Drones Demystified!

K. Alexis, C. Papachristos, Autonomous Robots Lab, University of Nevada, Reno

A. Tzes, Autonomous Robots & Intelligent Systems Lab, NYU Abu Dhabi

Drones Demystified!

Topic: State Estimation

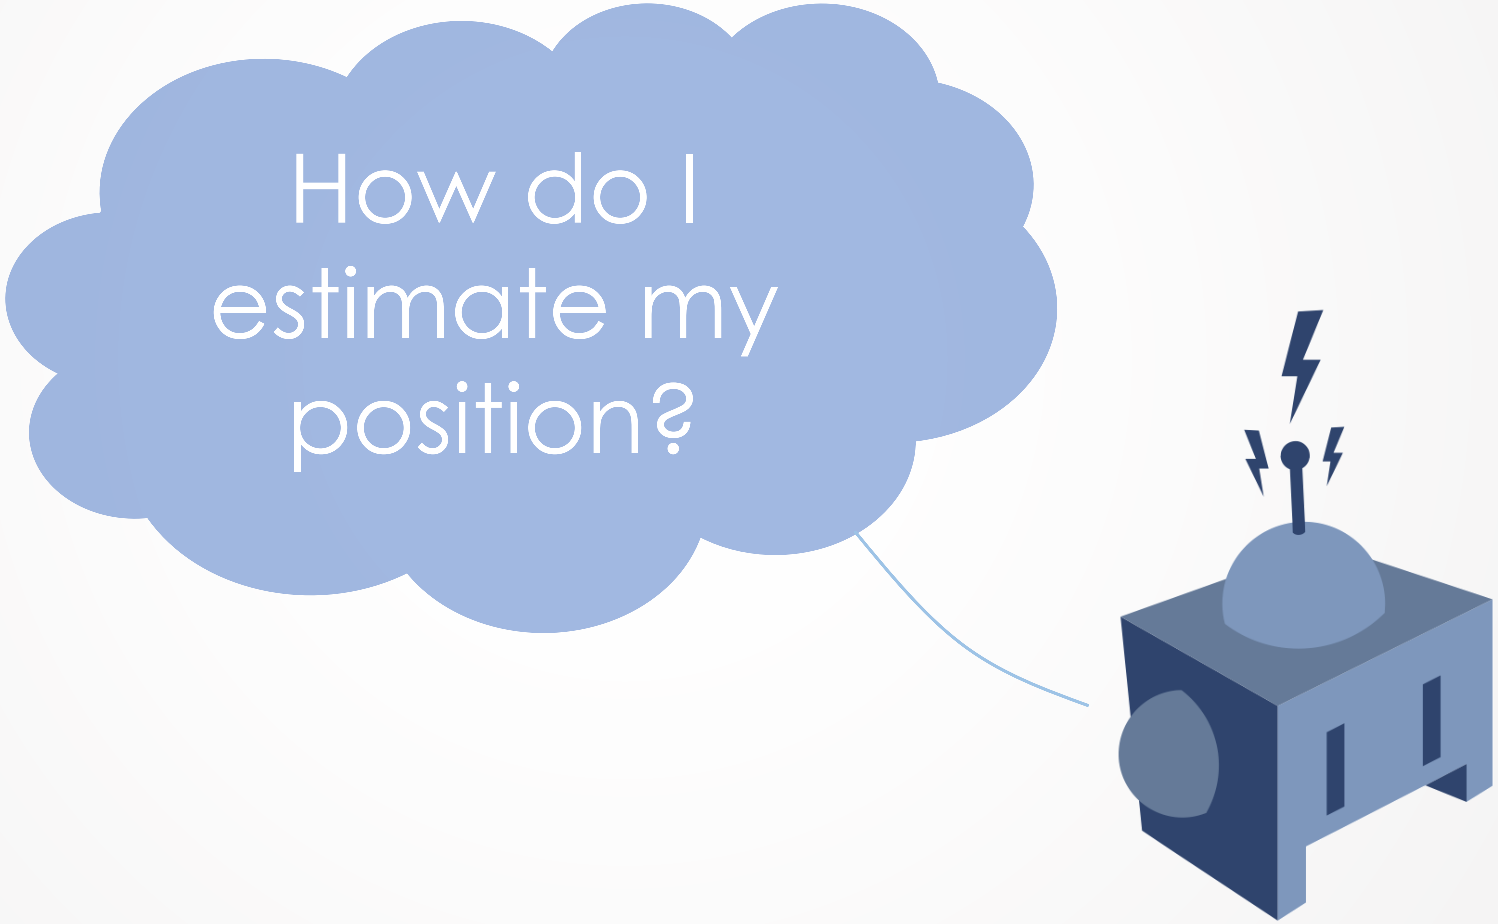
How do I estimate my position?

Drones Demystified!

