


Introduction to Artificial Intelligence

CS 171, Fall 2017

Introduction to Artificial Intelligence

Prof. Richard Lathrop

Introduction



People

- Instructor: Richard Lathrop, <rickl@uci.edu>
 - Office hours Wednesday 2-3pm, DBH 4224
- TAs: (Office hours TBD)
 - Jia Chen, <jiac5@uci.edu>
 - Sridevi Maharaj, <sridevi.m@uci.edu>
- Readers: (Office hours TBD)
 - Kyoungwon Kim <wonkim@uci.edu>
 - Zephyr Yao <zhihaoy1@uci.edu>
 - Dongxu Zhao <dongxuz1@uci.edu>
- Coding Project: Abdullah Younis <younisa@uci.edu>


Course outline

- <https://canvas.eee.uci.edu/courses/6753>
- Framed around three pillars of AI
 - Search
 - Logic
 - Learning (see also CS178)
- Project:
 - “Wumpus World” agent & tournament
 - Tournament Director: Abdullah Younis
 - Weekly Q&A sessions, Wed 8pm (planned)

Course outline

- Collaboration OK
- Grading
 - Optional Homeworks (5), not graded
 - Discussion participation (10%; 7 of 10)
 - Quizzes (20%)
 - Four in-class quizzes: 10/12, 10/26, 11/21, 12/5
 - Project (20%)
 - Wumpus World
 - Teams of 1 or 2
 - Several milestones through quarter; teams due 10/8
 - Midterm (25%)
 - In class, 11/7
 - Final (25%)
 - Cumulative
 - 12/15, 10:30-12:30 (note different time!)

What is AI?





?
=



?
=



What is AI?




What is AI?

- Competing axes of definitions:
 - Think v. Act
 - Human-like v. Rational
 - Often not the same thing
 - Cognitive science, economics, ...
- How to simulate human intellect & behavior by machine
 - Mathematical problems (puzzles, games, theorems)
 - Common-sense reasoning
 - Expert knowledge (law, medicine)
 - Social behavior
 - Web & online intelligence
 - Planning, e.g. operations research

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>

The Turing test

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

- Test requires computer to “pass itself off” as human

- Necessary?
- Sufficient?

- Requires:

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





Figure 1.1 The Turing test.

Act/Think Humanly/Rationally


- Act Humanly
 - Turing test
- Think Humanly
 - Introspection; Cognitive science
- Think rationally
 - Logic; representing & reasoning over problems
- Acting rationally
 - Agents; sensing & acting; feedback systems

Agents

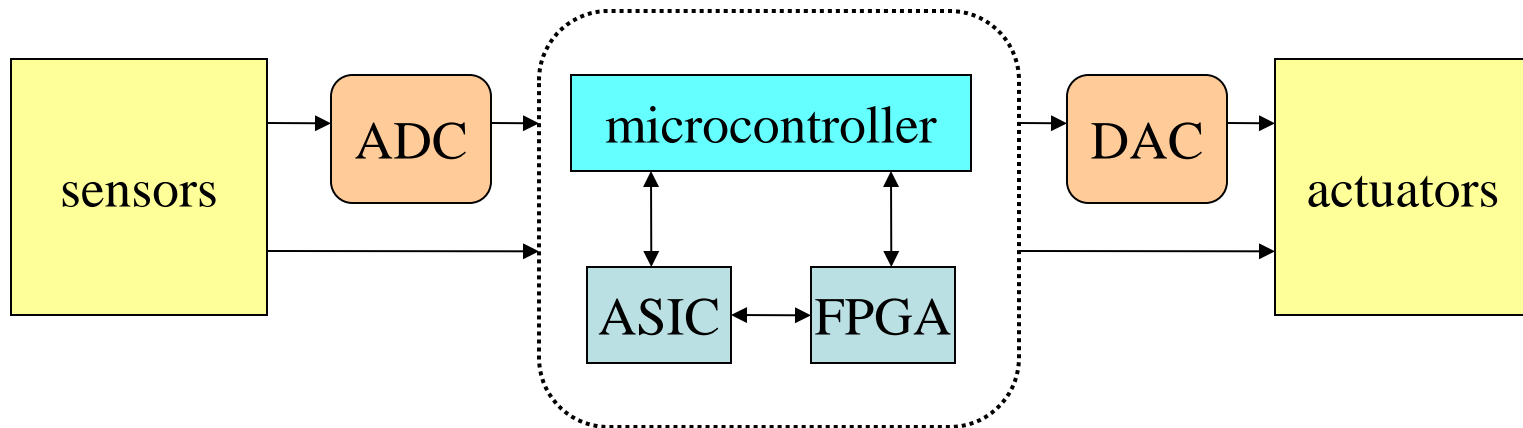
- An **agent** is anything that can be viewed as **perceiving** its **environment** through **sensors** and **acting** upon that environment through **actuators**
- Human agent:
 - Sensors: eyes, ears, ...
 - Actuators: hands, legs, mouth...
- Robotic agent
 - Sensors: cameras, range finders, ...
 - Actuators: motors




