CS 440/ECE 448 Lecture 10: Probability

Slides by Svetlana Lazebnik, 9/2016

Modified by Mark Hasegawa-Johnson, 2/2019

Outline

- Motivation: Why use probability?
  - Laziness, Ignorance, and Randomness
  - Rational Bettor Theorem
- Review of Key Concepts
  - Outcomes, Events
  - Random Variables; probability mass function (pmf)
  - Jointly random variables: Joint, Marginal, and Conditional pmf
  - Independent vs. Conditionally Independent events

Outline

- Motivation: Why use probability?
  - Laziness, Ignorance, and Randomness
  - Rational Bettor Theorem
- Review of Key Concepts
  - Outcomes, Events
  - Joint, Marginal, and Conditional
  - Independence and Conditional Independence

Motivation: Planning under uncertainty

- Recall: representation for planning
- **States** are specified as conjunctions of predicates
  - Start state: At(Me, UIUC) $\wedge$ TravelTime(35min,UIUC,CMI) $\wedge$ Now(12:45)
  - Goal state: At(Me, CMI, 15:30)
- **Actions** are described in terms of preconditions and effects:
  - Go(t, src, dst)
    - **Precond:** At(Me,src) $\wedge$ TravelTime(dt,src,dst) $\wedge$ Now($\leq$t)
    - **Effect:** At(Me, dst, t+dt)

Motivation: Planning under uncertainty

- Let action *Go(t)* = leave for airport at time *t*
  - Will *Go(t)* succeed, i.e., get me to the airport in time for the flight?
- Problems:
  - **Partial observability** (road state, other drivers' plans, etc.)
  - **Noisy sensors** (traffic reports)
  - **Uncertainty** in action outcomes (flat tire, etc.)
  - **Complexity** of modeling and predicting traffic
- Hence a purely logical approach either
  - Risks falsehood: "*Go(14:30)* will get me there on time," or
  - Leads to conclusions that are too weak for decision making:
    - *Go(14:30)* will get me there on time if there's no accident, it doesn't rain, my tires remain intact, etc., etc.
    - *Go(04:30)* will get me there on time

Probability

Probabilistic assertions summarize effects of

- **Laziness**: reluctance to enumerate exceptions, qualifications, etc. --- possibly a deterministic and known environment, but with **computational complexity limitations**
- **Ignorance**: lack of explicit theories, relevant facts, initial conditions, etc. --- environment that is **unknown** (we don’t know the transition function) or **partially observable** (we can’t measure the current state)
- **Intrinsically random phenomena** – environment is **stochastic**, i.e., given a particular (action,current state), the (next state) is drawn at random with a particular probability distribution

Outline

- Motivation: Why use probability?
  - Laziness, Ignorance, and Randomness
  - Rational Bettor Theorem
- Review of Key Concepts
  - Outcomes, Events
  - Joint, Marginal, and Conditional
  - Independence and Conditional Independence

Making decisions under uncertainty

- **Suppose the agent believes the following:**

  P(Go(deadline-25) gets me there on time) = 0.04

  P(Go(deadline-90) gets me there on time) = 0.70

  P(Go(deadline-120) gets me there on time) = 0.95

  P(Go(deadline-180) gets me there on time) = 0.9999

- **Which action should the agent choose?**
  - Depends on preferences for missing flight vs. time spent waiting
  - Encapsulated by a *utility function*
- **The agent should choose the action that maximizes the *expected utility*:**

  Prob(A succeeds) × Utility(A succeeds) + Prob(A fails) × Utility(A fails)

Making decisions under uncertainty

- More generally: the <u>expected utility</u> of an action is defined as:

$$\mathrm{E}[\mathrm{Utility}|\mathrm{Action}] = \sum_{outcomes} P(outcome|action)Utility(outcome)$$

- **Utility theory** is used to represent and infer preferences
- **Decision theory** = probability theory + utility theory

Where do probabilities come from?

