

UNCERTAINTY

In which we see what an agent should do when not all is crystal-clear.

Outline

- Uncertainty
- Probabilistic Theory
- Axioms of Probability
- Probabilistic Reasoning
- Independency
- Bayes' Rule
- Summary


2

Uncertainty

WUMPUS-World



Where are the traps?



There are no secure actions, but which one is best?


Uncertainty

3


Reasoning under Uncertainty




Actions 1, 2, 3




Successful: 2
Fails: 1, 2, 3




Successful: 2
Fails: 1, 3



Successful: 3
Fails: 1, 2



Successful: 1
Fails: 1, 2



Successful: 1, 3
Fails: 2

Uncertainty

4

Causes of Uncertainty


- Incomplete knowledge
- Agents unlikely to have complete knowledge about their environment
- Practical uncertainty: not enough data, e.g. noise
- Theoretical uncertainty: no complete theory exists, e.g. medical diagnosis

Uncertainty

5

Uncertainty and Vagueness

- Example:
Boundaries of spatial objects
 - uncertain boundaries: no known where exactly boundary is
 - vague boundary: boundary is not sharp but localized



Uncertain knowledge:
sonar records of a wall



Vague knowledge:
transition soil types

Uncertainty

6

Causes of Uncertainty

- Complex knowledge
 - A complete formalization in FOL is too complex to describe
- Qualification problem
 - Rules about application area are incomplete because there are too many preconditions

Uncertainty

7

Boundaries of Logics

- Technical Diagnosis
 - Input: Symptoms
 - Output: Sources of errors
- Diagnostic Reasoning

$$\forall p \text{ Symptom}(p, \text{no_sound}) \rightarrow \text{Failure}(p, \text{switch}) \vee \text{Failure}(p, \text{battery}) \vee \dots$$
- Example
 - Symptom: no sound from radio
 - Sources of errors: On/Off switch, batteries, ...
- Causal Reasoning

$$\text{Failure}(p, \text{switch}) \vee \text{Failure}(p, \text{battery}) \rightarrow \forall p \text{ Symptom}(p, \text{no_sound})$$

Uncertainty

8

Logics and Uncertainty

- Default logic
 - Rules are valid until contradicted

$$\forall p \text{ Symptom}(p, \text{no_sound}) \rightarrow \text{Failure}(p, \text{switch})$$
 - Additional knowledge can invalidate former derivations
 - Non-Monotonicity
- Fuzzy logic
 - Describes the degree of validity of a statement
 - $\text{rocky}(\text{location}) = 0.43$ i.e. 43% rock
 - Approach for vagueness, not uncertainty
 - Description of natural language

Uncertainty

9

Probability

- Probabilistic Statement
 - “There is a 70% chance of an empty battery if the portable CD player does not give a sound.”
- Combined uncertainty
 - From two different sources
 - “Unknown”: non-existing knowledge
 - “Incomplete”: existing knowledge too complex

Uncertainty

10

Question

- True or False?
 - Probabilistic Theory has the same ontological commitment as logics:
 - Facts hold or do not hold in the world

Uncertainty

11

Different AI Logics

- Logic languages
 - “Logics”
 - Characterized through language elements (logic constants)
 - Facts? Objects?
 - Time? Belief?
 - Ontological commitment (wrt. reality, what exists in world)
 - Epistemological commitment (what an agent believes about facts)

Logic language	Ontological commitment	Epistemological commitment
Propositional logic	Facts	True/False/Unknown
FOL	Facts Objects Relations	True/False/Unknown
Temporal logic	Facts Objects Relations, Times	True/False/Unknown
Probabilistic theory	Facts	Degree of belief 0..1
Fuzzy-Logic	Degree of truth	Degree of belief 0..1

Uncertainty

12

Uncertainty and Decisions

- Action A_t
 - Go to airport t minutes before flight
- Probabilistic verdicts through agent
 - A_{30} is the probability that In get the flight, 1%
 - A_{360} is the probability that In get the flight, 99%
- Which action does the agent select?
 - Does not depend only on probability, also on preferences (→ decision theory)
 - Decision theory = Probabilistic theory + Utility theory

