Autonomous Navigation for Flying Robots

Lecture 5.3:
Reasoning with Bayes Law

Jürgen Sturm
Technische Universität München

The State Estimation Problem

We want to estimate the world state $\mathbf{x}$ from

1. Sensor measurements $\mathbf{z}$ and
2. Controls (or odometry readings) $\mathbf{u}$

We need to model the relationship between these random variables, i.e.,

$$p(\mathbf{x} \mid \mathbf{z}) \qquad\qquad p(\mathbf{x}' \mid \mathbf{x}, \mathbf{u})$$

Causal vs. Diagnostic Reasoning

- $P(\mathbf{x} \mid \mathbf{z})$ is diagnostic
- $P(\mathbf{z} \mid \mathbf{x})$ is causal
- Diagnosic reasoning is typically what we need
- Often causal knowledge is easier to obtain
- Bayes rule allows us to use causal knowledge in diagnostic reasoning

Bayes Rule

- Definition of conditional probability

$$P(x, z) = P(x \mid z)P(z) = P(z \mid x)P(x)$$

- Bayes rule

$$P(x \mid z) = \frac{\overbrace{P(z \mid x)}^{\text{observation likelihood}}\overbrace{P(x)}^{\text{prior on world state}}}{\underbrace{P(z)}_{\text{prior on sensor observations}}}$$

Normalization

- Direct computation of $P(\mathbf{z})$ can be difficult
- Idea: Compute improper distribution, normalize afterwards

- Step 1: $L(x \mid z) = P(z \mid x)P(x)$

- Step 2: $P(z) = \sum_x P(z, x) = \sum_x P(z \mid x)P(x) = \sum_x L(x \mid z)$

- Step 3: $P(x \mid z) = L(x \mid z)/P(z)$

Background Knowledge

- Same derivation also works in the presence of background knowledge

$$P(x, y \mid z) = P(x \mid y, z) P(y \mid z) = P(y \mid x, z) P(x \mid z)$$

- Bayes rule with background knowledge

$$P(x \mid y, z) = \frac{P(y \mid x, z) P(x \mid z)}{P(y \mid z)}$$

Example: Sensor Measurement

- Quadrotor seeks the landing zone
- Landing zone is marked with many bright lamps
- Quadrotor has a brightness sensor

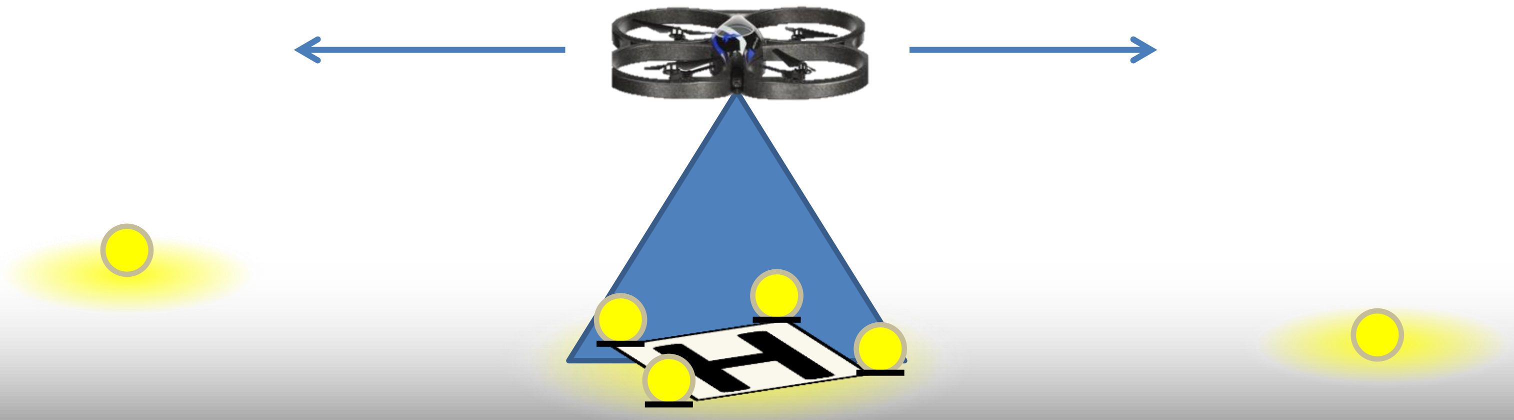

Example: Sensor Measurement

- Binary sensor $Z \in \{\text{bright}, \neg\text{bright}\}$
- Binary world state $X \in \{\text{home}, \neg\text{home}\}$
- Sensor model $P(Z = \text{bright} \mid X = \text{home}) = 0.6$

  $P(Z = \text{bright} \mid X = \neg\text{home}) = 0.3$
- Prior on world state $P(X = \text{home}) = 0.5$
- Assume: Robot observes light, i.e., $Z = \text{bright}$
- What is the probability $P(X = \text{home} \mid Z = \text{bright})$ that the robot is above the landing zone?

