Decision Making under Uncertainty

AI Class 10 (Ch. 15.1-15.2.1, 16.1-16.3)

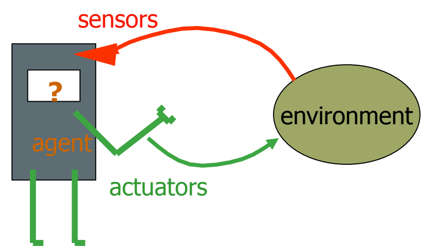

Material from Marie desJardin, Lise Getoor, Jean-Claude Latombe, Daphne Koller, and Paula Matuszek

Bookkeeping

- HW 3 out
  - Group work for non-programming parts!
  - Heavy on CSPs and probability
  - Forms groups today or in Piazza
- Soon: form project teams!

Today's Class

- Making Decisions Under Uncertainty
  - Tracking Uncertainty over Time
  - Decision Making under Uncertainty
- Project groups, part 1 ← ?

Introduction

- The world is not a well-defined place.
- Sources of uncertainty
  - Uncertain **inputs**: What's the temperature?
  - Uncertain (imprecise) **definitions**: Is Obama a good president?
  - Uncertain (unobserved) **states**: Where is the pit?
- There is uncertainty in **inferences**
  - If I have a blistery, itchy rash and was gardening all weekend I **probably** have poison ivy

Sources of Uncertainty

- Uncertain **inputs**
  - Missing data
  - Noisy data
- Uncertain **knowledge**
  - >1 cause → >1 effect
  - Incomplete knowledge of causality
  - Probabilistic effects

- Uncertain **outputs**
  - Default reasoning (even deduction) is uncertain
  - Abduction & induction inherently uncertain
  - Incomplete deductive inference can be uncertain
- Derived result is formally correct, but wrong in real world

Probabilistic reasoning only gives **probabilistic results** (summarizes uncertainty from various sources)

5

Reasoning Under Uncertainty

- People make successful decisions all the time anyhow.
  - How?
  - More formally: how do we do reasoning under uncertainty, with inexact knowledge?
- Step one: **understanding what we know**

6

MODELING UNCERTAINTY OVER TIME

7

States and Observations

- We don't have a continuous view of world
  - People don't either!
- We see things as a series of snapshots
- **Observations**, associated with **time slices**
  - t_1, t_2, t_3, …
- Each snapshot contains all variables, observed or not
  - $\mathbf{X}_t$ = (unobserved) state variables at time t; observation at t is $\mathbf{E}_t$
- This is **world state at time t**

8

Temporal Probabilistic Agent

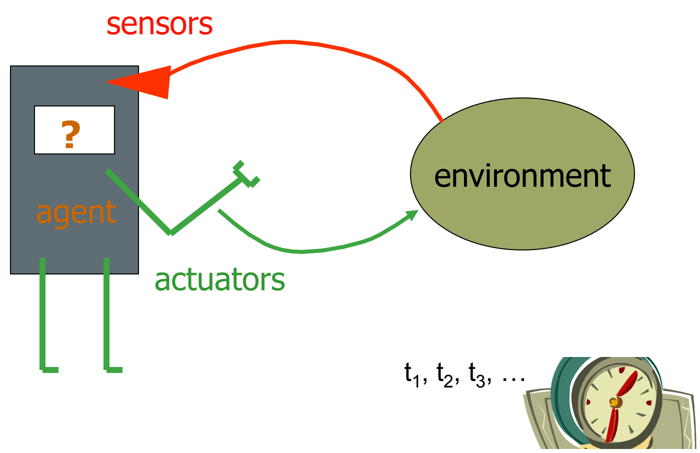

Time and Uncertainty

- The world changes
  - Examples: diabetes management, traffic monitoring
- Tasks: **track** it; **predict** it
- Basic idea:
  - Copy state and evidence variables for each time step
  - Model uncertainty in change over time
  - Incorporate new observations as they arrive

Time and Uncertainty

- Basic idea:
  - Copy state and evidence variables for each time step
  - Model uncertainty in change over time
  - Incorporate new observations as they arrive
- $\mathbf{X}_t$ = unobservable state variables at time t: $BloodSugar_t$, $StomachContents_t$
- $\mathbf{E}_t$ = evidence variables at time t: $MeasuredBloodSugar_t$, $PulseRate_t$, $FoodEaten_t$
- Assuming discrete time steps

11

States, Slightly More formally

- Process of change is viewed as series of snapshots
  - Time slices
  - Each describing the state of the world at a particular time
- Each time slice is represented by a set of random variables indexed by t:
  1. the set of unobservable **state variables** $\mathbf{X}_t$
  2. the set of observable **evidence variables** $\mathbf{E}_t$
- The observation at time t is $\mathbf{E}_t = e_t$ for some set of values e_t
- $\mathbf{X}_{a:b}$ denotes the set of variables from $\mathbf{X}_a$ to $\mathbf{X}_b$

12

Transition and Sensor Models

- Transition model
  - Models how the world changes over time
- Specifies a probability distribution
  - Over state variables at time t
  - Given values at previous times

$P(\mathbf{X}_t \mid \mathbf{X}_{0:t-1})$

How big can this get?

