

Announcements

Project 0: Python Tutorial.

Due on Monday 1/27. 11:59 pm.

Homework 0: Math Diagnostic.

Due on Wednesday 1/29. 11:59 pm.

Project 1: Search.

Out: longer than average. Get started.

Due Friday. 2/7. 11:59 pm.

Sections.

Start next week!

Office Hours: still working on rooms.

Department/university issues still being dealt with.

Check: signed up for Piazza and Gradescope?

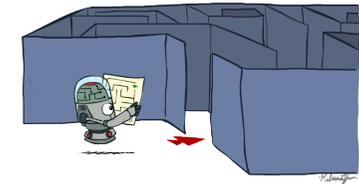
Pinned Post: AI in the news!

Lecture Attendance Link

Lecture Attendance Link: <http://bit.ly/2GEMokS>

CS 188: Artificial Intelligence

Search.



Today.

Agents that Plan Ahead

Search Problems.

Model world with state space.

Setting up state spaces.

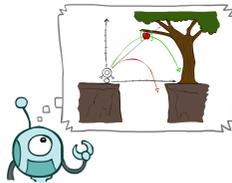
Maybe today.

Uninformed Search Methods:

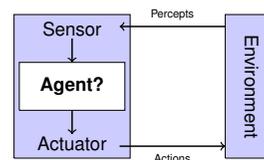
Depth-First Search

Breadth-First Search

Uniform-Cost Search



Frame of AI



An agent **perceives** its environment with **sensors** and **acts** on environment using **actuators**.

Car:

Sensors: camera, lidar, speed gauge, ..

Actuators: gas pedal, steering wheel, brake pedal, ...

Website, Program, ...

Input/Output: Outputs affect environment, which affects input, which outputs, ...

Rationality

A **rational agent** chooses actions that maximize **expected utility**.

Today: agents that have a goal, and a cost.

E.g., reach goal with lowest cost.

Later: agents have numerical utilities, rewards, etc.

E.g., takes action that maximizes total reward over time.

(Reward: largest total profit. or expected total profit.)

Agent Design

The environment largely determines the agent design.

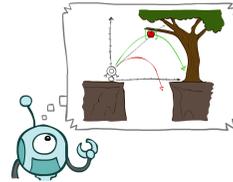
Fully/partially observable → agent request **memory** (internal state)

Discrete/ continuous → agent can/can't enumerate **all states**

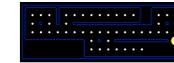
Stochastic/deterministic → agent deals with **contingencies**

Single-agent/multi-agent → agent may need to behave **randomly**.

Agents that Plan



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?

Examples:

Stove: hot, ouch!

Car: deer, brake!

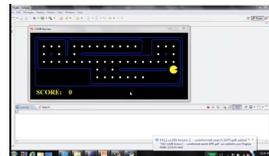
Go to Tahoe: skid in snow, ???

[Demo: reflex (L2D1 and L2D2)]

Video of Demo Reflex Optimal



Video of Demo Reflex Odd



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**.

Optimal or not optimal.

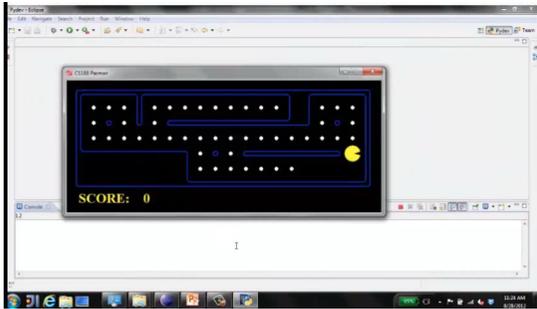
Complete or not.

Planning vs. replanning

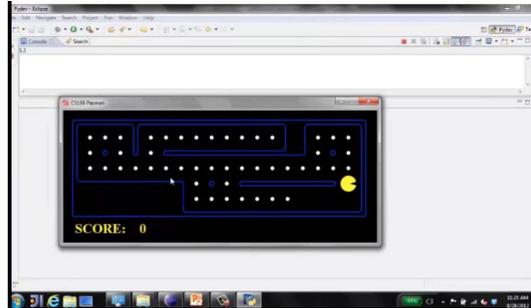
[Demo: nearest dot re-planning (L2D3), mastermind (L2D4)]

Video of Demo

Replanning



Video of Demo Mastermind



Search Problems



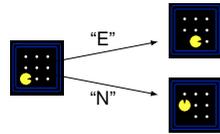
Search Problems

A **search problem** consists of:

A state space



A successor function (with actions, costs)



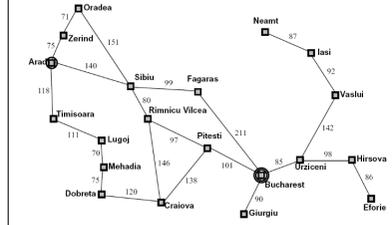
A start state and a goal test

A solution is a sequence of actions (a plan) which transforms the start state to a goal state

Search Problems Are Models



Example: Traveling in Romania



State space:
Cities
Successor function:
Roads: Neighboring city with cost = distance
Start state:
Arad
Goal test:
Is state == Bucharest?
Solution?

