Probability Foundations

Goal: Make statements and/or predictions about results of physical processes.

Even processes that seem to be simple at first sight may reveal considerable difficulties when trying to predict.

Describing real-world physical processes always calls for a simplifying mathematical model.

Although everybody will have some intuitive notion about probability, we have to formally define the underlying mathematical structure.

Randomness or chance enters as the incapability of precisely modelling a process or the inability of measuring the initial conditions.

- *Example*: Predicting the trajectory of a billard ball over more than 9 banks requires more detailed measurement of the initial conditions (ball location, applied momentum etc.) than physically possible according to Heisenberg's uncertainty principle.

Formal Approach on the Model Side

We conduct an experiment that has a set Ω of possible outcomes.
E. g.:

- Rolling a die $(\Omega = \{1, 2, 3, 4, 5, 6\})$
- Arrivals of phone calls $(\Omega = \mathbb{N}_0)$
- Bread roll weights $(\Omega = \mathbb{R}_+)$

Such an outcome is called an **elementary event**.

All possible elementary events are called the **frame of discernment** Ω (or sometimes **universe of discourse**).

The set representation stresses the following facts:

- All possible outcomes are covered by the elements of Ω. (**collectively exhaustive**).
- Every possible outcome is represented by exactly one element of Ω. (**mutual disjoint**).

Events

Often, we are interested in *higher-level* events
(e.g. casting an odd number, arrival of at least 5 phone calls or
purchasing a bread roll heavier than 80 grams)

Any subset $A \subseteq \Omega$ is called an **event** which **occurs**, if the outcome $\omega_0 \in \Omega$ of
the random experiment lies in A:

$$\text{Event } A \subseteq \Omega \text{ occurs} \quad \Leftrightarrow \quad \bigvee_{\omega \in A} (\omega = \omega_0) \;=\; \text{true} \quad \Leftrightarrow \quad \omega_0 \in A$$

Since events are sets, we can define for two events A and B:

- $A \cup B$ occurs if A or B occurs; $A \cap B$ occurs if A and B occurs.
- $\overline{A}$ occurs if A does not occur (i.e., if $\Omega \backslash A$ occurs).
- A and B are *mutually exclusive*, iff $A \cap B = \emptyset$.

A family of sets $\quad \mathcal{E} = \{E_1, \ldots, E_n\} \quad$ is called an **event algebra**, if the following conditions hold:

- The **certain event** Ω lies in $\mathcal{E}$.

- If $E \in \mathcal{E}$, then $\overline{E} = \Omega \backslash E \in \mathcal{E}$.

- If E_1 and E_2 lie in $\mathcal{E}$, then $E_1 \cup E_2 \in \mathcal{E}$ and $E_1 \cap E_2 \in \mathcal{E}$.

If Ω is uncountable, we require the additional property:

For a series of events $E_i \in \mathcal{E}, i \in \mathbb{N}$, the events $\displaystyle\bigcup_{i=1}^{\infty} E_i$ and $\displaystyle\bigcap_{i=1}^{\infty} E_i$ are also in $\mathcal{E}$.

$\mathcal{E}$ is then called a σ**-algebra**.

Side remarks:

Smallest event algebra: $\mathcal{E} = \{\emptyset, \Omega\}$

Largest event algebra (for finite or countable Ω): $\mathcal{E} = 2^{\Omega} = \{A \subseteq \Omega \mid \mathsf{true}\}$

Probability Function

Given an event algebra $\mathcal{E}$, we would like to assign every event $E \in \mathcal{E}$ its probability with a **probability function** $P : \mathcal{E} \to [0, 1]$.

We require P to satisfy the so-called **Kolmogorov Axioms**:

- $\forall E \in \mathcal{E} : 0 \leq P(E) \leq 1$

- $P(\Omega) = 1$

- For pairwise disjoint events $E_1, E_2, \ldots \in \mathcal{E}$ holds:

$$P(\bigcup_{i=1}^{\infty} E_i) = \sum_{i=1}^{\infty} P(E_i)$$

From these axioms one can conclude the following (incomplete) list of properties:

- $\forall E \in \mathcal{E} : P(\overline{E}) = 1 - P(E)$

- $P(\emptyset) = 0$

- If $E_1, E_2 \in \mathcal{E}$ are mutually exclusive, then $P(E_1 \cup E_2) = P(E_1) + P(E_2)$.

Elementary Probabilities and Densities

Question 1: How to calculate P?

Question 2: Are there "default" event algebras?

Idea for question 1: We have to find a way of distributing (thus the notion *distribution*) the unit mass of probability over all elements $\omega \in \Omega$.