- Sensor model
  - Models how evidence gets its values (sensor data)
    - E.g.: $BloodSugar_t \rightarrow MeasuredBloodSugar_t$

Markov Assumption

- **Markov Assumption**:
  - X_t depends on some finite (usually fixed) number of previous X_i's
- **First-order Markov process**: $P(X_t | X_{0:t-1}) = P(X_t | X_{t-1})$
  - k^{th} order: depends on previous k time steps

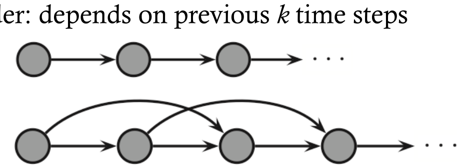

- **Sensor Markov assumption**: $P(E_t | X_{0:t}, E_{0:t-1}) = P(E_t | X_t)$
  - Agent's observations depend *only* on the actual current state of the world

Stationary Process

- Infinitely many possible values of *t*
  - Does each timestep need a distribution?
- Assume **stationary process**:
  - Changes in the world state are governed by laws that do not themselves change over time
  - Transition model $P(X_t | X_{t-1})$ and sensor model $P(E_t | X_t)$ are time-invariant, i.e., they are the same for all t

Complete Joint Distribution

- Given:
  - Transition model: $P(X_t | X_{t-1})$
  - Sensor model: $P(E_t | X_t)$
  - Prior probability: $P(X_0)$
- Then we can specify complete joint distribution of a sequence of states:

$$P(X_0, X_1, \ldots, X_t, E_1, \ldots, E_t) = P(X_0) \prod_{i=1}^{t} P(X_i \mid X_{i-1}) P(E_i \mid X_i)$$

Example

R_{t-1}	$P(R_t \mid R_{t-1})$
T	0.7
F	0.3

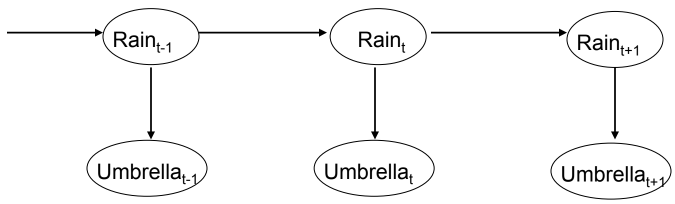

R_t	$P(U_t \mid R_t)$
T	0.9
F	0.2

This should look like a finite state automaton (since it is one)

17

Inference Tasks

- **Filtering** or monitoring: $P(X_t \mid e_1,\dots,e_t)$
  Compute the current belief state, given all evidence to date
- **Prediction**: $P(X_{t+k} \mid e_1,\dots,e_t)$
  Compute the probability of a future state
- **Smoothing**: $P(X_k \mid e_1,\dots,e_t)$
  Compute the probability of a past state (hindsight)
- **Most likely explanation**:
  $\arg\max_{x1,..xt} P(x_1,\dots,x_t \mid e_1,\dots,e_t)$
  Given a sequence of observations, find the sequence of states that is most likely to have generated those observations

18

Examples

- **Filtering:** What is the probability that it is raining today, given all of the umbrella observations up through today?
- **Prediction:** What is the probability that it will rain the day after tomorrow, given all of the umbrella observations up through today?
- **Smoothing:** What is the probability that it rained yesterday, given all of the umbrella observations through today?
- **Most likely explanation:** If the umbrella appeared the first three days but not on the fourth, what is the most likely weather sequence to produce these umbrella sightings?

