

# Review of terms and Bayes Rule

- $P(A)$  is the probability that  $A$  is true.
- $P(\neg A)$  is the probability that  $A$  is false.

$$P(A) + P(\neg A) = 1$$

- These probabilities are termed *prior probabilities* or *unconditional probabilities*
  - they don't depend on any other factors.

# OR and AND

- $P(A \vee B)$  is the unconditional probability that either A or B is true.
- $P(A \wedge B)$  is the unconditional probability that both A and B are true.

$$P(A \vee B) = P(A) + P(B) - P(A \wedge B)$$

- Joint Probabilities: The collection of the probabilities for the possible values of all statements (variables).

# Independence

- When two variables have no relationship.

$$P(A \wedge B) = P(A)P(B)$$

- Knowing about the value of B tells us nothing about the value of A
  - and vice-versa.
- *The probability that Dave will burp is independent of the background color on Joe's laptop.*

# Joint Probabilities Example

	<i>toothache</i>		$\neg$ <i>toothache</i>	
	<i>catch</i>	$\neg$ <i>catch</i>	<i>catch</i>	$\neg$ <i>catch</i>
<i>cavity</i>	.108	.012	.072	.008
$\neg$ <i>cavity</i>	.016	.064	.144	.576

- *toothache*: patient has a toothache
- *cavity*: patient has a cavity
- *catch*: probe *catches* in a hole in a tooth.

# Conditional Probability

- $P(A|B)$  is the probability that A is true given that B is true.
  - $P(A|B) = P(A \wedge B) / P(B)$
  - Can be derived from Joint Probabilities if we have them.

	<i>toothache</i>		$\neg$ <i>toothache</i>	
	<i>catch</i>	$\neg$ <i>catch</i>	<i>catch</i>	$\neg$ <i>catch</i>
<i>cavity</i>	.108	.012	.072	.008
$\neg$ <i>cavity</i>	.016	.064	.144	.576

$P(\text{cavity}|\text{toothache})$  can be derived from the highlighted areas:

$$(.108 + .012) / (.108 + .012 + .016 + .064)$$

# Conditional Probability: So What?

- Inference: deciding on the truth of a statement based on what we currently know.
  - Deductive Inference: based on solid, absolute rules of reasoning (never wrong).
    - *I know that Dave likes cookies, and I know that Oreos are cookies, therefore...*
  - Inductive Inference: based on observations. Could be right or wrong.
    - *I know that 82% of AI tests ask a question about Bayes Rule, so I infer there will be one on the test.*

# Bayes Rule

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

- useful when we know one conditional probability and need the other.
- We rarely have full joint probabilities!
  - we need tools to be able to compute probabilities based on the data we do have.

# Conditional Independence

- If two variables  $A, B$  are *independent* given any specific value of another variable  $C$ , we know:

$$P(A \wedge B|C) = P(A|C)P(B|C)$$

- This is *conditional independence*

$$P(\text{toothache} \wedge \text{catch}|\text{cavity}) = P(\text{toothache}|\text{cavity})P(\text{catch}|\text{cavity})$$

toothache and catch are independent given the presence or absence of a cavity (neither causes the other).

# Probability Distributions

- We often want to talk about (or compute with) an entire probability distribution (over some variable or variables).
  - For binary variables this means only two values.
  - Use uppercase variable name:  $P(\text{Cavity}) = \langle .2, .8 \rangle$
- We can refer to  $P(\text{Cavity}, \text{Toothache})$  as the full set of all probabilities of the combinations of cavity and toothache (4 values total).

# Conditional Independence of Distributions

$$P(X, Y|Z) = P(X|Z) P(Y|Z)$$

$$P(\text{Toothache}, \text{Catch}|\text{Cavity}) = P(\text{Toothache}|\text{Cavity}) P(\text{Catch}|\text{Cavity})$$

- This is stronger than the previous version

$$P(\text{toothache} \wedge \text{catch}|\text{cavity}) = P(\text{toothache}|\text{cavity}) P(\text{catch}|\text{cavity})$$

- This new version covers all combinations of values for Toothache and Catch.

# Importance of Conditional Independence

- Consider a general case:
  - We have 1 cause and  $n$  *effects*
    - Cavity causes toothache and catch.
  - The full joint probabilities will require  $2^{n+1}$  values
  - If  $n$  of the variables are conditionally independent, we can reduce this to  $n$  tables of 2 values each.
- Conditional Independence is sometimes assumed (even when it's not quite true), since this makes it feasible to deal with large numbers of variables.

