# Decision Making Under Uncertainty

(Ch. 15.1-15.2.1, 16.1-16.3)

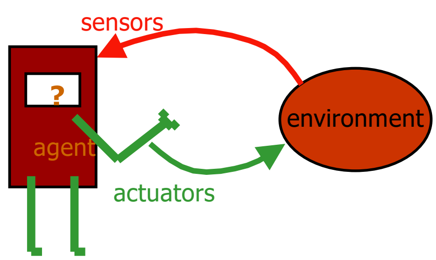


$X_t$ = unobserved
$E_t$ = observed

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

1

## Bookkeeping; reminders

- HW 1 grades posted soon
  - We have just gotten appropriate BB Ultra access
- HW 3 posted tonight
- All about projects
  - Project overview posted
- Probabilities over time
- Decision Theory

2

# Class Project

3

## About the Project

- Choosing a project
  - This will be up to you!
  - We would love to discuss your project ideas with you
- Deliverables
  - Project design
  - Phase I: working code, updated design
  - Phase II: final code
  - Final writeup

4

## About Group Work

- You **may** but do not have to do your project in a group of up to 4 people
- Highly recommended!
  - Gives you someone to work with, talk to, etc.
  - Gives you practice with teamwork, which is professionally important
- Expectation is for more ambitious (aka interesting) projects, **but…**
- All deliverables are the same
  - E.g., final paper is still 6-8 pages

5

## Project Ideas, Non-Exhaustively (1)

- Choose a game, and create an agent to play that game using artificial intelligence. Examples: chess, bridge, Minecraft, Wordle
- Develop an agent designed to interact intelligently with people in some context, for example, a chatbot or virtual assistant
- Develop an agent that (hypothetically) interacts with some real-world phenomenon, for example, the stock market
- Develop a recommender system for some existing corpus, for example, to recommend Netflix suggestions

6

## Project Ideas, Non-Exhaustively (2)

- Apply machine learning techniques to some existing corpus to draw conclusions, for example, a plagiarism detector, a COVID-19 outbreak predictor
- Develop some toolkit for solving a standard type of AI problem, or extending such a toolkit with new capabilities (a software development project)
- Use NLP to analyze documents and draw intelligent conclusions, for example, a resume analyzer, a spoiler detector
- Formulate, implement, and compare a novel solution to an existing problem
- Formulate, acquire data for, and apply a sufficient baseline for a novel task

7

## Deliverables: Project Design

- A written document **in AAAI conference format**
  - Author kit information available from the project description
  - Author kit includes templates for Word and LaTeX
- ~2 pages
- Include:
  - Idea: A description and motivation of the project
  - A description of the AI technique(s) you are going to use
  - A description of what you will implement in each phase
  - How your implemented system draws on ideas from the AI literature
  - Initial references
  - Your evaluation strategy

8

## Deliverables: Phase I

- Updated version of project plan (~3-4 pages)
  - Progress to date, evaluation of current functionality
- Code base
  - A working, but incomplete, version of your final project
    - Examples: it plays bridge, but chooses cards unintelligently; it reads in stock market data and proposes trades, but not well; it conducts a dialog with someone, but the utterances are gibberish
    - **What this means for your specific project can be discussed with us**
  - Must work on standard Linux systems
  - Include everything necessary to run your project, including a README and a dataset if appropriate

9

## Deliverables: Phase II

- Final code base
  - A complete system performing a task
  - Must work on standard Linux systems
  - Include everything necessary to run your project, including a README and a dataset if appropriate
  - May include evaluation-specific data, e.g., a set of sample stock market interactions

10

## Deliverables: Final Writeup

- AAAI format conference paper
- 6-8 pages, not counting references
  - Includes standard paper stuff like title and abstract
- Specific sections are **recommended** in the project description:
  - Introduction – description and motivation for the project
  - Related work – how your solution fits into the landscape
  - Approach – the core description of the work you did
  - Results – your evaluation strategy, description and analysis of results
  - Conclusion – final discussion of the work, future/follow-up work

11

## Forming Groups

- End of class today or beginning of class Thursday
- 2-4 people
  - Get together with your group and:
  - Everyone trade names and email addresses
  - One person email group member list to me & TA
- Pull up project description writeup
- Start talking about possible projects!