Probabilistic Theory

- How we see it in class
 - Probabilistic theory as an extension of propositional logic
 - Extension: the truth values are labeled with probabilities
- Note
 - Probability theory makes the same ontological commitment as logic, namely, that facts either do or do not hold in the world

Random Variable

- Basic Idea
 - $P(X=a)$ quantifies the probability that random variable X takes the value a
- Variables
 - Boolean
 - Range [TRUE,FALSE]
 - discrete
 - Range of finite set of boolean variables, e.g. weather [sunny,rainy,cloudy,snow]
 - continuous
 - Real values or subsets, z.B. [0,1]

Atomic Events

- Complete specification of the world state about which the agent is uncertain.
 1. Exclusive: only one statement true, $(c \wedge t)$ and $(c \wedge \neg t)$ is mutually exclusive
 2. Exhaustive: at least one of set of all aE is the case, i.e. disjunction of all aE is logically equivalent to TRUE
 3. Entailment: from aE we can entail the truth or falsehood of every proposition, e.g. $(c \wedge \neg t)$ entails $c = \text{TRUE}$, $c \Rightarrow t = \text{FALSE}$
 4. Propositions are logically equivalent to disjunction of all aE that entail the truth of the proposition. E.g. the proposition cavity is equivalent to the disjunction of the aE $c \wedge t$ and $c \wedge \neg t$.
- Example: Cavity (c) and tooth ache (t) has four atomic events (aE)

Prior Probability

- Probability
 - also: a priori probability, unconditional probability
 - $P(a) = 0.4$

means: probability associated with the proposition a is the degree of belief accorded to it in the absence of any other information
- P-Distribution
 - $\mathbf{P}(\text{Switch_on}) = (0.4, 0.6)$

$P(\text{Switch_on} = 1) = 0.4$
 $P(\text{Switch_on} = 0) = 0.6$
 - $\mathbf{P}(\text{Weather}) = (0.25, 0.15, 0.3, 0.3)$

$P(\text{Weather} = \text{sunny}) = 0.25$
 $P(\text{Weather} = \text{rainy}) = 0.15$
 ...

$$P(\text{Cavity} = \text{true}) = 0.1 \quad \text{or} \quad P(\text{cavity}) = 0.1$$

Probability Distribution

- A boolean random variable A
 - For all possible values a value of truth
- Probability distribution $\mathbf{P}(A)$
 - Denotes probability for all possible values for random variable A .

Value	Probability
A	$\mathbf{P}(A)$
0	$0.3 = P(\neg A)$
1	$0.7 = P(A)$

Note: similar notation!
 $P(A)$ short for $P(A=1)$
 $\mathbf{P}(A)$ P-Distribution of A

Joint Probability Distribution

- Multiple boolean random variables
- Cavity, toothache, catch
- For each possibility a defined probability
- $P(\neg \text{cavity} \wedge \neg \text{toothache} \wedge \neg \text{catch}) = 0.576$
- Joint distribution $\mathbf{P}(A,B,C)$
 - Gives probability of all possible value combinations of random variables
 - e.g. $\mathbf{P}(\text{Cavity}, \text{Toothache}, \text{Weather})$ can be represented by a $2 \times 2 \times 4$ table with 16 entries

	toothache		\neg toothache	
	catch	\neg catch	catch	\neg catch
cavity	0.108	0.012	0.072	0.008
\neg cavity	0.016	0.064	0.144	0.576


Conditional Probability

- Conditional probability
 - Agent has evidence (information) about former unknown variables
 - A priori probabilities no longer applicable
 - Notation: $P(A|B)$, A and B are any propositions
 - This is read as "the probability of A given that all we know is B."
 - Example.: $P(\text{cavity}|\text{toothache}) = 0.8$
- Product rule
 - conditional probabilities can be defined as unconditional probabilities:

$$P(a|b) = \frac{P(a \wedge b)}{P(b)}$$
 - which holds for all $P(b) > 0$
 - Can also be written as **product rule**: $P(a \wedge b) = P(a|b) P(b)$
 $P(a \wedge b) = P(b|a) P(a)$
 - $\mathbf{P}(X,Y) = \mathbf{P}(X|Y) \mathbf{P}(Y)$

