

CSI535 Introduction to A.I.

Overview of Lecture

What is A.I.?

The Agent Approach to A.I.

Course Overview

Course Logistics

What is Artificial Intelligence-1 ?

Giving intelligence to machines, but how?

“[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning . . .” (Bellman, 1978)	“The study of mental faculties through the use of computational models” (Charniak+McDermott, 1985)
“The study of how to make computers do things at which, at the moment, people are better” (Rich+Knight, 1991)	“The branch of computer science that is concerned with the automation of intelligent behavior” (Luger+Stubblefield, 1993)

Views of AI fall into four categories:


Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

What are the problems with each? Which is the best, Consider how you answer this question

Acting Like a Human

Turing (1950) “Computing machinery and intelligence”:

- ◇ “Can machines think?” → “Can machines behave intelligently?”
- ◇ Operational test for intelligent behavior: the Imitation Game



- ◇ Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- ◇ Anticipated all major arguments against AI in following 50 years
- ◇ Suggested major components of AI: knowledge, reasoning, language understanding, learning

Many prizes (\$\$\$) for passing the Turing test
... but no one even tries

Thinking Like a Human

- ◻ 1960s “cognitive revolution”: information-processing psychology replaced prevailing orthodoxy of behaviorism

Requires scientific theories of internal activities of the brain

- What level of abstraction? “Knowledge” or “circuits”?
- How to validate? Requires

- 1) Predicting and testing behavior of human subjects (top-down)
- or 2) Direct identification from neurological data (bottom-up)

Both approaches (roughly, Cognitive Science and Cognitive Neuroscience) are now distinct from AI

Laws of Thought: Thinking Rationally

Normative (or prescriptive) rather than descriptive

Aristotle: what are correct arguments/thought processes?

Several Greek schools developed various forms of logic:

notation and rules of derivation for thoughts;

may or may not have proceeded to the idea of mechanization

Direct line through mathematics and philosophy to modern AI

Problems: There are two ... Assumptions ... Computational limitations

Acting Rationally

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

Aristotle (Nicomachean Ethics):

Every art and every inquiry, and similarly every action and pursuit, is thought to aim at some good

Notion of limited rationality

Rational Agents

An agent is an entity that perceives and acts

This course is about designing rational agents

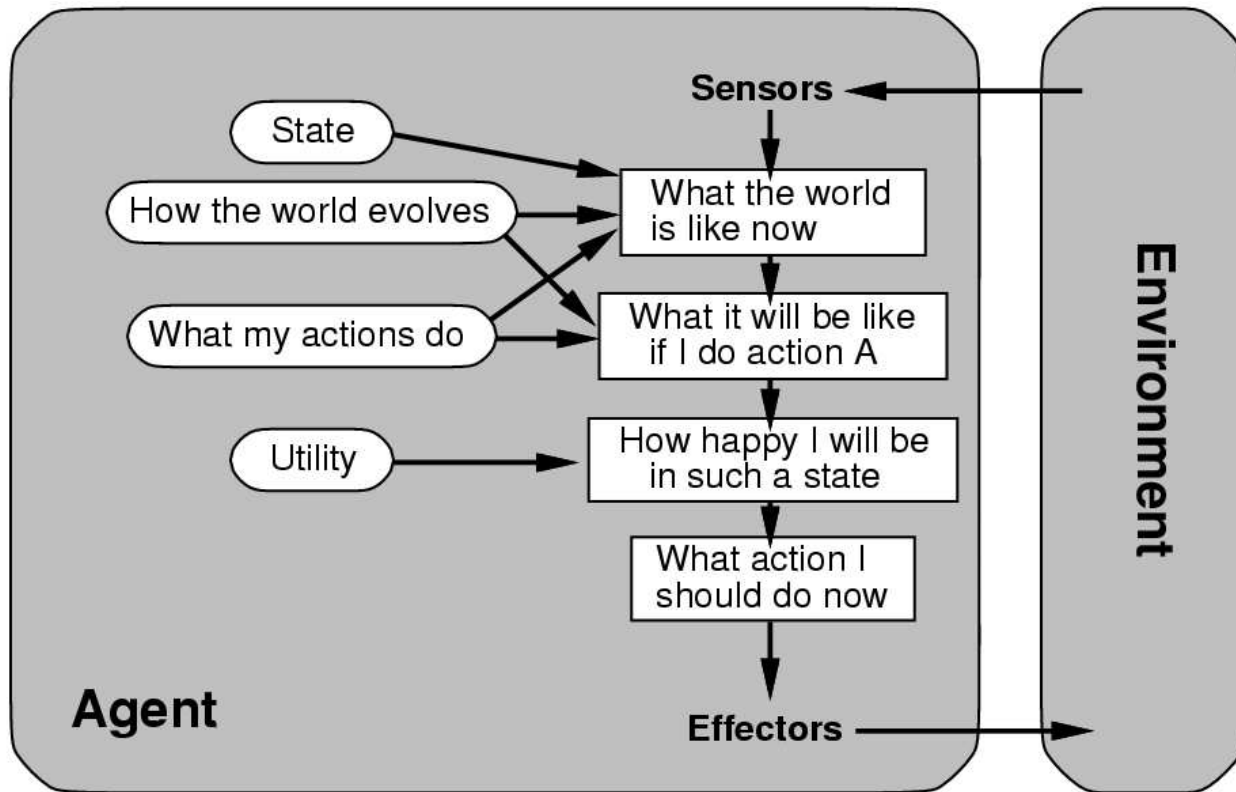
Abstractly, an agent is a function from percept histories to actions:

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

For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance

Caveat: computational limitations make perfect rationality unachievable
→ design best program for given machine resources

Agent Overview



Percepts??

Actions??

Goals??

Environment??

PAGE

Must first specify the setting for intelligent agent design

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

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

Actions?? steer, accelerate, brake, horn, speak/display, ...

Goals?? safety, reach destination, maximize profits, obey laws, passenger comfort, ...

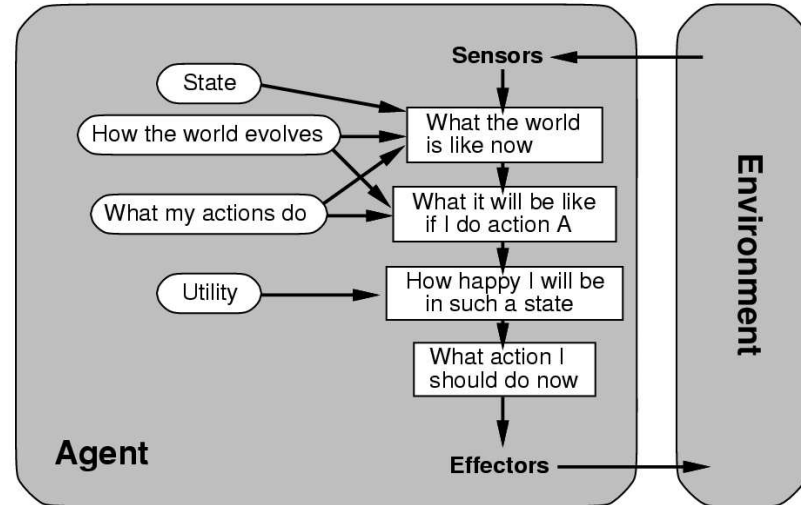
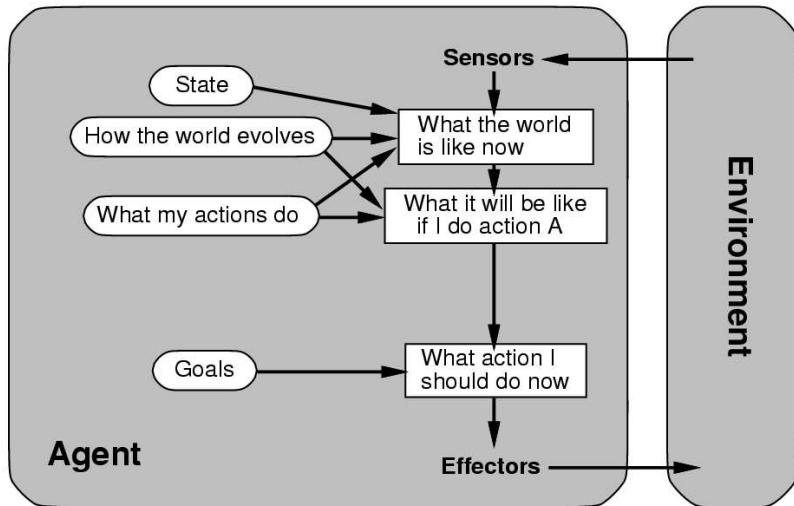
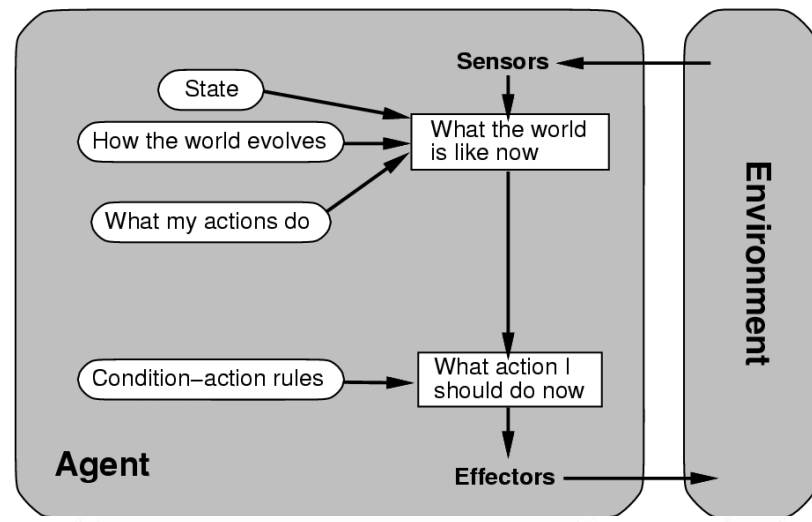
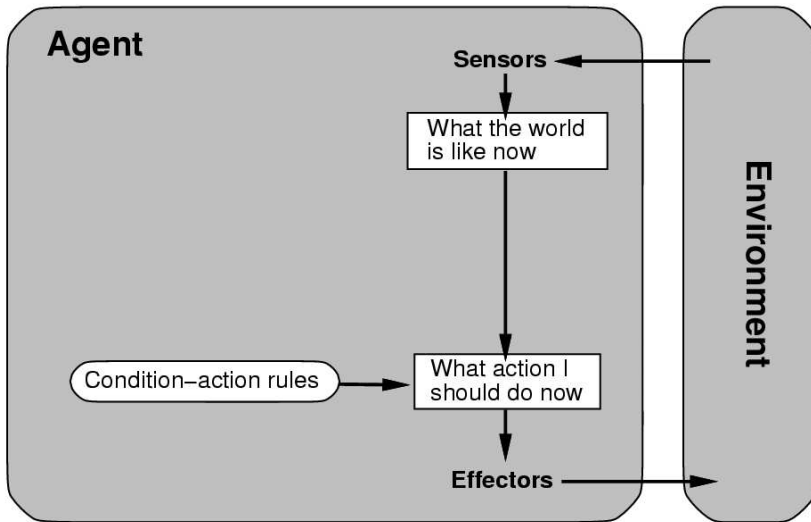
Environment?? US urban streets, freeways, traffic, pedestrians, weather, customers, ...