12

# Uncertainty

13

## Today's Class

- Making Decisions Under Uncertainty
  - Tracking Uncertainty over Time
  - Decision Making under Uncertainty
  - Decision Theory
  - Utility

14

## Why?

- The last couple of lectures looked at techniques to handle *uncertainty*
  - Bayesian networks
- The formalism is **static**, and so has limited ability to handle changing information
- Lots of reasoning tasks involve a **dynamic world**
  - Monitoring a patient
  - Tracking an airplane
  - Identifying the location of a robot
- This week we'll look at models that can handle such dynamic situations
  - Based on Bayesian networks

www.sci.brooklyn.cuny.edu/~parsons/courses/740-fall-2011/notes/lect07.pdf

15

## Introduction

- The world is not a well-defined place.
- Sources of uncertainty
  - Uncertain **inputs**: What's the temperature?
  - Uncertain (imprecise) **definitions**: Is UMBC a good school?
  - Uncertain (unobserved) **states**: What's the top card?
- There is uncertainty in **inferences**
  - If I have a blistery, itchy rash and was gardening all weekend I **probably** have poison ivy
    - Reasoning from **observations** to most likely **causes**

16

## Sources of Uncertainty

- Uncertain **inputs**
  - Missing data
  - Noisy data
- Uncertain **knowledge**
  - >1 cause → >1 effect
  - Incomplete knowledge of causality
  - Probabilistic effects
- Uncertain **outputs**
  - All uncertain:
    - Reasoning-by-default
    - Abduction & induction
    - Incomplete deductive inference
- Result is derived correctly but wrong in real world

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

17

## Reasoning Under Uncertainty

- People constantly make decisions anyhow.
  - Very successfully!
  - How?
  - More formally: how do we reason under uncertainty with inexact knowledge?
- Step one: **understanding what we know**

18

# Part I: Modeling Uncertainty Over Time

19

## States and Observations

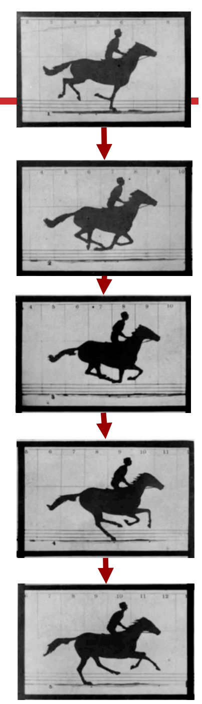

- Agents don't have a continuous view of world
  - People don't either!
- We see things as a series of snapshots:
- **Observations**, associated with **time slices**
  - $t_1, t_2, t_3, \ldots$
- Each snapshot contains all variables, observed or not
  - $\mathbf{X}_t$ = (unobserved) state variables at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

*Image: www.smithsonianmag.com/smithsonian-institution/how-19th-century-photographer-first-gif-galloping-horse-180970990*

20

## States and Observations, Intuitively

- So, we consider the world as a series of **time slices**
- Each slice contains some variables:
  - The set $\mathbf{X}_t$ which we can't observe; and
  - The set $\mathbf{E}_t$ which we can observe.
- At a given point in time we have an observation $\mathbf{E}_t = \mathbf{e}_t$
- What would be an example?



21

## Reasoning Over Time: Umbrella Example

- Consider you live and work in some location without a window
- You want to know whether it is raining
- Your only information is looking at whether somebody who comes into your office each morning is carrying an umbrella

Each day is one value of $t$

$\mathbf{E}_t$ contains the single variable $U_t$ (or $Umbrella_t$):
- Is the person carrying an umbrella?

$\mathbf{X}_t$ contains the single variable $R_t$ (or $Rain_t$):
- Is it raining?



22

# Temporal Probabilistic Agent

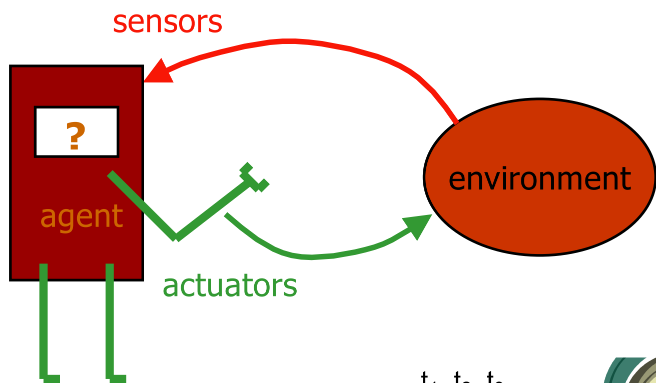


$t_1$, $t_2$, $t_3$, …

23

# Uncertainty and Time

- The world changes; we need to track and predict it
  - Examples: weather, diabetes management, traffic monitoring, stock markets, …
  - How does blood sugar change over time?
- Tasks: **track** changes; **predict** changes
- Basic idea:
  - For each time step, copy state and evidence variables
  - Model uncertainty in **change over time** (the Δ)
  - Incorporate new observations as they arrive

