Artificial Intelligence Class 2: Intelligent Agents

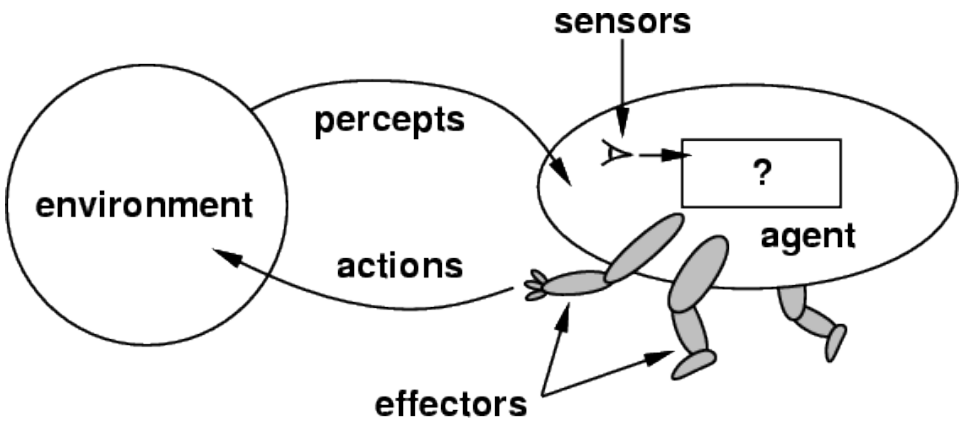

1

Bookkeeping

- Due last night:
  - Read academic integrity
  - Pre-reading for lecture

  } If you haven't done this, do!
- HW 1 (out at/by 11:59 PM)
  - Writing: 1 readings, 1 short (1-2pg) essay, ~6 questions
  - Coding: search
- Resources reminder
  - http://tiny.cc/ai-class
  - http://tiny.cc/ai-schedule
  - http://tiny.cc/ai-ai
  - **http://tiny.cc/ai-discord**

2

Today

- Defining an agent
- PEAS
- What is autonomy?
- Types of agents
- Defining environments
- Agent properties

3

Today's Class

- What's an agent?
  - Definition of an agent
  - Rationality and autonomy
  - Types of agents
  - Properties of environments

Agency is the capacity of individuals to act independently and to make their own free choices.

- Broadly: a thing that does something, with agency

4

What is an Agent?

- An intelligent agent is:
  - A (usually) autonomous entity which…
  - Observes an environment (the world)
  - Acts on its environment in order to achieve goals } Shows "agency"
- An intelligent agent **may** learn
  - Not always
  - A simple "reflex agent" still counts as an agent
- Behaves in a rational manner
  - Not "optimal"

6

How Do You Design an Agent?

- An intelligent agent:
  - Perceives its environment via **sensors**
  - Acts upon that environment with its **actuators** (or effectors)
- Properties:
  - Autonomous
  - Reactive to the environment
  - Pro-active (goal-directed)
  - Interacts with other agents via the environment

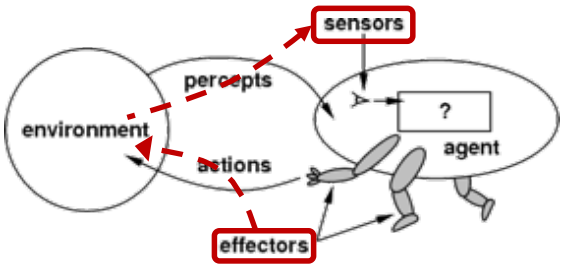

7

Human Sensors/Percepts, Actuators/Actions

- Sensors:
  - Eyes (vision), ears (hearing), skin (touch), tongue (gustation), nose (olfaction), neuromuscular system (proprioception), …
- Percepts: "that which is perceived"
  - At the lowest level – electrical signals from these sensors
  - After preprocessing – objects in the visual field (location, textures, colors, …), auditory streams (pitch, loudness, direction), …
- Actuators/effectors:
  - Limbs, digits, eyes, tongue, …
- Actions:
  - Lift a finger, turn left, walk, run, carry an object, …

The Point:

- Percepts and actions need to be *carefully defined*
- Sometimes at different levels of abstraction!