Environment Types

	Solitaire	Backgammon	Internet shopping	Taxi
<u>Accessible??</u>				
<u>Deterministic??</u>				
<u>Episodic??</u>				
<u>Static??</u>				
<u>Discrete??</u>				

And single agent versus multiagent

Agent Structures



Why so many A.I. Techniques

- Due to variation in PAGE. Consider:

Percepts Breeze, Glitter, Smell

Actions Left turn, Right turn,
Forward, Grab, Release, Shoot

Goals Get gold back to start
without entering pit or wumpus square

Environment

Squares adjacent to wumpus are smelly

Squares adjacent to pit are breezy

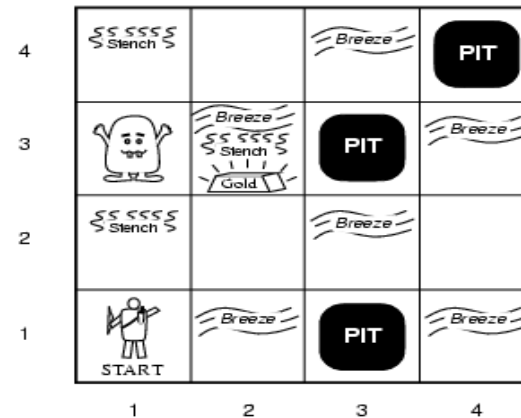
Glitter if and only if gold is in the same square

