"States" = cells on 2d grid

actions = attempts to move N, S, E, W.

* noisy dynamics
* rewards = feedback from environment
  - delayed vs. immediate
  - evaluative vs. instructive

More generally: how can autonomous agents learn from experience?

⇒ Markov decision processes, reinforcement learning.

Other "embodied" agents: elevators, helicopters

Other "embedded" agents: game-playing system, spoken dialog system.

$\boxed{\text{Themes of class}}$

1) Probabilistic models of uncertainty
2) Learning as optimizations
3) Power vs. tractable: how to develop compact representation of complex world?
4) Principles vs. heuristics: $\left.\begin{array}{l}\text{optimizations}\\ \text{calculations}\end{array}\right\}$ vs. rules-of-thumb
5) Synergies of AI: inference and learning, perception and action, theory and pratice.

9/28 $\boxed{\text{Motivation}}$

* Modeling of uncertainty
  - Inherent randomness (e.g., radioactive decay)
  - Gross statistical description of complex deterministic world. (e.g., coin toss)
* Probability acts as guardian of commonsense reasoning.
* Many empirical successes: robotics, language, speech, bioinformatics.

$\boxed{\text{Review}}$

* Discrete random variable X

  Domain of possible values $\{x_1, x_2, \cdots, x_m\}$

  Ex: month M $\{m_1 = \text{Jan}, m_2 = \text{Feb}, \cdots, m_{12} = \text{Dec}\}$.
* Unconditional (prior) probability $P(X = x_i)$
* Basic axioms: (i) $P(X = x_i) \geq 0$

  (ii) $\sum_i P(X = x_i) = 1$

  (iii) $P(X = x_i \text{ or } X = x_j) = P(X = x_i) + P(X = x_j)$ if $x_i \neq x_j$

  Probs add for union of mutually exclusive events.

* Conditional (or posterior) probabilities

$P(X = x_i \mid Y = y_j)$ prob that $X = x_i$ given $Y = y_j$.

In general, $P(X = x_i \mid Y = y_j) \neq P(X = x_i)$

Ex: W = weather $\{w_1 = \text{sunny}, w_2 = \text{rainy}\}$

$$\left.\begin{aligned} &P(W = \text{sunny}) = 0.9 \\ &P(W = \text{sunny} \mid M = \text{aug}) = 0.97 \\ &P(W = \text{sunny} \mid M = \text{jan}) = 0.83 \end{aligned}\right\} \text{ conditional dependence}$$

Ex: conditional independence

Day of week D $\{d_1 = \text{sun}, d_2 = \text{mon}, \cdots, d_n = \text{sat}\}$

$P(W = \text{sunny} \mid D = \text{tues}) = P(W = \text{sunny}) = 0.9$.

Also true: (i) $P(X = x_i \mid Y = y_j) \geq 0$

(ii) $\sum_i P(X = x_i \mid Y = y_j) = 1$

Note sum over i, not j. $\sum_j P(X = x_i \mid Y = y_j)$ nothing you can say about this.

* Joint probability

$P(X = x_i, Y = y_j)$ = Prob that $X = x_i$ <u>and</u> $Y = y_j$.

(*) Product rule

For all i, j: $P(X = x_i, Y = y_j) = P(X = x_i \mid Y = y_j)\, P(Y = y_j)$

$P(X = x_i, Y = y_j) = P(Y = y_j \mid X = x_i)\, P(X = x_i)$

* Marginalization

$P(X = x_i) = \sum_j P(X = x_i, Y = y_j)$

$P(X = x_i, Y = y_j) = \sum_k P(X = x_i, Y = y_j, Z = z_k)$

* Intuitively, easier for experts to assess conditional probabilities than joint probabilities.

* Shorthand notation.

(i) Implied universality

$P(X, Y) = P(X \mid Y)\, P(Y) = P(Y \mid X)\, P(X)$ implies that equality holds for all assignments x_i, y_j.

(*)

(ii) Implied assignment

$P(x, y, z) = P(X = x, Y = y, Z = z)$

* Generalized product rule.

$P(A, B, C, D, \cdots) = P(A)\, P(B \mid A)\, P(C \mid A, B)\, P(D \mid A, B, C) \ldots$

* Bayes rule:

From (*) $\longrightarrow$ $P(X \mid Y) = \dfrac{P(Y \mid X)\, P(X)}{P(Y)}$

* Generalized Bayes rule:

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

Example

* Binary random variables

B = burglary

E = earthquake

A = alarm.