- **Frequentism**
  - Probabilities are **relative frequencies**
  - For example, if we toss a coin many times,
    P(heads) is the proportion of the time the coin will come up heads
  - But what if we're dealing with an event that has never happened before?
    - What's the probability that the Earth will warm by 0.15*F this year?
- **Subjectivism**
  - Probabilities are **degrees of belief**
  - But then, how do we assign belief values to statements?
- A theoretical problem with Subjectivism:

Why do "beliefs" need to follow the laws of probability?

The Rational Bettor Theorem

- Why should the beliefs of a rational agent be consistent with the axioms of probability?
  - For example: why should $P(A) + P(\neg A) = 1$?
- Suppose an agent believes that $P(A)=0.7$, and $P(\neg A)=0.7$
- 1. Bet: if A occurs, agent wins \$100. If A doesn't occur, agent loses \$105.
  - Agent believes $P(A) = .7 > 100/(100+105) = .48$, so agent accepts this bet.
- 2. Bet: if ¬A occurs, agent wins \$100. If ¬A doesn't occur, agent loses \$105.
  - Agent believes $P(\neg A) = .7 > 100/(100+105) = .48$, so agent accepts this bet. Oops…
- **Theorem:** An agent who holds beliefs inconsistent with axioms of probability can be convinced to accept a combination of bets that is guaranteed to lose them money

Are humans “rational bettors”?

- Humans are pretty good at estimating some probabilities, and pretty bad at estimating others.
- What might cause humans to mis-estimate the probability of an event?
- What are some of the ways in which a “rational bettor” might take advantage of humans who mis-estimate probabilities?

Outline

- Motivation: Why use probability?
  - Laziness, Ignorance, and Randomness
  - Rational Bettor Theorem
- Review of Key Concepts
  - Outcomes, Events
  - Joint, Marginal, and Conditional
  - Independence and Conditional Independence

Events

- Probabilistic statements are defined over ***events*****, or sets of world states**
  - *A = "It is raining"*
  - *B = "The weather is either cloudy or snowy"*
  - *C = "I roll two dice, and the result is 11"*
  - *D = "My car is going between 30 and 50 miles per hour"*
- **An EVENT is a SET of OUTCOMES**
  - B = { outcomes : cloudy OR snowy }
  - C = { outcome tuples (d1,d2) such that d1+d2 = 11 }
- Notation: P(A) is the probability of the set of world states (outcomes) in which proposition A holds

Kolmogorov’s axioms of probability

- For any propositions (events) A, B
  - $0 \leq P(A) \leq 1$
  - $P(True) = 1$ and $P(False) = 0$
  - $P(A \vee B) = P(A) + P(B) - P(A \wedge B)$
    - Subtraction accounts for double-counting

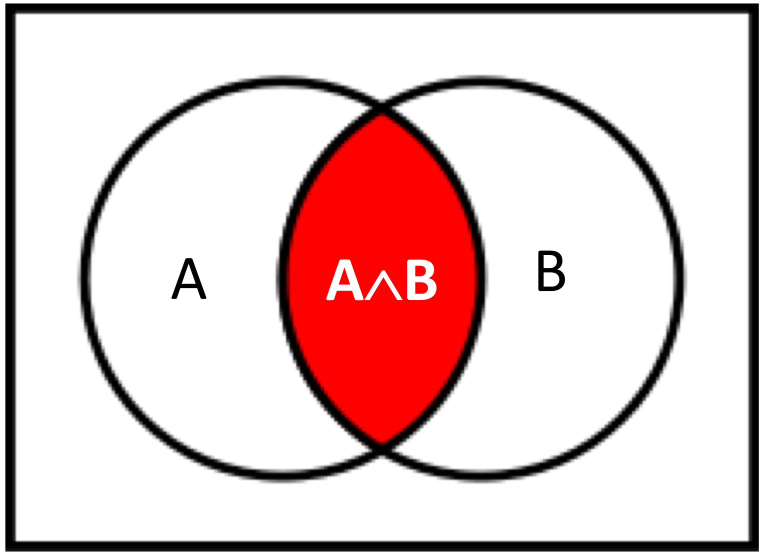

- Based on these axioms, what is $P(\neg A)$?
- These axioms are sufficient to completely specify probability theory for *discrete* random variables
  - For continuous variables, need *density functions*