8

Example: Automated Taxi

- Percepts: Video, sonar, speedometer, odometer, engine sensors, keyboard input, microphone, GPS, …
- Actions: Turn, accelerate, brake, speak, display, …
- Goals: Maintain safety, reach destination, maximize profits (fuel, tire wear), obey laws, provide passenger comfort, …
- Environment: U.S. urban streets, freeways, traffic, pedestrians, weather, customers, …
- Different aspects of driving may require different types of agent programs.

9

Rational agents

- An agent should try to "do the right thing," per:
  - What it can perceive
  - The actions it can perform
- The "right" thing is the one that causes the agent to be most successful, per some **performance measure**
  - An objective criterion for success of an agent's behavior
- Example: performance measures of a vacuum cleaner:
  - Amount of dirt
  - Time taken
  - Others?
  - Electricity used
  - Noise generated

10

Rationality (formally)

- An ideal rational agent, in every possible world state, does action(s) that maximize its expected performance
- Based on:
  - The percept sequence (world state)
  - Its knowledge (built-in and acquired)
- Rationality includes information gathering
  - If you don't know something, find out!
  - No "rational ignorance"
- Need a performance measure
  - False alarm (false positive) and false dismissal (false negative) rates, speed, resources required, effect on environment, constraints met, user satisfaction, …

11

PEAS

- Agents must have:
- **P**erformance measure
- **E**nvironment
- **A**ctuators
- **S**ensors
- Must first specify the setting for intelligent agent design

13

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

Image: www.roboticstomorrow.com/article/2018/10/piece-picking-robots/12618

14

PEAS

- Agent: Interactive English tutor
- **Performance measure:** Maximize student's score on test
- **Environment:** Set of students
- **Actuators:** Screen display (exercises, suggestions, corrections)
- **Sensors:** Keyboard

Image: www.pexels.com/photo/black-boy-watching-video-on-laptop-5905700/

15

PEAS: Class example

- Specifying the setting
- Consider designing an automated taxi driver:
- Performance measure?
- Environment?
- Actuators?
- Sensors?

Image: techcrunch.com/2021/06/03/waymos-driverless-taxi-service-can-now-be-accessed-on-google-maps/

16

PEAS: Class example

- Agent: Medical diagnosis system
- Performance measure:
- Environment:
- Actuators:
- Sensors:

17

Autonomy

- An autonomous system is one that:
  - Determines its own behavior
  - Not all its decisions are included in its design
- It is not autonomous if all decisions are made by its designer according to a priori decisions
- "Good" autonomous agents need:
  - Enough built-in knowledge to survive
  - The ability to learn
- In practice this can be a bit slippery

18

Some Types of Agent

- Table-driven agents
  - Use a percept sequence/action table to find the next action
  - Implemented by a (large) lookup table
- Simple reflex agents
  - Based on condition-action rules
  - Implemented with a production system
  - Stateless devices which do not have memory of past world states
- Agents with memory
  - Have internal state
  - Used to keep track of past states of the world

19

Some Types of Agent

- Agents with goals
  - Have internal state information, plus…
  - Goal information about desirable situations
  - Agents of this kind can take future events into consideration
- Utility-based agents
  - Base their decisions on classic axiomatic utility theory
  - In order to act rationally

20

(1) Table-Driven Agents

- Table lookup of:
  - Percept-action pairs mapping
  - Every possible state → best action
- Problems:
  - Too big to generate and store
  - Don't know non-perceptual parts of state
    - E.g., background knowledge
  - Not adaptive to changes in the environment
    - Must update entire table
  - No looping
    - Can't condition actions on previous actions/states

