

CPS 170: Artificial Intelligence

<http://www.cs.duke.edu/courses/spring09/cps170/>

Introduction to probability

Instructor: Vincent Conitzer

Uncertainty

- So far in course, everything deterministic
- If I walk with my umbrella, I **will not** get wet
- But: there is some chance my umbrella will break!
- Intelligent systems must take possibility of failure into account...
 - May want to have backup umbrella in city that is often windy and rainy
- ... but should not be excessively conservative
 - Two umbrellas not worthwhile for city that is usually not windy
- Need **quantitative** notion of uncertainty

Probability

- Example: roll two dice
- **Random variables:**
 - X = value of die 1
 - Y = value of die 2
- Outcome is represented by an ordered pair of values (x, y)
 - E.g., $(6, 1)$: $X=6, Y=1$
 - **Atomic event** or **sample point** tells us the **complete** state of the world, i.e., values of **all** random variables
- Exactly one atomic event will happen; each atomic event has a ≥ 0 probability; sum to 1
 - E.g., $P(X=1 \text{ and } Y=6) = 1/36$

6	1/36	1/36	1/36	1/36	1/36	1/36
5	1/36	1/36	1/36	1/36	1/36	1/36
4	1/36	1/36	1/36	1/36	1/36	1/36
3	1/36	1/36	1/36	1/36	1/36	1/36
2	1/36	1/36	1/36	1/36	1/36	1/36
1	1/36	1/36	1/36	1/36	1/36	1/36
	1	2	3	4	5	6

- An **event** is a proposition about the state (=subset of states)
 - $X+Y = 7$
- Probability of event = sum of probabilities of atomic events where event is true

Cards and combinatorics

- Draw a hand of 5 cards from a standard deck with $4 \cdot 13 = 52$ cards (4 suits, 13 ranks each)
- Each of the $\binom{52}{5}$ hands has same probability $1/\binom{52}{5}$
- Probability of event = number of hands in that event / $\binom{52}{5}$
- What is the probability that...
 - no two cards have the same rank?
 - you have a flush (all cards the same suit?)
 - you have a straight (5 cards in order of rank, e.g., 8, 9, 10, J, Q)?
 - you have a straight flush?
 - you have a full house (three cards have the same rank and the two other cards have the same rank)?

Facts about probabilities of events

- If events A and B are disjoint, then
 - $P(A \text{ or } B) = P(A) + P(B)$
- More generally:
 - $P(A \text{ or } B) = P(A) + P(B) - P(A \text{ and } B)$
- If events A_1, \dots, A_n are disjoint and exhaustive (one of them must happen) then $P(A_1) + \dots + P(A_n) = 1$
 - Special case: for any random variable, $\sum_x P(X=x) = 1$
- Marginalization: $P(X=x) = \sum_y P(X=x \text{ and } Y=y)$

Conditional probability

- We might know something about the world – e.g., “ $X+Y=6$ or $X+Y=7$ ” – given this (and **only** this), what is the probability of $Y=5$?
- Part of the sample space is eliminated; probabilities are renormalized to sum to 1

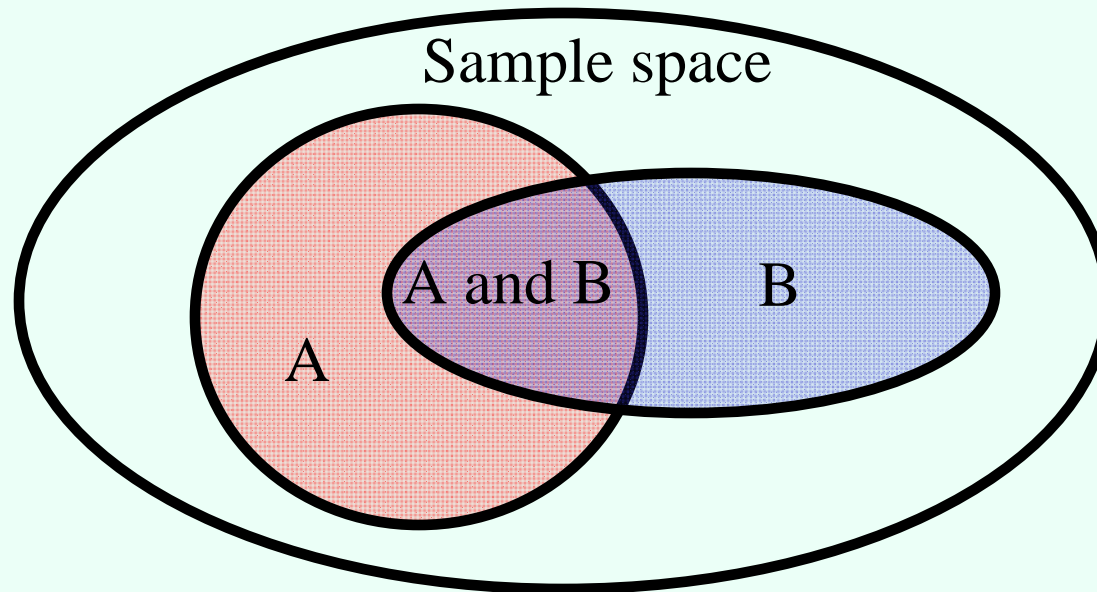
Y						
6	1/36	1/36	1/36	1/36	1/36	1/36
5	1/36	1/36	1/36	1/36	1/36	1/36
4	1/36	1/36	1/36	1/36	1/36	1/36
3	1/36	1/36	1/36	1/36	1/36	1/36
2	1/36	1/36	1/36	1/36	1/36	1/36
1	1/36	1/36	1/36	1/36	1/36	1/36
	1	2	3	4	5	6 X

