

“The mind is a neural computer, fitted by natural selection with combinatorial algorithms for causal and probabilistic reasoning about plants, animals, objects, and people.”

...

“In a universe with any regularities at all, decisions informed about the past are better than decisions made at random. That has always been true, and we would expect organisms, especially informavores such as humans, to have evolved acute intuitions about probability. The founders of probability, like the founders of logic, assumed they were just formalizing common sense.”

Steven Pinker, *How the Mind Works*, 1997, pp. 524, 343.

Learning Objectives

At the end of the class you should be able to:

- justify the use and semantics of probability
- know how to compute marginals and apply Bayes' theorem
- build a belief network for a domain
- predict the inferences for a belief network
- explain the predictions of a causal model

Using Uncertain Knowledge

- Agents don't have complete knowledge about the world.

Using Uncertain Knowledge

- Agents don't have complete knowledge about the world.
- Agents need to make (informed) decisions given their uncertainty.

Using Uncertain Knowledge

- Agents don't have complete knowledge about the world.
- Agents need to make (informed) decisions given their uncertainty.
- It isn't enough to assume what the world is like.
Example: wearing a seat belt.

Using Uncertain Knowledge

- Agents don't have complete knowledge about the world.
- Agents need to make (informed) decisions given their uncertainty.
- It isn't enough to assume what the world is like.
Example: wearing a seat belt.
- An agent needs to reason about its uncertainty.

Using Uncertain Knowledge

- Agents don't have complete knowledge about the world.
- Agents need to make (informed) decisions given their uncertainty.
- It isn't enough to assume what the world is like.
Example: wearing a seat belt.
- An agent needs to reason about its uncertainty.
- When an agent makes an action under uncertainty, it is gambling \implies probability.

- Probability is an agent's measure of belief in some proposition — **subjective probability**.

- Probability is an agent's measure of belief in some proposition — **subjective probability**.
- An agent's belief depends on its prior belief and what it observes.
- **Example:** An agent's probability of a particular bird flying
 - ▶ Other agents may have different probabilities
 - ▶ An agent's belief in a bird's flying ability is affected by what the agent knows about that bird.

Random Variables

- A **random variable** starts with upper case.
- The **domain** of a variable X , written $domain(X)$, is the set of values X can take. (Sometimes use “range”, “frame”, “possible values”).

Random Variables

- A **random variable** starts with upper case.
- The **domain** of a variable X , written $domain(X)$, is the set of values X can take. (Sometimes use “range”, “frame”, “possible values”).
- A tuple of random variables $\langle X_1, \dots, X_n \rangle$ is a complex random variable with domain $domain(X_1) \times \dots \times domain(X_n)$.
Often the tuple is written as X_1, \dots, X_n .

Random Variables

- A **random variable** starts with upper case.
- The **domain** of a variable X , written $domain(X)$, is the set of values X can take. (Sometimes use “range”, “frame”, “possible values”).
- A tuple of random variables $\langle X_1, \dots, X_n \rangle$ is a complex random variable with domain $domain(X_1) \times \dots \times domain(X_n)$.
Often the tuple is written as X_1, \dots, X_n .
- Assignment $X = x$ means variable X has value x .

Random Variables

- A **random variable** starts with upper case.
- The **domain** of a variable X , written $domain(X)$, is the set of values X can take. (Sometimes use “range”, “frame”, “possible values”).
- A tuple of random variables $\langle X_1, \dots, X_n \rangle$ is a complex random variable with domain $domain(X_1) \times \dots \times domain(X_n)$.
Often the tuple is written as X_1, \dots, X_n .
- Assignment $X = x$ means variable X has value x .
- A **proposition** is a Boolean formula made from assignments of values to variables or inequality (e.g., $<$, \leq, \dots) between variables and values.

Possible World Semantics

- A **possible world** specifies an assignment of one value to each random variable.
- A **random variable** is a function from possible worlds into the domain of the random variable.

Possible World Semantics

- A **possible world** specifies an assignment of one value to each random variable.
- A **random variable** is a function from possible worlds into the domain of the random variable.
- $\omega \models X = x$
means variable X is assigned value x in world ω .

Possible World Semantics

- A **possible world** specifies an assignment of one value to each random variable.
- A **random variable** is a function from possible worlds into the domain of the random variable.
- $\omega \models X = x$
means variable X is assigned value x in world ω .
- Logical connectives have their standard meaning:
 - $\omega \models \alpha \wedge \beta$ if $\omega \models \alpha$ and $\omega \models \beta$
 - $\omega \models \alpha \vee \beta$ if $\omega \models \alpha$ or $\omega \models \beta$
 - $\omega \models \neg\alpha$ if $\omega \not\models \alpha$
- Let Ω be the set of all possible worlds.