Shooting kills the wumpus if you are facing it

Shooting uses up the only arrow

Grabbing picks up the gold if in the same square

Releasing drops the gold in the same square



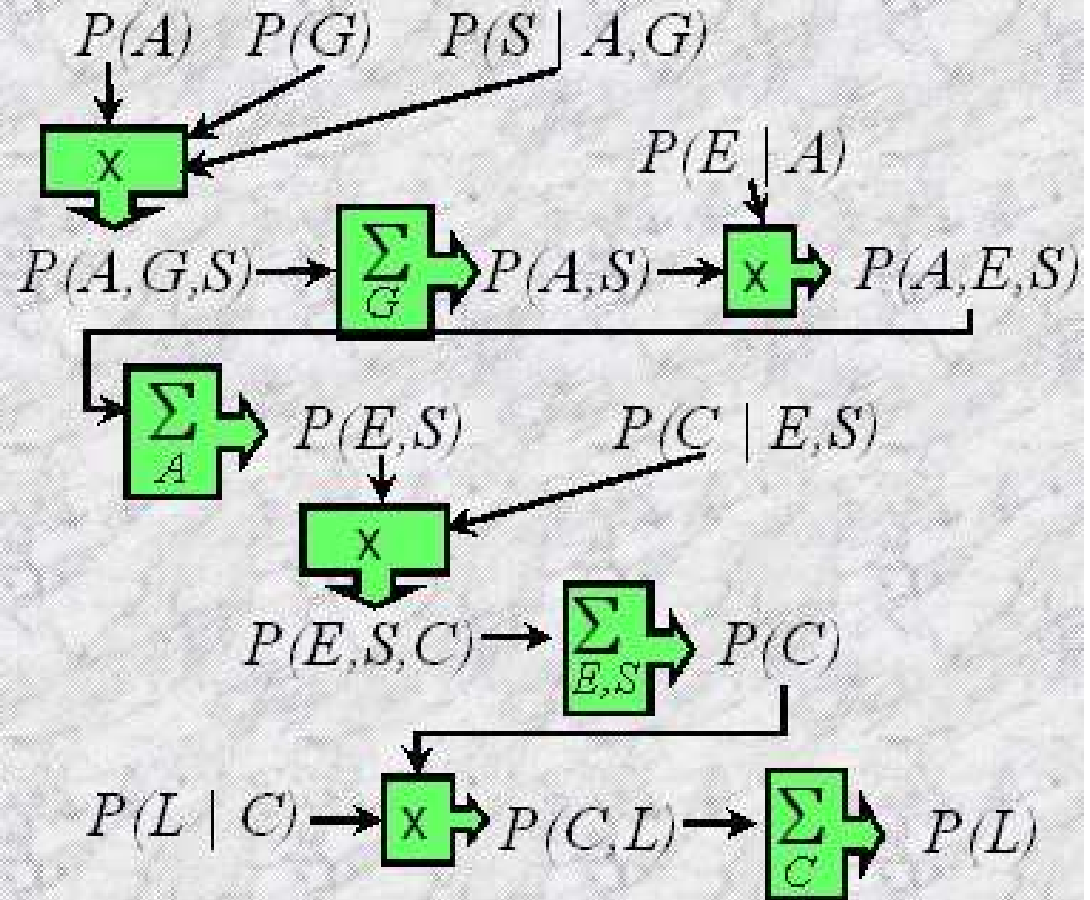

Requires Various Computational Techniques - 1

- Techniques for different agent components
- To be covered in the course:
 - Reasoning
 - Propositional and first order logic
 - Belief networks
 - Problem solving
 - Uninformed search, informed search
 - Constrained search (CSP)
 - Adversarial search (environment competes with agent)

Bayesian Belief Networks

- Combination of probabilistic modeling and DAGs
- Nodes on graph are propositional variables.
- Links represent apriori known causal dependencies.
- Reasoning by merging semantic models and evidence.
- Efficient representation of joint distribution

