COMP 4752
Computational Intelligence

Lecture 2

Agents, Actions, and Environments

Agents

- Perceives its environment through Sensors
  - Robot: Camera, Laser, Sonar
  - Human: 5 Senses
  - Games: Keyboard Input / RAM / Game State
- Acts on its environment through Actuators
  - Robots: Wheels, Tracks, Arms
  - Humans: Limbs, Tools
  - Games: Predefined Rules / Actions
- May have some knowledge about environment

Agents

Agents

- Percept
  - Agent's perceptual inputs at any instant
  - "Agent's current belief of the world"
  - Environment's current State
- Percept Sequence
  - Complete history of everything the agent has ever perceived so far

Agents

- *In general*, an agent's choice of action at any given instant can depend on the entire percept sequence to date
- This course: fully observable states, no history
- An agent's behaviour is described by the agent function that maps given percept sequence to an action

Examples

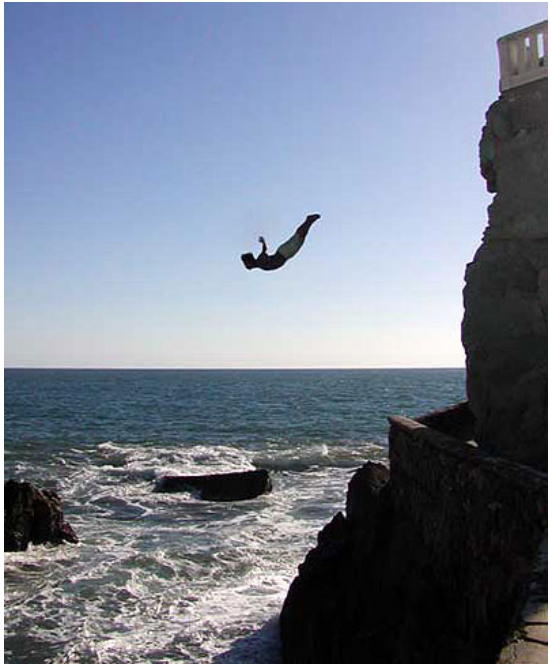

Examples


TSM
TL
0:11
26:04
TL PigJet
13
3690
TSM Santorin
11
TSM Dyrus
13
TSM WildTurtle
13
TSM Bjergsen
14
TECH
INSIDER


