Lecture 6

HRI, FALL '25

Andrea Bajcsy

Last Time:

- ☐ Probability Recap
- ☐ Bayes' Rule

This Time:

- ☐ Boltzmann Rationality
- ☐ Intent Inference & Expression

Last time: Introduced the idea + the utility of Bayes' Rule:

(goal: Y; action, trajectory: X)

$$P(Y|X) = \frac{P(X|Y)P(Y)}{P(X)} = \frac{P(X|Y)P(Y)}{\sum_Y P(X|Y)P(Y)}$$

($P(X|Y)$: action → goal)

KEY IDEA: Humans act in ways driven by their goals!

↳ robot can use this model to infer the H's goal / intent....

Q Where exactly does this probability model $P(X|Y)$ come from? (action: X, goal: Y)

A cognitive psychology, econometrics, data-driven modeling, common sense / heuristics ....

↳ econometrics: application of statistical methods for economic relations

Boltzmann Rational Model [Luce 1959, 1977; Ziebart 2010]

Assume that people are optimizing some reward function in their mind (e.g. dist to goal). This model says that people behave approximately optimally (i.e. approx. rationality) in pursuit of their goals:

action $(\leftarrow, \uparrow, \rightarrow, \downarrow)$

$$P(a \mid s; \text{goal}) = \frac{e^{Q(s,a;g)}}{\sum_{\bar{a} \in A} e^{Q(s,\bar{a};g)}}$$

$\bar{a} \in A$ ← set of all actions

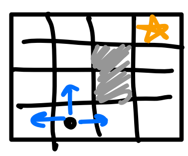

⇒ the higher the value of an action a, the exponentially more likely it is that the human will take this action

Another variation of this model you will see alot in the class is a "trajectory-level model":

cumulative reward $R(\xi) := \sum_{(s,a) \in \xi} r(s,a)$

recall, traj: $\xi = (s_0, a_0, s_1, a_1, \ldots)$

$$P(\xi \mid \text{goal}) = \frac{e^{R(\xi;\text{goal})}}{\sum_{\bar{\xi} \in \Xi} e^{R(\bar{\xi};\text{goal})}}$$

$\bar{\xi} \in \Xi$ ← set of all trajectories

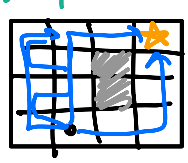

Why does intent matter?

⇒ intent <u>inference</u>: R wants to infer H's intent, so R observes the H and based on the observation sequence, R infers what the human is doing

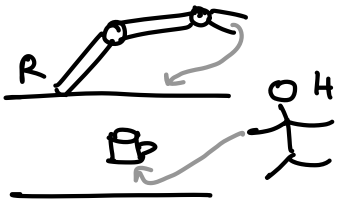

⇒ intent <u>expression</u>: H is also trying to infer the R's intent; so here, the R can choose its actions in a way that makes its own goals clear/"obvious" to H.

GOAL: If we want H & R's to work together seamlessly, they need to "<u>read</u>" each other's intentions and they need to "<u>express</u>" their own intentions clearly.

⇒ communication problem.

INTENT INFERENCE → "Integrating Human Observer Inference into Robot Motion Planning", Dragan and Srinivasa

<u>Goal</u>: Infer H's intent from their observed trajectory.

To start, assume we want to infer their (discrete) goal g and we observe a partial state trajectory:

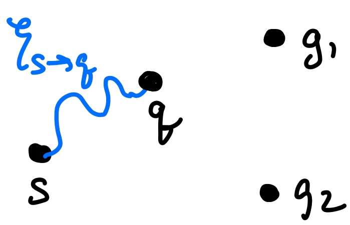

$$\xi_{s\to q} := (s, s_1, s_2, \ldots, q)$$

Q: What goal would make $\xi_{s\to q}$ likely?

A: $$\hat{g} = \arg\max_{g_i \in \{1,2\}} P(g_i \mid \xi_{s\to q}) \; \star$$

↳ maximum a posteriori estimate (MAP)

<u>How to solve?</u> Bayes' Rule! We know it is easier to reason abt. $P(\xi_{s\to q} \mid g)$, and so we can expand w/ Bayes Rule:

★ $P(g \mid \xi_{s \to q}) = \dfrac{P(\xi_{s \to q} \mid g) \cdot P(g)}{\sum_{\bar{g} \in \{1,2\}} P(\xi_{s \to q} \mid \bar{g}) \cdot P(\bar{g})}$

→ let's use Boltzmann Model here! (for $P(\xi_{s \to q} \mid g)$)

→ Prior (for $P(g)$)

INTENT EXPRESSION

GOAL: R has a goal, g^*, and it wants to choose a trajectory, ξ, such that H can easily infer g^*.

⇒ Now, roles are flipped! Human (observer) is computing $P(g \mid \xi_{s \to q})$ in their mind, and the robot picks $\xi_{s \to q}$ to maximize the probability that H infers g^*!

$$\xi^* := \arg\max_{\xi} \int P(g = g^* \mid \xi_{s \to \xi(t)})\, dt$$

- ξ^* ↑ express intent early and consistently!
- $P(g = g^* \mid \cdot)$ ↑ Human's posterior over robot's goal, evaluated @ true R goal g^*
- $\xi_{s \to \xi(t)}$ ← index into trajectory @ time t

VARIANTS:

- intent inference: $\arg\max_{g} P(g \mid \xi)$ (search over goals)
  ↳ = guess goal given trajectory
  → normalize over space of g: $\dfrac{\quad}{\sum_{g} \dots}$
- predictability: $\arg\max_{\xi} P(\xi \mid g)$ (search over traj's)
  ↳ = choose most likely traj. given goal
  → normalize over traj.: $\dfrac{\quad}{\sum_{\xi \in \Xi}}$ OR $\dfrac{\quad}{\int \dots d\xi}$
- legibility: $\arg\max_{\xi} P(g \mid \xi)$ (search over traj.)
  ↳ = choose traj s.t. goal is easily inferred!
  ↳ normalizing over goals: $\dfrac{\quad}{\sum_{g} \dots}$

Example:

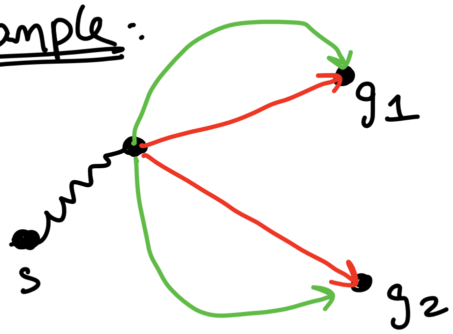

predictable = what you would expect if you knew the goal in advance

legible = what you'd do to help someone figure out your goal

predictability optimizes for efficiency & doesn't care if the human can guess the goal

legibility optimizes for communicativeness and may sacrifice efficiency in favor of helping observer (H) guess goal faster.