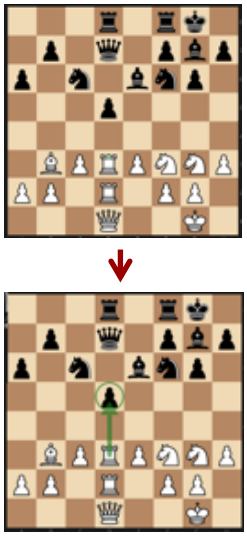

www.quora.com/How-do-you-know-if-your-chess-pieces-are-in-strategic-positions

21

(2) Simple Reflex Agents

- Rule-based reasoning
  - To map from percepts to optimal action
  - Each rule handles a collection of perceived states
    - "If your rook can capture a pawn, move the rook"
- Problems
  - Still usually too big to generate and to store
  - Still no knowledge of non-perceptual parts of state
  - Still not adaptive to changes in the environment
    - Change by updating collection of rules
  - Actions still not conditional on previous state

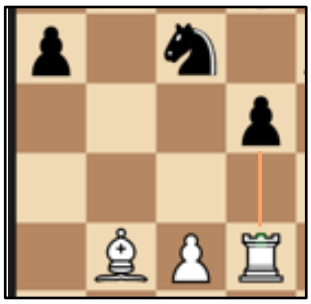

22

(1) Table-Driven/Reflex Agent

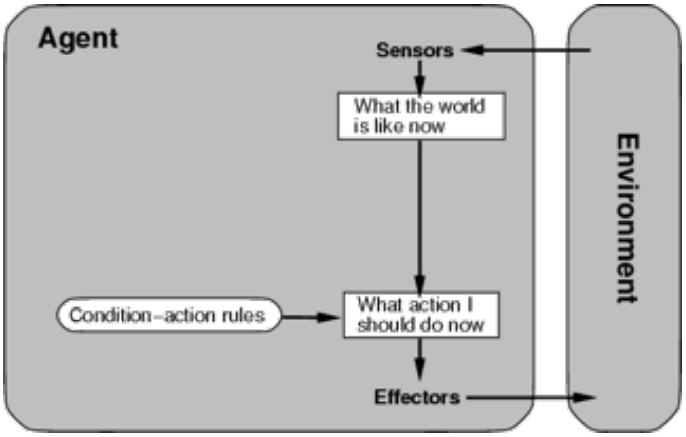

23

(3) Agents With Memory

- Encode "internal state" of the world
  - Used to remember the past (earlier percepts)
- Why?
  - Sensors rarely give the whole state of the world at each input
  - So, must build up environment model over time
  - "State" is used to encode different "worlds"
  - Different worlds can generate the same (immediate) percepts
- Requires ability to represent change in the world
  - Could represent just the latest state
  - But then can't reason about hypothetical courses of action

Image: www.roboticstomorrow.com/article/2018/10/piece-picking-robots/1261

24

(3) Architecture for an Agent with Memory

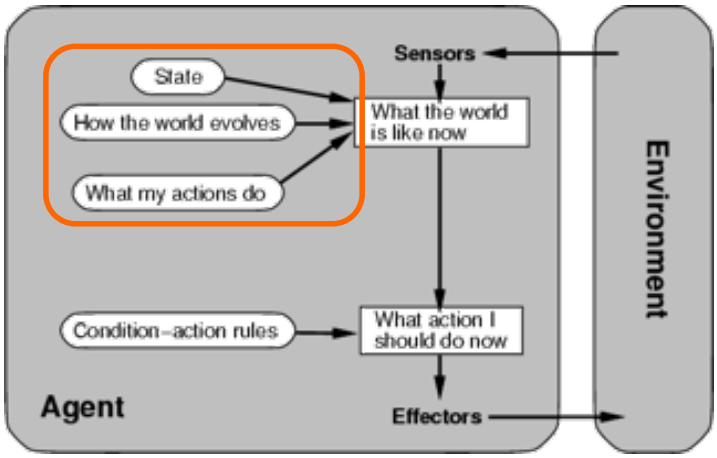

25

(4) Goal-Based Agents

- Choose actions that achieve a goal
  - Which may be given, or computed by the agent
- A goal is a description of a desirable state
  - Need goals to decide what situations are "good"
  - Keeping track of the current state is often not enough
- Deliberative instead of reactive
  - Must consider sequences of actions to get to goal
  - Involves thinking about the future
  - "What will happen if I do...?"