REPLAY
TL Piglet
13
TSM WildTurtle
13
TSM Bjergsen
14
4
12
9

Action

- An action taken by an agent in a given state transitions the state to another

State Transition Table

State \ Input	1	0
S_1	S_1	S_2
S_2	S_2	S_1

State Diagram

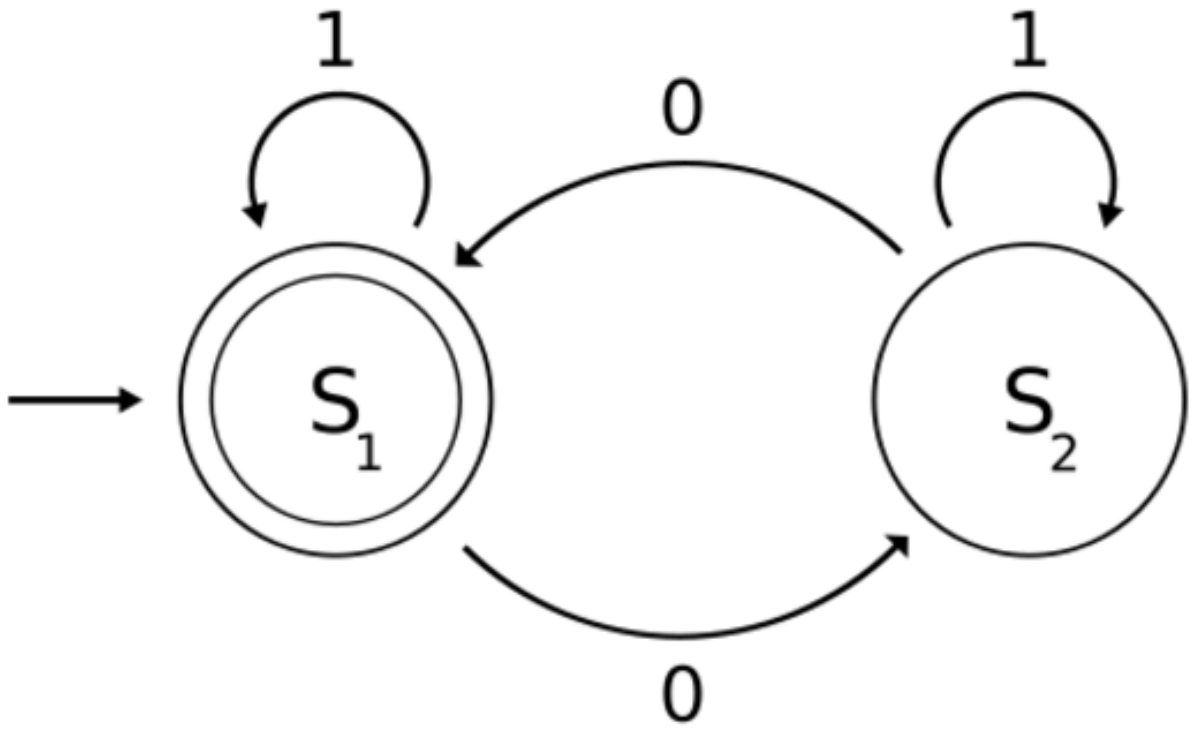

Action

- An action taken by an agent in a given state transitions the state to another
- State Transition Function (STF)
  - Successor Function, Table, Graph
- State S, Action A
- S’ = STF(S, A)

