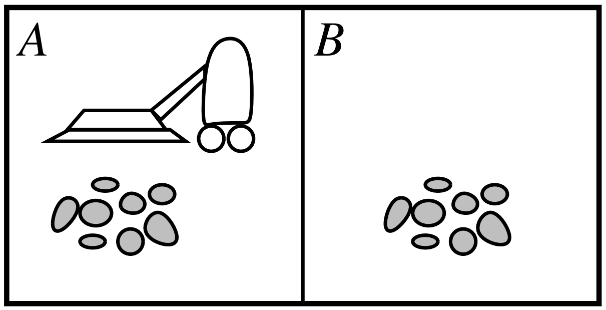

Rational Agents

- How can you measure the rationality of an agent?
  - What are the consequences of behavior?
  - Evaluate state of the *environment*
- Different measures result in different performance
  - Maximize dirt cleaned up
  - Maximize average cleanliness

Rationality

- Depends on:
  - The performance measure
  - The agent's prior knowledge of the environment
  - The actions that the agent can perform
  - The agent's percept/sensor sequence to date
- *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.*

Example: Vacuum world

- Performance:
  - 1 point per each clean square per time step
- Environment:
  - Environment known; dirt not
- Actions:
  - Left, right, suck
- Sensors:
  - 100% Reliable

Agent Type	Performance Measure	Environment	Actuators	Sensors
Robot soccer player	Winning game, goals for/against	Field, ball, own team, other team, own body	Devices (e.g., legs) for locomotion and kicking	Camera, touch sensors, accelerometers, orientation sensors, wheel/joint encoders
Internet book-shopping agent	Obtain re-quested/interesting books, minimize expenditure	Internet	Follow link, enter/submit data in fields, display to user	Web pages, user requests
Autonomous Mars rover	Terrain explored and reported, samples gathered and analyzed	Launch vehicle, lander, Mars	Wheels/legs, sample collection device, analysis devices, radio transmitter	Camera, touch sensors, accelerometers, orientation sensors, , wheel/joint encoders, radio receiver
Mathematician's theorem-proving assistant				

Environment types

- Fully observable vs. partially observable
- Single agent vs. multiagent
- Deterministic vs. stochastic
- Episodic vs. sequential
- Static vs. dynamic
- Discrete vs. continuous
- Known vs. unknown

Task Environment	Observable	Deterministic	Episodic	Static	Discrete	Agents
Robot soccer	Partially	Stochastic	Sequential	Dynamic	Continuous	Multi
Internet book-shopping	Partially	Deterministic*	Sequential	Static*	Discrete	Single
Autonomous Mars rover	Partially	Stochastic	Sequential	Dynamic	Continuous	Single
Mathematician's assistant	Fully	Deterministic	Sequential	Semi	Discrete	Multi

Agents as programs

- Agents take current percepts as input
  - Return an action to perform
  - Agent must remember full sequence of actions if necessary (Markov)
- Building full table of actions is not practical
  - Four basic agent types

Simple reflex agents

- Reflex agent only uses current percept
  - 4 possibilities vs 4^T including history
- Behavior is composed of if-then rules
  - if [status == dirty] then return suck
- Only works if environment is fully observable
  - Not if we need to correlate two percepts in time to know something

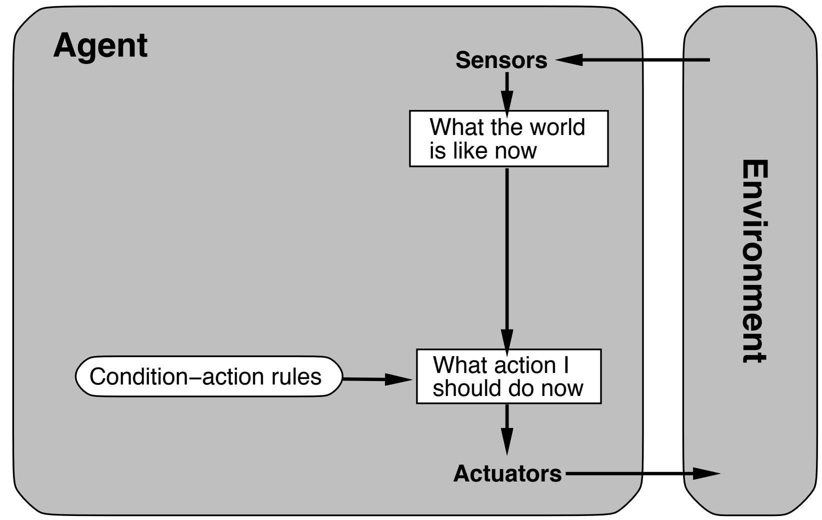

Model-based reflex agents

- Agent maintains internal state which reflects beliefs about the world
  - Requires model of environment & actions
- Upon receiving percept, agent updates state model
- Acts reflexively according to state model