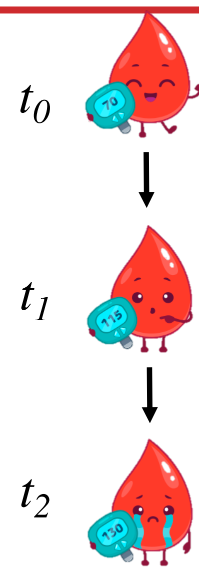


*Image: www.buzzrx.com/blog/how-are-dangerous-blood-sugar-levels-defined*

24

## Uncertainty and Time

- Basic idea:
  - Copy state and evidence variables for each time step
  - Model uncertainty in change over time
  - Incorporate new observations as they arrive
- $\mathbf{X}_t$ = unobserved/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

25

## States (more formally)

- Change is viewed as series of snapshots
  - Time slices/timesteps
  - Each describing the **state** of the world at a particular time
    - So we also refer to these as *states*
- Each time slice/timestep/state is represented as 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** $\mathbf{E}_t$

26

## Observations (more formally)

- Time slice (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** $\mathbf{E}_t$
- An **observation** is a set of observed variable instantiations at some timestep
- Observation at time *t*: $\mathbf{E}_t = e_t$
  - (for some values $e_t$)
- $\mathbf{X}_{a:b}$ denotes the set of variables from $\mathbf{X}_a$ to $\mathbf{X}_b$

27

## Back to Umbrellas

- State sequence starts at *t* = 0, and the interval between slices in general depends on the problem
  - For Umbrella, it is one day
  - In robot localization it is pretty arbitrary
- First piece of evidence arrives at *t* = 1
- So, the umbrella world is: $R_0, R_1, R_2, \ldots$
  $U_1, U_2, U_3, \ldots$
- *a* : *b* means the sequence of integers from *a* to *b*
- so $U_{2:4}$ is the sequence: $U_2, U_3, U_4$

www.sci.brooklyn.cuny.edu/~parsons/courses/740-fall-2011/notes/lect07.pdf

28

## Transition and Sensor Models

- We need to add two components to this backbone:
  - How the world evolves: the ***Transition model***
  - What the evidence tells us: the ***Sensor model***
- The transition model tells us: $\mathbf{P}(\mathbf{X}_t | \mathbf{X}_{0:t-1})$
- Or, the probability that it is raining today given the weather every previous day for as long as records have existed

www.sci.brooklyn.cuny.edu/~parsons/courses/740-fall-2011/notes/lect07.pdf

29

## Transition and Sensor Models, More Formally

- **So how do we model change over time?**
- Transition model
  - Models how the world changes over time
  - Specifies a probability distribution…
    - Over state variables at time *t*
    - Given values at previous times

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

This can get exponentially large!

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

30

# Markov Assumption(s)

We commonly assume a *first order* Markov process, where the current state depends only on the previous state

- **Markov Assumption**:
  - Make a *Markov assumption* that the value of the current state depends only on a finite fixed number of previous states.
  - $\mathbf{X}_t$ depends on some finite (usually fixed) number of previous $\mathbf{X}_i$'s
- **First-order Markov process**: $P(\mathbf{X}_t | \mathbf{X}_{0:t-1}) = P(\mathbf{X}_t | \mathbf{X}_{t-1})$

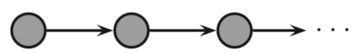

- $k^{th}$ order: depends on previous $k$ time steps

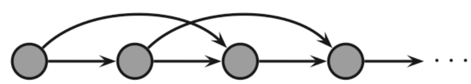

31

# Stationary Process

- Infinitely many possible values of *t*
  - Does each timestep need a distribution?
    - That is, do we need a distribution of what the world looks like at $t_3$, given $t_2$ AND a distribution for $t_{16}$ given $t_{15}$ AND …
- Usually circumvent this by assuming a **stationary process**:
  - Changes in the world state are governed by laws that do not themselves change over time
  - Transition model $P(\mathbf{X}_t | \mathbf{X}_{t-1})$ and sensor model $P(\mathbf{E}_t | \mathbf{X}_t)$ are time-invariant, i.e., they are the same for all *t*
  - Thus we only have one, general $\mathbf{P}(\mathbf{X}_t | \mathbf{X}_{t-1})$

32

# Markov Sensor Model

- **Sensor Markov assumption:** Agent's *observations* depend *only* on actual current state of the world
  - The evidence variables $\mathbf{E}_t$ *could* depend on lots of previous variables…
  - But we will assume the state is constructed in such a way that evidence only depends on the current state
- A Markov assumption for the sensor model:

$$P(\mathbf{E}_t | \mathbf{X}_{0:t-1}, \mathbf{E}_{0:t-1}) = P(\mathbf{E}_t | \mathbf{X}_t)$$

33

# Complete Joint Distribution

- Given:
  - Transition model: $P(\mathbf{X}_t | \mathbf{X}_{t-1})$
  - Sensor model: $P(\mathbf{E}_t | \mathbf{X}_t)$
  - Prior probability: $P(\mathbf{X}_0)$
- Then we can specify a **complete joint distribution** of a sequence of states:

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

- What's the joint probability of specific instantiations?