Kolmogorov's axioms of probability

- Axiom 1
 $0 \leq P(a) \leq 1$
- Axiom 2
 $P(\text{Tautology}) = 1$
 $P(\text{Contradiction}) = 0$
- Axiom 3
 $P(a \vee b) = P(a) + P(b) - P(a \wedge b)$



Using axioms of probability

Derivation of a variety of useful facts from the basic axioms. E.g., the familiar rule for negation follows by substituting $\neg a$ for b in axiom 3:

$$\begin{aligned}
 P(a \vee \neg a) &= P(a) + P(\neg a) - P(a \wedge \neg a) && \text{(by axiom 3 with } b = \neg a) \\
 P(\text{true}) &= P(a) + P(\neg a) - P(\text{false}) && \text{(by logical equivalence)} \\
 1 &= P(a) + P(\neg a) && \text{(by axiom 2)} \\
 P(\neg a) &= 1 - P(a) && \text{(by algebra)}
 \end{aligned}$$

Sum of all probabilities is always 1.


Are these axioms reasonable?

- Example
 - $P(a \wedge b) = 0.0$
 - Agent 2 choose to bet 4\$ on a, 3 on b and 2 on $\neg(a \vee b)$
- Theorem
 - If Agent 1 expresses a set of degrees of belief that violate the axioms of probability theory then there is a betting strategy for Agent 2 that guarantees that Agent 1 will lose money on every bet.
 - Proven by Finetti (1931)

Agent 1		Agent 2	
prop.	belief	bet	stakes
a	0.4	a	4 to 6
b	0.3	b	3 to 7
$a \vee b$	0.8	$\neg(a \vee b)$	2 to 8
Output for Agent 1			
$a \wedge b$	$a \wedge \neg b$	$\neg a \wedge b$	$\neg a \wedge \neg b$
-6	-6	4	4
-7	3	-7	3
2	2	2	-8
-11	-1	-1	-1

Probabilistic Inference

- Conditional probabilities
 $P(A|B) = P(A \wedge B) / P(B)$
 if $P(B) > 0$
- As product rule
 $P(A \wedge B) = P(A|B) P(B)$
 if $P(B) > 0$
 $P(A \wedge B) = P(B|A) P(A)$
 if $P(A) > 0$



Inferences with joint probabilistic distributions

- Inferences
 - Calculate posterior probabilities for given evidences
 - ako “knowledge base”
- Probabilities
 - Sum is 1
 - helps to calculate simple and complex propositions
 - take atomic events where proposition is true and then add probabilities

	toothache		¬toothache	
	catch	¬catch	catch	¬catch
cavity	0.108	0.012	0.072	0.008
¬cavity	0.016	0.064	0.144	0.576

Probabilistic Inference

25

Inferences with joint probabilistic distributions

	toothache		¬toothache	
	catch	¬catch	catch	¬catch
cavity	0.108	0.012	0.072	0.008
¬cavity	0.016	0.064	0.144	0.576

- For example, there are six atomic events in which $cavity \vee toothache$ holds:
 - $P(cavity \vee toothache) = 0.108 + 0.012 + 0.072 + 0.008 + 0.016 + 0.064 = 0.28$
- **Marginal probability**
 - Extract the distribution over some subset of variables, z.B. cavity
 - $P(cavity) = 0.108 + 0.012 + 0.072 + 0.008 = 0.2$

Probabilistic Inference

26

	toothache		¬toothache	
	catch	¬catch	catch	¬catch
cavity	0.108	0.012	0.072	0.008
¬cavity	0.016	0.064	0.144	0.576

Conditional Probabilities

- Example 1
 - Use product rule and take the probability distribution. Here: probability for having cavity given toothache:

$$P(cavity | toothache) = \frac{P(cavity \wedge toothache)}{P(toothache)} = \frac{0.108 + 0.012}{0.108 + 0.012 + 0.016 + 0.064} = 0.6$$