- If Ω is finite or countable a **probability mass function** p is used:

$$p : \ \Omega \to [0, 1] \quad \text{and} \quad \sum_{\omega \in \Omega} p(\omega) = 1$$

- If Ω is uncountable (i.e., continuous) a **probability density function** f is used:

$$f : \ \Omega \to \mathbb{R} \quad \text{and} \quad \int_{\Omega} f(\omega) \, \mathrm{d}\omega = 1$$

Idea for question 2 ("default" event algebras) we have to distinguish again between the cardinalities of Ω:

- Ω finite or countable: $\qquad\qquad \mathcal{E} = 2^{\Omega}$

- Ω uncountable, e. g. $\Omega = \mathbb{R}$: $\qquad \mathcal{E} = \mathcal{B}(\mathbb{R})$

$\mathcal{B}(\mathbb{R})$ is the **Borel Algebra**, i. e., the smallest σ-algebra that contains all closed intervals $[a, b] \subset \mathbb{R}$ with $a < b$.

$\mathcal{B}(\mathbb{R})$ also contains all open intervals and single-item sets.

It is sufficient to note here, that all intervals are contained

$$\{[a, b],]a, b],]a, b[, [a, b[\subset \mathbb{R} \mid a < b\} \subset \mathcal{B}(\mathbb{R})$$

because the event of a bread roll having a weight between $80\,\text{g}$ and $90\,\text{g}$ is represented by the interval $[80, 90]$.

Example: Rolling a Die

$$\Omega = \{1, 2, 3, 4, 5, 6\} \quad X = \mathrm{id}$$

$$p_1(\omega) = \tfrac{1}{6} \qquad\qquad F_1(x) = P(X \le x)$$

$$\sum_{\omega \in \Omega} p_1(\omega) = \sum_{i=1}^{6} p_1(\omega_i)$$

$$= \sum_{i=1}^{6} \frac{1}{6} = 1$$

$$P(X \le x) = \sum_{x' \le x} P(X = x')$$

$$P(a < X \le b) = F_1(b) - F_1(a)$$

$$P(X = x) \;=\; P(\{X = x\}) \;=\; P(X^{-1}(x)) \;=\; P(\{\omega \in \Omega \mid X(\omega) = x\})$$

Basics of Applied Probability Theory

Why (Kolmogorov) Axioms?

If P models an *objectively* observable probability, these axioms are obviously reasonable.

However, why should an agent obey formal axioms when modeling degrees of (subjective) belief?

Objective vs. subjective probabilities

Axioms constrain the set of beliefs an agent can abide.

Finetti (1931) gave one of the most plausible arguments why subjective beliefs should respect axioms:

"When using contradictory beliefs, the agent will eventually fail."

$P(A)$ designates the *unconditioned* or *a priori* probability
that $A \subseteq \Omega$ occurs if *no* other additional information is present.
For example:

$$P(\mathsf{cavity}) \ = \ 0.1$$

Note: Here, cavity is a proposition.

A formally different way to state the same would be via
a binary random variable Cavity:

$$P(\mathsf{Cavity} = \mathsf{true}) \ = \ 0.1$$

A priori probabilities are derived from statistical surveys or general rules.

Unconditional Probabilities

In general a random variable can assume more than two values:

$$P(\text{ Weather} = \text{sunny}) = 0.7$$
$$P(\text{ Weather} = \text{rainy}) = 0.2$$
$$P(\text{ Weather} = \text{cloudy}) = 0.02$$
$$P(\text{ Weather} = \text{snowy}) = 0.08$$
$$P(\text{Headache} = \text{true}) = 0.1$$

$P(X)$ designates the vector of probabilities for the (ordered) domain of the random variable X:

$$P(\text{Weather}) = \langle 0.7, 0.2, 0.02, 0.08 \rangle$$
$$P(\text{Headache}) = \langle 0.1, 0.9 \rangle$$

Both vectors define the respective probability distributions of the two random variables.

Conditional Probabilities

New evidence can alter the probability of an event.

Example: The probability for cavity increases if information about a toothache arises.

With additional information present, the a priori knowledge must not be used!

$P(A \mid B)$ designates the *conditional* or *a posteriori* probability of A *given* the sole observation (*evidence*) B.

$$P(\mathsf{cavity} \mid \mathsf{toothache}) \;=\; 0.8$$

For random variables X and Y $P(X \mid Y)$ represents the set of conditional distributions for each possible value of Y.

