Intelligent Agents

CE417: Introduction to Artificial Intelligence
Sharif University of Technology
Fall 2023

Soleymani

Some slides have been adopted from Klein and Abdeel, CS188, UC Berkeley and some slides from Zettlemoyer, CSE-573, Washington University.

Agents

- An agent is anything that can be viewed as
  - **Sensors**: perceive environment
  - **Actuators**: act upon environment

- Samples of agents
  - Human agent
    - Sensors: eyes, ears, and other organs for sensors
    - Actuators: hands, legs, vocal tract, and other movable or changeable body parts
  - Robotic agent
    - Sensors: cameras and infrared range finders
    - Actuators: various motors
  - Software agents
    - Sensors: keystrokes, file contents, received network packages
    - Actuators: displays on the screen, files, sent network packets

Agents & environments

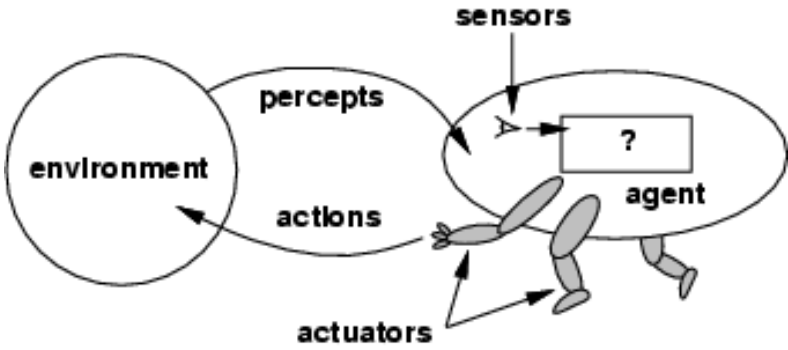

- Agent behavior can be described as an agent function that maps entire perception histories to actions:

$$f: P^* \rightarrow A$$

P^*: Percept sequence to date

A: Action set

Vacuum-cleaner world

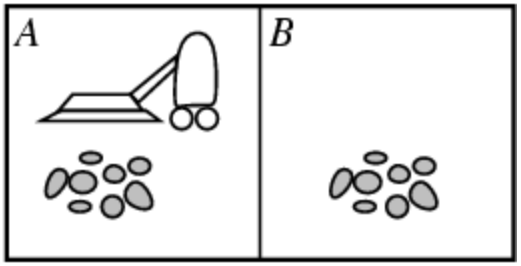

- Percepts: location and dirt/clean status of its location
  - e.g., [A,Dirty]

- Actions: *Left*, *Right*, *Suck*, *NoOp*

One simple rule implementing the agent function:
If the current square is dirty then suck, otherwise move to the other square

Rational agents

- "**do the right thing**" based on the perception history and the actions it can perform.

- **Rational Agent**: 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.

Performance measure

- Evaluates the sequence of environment states
- Vacuum-cleaner agent: samples of performance measure
  - ✖ Amount of dirt cleaned up
  - ☑ One point award for each clean square at each time step
    - Penalty for electricity consumption & generated noise
    - Mediocre job or periods of high and low activation?

Rational agents (vacuum cleaner example)

- Is this rational? If dirty then suck, otherwise move to the other square
  - Depends on
    - Performance measure, e.g., Penalty for energy consumption?
    - Environment, e.g., New dirt can appear?
    - Actuators, e.g., No-op action?
    - Sensors, e.g., Only sense dirt in its location?

Rationality vs. Omniscience

- Rationality is distinct from omniscience (all-knowing with infinite knowledge, impossible in reality)

- Doing actions in order to modify future percepts to obtain useful information
  - information gathering or exploration (important for rationality)
    - e.g., eyeballs and/or neck movement in human to see different directions

Autonomy

- An agent is autonomous if its behavior is determined by its own experience (with ability to learn and adapt)
  - Not just relies only on prior knowledge of designer
  - Learns to compensate for partial or incorrect prior knowledge
    - Benefit: changing environment
    - Starts by acting randomly or based on designer knowledge and then learns form experience
  - Rational agent should be autonomous

- Example: vacuum-cleaner agent
  - If dirty then suck, otherwise move to the other square
    - Does it yield an autonomous agent?
  - learning to foresee occurrence of dirt in squares

Task Environment (PEAS)

- Performance measure
- Environment
- Actuators
- Sensors

PEAS Samples…

- 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, display
  - Sensors: Cameras, sonar, speedometer, GPS, odometer, accelerometer, engine sensors, keyboard

PEAS Samples…

- Agent: Medical diagnosis system
  - Performance measure: Healthy patient, minimize costs
  - Environment: Patient, hospital, staff
  - Actuators: Screen display (questions, tests, diagnoses, treatments, referrals)
  - Sensors: Keyboard (entry of symptoms, findings, patient's answers)

PEAS Samples…

- Satellite image analysis system
  - Performance measure: Correct image categorization
  - Environment: Downlink from orbiting satellite
  - Actuators: Display of scene categorization
  - Sensors: Color pixel array

PEAS Samples…

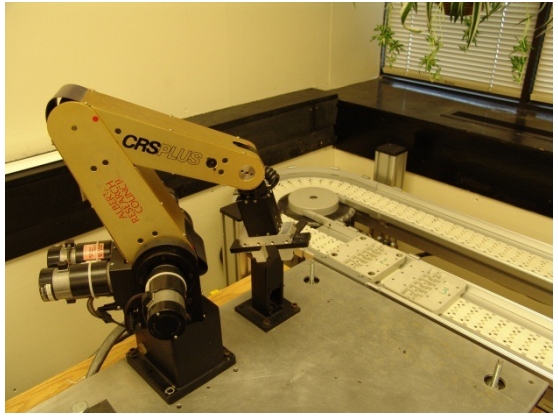

- Agent: Part picking robot
  - Performance measure: Percentage of parts in correct bins
  - Environment: Conveyor belt with parts, bins
  - Actuators: Jointed arm and hand
  - Sensors: Camera, joint angle sensors

PEAS Samples…

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

PEAS Samples…

- Agent: Pacman
  - Performance measure: Score, lives
  - Environment: Maze containing white dots, four ghosts, power pills, occasionally appearing fruit
  - Actuators: Arrow keys
  - Sensors: Game screen

