Using Uncertain Knowledge

- Agents don't have complete knowledge about the world.
- Agents need to make decisions based on 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 $\Longrightarrow$ probability.

Probability

- Probability is an agent's measure of belief in some proposition — subjective probability.
- Example: Your probability of a bird flying is your measure of belief in the flying ability of an individual based only on the knowledge that the individual is a bird.
  - Other agents may have different probabilities, as they may have had different experiences with birds or different knowledge about this particular bird.
  - An agent's belief in a bird's flying ability is affected by what the agent knows about that bird.

Numerical Measures of Belief

- Belief in proposition, f, can be measured in terms of a number between 0 and 1 — this is the probability of f.
  - The probability f is 0 means that f is believed to be definitely false.
  - The probability f is 1 means that f is believed to be definitely true.
- Using 0 and 1 is purely a convention.
- f has a probability between 0 and 1, doesn't mean f is true to some degree, but means you are ignorant of its truth value. Probability is a measure of your ignorance.

Random Variables

- A random variable is a term in a language that can take one of a number of different values.
- The domain of a variable X, written $dom(X)$, is the set of values X can take.
- A tuple of random variables $\langle X_1, \ldots, X_n \rangle$ is a complex random variable with domain $dom(X_1) \times \cdots \times dom(X_n)$. Often the tuple is written as $X_1, \ldots, X_n$.
- Assignment $X = x$ means variable X has value x.
- A proposition is a Boolean formula made from assignments of values to variables.

Possible World Semantics

- A possible world specifies an assignment of one value to each random variable.
- $w \models X = x$
  means variable X is assigned value x in world w.
- Logical connectives have their standard meaning:

  $w \models \alpha \wedge \beta$ if $w \models \alpha$ and $w \models \beta$

  $w \models \alpha \vee \beta$ if $w \models \alpha$ or $w \models \beta$

  $w \models \neg\alpha$ if $w \not\models \alpha$
- Let Ω be the set of all possible worlds.

Semantics of Probability: finite case

For a finite number of possible worlds:

- Define a nonnegative measure $\mu(w)$ to each world w so that the measures of the possible worlds sum to 1.
  The measure specifies how much you think the world w is like the real world.
- The **probability** of proposition f is defined by:

$$P(f) = \sum_{w \models f} \mu(w).$$

Axioms of Probability: finite case

Four axioms define what follows from a set of probabilities:

Axiom 1 $0 \leq P(f)$ for any formula f.

Axiom 2 $P(\tau) = 1$ if τ is a tautology.

Axiom 3 $P(f \vee g) = P(f) + P(g)$ if $\neg(f \wedge g)$ is a tautology.

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

Semantics of Probability: general case

In the general case we have a measure on sets of possible worlds, satisfying:

- $\mu(S) \geq 0$ for $S \subseteq \Omega$
- $\mu(\Omega) = 1$
- $\mu(S_1 \cup S_2) = \mu(S_1) + \mu(S_2)$ if $S_1 \cap S_2 = \{\}$.
  Or sometimes σ-additivity:

$$\mu(\bigcup_i S_i) = \sum_i \mu(S_i) \text{ if } S_i \cap S_j = \{\}$$

Then $P(\alpha) = \mu(\{w | w \models \alpha\})$.

Probability Distributions

- A probability distribution on a random variable X is a function $dom(X) \rightarrow [0, 1]$ such that

$$x \mapsto P(X = x).$$

  This is written as $P(X)$.
- This also includes the case where we have tuples of variables. E.g., $P(X, Y, Z)$ means $P(\langle X, Y, Z\rangle)$.
- When $dom(X)$ is infinite sometimes we need a probability density function...

Conditioning

- Probabilistic conditioning specifies how to revise beliefs based on new information.
- You build a probabilistic model taking all background information into account. This gives the 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.
Evidence e induces a new measure, μ_e, over possible worlds

$$\mu_e(\omega) = \begin{cases} \frac{1}{P(e)} \times \mu(\omega) & \text{if } \omega \models e \\ 0 & \text{if } \omega \not\models e \end{cases}$$

The conditional probability of formula h given evidence e is

$$\begin{aligned} P(h|e) &= \sum_{\omega \models h} \mu_e(w) \\ &= \frac{P(h \wedge e)}{P(e)} \end{aligned}$$

Chain Rule

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

Bayes' theorem

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

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

If $P(e) \neq 0$, you can 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:
  $P(symptom \mid disease)$
  $P(light\ is\ off \mid status\ of\ switches\ and\ switch\ positions)$
  $P(alarm \mid fire)$
  $P($*image looks like*  $\mid a\ tree\ is\ in\ front\ of\ a\ car)$
- and want to do evidential reasoning:
  $P(disease \mid symptom)$
  $P(status\ of\ switches \mid light\ is\ off\ and\ switch\ positions)$
  $P(fire \mid alarm)$.
  $P(a\ tree\ is\ in\ front\ of\ a\ car \mid$ *image looks like* $)$