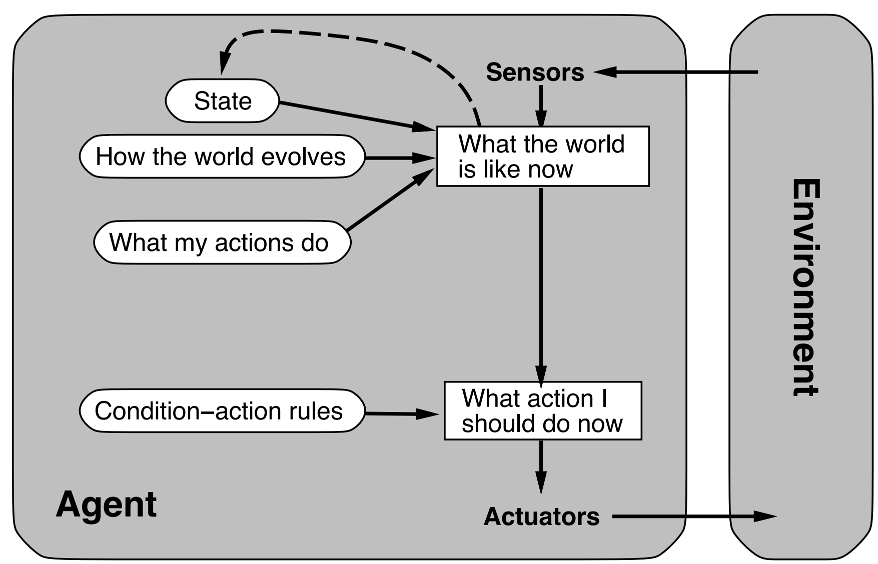

Goal-based agents

- Reflexive agents have rules for a single task
  - What if the task changes?
- Use model of the environment to predict the future
  - Find action sequence which converts the current state to the goal state
  - Partially encompasses *search* and *planning*

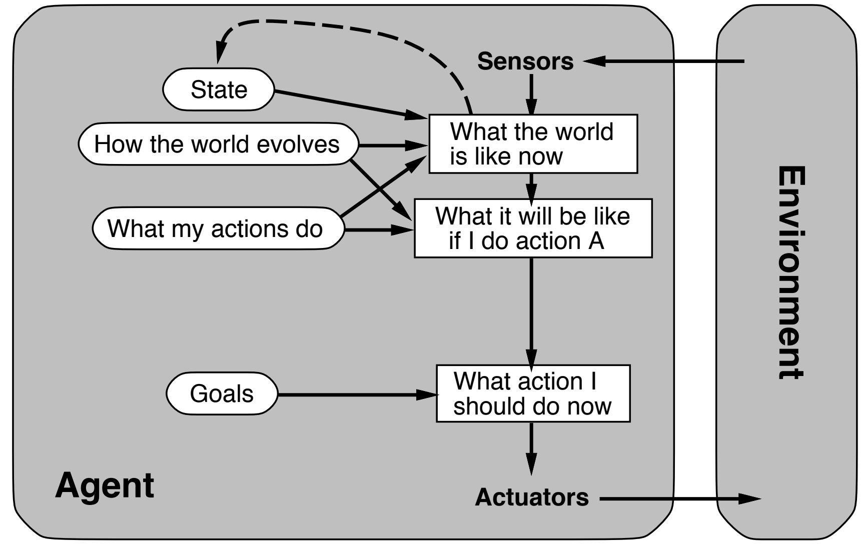

Utility-based agent

- Goals alone produce solutions but don’t measure solution quality
- Utility is a generic measure of quality
- Not required for rational behavior
  - But rational behavior can be described with utilities
- Most search/planning also uses utilities