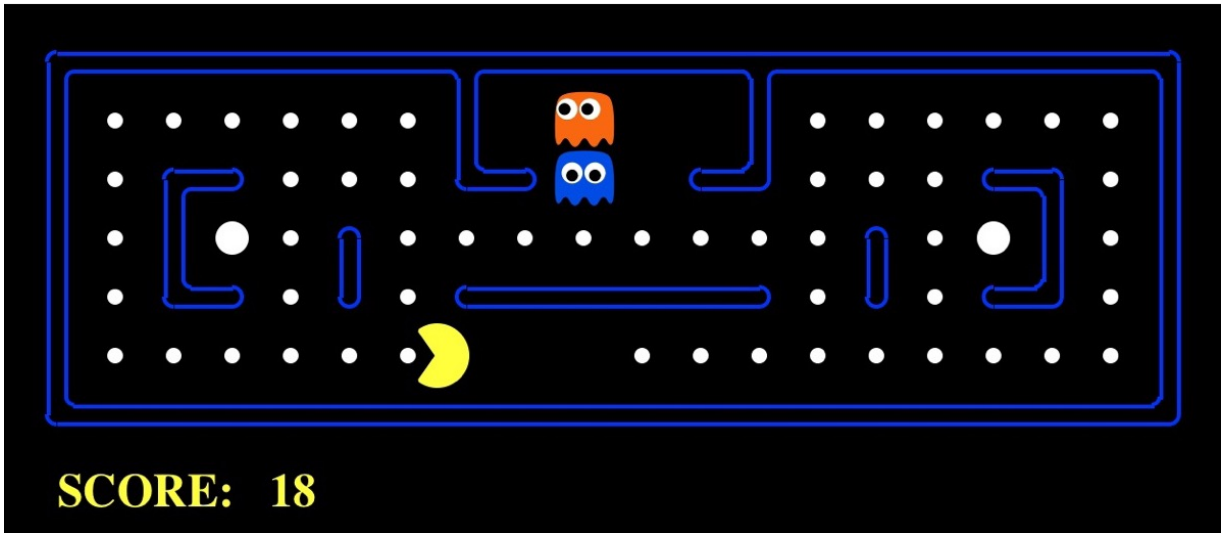

Environment types

- Fully observable (vs. partially observable): Sensors give access to the complete state of the environment at each time
  - Sensors detect all aspects relevant to the choice of action
  - Convenient (need not any internal state)
  - Noisy and inaccurate sensors or missing parts of the state from sensors cause partially observability

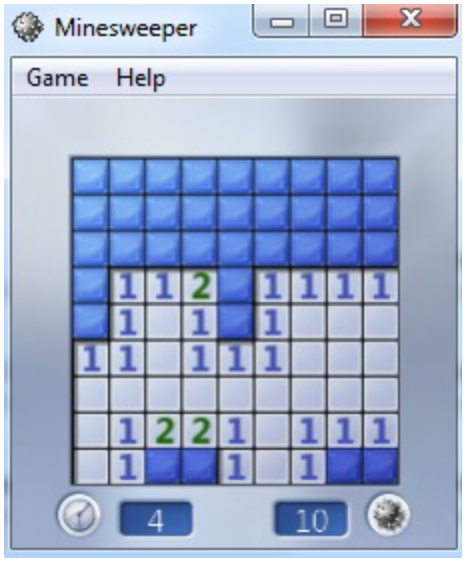

Fully observable vs. partially observable

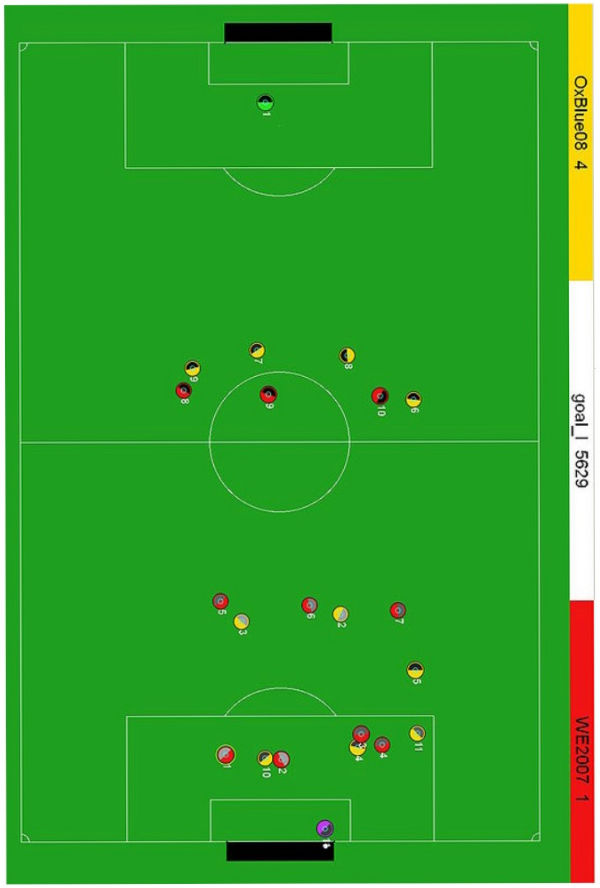

vs.

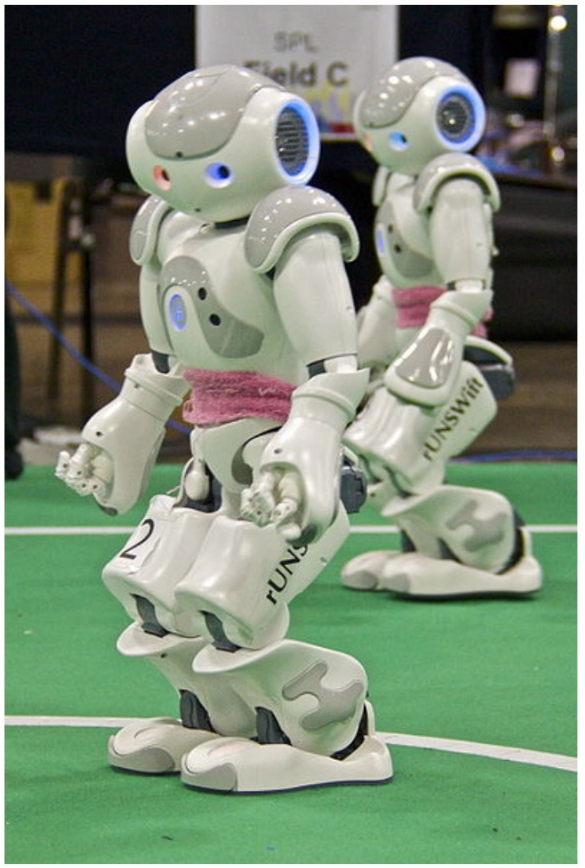

Environment types

- Deterministic (vs. stochastic): Next state is completely determined by the current state and the executed action
  - Partially observable environment could appear to be stochastic.
  - Environment is uncertain if it is not fully observable or not deterministic

vs.

