Artificial Intelligence Programming

Decision Making and Utility

Chris Brooks

Department of Computer Science
University of San Francisco

Making decisions

- At this point, we know how to describe the probability of events occuring.
  - Or states being reached, in agent terms.
- Knowing the probabilities of events is only part of the battle.
- Agents are really interested in maximizing performance.
- This means making the correct decisions as to what to do.
- Often, performance can be captured by *utility*.
- Utility indicates the relative value of a state.

Types of decision-making problem

- Single-agent, deterministic, full information, episodic
  - We've done this with the reflex agent
- Single-agent, deterministic, full information, sequential
  - We can use search here.
- Single-agent, stochastic, partial information, episodic
- Single-agent, stochastic, partial information, sequential
- multiple-agent, deterministic, full information, episodic

Types of decision-making problems

- Single-agent, deterministic, full information, episodic
  - We've done this with the reflex agent
- Single-agent, deterministic, full information, sequential
  - We can use search here.
- Single-agent, stochastic, partial information, episodic
  - We can use utility and probability here
- Single-agent, stochastic, partial information, sequential
  - We can extend our knowledge of probability and utility to a Markov decision process.
- multiple-agent, deterministic, full information, episodic (or sequential)
  - This is game theory

Expected Utility

- In episodic, stochastic worlds, we can use expected utility to select actions.
- An agent will know that an action can lead to one of a set S of states.
- The agent has a utility for each of these states.
- The agent also has a probability that these states will be reached.
- The *expected utility* of an action is:
- $\sum_{s \in S} P(s)U(s)$
- The principle of maximum expected utility says that an agent should choose the action that maximizes expected utility.

Example

- Let's say there are two levers.
  - Lever 1 costs \$1 to pull. With $p = 0.4$, you win \$2. With $p = 0.6$ you win nothing.
  - Lever 2 costs \$2 to pull. With $p = 0.1$ you win \$10. with $p = 0.9$ you lose \$1 (on top of the charge to pull).
- Should you a) pull lever 1 b) pull lever 2 c) pull neither?

Example

- EU(lever 1) = 0.4 * 1 + 0.6 * -1 = -0.2
- EU(lever 2) = 0.1 * 8 + 0.9 * -3 = 5.3
- EU(neither) = 0
- Lever 2 gives the maximum EU.
- TV digression - this is the choice contestants are faced with on "Deal or No Deal." The banker offers them a price slightly above the expected utility, and yet most contestants don't take it. Why?

Example: Vegas

- The typical roulette wheel has 38 numbers. (1-36, plus 0 and 00).
- 1-36 are either red or black.
- The payoff for betting on a single number is 35:1
  - In other words, if the number you picked comes up, you win $35. Otherwise, you lose $1.
- What is the expected utility of betting on a single number?

Example: Vegas

- The typical roulette wheel has 38 numbers. (1-36, plus 0 and 00).
- 1-36 are either red or black.
- The payoff for betting on a single number is 35:1
  - In other words, if the number you picked comes up, you win $35. Otherwise, you lose $1.
- What is the expected utility of betting on a single number?
- $\frac{1}{38} * 35 + \frac{37}{38} * -1 = -0.052$

Example: Vegas

- The typical roulette wheel has 38 numbers. (1-36, plus 0 and 00).
- 1-36 are either red or black.
- The payoff for betting on a single number is 35:1
  - In other words, if the number you picked comes up, you win $35. Otherwise, you lose $1.
- What if you decide to "spread the risk" and bet on two numbers?

Example: Vegas

- The typical roulette wheel has 38 numbers. (1-36, plus 0 and 00).
- 1-36 are either red or black.
- The payoff for betting on a single number is 35:1
  - In other words, if the number you picked comes up, you win $35. Otherwise, you lose $1.
- What if you decide to "spread the risk" and bet on two numbers?
- $\frac{2}{38} * 34 + \frac{36}{38} * -2 = -0.105$

Example: Vegas

- The typical roulette wheel has 38 numbers. (1-36, plus 0 and 00).
- 1-36 are either red or black.
- The payoff for betting on color is 1:1
  - In other words, if you bet 'red' and a red number comes up, you win $1. Otherwise, you lose $1.
- What is the expected utility of betting on 'red'?

Example: Vegas

- The typical roulette wheel has 38 numbers. (1-36, plus 0 and 00).
- 1-36 are either red or black.
- The payoff for betting on color is 1:1
  - In other words, if you bet 'red' and a red number comes up, you win $1. Otherwise, you lose $1.
- What is the expected utility of betting on 'red'?
- $\frac{18}{38} * 1 + \frac{20}{38} * -1 = -0.052$

Regarding Preferences

- In order for MEU to make sense, we need to specify some (hopefully reasonable) constraints on agent preferences.
- Orderability. We must be able to say that A is preferred to B, B is preferred to A, or they are equally preferable. We cannot have the case where A and B are incomparable.
- Transitivity. If $A \prec B$ and $B \prec C$, then $A \prec C$.
- Continuity. If $A \prec B \prec C$, then there is a scenario where the agent is indifferent to getting B and having a probability p of getting A and $1 - p$ chance of getting C.

Rational Preferences

- Monotonicity. If two actions A and B have the same outcomes, and I prefer A to B, I should still prefer A if the probability of A increases.
- Decomposability. Utilities over a sequence of actions can be decomposed into utilities for atomic events.
- These preferences are (for the most part) quite reasonable, and allow an agent to avoid making foolish mistakes.

Utility, Money, and Risk

- Utility comes from economics
  - Money is often used as a substitute for utility.
- Preferences for money behave oddly when dealing with small or large amounts.
- For example, you will often take more chance with small amounts, and be very conservative with very large amounts.
- This is called your ***risk profile***
  - convex - risk-seeking
  - concave, risk-averse
- Typically, we say that we have a ***quasilinear*** utility function regarding money.

Gathering information

- When an agent has all the information it can get and just needs to select a single action, things are straightforward.
  - Find the action with the largest expected utility.
- What if an agent can choose to gather more information about the world?
- Now we have a sequential decision problem:
  - Should we just act, or gather information first?
  - What questions should we ask?
    - Agents should ask questions that give them useful information.
    - "Useful" means increasing expected utility.
  - Gathering information might be costly, either in time or money.

Example

- An agent can recommend to a prospecting company that they buy one of n plots of land.
- One of the plots has oil worth C dollars; the others are empty.
- Each block costs C/n dollars.
- Initially, agent is indifferent between buying and not buying. (why is that?)