Semantics of Probability

Probability defines a measure on sets of possible worlds.
A **probability measure** is a function μ from sets of worlds into the non-negative real numbers such that:

- $\mu(\Omega) = 1$
- $\mu(S_1 \cup S_2) = \mu(S_1) + \mu(S_2)$
if $S_1 \cap S_2 = \{\}$.

Semantics of Probability

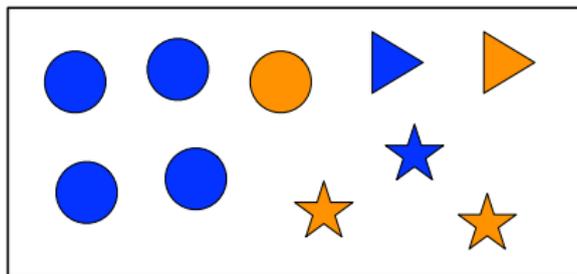
Probability defines a measure on sets of possible worlds.

A **probability measure** is a function μ from sets of worlds into the non-negative real numbers such that:

- $\mu(\Omega) = 1$
- $\mu(S_1 \cup S_2) = \mu(S_1) + \mu(S_2)$
if $S_1 \cap S_2 = \{\}$.

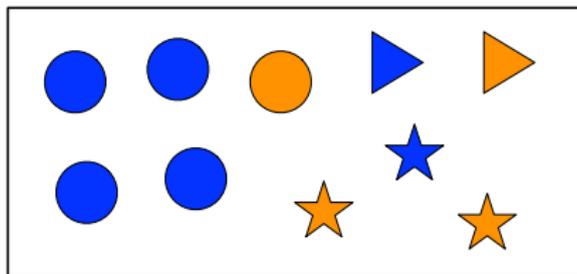
Then $P(\alpha) = \mu(\{\omega \mid \omega \models \alpha\})$.

Possible Worlds:



Suppose the measure of each singleton world is 0.1.

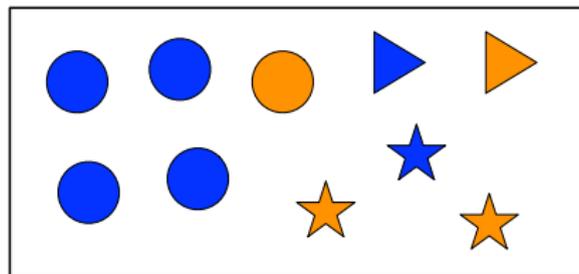
Possible Worlds:



Suppose the measure of each singleton world is 0.1.

- What is the probability of circle?

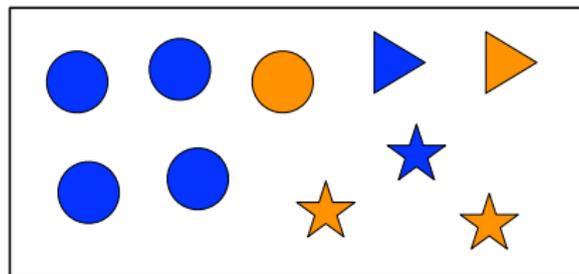
Possible Worlds:



Suppose the measure of each singleton world is 0.1.

- What is the probability of circle?
- What is the probability of star?

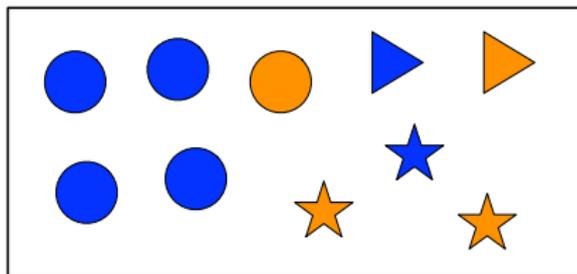
Possible Worlds:



Suppose the measure of each singleton world is 0.1.

- What is the probability of circle?
- What is the probability of star?
- What is the probability of triangle?

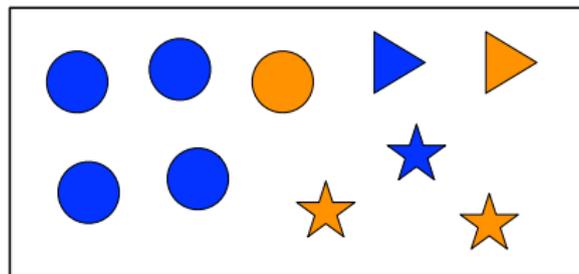
Possible Worlds:



Suppose the measure of each singleton world is 0.1.

