Last Time:

- ☐ RL summary
- ☐ POMDPs

Lecture 5

HRI, FALL '25

Andrea Bajcsy

This Time:

- ☐ Probability Recap
- ☐ Boltzmann Rationality

Uncertainty:

We are starting to see that alot of things in life are uncertain, including humans!

When we interact w/ other people, or the world, we may not know what their goals/intents/preferences/rewards are, even if we could perfectly observe their "physical" state!

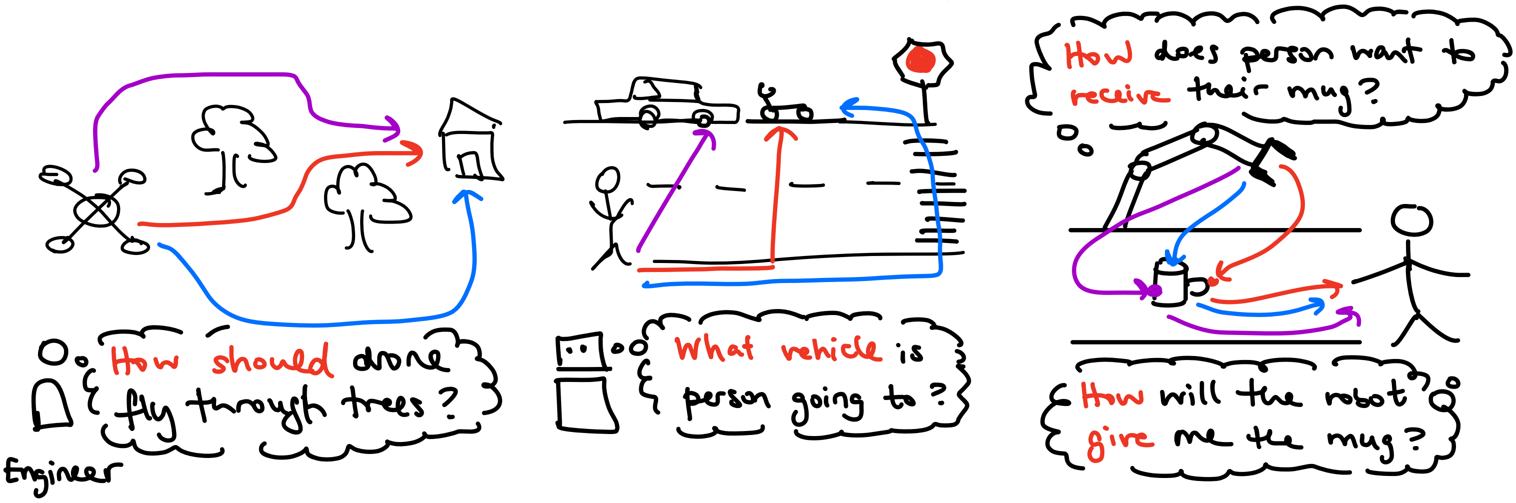

Let's model this type of uncertainty & recap some basic probability theory along the way.

Probability Recap

In probability theory, we want to make probabilistic assertions about possible "worlds" – possible human behavior, or environment responses ...

NOTATION:

- P(X) probability of random variable X

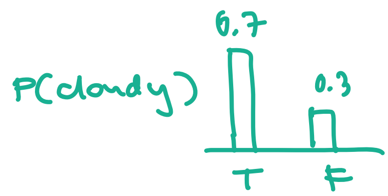

ex. cloudy ∈ {True, False}
↳ random variable

- $P(X, Y)$ probability of X and Y. Also called joint distribution

  "and"

  ex. $P(\text{cloudy}, \text{rain})$

- $P(X \mid Y)$ probability of x given Y. Also called conditional distribution.

  "given"

  ex. $P(\text{cloudy} \mid \text{rain})$

USEFUL RULES

- Product Rule (or chain rule):

$$P(x, y) = P(X \mid Y) \cdot P(Y) = P(Y \mid X) \cdot P(X)$$

ex. for it to be cloudy and rain (i.e. $P(\text{cloudy}, \text{rain})$), we need it to rain (i.e. $P(\text{rain})$) and we need it to be cloudy given the rain (i.e. $P(\text{cloudy} \mid \text{rain})$)

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

- Marginalization (i.e. "summing out"):

$$P(X) = \sum_{y \in Y} P(X, y)$$

sums up probabilities of each possible value of a target variable (i.e. Y in this case), thus taking them out of the equation.

ex.

	rain = T	rain = F
cloud = T	0.5	0.05
cloud = F	0.15	0.30

$P(\text{cloudy}, \text{rain})$

Cloud = T	0.55
cloud = F	0.45

marginalize out "rain"

$$P(\text{cloudy}) = \sum_{\text{rain} \in \{T, F\}} P(\text{cloudy}, \text{rain})$$

- <u>Independence</u>: when one random variable's probability doesn't depend on another.

$$P(X, Y) = P(X \mid Y)\, P(Y) \quad \leftarrow \text{via product rule}$$
$$= P(X) \cdot P(Y)$$

if X & Y are independent $P(X \mid Y) = P(X)$

<u>ex.</u> P(cloudy, heads) = P(cloudy) P(heads)

- <u>conditional independence</u>: if x and y are conditionally independent <u>given</u> Z then:

$$P(X, Y \mid Z) = P(X \mid Y, Z) \cdot P(Y \mid Z) \quad \leftarrow \text{Product Rule}$$
$$= P(X \mid Z) \cdot P(Y \mid Z)$$

if X & Y are indep. <u>given Z</u> then: $P(X \mid Y, Z) = P(X \mid Z)$

⇒ once we know Z, knowing Y doesn't tell me anything new abt. X.