Outcomes = Atomic events

- ***OUTCOME or ATOMIC EVENT:*** is **a complete specification of the state of the world**, or a complete assignment of domain values to all random variables
  - Atomic events are mutually exclusive and exhaustive

- E.g., if the world consists of only two Boolean variables *Cavity* and *Toothache*, then there are four outcomes:

  Outcome #1: $\neg Cavity \wedge \neg Toothache$
  Outcome #2: $\neg Cavity \wedge Toothache$
  Outcome #3: $Cavity \wedge \neg Toothache$
  Outcome #4: $Cavity \wedge Toothache$

Outline

- Motivation: Why use probability?
  - Laziness, Ignorance, and Randomness
  - Rational Bettor Theorem
- Review of Key Concepts
  - Outcomes, Events
  - Joint, Marginal, and Conditional
  - Independence and Conditional Independence

Joint probability distributions

- A ***joint distribution*** is an assignment of probabilities to every possible atomic event such that the probabilities sum to 1

Atomic event	**P**
$\neg$*Cavity* $\wedge$ $\neg$*Toothache*	0.8
$\neg$*Cavity* $\wedge$ *Toothache*	0.1
Cavity $\wedge$ $\neg$*Toothache*	0.05
Cavity $\wedge$ *Toothache*	0.05

  - Why does it follow from the axioms of probability that the probabilities of all possible atomic events must sum to 1?

Joint probability distributions $P(X_1, X_2, \ldots, X_N)$

- $P(X_1, X_2, \ldots, X_N)$ refers to the **probability of a particular outcome** (the outcome in which the events $X_1, X_2, \ldots,$ and X_N all occur at the same time)
- $P(X_1, X_2, \ldots, X_N)$ can also refer to the **complete TABLE**, with 2^N entries, listing the probabilities of X_1 either occurring or not occurring, X_2 either occurring or not occurring, and so on.
- This ambiguity, between the probability VALUE and the probability TABLE, will be eliminated next lecture, when we introduce random variables.

Joint probability distributions $P(X_1, X_2, \ldots, X_N)$

- Suppose we have a joint distribution of *N* random variables, each of which takes values from a domain of size *D:*
  - What is the size of the probability table?
  - Impossible to write out completely for all but the smallest distributions
  - We'll return to this when we talk about independence assumptions

Marginal distributions: from $P(X_1, \ldots, X_k, \ldots X_N)$ to $P(X_k)$

- Assume you are **given a joint distribution** (full table of outcomes) $P(X_1, \ldots, X_k, \ldots, X_N)$ and you want to compute $P(X_k)$
- By **summing** over all possible outcomes of $X_{i!=k}$ you can compute $P(X_k)$.
- This summation is called marginalization
- The resulting distribution is called a marginal probability (although it's just $P(X_k)$)

Marginal distributions: from $P(X_1, \ldots, X_k, \ldots X_N)$ to $P(X_k)$

- Assume you are **given a joint distribution** (full table of outcomes) $P(X_1, \ldots, X_k, \ldots, X_N)$ and you want to compute $P(X_k)$
- By **summing** over all possible outcomes of $X_{i!=k}$ you can compute $P(X_k)$.
- This summation is called marginalization
- The resulting distribution is called a marginal probability (although it's just $P(X_k)$)

Marginal probability distributions

- From the joint distribution p(X,Y) we can find the ***marginal distributions*** p(X) and p(Y)

P(Cavity, Toothache)	
¬Cavity ∧ *¬Toothache*	0.8
¬Cavity ∧ *Toothache*	0.1
Cavity ∧ *¬Toothache*	0.05
Cavity ∧ *Toothache*	0.05

P(Cavity)	
¬Cavity	?
Cavity	?

P(Toothache)	
¬Toothache	?
Toochache	?