Environment types

- Single agent (vs. multi-agent):
  - Crossword puzzle is a single-agent game (chess is a multi-agent one)
  - Is B an agent or just an object in the environment?
    - B is an agent when its behavior depends on A's behavior.
  - Multi-agent: competitive, cooperative

vs.

Environment types

- Discrete (vs. continuous): A limited number of distinct, clearly defined states, percepts and actions, time steps
  - Chess has finite number of discrete states, and discrete set of percepts and actions while Taxi driving has continuous states, and actions

vs.

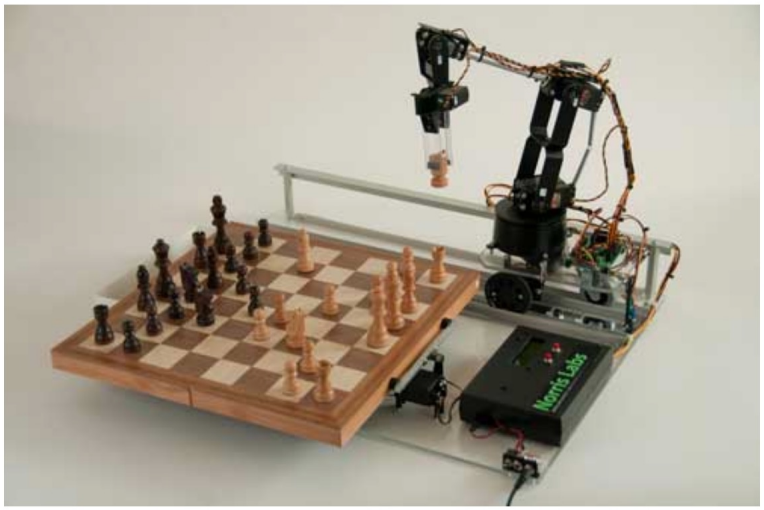

Environment types

- Episodic (vs. sequential): The agent's experience is divided into atomic "episodes“ where the choice of action in each episode depends only on the episode itself.
  - E.g., spotting defective parts on an assembly line (independency)
  - In sequential environments, short-term actions can have long-term consequences
    - Episodic environment can be much simpler

vs.

Environment types

- Static (vs. dynamic): The environment is unchanged while an agent is deliberating.
  - Semi-dynamic: if the environment itself does not change with the passage of time but the agent's performance score does.
  - Static (cross-word puzzles), dynamic (taxi driver), semi-dynamic (clock chess)

Environment types

- Known (vs. unknown): the outcomes or (outcomes probabilities for all actions are given.
  - It is not strictly a property of the environment
    - Related to agent's or designer's state of knowledge about "laws of physics" of the environment

- The real world is partially observable, multi-agent, stochastic, sequential, dynamic, continuous, (and unknown)
  - Hardest type of environment
  - The environment type largely determines the agent design

Pacman game

- Fully observable?
- Single-agent?
- Deterministic?
- Discrete?
- Episodic?
- Static?
- Known?

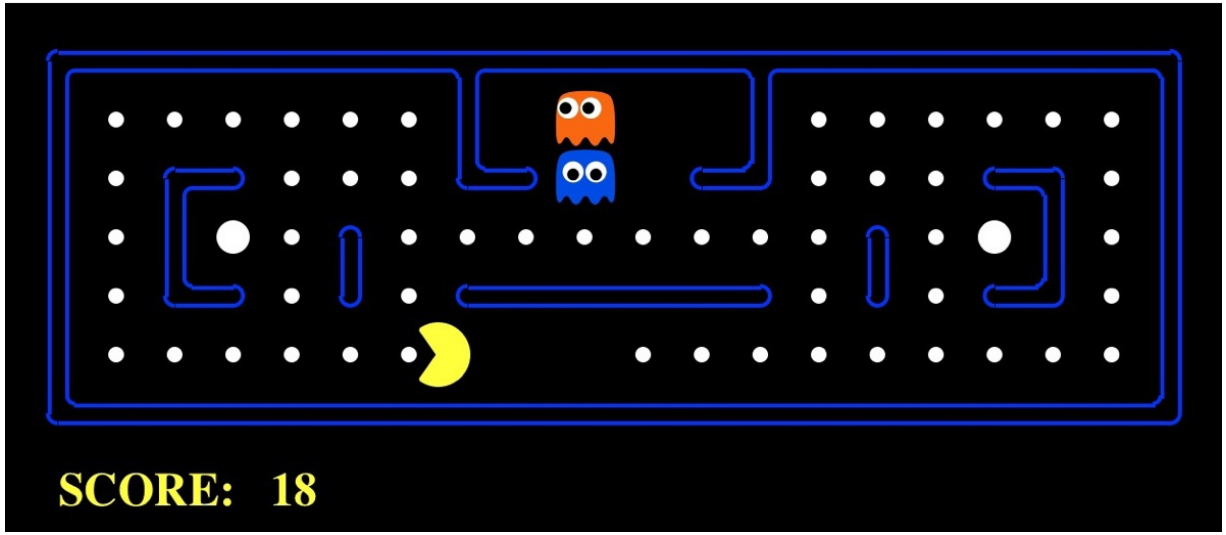

Environment types

Task Environment	Observable	Agents	Deterministic	Episodic	Static	Discrete
Crossword puzzle	Fully	Single	Deterministic	Sequential	Static	Discrete
Chess with a clock	Fully	Multi	Deterministic	Sequential	Semi	Discrete
Poker	Partially	Multi	Stochastic	Sequential	Static	Discrete
Backgammon	Fully	Multi	Stochastic	Sequential	Static	Discrete
Taxi driving	Partially	Multi	Stochastic.	Sequential	Dynamic	Continuous
Medical diagnosis	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Image analysis	Fully	Single	Deterministic	Episodic	Semi	Continuous
Part-picking robot	Partially	Single	Stochastic	Episodic	Dynamic	Continuous
Refinery controller	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Interactive. English tutor	Partially	Multi	Stochastic	Sequential	Dynamic	Discrete

