

Outline


- Syntax
- Semantics

Bayesian networks

- A simple, graphical notation for conditional independence assertions and hence for compact specification of full joint distributions
- Syntax:
 - a set of nodes, one per variable
 -
 - a directed, acyclic graph (link $\approx$ "directly influences")
 - a conditional distribution for each node given its parents:
$$P(X_i \mid \text{Parents}(X_i))$$
- In the simplest case, conditional distribution represented as a **conditional probability table** (CPT) giving the distribution over X_i for each combination of parent values

Example

- Topology of network encodes conditional independence assertions:




- *Weather* is independent of the other variables
- *Toothache* and *Catch* are conditionally independent given *Cavity*

Example


- I'm at work, neighbor John calls to say my alarm is ringing, but neighbor Mary doesn't call. Sometimes it's set off by minor earthquakes. Is there a burglar?
- Variables: *Burglary, Earthquake, Alarm, JohnCalls, MaryCalls*
- Network topology reflects "causal" knowledge:
 - A burglar can set the alarm off
 - An earthquake can set the alarm off
 - The alarm can cause Mary to call
 - The alarm can cause John to call

Example contd.



Compactness


- A CPT for Boolean X_i with k Boolean parents has 2^k rows for the combinations of parent values
- Each row requires one number p for $X_i = \text{true}$ (the number for $X_i = \text{false}$ is just $1-p$)
- If each variable has no more than k parents, the complete network requires $O(n \cdot 2^k)$ numbers
- I.e., grows linearly with n , vs. $O(2^n)$ for the full joint distribution
- For burglary net, $1 + 1 + 4 + 2 + 2 = 10$ numbers (vs. $2^5 - 1 = 31$)



Semantics

The full joint distribution is defined as the product of the local conditional distributions:

$$P(X_1, \dots, X_n) = \prod_{i=1}^n P(X_i \mid \text{Parents}(X_i))$$



e.g., $P(j \wedge m \wedge a \wedge \neg b \wedge \neg e)$

$$= P(j \mid a) P(m \mid a) P(a \mid \neg b, \neg e) P(\neg b) P(\neg e)$$

Constructing Bayesian networks

- 1. Choose an ordering of variables $X_1, \dots, X_n$
- 2. For $i = 1$ to n
 - add X_i to the network
 -
 - select parents from $X_1, \dots, X_{i-1}$ such that

$$P(X_i \mid \text{Parents}(X_i)) = P(X_i \mid X_1, \dots, X_{i-1})$$

This choice of parents ^{n} guarantees:

$$P(X_1, \dots, X_n) = \prod_{i=1}^n P(X_i \mid X_1, \dots, X_{i-1})$$

(chain rule)

$$= \prod_{i=1}^n P(X_i \mid \text{Parents}(X_i))$$

(by construction)