Conditional Probabilities

P(Weather | Headache) consists of the following table:

|	h $\hat{=}$ Headache = true	¬h $\hat{=}$ Headache = false
Weather = sunny	$P(\text{W = sunny} \mid \text{h})$	$P(\text{W = sunny} \mid \neg\text{h})$
Weather = rainy	$P(\text{W = rainy} \mid \text{h})$	$P(\text{W = rainy} \mid \neg\text{h})$
Weather = cloudy	$P(\text{W = cloudy} \mid \text{h})$	$P(\text{W = cloudy} \mid \neg\text{h})$
Weather = snowy	$P(\text{W = snowy} \mid \text{h})$	$P(\text{W = snowy} \mid \neg\text{h})$

Note that we are dealing with *two* distributions now!
Therefore each column sums up to unity!

Formal definition:

$$P(A \mid B) \; = \; \frac{P(A \wedge B)}{P(B)} \qquad \text{if} \quad P(B) > 0$$

$$P(A \mid B) \;=\; \frac{P(A \wedge B)}{P(B)}$$

Product Rule: $P(A \wedge B) \;=\; P(A \mid B) \cdot P(B)$

Also: $P(A \wedge B) \;=\; P(B \mid A) \cdot P(A)$

A and B are *independent* iff

$$P(A \mid B) \;=\; P(A) \qquad \text{and} \qquad P(B \mid A) \;=\; P(B)$$

Equivalently, iff the following equation holds true:

$$P(A \wedge B) \;=\; P(A) \cdot P(B)$$

Interpretation of Conditional Probabilities

Caution! Common misinterpretation:

"$P(A \mid B) = 0.8$ means, that $P(A) = 0.8$, given B holds."

This statement is wrong due to (at least) two facts:

$P(A)$ is *always* the a-priori probability,
never the probability of A given that B holds!

$P(A \mid B) = 0.8$ is only applicable as long as no other evidence except B is present.
If C becomes known, $P(A \mid B \wedge C)$ has to be determined.

In general we have:

$$P(A \mid B \wedge C) \neq P(A \mid B)$$

E. g. $C \rightarrow A$ might apply.

Joint Probabilities

Let $X_1, \ldots, X_n$ be random variables over the same framce of descernment Ω and event algebra $\mathcal{E}$. Then $\vec{X} = (X_1, \ldots, X_n)$ is called a *random vector* with

$$\vec{X}(\omega) \;=\; (X_1(\omega), \ldots, X_n(\omega))$$

Shorthand notation:

$$P(\vec{X} = (x_1, \ldots, x_n)) \;=\; P(X_1 = x_1, \ldots, X_n = x_n) \;=\; P(x_1, \ldots, x_n)$$

Definition:

$$P(X_1 = x_1, \ldots, X_n = x_n) \;=\; P\left(\left\{\, \omega \in \Omega \;\middle|\; \bigwedge_{i=1}^{n} X_i(\omega) = x_i \,\right\}\right)$$

$$=\; P\left(\bigcap_{i=1}^{n} \{X_i = x_i\}\right)$$

Joint Probabilities

Example: $P(\mathsf{Headache}, \mathsf{Weather})$ is the *joint probability distribution* of both random variables and consists of the following table:

	h $\,\hat{=}\,$ Headache = true	¬h $\,\hat{=}\,$ Headache = false
Weather = sunny	$P(\mathsf{W} = \mathsf{sunny} \,\wedge\, \mathsf{h})$	$P(\mathsf{W} = \mathsf{sunny} \,\wedge\, \mathsf{\neg h})$
Weather = rainy	$P(\mathsf{W} = \mathsf{rainy} \,\wedge\, \mathsf{h})$	$P(\mathsf{W} = \mathsf{rainy} \,\wedge\, \mathsf{\neg h})$
Weather = cloudy	$P(\mathsf{W} = \mathsf{cloudy} \wedge \mathsf{h})$	$P(\mathsf{W} = \mathsf{cloudy} \wedge \mathsf{\neg h})$
Weather = snowy	$P(\mathsf{W} = \mathsf{snowy} \,\wedge\, \mathsf{h})$	$P(\mathsf{W} = \mathsf{snowy} \,\wedge\, \mathsf{\neg h})$

All table cells sum up to unity.

Calculating with Joint Probabilities

All desired probabilities can be computed from a joint probability distribution.

	toothache	¬toothache
cavity	0.04	0.06
¬cavity	0.01	0.89

Example:
$$P(\text{cavity} \lor \text{toothache}) = P(\ \text{cavity}\ \land\ \text{toothache})$$
$$+\ P(\neg\text{cavity}\ \land\ \text{toothache})$$
$$+\ P(\ \text{cavity}\ \land\ \neg\text{toothache})\ =\ 0.11$$

Marginalizations:
$$P(\text{cavity}) = P(\ \text{cavity}\ \land\ \text{toothache})$$
$$+\ P(\ \text{cavity}\ \land\ \neg\text{toothache})\ =\ 0.10$$

Conditioning:
$$P(\text{cavity} \mid \text{toothache})\ =\ \frac{P(\text{cavity} \land \text{toothache})}{P(\text{toothache})}\ =\ \frac{0.04}{0.04 + 0.01}\ =\ 0.80$$

Problems

Easiness of computing all desired probabilities comes at an unaffordable price:

Given n random variables with k possible values each, the joint probability distribution contains k^n entries which is infeasible in practical applications.