- What is the probability of circle?
- What us the probability of star?
- What is the probability of triangle?
- What is the probability of orange?

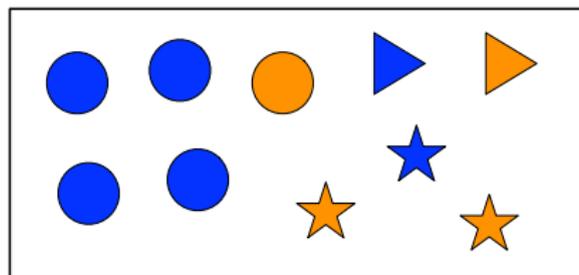
Possible Worlds:



Suppose the measure of each singleton world is 0.1.

- What is the probability of circle?
- What is the probability of star?
- What is the probability of triangle?
- What is the probability of orange?
- What is the probability of orange and star?

Possible Worlds:



Suppose the measure of each singleton world is 0.1.

- What is the probability of circle?
- What is the probability of star?
- What is the probability of triangle?
- What is the probability of orange?
- What is the probability of orange and star?
- What are the random variables?

Axioms of Probability (finite case)

Three axioms define what follows from a set of probabilities:

Axiom 1 $0 \leq P(a)$ for any proposition a .

Axiom 2 $P(\text{true}) = 1$

Axiom 3 $P(a \vee b) = P(a) + P(b)$ if a and b cannot both be true.

- These axioms are sound and complete with respect to the semantics.

Conditioning

- Probabilistic conditioning specifies how to revise beliefs based on new information.

Conditioning

- Probabilistic conditioning specifies how to revise beliefs based on new information.
- An agent builds a probabilistic model taking all background information into account.
This gives a **prior probability**.
- All other information must be conditioned on.
- If **evidence** e is the all of the information obtained subsequently, the **conditional probability** $P(h | e)$ of h given e is the **posterior probability** of h .

Semantics of Conditional Probability

- Evidence e rules out possible worlds incompatible with e .

Semantics of Conditional Probability

- Evidence e rules out possible worlds incompatible with e .
- Evidence e induces a new measure, μ_e , over possible worlds:

Semantics of Conditional Probability

- Evidence e rules out possible worlds incompatible with e .
- Evidence e induces a new measure, μ_e , over possible worlds:

$$\mu_e(S) = \begin{cases} & \text{if } \omega \models e \text{ for all } \omega \in S \end{cases}$$

Semantics of Conditional Probability

- Evidence e rules out possible worlds incompatible with e .
- Evidence e induces a new measure, μ_e , over possible worlds:

$$\mu_e(S) = \begin{cases} c \times \mu(S) & \text{if } \omega \models e \text{ for all } \omega \in S \end{cases}$$

Semantics of Conditional Probability

- Evidence e rules out possible worlds incompatible with e .
- Evidence e induces a new measure, μ_e , over possible worlds:

$$\mu_e(S) = \begin{cases} c \times \mu(S) & \text{if } \omega \models e \text{ for all } \omega \in S \\ & \text{if } \omega \not\models e \text{ for all } \omega \in S \end{cases}$$

Semantics of Conditional Probability

- Evidence e rules out possible worlds incompatible with e .
- Evidence e induces a new measure, μ_e , over possible worlds:

$$\mu_e(S) = \begin{cases} c \times \mu(S) & \text{if } \omega \models e \text{ for all } \omega \in S \\ 0 & \text{if } \omega \not\models e \text{ for all } \omega \in S \end{cases}$$

Semantics of Conditional Probability

- Evidence e rules out possible worlds incompatible with e .
- Evidence e induces a new measure, μ_e , over possible worlds:

$$\mu_e(S) = \begin{cases} c \times \mu(S) & \text{if } \omega \models e \text{ for all } \omega \in S \\ 0 & \text{if } \omega \not\models e \text{ for all } \omega \in S \end{cases}$$

We can show $c =$

Semantics of Conditional Probability

- Evidence e rules out possible worlds incompatible with e .
- Evidence e induces a new measure, μ_e , over possible worlds:

$$\mu_e(S) = \begin{cases} c \times \mu(S) & \text{if } \omega \models e \text{ for all } \omega \in S \\ 0 & \text{if } \omega \not\models e \text{ for all } \omega \in S \end{cases}$$

We can show $c = \frac{1}{P(e)}$.