Example

- Suppose we choose the ordering M, J, A, B, E




$$P(J \mid M) = P(J)?$$

Example

- Suppose we choose the ordering M, J, A, B, E

-



$$P(J \mid M) = P(J)?$$


No

$$P(A \mid J, M) = P(A \mid J)? \quad P(A \mid J, M) = P(A)?$$

Example

- Suppose we choose the ordering M, J, A, B, E

-



$$P(J \mid M) = P(J)?$$

No


$$P(A \mid J, M) = P(A \mid J)? \quad P(A \mid J, M) = P(A)? \quad \text{No}$$

$$P(B \mid A, J, M) = P(B \mid A)?$$

$$P(B \mid A, J, M) = P(B)?$$

Example

- Suppose we choose the ordering M, J, A, B, E
-



$$P(J \mid M) = P(J)?$$

No

$$P(A \mid J, M) = P(A \mid J)? \quad P(A \mid J, M) = P(A)? \quad \text{No}$$

$$P(B \mid A, J, M) = P(B \mid A)? \quad \text{Yes}$$


$$P(B \mid A, J, M) = P(B)? \quad \text{No}$$

$$P(E \mid B, A, J, M) = P(E \mid A)?$$

$$P(E \mid B, A, J, M) = P(E \mid A, B)?$$

Example

- Suppose we choose the ordering M, J, A, B, E
-



$$P(J \mid M) = P(J)?$$

No

$$P(A \mid J, M) = P(A \mid J)? \quad P(A \mid J, M) = P(A)? \quad \text{No}$$


$$P(B \mid A, J, M) = P(B \mid A)? \quad \text{Yes}$$

$$P(B \mid A, J, M) = P(B)? \quad \text{No}$$

$$P(E \mid B, A, J, M) = P(E \mid A)? \quad \text{No}$$

$$P(E \mid B, A, J, M) = P(E \mid A, B)? \quad \text{Yes}$$

Example contd.



- Deciding conditional independence is hard in noncausal directions
- (Causal models and conditional independence seem hardwired for humans!)
- Network is less compact: $1 + 2 + 4 + 2 + 4 = 13$ numbers needed

Summary


- Bayesian networks provide a natural representation for (causally induced) conditional independence
- Topology + CPTs = compact representation of joint distribution
- Generally easy for domain experts to construct

The Doorbell Problem

- The doorbell rang at 12'o clock in the midnight
 - Was someone there at the door?
 - Did Mohan wake up?
- Proposition 1: $\text{AtDoor}(x) \Rightarrow \text{Doorbell}$
- Proposition 2: $\text{Doorbell} \Rightarrow \text{Wake (Mohan)}$

Reasoning about Doorbell 1

- Given Doorbell, can we say $\text{AtDoor}(x)$, because $\text{AtDoor}(x) \Rightarrow \text{Doorbell}$?
- Abductive Reasoning
- But NO, the doorbell might start ringing due to some other reason e.g.
 - short circuit,
 - Wind
 - Animals



P implies Q is true only when P is true and then Q is true

Here Q is true but P may not be true then this is called Abductive reasoning

Reasoning about Doorbell 2

- Given Doorbell, can we say wake(Mohan), because Doorbell $\Rightarrow$ Wake(Mohan)?
- Deductive Reasoning
- Yes, only IF Proposition 2 is always true.
- However, in general Mohan may not always wake up, even if the bell rings.

Any Way Out?

- However, problems like that of the doorbell are very common in real life
- In AI, we often need to reason under such circumstances
- We solve it by proper modeling of *uncertainty* and *impreciseness*, and developing appropriate reasoning techniques.

Sources of Uncertainty - 1

- Implications may be weak

doorbell (0.8) $\Rightarrow$ wake(Mohan)



quantification of frequency
with which the rule applies

- Imprecise language like *often*, *rarely*, *sometimes*
 - Need to quantify these in terms of frequencies
 - Need to design rules for reasoning with these frequencies

Sources of Uncertainty - 2

- Precise information may be too complex
 - Too many antecedents or consequents
 - $\text{AtDoor}(x) \vee \text{ShortCkt} \vee \text{Wind} \dots \Rightarrow \text{Doorbell}$
- Incomplete knowledge
 - We may not know or guess all the possible antecedents or consequents
 - *'The bell rang due to some spooky reason'*

Sources of Uncertainty - 3

- **Conflicting information**

Experts often provide conflicting information:



quantification of measure of belief


Two doctors give different opinions about the patient depends on their belief system

- **Propagation of uncertainties**

- In absence of interdependencies propagation of uncertain knowledge increases the uncertainty of the conclusions


Tomorrow(sunny) [0.6], Tomorrow(warm) [0.8]

Tomorrow(sunny) $\wedge$ Tomorrow(warm) [?]



Degree of belief
in statement a implies b


Some certainty point 0.7



Strength of belief

Uncertainty gets
propagated

$$a \equiv 0.8$$



Propagation of
uncertainty

A Relook at MYCIN Rules

IF: (1) the stain of the organism is gram-positive, and
 (2) the morphology of the organism is coccus, and
 (3) the growth conformation of the organism is clumps.

THEN

there is suggestive evidence(0.7) that the identity of
the organism is staphylococcus.