Action Example

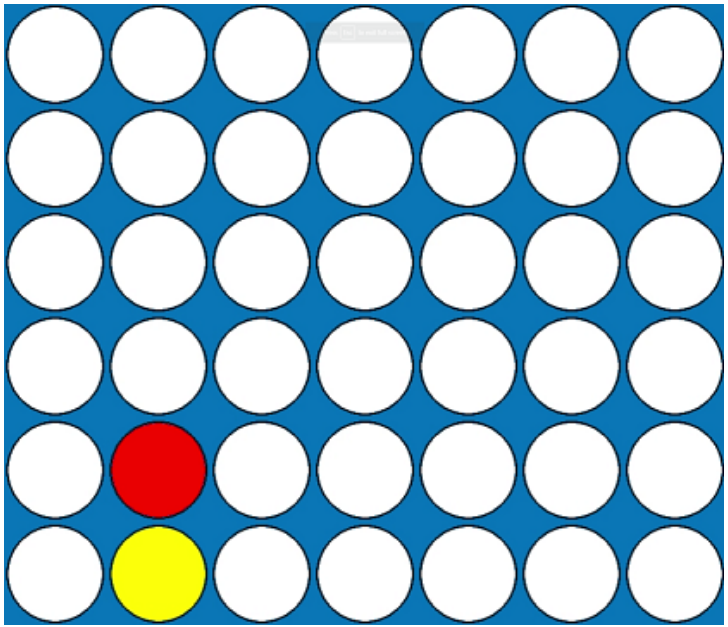

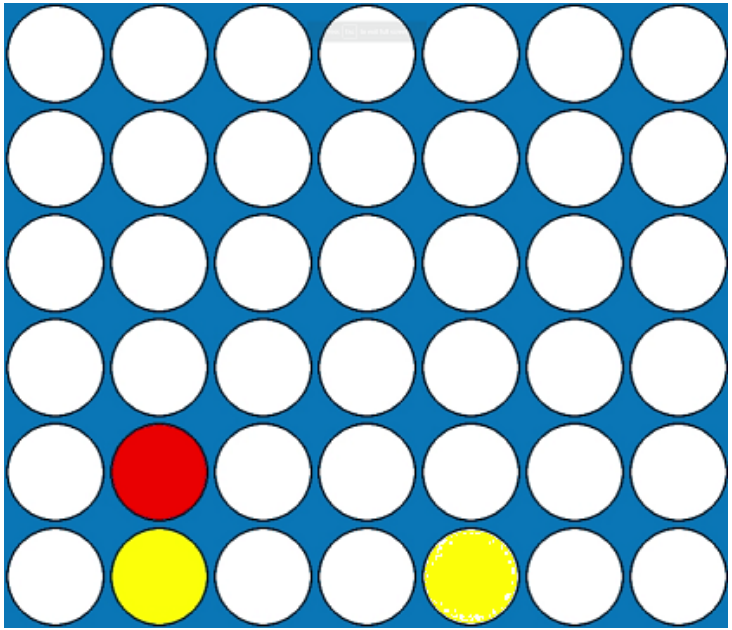

Policies

- Policy = Map from States to Actions
- Policy π, Optimal Policy π*

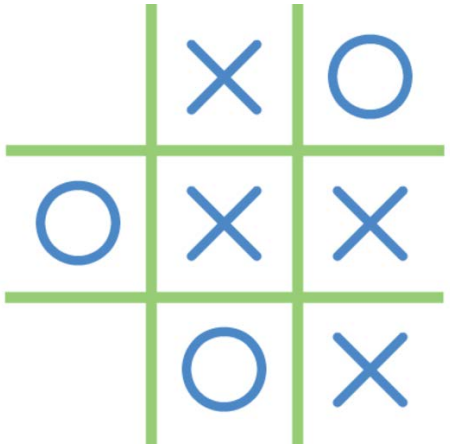

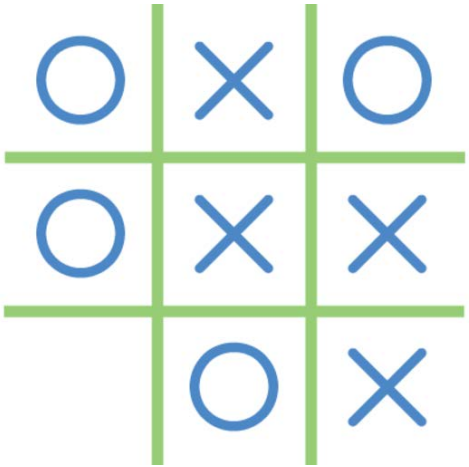

Rationality

- A rational agent is one that does the ‘right thing’ for every percent sequence
- What is the ‘right thing’
  - Makes the agent ‘most successful’
- Must define a measure of success
  - Performance Measure
- Rational agent has an Optimal Policy

Performance Measure

- Criterion for success of an agent
- Also called “Evaluation Function”
- Typical scenario
  1. Agent placed in an environment
  2. Agent performs sequence of actions
  3. Actions cause environment to go through a sequence of states
  4. If the sequence of states is desirable, the agent has performed well

Performance Measure

Performance Measure

Rationality

- Rationality depends on 4 things:
  - The performance measure that defines success
  - The agent’s prior knowledge of the environment
  - The actions the agent can perform
  - The agent’s percept sequence

Omniscience

- Rationality != Omniscience
- An omniscient agent:
  - Knows the exact outcome of actions
  - Knows the outcome of ‘randomness’
  - Knows the entire state of the environment
- Example:
  - Person bets $100 that a die will roll between 1-5
  - Die roll is a 6, person loses bet
  - Was the bet irrational?

