


Uncertainty

Russell and Norvig: Chapter 13
CMCS424 Fall 2003

based on material from Jean-Claude
Latombe, Daphne Koller and Nir Friedman

Uncertain Agent



An Old Problem ...



Types of Uncertainty

- ◆ **Uncertainty in prior knowledge**
E.g., some causes of a disease are unknown and are not represented in the background knowledge of a medical-assistant agent


Types of Uncertainty

For example, to drive my car in the morning:

- It must not have been stolen during the night
- It must not have flat tires
- There must be gas in the tank
- The battery must not be dead
- The ignition must work
- I must not have lost the car keys
- No truck should obstruct the driveway
- I must not have suddenly become blind or paralytic
- Etc...

Not only would it not be possible to list all of them, but would trying to do so be efficient?

Types of Uncertainty



- ◆ **Uncertainty in perception**
E.g., sensors do not return exact or complete information about the world; a robot never knows exactly its position

Types of Uncertainty

- ◆ Uncertainty in prior knowledge
E.g., so not rep medical and are ge of a
- ◆ Uncertainty
E.g., ad of preconditions, while these lists are in fact arbitrary long
- ◆ Uncertainty in perception

Sources of uncertainty:
1. Ignorance
2. Laziness (efficiency?)

What we call **uncertainty** is a summary of all that is not explicitly taken into account in the agent's KB

Questions

- ◆ How to represent uncertainty in knowledge?
- ◆ How to perform inferences with uncertain knowledge?
- ◆ Which action to choose under uncertainty?

How do we deal with uncertainty?

- ◆ Implicit:
 - Ignore what you are uncertain of when you can
 - Build procedures that are robust to uncertainty
- ◆ Explicit:
 - Build a model of the world that describe uncertainty about its state, dynamics, and observations
 - Reason about the effect of actions given the model

Handling Uncertainty

Approaches:

1. Default reasoning
2. Worst-case reasoning
3. Probabilistic reasoning

Default Reasoning

- ◆ Creed: The world is fairly normal. Abnormalities are rare
- ◆ So, an agent assumes normality, until there is evidence of the contrary
- ◆ E.g., if an agent sees a bird x , it assumes that x can fly, unless it has evidence that x is a penguin, an ostrich, a dead bird, a bird with broken wings, ...

Representation in Logic

- ◆ $BIRD(x) \wedge \neg AB_F(x) \Rightarrow FLIES(x)$
- ◆ Very active research field in the 80's
- ◆ $\rightarrow$ Non-monotonic logics: defaults, circumscription, closed-world assumptions
- ◆ Applications to databases
- ◆ ...

Default rule: Unless $AB_F(Tweety)$ can be proven True, assume it is False

But what to do if several defaults are contradictory?
Which ones to keep? Which one to reject?

Worst-Case Reasoning

- ◆ Creed: Just the opposite of Murphy's Law
- ◆ Uncertainty is defined by possible outcomes of an action or possible positions of a robot
- ◆ The agent assumes the worst case and chooses the actions that maximize the function in this case
- ◆ Example: Adversarial search



Probabilistic Reasoning

- ◆ Creed: The world is not divided between "normal" and "abnormal", nor is it adversarial. Possible situations have various likelihoods (probabilities)
- ◆ The agent has probabilistic beliefs – pieces of knowledge with associated probabilities (strengths) – and chooses its actions to maximize the expected value of some utility function

How do we represent Uncertainty?

We need to answer several questions:

- ◆ What do we represent & how we represent it?
 - What language do we use to represent our uncertainty? What are the semantics of our representation?
- ◆ What can we do with the representations?
 - What queries can be answered? How do we answer them?
- ◆ How do we construct a representation?
 - Can we ask an expert? Can we learn from data?

Probability

- ◆ A well-known and well-understood framework for uncertainty
- ◆ Clear semantics
- ◆ Provides principled answers for:
 - Combining evidence
 - Predictive & Diagnostic reasoning
 - Incorporation of new evidence
- ◆ Intuitive (at some level) to human experts
- ◆ Can be learned

Notion of Probability

You drive on Rt 1 to UM of the times there is a traffic light at PaintBranch & Rt 1. The next time you plan to drive there, you expect to see a traffic light with probability 0.7.