Joint -> Marginal by adding the outcomes

- From the joint distribution p(X,Y) we can find the ***marginal distributions*** p(X) and p(Y)
- To find p(X = x), sum the probabilities of all atomic events where X = x:

$$
\begin{aligned}
P(X=1) = {} & P(X=1, Y=1) \\
& + P(X=1, Y=2) \\
& + P(X=1, Y=3) \\
& + \cdots
\end{aligned}
$$

- This is called ***marginalization*** (we are *marginalizing out* all the variables except X)

Conditional distributions $P(X_k \mid X_i)$

- The **conditional probability** of event X_k, given event X_j, is the probability that X_k has occurred if you already know/assume that X_j has occurred.
- The conditional distribution is written $P(X_k \mid X_j)$.
- The probability that both X_j and X_k occurred was, originally, $P(X_j, X_k)$.
- But now you know/assume that X_j has occurred.
  So all of the other events are no longer possible.
  - Other events: probability used to be $P(\neg X_j)$, but now their probability is 0.
  - Events in which X_j occurred: probability used to be $P(X_j)$, but now their probability is 1.
- So we need to renormalize: the probability that both X_j and X_k occurred, GIVEN that X_j has occurred, is $P(X_k \mid X_j)=P(X_j, X_k)/P(X_j)$.

Conditional Probability: renormalize (divide)

- Probability of cavity given toothache:

  P(*Cavity = true* | *Toothache = true*)

- For any two events A and B, $P(A \mid B) = \dfrac{P(A \wedge B)}{P(B)} = \dfrac{P(A,B)}{P(B)}$

If we don't condition on A, the set of all possible events is this rectangle, so the whole rectangle has probability=1.

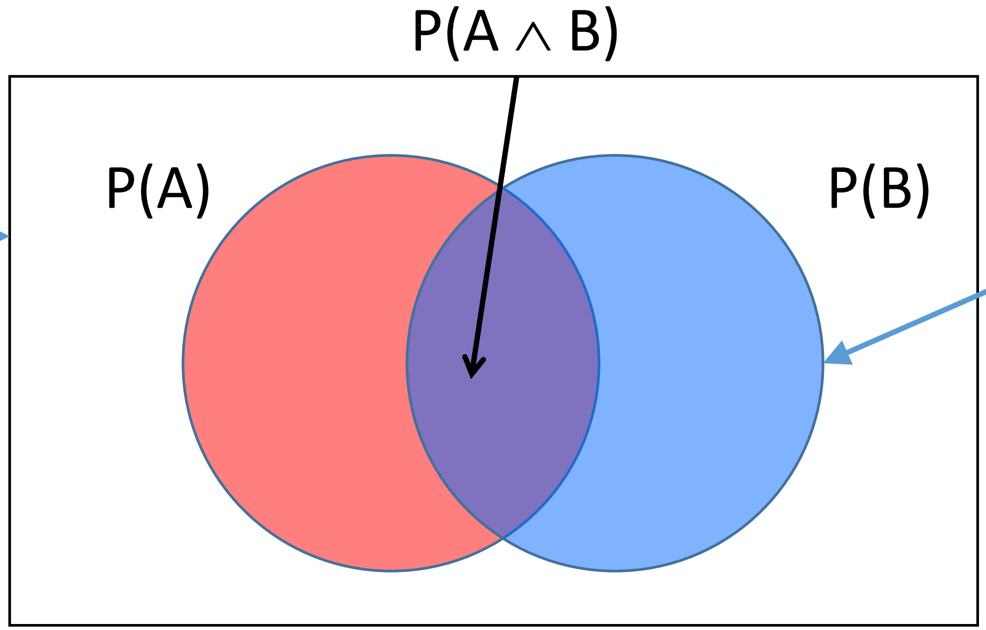

If we know/assume that B has occurred, the set of all possible events becomes restricted to the set of events in which B occurred. So we renormalize to make the area of this circle = 1.

Conditional probability

P(Cavity, Toothache)	
$\neg Cavity \wedge \neg Toothache$	0.8
$\neg Cavity \wedge Toothache$	0.1
$Cavity \wedge \neg Toothache$	0.05
$Cavity \wedge Toothache$	0.05

P(Cavity)	
$\neg Cavity$	0.9
$Cavity$	0.1

P(Toothache)	
$\neg Toothache$	0.85
$Toothache$	0.15

- What is $p(Cavity = true \mid Toothache = false)$?