Environments

Environments

- Task Environment = Problem
- Rational Agent = Solution
- The specification and properties of an environment greatly impact the design of a rational agent to act within it
- A State is a particular configuration of a given environment

Example Environment

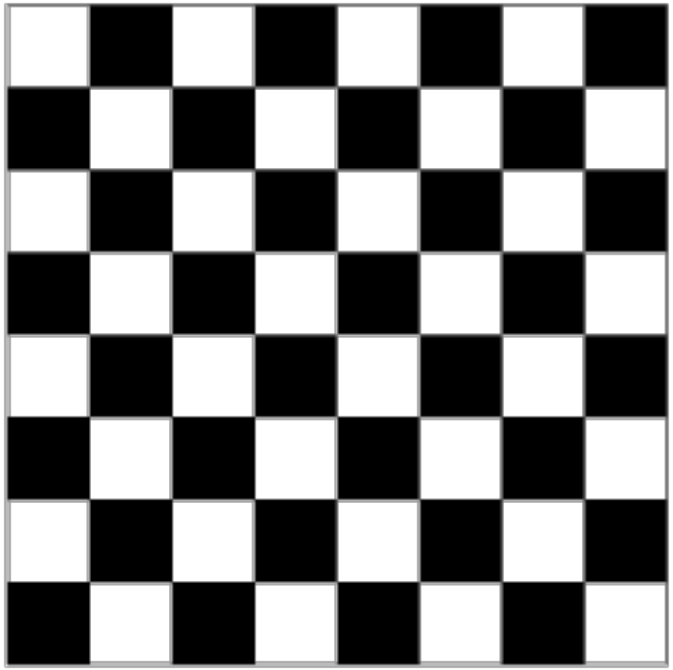

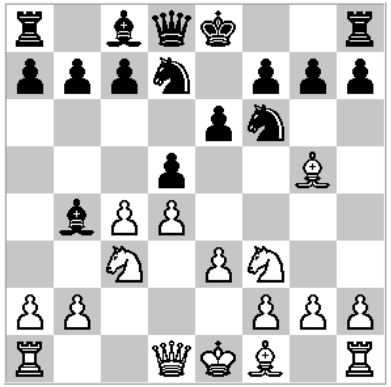

State / Action Space

- State Space
  - Number of possible configurations of the environment for a given problem
- Action Space
  - Number of actions possible from a state
  - Given as average or worst-case
- Used as an estimation of complexity

Environment Definition

- The rules / dynamics that govern it
  - Physics, Game Rules, etc
- The legal actions allowed at any State
- Initial / Goal States
  - Usually given by the problem instance

- Environments have several Properties

Fully vs. Partially Observable

- Fully Observable
  - Agent's sensors give it access to the complete state of the environment at any given time
  - Access to 'relevant' part of state for actions
  - Games: 'Perfect Information Game'
- Partially Observable:
  - Data is hidden or occluded
  - Sensors may be noisy
  - Games: 'Imperfect Information Game'

Fully vs. Partially Observable

Deterministic vs. Stochastic

- Deterministic
  - If the next state of the environment is completely determined by the current state and the action of the agent
  - Fully observable = No Uncertainty
- Stochastic
  - Uncertainty about the outcome of actions
  - Random nature to the environment
- Think from the point of view of the agent

Deterministic vs. Stochastic

Episodic vs. Sequential

- Episodic
  - Agent’s experience and actions are divided into atomic episodes
  - Agent perceives current state, takes an action
  - Next episode does not depend on previous
- Sequential
  - Current action can influence future decisions
  - Actions may have long-term consequences

