Autonomous Navigation for Flying Robots

Lecture 6.1: Bayes Filter

Jürgen Sturm
Technische Universität München

Markov Assumption

- Observations depend only on current state

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

- Current state depends only on previous state and current action

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

Markov Chain

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

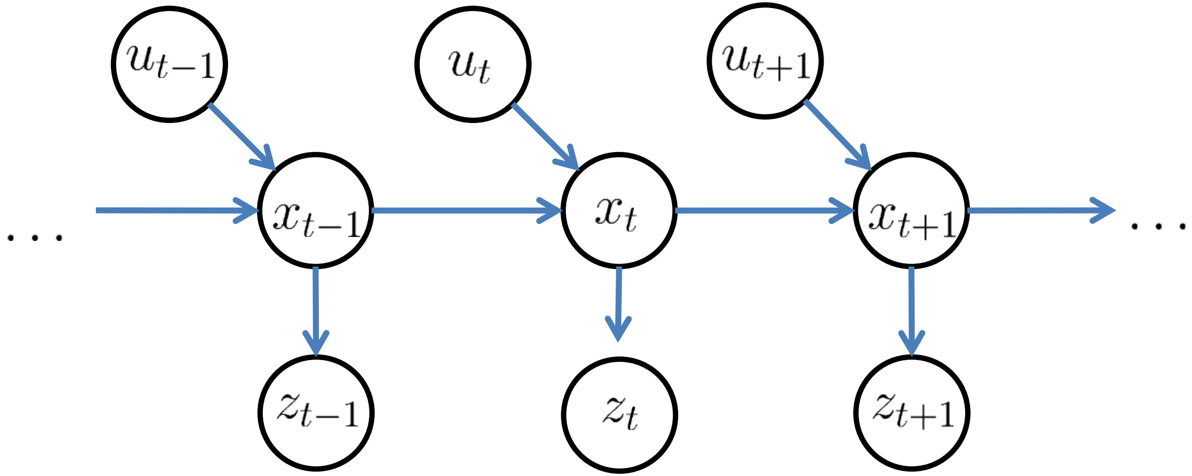

Underlying Assumptions

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

Bayes Filter

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

$$Bel(x_t) = P(x_t \mid u_1, z_1, \ldots, u_t, z_t)$$

Bayes Filter Algorithm

For each time step, do

1. Apply motion model

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

2. Apply sensor model

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

Notes

- Bayes filters also work on continuous state spaces (replace sum by integral)
- Bayes filter also works when actions and observations are asynchronous

Example: Localization

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

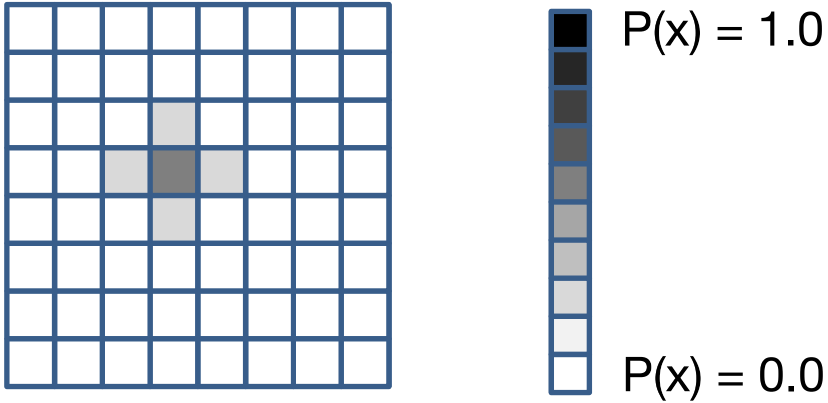

Example: Localization

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

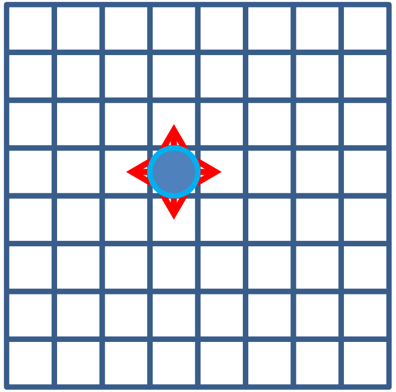

Example: Localization

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

$x_{t-1} =$ 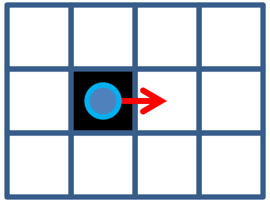 $, u = \text{east} \Rightarrow$ 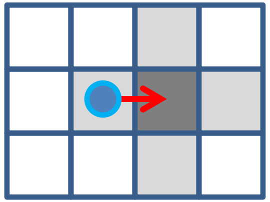