* Joint distribution

$$P(B, E, A) = P(B)\, P(E \mid B)\, P(A \mid B, E).$$

* Prior knowledge

$P(B=1) = 0.001 \longleftrightarrow P(B=0) = 1 - P(B=1) = 0.999$

$P(E=1 \mid B=0) = 0.002$
$P(E=1 \mid B=1) = 0.002$ $\Big\}$ $P(E=1 \mid B) = P(E=1) = 0.002$ (conditional independence)

B	E	P(A=1 | B, E)	P(A=0 | B, E)
0	0	0.001	⋮
1	0	0.94	1 − P(A=1 | B,E)
0	1	0.29	⋮
1	1	0.95	⋮

* Inferences

Do these probabilities capture commonsense reasoning?

Compare $P(B=1)$, $P(B=1 \mid A=1)$, and $P(B=1 \mid A=1, E=1)$?

$P(B=1) = 0.001$.

$P(B=1 \mid A=1) = P(A=1 \mid B=1)\, P(B=1) \,/\, P(A=1)$: Bayes rule.

$$\begin{aligned}
P(A=1 \mid B=1) &= \sum_{e \in \{0,1\}} P(A=1, E=e \mid B=1) && \text{: conditional marginalization} \\
&= \sum_e P(A=1 \mid E=e, B=1)\, P(E=e \mid B=1) && \text{: generalized product rule.} \\
&= \sum_e P(A=1 \mid E=e, B=1)\, P(E=e) && \text{: conditional independence} \\
&= P(A=1 \mid E=1, B=1)\, P(E=1) + P(A=1 \mid E=0, B=1)\, P(E=0) \\
&= (0.95)(0.002) + (0.94)(0.998) \\
&= 0.94002
\end{aligned}$$

$$P(A=1) = \sum_{e,b} P(A=1, E=e, B=b) \quad : \text{marginalization}$$
$$= \sum_{e,b} P(A=1 \mid E=e, B=b)\, P(E=e \mid B=b)\, P(B=b) \quad : \text{product rule}$$
$$= \sum_{e,b} P(A=1 \mid E=e, B=b)\, P(E=e)\, P(B=b) \quad : \text{conditional independence}$$
$$= \underbrace{(0.95)(0.002)(0.001)}_{\{e=b=1\}} + \underbrace{\cdots}_{\text{(other cases)}}$$
$$= 0.00252.$$

From Bayes rule: $P(B=1 \mid A=1) = P(A=1 \mid B=1)\, P(B=1) \,/\, P(A=1)$

$$= (0.94002)(0.001) / (0.00252)$$

$$\boxed{P(B=1 \mid A=1) = 0.37}$$

Compare to $\boxed{P(B=1) = 0.001}$

Now, compute: $P(B=1 \mid A=1, E=1) = \dfrac{P(A=1 \mid B=1, E=1)\, P(B=1 \mid E=1)}{P(A=1 \mid E=1)}$

Conditional Bayes rule

$P(B=1)$: conditional independence

$$= (0.95)(0.001) / P(A=1 \mid E=1)$$

$P(A=1 \mid E=1) = P(A=1, E=1) / \underbrace{P(E=1)}_{(0.002)}$: product rule

$$P(A=1, E=1) = \sum_{b} P(A=1, E=1, B=b) \quad : \text{marginalization}$$
$$= \sum_{b} P(A=1 \mid E=1, B=b)\, P(E=1 \mid B=b)\, P(B=b) \quad : \text{product rule}$$
$$= \sum_{b} P(A=1 \mid E=1, B=b)\, P(E=1)\, P(B=b) \quad : \text{conditional independence}$$

(summing over $b = \{0,1\}$ and substituting ...)

$$= 0.00058.$$

Pluggin in: $P(B=1 \mid A=1, E=1) = \dfrac{(0.95)(0.001)}{\left(\frac{0.00058}{0.002}\right)} = 0.0033.$

In sum: $P(B=1) = 0.001$

$P(B=1 \mid A=1) = 0.37 \uparrow$

$P(B=1 \mid A=1, E=1) = 0.0033 \downarrow$

Example of non-monotonic reasoning

$P(B=1 \mid A=1, E=1) < P(B=1 \mid A=1)$

The earthquake "explains away" the alarm, thus decreasing our belief in burglary.