Topic: State Estimation – Reasoning with Bayes Law

The State Estimation problem

- We want to estimate the world state x from:
  - Sensor measurements z and
  - Controls u

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

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

Causal vs. Diagnostic Reasoning

$P(\mathbf{x}|\mathbf{z})$ Is diagnostic

$P(\mathbf{z}|\mathbf{x})$ Is causal

- Diagnostic 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**:

$$P(x, z) = P(x|z)P(z) = P(z|x)P(x)$$

- **Bayes rule:**

$$P(x|z) = \frac{P(z|x)P(x)}{P(z)}$$

Observation likelihood

Prior on world state

Prior on sensor observations

Normalization

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

- **STEP 1:** $$L(x|z) = P(z|x)P(x)$$
- **STEP 2:** $$P(z) = \sum_x P(z,x) = \sum_x P(z|x)P(x) = \sum_x L(x|z)$$
- **STEP 3:** $$P(x|z) = L(x|z)/P(z)$$

Normalization

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

- STEP 1: $$L(x|z) = P(z|x)P(x)$$
- STEP 2: $$P(z) = \sum_x P(z,x) = \sum_x P(z|x)P(x) = \sum_x L(x|z)$$
- STEP 3: $$P(x|z) = L(x|z)/P(z)$$

Example: Sensor Measurement

- Quadrotor seeks the Landing Zone
- The landing zone is marked with many bright lamps
- The quadrotor has a light sensor.