Example: Sensor Measurement

- Sensor model $P(Z = \text{bright} \mid X = \text{home}) = 0.6$

  $P(Z = \text{bright} \mid X = \neg\text{home}) = 0.3$
- Prior on world state $P(X = \text{home}) = 0.5$
- Probability after observation (using Bayes)

$$P(X = \text{home} \mid Z = \text{bright})$$

$$= \frac{P(\text{bright} \mid \text{home})P(\text{home})}{P(\text{bright} \mid \text{home})P(\text{home}) + P(\text{bright} \mid \neg\text{home})P(\neg\text{home})}$$

$$= \frac{0.6 \cdot 0.5}{0.6 \cdot 0.5 + 0.3 \cdot 0.5} = \frac{0.3}{0.3 + 0.15} = 0.67$$

Combining Evidence

- Suppose our robot obtains another observation z_2 (either from the same or a different sensor)
- How can we integrate this new information?
- More generally, how can we estimate $p(x \mid z_1, z_2, \dots)$?

Combining Evidence

- Suppose our robot obtains another observation z_2 (either from the same or a different sensor)
- How can we integrate this new information?
- More generally, how can we estimate $p(x \mid z_1, z_2, \dots)$?

- Bayes formula (with background knowledge) gives us

$$P(x \mid z_1, \dots, z_n) = \frac{P(z_n \mid x, z_1, \dots, z_{n-1}) P(x \mid z_1, \dots, z_{n-1})}{P(z_n \mid z_1, \dots, z_{n-1})}$$

Recursive Bayesian Updates

$$P(x \mid z_1, \ldots, z_n) = \frac{P(z_n \mid x, z_1, \ldots, z_{n-1}) P(x \mid z_1, \ldots, z_{n-1})}{P(z_n \mid z_1, \ldots, z_{n-1})}$$

Markov Assumption:
z_n is independent of $z_1, \ldots, z_{n-1}$ if we know x

$$\Rightarrow \quad P(x \mid z_1, \ldots, z_n) = \frac{P(z_n \mid x) P(x \mid z_1, \ldots, z_{n-1})}{P(z_n \mid z_1, \ldots, z_{n-1})}$$

$$= \eta P(z_n \mid x) P(x \mid z_1, \ldots, z_{n-1})$$

$$= \eta_{1:n} \prod_{i=1,\ldots,n} P(z_i \mid x) P(x)$$

Example: Sensor Measurement

- Quadrotor seeks the landing zone
- Landing zone is marked with many bright lamps **and a visual marker**