Environment types

Task Environment	Observable	Agents	Deterministic	Episodic	Static	Discrete
Crossword puzzle	Fully	Single	Deterministic	Sequential	Static	Discrete
Chess with a clock	Fully	Multi	Deterministic	Sequential	Semi	Discrete
Poker	Partially	Multi	Stochastic	Sequential	Static	Discrete
Backgammon	Fully	Multi	Stochastic	Sequential	Static	Discrete
Taxi driving	Partially	Multi	Stochastic.	Sequential	Dynamic	Continuous
Medical diagnosis	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Image analysis	Fully	Single	Deterministic	Episodic	Semi	Continuous
Part-picking robot	Partially	Single	Stochastic	Episodic	Dynamic	Continuous
Refinery controller	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Interactive. English tutor	Partially	Multi	Stochastic	Sequential	Dynamic	Discrete

Agent program types

- Lookup table

- Basic types of agent program in order of increasing generality:
  - Reflexive
    - Simple reflexive
    - Model-based reflex agents
  - Planning-based agents
    - Goal-based agents
    - Utility-based agents
  - Learning-based agents

Reflex agents

- Reflex agents:
  - Choose action based on current percept (and maybe memory)
  - May have memory or a model of the world's current state
  - Do not consider the future consequences of their actions
  - Consider how the world is
- Can a reflex agent be rational?

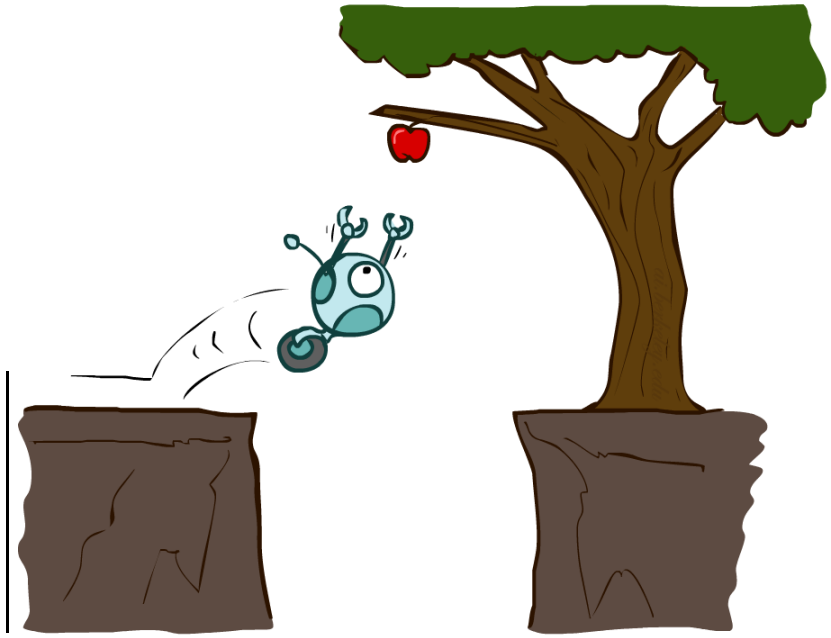

Reflex agents

- Simple, but very limited intelligence
- Simple reflex agents works only if the correct decision can be made on the basis of the current percept (fully observability)