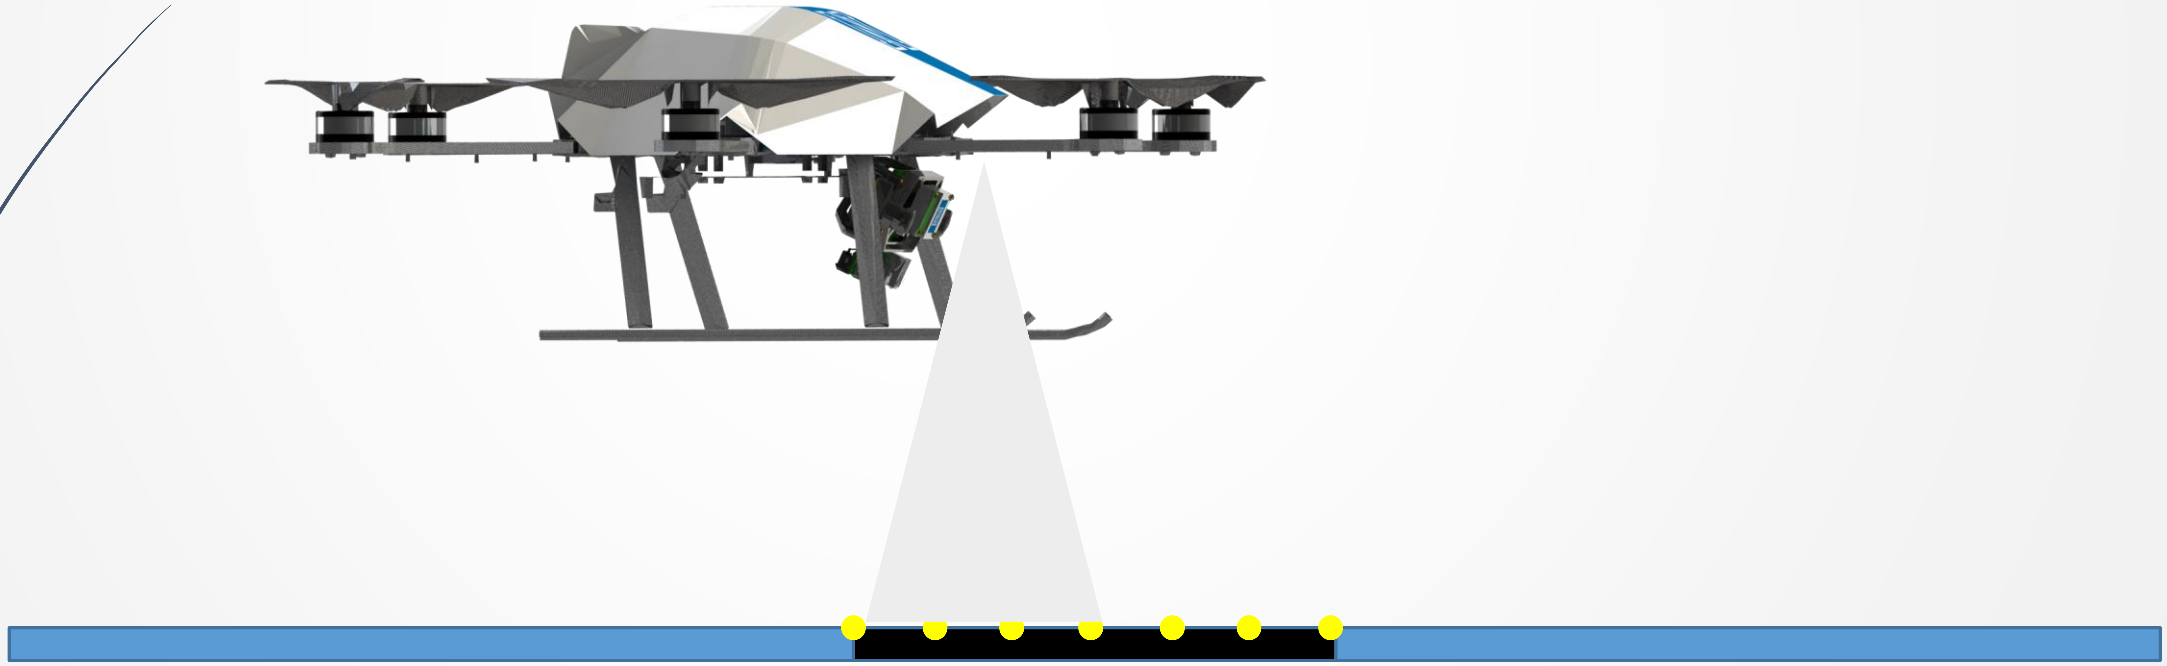

Example: Sensor Measurement

- Binary sensor $Z \in \{bright, \bar{bright}\}$
- Binary world state $X \in \{home, \bar{home}\}$
- Sensor model $$P(Z = bright | X = home) = 0.6$$ $$P(Z = bright | X = \bar{home}) = 0.3$$
- Prior on world state $P(X = home) = 0.5$
- Assume: robot observes light, i.e. $Z = bright$
- What is the probability $P(X = home | Z = bright)$ that the robot is above the landing zone.

Example: Sensor Measurement

- Sensor model: $$P(Z = bright | X = home) = 0.6$$ $$P(Z = bright | X = \bar{home}) = 0.3$$
- Prior on world state: $$P(X = home) = 0.5$$
- Probability after observation (using Bayes):

$$P(X = home | Z = bright) = \frac{P(bright|home)P(home)}{P(bright|home)P(home) + P(bright|\bar{home})P(\bar{home})} = \frac{0.6 \cdot 0.5}{0.6 \cdot 0.5 + 0.3 \cdot 0.5} = 0.67$$

Actions (Motions)

- Often the world is dynamic since
  - Actions are carried out by the robot
  - Actions are carried out by other agents
  - Or simply because time is passing and the world changes
- How can we incorporate actions?

Example actions

- MAV accelerates by changing the speed of its motors.
- The ground robot moves due to it being on an inclined terrain.

- Actions are never carried out with absolute certainty: leave a quadrotor hover and see it drifting!
- 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 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: $x \in \{\text{ground}, \text{air}\}$

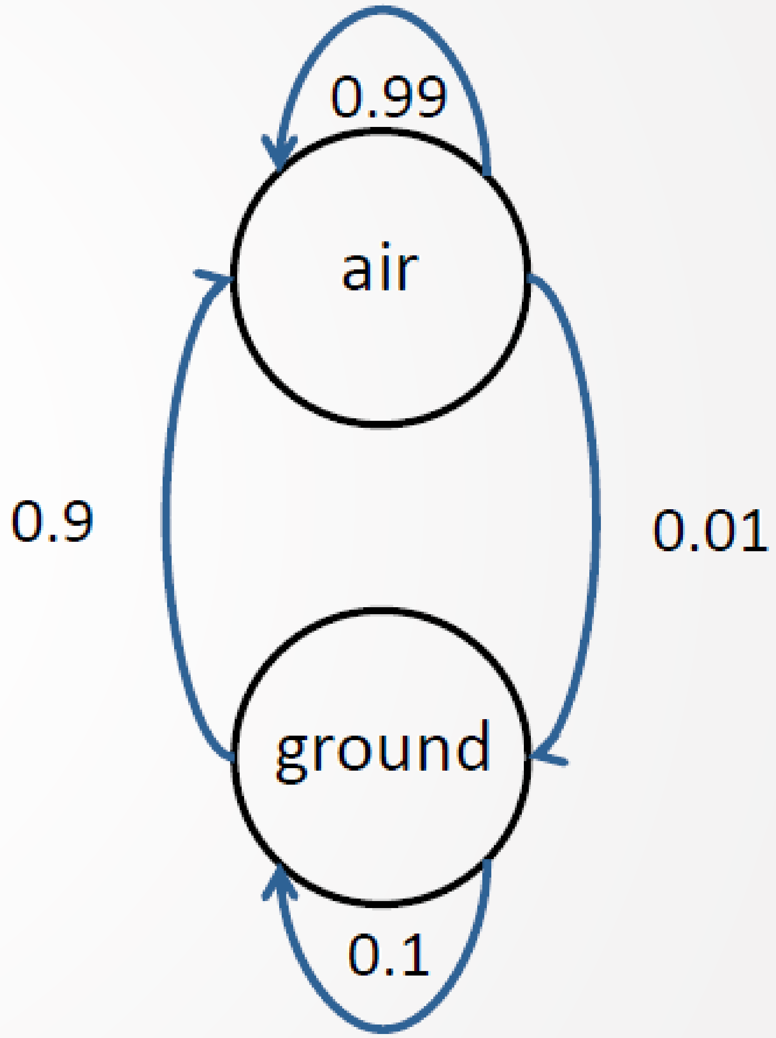

Integrating the Outcome of Actions

- 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 executes "take-off" action
- What is the robot's belief after one time step?

