22c:145 Artificial Intelligence

Uncertainty

- Reading: Ch 13. Russell & Norvig

Problem of Logic Agents

- Logic-agents almost never have access to the whole truth about their environments.
- A rational agent is one that makes rational decisions in order to maximize its performance measure.
- Logic-agents may have to either risk falsehood or make weak decisions in uncertain situation
- A rational agent's decision depends on relative importance of goals, likelihood of achieving them.
- Probability theory provides a quantitative way of encoding likelihood

Foundations of Probability

- Probability Theory makes the same ontological commitments as FOL
- Every sentence S is either true or false.
- The *degree of belief, or probability,* that S is true is a *number* P between 0 and 1.
- P(S) = 1 iff S is certainly true
- P(S) = 0 iff S is certainly false
- P(S) = 0.4 iff S is true with a 40% chance
- P(not A) = probability that A is false
- P(A and B) = probability that both A and B are true
- P(A or B) = probability that either A or B (or both) are true

Axioms of Probability

- All probabilities are between 0 and 1
- Valid propositions have probability 1. Unsatisfiable propositions have probability 0. That is,
  - P(A v : A) = P(true) = 1
  - P(A Æ : A) = P(false) = 0
  - P(: A) = 1 – P(A)
- The probability of disjunction is defined as follows.
  - P(A v B) = P(A) + P(B) – P(A Æ B)
  - P(A Æ B) = P(A) + P(B) – P(A v B)

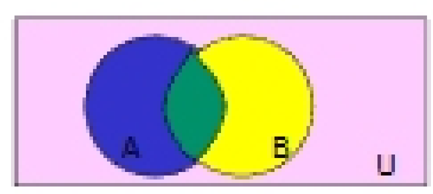

Exercise Problem I

Prove that

- P(A v B v C) =
  P(A) + P(B) + P(C) –
  P(A Æ B) – P(A Æ C) – P(B Æ C) +
  P(A Æ B Æ C)

How to Decide Values of Probability

P(the sun comes up tomorrow) = 0.999

- Frequentist
  - Probability is inherent in the process
  - Probability is estimated from measurements

Probs can be wrong!

A Question

Jane is from Berkeley. She was active in anti-war protests in the 60's. She lives in a commune.

- Which is more probable?
  1. Jane is a bank teller
  2. Jane is a feminist bank teller

A Question

Jane is from Berkeley. She was active in anti-war protests in the 60's. She lives in a commune.

- Which is more probable?
  1. Jane is a bank teller
  2. Jane is a feminist bank teller

1. A
2. A Æ B

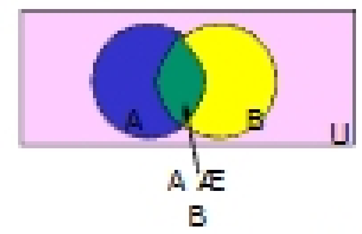

Conditional Probability

- P(A) is the unconditional (or prior) probability
- An agent can use unconditional probability of A to reason about A only in the absence of no further information.
- If some further evidence B becomes available, the agent must use the conditional (or posterior) probability:

  P(A|B)

  the probability of A given that the agent already knew that B is true.
- P(A) can be thought as the conditional probability of A with respect to the empty evidence:

  P(A) = P(A|).

Conditional Probability

1. P(Blonde) =
2. P(Blonde | Swedish) =
3. P(Blonde | Kenian) =
4. P(Blonde | Kenian Æ : EuroDescent) =

- If we know nothing about a person, the probability that he/she is blonde equals a certain value, say 0.1.
- If we know that a person is Swedish the probability that s/he is blonde is much higher, say 0.9.
- If we know that the person is Kenyan, the probability s/he is blonde much lower, say 0.000003.
- If we know that the person is Kenyan and not of European descent, the probability s/he is blonde is basically 0.
- **Computation:** P(A | B) = P(A Æ B)/P(B)

Random Variables

Variable	Domain
Age	{ 1, 2, …, 120 }
Weather	{ sunny, dry, cloudy, raining }
Size	{ small, medium, large }
Raining	{ true, false }

- The probability that a random variable X has value val is written as P(X=val)
- P: domain ! [0, 1]
  - Sums to 1 over the domain:
    - P(Raining = true) = P(Raining) = 0.2
    - P(Raining = false) = P(: Raining) = 0.8

Probability Distribution

- If X is a random variable, we use the **bold case P(X)** to denote a vector of values for the probabilites of each individual element that **X** can take.
- Example:
  - P(Weather = sunny) = 0.6
  - P(Weather = rain) = 0.2
  - P(Weather = cloudy) = 0.18
  - P(Weather = snow) = 0.02
- Then **P(**Weather**)** = <0.6, 0.2, 0.18, 0.02> (the value order of "sunny", "rain", "cloudy", "snow" is assumed).
- **P(**Weather**)** is called a probability distribution for the random variable Weather.

- Joint distribution: **P(**X1, X2, …, Xn**)**
  - Probability assignment to all combinations of values of random variables

Joint Distribution Example

	Toothache	:Toothache
Cavity	0.04	0.06
: Cavity	0.01	0.89

- The sum of the entries in this table has to be 1

Joint Distribution Example

	Toothache	:Toothache
Cavity	0.04	0.06
: Cavity	0.01	0.89

- The sum of the entries in this table has to be 1
- Given this table, one can answer all the probability questions about this domain
- P(cavity) = 0.1 [add elements of cavity row]
- P(toothache) = 0.05 [add elements of toothache column]

Joint Distribution Example

	Toothache	:Toothache
Cavity	0.04	0.06
: Cavity	0.01	0.89

- The sum of the entries in this table has to be 1
- Given this table, one can answer all the probability questions about this domain
- P(cavity) = 0.1 [add elements of cavity row]
- P(toothache) = 0.05 [add elements of toothache column]
- P(A | B) = P(A Æ B)/P(B) [prob of A when U is limited to B]

Joint Distribution Example

	Toothache	:Toothache
Cavity	0.04	0.06
: Cavity	0.01	0.89

- The sum of the entries in this table has to be 1
- Given this table, one can answer all the probability questions about this domain
- P(cavity) = 0.1 [add elements of cavity row]
- P(toothache) = 0.05 [add elements of toothache column]
- P(A | B) = P(A Æ B)/P(B) [prob of A when U is limited to B]
- P(cavity | toothache) = 0.04/0.05 = 0.8

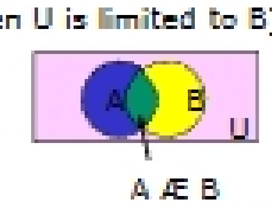

Joint Probability Distribution (JPD)

- A joint probability distribution **P**(X1, X2 …, Xn) provides complete information about the probabilities of its random variables.
- However, JPD's are often hard to create (again because of incomplete knowledge of the domain).
- Even when available, JPD tables are very expensive, or impossible, to store because of their size.
- A JPD table for n random variables, each ranging over k distinct values, has k^n entries!
- A better approach is to come up with conditional probabilities as needed and compute the others from them.

Bayes' Rule

- Bayes' Rule
  - P(A | B) = P(B | A) P(A) / P(B)
- What is the probability that a patient has meningitis (M) given that he has a stiff neck (S)?
  - P(M|S) = P(S|M) P(M)/P(S)

P(S|M) is easier to estimate than P(M|S) because it refers to causal knowledge:

- meningitis typically causes stiff neck.
- P(S|M) can be estimated from past medical cases and the knowledge about how meningitis works.
- Similarly, P(M), P(S) can be estimated from statistical information.