34

# Example: Is it raining, given umbrellas?

- Here are the sensor and observation models for the umbrella world:

| $R_{t-1}$ | $P(R_t \mid R_{t-1})$ |
|---|---|
| *t* | 0.7 |
| *f* | 0.3 |

Weather has a 30% chance of changing and a 70% chance of staying the same.

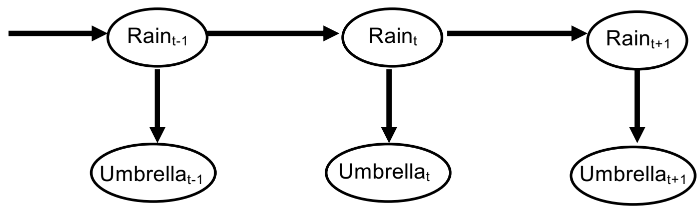


*(As before, arrows run from causes to effects)*

| $R_t$ | $P(U_t \mid R_t)$ |
|---|---|
| *t* | 0.9 |
| *f* | 0.2 |

If it's raining, the probability of someone carrying an umbrella is .9; if it's NOT raining, the probability of seeing an umbrella is .2

*Fully worked out HMM for rain: http://www.sci.brooklyn.cuny.edu/~parsons/courses/740-fall-2011/notes/lect07.pdf*

35

# Inference Tasks

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

36

## Inference Task Examples

- **Filtering**: What is the probability that it is raining today, given all of the umbrella observations up through today? $P(\mathbf{X}_t | e_1,\ldots,e_t)$
- **Prediction**: What is the probability that it will rain tomorrow, given all of the umbrella observations up through today? $P(\mathbf{X}_{t+k} | e_1,\ldots,e_t)$
- **Smoothing**: What is the probability that it rained yesterday, given all of the umbrella observations through today? $P(\mathbf{X}_k | e_1,\ldots,e_t)$
- **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? $\arg\max_{x1,..xt} P(x_1,\ldots,x_t | e_1,\ldots,e_t)$

37

## Filtering

- For each day $t$, $\mathbf{E}_t$ contains variable $U_t$ (whether the umbrella appears) and $\mathbf{X}_t$ contains state variable $R_t$ (whether it's raining)
- Compute the current belief state, given all evidence to date
- Maintain a **current** state estimate and update it
  - Instead of looking at all observed values in history
  - Also called **state estimation**
- Given result of filtering up to time $t$, agent must compute result at $t$+1 from new evidence $\mathbf{e}_{t+1}$:

$$P(\mathbf{X}_{t+1} \mid \mathbf{e}_{1:t+1}) = f(\mathbf{e}_{t+1},\ P(\mathbf{X}_t \mid \mathbf{e}_{1:t}))$$

… for some function $f$.

38

# Filtering

- A good algorithm for filtering will maintain a current state estimate and update it at each point.
- $P(X_{t+1}|e_{1:t+1}) = f(P(X_t|e_{1:t}), e_{t+1})$
- Where X is the random variable and e is evidence
- Saves recomputation.
- It turns out that this is easy enough to come up with.

http://www.sci.brooklyn.cuny.edu/~parsons/courses/740-fall-2011/notes/lect07.pdf

39

# Filtering

- We rearrange the formula for:
  - $P(X_{t+1}|e_{1:t+1})$
- First, we divide up the evidence:
  - $P(X_{t+1}|e_{1:t+1}) = P(X_{t+1}|e_{1:t}, e_{t+1})$
- Then we apply Bayes rule, remembering the use of the normalization factor α:
  - $P(X_{t+1}|e_{1:t+1}) = \alpha P(e_{t+1}|X_{t+1}, e_{1:t})\, P(X_{t+1}|e_{1:t})$
- And after that we use the Markov assumption on the sensor model:
  - $P(X_{t+1}|e_{1:t+1}) = \alpha P(e_{t+1}|X_{t+1})P(X_{t+1}|e_{1:t})$
  - The result of this assumption is to make that first term on the right hand side ignore all the evidence — the probability of the observation at $t + 1$ only depends on the value of $X_{t+1}$.