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

$$P(\text{True}) = P(A) + P(\neg A) - P(\text{False})$$

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

So:

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

- ◆ The probability of a proposition A is a number P(A) between 0 and 1

- ◆ $P(\text{True}) = 1$ and $P(\text{False}) = 0$
- ◆ $P(A \vee B) = P(A) + P(B) - P(A \wedge B)$

Axioms of probability

Frequency Interpretation

- ◆ Draw a ball from a urn containing n balls of the same size, r red and s yellow.
- ◆ The probability that the proposition A = "the ball is red" is true corresponds to the relative frequency with which we expect to draw a red ball $\rightarrow P(A) = ?$

Subjective Interpretation

There are many situations in which there is no objective frequency interpretation:

- On a windy day, just before paragliding from the top of El Capitan, you say "there is probability 0.05 that I am going to die"
- You have worked hard on your AI class and you believe that the probability that you will get an A is 0.9

Bayesian Viewpoint

- ◆ probability is "degree-of-belief", or "degree-of-uncertainty".
- ◆ To the Bayesian, probability lies subjectively in the mind, and can--with validity--be different for people with different information
- ◆ e.g., the probability that Bush will be reelected in 2004.
- ◆ In contrast, to the frequentist, probability lies objectively in the external world.
- ◆ The Bayesian viewpoint has been gaining popularity in the past decade, largely due to the increase computational power that makes many of the calculations that were previously intractable, feasible.

Random Variables

- ◆ A proposition that takes the value True with probability p and False with probability $1-p$ is a **random variable** with distribution $(p, 1-p)$
- ◆ If a urn contains balls having 3 possible colors – red, yellow, and blue – the color of a ball picked at random from the bag is a random variable with 3 possible values
- ◆ The **(probability) distribution** of a random variable X with n values $x_1, x_2, \dots, x_n$ is: $(p_1, p_2, \dots, p_n)$ with $P(X=x_i) = p_i$ and $\sum_{i=1, \dots, n} p_i = 1$

Joint Distribution

- ◆ k random variables $X_1, \dots, X_k$
- ◆ The **joint distribution** of these variables is a table in which each entry gives the probability of one combination of values of $X_1, \dots, X_k$
- ◆ Example:

	Toothache	¬Toothache
Cavity	0.04	0.06
¬Cavity	0.01	0.89

$P(\neg \text{Cavity} \wedge \text{Toothache})$

$P(\text{Cavity} \wedge \neg \text{Toothache})$

Joint Distribution Says It All

	Toothache	¬Toothache
Cavity	0.04	0.06
¬Cavity	0.01	0.89

◆ $P(\text{Toothache}) = ??$

◆ $P(\text{Toothache} \vee \text{Cavity}) = ??$

Conditional Probability

- ◆ Definition:
 $P(A|B) = P(A \wedge B) / P(B)$
- ◆ Read $P(A|B)$: **probability of A given B**
- ◆ can also write this as:
 $P(A \wedge B) = P(A|B) P(B)$
- ◆ called the **product rule**

Generalization

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

Bayes' Rule

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

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

Example

	Toothache	¬Toothache
Cavity	0.04	0.06
¬Cavity	0.01	0.89

Given:

- $P(\text{Cavity})=0.1$
- $P(\text{Toothache})=0.05$
- $P(\text{Cavity}|\text{Toothache})=0.8$

Bayes' rule tells:

- $P(\text{Toothache}|\text{Cavity})=(0.8 \times 0.05)/0.1$
=0.4

Representing Probability

- ◆ Naïve representations of probability run into problems.
- ◆ Example:
 - Patients in hospital are described by several attributes:
 - Background: age, gender, history of diseases, ...
 - Symptoms: fever, blood pressure, headache, ...
 - Diseases: pneumonia, heart attack, ...
- ◆ A probability distribution needs to assign a number to each combination of values of these attributes
 - 20 attributes require 10^6 numbers
 - Real examples usually involve hundreds of attributes

Practical Representation

- ◆ **Key idea** -- exploit regularities
- ◆ Here we focus on exploiting **conditional independence** properties

Example

- ◆ customer purchases: Bread, Bagels and Butter (R,A,U)

Bread	Bagels	Butter	$p(r,a,u)$
0	0	0	0.24
0	0	1	0.06
0	1	0	0.12
0	1	1	0.08
1	0	0	0.12
1	0	1	0.18
1	1	0	0.04
1	1	1	0.16