Hard to handle.

Hard to estimate.

Therefore:

1. Is there a more *dense* representation of joint probability distributions?

2. Is there a more *efficient* way of processing this representation?

   The answer is *no* for the general case, however, certain dependencies and independencies can be exploited to reduce the number of parameters to a practical size.

Stochastic Independence

Two events A and B are *stochastically independent* iff

$$P(A \wedge B) \;=\; P(A) \cdot P(B)$$

$$\Leftrightarrow$$

$$P(A \mid B) \;=\; P(A) \;=\; P(A \mid \overline{B})$$

Two random variables X and Y are *stochastically independent* iff

$$\forall x \in \text{dom}(X) : \forall y \in \text{dom}(Y) : \qquad P(X = x, Y = y) \;=\; P(X = x) \cdot P(Y = y)$$

$$\Leftrightarrow$$

$$\forall x \in \text{dom}(X) : \forall y \in \text{dom}(Y) : \qquad P(X = x \mid Y = y) \;=\; P(X = x)$$

Shorthand notation: $P(X, Y) = P(X) \cdot P(Y)$.
Note the formal difference between $P(A) \in [0, 1]$ and $P(X) \in [0, 1]^{|\text{dom}(X)|}$.

Conditional Independence

Let X, Y and Z be three random variables. We call X *and* Y *conditionally independent given* Z, iff the following condition holds:

$$\forall x \in \text{dom}(X) : \forall y \in \text{dom}(Y) : \forall z \in \text{dom}(Z) :$$
$$P(X = x, Y = y \mid Z = z) = P(X = x \mid Z = z) \cdot P(Y = y \mid Z = z)$$

Shorthand notation: $X \perp\!\!\!\perp_P Y \mid Z$

Let $\boldsymbol{X} = \{A_1, \ldots, A_k\}$, $\boldsymbol{Y} = \{B_1, \ldots, B_l\}$ and $\boldsymbol{Z} = \{C_1, \ldots, C_m\}$ be three disjoint sets of random variables. We call $\boldsymbol{X}$ *and* $\boldsymbol{Y}$ *conditionally independent given* $\boldsymbol{Z}$, iff

$$P(\boldsymbol{X}, \boldsymbol{Y} \mid \boldsymbol{Z}) = P(\boldsymbol{X} \mid \boldsymbol{Z}) \cdot P(\boldsymbol{Y} \mid \boldsymbol{Z}) \Leftrightarrow P(\boldsymbol{X} \mid \boldsymbol{Y}, \boldsymbol{Z}) = P(\boldsymbol{X} \mid \boldsymbol{Z})$$

Shorthand notation: $\boldsymbol{X} \perp\!\!\!\perp_P \boldsymbol{Y} \mid \boldsymbol{Z}$

Conditional Independence

The complete condition for $\boldsymbol{X} \perp\!\!\!\perp_P \boldsymbol{Y} \mid \boldsymbol{Z}$ would read as follows:

$$\forall a_1 \in \mathrm{dom}(A_1) : \cdots \forall a_k \in \mathrm{dom}(A_k) :$$
$$\forall b_1 \in \mathrm{dom}(B_1) : \cdots \forall b_l \in \mathrm{dom}(B_l) :$$
$$\forall c_1 \in \mathrm{dom}(C_1) : \cdots \forall c_m \in \mathrm{dom}(C_m) :$$
$$P(A_1 = a_1, \ldots, A_k = a_k, B_1 = b_1, \ldots, B_l = b_l \mid C_1 = c_1, \ldots, C_m = c_m)$$
$$= P(A_1 = a_1, \ldots, A_k = a_k \mid C_1 = c_1, \ldots, C_m = c_m)$$
$$\cdot P(B_1 = b_1, \ldots, B_l = b_l \mid C_1 = c_1, \ldots, C_m = c_m)$$

Remarks:

1. If $\boldsymbol{Z} = \emptyset$ we get (unconditional) independence.

2. We do not use curly braces ({}) for the sets if the context is clear. Likewise, we use X instead of $\boldsymbol{X}$ to denote sets.

Y

Group 1

Group 2

X

(Weak) Dependence in the entire dataset: X and Y dependent.

No Dependence in Group 1: X and Y conditionally independent given Group 1.

No Dependence in Group 2: X and Y conditionally independent given Group 2.