Episodic vs. Sequential

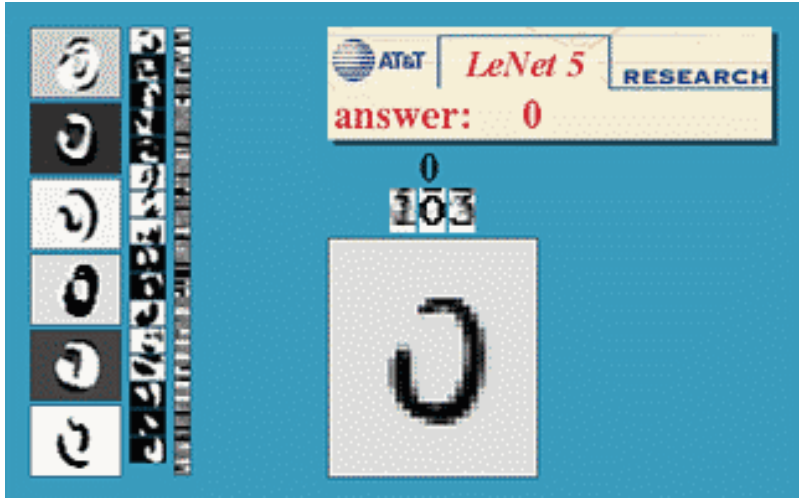

Dynamic vs. Static

- Dynamic
  - Environment may change while agent decides
  - Continuously asking agent what to do
- Static
  - Environment only changes on agent action
  - Decision making time doesn’t matter
- Games: Real-Time vs. Turn-Based

Dynamic vs. Static

Discrete vs. Continuous

- Distinction applied to the state of env.
- Determines the way time is handled
- Discrete
  - Finite number of distinct states
  - Discrete set of percepts and actions
- Continuous
  - Continuous time / actions

Discrete vs. Continuous

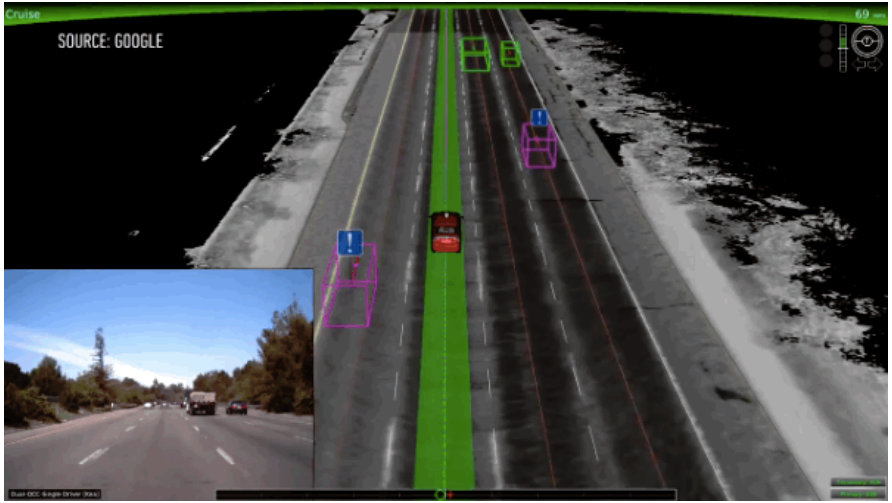

Single Agent vs. Multi-agent

- Single Agent
  - One agent or acting at a time
  - Ex: Puzzle solving, Path finding
- Multi-agent
  - Multiple agents acting together
  - May be cooperative or competitive