Agents and environments



Compare: Standard Embedded System Structure



Agents and environments




- The **agent function** maps from percept histories to actions:

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

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

Vacuum World



- **Percepts:** location, contents
 - e.g., [A, dirty]
- **Actions:** {left, right, vacuum,...}

Rational agents

- **Rational Agent:** For each possible percept sequence, a rational agent should select an action that is *expected* to maximize its *performance measure*, based on the evidence provided by the percept sequence and whatever built-in knowledge the agent has.
- **Performance measure:** An objective criterion for success of an agent's behavior (“cost”, “reward”, “utility”)
- *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 percepts & experience (with ability to **learn and adapt**) without depending solely on build-in knowledge

Task environment

- To design a rational agent, must specify task env.
- Example: automated taxi system

“PEAS”

- Performance measure

- Safety, destination, profits, legality, comfort, ...

- Environment

- City streets, freeways; traffic, pedestrians, weather, ...

- Actuators

- Steering, brakes, accelerator, horn, ...

- Sensors

- Video, sonar, radar, GPS / navigation, keyboard, ...

PEAS

- Example: 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

- Example: Agent = Part-picking robot
- **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. (If the environment is deterministic except for the actions of other agents, then the environment is **strategic**)
- **Episodic** (vs. **sequential**): An agent's action is divided into atomic episodes. Decisions do not depend on previous decisions/actions.

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.

How do we **represent** or **abstract** or **model** the world?

- **Single agent** (vs. **multi-agent**): An agent operating by itself in an environment. Does the other agent interfere with my performance measure?

task environm.	observable	determ./ stochastic	episodic/ sequential	static/ dynamic	discrete/ continuous	agents
crossword puzzle	fully	determ.	sequential	static	discrete	single
chess with clock	fully	strategic	sequential	semi	discrete	multi
poker						
back gammon						
taxi driving	partial	stochastic	sequential	dynamic	continuous	multi
medical diagnosis	partial	stochastic	sequential	dynamic	continuous	single
image analysis	fully	determ.	episodic	semi	continuous	single
partpicking robot	partial	stochastic	episodic	dynamic	continuous	single
refinery controller	partial	stochastic	sequential	dynamic	continuous	single
interact. Eng. tutor	partial	stochastic	sequential	dynamic	discrete	multi

task environm.	observable	determ./ stochastic	episodic/ sequential	static/ dynamic	discrete/ continuous	agents
crossword puzzle	fully	determ.	sequential	static	discrete	single
chess with clock	fully	strategic	sequential	semi	discrete	multi
poker	partial	stochastic	sequential	static	discrete	multi
back gammon						
taxi driving	partial	stochastic	sequential	dynamic	continuous	multi
medical diagnosis	partial	stochastic	sequential	dynamic	continuous	single
image analysis	fully	determ.	episodic	semi	continuous	single
partpicking robot	partial	stochastic	episodic	dynamic	continuous	single
refinery controller	partial	stochastic	sequential	dynamic	continuous	single
interact. Eng. tutor	partial	stochastic	sequential	dynamic	discrete	multi