PREMISE: (\$AND (SAME CNTXT GRAM GRAMPOS)
 (SAME CNTXT MORPH COCCUS)
 (SAME CNTXT CONFORM CLUMPS))

ACTION: (CONCLUDE CNTXT IDENT STAPHYLOCOCCUS TALLY
 .7)

Mycin: Certainty Factors

What do certainty factors mean?

It is an expert's estimate of degree of belief or disbelief in an evidence hypothesis relationship

$e \rightarrow h$

It is a subjective probability estimate provided by the expert from his/her experience.

Measure of Belief $MB(h,e)$

$$e_1 \rightarrow h \quad 0.6$$

Evidence does not support hypothesis


$$MD[h, e_1] = 0.6$$

$$e_2 \rightarrow h \quad 0.8$$

$$MB[h, e_1] = 0.7$$

$$\begin{aligned} CF[h, e] &= MB[h, e] - MD[h, e] \\ &= 0.8 - 0.6 = \underline{0.2} \end{aligned}$$

Certainty Factors



- (a) several rules contribute to one hypothesis
- (b) what is our belief in several propositions taken together?
- (c) what is our belief in the result of rule chaining?

Certainty Algebra

- heuristic (expert given) approach for reasoning with uncertainty
- let us introduce

measure of belief $MB(h|e)$

$1 > MB(h|e) > 0$ if $MD(h|e) = 0$

measure of disbelief $MD(h|e)$

$1 > MD(h|e) > 0$ if $MB(h|e) = 0$

certainty factor $CF(h|e)$

$CF(h|e) = MB(h|e) - MD(h|e)$

$MB(h|e)$ is read as measure of belief hypothesis given evidence


Certainty Factors

- Measure of belief: $MB[h,e]$
 - Measure of disbelief: $MD[h,e]$
- $$CF[h,e] = MB[h,e] - MD[h,e]$$

Additional Evidence

$$\begin{aligned} MB[h, e1 \wedge e2] &= 0 && \text{if } MD[h, e1 \wedge e2] = 1 \\ &= MB[h,e1] + MB[h,e2] \times (1 - MB[h,e1]) && \text{otherwise} \end{aligned}$$

$$\begin{aligned} MD[h, e1 \wedge e2] &= 0 && \text{if } MB[h, e1 \wedge e2] = 1 \\ &= MD[h,e1] + MD[h,e2] \times (1 - MD[h,e1]) && \text{otherwise} \end{aligned}$$



$$MB(h/e_1 \wedge e_2) = 0$$

$$\underline{\underline{MD(h/e_1 \wedge e_2)}} = 1$$

Only very recently, the scientific and engineering community has begun to recognize the utility of defining multiple types of uncertainty. In part the greater depth of study into the scope of uncertainty is made possible by the significant advancements in computational power we now enjoy. As systems become computationally better equipped to handle complex analyses, we encounter the limitations of applying only one mathematical framework (traditional probability theory) used to represent the full scope of uncertainty. The dual nature of uncertainty is described with the following definitions from [Helton, 1997]:

Aleatory Uncertainty – the type of uncertainty which results from the fact that a system can behave in random ways
also known as: Stochastic uncertainty, Type A uncertainty, Irreducible uncertainty, Variability, Objective uncertainty

Epistemic Uncertainty- the type of uncertainty which results from the lack of knowledge about a system and is a property of the analysts performing the analysis.
also known as: Subjective uncertainty, Type B uncertainty, Reducible uncertainty, State of Knowledge uncertainty, Ignorance

BAYES THEOREM

The diagram illustrates Bayes' Theorem with three boxes at the top. The left box, labeled 'Probability that claim **h** is true given evidence **e** and the background knowledge **b**', has an arrow pointing to the term $P(h | e, b)$. The middle box, labeled 'Prior probability that the hypothesis is true', has an arrow pointing to the term $P(h | b)$. The right box, labeled 'Likelihood of evidence if prior hypothesis is true', has an arrow pointing to the term $P(e | h, b)$. Below these, the equation is presented: $P(h | e, b) = \frac{P(h | b) \times P(e | h, b)}{P(h | b) \times P(e | h, b) + P(\neg h | b) \times P(e | \neg h, b)}$. A horizontal line separates the numerator from the denominator.

$$P(h | e, b) = \frac{P(h | b) \times P(e | h, b)}{P(h | b) \times P(e | h, b) + P(\neg h | b) \times P(e | \neg h, b)}$$

Read $\neg$ = NOT

h -> hypothesis

e -> evidence

b -> background knowledge

P -> Probability or Claim is true

Bayes' Theorem

as old as some vampires I know...