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

Drones Demystified!

Topic: State Estimation – Bayes Filter

Markov Assumption

- Observations depend only on current state

$$P(z_t | x_{0:t}, z_{1:t-1}, u_{1:t}) = P(z_t | x_t)$$

- Current state depends only on previous state and current action

$$P(x_t | x_{0:t}, z_{1:t}, u_{1:t}) = P(x_t | x_{t-1}, u_t)$$

Markov Chain

- A Markov Chain is a stochastic process where, given the present state, the past and the future states are independent.

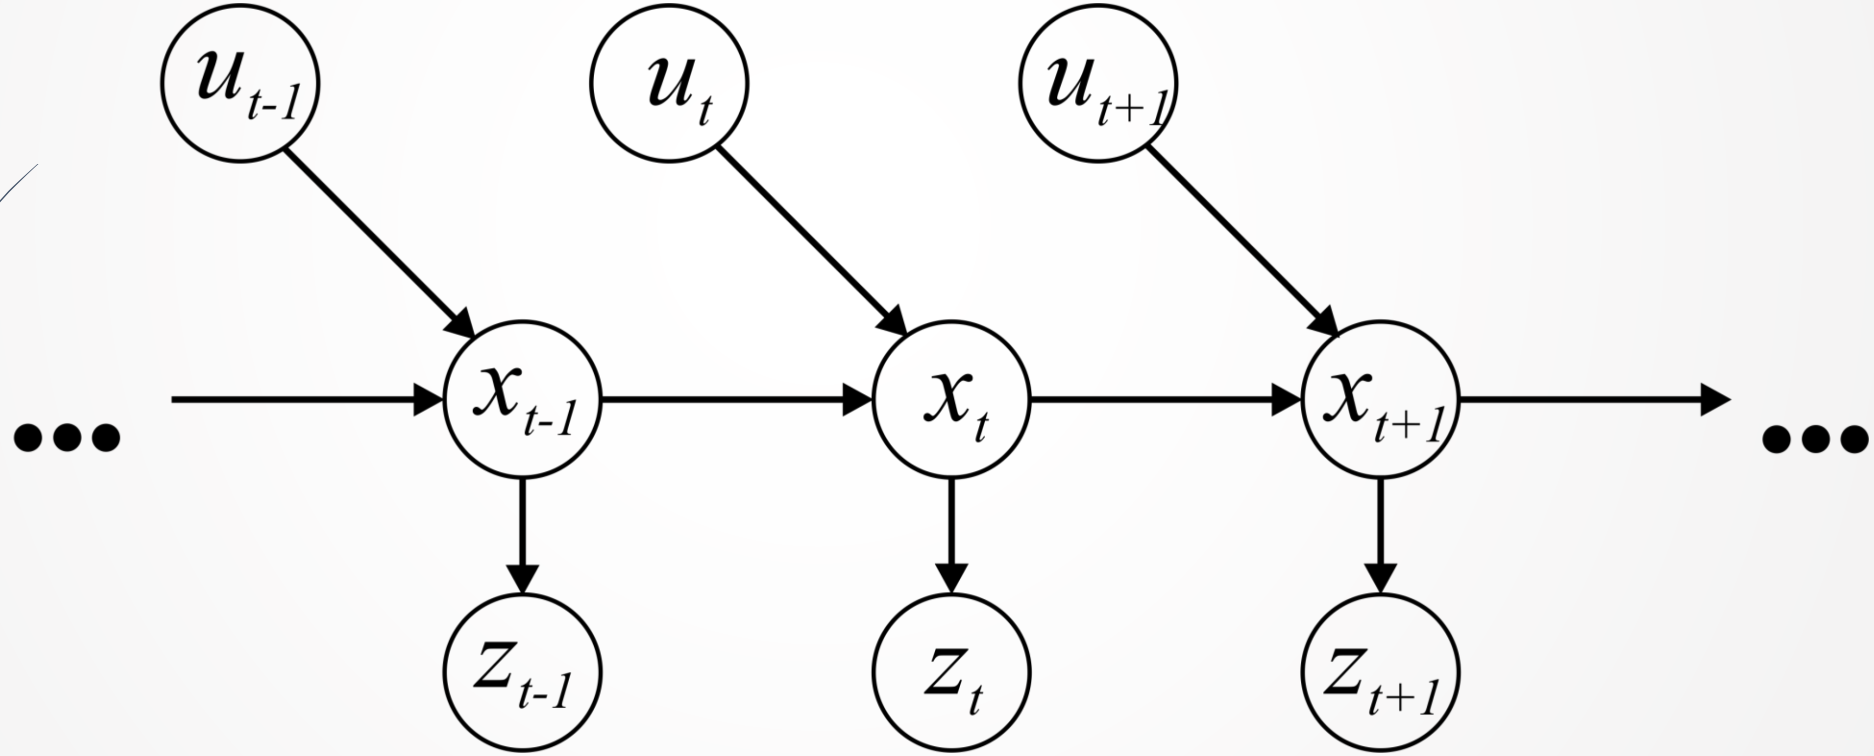

Underlying Assumptions

- Static world
- Independent noise
- Perfect model, no approximation errors

Bayes Filter

- **Given**
  - Sequence of observations and actions: z_t, u_t
  - Sensor model: $P(z|x)$
  - Action model: $P(x'|x, u)$
  - Prior probability of the system state: $P(x)$
- **Desired**
  - Estimate of the state of the dynamic system: x
  - Posterior of the state is also called belief:

$$Bel(x_t) = P(x_t|u_1, z_1, \dots, u_t, z_t)$$