- $\mathrm{dom}(G) = \{\mathsf{mal}, \mathsf{fem}\}$ Geschlecht (gender)
- $\mathrm{dom}(S) = \{\mathsf{sm}, \overline{\mathsf{sm}}\}$ Raucher (smoker)
- $\mathrm{dom}(M) = \{\mathsf{mar}, \overline{\mathsf{mar}}\}$ Verheiratet (married)
- $\mathrm{dom}(P) = \{\mathsf{preg}, \overline{\mathsf{preg}}\}$ Schwanger (pregnant)

p_{GSMP}		G = mal		G = fem	
		S = sm	S = $\overline{\mathsf{sm}}$	S = sm	S = $\overline{\mathsf{sm}}$
M = mar	P = preg	0	0	0.01	0.05
	P = $\overline{\mathsf{preg}}$	0.04	0.16	0.02	0.12
M = $\overline{\mathsf{mar}}$	P = preg	0	0	0.01	0.01
	P = $\overline{\mathsf{preg}}$	0.10	0.20	0.07	0.21

$$P(\mathsf{G}{=}\mathsf{fem}) \;=\; P(\mathsf{G}{=}\mathsf{mal}) \;=\; 0.5 \qquad\qquad P(\mathsf{P}{=}\mathsf{preg}) \;=\; 0.08$$
$$P(\mathsf{S}{=}\mathsf{sm}) \;=\; 0.25 \qquad\qquad\qquad P(\mathsf{M}{=}\mathsf{mar}) \;=\; 0.4$$

Gender and **Smoker** are not independent:

$$P(\mathsf{G}{=}\mathsf{fem} \mid \mathsf{S}{=}\mathsf{sm}) \;=\; 0.44 \;\neq\; 0.5 \;=\; P(\mathsf{G}{=}\mathsf{fem})$$

Gender and **Marriage** are marginally independent but
conditionally dependent given **Pregnancy**:

$$P(\mathsf{fem}, \mathsf{mar} \mid \overline{\mathsf{preg}}) \approx 0.152 \quad\neq\quad 0.169 \approx P(\mathsf{fem} \mid \overline{\mathsf{preg}}) \cdot P(\mathsf{mar} \mid \overline{\mathsf{preg}})$$

Bayes Theorem

Product Rule (for events A and B):

$$P(A \cap B) = P(A \mid B)P(B) \quad \text{and} \quad P(A \cap B) = P(B \mid A)P(A)$$

Equating the right-hand sides:

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

For random variables X and Y:

$$\forall x \forall y: \quad P(Y{=}y \mid X{=}x) = \frac{P(X{=}x \mid Y{=}y)P(Y{=}y)}{P(X{=}x)}$$

Generalization concerning background knowledge/evidence E:

$$P(Y \mid X, E) = \frac{P(X \mid Y, E)P(Y \mid E)}{P(X \mid E)}$$

$$P(\text{toothache} \mid \text{cavity}) = 0.4$$

$$P(\text{cavity}) = 0.1 \qquad P(\text{cavity} \mid \text{toothache}) = \frac{0.4 \cdot 0.1}{0.05} = 0.8$$

$$P(\text{toothache}) = 0.05$$

Why not estimate $P(\text{cavity} \mid \text{toothache})$ right from the start?

Causal knowledge like $P(\text{toothache} \mid \text{cavity})$ is more robust than diagnostic knowledge $P(\text{cavity} \mid \text{toothache})$.

The causality $P(\text{toothache} \mid \text{cavity})$ is independent of the a priori probabilities $P(\text{toothache})$ and $P(\text{cavity})$.

If $P(\text{cavity})$ rose in a caries epidemic, the causality $P(\text{toothache} \mid \text{cavity})$ would remain constant whereas both $P(\text{cavity} \mid \text{toothache})$ and $P(\text{toothache})$ would increase according to $P(\text{cavity})$.