Independence

- Two variables X and Y are **independent** if

$$P(x, y) = P(x)P(y)$$

- or, equivalently,

$$P(x | y) = P(x)$$

- in other words, Y carries no information about X

Example #1

Bread	Bagels	Butter	$p(r,a,u)$
0	0	0	0.24
0	0	1	0.06
0	1	0	0.12
0	1	1	0.08
1	0	0	0.12
1	0	1	0.18
1	1	0	0.04
1	1	1	0.16

Butter	$p(u)$
0	0.52
1	0.48

Bagels	$p(a)$
0	0.6
1	0.4

Bread	$p(r)$
0	0.5
1	0.5

Bagels	Butter	$p(a,u)$
0	0	0.36
0	1	0.24
1	0	0.16
1	1	0.24

Bread	Bagels	$p(r,a)$
0	0	0.3
0	1	0.2
1	0	0.3
1	1	0.2

$$P(a,u) = P(a)P(u)?$$

$$P(r,a) = P(r)P(a)?$$

Conditional Independence

- Let X , Y and Z be sets of random variables. X and Y are conditional independent given Z , denoted $I(X,Y|Z)$, iff, for all values of X , Y and Z

$$P(x, y | z) = P(x | z)P(y | z)$$

Example #2

Hotdogs	Mustard	Ketchup	$p(h,m,k)$
0	0	0	0.576
0	0	1	0.144
0	1	0	0.064
0	1	1	0.016
1	0	0	0.004
1	0	1	0.036
1	1	0	0.016
1	1	1	0.144

Mustard	$p(m)$
0	0.76
1	0.24

Ketchup	$p(k)$
0	0.66
1	0.34

Mustard	Ketchup	$p(m,k)$
0	0	0.58
0	1	0.18
1	0	0.08
1	1	0.16

$$P(m,k) = P(m)P(k)?$$

Example #2

H	M	K	$p(h,m,k)$
0	0	0	0.576
0	0	1	0.144
0	1	0	0.064
0	1	1	0.016
1	0	0	0.004
1	0	1	0.036
1	1	0	0.016
1	1	1	0.144

Mustard	Hotdogs	$p(m h)$
0	0	0.9
0	1	0.2
1	0	0.1
1	1	0.8

Ketchup	Hotdogs	$p(k h)$
0	0	0.8
0	1	0.1
1	0	0.2
1	1	0.9

$$P(m,k|h) = P(m|h)P(k|h)?$$

Mustard	Ketchup	Hotdogs	$p(m,k h)$
0	0	0	0.72
0	1	0	0.18
1	0	0	0.08
1	1	0	0.02
0	0	1	0.02
0	1	1	0.18
1	0	1	0.08
1	1	1	0.72

Example #1

Bread	Bagels	Butter	$p(r,a,u)$
0	0	0	0.24
0	0	1	0.06
0	1	0	0.12
0	1	1	0.08
1	0	0	0.12
1	0	1	0.18
1	1	0	0.04
1	1	1	0.16

Bread	Butter	$p(r,u)$
0	0	0.69...
0	1	0.29...
1	0	0.30...
1	1	0.70...

Bagels	Butter	$p(a,u)$
0	0	0.69...
0	1	0.5
1	0	0.30...
1	1	0.5

$$P(r,a|u) = P(r|u)P(a|u)?$$

Bread	Bagels	Butter	$p(r,a u)$
0	0	0	0.46...
0	1	0	0.23...
1	0	0	0.23...
1	1	0	0.08...
0	0	1	0.12...
0	1	1	0.17...
1	0	1	0.38...
1	1	1	0.33...

Car Example

- ◆ Three propositions:
 - Gas
 - Battery
 - Starts
- ◆ $P(\text{Battery}|\text{Gas}) = P(\text{Battery})$
Gas and Battery are independent
- ◆ $P(\text{Battery}|\text{Gas}, \text{Starts}) \neq P(\text{Battery}|\text{Starts})$
Gas and Battery are not independent given Starts

Summary

- ◆ Example 1: $I(X, Y|\emptyset)$ and not $I(X, Y|Z)$
- ◆ Example 2: $I(X, Y|Z)$ and not $I(X, Y|\emptyset)$
- ◆ conclusion: independence does not imply conditional independence!

Summary

- ◆ Types of uncertainty
- ◆ Default/worst-case/probabilistic reasoning
- ◆ Probability Theory
- ◆ Next:
 - Bayesian Networks
 - Making decisions under uncertainty