http://www.sci.brooklyn.cuny.edu/~parsons/courses/740-fall-2011/notes/lect07.pdf

40

## Filtering

- Let's look at that expression some more:
  - $P(X_{t+1}|e_{1:t+1}) = \alpha P(e_{t+1}|X_{t+1})P(X_{t+1}|e_{1:t})$
- The first term on the right updates with the new evidence and the second term on the right is a one step prediction from the evidence up to *t* to the state at *t* + 1
- Next we condition on the current state P(X):
  - $P(X_{t+1}|e_{1:t+1}) = \alpha P(e_{t+1}|X_{t+1}) \Sigma x_t \, P(X_{t+1}|x_t, e_{1:t})P(x_t|e_{1:t})$
- Finally, we apply the Markov assumption again:
  - $P(X_{t+1}|e_{1:t+1}) = \alpha P(e_{t+1}|X_{t+1}) \Sigma x_t \, P(X_{t+1}|x_t)P(x_t|e_{1:t})$
- We'll call the bit on the right $f_{1:t}$

http://www.sci.brooklyn.cuny.edu/~parsons/courses/740-fall-2011/notes/lect07.pdf

41

## Filtering

- $f_{1:t}$ gives us the required recursive update.
  - The probability distribution over the state variables at *t* + 1 is a function of the transition model, the sensor model, and what we know about the state at time *t*.
- Space and time constant, independent of *t*.
- This allows a limited agent to compute the current distribution for any length of sequence.

http://www.sci.brooklyn.cuny.edu/~parsons/courses/740-fall-2011/notes/lect07.pdf

42

# Recursive Estimation

- We use ***recursive estimation*** to compute $P(X_{t+1} \mid e_{1:t+1})$ as a function of $e_{t+1}$ and $P(X_t \mid e_{1:t})$

1. Project current state forward (t → t+1)
2. Update state using new evidence $\mathbf{e}_{t+1}$

$\mathrm{P}(\mathbf{X}_{t+1} \mid \mathbf{e}_{1:t+1})$ as function of $\mathbf{e}_{t+1}$ and $\mathrm{P}(\mathbf{X}_t \mid \mathbf{e}_{1:t})$:

$$\mathrm{P}(\mathbf{X}_{t+1} \mid \mathbf{e}_{1:t+1}) = \underline{\mathrm{P}(\mathbf{X}_{t+1} \mid \mathbf{e}_{1:t}, \mathbf{e}_{t+1})}$$

43

# Recursive Estimation

- $\mathrm{P}(\mathbf{X}_{t+1} \mid \mathbf{e}_{1:t+1})$ as a function of $\mathbf{e}_{t+1}$ and $\mathrm{P}(\mathbf{X}_t \mid \mathbf{e}_{1:t})$:

$$P(X_{t+1} \mid e_{1:t+1}) = P(X_{t+1} \mid e_{1:t}, e_{t+1}) \quad \text{dividing up evidence}$$
$$= \alpha P(e_{t+1} \mid X_{t+1}, e_{1:t})\, P(X_{t+1} \mid e_{1:t}) \quad \text{Bayes rule}$$
$$= \alpha P(e_{t+1} \mid X_{t+1})\, P(X_{t+1} \mid e_{1:t}) \quad \text{sensor Markov assumption}$$

- $\mathrm{P}(\mathbf{e}_{t+1} \mid \mathbf{X}_{1:t+1})$ updates with new evidence (from sensor)
- One-step prediction by conditioning on current state X:

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

44

# Recursive Estimation

- One-step prediction by conditioning on current state X:

$$P(\mathbf{X}_{t+1} \mid \mathbf{e}_{1:t+1}) = \alpha P(e_{t+1} \mid X_{t+1}) \sum_{x_t} \underbrace{P(X_{t+1} \mid x_t)}_{\text{transition model}} \underbrace{P(x_t \mid e_{1:t})}_{\text{current state}}$$

- …which is what we wanted!
- So, think of $\mathrm{P}(\mathbf{X}_t \mid \mathbf{e}_{1:t})$ as a "message" $f_{1:t+1}$
  - Carried forward along the time steps
  - Modified at every transition, updated at every new observation
- This leads to a recursive definition:

$$f_{1:t+1} = \alpha \text{ FORWARD}(f_{1:t}, e_{t+1})$$

45

# Filtering: Umbrellas example

- The prior is ⟨0.5, 0.5⟩. (R=*t*, R=*f*)
- We can first predict whether it will rain on day 1 given what we already know:
- $\mathbf{P}(\mathbf{R}_1) = \Sigma_{r_0} \mathbf{P}(R_1 | r_0)\, P(r_0)$
  $= \langle 0.7,0.3\rangle \times 0.5 + \langle 0.3,0.7\rangle \times 0.5$
  $= \langle 0.35,0.15\rangle + \langle 0.15,0.35\rangle$
  $= \langle 0.5,0.5\rangle$
- As we should expect, this just gives us the prior — that is the probability of rain when we don't have any evidence.

http://www.sci.brooklyn.cuny.edu/~parsons/courses/740-fall-2011/notes/lect07.pdf

46

## Filtering: Umbrellas example

- However, we have observed the umbrella, so that $U_1$ = *true*, and we can update using the sensor model:
- $\mathbf{P}(\mathbf{R}_1 | U_1) = \alpha \mathbf{P}(u_1 | R_1)\mathbf{P}(R_1)$
  $= \alpha\langle 0.9,0.2\rangle\langle 0.5,0.5\rangle$
  $= \alpha\langle 0.45,0.1\rangle$
  $\approx \langle 0.818,0.182\rangle$*
- So, since umbrella is strong evidence for rain, the probability of rain is much higher once we take the observation into account