Filtering

- Maintain a current state estimate and update it
  - Rather than looking at all percepts (observed values) in history
  - So, given result of filtering up to t, compute $t+1$ from $\mathbf{e}_{t+1}$
- We use **recursive estimation** to compute $P(X_{t+1} \mid e_{1:t+1})$ as a function of e_{t+1} and $P(X_t \mid e_{1:t})$
- We can write this as:

$$P(X_{t+1} \mid e_{1:t+1}) = P(X_{t+1} \mid e_{1:t}, e_{t+1})$$

Filtering 2

- $P(X_{t+1} \mid e_{1:t+1})$ as a function of e_{t+1} and $P(X_t \mid e_{1:t})$

$$
\begin{aligned}
P(X_{t+1} \mid e_{1:t+1}) &= P(X_{t+1} \mid e_{1:t}, e_{t+1}) \\
&= \alpha\, P(e_{t+1} \mid X_{t+1}, e_{1:t})\; P(X_{t+1} \mid e_{1:t}) \\
&= \alpha\, P(e_{t+1} \mid X_{t+1})\; P(X_{t+1} \mid e_{1:t}) \\
&= \alpha\, P(e_{t+1} \mid X_{t+1}) \sum_{x_t} P(X_{t+1} \mid x_t)\; P(x_t \mid e_{1:t})
\end{aligned}
$$

- This leads to a recursive definition:
  $f_{1:t+1} = \alpha\ \text{FORWARD}\ (f_{1:t}, e_{t+1})$

Filtering Example

$$P(X_{t+1} \mid e_{1:t+1}) = \alpha\, P(e_{t+1} \mid X_{t+1}) \sum_{X_t} P(X_{t+1} \mid X_t)\; P(X_t \mid e_{1:t})$$

R_{t-1}	$P(R_t \mid R_{t-1})$
T	0.7
F	0.3

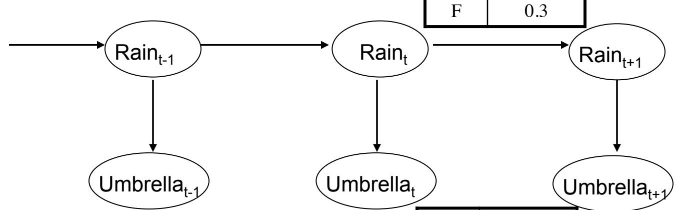

R_t	$P(U_t \mid R_t)$
T	0.9
F	0.2

What is the probability of rain on Day 2, given a uniform prior of rain on Day 0, U_1 = true, and U_2 = true?

Decision Making Under Uncertainty

Decision Making Under Uncertainty

- Many environments have multiple possible outcomes
- Some of these outcomes may be good; others may be bad
- Some may be very likely; others unlikely
- **What's a poor agent to do?**

Reasoning Under Uncertainty

- So how do we do reasoning under uncertainty and with inexact knowledge?
  - Heuristics
    - Mimic heuristic knowledge processing methods used by experts
  - Empirical associations
    - Experiential reasoning
    - Based on limited observations
  - Probabilities
    - Objective (frequency counting)
    - Subjective (human experience)

Non-deterministic vs. Probabilistic Uncertainty

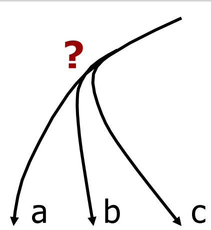

{a,b,c}

→ decision that is best for worst case

Non-deterministic model

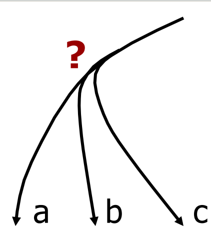

{$a(p_a)$, $b(p_b)$, $c(p_c)$}

→ decision that maximizes expected utility value

Probabilistic model

~ Adversarial search

Decision Theory

- Combine **probability** and **utility**

→ Agent that makes **rational** decisions
  - On average, lead to desired outcome

- Immediate simplifications:
  - Want most desirable immediate outcome (episodic)
  - nondeterministic, partially observable world
- Definition: result of an action *a* leads to outcome *s*':
  - RESULT(*a*) is a random variable; domain is possible outcomes
  - $P(\text{RESULT}(a) = s' \mid a, \boldsymbol{e})$

27

Expected Utility

- Goal: find best expected outcome
- Random variable X with:
  - n values $x_1,\ldots,x_n$
  - Distribution $(p_1,\ldots,p_n)$
- X is the state reached after doing an action A under uncertainty
- Utility function U(s) is the utility of a state, i.e., **desirability**