  $p(Cavity \mid \neg Toothache) = 0.05/0.85 = 1/17$

- What is $p(Cavity = false \mid Toothache = true)$?

  $p(\neg Cavity \mid Toothache) = 0.1/0.15 = 2/3$

Conditional distributions

- A conditional distribution is a distribution over the values of one variable given fixed values of other variables

P(Cavity, Toothache)	
¬Cavity ∧ ¬Toothache	0.8
¬Cavity ∧ Toothache	0.1
Cavity ∧ ¬Toothache	0.05
Cavity ∧ Toothache	0.05

P(Cavity | Toothache = true)	
¬Cavity	0.667
Cavity	0.333

P(Cavity|Toothache = false)	
¬Cavity	0.941
Cavity	0.059

P(Toothache | Cavity = true)	
¬Toothache	0.5
Toochache	0.5

P(Toothache | Cavity = false)	
¬Toothache	0.889
Toochache	0.111

Normalization trick

- To get the whole conditional distribution p(X | Y = y) at once, select all entries in the joint distribution table matching Y = y and renormalize them to sum to one

P(Cavity, Toothache)	
¬Cavity ∧ ¬Toothache	0.8
¬Cavity ∧ Toothache	0.1
Cavity ∧ ¬Toothache	0.05
Cavity ∧ Toothache	0.05

Select

Toothache, Cavity = false	
¬Toothache	0.8
Toochache	0.1

Renormalize

P(Toothache | Cavity = false)	
¬Toothache	0.889
Toochache	0.111

Normalization trick

- To get the whole conditional distribution p(X | Y = y) at once, select all entries in the joint distribution table matching Y = y and renormalize them to sum to one
- Why does it work?

$$P(x|y) = \frac{P(x,y)}{\sum_{x'} P(x',y)} = \frac{P(x,y)}{P(y)}$$

by marginalization

Product rule and chain rule

- Definition of conditional probability:

$$P(A \mid B) = \frac{P(A,B)}{P(B)}$$

- We can also obtain the joint from the conditional probability

$$P(A,B) = P(A \mid B)P(B) = P(B \mid A)P(A)$$

- More generally (the chain rule):

$$P(A_1,\ldots,A_n) = P(A_1)P(A_2 \mid A_1)P(A_3 \mid A_1,A_2)\ldots P(A_n \mid A_1,\ldots,A_{n-1})$$

$$= \prod_{i=1}^{n} P(A_i \mid A_1,\ldots,A_{i-1})$$

Product Rule Example: The Birthday problem

- We have a set of *n* people. What is the probability that two of them share the same birthday?
- Easier to calculate the probability that *n* people *do not* share the same birthday

$$P(B_1, \ldots, B_n \text{ distinct})$$
$$= P(B_1, B_2 \text{ distinct}) P(B_1, B_2, B_3 \text{distinct} | B_1, B_2 \text{ distinct}) \ldots$$
$$P(B_1, B_2, \ldots B_n \text{ distinct} | B_1, \ldots B_{n-1} \text{ distinct})$$

$$P(B_1, \ldots, B_n \text{ distinct}) = \left(\frac{364}{365}\right)\left(\frac{363}{365}\right) \ldots \left(\frac{365-n+1}{365}\right)$$

The Birthday problem

- For 23 people, the probability of sharing a birthday is above 0.5!