Simple reflex agents

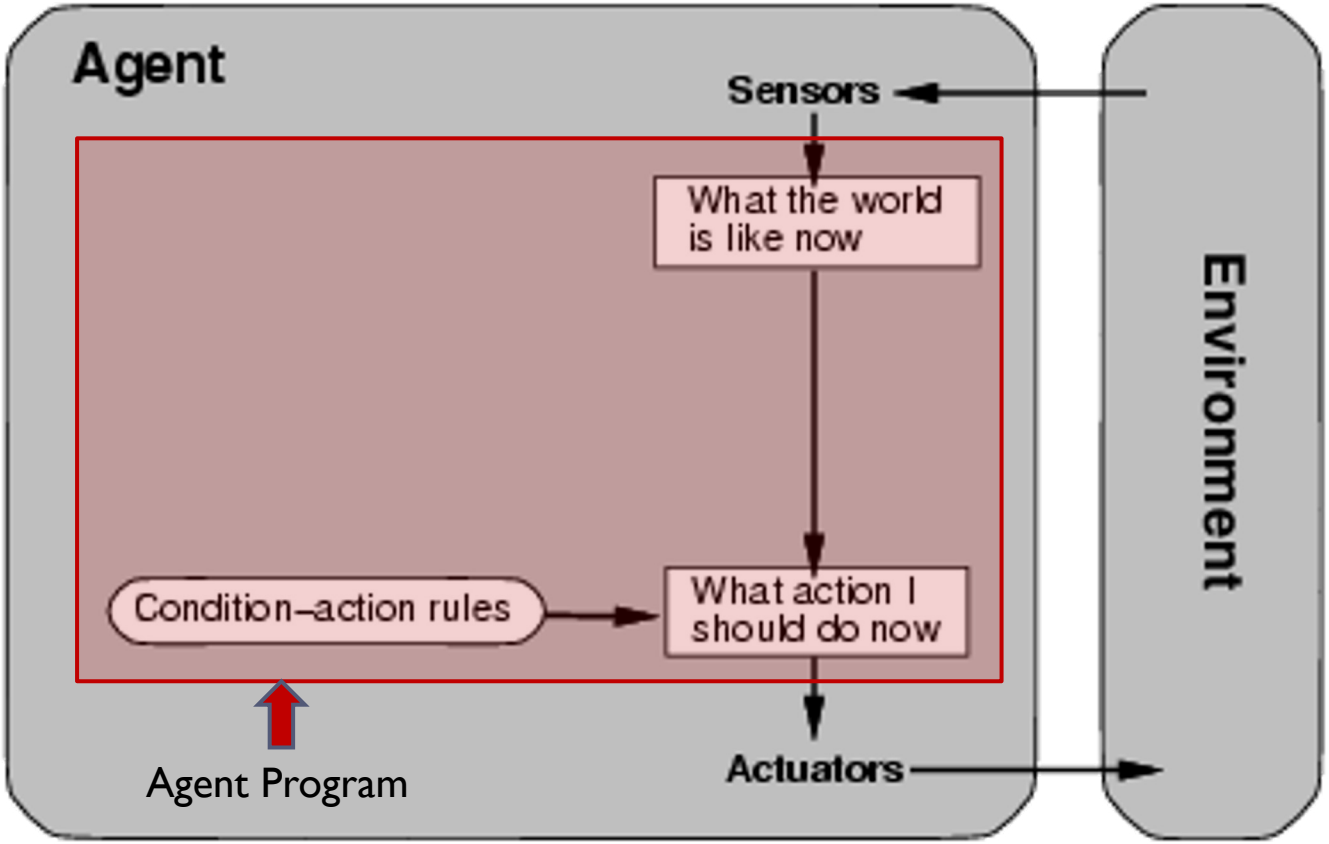

Video of demo reflex optimal

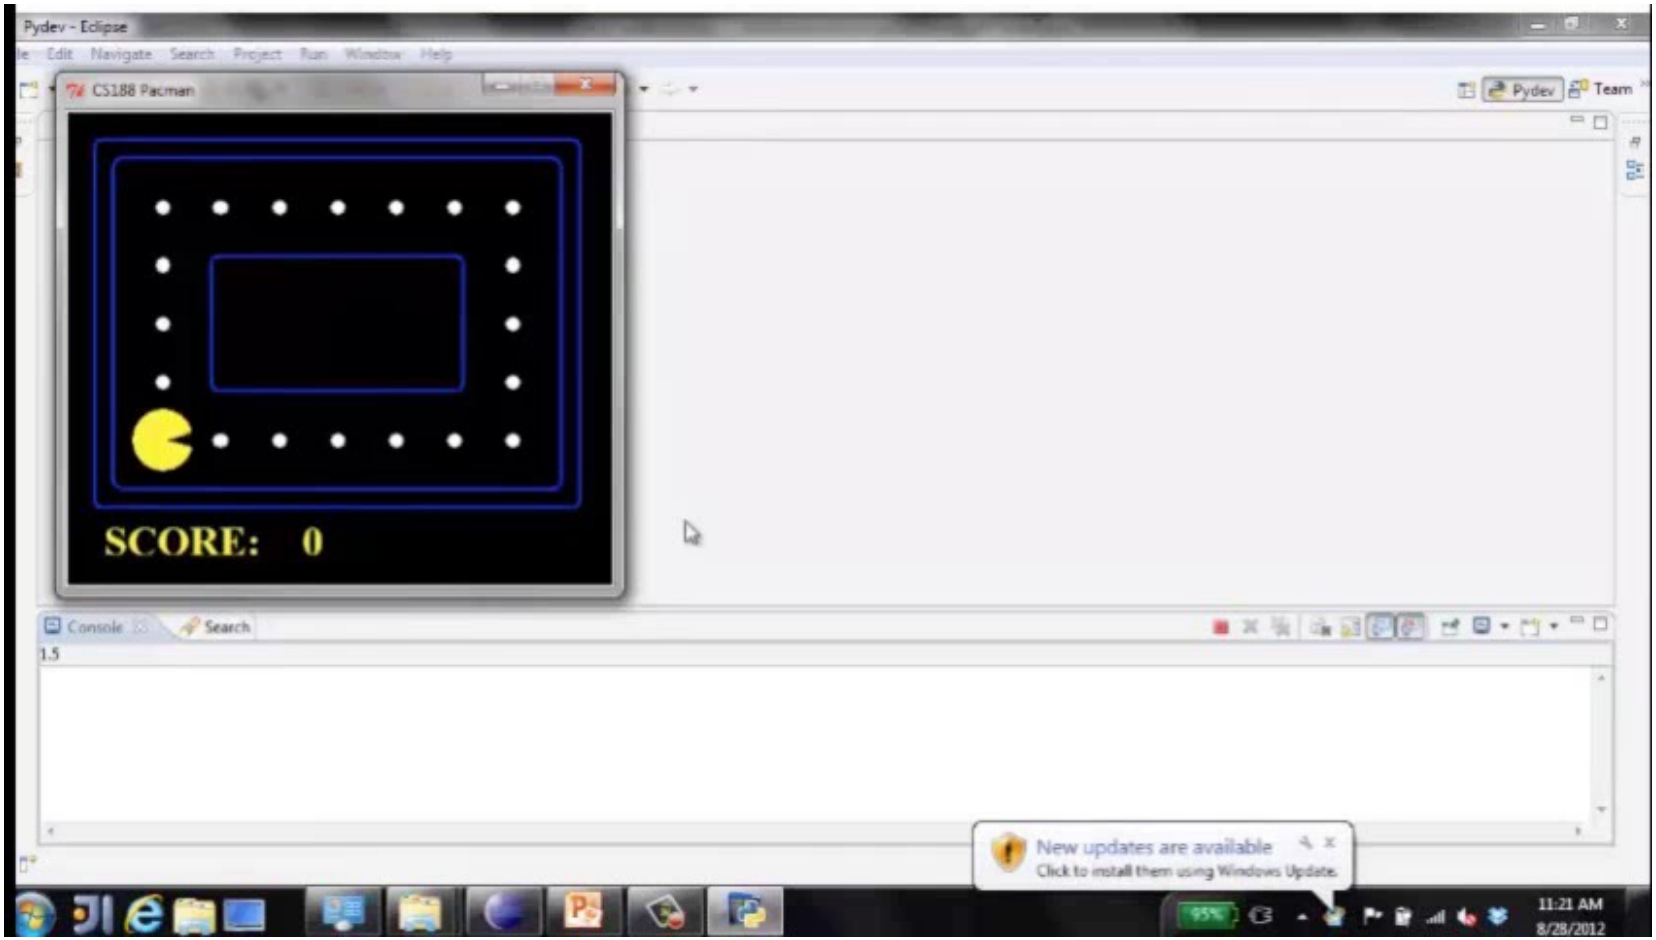

Video of demo reflex odd

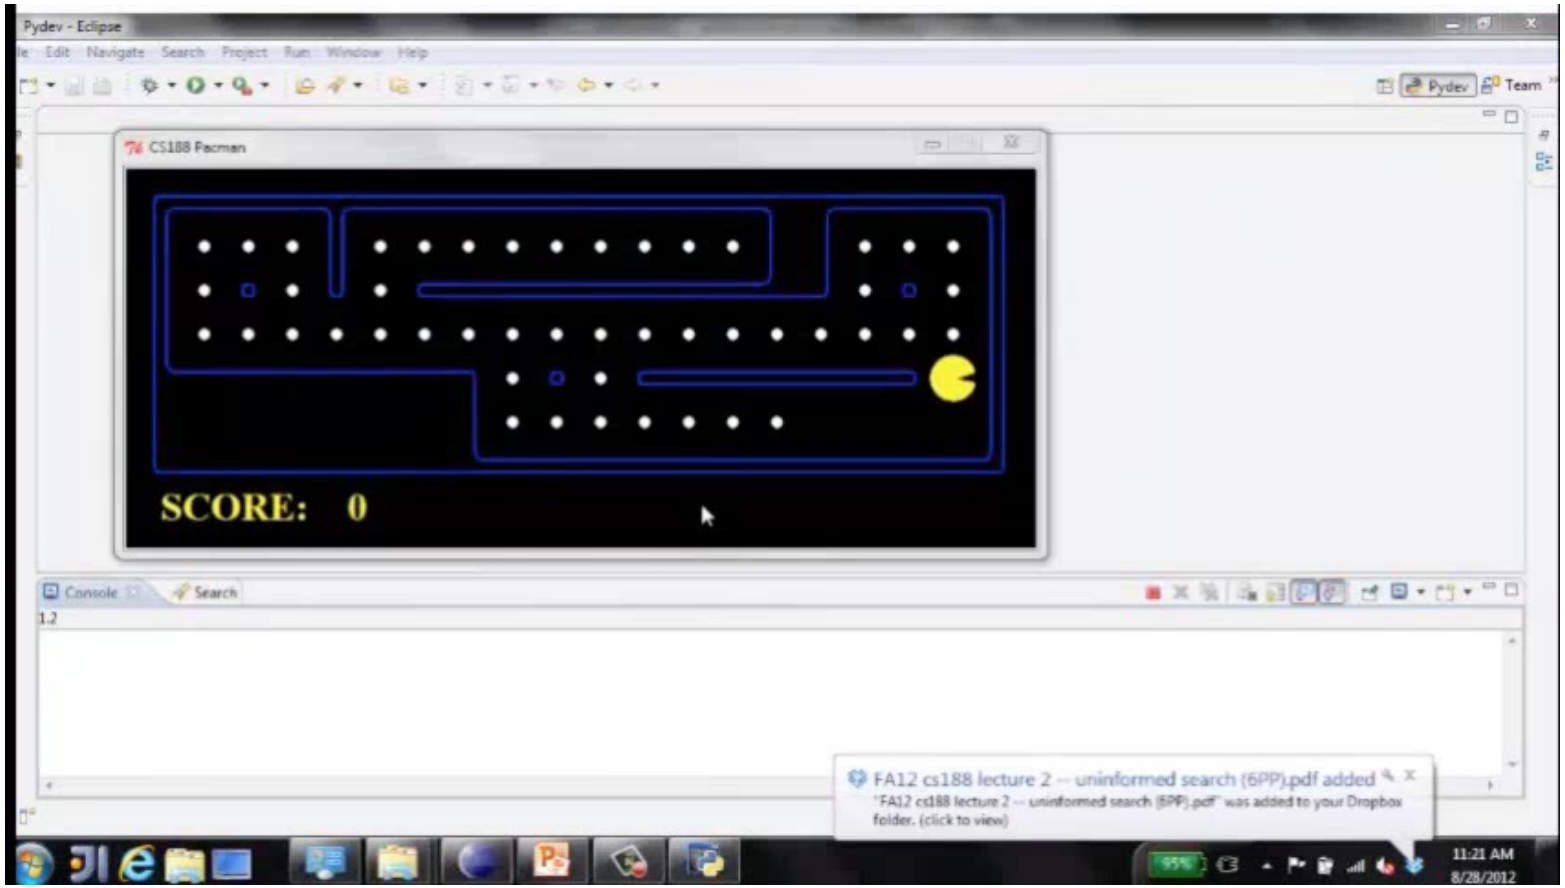

Agents that plan

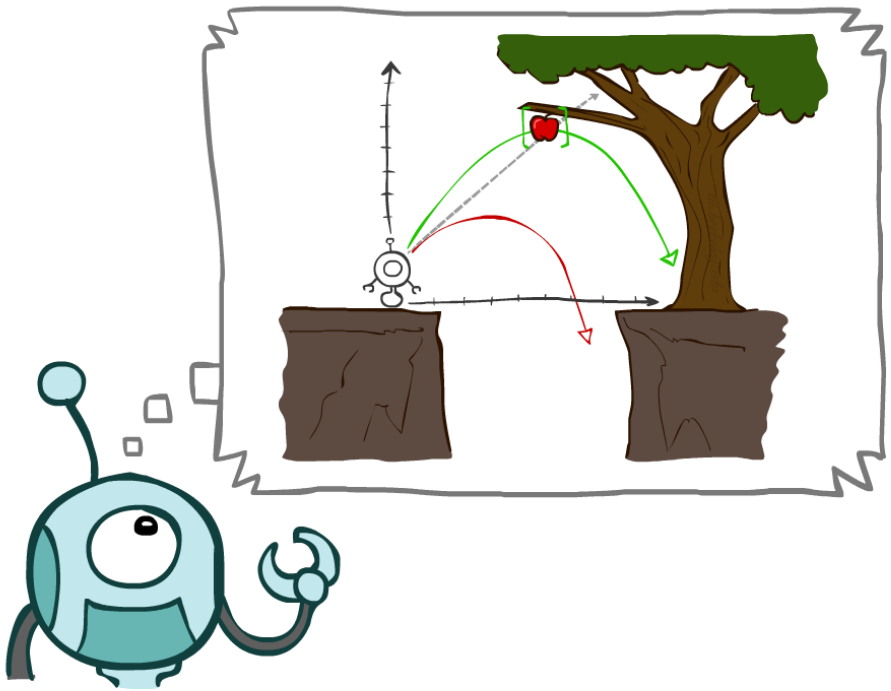

Goal-based agents

- Knowing about the current state is not always enough to decide what to do
- Situations that are desirable must be specified (goal)
  - According to performance measure
- Usually requires search and planning
  - to find action sequences achieving goal