<u>ex.</u> P(traffic, umbrella | rain) = P(traffic | rain) · P(umbrella | rain)

$\boxed{\text{BUT}}$ P(traffic, umbrella) ≠ P(traffic) P(umbrella)
b/c these two are <u>correlated</u> via the presence/absence of rain!

$P(X \mid Y) \Rightarrow P(\text{action} \mid \text{goal})$

BAYES THEOREM

ex. action ex. goal

What if we know the probability of X given Y, but we actually want to know probability of Y given X?

$P(Y \mid X) \Rightarrow P(\text{goal} \mid \text{action})$

Bayes' Theorem lets us establish a relationship btwn these two!

$$P(\underbrace{Y}_{\text{goal}} \mid \underbrace{X}_{\text{action}}) = \frac{P(\underbrace{X}_{\text{action}} \mid \underbrace{Y}_{\text{goal}}) \cdot P(\underbrace{Y}_{\text{goal}})}{P(\underbrace{X}_{\text{action}})}$$

Q How do we derive it?

$$P(X,Y) = P(Y|X) \cdot P(X)$$ ← product rule

⇒ divide by $P(X)$

$$\frac{P(X,Y)}{P(X)} = P(Y|X) \;\Rightarrow\; \frac{P(X|Y) \cdot P(Y)}{P(X)} = P(Y|X) \;\checkmark$$

⇒ product rule on numerator

Each component of Bayes's Rule has a name:

likelihood function or observation model

$$\underbrace{P(Y|X)}_{\text{posterior}} = \frac{\overbrace{P(X|Y)}^{\text{likelihood}} \cdot \overbrace{P(Y)}^{\text{prior}}}{\underbrace{P(X)}_{\text{normalizer}}}$$

This rule ends up being very useful for things like state estimation & inferring properties of human behavior!

Example: we want to infer which vehicle a pedestrian wants to move to (A or B). Let's use Bayes' Thm!

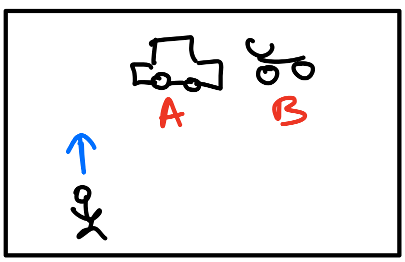

PRIOR

goal = A	goal = B
0.5	0.5

$P(goal)$

$\mathcal{A} = \{\uparrow, \rightarrow\}$

OBSERVATION MODEL

$P(a \mid goal)$

$P(a=\uparrow \mid g=A) = 0.8$

$P(a=\rightarrow \mid g=A) = 0.2$

$P(a=\uparrow \mid g=B) = 0.4$

$P(a=\rightarrow \mid g=B) = 0.6$

YOU OBSERVE $\boxed{a=\uparrow}$

Q COMPUTE POSTERIOR: $P(g \mid a=\uparrow) = ?$

$$P(goal \mid a=\uparrow) = \frac{P(a=\uparrow \mid goal)\, P(goal)}{P(a=\uparrow)}$$ // Bayes' Rule

$$= \frac{P(a=\uparrow \mid goal)\, P(goal)}{\sum_{goal \in \{A,B\}} P(a=\uparrow, goal)}$$ // marginalize over goals

$$= \frac{P(a=\uparrow \mid goal)\, P(goal)}{\sum_{goal \in \{A,B\}} P(a=\uparrow \mid goal)\, P(goal)}$$ // product rule

$P(goal)$ prior

goal = A	goal = B
0.5	0.5

$P(a \mid goal)$

$P(a=\uparrow \mid goal=A) = 0.8$
$P(a=\rightarrow \mid goal=A) = 0.2$
$P(a=\uparrow \mid goal=B) = 0.4$
$P(a=\rightarrow \mid goal=B) = 0.6$

$P(goal \mid a=\uparrow)$ posterior

goal = A	goal = B
$0.6\overline{6}$	$0.3\overline{3}$

goal A:

$$P(goal=A \mid a=\uparrow) = \frac{\overset{0.8}{P(a=\uparrow \mid goal=A)}\, \overset{0.5}{P(goal=A)}}{\underset{0.8}{P(a=\uparrow \mid goal=A)}\, \underset{0.5}{P(goal=A)} + \underset{0.4}{P(a=\uparrow \mid goal=B)}\, \underset{0.5}{P(goal=B)}}$$

$= 0.6\overline{6}$

goal B:

$P(goal=B \mid a=\uparrow) = 1 - P(goal=A \mid a=\uparrow) = 0.3\overline{3}$ ← law of total prob.