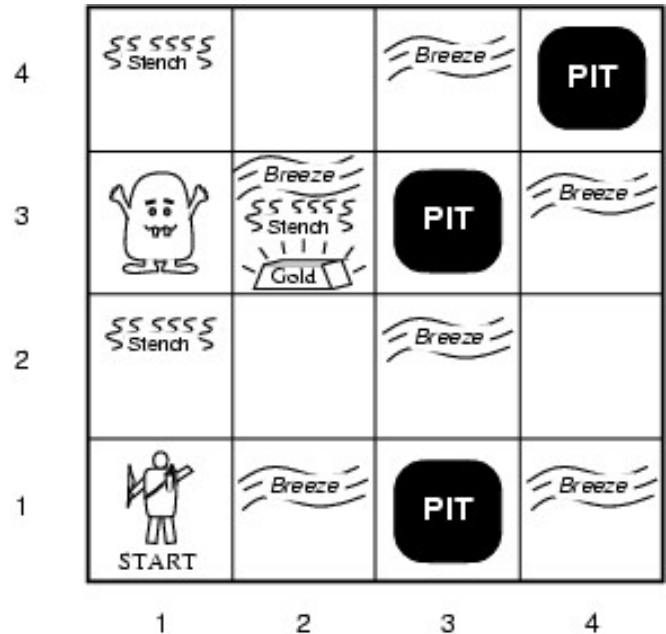
# Computing the Full Joint Distribution assuming 1 cause, n effects

$$P(\text{Cause}, \text{Effect}_1, \dots, \text{Effect}_n) = P(\text{Cause}) \prod_i P(\text{Effect}_i | \text{Cause})$$

- This distribution is known as a naive *Bayes* model
  - Only holds if the effects are conditionally independent.
  -

# Application: Wumpus World

- Relatively simple virtual world.
  - 4x4 grid of rooms.
  - Monster: The Wumpus
  - Danger: Bottomless Pits
  - Action: shoot an arrow
  - Treasure: Gold.
- We can look at using probabilities to guide an Agent through this environment.
- First we need to define the game (PEAS)



# Performance Environment Actuators Sensors

- Points are awarded and/or deducted:
  - find gold: +1000
  - death by Wumpus: -1000
  - death by endless pit: -1000
  - each action: -1
  - picking up the arrow: -10 points

# P<sub>erformance</sub> Environment A<sub>ctuators</sub> S<sub>ensors</sub>

- 4x4 grid.
- Agent starts at 1,1 (lower left).
- Agent starts facing to the right.
- The gold and the Wumpus are placed at random locations in the grid (can't be in start room).
- Each room other than starting room can be a bottomless pit with probability 0.2

# P<sub>erformance</sub> E<sub>nvironment</sub> Actuators S<sub>ensors</sub>

- Turn left 90°
- Turn right 90°
- Move forward into room straight ahead.
  - blocked by walls.
  - Eaten by a live Wumpus immediately.
  - Fall into pit immediately.
- Grab an object in the current room
- Shoot arrow in straight line (only allowed once).

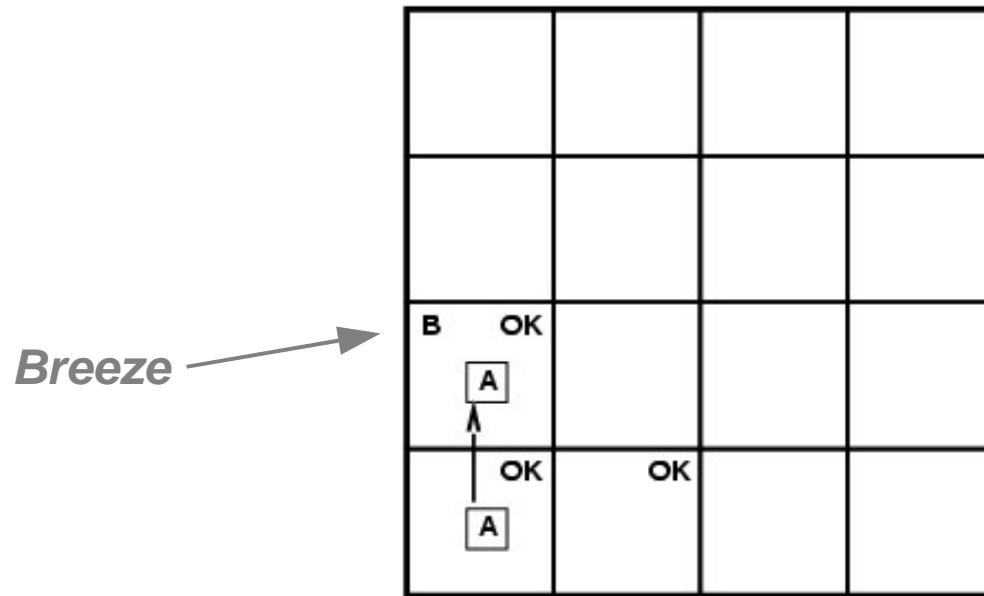
# P<sub>erformance</sub> E<sub>nvironment</sub> A<sub>ctuators</sub> Sensors

- Agent can smell, detect a breeze, see gold glitter, detect a bump (into a wall) and hear the wumpus scream when it is killed.
- Sensors:
  - *Stench*: a neighboring room holds the wumpus.
  - *Breeze*: a neighboring room holds a pit.
  - *Glitter*: The room the agent is in has some gold.
  - *Bump*: The agent just banged into a wall.
  - *Scream*: The Wumpus just died.