Single Agent vs. Multi-agent

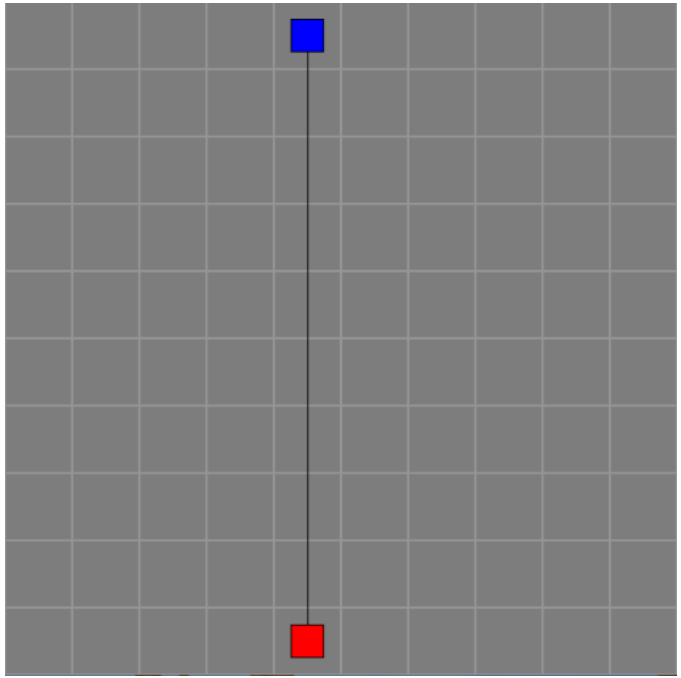

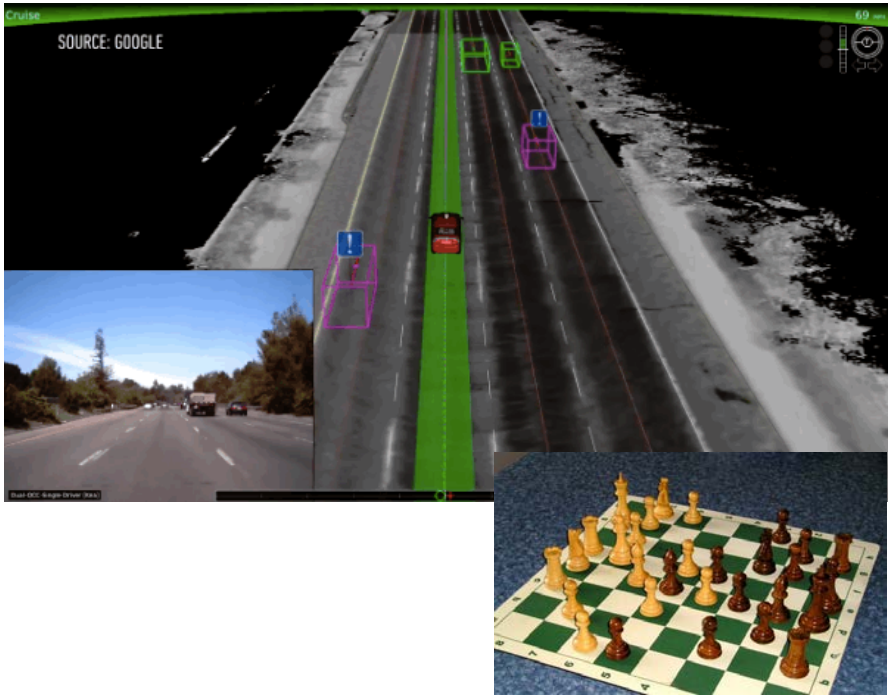

Complete vs. Incomplete Information

- Complete Information
  - All game rules / physics are known
  - All possible actions are known

- Incomplete Information
  - Game rules / physics may not be known
  - Actions may have to be discovered

Example Game Environments

Environ.	Episodic	Agents	Determ.	Observe	Static	Discrete
CROSSWORD	SEQUENTIAL	SINGLE	DETERM	FULLY	STATIC	DISCRETE
CHESS	SEQUENTIAL	MULTI	DETERM	FULLY	STATIC	DISCRETE
BACK GAMMON	SEQUENTIAL	MULTI	STOCHASTIC	FULLY	STATIC	DISCRETE
POKER	SEQUENTIAL	MULTI	STOCHASTIC	PARTIALLY	STATIC	DISCRETE
STARCRAFT	SEQUENTIAL	MULTI	STOCHASTIC	PARTIALLY	DYNAMIC	DISCRETE
REAL WORLD	SEQUENTIAL	MULTI	STOCHASTIC	PARTIALLY	DYNAMIC	CONTINUOUS

Recap

- Agent is in an Environment
- Environment has several Properties
- Agent Observes part/all of Environment
- Agents take Actions, which transition the State
- Performance Measure evaluates agent
- Rational Agent maximizes performance
- AI goal is to construct Rational Agents