** α is, as previously, just a normalizing constant that makes the probabilities add up to 1. We get it by dividing each element by the sum of both elements, e.g., 0.45/(0.45+0.1) ≈ 0.818.*

http://www.sci.brooklyn.cuny.edu/~parsons/courses/740-fall-2011/notes/lect07.pdf

47

## Filtering: Umbrellas example

- We can then carry out the same computation for Day 2, first predicting whether it will rain on day 2 given what we already saw:
- $\mathbf{P}(\mathbf{R}_2 | u_1) = \Sigma_{r_1} \mathbf{P}(R_2 | r_1)P(r_1 | u_1)$
  $= \langle 0.7,0.3\rangle \times 0.818 + \langle 0.3,0.7\rangle \times 0.182$
  $\approx \langle 0.627,0.373\rangle$
- So even without evidence of rain on the second day there is a higher probability of rain than the prior because rain tends to follow rain.
  - (In this model rain tends to persist.)

http://www.sci.brooklyn.cuny.edu/~parsons/courses/740-fall-2011/notes/lect07.pdf

48

# Filtering: Umbrellas example

- Then we can repeat the evidence update, $u_2$ ($U_2$ = *true*), so:
- $\mathbf{P}(\mathbf{R}_2|u_1,u_2) = \alpha\mathbf{P}(u_2|R_2)\mathbf{P}(R_2|u_1)$
  $= \alpha\langle 0.9,0.2\rangle\langle 0.627,0.373\rangle$
  $= \alpha\langle 0.565,0.075\rangle$
  $\approx \langle 0.883,0.117\rangle$
- So, the probability of rain increases again, and is higher than on day 1.

http://www.sci.brooklyn.cuny.edu/~parsons/courses/740-fall-2011/notes/lect07.pdf

49

# Filtering: Umbrellas example

- Put more succinctly:

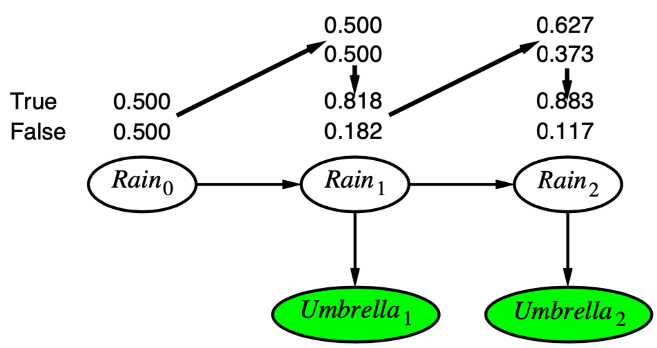


- We can think of the calculation as messages passed along the chain

http://www.sci.brooklyn.cuny.edu/~parsons/courses/740-fall-2011/notes/lect07.pdf

50

## Umbrellas, summarized

- $P(Rain_1 = t)$
  $= \Sigma_{Rain_0} P(Rain_1 = t \mid Rain_0)\, P(Rain_0)$
  $= 0.70 * 0.50 + 0.30 * 0.50 =$ **$0.50$**
- $P(Rain_1 = t \mid Umbrella_1 = t)$
  $= \alpha\, P(Umbrella_1 = t \mid Rain_1 = t)\, P(Rain_1 = t)$
  $= \alpha * 0.90 * 0.50 = \alpha * 0.45 \approx$ **$0.818$**
- $P(Rain_2 = t \mid Umbrella_1 = t)$
  $= \Sigma_{Rain_1} P(Rain_2 = t \mid Rain_1)\, P(Rain_1 \mid Umbrella_1 = t)$
  $= 0.70 * 0.818 + 0.30 * 0.182 \approx$ **$0.627$**
- $P(Rain_2 = t \mid Umbrella_1 = t, Umbrella_2 = t)$
  $= \alpha\, P(Umbrella_2 = t \mid Rain_2 = t)\, P(Rain_2 = t \mid Umbrella_1 = t)$
  $= \alpha * 0.90 * 0.627 \approx \alpha * 0.564 \approx$ **$0.883$**

51

# Part II: Decision Making Under Uncertainty

55

## Decision Making Under Uncertainty

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

56

## Reasoning Under Uncertainty

- How do we **reason** 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)

57

# Decision-Making Tools

- Decision Theory
  - Normative: how *should* agents make decisions?
  - Descriptive: how *do* agents make decisions?
- **Utility** and utility functions
  - Something's **perceived ability to satisfy needs or wants**
  - A mathematical function that ranks alternatives by utility