- Example 2
 - Compute the probability that there is no cavity, given toothache:

$$P(\neg cavity | toothache) = \frac{P(\neg cavity \wedge toothache)}{P(toothache)} = \frac{0.016 + 0.064}{0.108 + 0.012 + 0.016 + 0.064} = 0.4$$

Probabilistic Inference

27

	toothache		¬toothache	
	catch	¬catch	catch	¬catch
cavity	0.108	0.012	0.072	0.008
¬cavity	0.016	0.064	0.144	0.576

Normalization

- Constant
 - $1/P(toothache)$ always constant, no matter what value exists for *Cavity*
 - It can be viewed as a normalization constant for the distribution $P(Cavity|toothache)$, ensuring that it adds up to 1
 - We will use α to denote such constants
 - We can then write the two preceding equations in one
- $$P(Cavity|toothache) = \alpha P(Cavity, toothache)$$
- $$= \alpha [P(Cavity, toothache, catch) + P(Cavity, toothache, \neg catch)]$$
- $$= \alpha [\langle 0.108, 0.016 \rangle + \langle 0.012, 0.064 \rangle] = \alpha \langle 0.12, 0.08 \rangle = \langle 0.6, 0.4 \rangle$$

Probabilistic Inference

28

Algorithmic Analysis

- Joint probability distribution
 - n variables with max. k values
- Calculation complexity
 - Table size $O(k^n)$ is exponential in n
 - In the worst case: $O(k^n)$ calculation steps
- Problem in practice
 - We need a lot! of observations in order to get valid and reliable table entries!

Probabilistic Inference

29

Independency

- New variable: Weather
 - Joint probability distribution $P(Toothache, Catch, Cavity, Weather)$
 - Weather has 4 values = 32 entries
- Cavity is independant from weather (also marginal or absolute independency)
 - $P(Toothache, Catch, Cavity, Weather) = P(Toothache, Catch, Cavity) P(Weather)$
 - Reduction of entries: 8 + 4 instead of 32!
- In general
 - $P(X|Y) = P(X)$ or $P(Y|X) = P(Y)$ or $P(X, Y) = P(X) P(Y)$

Independency

30

Decomposition of Joint Distributions

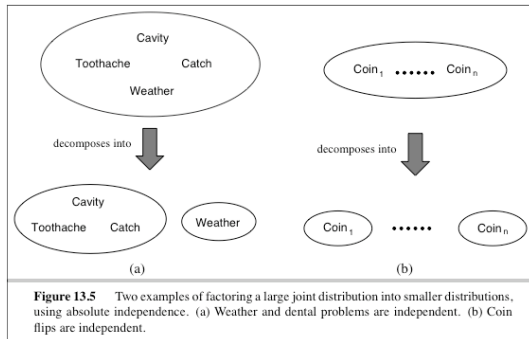


Figure 13.5 Two examples of factoring a large joint distribution into smaller distributions, using absolute independence. (a) Weather and dental problems are independent. (b) Coin flips are independent.

Independency

31

Bayes'-Rule

- Fundamental idea
 - Get around joint distributions
 - Calculate directly with conditional probabilities

- Recall the two forms of the product rule

$$P(a \wedge b) = P(a|b) P(b) \text{ for } P(b) > 0$$

$$P(a \wedge b) = P(b|a) P(a) \text{ for } P(a) > 0$$

$$P(b|a) = \frac{P(a|b)P(b)}{P(a)} \quad \mathbf{P}(Y|X) = \frac{\mathbf{P}(X|Y)\mathbf{P}(Y)}{\mathbf{P}(X)}$$

- Known as Bayes'-Rule

Bayes' Rule

32

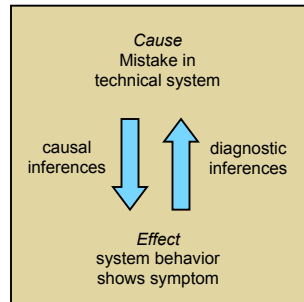
Bayes'-Rule

Why is this rule useful?