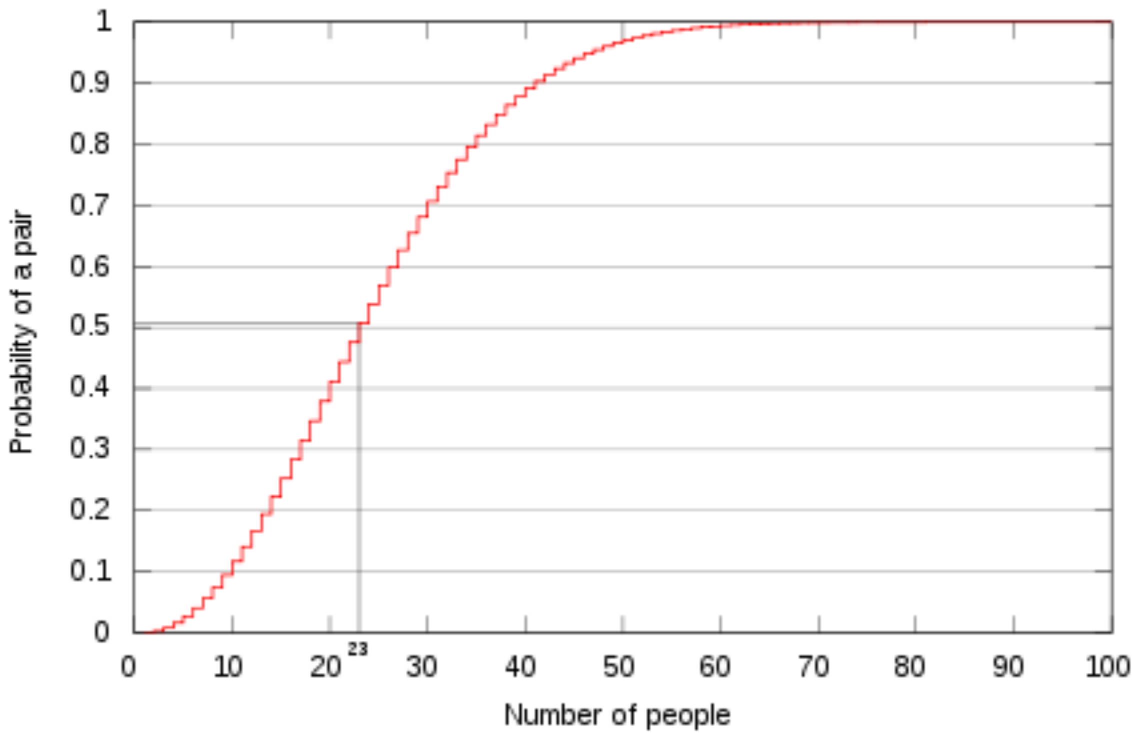

http://en.wikipedia.org/wiki/Birthday_problem

Outline

- Motivation: Why use probability?
  - Laziness, Ignorance, and Randomness
  - Rational Bettor Theorem
- Review of Key Concepts
  - Outcomes, Events, and Random Variables
  - Joint, Marginal, and Conditional
  - Independence and Conditional Independence

Independence ≠ Mutually Exclusive

- Two events A and B are *independent* if and only if $p(A \wedge B) = p(A, B) = p(A)\,p(B)$
  - In other words, $p(A \mid B) = p(A)$ and $p(B \mid A) = p(B)$
- We often make independence assumptions when designing models.
  - e.g., *Toothache* and *Weather* may be assumed to be independent
- Are two *mutually exclusive* events independent?
  - No! Quite the opposite! If A and B are mutually exclusive, and you know A happened, then you know that B *didn't* happen!! Also,
    $p(A \vee B) = p(A) + p(B)$ only if A, B mutually exclusive

Independence ≠ Conditional Independence

- Two events A and B are *independent* if and only if $p(A \wedge B) = p(A)\, p(B)$
  - In other words, $p(A \mid B) = p(A)$ and $p(B \mid A) = p(B)$
- **Conditional independence**: A and B are *conditionally independent* given C iff $p(A \wedge B \mid C) = p(A \mid C)\, p(B \mid C)$
  - Equivalent: $p(A \mid B, C) = p(A \mid C)$
  - Equivalent: $p(B \mid A, C) = p(B \mid C)$

Independence ≠ Conditional Independence

Cavity: Boolean variable indicating whether the patient has a cavity

Toothache: Boolean variable indicating whether the patient has a toothache

By William Brassey Hole(Died:1917)