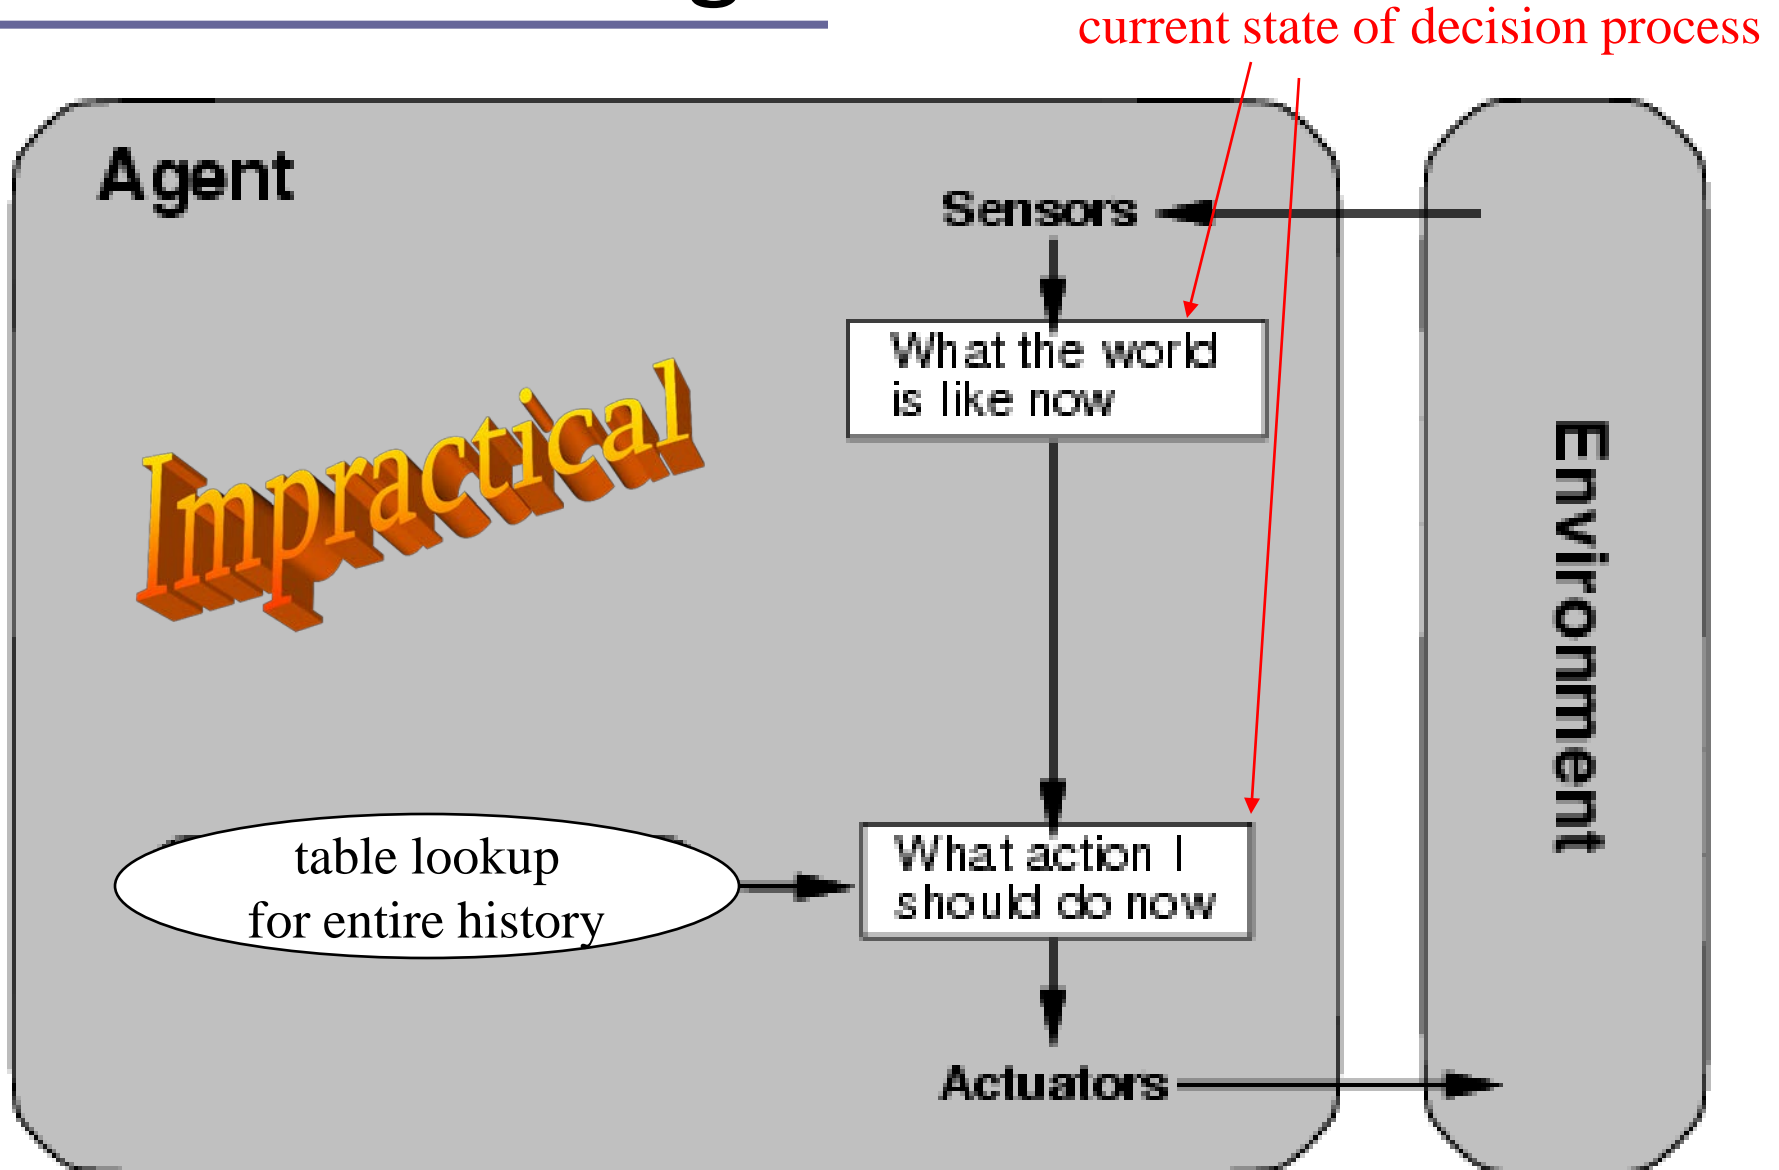
task environm.	observable	determ./ stochastic	episodic/ sequential	static/ dynamic	discrete/ continuous	agents
crossword puzzle	fully	determ.	sequential	static	discrete	single
chess with clock	fully	strategic	sequential	semi	discrete	multi
poker	partial	stochastic	sequential	static	discrete	multi
back gammon	fully	stochastic	sequential	static	discrete	multi
taxi driving	partial	stochastic	sequential	dynamic	continuous	multi
medical diagnosis	partial	stochastic	sequential	dynamic	continuous	single
image analysis	fully	determ.	episodic	semi	continuous	single
partpicking robot	partial	stochastic	episodic	dynamic	continuous	single
refinery controller	partial	stochastic	sequential	dynamic	continuous	single
interact. Eng. tutor	partial	stochastic	sequential	dynamic	discrete	multi

