CS 295A/395D: Artificial Intelligence

Potpourri of Unit 3

Prof. Emma Tosch

1 April 2022

The University of Vermont

First: Design pattern in AI/ML

Many tasks can be boiled down to alternating between:

1. Computing an expected value
2. Finding an argument (i.e., making a choice) that maximizes some function

For problems of any complexity, that function will be composed of other functions, including random variables.

Why expected values?

High level view:

Expected values are *summary information* and are useful when:

1. We don't know the point value (its value has either aleatory or epistemic uncertainty)
2. BUT…we know its distribution

Can make decisions on the basis of that summary information!

Why expected values?

High level view:

Expected values are *summary information* and are useful when:

1. We don't know the point value (its value has either aleatory or epistemic uncertainty)
2. BUT…we know its distribution

Can make decisions on the basis of that summary information!

Why expected values?

High level view:

Expected values are *summary information* and are useful when:

1. We don’t know the point value (its value has either aleatory or epistemic uncertainty)
2. BUT…we know its distribution

Can make decisions on the basis of that summary information!

Game theory: recipe for mixed strategies

1. Compute expression for E[P | Q = q] using P's utility function.
   1. Reminder: this will be over all possible choices for P.
   2. The pmf comes from P's choice.
   3. Q's choice is fixed here.
2. Compute expression for other values of Q=q.
3. Set expressions equal to each other and solve for p.

Why?

Utility functions

Recall: all agents are "rational"

- All agents have complete knowledge of the payoff matrix
- All agents seek to maximize their utility
- Assumption: all players are treating the game in a decision-theoretic manner

Can model game theory as decision theory

Given: simultaneous, no communication or coordination

- Treat other agent's actions random state
- Each state node encapsulates all of the uncertainty about player Q's actions
- Objective: use game theory to determine state distribution

Relation to minimax

Previously: search in a planning context

- Planning with logic (e.g. STRIPS, PADDL)
  - Search through change state as logical inference over limited language
  - Enforcing constraints
  - Challenge: finding a path efficiently
- Introduced notions of heuristics + cost. Difference?
  - Heuristics are estimates (used when it's okay to be slightly sub-optimal)
  - Costs assign value to state, used for ordering
- All deterministic; here, probabilistic

Minimax theorem in game theory

Subtlety in the player's objective:

- Minimize max loss?
- Maximize min gain?

Assume zero-sum game:

Player X maximizing its minimum gain is equivalent to minimizing its max loss.

$$\max_{x \in X} \min_{y \in Y} x^T A y = \min_{y \in Y} \max_{x \in X} x^T A y.$$

Recipe for easy solutions

Recall: Decisions are made locally

- Optimum vs. optimal
    - Optimum is global (something we cannot control)
    - Optimal is local (something we can control)
- Special case: saddle point for zero-sum games
    - Minimum between choices for P (here, between columns)
    - Maximum between choices for Q (here, between rows)

Utility functions: other models?

Recall: all agents are "rational"

- All agents have complete knowledge of the payoff matrix
- All agents seek to maximize their utility
- Assumption: all players are treating the game in a decision-theoretic manner

Formalizing with epistemic knowledge