Variable Elimination with loops

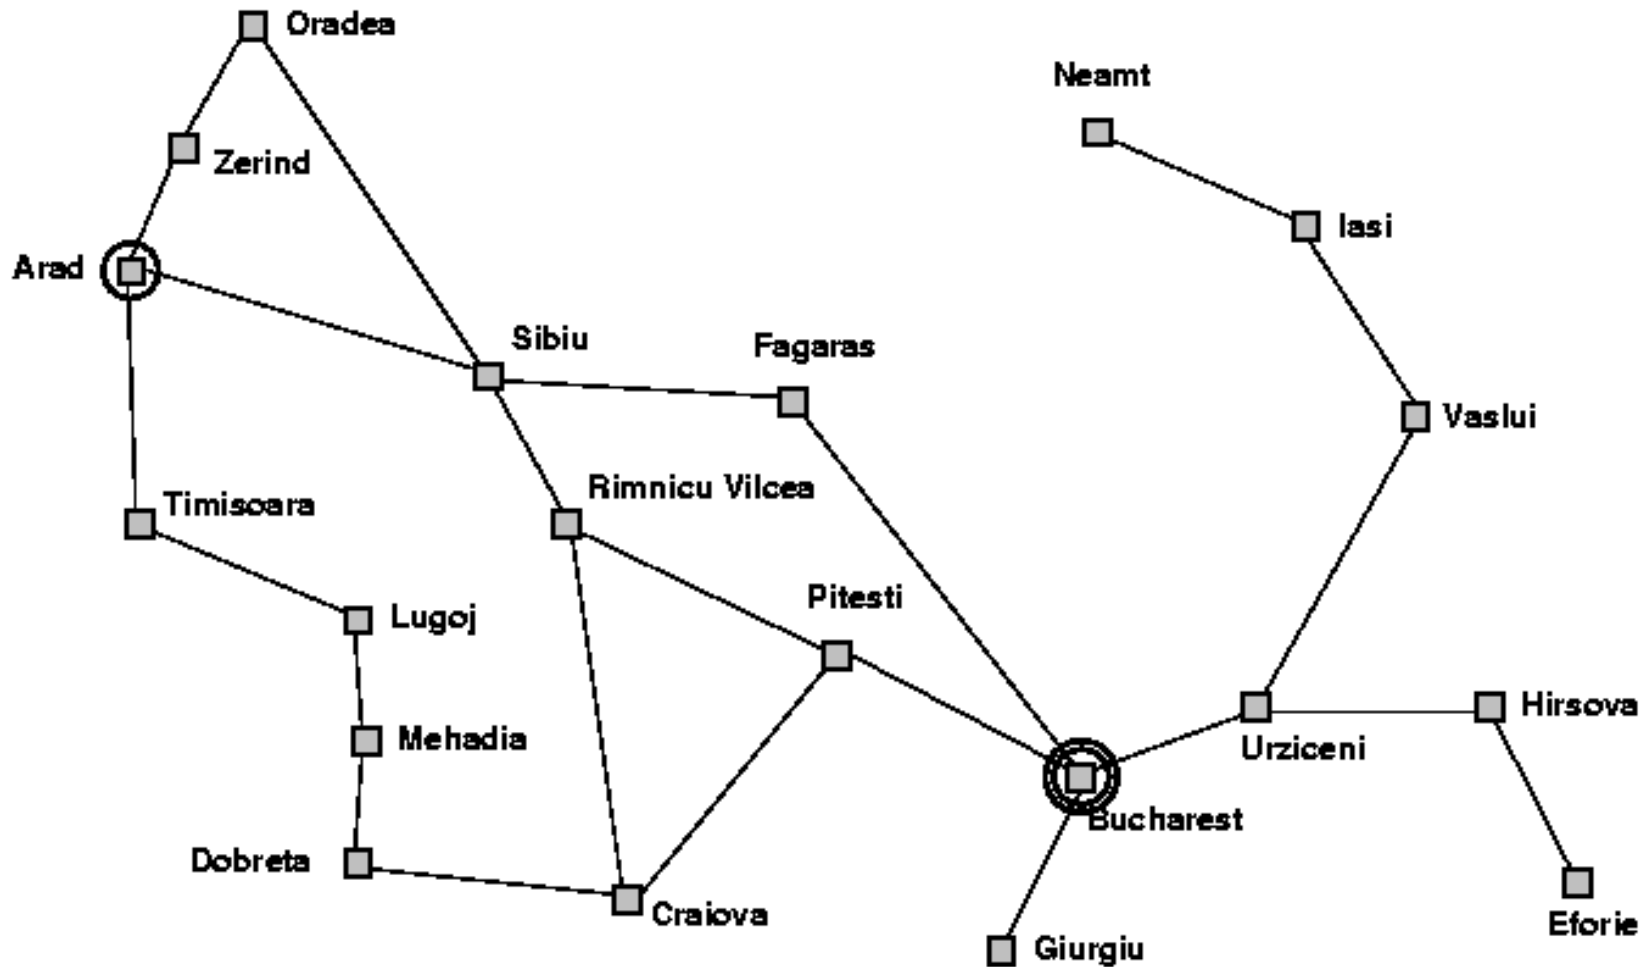


Complexity is exponential in the size of the factors

Direct World Representations

- Can compute any subset of propositions given another subset.
- Perform different types of reasoning
 - Prediction
 - Abduction
 - Explaining away
- Global semantics
- Local semantics exploit conditional independence

Map of Area



Eight State Puzzle

5	4	
6	1	8
7	3	2

Start State

1	2	3
8		4
7	6	5

Goal State

states??

operators??

goal test??

path cost??


What is a State and a Node?

A *state* is a (representation of) a physical configuration

A *node* is a data structure constituting part of a search tree

includes *parent*, *children*, *depth*, *path cost* $g(x)$

States do not have parents, children, depth, or path cost!



The **EXPAND** function creates new nodes, filling in the various fields and using the **OPERATORS** (or **SUCCESSORFN**) of the problem to create the corresponding states.

Requires Various Computational Techniques - 2

- To be covered in the course:
 - Planning and Natural Language Processing
 - Learning
 - Computational Learning Theory
 - Neural networks
 - Decision trees
 - Reinforcement learning

“A.I.” Faculty

- Murray – Automatic deduction and theorem proving. CSI 664
- Haas – Natural language processing, Prolog. CSI 636
- Berg – Machine learning, Neural networks, Bioinformatics.
- Davidson – Machine learning, probabilistic reasoning/learning, learning by compact encoding

Logistics and Next Lecture

- Text: Russell and Norvig – AI, 2nd Edition
- Mailing List, TA: Ke@cs.albany.edu
- Meeting times: MW 2:30-3:50 @ HU110
- Office hours: MW: 4-5pm
- Contact: davidson@cs.albany.edu, 518 442 5173
- 2 Assignments (25% each), Homeworks (10%), Final Exam (30%), Class participation (10%)
- Class web site : cs.albany.edu/~davidson, lectures
- Review of this material – Chapter 1 – (RN)
- Next classes: Agent Architectures – Chapter 2 – (RN)