Goal-based (Planning) agents

- Planning agents:
  - Ask “what if”
  - Decisions based on (hypothesized) consequences of actions
  - Must have a model of how the world evolves in response to actions
  - Must formulate a goal (test)
  - Consider how the world WOULD BE
- Planning vs. replanning

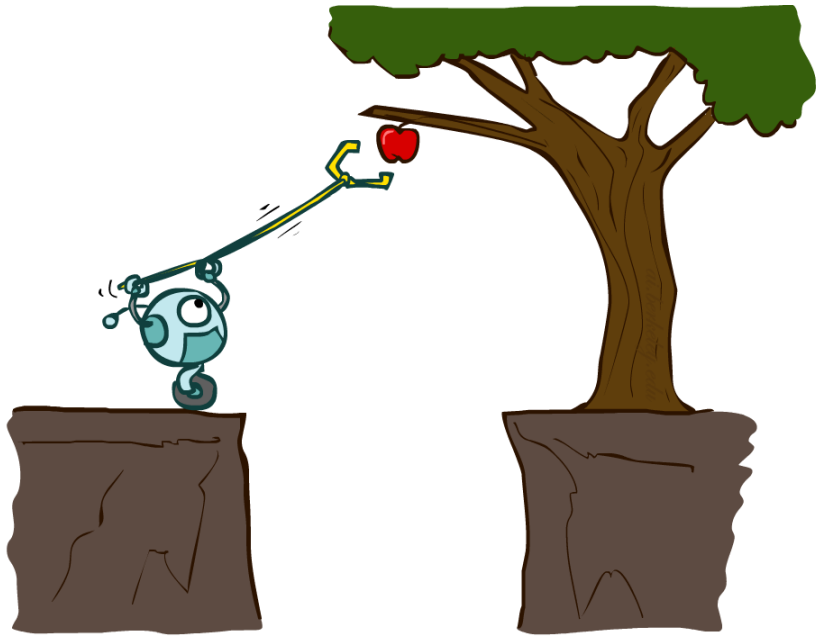

Video of demo mastermind

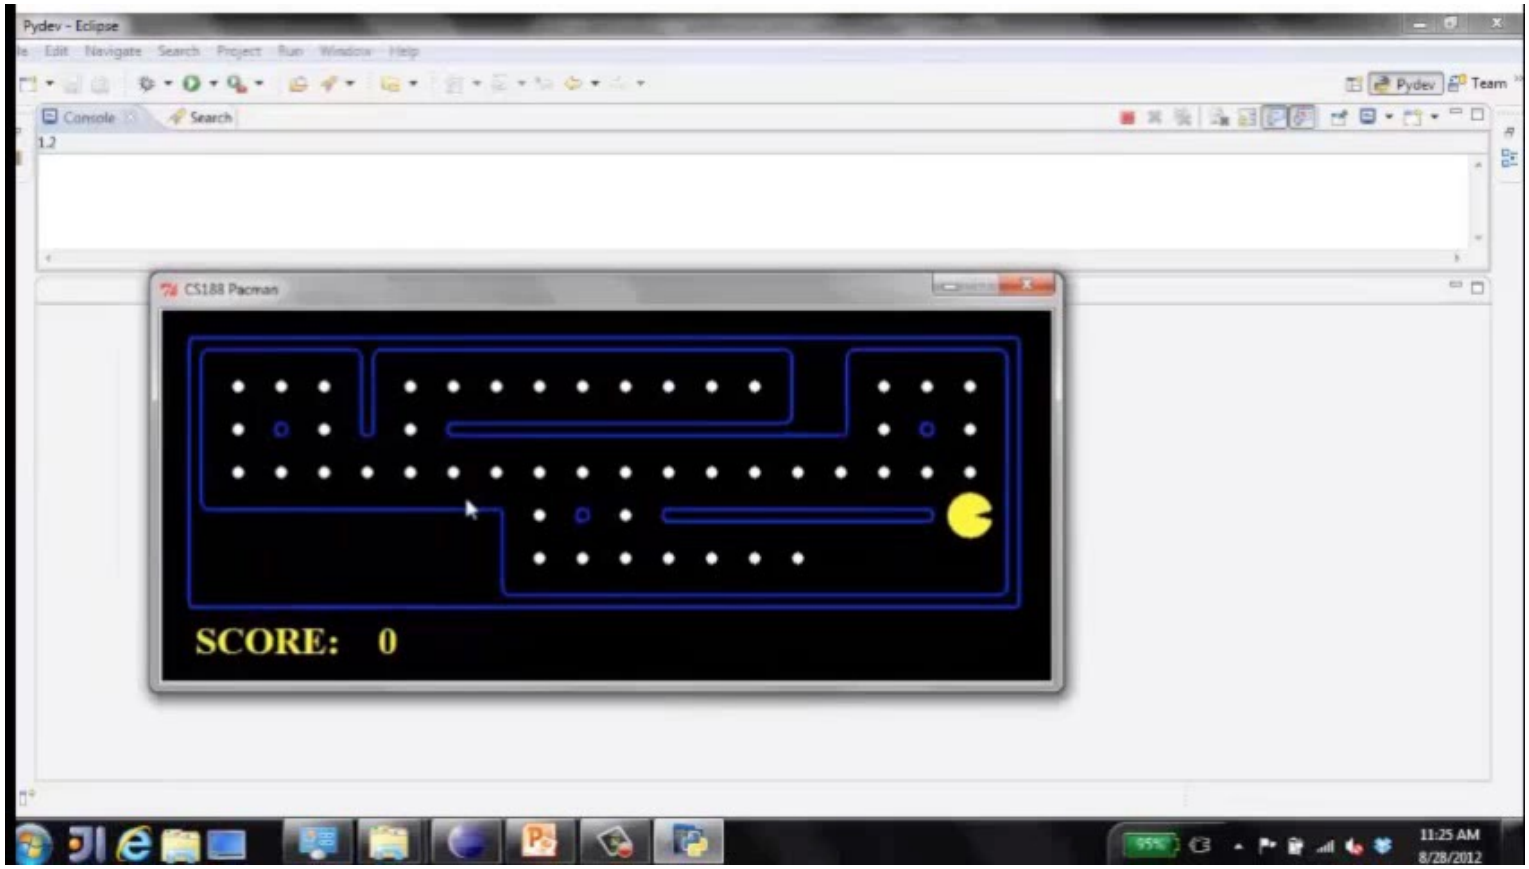

Video of demo replanning

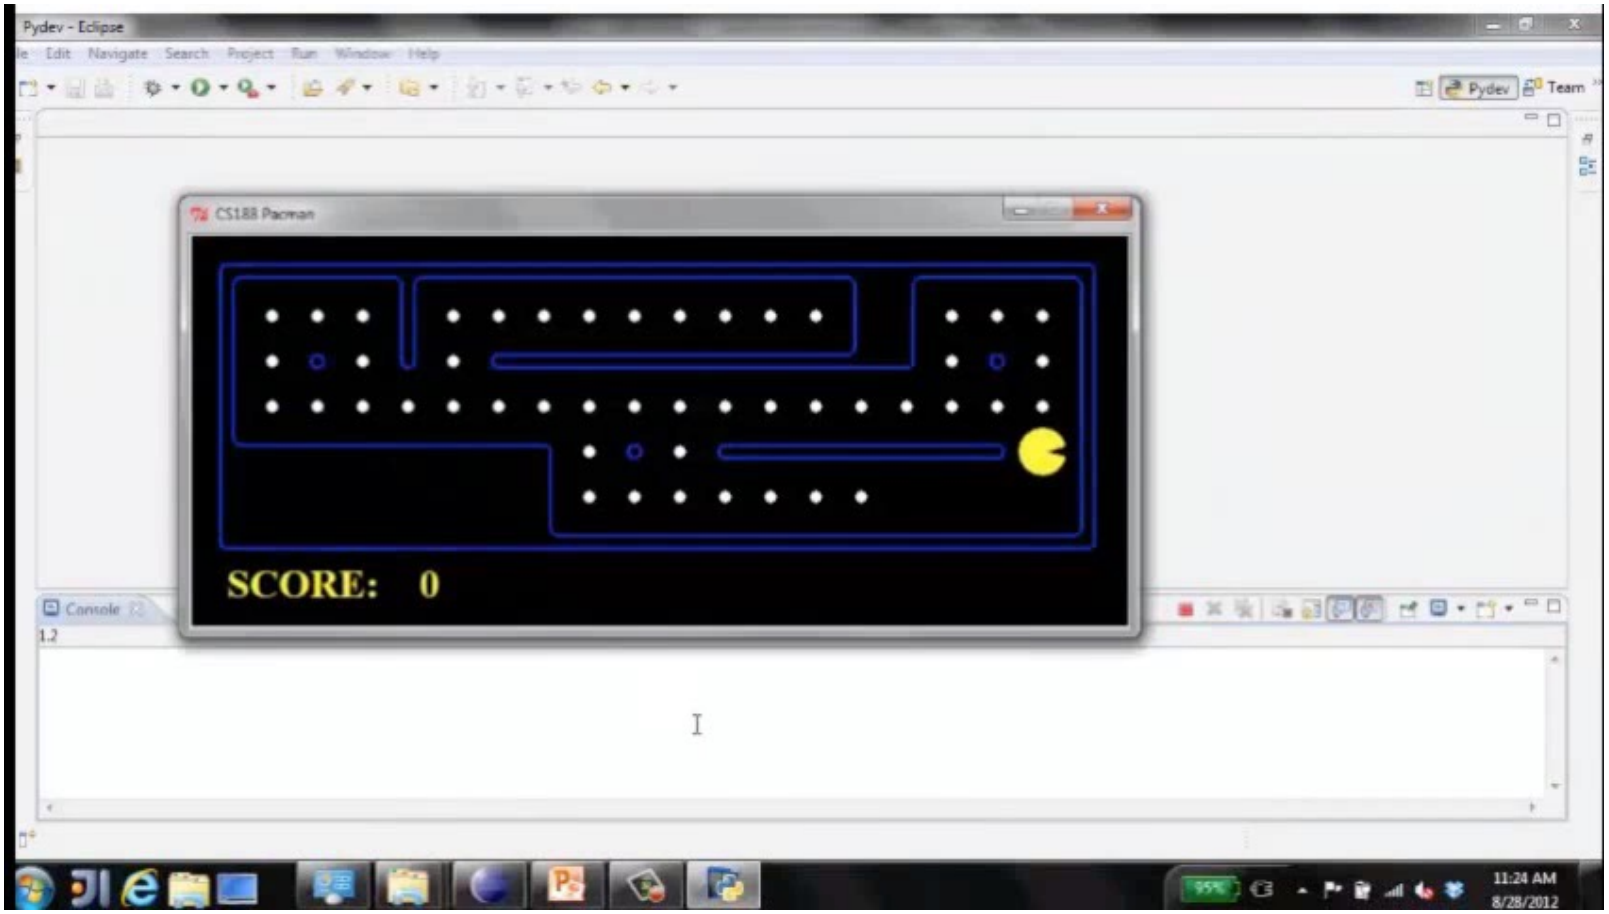

Goal-based agents vs. reflex-based agents

- Goal-based agents may be computationally less efficient but are more flexible
  - Example: going to a new destination
    - Goal-based agent: specifying that destination as the goal and solve a search problem
    - Reflexive agent: agent's rules for when to turn and when to go straight must be rewritten

Utility-based agents

- More general performance measure than goals
  - How happy would each world state make the agent?
- Advantages
  - Like goal-based agents show flexibility and learning advantages
  - Can trade-off conflicting goals (e.g. speed and safety)
  - Where none of several goals can be achieved with certainty
    - likelihood of success can be weighted by importance of goals
- Rational utility-based agent chooses the action that maximizes the expected utility of action outcomes

Learning-based agents

- Agent may not know the environment completely.
- Learning-based agents try to do the best by both exploring the environment and using the gathered information to decide rationally.