Bayes Filter Algorithm

- **For each time step, do:**
  - Apply motion model:

$$\overline{Bel}(x_t) = \sum_{x_t - 1} P(x_t | x_{t-1}, u_t) Bel(x_{t-1})$$

  - Apply sensor model:

$$Bel(x_t) = \eta P(z_t | x_t) \overline{Bel}(x_t)$$

  - η is a normalization factor to ensure that the probability is maximum 1.

Example: Localization

- Discrete state: $x \in \{1, 2, ..., w\} \times \{1, 2, ..., h\}$
- Belief distribution can be represented as a grid
- This is also called a **historigram filter**

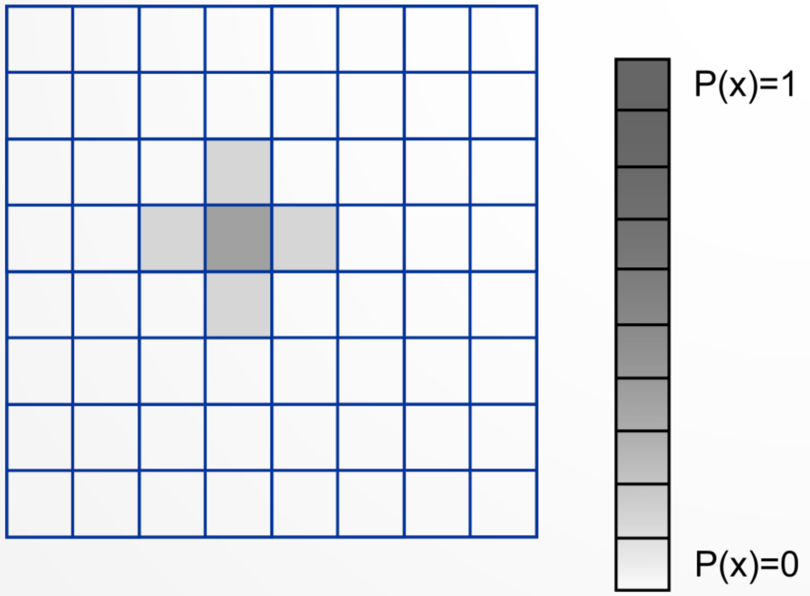

Example: Localization

- Action: $u \in \{north, east, south, west\}$
- Robot can move one cell in each time step
- Actions are not perfectly executed

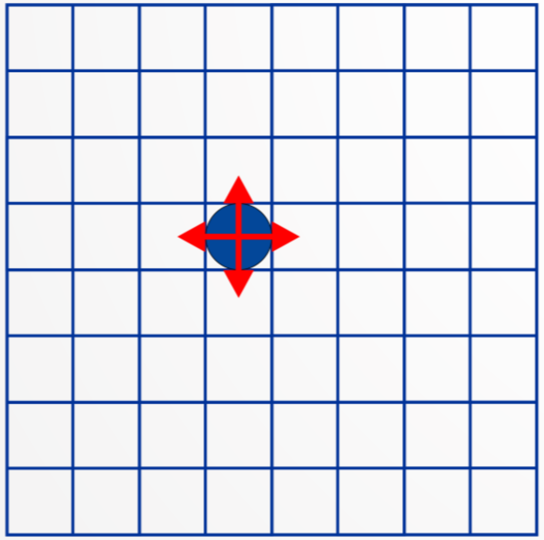

Example: Localization

- Action
- Robot can move one cell in each time step
- Actions are not perfectly executed
- Example: move east

$x_{t-1} =$ 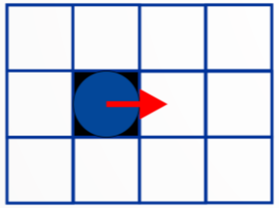 $u = east$ 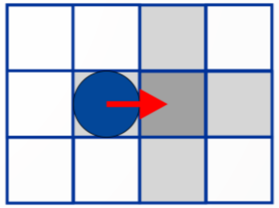