# Exploring a wumpus world

OK			
OK A	OK		


# Exploring a wumpus world




# Exploring a wumpus world




# Exploring a wumpus world




# Exploring a wumpus world




# Exploring a wumpus world



# Exploring a wumpus world



# Exploring a wumpus world



# Inference in the Wumpus World

- Deductive reasoning can easily get stuck
  - Only moves to *provably safe* states (rooms).
  - We will explore this when we discuss logic.
- Probabilistic reasoning can do better in this type of environment.
  - Which rooms are *likely* to be safe?
  - Which room is likely to provide the maximum payoff?

# Wumpus probabilistic inference.

- We can get the general idea by looking at just one of the many probabilities that a Wumpus agent might want to determine.
- We want to compute the probability that a room holds a bottomless pit.
  - Initially we know the probability  $P(P_{i,j})$  that room  $i,j$  holds a pit is: 0.2 (other than the start room).
  - We can adjust these probabilities as rooms are explored (and make some inferences).

# Wumpus Variables

- $P_{i,j}$  represents room  $i,j$  and is true if the room contains a pit.
- $B_{i,j}$  is true if a breeze is detected in room  $i,j$
- We assume *conditional independence* of each  $B_{i,j}$ 
  - the pits are the causes, the breezes the effects

# Wumpus joint probabilities

- Consider the situation after the agent has visited room  $_{1,2}$  and room  $_{2,1}$  and found a breeze in each.

- The full joint distribution is given by

$$P(P_{1,1}, \dots, P_{4,4}, B_{1,1}, B_{1,2}, B_{2,1})$$

which is:

$$P(B_{1,1}, B_{1,2}, B_{2,1} | P_{1,1}, \dots, P_{4,4}) P(P_{1,1}, \dots, P_{4,4})$$

Conditional probability of a breeze configuration given a pit configuration. Each of these is either 1 or 0.

Probability of a pit configuration. Assuming we have  $k$  pits, this is

$$0.2^k * 0.8^{16-k}$$

1,4	2,4	3,4	4,4
1,3	2,3	3,3	4,3
1,2 B OK	2,2	3,2	4,2
1,1 OK	2,1 B OK	3,1	4,1

# General Solution to $P_{1,3}$

- Assume we want to calculate the probability that room<sub>1,3</sub> has a pit.
  - $P(P_{1,3} | \text{everything we know})$ 
    - We know rooms 1,2 and 2,1 have a breeze
    - We know rooms 1,1 1,2 and 1,3 have no pits.
    - We need to sum over all possible configurations that are consistent with these know facts.
    - In general this will be a large computation...

# Computing $P_{1,3}$


$$P(P_{1,3} | \neg P_{1,1} \wedge \neg P_{1,2} \wedge \neg P_{2,1} \wedge \neg B_{1,1} \wedge B_{1,2} \wedge B_{2,1}) =$$

$$\sum_{\text{unknown}} \frac{P_{1,3} \wedge \neg P_{1,1} \wedge \neg P_{1,2} \wedge \neg P_{2,1} \wedge \neg B_{1,1} \wedge B_{1,2} \wedge B_{2,1}, \text{ unknown}}{\neg P_{1,1} \wedge \neg P_{1,2} \wedge \neg P_{2,1} \wedge \neg B_{1,1} \wedge B_{1,2} \wedge B_{2,1}}$$

We sum over all possible configurations of the unknown variables.

# Using conditional independence

- We can simplify the computations by assuming that observed breezes are conditionally independent of *other*, given *known*, *fringe* and *query*.
- In other words, the probability of a pit in room 1,3 is not dependent on the probability that there is a breeze in any of the rooms that are not adjacent to the rooms we have explored.



# Computing $P_{1,3}$

- The book has the complete derivation...

$$k P(P1,3) \sum_{fringe} P(\neg B_{1,1} \wedge B_{1,2} \wedge B_{2,1} | \neg P_{1,1} \wedge \neg P_{1,2} \wedge \neg P_{2,1} \wedge P1,3, fringe) P(fringe)$$

$$k = \frac{P(\neg P_{1,1} \wedge \neg P_{1,2} \wedge \neg P_{2,1} \wedge \neg B_{1,1} \wedge B_{1,2} \wedge B_{2,1})}{P(\neg P_{1,1} \wedge \neg P_{1,2} \wedge \neg P_{2,1})}$$

# Probability Quiz on Thursday

- What is the probability that there will be a probability quiz on Thursday?
  - Is this independent of the number of people who have given up  $\frac{1}{2}$  hour ago?
- Do we really need to understand all the math in the book to be able to use probabilities?
  - No, although if you did this seriously you would end up deriving all of this yourself (if you did a good job).
  - Using conditional independence simplifies computation significantly and might not be obvious.