A physician, after having estimated $P(\text{cavity} \mid \text{toothache})$, would not know a rule for updating.
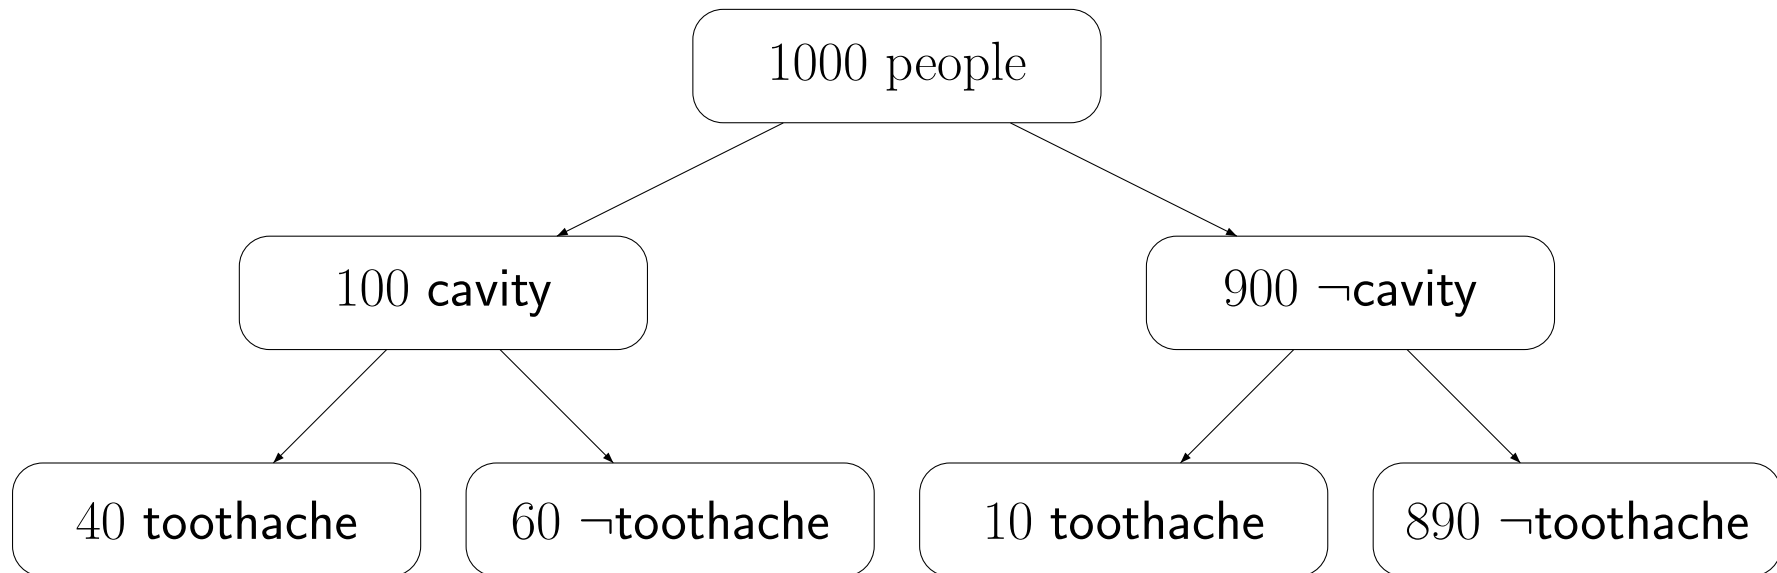
$$P(\text{toothache} \mid \text{cavity}) = 0.4 \qquad\qquad P(\text{cavity}) = 0.1$$

$$P(\text{toothache} \mid \neg\text{cavity}) = \tfrac{1}{90} \qquad\qquad P(\text{cavity} \mid \text{toothache}) = \frac{40}{40 + 10} = 0.8$$

```
                          ┌──────────────┐
                          │  1000 people │
                          └──────────────┘
                         /                \
                ┌──────────────┐    ┌──────────────┐
                │  100 cavity  │    │  900 ¬cavity │
                └──────────────┘    └──────────────┘
                 /          \         /          \
     ┌──────────────┐ ┌──────────────┐ ┌──────────────┐ ┌──────────────────┐
     │40 toothache  │ │60 ¬toothache │ │10 toothache  │ │890 ¬toothache    │
     └──────────────┘ └──────────────┘ └──────────────┘ └──────────────────┘
```

$$P(C \mid T) = \frac{P(T \mid C) \cdot P(C)}{P(T)} = \frac{P(T \mid C) \cdot P(C)}{P(T \mid C) \cdot P(C) + P(T \mid \neg C) \cdot P(\neg C)}$$

Relative Probabilities

Assumption:

We would like to consider the probability of the diagnosis GumDisease as well.

$$P(\text{toothache} \mid \text{gumdisease}) = 0.7$$
$$P(\text{gumdisease}) = 0.02$$

Which diagnosis is more probable?

