Lecture 1: Utility Theory and Decision Theory

Pascal Poupart

Decision Making under Uncertainty

- I give robot a planning problem: I want coffee
  - but coffee maker is broken: robot reports "No plan!"
- If I want more robust behavior – if I want robot to know what to do if my primary goal can't be satisfied – I should provide it with some indication of my *preferences over alternatives*
  - e.g., coffee better than tea, tea better than water, water better than nothing, etc.

Decision Making under Uncertainty

- But it's more complex:
  - it could wait 45 minutes for coffee maker to be fixed
  - what's better: tea now? coffee in 45 minutes?
  - could express preferences for <beverage,time> pairs

Preferences

- A *preference ordering* $\succcurlyeq$ is a ranking of all possible states of affairs (worlds) S
  - these could be outcomes of actions, truth assts, states in a search problem, etc.
  - s $\succcurlyeq$ t: means that state s is *at least as good as* t
  - s $\succ$ t: means that state s is *strictly preferred to* t
  - s~t: means that the agent is *indifferent* between states s and t

Preferences

- If an agent's actions are deterministic then we know what states will occur
- If an agent's actions are not deterministic then we represent this by lotteries
  - Probability distribution over outcomes
  - Lottery $L=[p_1,s_1;p_2,s_2;\ldots;p_n,s_n]$
  - s_1 occurs with prob p_1, s2 occurs with prob p_2,…

Preference Axioms

- **Orderability:** Given 2 states A and B
  - $(A \succ B) \sqsubseteq (B \succ A) \sqsubseteq (A \sim B)$
- **Transitivity:** Given 3 states, A, B, and C
  - $(A \succ B) \wedge (B \succ C) \Rightarrow (A \succ C)$
- **Continuity:**
  - $A \succ B \succ C \Rightarrow \exists p\ [p,A;1-p,C] \sim B$
- **Substitutability:**
  - $A \sim B \rightarrow [p,A;1-p,C] \sim [p,B;1-p,C]$
- **Monotonicity:**
  - $A \succcurlyeq B \Rightarrow (p \geq q \Leftrightarrow [p,A;1-p,B] \succcurlyeq [q,A;1-q,B]$
- **Decomposibility:**
  - $[p,A;1-p,[q,B;1-q,C]] \sim [p,A;(1-p)q,B;\ (1-p)(1-q),C]$

Why Impose These Conditions?

- Structure of preference ordering imposes certain "rationality requirements" (it is a weak ordering)
- E.g., why transitivity?
  - Suppose you (strictly) prefer coffee to tea, tea to OJ, OJ to coffee
  - If you prefer X to Y, you'll trade me Y plus $1 for X
  - I can construct a "money pump" and extract arbitrary amounts of money from you