- Thomas Bayes (1702 - 1761)
- Pierre-Simon Laplace (1749 - 1827)
- E.T. Jaynes (1922 - 1998)

BOLD STATEMENT...

- Bayes' Theorem is the mathematical model for all correct reasoning about empirical claims.
- Every time you reason correctly, you are following Bayes' Theorem, even if you don't know it. *And if you aren't following it, you aren't reasoning correctly.*
- Understanding Bayes' Theorem is therefore the key to understanding correct reasoning.
- **And this makes it a powerful tool you can use to master the universe.**

Empirical data and claims – Observational data and claims

Consider a simple example

- a. There is someone at the door ringing the doorbell. It will have 50% chance of that being a girl.
- b. But the probability that the person is a girl is 99% when you know the name is "Jaynes".
- c. The fact I knew is that out of 100 people, 99 people are girls by the name of "Jaynes". But one may be a boy also.

Applying Bay's theorem

$$\begin{aligned} \text{Probability of} \\ \text{"Jaynes" being a girl is true} &= \frac{\text{Jaynes that are girls}}{\text{Jaynes that are girls} + \text{Jaynes that are boys}} \\ &= 99\% \end{aligned}$$

Example from Jankiraman book-

Consider an incandescent bulb manufacturing unit. Here machine M1, M2 and M3 make 30%, 30% and 40% of the total bulbs. Of their output let's assume that 2%, 3% and 4% are defective. A bulb is drawn at random and is found defective. What is the probability that the bulb is made by machine M1, M2 or M3.

Solution – Let E1, E2 and E3 be the events that a bulb selected at random is made by machine M1, M2 and M3. Let Q denote that it is defective.

Prob(E1) = 0.3, Prob(E2) = 0.3 and Prob(E3) = 0.4 (Given Data)

Prob of drawing a defective bulb made by M1 = Prob (Q|E1) = 0.02

Prob of drawing a defective bulb made by M2 = Prob (Q|E2) = 0.03

Prob of drawing a defective bulb made by M3 = Prob (Q|E3) = 0.04


Therefore

$$\text{Prob (E1|Q)} = \frac{\text{Prob (E1)} * \text{Prob(Q|E1)}}{\text{Sum (1,3) Prob(Ei)} * \text{Prob (Q|Ei)}} = 0.3 * 0.02 / (0.03*0.2) + (0.03*0.3) + (0.04*0.4) = 0.1935$$

$$\text{Prob (E2|Q)} = \frac{\text{Prob (E2)} * \text{Prob(Q|E2)}}{\text{Sum (1,3) Prob(Ei)} * \text{Prob (Q|Ei)}} = 0.3 * 0.03 / (0.03*0.2) + (0.03*0.3) + (0.04*0.4) = 0.2903$$

$$\text{Prob (E3|Q)} = 0.5162$$

Disease: 10% have it, a test to detect it is 92% accurate,
and has a 5% false alarm rate.



1. You tested positive, what is the probability you have the disease?

$$\frac{.092}{.045 + .092} = .67$$

$$\frac{.008}{.855 + .008} = .0093$$

9907

2. Your friend tested negative, what is the probability your friend has the disease?

4.3 Dempster-Shafer Theory

- Dempster-Shafer theory is an approach to combining evidence
- Dempster (1967) developed means for combining degrees of belief derived from independent items of evidence.
- His student, Glenn Shafer (1976), developed method for obtaining degrees of belief for one question from subjective probabilities for a related question
- People working in Expert Systems in the 1980s saw their approach as ideally suitable for such systems.