Y						
6	1/11	0	0	0	0	0
5	1/11	1/11	0	0	0	0
4	0	1/11	1/11	0	0	0
3	0	0	1/11	1/11	0	0
2	0	0	0	1/11	1/11	0
1	0	0	0	0	1/11	1/11
	1	2	3	4	5	6 X

- $P(Y=5 \mid (X+Y=6) \text{ or } (X+Y=7)) = 2/11$

Facts about conditional probability

- $P(A | B) = P(A \text{ and } B) / P(B)$



- $P(A | B)P(B) = P(A \text{ and } B) = P(B | A)P(A)$
- $P(A | B) = P(B | A)P(A)/P(B)$
 - Bayes' rule

Conditional probability and cards

- Given that your first two cards are Queens, what is the probability that you will get at least three Queens?
- Given that you have at least two Queens (not necessarily the first two), what is the probability that you have at least three Queens?
- Given that you have at least two Queens, what is the probability that you have three Kings?

How can we scale this?

- In principle, we now have a complete approach for reasoning under uncertainty:
 - Specify probability for every atomic event,
 - Can compute probabilities of events simply by summing probabilities of atomic events,
 - Conditional probabilities are specified in terms of probabilities of events: $P(A | B) = P(A \text{ and } B) / P(B)$
- If we have n variables that can each take k values, how many atomic events are there?

Independence

- Some variables have nothing to do with each other
- Dice: if $X=6$, it tells us nothing about Y
- $P(Y=y \mid X=x) = P(Y=y)$
- So: $P(X=x \text{ and } Y=y) = P(Y=y \mid X=x)P(X=x) = P(Y=y)P(X=x)$
 - Usually just write $P(X, Y) = P(X)P(Y)$
 - Only need to specify $6+6=12$ values instead of $6*6=36$ values
 - Independence among 3 variables: $P(X, Y, Z) = P(X)P(Y)P(Z)$, etc.
- Are the events “you get a flush” and “you get a straight” independent?

An example without cards or dice

	Rain in Beaufort	Sun in Beaufort
Rain in Durham	.2	.1
Sun in Durham	.2	.5

*(disclaimer:
no idea if
these
numbers are
realistic)*

- What is the probability of
 - Rain in Beaufort? Rain in Durham?
 - Rain in Beaufort, given rain in Durham?
 - Rain in Durham, given rain in Beaufort?
- Rain in Beaufort and rain in Durham are **correlated**

A possibly rigged casino

- With probability $\frac{1}{2}$, the casino is rigged and has dice that come up 6 only $\frac{1}{12}$ of the time, and $\frac{1}{3}$ of the time

$Z=0$ (fair casino)

Y	$Z=0$ (fair casino)					
6	1/72	1/72	1/72	1/72	1/72	1/72
5	1/72	1/72	1/72	1/72	1/72	1/72
4	1/72	1/72	1/72	1/72	1/72	1/72
3	1/72	1/72	1/72	1/72	1/72	1/72
2	1/72	1/72	1/72	1/72	1/72	1/72
1	1/72	1/72	1/72	1/72	1/72	1/72
	1	2	3	4	5	6 X

$Z=1$ (rigged casino)

Y	$Z=1$ (rigged casino)					
6	1/96	1/144	1/144	1/144	1/144	1/288
5	1/48	1/72	1/72	1/72	1/72	1/144
4	1/48	1/72	1/72	1/72	1/72	1/144
3	1/48	1/72	1/72	1/72	1/72	1/144
2	1/48	1/72	1/72	1/72	1/72	1/144
1	1/32	1/48	1/48	1/48	1/48	1/96
	1	2	3	4	5	6 X

- What is $P(Y=6)$?
- What is $P(Y=6|X=1)$?
- Are they independent?