60% success rate, 10% to stay/move too far/ move one up/move one down

Example: Localization

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

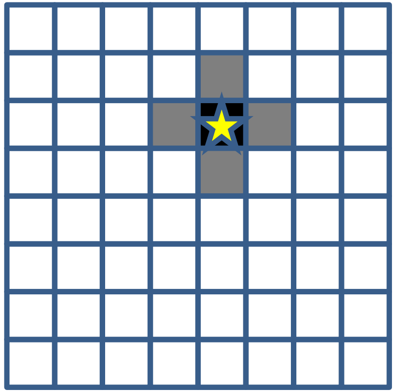

Example: Localization

- Let’s start a simulation run… (shades are hand-drawn, not exact!)

Example: Localization

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

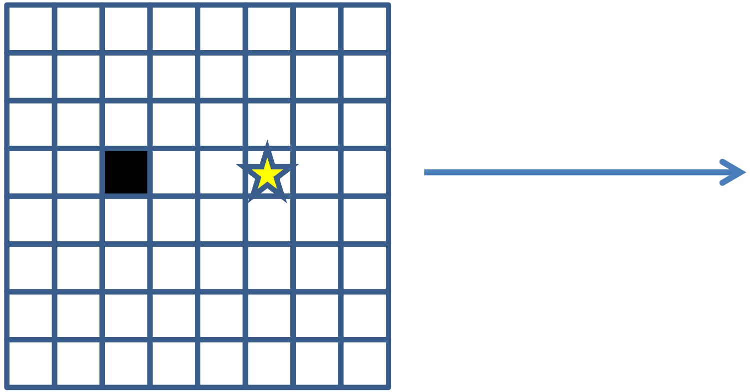

Example: Localization

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

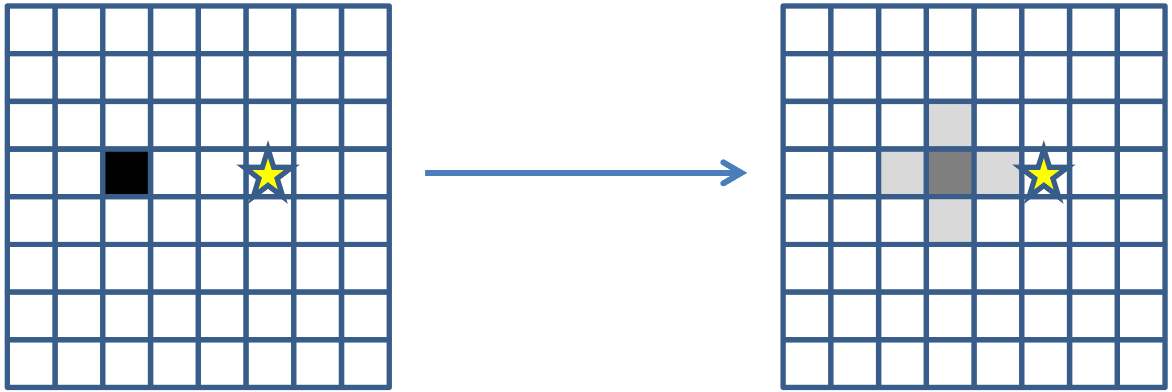

Example: Localization

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