4.3 Dempster-Shafer Theory

- Each fact has a degree of support, between 0 and 1:
 - 0 No support for the fact
 - 1 full support for the fact
- Differs from Bayesian approach in that:
 - Belief in a fact and its negation need not sum to 1.
 - Both values can be 0 (meaning no evidence for or against the fact)

4.3 Dempster-Shafer Theory


Set of possible conclusions: Θ

$$\Theta = \{ \theta_1, \theta_2, \dots, \theta_n \}$$

Where:

- Θ is the set of possible conclusions to be drawn
- Each θ_i is **mutually exclusive**: at most one has to be true.
- Θ is **Exhaustive**: At least one θ_i has to be true.

AI - Cycle



Difference Between database and knowledge base

Database can be lots of things, but usually it means a relational database such as Oracle, Microsoft Access, MySQL, etc. These are typically use for storing sets of related data such as accounting records, questions on yahoo answers, etc.

A Knowledge base is used in AI. When an AI algorithm tries to make a decision, it queries its knowledge base to determine how to act. Depending on the result it can then update the knowledge base. It is part of machine learning.

A database stores data - for example, personnel data, sales data etc. As they stand, simply raw data are not of much practical value, unless they can be transformed into information - for example you may be able to analyze the sales data and arrive at purchase patterns, so your company can leverage that information into profit. Now that data has become information. Now, the question to ask is: what kind of "expertise" did you apply in transforming the raw data into useful information? Can you store that "expertise", that "how to", in some place, so that somebody else or perhaps some automated process can use that "stored knowledge" to do future analysis? There, you have your knowledge base.

Deductive Reasoning / Modus Ponens

1. Deduce new information from logically related known information
2. A deductive argument offers assertions that lead automatically to a conclusion

For example if there is a dry wood, oxygen and a spark then there will be fire

Given : There is a wood, oxygen and a spark

We can deduce : There will be a fire

All men are Mortal, Socrates is a man therefore Socrates is mortal.

Inductive Reasoning

1. From a limited set of observations, we form a “generalization”

For example – Observation: All the crows that I have seen in my life are black

Conclusion: All crow are black

Abductive Reasoning

1. “Deduction is exact” in that the deductions follow a logically provable way from the basic axioms (sentences)
2. Abductive is a form of deduction that allows for plausible inference ie conclusion may be wrong.

For example Implication : She carries an umbrella if it is raining

Axiom : She is carrying an umbrella

Inference / Conclusion: It is raining

Analogical reasoning

1. Draw analogy between two situations looking for similarities and differences
For example When you say driving a truck is just like driving a car.

Commonsense Reasoning

1. Heuristic reasoning – gained thru experience, rule-of-thumb
for example If you are moving in a car when it is raining then reduce speed
Extensively used by doctors while diagnosing

Non-Monotonic Reasoning

1. Used when the facts of the case is not static and changes with time
For example – If the wind blow then the curtain sway
After sometime this wind stops blowing the truth no longer exist

Monotonic Reasoning

1. Used when the facts does not change over time
for example The sun rises in the east

Forward chaining is one of the two main methods of reasoning when using [inference rules](#) (in [artificial intelligence](#)) and can be described [logically](#) as repeated application of [modus ponens](#). Forward chaining is a popular implementation strategy for [expert systems](#), [business](#) and [production rule systems](#). The opposite of forward chaining is [backward chaining](#).

Forward chaining starts with the available [data](#) and uses inference rules to extract more data (from an end user for example) until a [goal](#) is reached. An [inference engine](#) using forward chaining searches the inference rules until it finds one where the [antecedent](#) (**If** clause) is known to be true. When found it can conclude, or infer, the [consequent](#) (**Then** clause), resulting in the addition of new [information](#) to its data.

Inference engines will [iterate](#) through this process until a goal is reached.

For example, suppose that the goal is to conclude the color of a pet named Fritz, given that he croaks and eats flies, and that the [rule base](#) contains the following four rules:

If X croaks and eats flies - **Then** X is a frog

If X chirps and sings - **Then** X is a canary

If X is a frog - **Then** X is green

If X is a canary - **Then** X is yellow

This rule base would be searched and the first rule would be selected, because its antecedent (**If** Fritz croaks and eats flies) matches our data. Now the consequents (**Then** X is a frog) is added to the data. The rule base is again searched and this time the third rule is selected, because its antecedent (**If** Fritz is a frog) matches our data that was just confirmed. Now the new consequent (**Then** Fritz is green) is added to our data. Nothing more can be inferred from this information, but we have now accomplished our goal of determining the color of Fritz.