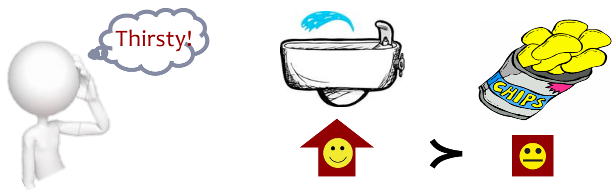


58

# What is Decision Theory?

- Mathematical study of strategies for optimal decision-making
  - Options involve different risks
  - **Expectations** of gain or loss
- The study of identifying:
  - The values, uncertainties and other issues relevant to a decision
  - The resulting optimal decision for a rational agent

| | Rain | No Rain |
|---|---|---|
| **Carries Umbrella** | Best | Bad |
| **No Umbrella** | Worst | Great |

59

## Decision Theory

- Combines **probability** and **utility** → Agent that makes **rational** decisions (takes rational actions)
  - On average, lead to desired outcome
- First-pass simplifications:
  - Want most desirable *immediate* outcome (episodic)
  - Nondeterministic, partially observable world
- Definition of **action**:
- An action *a* in state *s* leads to outcome *s'*, RESULT:
  - RESULT(*a*) is a random variable; domain is possible outcomes
  - $\mathrm{P}(\mathrm{RESULT}(a) = s' \mid a,\ e))$

And now this ties into understanding *states over time*

60

## Expected Value

- Expected Value
  - The **predicted future value** of a variable, calculated as:
  - The sum of all possible values
  - Each multiplied by the probability of its occurrence

A $1000 bet for a 20% chance to win $10,000?

EV = [20%($10,000) + 80%($0)] = $2000

61

# Satisficing

- Satisficing: achieving a goal sufficiently
  - Achieving the goal "more" does not increase utility of resulting state
  - Portmanteau of "satisfy" and "suffice"

Win a baseball game by 1 point now, or 2 points in another inning?

Full credit for a search is <= 3K nodes visited. You're at 2K. Spend an hour making it 1K?

Do you stop the coin flipping game at 1-0, or continue playing, hoping for 2-0?

At the end of semester, you can stop with a B. Do you take the exam?

You're thirsty. Water is good. Is more water better?

62

# 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

| | **Rain** | **No Rain** |
|---|---|---|
| **Carries Umbrella** | Best | Bad |
| **No Umbrella** | Worst | Great |

| | **Rain** | **No Rain** |
|---|---|---|
| **Carries Umbrella** | 5 | 1 |
| **No Umbrella** | 0 | 4 |

$$C \land R \succ \neg C \land \neg R \succ C \land \neg R \succ C \land \neg R$$

64

# Rational Agents

- Rationality (an overloaded word).
- A rational agent…
  - Behaves according to a **ranking over possible outcomes**
  - Which is:
    - Complete (covers all situations)
    - Consistent
    - Optimizes over strategies to best serve a desired interest
- Humans are none of these.

65

# Preferences

- An agent chooses among:
  - *Prizes* (A, B, etc.)
  - *Lotteries* (situations with uncertain prizes and probabilities)

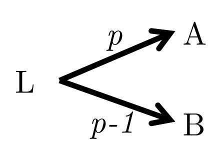


- Notation:
  - A ≻ B — A preferred to B
  - A ∼ B — Indifference between A and B
  - A ≻∼ B — B **not preferred** to A

66

# Expected Utility

- Goal: find best of expected outcomes
- 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
  - state = some state of the world at some timestep
- Utility function U(s) is the utility of a state, i.e., **desirability**

67

# Expected Utility

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

$$\mathrm{EU}[\mathrm{A}] = \mathrm{S}_{\mathrm{i}=1,\ldots,\mathrm{n}}\, \mathrm{P}(\mathrm{x_i}|\mathrm{A})\mathrm{U}(\mathrm{x_i})$$

68

## One State/One Action Example

- We start out in state 0. What's the utility of taking action A1?

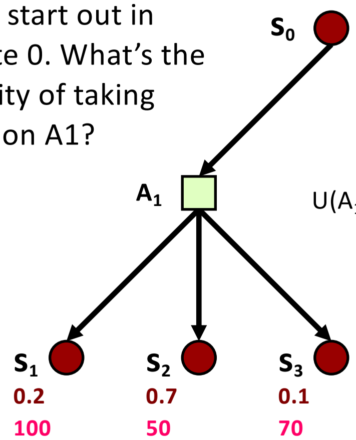


$U(A_1, S_0) = \underline{100 \times 0.2} + \underline{50 \times 0.7} + \underline{70 \times 0.1}$