- 60% success rate, 10% to stay/move too far/ move one up/ move one down

Example: Localization

- Binary observation: $z \in \{marker, \bar{marker}\}$
- One (special) location has a marker
- Marker is sometimes also detected in neighboring cells

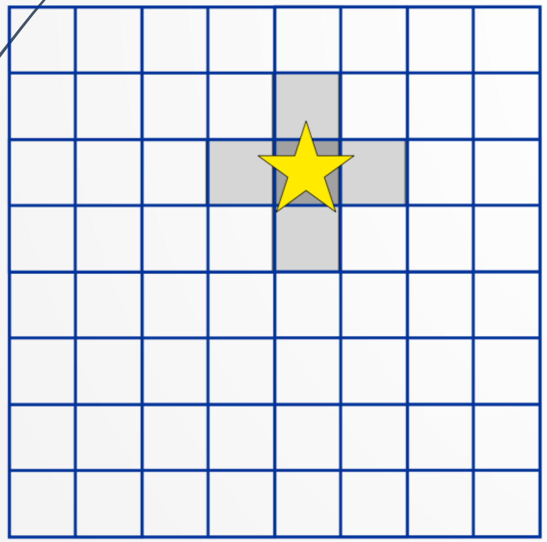

Example: Localization

- Let’s start a simulation run…

Example: Localization

- t=0
- Prior distribution (initial belief)
- Assume that we know the initial location (if not, we could initialize with a uniform prior)