27

(4) Architecture for a Goal-Based Agent

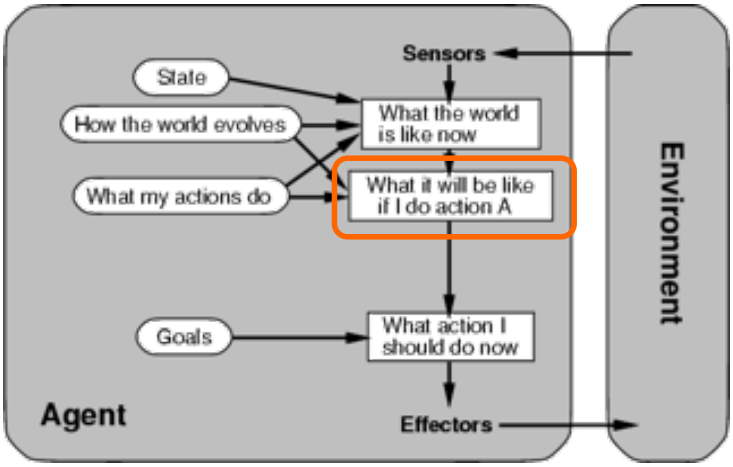

28

(5) Utility-Based Agents

- How to choose from multiple alternatives?
  - What action is best?
- What state is best?
  - Goals → crude distinction between "happy" / "unhappy" states
  - Often need a more general performance measure (*how* happy?)
- Utility function gives success or happiness **at a given state**
- Can compare choice between:
  - Conflicting goals
  - Likelihood of success
  - Importance of goal (if achievement is uncertain)

29

(4) Architecture for a complete utility-based agent

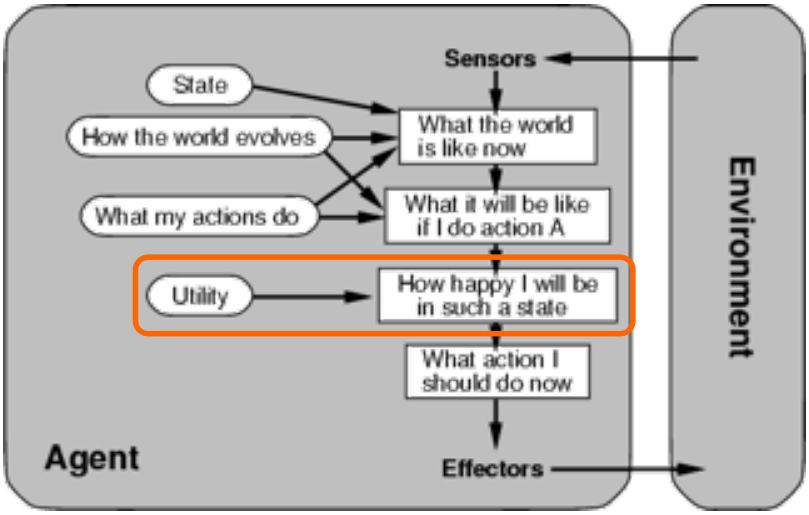

30

Properties of Environments

- Fully observable/Partially observable
  - If an agent's sensors give it access to the complete state of the environment, the environment is **fully observable**
  - Such environments are convenient
    - No need to keep track of the changes in the environment
    - No need to guess or reason about non-observed things
  - Such environments are also rare in practice

31

Properties of Environments

- Deterministic/Stochastic
  - An environment is **deterministic** if:
    - The next state of the environment is completely determined by
      - The current state of the environment
      - The action of the agent
  - In a stochastic environment, there are multiple, unpredictable outcomes
- In a fully observable, deterministic environment, the agent has no uncertainty.