If we are interested in *relative probabilities* only (which may be sufficient for some decisions), $P(\text{toothache})$ needs not to be estimated:

$$\frac{P(C \mid T)}{P(G \mid T)} = \frac{P(T \mid C)P(C)}{P(T)} \cdot \frac{P(T)}{P(T \mid G)P(G)}$$

$$= \frac{P(T \mid C)P(C)}{P(T \mid G)P(G)} = \frac{0.4 \cdot 0.1}{0.7 \cdot 0.02}$$

$$= 28.57$$

Normalization

If we are interested in the absolute probability of $P(C \mid T)$ but do not know $P(T)$, we may conduct a complete case analysis (according C) and exploit the fact that $P(C \mid T) + P(\neg C \mid T) = 1$.

$$P(C \mid T) = \frac{P(T \mid C)P(C)}{P(T)}$$

$$P(\neg C \mid T) = \frac{P(T \mid \neg C)P(\neg C)}{P(T)}$$

$$1 = P(C \mid T) + P(\neg C \mid T) = \frac{P(T \mid C)P(C)}{P(T)} + \frac{P(T \mid \neg C)P(\neg C)}{P(T)}$$

$$P(T) = P(T \mid C)P(C) + P(T \mid \neg C)P(\neg C)$$

Normalization

Plugging into the equation for $P(C \mid T)$ yields:

$$P(C \mid T) = \frac{P(T \mid C)P(C)}{P(T \mid C)P(C) + P(T \mid \neg C)P(\neg C)}$$

For general random variables, the equation reads:

$$P(Y\!=\!y \mid X\!=\!x) = \frac{P(X\!=\!x \mid Y\!=\!y)P(Y\!=\!y)}{\displaystyle\sum_{\forall y' \in \text{dom}(Y)} P(X\!=\!x \mid Y\!=\!y')P(Y\!=\!y')}$$

Note the "loop variable" y'. Do not confuse with y.

Multiple Evidences

The patient complains about a toothache. From this first evidence the dentist infers:

$$P(\mathsf{cavity} \mid \mathsf{toothache}) = 0.8$$

The dentist palpates the tooth with a metal probe which catches into a fracture:

$$P(\mathsf{cavity} \mid \mathsf{fracture}) = 0.95$$

Both conclusions might be inferred via Bayes rule. But what does the combined evidence yield? Using Bayes rule further, the dentist might want to determine:

$$P(\mathsf{cavity} \mid \mathsf{toothache} \wedge \mathsf{fracture}) = \frac{P(\mathsf{toothache} \wedge \mathsf{fracture} \mid \mathsf{cavity}) \cdot P(\mathsf{cavity})}{P(\mathsf{toothache} \wedge \mathsf{fracture})}$$

Problem:

He needs $P(\mathsf{toothache} \wedge \mathsf{catch} \mid \mathsf{cavity})$, i. e. diagnostics knowledge for all combinations of symptoms in general. Better incorporate evidences step-by-step:

$$P(Y \mid X, E) \; = \; \frac{P(X \mid Y, E) P(Y \mid E)}{P(X \mid E)}$$

Abbreviations:

C — cavity

T — toothache

F — fracture

Objective:

Computing $P(C \mid T, F)$ with just using information about $P(\,\cdot\, \mid C)$ and under exploitation of independence relations among the variables.

A priori: $\qquad\qquad P(C)$

Evidence toothache: $\qquad P(C \mid T) \quad = P(C) \dfrac{P(T \mid C)}{P(T)}$

Evidence fracture: $\qquad P(C \mid T, F) \quad = P(C \mid T) \dfrac{P(F \mid C, T)}{P(F \mid T)}$

Information about conditional independence
$$P(F \mid C, T) = P(F \mid C)$$

$$P(C \mid T, F) \; = \; P(C) \frac{P(T \mid C)}{P(T)} \frac{P(F \mid C)}{P(F \mid T)}$$

Seems that we still have to cope with symptom inter-dependencies?!

Multiple Evidences

Compound equation from last slide:

$$P(C \mid T, F) \;=\; P(C) \, \frac{P(T \mid C) \; P(F \mid C)}{P(T) \; P(F \mid T)}$$

$$=\; P(C) \, \frac{P(T \mid C) \; P(F \mid C)}{P(F, T)}$$

$P(F, T)$ is a normalizing constant and can be computed
if $P(F \mid \neg C)$ and $P(T \mid \neg C)$ are known:

$$P(F, T) = \underbrace{P(F, T \mid C)}_{P(F|C)P(T|C)} \; P(C) + \underbrace{P(F, T \mid \neg C)}_{P(F|\neg C)P(T|\neg C)} \; P(\neg C)$$

Therefore, we finally arrive at the following solution...

$$P(C \mid F, T) = \frac{\boxed{P(C)}\;\boxed{P(T \mid C)}\;\boxed{P(F \mid C)}}{\boxed{P(F \mid C)}\;\boxed{P(T \mid C)}\;\boxed{P(C)} + \boxed{P(F \mid \neg C)}\;\boxed{P(T \mid \neg C)}\;\boxed{P(\neg C)}}$$

Note that we only use causal probabilities $P(\,\cdot\,\mid C)$ together with the a priori (marginal) probabilities $P(C)$ and $P(\neg C)$.