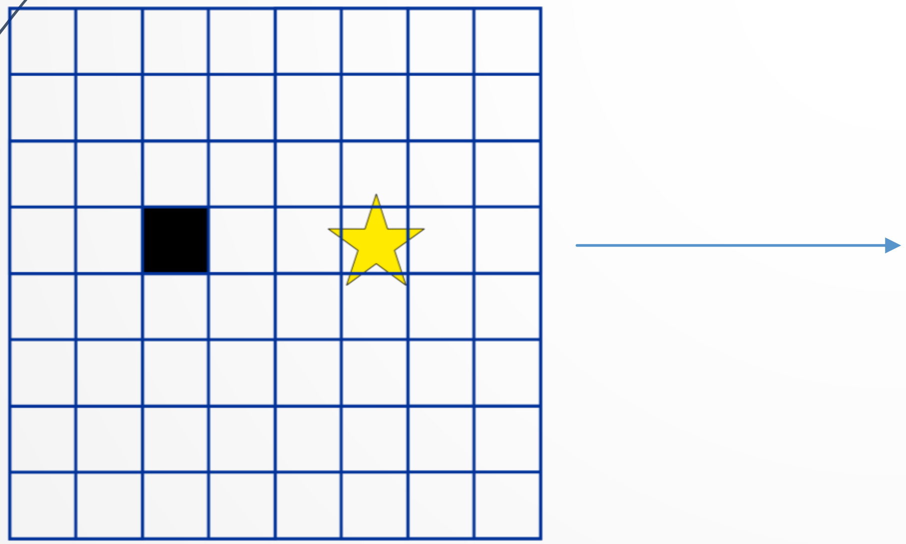

Example: Localization

- t=1, u =east, z=no-marker
- Bayes filter step 1: Apply motion model

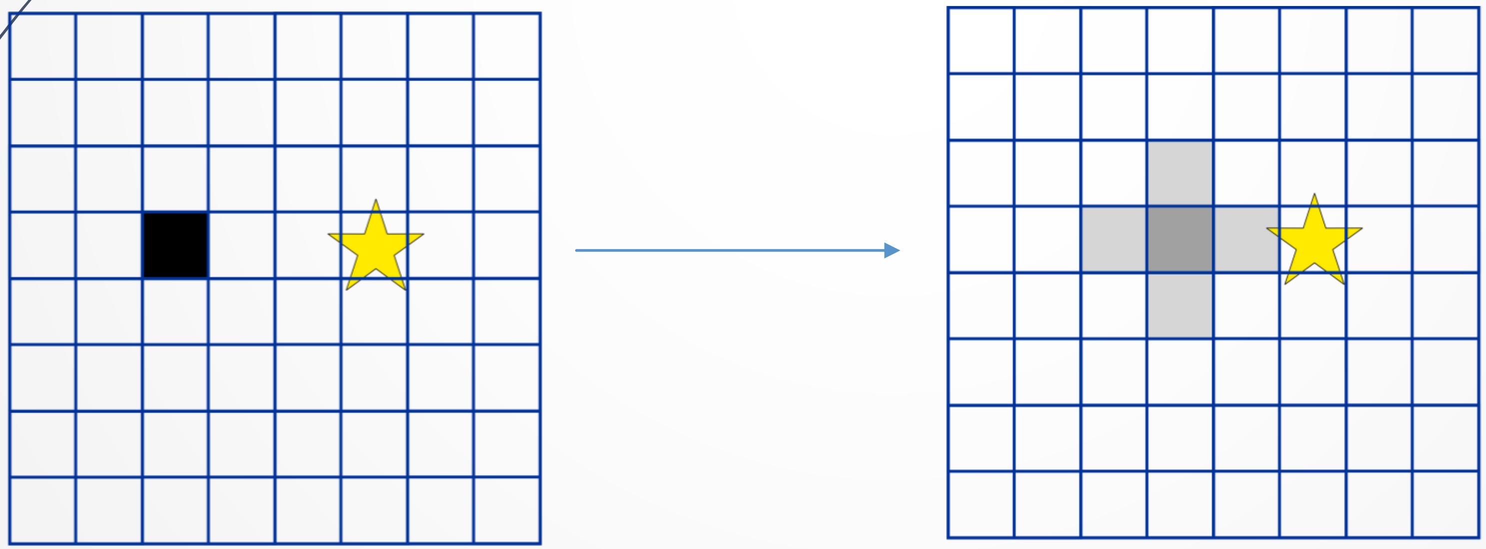

Example: Localization

- t=1, u =east, z=no-marker
- Bayes filter step 2: Apply observation model

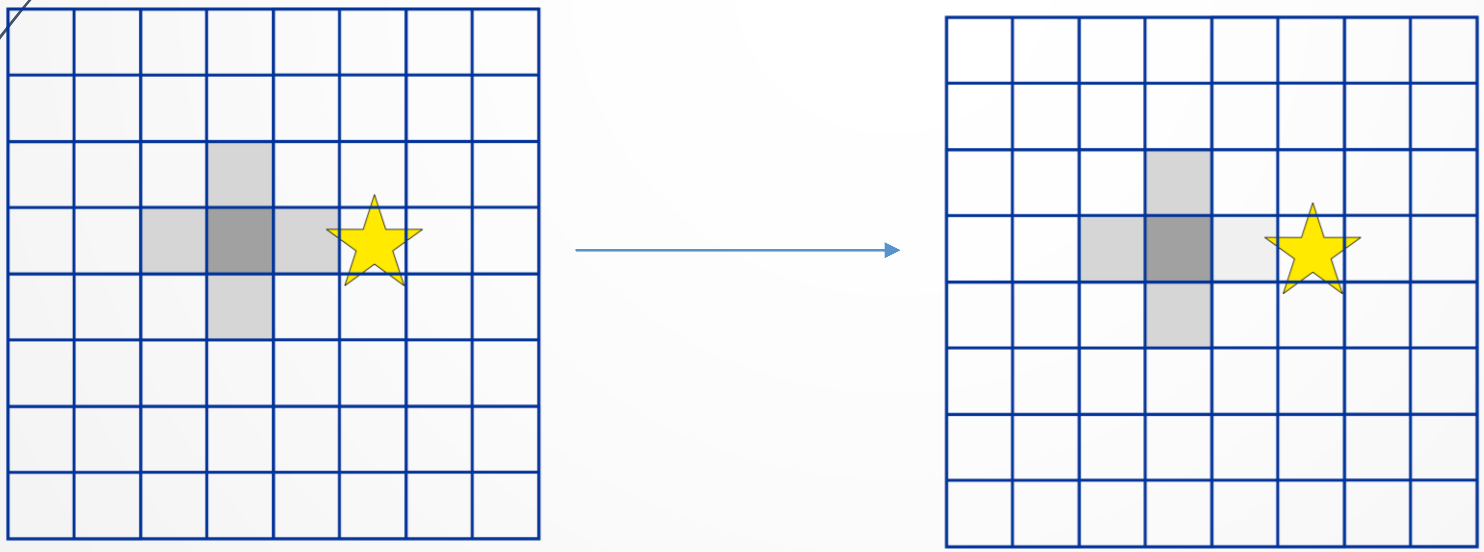

Example: Localization

- t=2, u =east, z=marker
- Bayes filter step 1: Apply motion model

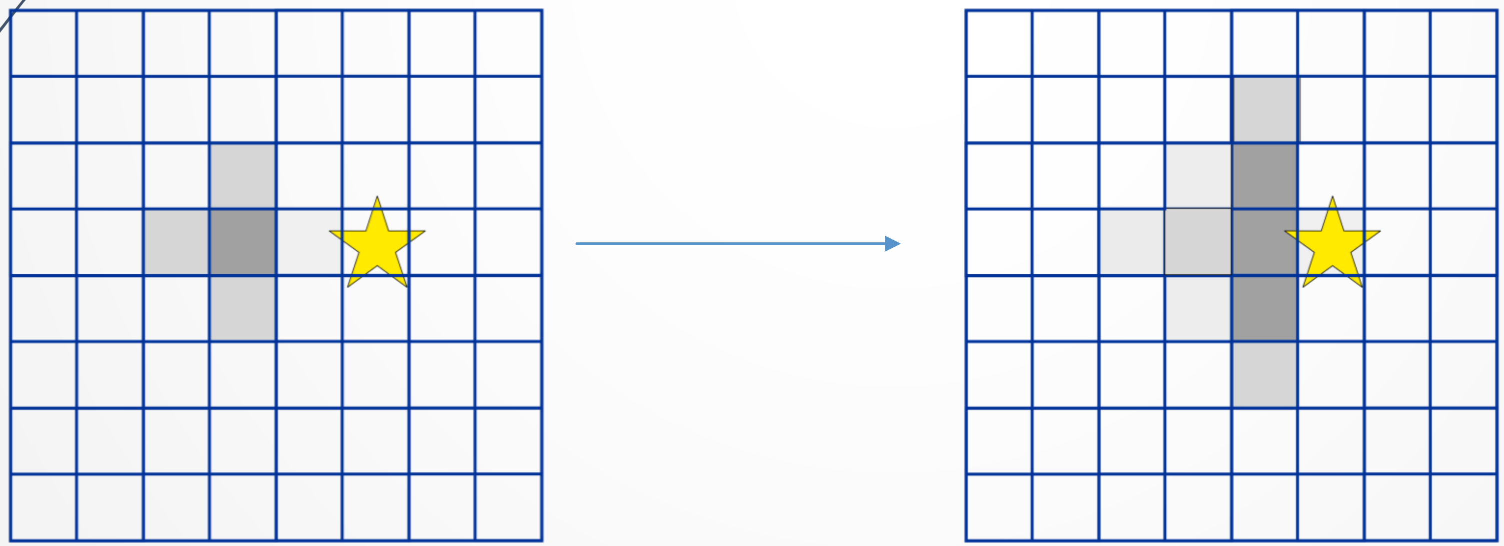

Example: Localization

- t=2, u =east, z=marker