Conditional independence

- Intuition:
 - the only reason that X tells us something about Y ,
 - is that X tells us something about Z ,
 - and Z tells us something about Y
- If we already know Z , then X tells us nothing about Y
- $P(Y | Z, X) = P(Y | Z)$ or
- $P(X, Y | Z) = P(X | Z)P(Y | Z)$
- “ X and Y are conditionally independent given Z ”

Medical diagnosis

- X: does patient have flu?
- Y: does patient have headache?
- Z: does patient have fever?
- $P(Y,Z|X) = P(Y|X)P(Z|X)$
- $P(X=1) = .2$
- $P(Y=1 | X=1) = .5, P(Y=1 | X=0) = .2$
- $P(Z=1 | X=1) = .4, P(Z=1 | X=0) = .1$
- What is $P(X=1|Y=1,Z=0)$?

*(disclaimer:
no idea if
these
numbers are
realistic)*

Conditioning can also introduce dependence

- X: is it raining?
 - $P(X=1) = .3$
- Y: are the sprinklers on?
 - $P(Y=1) = .4$
 - X and Y are independent
- Z: is the grass wet?
 - $P(Z=1 \mid X=0, Y=0) = .1$
 - $P(Z=1 \mid X=0, Y=1) = .8$
 - $P(Z=1 \mid X=1, Y=0) = .7$
 - $P(Z=1 \mid X=1, Y=1) = .9$

		<i>Not wet</i>	
		Raining	Not raining
No sprinklers	Sprinklers	.012	.056
	No sprinklers	.054	.378

		<i>Wet</i>	
		Raining	Not raining
No sprinklers	Sprinklers	.108	.224
	No sprinklers	.126	.042

- Conditional on $Z=1$, X and Y are **not** independent
- If you know $Z=1$, rain seems likely; then if you also find out $Y=1$, this “explains away” the wetness and rain seems less likely

Context-specific independence

- Recall $P(X, Y | Z) = P(X | Z)P(Y | Z)$ really means: for all x, y, z ,
- $P(X=x, Y=y | Z=z) = P(X=x | Z=z)P(Y=y | Z=z)$
- But it may not be true for *all* z
- $P(\text{Wet}, \text{RainingInLondon} | \text{CurrentLocation}=\text{New York}) = P(\text{Wet} | \text{CurrentLocation}=\text{New York})P(\text{RainingInLondon} | \text{CurrentLocation}=\text{New York})$
- **But not**
- $P(\text{Wet}, \text{RainingInLondon} | \text{CurrentLocation}=\text{London}) = P(\text{Wet} | \text{CurrentLocation}=\text{London})P(\text{RainingInLondon} | \text{CurrentLocation}=\text{London})$

Pairwise independence does not imply complete independence

- X is a coin flip, Y is a coin flip, $Z = X \text{ xor } Y$
- Clearly $P(X, Y, Z) \neq P(X)P(Y)P(Z)$
- But $P(Z|X) = P(Z)$
- Tempting to say X and Y are “really” independent, X and Z are “not really” independent
- But: X is a coin flip, Z is a coin flip, $Y = X \text{ xor } Z$ gives the exact same distribution

Monty Hall problem




image taken from http://en.wikipedia.org/wiki/Monty_Hall_problem

- Game show participants can choose one of three doors
- One door has a car, two have a goat
 - Assumption: car is preferred to goat
- Participant chooses door, but not opened yet
- At least one of the other doors contains a goat, the (knowing) host will open one such door (flips coin to decide if both have goats)
- Participant is asked whether she wants to switch doors (to the other closed door) – should she?

Expected value

- If Z takes numerical values, then the **expected value** of Z is $E(Z) = \sum_z P(Z=z) * z$
 - Weighted average (weighted by probability)
- Suppose Z is sum of two dice
- $E(Z) = (1/36)*2 + (2/36)*3 + (3/36)*4 + (4/36)*5 + (5/36)*6 + (6/36)*7 + (5/36)*8 + (4/36)*9 + (3/36)*10 + (2/36)*11 + (1/36)*12 = 7$
- Simpler way: $E(X+Y)=E(X)+E(Y)$ (always!)
 - **Linearity of expectation**
- $E(X) = E(Y) = 3.5$

Linearity of expectation...

- If a is used to represent an atomic state, then $E(X) = \sum_x P(X=x) \cdot x = \sum_x \left(\sum_{a: X(a)=x} P(a) \right) \cdot x = \sum_a P(a) \cdot X(a)$
- $E(X+Y) = \sum_a P(a) \cdot (X(a)+Y(a)) = \sum_a P(a) \cdot X(a) + \sum_a P(a) \cdot Y(a) = E(X)+E(Y)$

What is probability, anyway?

- Different philosophical positions:
- **Frequentism**: numbers only come from repeated experiments
 - As we flip a coin lots of times, we see experimentally that it comes out heads $\frac{1}{2}$ the time
 - Problem: for most events in the world, there is no history of **exactly** that event happening
 - Probability that the Democratic candidate wins the next election?
- **Objectivism**: probabilities are a real part of the universe
 - Maybe true at level of quantum mechanics
 - Most of us agree that the result of a coin flip is (usually) determined by initial conditions + classical mechanics
- **Subjectivism**: probabilities merely reflect agents' beliefs