28

Expected Utility

- X is state reached after doing an action A under uncertainty
- U(s) is the utility of a state ← **desirability**
- The expected utility of action A, given evidence EU(a | **e**), is average utility of outcomes (states in S), weighted by probability an action occurs:

$$EU[A] = S_{i=1,\ldots,n}\, p(x_i|A)U(x_i)$$

One State/One Action Example

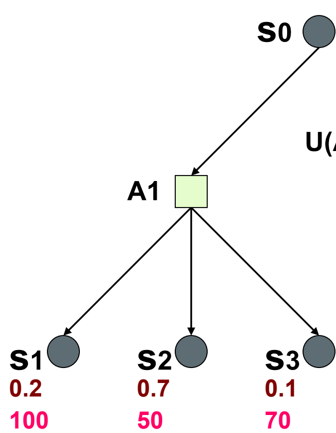

U(A1, S0) = 100 x 0.2 + 50 x 0.7 + 70 x 0.1
= 20 + 35 + 7
= 62

One State/Two Actions Example

- U (A1, S0) = 62
- U (A2, S0) = 74
- U (S0) = max_a\{U(a,S0)\}
  = 74

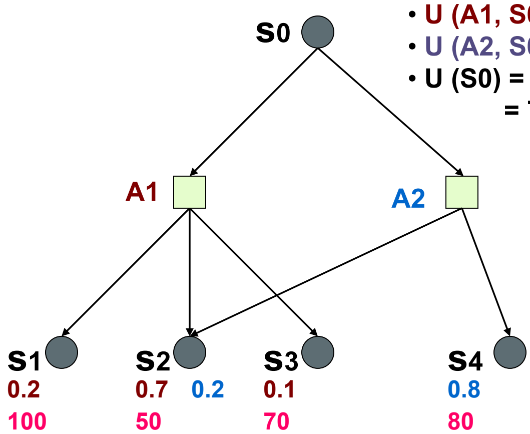

Introducing Action Costs

- U (A1, S0) = 62 – 5 = 57
- U (A2, S0) = 74 – 25 = 49
- U (S0) = max_a\{U(a, S0)\}
  = 57

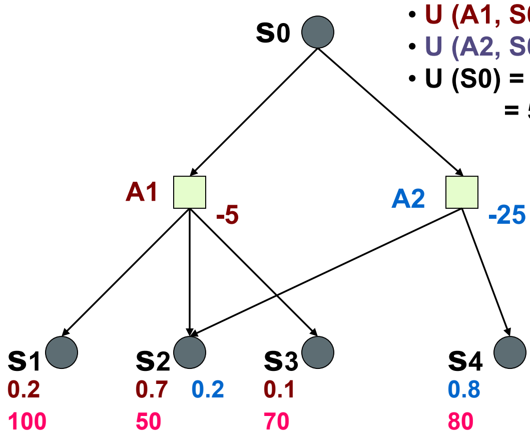

MEU Principle

- A **rational agent** should choose the action that maximizes agent's expected utility
- This is the basis of the field of **decision theory**
- The MEU principle provides a **normative criterion** for rational choice of action
- …AI is solved!

Not Quite…

- Must have a **complete** model of:
  - Actions
  - Utilities
  - States
- Even if you have a complete model, decision making is computationally **intractable**
- In fact, a truly rational agent takes into account the utility of reasoning as well (**bounded rationality**)
- Nevertheless, great progress has been made in this area
  - We are able to solve much more complex decision-theoretic problems than ever before

Axioms of Utility Theory

- Orderability
  $(A>B) \vee (A<B) \vee (A\sim B)$
- Transitivity
  - $(A>B) \wedge (B>C) \Rightarrow (A>C)$
- Continuity
  - $A>B>C \Rightarrow \exists p\ [p,A;\ 1-p,C] \sim B$
- Substitutability
  - $A\sim B \Rightarrow [p,A;\ 1-p,C]\sim[p,B;\ 1-p,C]$
- Monotonicity
  - $A>B \Rightarrow (p\geq q \Leftrightarrow [p,A;\ 1-p,B] \succsim [q,A;\ 1-q,B])$
- Decomposability
  - $[p,A;\ 1-p,\ [q,B;\ 1-q,\ C]] \sim [p,A;\ (1-p)q,\ B;\ (1-p)(1-q),\ C]$

Money Versus Utility

- Money $<>$ Utility
  - More money is better, but not always in a linear relationship to the amount of money
- Expected Monetary Value
- Risk-averse: $U(L) < U(S_{EMV(L)})$
- Risk-seeking: $U(L) > U(S_{EMV(L)})$
- Risk-neutral: $U(L) = U(S_{EMV(L)})$

Value Function

- Provides a ranking of alternatives, but not a meaningful metric scale
- Also known as an "ordinal utility function"
- Sometimes, only relative judgments (value functions) are necessary
- At other times, absolute judgments (utility functions) are required