What's in a State Space?

The world state includes every last detail of the environment.



Search state keeps only details needed for planning (abstraction).

Problem: Pathing

States: (x, y) location

Actions: NSEW

Successor: update location only

Goal test: is $(x, y) = \text{END}$?

Problem: Eat-All-Dots

States: $\{(x, y), \text{dot booleans}\}$

Actions: NSEW

Successor: update location and

possibly a dot boolean

Goal test: dots all false

Agent Design

The environment largely determines the agent design.

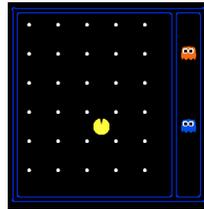
Fully/partially observable \rightarrow agent request **memory** (internal state)

Discrete/continuous \rightarrow agent can/can't enumerate **all states**

Stochastic/deterministic \rightarrow agent deals with **contingencies**

Single-agent/multi-agent \rightarrow agent may need to behave **randomly**.

State Space Sizes?



World state:

Agent positions: 120

Food count: 30

Ghost positions: 12

Agent facing: NSEW

How many World states?

$$120 \times (2^{30}) \times (12^2) \times 4$$

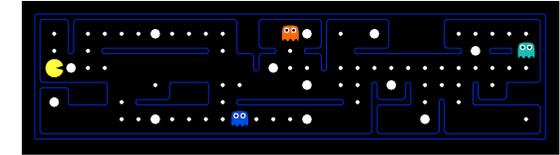
States for pathing?

120

States for eat-all-dots?

$$120 \times (2^{30})$$

Problem: Safe Passage

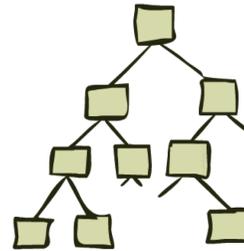


Problem: eat all dots while keeping the ghosts perma-scared?

What does the state space have to specify?

- agent position
- dot booleans
- power pellet booleans
- remaining scared time

State Space Graphs and Search Trees



State Space Graphs

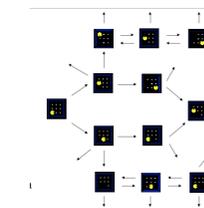
State space graph: A mathematical representation of a search problem

- ▶ Nodes are (abstracted) world configurations.
- ▶ Arcs represent successors (action results).
- ▶ The goal test is a set of goal nodes (maybe only one).

In a state space graph, each state occurs only once!

We can rarely build this full graph in memory (it's too big), but it's a useful idea.

E.g., replanning.



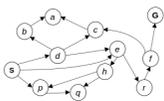
State Space Graphs

State space graph: A mathematical representation of a search problem

- ▶ Nodes are (abstracted) world configurations
- ▶ Arcs represent successors (action results)
- ▶ The goal test is a set of goal nodes (maybe only one)

In a state space graph, each state occurs only once!

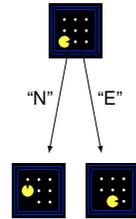
We can rarely build this full graph in memory (it's too big), but it's a useful idea.



Tiny search graph for a tiny search problem

Search Trees

This is now / start



Possible futures.

A search tree:

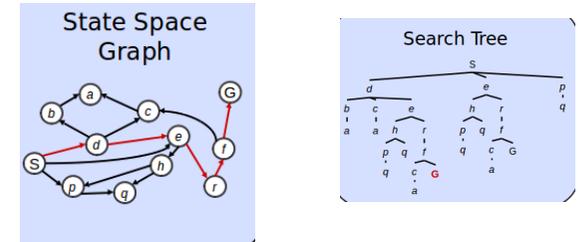
A "what if" tree of plans and their outcomes

The start state is the root node
Children correspond to successors

Nodes show states, but correspond to PLANS that achieve those states

For most problems, we can never actually build the whole tree

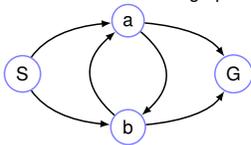
State Space Graphs vs. Search Trees



Each NODE in search tree is an entire PATH in state space graph.
We construct both on demand – and we construct as little as possible.

Quiz: State Space Graphs vs. Search Trees

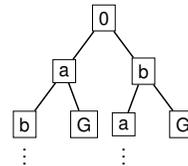
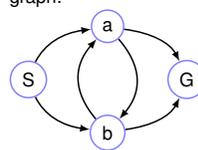
Consider this 4-state graph:



How large could search tree be?
 ∞

Quiz: State Space Graphs vs. Search Trees

Consider this 4-state graph:



How large could search tree be?
 ∞

Tree Search