$= 20 + 35 + 7$

$= 62$

69

## One State/Two Actions Example

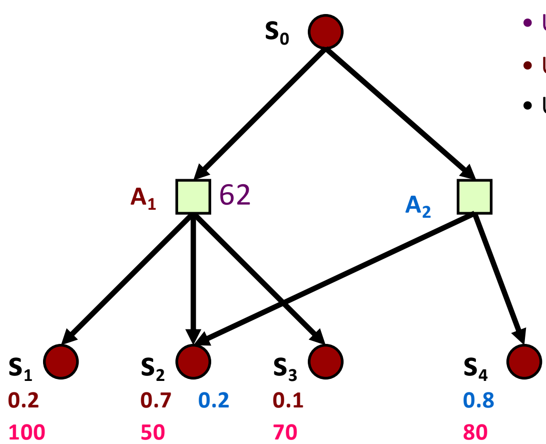


- $U(A_1, S_0) = 62$
- $U(A_2, S_0) = 74$
- $U(S_0) = \max_a\{U(a,S_0)\}$
  $= 74$

70

70

## Introducing Action Costs

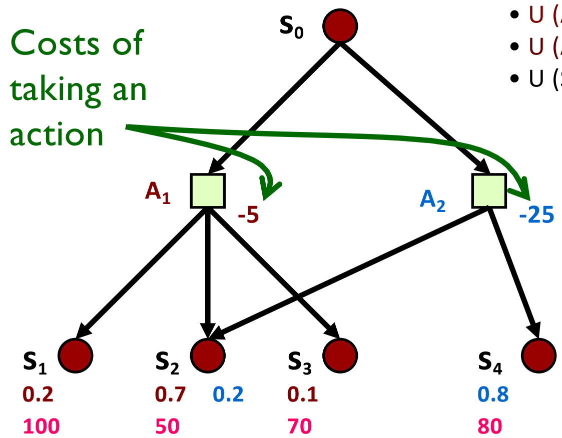


- U ($A_1$, $S_0$) = 62 – 5 = 57
- U ($A_2$, $S_0$) = 74 – 25 = 49
- U ($S_0$) = $\max_a$ {U(a, $S_0$)} = 57

71

## 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
- Decision-making is solved!
  - Not quite…

72

## Rational Preferences

- Preferences of a rational agent must obey constraints
  - Transitivity $(A \succ B) \wedge (B \succ C) \Rightarrow (A \succ C)$
  - Monotonicity $(A \succ B) \Rightarrow [p > q \Leftrightarrow [p, A; 1 - p, B] \succ [q, A; 1 - q, B])$
  - Orderability $(A \succ B) \vee (B \succ A) \vee (A \sim B)$
  - Substitutability $(A \sim B) \Rightarrow [p,A; 1 - p, C] \sim [p,B; 1 - p,C] )$
  - Continuity $(A \succ B \succ C \Rightarrow \exists p\ [p,A; 1-p,C] \sim B )$
- Rational preferences give behavior that **maximizes expected utility**
- Violating these constraints leads to irrationality
  - For example: an agent with intransitive preferences can be induced to give away all its money.

73

## 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

74

## Money

- Money does not behave as a utility function
  - That is, people don't maximize expected value of dollars.
- People are risk-averse:
  - Given a lottery L with expected monetary value EMV(L), usually U(L) < U(EMV(L))

Want to bet $1000 for a 20% chance to win $10,000?
[20%($10,000)+80%($0)] = $2000 > [100%($1000)]

- Expected Utility Hypothesis
  - rational behavior maximizes the expectation of some function u… which need not be monetary

76

## Money Versus Utility

- Money as 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: $U(L) = U(S_{EMV(L)})$

77

## Maximizing Expected Utility

- Utilities map **states** to **real numbers**.
  - Which numbers?
- People are **terrible** at mapping their preferences
  - Give each of these things a utility between 1 and 10:
    - Winning the lottery
    - Getting an A on an exam
    - Failing a class (you won't though)
    - Getting hit by a truck



## Maximizing Expected Utility

- Standard approach to assessment of human utilities:
  - Compare a state $A$ to a standard lottery $L_p$ that has
    - "best possible prize" $u^{\top}$ with probability $p$
    - "worst possible catastrophe" $u^{\perp}$ with probability $(1-p)$
  - adjust lottery probability $p$ until $A \sim L_p$

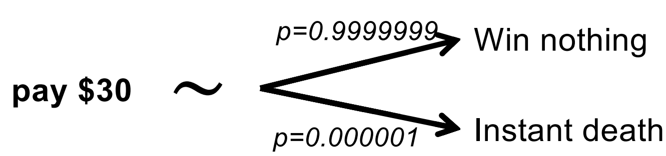

## Or, Less Grim…

- You are designing a cool new robot-themed attraction for Disneyworld!
- You could add a part that takes the project from $500M to $750M
- What piece of information do you need to decide whether this is the best action to take?

80