Multiple Evidences — Summary

Multiple evidences can be treated by reduction on

  a priori probabilities

  (causal) conditional probabilities for the evidence

  under assumption of conditional independence

General rule:

$$P(Z \mid X, Y) = \alpha \, P(Z) \, P(X \mid Z) \, P(Y \mid Z)$$

for X and Y conditionally independent given Z and with normalizing constant α.

Monty Hall Puzzle

Marylin Vos Savant in her riddle column in the New York Times:

You are a candidate in a game show and have to choose between three doors. Behind one of them is a Porsche, whereas behind the other two there are goats. After you chose a door, the host Monty Hall (who knows what is behind each door) opens another (not your chosen one) door with a goat. Now you have the choice between keeping your chosen door or choose the remaining one.

Which decision yields the best chance of winning the Porsche?

Monty Hall Puzzle

G You win the Porsche.

R You revise your decision.

A Behind your initially chosen door is (and remains) the Porsche.

$$
\begin{aligned}
P(G \mid R) &= P(G, A \mid R) + P(G, \overline{A} \mid R) \\
&= P(G \mid A, R)P(A \mid R) + P(G \mid \overline{A}, R)P(\overline{A} \mid R) \\
&= 0 \cdot P(A \mid R) + 1 \cdot P(\overline{A} \mid R) \\
&= P(\overline{A} \mid R) = P(\overline{A}) = \frac{2}{3}
\end{aligned}
$$

$$
\begin{aligned}
P(G \mid \overline{R}) &= P(G, A \mid \overline{R}) + P(G, \overline{A} \mid \overline{R}) \\
&= P(G \mid A, \overline{R})P(A \mid \overline{R}) + P(G \mid \overline{A}, \overline{R})P(\overline{A} \mid \overline{R}) \\
&= 1 \cdot P(A \mid \overline{R}) + 0 \cdot P(\overline{A} \mid \overline{R}) \\
&= P(A \mid \overline{R}) = P(A) = \frac{1}{3}
\end{aligned}
$$

Simpson's Paradox

Example: C = Patient takes medication, E = patient recovers

	E	$\neg E$	$\sum$	Recovery rate
C	20	20	40	50%
$\neg C$	16	24	40	40%
$\sum$	36	44	80	

Men	E	$\neg E$	$\sum$	Rec.rate	Women	E	$\neg E$	$\sum$	Rec.rate
C	18	12	30	60%	C	2	8	10	20%
$\neg C$	7	3	10	70%	$\neg C$	9	21	30	30%
	25	15	40			11	29	40	

$$P(E \mid C) \quad > \quad P(E \mid \neg C)$$
$$\text{but} \quad P(E \mid C, M) \quad < \quad P(E \mid \neg C, M)$$
$$P(E \mid C, W) \quad < \quad P(E \mid \neg C, W)$$

Philosophical topic, studied e.g. by Kant, Gärdenfors

Example for Focusing

- Prior knowledge: fair die

- New evidence: the result is an odd number

- Aposteriori knowledge via focusing: conditional probability

- Underlying probability measure did not change

Example for Revision

- Prior knowledge: fair die

- New evidence: weight near the 6

- Belief change via revision

- Underlying probability measure did change

Excursus: Causality vs. Correlation

Philosophical topic, studied e.g. by Aristoteles, still under discussion

Press acceleration pedal → car is faster   (causality)

Stork population high → human birthrate   (correlation, but no causality)

Visit doctor often → high risk of dying   (correlation, but no causality)

Probabilistic Reasoning

Probabilistic reasoning is difficult and may be problematic:

- $P(A \wedge B)$ is not determined simply by $P(A)$ and $P(B)$:
  $P(A) = P(B) = 0.5 \quad \Rightarrow \quad P(A \wedge B) \in [0, 0.5]$

- $P(C \mid A) = x, P(C \mid B) = y \quad \Rightarrow \quad P(C \mid A \wedge B) \in [0, 1]$
  Probabilistic logic is *not truth functional*!

Central problem: How does additional information affect the current knowledge?
I. e., if $P(B \mid A)$ is known, what can be said about $P(B \mid A \wedge C)$?

High complexity: n propositions $\rightarrow 2^n$ full conjunctives

Hard to specify these probabilities.

Summary

Uncertainty is inevitable in complex and dynamic scenarios that force agents to cope with ignorance.

Probabilities express the agent's inability to vote for a definitive decision. They model the degree of belief.

If an agent violates the axioms of probability, it may exhibit irrational behavior in certain circumstances.

The Bayes rule is used to derive unknown probabilities from present knowledge and new evidence.

Multiple evidences can be effectively included into computations exploiting conditional independencies.