- Causal experiences
C: cause, E: effect
- Diagnostic Inference

$$P(C|E) = \frac{P(E|C)P(C)}{P(E)}$$

This simple equation underlies all modern AI systems for probabilistic inference



Bayes' Rule

33

Technical Diagnosis

- Example
 - $C = \text{power_switch_off}$, $E = \text{no_sound_heard}$
- Bayes'-Rule:

$$P(C|E) = \frac{P(E|C)P(C)}{P(E)}$$

- Knowledge from causal experience
 - $P(E|C) = 0.98$
 - $P(C) = 0.01$ (mostly on)
 - $P(E) = 0.2$ (often defect)

Bayes' Rule

34

Technical Diagnosis

- Conditional probability

$$P(C|E) = \frac{P(E|C)P(C)}{P(E)} = \frac{0.98 * 0.01}{0.2} = 0.049$$
- Other independent probabilities

$$P(C) = 0.2 \text{ (more often switched off)}$$

$$P(C|E) = \frac{P(E|C)P(C)}{P(E)} = \frac{0.98 * 0.2}{0.2} = 0.98$$
- Depends highly on a priori probability $P(C)$!

Bayes' Rule

35

Combining evidences

- So far
 - an evidence of the form $P(\text{Effect}|\text{Cause})$
 - What if there is more than one evidence?
 $\mathbf{P}(\text{Cavity}|\text{toothache} \wedge \text{catch})$
 - Possible with joint distributions, but what if we have large problems?
- $\mathbf{P}(\text{toothache} \wedge \text{catch}|\text{Cavity}) = \mathbf{P}(\text{toothache}|\text{Cavity}) \mathbf{P}(\text{catch}|\text{Cavity})$
- Here also: significant reduction of algorithmic complexity $O(n)$ instead $O(2^n)$
- Conditional Independence
 - Toothache and Catch are not really independent
 - Probe goes into tooth that has cavity and catches the tooth
 - But: they are not directly dependent but via cavity

Bayes' Rule

36

Notes

- Conditional independence
 - allows “large” systems that rely on probabilities
 - analogue to independence
- Dentist domain
 - shows that a single cause can have a number of effects
 - theses are conditional independent (given cause)
 - this pattern is often seen

- Decomposition
 - ...of large probabilistic application areas important in AI
- Joint Probability Distribution

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


- These distributions are also called naive Bayes' Models (also Bayes' classifier)

Bayes' Rule and Conditional Independence

$$\begin{aligned} P(\text{Cavity} | \text{toothache} \wedge \text{catch}) \\ &= \alpha P(\text{toothache} \wedge \text{catch} | \text{Cavity}) P(\text{Cavity}) \\ &= \alpha P(\text{toothache} | \text{Cavity}) P(\text{catch} | \text{Cavity}) P(\text{Cavity}) \end{aligned}$$


This is an example of a **naive Bayes** model:

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



The total number of parameters is **linear** in n .

Wumpus World



$P_{ij} = \text{true}$ iff $[i, j]$ contains a pit

$B_{ij} = \text{true}$ iff $[i, j]$ is breezy

Include only $B_{1,1}, B_{1,2}, B_{2,1}$ in the probability model

Specifying the Probability Model

The full joint distribution is $P(P_{1,1}, \dots, P_{4,4}, B_{1,1}, B_{1,2}, B_{2,1})$

Apply product rule: $P(B_{1,1}, B_{1,2}, B_{2,1} | P_{1,1}, \dots, P_{4,4}) P(P_{1,1}, \dots, P_{4,4})$

(Do it this way to get $P(\text{Effect} | \text{Cause})$.)

First term: 1 if pits are adjacent to breezes, 0 otherwise

Second term: pits are placed randomly, probability 0.2 per square:

$$P(P_{1,1}, \dots, P_{4,4}) = \prod_{i,j=1,1}^{4,4} P(P_{i,j}) = 0.2^n \times 0.8^{16-n}$$

for n pits.