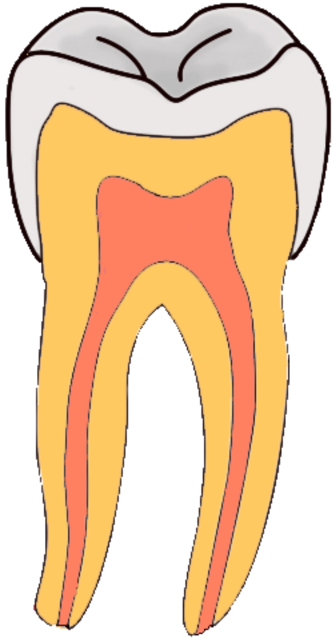

By Aduran, CC-SA 3.0

Catch: whether the dentist's probe catches in the cavity

By Dozenist, CC-SA 3.0

These Events are not Independent

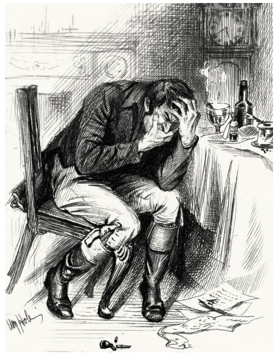

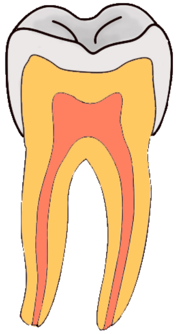

- If the patient has a toothache, then it's likely he has a cavity.
  Having a cavity makes it more likely that the probe will catch on something.

$$P(Catch \mid Toothache) > P(Catch)$$

- If the probe catches on something, then it's likely that the patient has a cavity. If he has a cavity, then he might also have a toothache.

$$P(Toothache \mid Catch) > P(Toothache)$$

- So Catch and Toothache are not independent

…but they are Conditionally Independent

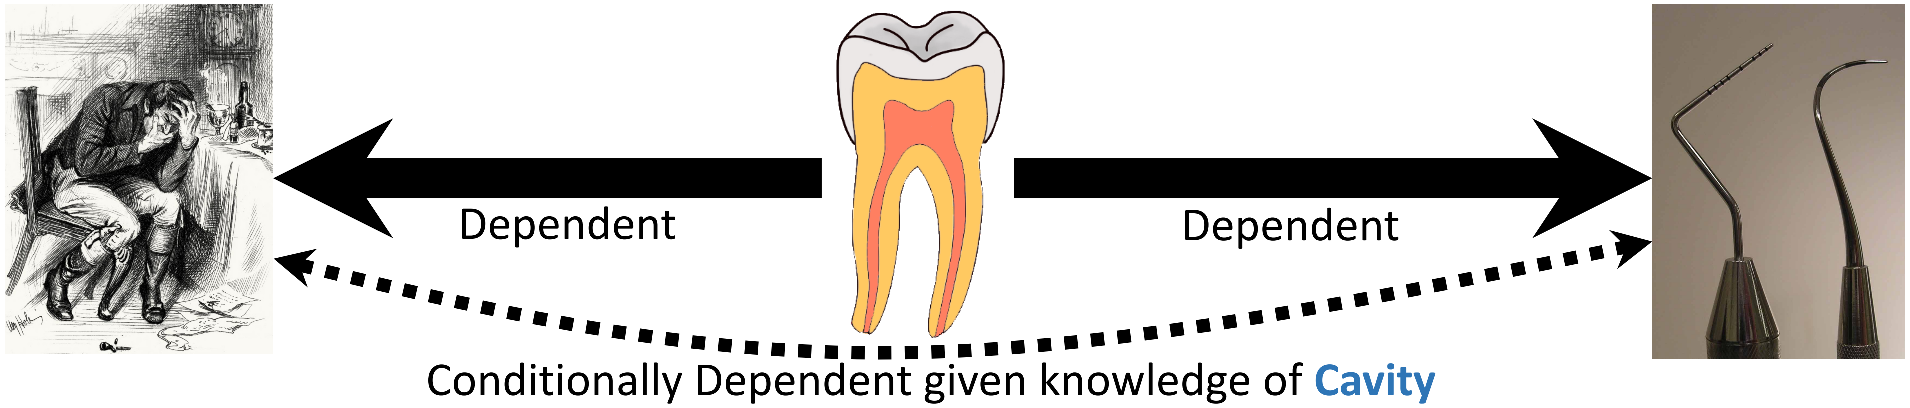

- Here are some reasons the probe might not catch, despite having a cavity:
  - The dentist might be really careless
  - The cavity might be really small
- Those reasons have nothing to do with the toothache!