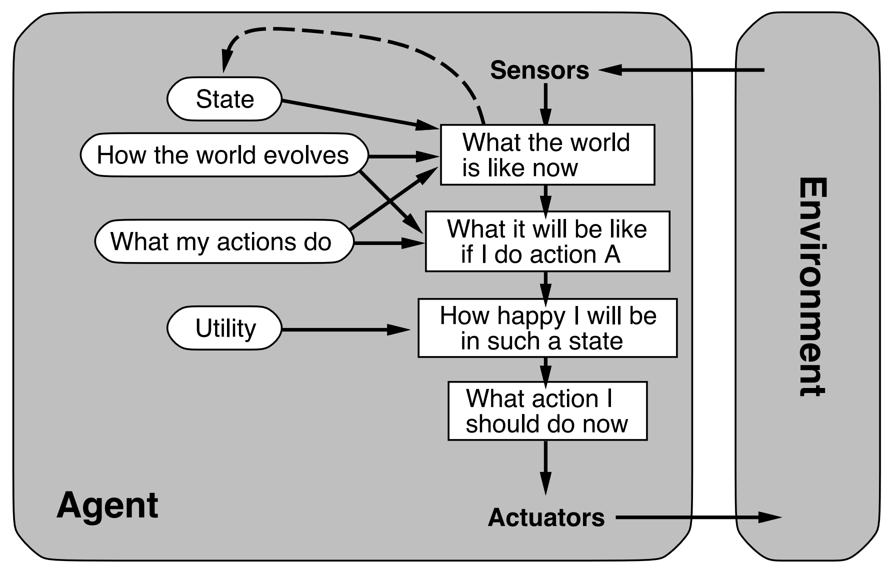

Learning

- Each approach can be enhanced with learning
  - "Critic" can provide feedback
  - Utilities can be the basis of rewards used for feedback
- Learning may change rules or their expected utilities

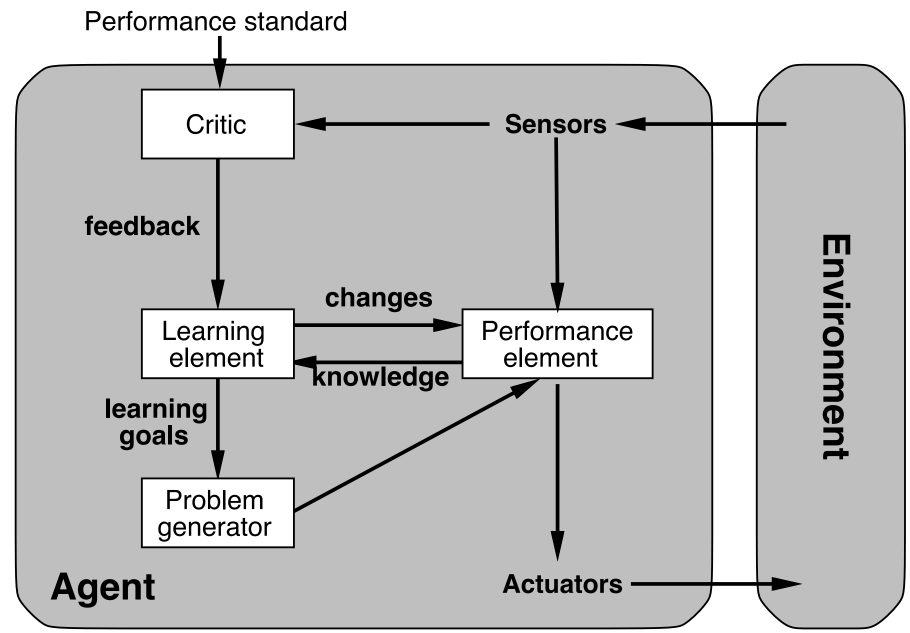

Environment/Agent Representation

- World can be *atomic* (opaque)
  - Black box operated on by actions
- World can be *factored*
  - Represented by variables and values
- World can be *structured*
  - Represented by the ideas of objects and their relationships

Summary

- AI is a field which is interested in *rational agents*
- Rational agents attempt to maximize their payoff
- Agents act in external environments
- Different agent architectures
- Different environment representations