Observations and Query

We know the following facts:

$$b = \neg b_{1,1} \wedge b_{1,2} \wedge b_{2,1}$$

$$\text{known} = \neg p_{1,1} \wedge \neg p_{1,2} \wedge \neg p_{2,1}$$

Query is $P(P_{1,3} | \text{known}, b)$

Define $\text{Unknown} = P_{ij}$ s other than $P_{1,3}$ and Known


For inference by enumeration, we have

$$P(P_{1,3} | \text{known}, b) = \alpha \sum_{\text{unknown}} P(P_{1,3}, \text{unknown}, \text{known}, b)$$

Grows exponentially with number of squares!

Using conditional independence

Basic insight: observations are conditionally independent of other hidden squares given neighbouring hidden squares



Define $\text{Unknown} = \text{Fringe} \cup \text{Other}$


$$P(b | P_{1,3}, \text{Known}, \text{Unknown}) = P(b | P_{1,3}, \text{Known}, \text{Fringe})$$

Manipulate query into a form where we can use this!

Using conditional independence (2)

$$\begin{aligned}
 P(P_{1,3}|known, b) &= \alpha \sum_{unknown} P(P_{1,3}, unknown, known, b) \\
 &= \alpha \sum_{unknown} P(b|P_{1,3}, known, unknown) P(P_{1,3}, known, unknown) \\
 &= \alpha \sum_{fringe} \sum_{other} P(b|known, P_{1,3}, fringe, other) P(P_{1,3}, known, fringe, other) \\
 &= \alpha \sum_{fringe} \sum_{other} P(b|known, P_{1,3}, fringe) P(P_{1,3}, known, fringe, other) \\
 &= \alpha \sum_{fringe} P(b|known, P_{1,3}, fringe) \sum_{other} P(P_{1,3}, known, fringe, other) \\
 &= \alpha \sum_{fringe} P(b|known, P_{1,3}, fringe) \sum_{other} P(P_{1,3}) P(known) P(fringe) P(other) \\
 &= \alpha P(known) P(P_{1,3}) \sum_{fringe} P(b|known, P_{1,3}, fringe) P(fringe) \sum_{other} P(other) \\
 &= \alpha' P(P_{1,3}) \sum_{fringe} P(b|known, P_{1,3}, fringe) P(fringe)
 \end{aligned}$$

Using conditional independence (3)



$$\begin{aligned}
 P(P_{1,3}|known, b) &= \alpha' \langle 0.2(0.04 + 0.16 + 0.16), 0.8(0.04 + 0.16) \rangle \\
 &\approx \langle 0.31, 0.69 \rangle
 \end{aligned}$$

$$P(P_{2,2}|known, b) \approx \langle 0.86, 0.14 \rangle$$

Summary

- Probability is a rigorous formalism for uncertain knowledge.
 - Uncertainty arises because of both laziness and ignorance. It is inescapable in complex, dynamic, or inaccessible worlds.
 - Uncertainty means that many of the simplifications that are possible with deductive inference are no longer valid.
 - Probabilities express the agent's inability to reach a definite decision regarding the truth of a sentence, and summarize the agent's beliefs.
 - Basic probability statements include **prior probabilities** and **conditional probabilities** over simple and complex propositions.
 - The **full joint distribution** specifies the probability of each complete assignment of values to random variables. It is usually too large to create or use in its explicit form.

Summary (2)

- The axioms of probability constrain the possible assignments of probabilities to propositions. An agent that violates the axioms will behave irrationally in some circumstances.
- When the full joint distribution is available, it can be used to answer queries simply by adding up entries for the atomic events corresponding to the query propositions.
- **Absolute independence** between subsets of random variables may allow the full joint to be factored into smaller joint distributions. This may greatly reduce complexity but seldom occurs in practice.
- **Bayes' rule** allows unknown probabilities to be computed from known conditional probabilities, usually in the causal direction. Applying Bayes' rule with many pieces of evidence may in general run into the same scaling problems as the full joint distribution.

Summary (3)

- **Conditional independence** brought about by direct causal relationships in the domain may allow the full joint to be factored into smaller, conditional distributions. The **naive Bayes** model assumes conditional independence of all effect variables given a single cause variable, and grows linearly with the number of effects.
- A wumpus-world agent can calculate probabilities for unobserved aspects of the world and use them to make better decisions than a purely logical agent.