Agent types

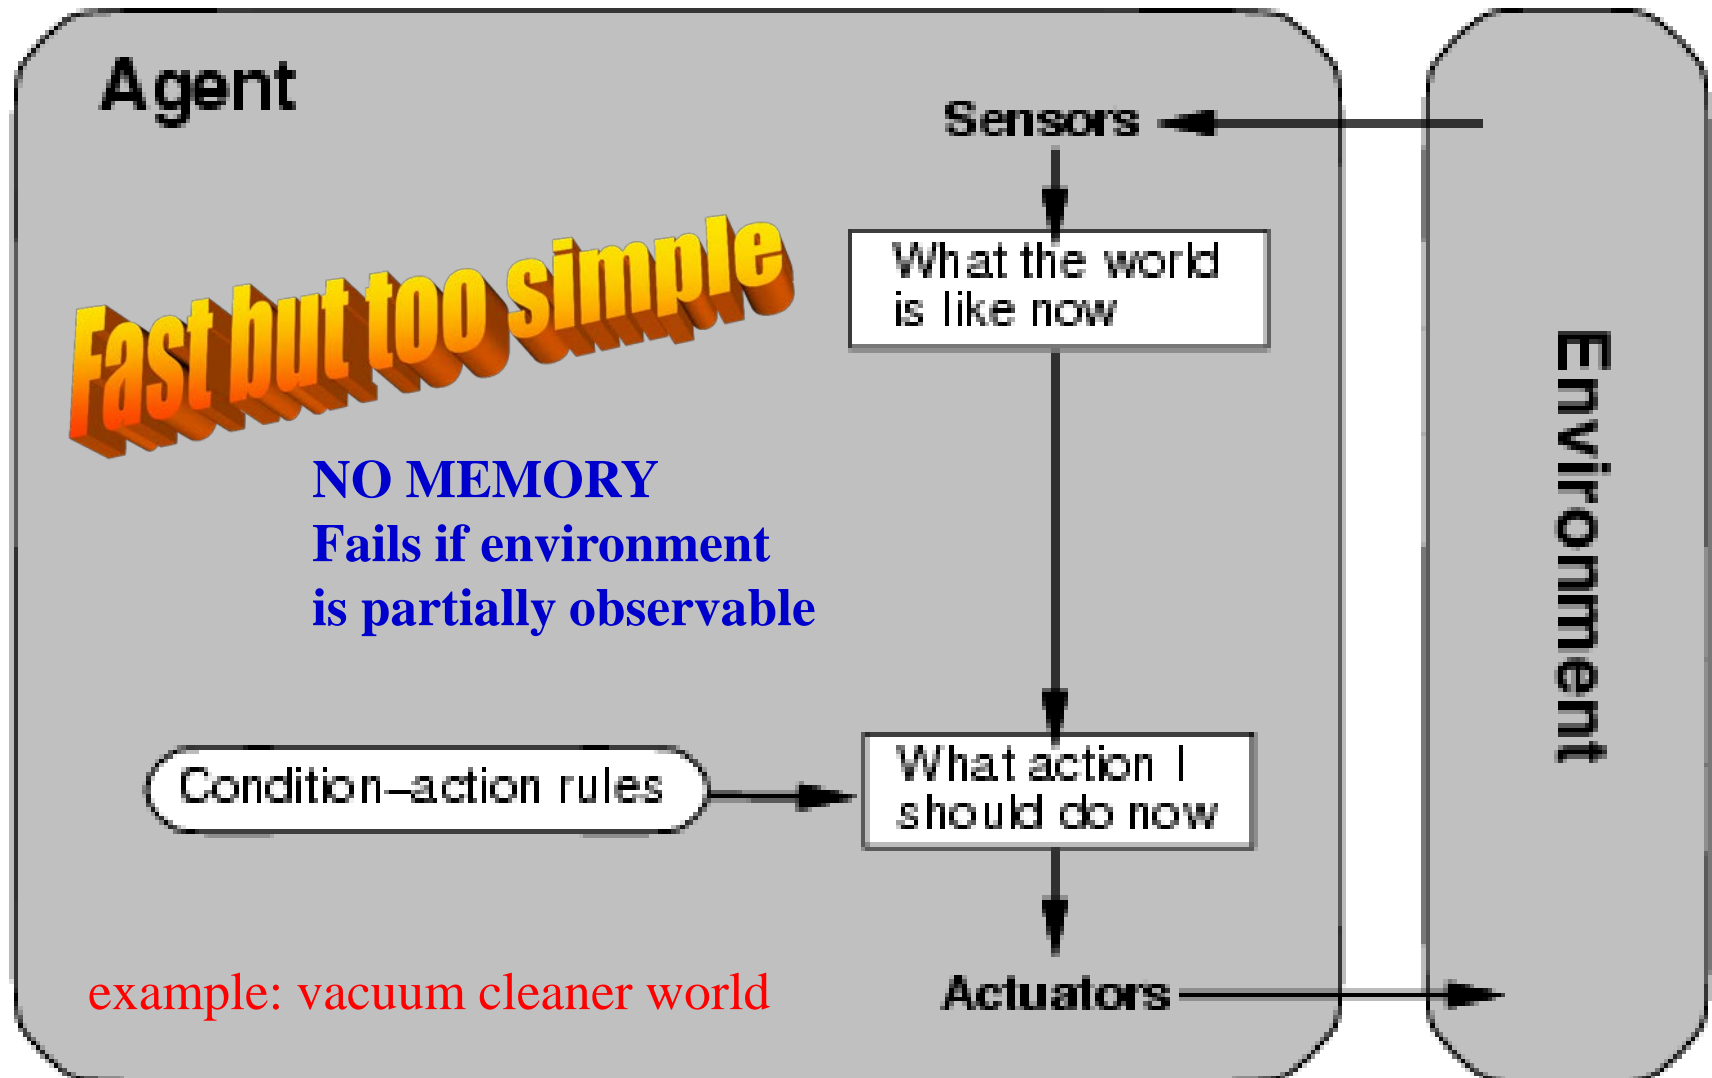
Five basic types in order of increasing generality:

- Table Driven agents
- Simple reflex agents
- Model-based reflex agents
- Goal-based agents
- Utility-based agents

Table Driven Agent.




Simple reflex agents



Model-based reflex agents


description of
current world state

Model the state of the world by:
modeling how the world changes
how its actions change the world




Goal-based agents

Goals provide reason to prefer one action over the other.
We need to predict the future: we need to plan & search



Utility-based agents


Some solutions to goal states are better than others.
Which one is best is given by a utility function.
Which combination of goals is preferred?



Learning agents

How does an agent improve over time?

By monitoring its performance and suggesting better modeling, new action rules, etc.



Summary

- **What is Artificial Intelligence?**
 - modeling humans' thinking, acting, should think, should act.
- **Intelligent agents**
 - We want to build agents that act rationally
 - Maximize *expected* performance measure
- **Task environment – PEAS**
 - Yield design constraints
- **Real-World Applications of AI**
 - AI is integrated in a broad range of products & systems
- **Reading**
 - Today: Ch. 1 & 2 in R&N
 - For next week: Ch. 3 in R&N (search)