- The conditional probability of formula h given evidence e is

$$P(h \mid e) =$$

Semantics of Conditional Probability

- Evidence e rules out possible worlds incompatible with e .
- Evidence e induces a new measure, μ_e , over possible worlds:

$$\mu_e(S) = \begin{cases} c \times \mu(S) & \text{if } \omega \models e \text{ for all } \omega \in S \\ 0 & \text{if } \omega \not\models e \text{ for all } \omega \in S \end{cases}$$

We can show $c = \frac{1}{P(e)}$.

- The conditional probability of formula h given evidence e is

$$\begin{aligned} P(h \mid e) &= \mu_e(\{\omega : \omega \models h\}) \\ &= \end{aligned}$$

Semantics of Conditional Probability

- Evidence e rules out possible worlds incompatible with e .
- Evidence e induces a new measure, μ_e , over possible worlds:

$$\mu_e(S) = \begin{cases} c \times \mu(S) & \text{if } \omega \models e \text{ for all } \omega \in S \\ 0 & \text{if } \omega \not\models e \text{ for all } \omega \in S \end{cases}$$

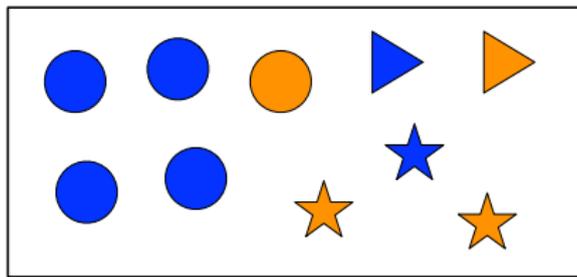
We can show $c = \frac{1}{P(e)}$.

- The conditional probability of formula h given evidence e is

$$\begin{aligned} P(h \mid e) &= \mu_e(\{\omega : \omega \models h\}) \\ &= \frac{P(h \wedge e)}{P(e)} \end{aligned}$$

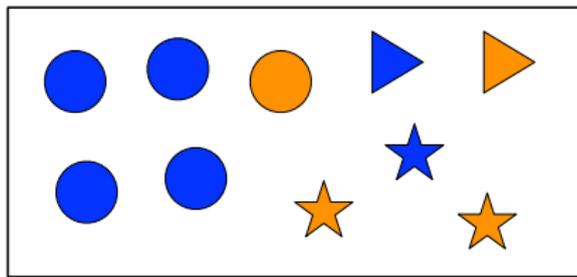
Conditioning

Possible Worlds:



Conditioning

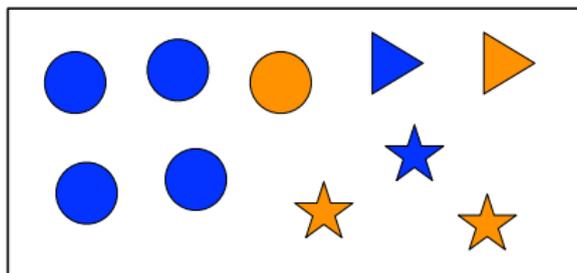
Possible Worlds:



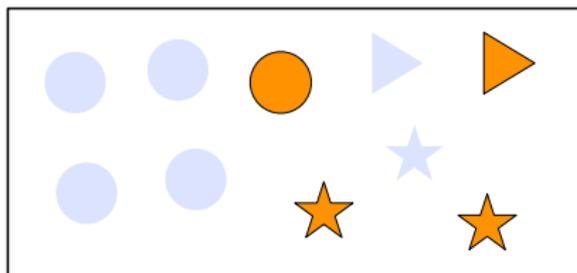
Observe $Color = orange$:

Conditioning

Possible Worlds:



Observe $Color = orange$:



Exercise

<i>Flu</i>	<i>Sneeze</i>	<i>Snore</i>	μ
true	true	true	0.064
true	true	false	0.096
true	false	true	0.016
true	false	false	0.024
false	true	true	0.096
false	true	false	0.144
false	false	true	0.224
false	false	false	0.336

What is:

- (a) $P(\textit{flu} \wedge \textit{sneeze})$
- (b) $P(\textit{flu} \wedge \neg \textit{sneeze})$
- (c) $P(\textit{flu})$
- (d) $P(\textit{sneeze} \mid \textit{flu})$
- (e) $P(\neg \textit{flu} \wedge \textit{sneeze})$
- (f) $P(\textit{flu} \mid \textit{sneeze})$
- (g) $P(\textit{sneeze} \mid \textit{flu} \wedge \textit{snore})$
- (h) $P(\textit{flu} \mid \textit{sneeze} \wedge \textit{snore})$