Because the data determines which rules are selected and used, this method is called [data-driven](#), in contrast to [goal-driven](#) backward chaining inference. The forward chaining approach is often employed by [expert systems](#), such as [CLIPS](#).

One of the advantages of forward-chaining over backward-chaining is that the reception of new data can trigger new inferences, which makes the engine better suited to dynamic situations in which conditions are likely to change.

Backward chaining starts with a list of [goals](#) (or a [hypothesis](#)) and works backwards from the [consequent](#) to the [antecedent](#) to see if there is [data](#) available that will support any of these consequents.^[2] An [inference engine](#) using backward chaining would search the [inference](#) rules until it finds one which has a consequent (**Then** clause) that matches a desired goal. If the antecedent (**If** clause) of that rule is not known to be true, then it is added to the list of goals (in order for one's goal to be confirmed one must also provide data that confirms this new rule).

For example, suppose that the goal is to conclude the color of my pet Fritz, given that he croaks

An Example of Backward Chaining.

1.If X croaks and eats flies – **Then** X is a frog

2.If X chirps and sings – **Then** X is a canary

3.If X is a frog – **Then** X is green

4.If X is a canary – **Then** X is yellow


This rule base would be searched and the third and fourth rules would be selected, because their consequents (**Then** Fritz is green, **Then** Fritz is yellow) match the goal (to determine Fritz's color). It is not yet known that Fritz is a frog, so both the antecedents (**If** Fritz is a frog, **If** Fritz is a canary) are added to the goal list. The rule base is again searched and this time the first two rules are selected, because their consequents (**Then** X is a frog, **Then** X is a canary) match the new goals that were just added to the list. The antecedent (**If** Fritz croaks and eats flies) is known to be true and therefore it can be concluded that Fritz is a frog, and not a canary. The goal of determining Fritz's color is now achieved (Fritz is green if he is a frog, and yellow if he is a canary, but he is a frog since he croaks and eats flies; therefore, Fritz is green).

Note that the goals always match the affirmed versions of the consequents of implications (and not the negated versions as in [modus tollens](#)) and even then, their antecedents are then considered as the new goals (and not the conclusions as in [affirming the consequent](#)) which ultimately must match known facts (usually defined as consequents whose antecedents are always true); thus, the inference rule which is used is [modus ponens](#).


Because the list of goals determines which rules are selected and used, this method is called [goal-driven](#), in contrast to [data-driven forward-chaining](#) inference. The backward chaining approach is often employed by [expert systems](#).

Programming languages such as [Prolog](#), [Knowledge Machine](#) and [ECLiPSe](#) support backward chaining within their inference engines

Boolean vs Fuzzy



Boolean vs Fuzzy



Logical Operators for Fuzzy

A	B	$\min(A,B)$
0	0	0
0	1	0
1	0	0
1	1	1

AND

A	B	$\max(A,B)$
0	0	0
0	1	1
1	0	1
1	1	1

OR

A	$1 - A$
0	1
1	0

NOT



AND
 $\min(A, B)$




OR
 $\max(A, B)$



NOT
 $(1 - A)$

Fuzzy System Structure



Fuzzy Set Representation

- Tall =
(0/5, 0.25/5.5, 0.7/6, 1/6.5, 1/7)
- **Numerator**: membership value
- **Denominator**: actual value of the variable

Fuzzy Rules

- If x is A then y is B

└──────────┘

premise or
antecedent

└──────────┘

conclusion or
consequent

- If hotel service is good then
tip is average

Fuzzy Rules

- If Speed is slow Then make the acceleration high
- If Temperature is low AND Pressure is medium Then make the speed very slow

Antecedent

Consequent

1. Fuzzify inputs



2. Apply OR operator (max)



3. Apply implication operator