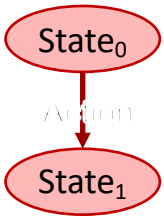

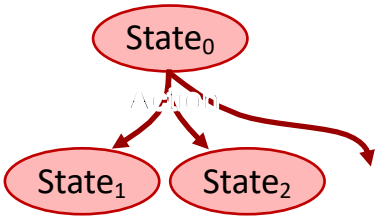

32

Properties of Environments II

- Episodic/Sequential
  - Episodic: subsequent episodes do not depend on what actions occurred in previous episodes
  - Sequential environment: Agent engages in a series of connected episodes
- Static/Dynamic
  - A static environment does not change while the agent is thinking
  - The passage of time as an agent deliberates is irrelevant
  - The agent doesn't need to observe the world during deliberation

33

Properties of Environments III

- Discrete/Continuous
  - If the number of distinct percepts and actions is limited, the environment is discrete, otherwise it is continuous.
    - Receives percepts describing the world one at a time
    - Maps this percept sequence to a sequence of discrete actions
- Single agent/Multi-agent
  - Whether the environment contains other intelligent agents.
  - In multi-agent environments, there are game-theoretic concerns (for either cooperative or competitive agents)
  - Single-agent environments are still more common.

34

Characteristics of Environments

	Fully observable?	Deterministic?	Episodic?	Static?	Discrete?	Single agent?
Solitaire						
Chess						
Taxi driving						
Internet shopping						
Medical diagnosis						

35

Characteristics of Environments

	Fully observable?	Deterministic?	Episodic?	Static?	Discrete?	Single agent?
Solitaire	**No**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**
Chess						
Taxi driving						
Internet shopping						
Medical diagnosis						

36

Characteristics of Environments

	Fully observable?	Deterministic?	Episodic?	Static?	Discrete?	Single agent?
Solitaire	No	Yes	Yes	Yes	Yes	Yes
Chess	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**No**
Taxi driving						
Internet shopping						
Medical diagnosis						

37

Characteristics of Environments

	Fully observable?	Deterministic?	Episodic?	Static?	Discrete?	Single agent?
Solitaire	No	Yes	Yes	Yes	Yes	Yes
Chess	Yes	Yes	Yes	Yes	Yes	No
Taxi driving	**No**	**No**	**No**	**No**	**No**	**No**
Internet shopping						
Medical diagnosis						

38

Characteristics of Environments

	Fully observable?	Deterministic?	Episodic?	Static?	Discrete?	Single agent?
Solitaire	No	Yes	Yes	Yes	Yes	Yes
Chess	Yes	Yes	Yes	Yes	Yes	No
Taxi driving	No	No	No	No	No	No
Internet shopping	**No**	**No**	**No**	**No**	**Yes**	**No**
Medical diagnosis						

39

Characteristics of Environments

	Fully observable?	Deterministic?	Episodic?	Static?	Discrete?	Single agent?
Solitaire	No	Yes	Yes	Yes	Yes	Yes
Chess	Yes	Yes	Yes	Yes	Yes	No
Taxi driving	No	No	No	No	No	No
Internet shopping	No	No	No	No	Yes	No
Medical diagnosis	**No**	**No**	**No**	**No**	**No**	**Yes**

40

Characteristics of Environments

	Fully observable?	Deterministic?	Episodic?	Static?	Discrete?	Single agent?
Solitaire	No	Yes	Yes	Yes	Yes	Yes
Chess	Yes	Yes	Yes	Yes	Yes	No
Taxi driving	**No**	**No**	**No**	**No**	**No**	**No**
Internet shopping	No	No	No	No	Yes	No
Medical diagnosis	**No**	**No**	**No**	**No**	**No**	**Yes**

→ Lots of (most?) real-world domains fall into the hardest cases! ←

41

Summary: Agents

- An agent:
  - Perceives and acts in an environment
  - Has an architecture
  - Is implemented by an agent program(s)
- An ideal agent:
  - Always chooses the "right" action
    - Which is, that which maximizes its expected performance
  - Given its percept sequence so far!
- An autonomous agent:
  - Uses its own experience to learn and make decisions
  - Not built-in knowledge, i.e., a priori world knowledge by the designer

42

Summary: Agents

- Representing knowledge is important for successful agent design
  - Percepts, actions and their effects, constraints, …
- The most challenging environments are:
  - Partially observable
  - Stochastic
  - Sequential
  - Dynamic
  - Continuous
  - Contain multiple intelligent agents

43