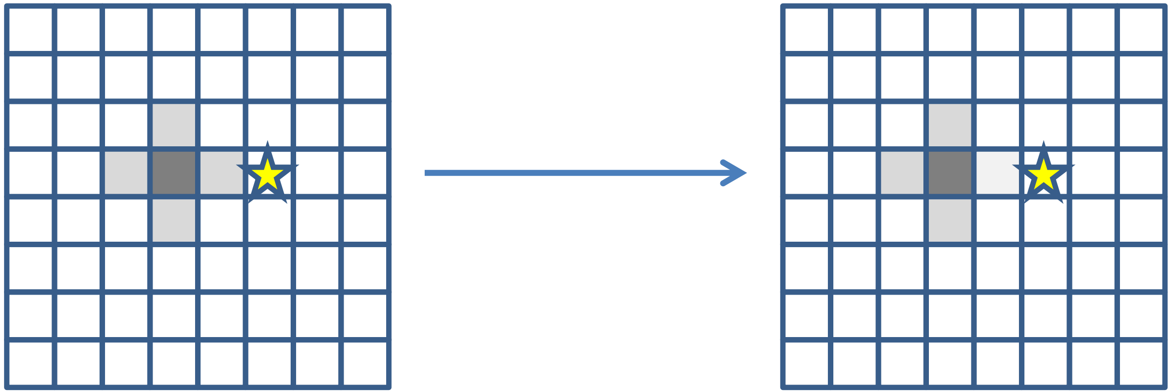

Example: Localization

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

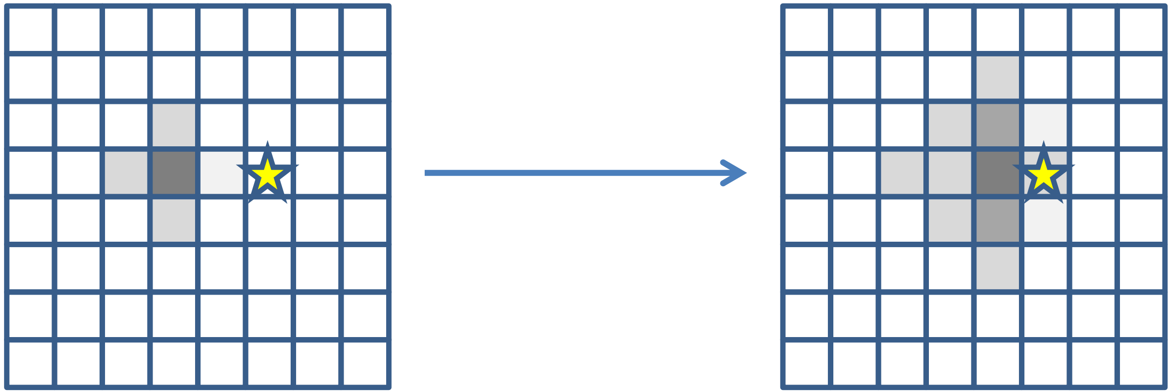

Example: Localization

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

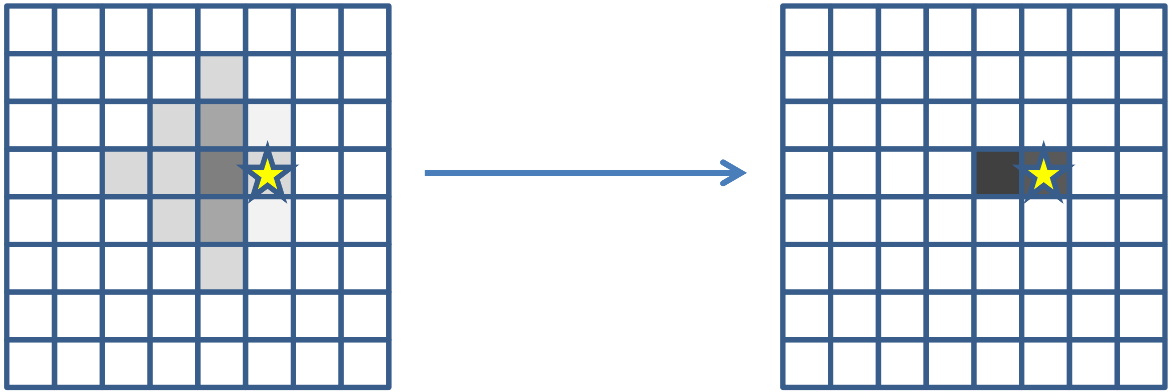

Lessons Learned

- Markov assumption allows efficient recursive Bayesian updates of the belief distribution
- Useful tool for estimating the state of a dynamic system
- Bayes filter is the basis of many other filters
  - Grid/histogram filter
  - **Kalman filter**
  - Particle filter
  - …