$$P(Catch|Cavity, Toothache) = P(Catch|Cavity)$$

- **Catch** and **Toothache** are **conditionally independent** given knowledge of **Cavity**

…but they are Conditionally Independent

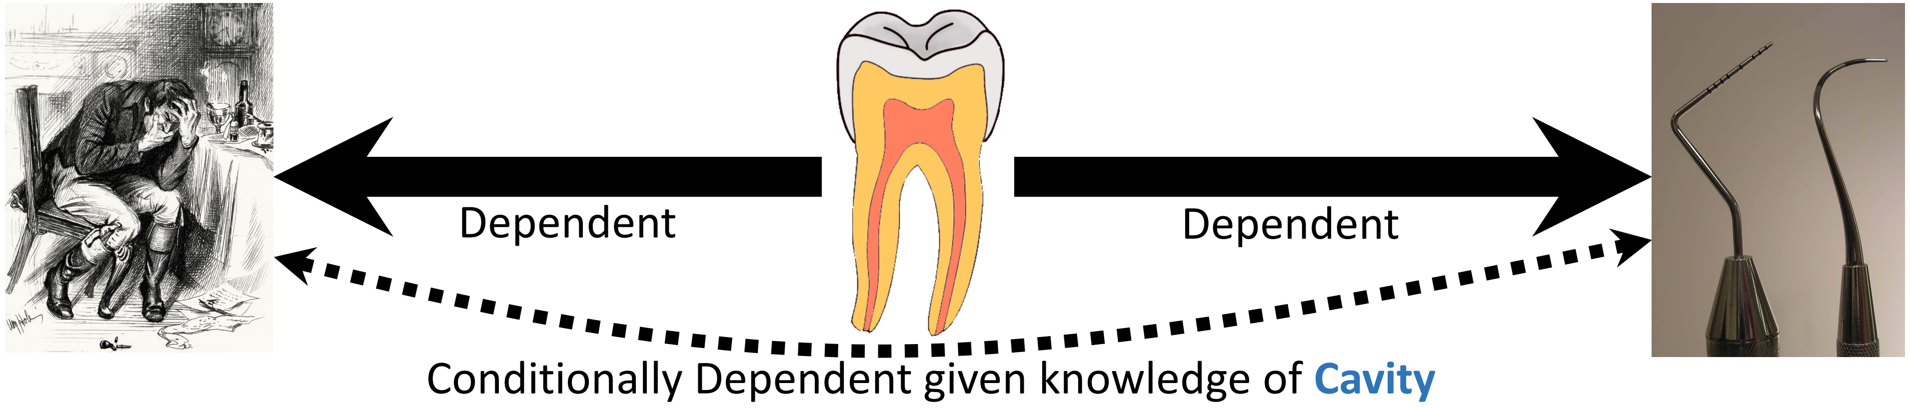

These statements are all equivalent:

$$P(Catch|Cavity, Toothache) = P(Catch|Cavity)$$

$$P(Toothache|Cavity, Catch) = P(Toothache|Cavity)$$

$$P(Toothache, Catch|Cavity) = P(Toothache|Cavity)\, P(Catch|Cavity)$$

…and they all mean that **Catch** and **Toothache** are **conditionally independent** given knowledge of **Cavity**

Outline

- Motivation: Why use probability?
  - Laziness, Ignorance, and Randomness
  - Rational Bettor Theorem
- Review of Key Concepts
  - Outcomes, Events
  - Random Variables; probability mass function (pmf)
  - Jointly random variables: Joint, Marginal, and Conditional pmf
  - Independent vs. Conditionally Independent events