- Bayes filter step 2: Apply observation model
- Question: where is the robot?

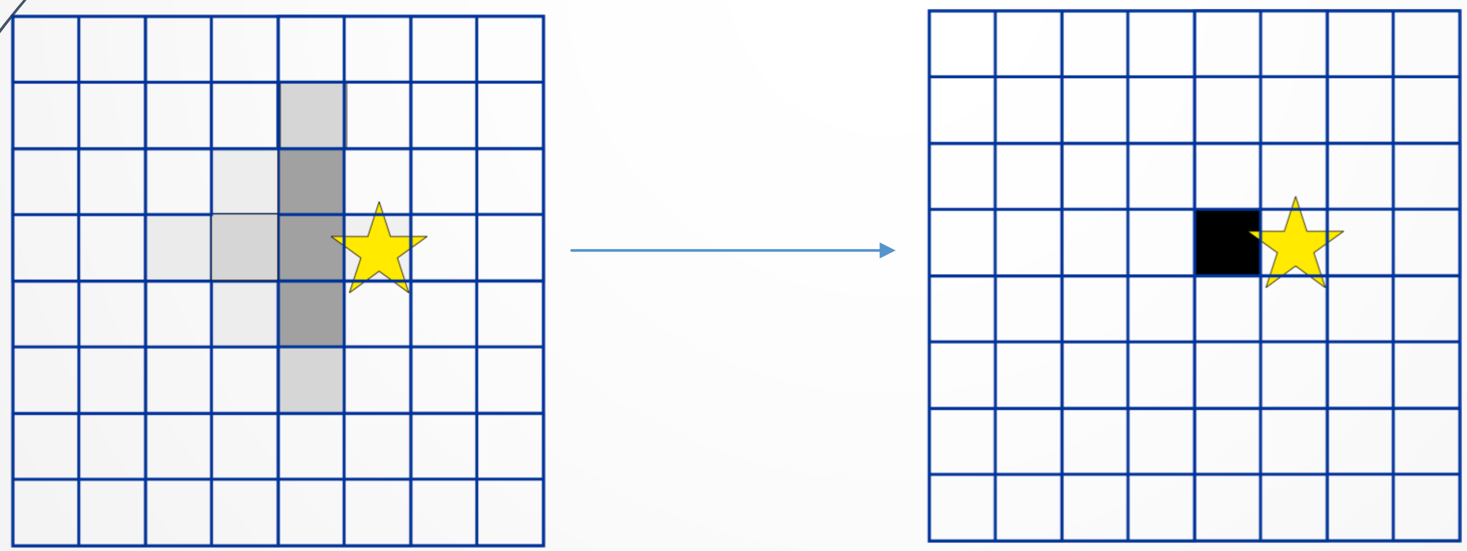

Find out more

- http://www.autonomousrobotslab.com/the-kalman-filter.html
- http://aerostudents.com/files/probabilityAndStatistics/probabilityTheoryFullVersion.pdf
- http://www.cs.unc.edu/~welch/kalman/
- http://home.wlu.edu/~levys/kalman_tutorial/
- https://github.com/rlabbe/Kalman-and-Bayesian-Filters-in-Python
- http://www.autonomousrobotslab.com/literature-and-links.html

Thank you!

Please ask your question!