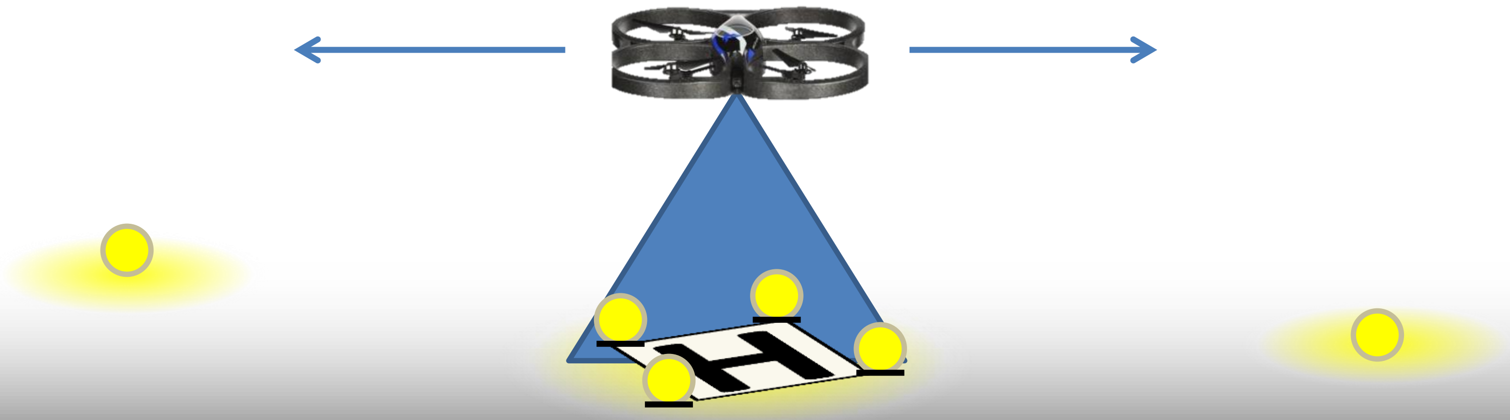

Example: Second Measurement

- Sensor model $P(Z_2 = \text{marker} \mid X = \text{home}) = 0.8$
  $P(Z_2 = \text{marker} \mid X = \neg\text{home}) = 0.1$
- Previous estimate $P(X = \text{home} \mid Z_1 = \text{bright}) = \boxed{0.67}$
- Assume robot does not observe marker
- What is the probability of being home?

$$P(X = \text{home} \mid Z_1 = \text{bright}, Z_2 = \neg\text{marker})$$

$$= \frac{P(\neg\text{marker} \mid \text{home}) P(\text{home} \mid \text{bright})}{P(\neg\text{marker} \mid \text{home}) P(\text{home} \mid \text{bright}) + P(\neg\text{marker} \mid \neg\text{home}) P(\neg\text{home} \mid \text{bright})}$$

$$= \frac{0.2 \cdot 0.67}{0.2 \cdot 0.67 + 0.9 \cdot 0.33} = \boxed{0.31}$$

The second observation lowers the probability that the robot is above the landing zone!

Actions/Controls (Motions)

- World state changes over time because of
  - actions carried out by the robot…
  - actions carried out by other agents…
  - or just time passing by…

  …change the world

- How can we incorporate actions?

Typical Actions

- Quadrotor accelerates by changing the speed of its motors
- Position also changes when quadrotor does “nothing” (and drifts..)

- Actions are never carried out with absolute certainty
- In contrast to measurements, actions generally increase the uncertainty of the state estimate

Action Models

- To incorporate the outcome of an action u into the current state estimate ("belief"), we use the conditional pdf

$$p(x' \mid u, x)$$

- This term specifies the probability that executing the action u in state x will lead to state x'

Example: Take-Off

- Action: $u \in \{\text{takeoff}\}$
- World state: $x \in \{\text{ground}, \text{air}\}$

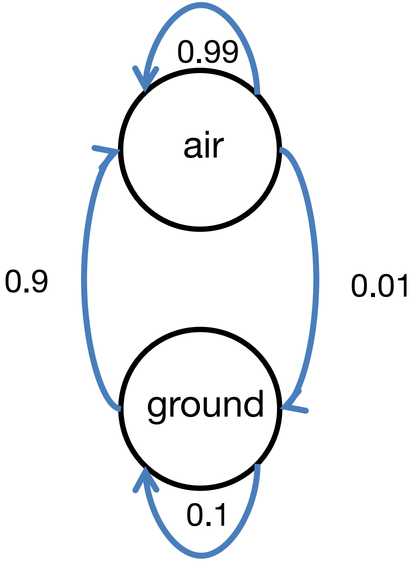

Integrating the Outcome of Actions

- Discrete case

$$P(x' \mid u) = \sum_x P(x' \mid u, x) P(x)$$

- Continuous case

$$p(x' \mid u) = \int p(x' \mid u, x) p(x) \mathrm{d}x$$

Example: Take-Off

- Prior belief on robot state: $P(x = \text{ground}) = 1.0$
  (robot is located on the ground)
- Robot executes "take-off" action
- What is the robot's belief after one time step?

$$
\begin{aligned}
P(x' = \text{ground}) = \sum_x P(x' = \text{ground} \mid u, x) P(x) \\
=& P(x' = \text{ground} \mid u, x = \text{ground}) P(x = \text{ground}) \\
+& P(x' = \text{ground} \mid u, x = \text{air}) P(x = \text{air}) \\
=& 0.1 \cdot 1.0 + 0.01 \cdot 0.0 = 0.1
\end{aligned}
$$

Lessons Learned

- Bayes rule
- Data fusion of sensor observations
- Data fusion of actions/motion commands

- Next:
  Bayes filter