Chain Rule

$$P(f_1 \wedge f_2 \wedge \dots \wedge f_n)$$
$$=$$

Chain Rule

$$\begin{aligned} &P(f_1 \wedge f_2 \wedge \dots \wedge f_n) \\ &= P(f_n \mid f_1 \wedge \dots \wedge f_{n-1}) \times \\ &\quad P(f_1 \wedge \dots \wedge f_{n-1}) \\ &= \end{aligned}$$

Chain Rule

$$\begin{aligned} & P(f_1 \wedge f_2 \wedge \dots \wedge f_n) \\ &= P(f_n \mid f_1 \wedge \dots \wedge f_{n-1}) \times \\ &\quad P(f_1 \wedge \dots \wedge f_{n-1}) \\ &= P(f_n \mid f_1 \wedge \dots \wedge f_{n-1}) \times \\ &\quad P(f_{n-1} \mid f_1 \wedge \dots \wedge f_{n-2}) \times \\ &\quad P(f_1 \wedge \dots \wedge f_{n-2}) \\ &= P(f_n \mid f_1 \wedge \dots \wedge f_{n-1}) \times \\ &\quad P(f_{n-1} \mid f_1 \wedge \dots \wedge f_{n-2}) \\ &\quad \times \dots \times P(f_3 \mid f_1 \wedge f_2) \times P(f_2 \mid f_1) \times P(f_1) \\ &= \prod_{i=1}^n P(f_i \mid f_1 \wedge \dots \wedge f_{i-1}) \end{aligned}$$

Bayes' theorem

The chain rule and commutativity of conjunction ($h \wedge e$ is equivalent to $e \wedge h$) gives us:

$$P(h \wedge e) =$$

Bayes' theorem

The chain rule and commutativity of conjunction ($h \wedge e$ is equivalent to $e \wedge h$) gives us:

$$P(h \wedge e) = P(h | e) \times P(e)$$

Bayes' theorem

The chain rule and commutativity of conjunction ($h \wedge e$ is equivalent to $e \wedge h$) gives us:

$$\begin{aligned}P(h \wedge e) &= P(h \mid e) \times P(e) \\ &= P(e \mid h) \times P(h).\end{aligned}$$

Bayes' theorem

The chain rule and commutativity of conjunction ($h \wedge e$ is equivalent to $e \wedge h$) gives us:

$$\begin{aligned}P(h \wedge e) &= P(h | e) \times P(e) \\ &= P(e | h) \times P(h).\end{aligned}$$

If $P(e) \neq 0$, divide the right hand sides by $P(e)$:

$$P(h | e) =$$

Bayes' theorem

The chain rule and commutativity of conjunction ($h \wedge e$ is equivalent to $e \wedge h$) gives us:

$$\begin{aligned}P(h \wedge e) &= P(h | e) \times P(e) \\ &= P(e | h) \times P(h).\end{aligned}$$

If $P(e) \neq 0$, divide the right hand sides by $P(e)$:

$$P(h | e) = \frac{P(e | h) \times P(h)}{P(e)}.$$

This is **Bayes' theorem**.

Why is Bayes' theorem interesting?

- Often you have causal knowledge:

$$P(\textit{symptom} \mid \textit{disease})$$

$$P(\textit{light is off} \mid \textit{status of switches and switch positions})$$

$$P(\textit{alarm} \mid \textit{fire})$$

$$P(\textit{image looks like } \img alt="stick figure" data-bbox="400 450 435 485" \mid \textit{a tree is in front of a car})$$

- and want to do evidential reasoning:

$$P(\textit{disease} \mid \textit{symptom})$$

$$P(\textit{status of switches} \mid \textit{light is off and switch positions})$$

$$P(\textit{fire} \mid \textit{alarm}).$$

$$P(\textit{a tree is in front of a car} \mid \textit{image looks like } \img alt="stick figure" data-bbox="780 740 815 775")$$

Exercise

A cab was involved in a hit-and-run accident at night. Two cab companies, the Green and the Blue, operate in the city. You are given the following data:

- 85% of the cabs in the city are Green and 15% are Blue.
- A witness identified the cab as Blue. The court tested the reliability of the witness in the circumstances that existed on the night of the accident and concluded that the witness correctly identifies each one of the two colours 80% of the time and failed 20% of the time.

What is the probability that the cab involved in the accident was Blue?

[From D. Kahneman, Thinking Fast and Slow, 2011, p. 166.]