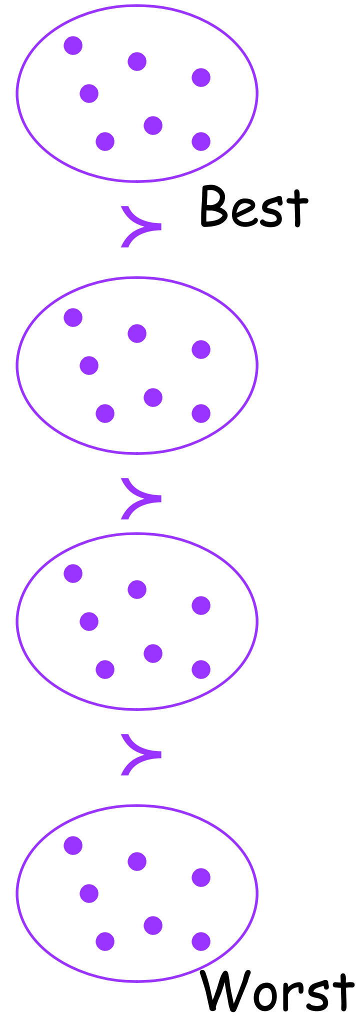

Decision Making under Uncertainty

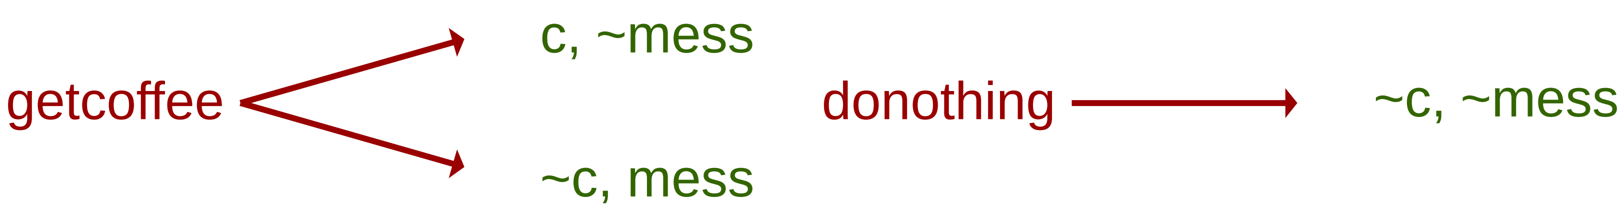

- Suppose actions don't have deterministic outcomes
  - e.g., when robot pours coffee, it spills 20% of time, making a mess
  - preferences: c, ~mess ≻ ~c,~mess ≻ ~c, mess
- What should robot do?
  - decision *getcoffee* leads to a good outcome and a bad outcome with some probability
  - decision *donothing* leads to a medium outcome for sure
- Should robot be optimistic? pessimistic?
- Really odds of success should influence decision
  - but how?

Utilities

- Rather than just ranking outcomes, we must quantify our degree of preference
  - e.g., how much more important is c than ~mess
- A *utility function* $U:S \rightarrow \mathbb{R}$ associates a real-valued *utility* with each outcome.
  - U(s) measures your *degree* of preference for s
- Note: U induces a preference ordering $\succcurlyeq_U$ over S defined as: $s \succcurlyeq_U t$ iff $U(s) \geq U(t)$
  - obviously $\succcurlyeq_U$ will be reflexive, transitive, connected

Expected Utility

- Under conditions of uncertainty, each decision d induces a distribution $\Pr_d$ over possible outcomes
  - $\Pr_d(s)$ is probability of outcome s under decision d

$$EU(d) = \sum_{s \in S} \Pr_d(s) U(s)$$

- The *expected utility* of decision d is defined

Expected Utility

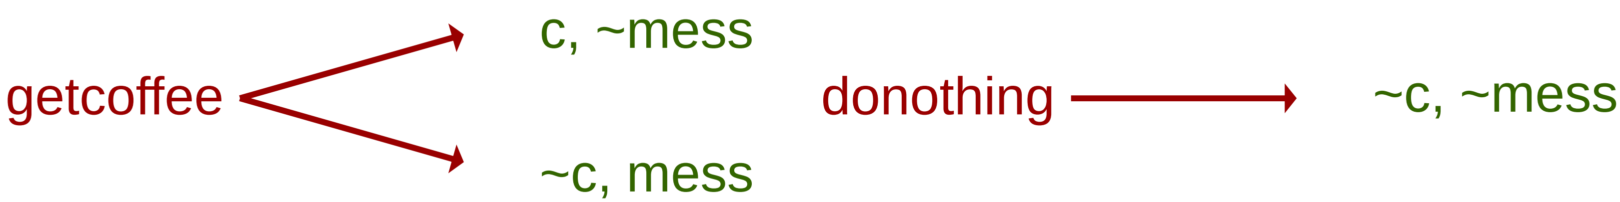

When robot pours coffee, it spills 20% of time, making a mess

If U(c,~ms) = 10, U(~c,~ms) = 5, U(~c,ms) = 0,

then EU(getcoffee) = (0.8)(10)+(0.2)(0)=8

and EU(donothing) = 5

If U(c,~ms) = 10, U(~c,~ms) = 9, U(~c,ms) = 0,

then EU(getcoffee) = (0.8)(10)+(0.2)(0)=8

and EU(donothing) = 9

The MEU Principle

- The *principle of maximum expected utility (MEU)* states that the optimal decision under conditions of uncertainty is that with the greatest expected utility.
- In our example
  - if my utility function is the first one, my robot should get coffee
  - if your utility function is the second one, your robot should do nothing

Decision Problems: Uncertainty

- A *decision problem under uncertainty* is:
  - a set of *decisions* D
  - a set of *outcomes* or states S
  - an *outcome function* $Pr : D \rightarrow \Delta(S)$
    - $\Delta(S)$ is the set of distributions over S (e.g., Pr_d)
  - a utility function U over S
- A *solution* to a decision problem under uncertainty is any $d^* \in D$ such that $EU(d^*) \succcurlyeq EU(d)$ for all $d \in D$

Expected Utility: Notes

- Utility functions needn't be unique
  - if I multiply U by a positive constant, all decisions have same relative utility
  - if I add a constant to U, same thing
  - *U is unique up to positive affine transformation*

- Where do utilities come from?
- How can we make decisions with imprecise/incomplete utilities

Decision Making with Imprecise/Incomplete Utilities

- Two common principles
  - Maximum *Expected* Expected Utilities
    - Distribution over utilities
    - Make decision that maximizes expectation w.r.t. outcome distribution and utility distribution
  - Minimax Regret
    - Set of possible utility functions
    - Make decision that minimizes the worst-case loss in expected utilities w.r.t. set of utilities